TECHNOLOGY BRIEF

October 1997

Compaq Computer Corporation

CONTENTS

Compaq CD Storage System

EXECUTIVE SUMMARY

New technologies are advancing the Compaq CD Storage System as a means of storing valuable data, reducing the amount of information stored on server disk drives. This technology brief describes an implementation of secondary storage. The CD Storage System is a near-line, host-attached server option. The discussion begins with the basic building blocks of the small computer systems interface (SCSI) and logical unit number (LUN) technologies. This brief demonstrates how the combination of SCSI and LUN technologies developed into an innovative solution to near-line, mass-storage and retrieval issues. It includes a high-level comparison of matrix architecture to standard serial SCSI architecture. This paper discusses the need for near-line storage technology and the direction Compaq is taking with the CD Storage System. Additionally, this paper describes some typical applications for the CD Storage System and the benefits provided by its array of compact disk-read only memory (CD-ROM) drives. Also included is a high-level description of the mass-storage management software bundled with every CD Storage System and the operating systems it supports.

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

COMPAQ

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 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 CD Storage System Technology

First Edition (October 1997)
Document Number ECG060.1097

INTRODUCTION

Customer demand to access ever-larger amounts of archived information grows daily. Compaq Computer Corporation understood this need and developed a mass-storage server option to meet it. The Compaq CD Storage System enables organizations with limited budgets or the need to provide multi-compact disk (CD) access to multiple users across the enterprise to expand their CD-storage systems as their business grows.

New technologies are advancing the CD Storage System as a means of storing valuable data, away from server disk drives. The CD Storage System is a near-line implementation that includes seven 12x CD-ROM drives totaling 4.5 gigabytes (GB) of data storage capacity. Up to seven CD Storage Systems can be configured into an array or matrix of devices that provides as many as 49 CD-ROM drives with a combined capacity of 31.5 GB – all attached to a single SCSI adapter port. The matrix architecture of CD-ROM storage devices in the CD Storage System allows flexible, cost-effective, secure access to valuable information.

A CD Storage System attaches directly to any Compaq server by way of the server's SCSI adapter port and cable attachments. Server-to-device communications occur by way of the small computer systems interface-logical unit number (SCSI-LUN) controller that is an integral part of each CD Storage System. The mass-storage management software that ships with the CD Storage System integrates the CD-ROM drives to the server.

The CD Storage System is available in both tower and 4U rack-mount form factors. The tower and rack solutions share field replacement units. Figure 1 is a photograph of the tower form factor.

Figure 1. The Compaq CD Storage System

Storage Classifications

Compaq recognizes two major categories of distributed-network storage: primary and secondary storage. Within these two categories, there are three distinct implementations.

Table 1: Storage Categories and Implementations

Storage Category	Implementation	Example	Media Type
Primary Storage	On-Line	Server Disk Drive	Non-removable
Secondary Storage	Off-Line	Archival Backup	Removable – Tape
	Near-Line	CD Storage System	Removable – CD-ROM

CD-server, a term often used instead of *near-line implementation*, is an industry buzzword mistakenly used to categorize a variety of near-line secondary storage devices. *CD-server* does not accurately describe all available forms of CD-storage products because there are significant cost and performance differences among all the products in the near-line category. Table 2 lists the different device classes of the near-line secondary storage implementation. The Compaq CD Storage System is a host-attached device that functions as an extension of the server.

Table 2: Near-Line Device Classes

Implementation Type	Device Classes	Comments
Near-Line Secondary Storage	Dedicated Server	Expensive, High-end Solution
	Network-attached	Moderately priced mini-server
	Host-attached	Extension of the Server

Standard SCSI

Server expansion slot "real estate" is expensive. The goal is to maximize the value associated with server expansion slots. Compaq uses SCSI because it is capable of connecting multiple peripheral devices to a single expansion board. For the CD Storage System, the SCSI interface is preferred over the integrated drive electronics (IDE) interface because SCSI allows for more expandability. Single channel IDE allows only two peripheral devices per expansion board, and enhanced IDE allows only four devices. The standard SCSI bus is an 8-bit parallel bus supporting up to seven separate devices – internal, external, or any combination.

For additional details about Compaq SCSI technology, refer to *Compaq Fast-Wide SCSI-2 Technology*, Compaq document number 101A/0395, and to *Compaq Wide-Ultra SCSI-3 Technology*, Compaq document number 510A/0697.

Standard SCSI devices attach serially to a host adapter via the SCSI bus; this configuration is called a chain. For attached devices to work on the standard SCSI bus, each device must have a unique identification (ID) number on the chain, ranging from zero to seven. A SCSI host adapter or SCSI controller expansion board plugs into the server. Because the host adapter uses one of the available

SCSI ID numbers, a maximum of seven peripheral devices can attach to the standard SCSI bus. The host adapter interprets requests from the computer and in turn, issues commands or requests and communicates to the other devices on the SCSI bus. When the host adapter translates requested information from, or issues a command for a specific device, it addresses that device by its specific SCSI ID number.

Figure 2 illustrates how unrelated storage devices attach to a standard SCSI bus and SCSI host adapter. In this illustration, the internal SCSI hard disks take two available SCSI ID numbers on the SCSI bus (the internal chain). The devices on the external SCSI bus (the external chain) must have their ID numbers set so as not to conflict with the internal devices. In addition, the last device in each chain must be terminated.

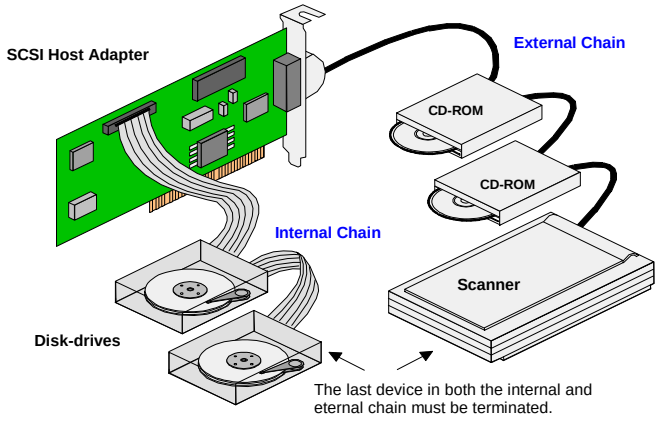

Figure 2. A standard SCSI host adapter with internal and peripheral devices attached

SCSI-LUN

Compaq uses SCSI-LUN technology in the CD Storage System to allow connection of more than seven peripheral devices to a single SCSI host adapter. The technology in the CD Storage System expands the connectivity of standard SCSI architecture by using an array or matrix of SCSI-LUN devices. LUN technology defines each SCSI device attached to the standard SCSI bus, including the SCSI host adapter, as one of eight logical units. A LUN identification number, ranging from zero to seven, identifies each SCSI device on the chain. With LUN technology, the standard SCSI bus retains the ability to support seven devices, and may contain a series of chained SCSI-LUN devices, replacing individually chained SCSI devices. If the devices replacing chained SCSI devices are SCSI-LUN controllers, as in Figures 3 and 5 below, then the SCSI ID number of each device on the bus may represent up to seven CD Storage System units.

As implemented in the CD Storage System, the SCSI-LUN ID number of the imbedded controller serves as the address for that individual unit. In this configuration, where the controllers replace a standard SCSI device on the SCSI bus, the SCSI host adapter accesses a SCSI-LUN device controller board, SCSI ID number *n*, and LUN zero. The value of n is the number assigned to that SCSI device on the standard SCSI bus, and that value varies by vendor. The CD Storage System units have an assigned value of SCSI ID 2, LUN 0.

Standard SCSI allows any two devices to communicate at one time. SCSI-LUN technology allows devices to communicate to other devices independently, addressing each via unique SCSI-LUN ID numbers. A further explanation of SCSI-LUN communications follows in the section “SCSI-LUN Device Communications.”

Figure 3 illustrates how storage devices attach to a standard SCSI bus and SCSI host adapter. In this illustration, the internal disk-drives have not changed from Figure 2 above. The external chained devices have been replaced with serially attached CD Storage System units. Cables attach to the rear of the CD Storage System, which comes equipped with two 50-pin narrow SCSI connectors and SCSI terminator. Each device on the external SCSI bus (the CD Storage System) should have the SCSI ID number set (located on the back panel) so as not to conflict with the internal devices, and devices at the end of the chain should be terminated.

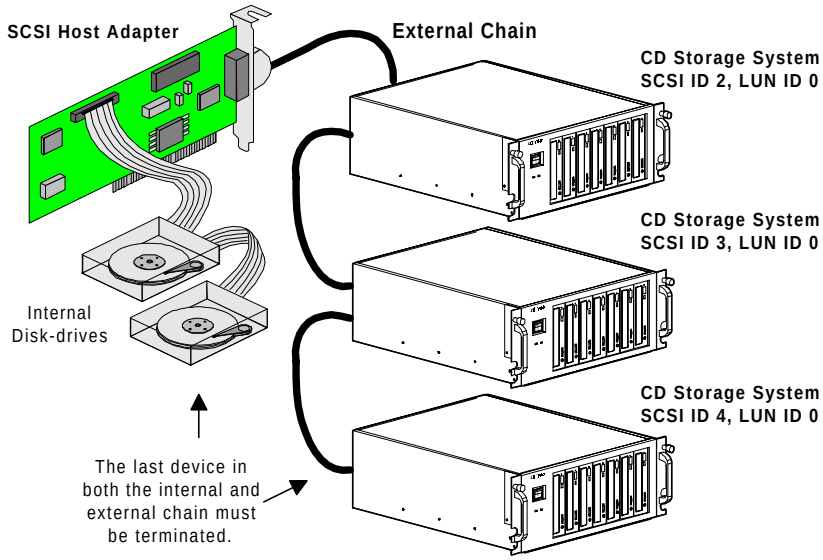

Figure 3. Externally attached SCSI devices – CD Storage Systems

Figure 4 is a block diagram of the serial architecture of the standard SCSI bus, with a SCSI host adapter and seven additional peripheral devices attached.

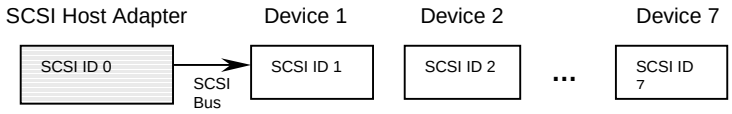

Figure 4. The standard SCSI bus with seven peripheral devices attached

Figure 5 is a block diagram of the unique addressing scheme of the SCSI-LUN matrix architecture. This block diagram illustrates how CD Storage Systems use SCSI-LUN controllers to attach to the SCSI host adapter via the SCSI host bus. The SCSI-LUN controller internal to each CD Storage System (Devices 1, 2, etc.) also connects to the SCSI CD-ROM drives via the internal SCSI-LUN device bus.

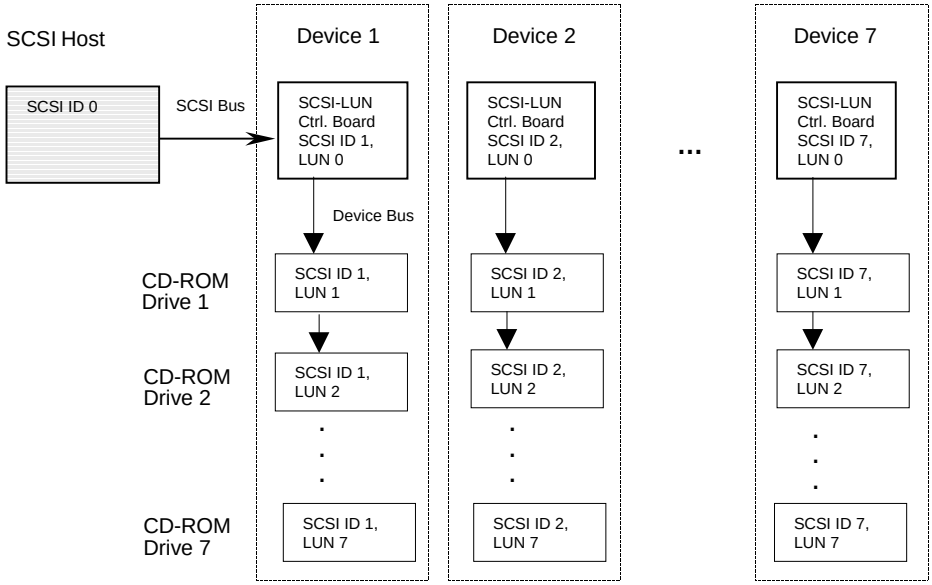

Figure 5. The SCSI-LUN matrix configuration with peripheral devices attached

SCSI-LUN Device Communications

The integral SCSI-LUN controller boards manage server-to-device communications. The process through which an array of CD Storage Systems and the SCSI-LUN control communication involves several steps. The SCSI-LUN controller allows the server, via the SCSI host adapter, to communicate to any of the attached devices on the host or device bus. The SCSI host accesses individual CD-ROM drives by addressing the SCSI and LUN ID number for any specific drive.

For a typical end-user wanting information from a specific CD-ROM stored in a CD Storage System, Figure 6 illustrates how it works.

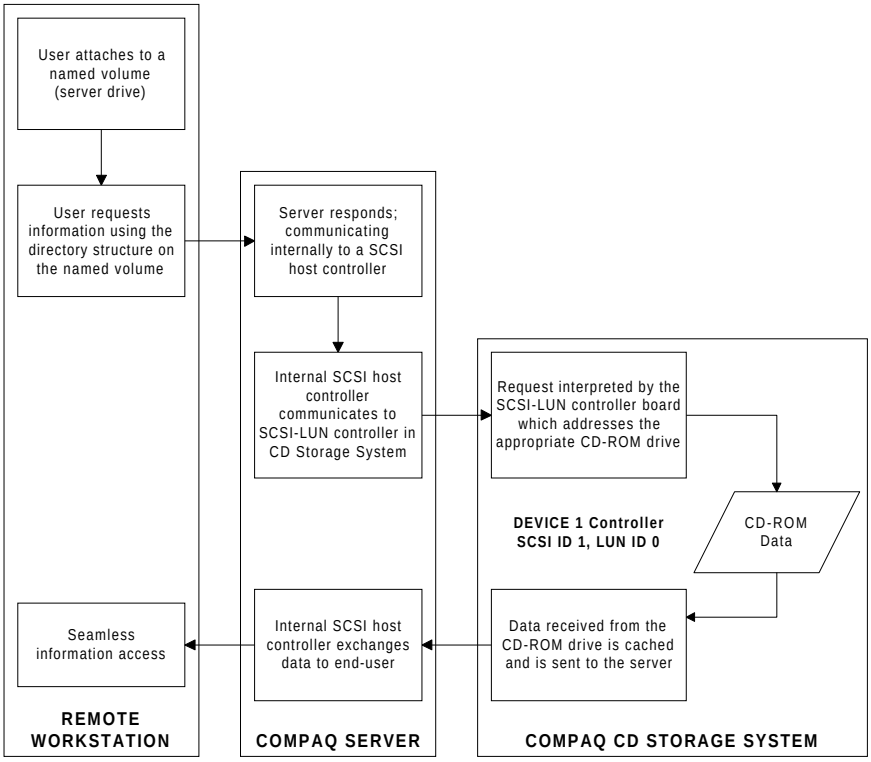

Figure 6. Flow chart of a remote workstation requesting CD-ROM data

From a remote workstation, the end-user attaches to the named volume and looks for the information. The user uses the same command statements (that is, point and click) as if it were any other network drive. The server responds by communicating to the CD Storage System using the internal SCSI host controller to the device SCSI-LUN controller, listed in the above example as the Device number 1 – SCSI ID 1, LUN ID 0. The CD Storage System's SCSI-LUN controller interprets the request, and the appropriate CD-ROM drive sends the data to the cache buffer. The server in turn processes the data directly to the requesting end-user.

While the SCSI CD-ROM device is processing the information requested to pass it back to the SCSI-LUN controller, the SCSI host adapter may access information from any other CD-ROM drive in the matrix. The host controller routinely issues multiple requests to the SCSI-LUN controllers and to the drive units individually. The devices process requests and commands independently.

Communication via a SCSI-LUN controller to a CD-ROM storage device is like that to the server disk drive, in that the data stored can be randomly accessed. The CD Storage System is currently configured so the CD-ROM drives are read-only devices, eliminating the ability to configure those devices into any level of redundant array of independent drives (RAID). Future considerations are explained in the section "The Future for the CD Storage System."

CD Storage System Configurations

A Compaq CD Storage System includes seven CD-ROM drives. Each CD-ROM drive acting independently provides 650 megabytes (MB) of storage capacity. Thus, a single CD Storage System has a combined storage capacity of 4.5 GB. Each CD Storage System uses true 12x speed CD-ROM drives, which transfer data at a constant linear velocity (CLV). CLV means that the data reads at the same rate no matter where it is on the media. Compaq chose this type of CD-ROM drive for use in the CD Storage System because it maximizes the density of data in each track and thereby allows full use of the available disk space. Please see the Appendix for explanations of true transfer rates, constant linear velocity, and constant angular velocity.

Figure 7 shows the basic configuration of the CD Storage System as a host-attached device.

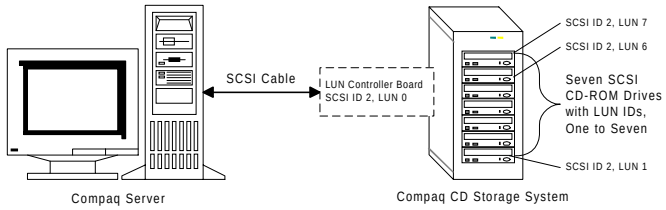

Figure 7. Basic configuration of a single Compaq CD Storage System

The CD Storage System uses a standard SCSI bus, incorporates LUN technology, and brings together standard SCSI CD-ROM devices. This expands standard SCSI architecture from a serial configuration to an array or matrix configuration. SCSI-LUN architecture makes it possible to chain up to seven CD Storage Systems (49 CDs for a combined capacity of 31.5 GB) into a matrix configuration that connects directly to a server using a single slot. Additional CD Storage Systems may be attached to the same server only if another SCSI host adapter is added. Figure 8 illustrates the maximum configuration of seven CD Storage Systems connected to a Compaq server via one server-SCSI adapter port.

To ensure the most stable combination of software and hardware, Compaq recommends chaining no more than seven CD Storage Systems in an array configuration. Furthermore, Compaq recommends using industry-standard SCSI repeaters. SCSI cable length limitations require use of a repeater if SCSI cabling connecting multiple units exceeds 18 feet (5.486 meters) in length. The repeater in Figure 8 is for illustration purposes only. The actual location of repeaters will vary depending upon individual customer requirements and physical layout of the plant. If physical layout and space limitations are a concern, maximum configurations may best be achieved using rack mount models in Compaq rack systems.

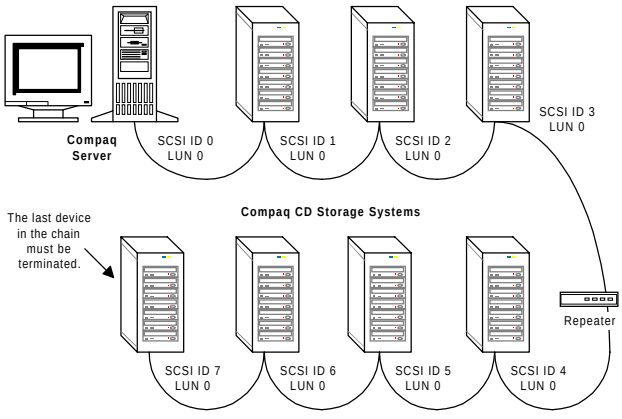

Figure 8. Maximum configuration for multiple CD Storage Systems

CD Storage System Software Support

The Compaq CD Storage System provides the customer with a secured distributed data source for the enterprise. The CD Storage System includes bundled configuration and mass-storage device management software called SCSI Express. The SCSI Express software is configurable for workgroup or departmental access to all or part of the CD-stored information.

The CD Storage System provides multiple device-drive access as a single named volume. Volume-naming conventions usually dictate reserving drives A through F, leaving only 20 other "named" possibilities. Rather than being limited to the volume-naming conventions and the English alphabet, the software packaged with every CD Storage System configures multiple devices as a single volume name. The software also configures to allow automatic assignment of a directory structure for each CD-ROM in the system.

To a remote workstation client, the CD Storage System has the appearance of a named-volume directory structure. The end-user simply attaches to a network drive, not unlike any other network drive on the server. This allows multiple users (totally configurable and controllable by the administrator) access to any CD-ROM drive. A server-attached console allows the administrator to look at a specific file in the directory tree on a named-volume server drive, or to request information directly from a specific CD-ROM drive.

The SCSI Express software solution is compatible with Microsoft Windows NT versions 3.51 and 4.0, Novell NetWare version 3.12 and IntranetWare. Remote administration of the CD Storage System is fully configurable within the SCSI Express management software for Windows NT and Windows 95 workstations. The Compaq Support Software CD will also be provided with each Compaq CD Storage System.

Windows NT Inspector – Support Functions:

- Manage and view the CD ROM media and all supported SCSI devices on your system
- Create/delete CD-ROM groups
- Add/delete or move media among CD-ROM groups
- Lock/un-lock media on your system, or share media (via copy feature)
- Update the media as new CDs are added
- View properties of the group, folder, and the media

Novell NetWare – CD-Utility Program

- View devices, media, volume, and configuration

NetWare Inspector – Client/Remote Utility

Troubleshooting – Maintenance Utility

- Configure host adapters
- Create SEAGENT account
- Add support options

Some network operating systems do not require SCSI Express software. SCO UnixWare 2.1x and SCO OpenServer 5.0.x and 3.0 operating systems do not require SCSI Express software, due to native Unix CD support. SunSoft Solaris for Intel Edition Version 2.5, 2.5.1, or 2.6 and IBM OS/2 version 2.x or later do not require the SCSI Express software due to native operating system support.

COMPAQ TECHNOLOGY SOLUTIONS – CD STORAGE SYSTEM

Businesses today generate huge amounts of data, and the volume of that data is growing exponentially. Moreover, end users require immediate access to that data. Typically, today's organizations need to group information into categories, with secured access for specific groups such as accounting, legal, and technical. The Compaq CD Storage System provides flexible, secure, cost-effective mass-storage and retrieval capabilities to meet the ever-increasing demand for access to near-line CD-stored data. Law libraries, educational reference material, health data banks, tax regulatory information, governmental archives, and data warehousing are only a few uses for the CD Storage System.

CD Storage System – Compaq Solutions to Customer Issues

Organizations with limited budgets or the need to provide multi-CD access to multiple users across the enterprise can benefit from the CD Storage System. With the CD Storage System, these organizations now have the ability to expand their CD systems as business grows. The CD Storage System offers many benefits:

- **Mass-storage data-retrieval capabilities** – The CD Storage System provides a centralized data source, mass-storage and retrieval capabilities of up to 31.5 GB per server-SCSI port with advanced SCSI-LUN hardware and software support. The CD Storage System provides access to all CD-ROM drives at power-up, assuring data availability.
- **Secure access to data by way of the server** – The CD Storage System with the advanced SCSI-LUN technology connects directly to the SCSI port on a Compaq server, and not directly to the network. The SCSI-LUN connection communicates information at the speed and with the security of the server-attached SCSI host adapter. The SCSI-LUN technology is SCSI-2 compliant and is compatible with Compaq SCSI adapters, providing continued investment protection. The CD Storage System also allows faster access to the data than that of conventional linear tape products.
- **Distributed source for the enterprise** – The CD Storage System easily configures to reduce or even eliminate multiple copies of software. Reducing the number of physical copies of software may also eliminate some software licensing and tracking issues. In some cases, purchasing a single copy of software and site licensing can actually reduce overall network support operating costs.
- **Cost effective information source** – Today the compact disk is a globally accepted medium for distributing information such as catalogs, databases, distribution software, or operating systems. By eliminating the need for organizations to purchase CD-ROM drives for every desktop, the CD Storage System represents a cost effective means of accessing information distributed and stored on compact disks. The CD Storage System provides lower support and installation cost, improved system storage manageability, and increased overall capabilities.

Typical Applications for the CD Storage System

Here are some typical customer applications and examples of how they could benefit from the CD Storage System:

- A leading banking organization has its own CD storage solution, complete with software developed in-house. Their solution involves a single-station, seven-device SCSI CD-ROM tower connected to a server. Instead of limitations of a single storage device, Compaq offers a multi-device mass-storage system. This organization would benefit from the CD Storage System SCSI-LUN technology. Their ability to expand successfully database capabilities and relieve volume-naming restrictions would be enhanced by the CD Storage System.
- An independent, not-for-profit, standards organization that publishes national standards uses a competitive product, a seven-drive SCSI CD-ROM device. This organization would benefit from the CD Storage System SCSI-LUN technology by expanding its database and reference library capabilities. The customer indicated that a rack-mounted version of the CD Storage System best fits their need and is the most cost-effective solution. Compaq has always made solutions available that would adapt well to a client's environment.
- Webmasters would benefit from the CD Storage System's ability to expand the capabilities of their Web server. Attaching several Compaq CD Storage Systems to a Web server overcomes the limitation of seven standard SCSI attached drives. If the CD Storage System were employed, the web site might contain up to 31.5 GB of shared information per SCSI port. With the CD Storage System, potential users could access Web sites and information seamlessly. Configuring the CD Storage System SCSI Express software allows the administrator to add CD-ROM drives as supplemental directories under a single named volume.

THE FUTURE FOR THE CD STORAGE SYSTEM

The Compaq CD Storage System is the first in the planned family of Compaq CD-ROM storage array systems. It allows the flexibility to upgrade to other technologies as they mature. Compaq looks to the future with possible features for the CD Storage System such as hot-pluggable devices, network attachment, and other technologies. Future drive technology may include the CD-recordable (CD-R), the CD-rewritable (CD-RW), the digital versatile disk (DVD), or magneto optical (MO). With the CD-R and CD-RW, the intent is to give end-users the ability to record or update data periodically on storage devices with removable media. The DVD and MO are possible options that would allow end-users to expand on the current business trends toward multimedia presentations and capacity. The direction and future products continue to solve the customer needs for mass-storage, greater configuration flexibility, and investment protection.

CONCLUSION

Compaq is a leader in secondary storage systems, with quality products such as Compaq Libraries (DLT and DAT); Compaq Tape Arrays, and individual tape drives that address a variety of customer needs for capacity, performance, and economy. Consistent with an established history of strong end-to-end solutions, Compaq incorporates powerful feature sets into the CD Storage System and all of its hardware product lines, delivering solutions to customer challenges. Compaq meets the CD-storage demands with the new CD Storage System that provides secure mass-storage and data-retrieval capabilities, while protecting current business investments.

APPENDIX – CD TECHNOLOGY BACKGROUND

The CD is a 12 cm/8 cm diameter plastic disk, on which a variety of data may be stored. It has a film layer of lacquer over a thin, highly polished aluminum plate on a polycarbonate substrate. The digitally stored data on CDs may range from audio or video reproductions to computer files. The Phillips and Sony Corporations first produced audio CDs and CD-drives in 1987. CD-ROM drives were introduced shortly thereafter but not generally adopted until 1994. Figure A-1 shows an exploded view of a CD and the laser mechanism used to read the data from the aluminum layer.

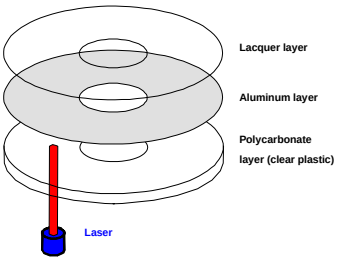

Figure A-1. The structure of a CD

The storage capacity of the standard CD is 650 megabytes (654.74 MB). Information on a CD is stored in a 3-mile (4.8-kilometer) long spiral of microscopic "pits" and "lands" that makes up a 2.2 micron wide "track." By comparison, if laid on the CD substrate, the width of a human hair would cover over 50 rows of pits. The track begins at the center of the disk and spirals its way to the outer edge, just like phonograph records (only in reverse). The pits are "burned" on the recording layer by bursts of a high-powered blue argon laser during optical encoding. The lands are the untouched areas between the pits. The process of burning a CD usually takes place on a master copy, which is in turn copied, or stamped. With today's CD-ROM technology, once a CD is burned or stamped into the substrate, the data is not susceptible to being over-written. This near-indestructible nature increases the life expectancy and value of the CD as a storage medium.

Figure A-2 illustrates the laser and sensor registering whether the beam of laser light is reflected by the land or absorbed by the pit. A CD-drive contains only one laser and sensor; two are shown here for demonstration purposes only. The sensor electronically translates the change in reflection as a binary system of ones and zeroes encoded on the aluminum data layer. The translation is in turn transmitted to other controlling devices, which communicate recorded data.

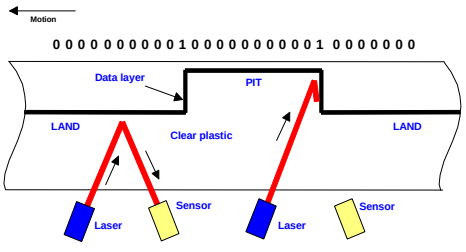

Figure A-2. Reading the CD device data structure

Transfer Rates, CLV, and CAV

First generation CD-drives were single-speed drives. They were slow, with data rates of only 150 kilobits/sec (150 KB/sec, sustained) and they were very expensive. Today, much faster and less expensive CD-drives are available. Their designated speeds are actually multipliers of the sustained transfer rate taken from the single-speed audio CD-drives. For example, a true 12x CD-drive has a "true" sustained transfer rate 12 times that of the audio CD: 12 times 150 KB/sec, or 1800 KB/sec.

Constant linear velocity (CLV) allows the head to read the same length of track at all times and radii (1.3 meters/sec). The data reads at the same rate independent of the media, but the velocity of the disk changes. In other words, the disk spins slower as the head moves to the outer edge of the disk. A CD-ROM using the CLV method spins from 539 revolutions per minute (RPM) at the inner edge, to 210 RPM at the outer edge. Depending on which track is being accessed, the density of bits (binary code) in each track can be made uniform by changing the speed of the disk. This allows the outer tracks to hold more data than the inner tracks and fully uses the available space on the disk.

Another method of reading data from a CD-drive uses a CAV to transfer data. CAV, measured in RPM, means that the read head sweeps the same angle, for the same amount of time, at all radii. In other words, the transfer rate for data varies substantially depending on where the laser reads the data on the media. The data reads slower on the inside tracks and at higher rates on the outside tracks, but the disk rotates at a constant RPM. The number of bits in each track is the same, but the density of bits varies because the inner tracks have smaller circumferences than the outer tracks. A so-called high speed CD-drive, such as the 16x or a 32x reads data at those speeds only on the outer tracks; these types of drives therefore are not considered "true speeds."

Since SCSI transfer rates are measured in a pure test environment and not all enterprise systems are created equal, then performance of SCSI-LUN technology affects everyone. Performance of the SCSI-LUN technology is dependent on variables; various and not necessarily quantifiable data. For example, performance depends on the class of server and other network devices used, the number of SCSI devices attached to that server-SCSI port, the total number of users accessing data at any given time, or the level of server activity. These are only some of the possibilities, and each variable will affect overall performance numbers. Throughout, Compaq's goal has been to maximize the sustained data transfer rate with a minimum of hardware and software overhead changes.