

February 2000
11YP-0200A-USEN

Prepared by Storage Products
Division

Compaq Computer Corporation

Contents

Executive Overview.....	3
Solution Description.....	3
Solution Components	4
Overview.....	4
RAID-levels supported	7
Number of Disks needed for a BCV.....	7
Number of volumes supported.....	8
Performance implications on the production system caused by the BCV.....	8
Expected performance of the BCV.....	9
Planning.....	11
Script Process	11
Backup Scripts	12
Methodology.....	13
Sizing and characterization data ..	13
User Application Impact.....	13
Backup Rates Table	14
Conclusion	15
Learn more	17

Technical Best Practices of Compaq Non-Disruptive Backup Solution for Windows NT

Abstract: The Compaq Non-Disruptive Backup Solution for Windows NT provides IT administrators with the necessary hardware, software, services and configuration information to overcome the restrictions imposed by the shrinking backup window on 24x7 business operations. The solution employs Compaq SANworks Enterprise Volume Manager, in conjunction with Compaq StorageWorks RAID subsystems, Enterprise Backup Solution, and DLT tape libraries to create point in time copies of critical data for off-line backup. This paper discusses the advantages and limitations of deploying the Non-Disruptive Backup Solution for the off-line backup of Oracle and Microsoft Exchange databases, employing both cloning and snapshot technologies. This complete solution offers a fully qualified custom configuration to overcome the complexities of ever changing backup strategies.

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.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Server and Workstation, Microsoft SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and intraNetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

Copyright ©1998 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.
Technical Best Practices of Compaq Non-Disruptive Backup Solution for Windows NT
White Paper prepared by Storage Products Division

First Edition (February 2000)
Document Number 11YP-0200A-USEN

I. Executive Overview

As the Internet drives companies from a conventional to an e-business model, backup windows have become a thing of the past. Data and applications must be available 24 hours a day, 7 days a week—driving businesses to create a backup infrastructure that doesn't interfere with the availability and performance of their primary storage systems. In addition, development of an enterprise-wide backup strategy has become increasingly complex.

The *Non-Disruptive Backup Solution* is a *Compaq NonStop*[™] eBusiness solution designed for businesses that need to run 24x7—especially those that rely on e-commerce, Oracle databases, and Microsoft® Exchange. This integrated storage solution solves the two major business problems associated with backup and recovery today:

- **Shrinking backup window**
- **Complexity of backup**

The *Compaq Non-Disruptive Backup Solution* is a combination of products and services that provides a complete backup infrastructure for both Storage Area Networks (SAN) and SCSI direct-attached secondary storage environments. Intended for users of *Compaq StorageWorks Raid Array RA8000* and *Enterprise Storage Array ESA12000* systems, it provides a complete tested and characterized backup solution. The solution manages the snapshot and cloning features of the *StorageWorks RAID* array using *Compaq SANworks Enterprise Volume Manager*, and includes customizable application and backup scripts for quick success. This solution initially supports Windows NT®, and in the second quarter of 2000 will also support Sun Solaris and *Compaq Tru64 UNIX*® operating environments.

"NT Server is now taking a place alongside mainframe and UNIX operating systems as a major platform for business critical applications. Compaq's Non-Disruptive Backup Solution helps ensure 24x7 availability of NT applications by eliminating the backup window and makes network bandwidth more available by moving high volume data transfer traffic from disk to tape off the LAN. A centralized backup strategy for NT storage can also be implemented to ensure backup consistency and restoration reliability"

—John Webster, Illuminata

II. Solution Description

The *Compaq Non-Disruptive Backup Solution for Windows NT* is a tested, fully integrated configuration of hardware, applications, backup software, and services that alleviates dependency on a specific time window for backup—while providing the consolidation, redundancy, and manageability benefits of tape automation. It provides a backup infrastructure for both SAN (dual loop or switched fabric) and SCSI (direct-attached) secondary storage environments—resulting in an end-to-end data protection environment for both primary and secondary storage.

In addition to being compatible with popular third-party backup software, the solution has been characterized with Oracle and Microsoft Exchange on *Compaq ProLiant* servers. The solution includes configuration and performance recommendations, so it can be deployed in a working environment quickly and easily.

The *Non-Disruptive Backup Solution* leverages the snapshot and cloning capabilities of the *StorageWorks HSG80 RAID* controller and uses *Compaq SANworks Enterprise Volume Manager (EVM)*, coupled with

various primary and secondary storage strategies. A key part of the *Non-Disruptive Backup Solution, EVM* is a web-based software management tool that facilitates controller-based clone and snapshot operations.

EVM lets you create, in background mode, independently addressable volumes for open system information storage. The volumes are virtual copies of active production volumes. This parallel processing enables you to redirect workloads and perform non-disruptive backups—while maintaining normal business operations. *EVM* can be used to meet business continuance requirements by minimizing application downtime. Because all backup operations are done on the clone or snapshot—rather than on the live data—the performance impact on the production system is minimized during the backup process.

III. Solution Components

- *Compaq SANworks Enterprise Volume Manager (EVM)* with scripts for third-party backup software and primary data applications
 - Third-party Backup Software: Legato NetWorker, VERITAS Backup Exec, and Computer Associates ARCserveIT (others to follow)
 - Applications: Oracle8 and Microsoft Exchange 5.5
- *Compaq StorageWorks RA8000 or ESA12000* with HSG80 and ACS 8.5 (F or S) controller firmware
- Choice of Compaq backup infrastructure:
 - Tape Library: TL891, TL895, ESL9326D
 - *Compaq Enterprise Backup Solution (EBS): EBS for Departments, EBS for Data Centers*
- SCSI or Fibre Channel (dual loop or switched fabric) interconnect to secondary storage
- *Compaq Batch Scheduler*
- Usability information: configuration, performance, and characterization data
- Windows NT backup server with Legato NetWorker, VERITAS Backup Exec, or Computer Associates ARCserveIT
- Tested on Oracle and Microsoft Exchange
- Standard Compaq 3-year warranty
- Single bundled part number for pre-configured entry-level, mainstream, or enterprise class solutions, or customized package to meet specific requirements
- Optional services, including 24x7 warranty uplift, installation and startup services, and more

IV. Designing your Solution

Overview

The *Compaq Non-Disruptive Backup Solution* makes it possible to take a copy—known as a Business Continuation Volume (BCV)—of a running Windows NT application's data and move it to another less busy machine or a dedicated machine to perform backup. This frees the original application host to quickly resume work and removes the backup overhead from that device. This process actually represents a minimally disruptive backup, since the backup operation does not conflict with the primary application.

Utilizing features incorporated within the *StorageWorks RA8000* controller firmware, the solution allows snapshot and clone technology to be controlled remotely. *Compaq SANworks Enterprise Volume Manager (EVM)* brings together this capability on any TCP/IP networked server.

The core program resides on a server located anywhere on the Internet network. The management software program can be controlled by web browsers—with username and password control—and can also be located anywhere on the Internet. The core *EVM* program communicates to its agents over TCP/IP. The communication to *Compaq StorageWorks* subsystems is via an agent running on a device that sits on the same Fibre Channel fabric as the *Compaq StorageWorks* disk subsystem. Other agents receive commands to execute jobs in application or backup servers that have backup jobs waiting to run. These servers and their applications must have tape devices available.

***EVM* configuration rules**

Complete configuration rules for an *EVM*-based storage network, the *EVM* server, and *EVM* clients are described in the *Compaq SANworks Enterprise Volume Manager* for Windows NT Installation Guide and online help.

Each *EVM* Network must have:

- 1 (and only 1) *EVM* server node. Although the *EVM* server can reside anywhere on the SAN, it is recommended that it is the same server as that which operates the backup applications software. This will make the backup job easier to manage.
- At least 1 *EVM* server or client node that is connected to an *EVM*-compatible *StorageWorks* subsystem.
- At least 1 *EVM* server or client node that has an HS Series Agent Version 2.2 installed.
- Access to a qualified web browser that supports access to port 2301.
- 1 *EVM*-compatible *StorageWorks* RAID Array 8000 (RA8000) or *Enterprise Storage Array* 12000 (ESA12000) subsystem, with SCSI-3 enabled and Command Console LUN (CCL) enabled.
- Each *EVM* network can have up to 9 *EVM* clients.
- Clone and snapshot availability are dependent on the version of Array Controller Software installed on the controller. Clone functionality is available on ACS V8.5F and V8.5S. Snapshot functionality is available on ACS V8.5S.

See ACS and solution documentation to determine operating system support and restrictions.

Supported Configurations

Typically, each *Non-Disruptive Backup Solution* configuration includes an application and backup server operating Oracle and/or Exchange databases behind Windows NT, a Fibre Channel-attached disk storage subsystem with HSG80 controller and ACS 8.5 firmware, backup application software, and an automated tape library. Since the *Non-Disruptive Backup Solution* leverages common qualified hardware and software, it can be implemented on top of existing configurations. Its support of SCSI, fibre loop and fibre switch secondary storage infrastructures allows for easy migration into SANs. The following chart provides a quick overview of compatible solutions configurations.

Compatibility Table for the Compaq Non-Disruptive Backup Solution							
OS	RAID Array	Application	Infrastructure	Backup Application	Tape Library		
Windows NT4.0 sp5	RA8000 with HSG80 and ACS 8.5 F/S	Oracle v8.05	Disk-Fibre Switch Tape-EBS Switch	Legato NetWorker v5.51	TL895		
				Legato NetWorker v5.51	ESL9326D		
			Disk-Fibre Switch Tape-Direct SCSI	Microsoft Exchange v5.5	VERITAS Backup Exec v7.3	TL895	TL891
						ESL9326D	TL895
					Computer Associates ARCserve/IT v6.61	TL895	TL891
						ESL9326D	TL895
	Disk-FCAL Tape-Direct SCSI	Microsoft Exchange v5.5	Legato NetWorker v5.51	TL895	TL891		
				ESL9326D	TL895		
			VERITAS Backup Exec v7.3	TL895	TL891		
				ESL9326D	TL895		
			Computer Associates ARCserve/IT v6.61	TL895	TL891		
				ESL9326D	TL895		
	Disk-FCAL Tape-EBS FCAL	Microsoft Exchange v5.5	VERITAS Backup Exec v7.3	TL895	TL891		
				ESL9326D	TL895		
			Computer Associates ARCserve/IT v6.61	TL895	TL891		
				ESL9326D	TL895		

V. Business Continuation Volumes (BCV)

There are two types of Business Continuation Volumes: snapshots and clones.

A BCV **clone** is created using controller-based mirroring to obtain data that is identical on both the clone and production volume. This means that the mirror copying of the data must happen before the clone can be split off. Therefore, normal operations must include an up-to-date clone of the database.

A BCV **snapshot** is created using a point-in-time copy in the controller. The snapshot volume is an empty volume where every disk-block is pointing to the original production volume. When data is written (therefore changed) to the original production volume, the changed data is copied over to the snapshot volume (copy on write). Access to the snapshot data is served by the original production disk, until the

writes have completed on the snapshot volume and the data has been migrated. The backup application sees the snapshot data as a single volume.

Supported Configurations (Controller-based)			
		Switch	Loop
8.5F	Clone	Supported	Supported
	Snapshot		
8.5S	Clone	Supported	
	Snapshot	Supported	

RAID-levels supported

Cloning uses controller-based mirroring as its engine, which automatically limits both the source (production) volume and the clone to JBOD, RAID levels 0, 1 and 0+1. Excluding RAID 5 from being used as the production volume is not as limiting as it sounds, since most likely RAID 0+1 has been chosen for performance reasons. The clone (not the original volume) consists of single disks striped together and is therefore unprotected. Given that the lifetime of a clone is only the time needed to make a backup, the probability of a disk failure is very low.

A **snapshot** does not impose any restrictions on what RAID level is used, on either the source (production) volume or the snapshot (backup) volume. The only restriction is that the snapshot volume must be the same size or larger, and it must be smaller than the 512 GB controller cache limit. This means that it is perfectly acceptable to have a RAID 0+1 environment for the production and use RAID 5 for the snapshot—which would guarantee that no disk failures occur while the backup of the snapshot is being copied to tape. EVM will only use single disks and stripe sets for the snapshot for performance and capacity reasons.

Number of Disks needed for a BCV

Both clone and snapshot require at least an equal amount of spare disks (and capacity) as the production volume.

Cloning requires that each clone disk be at least the size of the source (production) disks. Let's say, for example, that you have two RAID 0 volumes—one using three 9.1 GB disks and the other using six 4.3 GB disks. In order to clone the above one volume after the other, you would need three spare disks of at least 9.1 GB to clone the first set. Later, when the first backup is finished and the clone is removed, the second RAID set would require at least three additional 4.3 GB disks.

Snapshots require that the target LUN (RAID set) be at least the size as the source LUN. The disk requirements for EVM created snapshots are the same as for clones. Future releases of EVM will allow for fewer larger capacity disks for the snapshot volume.

It may seem that having a snapshot volume with so few disk spindles might cause a potential performance problem. Most requests caused by the backup job will be reads directed to the original production volume—not the snapshot. The exception would be the parts that already have been written to the production volume. This tends to reduce the risk that the small number of spindles would become a performance limitation. From the point of view of the backup workload, using a small number of spindles is acceptable as long as the write I/O rate on the original production volume is reasonably low.

Number of volumes supported

EVM supports a maximum of four production volumes that can be manipulated for one job. The controller also supports up to four snapshots at any given time—which means that if you already have four created, you must synchronize the creation of new ones or wait until some are deleted. There is also a restriction that you cannot make a snapshot if there is already one snapshot of that specific volume in place. Finally, you cannot make a snapshot of a snapshot.

Performance implications on the production system caused by the BCV

A **clone** must synchronize all its mirror-copies before the BCV is split off. This means that there is a performance penalty on the production volume while the clone is being synchronized (or copied). The copy process of the clone can be controlled to be either “normal” or “fast”. The difference is how much priority is given to the cloning process vs. user I/O operations. Using fast mode more than doubles the data rate of the copy process and imposes a little more overhead on production I/O operations. It is impossible to predict exactly how much this overhead will be, since it depends heavily on the competing production I/O workload. In a typical OLTP database or Exchange environment, the number of I/Os per second will be reduced 10-20% while the copy is being made. Obviously, if the OLTP workload is already saturating the controller, the addition of the clone copy will decrease performance. If not, these I/O rate differences may well go unnoticed by downstream clients of the user application.

Once the BCV clone is split off, there is no I/O performance penalty on the production volume beyond what would be expected of any single LUN. Prior to the split, mirror sets don't signal writes as complete until all copies are complete, unless controller writeback cache is on. If this is the case, then periodic cache flushes can take longer. This characteristic of any mirror set is made slightly more intrusive when you consider that during normal operations, three synchronized mirror copies are active—the intended two and the synchronized clone ready to be broken off.

A **snapshot** has no performance implications before it is created, and the actual creation is fast because no data is copied. A 512 MB controller cache size is suggested for snapshots. Smaller controller cache configurations may adversely affect snapshot performance. The recommended 512 MB cache configuration optimizes snapshot performance by avoiding accesses to the snap volume for reference to the original data by caching the "unchanged" flags from the BCVsnap.

Snapshots should be created using controller pairs instead of a single controller, even though only one controller may be used. The 512 MB cache requirement is for both controllers, even if the second controller is not used for snapshots, as controller pairs can not exist with different cache sizes. Snapshots will fail over in case of a single controller failure. Snapshots can be split across two controllers, but once created the snapshots can not move to the other controller like clones can.

After the BCV snapshot is created, there will be a performance penalty on some of the production volume writes. When data is written to the production volume, the old contents of the modified data blocks must be copied to the snapshot before the new data blocks can be written. This migration only occurs the first time a data block is written. In order to minimize the number of copy operations, the data is migrated over to the snapshot volume in 128-block transfers. After this initial migration, further writes to this 128-block area on the production volume will not cause any overhead to the original volume. The overhead for the production environment is difficult to estimate because it depends on the access pattern for the write operations of the application.

The worst case scenario is a completely random write pattern with very low locality— where every operation causes a migration of data to the snapshot volume until most of the data is migrated. This scenario is atypical.

When using a snapshot, all backup I/Os to the BCV snapshot will access the actual production volume for all unmodified data—therefore competing with the production workload. This may seem like a significant performance impact on the production volume when a snapshot is used. However, this would only be true for the extreme cases where write-intensive production applications are doing random operations at the same time the snapshot volume is being heavily accessed by backup. If this is the case, it is recommended that backup jobs run during off-peak hours when the I/O activity is low so that data migration over to the snapshot volume will be limited. In such environments, it would be reasonable to use large disks and even RAID5 for the snapshot volume as a very economical alternative without sacrificing performance.

Preliminary testing indicates a 30-35% I/O performance penalty on the production volume for 100% writes. With 50% reads, the I/O performance penalty is 15-20%. (The lower figure is without any activity on the snapshot volume and the higher is with a backup running). For a 100% read I/O load on the production volume, there is no significant I/O performance penalty to having a snapshot in place.

This storage performance difference relating to snapshots seems to filter up through the user applications. The backup rate to tape for the clone is similar to the standard online backup, whereas the backup rate for a snapshot is less because of the need for the controller to migrate modified data.

Expected performance of the BCV

A BCV **clone** is an independent volume and has no links to the original production volume. The performance seen when accessing the BCV is the same as accessing one of the original mirror sets.

Accesses to a BCV **snapshot** are going both to the original production volume (for unmodified data) and the snapshot volume for data that has been modified on the original. Therefore, performance depends on the production workload and on how much data has already migrated over to the snapshot volume. In addition, unlike the clone, the snapshot is limited to the same controller and unit offset so backup (from the backup host) and application accesses are through one controller and port.

Preliminary testing of sequential reads from a BCV snapshot indicates that with no activity on the original production volume and no data compression, it is possible to exceed the performance of a single DLT tape drive streaming during a backup.

Comparison of a BCV Clone vs. Snapshot

Characteristic	Snapshot	Clone
Persistence	Short lived—depends on availability of source volume at all times. Snapshot is lost if cache is destroyed.	Long lived—clone is available even if source volume is deleted. Clone volume is independent of cache.
Container type (of original)	Any container—RAID 5, RAID 0, RAID 1, RAID 0+1, JBOD.	Limited container—RAID 0, RAID 1 and RAID 0+1 and JBOD.
Container type and size (of BCV)	Any container—RAID 5, RAID 0, RAID 1, RAID 0+1, JBOD—size equal or bigger than original.	Limited container—RAID 0 — same number of spindles as original, size equal or bigger.
Flexibility	BCV snapshot must be on same controller and port as the source.	BCV clone is independent of the source location after being split off.
Initial construct time	Instantaneous	Depends on size of source volume (generally 40 minutes to 1 hour.)
Time to Split or Snap when container is always available	Instantaneous	Instantaneous—Start with 3-member Mirror sets. Recreate 3-member volume in 40 minutes to 1 hour.
Reconstruct source volume	Must copy snapshot to another volume to revert to source volume. Unit is available after copy is complete.	Reverse mirror operation. Clone is source of copy. Unit is available immediately during re-mirroring operation.
Impact on performance	Source volume and snapshot volume are both accessed during read and write I/O. Writes cause backend 128-block copies that limit front-end writes to 840/sec. This is made worse by split I/Os or RAID 5 read-modify-writes.	Source volume is accessed independently of clone volume for read and write I/O. Writes must complete on all 3-mirror sets unless writeback is on, in which case write cache flushes will take longer.

Planning

In planning the installation, it is imperative to understand the user application workload (Exchange or Oracle) that will ultimately benefit from removing the backup load. If this workload is very light, there may be little benefit to implementing the *Compaq Non-Disruptive Backup Solution*. If this workload is already saturated, again it will see little difference. If the workload lies between these two extremes, then the impact of removing the backup operation is mitigated by the impact of pausing the application (and in the case of the snapshot the reduced write performance).

The user application (Exchange or Oracle) is assumed to already be installed and will be migrated to the storage array. In addition, the backup application is installed and understood. Planning involves deciding on the level of redundancy (RAID level) and size of Logical Units (LUNs) required. This must be done with the understanding that any LUNs that are copied to BCVs require an equivalent amount of spare disks.

After the user application has been migrated to the RAID array, *EVM* can be installed. After the installation of *EVM*, the BCV process can be designed to fit the staffing and schedule requirements using the Backup Rates Table located in Section VII of this document.

Installation of the RAID Array 8000 or ESA12000 is assumed and installation guides for this are available at: http://www.compaq.com/products/storageworks/raidstorage/RAIDUser_docs.html.

For RAID controller planning purposes, consult:
ftp://ftp.compaq.com/pub/supportinformation/techpubs/user_reference_guides/aa-rhh6b-te.pdf.

VI. Sample Scripts

EVM requires scripts for both the database application and the backup application. *EVM* application scripts are made up of at least four files. Each *EVM* operation is divided into “Do” and “Undo” sets of commands. Required with each of these is a keyword file. Example script fragments are shown in the *Enterprise Volume Manager* and Oracle8 Best Practices documents and Best Practices for Exchange Database Management using *Compaq StorageWorks Enterprise Volume Manager*. Both documents are available at <http://www.compaq.com/products/storageworks/Storage-Management-Software/evmindex.html>.

EVM scripts are stored as ASCII files, but they are created and modified using *EVM*, a browser-based GUI. Directly editing the ASCII scripts incurs the significant risk that they will not be read by the *EVM* GUI. Not only are these scripts unique to each host, but they are sensitive to the sequence in which they are run.

Let’s say the backup host is expecting to backup a BCV named e:. If another BCV named d: has not yet been mounted to the backup host, then the new BCV will be called d: instead of e:. This will cause the backup to fail because it will be trying to backup e:, which does not exist. Similarly, if spare disks are needed in the controller to create a BCV and another BCV is still using them, then the job will fail.

Script Process

Scripts must be authored and maintained using *EVM*. The process requires the following decisions:

- Pick a name for this *EVM* script.
- Decide if the user application uses a BCV clone or snapshot.

- Decide which of the existing LUNs on that controller are to be copied to the BCV.
- Create the application commands to pause and to restart.
- Indicate the remote backup host where the BCV will be mounted and the backup will be done.
- Type in the backup host command to start the backup job of the newly mounted BCV.

Backup Scripts

These command lines are inserted into the *EVM* application wizard in order to call a specific backup application.

VERITAS Backup Exec

```
"<drive_letter> <path_name> bemcmd -o1 -j<backup_job_name>  
-w"
```

Legato Networker

```
"<drive_letter> <path_name> save -s<backup_server_name> -  
c<client_name> -g<group_name> -b<pool_name>"
```

Computer Associates ARCserveIT

```
"<drive_letter> <path_name> arcbatch mode=execute  
/h=<backup_server_name> /s=<job_name_path><job name>"
```

The user of this solution will need to be experienced with the user application (Exchange or Oracle) and the backup application. In addition, they will need to be familiar with command line scripts specific to the host operating system. This solution also assumes the pre-existence of the user application (Oracle or Exchange), backup applications, TCP/IP, and Microsoft Internet Explorer 4.0 browser.

VII. Performance and Sizing Characterization

The normal backup strategy for Windows NT, Oracle, or Exchange environment is to backup the application while it is running in an off-peak load condition. This means that the application is impacted by the work being done to do the backup. There are additional restrictions that impact performance, which are imposed on the application to support online backup. The *Compaq Non-Disruptive Backup Solution* removes those specific overheads while replacing them with others.

In the case of Microsoft Exchange 5.5, the tradeoff means the application is forcibly stopped and restarted. This can cause disruption to the client mail applications that are accessing the Exchange server at that moment. This disruption is directly related to the Microsoft application design and the APIs it provides.

In the case of Oracle8, there is an API that only momentarily forces a flush of cached transactions, which is less intrusive to the users.

When the running application goes back to normal operating mode, the backup work is typically done on another device (dedicated backup server) and the only impact is on the shared storage array. For situations where the application server CPU utilization is low, the backup application can run on the same server as the application.

Methodology

In order to quantify the difference between traditional (online hot) and non-disruptive backup, the following strategy was adopted. Tools were used to generate a repeatable workload for each application. Along with generating a repeatable workload, these tools measured client user-level performance. The choice of workload for Exchange and Oracle was governed by the factors of workload intensity and database size.

Unlike normal application benchmarks—which are meant to set new higher records—this workload was not intended to be the highest intensity the system can handle. This is because although this backup process is non or minimally disruptive, it is expected that MIS departments will continue to do backups during off-peak hours.

In addition, it is assumed that part of the reason that companies use the RA8000 is because of space requirements (besides the capability to do non-disruptive backup). Since this is the case, a significantly large database is required (well over 30 GB). The choice of workload in these benchmark tools also sets the size of the database.

The choice of baseline user application (Exchange or Oracle) workload is also critical to the performance test outcome. If too light a workload is selected, all the test alternatives might come out looking all the same. If the workload is already saturated on the application host, then there will be little difference—adding a backup workload to a saturated application workload is still saturated. Assuming the workload is somewhere between these two extremes, then some difference should be seen. These differences will be the sum of the benefit of removing the backup workload, the imposition of the changed storage environment, and the ripple effect of the pause of the application.

A workload that nearly saturates its host CPU guarantees that the additional online backup overhead would significantly impact the availability of the application to the user. Therefore, moving it off the application server would be a business advantage.

Workloads were run and baseline performance metrics were recorded by the tool. The workloads were then re-run with the normal online backup of the application running. The difference between these two is the impact of the normal backup on the user.

The workload was run again with each of the two types of BCVs in place to see what the effect of just having a BCV is to the application. Finally, each type of BCV is used to do a backup of each application using each backup application to see the impact of the *Non-Disruptive Backup Solution* methodology. The following describes the results seen in testing.

Sizing and characterization data

Performance metrics apply to two areas. The primary area is the impact to the running user application. The second area is the rate that the data can be backed up.

User Application Impact

We have to separate the two user applications, Oracle and Exchange. Exchange has an additional level of impact due to both the lack of an API that can flush cached transactions and uses a periodic cache flush mechanism. The *Compaq Non-Disruptive Backup Solution* (NDBUS) impacts the rate of backup, the ease of backup, and the availability of the application. The following chart summarizes these factors:

	Backup Performance	Backup Complexity	Application Availability
Microsoft Exchange	<p>The NDBUS improves backup performance over native backup for any of the following situations, alone or combined:</p> <ul style="list-style-type: none"> ▪ High application workload resulting in high saturation of the CPU ▪ Use of multiple streams to tape 	<ul style="list-style-type: none"> • Backup is simplified due to up front integration of primary storage, secondary storage, software, and services • The NDBUS separates the backup function from server operations, permitting dedicated staff for each. Allows staff to focus on their core competencies. 	<p>The high availability capabilities of the NDBUS will be leveraged by Microsoft with the introduction of Exchange 2000. Until that time, the NDBUS offers an advantage over manual, native backup which must be done nightly -- during which time the Exchange store is down.</p>
Oracle	<p>The NDBUS improves backup performance over native backup in any of the following situations, alone or combined:</p> <ul style="list-style-type: none"> ▪ High application workload resulting in high saturation of the CPU ▪ Use of multiple streams to tape <p>Performance increases once the BCVclone or BCVsnap process has completed.</p>	<ul style="list-style-type: none"> • The alternatives to the NDBUS are homegrown policies and procedures for data movement and data protection. The NDBUS provides established policies and procedures. 	<p>Application availability increases over native Oracle backup. Native backup is manual, and must be done nightly – during which time the Oracle database is down. The NDBUS is an automated, smoother backup operation with minimal disruption to the application</p> <p>Impact to user application performance:</p> <ul style="list-style-type: none"> ▪ Clone: during the copy operation ▪ Snapshot: during the initial backup operation

Backup Rates Table

GB/hr backup rates (for planning purposes).

Operation	Volume type	Rate
Backup	Native Hot	28-32 GB/hr
Backup	From Clone	25-35 GB/hr
Backup	From Snapshot	10-25 GB/hr
Restore	Direct to volume	30 GB/hr

We feel the difference between backup applications and tape device interconnects lie between these ranges. Changes to this might involve differences in compressibility of files, number and size of files, and directory structure. These ranges apply to Exchange or Oracle databases.

Conclusion

Until Microsoft Exchange 2000 offers APIs to allow for hot backup, Exchange's native backup might be the best choice—unless your user application has saturated its CPU. In this case, the *Compaq Non-Disruptive Backup Solution* would offer benefits by automatically offloading the backup application from your primary server. This choice would also maximize the productivity of your IT personnel by allowing them to specialize in either backup or Exchange.

With Oracle, the impact on the user application performance will be affected during the copy operation for cloning and during the initial backup operation for snapping.

If performance is your goal for a typical backup workload, cloning is a better solution because it yields improved backup performance.

IX. Questions and Answers

Q: How is this solution different from the combination of the component products?

A: This solution differs from Compaq's product offerings in that it provides a *usable*, integrated backup infrastructure that solves the two major customer problems of a shrinking backup window and backup complexity. The solution:

- Delivers tested, supported configurations specifically designed to minimize the backup window
- Includes performance, sizing, and characterization information for the most frequently used applications and backup software
- Provides validated *EVM* scripts for applications and backup software
- Offers optional value-added services that include support for third-party backup software products
- Encompasses all Compaq high-end automated tape products

Q: What is a clone?

A: A clone is a physical copy of a production volume. The clone is created as a "third mirror" of the volume. Once the clone is made, it acts as a mirror until the clone is broken off. Typically, the application will be quiesced before the clone is broken off to ensure point-in-time data integrity. If the clone is not quiesced, the clone will contain "crash consistent" data.

Q: What is a snapshot?

A: A snapshot is a virtual copy of a production volume. The snapshot is created instantly. Like the clone, the application must be quiesced prior to making the snapshot to ensure data integrity.

Q: When would I use a snapshot versus a clone?

A: Clones and snapshots can be used in similar ways. Both can be used for offline backup. However, there are some differences that should be considered prior to implementing the *Non-Disruptive Backup Solution*:

- Cloning makes a physical copy, so it offers the highest level of availability
- Snapshots are instantaneous, whereas cloning can take several minutes (depending on the size of the volume to be copied)
- Cloning does not support RAID 3/5, so snapshots are required if the environment requires these RAID levels.

Q: Is it necessary to use Compaq primary storage with this solution? Can I use storage from other vendors?

A: Only Compaq RA8000 and ESA12000 storage may be used with the *Non-Disruptive Backup Solution*. This is because they include the HSG80 controller, which is the heart of the snapshot/cloning functionality. Future Compaq products that use the HSG80 controller will be supported, but there are no plans to support storage from other vendors.

Q: Why isn't Microsoft SQL Server supported?

A: Microsoft's SQL Server, version 7.0 does not support offline backup. The next version of SQL Server will support offline backup, and the *Non-Disruptive Backup Solution* will support SQL Server once Microsoft includes this capability.

Q: How do I size my backup requirements?

A: This information can be found in the Sizing Guide, included with this Solution.

Q: How do I order the *Non-Disruptive Backup Solution*?

A: The *Non-Disruptive Backup Solution* is a combination of existing part numbers. Because customer requirements vary, there are flexible ways to order the solution:

1. Bundled Part Numbers: Three bundled part numbers for entry (#176944-001), departmental (#176945-001), and enterprise (#176946-001) environments are available.
2. Build-to-Order Part Numbers: Compaq Custom Systems will build a configuration to order, and integrate it prior to shipping. There are 3 modifiable part numbers available for this – they can be altered, per customer requirements. The CSS part numbers are: CT-NDBUS-01, CT-NDBUS-02, and CT-NDBUS-03.
3. The component parts may be ordered separately.

Learn more

For more information on the Non-Disruptive Backup Solution, please contact your Compaq representative, or visit the following website:

<http://www.compaq.com/products/storageworks/solutions/non-disruptive/wntindex.html>

© 2000 Compaq Computer Corporation

COMPAQ, the Compaq logo, and NonStop Registered in U.S. Patent and Trademark Office. ProLiant, SANworks, StorageWorks, and Tru64 UNIX trademarks of Compaq Information Technologies Group, L.P. Windows NT is a registered trademark of Microsoft Corporation in the United States and/or other countries. UNIX is a registered trademark in the US and other countries, licensed exclusively through X/Open Company, Ltd. All other product names mentioned herein may be trademarks or registered trademarks of their respective companies. Compaq Computer Corporation shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.