

April 9, 2002
16DP-0302A-WWEN

Prepared by Enterprise Storage Group
Storage Network Solutions

Compaq Computer Corporation

Contents

1.0	Management Summary	3
2.0	Solution Description.....	4
2.1	Solution Components	6
2.2	Features / Benefits.....	7
2.3	Business Value.....	8
2.4	Customer Target Market..	8
3.0	Performance Results	8
3.1	Overview.....	8
3.2	Test Description.....	9
3.3	Testing Environment.....	11
3.4	Performance Results	13
4.0	Supported Configurations...	16
4.1	Sizing.....	16
4.2	Solution Specific	
	Installation / Setup	
	Instructions	16
5.0	Operational Guidelines	20
5.1	Solution Management and	
	Maintenance	20
6.0	Compaq Storage Services	
	Offerings	20
6.1	Service, Support, and	
	Warranty	20
6.2	CarePaq Information for	
	Europe/Middle East/Africa	21
6.3	CarePaq Information for ...	
	North America.	22
7.0	Related Information	22
7.1	Related Websites	22
8.0	Appendix.....	23

StorageWorks Rapid Restore and Non-Disruptive Backup for Telecom Billing - SchlumbergerSema

Abstract: The StorageWorks Rapid Restore and Non-Disruptive Backup solution provides the ability to rapidly restore large quantities of critical data for SemaVision Business Support and Control System (BSCS) users across a unified Oracle database.

The solution employs the Enterprise Virtual Array technology to create copies of a defined data set, called Snapclones or Business Continuance Volumes (BCVs). You can use a BCV to facilitate an ultra-fast recovery of a database in the event of a failure. Additionally, it can be used to perform offline backups, data mining, and testing of new features, with minimal impact on the production server.

To assist the user, this guide contains the necessary procedures for using snapclones to backup and restore the application's key records. It also includes a reference platform and provides the results of extensive testing of the overhead and benefits of implementing snapclones.

This guide discusses the following:

- Advantages and considerations for creating and managing physical copies of a virtual volume in an Oracle environment.
- How to configure the BSCS Oracle database and deploy the Compaq StorageWorks Enterprise Virtual Array as part of the overall Customer Care and Billing application.
- How to recover a database in minutes, not hours!

Help us improve our technical communication. Let us know what you think about the technical information in this document. Your feedback is valuable and will help us structure future communications. Please send your comments to: ESG-SNS.SchlumbergerSema@compaq.com

Notice

The information in this publication is subject to change

Solutions Guide prepared by Enterprise Storage Group

16DP-0302A-WWEN First Edition ©2002 Compaq Information Technologies Group, L.P.

Compaq, AlphaServer, *SANworks*, StorageWorks, CarePaq, and Tru64 are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries.

SchlumbergerSema and SemaVision BSCS are registered trademarks of Schlumberger Limited.

Oracle is a registered trademark of Oracle Corporation.

Sequel and SQL are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Legato and Legato NetWorker are registered trademarks of Legato Systems, Inc.

All other product names mentioned herein may be trademarks of their respective companies.

All product names mentioned herein may be copyrighted by their respective companies

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

1.0 Management Summary

Telecommunications Service Providers of all sizes have come to depend on SchlumbergerSema for mission-critical Customer Care and Billing applications and worldwide commitment. SchlumbergerSema addresses the need of Telcos to add new services based on emerging technologies. Each expansion of the service mix adds data and may change business processes. To stay competitive, service providers need a flexible, accessible, available, and high-performance Customer Care and Billing (BSCS) application. The StorageWorks Rapid Restore and Non-Disruptive Backup solution was developed to optimize the availability of data by utilizing tools to make this data safe and readily available.

Compaq and SchlumbergerSema closely cooperated to define, test, and document this exciting solution to address these needs. The Compaq *StorageWorks* products used in this solution optimize the capabilities of a Storage Area Network (SAN) for the SchlumbergerSema Customer Care and Billing application. This *StorageWorks*- and *SANworks*- based solution is a flexible storage environment designed for rapid custom implementation and high performance. This solution employs the Compaq Enterprise Virtual Array management software and a SAN configuration to optimize overall system responsiveness and availability by providing for a minimally disruptive backup and rapid recovery of databases that might have become corrupt. To gain this competitive advantage, system administrators can create point-in-time copies, called Snapclones or Business Continuance Volumes (BCVs), of the database. They can then use these physical on-disk copies for backup or restore without compromising the BSCS databases.

The test results demonstrate that the StorageWorks Rapid Restore and Non-Disruptive Backup solution, built around SemaVision BSCS on an AlphaServer and Tru64 UNIX system platform and a StorageWorks SAN, provides the administrator with all the necessary tools to restore large quantities of critical data in minutes - not the usual hours - to restore from tape. Additionally, it shows that the user can backup the database with minimal interruption to the production server.

2.0 Solution Description

The StorageWorks Rapid Restore and Non-Disruptive Backup Solution is built around the SemaVision Customer Care and Billing application that is a major element within the SchlumbergerSema Business Support and Control Systems (BSCS) architecture. A Compaq AlphaServer ES45 and the TRU64 Unix operating system is the leading platform for deployment of SchlumbergerSema BSCS.

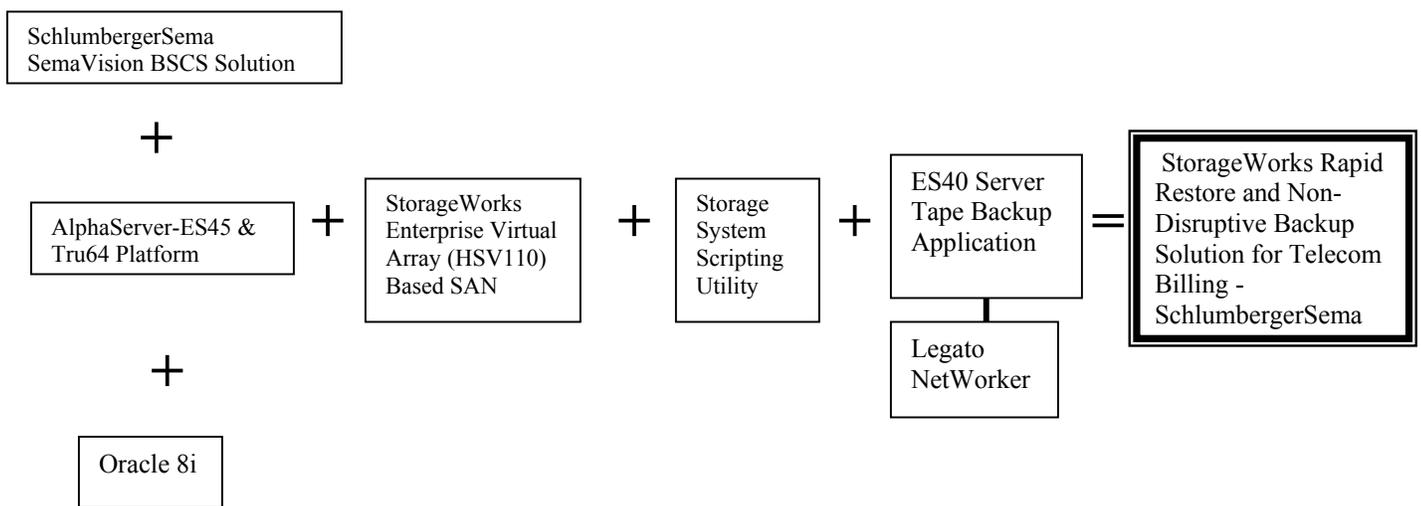


Figure 1 – Solution Building Blocks

The StorageWorks Enterprise Virtual Array (EVA) controller (HSV110) is a high performance fibre channel controller that provides the tool, called snapclones, to create physical copies of its virtual volumes. The virtually instantaneous snapclone operation performs a full background copy from the original virtual disk. At creation time the snapclone is available as a traditional snapshot, but once the background process is finished, it becomes a clone also known as a Business Continuation Volume (BCV). This BCV is ready to be pressed into service in a matter of a minute or two. It can be mounted on a backup server and used as the source for a full backup to tape without using the production system's resources. What's more, a BCV can be used as a rapid recovery source should corruption or accidental deletion occurs to the production database.

The Storage System Scripting Utility (SSSU) provides an interface that allows the administrator to issue simple commands or scripts to initiate and control the snapclone process. A shell script can be written to include the three basic steps; quiesce the production database, create the BCV, and resume database activity to create the BCV (See section 4.2).

The Legato Networker tape backup utility and two Storage DL891 tape libraries (each with 2 DLT drives) were used to backup the databases to tape. The time involved for this backup procedure provided us the reference time to assess the speed advantage when using the BCV to perform the same operation. The Legato utility was used to conduct these tests, but any supported tape backup utility could also be used.

The goal of the test parameters was to represent a service provider with 10-million subscribers. The server platform was configured with a Storage Area Network (SAN) of sufficient capacity to simulate a robust BSCS database. The performance of the server platform was tested with respect to call detail record (CDR) I/O rates.

This Solution Guide contains solution-specific information to provide the administrator with the necessary procedures for creating and managing snapclones for backup and restore applications for the BSCS database files, and for making informed management decisions using pre-developed performance data.

It also includes:

- The results of extensive testing highlighting the benefits of implementing snapclones.
- Considerations for creating and managing physical copies of a virtual volume in an Oracle environment.
- Performance results for a 10-million-subscriber database and the impact on the production server during snapclone backup and recovery activity.
- Instructions to deploy the BSCS Oracle database and the Compaq StorageWorks Enterprise Virtual Array as part SchlumbergerSema's overall Customer Care and Billing application.

2.1 Solution Components

SchlumbergerSema Telecom Solution

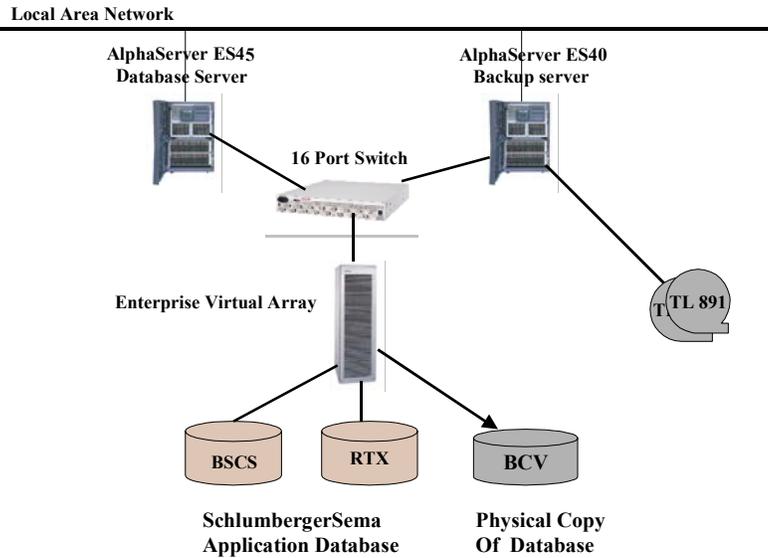


Figure 2
Figure 2

The main solution elements (Figure 2) are described below. A more complete BOM follows in section 3.3.

Hardware:

- Database Server

Machine:	Alpha Server ES45
OS-Version:	Tru64 UNIX V5.1A
# of CPU:	4
CPU-type:	EV6.8 (21264C)
CPU-speed:	1000 MHz
RAM:	32GB
File system:	AdvFs

- Backup Server

Machine:	Alpha Server ES40
OS-Version:	Tru64 UNIX V5.1A
# of CPU:	4
CPU-type:	EV6.8AL (21264B)

CPU-speed	833MHz
RAM	8GB
File system	AdvFs

System Software:

- Tru64 Unix V5.1A – Operating System Base License
- SAN Fibre Channel Dual Controller Base License
- SAN Enterprise Operating System Platform Kit for Tru64 UNIX
- SAN Enterprise Virtual Array Snapshot License

Application Software:

- SchlumbergerSema Business Support and Control Systems (BSCS), - SemaVision V 6.0
- Oracle V 8.1.7.2.1 - Database Server
- Legato Networker, V6.0.1 - Tape Backup Utility

Storage Area Network:

- Enterprise Virtual Storage Arrays - Model 2C12D with 6.4TB of disk storage
- Switched Fabric Based SAN
- *SANworks* Appliance

2.2 Features / Benefits

The solution provides the following benefits to SchlumbergerSema's customers.

- Best practices to maximize database availability during recovery operations.
- Best practices to maximize database availability by performing non-disruptive backup operations without utilizing the production server resources.
- Documented proof points of the rapid restore capability for BCSC databases can be used for business operations planning.
- A Compaq SAN offers greatly improved investment protection by providing available, flexible, and interoperable storage components.
- The ability to preconfigure SANs can shorten implementation time and therefore aid in the rapid deployment of new installations.

2.3 Business Value

The value of this solution becomes evident when you consider the following:

- The demands placed on the Telco carriers for the highest level of database availability are addressed by this non-disruptive backup and rapid restore solution.
- The solution incorporates the latest in SAN technology and tools to provide scalability of system resources to handle the explosive growth brought on by new integrated applications and database requirements.
- Larger and more efficient storage platforms are available to address increases in regulatory and business requirements for simultaneous access to customer profiles without loss of application performance levels.

2.4 Customer Target Market

High-growth telecommunications and service providers with 10 thousand to 20 million subscribers, such as:

- Operators and Service Providers implementing 2nd and 3rd generation technology
- Competitive local exchange carriers (CLECs)
- Incumbent local exchange carriers (ILECs)
- Application service providers (ASPs)
- Integrated service providers (ISPs)
- Customers with storage requirements greater than 1 TB
- Customers whose backup windows are rapidly shrinking because of increasing demand on system resources
- Customers whose business requires reduced database recovery time with on-disk based backup and replication

3.0 Performance Results

3.1 Overview

The fibre channel Enterprise controller (HSV110) as a key element of the Enterprise Virtual Array manages the cloning operation without using database server resources and provides no noticeable interruption to the application's users.

BCVs can be created without impacting the normal daily processing of billing data. To ensure a consistent view of the data, the application must be briefly quiesced. Tests show the time for BCV creation takes only minutes rather than the hours required for normal tape-based backup. Furthermore, the administrator can create the essential tape archives from the BCV on an independent server, completely eliminating the backup load on the database server.

These time savings that are made possible by these array-based operations benefit the users every day, but the most important time savings occur with a database failure that formally requires recovery from backup media. In conventional tape-based operations, many hours of recovery time are required to restore a large database such as BSCS from tape to disk. With a BCV copy of the database available, simply mounting the BCV in place of the original volume and restarting the database takes only a few minutes.

3.2 Test Description

Business Support and Control Systems

The tests were designed for a 10 million-subscriber environment but the solution can be scaled up or down. Although the tests created a complete BSCS database, the performance was observed while inserting rated records into the Rated Transaction Database (RTX) that holds the rated usage records and generated bills. This test was chosen because of its I/O intensive nature

The test database included 10-million subscribers with one contract per subscriber and 150 calls per month per subscriber/contract.

The test data set was 850,000 RTX (~247MB) records per Rated Load Handler (RLH). It consisted of 17 files, each containing 50,000 records. All the records had a unique customer_id/contract_id combination to insure tests were conducted on I/O to the array, not simply exercising cache memory.

The RTX dataset is placed in a tablespace distributed over the /u03 mount point. Before each dataset test, the test database was restored to consistent starting states. In order to represent a true production system, the RTX's data blocks were preallocated. In order to minimize disk contention the datasets were spread evenly across all available files.

Snapclone Tests

The purpose of snapclone tests was to identify the amount of time required to complete the snapclone synchronization and the performance impact on the active RLH session. The following three operational scenarios were tested:

Cold Snapclone - The database is shutdown during this clone operation so the test provides an understanding of the amount of time needed to synchronize the snapclone volume without database activity. Various tests were performed to understand the impact of the underlying HSV110 setup on the rate at which the snapclone synchronization completes. Although available to the backup server, no activity was performed on the snapclone during these tests.

Warm Snapclone - The database is open, but no RLH processes are active during this test case. This scenario provides information and a timeframe for production environments that have a small inactive period for this type of operation.

Hot Snapclone - The database is open and eight RLH processes are actively inserting Customer Detail Records (CDRs) into the database. This scenario provides the timeframe and performance impact of clone operations on the active database.

Note: Although the snapclone was not completely synchronized with the source volume, it was available for use immediately after the creation. There was no delay in utilizing the snapclone as a reporting database or as a source for the backup.

Tape Backup Tests

The purpose of the tape backup tests was to ascertain the impact of backing up the copy of the database to tape media. Two tests were performed:

- The first test was performed with the snapclone already completely synchronized, and
- The second was performed with a backup of a snapclone still synchronizing with the source database volume. Both backup tests utilized the Legato NetWorker utilizing four tape drives without compression enabled.

Note: In both cases, the snapclone database is a usable point-in-time snapshot of the active database. Therefore the snapclone can serve as the instant recovery vehicle, and the tape backup can function as the disaster recovery vehicle.

Cold backup

This cold backup test measured the time to backup the BCV copy of the database after it had completed physical synchronization with the source database. The intent was to determine a baseline for the time needed to write the clone copy of the database to tape utilizing four tape cartridges without compression.

Hot backup

This hot backup test measured the impact on both the backup time and the Rate Load Handler (RLH) throughput when performing the tape backup operation while, at the same time, creating a snapclone of the source database. The backup source is a BCV copy of the database, which is immediately made available (mounted) to the backup server. As is always the case, the snapclone is immediately available for use, regardless of whether the volume synchronization is complete.

3.3 Testing Environment

The following section describes the SchlumbergerSema Backup and Restore Solution and identifies the hardware, software, and firmware components. Configuration rules are explained where applicable.

Bill of Materials

Table 1 below lists a bill of materials that includes the products used in the solution.

<u>Item</u>	<u>Item Description</u>
1	SchlumbergerSema SemaVision BSCS V6.0
2	AlphaServer ES45 Production Server
3	AlphaServer ES40 Backup Server
4	Tru64 Unix Base Operating System License
5	Oracle 8i Enterprise Edition Release v 8.1.7.2.1
6	Enterprise Model 2C12D Virtual Storage Array
7	SAN VCS FC Dual (HSV110 Controller) Base License.
8	SAN Enterprise Sol (Solution) Platform Kit for TRU64 UNIX*
9	SAN VCS FC SNP (Snapshot) License
10	FC Hard Disk Drives, 36.4GB, 10K RPM
11	Fibre Channel HBA (switched fabric), KGPSA-BC
12	Fibre Channel Storage SAN Switch, 16 port
13	FC Optical Short Wave GBIC
14	Fibre Channel Optical Cable
15	TL891 Tape Library
16	Legato NetWorker Tape backup Utility
17	SAN Appliance

Table 1, Bill of Materials for SchlumbergerSema Backup and Restore Solution

* Includes the Storage Systems Scripting Utility.

(For Current Part Numbers, Versions, and Prices, contact your Sales Representative.)

UNIX Servers

The actual tests were conducted utilizing the AlphaServer ES45 running Tru64 UNIX V5.1A as described in the following sections.

ES45 – Database Server

Machine: Alpha Server ES45
OS-Version: Tru64 UNIX V5.1A
of CPU: 4
CPU-type: EV6.8 (21264C)
CPU-speed: 1000 MHz
RAM: 32GB
File system: AdvFs

ES40 – Backup Server

Machine: Alpha Server ES40
OS-Version: Tru64 UNIX V5.1A
of CPU: 4
CPU-type: EV6.8AL (21264B)
CPU-speed: 833 MHz
RAM: 8GB
File system: AdvFs

Both systems had 4-fiber channel I/O controllers distributed over the four PCI buses. Both the servers and all four ports of the HSV110 Enterprise Storage Array controllers were connected to a single 16-port fiber switch. The ES40 had two KZPSA adapters to allow the connection of two mini tape libraries (DLT892).

Enterprise Virtual Array

All testing was conducted with the Model # 2C12D, which is the 12-shelf version of the EVA with dual HSV110 controllers configured as follows:

A single disk group was configured using 167 36.4GB fibre channel disks. To support various tests, 8 virtual disks were created and presented to both the database server (ES45) and the backup server (ES40). The snapclone for the RTX database was presented to the backup server.

Database

Version: Oracle 8i Enterprise Edition Release 8.1.7.2.1
Parameters: See Appendix
Layout: 2-databases, BSCS and RTX.

- BSCS: Table DOCUMENT_ALL was in the BSCS database. The database consisted of approximately 130 data files each with 4GB, and a System Global Area (SGA) size of 2.1GB. The 4 redo-logs were 3GB each. There were 80 rollback segments. The size of the database was 650GB with 295GB free.
- RTX: The database consisted of about 70 data files, each having a size of 4GB, and a SGA size of 3.3GB. The 4 redo-logs were 400MB each and 32 rollback segments turned out to be sufficient. The size of the database was 294GB with 110GB free. The CDRs for the RLH test were inserted into a single clustered table called RTX_ST_001.

Direct I/O: Both databases worked in direct I/O mode where Oracle handled the file system like raw devices. This feature of Tru64 UNIX allows the system to achieve the performance of “raw” device mode but with the management ease of a file system based operation.

BSCS Application

BSCS Binaries: Date of make: Nov 12 2000
Release: V6.0 (bscs600_max01)
Level: INSENTIVE
Uname: OSF1 malta version T5.1, build 585 alpha
Oracle Version: V8.1.7.2.1

BSCS Setup:

Customer: 10 Million
Contracts: 1 contract per customer
Services: 10 services per contract
Rate cycles: 32
Bill cycle Groups: 10
Bill cycles: 10 with 1 million subscribers each

3.4 Performance Results

The 3 performance tests clearly show that BSCS V6.00 can be configured to run and perform well on a system the size of the ES45 (4 CPUs, 32GB RAM) and the 2C12D Enterprise Virtual Array. Rates of 4.9 million – 6 million RTX inserts per hour were observed during the various tests that were performed.

Note: The highest results obtained in these tests indicate that the test configuration would support insertion of 45 million Call Detail Records (CDRs) in 7.5 hours.

Cold Snapclone

The database was shutdown during this snapclone operation and testing provided an understanding of the amount of time needed to synchronize the snapclone volume without database activity.

- Test 1 used a single 800GB volume for the entire RTX database. **The snapclone synchronization was completed in 7.5 hours, which provides a synchronization rate of 1.77GB per minute.**
- Test 2 was conducted with the 800GB file domain composed of 2 – 400GB storage volumes on the HSV110 array. This test provided the impact of the number of volumes on the rate of

physically completing a snapclone. **The snapclone synchronization completed in 4 hours and 59 minutes or a rate of 2.6GB per hour.**

- Further tests revealed that the RTX database could be accommodated on a smaller volume, so the database was moved to a single 350GB volume. The time to physically synchronize this volume was measured without activity on the source database/volume. **The snapclone synchronization completed in 3 hours and 21 minutes or a rate of 1.7GB per minute.**

NOTE: If the amount of time allotted for the snapclone process to complete is limited, then consider using multiple AdvF domains. It is also important to size the virtual data volumes to closely fit the actual data, avoiding the inclusion of free space as part of a snapclone.

Warm Snapclone

The database was open, but RLH processes were inactive. This provided information and timeframe for production environments that have a small inactive period for this type of operation.

The open database test without activity produced the identical results to the cold snapclone tests described above.

Hot Snapclone

The database was open and 8 RLH processes were actively inserting CDR records into the database while the snapclone process completed synchronization.

Performance was reduced from 6.0 to 4.9 million inserts per hour, or by approximately 18.3% when compared to the best warm cache test runs without volume synchronization in process.

Note: Best practices would suggest that BCV synchronizations be scheduled during periods of low system activity, such as the middle of the night. This will significantly reduce the impact on production operations.

Tape Backup Tests

Cold Backup

The cold backup test measured the time to backup the BCV copy of the database after it had completed physical synchronization with the source database. The intent was to determine the baseline time required to write the BCV copy of the database to tape, utilizing all four tape drives with compression off.

The four simultaneous tape backup streams completed the task as follows:

<u>Backup Stream</u>	<u>Files saved</u>	<u>Gigabytes Saved</u>	<u>Duration (h:m:s)</u>
Backup Stream 1	24	70GB	2:13:23
Backup Stream 2	24	70GB	2:14:02
Backup Stream 3	24	71GB	2:18:02
Backup Stream 4	24	83GB	2:31:12
Total	96	294GB	2:31:12 or <i>271 Minutes</i>

This test concluded that the times were essentially the same across all 4 streams.

Hot Backup

The cold backup test measured the impact on the time to backup and the RLH throughput while synchronization of the snapclone was active. The backup source was a point in time snapclone copy of the database, which was immediately available and mounted to the backup server.

<u>Backup Stream</u>	<u>Files Saved</u>	<u>Gigabytes Saved</u>	<u>Duration (h:m:s)</u>
Backup Stream 1	24	70GB	2:13:49
Backup Stream 2	24	70GB	2:14:15
Backup Stream 3	24	71GB	2:23:10
Backup Stream 4	24	83GB	2:32:46
Total	96	294GB	2:32:46 or <i>272 Minutes</i>

These results demonstrate no significant impact (< .5%) on the production system of a tape backup made from a BCV on a backup system.

4.0 Supported Configurations

You can apply this Rapid Restore solution using snapclones for point-in-time BCVs of your important database files to any SchlumbergerSema BCSC configuration that uses Enterprise Virtual Array. SchlumbergerSema and Compaq support professionals can help you define the exact configuration required for your business. See Section 5 for details of available service programs.

This SchlumbergerSema solution requires that you use the *Storage System Scripting Utility* (SSSU) software to initiate the snapclone process, plus additional command scripts (included in section 4.2) to perform the sequence of steps needed to quiesce the production database, create the BCV, and resume database activity.

4.1 Sizing

The sizing of the customer's host system, and the storage requirements for the database, should be determined by SchlumbergerSema and Compaq support personnel in order to best suit the business volume of the operation. Storage space requirements will include the additional snapclone capacity required for each database volume selected for replication.

The administrator should also carefully analyze the need for more than one BCV set of the production databases. There will be a period each day when the previous BCV set is deleted but the new set is not fully resynchronized. With an Enterprise Array controller, the BCV is available for use even before it is fully synchronized. However, for maximum security of the data, keep the old BCV intact, until after the new one is completed.

In addition to the actual storage space required to fit the data and BCVs, sufficient available free space would enable the controller to work most efficiently. It is generally recommended that an EVA be configured to have about 20 percent excess capacity.

4.2 Solution Specific Installation / Setup Instructions

The process used to create the BCV of the online Oracle database in the test configuration is explained below. Although the referenced file names and host names are specific to this test configuration, these procedures illustrate the complete process.

- **Quiesce the Production Database:**

From an account on the database server that can access the database with the necessary database administrator (dba) privileges, execute the following Structured Query Language (SQL) script, `set_tbspac_backup.sql` to place all tablespaces in backup mode to ensure that the database tablespaces are in a consistent state and write activity to the database volumes has stopped.

During the test, this step took approximately 20 seconds and the specific scripts performed from the Oracle account were as follows:

```
#sqlplus system/manager@rtxdb2
sqlplus> @set_tbspac_backup.sql
where set_tbspac_backup.sql was as follows:
alter tablespace SYSTEM begin backup;
alter tablespace TEMP begin backup;
alter tablespace USERS begin backup;
alter tablespace RBS1 begin backup;
alter tablespace DATA begin backup;
exit;
```

- **Create the BCV**

Create the snapclone of the volume by using the SSSU utility and the following snapclone script, `mk_clone`. This script took an average of about 30 seconds to complete and must be executed from the root account:

```
#sssu < mk_clone
Where the mk_clone is as follows:
select manager eslsanapp01 username=administrator
password=Pageslsan
select cell SEMA_HSV
add copy cln_u03_16k_d31
storage="\VirtualDisks\Raid5_U03_16k_2\ACTIVE"
add lun 31 storage="\Virtual Disks\cln_u03_16k_d31\ACTIVE"
host="\Hosts\SEMA_ES40"
exit
```

Note: Although you can perform this operation using the SAN appliance interface scripts provide a better mechanism for repetitive scheduled operations.

- **Resume Database Activity**

Once the snapclone operation has been initiated, the tablespaces can be taken out of backup mode and database write activity can resume by executing the SQL script, `end_tbspac_backup.sql` from an account with dba privileges as follows:

In our test, this script took 9 seconds to complete.

```
sqlplus> @end_tbspac_backup.sql
Where end_tbspac_backup.sql is as follows:
alter tablespace SYSTEM end backup;
alter tablespace TEMP end backup;
alter tablespace USERS end backup;
alter tablespace RBS1 end backup;
alter tablespace DATA end backup;
exit;
```

- **Backup to Tape**

On the backup server (ES40), execute the script, `mk_mount_advfs_domain` to find the new snapclone device and enable the volume to be mounted as an AdvFS file system. Each time the Enterprise Virtual Array creates a new virtual disk, and when it is presented to a Tru64 host LUN, it is assigned a unique Worldwide ID (WWID) in accordance with Fibre Channel specifications. Tru64 sees this as a new SCSI device and assigns a new **'dsk*'** identity. The new identity must be associated with the existing dsk identity that points to the AdvFS domain defined for the expected mount point on the backup server.

This backup procedure can be performed immediately after the snapclone step has been completed.

The following script locates the newly created snapclone that was assigned to LUN 31 (reference the above script to *create the snapclone*). The hardware manager program "hwmgr" scans for the new LUN 31*. The disk special file manager "dsfmgr" is used to exchange the current "dsk" location to dsk31. The script then mounts the volume at /clone_u03.

```
mk_mount_advfs_domain
#!/usr/bin/csh
#
cd /dev/rdisk
hwmgr -scan scsi
file dsk*a | egrep -e 'LUN #31' | awk -F: '{print $1}' >
/tmp/d31
dsfmgr -e `sed 's/a/ /g' /tmp/d31` dsk31
rm /tmp/d31
mount -t advfs backup_u03#u03_16k /clone_u03
echo "done"
```

You can backup the snapclone copy of the database to tape using any tape backup software.

Optional BCV Usage

If the BCV is to be used only as a backup mechanism nothing further must be done.

However if you wish to open the snapclone database on the backup server, set the following items:

1. Copy the instance initialization file (initRTXDB2.ora).
2. Satisfy the mount points for the Oracle database files by mounting to the same name or by using a symbolic link.
3. Validate the oracle dump destination. In this instance, the Oracle disk resides in a different path, so you must create the proper directories on the backup server (audit_file_dest, background_dump_dest, etc.).
4. Since this database is operating in archive log mode, set the Oracle initialization parameter log_archive_dest to a valid location.

5.0 Operational Guidelines

5.1 Solution Management and Maintenance

The same storage space can be used repeatedly for the snapclones on a daily scheduled basis. When deleting snapclones, the time for the script to execute was 13.7 seconds. However, the snapclone name becomes free only after the reclaim space operation has been completed on the array. In our case it took just over 7 minutes to reclaim the 800GB of space in the background and make the snapclone name available for the next snapclone request. Keep this timing in mind when using scripts.

Applications that are bandwidth-sensitive may benefit by creating an AdvFS file domain from multiple volumes. You can assign one volume to path A and one to path B on the HSV110 storage array when presenting the LUNs to the database server. This will help balance the load over multiple paths to the AdvFS domain and allow striping of database files if desired.

Note: Multi-volume domains must be unmounted before creating a snapclone. It is not sufficient to simply pause Oracle; the entire domain must be unmounted with the UNIX “UMOUNT” command.

Configuring a single disk group, consisting of 167 disks behind the HSV110 controller, provides low enough latency and high enough bandwidth for the number of RLH processes in use during this test. If the number of RLH application streams increases, striping database files may be desired.

6.0 Compaq Storage Services Offerings

6.1 Service, Support, and Warranty

Software Product Services

- Stand-alone telephone support
- Rights to new license version
- Media and documentation updates

Hardware Product Services

- Installation services
- On-site maintenance (includes warranty support)

- Response time upgrades during the warranty period
- Post-warranty coverage
- RAID setup and performance consulting via statement of work

For additional hardware installation and maintenance information, please contact your local Customer Services Department.

Warranty Response Options

- Response Time - From next business day to same day 4-hours
- Coverage Time - From 5 days x 9 hours - to- 7 days x 24 hours
- Duration of Coverage - Minimum period of 1 year or multi-years

6.2 CarePaq Information for Europe/Middle East/Africa

CarePaqs exist for items of *StorageWorks* hardware typically attached to Compaq *alphaserver*[™] servers.

Additional CarePaq's are being developed to cover a broader range of *StorageWorks* hardware and software. Ask your local CarePaq Business Manager for the latest details.

CarePaq Priority Services Information for Europe/Middle East/Africa

CarePaq[™] Priority Services brings a totally new approach that removes the complexity from evaluating, planning and purchasing services on Compaq Alpha and Storage platforms. CarePaq Priority Services are easy to purchase and easy to use and cover all external storage products not covered under Server based CarePaqs.

Standard Offerings:

- CarePaq Priority Services provide an integrated hardware and software support that ensure consistent coverage across storage configurations.
- CarePaq Priority Service offers a rapid four-hour response with guaranteed work-through to completion for three years.
- CarePaq Priority Services have a standard call window of Monday to Friday 08:00 - 17:00 with uplifts to full 24x7 over three years with a four-hour response time and work through to completion.
- CarePaq Priority Services also provide software support, with a two-hour response time for telephone access to highly trained support personnel.

For more information on how to order CarePaq Priority Services in EMEA please contact your local Certified Integration Partner or Authorized Reseller.

CarePaq is defined as an upgrade to the product warranty attribute, available for a specific duration and hours of coverage.

CarePaq is not available for less than the product's warranty duration.

CarePaq is available for sale anytime during the warranty period for most products, but the commencement date will be the same as the Warranty Start Date (delivery date to end user customer). Proof of purchase may be required.

CarePaq services are prepaid.

Customized Services are also available. Contact your local Global Services Department.

6.3 CarePaq Information for North America.

CarePaq Priority Services Information for North America.

CarePaq/PSP is defined as an upgrade to the product warranty attribute, available for a specific duration and hours of coverage.

CarePaq/PSP is not available for less than the product's warranty duration.

CarePaq/PSP is available for sale anytime during the warranty period for most products, but the commencement date will be the same as the warranty start date (delivery date to end-user customer). Proof of purchase may be required.

CarePaq /PSP services are prepaid.

For additional USA CarePaq (hardware & software) information, as well as orderable part numbers, please refer to the URL listed below:

<http://www.compaq.com/services/CarePaq/index.html>

7.0 Related Information

7.1 Related Websites

The following website provide additional Information about SchlumbergerSema products and solutions

<http://www.slb.com/Hub/Docs/SchlumbergerSema/>

The following website more information on the SchlumbergerSema Customer Care and Billing - BCSC Solution:

http://www.schlumbergersema.com/prod_serv/intcustcare.htm

The following website provide additional Information about Compaq Telecom products and solutions

<http://telecom.compaq.com/>

The following websites provide additional Information about Compaq *StorageWorks* SAN products and solutions.

<http://www.compaq.com/products/storageworks/SAN/index.html>

<http://www.compaq.com/products/storageworks/san/entry/index.html>

The following website provide additional Information about Compaq AlphaServer products and solutions.

<http://www.compaq.com/alphaserver/>

The following websites provide additional Information about Compaq TRU64 Unix products and solutions.

http://www.compaq.com/products/quickspecs/10513_div/10513_div.PDF

8.0 Appendix

SchlumbergerSema-a Schlumberger Limited Company.

“SchlumbergerSema is a leading information technology services company that provides a unique combination of domain expertise and global capabilities delivered on a local basis. Together with Schlumberger Oilfield Services, it is one of two business segments of Schlumberger Limited, a global technology services company. Headquartered in New York City, SchlumbergerSema was formed in April, 2001, when Schlumberger Limited acquired Sema Plc., and combined it with its former Test & Transaction business and other recent acquisitions. Annual combined revenue of the organizations that now comprise SchlumbergerSema exceeded \$4 billion in 2000.

With over 30,000 people in more than 130 countries, SchlumbergerSema has a global network of local experts that provide best-in-class services and products. Placing strong emphasis on customer satisfaction, SchlumbergerSema offices are strategically located to ensure that the company's staff of experts is always easily available to provide customers with the highest level of support and service.”

The above quote is taken directly from the following SchlumbergerSema Web site:

<http://www.slb.com/Hub/Docs/SchlumbergerSema/>

Definitions

Business Continuance Volume (BCV): A storage volume created as a copy of a production data volume to be used for operations, such as backup and rapid recovery. Often used interchangeably with clone.

Snapclone: A process of creating an instant point-in-time independent copy of a virtual volume. A snapclone is immediately available for use, similar to a traditional snapshot while in the background a complete synchronized copy of the point-in-time data is made an independent virtual volume

Snapshot: A process of creating an instant point-in-time copy of a virtual volume in an array. A snapshot does not necessarily become a complete independent copy, as only data that changes on the original source volume is actually copied during the lifetime of a snapshot.

BSCS: “SemaVision BSCS (Business Support and Control System) is a set of applications and databases used to provide an end-to-end solution for telecom vendors. SemaVision BSCS provides convergent billing and supports multiple services and technologies. It is a carefully designed set of closely integrated components, working on a unified database.”

The above quote is taken directly from the following SchlumbergerSema Web site:

http://www.schlumbergersema.com/prod_serv/intcustcare.htm

RLH: RLH is the Rate Load Handler, which is responsible for inserting the rated UDR records into the RTX database.

RTX: RTX is the Rated Transaction Database-that holds the rated usage records and generated bills.

Oracle Parameter in init.ora as Used in the Testing:

Init.ora of the BSCS Database:

db_writer_processes	= 5
compatible	= 8.1.6
row_locking	= ALWAYS
db_block_size	= 8192
db_files	= 1022
db_name	= BMFB1
rollback_segments	= (RBS101, ... ,RBS180)
optimizer_mode	= CHOOSE
job_queue_interval	= 60
job_queue_processes	= 1
recovery_parallelism	= 12
db_block_buffers	= 200000
_spin_count	= 8192

```
db_block_lru_latches           = 10
disk_asynch_io                 = TRUE
db_file_multiblock_read_count  = 16
shared_pool_size               = 1048576000
shared_pool_reserved_size     = 16384000
log_checkpoint_interval        = 500000
log_checkpoint_timeout         = 3600
log_buffer                     = 8192000
max_rollback_segments         = 100
transactions_per_rollback_segment = 1
open_cursors                   = 4096
cursor_space_for_time          = TRUE
partition_view_enabled        = FALSE
serial_reuse                   = ALL
hash_join_enabled              = TRUE
sort_area_size                 = 8192000
sort_area_retained_size       = 8192000
sql_trace                      = TRUE
timed_statistics               = TRUE
distributed_transactions       = 300
open_links                     = 50
processes                      = 300
transactions                   = 300
dml_locks                      = 2000
enqueue_resources              = 600
optimizer_percent_parallel     = 0
parallel_min_servers           = 2
parallel_max_servers           = 64
```

Init.ora of the RTX Database:

```
db_writer_processes           = 5
db_block_lru_latches          = 10
compatible                     = 8.1.6
row_locking                   = ALWAYS
db_block_size                  = 16384
db_files                       = 1500
db_name                       = RTXDB1
db_file_multiblock_read_count  = 8
db_block_buffers               = 200000
shared_pool_size               = 52428800
log_checkpoint_interval        = 10000
processes                      = 50
parallel_min_servers           = 5
parallel_max_servers           = 40
log_buffer                     = 8192000
timed_statistics               = TRUE
```

```
sql_trace = TRUE
max_dump_file_size = 1024000
max_rollback_segments = 100
PARTITION_VIEW_ENABLED = TRUE
sort_area_size = 8192000
sort_area_retained_size = 8192000
rollback_segments = (RBS101, ... ,RBS132)

log_archive_dest = /work/oradata/RTXDB2/arch/arch_
job_queue_interval = 60
job_queue_processes = 1

log_archive_format = "RTXDB2_%t_%S.arc"
_log_archive_buffers = 5
```

```
_log_archive_buffer_size = 1024
```

Enterprise Virtual Array Controller (HSV110)

The Enterprise Virtual Array (HSV110) is the newest generation of the StorageWorks disk array family and is the first that is enabled by VersaStor technology. It is an ultra-high performance, ultra-high-capacity and ultra-high-availability "virtual" RAID (Vraid) storage solution for the upper-end of the enterprise marketplace that removes the time, space and cost boundaries of traditionally architected storage.

The Enterprise Virtual Array complements and extends the current MA8000 / EMA12000 / EMA16000 StorageWorks solutions and can co-exist in the same 1Gbit Open Fibre Channel SAN networks while providing future 2Gb/sec end-to-end Fibre Channel technology readiness.

Enterprise is designed for the data center where there is a critical need for improved storage utilization and scalability while meeting application-specific demands for consistently high transaction I/O and MB per second data rate performance, seamless capacity expansion, instantaneous replication, and simplified storage administration.

EVA supports 36GB and 72GB 10K rpm dual-ported 2Gb/sec FC disk drives at full rated transfer rates at initial release. An Enterprise controller pair will support up to 240 disk drives of single or mixed drive capacities, with 168 in the base cabinet.