August 1, 1996

Prepared by Systems Division

Compaq Computer Corporation

Compaq Server Array Clustering Strategy and Product Direction

CONTENTS

EXECUTIVE SUMMARY

To meet the growing demands of enterprise computing, Compaq has devised a clustering strategy for delivering true midrange-class availability, performance, and scalability via open server solutions. Central to this strategy is a technology alliance with industry partners who share the goal of setting the industry standard for clustering x86 servers operating under Microsoft Windows NT: Tandem Computer Incorporated and Microsoft Corporation.

Compaq is developing a new family of high availability clustering solutions called Server Arrays. Compaq Server Arrays will link two or more high-performance Compaq servers for both scalability and availability enhancements. These servers will be connected via a high-speed, very low-latency ServerNet interconnect developed with Tandem for server-to-server connections. The servers will be connected to storage subsystems via high-bandwidth, industry standard fibre channel technology. Server Arrays will incorporate new clustering software being developed by Microsoft and will be supported by database management systems such as Oracle Parallel Server.

Compaq Server Arrays will offer compelling advantages for enterprise computing:

- Highly available systems
- High performance
- Pay-as-you-go scalability
- Very attractive price points
- Data integrity
- Flexible growth options

Building upon its Recovery Server products, Compaq expects to make available the next product in the Server Array family, a 2-server failover system, in fourth quarter 1996. In 1997 Compaq expects to introduce not only a multi-server failover system with full Microsoft Wolfpack functionality, but also multi-server parallel processing systems.

This technical brief provides an overview of the Compaq Server Array clustering strategy and the family of Server Array solutions Compaq expects to offer in the next 18 months.

Please direct comments regarding this communication via internal bmail to this address:
Tech Com@HW Tech@Sys Hou
Compaq field personnel may also send comments via the internet to this address:
tech_com@bangate.compaq.com

COMPAQ

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL. THIS INFORMATION IS PROVIDED "AS IS." COMPAQ MAKES NO REPRESENTATION OR WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, OR STATUTORY, AND EXPRESSLY DISCLAIMS THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR PARTICULAR PURPOSE, GOOD TITLE AND NONINFRINGEMENT.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

Compaq Server Array Clustering Strategy and Product Direction
First Edition (August 1, 1996)
Document Number 098A/0896

INTRODUCTION

A revolution in business computing is underway. Explosive growth in corporate computing needs and technological changes are driving the convergence of the centralized and distributed computing models. As their computing needs grow and change, corporations are moving their operations to a decentralized business and a computing environment characterized by distributed processing, server consolidation, larger files and databases, and 7x24 operation. These market trends have significant ramifications for system availability, performance, and scalability.

Compaq has pioneered this *distributed enterprise*, a flexible, networked information infrastructure that delivers instant access to business-critical data and empowers local decision-makers. Compaq's breakthrough server technology, integration expertise, and key industry partnerships provide enterprise solutions that challenge traditional mainframe and proprietary midrange computing paradigms.

Compaq enterprise solutions include a strategy to deliver true midrange-class availability, performance, and scalability via open server solutions. This strategy evolved as Compaq sought the most effective ways to meet or exceed the levels of availability to which corporate customers were accustomed with proprietary midrange solutions—and to do that at more cost-effective price points.

This technical brief is an overview of the Compaq Server Array clustering strategy and the family of Server Array solutions Compaq expects to offer in the next 18 months.

WHAT IS CLUSTERING?

Compaq defines a cluster as a group of loosely coupled systems (that is, each having its own dedicated memory) that work collectively as a single system to provide fast, uninterrupted computing service. Clustering architecture provides the means for significantly increasing system scalability—a way to increase processing capacity and add I/O bandwidth—while achieving the highest possible levels of availability.

A cluster can address either or both of two distinct modes of operation: *failover* and *parallel processing*.

Failover

Failover solutions are primarily used to achieve high availability. They can provide continued user access to critical data on shared disk storage after a server failure. Failover alone does not increase scalability.

In the failover mode of operation, pairs of servers—one or both of which may include external storage—are connected by a server communication interconnect cable (Figure 1). Failover cluster solutions may use SCSI connections or high-speed LAN (local area network) connections for the server communication interconnect. External storage system(s) are connected to both servers via another interconnect. If one of the servers fails, its external storage system(s) automatically switch over to the paired server so that users can quickly resume accessing those files. The Compaq Recovery Server Option and Digital Clusters for Windows NT are examples of failover solutions. Most clustering solutions now available or soon to be available are capable of failover processing.

By automating the failure recovery process, the failover operating mode offers cost-effective solutions for such environments as remote locations that do not have on-site system administrators. It is important to remember, however, that all clients of a failed server and all applications running on it may experience brief downtime while the storage system(s) fail over to the paired server and software is restarted.

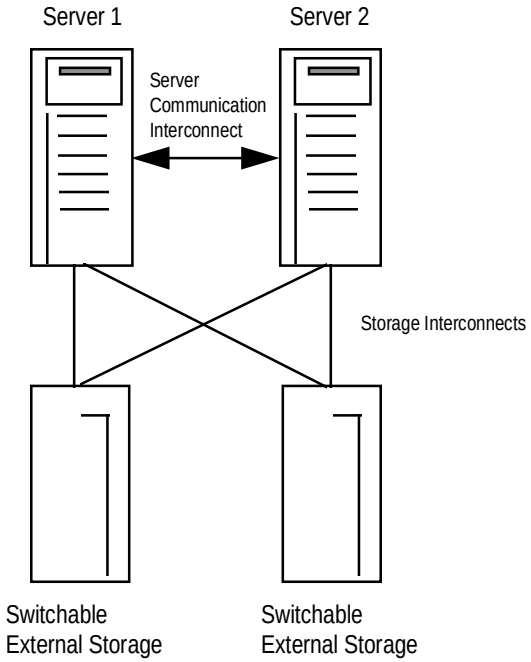

Figure 1. A server pair configured for operating in failover mode

Parallel Processing

Parallel processing solutions can provide the same availability benefits as failover solutions and have significant performance advantages. In a parallel processing environment, multiple instances of an application reside on two or more servers. Performance is improved, sometimes dramatically, because there are multiple paths to access requested data. Figure 2 illustrates a loosely coupled, parallel processing cluster.

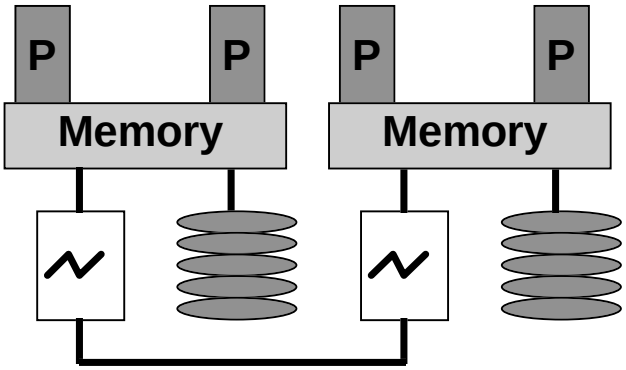

Figure 2. A loosely coupled, parallel processing cluster

Parallel processing can also improve availability. In the event of a server failure, only one instance of the application is lost; therefore, only a portion of the total application users are affected by the minimal downtime required to restart the lost instance. Availability can be truly maximized if requests to the failed server are immediately routed to another server so that all users continue to operate from other servers in the cluster and none even notices the server failure.

Truly distributed or parallel applications are responsible for coordinating the work to be done among the multiple, parallel instances of the application. They must address synchronization of access to data that is being updated. At the same time, they must effectively distribute the workload among the clustered servers and preserve data integrity in the event of a server failure. All of the leading database management systems (DBMS) have significant experience and technology that address these issues. Other software vendors (including Microsoft) are developing tools that solve these same issues in a more general (non-DBMS) environment.

Table 1 summarizes the key similarities and differences between a cluster operating in failover mode and a cluster operating in parallel processing mode.

TABLE 1: COMPARISON OF FAILOVER AND PARALLEL PROCESSING MODES

	Failover Mode	Parallel Processing Mode
High availability	yes	yes
Increased performance	no	yes
High scalability	no	yes
No. of servers in cluster	2 or more	2 or more
Effect of server failure	ALL clients of failed server and of all applications experience downtime.	A portion of the total application users experience minimal (perhaps not even noticeable) downtime.

COMPAQ SERVER ARRAY CLUSTERING STRATEGY

Understanding Compaq's Server Array clustering strategy requires a look at three factors: the core objectives for Compaq clustering solutions, Compaq's reasons for choosing the clustering approach to parallel processing, and the technologies being used to implement the strategy.

Objectives of Compaq Clustering Solutions

For its clustering solutions, Compaq set the following core objectives:

- **Performance:** meet or exceed midrange RISC performance.
- **Bandwidth:** high bandwidth; bandwidth must not limit performance.
- **Latency:** extremely low latency; significant improvement over traditional LAN interconnect.
- **Availability:** support for redundant interface boards and switches, transparent client re-connect, preservation of the transaction state, and hardware failure recovery.
- **Scalability:** support for small and large numbers of servers with the same cost-effective hardware.
- **Reliability:** data packet acknowledgment, self-checking, and fault tolerance.
- **Investment Protection:** industry-standard architecture that will last multiple generations.

- **Ease of Use:** a plug-and-play solution. A user should be able to plug an interconnect card into a PCI slot, connect it with the cabling provided with the card, load software via Compaq SmartStart, and then be off and running.

Again as a reminder, it is important to understand that a cluster may address two distinct modes of operation: failover and parallel processing. Because each offers advantages for different user environments, Compaq has elected to use both and to offer multiple solutions using the same core technologies.

Compaq's availability goal for parallel clustering solutions is absolutely minimum unscheduled downtime. In such a cluster, a server may fail or be shut down for planned maintenance, but a user's access to critical data should never be lost. The load of the downed server can be shifted to other servers in the cluster with no interruption in processing.

Why Choose Clustering?

In the past, clustering has been a viable enterprise solution available only for expensive, proprietary mainframes or minicomputers. The Open VMS Cluster from Digital Equipment Corporation (introduced in 1982 as the VAXcluster) was a pioneer in this category. Today the situation is changing, however. Previous barriers to open system clustering are falling because these key elements are now emerging:

- Very high-performance microprocessors
- Standard high-speed interconnects with higher bandwidth and reliability than traditional LANs
- Standard tools for distributed computing
- Effective system management software
- Cluster-aware software that will run across multiple servers in parallel
- Crossbar switches that can yield greater aggregate external bandwidth than the aggregate internal bandwidth of a single bus connecting the same number of processors inside a symmetrical multiprocessing system
- Greater reliability of the interconnect transports

For enterprise solutions, Compaq has chosen the open system clustering approach to parallel processing because it offers significant advantages in price/performance, availability, and scalability.

Price/Performance

Compaq servers are based on the Intel x86 processor architecture. As Figure 3 illustrates, performance of Compaq servers (as measured by the number of transactions processed per second) has improved year by year. When Compaq first transitioned from a uniprocessor environment to a dual-processor environment in 1989, performance by its servers climbed. With the transition from two processors to four processors in 1993, performance by its servers climbed again. In audited TPC-C benchmarks published June 3, 1996, the new ProLiant 5000 Server—powered by four 166-MHz Pentium® Pro processors—delivered midrange-level performance of 5676 tpmC with a price/performance value of $136/tpmC.[1] For a number of reasons, however, cost-effective performance scaling is currently limited to a four-processor system.

[1] Configuration comprised of ProLiant 5000 6/166, four processors, 2-gigabyte RAM, 391-gigabyte total storage, Microsoft Windows NT Server 4.0, Microsoft SQL Server 6.5.

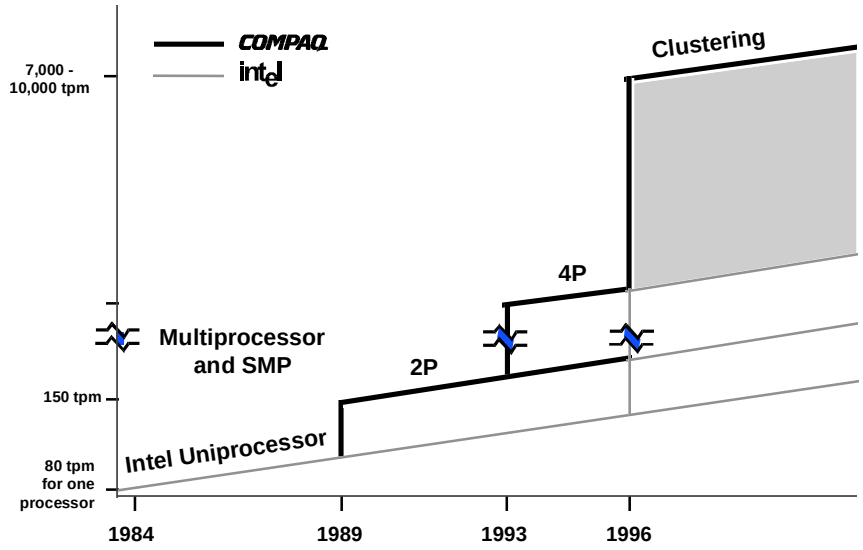

Figure 3. Performance growth of Compaq servers

Customers selecting a database have available to them proprietary UNIX alternatives that will scale beyond four processors. Some customers may assume that, as a matter of course, having more processors equates to greater performance. That misconception could result in an inaccurate assessment of true performance. In fact, TPC-C test results reported as of June 28, 1996, show that the ProLiant 5000 Server outperforms not only competing 4-processor servers from Digital and HP, but also 8-processor and 16-processor RISC systems from IBM and NCR respectively. Moreover, as Figure 4 indicates, the cost per transaction of the ProLiant 5000 is one half or less that of the other servers.

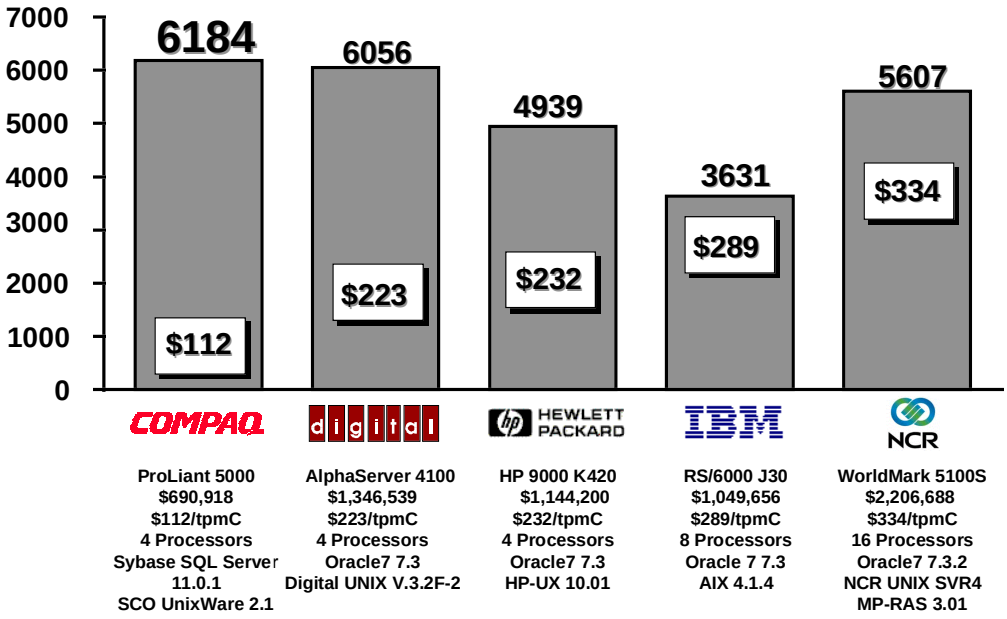

Figure 4: Comparison of price/performance in five mid-range class servers

Compaq believes that open system clustering offers the promise of circumventing what would otherwise be a scaling limitation to allow the same high level of performance at far lower cost to users. Clearly, then, one reason for clustering is to enhance performance.

Availability

Availability refers to the ability of users to access data stored on a system. As important as performance is, work does not get done if users lose access to the data they need. Availability has traditionally been a system level problem. Although the progressive transitions in Intel x86 Processor Architecture illustrated in Figure 3 pushed performance up, they did not improve availability. Responsibility for addressing availability has fallen to system providers. Among other things, Compaq has implemented Redundant Arrays of Inexpensive Disks (RAID), Error Checking and Correcting (ECC) memory, and redundant components to increase availability (Figure 5).

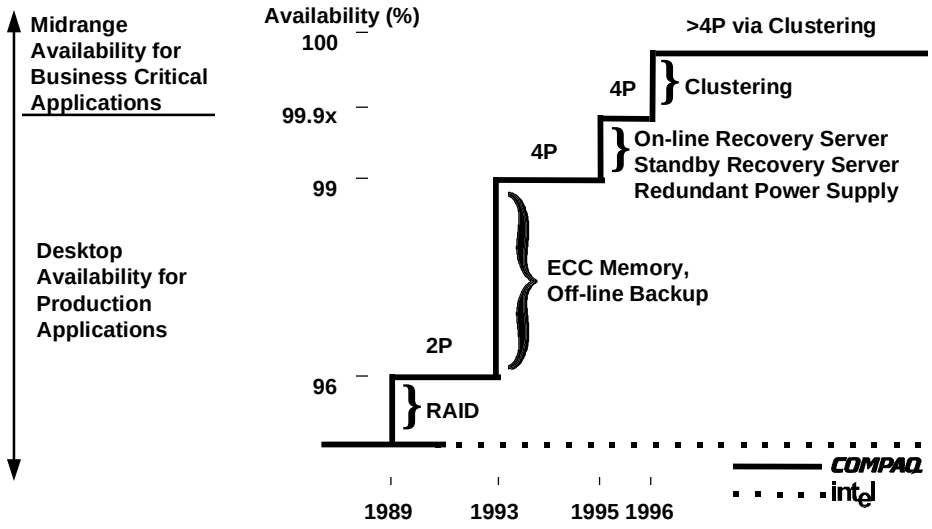

Figure 5. Increases in availability resulting from hardware advances

Today availability problems are most often software related. The availability levels indicated in Table 2 have been widely acknowledged in industry literature. Hardware advances have raised availability in conventional servers to the 99% level. Further improvements in availability through hardware enhancements have been possible only with the addition of high-availability, fault resilient, and fault tolerant features. Enterprise computing often requires the level of availability provided only in fault tolerant systems. Fault tolerant systems have the ability to recover from hardware problems without interrupting server performance. Some midrange developers offer fully redundant components and sometimes fully redundant systems. Usually, however, full redundancy can be accomplished only at substantial costs.

TABLE 2

Availability (%)	Total Downtime	Server Class
99	3.5 days/year	Conventional
99.9	8.5 hours/year	High Availability
99.99	1 hour/year	Fault Resilient
99.999	5 minutes/year	Fault Tolerant

Through the Compaq Recovery Server Option, Compaq offers two high-availability failover solutions that can address the needs of business critical environments. Compaq believes, however, that parallel processing is the high-availability solution for Compaq systems that will deliver proprietary midrange-level availability for business critical applications. A second reason for clustering technology, then, is to enhance availability.

Clustering competes with a number of other technologies in terms of performance and availability. Figure 6 indicates the relative strengths of Symmetrical Multiprocessing (SMP), Massively Parallel Processing (MPP), cluster, and fault tolerant solutions in terms of performance and availability. Clustering offers performance that surpasses that of SMP and competes well with MPP. In addition, clustering offers availability that far surpasses either SMP or MPP and approaches levels of true fault tolerance that are usually associated only with very expensive proprietary mainframes. The fact that clustering can significantly improve both performance and availability will make it very popular among enterprise-class customers and vendors alike.

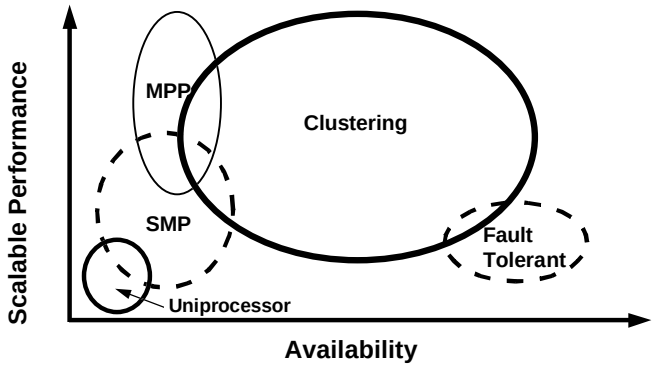

Figure 6. Relative strengths of SMP, MPP, cluster, and fault tolerant solutions[2]

Scalability

Scalability refers to the ability of a system (server, storage, network, or cluster) to grow or to expand. When used to describe clustering, the term is often used to measure performance. In this sense it refers to the available bandwidth of the interconnect and the effect on bandwidth as the number of users and the number or size of data packets grow.

Scalability of clusters will vary from vendor to vendor depending upon the architecture used. In fact many cluster solutions, including Clusters for Windows NT from Digital and Compaq On-Line Recovery Server (both 2-server failover systems), do not scale because they were designed to address high-availability issues. Compaq parallel processing cluster solutions will feature interconnect technology that provides essentially unlimited scalability.

[2] Illustration reproduced from the brochure "Advantage Cluster, Digital's UNIX Cluster: Building Highly Scalable, Highly Available Servers with Alpha AXP and DEC OSF/1," Digital Equipment Corporation, September 1994.

Manageability

Manageability refers to the ability of a system administrator to monitor performance and fault analysis information for the entire cluster by means of management applications and agents, and to control all the managed devices from a system console. Compaq's clustering solutions are the result of a close-knit industry partnership among several major hardware and software developers. Compaq solutions will combine the strengths of Compaq Insight Manager and Microsoft Windows NT management tools with additional capabilities that address challenges specific to cluster management. Since administrative systems available today can adequately manage 2-server systems and multi-server failover, Compaq is focusing on an administrative system that supports multi-server scalability.

Technology of Compaq Server Arrays

Compaq has chosen *Server Array* as the family name for its high-availability, high-performance clustering solutions. This name was chosen to show the system level analogy to current RAID technologies developed and proven within the mass storage arena. In storage systems, the desire is always for increases in availability, capacity, and performance. One approach for accomplishing this would be to develop specialized disk systems to provide those benefits. However, that approach to mass storage would require highly specialized drives, thereby sacrificing the price/performance benefits of being common with broadly available solutions. Instead, with RAID technology the desired characteristics of availability, capacity, and performance are provided to the system by presenting an architectural image of a specialized subsystem built from industry standard components.

In developing Server Arrays, Compaq is extending this concept to the system level by presenting to the end user the architectural image of a high-availability, high-capacity, high-performance system constructed from Compaq's industry standard servers and storage subsystems. Scalable Server Arrays will allow multiple instances of an application to reside and run concurrently on multiple connected servers, thereby increasing processing speed and availability.

From the beginning, an integral part of Compaq's plan to develop open-system enterprise computing solutions has been partnering with key industry leaders. In a news release dated October 16, 1995, Compaq announced the formation of a technology alliance with Tandem Computer Incorporated and Microsoft Corporation. With the common goal of establishing the industry standard for clustering x86 servers running Microsoft Windows NT, each partner in this alliance is contributing an important element of Compaq clustering solutions.

Servers

For its clustering solutions, Compaq will link cost-effective, high-performance ProLiant servers and/or ProSignia servers to create Server Arrays. With their Full-Spectrum Fault Management technology, today's Compaq ProLiant servers have a comprehensive set of features to deliver the levels of availability required for a broad range of applications. With features like hot-pluggable RAID, ECC memory, Uninterruptable Power Supplies, and Redundant Power Supplies, ProLiant servers are tolerant of a wide variety of system faults. Should a fault occur in a ProLiant server, Automatic Server Recovery will attempt to reboot the system automatically and bypass the subsystem that caused the failure.

Interconnects

For the critical interconnect component of its Server Arrays, Compaq is using not one, but two leading edge hardware technologies to capitalize on the advantages of each: fibre channel for device communication with shared system storage devices and ServerNet for efficient and reliable server-to-server communication (Figure 7).

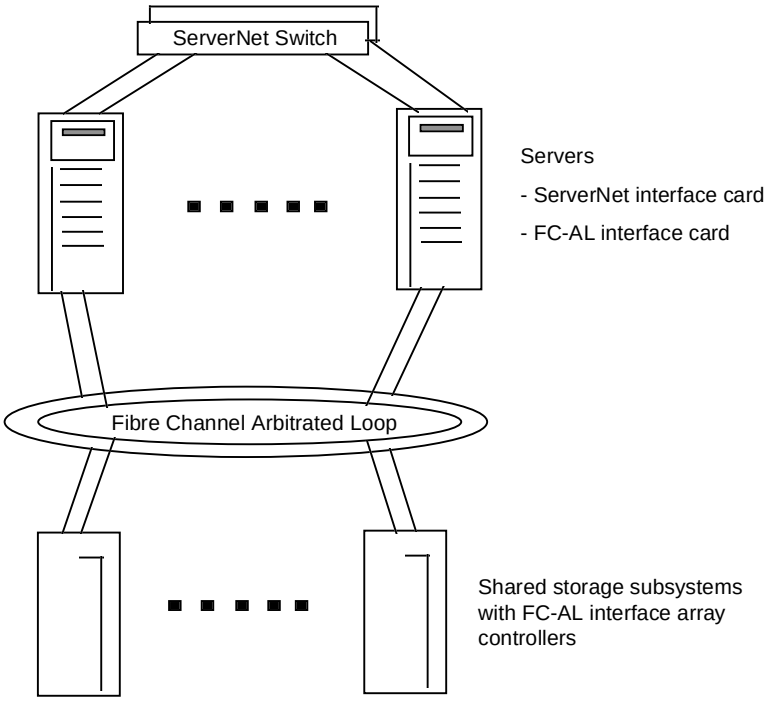

Figure 7. Compaq Server Arrays use fibre channel arbitrated loop for the shared storage interconnect and Tandem ServerNet for the server communication interconnect.

For the shared storage interface in its Server Arrays, Compaq will use fibre channel arbitrated loop (FC-AL). Fibre channel is a high-speed interconnect that will eventually attain speeds of up to 1 gigabit/second. It provides high burst bandwidth and is the emerging industry standard for storage interconnect. It is optimized for high-speed transfers of large amounts of data. Today it is capable of delivering 100 megabytes/second of data both into and out of a server simultaneously and over a distance of up to 10 kilometers.

Fibre channel already enjoys broad industry support, and it meets the technical goals of next generation server-to-storage interconnects. It offers these benefits for storage:

- Increased capacity through greater connectivity per server slot
- Greater availability through full redundancy and sharing
- Dynamic storage expansion via hot-pluggable storage subsystems
- Increased performance by means of the 100 megabytes/second interconnect
- Integration of primary and secondary storage
- Greater distance between system elements: 25 meters with copper cable and 10 kilometers with fibre
- Ease of use through smaller cables and connectors

The other core technology for Compaq clustering solutions—the server communication interconnect—has been developed with Tandem Computer Incorporated. With 22 years of experience in the high-availability, enterprise market, Tandem is the acknowledged leader in fault tolerant computing. Tandem developed ServerNet, its fifth generation availability technology, to become the foundation of its next generation of systems architectures and intends for it to become an industry standard.

ServerNet is bi-directional, point-to-point, packetized, router-based hardware, coupled with a light weight, reliable, messaging passing software interconnect. ServerNet provides an OS-independent, high-bandwidth, low-latency architecture for moving large amounts of data to multiple users at unprecedented speeds. Each ServerNet connection can be automatically redundant, with multiple levels of fault detection, isolation, and correction. A data packet acknowledgment feature ensures data integrity. ServerNet can provide aggregate throughput from 100 megabytes/second to 400 terabytes/second by means of cascading, wormhole routing and will support advances in technology far into the future.

Compaq views ServerNet as being significantly more advanced than any other available server communication interconnect technology. ServerNet also supports Compaq objectives for clustering solutions in these ways:

- **Scalable Throughput:** In developing ServerNet, Tandem created an architecutre that would be useful in systems ranging from two to a million servers. (While the applications and operating system software may not scale to match, the server communication interconnect will not be the bottleneck to performance.)
- **Low Latency Response:** A problem with most network-based clustering architecture is that latencies for remote accesses are often measured in 10's or even 100's of microseconds. Tandem designed ServerNet for extremely low latency by carefully tuning the hardware and software architectures.
- **Reliability:** The ServerNet design includes support for redundant interconnects. Redundant data paths ensure that a data packet gets to its destination, even in the event of a server failure.
- **Continuous Operation and Manageability of the Interconnect Mechanism:** The ServerNet design allows the system manager to scale the system and perform routine system maintenance without disrupting services.
- **Flexible Configurations:** The environments into which ServerNet will be placed are variable and will become even more so. They will range from 2-server, remote site systems for which availability is the driving requirement, to data center, many-server systems for which scalability is important, but availability is a must.
- **Low Cost:** To meet the wide range of configurations, low-cost hardware implementation is a requirement. Tandem used standard semiconductor and board technology, and through its partnership with Compaq, has continued to drive the costs of implementation even lower.
- **Open Systems Support:** By partnering with Compaq and Microsoft, Tandem has demonstrated its desire for ServerNet to be an industry standard.

At first glance, a clustering strategy that employs two distinct interconnect technologies may appear disjunct. In fact, however, there are significant reasons driving the union of these two interconnects:

- While ServerNet could certainly be used to connect all devices (servers, external storage, and external backup), there is significant storage industry momentum behind fibre channel. Disk drive manufacturers such as Seagate are building drives with fibre channel interfaces built in.

- The technical requirements for the two interconnects are different. Server communication interconnects require absolutely minimum latency. Message latency, even if measured in microseconds, can stall MP systems for hundreds to thousands of instructions per message. Latency requirements for storage systems, on the other hand, are paced by the speeds of the storage controller and drives themselves. Therefore, latencies measured at up to a millisecond are transparent to the overall operation. The more significant requirement for storage systems is maximized sequential bandwidth. While both ServerNet and fibre channel have excellent latency and bandwidth characteristics, each is optimized for its own particular use.

Separate and distinct interconnects for separate and distinct data communication tasks will provide both performance (no contention for available bandwidth) and high availability (no single point of failure) benefits.

Operating Systems and Clustering Software

Compaq is working on clustering solutions that will be broadly compatible with multiple operating systems. As a result of the collaborative effort between Microsoft, Tandem, and Compaq to develop a standard for Microsoft Windows NT clustering, Compaq clustering solutions will initially be released on the Microsoft Windows NT Server platform. Microsoft is enhancing that platform by developing new clustering software, codenamed Wolfpack. Wolfpack is a new development effort with new code written by Microsoft. Wolfpack has been architected to support a number of clustering interconnects, from legacy LAN to high-performance, fault-tolerant interconnects such as ServerNet, which will be available from Compaq.

Wolfpack is an open standard for Microsoft Windows NT clustering, and it is fully compatible with Compaq's clustering efforts. Microsoft is developing and deploying Wolfpack in phases. Initially Wolfpack will support 2-server failover suitable for all applications. Later it will support more than two servers. For more information about Wolfpack, access the Microsoft web site on the Internet. Then use the Microsoft search engine to search for Wolfpack.

As a result of its alliance with Tandem for software technology (announced April 23, 1996), Microsoft has been able to move forward quickly on integration of cluster technology into both Microsoft Windows NT and Microsoft Back Office. Partnerships such those formed by Compaq with Tandem for interconnect harware and Microsoft with Tandem for high availability software are an effective means for the industry to move forward in a compatible fashion.

Database Management

The most commercially known parallel enterprise application is Oracle Parallel Server (OPS). Originally developed for the VAXcluster and later ported to many proprietary RISC UNIX platforms, OPS has demonstrated that clustering can lead to both high-performance and high-availability DBMS solutions. Multiple instances of Oracle 7 database run concurrently across multiple servers to enhance DBMS performance significantly. Moreover, the multiple instances of Oracle 7 can act as backup for each other in the event of a server failure or shutdown, thereby greatly improving the DBMS availability.

Compaq is working closely with Oracle to provide the significant performance and availability benefits of ServerNet to OPS customers.

COMPAQ SERVER ARRAY PRODUCT ROADMAP

Compaq clustering solutions are fully aligned with Microsoft Wolfpack. Like Wolfpack itself, they are being developed and introduced in a phased approach.

Compaq Failover Solutions

Compaq introduced the first of its Server Array products, the Compaq Recovery Server Option, in mid-1995. The Recovery Server Option connects pairs of servers by means of a SCSI storage interconnect and a serial, server communication interconnect. One or both of the servers may include external storage. If so, the external storage system(s) are connected to both servers by means of SCSI cables and a switch box. The Recovery Server Option can be implemented in either of two ways: the Standby Recovery Server configuration or the On-Line Recovery Server configuration.

Standby Recovery Server

The Compaq Standby Recovery Server configuration (Figure 8) maximizes server availability for customers with file servers and application servers at remote locations or branch offices without extensive on-site technical expertise. This configuration allows one Compaq ProLiant or ProSignia server (either tower or rack-mounted models) to act as a standby for another similarly configured ProLiant or ProSignia server. The two servers are attached by SCSI interconnect to a common set of ProLiant Storage Systems that contain a single copy of the operating system, applications, and data. If the primary server fails, each ProLiant Storage System automatically switches over from the primary to the recovery server. The recovery server then boots, and the system is back on-line in minutes without administrator intervention.

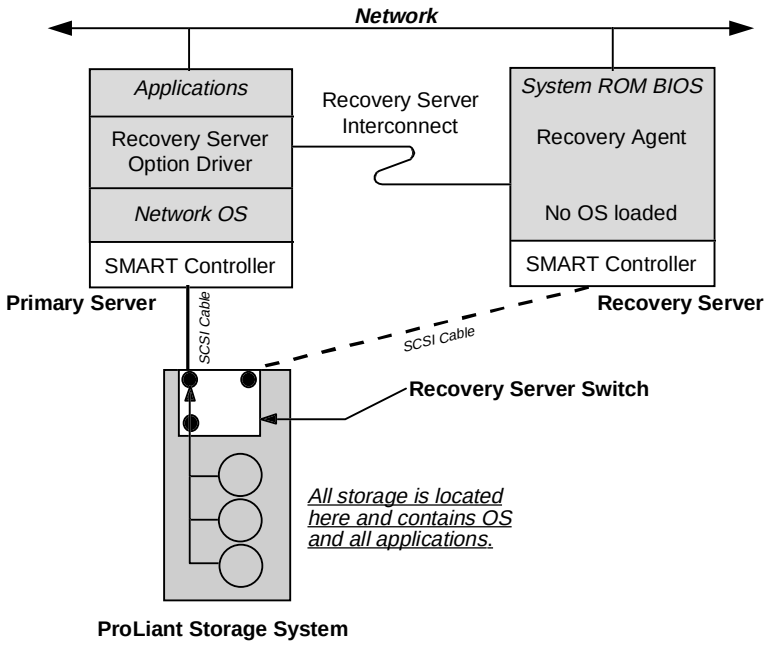

Figure 8. Normal operation of a Compaq Standby Recovery Server

The Standby Recovery Server is now supported under Microsoft Windows NT 3.5X, NetWare 3.12, and NetWare 4.X operating systems. It provides:

- Increased system availability as a result of automated, fast server recovery
- Fully automated server switchover, ideal for unattended, or *lights out*, operation
- Ability to schedule service on a failed server at the most convenient time
- Affordability, even for implementation at numerous operating locations

On-Line Recovery Server

The Compaq On-Line Recovery Server configuration (Figure 9) is a cost-effective means of increasing storage capacity and availability of business-critical applications for customers with numerous servers operating in the Microsoft Windows NT 3.5x environment. It pairs two independently operating ProLiant or ProSignia servers as hot (on-line) partners for each other.

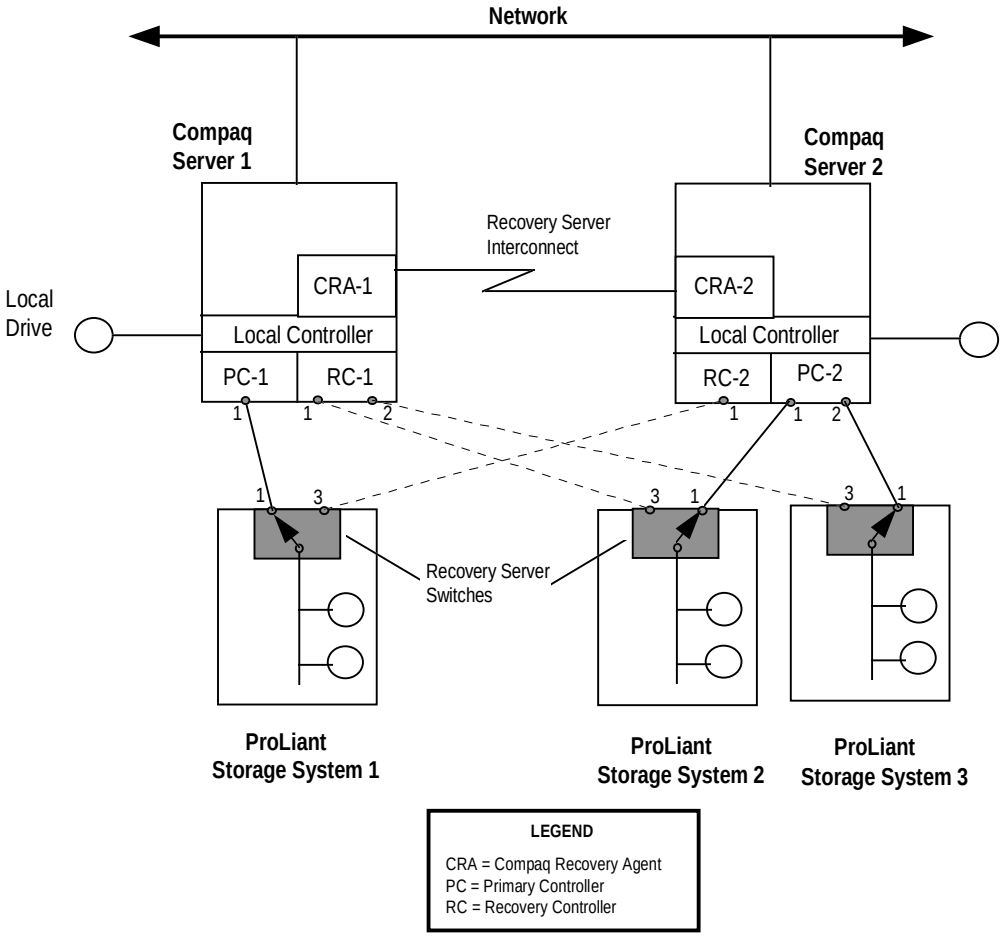

Figure 9. Normal operation of a Compaq On-line Recovery Server pair in which both servers have switchable external storage.

The On-Line Recovery Server allows any application to take advantage of this high-availability environment. The On-Line Recovery Server provides:

- Increased server availability
- Twice the capacity of the Compaq Standby Recovery Server
- Fully automated switchover with minimal interruption of client access to data
- Ability to pair independently operating servers with identical or different hardware configurations
- Ability to schedule server maintenance or replacement at a convenient time without interrupting user access to information
- Full compatibility with the comprehensive alert features of Compaq Insight Manager
- Cost-effective solution

Branch Office Server Array

In the fourth quarter of 1996, Compaq expects to introduce its first implementation of the high-reliability, high-bandwidth, low-latency ServerNet interconnect. This new high-availability solution will be a 2-server failover system operating under Microsoft Windows NT and will feature a Wolfpack API subset, SQL 6.5, and Oracle support. Like the On-Line Recovery Server, this new solution will directly link two active servers; however, with the ServerNet interconnect, there is no longer a requirement for using array controllers. The ServerNet high-speed transfer (50 megabytes/second) will provide higher availability because both paired servers will have immediate access to the other's data in the event of a failure. Moreover, the ServerNet packet acknowledgment feature will ensure data integrity.

Future Failover Enhancements

In the first half of 1997, Compaq expects to expand its new 2-server failover solution by adding more Wolfpack capabilities as they become available from Microsoft. By mid-1997 Compaq expects to introduce a multi-server failover system with full Wolfpack functionality, additional administrative tools, and fibre channel-attached external storage.

Compaq 2-server solutions, in which each server will require its own partner server, are ideal for branch offices. Compaq multi-server solutions will allow one server to provide failover capacity for multiple servers.

Compaq Parallel Processing Solutions

Future Compaq Server Arrays will offer true distributed enterprise-class clustering. These Server Arrays will deliver even higher availability, performance, and scalability. Compaq is using known techniques and industry partnerships to develop open-system parallel applications solutions rather than designing huge, proprietary solutions. This approach will speed development and result in more cost-effective solutions for customers.

In the first half of 1997, Compaq expects to introduce a 2-server parallel processing system operating under Microsoft Windows NT, with additional operating system support to be added later in the same year. In this 2-server system, the ServerNet cards in Server 1 will be directly linked by cable to ServerNet cards in Server 2.

In the second half of 1997, Compaq expects to introduce a solution that addresses scalability as well as availability. In this multi-server parallel processing system, the ServerNet cards in each server will be connected by cable to a ServerNet crossbar switch. External storage subsystems will be attached to the servers by means of a fibre channel interconnect.

PROJECTED SERVER ARRAY FUNCTIONALITY

Table 3 summarizes projected functionality for Compaq Server Array clusters during the next 18 months. Increasing functionality and investment protection will be achieved through the addition of capabilities as they become available.

TABLE 3. PROJECTED FUNCTIONALITY FOR COMPAQ SERVER ARRAY CLUSTERS

Timeframe	Failover Clusters	Parallel Processing Clusters
2H96	2-Server Failover	—
1H97	Expanded 2-Server Failover	2-Server Parallel Processing
2H97	Multi-Server Failover	Multi-Server Parallel Processing

CONCLUSION

Compaq enterprise solutions include a Server Array clustering strategy to deliver true midrange-class availability, performance, and scalability. Compaq has partnered with other industry leaders to set the industry standard for clustering x86 servers.

Compaq's broad strategy is to continue providing open systems that have more and more midrange functionality. Compaq will deliver products that link high-performance, cost-effective servers like the ProLiant 5000 and ProLiant 1500 into Server Arrays. The Compaq family of Server Array products will make midrange cluster-level availability accessible to customers whose cost-sensitivity is driven primarily by the large number of business-critical servers deployed at remote or branch locations.

Today clustering is an evolving technology, and the Compaq Server Array family of clustering solutions is evolving with it. The Server Array family began with current availability solutions found in the Recovery Server Option. It continues with a new 2-server failover solution that leverages the Tandem ServerNet technology and Microsoft Wolfpack compatibility. Then in 1997, as changes are made to the operating system and as parallel applications become more prevalent, it will grow to include a true enterprise clustering solution. The crossbar switch will be added to create a scalable cluster matrix, and a FC-AL shared storage solution will be introduced.

Compaq's Server Array family of clustering products will provide unprecedented performance and availability gains for customers with business critical data.