The Cluster Book

High Availability Segment
Industry Standard Server Group

COMPAQ

www.compaq.com/highavailability

Table of Contents

Subtab 3d – Novell NetWare Clusters

Subtab 3e – SCO UnixWare Clusters

Subtab 3f – ServerNet

Subtab 3g – Cluster Administrator

How To Use This Book

This High Availability Cluster Book has been designed to provide you with readily-available information about Compaq's Industry Standard clusters and to make your job easier. By providing product information in one concise format, this book can serve as a guide in determining which Compaq products best meet your and your customers' needs.

The cluster book is comprised of three main sections: Introduction, Executive Overview and the Compaq *ProLiant* Cluster Product Lines. Under the *ProLiant* Cluster Product Lines tab, you will find product line Microsoft PowerPoint presentations (with speaker's notes), Quick Specs and photos for each product. In addition, we have included select White Papers for your review. In the back of the book, you will find a CD containing the entire contents of the book, plus all of the PowerPoint presentations in their entirety.

Throughout the book, Web sites are referenced so that you will always have access to the latest information at hand. If you are not able to find the information you need, please send an email to **HACAL@compaq.com**.

Vince Gayman
Director of Marketing
Industry Standard High Availability & Clustering
Vince.Gayman@compaq.com

Welcome *to the ISSG Cluster Book*

Compaq and High Availability

Even with our long heritage of providing highly available systems to the industry, there has never been a more opportune time to be in the high availability business. With business requirements driving the need for continuous availability across a broad range of solutions, the Internet making business a 24 x 7 operation, and the cost of downtime continuing to escalate, customers are making high availability a requirement of most of their computing solutions. Analysts are now reporting tremendous opportunities for those that can deliver solutions to the high availability requirement with industry-standard building blocks.

The Industry Standard Server Group is making high availability and clustering a part of our offering for all our servers. In this book you will find detailed information on all our product offerings. This book's focus is on five key product lines:

- Packaged Clusters – The simple, affordable cluster
- Microsoft Clusters – High availability solutions for NT4 and Windows 2000
- Parallel Database Clusters – The scalable and available database cluster for NT and Windows 2000

- NetWare Clusters – Solutions for NetWare 4 and NetWare 5
- SCO NonStop Clusters – The solution for UnixWare 7
- *ServerNet II*
- Cluster Management

The Compaq Advantage and Strategy

The key advantage to the Compaq strategy is the leveraged model. The approach is to take the best components both internal and external and integrate them into a complete solution with the Compaq stamp of approval. The building blocks for this approach are the *ProLiant* servers, *StorageWorks* storage systems, optimized interconnect, Compaq Cluster Management tools and our key partners' clustering technology. The high availability team then tests a number of configurations, documents the configuration, the setup and operations, then certifies and releases the product to the field.

The competitive environment mandates that we continue to lead through our leveraged model. Our strategy is to focus on executing our business plan, identify changing market conditions and lead the industry in delivering Industry Standard Server high availability solutions.

Trends Influencing High Availability Strategy

Before we look at the external forces impacting our high availability strategy, we must recognize that we are driven by our Compaq Mission. Our mission, as stated below, resonates with the high availability story. To this end we strive to build the products that support our mission as a company and also meet the requirements our our target markets.

> **"Our mission is to ensure our customers' success in the 24 x 7 Internet world by delivering, with our partners, the best infrastructure, access, services and solutions…Everything to the Internet."**
>
> **—Michael Capellas**

Complementing our mission are trends influencing the demand for high availability. These trends can be grouped into two areas – technology and business.

Technology influencers push us to improve our product capabilities and partner relationships to deliver cutting-edge solutions to the high availability question. Some examples of technology influencers include:

- Storage area networks
- Dense rack mount servers
- Performance of standard servers (Intel chips, 8-Way, I/O)
- NT adoption for high availability systems

Business influencers impact the environment we and our customers compete in. Our customers and the products they use are driven to meet the new opportunities made available by these business forces. Some examples of business influencers include:

- Internet
- High availability for "all" applications
- Systems costs declining
- Globalization of business
- Consolidation of systems and storage
- Business advantages – service, extended hours
- Service level agreements

High Availability Market Dynamics

Both the technology and business trends we spoke about previously, create some dynamics in the market that enable us to position ourselves uniquely. The industry-standard market is continuing to demand and deliver availability needs that react into those areas reserved for the niche high availability market. At the same time, the niche high availability market is looking to gain excellent economies by leveraging industry standard technology to solve its availability requirements. For our customers, this means they must work in a very heterogeneous environment, change constantly, scale explosively, transact securely,

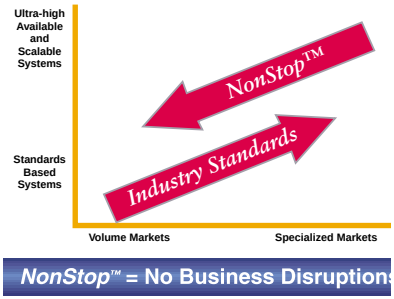

be anywhere at anytime and always be open (24 x 7), while utilizing industry-standard building blocks.

The Scope of High Availability

The spectrum of requirements for high availability is very broad. As the following pyramid chart depicts, there is a range of requirements from the low end (where NT and NetWare have been very strong and high availability was typically handled by single server configurations with limited high availability functionality) to the high end (where continuous operations were usually handled by proprietary systems with the highest levels of high availability implemented).

We are seeing two trends drive high availability across this spectrum. As previously mentioned, the requirements for high availability are moving down into the traditional NT/NetWare applications that may not have had the need for clustering, while at the same time NT is moving up into the applications area traditionally met by UNIX or proprietary systems. As this dynamic takes place, the need for high availability is a given for these customers' applications.

Early adopters are moving NT toward the high end of the pyramid into a role for database requirements. Our strategy at Compaq is to enable NT across this full spectrum of requirements, from the department all the way to the data center.

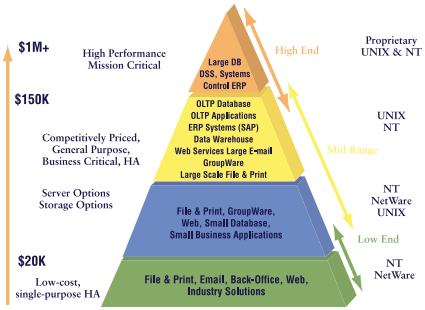

Top Drivers for Clustering

The Gartner Group has provided some insight into the key purchase criteria of the high availability market. Translated, the following chart is a prioritized list of key items that drive the need for high availability products. In summary, respondents want utility in availability and performance. Those who actively utilize high availability products are particularly interested in improving end-user satisfaction and reducing software failures. Those likely to implement more clusters in two years gave higher ratings for reduction of hardware (5.8) and software (5.9) failures.

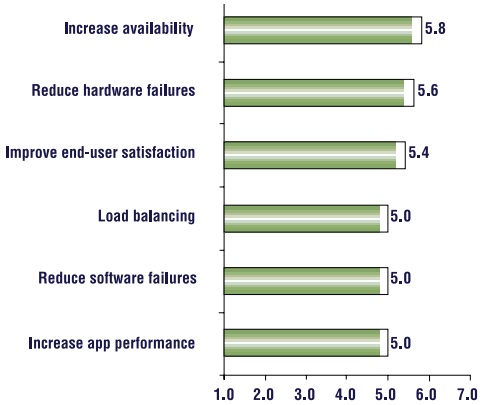

NOTE: Mean (average) responses – a scale of one to seven with one being "non-important" and seven being "very important."
Source: Gartner Group (November 1999)

The Opportunity is Real

The industry recognizes the real revenue opportunity for those that can deliver high availability solutions to the enterprise market. IDC segments the market into "availability levels" characterized by the amount of availability required of the prospective customer (AL4 being the highest level of availability required). In summary, IDC clearly articulates the nature of the availability market:

Growing the Availability Spectrum

Worldwide Revenue for Server Systems by Category, 1997-2003

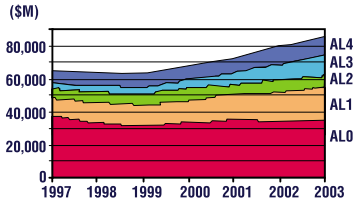

Source: IDC#21729, February 2000

"Impelled by improved technology and ever-increasing demands, servers and operating systems are more reliable than ever.... Transaction systems interact ever more intensively, and exploding Internet applications all demand servers that are extremely dependable, always online and expandable to meet growing requirements. These three criteria – reliability, availability and scalability (RAS) – epitomize what's needed from a highly available (HA) computing environment."

Source: IDC

Percentage of Worldwide Server revenue in AL2, AI3 and AL4 Categroies, 1997 – 2003

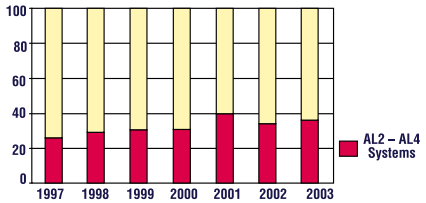

Source: IDC#21729, February 2000

The Whole Product Approach

This chart is a framework that reflects the broad scope of what it takes to move from a standard availability single server environment to a high availability clustered systems platform and then to a fault tolerant non-stop environment *(NonStop™ Himalaya™)*. This is only an illustration and by no means the comprehensive list of what it takes. The individual slices of this wheel illustrate the shared responsibilities required to make a system increasingly "available." In order to achieve the real gains you need to move out in all areas to achieve the next level of availability.

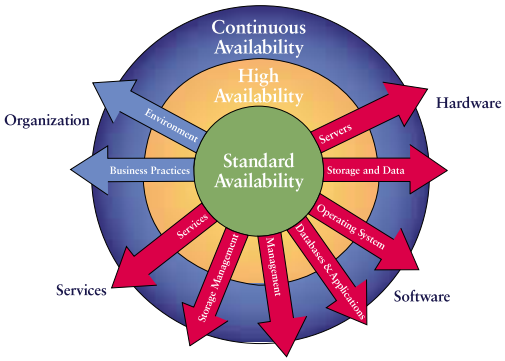

The inner ring denotes fault resilient solutions, such as redundant fans, power supplies, etc. Compaq has addressed each of the hardware and software components in its servers and has led the industry with such innovations as Hot Plug PCI.

The middle ring extends the availability through clustering technologies. Compaq ships more clustering solutions than any competitor and understand what it takes to truly achieve higher levels of availability

Finally the outer ring shows what *NonStop™ Himalaya™* customers require, an environment that never goes down. All of which is ensured by completely addressing all the necessary hardware, software, service and organizational issues.

Compaq wants to be the key IT partner that helps define and address the requirements of high availability customers. We are not just a platform provider. Compaq has leading partnerships helping key industry operating systems become more and more available. As an example, our work with Microsoft has helped harden this OS to be ready in the next two to three years for the fault-tolerant market.

Clustering for the Enterprise

Scalability
Availability
Manageability
Innovation

Compaq *ProLiant* servers are enabled for many high availability environments. In conjunction with strategic partners, Compaq offers a very broad range of functionality and high availability. Compaq delivers choice and power for enterprise computing, offering multiple servers from which to build their solutions. Compaq also partners with the leading software vendors providing broad OS and application support, giving customers the best solution to suit their business needs and IT environment.

What does power and choice mean to the customer?

- Flexibility to meet a broad range of requirements
- Doesn't lock customers into single answers
- Allows customers to choose the best solution to meet their needs
- An integrated building block approach to flexible industry standard clustering
- Strength and commitment: Compaq offers a proven combination of advanced industry-standard technologies delivering the highest levels of performance, capacity and availability
- Inter-operates within heterogeneous environments

The *ProLiant* Cluster Family

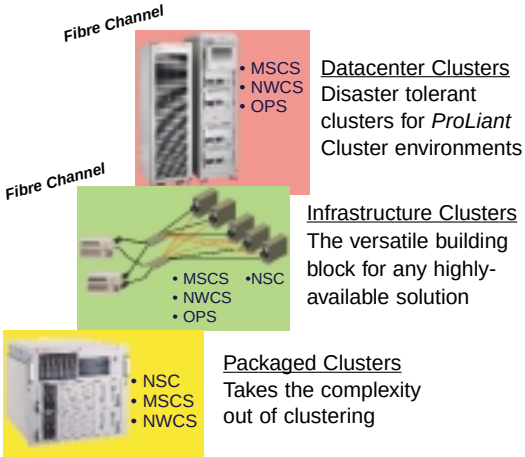

Legend

MSCS=	**Microsoft Cluster Server**
NWCS=	**NetWare Cluster Service**
OPS=	**Oracle Parallel Server**
NSC=	***NonStop*™ Cluster**

As mentioned previously, the portfolio of products covers everything from the low-end packaged clusters to the data center. Compaq brings together the servers, storage, software, service and manageability that deliver on the needs of the high availability customer.

Our value proposition is simple and resonates with our mission. We make sure our customers are successful in the 24 x 7 Internet world by delivering the best infrastructure, access, services and solutions required of the high availability market.

Services Completing the Clustering Solution

Compaq delivers a broad range of services and support capabilities that allow customers to leverage years of expertise in the areas of business critical solutions. The business critical services team for Compaq is part of a much larger organization committed to providing the highest levels of availability to customers from both direct support through Compaq and the building of expertise and services delivery through channel partners.

The list of services listed is not comprehensive. This an overview of the capabilities provided by Compaq services teams.

* **Business Critical Services** – In the Business Critical Services list below, the services are proactive capabilities to help customers plan, install and manage their systems to a high level of availability. This is very much a partnered approach between Compaq and customers.

Business Critical Services*

- Availability review
- Availability partnership
- Uptime guarantee
- Installation and startup services for Windows NT clusters

Business Critical Support Services Priority Service Plans (PSP) Range of Services

- On-site response times
- Call back response times
- Named engineers
- Technical advisors
- Software account manager
- Proactive notification on updates

Business Continuity Services

- Recover-all
- Hot site service

Industry Standard Clustering Solutions from Compaq

Compaq continues to extend the capabilities of our industry-standard clustering offerings. Through partnerships with Intel, Microsoft, Novell, Oracle and SCO, Compaq *NonStop™ Proliant* clusters offer customers industry standard technology that solves their business critical needs.

When compared with the competition some of our key differentiators include:

Program Scope

Compaq offers a full range of solutions – much broader than key competitors. Compaq covers all key partner operating systems. Compaq certifies almost its entire *ProLiant* server range and the full range of fibre channel storage systems. Ask customers to compare the certifications on the Microsoft Hardware Compatibility List for a good reflection of this strategy.

Packaged Cluster

Compaq is the only major vendor to offer a simple and affordable packaged cluster, integrated and packaged with everything necessary to start clustering.

Optimized building blocks

The server and storage strategy is fully in line with the customer need for high availability, using off-the-shelf, industry-standard technology to build fully redundant cluster configurations out of the best price performance components.

Cluster Management

A very key differentiator is the management tools that are fully integrated with *Compaq Insight Manager XE* to provide customers with a common interface and tool set to manage their clusters and servers.

More Proof of the Scalability of Oracle Parallel Server

Benchmark

On March 30, 2000, Compaq and Oracle once again demonstrated the enterprise class scalability and performance of the Microsoft NT operating environment. The combination of a Compaq Parallel Database Cluster and Oracle Parallel Server (OPS) running the standard benchmark for Oracle Applications resulted in previously unattainable results on NT. These results more proof of the scalability of Oracle Parallel Server in the NT environment and the Compaq commitment to advancing this technology.

This record benchmark was achieved using a two-node Compaq *ProLiant* 6400 Parallel Database Cluster with 10 *ProLiant* 6400 Application Servers, Compaq Fibre Channel storage technology and Compaq ServerNet as the high-speed cluster interconnect. Oracle Parallel Server Release 8.1.6 was used to merge the resources of the two database servers and apply them in parallel to the client requests generated by the benchmark tests. This configuration successfully recorded 3,284 concurrent users, a major breakthrough for an NT-based solution. The prior performance record in the NT environment, set with a single *ProLiant* 6400 server, had been 1,792 concurrent users. The new record converts to a scalability factor of 1.81, going from a single node to a two-node OPS cluster.

These results mean there is a much larger customer segment that will now have the additional option of an NT-Operating System when deploying Oracle Applications solutions because the performance boundaries have been pushed to new levels. At the same time, high availability is integral to the Compaq Parallel Database Cluster for OPS with automatic node recovery and client failover in the event of a server failure. Utilizing Compaq industry standard solutions will allow Compaq customers to achieve lower total cost of ownership at higher user counts.

This record benchmark was the result of the ongoing joint engineering efforts of Oracle and Compaq. Compaq has been the development platform for OPS on NT since its initial release in Oracle version 7.3. Since that time, Compaq has posted five record TPC-C benchmark results utilizing the parallel processing power of OPS, including the world record TPC-C benchmark of 101,657.17 tpm/C. As a result of these efforts, Compaq has established clear leadership in this growing market segment and has accumulated the internal expertise necessary to deliver real-world solutions based on Oracle Parallel Server.

For more detailed product information about Compaq Parallel Database Clusters see:
www.compaq.com/highavailability

For more details on the benchmark see:
www.oracle.com/apps_benchmark/html/results.html

For more information on Oracle Applications see:
www.oracle.com/applications/

First Windows NT Result to Break The 100,000 tpmC Barrier

Benchmark

February 2000
11XD-0200A-WWEN

Prepared by:
Industry Standard Server Division
Compaq Computer Corporation

Compaq Parallel Database Cluster with *ProLiant* 8500 Servers Delivers Best Price:Performance Utilizing Industry Standard Technologies and Best Performance with Industry Standard Servers

- Leading Cluster Price:Performance and Performance
- Heritage of leadership in industry standards
- For more information: **www.compaq.com/highavailability**

Abstract: On February 11, 2000, Compaq announced that the new Parallel Database Cluster (PDC/O2000) with *ProLiant* 8500 servers achieved 101,657.167 transactions per minute (tpmC) at just $35.68/tpmC. The PDC/O2000 is a six-node cluster of *ProLiant* 8500 servers, each with eight 550 MHz Pentium III Xeon processors running Oracle8i Enterprise Edition v8.1.6.0 with Oracle Parallel Server, Compaq Fibre Channel shared storage and Windows NT Server 4.0 Enterprise Edition. The results represent:

- The FIRST Windows NT result to break the 100,000 tpmC barrier
- Best Windows NT TPC-C Performance
- Best price:performance of any TPC-C result on a clustered configuration
- Largest-ever Oracle TPC-C performance on Microsoft Windows NT
- Best price:performance for any TPC-C cluster, demonstrating the TCO advantages of open system

First Windows NT Result to Break The 100,000 tpmC Barrier

Benchmark

Compaq Parallel Database Cluster with ProLiant 8500 Servers Delivers Best Price:Performance Utilizing Industry Standard Technologies and Best Performance with Industry Standard Servers

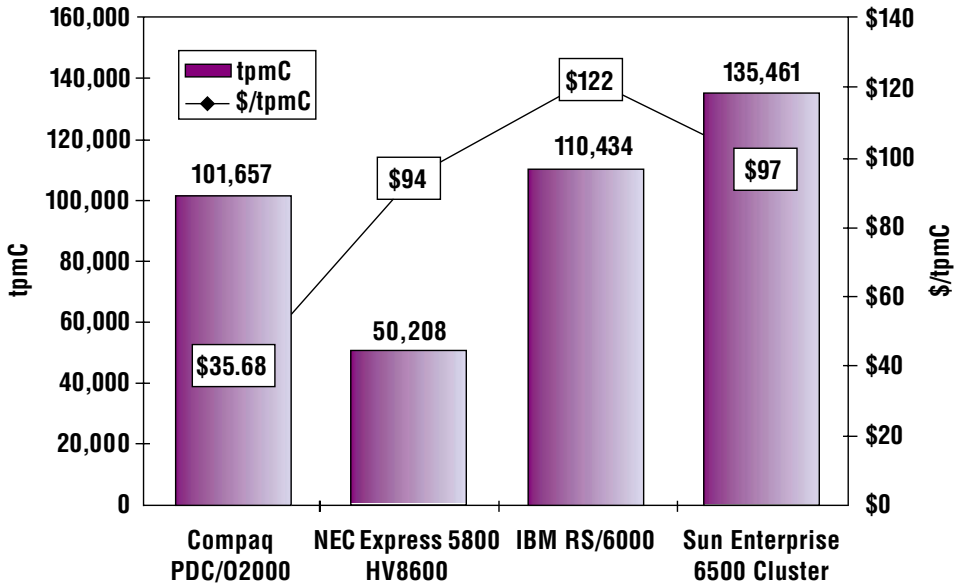

Figure 1: Cluster TPC-C Results as of December 1999

Company	# of Nodes	# of CPUs/Node	System	tpmC	$/tpmC	Operating System and Database	Total System Cost
Compaq	6	8	PDC/O2000 *ProLiant* 8500	101,657.167	$35.68	Microsoft Windows NT 4.0 and Oracle8i v8.1.6.0	$3,627,928.00
IBM	5	12	RS 6000 Enterprise Server S70	110,434.10	$122.44	IBM AIX 4.3.2 and Oracle8 v8.0.5	$13,521,883.00
Sun	4	24	Enterprise 6500 Cluster	135,461.40	$97.10	Sun Solaris 2.6 and Oracle8i Enterprise Edition v8.1.6.0	$13,153,324.00
NEC	4	8	Express5800 HV8600	50,208.43	$94.05	Microsoft Windows NT 4.0 and Oracle8i Enterprise Edition v8.1.6.	$4,723,164.00

Based on Published TPC-C Results as of February 14, 2000

First Windows NT Result to Break The 100,000 tpmC Barrier

Benchmark

This benchmark was performed on a six-node Parallel Database Cluster with the ServerNet interconnect option and ProLiant 8500 servers, each with eight 550MHz/2MB Pentium III Xeon processors and 4GB of memory sharing 9786.04GB of total disk storage.

Leading Cluster Price:
Performance and Performance

The Compaq PDC/O2000-*ProLiant* 8500 achieved this exceptional price:performance simulating a load of 85,680 concurrent users. The PDC/O2000-*ProLiant* 8500 price:performance is 71 percent lower than IBM, 63 percent lower than Sun and 62 percent lower than NEC cluster models.

In addition, the PDC/O2000-*ProLiant* 8500 performed OVER 100 PERCENT BETTER than any industry-standard clustered servers using Microsoft Windows NT 4.0.

Heritage of Leadership in Industry Standards

Compaq was the first company to demonstrate the power of industry standards with its six-node Parallel Database Cluster performance in April 1998. Working with its partners, Oracle and Microsoft, Compaq continues to lead performance in industry-standard, highly available and scalable cluster solutions.

This latest benchmark represents the largest Oracle TPC-C result ever achieved under Microsoft Windows NT and illustrates the synergy and commitment of Compaq's partnerships with Oracle and Microsoft. The Compaq Parallel Database family of products has the largest customer installed base of Oracle Parallel Server clusters in the industry as well as the largest and most comprehensive product offering. For more details on the PDC/O2000, visit the **www.compaq.com/highavailability** site on the Web.

These record-breaking results were achieved using Compaq's Fibre-Channel host adapter and storage array subsystem with 8-way *ProLiant* 8500 advanced technology and ServerNet Virtual Interface Architecture for highly reliable and scalable inter-cluster communications. The Compaq Fibre Channel Storage System is the first highly scalable, large bandwidth storage system designed specifically for industry standard based servers. With an external sixth-generation Smart Array RAID controller and a 100 MB/s Fibre Channel Arbitrated Loop (FC-AL) host interface, the Fibre Channel Storage System delivers exceptional scalability of capacity and performance.

For More Information

Full disclosure reports describing these benchmark results have been filed with the Transaction Processing Performance Council (TPC) and are available upon request. The full disclosure report describes the benchmark hardware and software configuration in detail, provides costs and lists the code actually used to perform the test. Similar reports from other vendors are the source of the price:performance comparisons provided above. Summaries of all tests are published each month by the TPC. Summaries are also posted on the Internet on the TPC's World Wide Web Server at **www.tpc.org**. With these benchmarks, customers can objectively compare the performance of different vendors' servers in specific areas such as database throughput in transactions per minute (tpmC) and cost per transactions per minute ($/tpmC).

This presentation is about our entry-level high availability cluster. We call it a “packaged cluster” because it is self contained with two server nodes and shared storage. The base unit can be ordered as a single part number for ease of ordering and set-up. The CL380 is a follow-on to the CL1850.

Compaq ProLiant CL380

Product Line/Speaker Notes

The Packaged Cluster is a entry level clusters offered in our ProLliant Cluster Product Family. The existing product in the Packaged Cluster is the CL1850. In June, 2000 we announced the CL380 which has the latest performance based on the DL380 server. The 10U form factor remains the same.

Compaq *ProLiant* CL380

Compact
- Small two-node cluster, including storage, packaged as a single tower or rack mounted system (10U)

Easy
- Plug and play clusters
- Standard server SKU & options
- Managed with Compaq Insight Manager
- Packaged offerings for service and support

Industry Standard
- Microsoft NT 4.0 Enterprise Edition, Windows 2000 Advanced Server, Novell NetWare and SCO
- Both servers active and running applications
- Servers backing each other up via cluster software

COMPAQ

www.compaq.com/highavailability

The CL380 is compact, easy to order and set up and is based on industry-standard products. The CL380 is only 10U (17.5″) high and contains shared storage and two independent server nodes. The unit can be ordered as rack or tower mount. The CL380 is an easy way to achieve the highest availability at entry-level pricing. Through CTO in North America, the unit can be fully configured with redundant components, drives and even the operating system. Users unpack, plug in, load application, set RAID and go. The CL380 will support SCO Unixware on 8/15/00.

Compaq ProLiant CL380

Product Line/Speaker Notes

Same look and feel as the CL1850 with the latest features of the new DL380 server. Changes in the CL380 include 550 MHz to 800 MHz processors, memory increases from 2GB to 8GB per cluster, 32 bit to 64 bit PCI card expansion and ROC (RAID on a chip), a standard feature embedded on the server system board to free up the PCI slot. SCO will be the latest OS support on 8/15/00.

Still the only packaged cluster available today! Our competition will ship a cluster in parts and leave the customer to build and configure the cluster. We build it for the customer. The cluster can be ordered direct from the Web through Compaq.com, a reseller or a major account team.

Compaq ProLiant CL380

Product Line/Speaker Notes

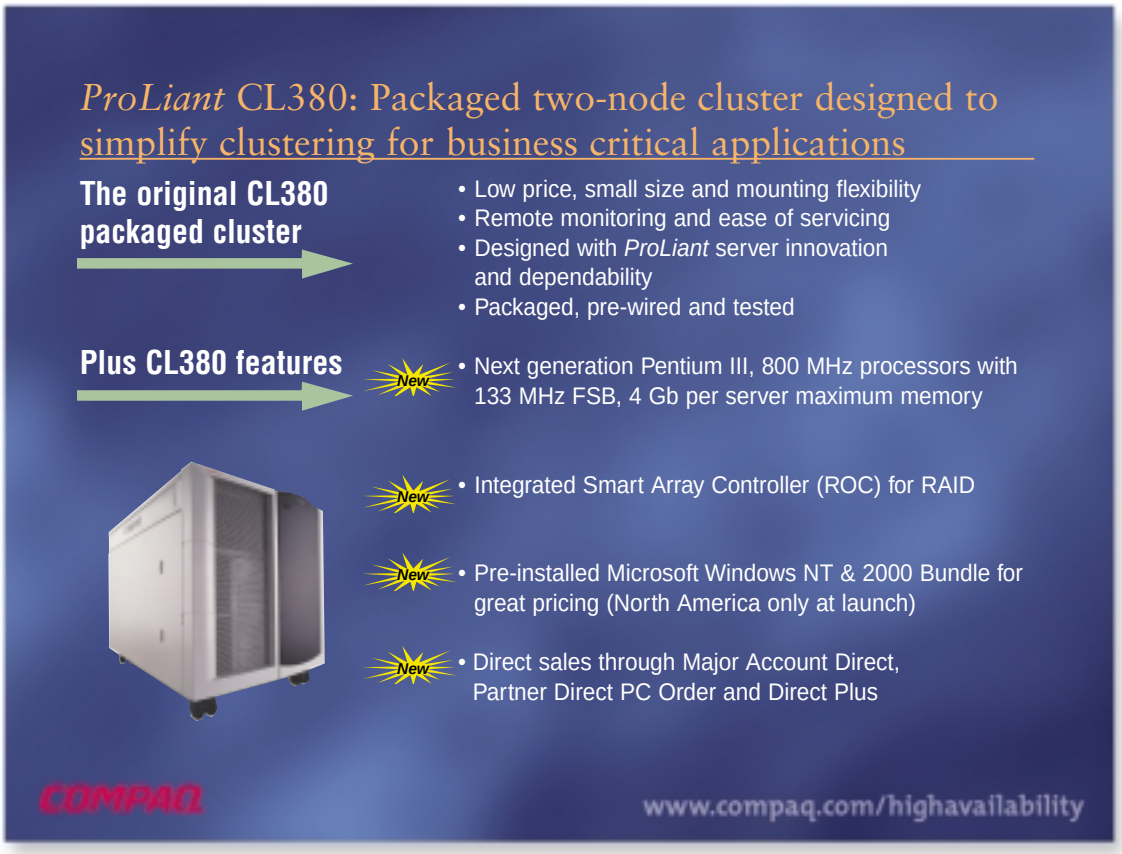

Our market segment is SMB and major accounts in the retail and financial markets. The low cost meets the needs of entry-level users with high-availability needs. Companies migrating to Internet and eBusiness demand high availability. Clustering is the only fully-redundant solution and we make it available to the entry-level market.

The CL380 is designed for small companies to run the enterprise infrastructures such as database, file/print and Web server. It also is great for a branch office/plant to run critical business applications or file/print, database or Web server. Many large companies with business critical applications for departmental usage choose the CL380.

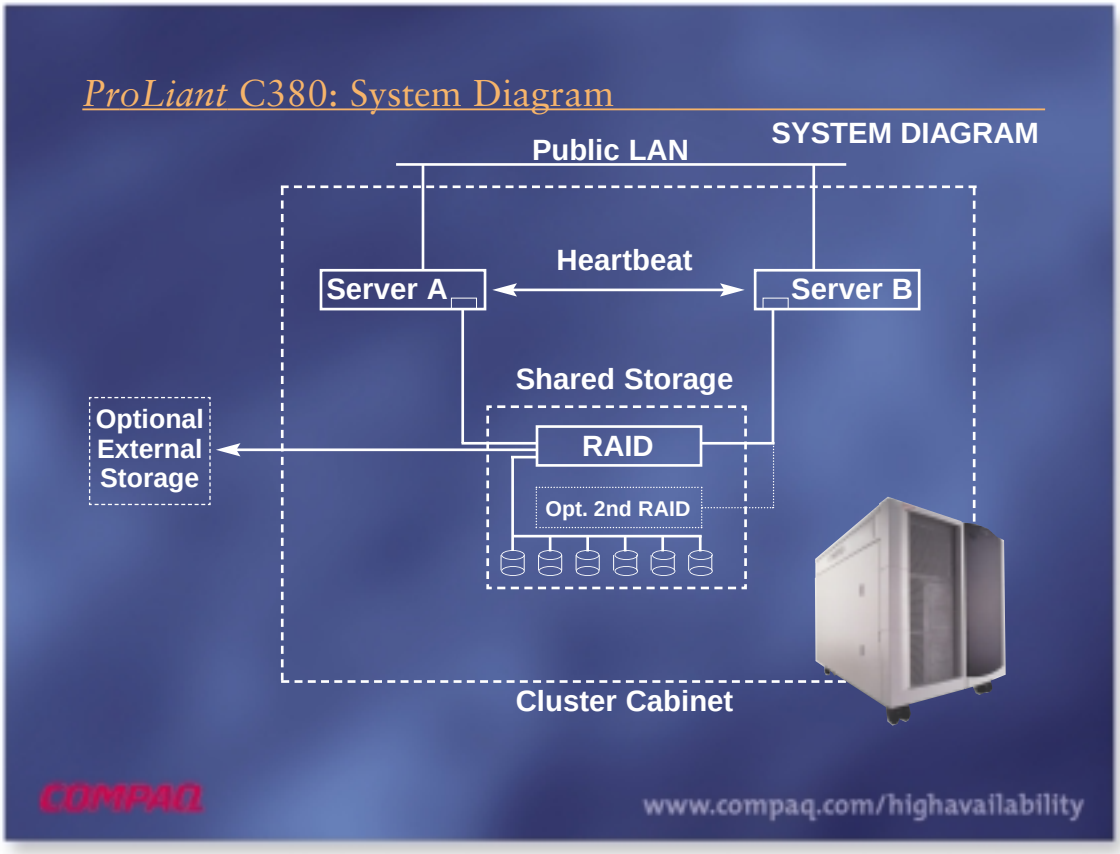

The system diagram shows the two independent server nodes connected together with a heartbeat for simple failover. The shared storage is connected to both server nodes and can be managed and partitioned by the RAID controller. The system is contained in a cabinet that is rack mountable. On the left is an optional storage cabinet for extended storage needs.

Compaq ProLiant CL380

Product Line/Speaker Notes

Here is the details of CL380 specification. Specific sections include the configuration features, cluster management, performance and reliability.

Compaq ProLiant CL380 Performance

for Microsoft Exchange Server 5.5 Enterprise Edition with Microsoft Cluster Server (NT 4.0)

White Paper

June 2000
12RB-0600A-WWEN

Prepared by:
High Availability Systems Engineering
Compaq Computer Corporation

Abstract: This document covers Microsoft SQL Server 7.0 Enterprise Edition performance optimization on a Compaq *ProLiant* CL380 running with Microsoft Cluster Server on the Windows NT Server 4.0, Enterprise Edition operating system. The test methodology assumed that the executable files, log files, data files and temporary database files were located on separate logical drives. An industry-standard benchmark was used to measure the impact that the RAID levels implemented combined with the location of the drives had on the performance results.

Introduction

The purpose of this document is to provide the optimal configuration for Microsoft SQL Server 7.0 running on a Compaq *ProLiant* CL380. Additional information regarding the *ProLiant* CL380 in Microsoft Windows 2000 Advanced Server, Novell NetWare and SCO UnixWare clusters, as well as performance White Papers for other database and messaging applications can be found at **www.compaq.com/highavailability**.

Expected Audience

This document is intended to help in the installation, configuration and administration of Microsoft Cluster Server with Microsoft SQL Server and assumes that the reader has working knowledge of the following:

- Installing and configuring Compaq *ProLiant* CL380
- Installing and configuring Microsoft Windows NT Server 4.0 Enterprise Edition
- Microsoft Windows NT Server interface and setting up and administering user accounts in a domain
- Installing and configuring Microsoft SQL Server 7.0 Enterprise Edition
- Installing, configuring and administering Microsoft Cluster Server
- Installing and running Benchmark Factory

Compaq *ProLiant* CL380

The Easiest, Most-Affordable Clustering Solution

The *ProLiant* CL380 is a uniquely packaged two-node cluster that is designed to simplify clustering for business critical applications. *ProLiant* CL380 consists of two Compaq server "nodes" and shared storage pre-packaged in a cost-effective, space efficient cabinet, giving customers the easiest most affordable clustering solution for NT and NetWare.

The *ProLiant* CL380 is ideal for:

- Remote systems requiring unattended high availability
- Branch Office industry applications – retail, health care, financial
- Dedicated function servers needing high availability Microsoft Exchange, Novell GroupWise, Lotus Domino, SQL Server and file/print
- Departmental applications that need high availability
- For space-constrained data centers and business offices

Benefits at a Glance

- Complete cluster housed in a single 10U cabinet, pre-wired and pre-tested, makes it easy to install and easy to order through a single part number.
- Convenient size, flexible configuration and integrated Compaq management tools mean easy deployment and rapid install.
- Compaq's newly-released Cluster Monitor (part of Compaq Insight Manager XE) and *Intelligent Cluster Administrator* enable Web-based remote monitoring and administration of the cluster across the Internet or your Intranet making *ProLiant* CL380 easy to mange and minimizing need for local technical support.
- Commonality with *ProLiant* servers and other industry-standard technologies mean low purchase price.
- On-line serviceability and easy access to hot swappable components makes *ProLiant* CL380 easy to service and reduces down-time, resulting in low cost of maintenance.
- Substantial in-system growth provides low cost incremental expansion of performance and storage capacity without the cost of system upgrade or change
- Ready for the Enterprise – Service and Support team up with the channel to deliver and support business critical requirements.

Key Features

- Innovative low-cost, space-saving design
- Small footprint packaging of two clustered *ProLiant* server nodes plus shared storage
- Unique chassis design offers deployment flexibility for tower or rack mounting
- Integrated switch to share keyboard, mouse and monitor between server nodes.

Serviceability

- On-line serviceability of all major components
- Hot Plug server and shared storage disks
- Hot Plug power supplies for shared storage
- Easy component access for ease of maintenance

Performance

- Two Intel Pentium III 800 MHz processors with 256KB of secondary (L2) cache per server
- Up to four gigabytes of ECC-protected, 133-MHz SDRAM system memory in each server node for faster clock speeds and increased system performance
- High performance 10,000 rpm SCSI disks
- Four PCI expansion slots
- Up to 252 GB of high performance SCSI storage

Reliability

- Cluster software provides the capability for application or server failover to maintain business operations in the event of hardware or software failures

- High performance RAID array controller for fault tolerant protection of shared data to maximize uptime
- Redundant power supplies for shared storage to provide continuous power after an internal power supply failure
- ECC memory can prevent memory errors from causing server failure
- 128MB ECC-protected SDRAM memory standard, expandable to 4GB
- Processor recovery for automatic reboot, fail-over processor redundancy
- Redundant NIC support with failover of heartbeat to maintain cluster
- Year 2000 support for smooth date transition
- Extensive variety of internal tape backup support

Manageability

- Web-based cluster management
- Compaq SmartStart installation utility
- Cluster verification utility
- Compaq Remote Insight Board option
- Color-coded cable receptacles

Factors such as backup/restore, topology and other issues should be considered when planning a deployment.

Real World Benchmarking

Typical Microsoft Exchange Benchmarking

When Exchange Benchmarks are published, the benchmarks have typically been conducted using the Microsoft Messaging Application Program Interface (MAPI) Messaging Benchmark. The MAPI Messaging Benchmark or MMB, was developed in an effort to move the focus away from users per server, so customers would not automatically base their deployments upon it. MMB is a benchmark workload based on the LoadSim canonical workload profile (a set of messaging actions) for what Microsoft had determined to be a "typical corporate e-mail user" (known as the Medium LoadSim User). The metric for MMB is a number (MMB) which represents a user transaction load attained during the benchmark run. In addition, a 95th percentile response time score, measured in milliseconds, is also reported. Microsoft audits and approves results from hardware vendors that desire to publish MMB scores for their server platforms. MMB and LoadSim have attained widespread acceptance in the Exchange Server space as the *de facto* standard for comparing Exchange Server performance.

What the Benchmarks Don't Address

Results from the MMB should be interpreted as only a benchmark for comparing messaging throughput of various servers and configurations and should not be confused with deployment recommendations. Factors such as backup/restore, topology and other issues should be considered when planning a deployment. It is important to understand that benchmarks are designed to give deployment implementation planners baseline references for understanding the capabilities of hardware platforms offered by vendors such as Compaq. When interpreting these benchmarks, keep several things in mind. First, many of the benchmarks are performed on configurations that are not customer deployable. For example, most vendors have published benchmarks based on disk subsystems configured as RAID 0 disk arrays. While RAID 0 does provide the highest levels of disk subsystem performance, it fails to provide any protection against data loss. Benchmarking also excludes issues, such as disaster recovery and information store size, and factors, such as directory replication and inter-server traffic. These issues are discussed on the following pages.

Backup/Restore

Since messaging and collaboration applications are quickly becoming mission critical to most organizations, the need for high availability and expedient disaster recovery is of the utmost importance. This situation has created a dilemma for many implementers. This is because today's hardware platforms combined with industry-leading software like Microsoft Exchange Server provide enterprise-class scalability supporting thousands of users per server from a performance viewpoint. However, due to high availability and disaster recovery concerns, many customers are choosing to deploy lower numbers of users per because as server user loads are increased, the information stores will grow in size. For example, with 1,000 users on a server with 30MB of mailbox storage allocated to each user, the information store would exceed 30GB (using a simplified calculation – many other factors are considered in calculating the size of the information store).

In many instances, the storage requirements for the information store will outweigh the I/O requirements (the focus of benchmarks). In addition, working space for information store maintenance, upgrades and other administrative activities should also be considered. Often, additional storage requirements for information store maintenance and other issues can be as high as 100 percent of the actual user mailbox cumulative total. If the MMB results were used to deploy a system, the server information stores could grow beyond disaster recovery capabilities.

The deployment implications present some serious challenges. First, can the backup mechanisms in place meet the requirements within the current backup or, more importantly, restore "window" (the time period available to perform backup activities)? Will the backup and, more importantly, the restore times provided meet IT service-level agreements? Second, is the increased vulnerability to system outage (20,000 users not able to work, versus 2,000 users) an acceptable risk? Finally, how much time will the server be unavailable while a restore is in progress?

User Workload

Another factor to consider is the actual workload used in the MMB. Although the Medium LoadSim User was believed to simulate a typical corporate e-mail user at the time it was defined, over the years the amount of e-mail received and the size and type of these messages has changed. Users are seeing more attachments, such as Word documents and PowerPoint presentations, and these attachments are growing larger in size. The Medium LoadSim User has an average message size of about 12KB, while many corporations today are seeing an average message size of over 50KB. Today's users are also often likely to be using the Outlook client, instead of the Exchange client used in the MMB. In discreet testing, the Outlook client was found to produce about a 10 percent decrease on the MMB result, as compared to using the Exchange client. All of these factors have an impact on the performance of the server, and thus the number of users that can be comfortably supported.

Real World Workload and Benchmarking

The purposes of Compaq's Real World Benchmark are to address the issues described above and to enable Compaq to deliver tested and deployable system configurations to the customer. The results of this benchmarking should provide better guidelines for deployment of Exchange mail-only servers.

- Only customer deployable systems are tested
- Taken into account are issues such as backup and restore of the database and protection of data against loss

- The workload is changed to more accurately reflect today's e-mail users
- The size of the database is increased to approximately 27MB per user
- No tuning is done to the Microsoft Exchange Server installation beyond the standard Exchange Performance Optimizer (PerfWiz) recommendations
- Each system allows amble room for growth and for peak usage periods
- The test verification guidelines are applied for the entire eight hours of the test, not for only the four-hour steady state, thus the increased burden of logins and logoffs are considered in sizing the system

Customer Deployable Configurations

This benchmark is conducted with only customer deployable systems that take into account protection against data loss, disaster recovery, information store size, multi-server configurations, etc.

All testing is conducted using RAID 5 for the Information Store, and RAID 1 for the OS, Page File, Exchange Log Files and Exchange DS/MTA Files, to ensure protection against data loss. Each system is configured with an average mailbox size of 27MB per user, and the backup capacity to do a backup of the Information Stores within a four-hour window.

Workload Changes

The Real World workload is a modification of the Medium LoadSim User workload. The Real World Workload has been designed to reflect the peak load on a Microsoft Exchange server. This peak load is then run for the entire length of the test. This workload uses the Outlook Client, increases the amount of mail sent and received, increases the size and frequency of attachments, increases the number of distribution lists, adds PowerPoint attachments, and allows for the transient user (a user who logs on and off a number of times during the course of a day). The result of these changes is that the average message size is 74.9KB, compared to 14.6KB for the Medium LoadSim User. The average number of messages sent during an eight-hour day is increased from 14.18 to 44.12. And the average messages received during the eight-hour day is increased from 66.3 to 161.90. See Appendix B for a more complete comparison of the Real World Workload to the Medium User Workload.

Tuning of the Exchange Server

When most MMB results are published, they include the tuning that was necessary on the Exchange Server to achieve those results. While a system manager may wish to tweak some of these parameters to increase the performance of the server, the amount of tuning required in the MMB testing would not be recommended for deployments. The Real World Benchmark does not tune the Exchange Server, but follows the tuning recommendations of the Exchange Optimizer Wizard (PerfWiz), which ships with the Exchange Server.

ProLiant CL380 Configuration

Table 1 details the configuration of the LoadSim clients used to simulate multiple Microsoft Exchange users generating the MMB workload for the 6650 MMB measurement.

ProLiant CL380 (Values are Per Server)	*Configuration*
Client CPU types and speeds	1P 800MHz Pentium III processor (per cluster node)
Memory	512 MB
Network Topology (100Base TX, Ethernet)	2 x 100 Base-TX
Network Controllers	Compaq NC3123/Compaq NC3163 embedded
Operating System name and version	Microsoft Windows NT Server 4.0 with SP6.0a
Microsoft Exchange Server	5.5 (Build 2187), with SP 3
Number of Shared Storage Drives	14 (Six drives in the internal CL380 Shared Storage Enclosure and eight additional drives in the StorageWorks 4214 Enclosure)

Table 1: ProLiant CL380 Configuration for Microsoft Exchange Server 5.5 EE

LoadSim Client

Table 2 details the configuration of the LoadSim clients used to simulate multiple Microsoft Exchange users generating the MMB workload for the 6650 MMB measurement.

ProLiant CL380 (Values are Per Server)	*Configuration*
Model	Compaq *ProLiant* 850R
Client CPU types and speeds	2P 200MHz Pentium Pro Processor
Number of clients	Five clients with 256MB RAM (100 users each)
Network Topology (100Base TX, Ethernet)	100 Base-TX
Network Controllers	Netflex3
Client network software name and version	Microsoft Windows NT Server 4.0 with SP6.0a
LoadSim version	5.5 (Build 2187)

Table 2: Configuration of LoadSim Clients for 500 Real World Users

Performance Data

Performance data for the Real World measurements are detailed in Table 3.

Summary	RAID 0 + 1	RAID 5
Supported Benchmark Load	500 Real World Users	500 Real World Users
Benchmark Profile	Real World User Load	Real World User Load
Protocol	Outlook MAPI	Outlook MAPI
Length of Test	Eight hours	Eight hours
Unless otherwise noted, values listed are averages over entire steady-state period		
Transaction Load (hourly)		
Messages Submitted	19,463	18,726
Message Recipients Delivered	79,255	77,146
Messages Sent	2,999	2,880
Transaction Load (per second)		
Message Opens/Sec	6.7	6.7
Folder Opens/Sec	1.6	1.6
Transaction Queues		
IS Send Queue Average Length	0.098	0.160
MTA Work Queue Average Length	0.056	0.062
Processor Utilization		
System Processor Utilization (percent)	9.194	8.699
System Processor Queue Length	0.203	0.145
System Context Switches/Sec	481.1	439.2
Memory Utilization		
Available Bytes	84.9 MB	86.2
Pages/Sec	0.433	0.456

Table 3: 500 Real World Users (measured during test run at steady state)

Note: *Performance Results were measured using Microsoft Windows NT Performance Monitor. Measurements were obtained by measuring averages for the period of steady-state activity (i.e., after all of the 500 users were successfully logged on). Tests measure the messaging throughput of a single-server, single-site topology.*

For deployment specific information contact a Microsoft or Compaq representative.
More information can be found at:
www.microsoft.com/exchange/

User Response Times

Table 4 details response times for various user actions during benchmark testing.

95th Percentile Response Time (in Milliseconds)

Client Actions	RAID 0+1	RAID 5
Read	141	141
Send	812	1422
Delete	125	62
Move	328	594
Submit	250	78

Table 4: User Response Times (Latencies) from Load Simulator

Descriptive Terms

Messages Submitted
Submit calls made by clients. This equates to total message sends by users.

Messages Sent
Messages that the Information Store sends to the MTA (not messages sent by clients). Normally all messages submitted by the clients are sent to the MTA, except in the case where all recipients are local mailboxes. In that case, since all the deliveries can be performed locally, no message is sent to the MTA.

Message Recipients Delivered
Separate mailboxes that messages have been delivered to. Think of this as the number of Reads that are "caused" by sending a message (one per recipient).

Message Opens/Sec
Messages accessed for reading per second.

Folder Opens/Sec
Folders opened for browsing per second.

IS Send Queue Average Length
Send Queue Size is the number of messages in the private information store's send queue.

MTA Work Queue Average Length
Work Queue Length is the number of outstanding messages in the Work Queue, which indicates the number of messages not yet processed to completion by the MTA.

Analysis of the Results

The Real World Benchmark is required to meet the same verification guidelines as the MAPI Messaging Benchmark (MMB). Two tests were run with the *ProLiant* CL380 in a max-storage configuration with the data disk being either RAID 0+1 or RAID 5 while the log files were always on a RAID 1 drive. These configurations were selected because they yield the highest amount of performance with a layer of fault tolerance. Table 5 lists the results of the Real World with a 500 user workload to simulate a branch office or remote location.

	Real World RAID 0+1	Real World RAID 5
Number of Simulated Users	500	500
Response Time (average in milliseconds in 95th percentile)	271	310
Messages Submitted (eight-hour steady state period)	19,463	18,726
Message Recipients Delivered (eight-hour steady state period)	79,255	77,146
Messages Sent (eight-hour steady state period)	2,999	2,880

Table 5: Performance Highlights (Compaq ProLiant CL380, (1) Pentium III 800-MHz)

Table 5 shows that the RAID 0+1 data drive configuration yielded overall higher performance statistics for the 500 user test. The RAID 0+1 configuration outperformed the RAID 5 configuration by a rate of 39 ms. The RAID 0+1 configuration also demonstrated higher performance numbers for message handling than the RAID 5 configuration. The RAID 0+1 configuration number of messages submitted was 737 more than the RAID 5 configuration. The RAID 0+1 total message recipients delivered was 2,109 more than the RAID 5 configuration. Finally, the RAID 0+1 configuration total messages sent was 119 more than the RAID 5 configuration. The charts that follow represent the differences between RAID 0+1 and RAID 5 performance pertaining to message handling.

Table 5 shows the impact that the RAID level used for the *ProLiant* CL380 Exchange 5.5 data drive has on the resulting response time. The chart shows that the RAID 0+1 data drive configuration has an average response time of 271 ms, while the RAID 5 data drive configuration has an average response time of 310 ms. The RAID 0+1 configuration outperformed the RAID 5 configuration by a rate of 39 ms.

Table 5 also shows the impact that the RAID level used for the *ProLiant* CL380 Exchange 5.5 data drive has on the resulting number of messages submitted. The chart shows that the RAID 0+1 data drive configuration has 19,463 messages submitted while the RAID 5 configuration has 18,726. The RAID 0+1 configuration out performed the RAID 5 configuration by a rate of 737 messages.

The same type of impact on the number of message recipients delivered is shown in Table 5. The impact that the RAID level used for the *ProLiant* CL380 Exchange 5.5 data drive has on the resulting number of message recipients delivered. The chart shows that the RAID 0+1 data drive configuration has 79,255 message recipients delivered while the RAID 5 configuration has 77,146. The RAID 0+1 configuration outperformed the RAID 5 configuration by a rate of 2,109 message recipients delivered.

Table 5 also shows the impact that the RAID level used for the *ProLiant* CL380 Exchange 5.5 data drive has on the resulting number of messages sent. The chart shows that the RAID 0+1 data drive configuration has 2,999 messages sent while the RAID 5 configuration has 2,880. The RAID 0+1 configuration out performed the RAID 5 configuration by a rate of 119 messages sent.

Table 6 details response times for various user actions during benchmark testing.

Client Actions	RAID 0+1 95th-Percentile Response Time (in Milliseconds)	RAID 5 95th Percentile Response Time (in Milliseconds)
Read	141	141
Send	812	1,422
Delete	125	62
Move	328	594
Submit	250	78

Table 6: User Response Times (Latencies) from Load Simulator

The configuration tested here, when compared to a similar configuration running the medium user load, would equate to approximately 3500 MMB. For the Real World scenario, the system tested had enough available resources to maintain a high performance rate for an eight-hour period and beyond.

Summary

Depending on the fault tolerance needs of an organization, each RAID set has its advantages and disadvantages. While providing fault tolerance with very good performance, RAID 0+1 support consumes more available storage than RAID 5. For example, a configuration that uses ten 9GB drives for the Exchange data drive, obtains a capacity of 45GB if using RAID 0+1 as the RAID level. High performance is obtained using the RAID 0+1 configuration. However, only 50 percent of the total usable storage space can be utilized since the remaining 45GB is used for RAID 0+1 fault tolerance support. If maximizing usable storage capacity is a top customer requirement, then RAID 5 is recommended. The same ten drive set formatted with RAID 5 provides approximately 81GB of storage. This is a capacity difference of approximately 36GB. Maximizing usable storage while retaining a level of fault tolerance is one factor that makes RAID 5 drive configurations an alternative that many customers choose. However, if the customer configuration decision is based solely on the need to optimize Microsoft Exchange performance, a RAID 0+1 configuration would be the recommended drive configuration alternative. The RAID 0+1 configuration provided better user response times, higher number of messages being handled and a solid fault tolerant solution. Each customer has unique needs and Microsoft Exchange requirements that must be analyzed before a RAID level configuration decision can be made.

Appendix A
Medium User versus Real World Test Parameters

The following is a chart comparing the Medium user load, as defined by Microsoft Load Simulator, and the Real World test parameters.

Medium User Parameters		Real World User Parameters	
Parameter	***Medium***	***Parameter***	***Real World***
Topology Properties		**Topology Properties**	
Security		***Security***	
Use a separate account for each Exchange user?	No	Use a separate account for each Exchange user?	No
Use one account for all Exchange users?	Yes	Use one account for all Exchange users?	Yes
Distribution Lists		***Distribution Lists***	
Use DLs?	Yes	Use DLs?	Yes
# DLs per site	30	# DLs per site	100
DL min/avg/max	2/10/20	DL min/avg/max	2/10/20
Test Properties		**Test Properties**	
Tasks		Tasks	
Send Mail		***Send Mail***	
# of times per day	4	# of times per day	7
Priority % High	0	Priority % High	1
Priority % Low	0	Priority % Low	0
Request Receipts % Delivery	0	Request Receipts % Delivery	0
Request Receipts % Read	0	Request Receipts % Read	0
Request Receipts % Both	0	Request Receipts % Both	0
Filename	***Weight***	***Filename***	***Weight***
Oups1k.msg	60	Oups1k.msg	37
Oups2k.msg	13	Oups2k.msg	18
Oups4k.msg	5	Oups4k.msg	14
Oups10kat.msg	5	Oups10kat.msg	0
OupsWDatt.msg	8	OupsWDatt.msg	7
OupsXLatt.msg	5	OupsXLatt.msg	7

Parameter	*Medium*	*Parameter*	*Real World*
Filename	*Weight*	*Filename*	*Weight*
OupsBMobj.msg	2	OupsBMobj.msg	10
OupsXLobj.msg	2	OupsXLobj.msg	0
McPP1Matt.msg	Na	McPP1Matt.msg	1
McPP100katt.msg	Na	McPP100katt.msg	5
McWD2Matt.msg	Na	McWD2Matt.msg	1
Recipients per Message	3	Recipients per Message1-5, avg	3
Add a Dist List to % msg sent	30	Add a Dist List to % msg sent	30
Save a copy in sent ?	Yes	Save a copy in sent ?	Yes
Process Inbox		**Process Inbox**	
Read new mail per day	12	Read new mail per day	12
Message actions		*Message actions*	
Reply	7	Reply	20
Reply All	5	Reply All	7
Forward	7	Forward	10
Delete	40	Delete	100
Move	20	Move	30
Copy	0	Copy	0
Read note delay min/avg/max	1.0/1.0/1.0	Read note delay min/avg/max	1.0/1.0/1.0
Load % of attachments	25	Load % of attachments	75
Accept % of meeting requests	70	Accept % of meeting requests	70
Browse Mail		**Browse Mail**	
Browse mail per day	15	Browse mail per day	15
Public Folder Post		*Public Folder Post*	
Not tested	Not tested	Not tested	Not tested
Browse Public Folders		*Browse Public Folders*	
Not tested	Not tested	Not tested	Not tested
Free/Busy		*Free/Busy*	
Update schedule times per day	Not tested	Update schedule times per day	4
Update free/busy Information?	Not tested	Update free/busy Information?	No
Schedule Size(kb) min/max/avg	Not tested	Schedule Size(kb) min/max/avg	5/40/22

Parameter	Medium	Parameter	Real World
Request Meetings		*Request Meetings*	
Make new meetings per day	0.10	Make new meetings per day	2
Meeting length(hrs) min/avg/max	1/2/4	Meeting length(hrs) min/avg/max	1/2/7
Attendees min/avg/max	1/3/20	Attendees min/avg/max	1/5/40
Add a DL % of the tim	10	Add a DL % of the time	20
Make Appointments		*Make Appointments*	
New appts per day	0.2	New appts per day	4
Appt length min/avg/max	1/2/5	Appt length min/avg/max	1/3/9
% recurring appointments	15	% recurring appointments	15
% all day events	5	% all day events	5
Browse Calendar		*Browse Calendar*	
# of times per day	3	# of times per day	6
Journal Mail Items		Journal Mail Items	
# of times per day	Not tested	# of times per day	Not tested
Journal Applications		*Journal Applications*	
Activity # of times per day	1.5	Activity # of times per day	3
Logoff		*Logoff*	
# of times per day to log off	Not tested	# of times per day to log off	3
Always keep connection?	Not tested	Always keep connection?	No
Empty deleted items?	Not tested	Empty deleted items?	Yes
Browse Contacts		*Browse Contacts*	
# of times per day	Not tested	# of times per day	10
Create Contact		*Create Contact*	
# times/day to make new contact	Not tested	# times/day to make new contact	1.4
Test/Logon		*Test/Logon*	
Logon immediately at the very beginning of the test?	Yes	Logon immediately at the very beginning of the test?	Yes
Log off at the end of each simulated day?	No	Log off at the end of each simulated day?	Yes
Empty Deleted Items folder on logoff?	Yes	Empty Deleted Items folder on logoff?	Yes
Test Report Approximate message traffic, per user, per day		Test Report Approximate message traffic, per user, per day	
Total Received	66.3	Total Received	161.90
Reply	3.76	Reply	20.56

Parameter	*Medium*	*Parameter*	*Real World*
Reply all	2.67	Reply all	6.48
Forward	3.76	Forward	10.08
Total Submitted	14.18	Total Submitted	44.12
Average # of recipients per message (all msgs)	4.68	Average # of recipients per message (all msgs)	3.67
Approx receipts requested, per user, per day		*Approx receipts requested, per user, per day*	
Read receipts	0	Read receipts	0
Delivery receipts	0	Delivery receipts	0
Initialization		**Initialization**	
Mailbox Setup		*Mailbox Setup*	
# messages in Inbox	4	# messages in Inbox	55
# messages in Deleted Items	1	# messages in Deleted Items	1
# of new folders	40	# of new folders	10
# of messages in new folders	5	# of messages in new folders	55
Calendar Setup		*Calendar Setup*	
# of appointments	25	# of appointments	25
Contact Setup		*Contact Setup*	
# of contacts	64	# of contacts	64

- *mcWD2Matt.msg is a new message file with a 2MB Word attachment*
- *mcPP100kat.msg is a new message file with a 100KB PowerPoint presentation attachment*
- *mcPP1Mat.msg is a new message file with a 1MB PowerPoint presentation attachment*

Notice

Compaq ProLiant CL380 Performance Optimization

for Microsoft SQL Server 7.0 Enterprise Edition with Microsoft Cluster Server (NT 4.0)

White Paper

June 2000
12RD-0600A-WWEN

Prepared by:
High Availability Systems Engineering

Compaq Computer Corporation

Abstract: This document covers Microsoft SQL Server 7.0 Enterprise Edition performance optimization on a Compaq *ProLiant* CL380 running with Microsoft Cluster Server on the Windows NT Server 4.0 Enterprise Edition operating system. The test methodology assumed that the executable files, log files, data files and temporary database files were located on separate logical drives. An industry-standard benchmark was used to measure the impact that the RAID levels implemented combined with the location of the drives had on the performance results.

Introduction

The purpose of this document is to provide the optimal configuration for Microsoft SQL Server 7.0 running on a Compaq *ProLiant* CL380. Additional information regarding the *ProLiant* CL380 in Microsoft Windows 2000 Advanced Server, Novell NetWare and SCO UnixWare clusters as well as performance White Papers for other database and messaging applications can be found at **www.compaq.com/highavailability**.

Expected Audience

This document is intended to help in the installation, configuration and administration of Microsoft Cluster Server with Microsoft SQL Server and assumes that the reader has working knowledge of the following:

- Installing and configuring Compaq *ProLiant* CL380
- Installing and configuring Microsoft Windows NT Server 4.0 Enterprise Edition
- Installing and configuring Microsoft SQL Server 7.0 Enterprise Edition
- Installing, configuring and administering Microsoft Cluster Server
- Installing and running Benchmark Factory

Compaq ProLiant CL380 Performance Optimization

for Microsoft SQL Server 7.0 Enterprise Edition
with Microsoft Cluster Server (NT 4.0)

White Paper

Compaq *ProLiant* CL380

The Easiest Most Affordable Clustering Solution

The *ProLiant* CL380 is a uniquely packaged two-node cluster that is designed to simplify clustering for business critical applications. *ProLiant* CL380 consists of two Compaq server "nodes" and shared storage pre-packaged in a cost-effective, space-efficient cabinet, giving customers the easiest, most affordable clustering solution for NT.

The *ProLiant* CL380 is ideal for:

- Remote systems requiring unattended high availability
- Branch Office industry applications – retail, health care, financial
- Dedicated function servers needing high availability – SQL Server, Microsoft Exchange, Novell GroupWise, Lotus Domino and file/print
- Departmental applications that need high availability
- Ideal for space constrained data centers and business offices

Benefits at a Glance

- Complete cluster housed in a single 10U cabinet, pre-wired and pre-tested, makes it easy to install and easy to order through a single part number.
- Convenient size, flexible configuration and integrated Compaq management tools mean easy deployment and rapid install.
- Compaq's newly released *Cluster Monitor* (part of Compaq *Insight Manager XE)* and Compaq *Intelligent Cluster Administrator* enable Web-based remote monitoring and administration of the cluster across the Internet or your intranet, making *ProLiant* CL380 easy to manage and minimize the need for local technical support.
- Commonality with *ProLiant* servers and other industry-standard technologies means low purchase price.
- On-line serviceability and easy access to hot swappable components makes the *ProLiant* CL380 easy to service and reduces downtime, resulting in low cost of maintenance.
- Substantial in-system growth provides low cost incremental expansion of performance and storage capacity without the cost of system upgrade or change.
- Ready for the Enterprise – Service and Support team up with the channel to deliver and support business-critical requirements.

Key Features

- Innovative low-cost, space-saving design
- Small footprint packaging of two clustered *ProLiant* server nodes plus shared storage
- Unique chassis design offers deployment flexibility for tower or rack mounting
- Integrated switch to share keyboard, mouse and monitor between server nodes

Serviceability

- On-line serviceability of all major components
- Hot Plug server and shared storage disks
- Hot Plug power supplies for shared storage
- Easy component access for ease of maintenance

Performance

- Two Intel Pentium III 800 MHz processors with 256 KB of secondary (L2) cache per server
- Up to 4GB of ECC-protected, 133-MHz SDRAM system memory in each server node for faster clock speeds and increased system performance
- High performance 10,000 rpm SCSI disks
- Four PCI expansion slots
- Up to 252GB of high performance SCSI storage

Reliability

- Cluster software provides the capability for application or server failover to maintain business operations in the event of hardware or software failures

- High performance RAID array controller for fault tolerant protection of shared data to maximize uptime
- Redundant power supplies for shared storage to provide continuous power after an internal power supply failure
- ECC memory can prevent memory errors from causing server failure
- 128MB ECC-protected SDRAM memory standard, expandable to 4GB
- Processor recovery for automatic reboot, fail-over processor redundancy
- Redundant NIC support with failover of heartbeat to maintain cluster
- Year 2000 support for smooth date transition
- Extensive variety of internal tape backup support

Manageability

- Web-based cluster management
- Compaq *SmartStart* installation utility
- Cluster verification utility
- Compaq *Remote Insight Board* option
- Color-coded cable receptacles

Ensuring Data Integrity

Ensuring data integrity is the first step in planning a database server. This process requires that correct fault tolerance levels be determined and adequate backup and recovery procedures be developed.

This section provides an overview of Redundant Array of Inexpensive Disks (RAID) and provides information that will help you select the appropriate RAID levels. This section also includes an overview of Compaq RAID products and sample storage configurations.

RAID Basics

Many databases use a Redundant Array of Inexpensive Disks (RAID) to improve performance and boost reliability. Three RAID levels are used in database environments: (1) Drive Striping (RAID 0), (2) Drive Mirroring (RAID 1) and (3) Distributed Data Guarding (RAID 5). Table 1 describes these levels.

Table 1. RAID Levels

RAID Level Drive Striping (RAID 0)

Description

RAID 0 is known as stripe sets because data is simply distributed, or striped, across all of the drives in the array. This RAID level does not provide data redundancy and therefore provides no fault protection against data loss.

Striping involves dividing the total storage available into some number of evenly sized blocks. The block size, called the stripe size, is typically no smaller than a disk sector (commonly 512 bytes) and usually some geometric multiple of the minimum size (1KB, 2KB, 4KB, 8KB, etc.). The first block is stored on the first drive of the array, the second block on the second drive and so on. In any RAID 0 stripe set, block M (starting with 0) in a set of N drives would be stored as block M/N on drive M mod N, where the drive numbering starts with 0.

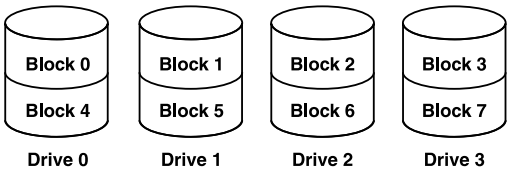

If any drive in the array fails, the entire array will fail. This configuration provides high performance at a low cost.

Drive Mirroring (RAID 0+1)

In the past, drive mirroring allowed a single drive to be duplicated onto another drive. Every write operation was applied to both drives, and all read operations accessed the primary drive. If the primary drive failed, read operations were issued against the mirror drive. Current implementations combine striping with mirroring to deliver improved performance and manageability. This arrangement is referred to as RAID 0+1.

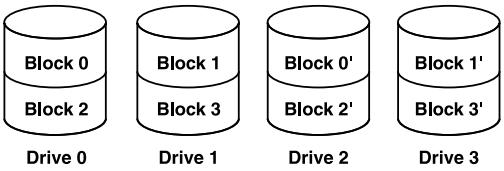

A RAID 0+1 array remains accessible when any drive in the array fails. The array can survive multiple drive failures provided that one drive of each mirrored pair of drives remains operational.

RAID 0+1 requires twice the number of drives as RAID 0 to provide the same storage capacity. Depending on the application's I/O profile and the RAID implementation specifics, a RAID 0+1 array can provide better performance (using twice the number of drives) than RAID 0.

Distributed Data Guarding (RAID 5)

This RAID level uses striping (similar to RAID 0) and parity to provide fault tolerance. Within each stripe, one of the drives stores the parity information for the remainder of the stripe. The parity information alternates among the drives from stripe to stripe, ensuring that I/O operations are evenly distributed across the array. RAID 4, Data Guarding, is similar to RAID 5 except that all the parity information is stored on one of the drives, resulting in uneven I/O distribution and poorer performance.

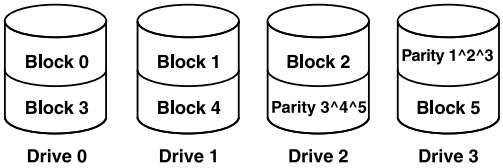

The drive array remains accessible when any drive in the array fails, but the array is no longer fault tolerant; data loss will occur if a second drive fails.

When considering storage capacity, remember that Distributed Data Guarding requires one more drive than RAID 0 (RAID 5 requires at least three drives). RAID 5 write performance is much lower than both RAID 0 or RAID 1 because for each write required for I/O operations, old data and parity information must be read and the new data and updated parity information must be written. At a 2:1 Read:Write ratio, a RAID 5 configuration generates twice the number of I/Os as a RAID 0 configuration and 50 percent more I/Os than RAID 1. In Read Only workloads, RAID 5 delivers roughly the same performance as RAID 0.

Test Methodology

Benchmark Factory is the benchmark tool selected to collect SQL Server 7.0 performance data to determine the optimal Compaq *ProLiant* CL380 configuration. Benchmark Factory is a software suite designed to develop synthetic server benchmarks that measure the performance of your server software and hardware. The selected Benchmark Factory TPC-C workload returns its results as transactions per second (TPS). Each client totals the number of transactions executed and divides that number by the time it took to do the work. For the overall server throughput, Benchmark Factory combines the entire client throughput to create the number of transactions per second for the server's total system throughput.

Benchmark Factory uses an industry-standard methodology to collect accurate timing statistics about the server. Both hardware and software vendors, as well as prominent professional publications, use this methodology, which has been developed over a 10-year period. The methodology uses defined test phases, where each phase has a specific purpose in providing accurate and reproducible results. The phases are described below:

- Quiet Time – the time to wait before executing transactions against the server. Quiet Time allows the network to stabilize and the clients to resynchronize.
- RampUp Time – the fixed amount of time to wait before timing transactions. Allowing the transactions to execute for a period of time before actually collecting statistics gives the system under test time to reach a steady state. RampUp Time also limits skewed results caused by an inconsistent server load when the first few clients start.
- Timing Phase – the fixed amount of time the statistics are measured for the transaction mix. Factors that affect the Timing Phase are server configuration, number of virtual users and the type of workload.
- RampDown Time – the fixed amount of time at the end of the execution period to ensure that statistics are collected with a consistent load. RampDown Time assures that your test results are not skewed because the server load for the last few clients was very light, thus resulting in an artificially high score.
- Think Time – the delay between transaction executions, used to simulate the time that a user is "thinking" about the results of the transaction.
- Keying Time – the additional constant delay used to simulate time spent entering data.

ProLiant CL380 Configuration

Below are the components used for the performance testing:

Hardware Components

1 – Compaq *ProLiant* CL380 (2 server nodes)
2 – CR3500 RAID controller
4 – 800 MHz Intel Processor (2 per server)
2 – 1 GB Memory module (1 per server)
1 – RA4214 Storage unit
16 – 9.1 GB Ultra2 drive (1 per server) (14 shared storage)

Table 1 – Hardware Component

Software Components

SmartStart 4.70
Windows NT Server 4.0 Enterprise Edition
Microsoft Cluster Server 1.0
SQL Server 7.0 Enterprise Edition
Microsoft Service Pack 6a
Microsoft SQL Server Service Pack 1
Benchmark Factory build 238
CR3500 Array Configuration Utility

Table 2 – Software Components

Internal Shared Storage

It is important to understand the SCSI ID numbering that is used in the shared storage area of the *ProLiant* CL380. The CR3500 Array Configuration Utility uses the SCSI IDs when configuring physical drives into drive arrays, so it is important to relate the SCSI ID with the physical drive position. The SCSI IDs addressing the six internal shared storage drives are SCSI IDs 9-14. The remaining SCSI IDs with the exception of SCSI IDs 6 and 7, which are reserved for use by the disk side bus connections to the CR3500 Controllers, are used for representing drives in the external shared storage cabinet.

External Shared Storage

When more storage is needed than can be provided by the six internal shared storage drives, a StorageWorks Enclosure 4214 may be added. This extends the disk side bus to provide more storage capacity. Figure 1 describes the SCSI IDs and the relative placement in both the internal and external shared storage areas.

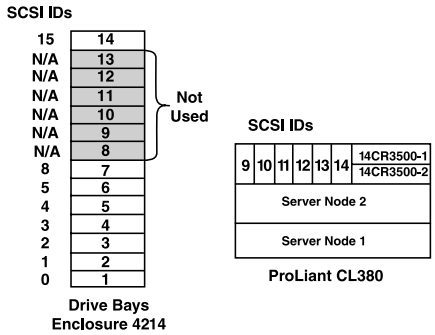

Figure 1 – SCSI ID Numbering

The optimal SQL Server 7.0 configuration assumes that the executable files, log files, data files and temporary database files will be on separate logical drives. The executables and the log files reside on separate RAID 1 sets. RAID 1 is used to provide fault tolerance. The remaining two variables are the logical drive locations for the data and temp db files. Seven configurations for each RAID set (RAID 0, 0+1, 5) are used to validate the optimal configuration. Each configuration focuses on the placement of the data and temp db files. On the next page are the details of each configuration used in this benchmark exercise.

Configuration	File Type	RAID Type	SCSI ID(s)	Host Port
1	Executables *	RAID 1	8,9	1
	Log files	RAID 1	10,11	1
	Data files	RAID 0, 0+1, 5	0,1,2,3,4,5	0
	Temp db	RAID 0, 0+1, 5	12,13,14,15	0
2	Executables *	RAID 1	8,9	0
	Log files	RAID 1	10,11	0
	Data files	RAID 0, 0+1, 5	0,1,2,3,4,5	1
	Temp db	RAID 0, 0+1, 5	12,13,14,15	1
3	Executables *	RAID 1	8,9	1
	Log files	RAID 1	10,11	1
	Data files	RAID 0, 0+1, 5	0,1,2,3,4,5	0
	Temp db	RAID 0, 0+1, 5	12,13,14,15	1
4	Executables *	RAID 1	8,9	0
	Log files	RAID 1	10,11	0
	Data files	RAID 0, 0+1, 5	0,1,2,3,4,5	1
	Temp db	RAID 0, 0+1, 5	12,13,14,15	0
5	Executables *	RAID 1	0,1	1
	Log files	RAID 1	2,3	1
	Data files	RAID 0, 0+1, 5	4,5,8,9,10,11	0
	Temp db	RAID 0, 0+1, 5	12,13,14,15	0
6	Executables *	RAID 1	8,9	1
	Log files	RAID 1	10,11	1
	Data / Temp db files **	RAID 0, 0+1, 5	0,1,2,3,4,5,12,13,14,15	0
7	Executables *	RAID 1	8,9	1
	Log files	RAID 1	10,11	1
	Data / Temp db files ***	RAID 0, 0+1, 5	0,1,2,3,4,5,12,13,14,15	0

* *Two logical partitions 100 MB for quorum drive, remainder used for executables.*
** *Two logical partitions using an 80 percent for data and 20 percent for temp db. *** one logical partition.*
The SCSI IDs above are used according to their priority as defined by the SCSI-2 interface protocol.

Table 3 – Configurations

Table 4 summarizes each SCSI ID and its priority level as defined by the SCSI-2 interface protocol.

SCSI ID	Priority	SCSI ID	Priority
7	1	5	9
6	2	14	10
5	3	13	11
4	4	12	12
3	5	11	13
2	6	10	14
1	7	9	15
0	8	8	16

Table 4 – SCSI ID and Priority Level

Benchmark Factory Settings

The following settings are used in the performance benchmark runs:

Benchmark Factory Phase	Duration
Quiet Time	30 seconds
RampUp Time	5 minutes
Timing Phase	10 minutes
RampDown Time	15 minutes
Think Time (Absolute Value)	750 milliseconds
Keying Time (Absolute Value)	750 milliseconds

Results

Each RAID set is summarized in chart form. The charts represent the transactions per second measured during the Benchmark Factory test runs at user intervals from 50 to 500, stepping by 50 users per data-collection interval. Appendix A contains the individual results for each RAID set. Refer to Table 3 for the details of each test configuration.

When examining the test results that follow, note how the combination of RAID level, SCSI ID of drives used in the RAID set and the host port designated to control the drives contribute to the overall benchmark performance results.

RAID 5

Refer to Table 3 for the details of each test configuration used in the RAID 5 tests.

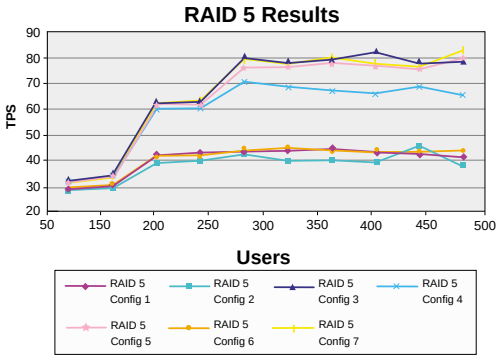

Figure 2 – RAID 5 Results

Configuration 7 provided the highest transaction per second results because the 10 physical drives being controlled by the highest priority CR3500 controller are set up as one logical partition. This partition serves as the storage location for the highly accessed database files as well as the temporary database files. The executables and log files are controlled via a separate, lower priority host port.

Configurations 3 and 4 are very similar in performance capabilities – 78.53 TPS and 79.96 TPS, respectively – making Configuration 4 perform only slightly better. Both configurations consist of two RAID 1 sets for executables and log files controlled by one CR3500 controller. Configuration 3 uses the same lower priority host port to also control a four-drive RAID 5 set used for temporary database files. The higher priority host port is used to handle only the data drive in Configuration 3. Configuration 4 uses the higher priority host port to control the executables, logs and temporary database files. The data drive is controlled by the lower priority host port. Configuration 4 with three drives controlled by the higher priority host port slightly outperforms Configuration 3.

Configuration 6 also has 10 physical drives being controlled by the highest priority CR3500 controller. These drives are broken into two logical partitions, which seems to introduce a level of contention on the CR3500 controller since both drives are being heavily accessed at the same time. The executables and log files are still controlled via a separate host port, using the lower priority CR3500. While the configuration difference seems slight, note that Configuration 6 produces a TPS result of 43.79, which is only slightly more than 50 percent of the TPS for Configuration 7, which is 82.98.

Configurations 1 and 2 both introduce a level of contention on a single host port because one controller is handling two RAID 5 sets. One RAID 5 set is made up of six disks and used for the data files and the second RAID 5 set is made up of four disks and used for the temporary database files. The second host port is controlling two RAID 1 sets, one for the executables and one for the log files. Configurations 1 and 2 are very similar in performance capabilities – 41.12 TPS and 37.48 TPS, respectively – making Configuration 1 perform only slightly better. This edge in performance is due to the higher priority host port being assigned to the RAID 5 sets that are the more active database and temporary database drives in Configuration 1.

RAID 0+1

Refer to Table 3 for the details of each test configuration used in the RAID 5 tests.

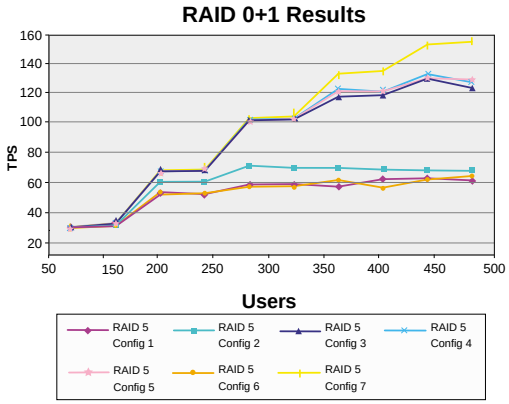

Figure 3 – RAID 0+1 Results

Configuration 7 provided the highest transactions-per-second results, 138.9, because the 10-disk RAID 0+1 set being controlled by the highest priority CR3500 controller is set up as one logical partition. This partition serves as the storage location for the highly accessed database files as well as the temporary database files. The two RAID 1 sets storing executables and log files are controlled via a separate, lower priority host port.

Configurations 3 and 4 are fairly similar in performance capabilities – 112.26 TPS and 117.13 TPS, respectively – making Configuration 4 perform only slightly better. Both configurations consist of two RAID 1 sets for executables and log files controlled by one CR3500 controller. Configuration 3 uses the same lower priority host port to also control a four-drive RAID 0+1 set used for temporary database files. The higher priority host port is used to handle only the

data drive in Configuration 3. Configuration 4 uses the higher priority host port to control the executables, logs and temporary database files. The data drive is controlled by the lower priority host port. Configuration 4 with three drives controlled by the higher priority host port slightly outperforms Configuration 3. Configuration 5 falls in between Configuration 3 and 4 with a TPS rate of 117.13.

Configuration 6 also has 10 physical drives being controlled by the highest priority CR3500 controller. These drives are broken into two logical partitions, which seems to introduce a level of contention on the CR3500 controller since both drives are being heavily accessed at the same time. The executables and log files are still controlled via a separate host port, using the lower priority CR3500. While the configuration difference seems slight, note that Configuration 6 produces a TPS result of 58.88, which is less than 50 percent of the TPS for Configuration 7, which is 153.19.

Configurations 1 and 2 both introduce a level of contention on a single host port because one controller is handling two RAID 0+1 sets. One RAID 0+1 set is made up of six disks and used for the data files and the second RAID 0 set is made up of four disks and used for the temporary database files. The second host port is controlling two RAID 1 sets, one for the executables and one for the log files. Configurations 1 and 2 are very similar in performance capabilities – 86.57 TPS and 96.18 TPS, respectively – making Configuration 2 perform only slightly better.

RAID 0

Refer to Table 3 for the details of each test configuration used in the RAID 5 tests.

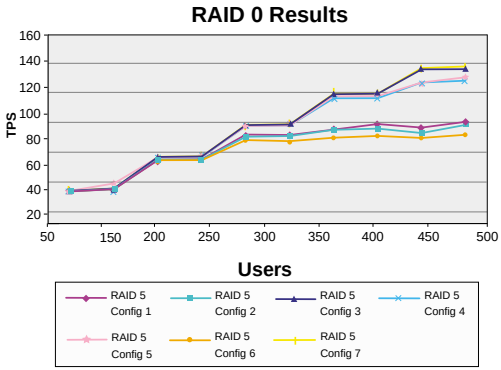

Figure 4 – RAID 0 Results

Configuration 7 provided the highest transaction-per-second results, 153.19, because the 10-disk RAID 0 set being controlled by the highest priority CR3500 controller is set up as one logical partition. This partition serves as the storage location for the highly accessed database files as well as the temporary database files. The two RAID 1 sets storing executables and log files are controlled via a separate lower-priority host port.

Configurations 3 and 4 are fairly similar in performance capabilities – 150.5 TPS and 142.61 TPS, respectively – making Configuration 3 perform only slightly better. Both configurations consist of two RAID 1 sets for executables and log files controlled by one CR3500 controller. Configuration 3 uses the same lower-priority host port to also control a four-drive RAID 0 set used for temporary database files. The higher-priority host port is used to handle only the data drive in Configuration 3. Configuration 4 uses the higher-priority host port to control the executables, logs and temporary database files. The data drive is controlled by the lower priority host port. Configuration 3 with the data drive controlled by the

higher priority host port slightly outperforms Configuration 4. Separating the data drive on Configuration 4 so it is controlled solely by one host port is good, but assigning it to the lower priority host port prevents this configuration from outperforming Configuration 3. Configuration 5 falls just below Configuration 4 with a TPS rate of 138.27. Assigning both the database file drive and the temporary database file drive to the same higher priority host port combined with the SCSI ID change for that one run results in Configuration 5 performing slightly worse than Configuration 4.

Configuration 6 also has 10 physical drives being controlled by the highest priority CR3500 controller. These drives are broken into two logical partitions, which seems to introduce a level of contention on the CR3500 controller since both drives are being heavily accessed at the same time. The executables and log files are still controlled via a separate host port, using the lower priority CR3500. While the configuration difference seems slight, note that Configuration 6 produces a TPS result of 99.17, which is more than 50 percent of the TPS for Configuration 7, which is 153.19.

Configurations 1 and 2 both introduce a level of contention on a single host port because one controller is handling two RAID 0 sets. One RAID 0 set is made up of six disks and used for the data files and the second RAID 0 set is made up of four disks and used for the temporary database files. The second host port is controlling two RAID 1 sets, one for the executables and one for the log files. Configurations 1 and 2 are very similar in performance capabilities – 86.57 TPS and 96.18 TPS, respectively – making Configuration 2 perform only slightly better.

All RAID Levels

Refer to Table 3 for the details of each test configuration used in the RAID 5 tests.

Based on the results, the optimally performing configuration is Configuration 7. Below is a summary graph of Configuration 7 for each RAID set.

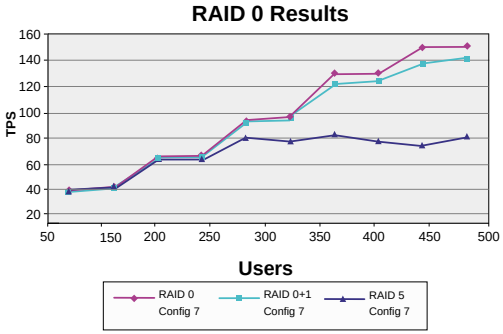

Figure 5 – Configuration 7 Results for All RAID Levels

The RAID 0 configuration performed the best, providing the highest transactions per second of any configuration. There are instances where performance is more important than fault tolerance. In this case, a decision to implement a RAID 0 configuration may be reasonable. However, remember that RAID 0 provides no fault tolerance and therefore no high availability.

The RAID 0+1 configuration provided the next highest transactions-per-second results. Note that the RAID 0+1 and the RAID 0 results indicate that they may scale even higher with a benchmark load greater than 500 Benchmark Factory users. However, the RAID 5 curve seems to flatten out at a load of around 250 Benchmark Factory users. The RAID 5 test results began to fall below the RAID 0 and RAID 0+1 results at around 250 Benchmark Factory users. This trend

continued until the RAID 5 result was roughly 60 transactions per second lower than RAID 0+1 and about 75 transactions per second lower than the RAID 0 data point at 500 Benchmark Factory users.

Summary

Depending on the fault-tolerance needs of the organization, each RAID set has its advantages and disadvantages. While RAID 0 provides the best performance, it does not provide fault tolerance. If one of the drives fails, the entire RAID set has to be restored. RAID 0+1 provides fault tolerance, but consumes the most available storage space because of its implementation. RAID 5 does not perform better than RAID 0+1, but RAID 5 provides fault tolerance and more available storage space than RAID 0+1. For example, ten 9GB drives would yield a total of about 81GB of formatted storage space using RAID 5 and about 45GB of formatted storage space using RAID 0+1. Thorough work should be performed when determining a SQL server configuration to ensure the best design.

Using the previous results, it is apparent that Configuration 7 is the best choice for a SQL Server 7.0 implementation on a Compaq *ProLiant* CL380.

Referenced Documents

- Tuning *ProLiant* Servers with SQL Server 7.0 (June 1999)
- Configuring *ProLiant* Servers for Microsoft SQL Server 7.0 (June 1999)
- SCSI-2 Technical Specification Working Draft (1993)

Appendix A –
Results of Each RAID Set

USERS	RAID 0 Config 1	RAID 0 Config 2	RAID 0 Config 3	RAID 0 Config 4	RAID 0 Config 5	RAID 0 Config 6	RAID 0 Config 7
50	31.31	31.42	31.99	31.97	31.95	31.52	32.02
100	33.54	33.6	34.17	34.14	34.09	33.64	34.18
150	61.79	62.39	64.94	64.61	64.77	62.34	65.11
200	61.86	62.42	65.46	65.19	64.95	62.67	65.47
250	81.39	84.69	96	95.09	95.09	86.77	96.26
300	80.35	85.34	96.56	95.52	95.75	86.34	96.75
350	83.7	91.41	126.11	124.05	122.14	91.6	126.53
400	85.44	92.56	126.57	124.23	122.22	97.02	126.6
450	83.47	88.25	150.33	137.51	137.54	93.57	151.54
500	86.57	96.18	150.5	142.61	139.27	99.17	153.19

RAID 0+1 USERS	RAID 0+1 Config 1	RAID 0+1 Config 2	RAID 0+1 Config 3	RAID 0+1 Config 4	RAID 0+1 Config 5	RAID 0+1 Config 6	RAID 0+1 Config 7
50	31.8	31.79	31.8	31.82	31.84	31.41	31.96
100	32.44	32.87	33.97	33.97	33.99	32.62	34.12
150	50.67	58.06	64.11	64.1	64.1	52.14	64.7
200	51.49	58.16	64.37	64.43	64.53	50.94	65.17
250	55.22	67.39	93.57	93.34	94.01	56.53	94.84
300	55.52	66.12	94.05	94.39	94.48	56.77	95.81
350	59.11	66.1	107.19	110.39	111.72	55.27	120.47
400	54.69	65.12	108.08	110.28	110.1	59.59	122.06
450	59.4	64.66	117.63	117.85	120.25	60.11	137.33
500	61.42	64.4	112.26	117.13	115.53	58.88	138.9

RAID 5 USERS	RAID 5 Config 1	RAID 5 Config 2	RAID 5 Config 3	RAID 5 Config 4	RAID 5 Config 5	RAID 5 Config 6	RAID 5 Config 7
50	28.6	28.28	31.39	31.32	31.15	29.35	31.35
100	29.84	29.04	33.51	33.61	33.21	30.3	33.54
150	41.7	38.87	62.23	62.04	60.13	41.88	62.31
200	41.91	39.79	62.84	61.7	60.31	42.97	63.37
250	43.75	42.28	79.99	76.24	70.72	43.35	79.64
300	44.82	39.73	78	76.42	68.75	43.66	77.68
350	43.82	39.95	79.37	78.08	67.3	44.38	80.23
400	43.02	39.1	82.26	76.99	66.15	43.22	77.72
450	42.43	45.66	77.79	75.55	68.81	43.26	76.58
500	41.12	37.48	78.53	79.96	65.41	43.79	82.98

Notice

Compaq ProLiant CL380 Performance Optimization for Microsoft SQL Server 7.0 Enterprise Edition with Microsoft Cluster Server (NT 4.0)
White Paper prepared by High Availability Systems Engineering
First Edition (June 2000)
Document Number 12RD-0600A-WWEN

Compaq ProLiant CL380 Performance

for Lotus Domino R5
with Microsoft Cluster Server (NT 4.0)

White Paper

June 2000
12RC-0600A-WWEN

Prepared by:
High Availability
Systems Engineering
Compaq Computer
Corporation

Abstract: This document covers Lotus Domino R5 performance on a Compaq *ProLiant* CL380 running with Microsoft Cluster Server on Windows NT Server 4.0, Enterprise Edition operating system. Included is a performance and resource utilization comparison of Domino R5 running on a standalone server and Domino R5 running with a Microsoft Cluster Server two-node, active-passive cluster configuration. The paper also addresses the performance impact of changing the *ProLiant* CL380 cluster configuration while holding the user load to a constant simulation of 500 users.

Introduction

This document will specifically address performance management for *ProLiant* CL380 clusters operating with Microsoft Windows NT Server 4.0, Enterprise Edition. Additional information regarding the *ProLiant* CL380 in Microsoft Windows 2000 Advanced Server, Novell NetWare and SCO UnixWare clusters as well as performance White Papers for other messaging and database applications can be found at **www.compaq.com/highavailability**.

Compaq ProLiant CL380 Performance

for Lotus Domino R5
with Microsoft Cluster Server (NT 4.0)

White Paper

Expected Audience

This document is intended to help in the installation, configuration and administration of Microsoft Cluster Server with Lotus Domino R5 and assumes that the reader has working knowledge of the following:

- Installing and configuring Compaq *ProLiant* CL380
- Installing and configuring Microsoft Windows NT Server 4.0 Enterprise Edition
- Installing and configuring Lotus Domino R5 Enterprise Edition
- Installing, configuring and administering Microsoft Cluster Server
- Installing and running industry-standard Domino R5 mail benchmark

Compaq *ProLiant* CL380 Product Overview

Making clustering easy and affordable for businesses of all sizes, the Compaq *ProLiant* CL380 is a uniquely packaged two-node cluster that is designed to simplify clustering for business critical applications such as Domino R5 messaging.

The *ProLiant* CL380 is positioned to meet the requirements of business applications in branch offices, remote locations or departmental computing where space is at a premium, technical support is minimal and simplicity of operation is important. It is ideal for environments that use Microsoft clustering products for Lotus Domino messaging applications. The *ProLiant* CL380 consists of two Compaq server "nodes" and shared storage pre-packaged in a cost-effective, space efficient cabinet giving customers the easiest, most affordable clustering solution for Windows NT.

Benefits at a Glance

- Complete cluster housed in a single 10U cabinet, pre-wired and pre-tested makes it easy to install and easy to order through a single part number.
- Convenient size, flexible configuration and integrated Compaq management tools (easy deployment and rapid install)
- Compaq *Cluster Monitor* (part of Compaq *Insight Manager XE)* and Compaq *Intelligent Cluster Administrator* enable Web-based remote monitoring and administration of the cluster across the Internet or your intranet (easy to manage and minimizes the need for local technical support)
- Low purchase price (affordability provides a high availability solution alternative for cost-sensitive projects)
- On-line serviceability and easy access to hot swappable components makes *ProLiant* CL380 easy to service and reduces downtime, resulting in low cost of maintenance.
- Substantial in-system growth provides low cost incremental expansion of performance and storage capacity without the cost of system upgrade or change
- Ready for the Enterprise – Service and Support team up with the channel to deliver and support business critical requirements.

Key Features

Innovative low-cost, space-saving design

- Small footprint packaging of two clustered *ProLiant* server nodes plus shared storage
- Unique chassis design offers deployment flexibility for tower or rack mounting
- Integrated switch to share keyboard, mouse and monitor between server nodes.

Serviceability
- On-line serviceability of all major components
- Hot Plug server and shared storage disks
- Hot Plug power supplies for shared storage
- Easy component access for ease of maintenance

Performance
- Two Intel Pentium III 800 MHz processors with 256 KB of secondary (L2) cache per server
- Up to 4GB of ECC-protected, 133-MHz SDRAM system memory in each server node for faster clock speeds and increased system performance
- High performance 10,000 rpm SCSI disks
- Four PCI expansion slots, one 32-bit PCI slot and three 64-bit PCI slots (two of the four I/O slots are populated with devices in the standard configuration)
- Up to 252GB of high performance SCSI storage

Reliability
- Cluster software provides the capability for application or server failover to maintain business operations in the event of hardware or software failures
- High performance RAID array controller for fault-tolerant protection of shared data to maximize uptime
- Redundant power supplies for shared storage to provide continuous power after an internal power

Supply Failure
- ECC memory can prevent memory errors from causing server failure
- Processor Recovery for automatic reboot, fail-over processor redundancy
- Redundant NIC support with failover of heartbeat to maintain cluster
- Year 2000 support for smooth date transition
- Extensive variety of internal tape backup support

Manageability
- Web-based cluster management
- Compaq *SmartStart* installation utility
- Cluster verification utility
- Remote Insight Board option
- Color-coded cable receptacles

Compaq Advantage
- Compaq simplifies clustering in a low cost, easy-to-use cluster platform
- Provides for remote cluster monitoring and management – minimizes need for local technical support
- Ready for the Enterprise – Service and support team up with the channel to deliver and support business critical requirements.
- Complete family of *ProLiant* Clusters to address a range of high availability needs in businesses of all sizes

Fault Tolerant Features

Reducing single points of failure is a key consideration for providing the highest levels of availability for your *ProLiant* CL380. Because the *ProLiant* CL380 is highly available, rather than continuously available, it is necessary to understand which parts of the system are vulnerable to faults. Depending on your needs, you may leave all vulnerable areas alone, accepting the risk associated with a potential fault. Or, if the risk of fault is unacceptable for a given component, you may employ redundant components to minimize or remove the single points of failure.

The *ProLiant* CL380 provides many fault tolerant features, which help eliminate potential single points of failure. These features include:

- RAID protected disk subsystems
- Redundant CR3500 controllers
- On-line spare drives
- Redundant shared storage power supplies
- Redundant shared storage fans
- Redundant NICs

Supported RAID Levels

The *ProLiant* CL380 will support up to four internal drives per server node. Up to six drives are supported in the internal shared storage area. Up to eight additional shared drives can be added with the use of an optional *StorageWorks* Enclosure Model 4214/4314, for a maximum of 14 shared storage drives.

The *ProLiant* CL380 supports the following RAID levels:

- JBOD — "Just a Bunch Of Disks," individual drives
- RAID 0 — Striped physical disk group
- RAID 1 — Mirrored physical disk group
- RAID 0+1 — Striped, mirrored physical disk group
- RAID 4 — Striped physical disk group with fixed parity drive (CR3500 Controller only)
- RAID 5 — Striped physical disk group with floating parity drive

These configuration options provide maximum flexibility for the customer to configure their logical drives in a way that fits each application's need for speed, availability and capacity. The fault tolerance method you choose affects the amount of available disk storage capacity and performance of your drive array. If you require a fault tolerant system for critical messaging data, use RAID 5 for maximum storage space efficiency or RAID 0+1 or RAID 1 if I/O performance is more important. If you will be storing non-critical data and both space and performance are important, RAID 0 offers the best of both parameters. However, RAID 0 has no data protection and you will have to rely on backups in case of hardware failure.

Drive arrays in the *ProLiant* CL380 shared storage area have several important attributes:

- A single array cannot span more than one Shared Storage RAID Controller CR3500.
- A single Shared Storage RAID Controller CR3500 can control multiple arrays.
- RAID 1 arrays must have at least two physical drives.
- RAID 0+1 arrays must have at least four physical drives.
- RAID 5 arrays must have at least three physical drives.
- For maximum space efficiency, all drives in a single array should be the same capacity.

More information on drive arrays and RAID levels can be found in Appendix A, "Hard Drive Arrays" in the *ProLiant* CL380 Clustering Primer, which can be found on the *ProLiant* CL380 Documentation CD.

Note: Compaq recommends using RAID 1, RAID 0+1 or RAID 5 for all drives in the shared storage area. Use of drives configured for JBOD, RAID 0 or RAID 4 should be avoided.

Redundant CR3500 Controllers

The *ProLiant* CL380 comes standard with one CR3500 controller, and support for an optional second one. This provides a higher level of availability to the overall system to protect against the failure of one of the CR3500 controllers. If one of the CR3500 controllers fails, the other CR3500 controller will take control of its drives. The two controllers also have access to

each other's write cache memory. The surviving controller will read the data in the failed controller's write cache and continue writing it to disk without interrupting data access. The failed controller can be replaced without powering down the Shared Storage RAID subsystem and without interrupting operation. When a failed controller is replaced, the firmware automatically resets the new controller to the same disk configuration. The replacement controller should have the same firmware version. If the new controller does not have the same firmware version, the firmware should be updated.

Note: The cache memory on each CR3500 controller is mirrored to another memory bank on the same controller. This provides an extra copy of the data in the write cache in case of a memory chip failure. The CR3500 controller does not use an on-board battery to provide backup power to the memory chips in case of a power failure. In the case of a power loss, all data in the write cache would be lost. For this reason, it is imperative to use a UPS to provide backup power for the ProLiant CL380. The use of a properly sized UPS provides sufficient time for all results in both the server and the CR3500 controller cache memory to be written to the disks.

Note: A manual process must be completed before removing an active CR3500 controller from the shared storage system. Removing the CR3500 controller while it is still active will make the system inoperable and may cause loss of data. The controller should be failed over (using the CR3500 Configuration Utility) if two controllers are being used, or the shared storage system should be shut down before removing a single controller.

Note: Compaq recommends the use of a second CR3500 Controller to provide a higher level of fault tolerance. Redundant CR3500 Controllers allow the system to continue to operate uninterrupted even if one of the CR3500 Controllers fails.

If using two controllers, performance can be optimized by having one controller responsible for a portion of the shared storage drives. The second controller can than be assigned the responsibility of the remaining shared storage drives. By intelligently splitting the drives across the two controllers, the I/O load can be optimally distributed resulting in higher performance.

Domino R5 High Availability Features

Domino Reliability

Domino continues to set the standard in reliability and availability with R5's enhanced load balancing and fail-over for Notes Clients and Web browsers. Domino R5 gives you the broadest range of choices for high availability solutions:

- Failsafe directory replication ensuring availability whether you're on the network or working disconnected
- Message tracking and reporting to ensure reliable message delivery
- Expert analysis tools for increased intelligence of server deployment and tasks
- Multi-threaded router to ensure there are no bottlenecks
- Capacity to support up to six clustered servers even across different platforms
- Geographically dispersed clustering for disaster recovery
- Support for both active/passive and active/active levels of MS cluster server

Domino Clustering Alternatives

Compaq provides customers who require a highly available Domino R5 messaging solution with two deployment alternatives:

- **Microsoft Cluster Services and Domino R5 –** Domino Server runs as a resource under Microsoft Cluster Services, which monitors the resource, providing a failover solution in the event of a hardware or operating system problem
- **Domino R5 Clustering –** Domino databases are replicated from one server to another to help ensure that the database will always be available to the client

Microsoft Cluster Server (MSCS)

Microsoft Cluster Server (MSCS) is a product offered in Windows NT Enterprise Edition. MSCS is an operating system level clustering product bringing high availability to Windows NT servers and applications. MSCS for Windows NT currently supports two-node configurations and requires certified system configurations. More than a standby product, MSCS allows both nodes in the cluster to be actively running. In case of a failure, the protected application would be failed over to the second node that would have to assume the additional workload.

Microsoft provides Microsoft Cluster Server as an added component to Windows NT Enterprise Edition. Microsoft Cluster Server is software that runs on top of NT Enterprise Edition to provide operating system level clustering. Two major implementations are available when installing Domino as an application to be controlled by MSCS. Using an active-passive implementation provides support for Domino running on only one of the clustered nodes. Using an Active-Active implementation provides support for one or more instances of Domino running on either or both clustered nodes.

It is not the intent of this paper to give a detailed explanation of MSCS. The reader should have a base line understanding of the MSCS product. A brief description of key concepts and terms is provided to ensure that the reader understands the description of the test configuration provided later in the paper. Additional reading on this subject can be found at **www.microsoft.com** by searching for Microsoft Cluster Server.

Resources

Resources are the lowest entity of the cluster and can be defined as a single manageable unit. MSCS resources includes entities such as:

- Disks
- IP addresses
- Network names
- Applications

The cluster manager has the responsibility of maintaining the state of these resources. The cluster manager must keep track to know if the resource is online or offline and which server is currently operating properly. To prevent the need to code a complex cluster manager that would have to recognize any possible resource that could ever be defined, Microsoft has simply enabled the cluster manager to make a limited number of calls to the resource (online, offline, isalive, looksalive, terminate…). Each resource has associated resource Dynamic Link Libraries (DLLs). These resource DLLs contain the rules to manage the resource. A disk would have a different set of rules than a network name resource or an application resource used to control Domino. Within these DLLs are the corresponding entry points for the cluster calls. To bring a resource online, the cluster manager issues an online command branching to the online routine within that resource's DLL. Within this DLL would be the code required to bring that specific resource online. Other calls work similarly.

Grouping of Resources

With a properly managed resource, it is necessary to group these resources together. This grouping of dependent resource forms "virtual servers." Clients no longer communicate with the physical server in a cluster; they communicate with virtual servers. This group or virtual servers could be running on any server within the cluster at any point in time.

An example of grouping resources together to form a virtual server would be made up of the following:

- **an IP resource**
  the IP address to which the network name would resolve
- **a network name resource**
  the virtual server name that the clients would know to use to access this application
- **a physical disk resource**
  the data and, possibly the program files, that the application would use
- **an application resource grouped together**
  the application that the clients access.

All of these resources will run on the same server. If the group was moved to another server, then all the resources in this group would move together.

Active-Passive Implementation - MSCS with Domino Single Instance

First, the MSCS installation using a single Domino Server instance will be discussed. MSCS provides a cluster administration tool that is used to set up the cluster. Cluster administrator is also used to create entities that can be managed by the cluster. For example, when setting up Domino Server to run as a cluster aware application, new resources must be defined that will make up the new Domino Server group. The Domino Server resource group is made up of the following resources:

- Shared Disk = data drive for Domino Server
- IP Address = IP address used to access this specific Domino Server
- Network Name = Virtual server name given to this group
- Generic Service = The NT service name of the Domino Server

These four resources are based on standard resource types provided with MSCS. During the Domino installation process, the option specifying to run Domino as an NT Service must be selected.

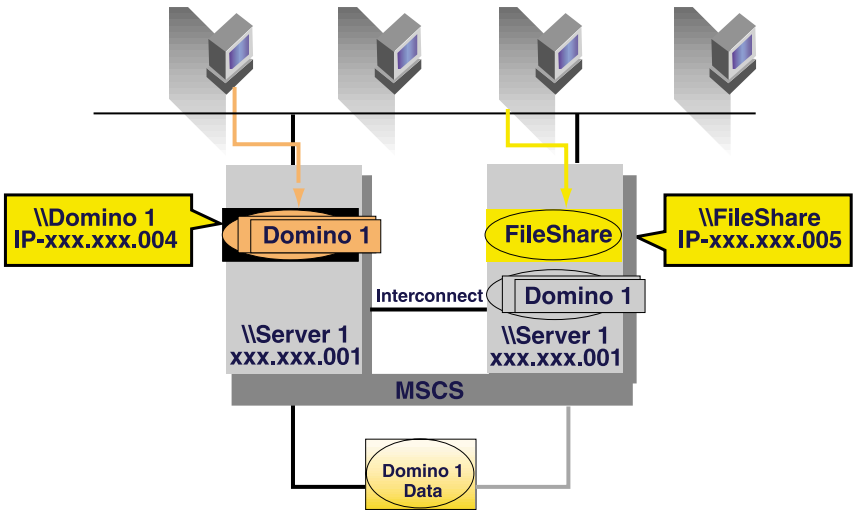

The two cluster members share the same shared storage subsystem RAID set for the data directory specified during the Domino Server installation. Microsoft Cluster Services provides a mechanism to ensure that at any given point in time, only one of the cluster nodes has control of this Domino data drive, which is part of the Domino resource group. In MSCS, the Domino resource group will fail over if the Domino Server can no longer respond to health checks issued by MSCS. MSCS monitors the Domino resources, and if a resource becomes unresponsive to health checks, the Domino server resource group fails over to the other cluster node. Thus, protection against hardware or software problems is provided via fail over. However, it is important to note that if the database is inaccessible because the disk drive or RAID set is down, failover cannot occur. The Domino IP, Domino Disk and Domino Service resources will be stopped on the problematic node and must be able to be started on the second node. If the disk cannot be started due to a storage-related hardware problem, failover cannot occur. This is because a single copy of the database is maintained on the shared storage disk subsystem.

Active-Active Implementation
MSCS with Multiple Domino Instances

Another implementation alternative when using MSCS involves installing multiple instances of Domino Server so that both NT Servers could be running one or more instances of Domino at the same time. This is a very attractive solution for customers because both nodes can now be productively running the Domino Server application. In other words, this allows the customer to run Domino Server on both servers in the cluster rather than having nothing or only file shares running on the second node within the cluster.

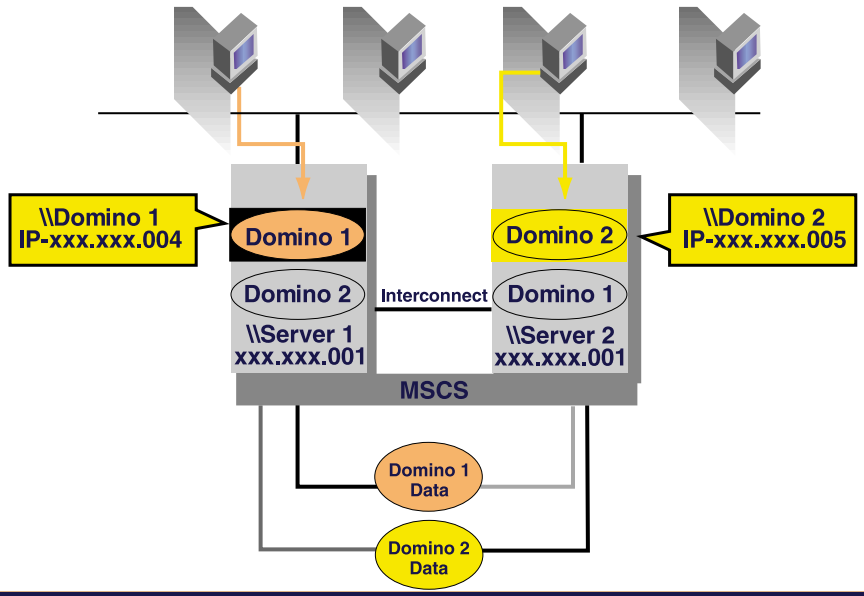

In an Active-Active implementation of Domino, both NT Servers are configured so that at any point in time either can be running two instances or both can be running one instance of Domino Server. For example, initially \\Server1 is running Domino1 and \\Server2 is running Domino2. At any time the administrator could manually move the Domino1 group of resources from \\Server1 to \\Server2 so that scheduled maintenance could be completed without interrupting the user access to data on \\Server2.

Using this Active-Active implementation model, the two Domino Server instances and their corresponding shared disk drives can be owned by either server or one can be owned by each server. For example, \\Server1 is currently shown as the owner in control of the Domino Server group named Domino1 while \\Server2 is shown as the owner in control of the Domino Server group named Domino2. Ownership of the Domino resource group means control of the other resources that are defined to make up that group. In this example, the Domino1 data shared drive, network name Domino1 and the IP address xxx.xxx..004 are owned by \\Server1. The Domino2 data shared drive, network name Domino2, and IP address xxx.xxx.005 are being controlled by \\Server2 because these resources make up the Domino Server group that is currently being controlled by \Server2.

Domino R5 Clustering

Domino clustering is an advanced server feature that enables as many as six Domino Servers located in the same domain and connected by a local area network to be logically grouped into a cluster that is accessible by clients. These clusters of Domino Servers provide high availability of data and services and support a large number of users. Enabling and running clustering requires a Domino Enterprise Server license. Domino server clusters provide the following benefits:

- **Failover** – Clusters provide failover protection for business critical databases and servers. Users continue to have access to a database with failover even when a server goes down. This is accomplished by redirecting the users database requests to another server member of the cluster that is available at the time of the database request.
- **Workload Balancing** – Server workload balancing provides support for heavily used servers to pass requests over to other servers that are cluster members so that work is more evenly distributed among cluster members. Optimal system performance can be achieved using workload balancing.
- **Cluster Replication** – Setting up database replicas on two or more servers in a cluster is required to effectively utilize clustering capabilities. Replication is then used to make certain that all changes made to a database or to the setup of the cluster itself is immediately passed to other databases and servers in the cluster. This is called event-driven replication, meaning that as soon as an event causes a change to a database to be made, the replication process immediately synchronizes the databases rather than waiting for scheduled replication to perform this synchronization. Thus, cluster replication helps provide high availability of information.

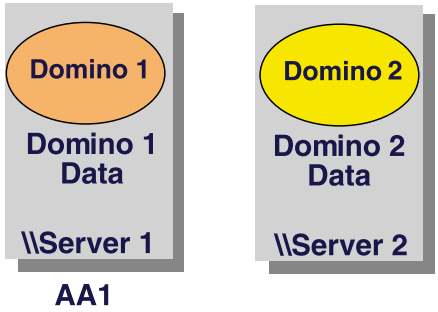

These two servers could be mail servers that have 500 replica mail databases to provide failover capabilities to those 500 Notes Client mail users. The two servers could also be application servers that are responsible for business-critical application databases such as an important Human Resource application. This HR App database may have replica copies on both Domino servers data drives. Special cluster-related NOTES.INI parameters may be used to help ensure that the user workload is evenly balanced across the two cluster members. The details regarding achieving this failover and workload balancing are outside the scope of this paper. Additional reading on this subject can be found at **www.lotus.com** by searching for Domino Clustering.

ProLiant CL380 Performance for Domino R5 Active-Passive Cluster

Included is a performance and resource utilization comparison of Domino R5 running on a standalone server and Domino R5 running in an active-passive configuration in a Microsoft Cluster Server two-node cluster configuration. The scope of the performance testing for the *ProLiant* CL380 was limited to a characterization of Domino R5 performance for an MSCS active-passive implementation because of the target market of the CL380. Test results indicate that a single instance of Domino R5 running on the *ProLiant*

CL380 can easily support 500 Notes mail users. Since the *ProLiant* CL380 packaged cluster is targeted for branch offices, remote locations and small to medium size businesses, the decision to recommend an active-passive implementation was made for the following reasons:

- An active-passive implementation can more than readily handle a user load of 500 users
- Implementing a single Domino Server helps keep the administration procedures as simple as possible
- Implementing a single Domino Server helps keep the maintenance and day-to-day operational procedures as simple as possible

While holding the user load to a constant simulation of 500 users, the following configuration variations were included in the tests:

- **Varying number of processors** Single and dual 800 MHz Pentium III
- **Varying system memory** 1 GB, 512 MB, 384 MB and 256 MB
- **Changing RAID levels of domino data drives** RAID 0+1 and RAID 5
- **Varying total number of shared storage drives** 6 and 14

Several tests were run and their results are included to address the performance impact of changing the *ProLiant* CL380 cluster configuration. The results are presented in table form and chart form using eight scenarios to examine the data and compare and contrast performance results.

Test Methodology

This section outlines the overall test methodology that we used for our Domino R5 mail test scenarios. It includes information about the system configurations, the workloads and our evaluation scenarios.

System configurations:
To run the scenarios, we used the following configurations:

Servers

- CPUs: Compaq *ProLiant* CL380 Nodes with 1-2 Pentium III/800 MHz processors
- Memory: 1 GB, 512 MB, 384 MB or 256 MB as designated
- Page File: Size set using NT 4.0 defaults for each memory configuration
- Network: 100Mbit Ethernet (private)

- OS: Windows NT 4.0, Enterprise Edition, Service Pack 6
- Domino: Release 5.03
- CR3500 Controllers – Redundant controllers used for every test with the load balanced across the controllers to maximize system performance while providing redundancy.

For six drive tests, logical drives E: and F: were controlled by one controller while drive G: was controlled by the other controller. For 14 drive tests, drives E:, F: and G: were controlled by one controller, while drive H: was controlled by the other controller.

Client

- CPUs: Compaq *ProLiant* 850R with one Pentium Pro 200MHz processor
- Memory: 128 MB RAM
- One system disk

- OS: Windows NT Server 4.0
- Domino: Release 5.03
- Industry-standard R5 Mail workload benchmark

Simulation Workloads

For all test scenarios, the industry standard R5 Mail workload was run against the system under test for two hours. The workload simulated 500 R5 Mail clients connecting to their Domino R5.03 mail databases. These clients were simulated from one load generating client, a *ProLiant* 850R, for each test case performed. The performance data desired required that Windows NT PerfMon be used to capture performance data to logs during the test runs.

The R5 Mail workload models an active user on a client reading and sending mail, scheduling an appointment, and sending a calendar and schedule invitation. The average user will execute this test script four times per hour. During each iteration of the script, the following activity occurs:

- Five documents are read
- Two documents are updated
- Two documents are deleted
- One view scrolling operation is exercised
- One database is opened and closed
- One view is opened and closed
- Three messages made up of one memo to three recipients is sent
- Three lookups against the Domino Directory is made
- One Calendar and Scheduling appointment is accepted
- One Calendar and Scheduling appointment is sent to recipients every 90 minutes

The normal message size for the test cases is designated as 10,000 bytes. The number of mail notes per user is set to 100. The number of message recipients is designated as three. The R5 Mail workload was run against the system under test (Domino 1 running on Node1 of the cluster) for each test configuration of the eight scenarios described in the section that follows.

Scenario 1
Overhead of MSCS

In the first test scenario, two configurations were tested to determine the overhead, or cost of system resources, for running MSCS. The first test was run without MSCS and the second test was run with MSCS installed using an Active-Passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1.

Configuration 1
Systems – Two Processors, 1GB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, no MSCS, active system under test

Data Directory – 12 Shared Storage Drives

Logical	*Size*	*RAID Level*	*Usage*
Drive G:	26019 MB	RAID 0+1	Data: Link Files, 150 Mail Files
Drive H:	26019 MB	RAID 0+1	350 Mail Files
Node 2 – NT 4.0 Enterprise Edition, no MSCS, logged in and sitting idle			

Configuration 2
Systems – 2 Processors, 1GB RAM, RAID 0+1

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 14 Shared Storage Drives, 12 Data related drives, 2 drives for Quorum and Transaction logging

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	26019 MB	RAID 0+1	Data: Link Files, 150 Mail Files
Drive H:	26019 MB	RAID 0+1	350 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 Standalone	*Configuration 2 MSCS*	*Difference*
Transactions per minute	710	730	20
Response time (ms)	56	58	2
Avg. processor utilization	5.6%	5.9%	.3%
Avg. available memory	532	531	1
Avg. disk queue length G:	.229	.248	.019
Avg. disk queue length H:	.485	513	.028

Table 1. Domino Running on Standalone Server versus Domino Running with MSCS

The system resource utilization of the system running Domino R5 in a standalone server environment was nearly identical to the system configured with Microsoft Cluster Server running Domino R5. The response time is also nearly identical, so the conclusion can be made that the cost of high availability, in the form of application failover capabilities of Domino running with MSCS, is minimal at a 500 Notes client user load.

Scenario 2
Impact of Number of Processors

In the second test scenario, two configurations were tested to determine the performance impact of the number of processors on system resources and test results when running a 500 user simulation. The tests were run with MSCS installed using an active-passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1.

Configuration 1
Systems – Two Processors, 512 MB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
TDrive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2
Systems – 1 Processor, 512 MB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 2P	*Configuration 2 1P*	*Difference*
Transactions per minute	772	763	9
Response Time (ms)	76	109	33
Avg. Processor Utilization	5.6%	11.5%	5.9%

Table 2. Impact of Number of Processors on Performance

As expected, the performance results of the dual processor configuration demonstrated faster user response time and lower overall system processor utilization than the single processor configuration.

Note: These tests were run to compare the processor usage characteristics of Domino R5 mail server and should not be interpreted as a recommendation to run Domino with a single processor system. When making a configuration decision regarding the number of processors to use, customers are reminded to take into consideration other applications and utilities that may also be running on the *ProLiant* CL380. These test results demonstrate processor and resource utilization rates of a *ProLiant* CL380 running Domino R5 with no supplemental applications or utilities.

Scenario 3
Impact of Memory Using RAID 0+1

In the third test scenario, three configurations were tested to determine the performance impact of memory on system resources and test results when running a 500 user simulation. The tests were run with MSCS installed using an active-passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1.

Configuration 1
Systems – 1 Processors, 512 MB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2
Systems – One Processor, 384 MB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 3
Systems – One Processors, 256 MB RAM, RAID 0+1
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 512 MB	*Configuration 2 384 MB*	*Configuration 3 256 MB*
Transactions per minute	770	744	761
Response time (ms)	85	91	104
Invitations sent	710	738	665
Invitation responses	173	175	173
Avg. available memory	165	93	18

Table 3. Impact of Memory on Performance Using RAID 0+1

As is evident in the table above, the RAID 0+1 configuration with 512MB of RAM produced the best client response time experience. The test results confirmed that an increase in memory should improve the user response time. The 512 MB configuration outperformed the 384 MB configuration by a rate of 6 ms. The 384 MB configuration out performed the 256 MB configuration by a rate of 13 ms. The 512 MB configuration out performed the 256 MB configuration by a rate of 19 ms.

As is evident in the chart above, the RAID 0+1 configuration with 512MB of RAM produced the best client response time experience. The test results confirmed that an increase in memory should improve the user response time. The 512 MB configuration out performed the 256 MB configuration by a rate of 19 ms.

Scenario 4
Impact of Memory Using RAID 5

In the fourth test scenario, three configurations were tested to determine the performance impact of memory on system resources and test results when running a 500 user simulation. The *ProLiant* CL380 was configured identically except for the amount of memory and the corresponding page file size appropriate for each memory configuration. The tests were run with MSCS installed using an active-passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1.

Configuration 1
Systems – 1 Processors, 512 MB RAM, RAID 5 Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 5	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

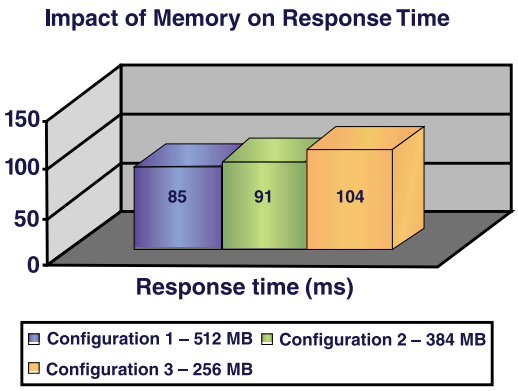

Configuration 2
Systems – One Processor, 384 MB RAM, RAID 5
Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 5	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 3
Systems – One Processor, 256 MB RAM, RAID 5
Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 6 Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346MB	RAID 5	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 512 MB	Configuration 2 384 MB	Configuration 3 256 MB
Transactions per minute	769	774	769
Response time (ms)	108	121	133
Invitations sent	709	715	639
Invitation responses	179	172	128
Total mail routed	3607	3898	4094
Total mail delivered	3607	3898	4094
Avg. available memory	172	93	20

Table 4. Impact of Memory on Performance Using RAID 5

As is evident in the table above, the RAID 5 configuration with 512MB of RAM produced the best client response time experience. The test results confirmed that an increase in memory should improve the user response time. The 512 MB configuration outperformed the 384 MB configuration by a rate of 13 ms. The 384 MB configuration outperformed the 256MB configuration by a rate of 12 ms. The 512MB configuration out performed the 256MB configuration by a rate of 25 ms.

As is evident in the chart above, the RAID 5 configuration with 512MB of RAM produced the best client response time experience. The test results confirmed that an increase in memory should improve the user response time. The 512MB configuration out performed the 256MB configuration by a rate of 25 ms.

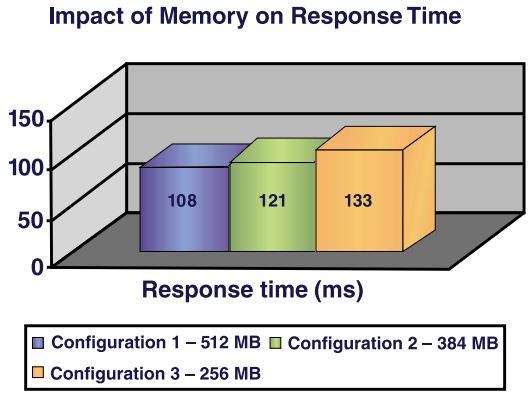

Scenario 5
Impact of RAID Level Used for Shared Storage Subsystem – 6 Drives

In the fifth test scenario, two configurations were tested to determine the performance impact of the RAID level used on system resources and test results when running a 500 user simulation. The RAID level is specifically referring to the level used for the Domino data drive in the shared storage subsystem. The *ProLiant* CL380 with six total shared storage drives was configured identically except for the RAID level used when configuring drive G:, the Domino data directory. Both configurations used two drives in a RAID 1 mirror for drive E:, serving as the MSCS quorum, and drive F:, storing the Domino transaction logs. Thus, both tests used the same configuration for drives E: and F:. As stated, the Domino data directory configuration was different. The first configuration used RAID 0+1 for the remaining four drives and created drive G: to use as the Domino data directory to store the mail databases. The second configuration used RAID 5 for the remaining four drives and created drive G: to use as the Domino data directory to store the mail databases. The tests were run with MSCS installed using an Active-Passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1 of the cluster.

Configuration 1
Systems – One Processor, 384 MB RAM, RAID 0+1 Domino data drive
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2
Systems – One Processor, 384 MB RAM, RAID 5 Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	26019 MB	RAID 5	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 RAID 0+1	Configuration 2 RAID 5
Transactions per minute	744	774
Response time(ms)	91	121
Invitations sent	738	715
Invitation responses	175	172
Total mail routed	3899	3899
Total mail delivered	3899	3899
Avg. processor utilization	12.9%	12.4%
Max. processor utilization	22.6%	20.4%
Avg. available memory	93	93
Avg. disk queue length G:	1.276	2.098

Table 5. Impact of RAID level on performance using six drives

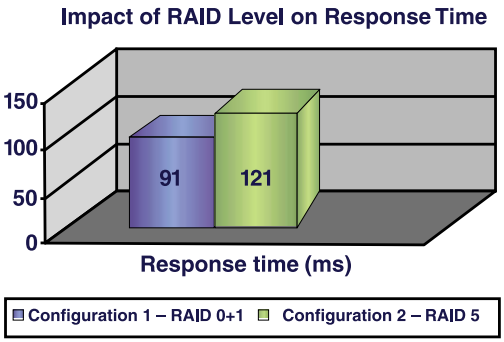

The RAID level used for the Domino data drives of the two configurations resulted in little to no impact on processor utilization rates of the overall system or to the average available memory. The greatest measurable differences are the impact of the RAID level on the response time and the disk queue lengths of the Domino data drive G:. The differences in response time and average disk queue lengths are depicted in the two charts that follow.

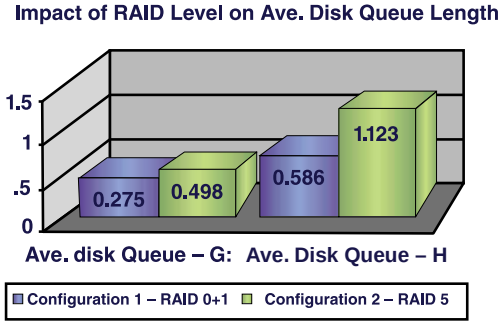

As is evident in the chart above, the RAID 0+1 configuration produced the best client response time experience. The test results consistently confirmed that using the same number of drives in the disk subsystem of a Domino data configuration, a RAID 0+1 configuration would provide the faster user response time. The six-drive RAID 0+1 configuration out performed the six-drive RAID 5 configuration by a rate of 30 ms.

The average disk queue length is an indicator of how responsive the disk subsystem is to the demanded I/O requests being generated by the benchmark. A general rule of thumb is that the queue length should be no more than the total number of drives in the RAID set plus one. Using this guideline, the performance of both configurations is acceptable. However, the RAID 5 configuration has a disk queue length that is 64.4 percent longer than the RAID 0+1 configuration. These differences in disk queue length help contribute to the variations in user response times shown in the previous chart.

Scenario 6
Impact of RAID Level Used for Shared Storage Subsystem – 14 Drives

In the sixth test scenario, two configurations were tested to determine the performance impact of the RAID level used on system resources and test results when running a 500 user simulation. The RAID level is specifically referring to the level used for the Domino data drive in the shared storage subsystem. The *ProLiant* CL380 with 14 total shared storage drives was configured identically except for the RAID level used when configuring drive G:, the Domino data directory for 150 users and drive H:, the mail database directory for the remaining 350 users. The *ProLiant* CL380 configuration used two drives in a RAID 1 mirror for drive E:, serving as the MSCS quorum, and drive F:, storing the Domino transaction logs.
The first test configuration used RAID 0+1 for the next six drives and created drive G: to use as the Domino data directory. The data directory stored 150 mail databases and contained link files for 350 additional mail databases residing on drive H:. RAID 0+1 was used for the remaining six drives to create drive H:, which stored the mail databases for 350 users.
The second test configuration used RAID 5 for the next six drives and created drive G: to use as the Domino data directory. The data directory stored 150 mail databases and contained link files for 350 additional mail databases residing on drive H:. RAID 5 was used for the remaining six drives to create drive H:, which stored the mail databases for 350 users.

The tests were run with MSCS installed using an Active-Passive Domino Cluster implementation. All R5 Mail workload activity was generated against the single Domino Server running on Node 1 of the cluster.

Configuration 1
Systems – 1 Processors, 512 MB RAM, RAID 0+1
Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 14 Shared Storage Drives, 12 Data related drives, two drives for Quorum and Transaction logging:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction log
Drive G:	26019 MB	RAID 0+1	Data: link files, 150 mail files
Drive H:	26019 MB	RAID 0+1	350 mail files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2
Systems – One Processor, 512 MB RAM, RAID 5
Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 14 Shared Storage Drives, 12 Data related drives, 2 drives for Quorum and Transaction logging:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction log
Drive G:	43365 MB	RAID 5	Data: link files, 150 mail files
Drive H:	43365 MB	RAID 5	350 mail files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 RAID 0+1	*Configuration 2 RAID 5*
Transactions per minute	744	774
Response time(ms)	91	121
Invitations sent	738	715
Invitation responses	175	172
Total mail Routed	3899	3899
Total Mail Delivered	3899	3899
Avg. Processor Utilization	12.9%	12.4%
Max. Processor Utilization	22.6%	20.4%
Avg. Available Memory	93	93
Avg. Disk Queue Length G:	1.276	2.098

Table 6. Impact of RAID Level on Performance Using 14 Drives

The RAID level used for the Domino data drives of the two configurations resulted in little impact on processor utilization rates of the overall system or to the average available memory. The greatest measurable differences are the impact of the RAID level on the response time and the disk queue lengths of the Domino data drive G: and Domino mail database drive H:. The differences in response time and average disk queue lengths are depicted in the two charts that follow.

As is evident in the chart above, the RAID 0+1 configuration produced the best client response time experience. The test results consistently confirmed that using the same number of drives in the disk subsystem of a Domino data configuration, a RAID 0+1 configuration would provide the faster user response time. The 14-drive RAID 0+1 configuration out performed the 14-drive RAID 5 configuration by a rate of 32 ms.

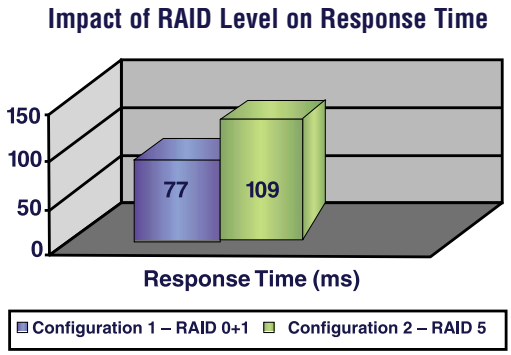

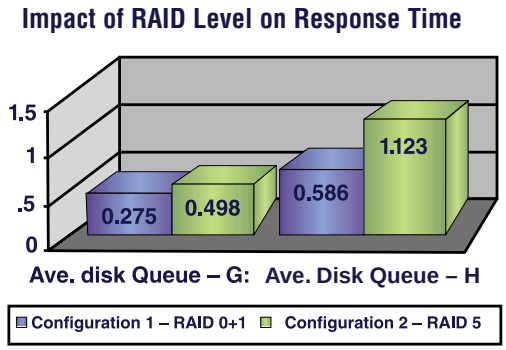

As previously stated, the average disk queue length is an indicator of how responsive the disk subsystem is to the demanded I/O requests being generated by the benchmark. A general rule of thumb is that the queue length should be no more than the total number of drives in the RAID set plus one. Using this guideline, the performance of both configurations is acceptable. However, the RAID 5 configuration has a disk queue length for drive G: that is 81.1 percent longer than the RAID 0+1 configuration. The RAID 5 configuration has a disk queue length for drive H: that is 91.6 percent longer than the RAID 0+1 configuration. These differences in disk queue length help contribute to the variations in user response times shown in the previous chart.

Scenario 7
Impact of Number of Disks in Shared Storage Using RAID 0+1

In the seventh test scenario, two configurations were tested to determine the performance impact on system resources and test results when running a 500 user simulation. This scenario measured the performance impact due to the number of drives in a RAID set used for the Domino data directory. The RAID level is specifically referring to the level used for the Domino data drive in the shared storage subsystem. Configuration 1 is the same test configuration used as Scenario 5 Configuration 1, while configuration 2 is the same test configuration used as Scenario 6 Configuration 1.

Configuration 1

Systems – One Processor, 512 MB RAM, RAID 0+1 Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction Log
Drive G:	17346 MB	RAID 0+1	Data: 500 Mail Files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2

Systems – One Processor, 512 MB RAM, RAID 0+1 Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 14 shared storage drives, 12 data related drives, two drives for quorum and transaction logging

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction log
Drive G:	26019 MB	RAID 0+1	Data: link files, 150 mail files
Drive H:	26019 MB	RAID 0+1	350 mail files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 6 Drives	*Configuration 2 14 Drives*
Transactions per minute	770	769
Response time (ms)	85	77
Total mail routed	3952	4066
Total mail delivered	3952	4066
Avg. processor utilization	13.2%	12.1%
Max. processor utilization	22.7%	21.6%
Avg. available memory	165	173
Avg. disk queue length G:	1.18	.275
Avg. disk queue length H:	NA	.586

Table 7. Impact of the Number of Disks in Shared Storage (Six versus 14) RAID 0+1

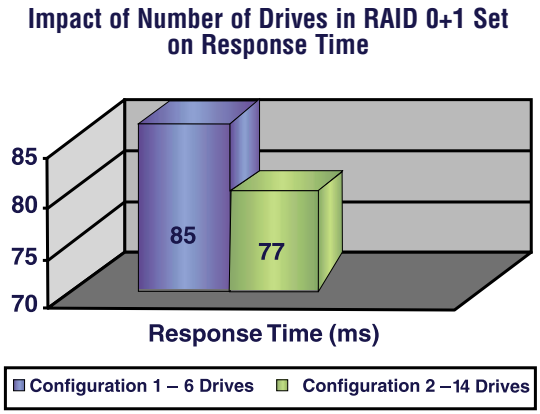

The number of drives used to create the Domino data directory of the two configurations resulted in little to no impact on processor utilization rates of the overall system and a marginal difference in the average available memory. The greatest measurable differences are the impact of the RAID level on the response time and the disk queue lengths of the Domino data drive G:. The differences in response time are depicted in the chart that follows.

The number of drives used for the Domino data directory directly impact the client response time. As is evident in the chart, the 14-drive configuration produced the best client response time experience. The test results consistently confirmed that using more drives in the disk subsystem of a Domino data configuration would provide the faster user response time. The 14-drive configuration out performed the six drive configuration by a rate of 8 ms.

Scenario 8
Impact of Number of Disks in Shared Storage Subsystem Using RAID 5

In the eighth test scenario, two configurations were tested to determine the performance impact on system resources and test results when running a 500 user simulation. This scenario measured the performance impact due to the number of drives in a RAID set used for the Domino data directory. The RAID level specifically refers to the level used for the Domino data drive in the shared storage subsystem. Configuration 1 below uses the same disk configuration as Scenario 5 Configuration 2, while Configuration 2 below uses the same disk configuration as Scenario 6 Configuration 2.

Configuration 1
Systems – One Processor, 384 MB RAM, RAID 5 Domino data drive
Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – Six Shared Storage Drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction log
Drive G:	17346 MB	RAID 0+1	Data: 500 mail files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Configuration 2

Systems – One Processors, 384 MB RAM, RAID 5 Domino data drive

Node 1 – NT 4.0 Enterprise Edition, Domino R5.03 Server, MSCS cluster with Domino resource group, active system under test

Data Directory – 14 Shared Storage drives: 12 Data related drives, two Quorum and Transaction logging drives:

Logical	*Size*	*RAID Level*	*Usage*
Drive E:	100MB	RAID 1	Quorum
Drive F:	8573MB	RAID 1	Transaction log
Drive G:	43365 MB	RAID 5	Data: link files, 150 mail files
Drive H:	43365 MB	RAID 5	350 mail files

Node 2 – NT 4.0 Enterprise Edition, MSCS member node, logged in and sitting idle

Results

	Configuration 1 6 Drives	*Configuration 2 14 Drives*
Transactions per minute	774	775
Response time (ms)	121	127
Total mail routed	3899	3737
Total mail delivered	3899	3737
Invitations sent	715	679
Invitation responses	172	132
Avg. processor utilization	12.4%	11.8%
Max. processor utilization	20.4%	19.3%
Avg. available memory	93	85
Avg. disk queue length G:	2.098	.546
Avg. disk queue length H:	NA	1.218

Table 8. Impact of the Number of Disks in Shared Storage (6 versus 14) RAID 5

The number of drives used to create the Domino data directory of the two configurations resulted in little to no impact on processor utilization rates of the overall system and a marginal difference in the average available memory. The greatest measurable differences are the impact of the RAID level on the response time and the disk queue lengths of the Domino data drive G:. The differences in average disk queue length for drive G: are depicted in the chart that follows.

The increase in the number of drives used for the Domino data directory directly impact the average length of the disk queue for drive G:. As is evident in the chart above, the 14-drive configuration produced the shorter disk queue length due to the I/O efficiency of this configuration. The I/O efficiency improvement can be contributed to additional drive spindles included in the fourteen drive RAID 5 set for drive G:.

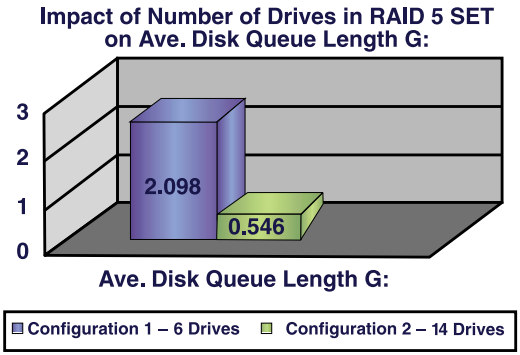

Summary

The objective of this performance document is to provide customers with an indication of how the *ProLiant* CL380 could perform with a 500 real world user load on the Domino R5 mail server. A comparison of a standalone system running Domino R5 was compared to Domino R5 with Microsoft Cluster Server. The test scenarios provided a comparison of using RAID 0+1 and RAID 5 for the Domino data directory. Also examined was the impact of the number of processors and the amount of memory on Domino R5 mail performance. The final testing included collecting performance data for both internal 6-drive and external 14-drive configurations for the *ProLiant* CL380.

To summarize the findings:

- Minimal overhead for Microsoft Cluster Server was observed
- Multiple processors improved client response time
- Memory upgrades resulted in faster response times
- RAID 0+1 provided shorter disk queues and better performance than RAID 5
- Additional drives in a RAID set generally improved overall Domino performance and system resource utilization

Compaq ProLiant CL38O Performance

for Lotus Domino R5
with Microsoft Cluster Server (NT 4.0)

White Paper

Notice

Compaq ProLiant CL380 Performance for Lotus Domino R5 with Microsoft Cluster Server (NT 4.0)
White Paper prepared by High Availability Systems Engineering
First Edition (June 2000)
Document Number 12RD-0600A-WWEN

WHAT'S NEW

- Intel Pentium III 800 MHz processor per server (dual processor capability per server)
- RCC LE 3.0 chipset with 133-MHz front side bus
- Integrated Smart Array Controller
- 64-bit slots for 64-bit expansion card support
- Simplified system board removal for fast and easy servicing

AT A GLANCE

- Two 2P clustered servers plus shared storage in a single cabinet
- Shared storage system supports up to six 1″ wide Ultra2/Ultra3 SCSI 10,000 rpm hot pluggable drives (up to 109.2 GB)
- 256-KB level 2 ECC cache per processor
- 128-MB PC133-MHz registered ECC SDRAM memory, expandable to 4 GB (per server)
- Four expansion slots, three 64-bit PCI, one 32-bit PCI (per server)
- Compaq NC3163 Fast Ethernet NIC (embedded) PCI 10/100 WOL for heartbeat monitoring
- Compaq NC3123 Fast Ethernet NIC PCI10/100 Mbps WOL full-duplex Ethernet for public LAN (occupies one PCI slot per server)
- Integrated dual channel wide Ultra2 SCSI adapter – utilized for server boot (per server)
- Compaq 64-Bit Dual Channel Wide Ultra2 SCSI Adapter – utilized for shared storage (occupies PCI slot)
- Two 1″ wide Ultra 2/Ultra3 SCSI hot plug drive bays (per server)
- Rack or tower form factor – 10U
- Removable servers, hot plug disks, power supplies and RAID controllers for easy serviceability and toolless access to major components
- Integrated Remote Console (IRC)
- Automatic Server Recovery-2 (ASR-2)
- Protected by Compaq Services, including a three-year, on-site limited warranty and extended pre-failure warranty which covers processors, memory and disk drives on the server. Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

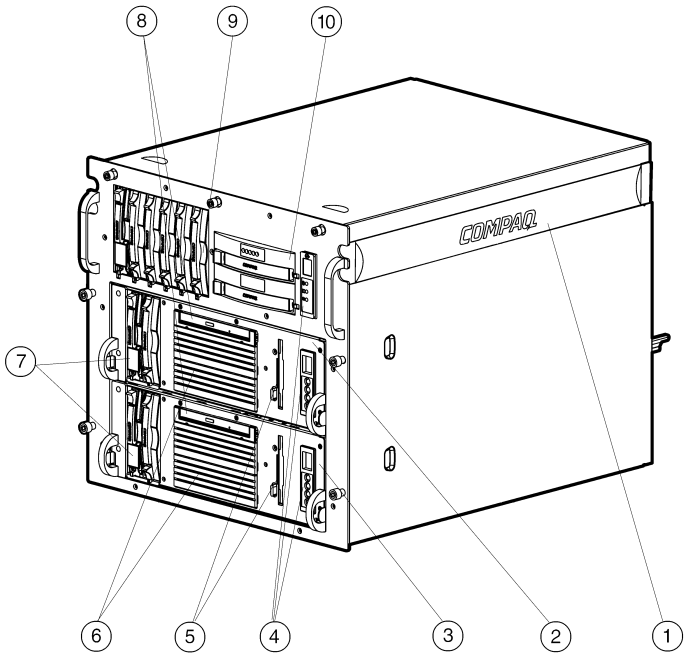

1 Side Rails
2 Server 1
3 Server 2
4 Protected power switches with front panel status LEDs
5 1.44MB disk drive
6 One full height or two half height removable media bays
7 4 x 1″ wide Ultra2/Ultra3 SCSI hot plug drive bays
8 Low profile high speed IDE CD ROM
9 6 x 1″ wide Ultra2/Ultra3 SCSI hot plug drive bays for shared storage
10 Shared storage RAID controllers (CR3500)

Processor (per server)
Intel Pentium III processors 800 MHz
Cache Memory (per server)
256-KB level 2 writeback cache per processor
Upgradability (per server)
Upgradable to dual processing
Memory (per server)
Standard 128 MB (PC133MHz
Registered ECC SDRAM memory) Maximum 4 GB

Network Controller (per server)
Compaq NC3163 Fast Ethernet NIC (embedded)
PCI 10/100 WOL for heartbeat monitoring
Compaq NC3123 Fast Ethernet NIC PCI 10/100 Mbps
WOL full-duplex Ethernet for public LAN
(occupies one PCI slot per server)
Expansion slots (per server)

I/O (Total)	4, 2 available
64-bit PCI	3, 2 available
32-bit PCI	1, 0 available

Storage Adapters (per server)
64-Bit Dual Channel Wide Ultra2 Adapter occupying a PCI slot (interface to shared storage system)
Integrated Dual Channel Wide Ultra2 SCSI adapter (utilized for server boot)

Storage Diskette Drive (per server)	1.44 MB
CD-ROM High Speed IDE CD-ROM Drive (per server)	None
Hard Drives Maximum Shared Internal Storage (6 x 18.2 GB 1″ Ultra2/Ultra3 drives)	109.2 GB
Maximum Non-Shared Internal Storage (internal drive cage with optional 2 x 1″ Ultra2/Ultra3 drive cage, per server)	72.8 GB

Shared Storage
(total for cluster server system)
Shared Storage RAID CR3500 controller. One ships standard. Up to two controllers in the shared storage. 9.1-GB or 18.2-GB pluggable wide Ultra2/Ultra3 SCSI 1″ drives. None ship standard. Up to six drives in the cluster server cabinet.

Interfaces (per server)

Parallel	1
Serial	2
Pointing Device (mouse)	1
Graphics	1
Hot Plug Keyboard	1
External SCSI (for tape only)	1
Network RJ-45	1

Graphics (per server)
Integrated ATI Rage IIC Video Controller with 4MB Video Memory
Form Factor (total cluster)
Tower or Rack (10U), (17.65-inch)

Compaq Software
Compaq *Insight Manager*
The comprehensive system management tool that monitors the operation of Compaq servers, workstations and clients
Compaq *SmartStart*
The server configuration and software integration tool from Compaq
Array Configuration Utility (ACU)
with Smart Array Controllers
Same easy-to-use yet powerful Graphical User Interface (GUI) configuration utility and easy-to-navigate wizards for all Smart Array Controllers
CR3500 Configuration Utility
An easy-to-use yet powerful GUI configuration utility for the CL380 shared storage RAID controller(s)

MODELS
Compaq *ProLiant* CL380 6/800 Models

Available in multiple configurations

Intelligent Manageability

Automatic Server Recovery-2 (ASR-2)
Compaq Remote Insight Board (PCI) (optional)
Compaq Remote Insight Light-Out Edition (optional)
Compaq Insight Manager
Integrated Remote Console (IRC)
Hardware Remote Console
Full Remote Server Reboot Capability
Reset and Failure Sequence Replay
Host Hardware Independent
Operating System Independent
Communication is out-of-band
Integrated Management Log
Server Health Logging
Revision History Table
Automatic Revision Tracking
Drive Parameter Tracking (with Smart Array Controllers)
Dynamic Sector Repairing (with Smart Array Controllers)
Off-Line Backup Processor capability
Pre-Failure Warranty
(covers processors, server hard drives, and memory)

Security

Power-on password
Keyboard password
Diskette drive control
Diskette boot control
QuickLock, Network Server Mode
Parallel and serial interface control
Administrator's password
Disk configuration lock

Power Supply

225 Watt, CE Mark Compliant

OS Support

Microsoft Windows NT
Microsoft Windows 2000
Novell NetWare
SCO UnixWare

Service and Support

Compaq Services provides a wide range of service, including a three-year, limited warranty fully supported by a worldwide network of resellers and service providers, lifetime toll-free 7 x 24 hardware technical phone support, and pre-failure warranty coverage of hard drives on the server, memory and processors. Disaster recovery services are available through partnership with Comdisco.

Note: Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

Models (per server)

6/800 Model 1 Tower 164607-xx1

Processor(s)
Intel Pentium III processors 800 MHz

Network Controller
NC3163 Fast Ethernet NIC (embedded) PCI WOL 10/100 and NC3123 Fast Ethernet NIC PCI 10/100 Mbps WOL

Cache Memory
256-KB level 2 writeback cache per processor

Storage Adapter
64-bit Dual Channel Wide Ultra2 Adapter
and Integrated Dual Channel Wide Ultra2 SCSI Adapter

Memory
PC133-MHz registered ECC SDRAM
128MB standard, 4GB maximum

Internal Storage
Nonshared= 72.8GB per server
Shared= 109.8GB per server

Optical Drive
High Speed IDE CD-ROM Drive

Form Factor
Tower

6/800 Model 1 Rack 164608-xx1

Processor(s)

Intel Pentium III processors 800 MHz

Network Controller

NC3163 Fast Ethernet NIC (embedded)
PCI WOL 10/100 and NC3123
Fast Ethernet NIC PCI 10/100 Mbps WOL

Cache Memory

256-KB level 2 writeback cache per processor
Storage Adapter 64-bit Dual Channel Wide Ultra2 Adapter and Integrated Dual Channel Wide Ultra2 SCSI Adapter

Memory

PC133-MHz registered ECC SDRAM
128MB standard, 4GB maximum

Internal Storage

Nonshared= 72.8GB per server
Shared= 109.8GB per server

Optical Drive

High speed IDE CD-ROM drive

Form Factor

Rack (10U), (17.65-inch)

Country Code Key

XX=00	US
XX=29	Japan
XX=37	APD
XX=42	EURO
XX=AA	PRC

Processor

P800 Pentium III 256K Processor Option Kit
161084-B21

Memory (DIMMs)

1024-MB PC133MHz Registered ECC SDRAM DIMM
Memory Option Kit 128280-B21
512-MB PC133MHz Registered ECC SDRAM DIMM
Memory Option Kit 128279-B21
256-MB PC133MHz Registered ECC SDRAM DIMM
Memory Option Kit 128278-B21
128-MB PC133MHz Registered ECC SDRAM DIMM
Memory Option Kit 128277-B21
64-MB PC133MHz Registered ECC SDRAM DIMM
Memory Option Kit 128276-B21

Optical Drives

DVD-ROM Drive Option Kit (Opal)
(The DVD-ROM Option Kit cannot be used in conjunction with the standard low-profile IDE CD-ROM drive. DVD-ROM Drive Option Kit only supported by Windows NT or Windows 2000 OS.) 388481-B21

Internal Storage

Wide Ultra2/Ultra3 SCSI Hard Drive Option Kit (2 x 1″) (Opal)
(The *ProLiant* CL380 Server nodes can support up to four 1″ Ultra2/Ultra3 hot plug drives on a single SCSI channel with the purchase of the optional Wide Ultra2/Ultra3 SCSI Hard Drive Cage Option Kit [2 x 1″]) 175254-B21

Hard Drives

Wide Ultra3 SCSI – Universal Hot Plug

18.2-GB Wide Ultra3 SCSI 10,000 rpm drive (1″)
142673-B22
9.1-GB Wide Ultra3 SCSI 10,000 rpm drive (1″)
142671-B22

Note: Wide Ultra3 drives require an additional wide Ultra3 SCSI adapter or controller to support wide Ultra3 speeds.

Wide Ultra2 SCSI – Universal Hot Plug

18.2-GB Wide Ultra2 SCSI 10,000 rpm drive (1″)
128418-B22
9.1-GB Wide Ultra2 SCSI 10,000 rpm drive (1″)
328939-B22

Wide Ultra2 SCSI – Non-Hot Plug

18.2-GB Wide Ultra2 SCSI drive (1″)	388143-B21
9.1-GB Wide Ultra2 SCSI drive (1″)	349526-B21

Note: Non-hot plug wide Ultra2 drives are supported only in the removable media bays.

Note: Non-hot plug wide Ultra2 drives require an additional controller.

Storage Controllers

Shared Storage RAID CR3500 Controller (The Shared Storage RAID CR3500 Controller is supported only in the Shared Storage System.)		**388332-B21**
Integrated Smart Array Controller Module		**128293-B21**
Smart Array 431 Controller		**127695-B21**
Smart Array 431 Controller (Japan)		**127695-291**
Smart Array 221 Controller		**388099-B21**
Smart Array 221 Controller (Japan)		**388099-291**
Smart Array 3200 Controller		**295643-B21**
Smart Array 3200 Controller (Japan)		**295643-291**
Smart Array 4200 Controller		**295636-B21**
Smart Array 4200 Controller (Japan)		**295636-291**
64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter		**154457-B21**
64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter (Japan)		**154457-291**
64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter		**129803-B21**
64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter (Japan)		**129803-B21**
Wide-Ultra SCSI-3 SCSI Adapter		**272514-001**

Communications

Compaq NC3123 Fast Ethernet NIC PCI 10/100 WOL		**174830-B21**
Compaq NC3122 Fast Ethernet NIC PCI Dual 10/100		**317450-B21**
Compaq NC3134 Fast Ethernet NIC 64 PCI Dual Port 10/100		**138603-B21**
Compaq NC3135 Fast Ethernet Module Dual 10/100 Upgrade Module for	**NC3134 and NC3131**	**138604-B21**
Compaq NC6132 Gigabit Module 1000 SX Upgrade Module for	**NC3134 and NC3131**	**338456-B23**
Compaq NC6133 Gigabit Module 1000 LX Upgrade Module for	**NC3134 and NC3131**	**338456-B24**
Compaq NC3133 Fast Ethernet Module 100 FX Upgrade Module for	**NC3134 and NC3131**	**338456-B25**

Compaq NC6134 Gigabit NIC 64 PCI 1000SX	174818-B21
Compaq NC4621 Token Ring NIC PCI 4/16 WOL	379956-B21
IBM 16/4 Token Ring PCI Special	133749-B21

Internal Tape Drives

Internal 20/40-GB DAT DDS-4 Tape Drive	157769-B21
Internal 12/24-GB DAT Drive (Opal)	295513-B22
Internal 12/24-GB DAT Drive (Opal) (Japan)	295513-292
Internal 4/8-GB DAT Drive (Opal)	295353-B22
Internal 4/8-GB DAT Drive (Opal) (Japan)	295353-292
Internal 35/70-GB DLT Drive	242520-B21
Internal 35/70-GB DLT Drive (Japan)	242520-291
Internal 20/40-GB DLT Drive (Opal)	340743-B22
Internal 20/40-GB DLT Drive (Opal) (Japan)	340743-292
20/40-GB DAT 8 Cassette Autoloader – Internal	166504-B21
Internal AIT 35-GB Tape Drive	388504-B21
Internal AIT 50-GB Tape Drive	157766-B21
4-GB DDS-2 DAT Cassette	137611-001
DDS-4 Media (10 pack)	152842-001
12/24-GB DDS-3 DAT eight Cassette Pre-Filled Magazine (seven data cassettes, one cleaning cassette)	134117-B21
20/40-GB DAT DDS-4 eight Cassette Magazine with Media	171532-B21
12-GB DDS-3 DAT Cassette (10 Pack)	295515-B21
DAT Cleaning Cassette	242781-001
AIT 35-GB Data Cassette (five pack)	402371-B21
AIT 50-GB Data Cassette (five pack)	152841-001
AIT Cleaning Cassette	402374-B21
DLT IV Drive Cartridge (single)	295194-B21
DLT IV Drive Cartridge (seven pack)	295192-B21
DLT Cleaning Cartridge	199704-001

Redundant Options

Compaq Hot Plug Redundant Power Supply (48 volt DC) **328065-B21**

(The* ProLiant *CL380 comes standard with a 110 volt AC power supply. Each 48 volt DC option kit [328065-B21] contains one power supply. Therefore, to convert to DC power supplies, four option kits must be purchased.)

Management Options

Compaq *Remote Insight* Lights-Out Edition (NA, LA) **157866-001**

Note: The Compaq Remote Insight Lights-Out Edition must be installed in a 32-bit PCI slot in order to function correctly.

Compaq *Remote Insight* Lights-Out Edition (EURO) **157866-021**

Note: The Compaq Remote Insight Lights-Out Edition must be installed in a 32-bit PCI slot in order to function correctly.

Compaq *Remote Insight* Lights-Out Edition (APD, Japan) **157866-371**

Note: The Compaq Remote Insight Lights-Out Edition must be installed in a 32-bit PCI slot in order to function correctly.

Compaq *Remote Insight* (PCI) LAN + NA Modem **294013-002**

Compaq *Remote Insight* Board (PCI) LAN/Serial without modem **294013-B24**

Compaq *Remote Insight* Board (PCI) LAN only **294013-B25**

Monitors

TFT8020 18″ Flat Panel Monitor (Opal)	326100-001
TFT8020 18″ Flat Panel Monitor (Carbon)	159244-001
TFT8020 18″ Flat Panel Monitor (Opal, bulk pack, two monitor heads per box)	173610-001
TFT5010 15″ Flat Panel Monitor (Opal)	104741-001
TFT5010 15″ Flat Panel Monitor (Carbon)	154486-001
TFT5000R Rackmount Flat Panel	120207-001
TFT5000R Rackmount Flat Panel (Int'l)	120207-B31
TFT5000R Rackmount Flat Panel (Japan)	120207-291
P1610 Color	305708-001
P1100 Color	325606-001
P900 Color	386326-001
P700 Color	380207-001
V1100 Color Monitor	386472-001
V700 Color (Opal)	325800-001
V500 Color	325900-001
S510 Color	168636-001
S710 Color	152991-001
S910	138485-001

Note: Monitors larger than 17″ may be too heavy for use in rack systems.

External Tape Drives

External 20/40-GB DAT DDS-4 Drive	157770-001
External 20/40-GB DAT DDS-4 Drive (Int'l)	157770-B31
External 20/40-GB DAT DDS-4 Drive (Japan)	157770-291
External 35/70-GB DLT Drive	242521-B21
External 35/70-GB DLT Drive (Japan)	242521-291
External 20/40-GB DLT Drive (Opal) 3	40744-B22
External 20/40-GB DLT Drive (Opal) (Japan)	340744-292
20/40-GB DAT 8 Cassette Autoloader – External	166505-001
20/40-GB DAT 8 Cassette Autoloader – External (Int'l)	166505-B31
20/40-GB DAT 8 Cassette Autoloader – External (Japan)	166505-291
External AIT 35-GB Tape Drive	388507-001
External AIT 35-GB Tape Drive (Int'l)	388507-B31
External AIT 35-GB Tape Drive (Japan)	388507-291
External AIT 50-GB Tape Drive	157767-001
External AIT 50-GB Tape Drive (Int'l)	157767-B31
External AIT 50-GB Tape Drive (Japan)	157767-29
4-GB DDS-2 DAT Cassette	137611-001
DDS-4 Media (10 pack)	152842-001
12/24-GB DDS-3 DAT eight Cassette Pre-Filled Magazine (seven data cassettes, one cleaning cassette)	134117-B21
20/40-GB DAT eight Cassette Magazine with Media	171532-B21
12-GB DDS-3 DAT Cassette (10 Pack)	295515-B21
DAT Cleaning Cassette	242781-001
AIT 35-GB Data Cassette (five pack)	402371-B21
AIT 50-GB Data Cassette (five pack)	152841-001
AIT Cleaning Cassette	402374-B21
DLT IV Drive Cartridge (single)	295194-B21
DLT IV Drive Cartridge (seven pack)	295192-B21
DLT Cleaning Cartridge	199704-001

Shared External Storage

Compaq *StorageWorks* Enclosure Model 4314T (tower) (U.S.)	190210-001
Compaq *StorageWorks* Enclosure Model 4314T (tower) (Japan)	190210-291
Compaq *StorageWorks* Enclosure Model 4314T (tower) (Int'l)	190210-B31
Compaq *StorageWorks* Enclosure Model 4214T (tower) (U.S.)	124875-001
Compaq *StorageWorks* Enclosure Model 4214T (tower) (Japan)	124875-291
Compaq *StorageWorks* Enclosure Model 4214T (tower) (Int'l)	124875-B31
Compaq *StorageWorks* Enclosure Model 4314R (rack-mountable) (U.S.)	190209-001
Compaq *StorageWorks* Enclosure Model 4314R (rack-mountable) (Japan)	190209-291
Compaq *StorageWorks* Enclosure Model 4314R (rack-mountable) (Int'l)	190209-B31
Compaq *StorageWorks* Enclosure Model 4214R (rack-mountable) (U.S.)	103381-001
Compaq *StorageWorks* Enclosure Model 4214R (rack-mountable) (Japan)	103381-291
Compaq *StorageWorks* Enclosure Model 4214R (rack-mountable) (Int'l)	103381-B31
Compaq *StorageWorks* Enclosure Model 4354R (rack-mountable)	190211-001
Compaq *StorageWorks* Enclosure Model 4354R (rack-mountable) (Int'l)	190211-B31
Compaq *StorageWorks* Enclosure Model 4354R (rack-mountable) (Japan)	190211-291
Compaq *StorageWorks* Enclosure Model 4254R (rack-mountable)	138151-001
Compaq *StorageWorks* Enclosure Model 4254R (rack-mountable) (Int'l)	138151-B31
Compaq *StorageWorks* Enclosure Model 4254R (rack-mountable) (Japan)	138151-291

Note: The Compaq StorageWorks enclosure 4200 and 4300 family only supports the wide Ultra2/Ultra3 1" hot pluggable hard drives.

StorageWorks Enclosure 4200 Redundant Power Supply Option	119826-B21
StorageWorks Enclosure 4200 Dual Bus Module Option	119829-B21
StorageWorks Enclosure 4200 Dual Bus Module Option (Japan)	119829-291
StorageWorks Enclosure 4200 Tower-to-Rack Conversion Kit	105213-B21
StorageWorks Enclosure 4200 Ultra3 Single Bus Module Option	190212-B21
StorageWorks Enclosure 4200 Ultra3 Dual Bus Module Option	190213-B21
6-ft Wide to VHDCI Cable	313375-001
12-ft Wide to VHDCI Cable	313375-002
39-ft (12m) VHDCI to VHDCI SCSI Cable	105214-B21

Non-Shared External Storage

SCSI Storage Expander II	332609-001
SCSI Storage Expander II (Int'l)	332609-002

Rack-Mountable SCSI Storage Expander II	332607-001
Rack-Mountable SCSI Storage Expander II (Int'l)	332607-B31
ProLiant Storage System Tower Model UE	348710-002
ProLiant Storage System Tower Model UE (Int'l)	348710-B31
ProLiant Storage System Tower Model UE (Japan)	348710-291
ProLiant Storage System Tower Model U2	304110-B22
ProLiant Storage System Tower Model U2 (Int'l)	304114-B32
ProLiant Storage System Tower Model U2 (Japan)	304114-292
Rack-Mountable *ProLiant* Storage System Model UE	348700-001
Rack-Mountable *ProLiant* Storage System Model UE (Int'l)	348700-B31
Rack-Mountable *ProLiant* Storage System Model UE (Japan)	348700-291
Rack-Mountable *ProLiant* Storage System Model U2	304100-B22
Rack-Mountable *ProLiant* Storage System Model U2 (Int'l)	304104-B32
Rack-Mountable *ProLiant* Storage System Model U2 (Japan)	304104-292

Note: The Compaq ProLiant storage system family supports only the wide-ultra SCSI-3 1" or 1.6" hot pluggable hard drives.

ProLiant Storage System Model F1 or U1 to Model UE Upgrade Kit	348718-B21
ProLiant Storage System Model F1 or U1 to Model UE Upgrade Kit (Japan)	348718-291
ProLiant Storage System Model F2 or U2 to Model UE Upgrade Kit	348706-B21
ProLiant Storage System Model F2 or U2 to Model UE Upgrade Kit (Japan)	348706-291
ProLiant Storage System/U Redundant Power Supply	224206-001
ProLiant Storage System/U Redundant Power Supply (Int'l)	224214-001
Conversion Kit, SB to DB/U	304116-B21
Conversion Kit, DB to SB/U	304115-B21

Tape Arrays and Libraries

TA1000 AIT Tape Array, Rackmount (U.S.)	175005-001
TA1000 AIT Tape Array, Rackmount (Int'l)	175005-B31
TA1000 AIT Tape Array, Rackmount (Japan)	175005-291
DLT Tape Array Model 0	199860-001
DLT Tape Array II Model 35/70-4 no software	295165-001
DLT Tape Array II Model 35/70-4 no software (Int'l)	295165-B31
DLT Library 15 Cartridge Model 35/70-1	350365-001
DLT Library 15 Cartridge Model 35/70-2 (Int'l)	350365-002
DLT Library 15 Cartridge Model 35/70-1 (Japan)	350365-291

DLT Library 15 Cartridge Model 35/70-2 (Japan)	350365-292
DLT Library 15 Cartridge Model 20/40-1	350362-001
DLT Library 15 Cartridge Model 20/40-2	350362-002
DLT Library 15 Cartridge Model 20/40-1 (Int'l)	350362-B31
DLT Library 15 Cartridge Model 20/40-1 (Japan)	350362-291
DLT Library 15 Cartridge Model 20/40-2 (Int'l)	350362-B32
35/70 DLT Library Ready Drive	350367-B21
35/70 DLT Library Ready Drive (Japan)	350367-291
20/40 DLT Library Ready Drive	350364-B21
20/40 DLT Library Ready Drive (Japan)	350364-291
DLT Library Tabletop Conversion Kit	350384-B21
SSL2020 AIT Mini-Library 1 drive, 20 slot, Table Top	175195-B21
SSL2020 AIT Mini-Library 2 drive, 20 slot, Table Top	175195-B22
SSL2020 AIT Mini-Library 1 drive, 20 slot, Rackmount	175196-B21
SSL2020 AIT Mini-Library 2 drive, 20 slot, Rackmount	175196-B22
TL891 DLT library unit, 1 drive Table Top	120875-B21
TL891 DLT library unit, 2 drive Table Top	120875-B22
TL881 DLT Library, 1 drive Table Top, FWD	128667-B21
TL881 DLT Library, 2 drive Table Top, FWD	128667-B22
TL881 DLT Library, 1 drive Table Top, FWSE	128668-B21
TL881 DLT Library, 2 drive Table Top, FWSE	128668-B22
TL891 DLT library unit, 1 drive Rack	120876-B21
TL891 DLT library unit, 2 drive Rack	120876-B22
TL881 DLT Library, 1 drive Rack	128669-B21
TL881 DLT Library, 2 drive Rack	128669-B22
MiniLibrary Data Unit	128670-B21
10 Slot Magazine (for TL891)	120881-B21
Mini-Library Expansion Unit	120877-B21
Add-On 35/70GB Drive for TL891	120878-B31
Add-On 20/40GB Drive for TL881	128671-B21
DLT IV Data Cartridges (single pack)	295194-B21
DLT IV Data Cartridges (seven pack)	295192-B21
DLT Cleaning Cartridge	199704-001

DLT Tape Array Cable Kit	242427-001
DLT Magazine with five Type IV DLT Tape	295168-B21
10 Slot Magazine	120881-B21
Add-On 35/70-GB Drive for TL891	120878-B31
Add-On 20/40-GB Drive for TL881	128671-B21
MiniLibrary pass-through Mechanism Kit	120880-B21
MiniLibrary Table Top to Rack Mount Conversion Kit, Metric/RETMA racks	120882-B21
MiniLibrary Table Top to Rack Mount Conversion Kit, Compaq racks	120883-B21
SSL2020 AIT Library Pass Thru with Transport	175312-B21
SSL2020 AIT Library Pass Thru Extender	175312-B22
AIT 50GB Drive Add-On LVD Drive for SSL2020 AIT Library	175197-B21
19 Slot Magazine for SSL2020 AIT Library	175198-B21
Tabletop to Rackmount Upgrade Kit for SSL2020 AIT Library, Compaq Rack	175199-B21
Tabletop to Rackmount Upgrade Kit for SSL2020 AIT Library, RETMA racks	175199-B21
AIT 50-GB Data Cassette (five pack)	152841-001
AIT Cleaning Cassette	402374-B21
DLT7000 drive upgrade kit	349351-B21

Note: The DLT 15 Cartridge Library Model 20/40 or Model 35/70 can be mounted in a Compaq rack. To use any model of the library on a tabletop, the customer would purchase one of the Library Models listed above and would also order the Tabletop Conversion Kit (350384-B21) to provide metal top and side covers. Refer to the Compaq DLT Libraries QuickSpec for details.
DLT Library Model 20/40 or 3570 includes one 20/40 DLT data cartridge or one 35/70 DLT data cartridge and three empty five-cartridge magazines. US 350362-001/-002 (1 drive/2 drive) includes U.S. power cord, US 350365-001/-002 (1 drive/2 drive) includes U.S. power cord, Int'l -B31/B32 (1 drive/2 drive) CTO power cord for country ordering the unit.

Note: The TL891 library unit (1 & 2 drive) or the TL881 library unit (1 & 2 drive) require country kits (349386-001, 349386-011, 349386-AR1, 349386-021, 349386-031, 349386-111, 349386-061, 349386-081or 349386-291), SCSI Cable kits (328215-001, 328215-002, 328215-003 or 328215-004) and ProLiant Connectivity Kits (113395-001 or 113396-001). Please refer to the TL891/881 Automated DLT Tape Library Solution QuickSpecs for further details.

Uninterruptible Power Systems

UPS Model T2000 (100-127 volt) (U.S.)	242688-005
UPS Model T2000j (100-127 volt) (Japan)	242688-295
UPS Model T2400h (200-240 volt) (Int'l)	242688-006
UPS Model T2400h-NA (200-240 volt) (NA)	242688-007
Rack-Mountable UPS R3000 (100-127 volt)	242705-001
Rack-Mountable UPS R3000j (100-127 volt) (Japan)	242705-291
Rack-Mountable UPS R3000h (208-240 volt) (Int'l)	242705-002
Rack-Mountable UPS R3000h-NA (208-240 volt) NDC (non-detached cord)	242705-003
Rack-Mountable UPS R6000 (208 VAC default) (NA)	347207-001
Rack-Mountable UPS R6000 (230 VAC default) (Int'l)	347207-B31
Rack-Mountable UPS R6000 (200 VAC default) (Japan)	347207-291
R6000 Extended Runtime Module	347224-B21
SNMP-EN Adapter	347225-B21
High to Low Voltage Transformer (250VA)	388643-B21
Multi-Server UPS Card	123508-B21
Scalable UPS Card	123509-B21
Rack and Tower Conversion Kits	
ProLiant CL380 Rack to Tower Conversion Kit	388333-B21
ProLiant CL380 Tower to Rack Conversion Kit	388334-B21

Please see the Compaq Rack Builder for configuration assistance at

www.compaq.com/products/storageworks/rack_accessories/rackbuilderpro2.html
or download it at www.compaq.com/support/files/server/us/download/6956.html

Compaq Rack 9000 Series (Opal)

Compaq Rack 9142 (42U) – Pallet	120663-B21
Compaq Rack 9142 (42U) – Shock Pallet	120663-B22
Compaq Rack 9142 (42U) – Crated	120663-B23
Compaq Rack 9136 (36U) – Pallet	120664-B21
Compaq Rack 9136 (36U) – Shock Pallet	120664-B22
Compaq Rack 9136 (36U) – Crated	120664-B23

Compaq Rack 9122 (22U) – Pallet	120665-B21
Compaq Rack 9122 (22U) – Shock Pallet	120665-B22
Compaq Rack 9122 (22U) – Crated	120665-B23

Note: – B21 part numbers are for all Geos
– B22 part numbers are used to ship configured racks (by custom systems, VARs and Channels)
– B23 part numbers are used for air shipments

Rack Options for Compaq Rack 9000 Series

Baying Kit	120669-B21
9142 Side Panel Kit	120670-B21
9136 Side Panel Kit	120671-B21
Ballast Option Kit	120672-B21
Stabilizer Option Kit	120673-B21
Rack Rail Adapter Kit (25″ depth)	120675-B21
Fan (110 VAC) Roof Mount	120677-B21
Fan (220 VAC) Roof Mount	120678-B21
Grounding Rack Option Kit	120682-B21
Sliding Shelf Kit	332564-B21

Rack Options for Compaq Rack 9000 Series

Monitor Utility Shelf	303606-B21
Rack Blanking Panel Kit (Beach Gray) (U.S.)	169940-001
Rack Blanking Panel Kit (Opal) (U.S.)	169940-B21
Compaq Networking Cable Management Kit	292407-B21
Compaq Networking Recessed Rail Management Kit	292406-B21
High Capacity Cable Management Arm Kit	124711-B21
Compaq Power Distribution Unit-low voltage	295363-001

Compaq Power Distribution Unit-high voltage (NA)	295363-002
Compaq Power Distribution Unit – 30A (NA only)	295363-003
Compaq Power Distribution Unit – 16A (NA, Int'l)	295363-B21
Compaq Power Distribution Unit-high voltage PDUh-NDC (non-detached cord) (Int'l)	295363-B32
Server Console Switch 1 x 2 port (100-230 VAC)	120206-001
Server Console Switch 1 x 2 port (100-230 VAC) (Int'l)	120206-B31
Server Console Switch 1 x 2 port (100-230 VAC) (Japan)	120206-291
Server Console Switch 1 x 4 port (100-230 VAC)	400336-001
Server Console Switch 1 x 4 port (100-230 VAC) (Int'l)	400336-B31
Server Console Switch 1 x 4 port (100-230 VAC) (Japan)	400336-291
Server Console Switch 1 x 8 port (100-230 VAC)	400337-001
Server Console Switch 1 x 8 port (100-230 VAC) (Int'l)	400337-B31
Server Console Switch 1 x 8 port (100-230 VAC) (Japan)	400337-291
Server Console Switch 2 x 8 port (100-230 VAC)	400338-001
Server Console Switch 2 x 8 port (100-230 VAC) (Int'l)	400338-B31
Server Console Switch 2 x 8 port (100-230 VAC) (Japan)	400338-291
Server Console Switch 2 x 8 port (48VDC)	400542-B21
CPU to Server Console Cable, 12′	110936-B21
CPU to Server Console Cable, 20′	110936-B22
CPU to Server Console Cable, 40′	110936-B23
CPU to Server Console Cable, 3′	110936-B24
CPU to Server Console Cable, 6′	110936-B25
CPU to Server Console Cable (Plenum Rated) 20′	149363-B21
CPU to Server Console Cable (Plenum Rated) 40′	149364-B21
Switch Box Connector Kit (115V)	144007-001
Switch Box Connector Kit (230V)	144007-002
Switch Box Connector Kit (high voltage) (Int'l)	144007-B33
Rack Internal Trackball Keyboard (Opal) (UK)	185152-036
Rack Internal Trackball Keyboard (Opal) (Germany)	185152-046

Rack Internal Trackball Keyboard (Opal) (France)	185152-056
Rack Internal Trackball Keyboard (Opal) (Italy)	185152-066
Rack Internal Trackball Keyboard (Opal) (Spain)	185152-076
Rack Internal Trackball Keyboard (Opal) (Denmark)	185152-086
Rack Internal Trackball Keyboard (Opal) (Norway)	185152-096
Rack Internal Trackball Keyboard (Opal) (Sweden/Finland)	185152-106
Rack Internal Trackball Keyboard (Opal) (Switzerland)	185152-116
Rack Internal Trackball Keyboard (Opal) (Portugal)	185152-136
Rack Internal Trackball Keyboard (Opal) (Belgium)	185152-186
Rack Internal Trackball Keyboard (Opal) (Japan)	185152-296
Rack Internal Trackball Keyboard (Opal) (U.S.)	185152-406
Rack Internal Trackball Keyboard (Opal) (Int'l)	185152-B36
1U Rack Keyboard Drawer	338056-B21
Keyboard/Monitor/Mouse extension cables	169989-001

Rack Options for Third Party Rack Cabinets

Depth Adjustable Fixed Rails	332558-B21

Other

Internal Wide Ultra3 Cable Option Kit, HP	166390-B21
Compaq Intelligent Cluster Administrator V1.5	122284-B22
Compaq Intelligent Cluster Administrator V1.5 (Japan)	122284-292
Compaq Intelligent Cluster Administrator V2.0	168843-B21

Service Upgrades (*CarePaqs*)

Hardware Services 4-Hour On-site Service

4-Hour On-site Service, 9-Hour x 5-Day Coverage, 3 Years	(2-5-2 Part Number for U.S.)	FM-CW4HR-36
4-Hour On-site Service, 9-Hour x 5-Day Coverage, 3 Years	(2-5-2 Part Number for Canada)	FP-CW4HR-36
4-Hour On-site Service, 24-Hour x 7-Day Coverage, 3 Years	(2-5-2 Part Number for U.S.)	FM-CW724-36
4-Hour On-site Service, 24-Hour x 7-Day Coverage, 3 Years	(2-5-2 Part Number for Canada)	FP-CW724-36

Installation & Start-up Services

Install NT Server 4.0 and Microsoft Cluster Server on 2 servers	(2-5-2 Part Number for U.S.)	FM-WNTBM-IN
Install Microsoft Cluster Server on 2 servers	(2-5-2 Part Number for U.S.)	FM-WNTBN-IN
Hardware Installation	(2-5-2 Part Number for U.S.)	FM-CWINS-IN
Hardware Installation	(2-5-2 Part Number for Canada)	FP-CWINS-IN
Hardware Installation	(6-3 Part Number for EMEA)	Quote
O/S & Hardware Installation	(6-3 Part Number for U.S.)	
(Hardware installation plus two days O/S start-up activity [installation, configuration, network access and orientation.])		Quote
O/S & Hardware Installation (6-3 Part Number for EMEA)		
(Hardware installation plus two days O/S start-up activity [installation, configuration, network access and orientation.])		Quote

Business Solutions Priority Service Plan – Priority Level

5 x 9 4-hour H/W, 5 x 9 two-hour O/S (NT and NetWare)	Three Years	(2-5-2 Part Number for U.S.)	FP-CX106-36
5 x 9 4-hour H/W, 5 x 9 two-hour O/S (NT and NetWare)	Three Years	(2-5-2 Part Number for Canada)	FP-CZ106-3

Business Solutions Priority Service Plan – Priority 24 Level

7 x 24 4-hour H/W, 7 x 24 two-hour O/S (NT and NetWare)	Three years	(2-5-2 Part Number for U.S.)	FP-CX206-36

Business Solutions Priority Service Plan – Priority Plus Level

7 x 24 4-hour H/W, 7 x 24 one-hour O/S (NT & NetWare), SPS account review	Three Years	(2-5-2 Part Number for U.S.)	FP-CX306-36
7 x 24 4-hour H/W, 7 x 24 one-hour O/S (NT & NetWare), SPS account review	Three Years	(2-5-2 Part Number for Canada)	FP-CZ306-36

Business Solutions Priority Service Plan – Priority Premier Level

7 x 24 two-hour H/W, 7 x 24 one-hour O/S (NT and NetWare), SPS Account Review, Proactive Management Review, named S/W rep	Three Years	(2-5-2 Part Number for U.S.)	FP-CX506-36
7 x 24 two-hour H/W, 7 x 24 one-hour O/S (NT & NetWare), SPS Account Review, Proactive Management Review, named S/W rep	Three Years	(2-5-2 Part Number for Canada)	FP-CZ506-36

Business Solutions Priority Service Plan – Priority Executive Premier Level

7 x 24 two-hour H/W, 7 x 24 30-minute O/S (NT and NetWare), named S/W rep, Proactive Management Review, SPS Account Review	Three Years	(2-5-2 Part Number for U.S.)	FP-CX706-36
7 x 24 two-hour H/W, 7 x 24 30-minute O/S (NT & NetWare), Proactive Management Review, named S/W rep, SPS Account Review	Three Years	(2-5-2 Part Number for Canada)	FP-CZ706-36

Note: For more information, customer/resellers may contact **www.Compaq.com/services**

STANDARD MEMORY (PER SERVER)

Compaq **ProLiant** ***CL380 128-MB PC133-MHz Registered ECC SDRAM installed (1 x 128-MB SDRAM DIMM)***

Standard	Slot			
	1	2	3	4
128 MB	128 MB	Empty	Empty	Empty

STANDARD MEMORY PLUS OPTIONAL MEMORY (PER SERVER)

Up to 3200-MB memory is available with the optional installation of PC133-MHz Registered ECC SDRAM DIMM memory kits.

Up To	Slot			
	1	2	3	4
3200 MB	128 MB	1024 MB	1024 MB	1024 MB

STANDARD MEMORY REPLACED WITH OPTIONAL MEMORY (PER SERVER)

Up to 4096 MB of memory is available with the removal of the standard 128-MB of memory and the optional installation of 133-MHz Registered ECC SDRAM memory kits.

UpTo	Slot			
	1	2	3	4
4096 MB	1024 MB	1024 MB	1024 MB	1024 MB

Note: Chart does not represent all possible configurations.

Following are memory options available from Compaq:

• 1024-MB PC133-MHz Registered ECC SDRAM DIMM Memory Option Kit	128280-B21
• 512-MB PC133-MHz Registered ECC SDRAM DIMM Memory Option Kit	128279-B21
• 256-MB PC133-MHz Registered ECC SDRAM DIMM Memory Option Kit	128278-B21
• 128-MB PC133-MHz Registered ECC SDRAM DIMM Memory Option Kit	128277-B21
• 64-MB PC133-MHz Registered ECC SDRAM DIMM Memory Option Kit	128276-B21

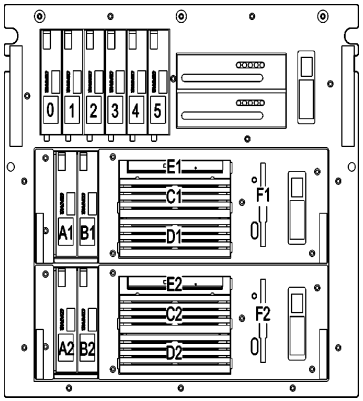

Key:

0-5- 6 x 1″ wide Ultra2/Ultra3 SCSI Drive Cage for internal shared storage

A1, A2, B1, B2-2 x 1″ Ultra2/Ultra3 drive cage

C1, C2,D1, D2 – Two available half height bays for tape backup storage

E1, E2-Low-profile High Speed IDE CD-ROM

F1, F2-3.5″ diskette drive

DRIVE SUPPORT

	Quantity Supported	*Position Supported*	*Controller*
Removable Media			
1.44-MB Diskette Drive	Up to 1	D1, D2	Integrated
Low-profile IDE CD-ROM Drive	Up to 1	C1, C2	Integrated IDE
DVD-ROM Option Kit	Up to 1	C1, C2	Integrated IDE
Hard Drives			
Wide Ultra3 Hot Pluggable Drives			
1-inch			Integrated Dual Channel Wide Ultra2 SCSI Adapter
18.2-GB 10,000 rpm	Up to 6	0-5	64-bit Dual Channel Wide Ultra2 SCSI Adapter
9.1-GB 10,000 rpm			(occupying a PCI slot)
	Up to 4	A1, A2,	Integrated Smart Array Controller Module
	(2 per server)	B1, B2	Smart Array 221 Controller
			Smart Array 431 Controller
			Smart Array 3200 Controller
			Smart Array 4200 Controller
			64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter
			64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter

	Quantity Supported	Position Supported	Controller
Wide Ultra2 Hot Pluggable Drives 1 inch 18.2GB 10,000 rpm 9.1GB 10,000 rpm	Up to 6 Up to 4 (2 per server)	0-5 A1, A2, B1, B2	Integrated Dual Channel Wide Ultra2 SCSI Adapter 64-bit Dual Channel Wide Ultra2 SCSI Adapter (occupying a PCI slot) Integrated Smart Array Controller Module Smart Array 221 Controller Smart Array 431 Controller Smart Array 3200 Controller Smart Array 4200 Controller 64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter 64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter
Wide Ultra2 SCSI – Non-Hot Plug 1 inch 18.2GB 9.1GB	Up to 2 (per server)	C1 & D1, C2 & D2	Integrated Dual Channel Wide Ultra2 SCSI Adapter 64-bit Dual Channel Wide Ultra2 SCSI Adapter (occupying a PCI slot) Integrated Smart Array Controller Module Smart Array 3200 Controller Smart Array 221 Controller Smart Array 4200 Controller Smart Array 431 Controller 64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter 64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter
Shared External Storage *StorageWorks* Enclosure 4200/4300 Family (supports Wide Ultra2/Ultra3 1″ drives only)	Up to 1	External	Shared Storage RAID CR3500 Controller
Non-Shared External Storage *StorageWorks* Enclosure 4200/4300 Family (supports Wide Ultra2/Ultra3 1″ drives only)	Up to 4	External	Smart Array 221 Controller Smart Array 431 Controller Smart Array 3200 Controller 64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter 64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter

DRIVE SUPPORT

	Quantity Supported	Position Supported	Controller
ProLiant Storage System Model UE *ProLiant* Storage System Model U2 (supports Wide-Ultra SCSI-3 1″ or 1.6″ drives only)	Up to 4	External	Integrated Dual Channel Smart Array 3200 Controller Smart Array 4200 Controller Smart Array 221 Controller Wide Ultra2 SCSI Adapter Wide Ultra SCSI-3 SCSI Adapter 64 bit/66 MHz Dual Channel Wide Ultra3 SCSI Adapter
SCSI Storage Expander II	Up to 3	External	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide Ultra SCSI-3 Adapter
Maximum Shared Storage Capacity Internal 109.2 GB External Total			(Ultra2/Ultra3 SCSI Attached) (6 x 18.2GB Wide Ultra3 hard drives) 145.6 GB (1 shared storage system: 12 x 18.2GB Wide Ultra3 1″ hard drives) 254.8 GB
Maximum Non-Shared Storage Capacity Internal External Total		 72.8 1.164 TB 1.237 TB	(*ProLiant* Storage System Attached) (4 x 18.2GB Wide Ultra3 hard drives – both servers) (4 *ProLiant* Storage Systems x 8 x 36.4GB Wide Ultra SCSI-3 1.6″ hard drives)
Tape Drives Internal 12/24GB DAT Internal 20/40GB DAT DDS-4 Internal 4/8-GB DAT	Up to 4	A1, B1, A2, B2	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide Ultra SCSI-3 Adapter 64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter 64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter

DRIVE SUPPORT

	Quantity Supported	Position Supported	Controller
20/40-GB DAT DDS-4 8 Cassette Autoloader (internal) Internal 35/70-GB DLT Internal 20/40-GB DLT	Up to 4 1 Drive	A1, B1, A2, B2	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide-Ultra SCSI-3 Adapter 64 bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter 64 bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter
DLT Tape Array II Model 35/70	1 to 4	External	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide Ultra SCSI-3 Adapter
35/70-GB DLT 20/40-GB DLT	Up to 3	External	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide Ultra SCSI-3 Adapter 64-bit/6-6MHz Dual Channel Wide Ultra3 SCSI Adapter 64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter
DLT 15 Cartridge Library Model 35/70-1 DLT 15 Cartridge Library Model 35/70-2 DLT 15 Cartridge Library Model 20/40-1 DLT 15 Cartridge Library Model 20/40-2		1 Library per SCSI channel	Integrated Dual Channel Wide Ultra2 SCSI Adapter Wide-Ultra SCSI-3 Adapter
TL891 DLT Library (1 – 2 Drive) TL881 DLT Library (1 – 2 Drive)	2 drives per SCSI channel	External	Dual Channel Wide Ultra SCSI-3 Adapter Wide Ultra SCSI Adapter

SUPPORTED TAPE DRIVE OPTIONS

Tape Drive	*Operating System*	*Key*
Internal 12/24-GB DAT Drive	Microsoft Windows NT 4.0	AS, NT, SE
Internal 20/40-GB DAT DDS-4 Drive	Microsoft Windows NT 4.0	AS, NT, SE
Internal 4/8-GB DAT Drive	Microsoft Windows NT 4.0	AS, NT, SE
AIT 35-GB Tape Drive	Microsoft Windows NT 4.0	AS, NT, SE
AIT 50-GB Tape Drive	Microsoft Windows NT 4.0	AS, NT, SE
12/24-GB DAT DDS-3 8 Cassette autoloader	Microsoft Windows NT 4.0	AS, NT, SE
20/40-GB DAT DDS-4 8 Cassette autoloader	Microsoft Windows NT 4.0	AS, NT, SE
DLT Tape Array II Model 35/70	Microsoft Windows NT 4.0	LG
35/70-GB DLT Drive	Microsoft Windows NT 4.0	AS, NT, SE, LG
DLT 15 Cartridge Library Model 35/70-1	Microsoft Windows NT 4.0	AS, NT, SE
DLT 15 Cartridge Library Model 35/70-2	Microsoft Windows NT 4.0	AS, NT, SE
20/40-GB DLT Drive	Microsoft Windows NT 4.0	AS, NT, SE
DLT 15 Cartridge Library Model 20/40-1	Microsoft Windows NT 4.0	AS, NT, SE
DLT 15 Cartridge Library Model 20/40-2	Microsoft Windows NT 4.0	AS, NT, SE
TL891 DLT Library (1 – 2 Drive)	Microsoft Windows NT 4.0	AS, NT, SE
TL881 DLT Library (1 – 2 Drive)	Microsoft Windows NT 4.0	AS, NT, SE

Keys to Chart

AS Computer Associates ARCserve
JS Computer Associates JETserve
AO Computer Associates ARCserve Open
LG Legato Networker Software from Legato
ST Sytos Plus software
SE Seagate Software Backup Exec
NT Microsoft Windows NT support for Compaq tape drives included in operating system

TECHNICAL SPECIFICATIONS

Dimensions (HxWxD)		
Rack Form Factor	17.35 x 17.50 x 22.00 in	
Tower Form Factor	19.75 x 17.50 x 22.00 in	
Weight		
Two servers, one shared storage system, cabinet	125 lbs	
Input Requirements (per power supply)		
Range Line Voltage	115 VAC/230 VAC	
Nominal Line Voltage	100 VAC, 120 VAC, 220 VAC, 240 VAC, 250 VAC	
Line Frequency	50 to 60 Hz	
Input Power	90 VAC/264 VAC	
Heat	41° to 122°F/5° to 50° C	
Operating Temperature Range	–40° to 185°F/–40° to 85°C	
Output Power	225W	
Operating Temperature	50° to 93°F/10° to 35°C	
Shipping Temperature	–22° to 140°F/–30° to 60°C	
Relative Humidity (non-condensing)		
Operating	20% to 80%	
Non-operating	5% to 90%	
Maximum Wet Bulb Temperature	101.7°F/38.7°C	
Acoustic Noise	NPEL (BELS)	AVERAGE SPL (dBA)
Idle (Fixed Disk Drives Spinning)	6.35	49.1
Operating (Random Seeks to Fixed Disks)	6.43	49.8

TECHNICAL SPECIFICATIONS

Diskette Size	3.5 in
LED Indicators (front panel)	Green
Read/Write Capacity per Diskette (high/low density)	1.44 MB/720 KB
Drive Supported	One
Drive Height	One-third
Drive Rotation	300 rpm
Transfer Rate (high/low)	500/250 Kb/s
Bytes/Sector	512
Sectors/Track (high/low)	18/9
Tracks/Side (high/low)	80/80
Access Times	
Track-to-Track (high/low)	3/6 ms
Average (high/low)	94/169 ms
Settling Time	15 ms
Latency Average	100 ms
Cylinders (high/low)	80/80
Read/Write Heads	Two

TECHNICAL SPECIFICATIONS

Hot Plug Wide Ultra3 SCSI Hard Drives

18.2 GB 10K	**Capacity**	18209.3 MB	**Seek Time** (typical reads, including settling)	
	Height	1 in/25.4 mm	Single Track	0.8 ms
	Width	3.5 in/88.9 mm	Average	5.2 ms
	Interface	Wide Ultra3 SCSI	Full-Stroke	12 ms
	Transfer Rate	Rotational Speed		10,000 rpm
	Synchronous			
	(Maximum)	160 MB/s	**Physical Configuration**	
			Bytes/Sector	512
			Logical Blocks	35,565,080
			Operating Temperature	50º to 95ºF/10º to 35ºC

9.1 GB 10K	**Capacity**	9100.0 MB	**Seek Time** (typical reads, including settling)	
	Height	1 in/25.4 mm	Single Track	0.8 ms
	Width	3.5 in/88.9 mm	Average	5.0 ms
	Interface	Wide Ultra3 SCSI	Full-Stroke	12 ms
	Transfer Rate	Rotational Speed		10,000 rpm
	Synchronous			
	(Maximum)	160 MB/s	**Physical Configuration**	
			Bytes/Sector	512
			Logical Blocks	17,773,524
			Operating Temperature	50º to 95ºF/10º to 35ºC

Hot Plug Wide Ultra2 SCSI Hard Drives

18.2 GB 0 K	**Capacity**	18209.3 MB	**Seek Time** (typical reads, including settling)	
	Height	1 in/25.4 mm	Single Track	0.8 ms
	Width	3.5 in/88.9 mm	Average	5.7 ms
	Interface	Wide Ultra2 SCSI	Full-Stroke	12.2 ms
	Transfer Rate	Rotational Speed	10,000 rpm	
	Synchronous			
	(Maximum)	80 MB/s		
			Physical Configuration	
			Bytes/Sector	512
			Logical Blocks	35,565,080
			Operating Temperature	50º to 95ºF/10º to 35ºC

TECHNICAL SPECIFICATIONS

9.1 GB	**Capacity**	**Seek Time**	
10K	9100.0 MB	(typical reads, including settling)	
Height	1 in/25.4 mm	Single Track	0.8 ms
Width	3.5 in/88.9 mm	Average	5.4 ms
Interface	Wide Ultra2 SCSI	Full-Stroke	12.2 ms
Transfer Rate	Rotational Speed	10,000 rpm	
Synchronous			
(Maximum)	80 MB/s		
		Physical Configuration	
		Bytes/Sector	512
		Logical Blocks	17,773,524
		Operating Temperature	50º to 95ºF/
			10º to 35ºC

18.2 GB	**Capacity**	**Seek Time**	
	8209.3 MB	(typical reads, including settling)	
Height	1 in/25.4 mm	Single Track	0.8 ms
Width	3.5 in/88.9 mm	Average	7.0 ms
Interface	Wide Ultra2 SCSI	Full-Stroke	15 ms
Transfer Rate	Rotational Speed		7200 rpm
Synchronous			
(Maximum)	80 MB/s		
		Physical Configuration	
		Bytes/Sector	512
		L ogical Blocks	35,565,080
		Operating Temperature	50º to 95ºF/
			10º to 35ºC

9.1 GB	**Capacity**	**Seek Time**	
	9100.0 MB	(typical reads, including settling)	
Height	1 in/25.4 mm	Single Track	0.8 ms
Width	3.5 in/88.9 mm	Average	7.1 ms
Interface	Wide Ultra2 SCSI	Full-Stroke	16 ms
Transfer Rate	Rotational Speed	7200 rpm	
Synchronous			
(Maximum)	80 MB/s		
		Physical Configuration	
		Bytes/Sector	512
		Operating Temperature	Logical
		Blocks	17,773,524
			50º to 95ºF/
			10º to 35ºC

Low-Profile
24X Max
IDE CD-ROM
Drive

TECHNICAL SPECIFICATIONS

Disk	
Applicable Disk	CD-DA, CD-ROM (Mode 1 and 2)
CD-XA, CD-I (Mode 2, Form 1 and 2)	
CD-I Ready, CD Extra, Video CD, CD-Bridge	
Photo CD (Single and Multi-session)	
CD-WO	
Capacity	550 MB (Mode 1, 12 cm)
640 MB (Mode 2, 12 cm)	
Diameter	4.7 in, 3.15 in/12 cm, 8 cm
Rotational Speed	4200 rpm maximum
Center Hole	0.6 in/15 mm diameter
Thickness	1.2 mm
Track Pitch	1.6 m
Block Size	
Mode 0	2,368, 2,352 bytes
Mode 1	2,352, 2,340, 2,336, 2,048 bytes
Mode 2	2,352, 2,340, 2,336, 2,048 bytes
Interface	IDE (ATAPI)
Access Times (typical)	
Random	< 140 ms
Full-Stroke	< 300 m
Data Transfer Rate	
Sustained	150 KB/s (sustained 1X)
Burst	2100 to 4800 KB/s
Cache Buffer	128 KB
Start-up Time (typical)	< 10 seconds
Stop Time	< 5 seconds
Operating Conditions	
Temperature	41° to 120°F/5° to 55°C
Humidity	10 percent to 80 percent
Dimensions	
(HxWxD, maximum)	0.51 x 5.24 x 5.2 in/12.7 x 131 x 130 mm
Weight	< 0.75 lb/< 340 g

TECHNICAL SPECIFICATIONS

Drives Supported	Up to 6 Wide Ultra2/Ultra3 SCSI Hard Drives (internal)
SCSI Electrical Interface	Low Voltage Differential
Data Transfer Method	32-bit PCI bus-master
SCSI Channel Transfer Rate (peak)	80 MB/s (internal)
Maximum Transfer Rate per PCI Bus (peak)	133 MB/s
SCSI Termination	Active Termination
SCSI Connectors	1 external, (68-pin) (Tape only) 2 internal, (68-pin) (Port 1 – Tape only; Port 2 – RAID 0, 1, 0+1, 5 or Tape)
Cache	8-MB Read Cache (ECC protected)
Logical Drives Supported	32
Software Upgradable Firmware	Yes
Reliability Features	
Distributed Data Guarding (RAID 5)	Yes
Data Guarding (RAID 1)	Yes
Drive Mirroring (RAID 0)	Yes
Drive Stripping (RAID 0+1)	Yes
Automatic Data Recovery	Yes

Drives Supported	Up to 15 SCSI devices per channel
Data Transfer Method	64-bit PCI bus-master
SCSI Channel Transfer Rate	160 MB/s (80 MB/s per channel)
Maximum Transfer Rate per PCI Bus (peak)	133 MB/s
SCSI Termination	Active Termination
SCSI Connectors	2 external (80-pin), 2 internal (68-pin)

Compaq
NC3163 Fast
Ethernet NIC
Embedded
10/100 WOL

TECHNICAL SPECIFICATIONS

Network Interface Compatibility	10Base-T/100Base-TX IEEE 802.3/802.3u compliant
Data Transfer Method	32-bit bus-master PCI
Network Transfer Rate	
10Base-T (Half-Duplex)	10 Mb/s
10Base-T (Full-Duplex)	20 Mb/s
100Base-TX (Half-Duplex)	100 Mb/s
100Base-TX (Full-Duplex)	200 Mb/s
Connector	RJ-45
Cable Support	
10Base-T	Categories 3, 4 or 5 UTP; up to 328 feet (100 meters)
100Base-TX	Category 5 UTP; up to 328 feet (100 meters)

TECHNICAL SPECIFICATIONS

Controller Chip	ATI RAGE IIC PCI
Video DRAM	4-MB Video SGRAM
Data Transfer Method	32-bit PCI

Support Resolution	Supported Color Depths:
640 x 480	16.7M, 64K, 256, 16
800 x 600	16.7M, 64K, 256, 16
1024 x 768	16.7M, 64K, 256, 16
1152 x 864	16.7M, 64K, 256, 16
1280 x 1024	16.7M, 64K, 256, 16
1600 x 1200	64K, 256, 16

Connector	VGA

Drives Supported	Up to 14 drives (The *ProLiant* CL380 will support up to four internal hard drives [both servers] or 12 external hard drives using the *ProLiant* Storage System Model UE storage enclosure.)

Data Transfer Method	32-bit PCI bus-master
Maximum Transfer Rate per PCI Bus (peak)	133 MB/s
Simultaneous Drive Transfer Channels	One
Total Transfer Rate	80 MB/s per channel
Software Upgradable Firmware	Yes
Cache	6-MB of ECC protected read cache

Reliability Features	
Distributed Data Guarding (RAID 5)	Yes
Drive Mirroring (RAID 1)	Yes
Drive Stripping (RAID 0, 0 + 1)	Yes
Automatic Data Recovery	Yes

TECHNICAL SPECIFICATIONS

Protocol	Wide Ultra3 SCSI
SCSI Electrical Interface	Low Voltage Differential (LVD)
Drives Supported	Up to 14 Drives Wide Ultra3 or Wide Ultra2 SCSI hard drives
SCSI Port Connectors	Single External/Internal SCSI port
Data Transfer Method	64-bit PCI bus-master
Maximum Transfer Rate on System Bus (peak)	266 MB/s
Simultaneous Drive Transfer Channels	One
Channel Transfer Rate	160 MB/s per channel
Software Upgradable Firmware	Yes
Read-Only Cache	16-MB Cache Performance Engine
Logical Drives Supported	32
Reliability Features	
Distributed Data Guarding (RAID 5)	Yes
Drive Mirroring (RAID 1)	Yes
Drive Striping (RAID 0, 0+1)	Yes
Automatic Data Recovery	Yes
Disk Drive Protocol Support	Wide Ultra3 and Wide Ultra2 SCSI
Enclosure Support	New *StorageWorks* Enclosure 4200/4300 Family (Models 4214, 4254, 4314 and 4354)

TECHNICAL SPECIFICATIONS

Smart Array 3200 Controller

Drives Supported	Up to 15 drives per channel (The *ProLiant* CL380 will support up to four internal hard drives [both servers] or 12 external hard drives using the *ProLiant* Storage System Model UE storage enclosure.)
Data Transfer Method	32-bit PCI bus-master
Maximum Transfer Rate per PCI Bus (peak)	133 MB/s
Simultaneous Drive Transfer Channels	Two
Total Transfer Rate	160 MB/s (80 MB/s per channel)
Software Upgradable Firmware	Yes
Cache	64-MB Array Accelerator
Reliability Features	
Distributed Data Guarding (RAID 5)	Yes
Data Guarding (RAID 4)	Yes
Drive Mirroring (RAID 1)	Yes
Drive Stripping (RAID 0)	Yes
Automatic Data Recovery	Yes

Smart Array 4200 Controller

Protocol	Wide Ultra2 SCSI
SCSI Electrical Interface	Low Voltage Differential
Drives Supported (maximum)	56 drives
Data Transfer Method	
Maximum Transfer Rate on PCI Bus (peak)	64-bit PCI bus-master
Simultaneous Drive Transfer Channels	Four
Total Channel Transfer Rate	320 MB/s
Software Upgradable Firmware	Yes
Cache	64 MB battery-backed and removable
Logical Drives Supported	32
Reliability Features	
Cache Battery Backup	Yes
Online Capacity Expansion	Yes
Logical Drive Capacity Extension	Yes
Online RAID Level Migration	Yes
Online Stripe Size Migration	Yes
Automatic Data Recovery	Yes
Distributed Data Guarding (RAID 5)	Yes
Data Guarding (RAID 4)	Yes
Drive Mirroring (RAID 1)	Yes
Drive Striping (RAID 0, 0 + 1)	Yes

TECHNICAL SPECIFICATIONS

64-bit/66-MHz Single Channel Wide Ultra3 SCSI Adapter

Drives Supported	Up to 14 SCSI devices
Storage Capacity	254.8 GB
Data Transfer Method	64-bit/66-MHz PCI bus-master
Maximum SCSI Channel Transfer Rate	160 MB/s per channel
Maximum Transfer Rate per PCI Bus (peak)	532 MB/s per channel
SCSI Protocols	Wide Ultra3 SCSI Wide Ultra2 SCSI Wide-Ultra SCSI-3 Fast-Wide SCSI-2 Fast SCSI-2
Electrical Protocol	Low Voltage Differential (LVD)
SCSI Termination	Active Termination
External SCSI Connectors	One 68-pin Wide-Ultra SCSI-3 connector
Internal SCSI Connectors	Two 68-pin Wide-Ultra SCSI-3 connectors (one LVD and one SE) One 50-pin Narrow Ultra connector

64-bit/66-MHz Dual Channel Wide Ultra3 SCSI Adapter

Drives Supported	Up to 28 SCSI devices (14 per channel)
Storage Capacity	1.019 TB
Data Transfer Method	64-bit/66-MHz PCI bus-master
Maximum SCSI Channel Transfer Rate	160 MB/s per channel
Maximum Transfer Rate per PCI Bus (peak)	532 MB/s per channel
SCSI Protocols	Wide Ultra3 SCSI Wide Ultra2 SCSI Wide-Ultra SCSI-3 Fast SCSI-2
Electrical Protocol	Low Voltage Differential (LVD)
SCSI Termination	Active Termination
External SCSI Connectors	Two 68-pin VHDCI connectors
Internal SCSI Connectors	Two 68-pin Wide-Ultra SCSI-3 connectors One 50-pin Narrow Ultra connector

TECHNICAL SPECIFICATIONS

RAID CR3500 Controller

Drives Supported	Internal SCSI IDs are 9, 10, 11, 12, 13, 14. External storage is a split bus with SCSI IDs 0, 1, 2, 3, 4, 5 and 8 on bus A and bus B, or single bus with SCSI IDs 0, 1, 2, 3, 4, 5, 8, 15.
Data Transfer Method	Wide Ultra2 SCSI
Maximum Transfer Rate per PCI Bus (peak)	40 MB/s
Simultaneous Drive Transfer Channels	One
Total Transfer Rate	40 MB/s
Software Upgradable Firmware	Yes
Cache	128 MB
Reliability Features	
Distributed Data Guarding (RAID 5)	Yes
Data Guarding (RAID 4)	Yes
Drive Mirroring and Stripping (RAID 0+1)	Yes
Drive Mirroring (RAID 1)	Yes
Drive Stripping (RAID 0)	Yes
Just a Box Of Disks (JBOD)	Yes
Automatic Data Recovery	Yes

Compaq Rack 9142 (42U) (optional)

Dimensions	**Total Cabinet**	**Shipping**	**Color**
Height	78.7 in/2,000 mm	83.38 in/2,168.65 mm	
Depth	35.8 in/909 mm	48 in/1,219.2 mm	Opal
Width	23.7 in/603 mm	32 in/812.8 mm	
Weight	253 lbs/114.84 kg	325 lbs/147.52 kg	

Compaq Rack 9136 (36U) (optional)

Dimensions	**Total Cabinet**	**Shipping**	**Color**
Height	68.6 in/1,742 mm	75.25 in/1,911.35 mm	
Depth	35.8 in/909 mm	48 in/1,219.2 mm	Opal
Width	23.7 in/603 mm	32 in/812.8 mm	
Weight	165 lbs/75 kg	234 lbs/106 kg	

TECHNICAL SPECIFICATIONS

Compaq Rack 9122 (22U) (optional)

Dimensions	Total Cabinet	Shipping	Color
Height	43 in/1,092 mm	52.25 in/1,327.15 mm	
Depth	35.8 in/909 mm	48 in/1,219.2 mm	Opal
Width	24 in/610 mm	32 in/812.8 mm	
Weight	176 lbs/79.89 kg	225 lbs/102.13 kg	

OS Driver Support

Operating Systems	Current Drivers	Driver Location (Punch-out diskettes only)
Microsoft Windows NT Server, Enterprise Edition 4.0	NTSSD 2.17a	SmartStart CD and www.compaq.com
Microsoft Windows 2000 Advanced Server	Support Software Bundle for Windows 2000 5.02a	SmartStart CD and www.compaq.com
Novell NetWare NCS 5.1	NTSSD 5.50	SmartStart CD and www.compaq.com
Novell NetWare NHS 4.2 w/SS 4.70	NTSSD 5.50	SmartStart CD and www.compaq.com


Compaq *ProLiant*
NT Clusters
COMPAQ
www.compaq.com/highavailability

Product Line/Speaker Notes

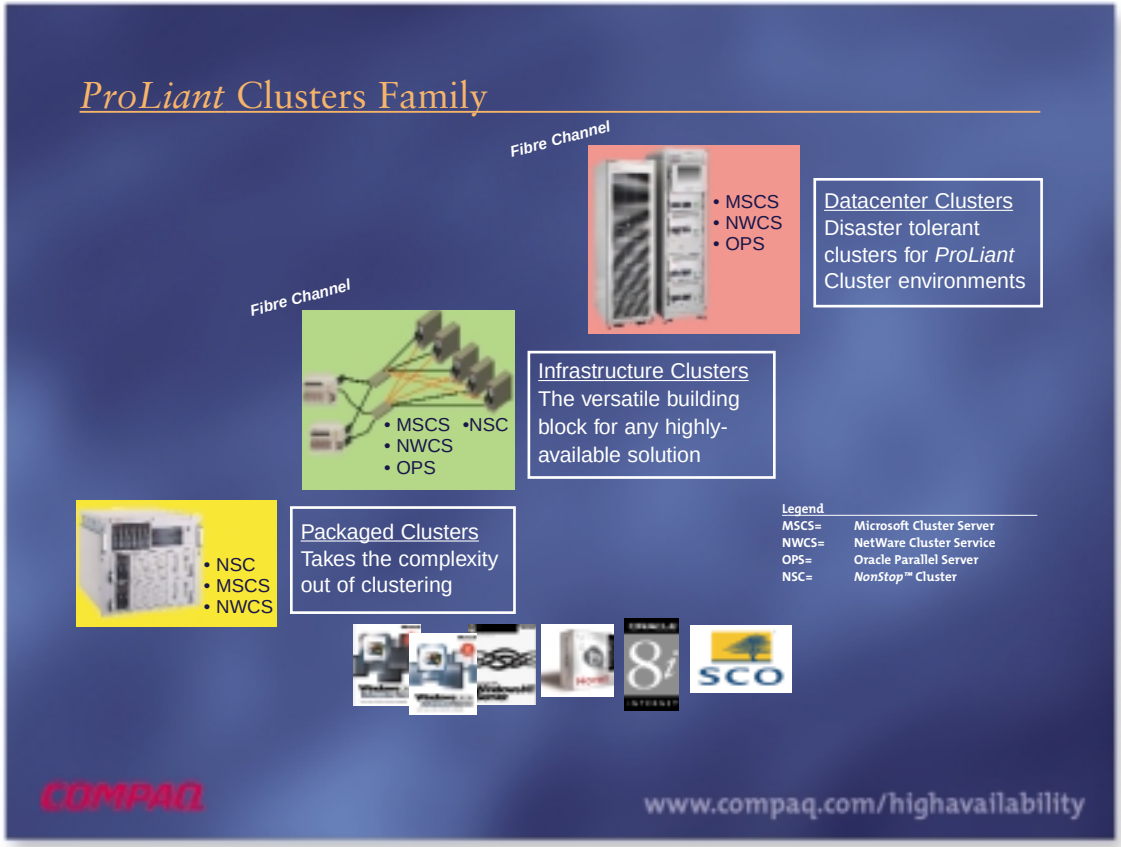

Microsoft has positioned the new suite of Windows 2000 operating systems differently than NT4. Each server-based operating system will be supported in cluster environments tailored to the capabilities of each version. Each of these operating systems carries strengths with it in various markets and application areas and provides a level of flexibility that has not been available before.


Microsoft Cluster Server and Microsoft Cluster Service – Features
• Functions
• Automatically detects, isolates and recovers from failures in:
— hardware
— operating system
— application software
• Provides a single point of cluster management
• Provides an industry-standard application program interface
• Features
• Two-node failover functionality
• Active/Active or Active/Standby
• Industry Standard API
• Cluster alias and IP address failover
• Multiple failover objects
• Automatic failover/failback
• Operator-controlled failover/failback possible
• Single system administration GUI
• "Cluster-aware" application support
• Generic application support
COMPAQ
www.compaq.com/highavailability

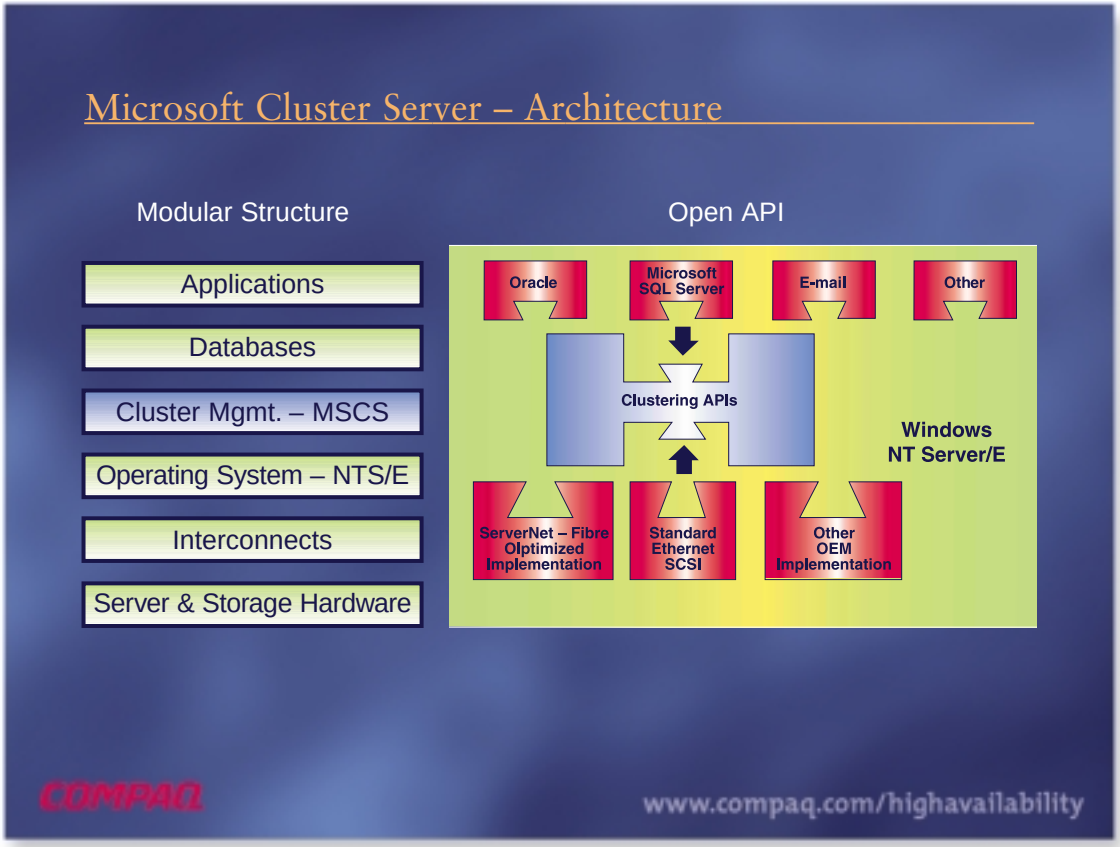

Architectural representation

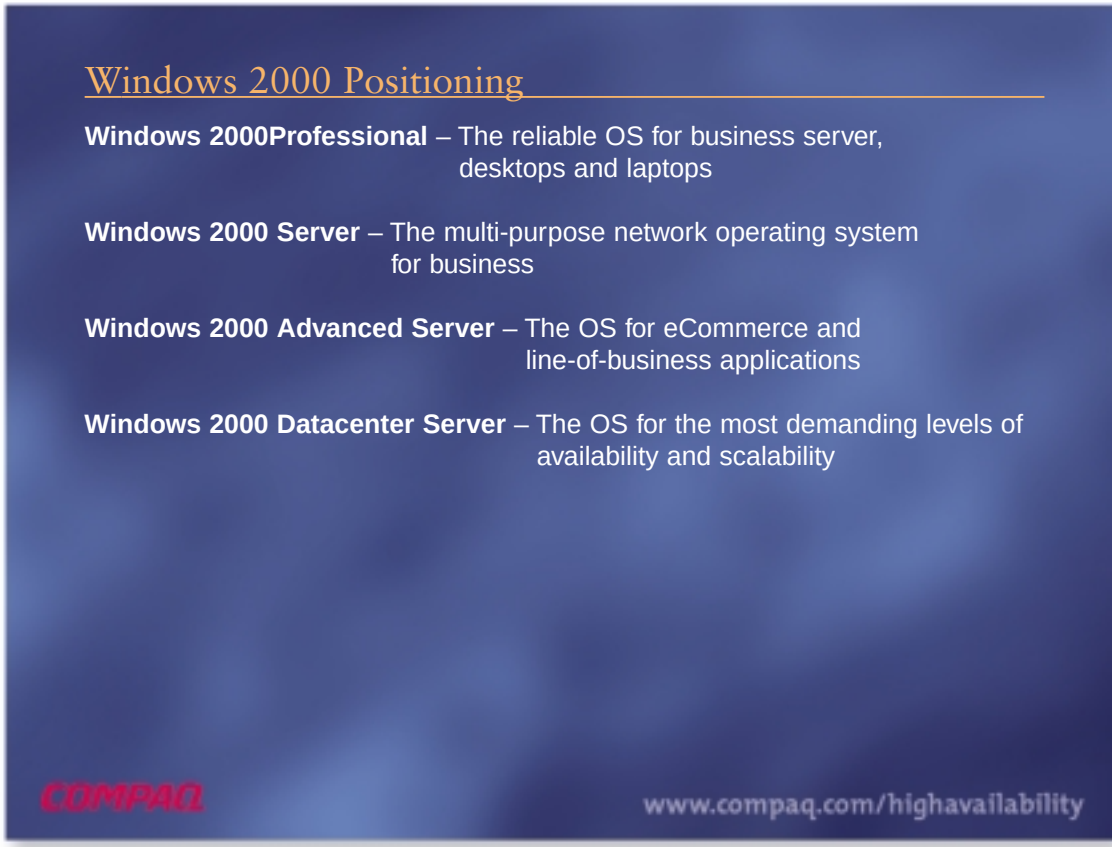

Microsoft has positioned the new suite of Windows 2000 operating systems differently than NT4. The server-based operating systems will be supported in cluster environments tailored to the capabilities of each version. Each of these operating systems carry with it strengths in various markets and application areas, and provide a level of flexibility not available before.

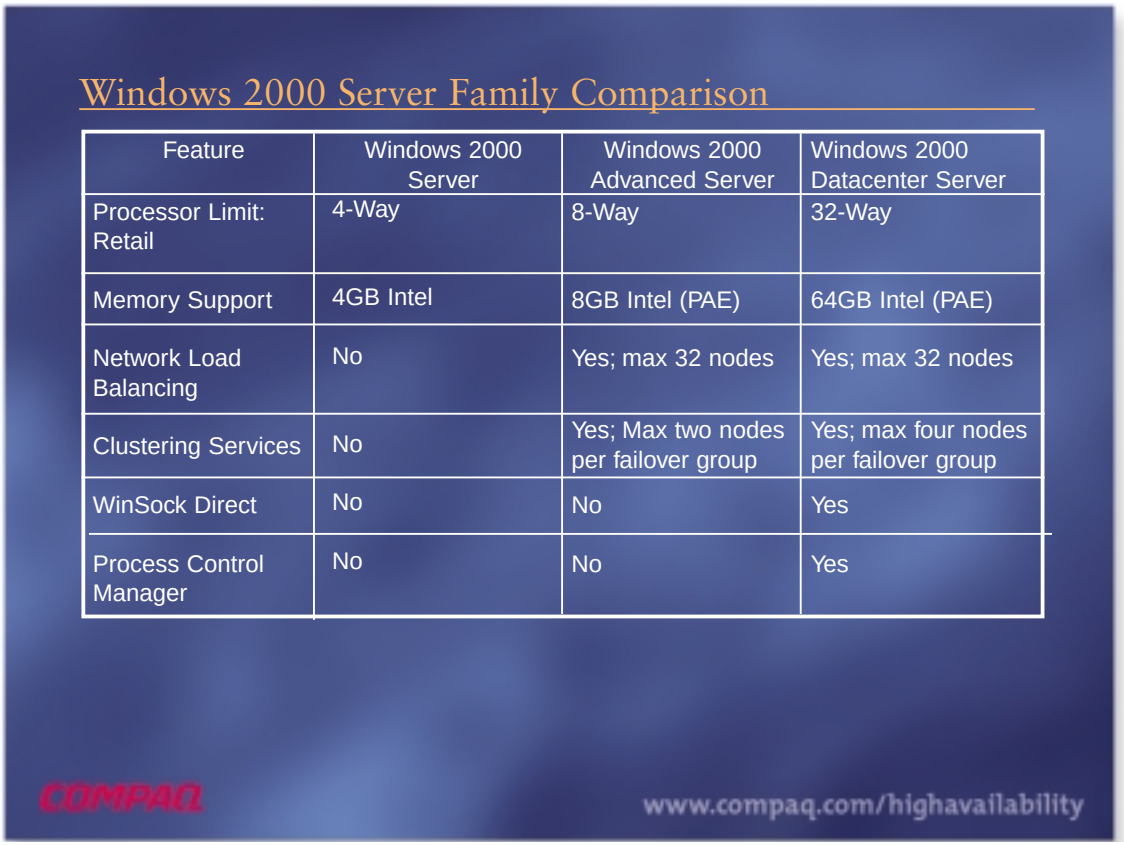

This table represents a comparison of features and functionality across the Windows 2000 product line. Note the significant feature differences between Server, Advanced Server and Datacenter Server. Compaq high availability products will include support for each of these operating system types.

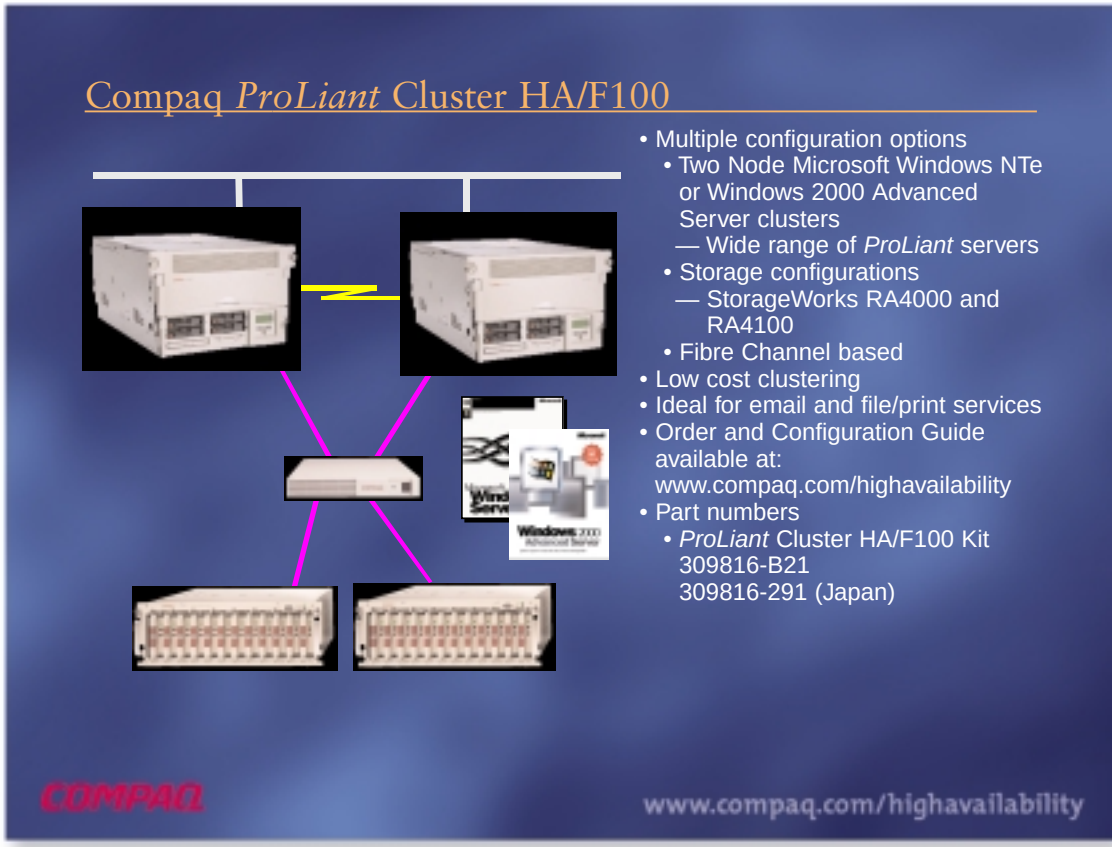

The Compaq *ProLiant* Cluster HA/F100 is a single-loop configuration which allows for MSCS Server or application failover and failback in the event an application or server component fails.

Product Line/Speaker Notes

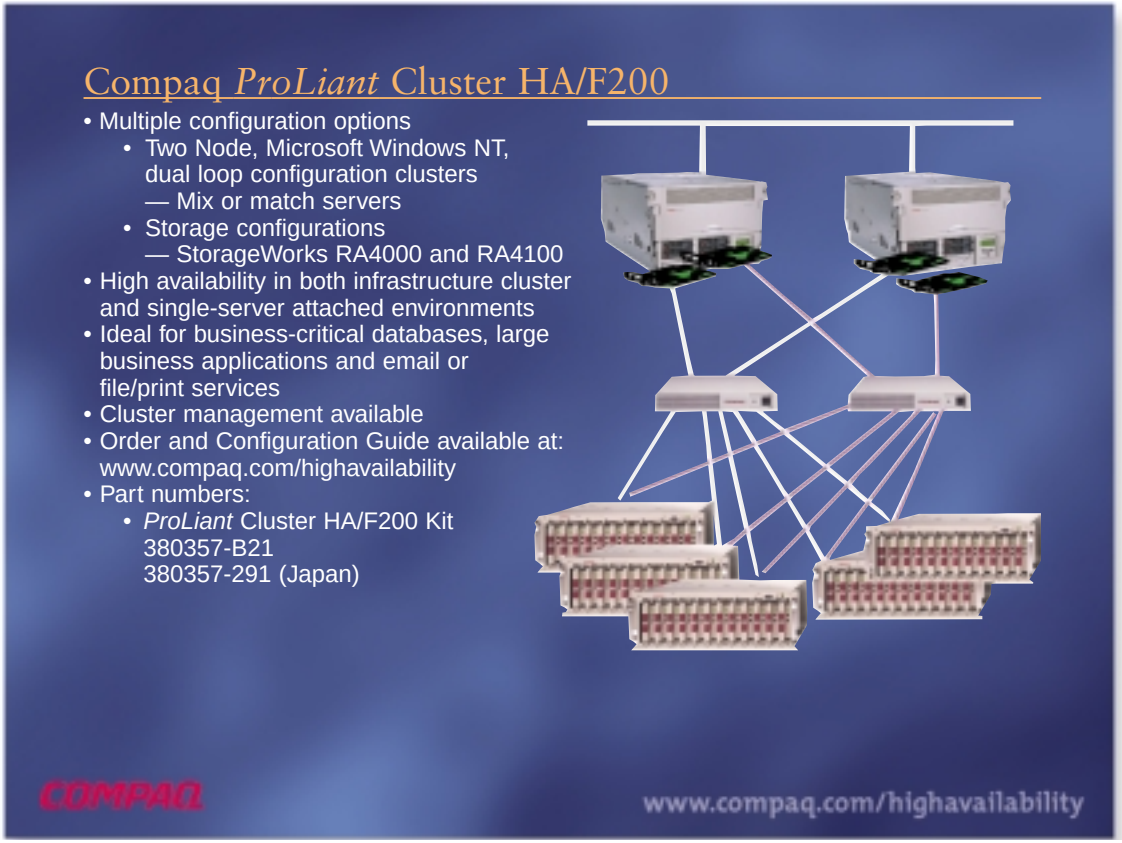

The Compaq *ProLiant* Cluster HA/F200 provides for the full functionality of MSCS in a totally redundant environment in which if any single component fails, the access to storage and application data is still maintained via an alternate path.

The Compaq *ProLiant* Cluster HA/F500 is the ultimate Windows 2000/NT cluster with No-Single-Points-of-Failure in a high-end solution in which Disaster Tolerant configurations can assure access to data in the event of a failure within a matter of minutes…not hours or days.

Windows 2000 products will be released in four versions. Windows 2000 Professional, Windows 2000 Server, Windows 2000 Advanced Server, and Windows 2000 Datacenter Server. Windows 2000 Advanced Server and Datacenter Server are the versions that support *ProLiant* clustering. Compaq will provide a full suite of clusters for Windows 2000 Advanced Server as a logical follow-on to NT 4.0 Enterprise edition. These solutions will cover the entire Compaq StorageWorks storage subsystems. Datacenter Server clusters will be focused on high-end enterprise server and high-end *StorageWorks* storage. Windows 2000 Clusters will continue to fall within the current cluster naming guidelines, CL1850, HA/F100, HA/F200, HA/F500 clusters. For detailed information related to Windows 2000 support please go to each of these product sites.

Compaq ProLiant Clusters & Windows 2000

Frequently asked questions about Compaq participation in Microsoft's Windows 2000 Advanced Server and Datacenter Server programs.

Q. Which versions of Windows 2000 support clustering?

A. Windows 2000 Advanced Server and Windows 2000 Datacenter Server are the versions of Windows 2000 operating system that include cluster support.

Q. What is the difference between Windows 2000 Advanced Server and Datacenter Server? When should I use each?

A. Windows 2000 Advanced Server is the basic replacement for Windows NT Server 4, Enterprise Edition, in that it will support basic clustering functionality, including all of the currently supported cluster configurations for NT/EE 4.0. Datacenter Server is intended for enterprise applications in which the highest levels of integration and testing (Gold Certification) will be required. Datacenter Server clusters are currently in the process of being defined.

Q. Which Compaq *ProLiant* Cluster solutions support Windows 2000? Does the *ProLiant* CL380 support Windows 2000?

A. We fully intend to support all fibre channel clusters as well as the *ProLiant* CL380 on Windows 2000.

Q. How can I find out which Compaq *ProLiant* Cluster configurations are certified by Microsoft?

A. Supported cluster configurations are found on the Compaq High Availability Web at **www.compaq.com/highavailability**.

Q. When will dual loop/fabric support be available for the F200 and F500 *ProLiant* Clusters?

A. Current plans are for 3Q '00 for both of these programs.

Q. Is the Compaq Cluster Verification Utility supported with Windows 2000 clusters?

A. Yes, it will be supported as each of the cluster types are supported on Windows 2000 Advanced Server.

Q. Are the Compaq *Cluster Monitor* and *Intelligent Cluster Administrator* supported with Windows 2000 clusters?

A. Both of these pieces of cluster management software will be supported on *ProLiant* Clusters.

Q. Can I upgrade my existing Compaq *ProLiant* Cluster solutions to Windows 2000 clusters? What do I need to do?

A. You will be able to migrate your fibre channel cluster solutions currently running on Windows NT 4.0 to Windows 2000. As each of the products are released, documentation for migration will be

included with the new kits. This information will also be available on our external web site in the form of technical white papers.

Q. Is Compaq participating in the Microsoft Windows 2000 beta program?

A. Yes. Microsoft is currently managing support of their beta programs via web access and news groups.

Q. Where can I find information relative to Compaq Servers and Windows 2000?

A. Compaq has a Win 2000 web site at **www.compaq.com/partners/microsoft/windows2000/support/index.html**

Q. What are the configurations that will be supported for Datacenter Server once it is fully released?

A. When Datacenter is released, we will support the components as defined by Microsoft. The server must be eight-way capable or more. Compaq *ProLiant* 8500 servers and Compaq *StorageWorks* RA/MA8000/ESA/EMA12000 will be the storage systems supported. Compaq has also announced a 32-way system that will support Datacenter.

Q. What if I want to get started now to test Datacenter Server four-node clusters?

A. Compaq servers and storage that have been certified as MSCS clusters and reside on the current Microsoft cluster HCL web site should work with Datacenter Server version of Microsoft Windows 2000 operating system in a four-node single loop cluster environment with the drivers identified by Compaq on our web site.

What's New in Microsoft 2000 Cluster Server

White Paper

Introduction

First designed for the Windows NT Server 4.0 operating system, cluster service is substantially enhanced in the Windows 2000 Advanced Server operating system. Windows 2000 Advanced Server provides system services for server clustering as a standard part of the product. A server cluster is a set of independent servers (referred to as nodes) that are managed together. The objective of clustering is to provide very high levels of application and data availability.

The cluster service technology built into Windows 2000 Advanced Server allows two servers to be connected into a cluster for higher availability and easier manageability of server resources. The two servers do not have to be the same size or have the same configuration.

Cluster service is one of two complementary Windows clustering technologies provided as extensions to the base Windows 2000 and Windows NT operating systems. The other clustering technology, Network Load Balancing, complements cluster service by supporting highly available and scalable clusters for front-end applications and services such as Internet or intranet sites, Web-based applications, media streaming and Microsoft Terminal Services.

This document gives an overview of the enhancements of cluster service in Windows 2000 Advanced Server. Additionally, improvements made in Windows 2000 Advanced Server's scalability, reliability and availability enhance cluster service. This document will also explain some features of the Cluster Tool utility from the Windows 2000 Resource Kit, as well as discuss future directions for cluster service.

What's New in Microsoft 2000 Cluster Server

White Paper

Cluster Service for Windows 2000 Features

Installation and Administration
The cluster service may be installed either during or after the installation of Windows 2000AS. The shared disks must be configured, that is RAID sets created and formatted with NTFS, before installing the cluster service.

The installation of cluster service is done through the Add/Remove Programs utility in Control Panel. Select the Add/Remove Windows Components button on the left panel and select Cluster Service to install.

A cluster may be administered with Cluster Administrator from either of the cluster nodes. From a client, a cluster may be administered through the Microsoft Management Console (MMC). Use the Console to connect to the cluster or a cluster node, and then manage the cluster service. The MMC cluster snap-in has to be installed first. This is done by executing ADMINPAK.MSI from the Windows 2000 Advanced Server CD.

Added Features

Support for Rolling Upgrades
Administrators can easily take a server offline for maintenance, permitting "rolling upgrades" of system and application software. There are two major advantages to a rolling upgrade. First, service outages are very short during the upgrade process. Second, you do not have to recreate your cluster configuration. The configuration will remain intact during the upgrade process.

Support for Active Directory and MMC Integration
Cluster service for Windows 2000 uses the Active Directory service to publish information about clusters. Integration with Microsoft Management Console (MMC) allows administrators to monitor the status of all resources in the cluster.

Recovery from Network Failures
Cluster service for Windows 2000 implements a sophisticated algorithm to detect and isolate network failures and to improve failure recovery actions. It can detect a number of different states for network failures and then use the appropriate failover policy to determine whether to fail over the resource group.

Health Monitoring
Cluster service monitor the health of standard applications and servers and can automatically recover mission-critical data and applications from many common types of failure – usually in under a minute. The unit of monitoring and failover is a service or an application. The cluster-administration console can also be used to move workloads within the cluster to balance processing loads or to unload servers for planned maintenance or testing, without taking important data and applications offline for any significant period.

Plug and Play Support for Networks and Disks
The Plug and Play technology built into Windows 2000 allows the cluster service to detect the addition and removal of network adapters, TCP/IP network stacks and shared physical disks.

WINS, DFS and DHCP Support
Cluster service now supports Windows Internet Name Service (WINS), Dynamic Host Configuration Protocol (DHCP) and the Distributed File Services as cluster-aware resources that support failover and automatic recovery. A file share resource can now serve as a

distributed file system (DFS) root or share its folder subdirectories for efficient management of large numbers of related file shares.

Com Support for the Cluster API

Cluster Service of Windows 2000 Advanced Server includes a standard, cross-platform API for developing and supporting cluster-aware applications. This API can be used to create scalable, cluster-aware applications that can automatically balance loads across multiple servers within the cluster and can be accessed by the Windows Scripting Host to control cluster behavior and automate many cluster administration tasks.

Windows 2000AS Features That Enhance Clustering

Scalability

Windows 2000 Advanced Server supports servers that can be expanded with multiple processors. Up to eight processors are supported. An additional memory expansion of up to 8 gigabytes is also supported in Windows 2000 Advanced Server.

Reliability

Kernel-Mode Write Protection

To protect each section of the operating system from errors in other sections, Windows 2000 adds write protection for code and read-only subsections of the kernel and device drivers, just as Microsoft Windows NT always has for user-mode programs and dynamic-link libraries (DLL).

To provide this new protection, hardware memory mapping marks the memory pages containing code, assuring they cannot be overwritten, even by the operating system. This prevents kernel-mode software from silently corrupting other kernel-mode code.

Windows File Protection

On versions of Windows before Windows 2000, installing software in addition to the operating system may overwrite shared system files such as DLL and executable files.

Windows File Protection verifies the source and version of a system file before it is installed. This prevents the replacement of protected system files such as .sys, .dll, .ocx, .ttf, .fon and .exe files.

Driver Signing Feature

Driver Signing is included in Windows to help promote driver quality by allowing Windows 2000 and Windows 98 to notify users whether or not a driver they are installing has passed the Microsoft certification process.

This certification proves to users that the drivers they employ are identical to those Microsoft has tested and notifies users if a driver file has been changed after the driver was put on the Hardware Compatibility List.

Availability

Service Pack Slipstreaming

Service Pack (SP) media can now be easily slipstreamed into the base operating system, which means customers do not have to reinstall SPs after installing new components.

Fewer Maintenance Reboots

Many of the configuration changes that require Windows NT 4.0 to be rebooted no longer require reboots with Windows 2000. Examples of some of the configuration changes that no longer require reboots include the following:

- Extending or mirroring NTFS volumes
- Installing or removing Plug and Play devices
- Adding or removing network protocols or services

- Changing IP or IPX parameters
- Changing network-binding orders
- Adding a new pagefile or increasing its initial or maximum size
- Installing Internet Information Server, SQL Server 7.0, Exchange 5.5 or Microsoft Transaction Services.

Improved Diagnostics

Kernel-Only Crash Dumps – In addition to full-memory crash dumps, Windows 2000 Server supports kernel-only crash dumps to allow faster reboots for systems with large amounts of physical memory.

Faster CHKDSK – Performance of CHKDSK in Windows 2000 has been enhanced significantly. Since so many factors affect CHKDSK performance, it is difficult to quantify these improvements. In some configurations, it is more than 10 times faster under Windows 2000 than it was under Windows NT 4.0.

MSINFO – The MSINFO tool provides information that support staff can use to troubleshoot system problems. MSINFO can be used in a couple of ways. During telephone support, an engineer can ask a user to run MSINFO and relay relevant information. Alternatively, the user can use MSINFO to generate system information that can be saved and then sent to a support engineer.

Faster Recovery and Restart

Recovery Console – This utility allows users to read/write NTFS volumes using the Windows 2000 boot floppy. In addition, it provides a way for administrators to access and recover a Windows 2000 installation (FAT, FAT32, NTFS) with a set of specific commands while preserving Windows 2000 security.
Safe Mode Boot – To help users and administrators diagnose system problems such as errant device drivers, the Windows 2000 operating system can be started using Safe Mode Boot.
Kill Process Tree – This utility allows Task Managers to stop not only a single process, but also any processes created by that parent process.

Recoverable File System – The Windows 2000 file system (NTFS) is highly tolerant of disk failures because it logs all disk I/O operations as unique transactions. In case of a disk failure, the file system can quickly undo or redo transactions as appropriate when the system is brought back up. This reduces unavailable time since the file system can quickly return to a known, functioning state.

Storage Management

Remote Storage Services (RSS) –
Automatically monitors the amount of space available on a local hard disk. When the free space on a primary hard disk dips below the needed level, RSS automatically removes local data that has been copied to remote storage, providing the free disk space needed.

Removable Storage Manager (RSM) –
Presents a common interface to robotic media changers and media libraries. It allows multiple applications to share local libraries and tape or disk drives and controls removable media within a single-server system.

Disk Quotas –
Windows 2000 Server supports disk quotas for monitoring and limiting disk space use on NTFS volumes.

Dynamic Volume Management –
Allows online administrative tasks to be performed without shutting down the system or interrupting users.

Windows 2000 Resource Kit Cluster Tool Utility

Backup Cluster
The Cluster Tool Configuration Backup Wizard allows you to backup the current configuration for an NT Cluster.

The configuration data that is backed up includes the Groups, Resources and Resource Types for the cluster. The backup data is stored in XML format, allowing flexible access and reporting capabilities.

Restore Cluster
The Cluster Tool Configuration Restore Wizard allows you to restore the configuration for an NT Cluster from a backup file created with the Configuration Backup Wizard.

The configuration data that is restored may include the Groups, Resources and Resource Types for the cluster. Each item may be individually selected for restore. Additionally, the state in which resources are restored (online, offline, original state) may be selected.

Migrate Resources
The Cluster Tool Migration Wizard helps simplify the process of moving cluster resources from a stand-alone NT Server (single-node cluster) to an NT cluster (two-node cluster).

Future Directions

As Windows-based products evolve, the future development of cluster service will focus on the following important areas:

- Certification and support for even larger multi-node cluster configurations.
- Easier installation and verification of cluster configurations, including support for new types of hardware.
- Simpler, more powerful management of cluster-based applications and services, including continued focus on scripted, remote and "lights out" management.
- Extension of cluster-based availability and scalability benefits to more system services.
- Tighter integration of the infrastructure and interfaces of all Windows-based clustering technologies to enhance performance, flexibility and manageability.

MODELS

Compaq *ProLiant* Cluster/S100 Kit
340673-B21
340673-291 (Japan)

Compaq *ProLiant* Cluster/S100 provides high availability for applications and data in business-critical environments. The Compaq *ProLiant* Cluster is a hardware and software solution designed for any customer installation with significant uptime requirements.

Description

The Compaq *ProLiant* Cluster/S100 is a two-node cluster solution which provides high availability for applications and data in a business-critical environment.

The Compaq *ProLiant* Cluster/S100 takes full advantage of Compaq's industry-leading server products and SCSI-based storage systems in conjunction with Compaq SMART Array controllers for hardware RAID. The Compaq *ProLiant* Cluster/S100 utilizes the proven technology inherent in the Recovery Server Option (RSO) switch along with Microsoft Cluster Server (MSCS) software and Compaq enabling software to provide this high availability cluster solution.

The Compaq *ProLiant* Cluster/S100 Kit provides the essential software, cabling and documentation components required to cluster and manage a wide range of servers, with SMART Array controllers and external storage systems in a highly available configuration, managed by MSCS (ordered separately).

STANDARD FEATURES

High Availability at a Minimum Cost

The Compaq *ProLiant* Cluster/S100 provides high availability for applications and data in business-critical environments. Because it is built from industry-standard components, the Compaq *ProLiant* Cluster/S100 can be implemented at a much lower cost than traditional proprietary cluster solutions, without compromising availability. The Compaq *ProLiant* Cluster/S100 utilizes Compaq industry-leading servers, reliable Compaq *ProLiant* Storage Systems, proven RSO switches, and industry-standard cluster management software to provide an integrated high availability solution.

Investment Protection

Both existing and new Compaq servers and storage systems are certified in the Compaq *ProLiant* Cluster/S100 configurations. This means that customers can achieve high availability systems through clusters, which are assembled using their existing servers, storage and SMART Array controllers.

Integrated Solutions

Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq servers, interconnects, *ProLiant* Storage Systems, Microsoft and Compaq software, and Compaq integration documentation have all been thoroughly tested in cluster configurations. Compaq's close relationship with application partners has resulted in the development of cluster applications Technotes, which assist in the rapid deployment of major business applications and database products on Compaq *ProLiant* Clusters.

Cluster Management

The Compaq *ProLiant* Cluster/S100 configurations include Compaq Setup Utility software to enable rapid installation and cluster verification, configuration and deployment throughout the enterprise.

COMPONENTS

Compaq *ProLiant* Cluster/S100 Kit

12′ Ethernet Crossover Cable
Compaq *ProLiant* Cluster Series S Model 100 User Guide
Compaq *ProLiant* Cluster S/100 Software TechNotes:

- Microsoft IIS 3.0 Service Failover Using MSCS
- Microsoft IIS 3.0 Resource Failover Using MSCS
- Microsoft File and Print Service Failover Using MSCS

Servers
Most models of the Compaq *ProLiant* servers are supported in a cluster. Check **www.compaq.com/highavailability** for the latest list of supported servers, controllers and storage combinations.

Interconnect
Compaq strongly recommends the use of a private network connection between cluster nodes. In a two-node cluster, Compaq recommends Ethernet. An Ethernet crossover cable is included in the Compaq *ProLiant* Cluster/S100 Kit for this purpose. Compaq ServerNet is also an acceptable interconnect. See Compaq ServerNet PCI Adapter QuickSpec for ordering and detail information.

Storage
The following systems are supported in the *ProLiant* Cluster/S100 product configuration:
- Compaq *ProLiant* Storage Systems
- Compaq *ProLiant* Storage System/F1 with options
- Compaq *ProLiant* Storage System/U1 with options
- Compaq *ProLiant* Storage System RSO/U

Software
Compaq *ProLiant* Cluster/S100 Utility

Supported Operating Systems
Microsoft software is not included in the Compaq *ProLiant* Cluster/S100 Kit. Microsoft NT Server 4.0 Enterprise Edition is required. This product will not be supported under Windows 2000 operating system.

Options
Compaq Recovery Server Option (RSO) with firmware revision level 2.0B or later (U6, reference number 213812-001 on the ROM indicates an upgrade is needed) is a prerequisite to the installation of the Compaq *ProLiant* Cluster/S100. See server and storage options for the appropriate Recover Server Option for your particular storage system.

Order RSO 213817-001/291 for Model 1 Storage, 304117-B21/291 for F1 or U1. RSO is included with *ProLiant* Storage System RSO/U, 387256-001/291/B31 or 387490-002/292/B32.

Compaq SMART-2 (2e, 2p, SL, DH) Array Controllers with firmware 1.94 or later (Options ROMPaq 2.39 or later) are required.

Smart Array 3200 and Smart Array 221 are supported in like server configurations.

CL14U Cluster Rack Kit (347272-B21) is supported.

The Compaq *ProLiant* Cluster/S100 Kit
340673-B21
340673-291 (Japan)

Service and Support
Compaq servers and storage systems are protected by Compaq Services, including a three-year Limited Warranty[1], 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-Failure Warranty applies to certain hard drives, memory and processors for servers monitored by Compaq *Insight Manager 2.0* or higher. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

MODELS

Compaq *ProLiant* Cluster HA/F100 Kit
309816-B21
309816-291

Compaq *ProLiant* Cluster Solutions provide high availability for applications and data in a business-critical environment. The Compaq *ProLiant* Cluster HA/F100 is a hardware and software solution designed for any customer installation with significant uptime requirements.

Description

Compaq *ProLiant* Cluster solutions provide high availability for applications and data in a business-critical environment.

The Compaq *ProLiant* Cluster HA/F100 utilizes Compaq's industry-leading server products, Compaq Fibre Channel Storage, Ethernet or ServerNet interconnect and Compaq's leading installation and systems management utilities. The Compaq *ProLiant* Cluster HA/F100 utilizes the industry-standard Microsoft Cluster Server software for cluster operation and management. Compaq provides a cluster kit, which contains the components that are required to connect two Compaq servers and Compaq Fibre Channel Storage systems together in a cluster. This cluster kit contains documentation and an interconnect cable.

STANDARD FEATURES

Maximum Availability, Minimum Cost

Because they are built from industry-standard components, Compaq *ProLiant* Clusters deliver high levels of application availability at a much lower cost than traditional, proprietary cluster solutions. The Compaq *ProLiant* Cluster HA/F100 utilizes Compaq industry-leading servers, the fast and reliable Compaq Fibre Channel Storage System, cluster node interconnect options. Compaq installation and systems management software provide an integrated high-availability solution.

Investment Protection

Both existing and new Compaq servers are certified for Compaq *ProLiant* Cluster HA/F100 configurations. This means that customers can build clusters using existing servers or by mixing old and new servers.

Integrated Solutions

Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq servers, interconnect options, system management software and implementation documentation have all been thoroughly tested in cluster configurations. Compaq's close relationships with application partners have resulted in the development of cluster application TechNotes, which assist in the rapid deployment of major business applications on Compaq *ProLiant* Clusters.

Sophisticated Cluster Management

Compaq *ProLiant* Cluster HA/F100 configurations utilize Compaq SmartStart and Compaq Insight Manager software to enable rapid installation, configuration and deployment of clustered environments throughout the enterprise. Compaq SmartStart and Compaq Insight Manager are now "cluster enabled," allowing easy installation and sophisticated cluster and systems management, either locally or from remote locations. Compaq Cluster Verification Utility Software is used to verify the cluster configuration that is deployed.

COMPONENTS

Servers

Most models of the *ProLiant* Servers are supported in a cluster. Some server mixes are also supported. Check **www.compaq.com/highavailability** for the latest list of supported servers and server combinations.

Interconnect

Compaq strongly recommends the use of a private network connection between cluster nodes. An Ethernet crossover cable is included in the Compaq *ProLiant* Cluster HA/F100 Kit for this purpose. ServerNet interconnect technology is also supported in cluster configurations.

Storage

Compaq *ProLiant* Cluster HA/F100 configurations have been certified with Compaq *StorageWorks* RAID Array 4000 and RAID Array 4100 Storage Systems (RA4000/RA4100), previously know as Fibre Channel Storage Systems, as shared storage. Both seven-port and 12-port hubs are supported as the server to storage interconnect; however, a maximum of five RA4000 or RA4100 Storage Systems per cluster are supported.

Supported Operating Systems

Microsoft NT Server 4.0 Enterprise Edition (includes Microsoft Cluster Server)
(Certain restrictions and exclusions apply. Consult the Compaq *Customer Support Center at 1-800-345-1518 for details.)*
Microsoft Windows 2000 Advanced Server

Compaq *ProLiant* Cluster HA/F100 Kit

- 12′ Ethernet Crossover Cable
- Compaq *ProLiant* Clusters HA/F100 and HA/F200 Administrator Guide
- Compaq *ProLiant* Cluster HA/F100 for Microsoft Windows 2000 Advanced Server Supplement Guide
- Compaq Cluster Verification Utility software
- Compaq *ProLiant* Cluster HA/F100 Installation Poster

OPTIONS

All server options are supported, except Recovery Server Option and the Compaq *ProLiant* Cluster/S100 Kit. All Fibre Channel Storage options are supported, except long-wave GBICs.

Compaq *ProLiant* Cluster HA/F100 Kit	**309816-B21 or 309816-291 (Japan)**
CL14U Cluster Rack Kit	**347272-B21**
100′ Ethernet Cable	**309835-B21**
ServerNet Controller	**309833-B22**
ServerNet Cable (4m)	**309834-B23**
ServerNet Cable (30m)	**309834-B22**

Service and Support

Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty[1], 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-Failure Warranty[1] applies to certain hard drives, memory and processors for servers monitored by Compaq Insight Manager 2.0 or later. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

MODELS

Compaq *ProLiant*
Cluster HA/F200 Kit
380357-B21
38035az7-291 (Japan)

Compaq *ProLiant* Cluster HA/F200 provides a high level of availability in both cluster and single server attached environments. This kit offers cluster capability on virtually all of Compaq's entry, midrange and high-end *ProLiant* servers. The Compaq Redundancy Manager software for a dual loop with No Single Point of Failure server to storage interconnect configuration capability in a clustered environment and dual loop single server configuration, along with intelligent cluster management, make this solution a fit for applications and data access protection in a business-critical environment. The HA/F200 solution is designed for any customer installation with significant uptime requirements. The HA/F200 provides dual Fibre Loop access to storage, the Compaq Redundancy Manager (a load balancing routine for improved I/O), *Intelligent Cluster Administrator* for enhanced cluster administration and the Cluster Configuration Management-Cluster Verification Utility for cluster configuration assurance.

Description

Compaq *ProLiant* Cluster HA/F200 provides high availability for applications and data in business-critical NT environments, in both clustered and single server to storage configurations.

The Compaq *ProLiant* Cluster HA/F200 uses Compaq's industry-leading *ProLiant* servers, Compaq *StorageWorks* RAID Array 4000 (RA4000), (formerly the Fibre Channel Storage System), Ethernet or Netserver to server interconnect and Compaq's leading installation and systems management utilities.

The Compaq *ProLiant* Cluster HA/F200 uses the industry-standard Microsoft Cluster Server software included in Microsoft NT Server Enterprise Edition operating system for cluster operation and failover management. The addition of intelligent cluster management and improved systems management software included in the HA/F200 kit, make this solution a must for consideration in many Windows NT environments.

Compaq provides a cluster kit, which contains the components required to connect and manage two Compaq *ProLiant* servers and the *StorageWorks* RAID Array 4000/4100 together in a cluster. The HA/F200 cluster kit contains cluster enabling I/O management software, Web-based Cluster Administration software, cluster configuration verification software and an interconnect cable. Microsoft NT Server Enterprise Addition needs to be purchased separately from your local Microsoft provider.

The Compaq Cluster HA/F200 Kit also delivers a higher level of availability in single-server environments through the use of the Compaq Redundancy Manager, which supports multi-path I/O management software included in the kit. This software allows for a redundant fibre channel loop (two paths through the fibre channel hub) utilizing dual RAID controllers in the RA4000/4100 and dual HBAs in the servers connected to one or more RA4000/4100 storage subsystems. This can be implemented in a standard Microsoft NT environment and does not require Microsoft's Server Enterprise Addition operating system.

STANDARD FEATURES

Maximum Availability, Minimum Cost

Because Compaq *ProLiant* Clusters are built from industry-standard components, Compaq delivers high levels of application availability at a much lower cost than traditional, proprietary cluster solutions. The Compaq *ProLiant* Cluster HA/F200 uses Compaq industry-leading *ProLiant* servers, the fast and reliable Compaq *StorageWorks* RAID Array 4000 (RA4000) and cluster node interconnect options, ServerNet or Ethernet. The HA/F200's dual loop configuration will provide another level of availability above the previously released F100 cluster. The F100 will continue to be offered as a fibre channel entry-level solution for customers interested in a higher level of availability than single server environments provide, but choose not to move to the highest level of availability found in the HA/200 cluster configuration.

Investment Protection

Both existing and new Compaq *ProLiant* servers are certified for Compaq *ProLiant* Cluster HA/F200. This means that customers can build clusters using existing servers or by mixing old and new servers. Customers who purchased the F100 cluster from Compaq can easily migrate to this dual loop solution simply by purchasing the new HA/F200 kit and the additional hardware, hubs and cables to add the second loop to that configuration.

Integrated Solutions

All Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq *ProLiant* Servers, interconnect options, system management software, flexibility in configuration support and implementation documentation have all been thoroughly tested in cluster configurations. Compaq's close relationships with application partners have resulted in the development of cluster application TechNotes, which assist in the rapid deployment of major business applications on Compaq *ProLiant* Clusters.

Sophisticated Cluster Management

Compaq *ProLiant* Cluster HA/F200 configurations utilize Compaq *SmartStart*, Compaq *Insight Manager, and Insight Manager XE – Cluster Monitor*, and the *Intelligent Cluster Administrator* software offerings to enable rapid installation, configuration monitoring and management, as well as deployment of clustered environments throughout the enterprise. Compaq *SmartStart* and Compaq *Insight Manager* are now "cluster enabled," allowing easy installation and sophisticated cluster and systems management, either locally or from remote locations.

COMPONENTS

Servers

A broad range of *ProLiant* server models are supported in the HA/F200 configurations, in both matched-pair and mixed-pair combinations. Please check Compaq's High Availability Website for an up-to-date listing of certified HA/F200 configurations at **www.compaq.com/highavailability.**

Interconnect

Compaq strongly recommends the use of a private network connection between cluster nodes. An Ethernet crossover cable is included in the Compaq *ProLiant* Cluster HA/F200 Kit, and an extended 100 length Ethernet cable kit is available for this purpose. ServerNet interconnect technology is also supported in cluster configurations as an optional interconnect.

Storage

Compaq *ProLiant* Cluster HA/F200 configurations have been certified with the Compaq *StorageWorks* RAID Array 4000 (RA4000) as shared storage.

Supported Operating Systems

Microsoft NT Server 4.0 Enterprise Edition
(includes Microsoft Cluster Server)

Compaq *ProLiant* Cluster HA/F200 Kit Contents

12′ Ethernet Crossover Cable
Compaq *ProLiant* Cluster HA/F100 and HA/F200 Administrator Guide
Compaq Redundancy Manager software
Compaq *Intelligent Cluster Administrator* software
Compaq Cluster Verification Utility (CCVU) software
Compaq *ProLiant* Cluster HA/F200 installation poster
Compaq Multimedia CD Supplement

OPTIONS

Most server options and the *StorageWorks* RA4000/4100 options are supported in a cluster. Refer to the Compaq High Availability Website at **www.compaq.com/highavailability/** for ordering and configuration guides, release notes and restrictions.

Compaq *ProLiant* Cluster HA/F200 Kit	**380357-B21**
CL14U Cluster Rack Kit	**347272-B21**
100′ Ethernet Cable	**309835-B21**
ServerNet Controller	**309833-B22 or 309833-292 (Japan)**
ServerNet Cable (4 m)	**309834-B23**
ServerNet Cable (30 m)	**309834-B22**

Service and Support

Compaq servers and storage systems are protected by Compaq Services including a three-year Limited Warranty[1], 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-Failure Warranty[1]

applies to certain hard drives, memory and processors for servers monitored by Compaq *Insight Manager v2.0* or later. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

MODELS

Compaq *ProLiant* Cluster HA/F500

Enhanced Cluster DT Kit
164227-B21

Enhanced Cluster Kit
379937-B23

Basic Cluster Kit
103250-B23

Compaq *ProLiant* Cluster HA/F500 is Compaq's high-end Windows two-node cluster solution that provides the highest levels of availability in a No-Single-Point-of-Failure cluster environment. It is ideal for customers' business critical applications and data access requirements.

Compaq ProLiant Cluster HA/F500

The HA/F500 is a cluster platform comprised of *ProLiant* servers, *StorageWorks* storage, fibre channel interconnects and Compaq software, running on a Microsoft Windows Operating System. This solution offers high availability through the use of redundant hardware, software, and manageability tools found in the servers and storage subsystem. This platform features the highest level of availability, and can be configured to provide an environment with No-Single-Point-of-Failure.

The cluster kits support a variety of configurations including single path, dual path and disaster-tolerant configurations with a separation of up to 10 Km between nodes. The Basic Cluster Kit provides a base set of documentation and a crossover cable for those customers interested in utilizing their existing hardware. The Enhanced Cluster Kit provides both hardware and software pieces necessary to configure a no-single-point-of-failure cluster solution and the Enhanced Cluster Kit provides additional documentation and software that has been tested and certified to work. In the disaster tolerant configurations show later in this document. The complete solution consists of the following major components:

- Compaq *ProLiant* high-end and high-density Servers
- Compaq *StorageWorks* Optical Fibre Channel External RAID storage subsystem (RA or MA8000, ESA or EMA12000) with orderable controller software for hub or switch support
- Optical Fibre Channel interconnect hardware including (host bus adapters, cables and GBIC, hubs or switches)
- Microsoft Operating System (purchased separately)
- Compaq *ProLiant* Cluster HA/F500 Basic, Enhanced or Enhanced DT Kit.

Supported servers for this solution can be found at the High Availability section of Compaq's web site at **www.compaq.com/highavailability**. Some examples of the possible system configurations and connectivity block diagrams are found at the end of this document.

COMPONENTS AND FEATURES

Highest Level of Availability

The *ProLiant* Cluster HA/F500 is an integrated solution that consists of Compaq's high-end or high-density *ProLiant* servers, Compaq *StorageWorks* storage subsystems and fibre channel server-to-storage interconnect hardware and software. This unique combination of highly available server and storage hardware, redundant connections and the failover capabilities the Microsoft Windows Operating System gives the customer the highest levels of data access and applications availability found in the market today.

Highly Available Server Hardware

Compaq high-end and high density servers incorporate high availability and manageability features such as PCI hot plug, redundant hot plug fans, redundant processor power modules, redundant NIC support, hot plug dual port 10/100 TX NIC and redundant hot plug power supplies.

Redundant connections between server and storage system

The *ProLiant* Cluster HA/F500 solution supports dual fibre channel hubs or switches and dual active/active RAID controllers that provide total hardware redundancy when accessing the external storage subsystem. Dual host bus adapters (HBAs) installed in the servers support redundant server to storage connections, with static load balancing and auto failover in the event of an HBA or optical link failure, through Compaq Secure Path software, provided as part of the Enhanced and Enhanced DT cluster kits.

Single Point of Control Management through *Intelligent Cluster Administrator* and *Insight Manager XE - Cluster Monitor*

Compaq's Single Point of Control cluster monitor and administrator offerings provide enhanced cluster management and control via a Web-based interface. The *Intelligent Cluster Administrator* provides industry-leading Web-based cluster administration and cluster configuration management. Optionally, Insight Manager XE - Cluster Monitor provides common data collection and repository for your enterprise servers, desktops and clusters. The Cluster Monitor provides real time feedback on cluster events giving the administrator early warning capability for cluster monitor points to alleviate potential business-critical outages due to cluster failover.

- Failover capability of Windows Operating System
  The Windows Operating System provides for fully automated detection, failover and failback at the hardware and application level in order to maintain client access in the case of application or hardware failure. The operating system manages the components of the cluster as resources and maintains availability to the attached network. This functionality is provided via MSCS, the accepted industry-standard NT cluster management software. It is provided as part of Windows NT Enterprise Server (NTS/E) V 4.0 operating system, and via Microsoft Windows 2000 Advanced Server and Datacenter Server operating systems.

Highly Available Storage

The RA8000/ESA12000 or MA8000/EMA12000 storage systems from Compaq *StorageWorks* are high-capacity storage subsystems that supports single or redundant RAID controllers, which manage up to 72 disks per controller or controller pair. Each two-node cluster configuration, with the use of hubs (FC-AL), will support two pairs of controllers for a maximum raw storage capacity configuration of 2.621 TB per cluster (two storage subsystems x 72 drives x 18.2 GB= 2.261 TB). The use of fibre channel switches (FC-Switch), will support four pairs of controllers for a maximum raw storage capacity configuration of 5.242 TB per cluster (four storage subsystems x 72 drives x 18.2 GB= 5.242). Note that the capacities increase linearly as higher capacity drives are added and supported by the storage subsystems. The dual RAID controllers can also be configured to support two independent clusters on a single storage subsystem utilizing a hub (FC-AL), and up to four independent clusters utilizing the switch configuration (FC-Switch).

High availability storage components include redundant power supplies, fans, controllers, host bus adapters, cache battery backup, hot global spare drives and multi-level RAID architecture (0, 1, 3/5, 1+0).

High-performance features include 64 MB to 512 MB of Mirrored Write Back Cache memory per controller. The RA8000/MA8000 ships with 64 MB per controller installed and the ESA12000/EMA12000 ships with 256 MB per controller installed. Additional memory modules for upgrade to 512 MB can be ordered separately. A unique read-ahead algorithm can boost performance for certain applications by up to 40 percent. Dual host bus adapters, which are static load balanced, can also increase system performance.

Manageability through *StorageWorks* Command Console (SWCC) provides a graphical user interface (GUI) to setup/configure, monitor, and troubleshoot storage subsystems. The SWCC works in conjunction with Compaq Insight Manager (CIM) and SmartStart.

Disaster Tolerant solutions, through the use of Data Replication Manager (DRM) software, can take advantage of the fibre channel switch fabric to implement: (a) Online, real-time local and remote data replication, (b) Data replication to remote site at fibre channel speeds (100 MB/sec), (c) Synchronous mirroring, (d) Remote site at extended distances (10 Km using the enhanced DT cluster configuration) from production data center, (e) Host independent copy operation, (f) Automatic or manual recovery from site failure.

Based on Industry-Standard Hardware and Software

Because the solution is based on industry-standard hardware, it can be implemented at a much lower cost than proprietary RISC/UNIX-based cluster solutions without compromising availability.

Investment Protection

The HA/F500 combines state-of-the art technology and industry standards-based components to provide a long-term business solution. The implementation of fibre channel arbitrated loop (FC-AL) and fibre channel switch (FC-Switch) technology, in conjunction with high speed processing power and system and network interconnects, position the user to take advantage of new software that optimizes the capabilities of much of the current hardware. The Compaq *ProLiant* Cluster HA/F500 can be configured utilizing many of the standard servers and components that are available now and may already be on-site. This means that customers can achieve high availability systems through clusters that they have assembled using their existing servers.

Integrated Solutions

Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq *ProLiant* Servers, Fibre Channel Hub and Switch Interconnect, Compaq *StorageWorks* storage Systems, Microsoft and Compaq Software and Compaq integration documentation have all been thoroughly tested in cluster configurations. Compaq's close relationship with application partners resulted in the development of cluster applications certified for use on Compaq hardware. Compaq uses this expertise to help customers design cluster configurations that will meet their business needs and application requirements.

The Compaq *ProLiant* Cluster HA/F500 Enhanced DT Kit

- Compaq *ProLiant* Cluster HA/F500 Administrator Guide
- Compaq *ProLiant* Cluster HA/F500 Disaster Tolerant Supplement
- Compaq *ProLiant* Cluster HA/F500 Enhanced DT Poster
- Compaq Cluster Verification Utility (CCVU) Software
- Web-based *Intelligent Cluster Administrator* Software
- Compaq *StorageWorks* Secure Path software
- Kit contents card

The Compaq *ProLiant* Cluster HA/F500 Enhanced Kit

- 12′ Ethernet Crossover Cable
- Compaq *ProLiant* Cluster HA/F500 Administrator Guide
- Compaq *ProLiant* Cluster HA/F500 Enhanced Configuration Poster
- Compaq *ProLiant* Cluster HA/F500 Basic Configuration Poster
- Compaq Cluster Verification Utility (CCVU) Software
- Web-based *Intelligent Cluster Administrator* Software
- Compaq *StorageWorks* Secure Path software
- Kit contents card

The Compaq *ProLiant* Cluster HA/F500 Basic Kit

- 12′ Ethernet Crossover Cable
- Compaq *ProLiant* Cluster HA/F500 Administrator Guide
- Compaq *ProLiant* Cluster HA/F500 Basic Configuration Poster
- Compaq Cluster Verification Utility (CCVU) Software
- Kit contents card

Microsoft Windows NT Enterprise Server (NTS/E) or Microsoft Windows 2000 Advanced Server or Datacenter Server Operating System

Microsoft Operating System is required and is not included in the cluster kit (order directly from Microsoft or qualified distributor). The Basic and Enhanced clusters are supported under both Windows NT, Windows 2000 Advanced Server and Datacenter Server. Enhanced DT clusters are currently only supported under Windows NT.

Servers
Compaq *ProLiant* Cluster HA/F500 is intended for use with Compaq *ProLiant* high-end or high-density servers. For supported configurations, refer to the Compaq High Availability web site at **www.compaq.com/highavailability**.

Heartbeat Interconnect
Compaq strongly recommends the use of a private network connection between cluster nodes. A 12′ Ethernet crossover cable is included in the Compaq *ProLiant* Cluster HA/F500 Kits for this purpose. An optional 100′ Ethernet cable (309835-B21) can be purchased separately. Note that this cable is not included in the Enhanced DT Kit, as the server to server interconnect is via an extended distance.

Required Base Fibre Channel Cluster Component (Options)
Fibre Channel PCI HBA (380574-001)
Fibre Channel Storage Hub 7 (234453-001/-B31/-291) or Hub 12 (242795-B21)
Fibre Channel Switch 16 port (158223-B21) or eight port (158222-B21). See switch documentation for features and functionality of each of these switches and hubs.
Required purchase of one OS software/platform kit 380551-001 with the storage system per cluster. Disaster tolerance required the DT Software Platform kit for NT/Intel (128696-B21.)

Controller Software
ACS V8.5F kit for switched or FC-AL fibre channel (128697-B21)
ACS V8.5P kit for Data Replication (128698-B21).
Fibre Cable Options (234457-B21/-B22/-B23/-B24/-B25)

(See individual QuickSpecs for detailed-information relative to each of the above basic cluster components.)

The Compaq *ProLiant* Cluster HA/F500 Kit

Enhanced DT Kit	164227-B21	(International)
Enhanced Kit	379937-B23	(International)
Basic Kit	103250-B23	(International)

Service and Support

Compaq servers and storage systems are protected by Compaq Services, including a three-year limited warranty[1], 7 x 24 hardware technical phone support and on-line support through CompuServe, Prodigy, America Online and the Internet. Pre-Failure Warranty applies to certain hard drives, memory and processors for servers monitored by Compaq Insight Manager 2.0 or later. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider. Compaq has recently released extended services specific to High Availability such as Consulting Services, Business Critical Services, etc. Compaq maintains a global services and support network, which includes a large number of NT trained persons throughout the world. Go to Compaq's Web site for more detailed information.

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

HA/F500 Single Loop

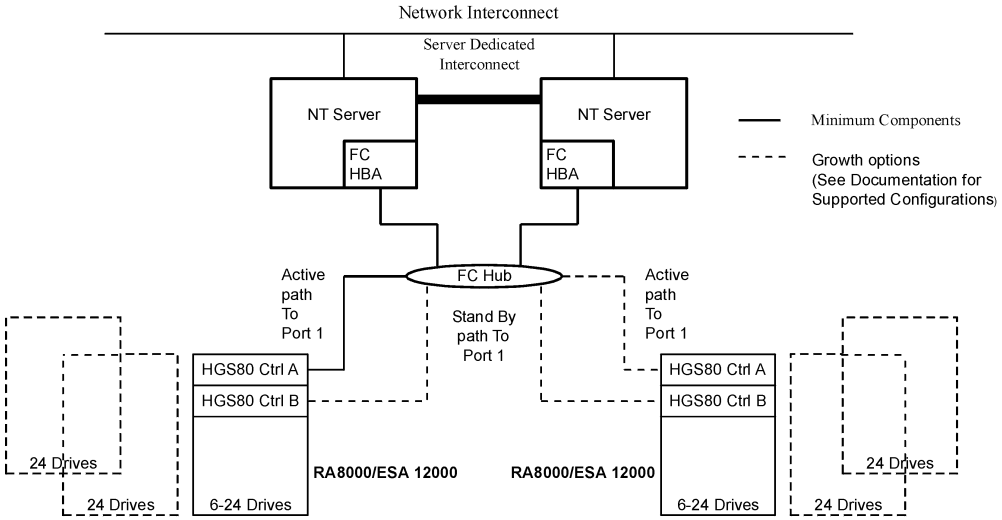

HA/F500 Double Loop NSPOF

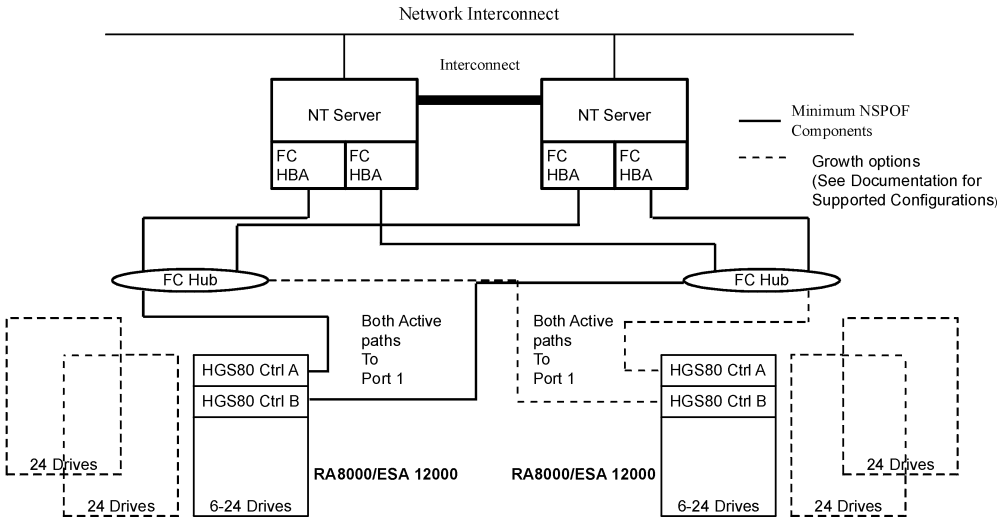

HA/F500 Dual Cluster Configuration

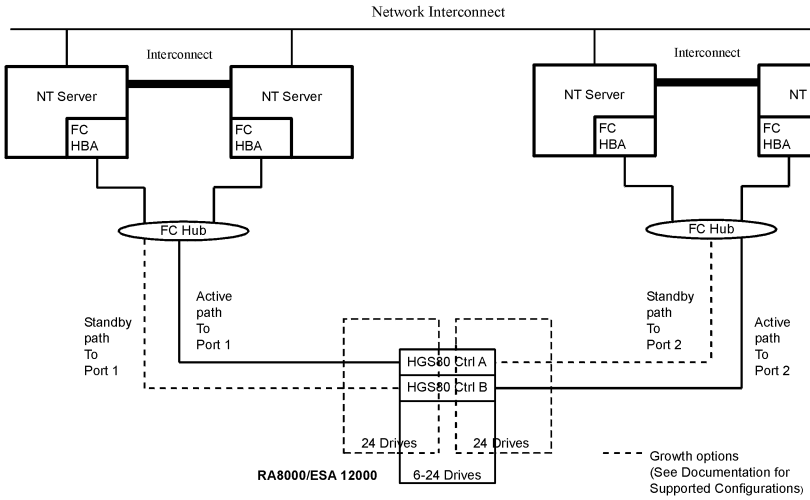

HA/F500 Dual NSPOF Cluster Configuration

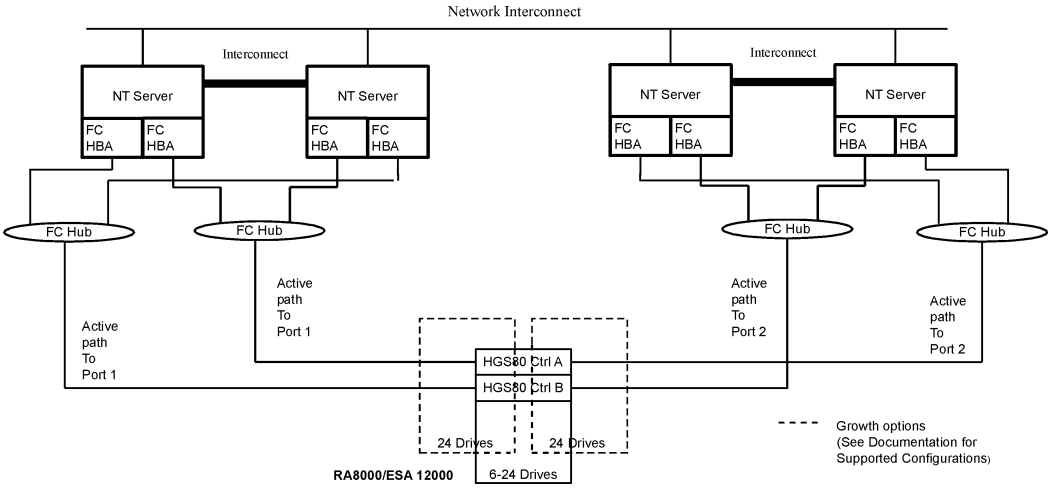

Switch Configurations

HA/F500 Two Independent Clusters

Single Bus

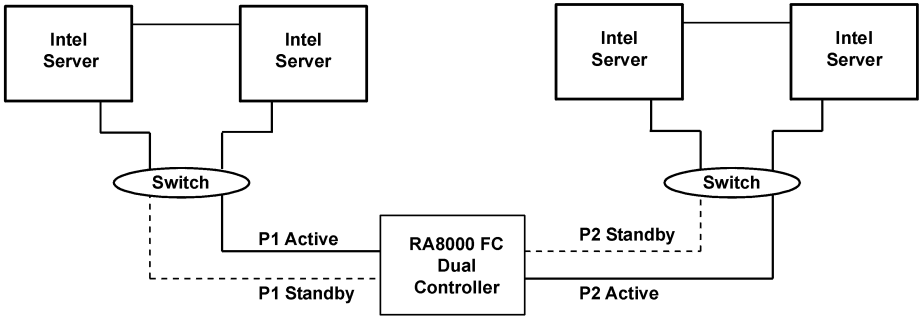

- Up to 16 LUNs/8 per Cluster
- Up to 72 Disks
- Up to 500 meters per segment

- Dual redundant Controllers, Transparent Failover Mode
- ACS 8.4F

Multiple Bus - NSPOF

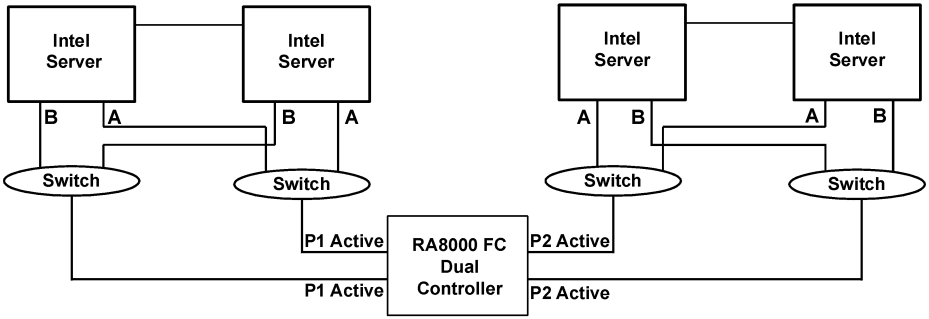

- Up to 16 LUNs/8 per Cluster
- Up to 72 Disks
- Up to 500 meters per segment

- Dual redundant Controllers, Multiple-bus Failover Mode
- NT Multi-path driver Secure path v2.1
- ACS 8.4F

MSCS Intel Cluster - 4 Controllers, Multiple Bus - NSPOF

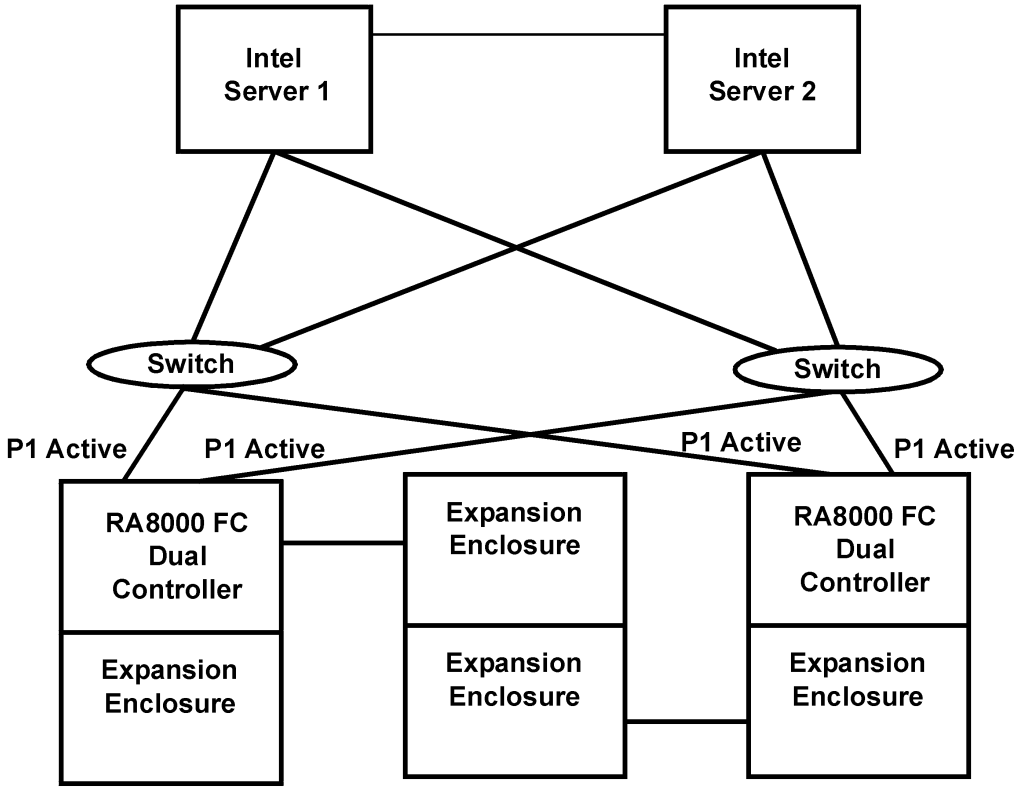

- Up to 16 LUNs
- Configure 4 LUNs/Adapter for load balancing
- Up to 144 Disks
- Up to 500 meters per segment

- Dual redundant Controllers, Multiple-bus Failover Mode
- NT Multi-path driver Secure path v2.1
- Uses Port 1 of each controller/4 LUNs /port
- ACS 8.4F

Dual-Switch Multi-Host

"Four Cluster" Configuration
Dual-Loop Multi-Host Shared Storage
Multiple Bus, 4 Controllers - NSPOF

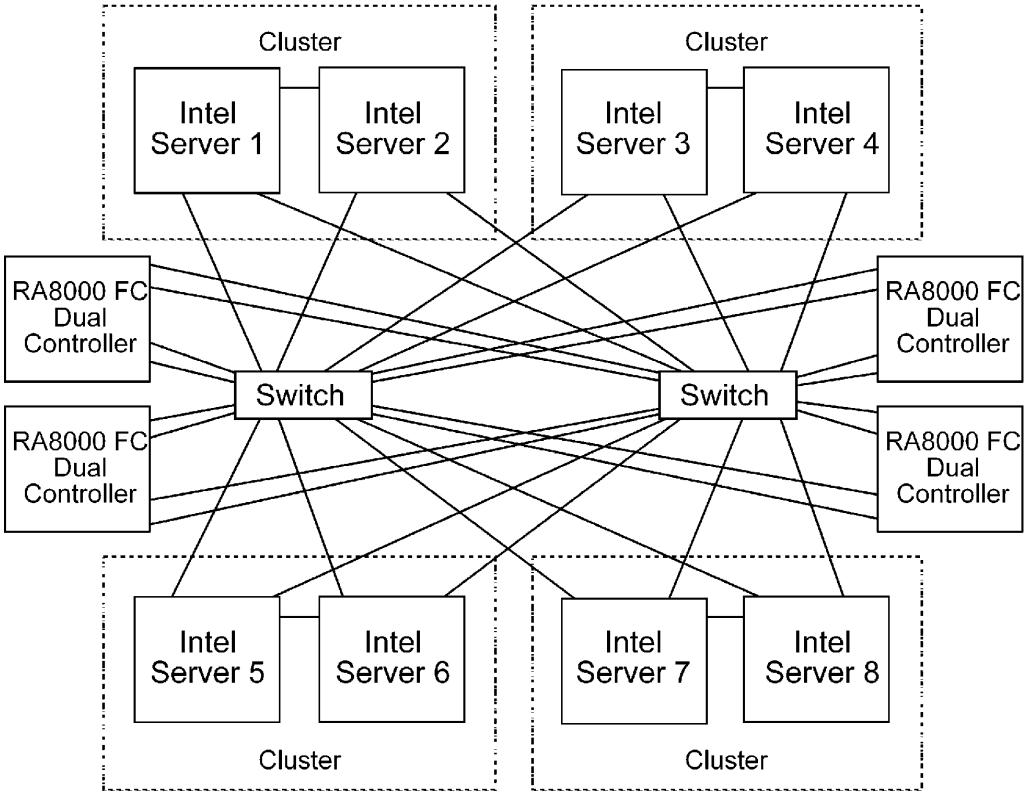

- Up to 16 LUNs/ per Cluster
- Up to 288 Disks
- Up to 500 meters per segment

- 4 Dual Redundant Controller Pairs Mode
- Multiple-Bus Failover Mode
- NT Multi-path driver Secure path
- ACS

HA/F500

Enhanced Disaster tolerance
Basic Configuration

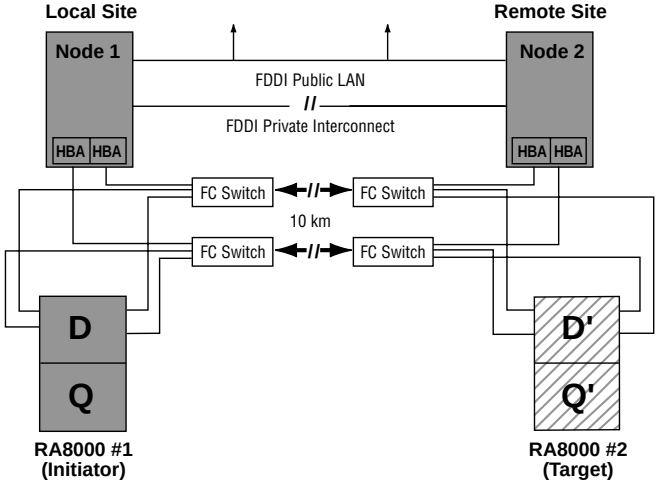

Compaq Parallel Database Clusters

Power for the Data Center!

COMPAQ

www.compaq.com/highavailability

Compaq Parallel Database Clusters

Product Line/Speaker Notes

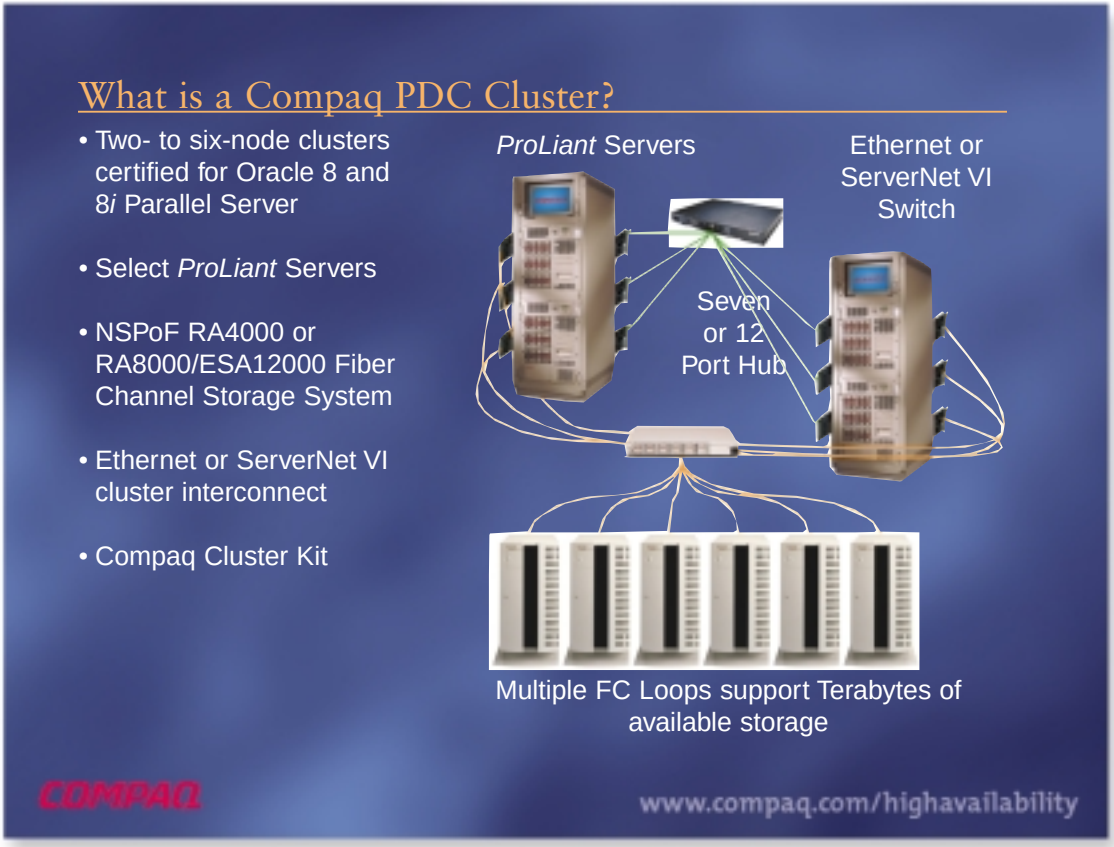

The Compaq Parallel Database Cluster for Oracle8*i* Parallel Server delivers high availability and scalability of Oracle database environments, using industry-standard components and software. It provides significantly improved uptime and lower operating costs. Since the PDC cluster solution is based on industry-standard hardware, it can be implemented at a much lower cost than RISC/UNIX-based cluster solutions without compromising availability, making it the perfect solution for business critical applications such as data warehousing, Web back-end databases, custom database application development, packaged applications and database migration from other platforms.

The Compaq Parallel Database Cluster for Oracle8*i* Parallel Server allows a two- to six-node cluster of Windows NT-based Compaq servers to simultaneously share a single Oracle database. Certified for Oracle8*i* Server Release and Oracle8*i* Parallel Server option, the PDC cluster integrates Compaq's unique Operating System Dependent (OSD) software modules with Compaq's industry standard, fault-resilient hardware:

- Compaq *ProLiant* Servers (Please refer to the PDC Certification Matrix for supported *ProLiant* Server models in the PDC configurations)
- PDC clusters support the Compaq *ProLiant* RA8000/ESA12000 or the RA4000 Storage Subsystems as shared storage

- PDC clusters support either Ethernet or ServerNet VI dedicated cluster Interconnects

Compaq provides an Enhanced Cluster Kit for the PDC cluster configurations, which contains the Compaq OSD modules for Oracle Parallel Server Release 8.0.5 and 8.1.5, and an administrator guide.

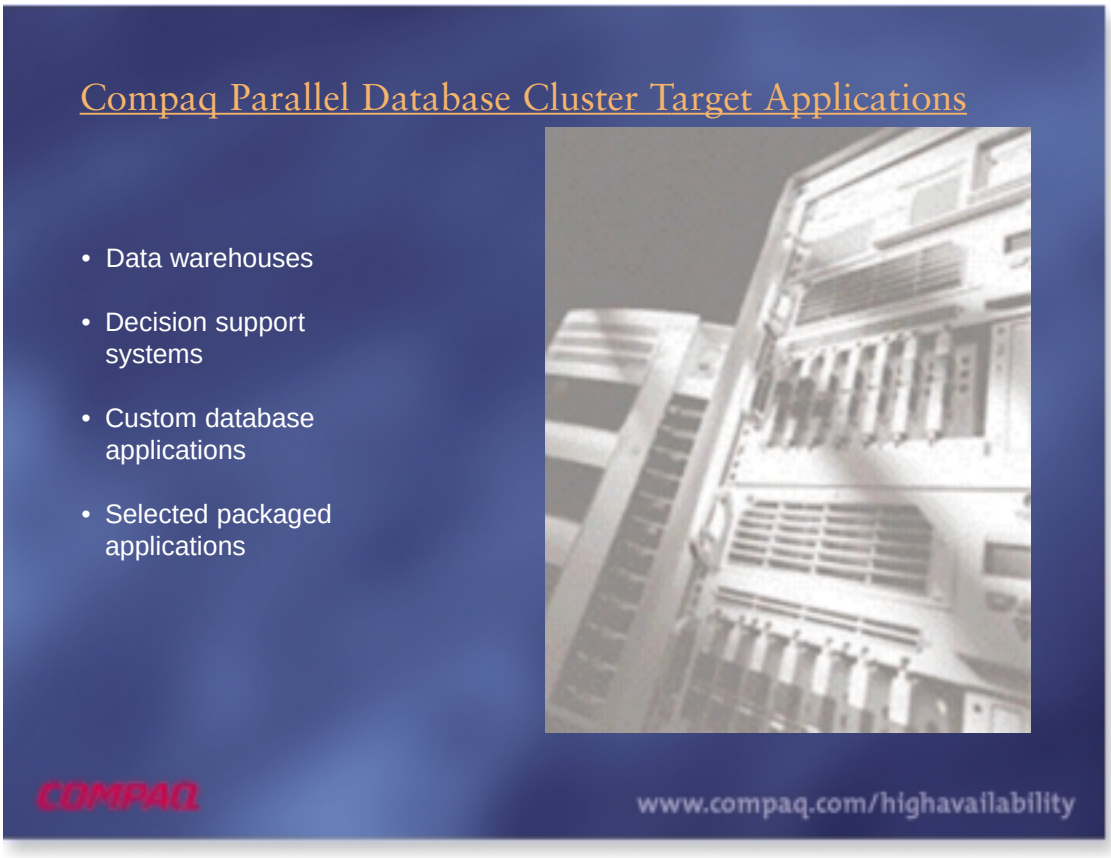

Compaq PDC Cluster Target Applications:

Oracle8*i* is designed for the applications listed above.

Oracle8*i* Parallel Server is a good fit for these applications if:

- High availability is a priority.
- The databases are mostly read-only with minimal update writes.
- The data is highly partitioned.

Feb. 11th, 2000, Compaq Computer Corporation submitted the record-breaking TPC-C benchmark result on the six-node Compaq Parallel Database Cluster Model PDC/O2000, with the new 8-way Profusion-based *StorageWorksProLiant* servers and Compaq Fibre Channel Storage Subsystem using Oracle8*i* Enterprise Server V8.1.6 with the Parallel Server option.

The 101657.17 tpmC number, achieved with a six-node *ProLiant* 8500 cluster, Fibre Channel Storage and Oracle's advanced Parallel Server technology, crushed the previous Compaq record of 99274.90 tpmC. This is the first time a Windows/NT platform has broken through the 100,000 tpmC barrier.
By taking a position in the top ten tpmC postings for all operating systems, this dramatic performance breakthrough on Windows/NT reaffirms Compaq and Oracle's combined leadership in providing reliable, highly-scalable and cost-effective parallel database, clustering and SMP technology under both Windows/NT and Unix.

More detailed information and a performance brief about this benchmark is available at the following Compaq URLs:

www.compaq.com/products/servers/benchmarks/index.html

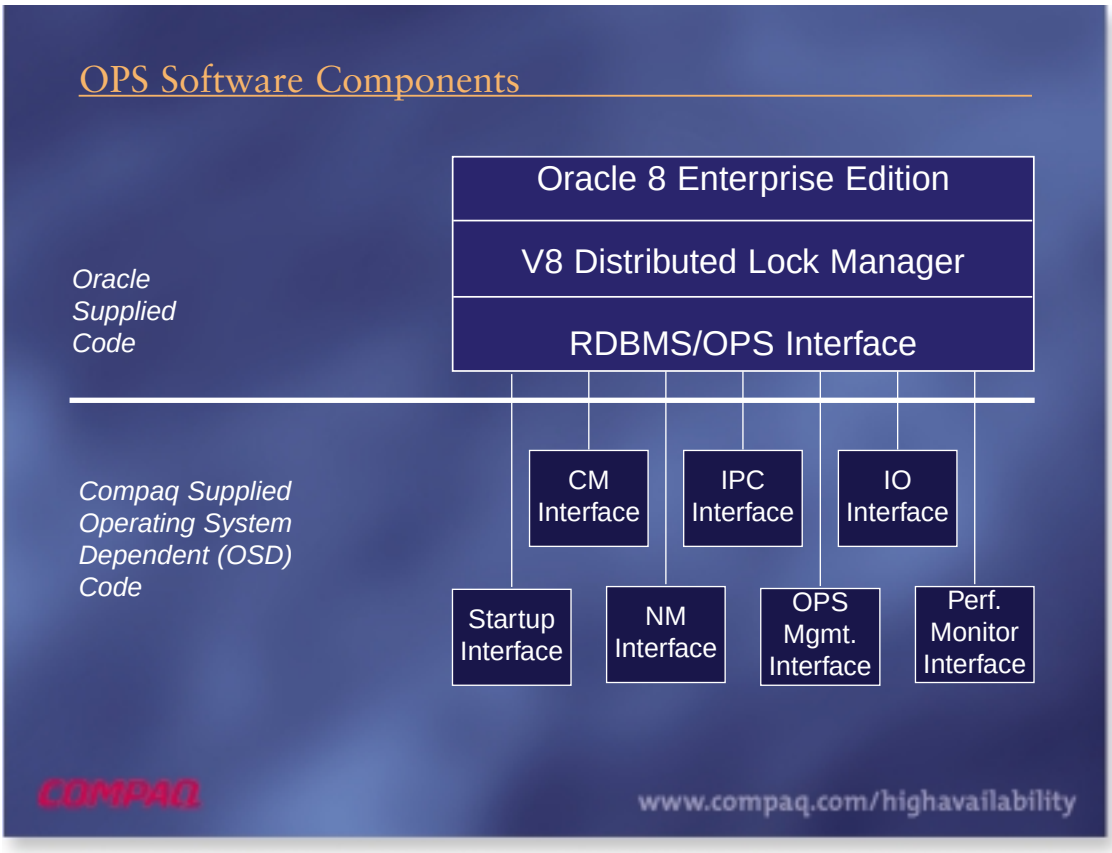

Oracle8*i* Parallel Server, when used as an add-on to Oracle8*i* Enterprise Edition for Windows NT, provides scalable and reliable application availability across multiple nodes in an NT cluster. It delivers the availability, scalability, and performance needed in today's business-critical online transaction processing systems.

With Oracle8*i* Parallel Server, a single database is accessed by multiple nodes in the cluster. Each node (a Windows NT server) runs an Oracle instance operation against a common database. The database resides on shared disks physically accessible by all nodes in the cluster. In the event of a node failure, Oracle8 Parallel Server eliminates the failed node from the cluster with minimal impact on surviving nodes.

What Are Compaq Oracle8 Parallel Server OSDs?

Oracle8*i* Parallel Server Operating System Dependent (OSD) modules are dynamic link libraries (DLLs) that implement Application Program Interfaces (APIs) specified by Oracle. The OSD software is distinct from the Oracle software. OSDs are required to host Oracle8*i* Parallel Server under the Windows NT operating system on a hardware vendor's platform. OSD are also required for the proper operation of the cluster. Oracle requires that each hardware vendor

making Oracle8 Parallel Server available on that vendor's platform develop OSD components specific to the platform. Hardware vendors also must ensure that their OSD modules adhere to Oracle's Oracle8 Parallel Server cluster interface specifications.

Oracle 8*i*, the database for Internet computing, changes the way information is managed and accessed to meet the demands of the Internet age, while providing significant new features for traditional applications. Oracle 8*i* Parallel Server Release 8.1.6 is the clustered option of the Oracle8i database based on shared-disk cluster architecture.

Key features of Oracle 8*i* Parallel Server Release 8.1.6 include:

- Scalability and performance architecture improvements
- Networking enhancements
- Back-up/restore/recovery improvements
- Administration improvements

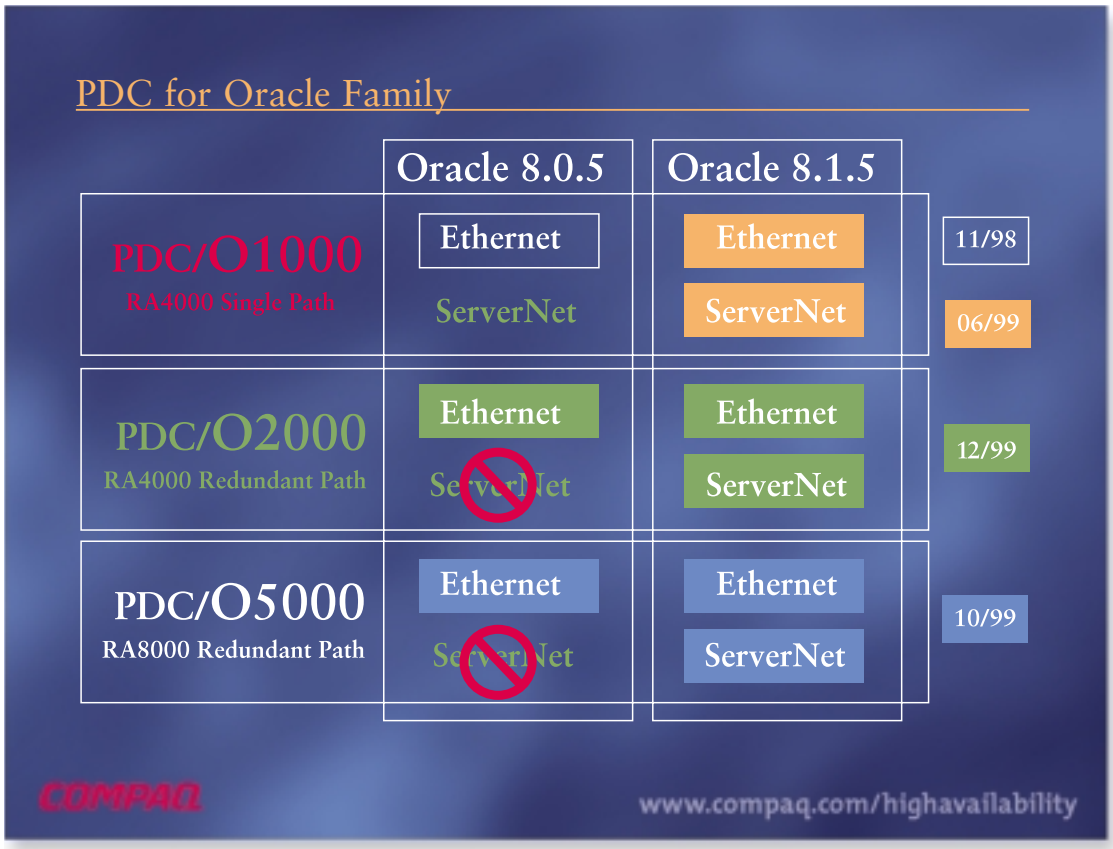

The primary difference between the PDC/O1000, PDC/O2000 and PDC/O5000 is the storage subsystem utilized. The PDC/O1000 and PDC/O2000 utilize the Compaq RAID Array RA4000 Storage Subsystem for optimal price performance.

The PDC/O1000 configuration is limited to a single IO path per storage array, creating a single point of failure.

The PDC/O2000 configuration supports multiple redundant IO paths per storage array creating a No Single Point of Failure configuration.

The PDC/O5000 is the high-end performance and scalability configuration which utilizes the RA8000/ESA12000 Storage System. This configuration features fully redundant IO paths allowing for a No Single Point of Failure implementation.

The PDC/O1000, PDC/O2000 and PDC/O5000 for OPS 8.0.5 support the Ethernet cluster interconnect only.

The PDC/O100, PDC/O2000 and PDC/O5000 for OPS 8.1.5 support either Ethernet and ServerNet VI as cluster interconnect.

Product Line/Speaker Notes

This slide provides detailed features and product description of the Compaq PDC/O1000 and PDC/O2000 clusters.

This slide provides detailed features and product description of the Compaq PDC/O5000 cluster.

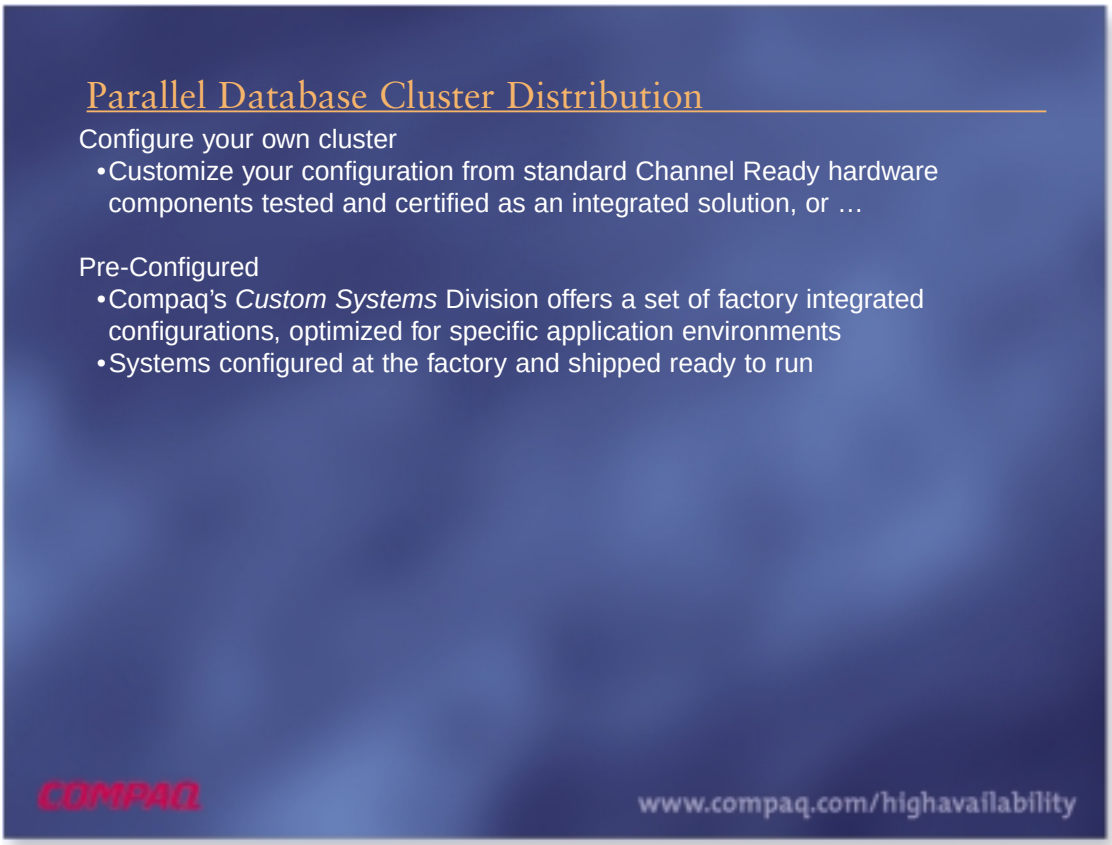

Customers can order the Compaq PDC cluster components and specific Compaq PDC Cluster Kit from channel partners or resellers or Compaq Custom Systems.

Compaq Custom Systems can deliver a pre-configured PDC/O5000 cluster to a customer or channel location. For additional hardware and software configuration information, contact your local Custom Systems representative. For a complete understanding of Custom Systems capabilities and list of local offices see the following web location:

www.digital.com/customsystems

The Compaq Customer Services organization has participated in the definition and delivery of supporting services for the PDC product. Please call 1-800-289-9052 to request the details on available services for the hardware, the software and the clusters. You may also access the Web at **www-mcs.shr.dec.com/release/allcsps.htm** for the associated release material on services.

Compaq Custom Systems can deliver a pre-configured PDC cluster to a customer or channel location. For additional hardware and software configuration information, contact your local Custom Systems representative. For a complete understanding of Custom Systems capabilities and list of local offices see the following web location: **www.digital.com/customsystems**

MODELS

Compaq Parallel Database Cluster Model PDC/O1000 for Oracle8*i* Parallel Server Release 8.1.5
PDC/O1000
Ethernet Kit – Basic
103424-B21
103424-291 (Japan)
PDC/O1000
***ServerNet* VI Kit – Basic**
103425-B21
103425-291 (Japan)
Compaq Parallel Database Cluster Model PDC/O1000 for Oracle8 Parallel Server Release 8.0.5
PDC/O1000
Ethernet Kit – Basic
382120-B21
382120-291 (Japan)

Standard Features

COMPAQ PARALLEL DATABASE CLUSTER MODEL PDC/O1000 FOR ORACLE8 PARALLEL SERVER

Compaq Parallel Database Cluster solutions provide high availability, scalability, and manageability for applications and data in a business-critical environment. The Compaq Parallel Database Cluster Model PDC/O1000 for Oracle Parallel Server is a hardware and software configuration specifically designed and tested for use with the Oracle Server and Oracle Parallel Server option in support of database opportunities such as data warehousing, Web back-end databases, custom database application development, packaged applications and database migration from other platforms.

Description

The Compaq Parallel Database Cluster Model PDC/O1000 for Oracle Parallel Server (PDC/O1000) allows a two- to six-node cluster of Windows NT-based Compaq servers to simultaneously share a single Oracle database. Certified for Oracle8 Parallel Server option Release 8.0.5 and Oracle8i Parallel Server option Release 8.1.5, PDC/O1000 integrates Compaq's unique Operating System Dependent (OSD) enabling software with Compaq's industry standard, fault-resilient hardware:

- Compaq *ProLiant* Servers (Please refer to the PDC/O1000 Certification Matrix for supported *ProLiant* Server models in the PDC/O1000 cluster configurations)
- Compaq RAID Array RA4000 as shared storage
- Ethernet or ServerNet VI dedicated cluster Interconnects

Compaq provides an Ethernet Cluster Kit or a ServerNet VI Cluster Kit for the PDC/O1000 configurations which contain the Compaq OSD modules for Oracle Parallel Server and an administrator guide.

The PDC/O1000 configuration supports multiple Fibre Channel Arbitrated Loops (FCAL) with Short Wave GBICs and multi-mode Fibre optic cabling.

Maximum Availability, Minimum Cost

The PDC/O1000 exploits fully redundant and optimized Windows NT Server-based Compaq *ProLiant* servers and other highly reliable, industry-standard hardware. It provides significantly improved uptime and lower operating costs than multiple single-server or proprietary database implementations.

Increased Throughput and Simplified Management with Oracle Software

All nodes within the PDC/1000 Cluster share a single Oracle database. Oracle Parallel Server software increases throughput to the Oracle database by coordinating each server's access to the data and allows the entire cluster to be managed from a single workstation – providing increased performance and simplified management of the database.

Specialized Capabilities for Clusters

The PDC/O1000 utilizes either Ethernet or ServerNet VI technology as the dedicated cluster Interconnect.
The low-latency, high-bandwidth ServerNet VI technology provides high performance for distributed lock manager response and dedicated cluster interconnect. In addition, ServerNet VI technology offers a low-overhead protocol stack, which offloads processor resources. These factors allow ServerNet VI technology to meet the high-speed, low-latency requirement for an ideal Oracle Parallel Server cluster interconnect.

Automatic Online Recovery

In the event of a node failure, the remaining nodes can continue to have access to the database. Meanwhile, users on the failed node can switch to another node and continue processing. The remaining nodes automatically recover all of the transactions performed by the applications before the failure occurred, ensuring data integrity.

Data Integrity and Database Consistency

Each copy of the Oracle database is allocated its own mirrored, multi-segment redo logs that store transactions not yet written to the database. If a node or an Oracle copy fails, the remaining nodes use the redo log from the failed node to bring the database back to a consistent state.

HARDWARE CONFIGURATION ELEMENTS

- Compaq *ProLiant* Servers (two to six nodes of the same model. Please refer to the PDC/O1000 and Certification Matrix for supported *ProLiant* Server models in the PDC/O1000 cluster configurations.
- Compaq RAID Array RA4000 as shared storage
- 100BaseTX Ethernet Interconnect or ServerNet VI Interconnect.

COMPONENTS

Servers

The PDC/O1000 cluster supports two to six nodes of Compaq *ProLiant* server. Please refer to the PDC/O1000 Certification Matrix for supported *ProLiant* Server models in the PDC/O1000 configurations. PDC/O1000 Cluster must be identically configured, including:

- All the servers within a PDC/O1000 cluster must belong to the same server model.
- Minimum memory: Each *ProLiant* server in the PDC/O1000 Cluster must have a minimum of 256 MB of memory and a minimum of 512 KB L2 cache per processor.
- Balanced servers: The servers within a PDC/O1000 cluster must be balanced. That is, each server must have the same type and number of processors, the same amount of primary and secondary cache and the same type and amount of memory.
- Same server model: Using more than one server model in a PDC/O1000 cluster configuration is currently not supported.

Storage

The PDC/O1000 cluster configurations have been certified with Compaq RAID Array RA4000 as shared storage. This innovative storage solution from Compaq provides the capacity, scalability and performance needed for the explosive data growth in the enterprise. The Compaq *StorageWorks* Raid Array can scale from several gigabytes to over 1.6 terabytes per server slot allowing you to purchase only the storage capacity you need today and feel confident that you have enough bandwidth to easily expand in the future. For additional information on *StorageWorks* RAID Array 4000 please see:
www.compaq.com/products/storageworks/raidstorage/FibreChannelSSindex.html

The PDC/O1000 cluster configuration supports multiple Fibre Channel Arbitrated Loops (FCAL) with Short Wave GBIC and multi-mode Fibre optical cable. The RA4000s are connected using a seven-port or 12-port Compaq Fiber Channel Storage hub.

Cluster Interconnect

The PDC/O1000 cluster support the following cluster interconnects:

- Ethernet PCI adapters in conjunction with a 100BaseTX Ethernet hub/switch OR
- ServerNet VI Interconnect with ServerNet Switches.

Cluster Kits

The PDC/O1000 cluster requires the PDC/O1000 cluster Kit which contains:

- Compaq's OSD software components
- User installation and Administrators Guide for Cluster Specific components, including site preparation information.

Supported Operating Systems

The PDC/O1000 supports Windows NT Server Enterprise Edition and Windows NT Server 4.0 with Service Pack 3 or Service Pack 4.

Service and Support

Compaq Customer Services organization has participated in the definition and delivery of supporting services for the PDC/O1000 product. Please call 1-800-289-9052 to request the details on available services for the hardware, the software and the clusters. Or you may access the Web at **www-mcs.shr.dec.com/release/allcsps.htm** for the associated release material on services. Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty[1], 7 x 24 hardware technical phone support and on-line support through CompuServe, Prodigy, America Online and the Internet. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

Compaq Custom Systems

Compaq Custom Systems can deliver a pre-configured PDC/O1000 cluster to a customer or channel location. For additional hardware and software configuration information, contact your local Custom Systems representative. For a complete understanding of Custom Systems capabilities and list of local offices see the following web location:
www.digital.com/customsystems

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

Standard Features

MODELS

Compaq Parallel Database Cluster Model PDC/O2000 for Oracle8*i* Parallel Server Release 8.1.5
PDC/O2000 Ethernet Kit - Enhanced
162950-B21

PDC/O2000 *ServerNet* VI Kit - Enhanced
162951-B21

Compaq Parallel Database Cluster Model PDC/O2000 for Oracle8 Parallel Server Release 8.0.5
PDC/O2000 Ethernet Kit - Enhanced
162952-B21

COMPAQ PARALLEL DATABASE CLUSTER MODEL PDC/O2000 FOR ORACLE PARALLEL SERVER

Compaq Parallel Database Cluster solutions provide high availability, scalability, and manageability for applications and data in a business-critical environment. The Compaq Parallel Database Cluster Model PDC/O2000 for Oracle Parallel Server is a hardware and software configuration specifically designed and tested for use with the Oracle Server and Oracle8 Parallel Server option in support of database opportunities such as data warehousing, Web back-end databases, custom database application development, packaged applications and database migration from other platforms.

Description

The Compaq Parallel Database Cluster Model PDC/O2000 for Oracle Parallel Server (PDC/O2000) allows a two- to six-node cluster of Windows NT-based Compaq servers to simultaneously share a single Oracle database. Certified for Oracle8 Parallel Server option Release 8.0.5 and Oracle8i Parallel Server option Release 8.1.5, PDC/O2000 integrates Compaq's unique Operating System Dependent (OSD) enabling software with Compaq's industry standard, fault-resilient hardware:

- Compaq *ProLiant* Servers (Please refer to the PDC/O2000 Certification Matrix for supported *ProLiant* Server models in the PDC/O2000 cluster configurations)
- Compaq RA4000 Storage subsystem as shared storage.
- Ethernet or ServerNet VI dedicated cluster Interconnects

Compaq provides an Enhanced Ethernet Cluster Kit or an Enhanced ServerNet VI Cluster Kit for the PDC/O2000 configurations, which contain the Compaq OSD modules for Oracle Parallel Server, an administrator guide and configuration posters.

The PDC/O2000 configurations have been certified with Compaq RAID Array RA4000 as shared storage. All the active components of the Compaq RAID Array RA4000 are redundant and hot replaceable which provides high system availability. Additionally, Dual HBA fail-over support through the use of Compaq Redundancy Manager software provides redundant data paths from the host to ensure higher availability and No Single Point of Failure. The PDC/O2000 configuration supports multiple redundant Fibre Channel Arbitrated Loops (FCAL) with Short Wave GBICs and multi-mode Fibre optic cabling.

Maximum Availability, Minimum Cost

The PDC/O2000 exploits fully redundant and optimized Windows NT Server-based Compaq *ProLiant* servers and other highly reliable, industry-standard hardware. It provides significantly improved uptime and lower operating costs than multiple single-server or proprietary database implementations.

Increased Throughput and Simplified Management with Oracle Software

All nodes within the PDC/2000 Cluster share a single Oracle database. Oracle Parallel Server software increases throughput to the Oracle database by coordinating each server's access to the data and allows the entire cluster to be managed from a single workstation – providing increased performance and simplified management of the database.

Specialized Capabilities for Clusters

The PDC/O2000 utilizes either Ethernet or ServerNet VI technology as the dedicated cluster interconnect. The low-latency, high-bandwidth ServerNet VI technology provides high performance for distributed lock manager response and dedicated cluster interconnect. In addition, ServerNet VI technology offers a low-overhead protocol stack, which offloads processor resources. These factors allow ServerNet VI technology to meet the high-speed, low-latency requirement for an ideal Oracle Parallel Server cluster interconnect.

Automatic Online Recovery

In the event of a node failure, the remaining nodes can continue to have access to the database. Meanwhile, users on the failed node can switch to another node and continue processing. The remaining nodes automatically recover all of the transactions performed by the applications before the failure occurred, ensuring data integrity.

Data Integrity and Database Consistency

Each copy of the Oracle database is allocated its own mirrored, multi-segment redo logs that store transactions not yet written to the database. If a node or an Oracle copy fails, the remaining nodes use the redo log from the failed node to bring the database back to a consistent state. The Oracle Distributed Lock Manager ensure data consistency and supports multi-initiator access to the storage environment.

HARDWARE CONFIGURATION ELEMENTS

- Compaq *ProLiant* Servers (two to six nodes of the same model. Please refer to the PDC/O2000 Certification Matrix for supported *ProLiant* Server models in the PDC/O2000 cluster configurations.
- Compaq RAID Array RA4000 as shared storage
- 100BaseTX Ethernet Interconnect or ServerNet VI Interconnect.

COMPONENTS

Servers

The PDC/O2000 cluster supports two to six nodes of Compaq *ProLiant* server. Please refer to the PDC/O2000 Certification Matrix for supported *ProLiant* Server models in the PDC/O2000 configurations.

PDC/O2000 Cluster must be identically configured, including:

- All the servers within a PDC/O2000 cluster must belong to the same server model.
- Minimum memory: Each *ProLiant* server in the PDC/O2000 Cluster must have a minimum of 256 MB of memory and a minimum of 512 KB L2 cache per processor.
- Balanced servers: The servers within a PDC/O2000 cluster must be balanced. That is, each server must have the same type and number of processors, the same amount of primary and secondary cache and the same type and amount of memory.

Storage

The PDC/O2000 cluster configurations have been certified with Compaq RAID Array RA4000 as shared storage. This innovative storage solution from Compaq provides the capacity, scalability and performance needed for the explosive data growth in the enterprise. The *StorageWorks* RAID Array can scale from several gigabytes to over 1.6 terabytes per server slot allowing you to purchase only the storage capacity you need today and feel confident that you have enough bandwidth to easily expand in the future. For additional information on Compaq *StorageWorks* RAID Array 4000 please see:
www.compaq.com/products/storageworks/raidstorage/FibreChannelSSindex.html

The PDC/O2000 cluster configuration supports multiple Fibre Channel Arbitrated Loops (FCAL) with Short Wave GBIC and multi-mode Fibre optical cable. The RA4000s are connected using a seven-port or 12-port Compaq Fiber Channel Storage hub.

Cluster Interconnect

The PDC/O2000 cluster support the following cluster interconnects:

- Ethernet PCI adapters in conjunction with a 100BaseTX Ethernet hub/switch OR
- ServerNet VI Interconnect with ServerNet Switches.

Cluster Kits

The PDC/O2000 cluster requires the PDC/O2000 cluster kit, which contains:

- Compaq's OSD software components
- User installation and Administrators Guide for Cluster Specific components, including site preparation information
- PDC/O2000 Configuration Poster.

Supported Operating Systems

The PDC/O2000 supports Windows NT Server Enterprise Edition and Windows NT Server 4.0 with Service Pack 3, 4 or 5.

Service and Support

Compaq Customer Services organization has participated in the definition and delivery of supporting services for the PDC/O2000 product. Please call 1-800-289-9052 to request the details on available services for the hardware, the software and the clusters. Or you may access the Web at **www-mcs.shr.dec.com/release/allcsps.htm** for the associated release material on services.
Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty[1], 7 x 24 hardware technical phone support and on-line support through CompuServe, Prodigy, America Online and the Internet.

For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

Compaq Custom Systems

Compaq Custom Systems can deliver a pre-configured PDC/O2000 cluster to a customer or channel location. For additional hardware and software configuration information, contact your local Custom Systems representative. For a complete understanding of Custom Systems capabilities and list of local offices see the following Web location:
www.compaq.com/solutions/customsystems/proliant/database.html

1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 0for details.

MODELS

Compaq Parallel Database Cluster Model PDC/O5000 for Oracle8*i* Parallel Server Release 8.1.6
PDC/O5000
Cluster Kit – Enhanced
177927-B21

Overview

Compaq Parallel Database Cluster solutions provide high availability, scalability and manageability for applications and data in a business-critical environment. The Compaq Parallel Database Cluster Model PDC/O5000 for Oracle8*i* Parallel Server Release 8.1.6 is a hardware and software configuration specifically designed and tested for use with the Oracle8*i* Server Release 8.1.6 and Oracle8*i* Parallel Server Option Release 8.1.6 in support of database opportunities such as data warehousing, Web back-end databases, custom database application development, packaged applications and database migration from other platforms.

Description

The Compaq Parallel Database Cluster Model PDC/O5000 for Oracle8*i* Parallel Server Release 8.1.6 (PDC/O5000) allows a two- to six-node cluster of Windows NT-based Compaq servers to simultaneously share a single Oracle database. Certified for Oracle8*i* Server Release 8.1.6 and Oracle 8*i* Parallel Server Option Release 8.1.6, the PDC/O5000 cluster integrates Compaq's unique Operating System Dependent (OSD) software modules with Compaq's industry standard, fault-resilient hardware:

- Compaq *ProLiant* Servers (Please refer to the PDC/O5000 Certification Matrix for supported *ProLiant* Server models in the PDC/O5000 configurations)
- PDC/O5000 supports the Compaq *StorageWorks* Modular Array MA8000 and Enterprise Modular Array EMA12000 (MA8000/EMA12000). The PDC/O5000 also supports the Compaq RAID Array 8000 and Enterprise Storage Array 12000(RA8000/ESA12000) Storage Subsystems as shared storage
- PDC/O5000 support either Ethernet or ServerNet VI dedicated cluster interconnects

Compaq provides an Enhanced Cluster Kit for the PDC/O5000 configurations, which contains the Compaq OSD modules for Oracle Parallel Server and an administrator guide.
The PDC/O5000 supports multiple redundant Fibre Channel Arbitrated Loops (FC-AL) with up to two Compaq StorageWorks MA8000/EMA12000 or RA8000/ESA12000 Storage Subsystems per redundant FC-AL.

Maximum Availability, Minimum Cost

The PDC/O5000 exploits fully redundant and optimized Windows NT Server-based Compaq *ProLiant* servers and other highly reliable, industry-standard hardware. It provides significantly improved uptime and lower operating costs than multiple single-server or proprietary database implementations.
The PDC/O5000 configurations have been certified with Compaq StorageWorks MA8000 and EMA12000. All MA8000/EMA12000 models provide redundant cooling, N+1 power redundancy and environmental monitoring. Drives and most solution components are hot swappable. The MA8000/EMA12000 FC redundant architecture eliminates single points-of-failure from server to storage in clustered or single server configurations. For additional protection, dual host bus adapter (HBA) failover software is available on many platforms.
The PDC/O5000 configurations have also been certified with Compaq RA8000 and ESA12000 Storage Subsystems as shared storage. All the active components of the Compaq RA8000 and ESA12000 Storage Subsystem are redundant and hot replaceable which provides high system availability and No Single Point of Failure. Additionally, Dual HBA fail-over support through the use of Secure Path software provides redundant data paths from the host to ensure No Single Point of Failure.

Increased Throughput and Simplified Management with Oracle Software

All nodes within the PDC/O5000 Cluster share a single Oracle database. Oracle Parallel Server software increases throughput to the Oracle database by coordinating each server's access to the data and allows the entire cluster to be managed from a single workstation – providing increased performance and simplified management of the database.

Specialized Capabilities for Clusters

The PDC/O5000 utilizes either Ethernet or ServerNet VI technology as the dedicated cluster interconnect. The low-latency, high-bandwidth ServerNet VI technology provides high performance for distributed lock manager response and dedicated cluster interconnect. In addition, ServerNet VI technology offers a low-overhead protocol stack, which offloads processor resources. These factors allow ServerNet VI technology to meet the high-speed, low-latency requirement for an ideal Oracle Parallel Server cluster interconnect.

Automatic Online Recovery

In the event of a node failure, the remaining nodes can continue to have access to the database. Meanwhile, users on the failed node can switch to another node and continue processing. The remaining nodes automatically recover all of the transactions performed by the applications before the failure occurred, ensuring data integrity.

Data Integrity and Database Consistency

Each copy of the Oracle database is allocated its own mirrored, multi-segment redo logs that store transactions not yet written to the database. If a node or an Oracle copy fails, the remaining nodes use the redo log from the failed node to bring the database back to a consistent state.

HARDWARE CONFIGURATION ELEMENTS

- Compaq *ProLiant* Servers (two to six nodes of the same model. Please refer to the PDC/O5000 Certification Matrix for supported *ProLiant* Server models in the PDC/O5000 configurations)
- Compaq StorageWorks MA8000 and EMA12000 or Compaq RA8000 and ESA12000 Storage Subsystems
- 100BaseTX Ethernet interconnect or ServerNet VI interconnect.

COMPONENTS

Servers

The PDC/O5000 supports two to six nodes of Compaq *ProLiant* server. Please refer to the PDC/O5000 Certification Matrix for supported *ProLiant* Server models in the PDC/O5000 configurations.
Note that all the *ProLiant* servers in the PDC/O5000 Cluster must be identically configured, including:

- Minimum memory: Each *ProLiant* server in the PDC/O5000 Cluster must have a minimum of 256 MB of memory and a minimum of 512 KB L2 cache per processor.
- Balanced servers: The servers within a PDC/O5000 cluster must be balanced. That is, each server must have the same type and number of processors, the same amount of primary and secondary cache, and the same type and amount of memory.
- Server mixes are currently not supported in the PDC/O5000 Cluster configuration

Storage

The PDC/O5000 configurations have been certified with Compaq StorageWorks MA8000 and EMA12000. MA8000/EMA12000 transitions the RA8000/ESA12000 storage solutions to the new StorageWorks packaging, maintains the support of the HSG80 controller (and associated features) and adds the benefits of universal drive support, higher density of drives per enclosure,

flexible configurations and higher capacity per cabinet. All MA8000/EMA12000 models provide redundant cooling, N+1 power redundancy and environmental monitoring. Drives and most solution components are hot swappable. Each solution can be configured with dual redundant controllers that can operate in dual active mode. Each controller has a pair of FC host ports. In the event of a path failure, the controllers can automatically failover to the remaining path.
The MA8000/EMA12000 FC redundant architecture eliminates single points of failure from server to storage in clustered or single server configurations. For additional protection, dual host bus adapter (HBA) failover software is available on many platforms.
The PDC/O5000 configurations have also been certified with Compaq RA8000 and ESA12000 Storage Subsystems as shared storage. All the active components of the Compaq RA8000 and ESA12000 Storage Subsystem are redundant and hot replaceable which provides high system availability and No Single Point of Failure. Additionally, dual HBA fail-over support through the use of Secure Path software provides redundant data paths from the host to ensure No Single Point of Failure.

For additional information on the RA8000/ESA12000 Storage Subsystem please see:
www.compaq.com/products/storageworks/RA8000/RA8000index.html

The PDC/O5000 cluster configuration supports multiple redundant Fibre Channel Arbitrated Loops (FC-AL) with Short Wave GBIC and multi-mode Fibre optical cable. The RA8000/ESA12000 is a high capacity storage subsystem that supports redundant RAID controllers and up to 72 disks per controller pair. The RA8000/ESA12000 storage subsystem can scale from several GB to 2.621 TB per server slot.

Interconnect
The PDC/O5000 support the following cluster interconnects:

- Ethernet PCI adapters in conjunction with a 100 BaseTX Ethernet hub/switch or
- ServerNet VI Interconnect with ServerNet Switches.

Cluster Kits
The PDC/O5000 requires the PDC/O5000-8.1.6 Cluster Kit – Enhanced (177927-B21) which contains:

- Compaq's OSD software components
- Multi-path software (Secure Path) – Compaq StorageWorks Secure Path is a high availability software product providing continuous data access for the Compaq StorageWorks MA8000/EMA12000 or the RA8000/ESA12000. Secure Path effectively eliminates controllers, disk drives, interconnect hardware and host bus adapters as single points of failure
- User installation and Administrators Guide for Cluster Specific components, including site preparation Information

Supported Operating Systems
The PDC/O5000 supports Windows NT Server Enterprise Edition and Windows NT Server 4.0 with Service Pack 3 or Service Pack 4.

Service and Support
Compaq Customer Services organization has participated in the definition and delivery of supporting services for the PDC/O5000 product. Please call 1-800-289-9052 to request the details on available services for the hardware, the software and the clusters. Or you may access the Web at
www-mcs.shr.dec.com/release/allcsps.htm
for the associated release material on services.
Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty[1], 7 x 24 hardware technical phone support and on-line support through CompuServe,

Prodigy, America Online and the Internet.

For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

Note: 1 Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

Compaq Custom Systems

Compaq Custom Systems can deliver a pre-configured PDC/O5000 cluster to a customer or channel location. For additional hardware and software configuration information, contact your local Custom Systems representative. For a complete understanding of Custom Systems capabilities and list of local offices see the following Web location:
www.digital.com/customsystems

Overview: *ProLiant* Clusters for NetWare

A complete range of NetWare cluster solutions
- Low End/Packaged Cluster
  - Low Cost, Small Size, Simplified Clustering to
- High End/Enterprise Solutions
  - Large Size, High Performance

Industry-leading reliability and functionality
- NetWare Clusters
- Compaq Servers and Storage

Integrate with Compaq Service and Support offerings

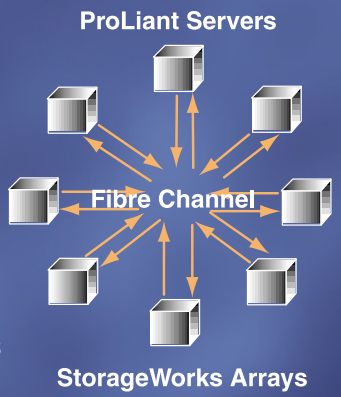

COMPAQ

www.compaq.com/highavailability

ProLiant Clusters for NetWare provide the capability, availability and reliability customers demand for their infrastructures and data centers.

These clusters are based on Novell's innovative NetWare Clustering Services deployed on the platform the world depends on most – the Compaq *ProLiant* servers and *StorageWorks* SAN storage.

These solutions extend from the innovative 2-node CL380 to multiple servers and storage boxes based on Compaq's industry-leading fibre channel solution. Compaq and Novell management tools give customers efficient manageability of the hardware, software and cluster system.

Compaq's global Customer Services offering complete the package by putting the skilled people where and when you need them.

Overview: Proliant Clusters for NetWare

Product Line/Speaker Notes

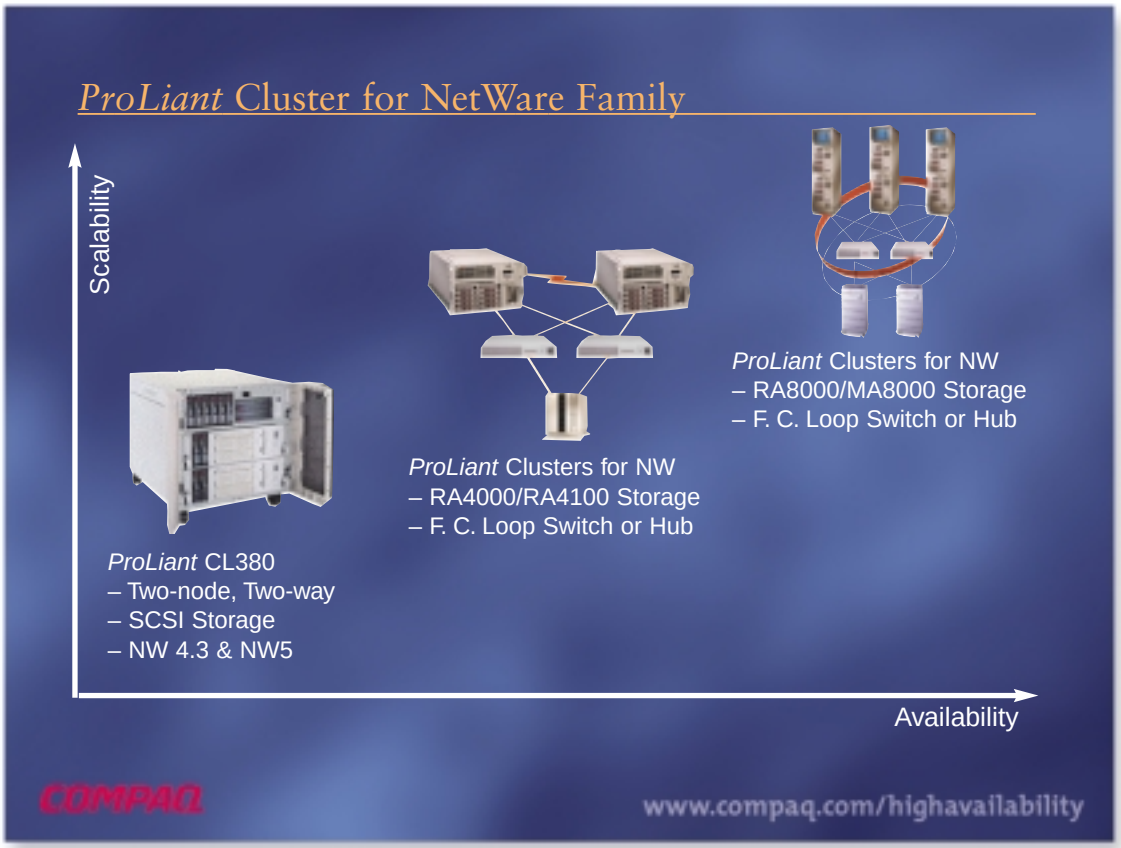

Packaged Clusters
The easiest and most affordable clustering solution – the CL380 Packaged Cluster

For customers who have a need for higher availability but don't have the infrastructure or time to set up a clustering or computing environment; for example, remote offices.

Infrastructure Clusters
The integrated building block approach to flexible industry standard clustering – across a range of servers, storage, and fibre channel interconnects which make up a *ProLiant* Cluster for NetWare 5.

Tight integration of the platform and clustering technology to address reliability and availability requirements.

A spectrum of solution capabilities that address a range of requirements: entry level SAN, NSPoF solutions, etc.

Surrounding the technology with service and support, management, best practices information.

Data Center Clusters
Proven combination of advanced industry standard technologies delivering the highest

levels of performance, capacity and availability.

Going beyond infrastructure for customers with the most demanding requirements:

Large multi-node clusters, high performance/ high-end Storage Area Networks, consolidation of servers/ applications/storage, simplified system management, business critical service offerings

Each of these solutions have "the Compaq advantage"

- Three years of joint design and development work
- NHAS and NCS were developed on Compaq platforms
- Compaq *ProLiants* and *StorageWorks* Storage Systems as a foundation provide investment protection for future Novell Clustering solutions
- Industry leading cluster management tools
- Broadest range of systems certified for Novell clustering
- Compaq's Fibre Channel technology is a key solution enabler
- Compaq Services completes the ideal clustering solution with specialized CarePaqs and Priority Service Packages

Overview: Proliant Clusters for NetWare

Product Line/Speaker Notes

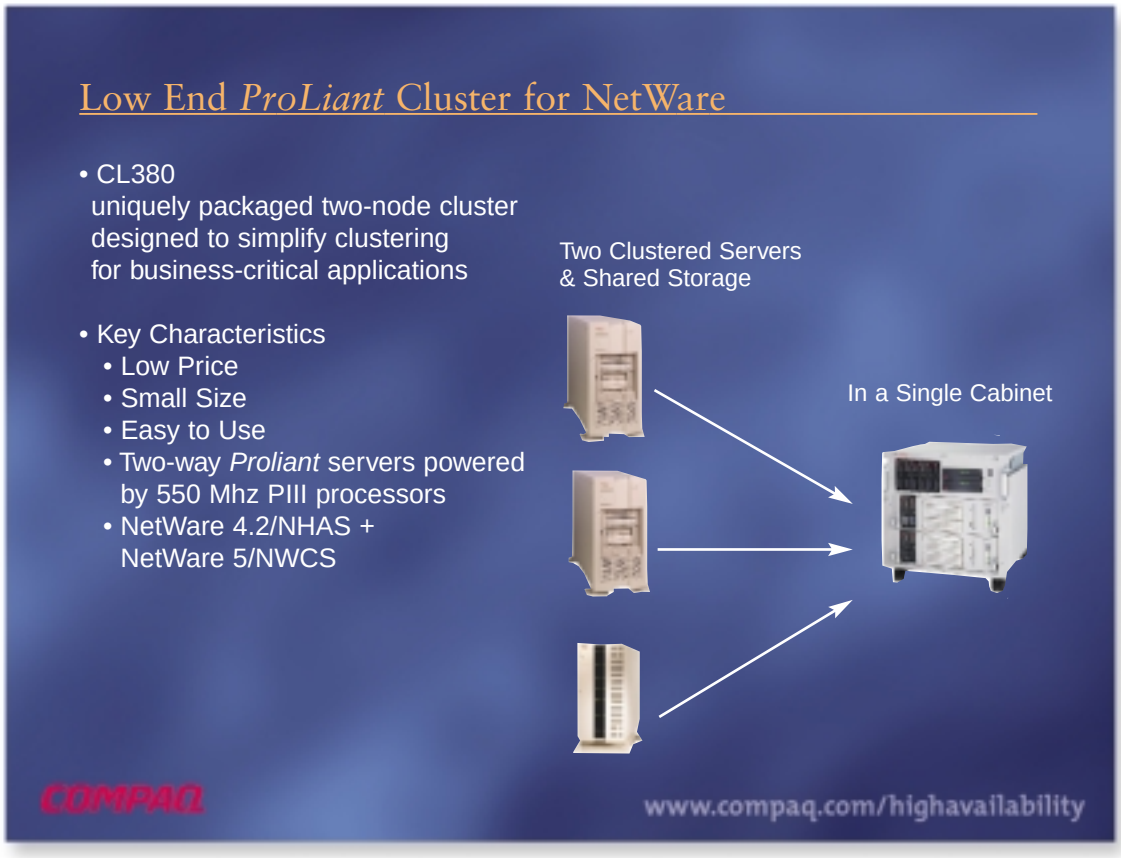

The CL380 makes clustering easy by it's innovative – pre-packaged, pre-wired, pre-tested design.

- A single SKU specifies the system making the CL380 easy to order and ship.
- Deployment is easy, no components or wires to install

The CL380 provides a low-cost high availability solution making it economical to extend clustering protection to a wider range of important business functions such as e-mail, file & print, etc.

The small size and configuration flexibility of the CL380 make it ideal for environments such as branch offices where space is constrained.

• Redundancy, cluster failover and remote monitoring allow the CL380 to operate without the need for on-site technical support. So in a branch office, the CL380 can be used to run critical business applications with the knowledge that a failure will not bring the system down. Remote monitoring allows for the detection of latent problems which can be corrected before having an impact on the performance of the system.

ProLiant Clusters for NetWare 5 can be configured to meet the infrastructure computing needs of an enterprise.

- File & Print and E-Mail are traditional strengths of the NetWare NOS.
- Small Storage Area Networks (SANs) provide for the consolidation of storage or the implementation of a disk "farm" to more effectively manage storage resources

High Availability can be extended to these types of applications by clustering the *ProLiant* servers that best fit your requirements.

- New *ProLiant* servers can usually be clustered with previous *ProLiant* models to allow a high degree of investment protection
- The ability to mix different server models is another means of protecting your previous investment is computing resources

Dual Fibre Channel Loops or Loop Switches can be configured with the RA4000 or RA4100 to provide for No Single Point of Failure in the path from Server to Storage.

Overview: Proliant Clusters for NetWare

Product Line/Speaker Notes

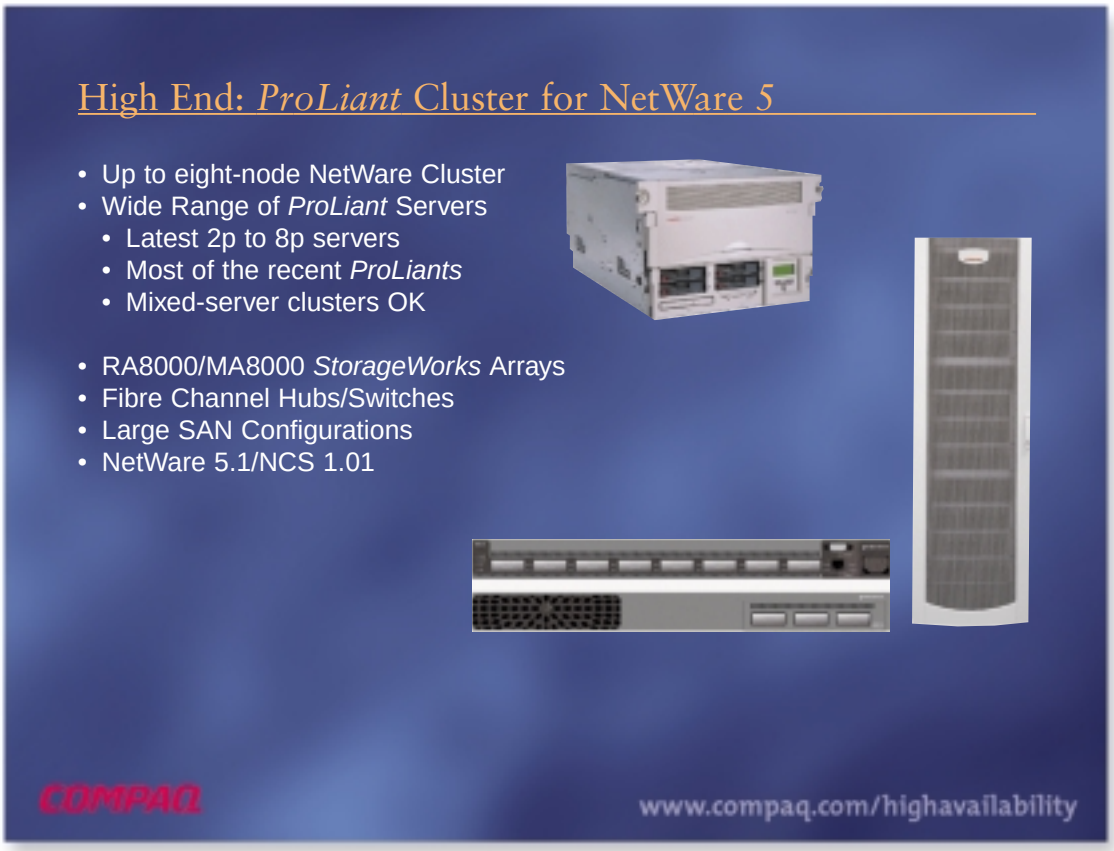

Data Center requirements typically mean:

- Large multi-node clusters
- High capacity storage
- High performance interconnects from servers to storage

ProLiant Clusters for NetWare 5 have Data Center capabilities:

- Eight-node clusters (with all nodes active)
- Up to 8-Way *ProLiant* Servers using the latest Intel chip technology
- *StorageWorks* MA8000 family storage with 20TB storage capacity
- Fibre Channel Switches to facilitate large clustered SAN environments

Compaq *Insight Manager/XE* (CIM/XE) is well known for providing the ability to discover, monitor and manage server resources. Resources can be identified, problems can be identified and operational characteristics can be monitored to determine that resources should be increased or shifted.

StorageWorks Command Console allows the set-up and monitoring of storage resources utilizing a robust and intuitive GUI interface.

Novell's ConsoleOne is used to monitor the operation of the cluster and indicate if a problem has occurred, the source of the problem and the status of the cluster resources.

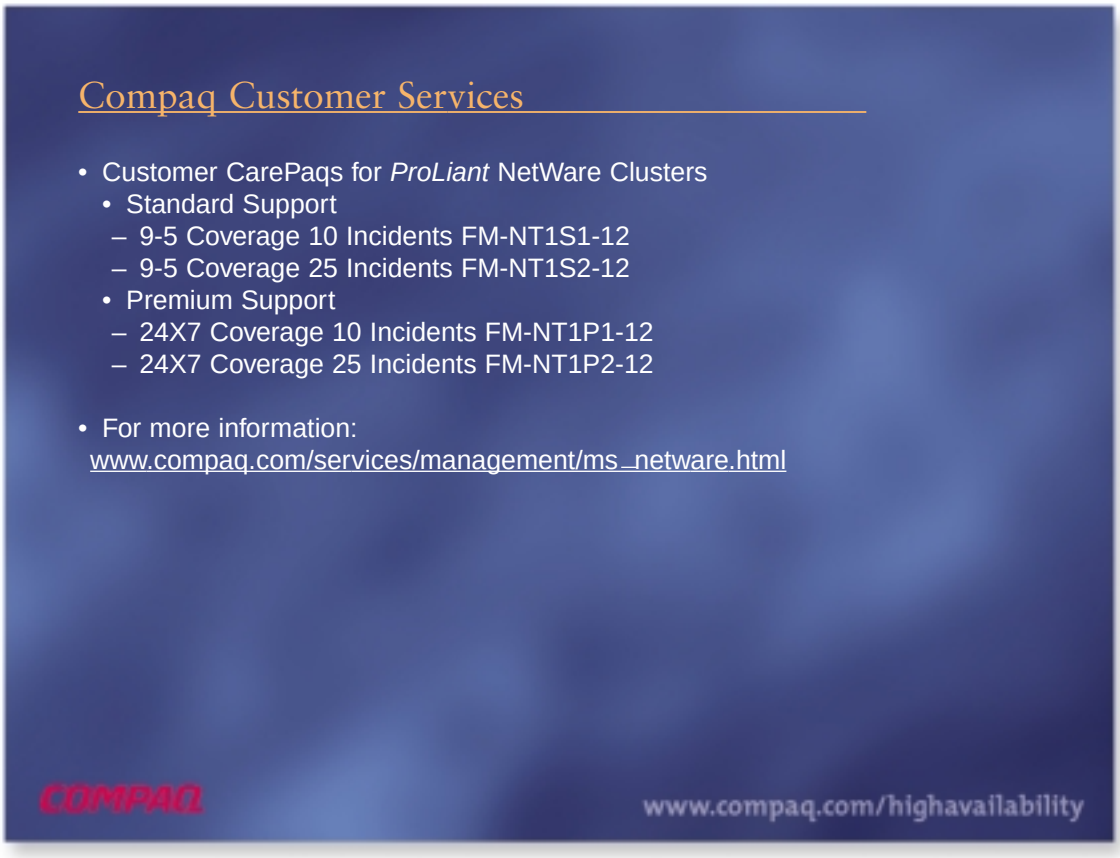

Compaq's Customer Services organization can provide the right resources at the right time to get your cluster running and to keep it running.

CarePaqs are available to provide standard 9-to-5 support

CarePaqs are also available to meet the needs of around the clock operations with coverage 24 hours a day, seven days a week

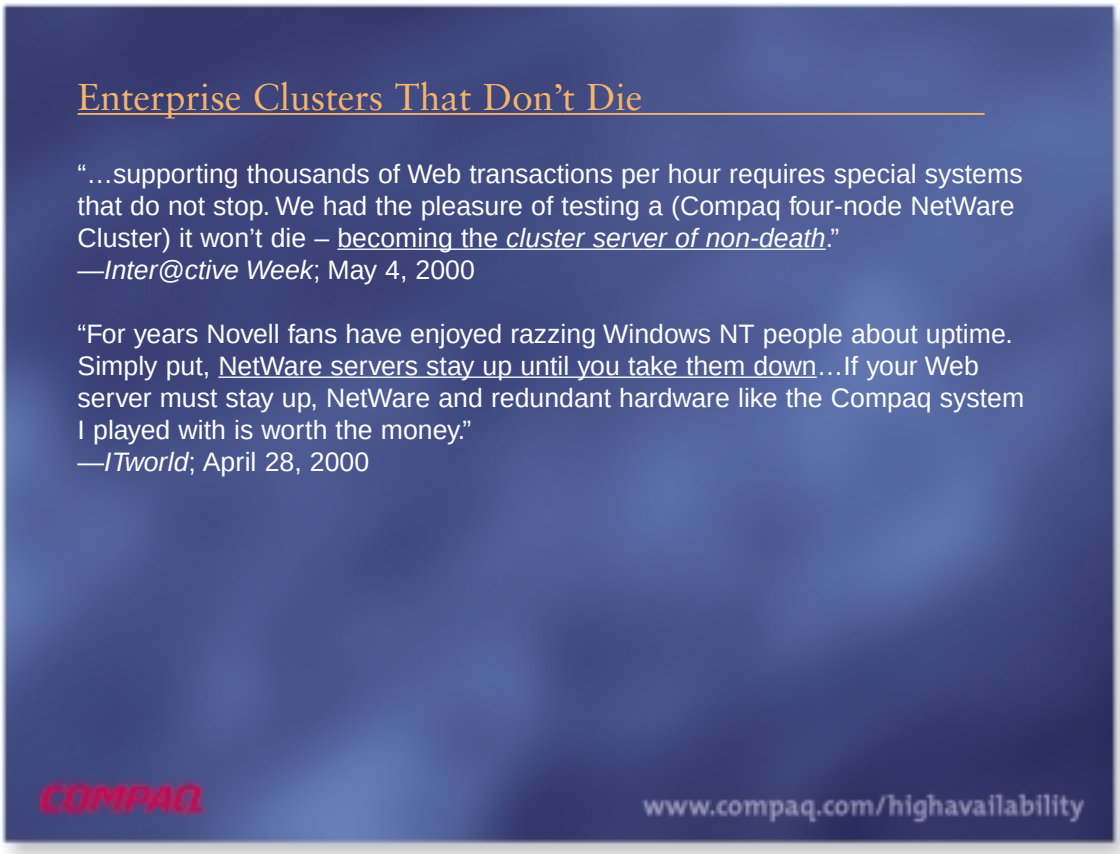

Computer industry publications have tested and verified that Novell NetWare, and NetWare Clusters, in conjunction with Compaq *ProLiant* servers provide industry-leading capabilities

The *Inter@ctive Week* and *ITworld* articles were based on a cluster of four *ProLiant* 6400's running NetWare 5.1 and NetWare Cluster Server 1.01.

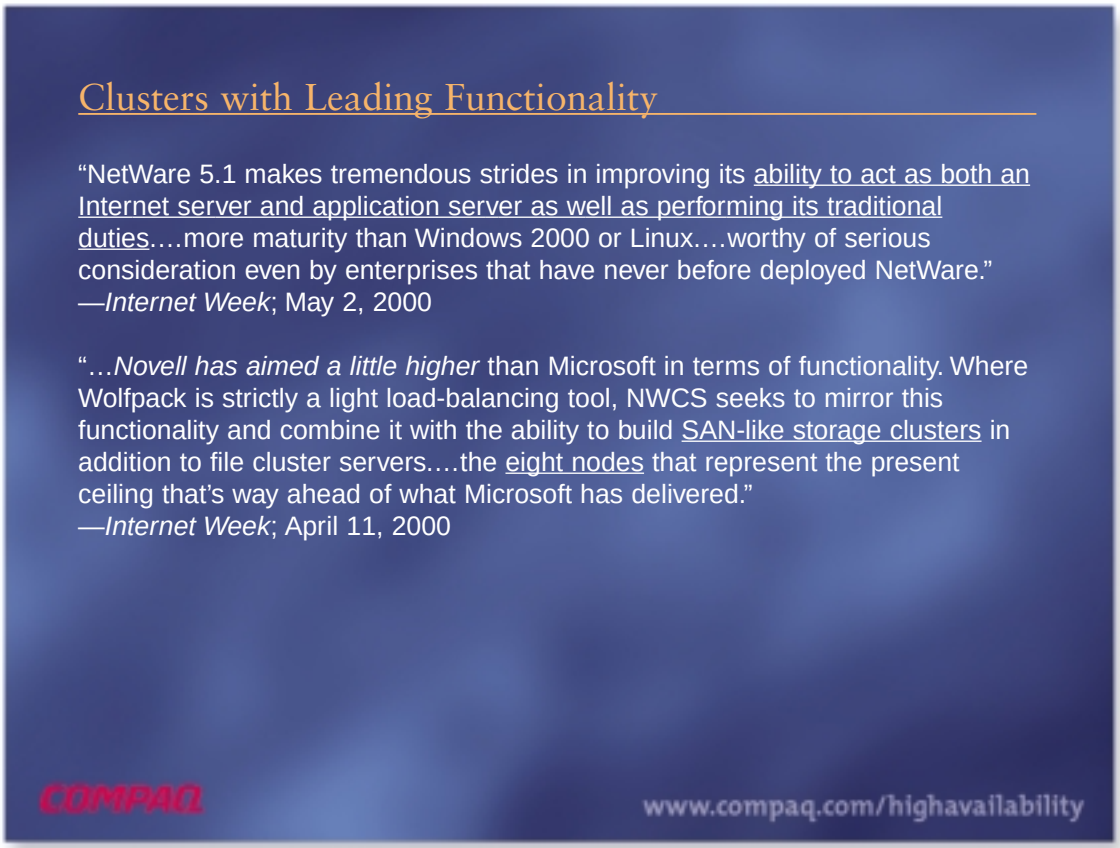

The April 11 *Internet Week* article is based on tests of NW 5.1 and NWCS 1.01 running on four *ProLiant* 1600's. The reviewer is also very positive about Novell's ConsoleOne cluster management software.

The May 2nd *Internet Week* article is based on tests with the *ProLiant* 1600 and (separately) on an IBM NetFinity 5500. The article also has positive comments about Novell's NDS and ZENworks as well as discussing NetWare for Web servers.

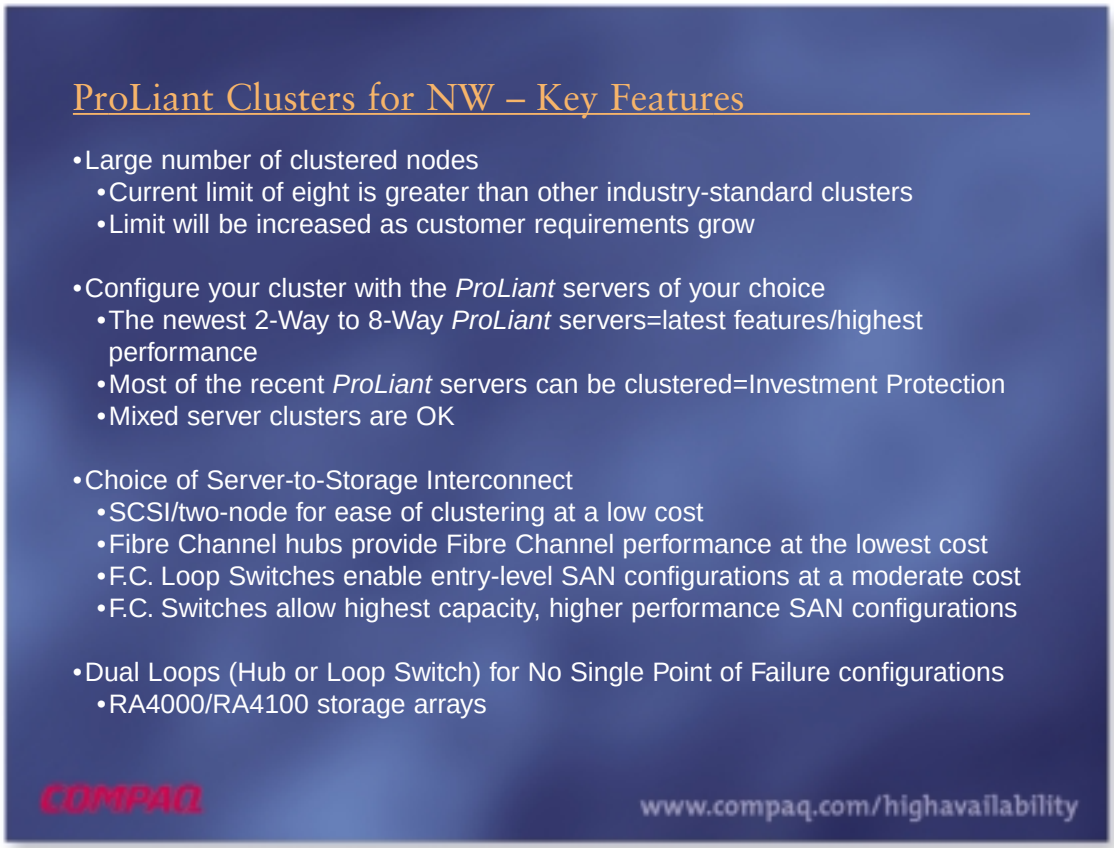

ProLiant Clusters for NetWare offer several key features:

- NWCS will support up to 32 clustered nodes. Our limit today is set at eight nodes since this is the number we've subjected to rigorous testing so far. We're continuing the testing on larger clusters and will expand our limit as our customers move to the larger number of nodes.
- Pick the *ProLiant* servers of your choice – new, old, mixed types
- Interconnects are available for the price – performance you need
- For environments that require the greatest levels of protections *ProLiant* Clusters for NW with the RA4000/41000 can be configured with dual loops or loop switches for NSPOF

Overview: Proliant Clusters for NetWare

Product Line/Speaker Notes

ProLiant Clusters for NetWare address the real world problems faced by our customers:

- Run business applications with outstanding performance
- Keep critical computing resources on-line and available to run the business
- Reduce computing costs while simplifying the computing environment
- Transparently add or upgrade resources to accommodate new technology and growth.

The combination of Data Center, Infrastructure and Packaged Clusters from Compaq provide the most complete offering of Industry standard clustering available.

ProLiant Clusters for NetWare provide the best performance, range of choices and value for your business.

MODELS
ProLiant **Cluster for NetWare 5 Novell Cluster Server**
Select *ProLiant* Servers
StorageWorks **storage arrays**
Fibre Channel server-to-storage interconnect

Overview

Compaq *ProLiant* Cluster for NetWare 5 provides high availability for applications and data in business-critical Novell NetWare 5 environments with support for up to 8 cluster nodes.

The Compaq *ProLiant* Cluster for NetWare 5 uses Compaq's industry-leading *ProLiant* servers, Compaq *StorageWorks* Raid Arrays and Fibre Channel interconnects from server to storage. These configurations have been fully tested by Compaq and certified by Novell for use in NetWare Cluster Services for NetWare 5 environments.

The Compaq *ProLiant* Cluster for NetWare 5 utilizes NetWare Cluster Services for NetWare 5 software with Novell NetWare 5 for cluster operation and failover management.

The Novell NetWare 5 operating system and the NetWare Cluster Services software need to be purchased separately from your local Novell authorized reseller.

STANDARD FEATURES

Maximum Availability, Minimum Cost

Because Compaq *ProLiant* Clusters are built from industry-standard components, Compaq delivers high levels of application availability at a much lower cost than traditional, proprietary cluster solutions. The Compaq *ProLiant* Cluster for NetWare 5 uses Compaq industry-leading *ProLiant* servers, the fast and reliable Compaq *StorageWorks* Raid Array storage family and Fibre Channel Loop and Switch interconnects.

Investment Protection

A select set of existing and future Compaq *ProLiant* servers will be certified for Compaq *ProLiant* Cluster for NetWare 5. This means that customers can build clusters using existing servers or by mixing old and new servers. Customers who purchased the Compaq Fibre Channel Storage System, (now called the RA4000), for use with NHAS clustering for NetWare 4 can easily migrate to the NetWare Cluster Services For NetWare 5 solution. Multi-path I/O is supported as well as the option to configure dual Fibre Channel Loops with the RA4000/RA4100.

Integrated Solutions

All Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution for business-critical environments. Compaq *ProLiant* servers, interconnect options, system management software, flexibility in configuration support, and implementation documentation have all been thoroughly tested in cluster configurations. Compaq's close relationship with application partners, and expertise with the Novell environment has resulted in the development of tools, papers and solutions, which assist in the rapid deployment of major business applications on Compaq *ProLiant* Clusters. For more information, visit the following Web site: **www.compaq.com/highavailability**.

Sophisticated System Management

Management tools from Compaq and Novell are integrated to provide management of the servers, storage and cluster system. Compaq *ProLiant* Cluster for NetWare 5 configurations utilize Compaq *SmartStart* and Compaq *Insight Manager* for rapid installation, configuration monitor and management of the *ProLiant* Servers. Compaq tools, including Compaq Insight Manager and *StorageWorks* Control Console, are combined with Novell's CPQ OnLine utility provide for detailed management and configuration of the storage arrays. Finally, ConsoleOne from Novell provides for comprehensive monitoring and management of the clustered systems.

COMPONENTS

Servers

A broad range of *ProLiant* server models are supported in the *ProLiant* Cluster for NetWare 5 configurations, in both matched-pair and mixed-pair combinations. Please check the Compaq

High Availability Web site for an up to date listing of certified *ProLiant* Cluster for NetWare 5 configurations at **www.compaq.com/highavailability**.
Detailed information on the specific servers can be found at
www.compaq.com/products/servers/platforms.

Interconnect

Both Fibre Channel Loops and Fibre Channel Switches are used to provide high performance, scalable interconnects from the servers to the storage arrays. Detailed information on the Fibre Channel devices can be found at
www.compaq.com/products/storageworks/raidstoragesystems.html .

Storage

Compaq *ProLiant* Cluster for NetWare 5 configurations have been certified with the Compaq *StorageWorks* Raid Array family including the RA4000, RA4100, RA8000/ESA12000 and the MA8000/EMA12000. Detailed information for configuring these subsystems can be found at
www.compaq.com/products/storageworks/raidstoragesystems.htm.

Supported Operating Systems

Novell NetWare 5

OPTIONS

Most server options and the *StorageWorks* options are supported in a cluster. Refer to the Compaq high availability web sites at
www.compaq.com/highavailability and
www.compaq.com/products/storageworks/raidstoragesystems.html for ordering and configuration guides, release notes, and restrictions.

Service and Support

Compaq servers and storage systems are protected by Compaq Services including a three-year, limited warranty[1], 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-Failure Warranty applies to certain hard drives, memory and processors for servers monitored by Compaq *Insight Manager 2.0* or later. For more information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

1 Note: Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

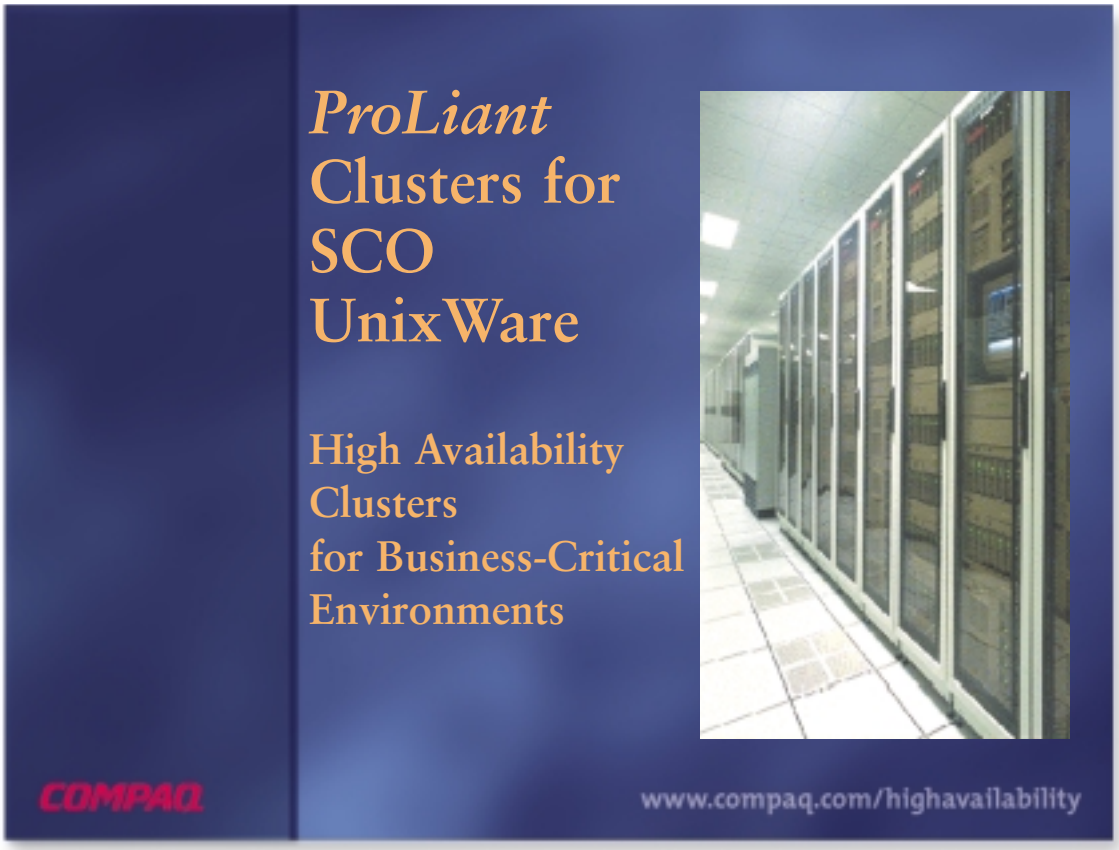

This presentation is about the Compaq *ProLiant* Clusters for SCO UnixWare product line.

- These clusters offer industry standard UnixWare 7 running on a cluster of *ProLiant* servers enabled by UnixWare *NonStop*™ Clusters software.
- All *ProLiant* Clusters for SCO UnixWare 7 configurations are ideal for high availability needs in business-critical environments where industry-standard servers are needed.
- *ProLiant* Clusters for SCO UnixWare is available in cluster kits that allow a wide variety of two- to six-node clustered, highly available and application-scaling solutions to be built.

The clusters are based on Compaq *ProLiant* servers and options, StorageWorks and UnixWare 7 *NonStop*™ Clusters software.

This software is developed by Compaq, licensed to SCO and can be purchased from an authorized SCO reseller.

ProLiant Clusters for SCO UnixWare bring industry-leading cluster technology to industry-standard servers for the UNIX community.

The solution provides greater than 99.99 percent availability for both planned and unplanned downtime at an unmatched price and price/performance level.

- This technology is further enhanced for application uptimes with the new Online Software Upgrade utility that allows new versions of the operating system to be installed and tested while the system is operating and performing productive work.
- This eliminates one of the key planned downtime activities for software upgrades and allows the cluster to continue to operate on business-critical applications while bringing in a new version of the operating system and applications.

ProLiant Clusters for SCO UnixWare running UnixWare *NonStop*™ Clusters software was announced November, 1999 and offered a range of two- to six-node highly available application-scaling solutions.

Additionally, UnixWare *NonStop*™ Clusters software was developed by Compaq and is licensed to SCO. It is purchased along with UnixWare 7 from an authorized SCO reseller or distributor.

This clustered solution offers:

- The most advanced single system image (SSI) clustering technology. Multiple nodes appear as a single SMP system to users, administrators and applications.
- Single-file system provides ease of use for programming, management, administration and extensions by providing a single name space; single root file system; single device view; inter-node streams, pipes and semaphores; single password file, single IP host table; and network address, cluster-wide time synchronization and a single administrative command set.
- Only clustering technology with active process migration with automatic load balancing.
- Most standalone UnixWare applications can be ported without modification, as can Sun Solaris and HP applications. Applications need not be cluster-aware.

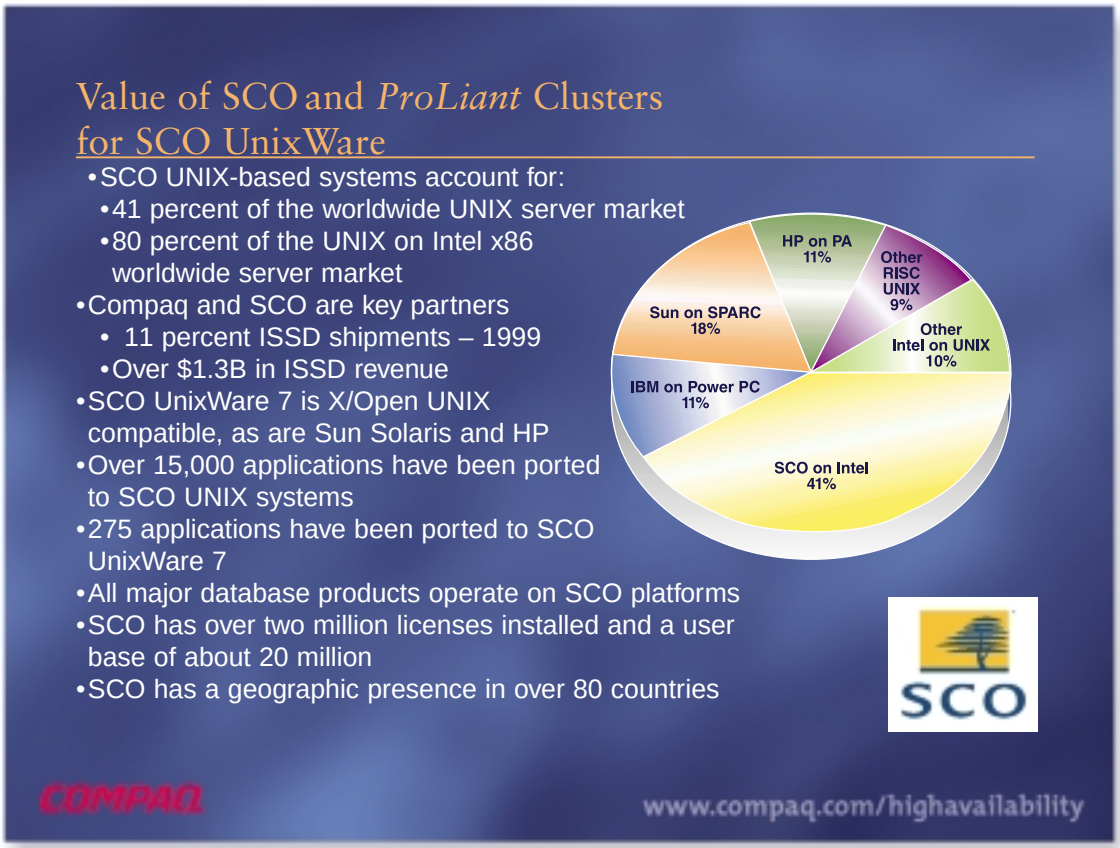

This slide shows that SCO UNIX based systems account for:

- 41 percent of the worldwide UNIX server market and 80 percent of the UNIX on Intel x86 worldwide server market.

Compaq and SCO are key partners

- *ProLiants* shipping with SCO operating systems accounted for 11 percent of ISSD shipments in 1999 for over $1.3B in ISSD revenue.
- Today, Compaq is the largest supplier of systems running SCO UNIX operating systems, with a greater than 40 percent share of the worldwide SCO UNIX market.

Additionally:

- SCO UnixWare 7 is X/Open UNIX compatible, as are Sun Solaris and HP, so applications can be easily ported to SCO UnixWare 7.
- Over 15,000 applications have been ported to SCO UNIX systems and SCO UNIX (OpenServer and UnixWare) is available in six language variants, including Japanese.
- 275 applications have been ported to SCO UnixWare 7.
- SCO ships over 300,000 UNIX-based systems per year.
- All major database products operate on SCO platforms.
- SCO has over two million licenses installed and a user base of about 20 million.
- SCO has geographic presence in over 80 countries.
- SCO sees huge UnixWare 7 *NonStop*™ Clusters software growth.
- SCO UnixWare 7 is the fastest growing commercial UNIX operating system in the world with a 59.8 percent CGR.
- UnixWare *NonStop*™ Clusters software is the path to IA64 and Monterey

Customers Need Clustering

- Resilience to application failure—Fault tolerance
  - Needed for 24 x 7 operation
- Consolidation – fewer but larger servers
  - Lower administration and maintenance costs
  - Reduced complexity and risk
- Manageability
  - Centralized administration and reduced operator costs
- Scalability
  - Can easily expand to meet business needs

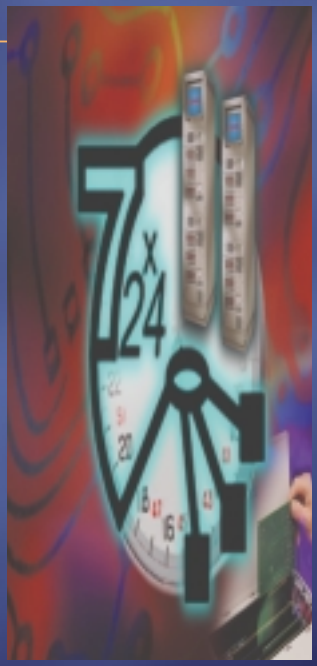

COMPAQ

www.compaq.com/highavailability

With ProLiant Clusters for SCO UnixWare, Compaq offers resilience to failure that is required for 24 x 7 applications and operations.

- High availability is provided via the OnLine Software Upgrade (OSU) capability that is a component of the UnixWare 7 *NonStop*™ Clusters software. This feature allows users to add new software and hardware to the cluster without taking the cluster down.
- With OSU, availability on *ProLiant* Clusters is over 99.99 percent and in some cases approaches 99.995 percent, or 26 minutes of downtime per year.
  - Unlike most other measurements of availability, this number includes hardware, operating system and planned and unplanned outages.

Let's discuss what the very high availability number means. It is based on SCO's availability modeling of hardware combined with the Operating System and clustering software.

- Our high availability capability is not intended to imply a guaranteed service uptime. However, Compaq will provide its customers of this product with a Global Services Single Point of Control contract.

Larger servers and fiber-connected storage are supported for optimum performance.

Single System Image (SSI) management allows centralized administration, thus reducing operator costs.

And with ServerNet, the cluster can grow to six nodes to support business growth.

The main target markets for this product line are:

- Existing OpenServer/UnixWare customers because there is a huge installed base.
- Low-end ISPs and ASPs who are familiar with LINUX or Sun Solaris who require more robustness at a low cost.
- RISC sites wanting to move to Intel IA32/64 for cost reasons and complexity reduction, as well as needing a uniform hardware platform.
- Organizations who want clustering and dislike NT/ Win2000. SCO has a proven track record in the same markets NT/Windows 2000 is sold in, particularly in the SMB market.
- Many customers need larger systems with easier manageability.
- Many of these users just described are looking for greater applications and data availability.

Typical customers for this product line include:

- Command/Control – for city police, fire and ambulance departments, as well as security services.
- SMB customers – for back office and speciality applications, as well as start-up ISPs and ASPs.
- Replicated Sites – in the retail, government (such as customs departments) and services industries (such as railways).
- Telecommunications industry – so that they can rapidly deploy then scale new services as needed for non-switch applications, such as short messaging services and over-the-air-activation of cellular services. Additionally, telecom companies can use this product for call centers.
- Finance industry – for use in banking as well as in stock exchange share-dealing.
- Health care industry – for back office and call center applications.

This slide summarizes key markets where *ProLiant* Clusters for SCO UnixWare are targeted.
While the product line is suitable for applications at the high end, this product is being targeted at the mid-range, departmental and small workgroup markets. Success has come through applications partnerships with SCO and ISVs who sell key applications in these areas.

- In the mid range, the applications typically found are database as well as ISV applications driven solutions. In some cases, the application is a custom in-house development.
- At the departmental level and retail distributed levels. Key retail, replicated (or distributed) applications are found in these target markets:

Pharmacy
Department Store
Supermarkets
Food Service
Hospitals
Automotive
Gas/Convenience

The small business and workgroup level can be defined as having between five and 500 employees and up to $500 million per year in revenue. Typically these organizations solutions are driven by ISVs. This is where retail point-of-sale applications predominate.

This slide depicts how *ProLiant* Clusters for SCO UnixWare brings customers enterprise class availability features, even at the lower end of the market and extending into the data center

- The different cluster kits offer a variety of cluster configurations, price points and install options.
  - The basic *ProLiant* Clusters for SCO UnixWare kit allows two- to six-node clusters connected by either ServerNet or Ethernet.
  - Clusters using the ServerNet interconnect capability can be scaled to six nodes and are suitable for data center applications.
  - Clusters using Ethernet as the interconnect are two nodes and are suitable as infrastructure clusters since they can also support a wide variety of FCAL storage capabilities.
- The Quick Install Cluster kits include unique quick install packages with UnixWare 7 *NonStop*™ Clusters software and the UnixWare 7 operating system. It offers media kits designed for a two-node Ethernet IP *ProLiant* cluster.
  - These clusters are smaller and very competitively priced. They are very suitable to entry level clusters for branch locations and other replicated sites.
- UnixWare 7 *NonStop*™ Clusters software and UnixWare 7 licenses and support for the software are purchased separately through authorized SCO resellers.

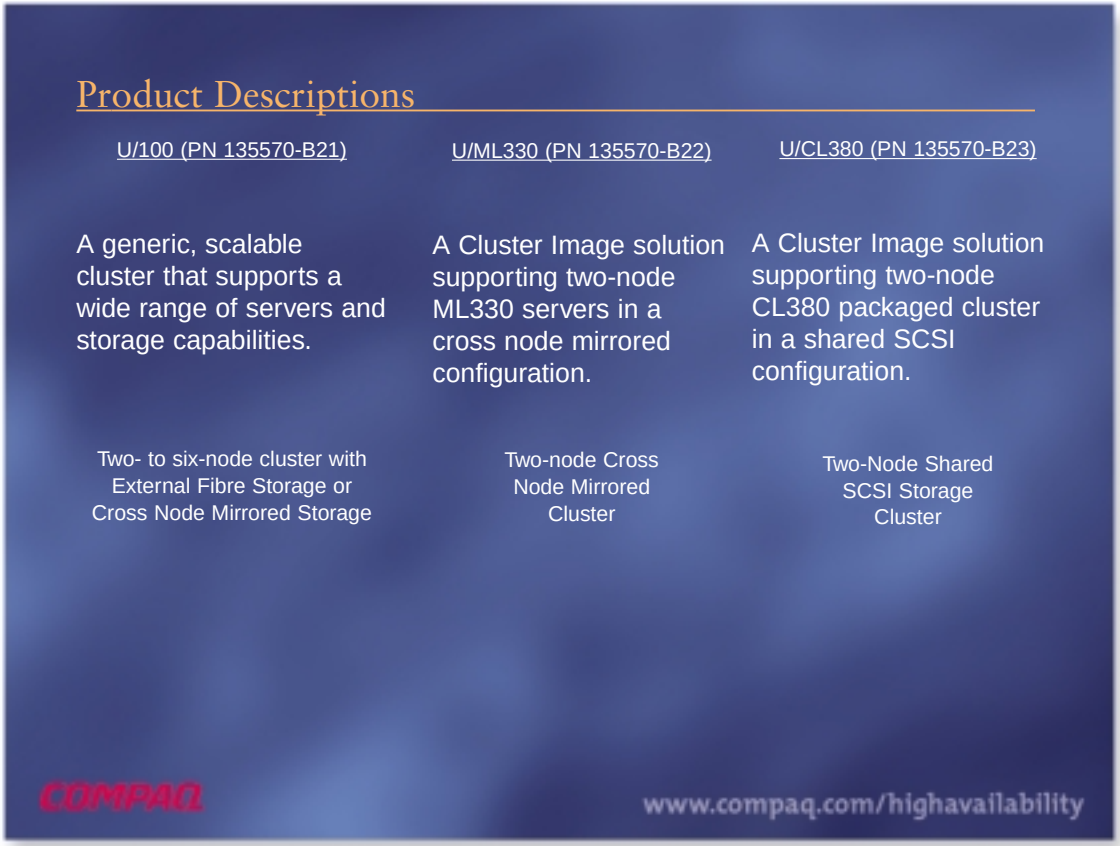

This slide is optional because it depicts new products that will begin shipping in July, 2000. These products are:

- The U/100 is a generic scalable cluster that supports a wide range of servers and storage capabilities. It is scalable from two to six nodes, using ServerNet or can be configured with two nodes only with the Ethernet IP node interconnect option. It will begin shipping on July 17, 2000.
- The U/ML330 is a cluster image kit supporting two-node, ML330 servers in an Ethernet cross-node mirrored configuration. It will begin shipping on July 21, 2000.
- The U/CL380 is a cluster image kit supporting the two-node CL380 packaged cluster in a shared SCSI configuration. It will begin shipping August 18, 2000.

The Quick Install clusters are designed to allow users to install the software for the cluster and be productive within an hour. ProLiant Clusters for SCO UnixWare offer many advantages.

- First, they offer the industry's most advanced single-system image clustering technology, where multiple nodes appear as a single SMP system to users, administrators and applications. This comprehensive SSI simplifies system administration tasks with a consistent, intuitive view of the cluster.
- It provides a single cluster file system; a single device view; inter node streams, pipes and

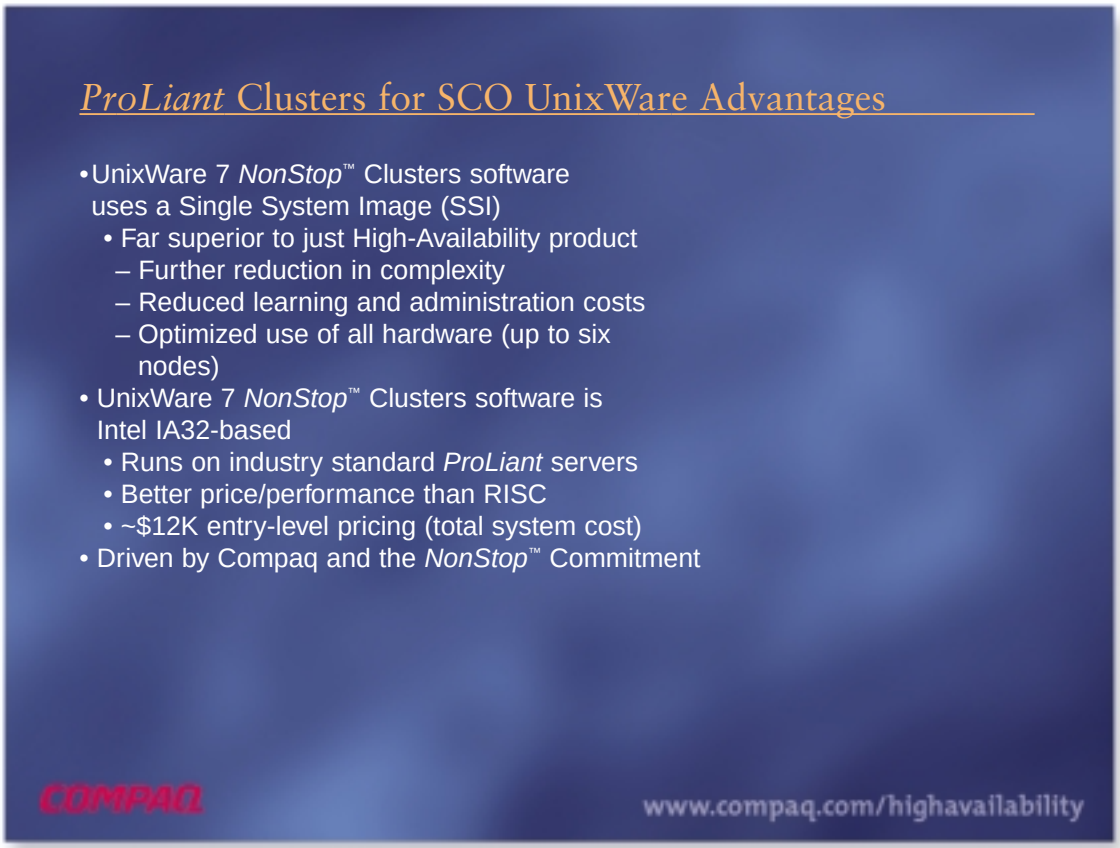

semaphores; a single password file; single IP host table and network address; and more. Inter-process relationships across server nodes are preserved.

- This technology also supports active process migration with automatic load balancing, which provides the most comprehensive and easy to manage cluster in the industry. System administrators familiar with standard SCO admin tools need not learn new interfaces as the cluster management interface is a GUI extension to existing SCOAdmin tools and utilities, as well as integrated into *Compaq Insight Manager* and *Insight Manager* Web-enabled utilities for clusters.

This cluster is based on industry-standard *ProLiant* servers and options for unmatched price and price/performance for a highly available solution.

- Wide range of *ProLiant* server support from the uniprocessor ML330 to the 8-Way SMP DL750.
- Low cost, entry two-node cluster features include Ethernet node connectivity, internal drive cross node mirroring and *ProLiant* 300 Series family support provide the lowest cost cluster in the *ProLiant* Cluster family.

This product line has been developed from Compaq's 20-plus years of experience in delivering *NonStop*™ capabilities to the marketplace.

Summary

Leading-edge features:
- Unsurpassed price/performance
- Scalability and reliability
- The most advanced SSI technology in the industry
- Technology migration at rapid speed
- Easy, fast cluster installation
- High availability
- Broad range of servers and storage

Best fit for customers who have:
- HA requirements for departmental and distributed applications
- HA requirements and restricted budget
- HA requirements for space constrained areas
- Customers looking for easy installation and management

World-wide service and support

COMPAQ

www.compaq.com/highavailability

As we have discussed, Compaq and SCO are the world's volume leaders for industry-standard servers and open UNIX operating systems. In summary, this new release of *Compaq ProLiant* Clusters for SCO UnixWare offers:

- Unsurpassed price/performance when combined with *ProLiant* servers' Intel economics.
- Scalability and reliability without traditional clustering management penalties.
- The most advanced SSI cluster technology in the industry.
- Technology migration at rapid speed.
- Cluster high availability.
- A range of Compaq *ProLiant* Servers & Storage.
- A range of Compaq StorageWorks capability from 9.1GB cross-node mirroring configurations to 9.6TB of RA4000/RA4100 Fibre Channel Array storage.
- Additionally, Compaq offers cluster installation improvements along with SCO's software through its cluster images. This software allows users to install the cluster and be productive on it within an hour.

This product line (especially at the low end, entry-level area) is the best fit for:

– Customers with high-availability requirements for departmental and distributed applications
– Customers with high-availability requirements and restricted budget
– High-availability requirements for space constrained areas
– Customers looking for easy installation and management

This product line offers investment protection like no other clustering

technology, scalability to six nodes and a wider selection of hardware with low-cost servers and Ethernet two-node interconnectivity.

World-wide service and support are available as *CarePaqs* for Compaq's hardware.

Compaq has many satisfied customers running *ProLiant* Cluster for SCO UnixWare and this release will add many more who have requirements for low-cost high-availability UNIX clusters.

HRG Assessment: Compaq ProLiant Clusters Running NonStop™ Clusters for SCO UnixWare Software

White Paper

For more than five years Harvard Research Group (HRG) has tracked the development of high availability (HA) clusters. Over that time HRG has conducted thousands of interviews with information technology (IT) professionals. Many of these professionals have expressed a serious lack of confidence in high availability clustering solutions, particularly in the Intel standard server market. HRG believes that clustering for increased levels of availability, scalability and manageability is exactly where the industry is heading. In today's fast-paced global 7x24 business environment, systems need to be easily expandable and able to accommodate ever-increasing performance requirements, yet require little or no additional management overhead. However, HRG believes that the clustering market will experience its true growth potential only once clustering solutions can uniformly provide acceptable and proven levels of scalability, manageability and availability. Earlier this year, Compaq Computer and SCO® committed themselves to using their combined forces to develop a solution that delivers application and data availability, scalability and manageability in the Intel standard server SCO UnixWare environment. That solution is Compaq *ProLiant* Clusters running SCO UnixWare 7 *NonStop™* Clusters software.

The recent Compaq/SCO offering, combining Compaq's ProLiant *servers and storage with UnixWare 7.1* Non-Stop™ *Clusters software, provides one of the better HA UNIX-based data and cluster management solutions available today. Originally developed by Compaq's Business Critical Server Division (formerly the Compaq Tandem division) which has 20-plus years of expertise in high-performance clusters, Compaq has produced an industry-standard clustered server solution that offers a unique combination of application/data recovery, scalability and full Single System Image (SSI) functionality. Any company seeking a good high availability solution should take a long look at what Compaq and SCO have developed.*

Many of the HA system users interviewed by Harvard Research Group (HRG) during 1998 and 1999 expressed concern over their ability to manage clusters of Intel-based PC servers. Based on these interviews HRG is convinced that IT managers are not only in need of a good HA clustering solution for their commodity servers but are increasingly seeking solutions that are scalable and easy to manage. Additionally, IT managers expressed the need to "virtulize" their system environment and manage it as if it were a single logical system. While most clusters do a good job of addressing availability issues, many do very little to resolve the need for scalability and manageability. Clusters typically bring additional complexity and unfamiliar management tools which further add to the confusion and anxiety expressed in the interviews.

Harvard Research Group defines a cluster as a loosely coupled collection of servers and storage devices managed by a cluster software module as a single logical system. The computers that comprise a cluster are called nodes. Nodes can be either uni-processor (UP) or symmetric multiprocessing (SMP) systems. Each node runs its own copy of the operating system kernel and has its own memory and I/O devices. The primary benefits of a cluster are availability and scalability.

Most clusters currently available for Intel-based servers provide basic availability via failover, single system image only for the applications, and limited scalability. Compaq's collaboration with SCO to deliver UnixWare 7 *NonStop™* Clusters software has expanded the cluster definition by allowing the cluster to appear to be a single system to the user, administrator and application. This robust implementation of a single system image (SSI), not available with other UNIX clusters, differentiates other cluster offerings running on *ProLiant* Clusters by allowing the administrator to manage a single resource, rather than a collection of

server nodes. Additionally, the SSI model makes it easier for application developers to port their software from noncluster environments to the SCO UnixWare 7 *NonStop*™ Clusters environment and allows more efficient use of the resources available in the cluster. For ease of growth without losing application execution time due to planned outages, SCO UnixWare 7 *NonStop*™ Clusters software supports transparent online maintenance and online addition of nodes to the cluster.

Compaq's SSI implementation, licensed to SCO for use with UnixWare 7 *NonStop*™ Clusters software, layers clustered services over each standard UNIX service that needs to participate in the SSI, such as a file system service or UNIX System V interprocess communications services. The clustered service maintains the standard service call interface for the service, so that applications do not need to be altered, recompiled or relinked. SSI also provides transparent access to all OS resources cluster-wide using a single namespace.

SSI helps reduce administrative overhead by providing a consistent, intuitive view of the cluster and allows multiple nodes to appear as a single symmetric multiprocessor (SMP) system to users, administrators and applications. Other significant features of SCO UnixWare 7 *NonStop*™ Clusters software include:

- SSI functionality that includes applications, operating system and devices.
- Expandability that supports the addition of up to six eight-way SMP server nodes.
- Ability to run applications out of the box without the need to be cluster aware.
- Uses same administrative commands as a single system.
- Ability to move system process to another node without stopping and restarting the process.
- Capability to configure a highly-available cluster with small, manageable fault zones.

Compaq's Integrity XC

The NonStop™ Clusters technology is also being used in Compaq's Integrity XC, a two- to 12-node packaged system of ProLiant *industry-standard servers,* ServerNet *system area network (SAN), and NonStop™ Clusters for SCO UnixWare software. It is marketed directly to telecommunications partners and end users by the Compaq Telecommunications Platform Division.*

Product Overview

Compaq *ProLiant* Clusters for SCO UnixWare 7 is a two- to six-node clustered solution for high availability and application scaling. The *NonStop*™ Clusters software is certified to run on industry-standard Compaq *ProLiant* servers, spanning the range from entry Compaq PL800 to the high-end Compaq 8-Way PL8500 SMP. Compaq will make available *ProLiant* Clusters for UnixWare 7 (cluster kits) to the distribution channel, starting in November 1999.

The SCO UnixWare 7 *NonStop*™ Clusters software environment provides a replicated operating system, which continues membership services as if the hardware were replicated. The software employs standard UNIX system application program interfaces (APIs) for programming, networking and management. It is binary compatible with SCO UnixWare 7 and transparent to existing applications. Therefore, users can continue to run existing applications and libraries after installing the SCO UnixWare 7 *NonStop*™ Clusters software. Compaq's enhancements consist of key cluster "hooks" placed in the standard SCO UnixWare 7.1 operating system that allow most hardware adapters and software applications developed for use on SCO UnixWare
to run without change.

Availability

To ensure availability, SCO UnixWare 7 *NonStop*™ Clusters software provides clustered services over each participating standard UNIX system service (see Figure 1). The clustered services maintain the standard service call interface, so upper levels of the operating system need not change. Applications access clustered services through standard UNIX system libraries, which in turn access clustered services through the service call interface. The clustered service determines whether a request can be handled locally or must be forwarded to another node. If the request is passed to another node, it uses an internode communication system to communicate to the service peer on another node. The request is then handled by the standard UNIX system service on the targeted node.

System availability after the loss of cluster resources or nodes is supported through automatic restart on failure and failover of nodes or resources. The cluster software enables recovery of failed applications and provides tightly defined fault coverage to ensure data integrity. Fault coverage within the SCO UnixWare 7 *NonStop*™ Clusters environment can be defined down to the device, process, system and application. Most other clustering technologies are limited to just the system and the application. This means that applications are assured of the continued availability of resources (I/O, Kernel State), even if the application's specific node fails. SCO UnixWare 7 *NonStop*™ Clusters software creates a "hardened" distributed environment, which allows the cluster to continue to provide a fully functional operating environment even after the loss of cluster resources (nodes, disks, adapters or operating system services).

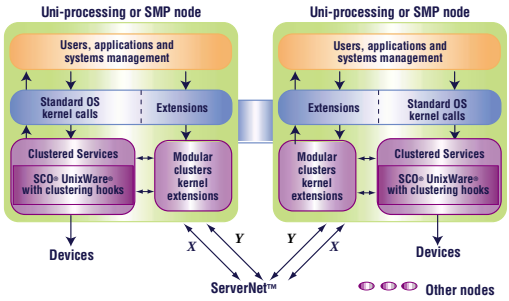

Figure 1 - SCO UnixWare 7 *NonStop*™ Clusters software architecture

Availability is further enhanced by the Compaq's ServerNet technology. The ServerNet system area network (SAN) is a fault-tolerant SAN technology that offers low-latency and high performance. Its core services provide a memory-like read/write interface for access among connected nodes. These services allow the upper layers to open physical connections among nodes and provide low-level support for the ServerNet interface devices. ServerNet technology also provides basic data mapping and movement services, as well as error handling.

Shared Storage

The cluster environment also includes dual-attached storage devices that allow a second node to access a device, even if the primary controller node has failed. To tolerate the loss of a disk, the cluster environment supports the ability to mirror disks in the external storage modules.

SCO UnixWare 7 *NonStop*™ Clusters software also supports a feature called Cross-Node Mirroring, which allows the mirroring of applications and data between two nodes using internal disks, thus eliminating the need for shared external storage. The internal disks of one node are mirrored on the second node using ServerNet technology. In this way, the backup node always has a current copy of the cluster's system disks and file systems for failover access. Cross-Node Mirroring is well suited for small replicated clusters. Applicable in two-node entry clusters, this feature gives *ProLiant* Clusters for UnixWare 7 a very cost-effective entry-level cluster configuration.

Scalability

Compaq's SSI approach to cluster management allows servers within the cluster to be scaled as needed without incurring significant planned downtime. In addition, hardware investment is protected because nodes are not required to contain the same number of processors, processor speeds, memory capacities or internal disk storage capacities. Using Compaq *ProLiant* 8-Way servers, a six-node cluster configuration can attach up to 33 Compaq RAID arrays, 48 processors and 32GB of memory. SCO UnixWare 7 *NonStop*™ Clusters software architecture can also provide enhanced performance scalability (compared with nonclustered systems) for large single applications and for workloads consisting of a mix of applications. Typical workloads that scale well are applications that allow work to be distributed across nodes or applications where the amount of internode communications overhead is small, compared to the amount of local work done on the nodes. There can be as much as one-to-one performance scaling, depending on the structure of the applications.

Manageability

With the single system image (SSI) capability, both applications and users have the ability to view the multiple nodes of the cluster as a single logical system. The SSI technology allows the cluster to be managed as a single resource, not a collection of systems. This gives operations and development a consistent, intuitive and familiar view of the entire cluster and its resources. The administrator works with the familiar interfaces and commands of the standard UNIX system administration tools, and they see a single root file system, one set of users and one password file, a single IP host table, and a single network address. As a result, there is no need for new or specialized system management software or additional administration staff to manage the cluster. The addition of resources to any node in the cluster is immediately visible to the cluster as a whole, making the cluster act as a "virtual computer."

Significant Single System Image (SSI) Features

SSI concepts can be applied to applications, specific cluster components or the entire cluster. As such, some product definitions use the term SSI when referring to (for example) an application such as an application oriented cluster where the application has a single image like Oracle Parallel Server clusters. Compaq has implemented "full" SSI cluster technology: for example, one that supports SSI concepts across all elements in the cluster.

The goal of a full SSI cluster is simply to group a number of physically separate servers (or nodes) into what user, application and middleware processes see as a single operating system and server. In particular, these processes should be able to operate in the same way – with no modifications – as they would in a single non-clustered server.

In "full" SSI clusters, the clustering software provides users and applications with a single, consistent view of the underlying platform, regardless of which "node" is physically hosting the process at any one time. SSI clusters achieve this through a number of means:

- Single File Systems: SSI clusters implement a single set of file systems; that is, there is one root file system that all nodes in the cluster use. Devices, log files and network interfaces can all be named and used consistently and uniquely all over the cluster.
- A clustered file system (CFS) layer transparently provides data reliability and integrity for the cluster.
- Shared kernel objects: a user issuing the UNIX ps command is able to see all process on the cluster, not just processes on the one node. Similarly for pipes, sockets and devices.
- Binary compatibility with the non-clustered operating system implementation can be maintained: no application modifications should be necessary.
- An instance of the operating system can execute on each node in the cluster, giving absolute performance advantages for applications, as they don't need to traverse the cluster to access operating system services.
- No one single node in the cluster is designated a permanent "master" node.
- Processes are able to actively migrate from one node to another. In other words, an executing process can pause on one node, migrate to another node and continue execution from the next instruction without restarting.
- Applications can be balanced across the cluster; this load balancing is transparent to the application and user.
- Network traffic can be balanced across multiple interfaces in the cluster. Client connections can be balanced across the cluster, and node failures do not necessarily impact client connections.

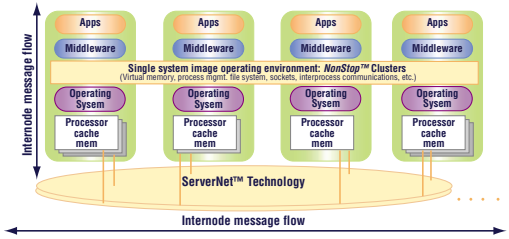

Figure 2 - SCO UnixWare 7 *NonStop*™ Clusters Single

Systems Image

Cluster membership services

Cluster membership services ensure a consistent cross-cluster view of the status of each node and prevent false failover notice. They help to initiate recovery actions after the loss of a node. To ensure a consistent view of node status, the services track and manage the appearance and disappearance of nodes within the cluster. Tasks include detection of node failure, potential node failure or node disappearance; node up and down notification to registered members; and execution of a regroup operation.

Cluster file system

The cluster file system (CFS) environment enhances manageability and ease of maintenance. Because all devices are visible to all nodes, every node can see every device – not just the system disk, but the entire cluster. In addition, because all devices are cluster resources, they can be maintained as pooled resources or partitioned as separately managed devices. For example, local disks and storage devices can be supported concurrently throughout the cluster. The CFS manages the naming tree and provides a single file system hierarchy that is seen identically by all processes on all nodes at all times. When a physical file system mount is performed on any node, it is immediately visible to all processes on all nodes at the same mount point.

The CFS also handles access to data and is layered over the base physical file system to provide coherent distributed access. At any given time, only one node is managing the file systems. The node maintains cache and attribute coherency, so any process can access any file without being cluster aware. Even mapped file operations can be performed coherently across nodes and also provide UNIX system semantics for file access, improve caching operations and ensure cache coherency at all client nodes. The CFS environment maintains file coherency for open files shared among processes on multiple nodes, file attributes and file times, and advisory and mandatory record locks.

Clusterwide devices and pipes

To enable transparent clusterwide access to resources and devices, the SSI provides distributed access to special files that are managed on other nodes, a feature not typically offered by distributed file systems. This works not only for STREAMS special files such as pipes and terminals, but also for mass storage devices such as disk drives. UnixWare 7.1 *NonStop*™ Cluster software also has file system support for transparently accessing remote devices and pipes.

Clusterwide swap and paging service

NonStop™ Clusters software offers a single common view of the aggregated swap and paging space of the cluster so any or all nodes can add swap partitions in the standard SCO UnixWare fashion, and nodes do not need to have a local swap partition.

Clusterwide process management

A common set of standard SCO UnixWare process management functions is available across all of the nodes in a cluster. Processes can execute on any of the available nodes, and there are full SCO UnixWare-compliant semantics for processes, including maintenance of interprocess relationships, visibility of all processes on all nodes and load leveling. Threads (lightweight processes) can be migrated so that all threads of the process are moved at the same time as the process.

Automatic process migration and load balancing

NonStop™ Clusters' Single System Image software allows administrators to migrate applications, operating system objects and processes between cluster nodes. Migrating into default or predefined operating areas balances the load between nodes. Migration can be automatic and the load can be balanced with other tasks or through a failover to specific node groups. Any node can be selected for failover. Automatic process migration and load balancing are available during normal operation or after application restart or failover. This promotes efficient cluster operation without the need for a dedicated standby node.

Competitive alternatives

Today, almost every major vendor has a failover cluster offering. This includes UNIX-based offerings from IBM, Hewlett-Packard and Sun Microsystems. These offerings are all based upon individual vendor implementations of UNIX and proprietary RISC processors. Only SCO UnixWare 7 *NonStop*™ Clusters software is based upon standard Intel processors and can be used in a heterogeneous cluster environment. SCO and its Open Server and UnixWare operating systems dominate IA UNIX units shipments with over 78 percent of the market.

For SCO UnixWare 7 users, *NonStop*™ Clusters software is the best game in town. Furthermore, the non-Intel, more traditional UNIX-based clustering solutions have focused mostly on availability. They lack robustness in the areas of scalability and manageability and have not implemented the depth of single system image functionality that Compaq and SCO have developed with SCO UnixWare 7 *NonStop*™ Clusters software.

While the other UNIX cluster vendors have been incrementally adding SSI-like functionality to their offerings, none currently match up to Compaq *Proliant* Clusters for SCO UnixWare 7. The following table shows the features of Compaq *Proliant* Clusters running *NonStop*™ Clusters for UnixWare software that HRG believes sets them apart from the pack.

Compaq *NonStop*™ Clusters Software Competitive Comparison

Compaq NonStop™ Clusters Feature	*Compaq NonStop™ Clusters software*	*Other UNIX Clusters*
Single system programming model	Yes	No
Disks mirrored internally across nodes	Yes	No
Active process migration	Yes	No
Automatic load balancing	Yes	No
Single Administrative command set	Yes	No
Single file system namespace	Yes	IBM Only
Single device name space	Yes	No
Single IPC namespace	Yes	No
Single root file system	Yes	No
Single IP and port space	Yes	No

Summary and Conclusions

Compaq *Proliant* Clusters running SCO UnixWare 7 *NonStop*™ Clusters offers industry leading scalability and availability for commercial servers while maintaining compatibility with and use of the standard UNIX system commands. With the SSI architecture there is nothing additional needed to manage the cluster. No new system management software or additional administration staff is required to manage the cluster, in either operations or development. SCO UnixWare 7 *NonStop*™ Clusters software gives a consistent, intuitive and familiar view of the entire cluster and its resources. It minimizes planned outages with rolling upgrade capabilities and supports tailored availability deployment options, including the ability for users to deploy continuous-availability (AE-4) solutions.

In particular, the management software provides unique ease of use and administration benefits, such as:

- Active process migration with no need to stop and restart applications and services.
- Automatic load balancing for optimized CPU utilization.
- Same application programming interface (API)/application binary interface (ABI) as base UnixWare 7.1 Operating System.
- Run most standard UnixWare 7.1 applications unchanged.
- Ability to implement cluster-aware applications that scale across multiple nodes.
- Ability to build software continuous availability into the application.
- Easy management via Compaq's Insight Management and Remote Insight Management facilities.

HRG Assessment: Compaq ProLiant Clusters Running NonStop™ Clusters for SCO UnixWare Software

White Paper

The Compaq *ProLiant* Clusters for SCO UnixWare 7 *NonStop*™ Clusters software provides robust application and data availability, a highly reliable operating environment, significant scalability and simplified management. In HRG's opinion, Compaq's 20-plus years of expertise in developing high-performance clusters (obtained through its acquisition of Tandem) has brought *NonStop*™ capabilities to clusters based on industry-standard, Intel-based servers. Their product provides an industry-standard clustered solution that offers a unique combination of application and data recovery, application scalability and single system image (SSI) functionality. Any company seeking a good high availability solution is strongly encouraged to consider what Compaq and SCO have to offer.

HRG believes that Compaq's introduction of *ProLiant* Clusters for SCO UnixWare 7 will provide an exciting addition to Compaq's suite of cluster solutions for Intel industry standard servers. Compaq now offers a full range of X86 offerings with their partners Microsoft, Novell and SCO, thus providing their customers the ability to standardize on *ProLiant* servers and storage throughout the enterprise.

Compaq *ProLiant* Clusters for SCO UnixWare 7 offers some significant configuration advantages. They are available in both tower and rack-mounted configurations that provide flexibility as well as a range of server functionality from the 1-2P PL800 as an entry server to the high performance PL8500 8P server. The ability to mix and match servers within the cluster allows the users to build sophisticated configurations for maximizing application performance and data throughput, while minimizing investment in non-essential hardware. Customers can start with a pair of Compaq *ProLiant* PL800 servers using Cross Node Mirroring at an estimated entry price of approximately $8,000 and then scale to six servers without losing any application uptime. This provides users added flexibility to make better business decisions.

In summary, HRG is impressed with Compaq's offering and expect to see SCO grow into the enterprise space. Compaq *ProLiant* Clusters SSI, price and scalability features combined with industry standard volume economics provide state-of-the-art functional benefits strong enough to challenge any cluster offering.

User Defined Availability Environment Classifications

HRG has defined availability in terms of the impact a system being unavailable to perform work has on the activity of the business and consumer (end user) of the service, rather than the technologies used to achieve it. The five Availability Environment Classifications (AEC) below define availability in terms of the impact on the both the business and the end user or consumer:

- Fault Tolerant (AE-4) – Business functions that demand continuous computing and where any failure is transparent to the user. This means no interruption of work, no transactions lost, no degradation in performance and continuous 24x7 operation.
- Fault Resilient (AE-3) – Business functions that require uninterrupted computing services, either during essential time periods or during most hours of the day and most days of the week throughout the year. This means that the user stays on-line. However, current transaction may need restarting and users may experience performance degradation.
- High Availability (AE-2) – Business functions that allow minimally interrupted computing services, either during essential time periods or during most hours of the day and most days of the week

throughout the year. This means user will be interrupted but can quickly relog on. However, they may have to rerun some transactions from journaled files and they may experience performance degradation.

- Highly Reliable (AE-1) – Business functions that can be interrupted as long as the integrity of the data is ensured. To the user, work stops and uncontrolled shutdown occurs. However, data integrity is ensured.
- Conventional (AE-0) – Business functions that can be interrupted and where the integrity of the data is not essential. To the user, work stops and uncontrolled shutdown occurs. Data may be lost or corrupted.

Disaster Recovery capability is a horizontal availability feature that is applicable to any of the Availability Environments (AEs). It provides for remote backup of the information system and makes it safe from disasters such as an earthquake, fire, flood, hurricane, power failure, vandalism or act of terrorism.

Harvard Research Group is an information technology market research and consulting company. The company offers highly focused market research and consulting services to vendors and users of computer hardware, software and services. For more information, contact Harvard Research Group as follows:

Harvard Research Group™
Box 297
Harvard, MA 01451 USA

Tel. (978) 263-3399
Fax (978) 263-0033

E-mail hrg@hrgresearch.com
www.hrgresearch.com

MODELS

Compaq *ProLiant* Clusters for SCO UnixWare U/ML330 kit 155370-B22

The U/ML330 kit is designed to enable two-node *ProLiant* ML330-based clusters and includes a Quick Cluster Install kit and SCO software media for rapid and simplified installation. (Note: UnixWare 7.1.1 and *NonStop™* Clusters 7.1.1b+IP software license keys are purchased separately from a SCO reseller or distributor prior to production use.)

The added *ProLiant* Clusters for SCO UnixWare Quick Install Image CDs, Install Wizard, *NonStop™* Clusters Verification and UPS Power Management Utilities, ML330 Cluster Reference Guide and ML330 Cluster configuration poster provide the customer with the tools for a rapid and successful configuration and installation.

Description

Compaq *ProLiant* Clusters for SCO UnixWare 7 U/ML330 kit is designed for a two-node *ProLiant* ML330-based high availability cluster configuration. The U/ML330 kit supports clusters with the Ethernet IP node-to-node interconnect and the *NonStop™* Cross Node Mirrored reliable storage option.

The kit includes an unique Quick Install package with an installation wizard, ML330-specific installation poster and cluster reference guide, Ethernet node interconnect and cluster integrity cables and UnixWare 7.1.1 and *NonStop™* Clusters 7.1.1b+IP media kits – all designed for a quick and simplified installation.

Note: UnixWare 7.1.1 and NonStop™ *Clusters 7.1.1b+IP software license keys are purchased separately from a SCO reseller or distributor.*

The kit is ideal for entry-level applications, replicated sites and customers with limited on-site computer expertise. Compaq *ProLiant* Clusters for SCO UnixWare 7 U/ML330 kit and Compaq hardware are combined with the industry-standard SCO UnixWare 7.1.1 Operating System and *NonStop™* Clusters 7.1.1b+IP software available from a SCO reseller or distributor.

The two-node cluster solution is offered in a tower configuration. A wide variety of disk, tape backup, network interface controllers and remote management options are supported.

STANDARD FEATURES

Rapid and Simplified Installation

The U/ML330 kit includes one Quick Install Cluster Image CD per node for rapid and simplified installation. The necessary software is contained on the set of Quick Install CDs, including UnixWare 7 and *NonStop™* Clusters 7 already configured for immediate cluster boot. (Note: UnixWare 7.1.1 and *NonStop™* Clusters 7.1.1b+IP software license keys are purchased separately from a SCO reseller or distributor.) The Quick Install kit is designed for a successful installation of the operating and cluster software within approximately one hour. An install wizard is provided allowing customer-specific parameters and input of operating and cluster software licenses. An ML330 cluster Reference Guide, install poster and cluster verification utilities provide additional tools for successful configuration, installation and operation.

Maximum Availability, Minimum Cost

Because Compaq *ProLiant* Clusters are built from industry-standard components, Compaq delivers high levels of application availability at a much lower cost than traditional, proprietary cluster solutions. The *ProLiant* ML330 offers the latest technology at the most affordable price. With the addition of two ML330 servers and two 9.1GB or greater hard disk drives, the cluster kit includes all the hardware necessary to configure a *ProLiant* Cluster for SCO UnixWare. The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/ML330 kit utilizes the very economical Ethernet IP node interconnect and Cross Node Mirroring reliable storage options. No cluster interconnect hardware is required as the embedded Ethernet included with the ML330 server is supported. The cross node mirroring provides completely redundant and reliable storage and is included with all Compaq *ProLiant* Clusters for SCO UnixWare 7. Only one hard drive per server is necessary to implement this storage solution.

Integrated Solutions

All Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution. The solution brings industry leading cluster technology to industry-standard servers for the UNIX community. Compaq *ProLiant* servers,

interconnect options, system management software, flexibility in configuration support and implementation documentation have all been thoroughly tested in cluster configurations. *NonStop™* Clusters software is developed by Compaq for UnixWare and has resulted in a very close partnership with SCO. The relationship between Compaq and application partners has resulted in the development of cluster application TechNotes, which assist in the rapid deployment of major business applications on Compaq *ProLiant* Clusters.

Sophisticated Cluster Management

In addition to the Quick Install CDs, all Compaq *ProLiant* Clusters for SCO UnixWare 7 utilize Compaq *SmartStart* and Compaq *Insight Manager* software to further enable simplified installation, configuration and management. Compaq *SmartStart* and Compaq *Insight Manager* are "cluster enabled," allowing easy and sophisticated cluster and systems management. *ProLiant* Clusters for SCO UnixWare 7 also support the Compaq Remote Insight Board for remote monitoring, system administration and diagnostics, providing yet another level of cluster management.

COMPONENTS

Servers

The U/ML330 kit supports any two *ProLiant* ML330 servers; the two servers must be identical. Please check the Compaq High Availability Web site for an up-to-date listing of certified *ProLiant* Clusters for SCO UnixWare 7 configurations at **www.compaq.com/solutions/enterprise/highavailability/sco/cpqnsc-pd.html**.

Interconnect

The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/ML330 kit offers a direct Ethernet IP connection. Using the kit's Ethernet interconnect cable, the Ethernet included as standard with the *ProLiant* ML330 is used for the interconnect hardware. The Ethernet interconnect is configured as a private network on its own subnet. A second physical Ethernet connection is required for the public local area network (LAN).

Storage

Cross Node Mirroring reliable storage is included with every ML330-based cluster. The Cross Node Mirroring option mirrors the internal drives in the two nodes providing a reliable storage option. It requires only one internal hard disk drive per server. Additional internal drives may be added to increase storage capacity. An ML330 cluster currently supports from 9.1GB to 91GB of reliable mirrored storage using 1" 18.2GB drives. A tape back up drive per node may be added to generate even a third copy of the data.

Supported Operating Systems

SCO UnixWare 7.1.1 (all editions) and *NonStop™* Cluster software 7.1.1b+IP.

Compaq ProLiant Clusters for SCO UnixWare 7 U/ML330 Kit Contents

- Compaq *ProLiant* Clusters for SCO UnixWare 7 ML330 Quick Install CD Set Version 1.0 with Install Wizard
- Compaq *ProLiant* Clusters for SCO UnixWare 7 ML330 Install Guide (in English)
- Cluster Verification and UPS Power Management Utilities Version 1.2 (included on the Quick Install CDs)
- Compaq *ProLiant* Cluster for SCO UnixWare 7 *NonStop™* Clusters ML330 configuration poster (in English)
- SCO UnixWare 7.1.1 media kit (in English)
- NonStop Clusters 7.1.1b+IP media kit (in English)
- Cluster integrity serial cable (12')
- Ethernet interconnect cable (12')
- Cluster node labels

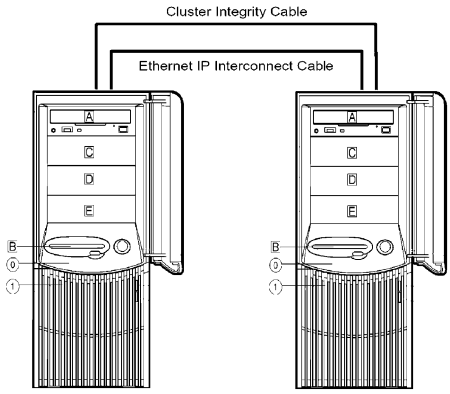

Non-Hot Plug Drive Cage

A	3.5″ diskette drive
B	32X Max IDE CD-ROM
C, D, E	Removable media bays (1.6″ each)
0, 1	Two 1″ non-hot plug drive bays

OPTIONS

Most server and rack options are supported in the cluster. Refer to the Compaq high availability Web site at **www.compaq.com/solutions/enterprise/high availability/sco/cpqnsc-pd.html** for the latest list of certified options, supported configurations, release notes and restrictions.

Compaq ProLiant Clusters for SCO UnixWare U/ML330 Kit
155370-B22

Note: One kit required per cluster.

ProLiant Model ML330 Servers

Note: Any two identical ProLiant ML330 servers required for each cluster.

Model ML330T01 P667-256K	***171469-xx1***
Model ML330T01 P667-256K 9.1GB	***171469-xx2***
Model ML330T01 P733-256K	***157797-xx1***
Model ML330T01 P733-256K 9.1GB	***157797-xx2***

Country Code Key

XX=00	**US**
XX=01	**Australia**
XX=02	**EURO**
XX=03	**UK**
XX=04	**Germany**
XX=05	**France**
XX=29	**Japan**
XX=37	**APD**
XX=AA	**PRC**

Service and Support

Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty, 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-failure warranty is included and applies to certain hard drives, memory and processors for servers monitored by Compaq 2.0 or higher. For information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

Compaq *Carepaq* service upgrades are available for your Compaq server and storage hardware including:

CarePaq Category	Service Name	Service Description
Hardware services	Next business day on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	24-hour x 7-day coverage
Installation & start-up services	Hardware installation	

For specific part numbers and additional information, contact your Compaq Authorized Service Provider or see the *Carepaq* server options listed at **www.compaq.com/services/carepaq/cp_matrix_servers.html** and storage at **www.compaq.com/services/carepaq/cp_matrix_storage.html**

Note: Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

MODELS

Compaq *ProLiant* Clusters for SCO UnixWare U/100 kit 155370-B21

The U/100 kit is designed to enable a wide variety of two- to six-node, highly available, and application scaling cluster solutions. The kit offers cluster capability on virtually all Compaq entry, midrange and high-end *ProLiant* servers. The added Enhancements CD (with *NonStop™* Cluster Verification, ServerNet Verification and UPS Power Management Utilities), Cluster Reference Guide and Configuration Poster provide the user with systems verification and installation tools to ensure a successful configuration for either an entry *NonStop™* Cross Node Mirrored Storage or an External Fibre Channel Storage implementation.

Description

Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 Kit is designed to enable highly available two- to six-node scalable clusters connected by *ServerNet*, or two-node Ethernet IP connected solutions. The kit supports nearly the entire range of Compaq's industry-leading *ProLiant* servers, Compaq *StorageWorks* RAID Array 4000 (RA4000) and RAID Array 4100 (RA4100), *ServerNet* or Ethernet IP node-to-node interconnect and Compaq installation and systems management utilities.

Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 Kit and Compaq hardware are combined with the industry-standard SCO UnixWare 7.1.1 Operating System and *NonStop™* Clusters 7.1.1b+IP software available from a SCO provider. *NonStop™* Clusters software, developed by Compaq for distribution by SCO, provides up to 99.995 percent system availability for applications and data in business-critical SCO UnixWare environments.

The single processor *ProLiant* ML330 to the 8-Way SMP *ProLiant* 8500 may be mixed and matched in a wide variety of combinations to configure and scale a cluster from two to six nodes and two to 64 processors. *NonStop™* Cross Node Mirrored or *StorageWorks* RA400/RA4100 External Fibre Channel Reliable Storage options are available. Clusters are offered in either tower or rack-mounted configurations and can be configured with no single point of failure. A wide variety of disk, tape backup, network interface controllers, remote management and rack options are supported.

STANDARD FEATURES

Maximum Availability, Minimum Cost

Because Compaq *ProLiant* Clusters are built from industry-standard components, Compaq delivers high levels of application availability at a much lower cost than traditional, proprietary cluster solutions. The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 kit utilizes industry-leading Compaq *ProLiant* servers, the fast and reliable fibre-connected Compaq *StorageWorks* RAID Array 4000 (RA4000) and RAID Array 4100 (RA4100), and *ServerNet* Fault Tolerant Node interconnect for fast reliable and consistent inter-node cluster communications and mirrored data writes. The *ServerNet* solution supports from two to six node configurations with no single point of failure providing another level of availability not offered on many other cluster systems. Low cost Ethernet IP node interconnect and cross node mirroring reliable storage options are also available for two-node configurations.

Investment Protection

Both existing and new Compaq *ProLiant* servers are certified for use with the Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 kit. Customers can build clusters using existing servers or by mixing old and new servers. This means that an existing cluster can be upgraded in the future without the need to replace hardware. Customers who purchased certified *ProLiant* servers and storage for non-clustered applications can migrate to this highly available and no single point of failure solution by purchasing a *ProLiant* Cluster for SCO UnixWare 7 U/100 Kit, node interconnect hardware, and UnixWare 7 and *NonStop™* Cluster software licenses and media kits.

Integrated Solutions

All Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution. The solution brings industry-leading cluster technology to industry-standard servers for the UNIX community. Compaq *ProLiant* servers, interconnect options, system management software, flexibility in configuration support and implementation documentation have all been thoroughly tested in cluster configurations. *NonStop™* Clusters software is developed by Compaq for UnixWare and has resulted in a very close partnership with SCO. The relationship between Compaq and application partners has resulted in the development of cluster application TechNotes, which assist in the rapid deployment of major business applications on Compaq *ProLiant* Clusters.

Sophisticated Cluster Management

Compaq *ProLiant* Clusters for SCO UnixWare 7 utilize Compaq *SmartStart* and Compaq *Insight Manager* software to enable simplified installation, configuration and management. Compaq *SmartStart* and Compaq *Insight Manager* are "cluster enabled," allowing easy installation and sophisticated cluster and systems management. *ProLiant* Clusters for SCO UnixWare 7 also support the Compaq Remote Insight Board for remote monitoring, system administration and diagnostics, providing yet another level of cluster management.

COMPONENTS

Servers

A very broad range of *ProLiant* server models are supported, in both matched-pair and mixed-pair combinations. Please check the Compaq High Availability Web site at **www.compaq.com/highavailability/sco** for an up-to-date listing of certified *ProLiant* Clusters for SCO UnixWare 7 configurations, restrictions and releases notes.

Interconnect

The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 Kit offers two different node-to-node interconnect options. One *ServerNet* low latency, high bandwidth PCI host bus adapter may be used per node for high performance two to six node configurations. Also, if *ServerNet*-connected clusters are being configured for three to six nodes, the fault-redundant *ServerNet* switch pair is required. For entry-level, two-node configurations, a direct Ethernet IP connection is a low-cost alternative. Either the Ethernet included as standard with the *ProLiant* nodes or one of the certified Compaq fast Ethernet controllers may be used for the interconnect hardware. Please check the Compaq High Availability Web site at **www.compaq.com/highavailability/sco** for an up-to-date listing of certified interconnect configurations and Ethernet controllers. The Ethernet interconnect is configured as a private network on its own subnet. A second physical Ethernet connection is required for the public local area network (LAN).

Storage

The U/100 kit supports both internal disk storage clusters using the Cross Node Mirroring reliable storage function as well as external storage using the Compaq *StorageWorks* RAID Array 4000 (RA4000) or the RAID Array (RA4100) in a redundant, reliable implementation. The Cross Node Mirroring option mirrors the internal drives in the first two nodes and may be implemented on any size cluster. The *StorageWorks* RAID Array option is a high performance reliable storage solution providing up to 291.2GB of storage, using 1.6″ 36.4GB drives. Please check the Compaq High Availability Web site at **www.compaq.com/highavailability/sco** for an up-to-date listing of certified storage configurations.

Supported Operating Systems

SCO UnixWare 7.1.1 (all editions) and *NonStop™* Cluster software 7.1.1b+IP.

Compaq *ProLiant* Clusters for SCO UnixWare 7 U/100 Kit Contents

- Compaq *ProLiant* Clusters for SCO UnixWare 7 Install Guide (in English)
- Compaq *ProLiant NonStop™* Clusters Enhancements CD Version 1.2 (with the Cluster, *ServerNet* Verification and UPS Power Management Utilities)
- Compaq *ProLiant* Cluster for SCO UnixWare 7 *NonStop™* Clusters configuration poster (in English)
- Cluster Integrity serial cable (12′)
- Ethernet interconnect cable (12′)
- Cluster node labels

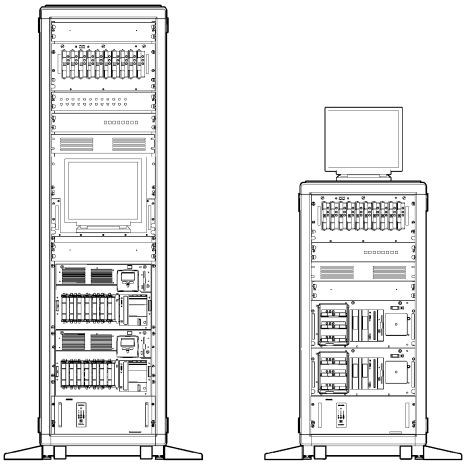

OPTIONS

Compaq ProLiant Clusters for SCO UnixWare U/100 Kit
155370-B21

ServerNet Node Interconnect

ServerNet PCI Adapter
309833-B22, 309833-292 (Japan)

ServerNet Thin Cable (4m)
309834-B23

ServerNet Cable (30m)
309834-B22

ServerNet I 6-Port Switch
327337-B31, 327337-291 (Japan)

ServerNet I Switch Slide Mount Rail Kit
327724-B21

Ethernet Interconnect

Note: The Ethernet included with the server is supported and is the recommended node-to-node interconnect. Alternately, one of the optional listed fast Ethernet network interface controllers may be added; one per server is required. The Ethernet node interconnect is configured as a private network on its own subnet. A second physical Ethernet connection is required for the public local area network (LAN).

Compaq NC3131 Fast Ethernet NIC 64 Dual Base 10/100 PCI Adapter	***338456-B21***
Compaq NC3134 Fast Ethernet NIC 64 Dual Port 10/100 PCI Adapter	***138603-B21***
Compaq NC3135 Fast Ethernet Upgrade Module, Dual 10/100	***138604-B21***

Service and Support

Compaq servers and storage systems are protected by Compaq Services, including a three-year limited warranty, 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-failure warranty is included and applies to certain hard drives, memory and processors for servers monitored by Compaq *Insight Manager 2.0* or higher. For information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

Compaq *CarePaq* service upgrades are available for your Compaq server and storage hardware including:

CarePaq Category	*Service Name*	*Service Description*
Hardware services	Next business day on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	24-hour x 7-day coverage
Installation & start-up services	Hardware installation	

For specific part numbers and additional information, contact your Compaq Authorized Service Provider or see the CarePaq server options listed at **www.compaq.com/services/carepaq/cp_matrix_servers.html** and storage at **www.compaq.com/services/carepaq/cp_matrix_storage.html**

Note: Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

MODELS

Compaq *ProLiant* Clusters for SCO UnixWare U/CL380 kit 155370-B23

The U/CL380 kit is designed to enable *ProLiant* CL380-based clusters and includes a Quick Install kit and SCO software media for rapid and simplified installation. (Note: UnixWare 7.1.1 and *NonStop™* Clusters 7.1.1b+IP software licenses are purchased separately from a SCO reseller or distributor.)

The added *ProLiant* Clusters for SCO UnixWare Quick Install Image CDs, Install Wizard, *NonStop™* Clusters Verification and UPS Power Management Utilities, CL380 Cluster Reference Guide and CL380 Cluster configuration poster provide the customer with the tools for a rapid and successful configuration and installation.

Description

Compaq *ProLiant* Clusters for SCO UnixWare 7 U/CL380 kit is designed for a *ProLiant* CL380-based high availability cluster configuration. The U/CL380 kit supports clusters with the Ethernet IP node to node interconnect and the CL380's shared SCSI RAID storage system.

The kit includes a unique Quick Cluster Install Package with an installation wizard, CL380-specific installation poster and cluster reference guide, cluster integrity cable and UnixWare 7 and *NonStop™* Clusters media kits – all designed for a quick and simplified installation.

Note: UnixWare 7.1.1 and NonStop™ *Clusters 7.1.1b+IP software licenses are purchased separately from a SCO reseller or distributor.*

The kit is ideal for entry-level applications, replicated sites and customers with limited on-site computer expertise. Compaq *ProLiant* Clusters for SCO UnixWare 7 U/CL380 kit and Compaq hardware are combined with the industry-standard SCO UnixWare 7.1.1 Operating System and *NonStop™* Clusters 7.1.1b+IP software available from a SCO reseller or distributor. The two-node cluster solution is offered in either a pre-package tower or rack-mounted configuration. A wide variety of disk, tape backup, network interface controllers, remote management and rack options are supported.

STANDARD FEATURES

Rapid and Simplified Installation

The U/CL380 kit includes one Quick Install Cluster Image CD per node for rapid and simplified installation. The necessary software is contained on the set of Quick Install CDs, including UnixWare 7 and *NonStop™* Clusters 7 already configured for immediate cluster boot. The U/ML330 kit also contains a Quick Install Configuration Poster specifically targeted for CL380 clusters utilizing the Ethernet IP node interconnect method.

The Quick Install kit is designed for a successful installation of the operating and cluster software within approximately one hour. An install wizard is provided allowing customer-specific parameters and input of operating and cluster software licenses. A CL380 cluster Reference Guide, install poster and cluster verification utilities provide additional tools for successful configuration, installation and operation.

Pre-packaged Cluster Design

The *ProLiant* CL380 server is a unique pre-packed two-node system in a single, space efficient cabinet. In addition to the internal storage, the integrated and pre-wired reliable RAID storage system supports up to six 1.0″ drives. The CL380 includes both embedded and add-in 10/100 fast Ethernet network interface controllers per server for node-to-node interconnectivity and local area network access, respectively. The addition of one 9.1GB or greater internal drive per node (two of total per CL380 server) provides all the necessary hardware required for a Compaq *ProLiant* Cluster for SCO UnixWare. The CL380 is a packaged two-node configuration fully integrated into the included tower or rack-mounted cabinet.

Maximum Availability, Minimum Cost

Because Compaq *ProLiant* Clusters are built from industry-standard components, Compaq delivers high levels of application availability at a much lower cost than traditional proprietary cluster solutions. The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/CL380 kit also utilizes the very economical Ethernet IP node interconnect and the CL380's included reliable RAID storage system. No cluster interconnect hardware is required as the embedded Ethernet included with the CL380 server is supported.

Integrated Solutions

All Compaq *ProLiant* Clusters integrate hardware and software to provide a total solution. The solution brings industry-leading cluster technology to industry-standard servers for the UNIX community. Compaq *ProLiant* servers, interconnect options, system management software, flexibility in configuration support and implementation documentation have all been thoroughly tested in cluster configurations. *NonStop™* Clusters software is developed by Compaq for UnixWare and has resulted in a very close partnership with SCO. The Compaq relationship with application partners has resulted in the development of cluster application TechNotes, which assists in the rapid deployment of major business applications on Compaq *ProLiant* Clusters.

Sophisticated Cluster Management

In addition to the Quick Install CDs, all Compaq *ProLiant* Clusters for SCO UnixWare 7 utilize Compaq SmartStart and Compaq Insight Manager software to further enable simplified installation, configuration and management. Compaq SmartStart and Compaq Insight Manager are "cluster enabled," allowing easy and sophisticated cluster and systems management. *ProLiant* Clusters for SCO UnixWare 7 also support the Compaq Remote Insight Board for remote monitoring, system administration and diagnostics, providing yet another level of cluster management.

COMPONENTS

Servers

The U/CL380 kit supports any *ProLiant* CL380 server. Please check the Compaq high availability Web site for an up-to-date listing of certified *ProLiant* Clusters for SCO UnixWare 7 configurations at **www.compaq.com/solutions/enterprise/highavailability/sco/cpqnsc-pd.html**.

Interconnect

The Compaq *ProLiant* Clusters for SCO UnixWare 7 U/CL380 kit offers a direct Ethernet IP connection. Using the Ethernet interconnect cable and embedded fast Ethernet included with the CL380 server, all the hardware is provided for node-to-node connectivity. The Ethernet interconnect is configured as a private network on its own subnet. The CL380's included second Ethernet adapter is used for the public local area network (LAN).

Storage

The SCSI RAID storage system included with the *ProLiant* CL380 is utilized for shared and reliable storage. Up to six 1.0" hot-pluggable drives are supported for a maximum of 109.2GB of storage, using 18.2GB drives. A tape back up drive per node may be added to generate even additional data protection.

Supported Operating Systems

SCO UnixWare 7.1.1 (all editions) and *NonStop™* Cluster software 7.1.1b+IP.

Compaq ProLiant Clusters for SCO UnixWare 7 U/CL380 Kit Contents

- Compaq *ProLiant* Clusters for SCO UnixWare 7 CL380 Quick Install CD Set Version 1.0 with Install Wizard
- Compaq *ProLiant* Clusters for SCO UnixWare 7 CL380 Install Guide (in English)
- Cluster Verification and UPS Power Management Utilities Version 1.3 (installed on the Quick Install CDs)
- Compaq *ProLiant* Clusters for SCO UnixWare 7 CL380 Quick Start Configuration Poster (in English)
- SCO UnixWare 7.1.1 Media kit (in English)
- *NonStop™* Clusters 7.1.1b+IP Media kit (in English)
- Cluster Integrity Serial Cable (12FT)

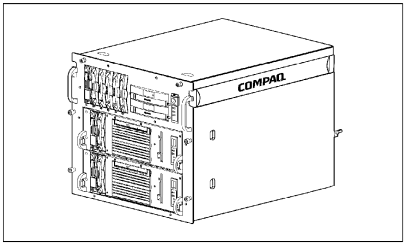

Figure 1-1. Compaq ProLiant CL380 rack model

OPTIONS

Most server and rack options are supported in the cluster, refer to the Compaq high availability Web site at **www.compaq.com/solutions/enterprise/highavailability/sco/cpqnsc-pd.html** for the latest list of certified options, supported configurations, release notes and restrictions.

Compaq ProLiant Clusters for SCO UnixWare U/CL380 Kit	***155370-B23***

Note: One kit required per cluster.

ProLiant Model CL380 Servers

Note: Any one ProLiant CL380 server (includes two nodes) and at least one internal hard driver per node required for each cluster.

Model CL380 6/800 (tower)	***164607-xx1***
Model CL380 6/800 (rack)	***164608-xx1***

Country Code Key

XX=00	**US**
XX=29	**Japan**
XX=37	**APD**
XX=42	**EURO**

Service and Support

Compaq servers and storage systems are protected by Compaq Services including a three-year limited warranty, 7 x 24 hardware technical phone support and online support through CompuServe, Prodigy, America Online and the Internet. Pre-failure warranty is included and applies to certain hard drives, memory and processors for servers monitored by Compaq Insight Manager 2.0 or higher. For information on support offerings for your Compaq hardware, contact your Compaq Authorized Service Provider.

For specific part numbers and additional information, contact your Compaq Authorized Service Provider or see the CarePaq server options listed at **www.compaq.com/services/carepaq/cp_matrix_servers.html** and storage at **www.compaq.com/services/carepaq/cp_matrix_storage.html**

Note: Certain restrictions and exclusions apply. Consult the Compaq Customer Support Center at 1-800-345-1518 for details.

Compaq *CarePaq* service upgrades are available for your Compaq server and storage hardware including:

CarePaq Category	*Service Name*	*Service Description*
Hardware services	Next business day on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	9-hour x 5-day coverage
Hardware services	4-hour on-site service	24-hour x 7-day coverage
Installation & start-up services	Hardware installation	


Compaq's ServerNet II
System Area Network
COMPAQ
www.compaq.com/highavailability

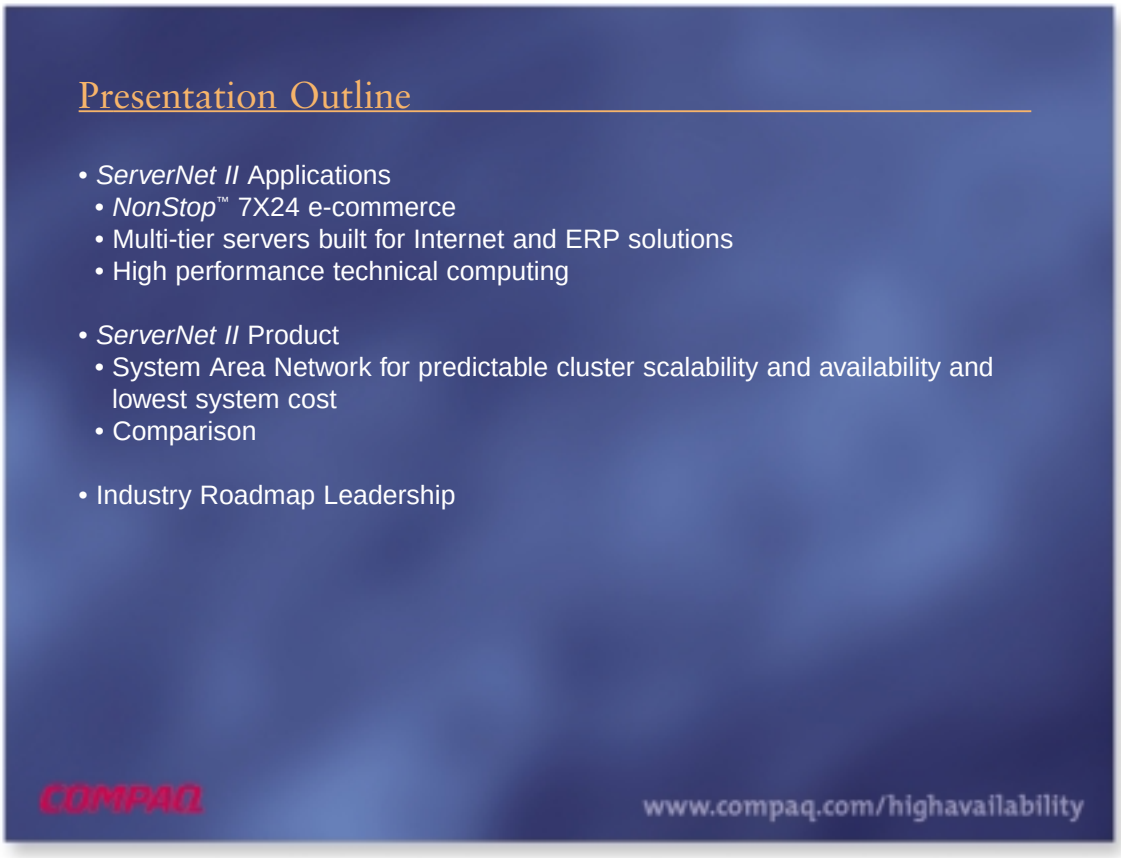
Presentation Outline
• ServerNet II Applications
• NonStop™ 7X24 e-commerce
• Multi-tier servers built for Internet and ERP solutions
• High performance technical computing
• ServerNet II Product
• System Area Network for predictable cluster scalability and availability and lowest system cost
• Comparison
• Industry Roadmap Leadership
COMPAQ
www.compaq.com/highavailability

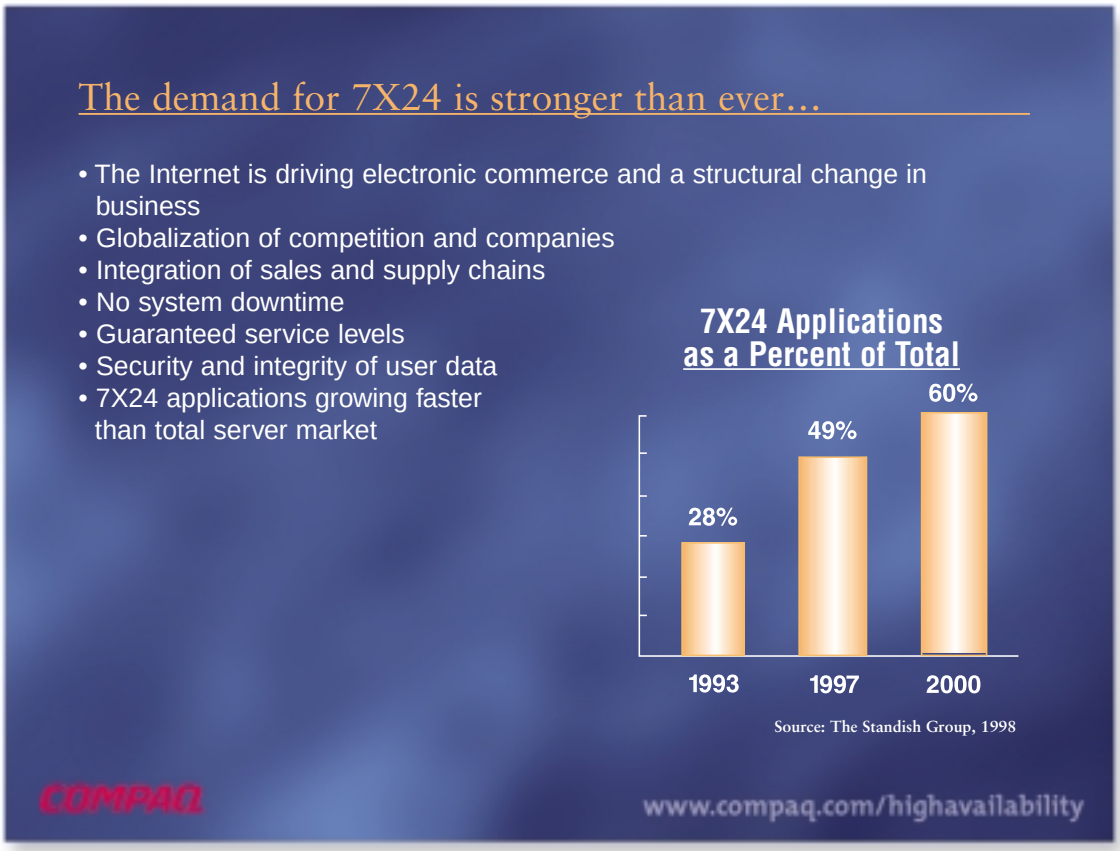

The global and 7x24 demands of the Internet and e-commerce are challenging the scalability, reliability, availability and performance of servers.

Today's computing model is becoming more distributed as companies work to meet the unceasing demands of the Internet economy.

Customers are increasingly appreciating the value of zero downtime. The cost of downtime varies by industry but can be very high depending on the application. In the case of certain Internet companies, for example, eBay, downtime has even impacted market capitalization.

These customer requirements drive cluster market opportunity, especially in the 7x24 applications as noted by the Standish Group study above. Market data from the Gartner/Dataquest Clustering forecast published July 1999 indicates an 82 percent CAGR (1998-2003) for NT clusters and 41 percent CAGR for Unix clusters.

Applications require clustering capabilities to achieve the highest availability and scalability levels demanded by 7x24 applications.

These applications exhibit the following characteristics:

- Require low latency data interchange
- Generate high inter-node transaction traffic
- Utilize distributed data
- Require extended (7x24x365) operation

These characteristics apply to e-commerce applications, Application Server Providers (ASP), Internet Service Providers (ISP), Enterprise Resource Planning (ERP), real-time (zero latency) data warehousing/data mining applications and high performance scientific, engineering and educational applications.

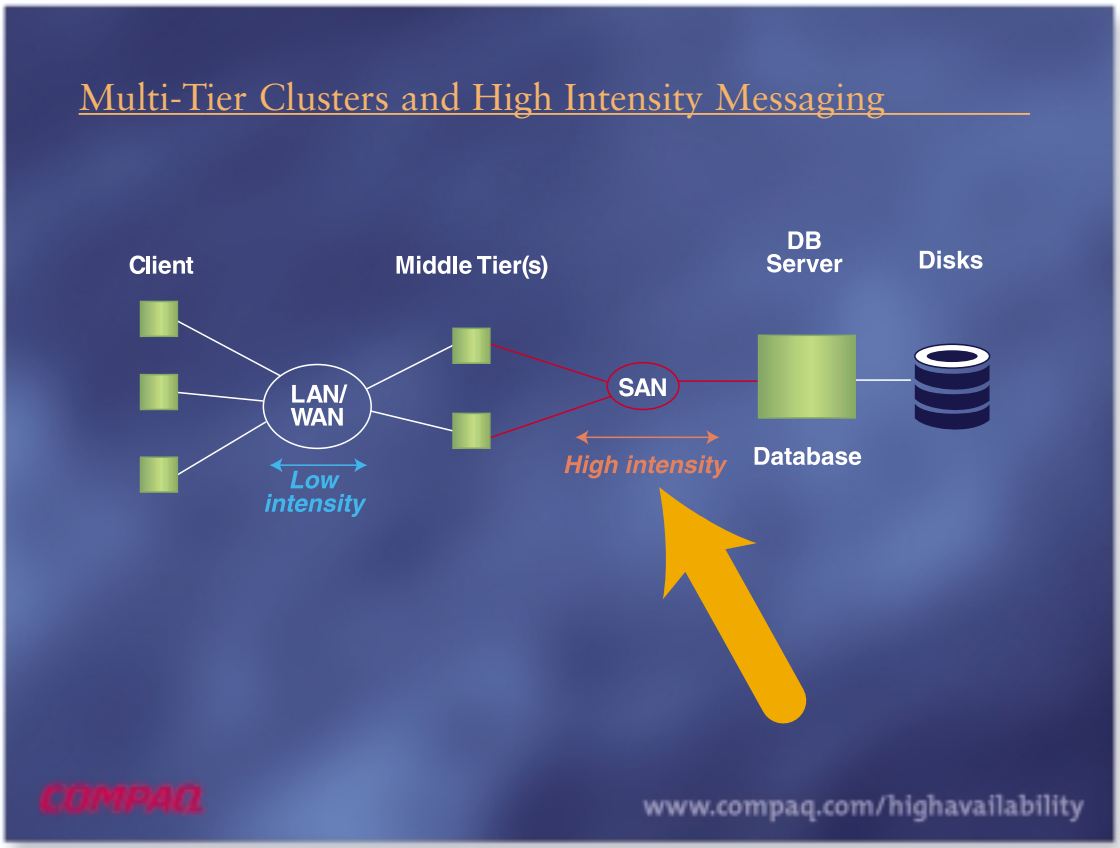

- Customers deploying three-tier, high intensity messaging applications are finding that the network overheads are reducing their ability to support more users.
- By using a VI-compliant interconnect like ServerNet, the high intensity messaging overhead (CPU utilization) is significantly reduced. The release of CPU cycles to the applications allows more users to be supported.
- In cases where the networking overhead is high, adding additional servers will not yield an improvement in the number of users that can be supported. Only a VI compliant interconnect like ServerNet can yield results.

Example setup for SAP Sales Distribution (SD) benchmark:

- Compaq PL5500R 4P Xeon 1MB L2
- 19 Application servers (62 percent utilized)
- One Oracle 8.0.5 Database server (99 percent utilized)

SAN traffic flows:

- Client to middle tier traffic is low intensity
- Middle tier to database traffic is high intensity SQL*Net
- For OPS, traffic flows among nodes in database cluster

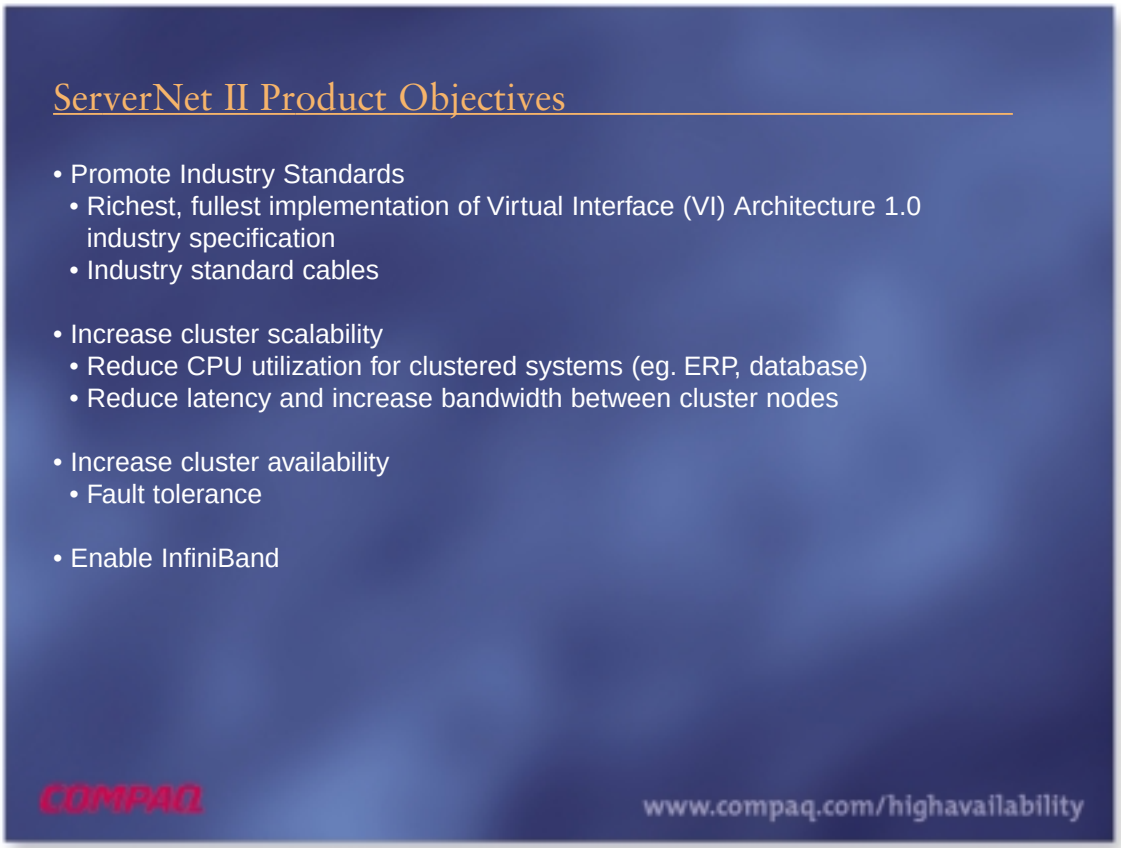
ServerNet II Product Objectives
• Promote Industry Standards
• Richest, fullest implementation of Virtual Interface (VI) Architecture 1.0 industry specification
• Industry standard cables
• Increase cluster scalability
• Reduce CPU utilization for clustered systems (eg. ERP, database)
• Reduce latency and increase bandwidth between cluster nodes
• Increase cluster availability
• Fault tolerance
• Enable InfiniBand
COMPAQ
www.compaq.com/highavailability

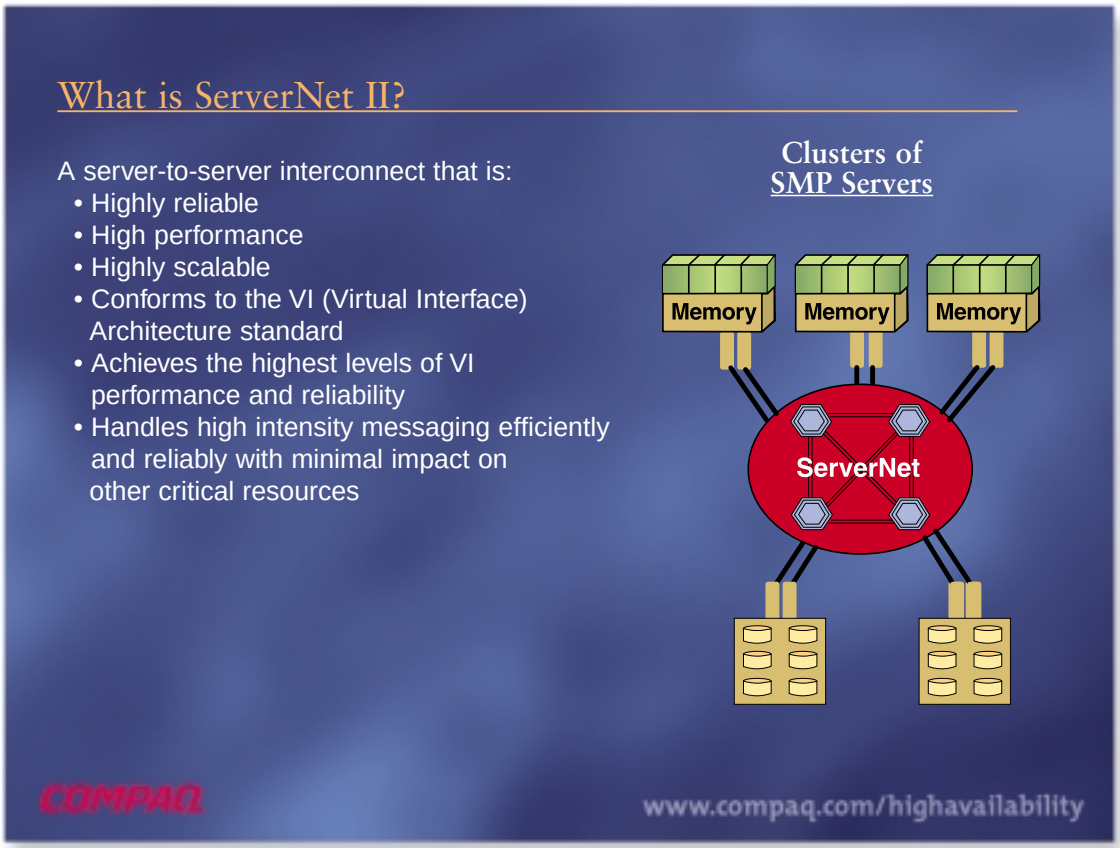

- ServerNetII is a high performance, low latency network used to interconnect servers with other servers in a System Area Network (SAN).

- System Area Networks are reliable, high performance, point-to-point networks for communications between independent servers within a datacenter.

- *ServerNet II* conforms to the industry-specified Virtual Interface (VI) specification and implements the highest levels of performance and reliability.

- *ServerNet II* was designed for NonStop Internet and data center environments. It adds robustness, scalability and flexibility to the business computing cluster by handling high-intensity messaging efficiently and reliably with minimal impact to other critical systems resources.

- With *ServerNet II*, it is possible to put together a cluster of standards based systems and achieve mainframe class computing at very low cost.

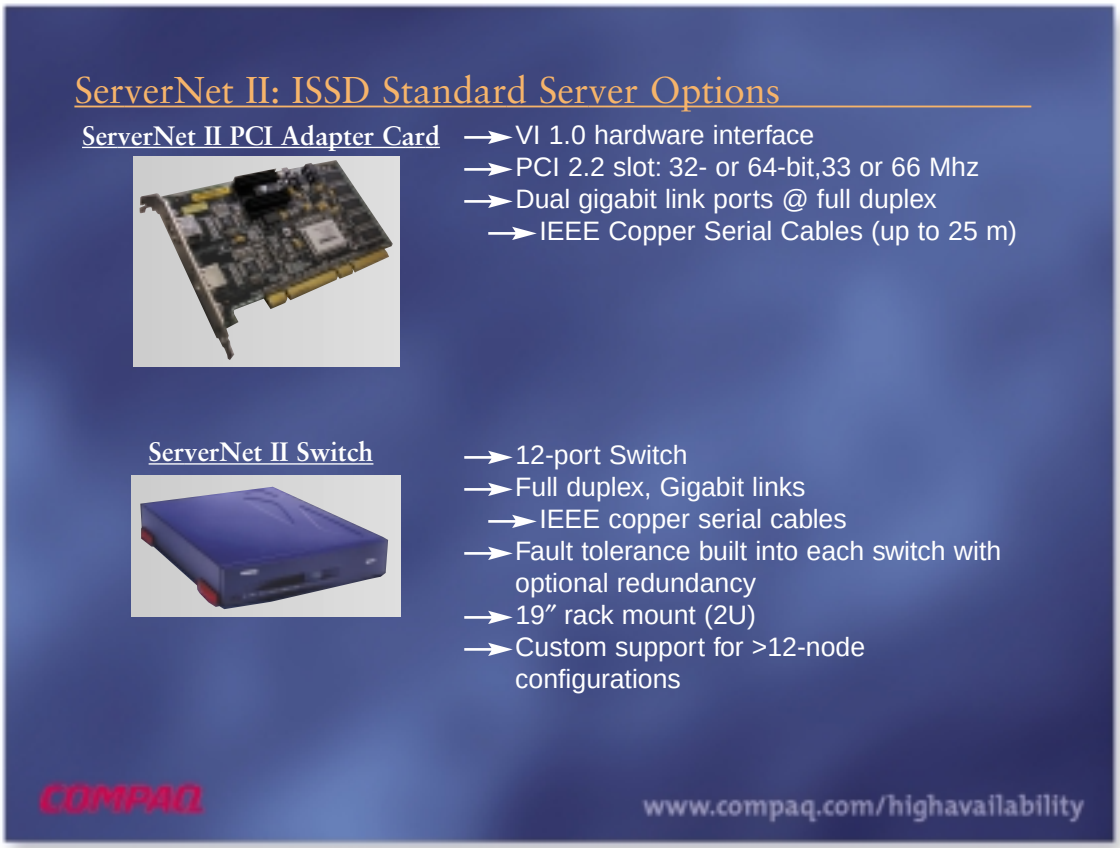

The ServerNetII PCI adapter fits into a standard 32- or 64-bit PCI bus running at either 33 or 66 MHz speeds.

The Compaq-designed ASIC implements the VI spec at the highest reliability and performance levels. The ServerNet II PCI Adapter delivers advantages over traditional networking providing more efficient (lower) CPU utilization, improved (lower) latency and greater bandwidth.

The *ServerNet II* PCI Adapter is used in a point-to-point configuration that can be expanded to very large networks utilizing the Compaq *ServerNet II* Switch.

The *ServerNet II* Switch requires 2U in a 19″ Rack. The *ServerNet II* Switch supports up to twelve connections. Specialized *ServerNet II* Switch configurations exceeding twelve connections can be created from multiple switches cascaded into symmetric or asymmetric topologies.

A fault-tolerant network is achieved by configuring redundant switch fabrics to each of the two ports of a *ServerNet II* PCI Adapter and utilizing two switches per fabric.

Each of the twelve connections support point-to-point

data links of 1.25 + 1.25 gigabits per second of bi-directional bandwidth giving an aggregate switch bandwidth of 15 gigabits per second. Effective bandwidth can be increased between *ServerNet II* Switches to 5.0 + 5.0 gigabits per second in specialized configurations by merging sets of parallel links.

The *ServerNet II* Switch path formation latency for up to twelve incoming packets switched (non-blocking) to twelve outgoing routes is approximately 300 nanoseconds. Routing tables support arbitrary network topologies.

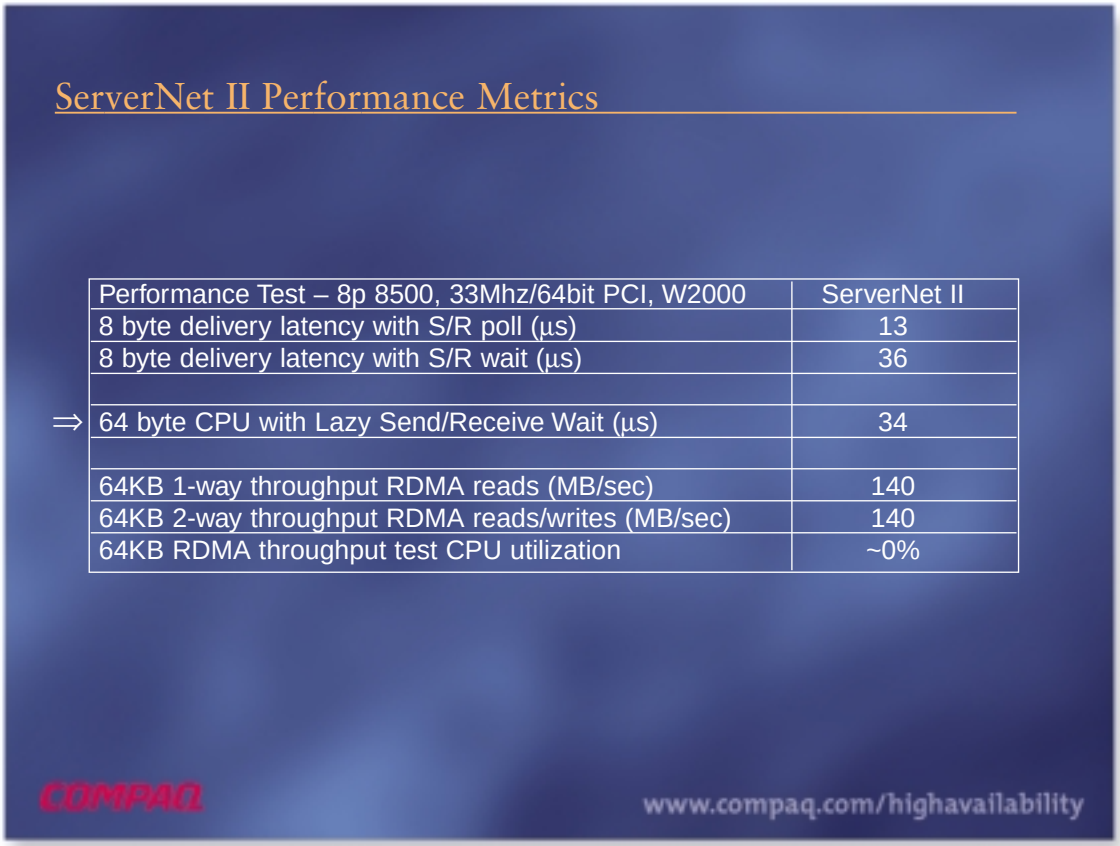

The key performance parameters for VI are:

- Latency i.e. message delivery time measured in microseconds
- Throughput measured in megabytes per second
- CPU Utilization: the fewer CPU cycles consumed by clustering, the more CPU cycles are available for the application.
- The *ServerNet II* PCI Adapter delivers advantages over traditional networking providing more efficient (lower) CPU utilization, improved (reduced) latency and greater bandwidth.
  - Delivered, small message latency approaches 10 microseconds.
  - Data can be striped across both ports in the dual ported *ServerNet II* PCI Adapter for simultaneous sustainable transfers of 180MB/s of applications data in 66/64 PCI configurations.

The VIPL (VI Programming Library) interface from ServerNet requires less than 2 percent of the cumulative CPU for sending and receiving at the maximum bandwidth. While traditional networking consumes CPU resources in a manner that is fairly proportional with an increase in message size, ServerNet CPU requirements remain constant as message size increases.

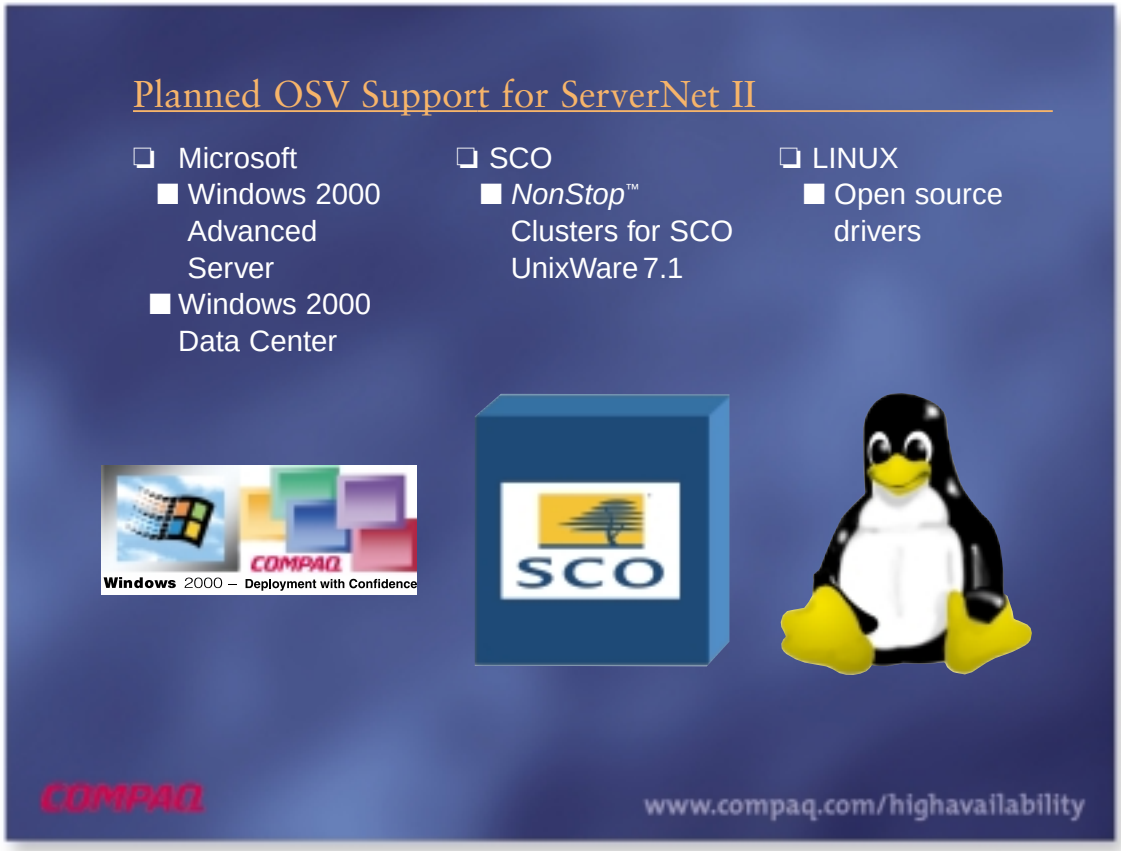

The *ServerNet II* products will be used in the following solutions:

- Windows 2000 Advanced Server:
  - –Microsoft SQL Server 2000 (third quarter, 2000)
  - –SAP (Replicated Enqueue) (third quarter, 2000)
  - –Oracle OPS (fourth quarter, 2000)
    Oracle Net8 (first half, 2000)

Future:

- –Other Microsoft applications using VI
- –Customer and 3rd party applications using the WSD API or the VIPL (VI Programming Library) interface

- Windows 2000 Datacenter:
  Microsoft Windows Sockets Direct (WSD API) (third quarter, 2000)

Future: other Microsoft applications using WSD API

- SCO UNIX: *NonStop*™ Clusters
  (fourth quarter, 2000)

- LINUX: open source clusters: Beowulf Clusters and customer apps (third quarter, 2000)

Specific *ServerNet II* competitors:
(**+** indicates strength; **–** indicates weakness)

Gigabit Ethernet

- **+** Massive industry momentum, well-understood by customers
- **–** Gigabit Ethernet cost per port is more than *ServerNet II* due to the high cost of Gigabit switches
- **–** VI is not available for Gigabit Ethernet
- **–** Not as robust as ServerNet for high-intensity traffic due to ServerNet fault-tolerant features and reliability guarantees

GigaNet

- **+** Backed by Intel and first-to-market with VI hardware
- **–** GigaNet cost per port is more than ServerNetII.
- **–** Not as robust as ServerNet for high-intensity traffic due to ServerNet fault-tolerant features and reliability guarantees
- **–** Start-up with thin resources

Compaq was the first to introduce the concept of a System Area Network for industry standard servers base on ServerNet.

With VI, industry specifications were developed for IPC (Inter Processor Communications) by Compaq, Intel and Microsoft.

With InfiniBand, we now have truly open standards that will allow any vendor to innovate within the InfiniBand standard, much as has been done within standards such as PCI and Fibre Channel

Strategically, Compaq has built a core competency around I/O technologies, including bus technology, storage I/O and system interconnects.
As we move toward InfiniBand as the next generation of I/O technology, Compaq will leverage our expertise around VI and the ServerNet core to build the next generation of products.

Compaq ServerNet II SAN Interconnect

for Scalable Computing Clusters

White Paper

June 2000

Compaq Computer Corporation
Prepared by
ISSD Technology Communications

EXECUTIVE SUMMARY

This paper introduces Compaq's *ServerNet II* system area network (SAN) interconnect for use in scalable business computing clusters. Clusters generate intense server-to-server messaging that can overwhelm traditional network technologies. The software protocol overhead generated by traditional network technologies impedes server-to-server communication and limits the scalability, availability and flexibility of clusters. In contrast, the *ServerNet II* SAN interconnect is a high-bandwidth, low-latency network technology that is proven to handle intense messaging in server clusters. A *ServerNet II* SAN allows applications to be distributed over multiple servers, which increases cluster scalability through greater parallelism and higher availability. The reader should be familiar with the **Virtual Interface Architecture** and **Compaq ServerNet technology**.

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

Compaq ServerNet II SAN Interconnect

for Scalable Computing Clusters

White Paper

INTRODUCTION

The enterprise computing systems of many businesses were built around centralized proprietary computers, which are expensive to manage and upgrade. The networks that evolved around these centralized systems use traditional network technologies, such as Ethernet and TCP/IP, which are designed for heterogeneous computing environments. Today, businesses can satisfy their computing goals with a collection of low-cost, industry-standard servers distributed within a cluster. A cluster is a group of two or more interconnected servers that act as a single computing unit. Servers in a cluster can be physically distributed across various locations, yet to a user the cluster appears as a single, unified computing resource. In a parallel application cluster, for example, each server simultaneously runs a copy of an application and the operating system. Client requests are divided among the servers' CPUs (central processing units), and the servers exchange information about the portion of the client's request they are processing. If one server fails, the jobs it was processing are distributed across the remaining servers. Consequently, clustered servers collectively generate tens of thousands of messages per second to keep track of the jobs they are executing.

The high volume of server-to-server messaging in a cluster requires a very efficient software communication interface and a highly reliable, high-speed, low-latency hardware interconnect. Without an efficient server-to-server interconnect, the performance and scalability of a cluster is severely limited. Traditional network technologies produce excessive software overhead, due to their heterogeneous environments, and are too inefficient for intense server-to-server communication. Therefore, in December 1997 Compaq and other industry leaders developed the Virtual Interface (VI) Architecture specification for a distributed messaging interface that allows more efficient server-to-server communication through a system area network (SAN). The VI Architecture provides a common software and hardware interface standard so that customers can choose the optimum SAN interconnect (Compaq *ServerNet II*, gigabit Ethernet, cLAN, ATM or others).

The Compaq *ServerNet II* SAN interconnect is the most complete industry-standard implementation of the VI Architecture specification. Measurement of CPU utilization under certain conditions shows that *ServerNet II* provides three times the performance of gigabit Ethernet and TCP/IP. As a result, *ServerNet II* is quickly emerging as the interconnect of choice for implementing clusters of industry-standard servers. Operating system vendor and independent software vendor support for *ServerNet II* includes Windows 2000 Data Center, Windows 2000 Advanced Data Server, Linux and SCO UnixWare 7.1.

TRADITIONAL NETWORKS VERSUS SYSTEM AREA NETWORKS

Clustered servers can communicate with each other by using a traditional network or a SAN. Traditional network technologies allow several different types of devices to communicate with each other in complex, heterogeneous networks. Network applications manage communications between the devices using multipoint protocols like TCP/IP and network interface controllers. These protocols generate a large amount of software overhead (error checking and control information) to ensure that messages are sent and received reliably. This software overhead continually interrupts the servers' CPUs. Consequently, the high volume of server-to-server messaging in a cluster can overwhelm traditional network protocols like TCP/IP. In fact, clusters that use traditional networks for messaging can lose as much as 20 percent to 30 percent of their capability during intense messaging.

This loss of system capability results in slower response times for users and reduces the scalability, availability and flexibility of clusters. When you consider the unpredictable nature of Internet traffic, it becomes apparent that traditional network technologies cannot handle server-to-server messaging in e-commerce environments.

In contrast, a SAN is a high-performance network that can handle intense messaging efficiently. The efficiency of SANs allows data to be distributed anywhere within the cluster, which increases scalability through greater parallelism and higher availability. Initially, only proprietary SAN technologies could handle the large amount of messaging traffic in a cluster. Therefore, in December 1997, Compaq, Intel, Microsoft and other industry leaders developed the Virtual Interface (VI) Architecture specification to promote an industry-standard architecture for server-to-server communication.[1] Unlike traditional networking technologies, the VI Architecture allows distributed applications to move data between clustered servers without invoking operating system functions. As a result, operating systems perform significantly better by avoiding unnecessary transitions to and from user applications and by avoiding excessive CPU interrupts.

Relational database, operating system and enterprise resource planning (ERP) products were among the first to leverage affordable VI Architecture technology to add robustness and predictability to industry-standard clusters. ERP applications are configured using multiple tiers to achieve sufficient scale (see Figure 1). The top tier consists of the ERP web or graphic user interface, the middle tier contains application servers, and the bottom tier contains relational database servers. ERP applications place great communications demands on the database servers, resulting in intense messaging. Message intensity within the server cluster is one to two orders of magnitude greater than WAN/LAN (or wireless) messaging destined outside the cluster. Request/reply transactions consisting of SQL requests and SQL result-sets (messages to and from the database) account for most of the messaging traffic. Recent large SAP sales and distribution benchmarks require the database to manage message profiles exceeding 32,000 messages per second with an average message size just over 1,600 bytes.

[1] *For more information, see technology brief Virtual Interface Architecture for System Area Networks, document number 0184-0699-A.*

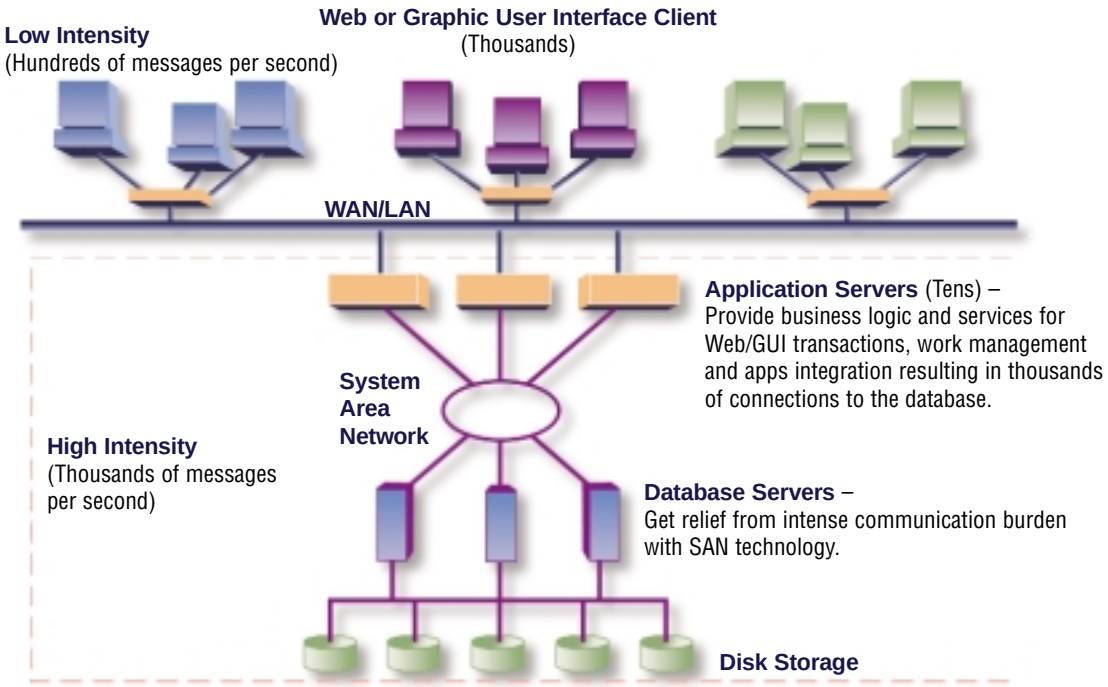

Figure 1. Multi-tier architecture.

ServerNet II SAN Technology Improves the Response of Clusters

ERP and database applications usually operate efficiently if powerful CPUs and large memory resources are available to them. Distributed ERP applications accommodate more users (scale) by adding applications servers (Figure 1), which increases the need for SQL services from the database. Overall system performance depends upon the ability to keep the database busy. However, extreme software overhead generated by traditional network technologies forces the database to let the operating system functions and heavy communication protocols use a substantial portion of the CPU resources. With *ServerNet II*, database CPU utilization increases more than three times that of gigabit Ethernet and TCP/IP with the messaging loads just described. Published benchmark results demonstrate the value of *ServerNet II* SAN technology by increasing the number of SAP SD users supported by the system by more than 30 percent.

Parallel relational database architectures rely on clustering multiple nodes to support very large databases and increase database availability. Queries, updates and utility functions are distributed across multiple servers within a database cluster to maximize parallelism and to reduce the impact of a failure. Parallel relational databases rely on message systems to distribute processes, data and database control information to database instances within the cluster. Parallel databases demand more than CPU efficiencies from the messaging system. Low message latency allows faster synchronization of parallel operations, which results in lower response times to user queries. CPU utilization and latency measurements with *ServerNet II* messaging show a three-fold improvement over gigabit Ethernet and TCP/IP with the messaging loads just described. Published database benchmark results demonstrate the value of SAN technology and efficient messaging to the database cluster system.

COMPAQ *SERVERNET II* TECHNOLOGY

Compaq enterprise servers have long used ServerNet SAN technology to provide systems with superior scalability, availability and flexibility. ServerNet is the internal cluster interconnect within several Compaq *NonStop™* products. Compaq Himalaya systems with ServerNet SAN messaging provide *NonStop™* transaction services to automatic teller machines, stock exchanges and credit card processing for most financial communities and consumers worldwide. Compaq Integrity products use ServerNet to achieve superior UNIX availability to the telecommunications industry. *NonStop™* Cluster products based upon Compaq *ProLiant* servers and the SCO operating system rely upon ServerNet for superior scalability and availability features. *ProLiant* Parallel Database clusters use ServerNet for scalability.

The Compaq *ServerNet II* SAN interconnect is the most complete industry-standard implementation of the VI Architecture specification. Compaq *ServerNet II* is a low-latency, distributed memory fabric that uses low-cost, industry-standard components to provide highly reliable server-to-server communication. *ServerNet II* technology leverages several features of the VI Architecture that allow distributed applications to operate independent of the operating system and avoid unnecessary CPU interrupts. CPU utilization measurements of small message (64 bytes) latency show that *ServerNet II* messaging delivers over three times the performance of gigabit Ethernet and TCP/IP. Because of this high performance, *ServerNet II* SAN configurations can scale to tens of servers within a cluster.

ServerNet II Hardware Components

As shown in Figure 2, the basic hardware components for a cluster of two servers include *ServerNet II* PCI Adapters and ServerNet cables. Each *ServerNet II* PCI Adapter has two ports (X and Y). The respective X and Y ports on the adapters can be linked to create redundant connections. If either the X or Y connection should fail, data is seamlessly transferred by the other connection, thus improving reliability. *ServerNet II* also provides high reliability through self-checking built into the application-specific integrated circuits (ASICs) and hardware routing protocol. The hardware routing protocol ensures end-to-end, reliable data transfer.

ServerNet II uses industry-standard 64 bit, 66MHz PCI components that contain several features of the VI Architecture to move data efficiently. The link media for *ServerNet II* is the same physical media used by gigabit Ethernet and Fibre channel (Serial 1000Base-CX up to 25 meters). Links are bi-directional with speeds of 1.25 + 1.25 gigabits per second (Gb/s).

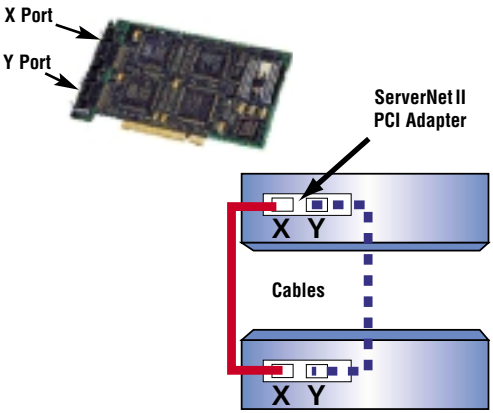

Figure 2. Redundant, direct *ServerNet II* interconnect for a cluster of two servers.

A *ServerNet II* switch is required for clusters of more than two servers. The ServerNet switch is a 2U (3.5″) high, 19″-wide, rack-mounted device with connections for 12 ServerNet cables. Figure 3 shows a SAN consisting of six servers connected to a *ServerNet II* switch. A second switch and cables can be added for increased fault tolerance. As more nodes and data paths are added to the cluster, the aggregate bandwidth of the *ServerNet II* interconnect increases. Multiple switches can be cascaded to support as much bandwidth as needed.

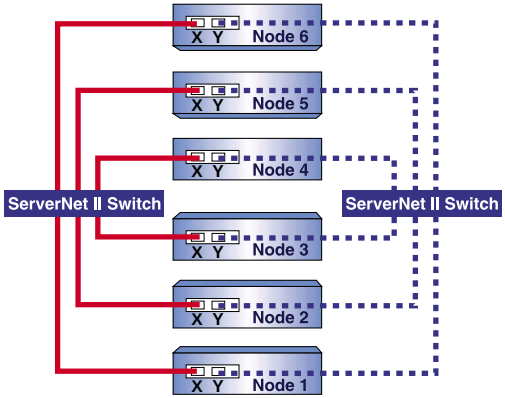

Figure 3. Six-node cluster with redundant *ServerNet II* switches.

Low Latency Techniques Improve Performance

ServerNet II uses two techniques to ensure low-latency data transmission: wormhole routing and the push/pull approach. In wormhole routing, *ServerNet II* retrieves the data and divides it into 512-byte packets. The destination address is added to the front end of each packet, which allows the switch to route all packets to the destination node. If necessary, packets from the same data file can travel different paths to reach the destination. 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.

The push/pull (write/read) approach of *ServerNet II* allows the burden of data movement to be absorbed by either the source or target server. At the beginning of a push (write) transaction, the source notifies the destination to allocate enough buffers to receive a large message. Before sending the data, the source waits for acknowledgment from the destination that the buffers are available. To pull (read) data, the destination allocates buffers before it requests data. Then, it transfers the data through the *ServerNet II* PCI Adapters without interrupting the OS or application.

ServerNet II Performance Summary

The efficiency of messaging is defined in terms of CPU utilization, latency and bandwidth. *ServerNet II* communication protocols are implemented in native hardware to reduce CPU utilization to a small fraction of that experienced with traditional protocols such as TCP/IP. *ServerNet II* protocols also reduce the typical operating system services necessary to support traditional protocols. For small messages (64 Bytes or less), latency and CPU measurements show that *ServerNet II* provides a three-fold improvement over Gigabit Ethernet and TCP/IP. The latency of small messages approaches 10 microseconds. For large messages (16KBytes or greater), measurements show that each *ServerNet II* link delivers a bi-directional bandwidth of 180 MB/s while consuming less than 2 percent of CPU resources. By comparison, gigabit Ethernet bandwidth can approach 100 MB/s, but only by consuming intolerable amounts of CPU resources in both the source and the destination nodes. The table below summarizes results of various *ServerNet II* hardware tests performed using a Compaq *ProLiant* 8500 server configured with eight 500-MHz processors. The server was processing 512-byte bursts over a 64-bit, 66-MHz PCI bus.

ServerNet II Performance Summary*

Performance Test:	Measured
8 byte delivery latency with Send/Receive poll	12 µs
8 byte delivery latency with Send/Receive wait	32 µs
64 byte CPU cost with Send/ Receive wait	27 µs
64 byte CPU cost with Lazy Send/ Receive wait	29 µs
64K 1-way throughput RDMA writes (1VI-4 VIs)	92-132 MB/s
64K 1-way throughput RDMA reads (1 VI-4 VIs)	129-134 MB/s
64K 2-way throughput RDMA (reads-writes)	181-194 MB/s
64K RDMA throughput test CPU utilization	~0%

* *This number was measured on 33MHz, 32-bit PCI with a 500MHz CPU.*

ServerNet II Leverages VI Architecture Features to Improve Scalability

High levels of scalability require the ability to distribute data anywhere within the cluster without significantly penalizing overall cluster performance. This data location independence improves scalability through greater parallelism and higher availability. *ServerNet II* leverages some features of the VI Architecture to greatly increase the scalability of parallel applications over that achieved using traditional network technologies. *ServerNet II* SANs allow CPUs to spend more time doing useful work by not making them wait on messaging. This allows *ServerNet II* configurations to scale to tens of servers within a cluster and more.

ServerNet II PCI adapters conform to the VI Architecture specification and implement:

- Reliable reception
- Remote DMA read and write
- Thread-safe VI provider
- 64,536 VI queues
- Hardware VI completion queues
- Hardware flow control
- Full VI ptag support

The VI Programming Library (VIPL) interface is a user-mode interface for programming Application Programming Interfaces (APIs) that are typically exported by operating systems.

Fault Tolerance at All Levels

Compaq *ServerNet II* technology has fault-tolerance built into multiple architectural levels. Each X and Y port in a *ServerNet II* adapter connects to a unique fabric in a ServerNet SAN. Messages are assigned a port and routed to the destination by one of the ServerNet fabrics. If a failure occurs in the port, its associated link, or fabric, the affected messages are automatically and transparently restarted through the second port. *ServerNet II* hardware guarantees reliable and ordered message delivery. Fault-tolerant features within a ServerNet Switch include duplicate state machines, path disables, link keep-alive protocol, protocol checkers link-level flow control, transaction timeout counters and inclusion of both source and destination addresses within each packet.

Operating System Vendor and Independent Software Vendor Support

OSV/ISV solutions for *ServerNet II* products will include:

Windows 2000 Advanced Server and Datacenter Server:

- Microsoft SQL Server, Microsoft Windows Sockets Direct (WSD) API, COM+ and other Microsoft applications using WSD or the VIPL.
- Oracle OPS and Oracle Net8
- SAP
- Customer and third party applications using the WSD API or the Virtual Interface Primitive Library interface
- ERP applications using WSD or the VIPL
  SCO Unix NonStop Clusters
- Linux Beowolf Clusters and third-party applications

THE FUTURE WITH INFINIBAND

Although *ServerNet II* SAN architectures provide excellent performance for today's parallel application clusters, distributed computing systems still require multiple interconnects because interprocessor communication is different from I/O communication. Distributed computing systems consist of multiple function-specific fabrics for interprocessor communication, IP networks and storage area networks. A single I/O architecture is needed both to simplify existing interconnect methodologies and to provide a flexible framework for proven computing and storage technologies to evolve. Therefore, in 1999 Compaq and other industry leaders[2] formed the InfiniBand℠ Trade Association, **www.infinibandta.org**. to develop a new industry-standard specification for a channel-based, switched fabric I/O interconnect. The goal of the InfiniBand Trade Association is to combine the best traits of existing interconnect technologies into a single specification called the InfiniBand Architecture.

The InfiniBand Architecture will reduce the complexity of distributed computing systems by using a single communication fabric for both I/O and interprocessor communication. During 2000, Compaq *ServerNet II* will provide customers with a bridge to InfiniBand by implementing InfiniBand verbs over *ServerNet II*. This will allow ISVs and early adopters to begin porting their applications to InfiniBand before InfiniBand hardware is available.

SUMMARY

Compaq's *ServerNet II* SAN interconnect addresses the need for efficient messaging in scalable computing clusters. *ServerNet II* leverages several industry-standard features of the VI Architecture to handle intense messaging efficiently and reliably with minimal impact to other critical system resources. These features add robustness to a cluster and allow it to scale to tens of servers.

For the long term, Compaq is investing in the InfiniBand Architecture as a common I/O switched-fabric interconnect for the industry. Compaq intends to emulate InfiniBand verbs using *ServerNet II* during 2000 for early InfiniBand development partners.

[2] *Steering committee companies include Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft and Sun Microsystems. Sponsoring companies include Adaptec Inc., Cisco Systems Inc., Fujitsu-Siemens, Hitachi, Lucent Technologies, NEC, Nortel Networks Corp and 3Com. There are also over 120 member companies.*

NOTICE

The information in this publication is subject to change without notice and is provided "as is," without warranty of any kind. The entire risk arising out of the use of this information remains with the recipient. In no event shall Compaq be liable for any direct, consequential, incidental, special, punitive or other damages whatsoever (including without limitation, damages for loss of business profits, business interruption or loss of business information), even if Compaq has been advised of the possibility of such damages.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

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 *ServerNet II* SAN Interconnect
for Scalable Computing Clusters
First Edition (June 2000)
Document Number TC000602WP

MODELS

Compaq *ServerNet II* Switch
12 ports, Copper Serial
174739-001 (US)
174739-B31 (Int'l)

***ServerNet II* Cables**
***ServerNet II* 5-meter Copper Serial Cable Kit**
174737-B21
***ServerNet II* 10-meter Copper Serial Cable Kit**
174738-B22

***ServerNet II* is a high performance, low latency network used to interconnect servers with other servers in a System Area Network.**

The *ServerNet II* Switch is a 12-port, non-blocking switch used to connect from two to twelve Compaq servers in a cluster.

The *ServerNet II* Switch is available in a 2U Rack Mountable package.

Fault tolerant networks are standard options at installation. Specialized configurations of multiple switches are available.

Usage

The *ServerNet II* Switch (ServerNet Switch) is used to connect Compaq servers together in a ServerNet System Area Network to achieve new levels of server capability and flexibility. The ServerNet Switch efficiently and reliably routes packets for the purpose of distributing work among a collection of servers.

The ServerNet System Area Network is a point-to-point, packet-switched, wormhole-routed, robust fabric for connecting servers. Compaq *ServerNet II* PCI Adapter and Drivers within each server provide highly efficient messaging facilities to Operating Systems, Middleware and Applications for each server within a System Area Network.

Configuration Flexibility

The ServerNet Switch requires 2U in a 19″ Rack. Mounting brackets are included. The ServerNet Switch supports up to 12 connections (see Connectors and Interface Cables). Specialized ServerNet Switch configurations exceeding 12 connections can be created from multiple switches cascaded into symmetric or asymmetric topologies.

Fault Tolerant

Fault tolerance is designed into individual switches, including features such as duplicate-state machines, path disables, link keep-alive protocol, protocol checkers, link-level flow control, transaction timeout counters, inclusion of both source and destination addresses within each packet.

A fault tolerant network is achieved by configuring redundant switch fabrics to each of the two ports of a *ServerNet II* PCI Adapter.

Switching Bandwidth and Efficiency

Each of the 12 connections support point-to-point data links of 1.25 + 1.25 gigabits per second of bi-directional bandwidth giving an aggregate switch bandwidth of 15 gigabits per second. Effective bandwidth can be increased between ServerNet Switches to 5.0 + 5.0 gigabits per second in specialized configurations by merging sets of parallel links.

The ServerNet Switch path formation latency for up to 12 incoming packets switched (non-blocking) to 12 outgoing routes is approximately 300 nanoseconds. Routing tables support arbitrary network topologies.

Manageability

ServerNet Switch control is contained within ServerNet links. No additional networking hardware is involved to initialize and control *ServerNet II*. Switch-related host software drivers support installation and industry-standard management tools such as SNMP for enterprise systems management access.

Compatibility and Integration Testing

The ServerNet Switch has been tested on a variety of Compaq *ProLiant* and Alpha-based servers and workstations.

System Area Network Architecture

Servers that have been tested with the *ServerNet II* Switch are the PL1600, PL1850R, PL3000, PL5500, PL6400R, PL6500, PL8000, PL8500, DL380, DL580, ML370, ML530 and the ML570.

Warranty

Same as the connected servers.

SPECIFICATIONS

Signaling and Interfaces	
ServerNet Links	125 MHz, packet payload up to 512 bytes
Connectors	12 eight-pin Berg MetaGig connectors
Interface Cables	Cables are IEEE industry standard – Serial 1000Base-CX (up to 25 meters)
Safety and Regulatory Compliance	FCC Class A
CISPR 22/EN 55022	
C.R.C., c.1374 (Canadian Radio Interference)	
CE Mark	
New Zealand Radio Frequency Service, RFS49	
EN 50082-1, IEC 801-2,3,4	
Council Directive 89/336/EEC	
IEC 950/EN 60 950	
UL 1950	
CSA 22.2-950; VDE 0805/IEC950	

Physical	
Power and Cooling	
Power supply	+110/220V AC (+ or-5%)
50 Hz / 60 Hz	
Cooling	12V DC fan
Switch Chassis Physical	
Dimensions (HxWxD)	3.46 x 16.85 x 23.25 in/8.8 x 42.8 x 59.0 cm
Power	
Low Range – USA, Japan	100 to 120 VAC 50 Hz/60 Hz
High Range – Europe, US Industrial	200 to 240 VAC 50 Hz/60 Hz
Environmental Requirements	
Operating temperature	41° to 131°/5° to 55°C – TDM Class A
Relative humidity	10% to 85%
Maximum thermal drift (%/hr)	51.8°F/11°C – TDM Class A
Thermal shock (Delta °)	68°F/20°C – TDM Class A

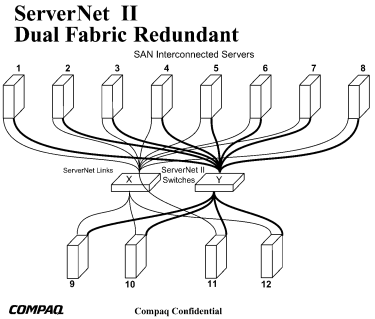

MODELS

Compaq *ServerNet II* PCI Adapter
Dual ported-copper serial interfaces
174736-B21

***ServerNet II* cables**
***ServerNet II* 5-meter copper serial cable kit**
174737-B21
***ServerNet II* 10-meter copper serial cable kit**
174738-B22

VI industry-specified compliant hardware and driver support for efficient server-server messaging. A PCI adapter supporting the Compaq *ServerNet II* interface and Compaq's *ServerNet II* system area network.

Usage

The *Servernet II* PCI Adapter is used for CPU-efficient, high speed, low latency messaging provided by a *Servernet* systems area network. A *Servernet II* PCI Adapter ("*Servernet* Adapter") allows efficient server-to-server communications between operating systems, middleware and applications distributed over two or more servers connected in the systems area network.

The *Servernet* system area network is a point-to-point, packet-switched, wormhole-routed, robust fabric for connecting a collection of industry standard servers to work together to achieve new levels of performance and scalability. The *Servernet* Adapter has two active ports for fault tolerance and additional performance.

Configuration Flexibility

The *Servernet* Adapter is used in a point-to-point configuration that can be expanded to very large networks utilizing the Compaq *Servernet II* Switch.

Fault Tolerant

Each port in a *Servernet* Adapter connects to one of two possible unique fabrics in a *Servernet* system area network. Messages are assigned a port and routed by a *Servernet* fabric to the destination. If a failure occurs in the port or its associated link or that fabric, then the affected messages are automatically and transparently restarted through the second port. Hardware transaction timers and guarantees ensure once, and only once, delivery by the reliability characteristics of *Servernet*.

Performance

The *Servernet* Adapter delivers advantages over traditional networking providing more efficient (lower) CPU utilization, improved (lower) latency and greater bandwidth.

- Delivered, small message latency approaches 10 microseconds.
- Data can be striped across both ports in the dual-ported *Servernet* Adapter for simultaneous sustainable transfers of 180MB/s of applications data. Bandwidth is limited by 66/64 PCI configurations.
- The VIPL interface from *Servernet* requires less than 2 percent of the cumulative CPU for sending and receiving at the maximum bandwidth. While traditional networking consumes CPU resources in a manner that is fairly proportional with an increase in message size, *Servernet* CPU requirements remain constant as message size increases.

Virtual Interface (VI) Architecture Features

The *Servernet* Adapter conforms to the Virtual Interface (VI) Architecture specification **www.viarch.org** and implements:

- Reliable reception
- Remote DMA read and write
- Thread-Safe VI provider
- 64,536 VI queues
- Hardware flow control
- Full ptag support

The VIPL interface is a user-mode interface for programming APIs typically exported from operating systems.

Compatibility and Integration Testing

The *Servernet II* PCI Adapter has been tested on a variety of *Compaq ProLiant* and Alpha-based servers and workstations.

System Area Network Architecture

Servers that have been tested with the *Servernet II* PCI adapter are the PL1600, PL1850R, PL3000, PL5500, PL6400R, PL6500, PL8000,

PL8500, DL380, DL580, ML370, ML530 and the ML570.

OS Driver Support
- Microsoft Windows 2000
- Linux

Warranty
Same as the server

SPECIFICATIONS

Signaling and Interfaces	
PCI	One PCI Rev 2.1 compliant bus-master slot (66MHz/64 bits, 33MHz/64 bits or 33MHz/32 bits)
Servernet Links	125MHz, 512 byte packets
Connectors	Two eight-pin Berg MetaGig connectors
Interface Cables	Cables are IEEE industry-standard Serial 1000Base-CX (up to 25 meters)
Power	
From host	+5.0 V DC 5 percent
On card	+5.0 V, +3.3 V and +2.5V
Physical (HxLxD)	5.0 x 7.5 x 0.87 in/12.7 x 19.1 x 2.2 cm
Safety and Regulatory Compliance	FCC Class A CISPR 22/EN 55022 C.R.C., c.1374 (Canadian Radio Interference) CE Mark New Zealand Radio Frequency Service, RFS49 EN 50082-1, IEC 801-2,3,4 Council Directive 89/336/EEC IEC 950/EN 60 950 UL 1950 CSA 22.2-950 VDE 0805/IEC950

SPECIFICATIONS

Environmental Requirements	
Operating	
Temperature	5º to 55ºC / 41º to 131ºF
Humidity	10 percent to 85 percent (non-condensing)
Altitude	10,000 ft/3,048 m
Non-Operating	
Temperature	–40º to 150ºF/ –40º to 66ºC
Humidity	10 percent to 95 percent (non-condensing)
Altitude	13,1233 ft/40,000 m
Shipping Temperature	-40º to 66º C / -40º to 150º F

ServerNet II PCI ADAPTER

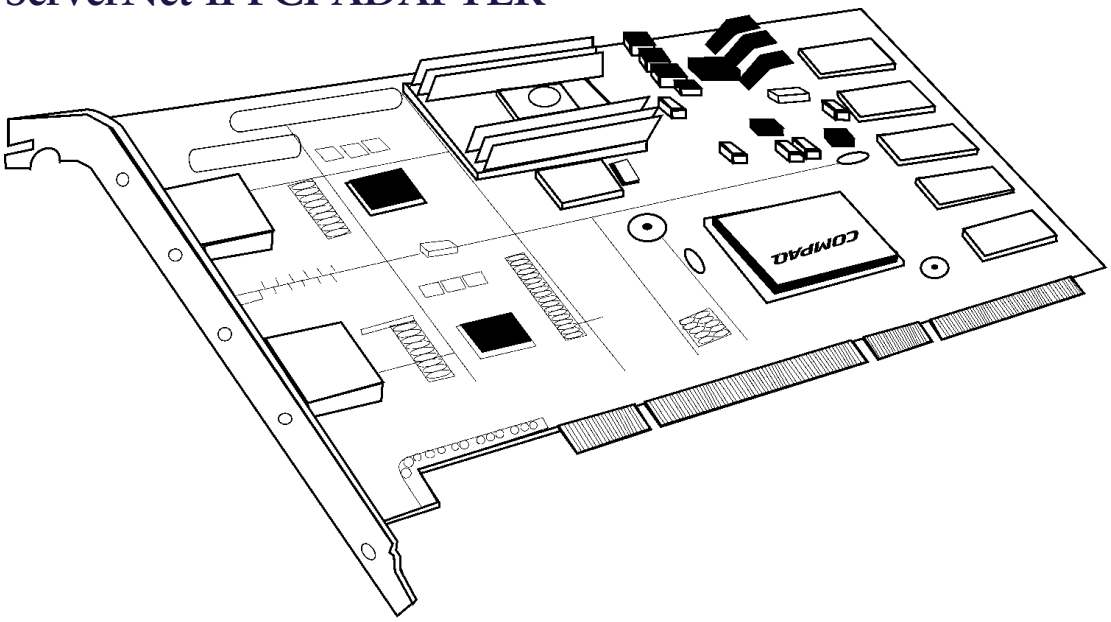

Managing for Availability

Compaq Intelligent Cluster Administrator v2.0

Product Line/Speaker Notes

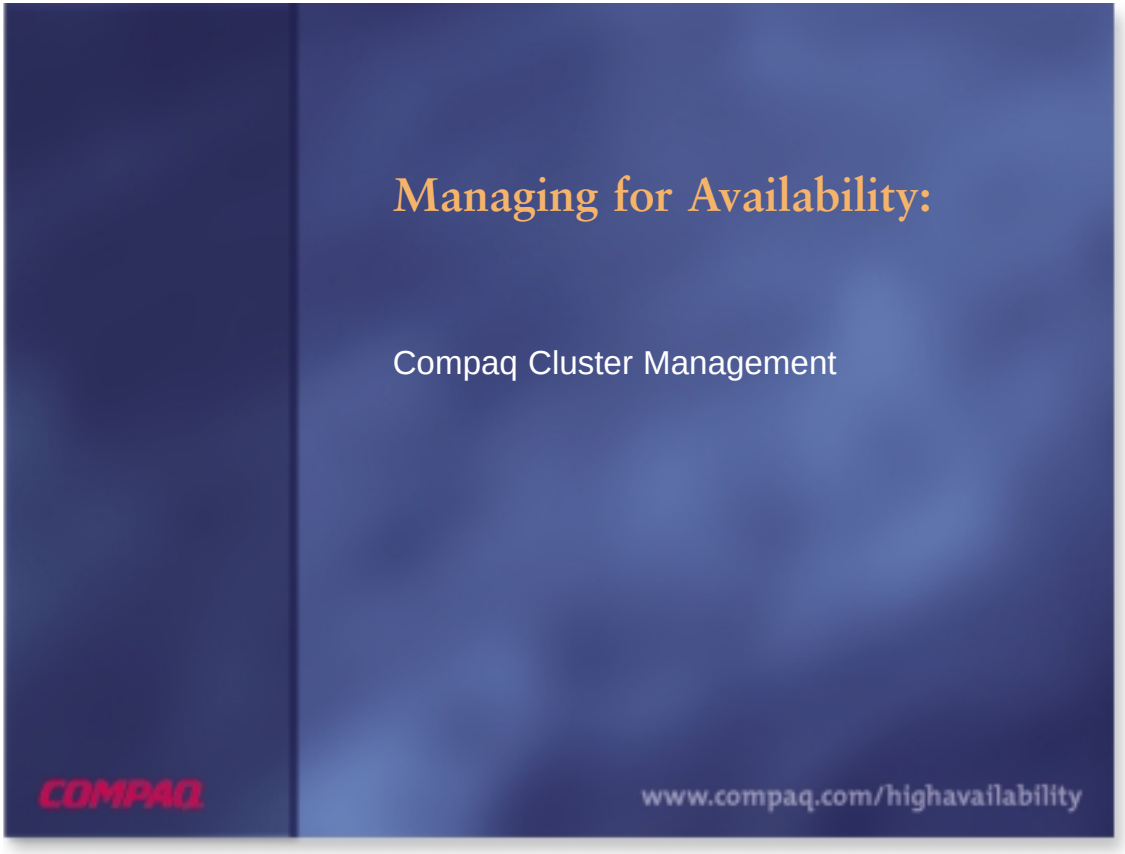

Welcome to the world of Compaq *Intelligent Cluster Administrator*. Compaq is committed to a comprehensive cluster management strategy across all its cluster platforms. Compaq Intelligent Cluster Administrator provides a consistent, intuitive and easy-to-use set of interfaces to monitor, manage and extend your application-critical systems.

Managing for Availability

Compaq Intelligent Cluster Administrator v2.0

Product Line/Speaker Notes

Cluster administration is complex, frustrating and time consuming. Through experience and innovation, Compaq is developing sophisticated cluster management and administration tools that are intuitive, easy to perform and extensible across all facets of the enterprise, thereby reducing the costs of ownership of industry-standard clusters.

Key Messages

As the provider of the best managed servers, systems and network devices in the industry, Compaq has built knowledge and expertise into its cluster administration tools to accelerate deployment, simplify operation and minimize resource investment in industry-standard clusters. Compaq has integrated the monitoring of clusters with the monitoring of servers, systems and network devices. Compaq is enhancing Compaq Insight Manager – XE to provide a browser-based cluster monitor subsystem with advanced cluster status notification and anytime, anywhere access to NT MSCS clusters.

Compaq is also delivering advanced and extensible browser-based tools for the administration and management of clusters that provide a single point of control for consistent administration and rapid deployment of clusters throughout the enterprise.

The Compaq cluster management offerings will allow the business manager and IT manager to focus on managing critical business issues instead of managing systems and clusters.

Create preference for Compaq products by providing the best customer experience in deploying and operating Compaq systems through highly-valued management products, platform features, services and partnerships.

Compaq will continue to deliver the most manageable systems by designing the best tools to deploy, operate and maintain those systems into all Compaq products.

Continue to drive and develop to emerging industry standards.

And, continue to provide a portfolio of world-class professional services.

Standards-based management of SNMP, DMI and HTTP devices

Automatic discovery of Compaq and third-party devices.

Also discovers devices with HTTP Web servers (Web-based agents) and provides a hot link to the Web server on the device.

Browser-based access with higher security than SNMP.

Uses a SQL server database as a repository. The schemas are published for use with third-party reporting programs.

Comes out of the box with pre-configured queries.

Integrated Cluster Monitor

Key features:

- Creates Compaq *Insight Manager XE* notifications to escalate or document cluster problems
- Display utilization or status data for specific cluster attributes
- Browse cluster and component status in a tree hierarchy
- Monitor cluster status
- Investigate the sources of specific alerts

COMPAQ

www.compaq.com/highavailability

Insight Manager XE Cluster Monitor provides the system administrator with a one-stop management control point for clusters. This common repository and data collection method will make the administrator's task simpler and faster in that they need to know one service to manage all the devices in their enterprise, thereby reducing administrative costs…and it is a basic component of XE.

Browser-based for anytime, anywhere access allows the administrator complete flexibility without having to dedicate specific processor or console resources to administration of servers, desktops or clusters.

Making cluster management less stressful by providing a simple, easy-to-view interface, with extensive problem definition and remedy information displayed on event or device selection. This, combined with intuitive color-coded identification of cluster "hot spots" for rapid problem determination and resolution, will make the administrator's job more efficient and therefore less time-consuming.

Decreased cluster downtime through customized fault management and performance threshold controls which can be set to meet specific cluster processing needs will minimize "hung" systems.

Understand enterprise systems control points by providing a history log and configuration control points through the Insight Manager – XE repository and the separation of administrative responsibilities through our level of change authority mechanism, we can protect your business information assets and processing capability from unauthorized access.

Key Features

- Web-based for ease of utilization with anytime, anywhere access
- Customizable to meet administrator workload or job responsibility
- Selectable monitor points for cluster monitor customization
- Application – all exchange clusters in the enterprise
- Security levels for administration control and access
- Problem window and topology tree-event selection for micro- or macromanagement
- Optional linkage to *Intelligent Cluster Administrator* for Single Control Point
- Event problem description and potential remedy displayed for efficiency
- Drill down to problem device for further analysis and repair

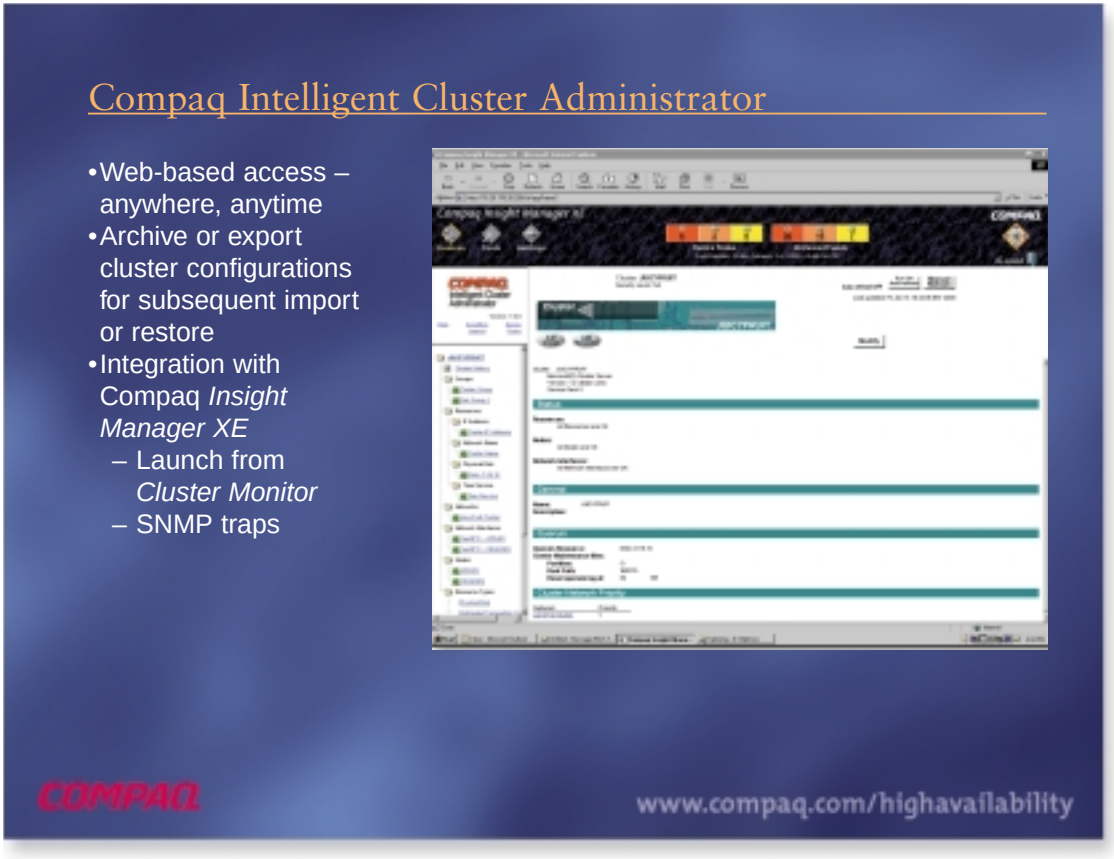

Compaq *Intelligent Cluster Administrator* allows you to view information from a Web browser. Browsers supported include Microsoft Internet Explorer 4.01 with Service Pack 1 or Microsoft Internet Explorer 5.0 and Netscape Communicator 4.5 or higher.

Intelligent Cluster Administrator provides a default archive policy which automatically archives the active cluster configuration before an import operation.

Compaq provides a comprehensive management solution by integrating *Intelligent Cluster Administrator* with Compaq Insight Manager XE – Cluster Monitor. *Intelligent Cluster Administrator* can be launched from Compaq Insight Manager XE – Cluster Monitor. *Intelligent Cluster Administrator* can also be configured to send SNMP traps to Compaq Insight Manager XE.

All Compaq customers who previously registered earlier versions of Compaq *Intelligent Cluster Administrator* will receive a free upgrade to *Intelligent Cluster Administrator* version 2.0. This "free upgrade" offer may not apply to future releases of Compaq *Intelligent Cluster Administrator*.

Managing for Availability

Compaq Intelligent Cluster Administrator v2.0

Product Line/Speaker Notes

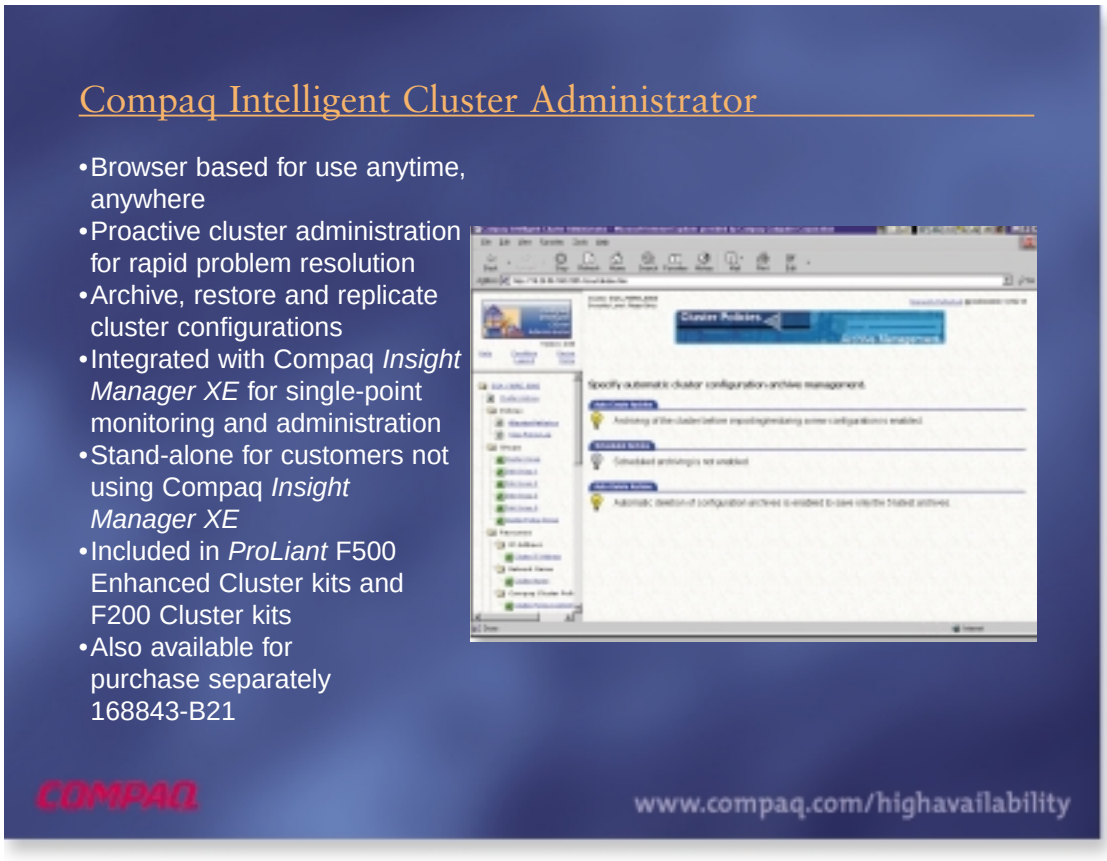

Compaq cluster administration and monitoring tools reduce system management costs, improve administrator efficiency in dealing with clusters and provide for early warning and potential rapid response workarounds, thereby minimizing system outages.

Compaq *Intelligent Cluster Administrator* is a Web-based anywhere, anytime utility for ease of administration of Microsoft MSCS clusters. It provides the ability to proactively monitor and administer an active cluster, manage cluster configuration archives and export or import cluster configurations. It can be used in a stand-alone fashion or integrated with the cluster monitor within Compaq Insight Manager XE thus enabling administrators to respond to an alert by launching Compaq *Intelligent Cluster Administrator* directly from the cluster monitor.

Specific administration capabilities include:

- Display status
- Modify properties of cluster objects
- Change the state of cluster objects (pause, resume, bring online, take offline)
- Add groups and resources and establish dependencies
- Assign failover policies for cluster resources
- Failover resources and nodes

The cluster configuration archive function includes the capability to maintain, compare and report on the current and/or archived configuration. You can also export and import cluster configurations to and from an external file to facilitate the replication of standard cluster configurations throughout an organization.

Compaq *Intelligent Cluster Administrator* is included in Compaq *ProLiant* Enhanced Cluster kits or it can be purchased separately (part number 168843-B21).

- Stop and start cluster services

Compaq *Intelligent Cluster Administrator* version 2.0 is well-suited for making modifications to an existing cluster configuration. Cluster policy management sets *Intelligent Cluster Administrator* apart from similar cluster management tools.

Accessing Compaq Intelligent Cluster Administrator

From Compaq *Insight Manager*, you are only three clicks away from comprehensive cluster management

- Click Cluster Monitor on the *Insight Manager XE* Overview page
- Click Cluster in Treeview
- Click Cluster Admin.

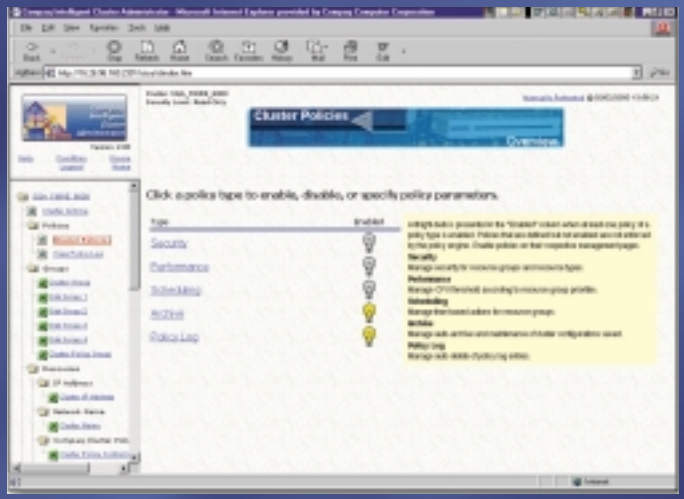

COMPAQ

www.compaq.com/highavailability

To start Compaq *Intelligent Cluster Administrator* from within Compaq Insight Manager XE requires only three mouse clicks (review instructions in slide).

The Compaq *Intelligent Cluster Administrator* can also be accessed by launching Microsoft Internet Explorer or Netscape Communicator and entering the address of the target machine in the URL field. The Compaq Web-based Management Device home page for the target machine will appear.

MODELS

Compaq *Intelligent Cluster Administrator v2.0*
168843-B21 International

As a stand-alone utility for Compaq HA/F100, HA/F500 and CL380 customers.

***Compaq Intelligent Cluster Administrator v2.0* is included in the *ProLiant* Cluster HA/F500 Enhanced Cluster Kit:**
379973-B23 (International)
It is also a component of the *ProLiant* Cluster HA/F200 Cluster Kits:
380357-B21
380357-291 (Japan)

- **A Web-based administration utility for Compaq *ProLiant* Clusters supporting Microsoft Cluster Service**
- **Provides policy management which allows the system administrator a way to define conditions so that the system will automatically take corrective action when criteria are met**
- **Enhanced Application Availability wizard that leads users through the process of making the cluster aware of applications not designed to be highly available in MSCS.**
- **Provides Web-based cluster administration tools for ease of cluster administration and control from anywhere at any time**
- **Supports Web access for local or remote administration**
- **Support for Windows 2000 Advance Server**

Overview

Compaq Intelligent Cluster Administrator is a Web-based anytime, anywhere utility for ease of cluster administration. This Compaq innovation supports Compaq *ProLiant* Clusters using Microsoft Cluster Service (MSCS) with Enterprise Edition Version 4.0 SP5 or higher or Windows 2000 Advanced Server. The Intelligent Cluster Administrator can be used stand-alone or integrated with Compaq Insight Manager XE – Cluster Monitor.

The Intelligent Cluster Administrator is an advanced cluster utility that provides Web-enabled cluster administration and Single Point of Control for NT MSCS clusters. The Intelligent Cluster Administrator is one of Compaq's initial steps for its enterprise cluster administration and management strategy, designed to promote industry-standard product concepts in what has been a proprietary cluster management environment.

Product Description

The first release provided ease of use and efficient cluster administration for NT MSCS clusters, enabling the administrator to have a Single Point of Control for all the clusters in the enterprise. This means that from an Internet Explorer or Netscape browser, the administrator could point to a cluster URL address and utilizing Compaq's Web-enabled management agents, monitor and manage a Compaq *ProLiant* Cluster. From the browser, the administrator can check for any cluster destabilizing conditions like CPU or storage utilization thresholds or application slowdowns and dynamically reallocate cluster resources to meet processing demands. The administrator can also perform work-arounds for cluster resources that may be failing to continue cluster operations without causing the cluster to failover. This can all be done from the browser, thereby eliminating physically going to the cluster to use the local console or the LAN console to support the cluster.

Compaq Intelligent Cluster Administrator 2.0 includes Policy Management – providing functions to simplify and extend the monitoring and management of Microsoft Cluster Service. Policies are set up by defining conditions so the Intelligent Cluster Administrator will automatically take corrective action when conditions are met. Additional functions included in Version 2.0 are:

- Pre-configured cluster images which are selectable by the administrator to ease in initialization of a MSCS cluster
- Enhanced Application Availability Wizard (EAA)
- Generic Script Resource Type allowing for command-line execution
- Security policies
- Performance policies
- Automatic cluster archive
- Policy log management

Intelligent Cluster Administrator – state-of-the-art ocluster configuration management facilities

- Analyze and track changes to a specific cluster configuration
- Time-based policies for scheduling system administration tasks to be performed at specific intervals
- Restore, update or replicate cluster configurations
- Add, change, delete configuration log entries
- Compare an existing cluster configuration to a log configuration for differences
- Configuration policies check on the advisability of configuration changes
- Verify any configuration log cluster with any current cluster configuration
- Designate a cluster log configuration as the "ideal" configuration for optimal cluster state
- Export any current or cluster log configuration
- Import any current or logged configuration to the same or another physical cluster configuration

- Select from a series of preconfigured cluster configurations for ease of initialization of the MSCS cluster

The Intelligent Cluster Administrator frees the administrator from tedious control processes and time-consuming administration by allowing the administrator the ability via this flexible Web-based utility to:

- Create new cluster configurations
- Model them against existing configurations
- Deploy the new, changed or old configuration to the same cluster
- Point the cluster configuration to another physical cluster anywhere in the enterprise, making cluster administration faster and easier

Summary

Cluster administration can be accomplished without having to train the staff by allowing a few administrators to manage all the clusters in the enterprise. This is Single Point of Control for cluster operational availability, integrity and management. This, coupled with browser-based anytime, anywhere access, is a testament to Compaq's management concept for ease of use, extensibility and functional enhancements for enterprise systems management.

Compaq *Insight Manager XE* – a complementary management interface for cluster monitoring

Compaq Insight Manger XE – Cluster Monitor is a complementary application that works with the industry-standard management agents. It uses a browser as the user interface, providing easy access at anytime from anywhere on the Intranet. The management application, Compaq Insight Manager XE consists of three components:

- Management Services that provide a core set of capabilities – such as auto-discovery, identification, data collection, central repository, event management, basic notification and secure access – that are used by add-ins from Compaq, a Management Solutions partner and/or an end-user of Compaq Insight Manager XE
- Systems Manager that provides real-time Web access for data on managed Compaq systems, notification of events and the capability to perform management operations on a group of Compaq systems
- Cluster Monitor that provides real-time event monitoring and display of cluster status ensuring the administrator can react quickly to potentially destabilizing cluster events

Compaq Insight Manager XE features on-line HTML documentation and robust context-based help. Insight Manager XE continues to build on the Compaq commitment to superior systems management capabilities that lower the total cost of ownership by delivering a leadership management product.

Intelligent Cluster Administrator Benefits

Reduces system management cost

This intuitive, easy-to-use tool is Web-based for anytime, anywhere operation. It allows the administrator to handle cluster configuration management and deployment from anywhere in the enterprise, as long as a browser is available. The information is presented to the administrator in an intuitive interface consistent with Compaq Insight Manager XE subsystems.

Improves operational efficiency and effectiveness

Manage your business – the way you want, when you want, where you want – with Web browser access to the Intelligent Cluster Administrator. Intelligent Cluster Administrator allows for rapid cluster modeling, operational monitoring and local or remote deployment, thus eliminating critical time lost due to travel or console based configuration management.

The Single Point of Control provided by the Intelligent Cluster Administrator allows increased operational efficiencies over traditional LAN- or console-based configuration tools.

Intelligent Cluster Administrator Version 2.0 provides value-addition by simplifying and extending the monitoring and managing of Microsoft Cluster Service. The Policy Management function differs from traditional management in that it is proactive.

Decreases application downtime
By knowing what cluster resources are required for each application, the administrator can dynamically reconfigure the cluster to eliminate failures and maintain operational efficiency for portions of the cluster without having the entire cluster go offline. The Intelligent Cluster Administrator will store optimal cluster configurations or special configurations based on operational needs, so the administrator can choose whichever configuration fits the need for that specific cluster to minimize application processing down time.

Become immediately productive
By working with an existing cluster configuration via the "export utility," the administrator can modify the configuration and "restore" it to its original cluster or "import" the cluster to another physical cluster. This provides rapid deployment of MSCS clusters across the enterprise, as well as lowers the learning curve by simplifying the setup and configuration, making available the management capabilities with limited investment.

User Interface Requirements

Netscape Communicator Version 4.5 or higher
Microsoft Internet Explorer 4.01 SP1 or Microsoft Internet Explorer 5.0 or higher

MODELS

Compaq *Insight Manager XE Version 2.0*

Compaq *Insight Manager XE v2.0* will be available on the Compaq Management CD starting with release 4.80. The Compaq Management CD ships with Compaq *ProLiant* and *Prosignia* servers and is included in the *SmartStart* subscription.

For additional information on *Compaq Insight Manager XE v2.0* or *SmartStart*, visit www.compaq.com/ manage.

Key Differentiators:

- Provides Web-based management to a broad base of Compaq and third-party devices including Intel and Alpha servers, clusters, desktops, workstations, portables, storage and networking products
- Support for heterogeneous operating environments – Microsoft Windows 2000, Microsoft Windows NT, Novell NetWare, SCO Unix, OS/2, OpenVMS and Tru64 UNIX
- Automatically discovers third-party HTTP Web server devices
- Dynamic queries help manage systems consistent with how they are used, either as a single device or part of a logical group
- SQL-compatible repository for discovered devices, alarms, asset and status reporting
- Supports Microsoft SQL Server 6.5, 7.0 and Microsoft Data Engine (MSDE)
- Integrates with Enterprise Management products. Support for HP OpenView, Tivoli NetView, CA Unicenter, Tivoli EC and SMS

Overview

Compaq *Insight Manager XE* leverages the power of the Internet to provide Web-based systems management and is a key enabler of the Compaq vision for virtual presence. Compaq *Insight Manager XE* reduces systems management cost, improves operational efficiency and effectiveness, and minimizes systems downtime. It provides device management capabilities that consolidate and integrate management data from Compaq and third-party devices using SNMP, DMI and HTTP. With Compaq *Insight Manager XE* you can monitor and manage groups of servers, clients, clusters and networking products anywhere, anytime from a standard Web browser.

With its integrated Cluster Monitor, Compaq *Insight Manager XE* provides administrators with a single monitoring point for clusters and stand-alone systems. In addition to the monitoring capabilities provided for MSCS clusters, version 2.0 adds support for the identification of OpenVMS, and Tru64 UNIX clusters, and is extensible for future support of Novell and SCO clusters. The Cluster Monitor aggregates system data and presents it to administrators as a single view of their cluster configurations. Compaq also offers the Compaq *Intelligent Cluster Administrator* that can be launched from Cluster Monitor to provide a single point of monitoring and control for MSCS clusters.

Compaq Insight Manager Windows console customers ready to upgrade to Compaq *Insight Manager XE v.2.0* can use the migration tool included on the Compaq Management CD. Once upgraded to Compaq *Insight Manager XE*, customers will be able to manage new Compaq and third-party devices without having to constantly upgrade their management application. See product specifications for minimum system requirements.

New in Version 2.0

Version 2.0 builds on Compaq Insight Manager version 1.1 by adding Version Control capabilities, more flexible support for third-party devices and SNMP traps, additional data collection capabilities, integrated paging capabilities and enhanced performance.

Compaq Insight Manager XE Version 2.0 Features

Based on industry standards – Offers consistent management of Compaq and non-Compaq devices using DMI V2, SNMP or HTTP protocols.
Accessibility to management information from anywhere in the IT environment – The Web-based browser provides full access from anywhere on the Intranet, eliminating the need for a dedicated Compaq Insight Manager management console.
Pre-defined queries and management tasks – Enables immediate access to management information.
System-wide disk threshold setting – Set disk threshold across a single system or group of Compaq servers.
Integrated alphanumeric paging and task scheduling – Increases mobility while allowing the administrator to respond quickly to alerts and alarms generated by the Compaq Management Agents.
Historical data collection – Selected information is collected and stored independent of source (DMI, SNMP).
ODBC database support – Lets users quickly and easily access data for reporting with industry-standard reporting applications.
Cluster Events Prioritization – For quick identification and resolution of potentially destabilizing events.
Version Control – Checks versions of software and firmware and indicates whether you should upgrade and why.

Password authentication – Provides a higher level of security than standard browser password authentication.
SNMP extensions toolkit – Contains SNMP MIB registration, SNMP MIB upload, trap editor, trap categories editor and command line utility to support the management of new devices without having to upgrade management application.
Integration with Novell NetWare 5.1 management – Discovers and links direct to the NetWare management portal on port 8008.

Compaq Insight Manager XE Version 2.0 Benefits

Easy migration from Windows console – Windows console customers (Compaq Insight Manager 4.7 or greater) can easily export host files for import into Compaq *Insight Manager XE* to create a managed device list. This task is easily accomplished using the Migration Tool provided on the Compaq Management CD.
Manages new devices without upgrading to a new version of Compaq Insight Manager – The SNMP Extension Tool provides the ability to add, discover and identify new devices. It updates the MIB database to the latest MIB revisions allowing you to discover and identify new devices.
Identifies outdated software or firmware – Version Control compares system software and firmware against a current database and indicates whether you should upgrade and why. It checks versions of Compaq OS drivers, Management Agents, Compaq Utilities and firmware.
One-stop cluster monitoring and administration utility – Integration of the Compaq Intelligent Cluster Administrator with the Cluster Monitor in Compaq *Insight Manager XE* provides a single interface to monitor, modify and deploy clusters throughout the enterprise.
Immediate use of the management application – Right out of the box, Compaq *Insight Manager XE* provides common system views with one click access from the main page. Even when you first use it, the application knows how to discover and report on your systems. Additional queries and events can be customized and defined by you without programming.
Automated device discovery process – Automatically discovers devices with internal HTTP Web servers. Links are provided directly to Web servers used for configuring the device or for monitoring status (networking devices, printers or third-party devices with Web agents).

Compaq Management Agents

The Compaq Management Agents expose the instrumentation of Compaq hardware and subsystems. They provide in-depth subsystem, status and fault information on servers, workstations, desktops and portables. The Compaq Management Agents communicate directly with Compaq *Insight Manager XE* and are the same agents used with the Compaq Insight Manager Windows console.

The Web-enabled Compaq Management Agents have the option to announce themselves to Compaq *Insight Manager XE* when they are powered up, limiting the need to perform auto-discovery. The Web-enabled Compaq Management Agents are accessible either directly through a browser or through Compaq *Insight Manager XE*. Most Compaq Management Agents provide the ability to browse directly through Compaq systems for real-time system, subsystem and component information and status.

Compaq *Insight Manager XE* and the Compaq Management Agents allow you to browse directly to system information and launch Compaq Insight Manager LC, Compaq Intelligent Cluster Administrator or non-Compaq Web services running on a device.

Web-enabled agents are now available for systems running Microsoft Windows 2000 or NT, Novell NetWare, SCO UnixWare, UnixWare 7 and SCO OpenServer 5.

SPECIFICATIONS

Minimum Hardware and Software Requirements for Compaq Insight Manager XE

Hardware	Compaq *Prosignia* or *ProLiant* server
System memory	128MB of RAM with Microsoft SQL Server or MSDE on same server –or– 96MB of RAM with Microsoft SQL Server running on a remote system
Disk space	200MB for Compaq *Insight Manager XE* software 450MB for the database server drive and database log
Server operating system	Microsoft Windows NT Server 4.0, with Service Pack 5 or later Microsoft Windows 2000 with SNMP Hot Fix Q253302 (available on the Compaq Management CD) –or– Microsoft Windows 2000 with Service Pack 1
Server software	TCP/IP installed SNMP services installed and active Microsoft Internet Explorer 4.01 with Service Pack 2 or Microsoft Internet Explorer 5.0 Compaq Server Support for Microsoft Windows NT 4.0 2.12c or later JVM 3188 or later, found at **www.microsoft.com/java**
Database	Microsoft Data Engine (MSDE) included on Management CD –or– Microsoft SQL Server 6.5 with Service Pack 5a –or– Microsoft SQL 7 with Service Pack 1 or later

Minimum Hardware and Software Requirements for the Browser

Web browser	Microsoft Internet Explorer 4.01 (with Service Pack 2) or Microsoft Internet Explorer 5.0 or later JVM 3188 or later for Microsoft Internet Explorer 4.01 Service Pack 2 or Microsoft Internet Explorer 5.0. You can find this software at **www.microsoft.com/java** and **www.microsoft.com/ie**
System memory	32 MB RAM for Microsoft Windows 95 or Microsoft Windows 98 48 MB RAM for Microsoft Windows NT

Requirements for Compaq Management Agents

Supported Operating Systems on Compaq *Prosignia* and *ProLiant* servers:	
Novell:	• NetWare 3.12, 3.20, 4.11, 4.20, 5, intraNetWare, intraNetWare for Small Business
Microsoft:	• Windows 2000, Windows NT 3.51 and 4.0
SCO:	• UnixWare, 2.1.2. 2.1.3, 7, 7.0.1, 7.1, 7.1.1 OpenServer Release 5.04, 5.05
IBM:	• OS/2 Warp Version 3, Version 4

Supported Operating Systems on Compaq AlphaServers:

Compaq:	• Tru64 UNIX V4.0F • OpenVMS 7.1, 7.1-1H1, 7.1-1H2, 7.1-1H3, 7.1-2 and 7.2

Supported Operating Systems on Compaq Workstations:

Supported Models:	• Compaq Professional Workstation
Supported Operating Systems:	• Microsoft Windows 2000, Windows NT Workstation 4.0

Supported Operating Systems on Compaq Desktops:

Supported Models:	• Compaq Deskpro
Supported Operating Systems:	• Microsoft Windows 2000, Windows NT Workstation 4.0, Windows 95, Windows 98

Supported Operating Systems on Compaq Portables:

Supported Models:	• Compaq Armada, LTE 5000 (Win 95 & Win NT)
Supported Operating Systems:	• Microsoft Windows 2000, Windows NT Workstation 4.0, Windows 95, Windows 98

Supported Uninterruptible Power Supplies:

Supported Models:	• All Compaq UPS Models: T700, T700h, T1000, R1500h, R3000, R3000h T1000h, T1500, T1500h, T2000, T2400h, R1500,
Supported Operating Systems:	• Microsoft Windows 2000 and NT 3.51, 4.0

Download the latest Compaq Management Agents for servers or clients from:
www.compaq.com/support/files/index.html

For more information about OpenVMS agents and installation, please refer to:
www.openvms.digital.com/openvms/products/mgmt_agents/index.html

For more information about Tru64 UNIX agents and installation, please refer to:
www.unix.digital.com/faqs/publications/pub_page/pubs_page.html

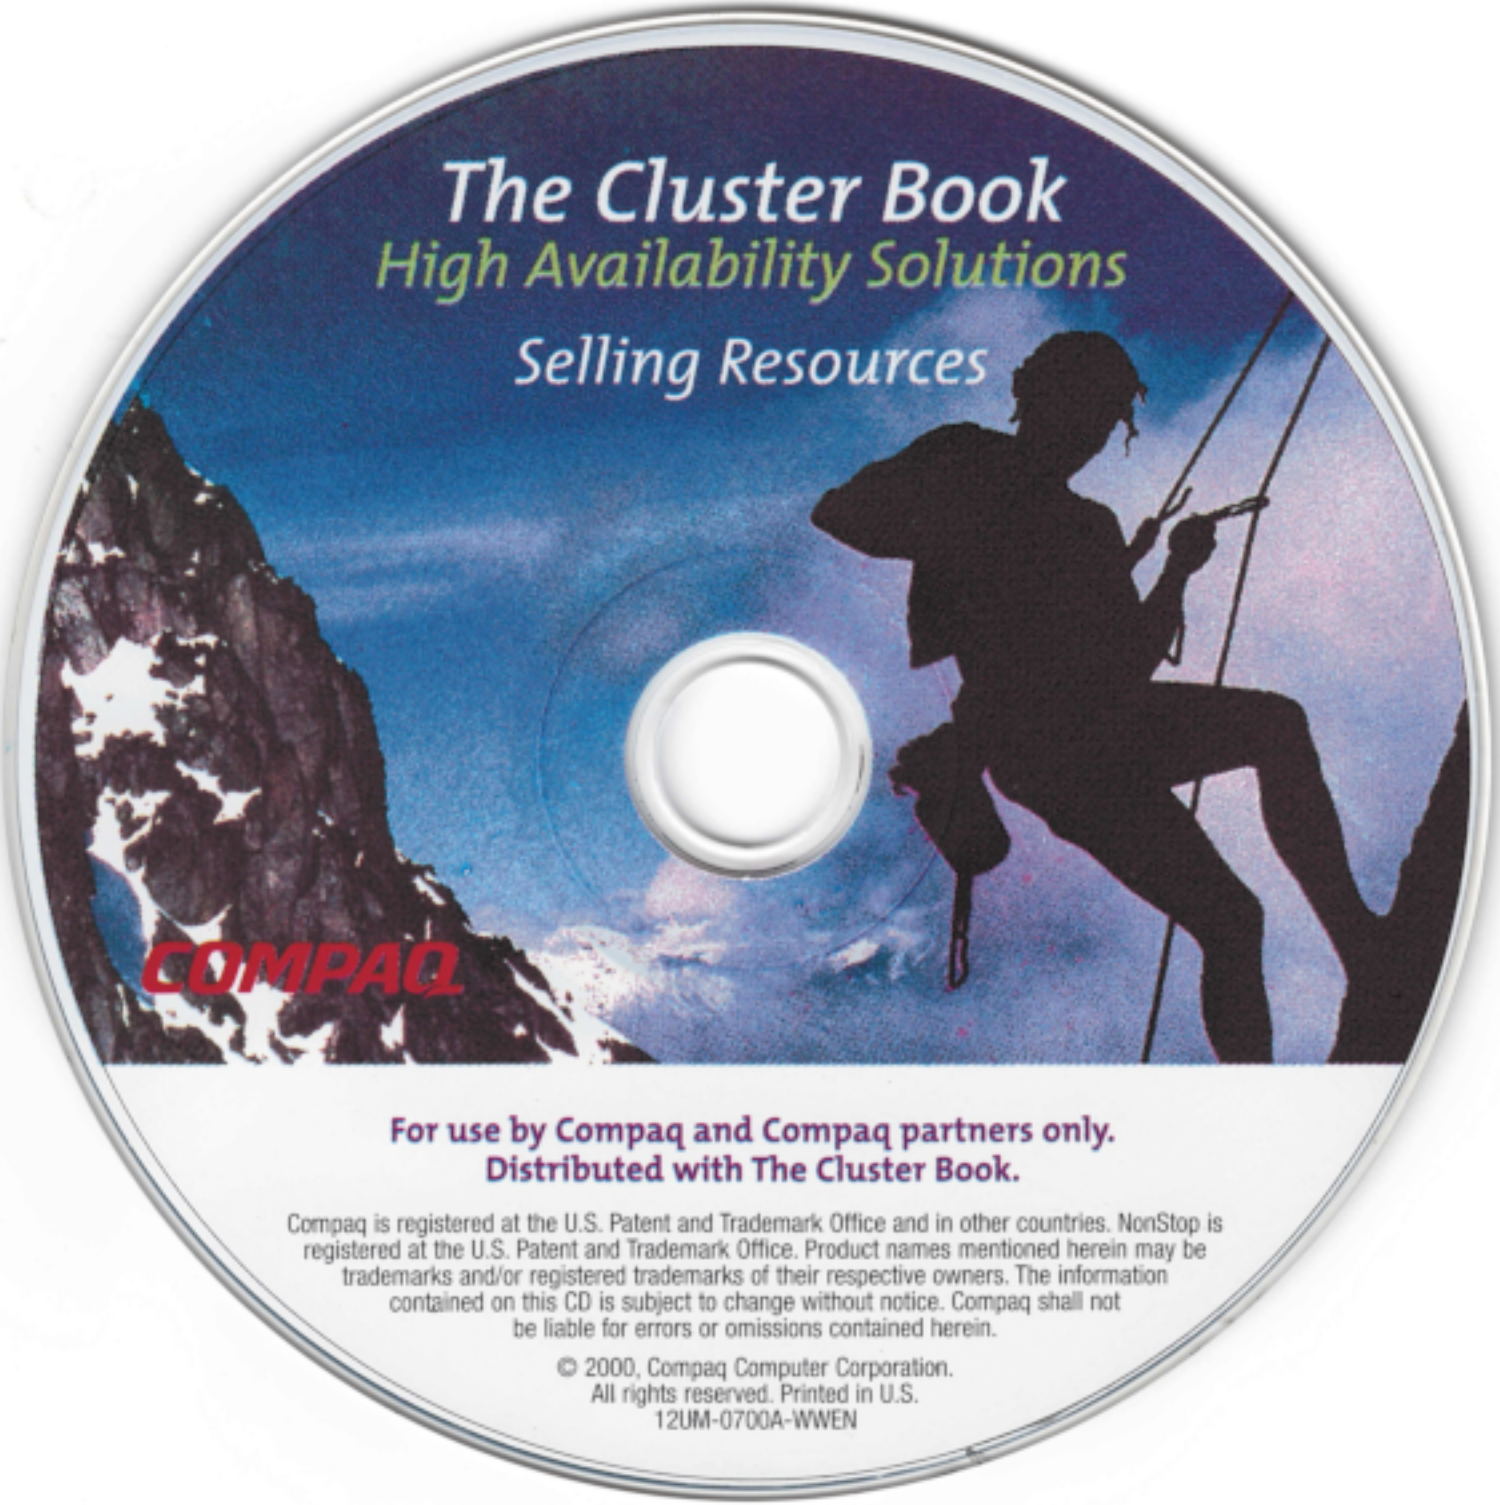

For use by Compaq and Compaq partners only.
Distributed with The Cluster Book.

12UM-0700A-WWEN