

Julu 2002
1638-0102A-WWEN

Prepared by Enterprise Storage Group
Compaq Computer Corporation

Contents

1.0 Management Summary	3
2.0 Solution Description	4
2.1 Solution Components	5
2.2 Features / Benefits	6
2.3 Key Supporting Products	7
2.4 Business Value	7
2.5 Customer Profile	7
2.6 Application Best Fits	8
2.7 Cost / Value Analysis	8
2.8 Positioning with other Compaq Products or Solutions	12
3.0 Performance	13
3.1 Test Description	13
3.2 Test Configuration	13
3.3 Performance Results	16
4.0 Supported Configurations ...	20
4.1 Configuration Information	21
4.2 Sizing	21
4.3 Partner Content	22
4.4 Compatibility	23
4.5 Solution Specific Installation / Setup Instructions	23
4.6 Service	25
5.0 Operational Guidelines	26
5.1 Operation of Solution	28
5.2 Solution Management and Maintenance	35
5.3 On-Going Service and Support	36
6.0 Related Information	37
6.1 Web / Documentation Links	37
6.2 Related Offerings	37
6.3 Feedback	37
Appendix	38

Virtualized Storage Management for SQL Server 2000

StorageWorks Virtualization and Recovery Solution for Microsoft SQL Server 2000 Environments

Abstract

This Solution Guide provides tools and best practices to increase application availability, using the proven virtualization technology, *SANworks* Virtual Replicator by Compaq.

It is intended for use by experienced Microsoft SQL Server 2000 administrators.

In this guide, administrators learn to use snapshot technology to recover a SQL Server environment in far less time than with conventional tape restore methods.

Additionally, this guide describes how virtual storage pools can be centrally managed and expanded, without impacting users or application availability. The solution can be applied as part of a standard disaster recovery practice.

Notice

Virtualized Storage Management for SQL Server 2000
Solutions Guide prepared by Enterprise Storage Group

1638-0102A-WWEN Third Edition ©2002 Compaq Information Technologies Group, L.P.

Compaq, the Compaq logo, VersaStor and ProLiant are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries.

Microsoft, Windows, Windows 2000, Windows 2000 Server and Workstation, Microsoft SQL Server for Windows 2000, Microsoft Exchange for Windows 2000 are trademarks and/or registered trademarks of Microsoft Corporation in the U.S. and/or other countries.

Pentium is a registered trademark of Intel Corporation in the U.S. and/or other countries.

All other product names mentioned herein may be trademarks of 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

Successful businesses utilizing database servers face similar challenges, including the continuous growth of database volumes with critical data, shrinking backup windows, and time-crunched database administrators. The Virtualized Storage Management solution for SQL Server 2000 introduces new tools and strategies for administrators to enhance data availability and streamline storage management. This easy to deploy solution provides a cost-effective method for maximizing data availability through fast recovery and on-line volume growth.

Using snapshot technology, data protection is increased with a time effective recovery alternative to standard tape restores. Administrators maximize operational efficiencies by dynamically growing storage pools without affecting application availability.

SANworks Virtual Replicator enables administrators to manage a single storage infrastructure independent of the physical disk drives and associated disk arrays. *SANworks* Virtual Replicator uses virtual server storage technology consistent with the Compaq VersaStor™ initiative.

This solution provides an integrated infrastructure addressing the challenges of unpredictable growth and disk utilization efficiency to restore the integrity of a database for as little as the cost of the *SANworks* Virtual Replicator licenses.

Testing has shown the following advantages:

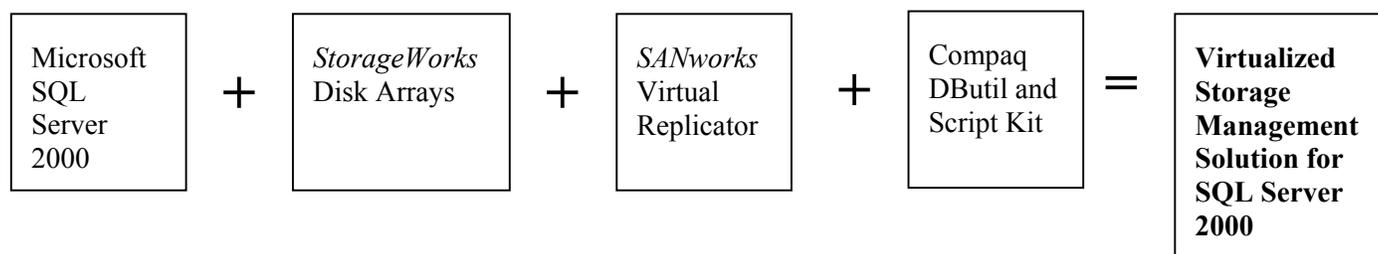
- Snapshot-based restores reduce downtime as compared to tape-based restores. For example, snapshots with changes of 25% can be used to restore the SQL Server database in just 32 minutes – representing over a **four-fold decrease in downtime** as compared to a similar restore from tape. Databases with less activity can represent even shorter recovery times.
- Dynamic online storage growth can be achieved without downtime, as compared to **more than one hour of downtime using standard volume growth practices**.
- Increasing virtual storage pools takes **less than four minutes**. This is a significant time savings, compared with standard volume expansion procedures, which could take hours.

"The requirement for continuous, 24x7 access to our data has added a new dimension to backup operations. We can't afford downtime or access interruptions to process data backup. With the snapshot capability of SANworks Virtual Replicator [the core component in the Virtualized Storage Management solutions set] we can quickly and easily take a space efficient, point in time replica of the data on our production volume..."

- Richard A Baldwin, President & CEO, Nth Generation

2.0 Solution Description

The Virtualized Storage Management for SQL Server 2000 Solution is a fully integrated configuration of hardware and software that provides host-based virtualization for online capacity expansion and snapshot technology for rapid database restoration, providing minimal downtime for Microsoft SQL Server 2000. To achieve these results, this solution for SQL Server 2000 leverages the capabilities of Compaq *StorageWorks* RAID arrays, along with Compaq *SANworks* Virtual Replicator, a special utility, and user customizable scripts as illustrated below.



This solutions guide is intended to give database administrators the best practices, procedures, and tools needed to:

- Create database snapshots
- Perform snapshot based database recoveries
- Dynamically expand virtual storage pools
- Streamline the solution implementation
- Make informed management decisions using performance data guidelines

The solution is a tested configuration of hardware, software, and services. Procedures for installation, configuration, integration, and operation within a typical application environment are included, for a smooth and predictable deployment experience. Performance and operational guidelines, not usually found within individual product documentation, are also available.

This guide includes the following solution specific information:

- Step by step installation instructions
- Validated scripts for the snapshot creation of a SQL Server 2000 database
- Best practice information, cautions, and areas of consideration
- Performance, sizing, and characterization information for a typical customer environment
- Discussion relating to the integration of this solution with regular online backups
- A consolidated list of available tools and other white papers for reference

2.1 Solution Components

The Virtualized Storage Management solution for SQL Server 2000 consists of four components:

SQL Server 2000

SQL Server 2000 may be configured on any supported Windows 2000 server with multiple network clients, and may utilize a MSCS configuration for improved application availability. The application servers own the database, which is stored on virtual disks, sharing disk capacity across the SAN or network.

SANworks Virtual Replicator 3.0

To provide the storage virtualization function and the ability to create snapshots of virtual disks, each application host requires a *SANworks* Virtual Replicator license. In a clustered environment, each node connected to the server instance also requires a *SANworks* Virtual Replicator license.

Disk Storage

Although the configuration tested used the *StorageWorks* MSA1000 fibre based storage array, the solution is neither limited to this controller nor is it a prerequisite. The solution configuration can operate using direct attached storage or any SAN-based fibre array controller. All disk storage devices on the network may be considered as part of the available storage pool and may be used in the creation of virtual disks. The solution can operate with multiple storage devices, as long as they are available to the specific SQL Server 2000 for storage of the associated database. Existing disk storage may be utilized without any modifications.

Tape Backup

A tape backup device and suitable backup software configuration is not required for this solution, but is highly recommended. It is assumed that the user already operates a suitable tape backup configuration and has a pre-existing backup policy and schedule. This solution easily integrates into any pre-existing backup configuration.

Consideration: Although this solution provides procedures for the use of a Virtual Replicator snapshot as a source for data restoration, a snapshot may only be used to recover from specific failure conditions. It is recommended that you continue to perform periodic tape backups, either directly from the application database or from the snapshot itself.

Solution Package

The specially developed tools provided by this solution are necessary to guarantee interoperability between the application, software, and hardware components.

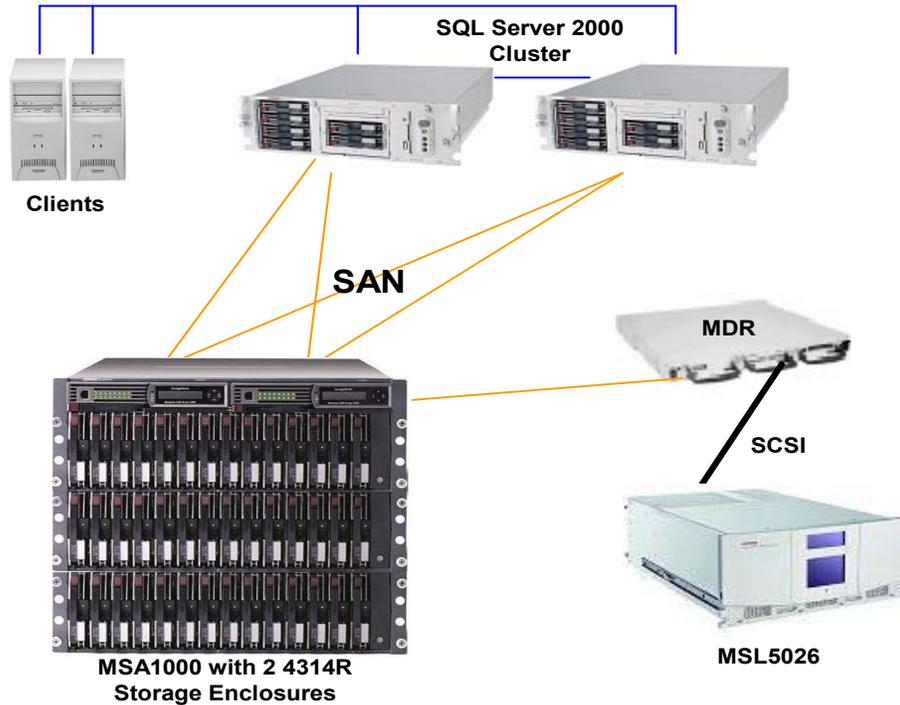
Script Kit – automates operation for designated application servers and database

DButil – provides the necessary interface to SQL Server 2000 to pause and resume the application, allowing for consistency during the snapshot process

Solution Guide – Consolidates information to streamline implementation and operation

Figure 1 represents a typical solution configuration. The disk array configuration may be modified to meet your specific requirements and is only shown as an example. The tape library / backup configuration may also be modified to meet your requirements. The Storage Area Network configuration, although recommended for improved performance, is also not a core requirement for successful operation of this solution. Additional configuration guidelines are presented in Section 4.0, Supported Configurations.

Figure 1 – Typical Solution Configuration



2.2 Features / Benefits

The Virtualized Storage Management for SQL Server 2000 solution provides the following benefits to administrators.

- Simplifies management of data and storage with easy to use, automated tools to increase productivity of IT staff.
- Provides quick and efficient snapshots that are created in seconds, making use of storage capacity only when production data changes.
- Does not require significant investments in hardware or software and can be applied directly to existing environments in just a few hours.
- Can be applied as an integral part of a standard disaster recovery practice that provides timely restoration of corrupted database transactions. Using snapshots, database administrators can recover the SQL Server environment without using conventional tape restore methods.
- Allows virtual storage pools to be centrally managed and expanded effortlessly, without impact to online users or application availability.

- “Solution in a Box” – A complete solution – fully tested, integrated, and documented, providing a simple, predictable deployment.
- With do-it-yourself scripting, database administrators can modify scripting to best fit their business requirements without the need for an onsite service call – Installation and Startup Services are also available from Compaq Global Services.
- Leverages Compaq’s depth of knowledge and experience with Microsoft deployments. There is “peace of mind” knowing that the solution is backed by the industry experts.
- Applicable across the entire range of *StorageWorks* arrays, in SAN and direct attached environments. Select the hardware and network configuration that meets your business need, and be prepared for future expansion.

2.3 Key Supporting Products

This solution requires *SANworks* Virtual Replicator 3.0.

See Section 4.0, Supported Configurations, for complete configuration information.

2.4 Business Value

Solution

All the necessary installation instructions, scripts, and operational best practices are provided at no cost to facilitate a quick, trouble free deployment without the need for in-house development of scripts or best practices. The solution requires no additional products and/or services beyond those required to deploy *SANworks* Virtual Replicator.

Most importantly, even though the solution is automated, you can customize the scripts without a mandatory and costly onsite service engagement. Open systems tools are leveraged whenever possible and Compaq specific tools are provided only where specific functionality is required.

Online volume growth

With online volume growth, you no longer have to over allocate disk storage capacity at the original deployment phase of an application. You can avoid inefficient capacity utilization by allocating only the required number of disks for the planning period and then dynamically grow the data volumes as the database grows. The online volume growth feature also allows you to respond quickly to unpredictable capacity growth without affecting database availability.

Snapshots

Administrators typically rely on tape restores for the recovery of files or databases. By integrating a snapshot recovery strategy into an existing tape backup policy, restoration times can be more than cut in half, directly improving application availability.

2.5 Customer Profile

You can implement the solution where SQL Server 2000 is currently installed or planned. Although tested specifically on Compaq ProLiant servers and *StorageWorks* SAN based storage

hardware, the solution can be administered in any open systems environment where *SANworks* Virtual Replicator is supported.

The solution is specifically targeted towards medium size enterprises (SQL Server database sizes up to 2TB) requiring improved access to data, simplified management for rapidly growing databases, and/or quick recovery from data corruption or frequently lost files. The solution offers the highest value to those companies implementing entry level or departmental level array based SANs, building a foundation for storage virtualization.

It is ideally suited for those customers considering the *StorageWorks* MSA 1000 as part of an entry-level SAN configuration.

2.6 Application Best Fits

This solution is best applied in situations that require:

- Quick restores or database restoration from database corruption, lost or deleted data, frequent user restore requests, or virus damage
- An additional data protection scheme to provide 24 hour database protection, where the time to perform more frequent tape backups is not available
- More efficient use of disk resources. By allowing up to 8 logical units in a storage pool, this solution can allow for a smaller initial disk investment and smaller planning increments, which leads to more efficient disk allocation.
- Strategies to respond to unplanned or frequent volume growth or to cover situations where disk utilization is consistently under-forecasted
- Snapshot capabilities that span multiple disk controllers
- A solution that provides both online volume growth and snapshots all within one single package

2.7 Cost / Value Analysis

There are many options to consider when maintaining a SQL Server storage environment. Each alternative introduces additional costs related to software, installation and configuration, and management. Each alternative also has an inherent value, which includes the following:

- Increase in application availability (limited disruption to the user)
- Ability to respond to failure events (quick database recovery)
- Ease of manageability (minimum daily attention required by the IT administrator)
- Increase in storage utilization efficiency (minimization of the number of unallocated disk drives supporting the application)

Table 1 – Cost / Value Summary

Implementation Scenario	Cost		Value	
	Initial Implementation	Continued Management	Initial Implementation	Continued Management
Standard	Inefficient use of storage resources during initial database deployment	Excessive application non-availability during restores Application non-availability during unplanned volume expansion activities	Easy deployment No additional hardware or software required	Fits into existing management practices
Microsoft Dynamic Disk	Ships with Windows 2000	Excessive application non-availability during restores	Easy deployment No additional hardware or software required	Ability to respond to unplanned storage growth Improved storage efficiencies
Controller - based Cloning	Cost of controller based cloning software licenses Requires allocation of 2x disk capacity	Heterogeneous environments may require multiple management interfaces		Ability to offload backup to non-application server Ability to develop additional restore strategies for improved application availability
SANworks Virtual Replicator	Cost of host based <i>SANworks</i> Virtual Replicator software licenses Cost of testing and developing database integrated scripts and policies	Minimal	Requires less disk capacity than clone implementations	Ability to respond to unplanned storage growth Improved storage efficiencies Ability to develop additional restore strategies for improved application availability
Virtualized Storage Management Solution	Solutions package is a free downloadable enhancement to <i>SANworks</i> Virtual Replicator	Minimal	All of the above, plus all installation instructions, scripts, and best practices are provided	All management and maintenance policies are provided

Standard Implementation

Database administrators rely on standard tape backup strategies to recover from database failure events. Regardless of the level of the database failure, the only recovery strategy available to the administrator is to restore the database or database components from tape. Even an insignificant failure event requires a significant restoration period, over which time the application availability would be compromised.

Database administrators rely on typical storage capacity planning techniques, where storage capacity equal to the maximum database size over a particular planning period, plus some margin, is allocated to the application. This technique requires the purchase of additional storage capacity at the beginning of the planning period, regardless of the forecasted utilization, which can result in over allocation and under utilization of storage.

Our experience shows that about 30 percent of the allocated storage is actually utilized at the outset of the planning period. Although this margin provides room for forecasted database growth and unforeseen spikes in database size, it is clearly a bad investment practice, especially with the historical decrease in storage pricing over time.

There is always a trade off between over allocation of storage resources at the start of a planning period and the impact of downtime for unplanned volume expansion. In the event that the database grows beyond its allocated capacity, new and larger volumes have to be created. The database needs to be migrated to the larger volumes outside of the planned maintenance period. This practice directly impacts application availability during the data migration to the new volume.

Microsoft Dynamic Disk Implementation

The Microsoft Windows 2000 operating system offers the administrator the additional capability of storage virtualization through the use of Dynamic Disks. With Dynamic Disks, the administrator can create volumes that span multiple disks, increase the size of an existing volume, or adding a mirrored volume without a reboot. This enables improved disk utilization by providing the administrator the ability to increase disk volume size as the need arises during the planning period. It also provides the administrator with a tool to respond to unforeseen or under-forecasted storage growth.

Consideration: Microsoft Windows 2000 Advanced Server and Datacenter do not provide support for dynamic disks in a server cluster (MSCS cluster) environment. Refer to Microsoft Tech Note: [Microsoft Cluster Service Installation Resources, Q259267](#).

Consideration: Dynamic disks must be originally configured as such (1 MB of configuration data must be stored during initial configuration at the end of each dynamic disk). Existing disks, not already configured as dynamic disks, cannot be imported into an existing dynamic disk pool.

The Microsoft Dynamic Disk implementation does not address the previously described issue associated with relying on tape for database restoration.

Controller-based Clone Implementation

Administrators also use array controller-based cloning software to create business continuance volumes (BCVs). This technique provides a controller-based option for both database backup and database restores. When integrated with existing tape-based backup and restore policies, a clone-based backup can improve application availability by off loading the backup function to a non-application server. Clones may also be used for quick restoration for certain database failures, improving application downtime as compared to typical restoration from tape.

Cloning software packages are typically controller specific and therefore are not supported across multiple storage configurations or manufacturers' products. Cloning requires a significantly greater investment in disk resources, as you must allocate 50% more capacity for mirrored volumes, as compared to potentially lower levels for snapshots.

Clone implementation does not address the previously described issue associated with the inability to dynamically grow data volumes. Clone based solutions do not offer storage virtualization capabilities, such as those features found in Microsoft Dynamic Disk or *SANworks* Virtual Replicator.

Caution: Compatibility between cloning software and the Dynamic Disk function should be thoroughly investigated and tested. Support for the cloning of volumes configured as Dynamic Disks should be verified.

Implementation of *SANworks* Virtual Replicator

With *SANworks* Virtual Replicator, database administrators dynamically allocate additional disk capacity and can also restore directly from snapshots of production volumes, all in one package. Operating as a host-based application, this solution can span multiple storage configurations and can be implemented across multiple storage manufacturers' products.

Online Volume Growth - With the ability to dynamically allocate storage capacity, the administrator no longer has to over allocate capacity. More importantly, situations that require the allocation of additional capacity between planned maintenance periods do not have to result in downtime. The unplanned growth of existing database volumes normally requires the application to be unavailable for up to one hour, whereas using the online volume growth capability of *SANworks* Virtual Replicator, application downtime is reduced to zero.

Snapshots – Snapshots are used to augment conventional disaster recovery plans that rely on restoration from local or remote tape backups. Snapshots are easy to create and discard as necessary. Since snapshots take only minutes to create, they can be recreated frequently providing a current base from which to restore. Snapshot volumes typically utilize less capacity than clone volumes, offering lower disk utilization costs.

Restoration from snapshot volumes can reduce recovery times, which directly affects application availability. Snapshots can be functionally equivalent to conventional differential backups and can form an integral part of a backup and recovery strategy to provide highly available databases environments. *SANworks* Virtual Replicator supports incremental backup to maximize efficient backup operations with industry standard backup utilities.

Virtualized Storage Management Solution Implementation

Beyond the cost of purchasing software licenses, the implementation of any software solution also requires the cost of installation, testing, integration, and deployment into the production database environment. In order to implement any software-based solution, the database administrator must find the time and resources to develop and debug the required tools and scripts for integration with the application.

This solution does not require additional or specialized hardware, software, or services, making for a very cost-effective implementation. This Solutions Guide provides the administrator with installation instructions, scripts, and best practices required for a quick and successful implementation of the *SANworks* Virtual Replicator software within a SQL Server 2000 environment. Using this document, the times needed to create, test, and deploy the solution is significantly reduced, by as much as 50%, by leveraging Compaq expertise.

2.8 Positioning with other Compaq Products or Solutions

If you require the most rapid restore capabilities possible in the industry, Compaq also offers the *Rapid Restore for SQL 2000 solution*. This solution, based on *SANworks* Enterprise Volume Manager by Compaq, provides extremely rapid recovery through cloning technology. Instead of the virtual copies that a snapshot provides, clones are an entire physical replica or third mirror of the database. Testing for this solution has shown that the environment can be restored in a matter of just minutes. This solution works with the entire family of *StorageWorks* RAID arrays as supported by *SANworks* Enterprise Volume Manager. For more information, or to obtain the solutions guide, see:

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

3.0 Performance

A test configuration was developed to characterize the key features of *SANworks* Virtual Replicator within a SQL2000 environment. The impact to the SQL application, during the creation of multiple virtual disks and the on-line dynamic growth of a SQL database volume, was measured. The performance impact of using a *SANworks* Virtual Replicator created snapshot for recoveries was also determined. Baseline performance results were also captured, allowing you the ability to compare these results to other potential configurations.

3.1 Test Description

Solution characterization testing was based on a low to mid-range SQL Server 2000 environment operating *SANworks* Virtual Replicator 3.0 software. A TPC-C load simulator was used to produce a relatively high-level user load. Performance information was captured and used as a baseline for later comparison.

Specific server based and SQL application based performance indicators were monitored while the following test plan was performed:

- Baseline configuration without *SANworks* Virtual Replicator
- *SANworks* Virtual Replicator virtual disk baseline without a snapshot
- *SANworks* Virtual Replicator snapshot creation
- *SANworks* Virtual Replicator Snapback (restore of database form a snapshot)
- *SANworks* Virtual Replicator online volume growth (compared to traditional LUN growth)
- *SANworks* Virtual Replicator snapshot tape backup/restore using a third party backup application

3.2 Test Configuration

The Virtualized Storage Management test recreates a subset of a SQL Server 2000 user environment. The test configuration consisted of clustered (active/passive) SQL Server 2000 application servers attached via fibre channel to a *StorageWorks* MSA1000 storage array. The storage array was configured using RAID 1+0 for optimum data protection.

The basic operating components of the configuration include:

Clients

Deskpro Workstations running a TPC-C benchmark load generator application, to simulate users.

Application Servers

Clustered ProLiant servers loaded with Microsoft SQL Server 2000, Compaq *StorageWorks* Secure Path 3.1B for Workgroups (*StorageWorks* Secure Path 3.1a for *StorageWorks* HSG80 array configurations), and *SANworks* Virtual Replicator 3.0 software.

Enterprise Storage

StorageWorks MSA1000 based disk array with 1.5TBs of disk storage capacity.

Backup – Enterprise Backup Solution

SAN attached *StorageWorks* MSL5026SL SDLT library and VERITAS Backup Exec 8.6 backup software.

Application Server Configuration

(2) Compaq ProLiant DL380

Hardware

(2) Pentium III, 1 GHz Processors – 512k level 2 cache (256 per processor)

2 GB RAM

(2) 18 GB Ultra3, 15k rpm drives (Operation System and SQL Application)

Software

Microsoft Windows 2000 Advanced Server (Service Pack 2)

Microsoft Cluster Server – MSCS

StorageWorks Secure Path 3.1b for Workgroups

Microsoft SQL 2000 Enterprise Server (Service Pack 1)

SANworks Virtual Replicator 3.0 (with Import Unit Patch for clusters)

StorageWorks Modular SAN Array 1000 Support Software

Enterprise Storage Configuration

StorageWorks MSA1000 Enterprise Array with Dual Controllers

(2) 2 Gigabit *StorageWorks* MSA Fabric Switches

Dual port fibre I/O module

(2) 4314R disk shelf (single bus)

36 GB 10000rpm U3 disk drives

Database Configuration

SQL Server 2000 database size of 200 GB spread over 3 virtual disks of a total capacity of 212 GB.

Two pools were used, one for the data volumes and one for the logs.

Load Generating Application

Microsoft TPC-C Benchmark Kit for SQL

Table 2 - Hardware Part Numbers / Solution Configuration

Server Hardware	Part Number
ProLiant DL380 servers –	193706-001
Dual Pentium III 1GHz Processors, 512K level 2 cache (256K per processor)	187602-B21
2GB RAM Memory	128280-B21
Thermal Upgrade Kit	210818-B21
Two 18GB Ultra 3, 15K rpm drives in each server	188122-B22
FCA 2101 (LP952) HBA's in each server	245299-B21
NC6134 Gbit NIC card	174818-B21

Storage Hardware	
<i>StorageWorks</i> MSA 1000 Storage Array	201723-B21
<i>StorageWorks</i> MSA 1000 Controller	281231-B21
<i>StorageWorks</i> MSA Fibre Channel I/O Module	218960-B21
(2) <i>StorageWorks</i> MSA Fabric Switch 6	218232-B21
(2) <i>StorageWorks</i> Enclosure Model 4214R	190209-B31
(42) 36GB Ultra 3, 10K rpm drives	176496-B22
<i>StorageWorks</i> Modular Data Router	163083-B21
<i>StorageWorks</i> MSL5026SL RM SDLT Minilibrary	231892-B22
30m LC-SC Multi-Mode Fibre Channel Cable	221691-B26
15m SC-SC Multi-Mode Fibre Channel Cable	221691-B24
5m LC-LC Multi-Mode Fibre Channel Cable	221692-B22
(2) 8 Port D-Link Gigabit Switch (network)	Third Party

Table 3 - Software Part Numbers / Solution Configuration

Software	Part Number
Windows 2000 Advanced Server with SP2	Third Party
Microsoft Cluster Server - MSCS	Included
Microsoft SQL Server 2000 Enterprise Edition with SP1	Third Party
<i>SANworks</i> Secure Path 3.1b for Workgroups	213076-B21
<i>SANworks</i> Virtual Replicator 3.0 CD/Media (license required)	261769-B21
<i>SANworks</i> Virtual Replicator 3.0 (5 license)	261775-B21
MSA 1000 Support Software (Included with MSA 1000)	253691-001
TPC-C Benchmark Kit for SQL (load generator)	Third Party
VERITAS Backup Exec Software	Third Party

Consideration: This solution was tested on a specific configuration as documented in this section, using a SQL Server 2000 environment with typical database sizes and user loads. The best practices, considerations, and performance information may be applied to similar environments, but results may differ from those published.

3.3 Performance Results

Test Characterization Data

The data presented was captured using Windows Performance Monitor and the metrics were recorded for the duration of each test. The results were compared to Test 1 (Baseline configuration without *SANworks* Virtual Replicator), and summarized.

Table 4 – System Monitor Performance Counters

Metric (Note: data represents averages over test duration)	(1) Baseline without SWVR MSA 1000	(2) Baseline with SWVR MSA 1000	(3) Snapshot (backup) MSA 1000	(4) Snapback (restore) MSA 1000	(5) Online Volume Growth MSA 1000
Test Duration	4 hrs	4 hrs	3 min 15 sec	1.5 hour	2 min
Transaction/sec	62	71	0	0	69
Active Transactions	38	40	50	0	45
% Processor Time	52	52	49	2.17	54
Processor Queue Length	2.1	2.5	1	0.9	3.4
Memory Available Bytes MB	331	320	266	268	248
Pool Copyouts	N/A	N/A	1246956	305624	0
Log Snapshot Delspace MB	N/A	N/A	0	716	N/A
Data1 Snapshot Delspace MB	N/A	N/A	0	14207	N/A
Data2 Snapshot Delspace MB	N/A	N/A	0	40474	N/A

Test 1: Baseline without *SANworks* Virtual Replicator installed

The baseline test involved a 4 hour load simulation on a SQL Server 2000 database of 208 GB. This test was run without *SANworks* Virtual Replicator loaded to generate a performance counter baseline for comparison against the results of subsequent tests.

Test 2: Baseline with *SANworks* Virtual Replicator installed

The SQL database was imported into the virtual pool created by *SANworks* Virtual Replicator. The same 4 hour load was run and any performance impacts were recorded.

Result: Running off a virtual disk results in NO impact on system or database performance.

Conclusion: *SANworks* Virtual Replicator application does not impact the SQL Server 2000 environment.

Test 3: Snapshot impact

Using the tools and procedures provided by this solution package, snapshots were created for three virtual disks (two database volumes and one log file).

Result: The entire snapshot creation process took ONLY 3 minutes and 15 seconds. Most importantly, it was demonstrated that user transactions were not impacted. Active transactions are not lost, they are queued until the snapshot process is complete.

Conclusions: Snapshot creation occurs quickly and is transparent to the SQL Server user. Users do not have to reconnect to the application. The database is not taken offline, no transactions are lost, and system performance is not affected.

Test 4: Restore from Snapshot using Snapback

The database was restored using both the *SANworks* Virtual Replicator Snapback feature and a traditional tape restore. The tape restore time was used as a baseline for comparison and resulted in a constant restore time of over 2 hours. Snapback restore times varied based on the amount of copy-outs performed while the snapshot was open (see Table 5). During a copy-out, old data from the parent disk is copied to the snapshot disk prior to database writes.

Although the database may not show a visible change in size, writes to the database will potentially cause an increase in copy-outs thus increasing the size of the snapshot volume. Therefore, the longer a snapshot remains open, the longer it takes to restore.

Consideration: High levels of database activity, coupled with long periods after snapshot creation, will cause the snapshot volume to increase significantly and consequently consume more virtual pool space. It is important to monitor the snapshot growth to determine the best period between snapshots for your specific environment.

Results: Snapback times ranged from 1 minute (1% change in virtual disk) to 60 minutes (45% change in virtual disk).

Conclusion: Restoring from a snapshot whose percentage of the virtual volume is less than 25% is extremely time efficient when compared to tape-based restore.

Best Practice: If the snapshot is to be used for restoration purposes, it is suggested that a snap period be established that produces a snapshot size, as a percentage of the virtual volume, that is less than 25 percent. This is recommended in order to take advantage of the benefits of snapshots for rapid restoration. You may wish to perform your own snapshot and tape restore tests to determine a snapshot volume limit for your own specific environment.

To monitor the snapshot percentage of the virtual volume, check the “delspace” value for the virtual disk. The ”delspace” is displayed in the Snap Manager interface and represents the actual data space that the associated snapshot uses in the virtual pool. If the snapshot percentage continually exceeds the 25 percent suggested limit within the period between snapshots, the following alternatives are suggested:

1. Reduce the time between snapshots
2. Consider an alternative rapid recovery method (Rapid Restore Solution for SQL Server 2000)
3. Rely on tape based restores

Table 5 and Figure 2 show the backup and restore times for a 56 GB SQL database file using various methods. The graph has a comparison of the restore times to show how the Snapback feature in *SANworks* Virtual Replicator can minimize the restoration process. The restore from Snapback, derived from a snapshot size that was approximately 25% of the 56 GB parent virtual disks, resulted in a significant improvement over typical database restore alternatives.

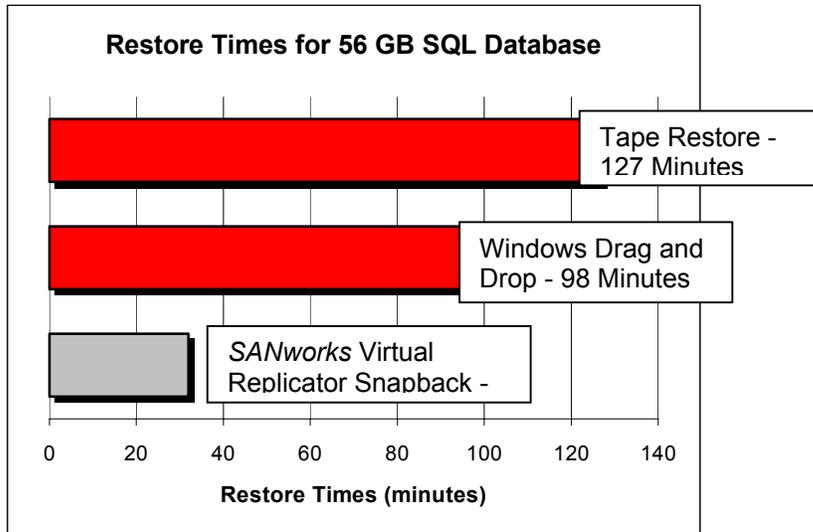
Table 5 - Backup and Restore Times and Data Rates Using Various Methods

Backup Type	Backup Time (min)	Backup Rate (GB/Hr)	Restore Time (min)	Restore Rate (GB/Hr)	Data size (GB)	% Virtual Disk Change	Delspace (GB)
WIN2K Drag/Drop copy	N/A	N/A	98	9	56	25%	14.005
Snapshot of all Database Volumes	3.25	3840	N/A	N/A	208	N/A	N/A
Snapback after 10 Min	N/A	N/A	1	38	56	1%	0.639
Snapback after 1hr	N/A	N/A	32	27	56	25%	14.207
Snapback after 4hr	N/A	N/A	60	25	56	45%	24.876
SQL Native Backup/Restore ¹	20	168	20	168	56	N/A	N/A
VERITAS Tape Full Backup/Restore	80	42	127	26	56	N/A	N/A

¹ Clone-based recovery methods require considerably larger disk capacities and represent a more hardware intensive alternative. Tape should always be used as a final restore option and as the most suitable option for catastrophic failures. Although the restore time using SQL native restore feature was faster than the Snapback process, there are certain factors to consider:

- The native backup process used to create the disk copy takes a considerably longer time than the 2 minutes required to create a snapshot.
- During the native backup creation period, user performance will be degraded.
- The native backup disk copy requires the allocation of additional disk space equal to the database volume.
- Using the SQL native backup feature during periods of heavy user activity will have an impact on both server and storage performance.

Figure 2: Restore Times vs. Restore Method



Test 5: Online Volume Growth vs. Traditional LUN Growth

The data volumes were increased by 30GB while a full user load was running on the production server. This result was compared to a traditional LUN expansion where a new LUN was created from a separate disk and a copy of the database was written to that LUN.

Result: The entire *SANworks* Virtual Replicator online volume growth process took four minutes to complete as compared to 60 minutes for the traditional LUN expansion technique. During this online volume growth operation, there was no impact on the SQL application.

Conclusion: *SANworks* Virtual Replicator online volume growth occurs within minutes and is transparent to the SQL Server user. The database is not taken offline, no transactions are lost, and system performance is not affected. Dynamic on-line storage growth can be achieved with less than 4 minutes of administrative time, enabling additional capacity to on-line database volumes in less than 7% of the time that it typically takes without the solution.

Figure 3 – Comparison of Disk Growth Times*

Process	Traditional	SANworks Virtual Replicator
Preparing Disk Hardware		
Add disks to Chassis	20	20
Configure RAID Array	30	30
Create Logical Drive	10	10
Total Time for Preparation	60 minutes	60 minutes
Volume Expansion		
Present Logical Drive to O/S	1	1
Partition and Format disk	5	N/A
Add disk to Virtual Pool	N/A	1
Grow Virtual Disk	N/A	2
Shut down databases. No user access.	2	N/A
Copy databases and Logs to new LUN. (I/O is 20MB/s, 60GB of data)	50	N/A
Start databases	2	N/A
Total Time for Volume Growth	60 minutes	4 minutes

NOTE: All components were already in place to run this test (no logistical issues). The level of database administrator experience can affect the results.

4.0 Supported Configurations

Note: In order to continually provide improved features and functions, the hardware and software products used in this solution are subject to change. Although it is our objective that all newly released products be compatible with previously released solutions, Compaq makes no warranty of compatibility of any kind. As always, it remains the responsibility of the user to refer to the specific component documentation when applying component versions that do not match those stated in this guide.

4.1 Configuration Information

Minimum Requirements

The configuration must include:

- at least one Windows 2000 host operating a SQL Server 2000 database
- one instance of *SANworks* Virtual Replicator 3.0
- a storage volume equal to at least 1.3 times the database volume, viewable by the application server

Minimum application host requirements:

- Pentium class Processor
- 128 MB Memory
- 40 MB available disk space for full installation

Consideration: Use a storage volume size of 1.3 times the database size only as a guideline. Carefully consider the actual volume size after analysis of the number of snaps, the size of the database changes, and the period between snapshots. Although not recommended or likely, the snapshot volume can consume up to the same size as the virtual disk. *SANworks* Virtual Replicator provides the user with a warning when there is 5 percent of the volume left for snapshot writes.

Maximum Requirements

The virtual storage pool viewable to the Windows 2000 server is limited to 2 TBs per virtual disk. Refer to *SANworks* Virtual Replicator Quick Spec for more detail. See Section 6.1, Web / Documentation.

4.2 Sizing

Snapshot Planner

The *SANworks* Virtual Replicator application includes a tool that can be used to plan snapshot allocations, called Snapshot Planner. You can download the *SANworks* Virtual Replicator kit from

<http://www.compaq.com/products/sanworks/softwaredrivers/vr/>

This kit includes the Snapshot planner tool which can be used prior to purchasing the software license. This tool needs to be used prior to loading and configuring *SANworks* Virtual Replicator to help predict the overhead of snapshots in terms of disk space and disk I/O. Snapshot Planner tracks actual reads and writes to existing drives and determines how much disk space a snapshot of each drive requires. The Planner also projects any additional disk I/O that the snapshot may generate. For more information on how to use Snapshot Planner, refer to the *SANworks* Virtual Replicator 3.0 System Administrators Guide.

The virtual storage pool size depends on the size of the virtual disk that is required for the database volume. You should plan to configure a pool with enough capacity to allow space for creating a primary virtual disk for the data volume, space for copyout activity associated with two snapshots, and space for overhead.

Best Practice: Run Snapshot Planner before installing and setting up the *SANworks* Virtual Replicator environment. After running Snapshot Planner, the tool should be uninstalled prior to installation of *SANworks* Virtual Replicator.

Related sizing information may be obtained from the following location:

Active Answers Microsoft SQL Server Sizing Tools

<http://activeanswers.compaq.com/ActiveAnswers/Render/1,1027,519-6-100-225-1,00.htm>

4.3 Partner Content

The following third party products are required for the successful implementation of this solution:

- Microsoft Windows 2000 Server or Advanced Server
- Microsoft SQL Server 2000
- Microsoft Internet Explorer
- Microsoft Cluster Server (MSCS) (optional)
- Microsoft Task Scheduler or related product (optional)
- Tape backup application (optional)

Consideration: Compaq and Microsoft are long-standing partners. With their combined strengths, they offer comprehensive support, integrated services, recognized expertise, and complete solutions. More information on this partnership can be found on the Compaq/Microsoft Global Alliance (<http://www.compaq.com/partners/microsoft/index.html>) web site.

Consideration: During implementation, you can back up the snapshot using the SAN-based Compaq Enterprise Backup Solution (EBS). Check this link for a list of EBS Software Partners who have developed a tightly integrated and tested data protection solution that combines the partner backup application with Compaq *StorageWorks* hardware and *SANworks* software.
(<http://www.compaq.com/products/storageworks/ebs/partners.html>)

4.4 Compatibility

This solution operates with all array controllers and direct attached storage configurations viewed as physical devices by a Windows 2000 Server or Advanced Server operating system. It is supported behind MSCS configurations and operation in a redundant failover configuration. The solution has been tested using the *StorageWorks* MSA 1000 array controller, but can be applied across any of the *StorageWorks* arrays.

Consideration: Note that the configurations demonstrated in this guide do not represent a complete qualification and are only a demonstration of performance or functionality for the as tested condition.

Consideration: Any Windows 2000 server platform used in this solution must comply with the hardware and software prerequisites stated in the *SANworks* Virtual Replicator Quick Spec:
<http://www.compaq.com/products/sanworks/vr/index.html>

All other solution hardware components are supported as described in their associated published Quick Specs.

If integrating with a SAN-based backup configuration, please refer to the existing Compaq Enterprise Backup Solution Compatibility table:
<http://www.compaq.com/products/storageworks/ebs/EBScompatmatrix.html>

4.5 Solution Specific Installation / Setup Instructions

You can follow the installation procedures provided in the documentation associated with each component product. In addition, you should be familiar with the information found in the *SANworks* Virtual Replicator Installation Guide.

Caution: Many of the commands and procedures described require database administrator privileges and can be disruptive to database operation if used incorrectly. If you are not familiar with these commands and functions, consult your SQL Server documentation before proceeding. As with any change to your database configuration or backup plan, it is important that you test these changes thoroughly before implementing them in a production environment.

If using Microsoft Cluster Server in this solution, you need to install *SANworks* Secure Path. Secure Path manages the path that the server uses to connect to the storage. It is recommended that you install Secure Path before installing *SANworks* Virtual Replicator. Refer to the Secure Path Installation Guide for installing Secure Path within a clustered environment.

Once you have determined the pool size and virtual disk sizes you need (refer to Section 4.2, Sizing), you can install *SANworks* Virtual Replicator. For clustered configurations, be sure to install the first instance of *SANworks* Virtual Replicator on the standby or passive nodes within the cluster. You will need one Virtual Replicator license for each node in your cluster.

Consideration: It is highly recommended that you download and install the *SANworks* Virtual Replicator 3.0 Import Unit patch. This patch has specific cluster fixes and must be installed if you intend to import existing partitioned volumes into your storage pools. Click on the following link to obtain the patch.

<http://www.compaq.com/products/sanworks/softwaredrivers/vr/vr30import.html>

Once you have installed Virtual Replicator and set up the software licenses on all passive nodes, move cluster resources to those nodes and install *SANworks* Virtual Replicator on the remaining nodes. You should now be able to use the SnapManager interface to configure your storage pools. Once the storage pools are configured, you are ready to use *SANworks* Virtual Replicator within your SQL environment.

DButil SQL Utility

In order to ensure data consistency, before creating a snapshot, you need to queue all disk writes to the SQL database. A Compaq developed utility called DButil is provided to perform the necessary SQL database pause and resume functions. DButil provides a command line interface to perform the suspend and resume operations. DButil uses the Microsoft supported interface for SQL Server 2000 Snapshot Backup and Restore.

Best Practice: You must copy DButil on to each node in a SQL Server cluster. If the cluster should fail over, you must place the scripts on the failover node and recreate the schedules.

You can download DButil v1.0 from the Compaq Virtualized Storage Management Solution web site at the following link.

<http://www.compaq.com/products/storageworks/solutions/vsm-sql2000/scripts.html>

DButil must be placed on each node that connects to the same SQL server database. After you place DButil in a folder of your choice, you are ready to create snapshots.

4.6 Service

The following optional service is available to help the user in the design, installation, configuration, and operation of this solution:

***SANworks* Virtual Replicator Installation and Startup Service**

Part Numbers: 269991-002 (QR-SMSVR-IS)

Compaq CarePaq Part Numbers: 236027-002 (FP-CSTAR-09)

Service Description Summary

The *SANworks* Virtual Replicator Installation and Startup Service provides the following:

- Optionally implements the Virtualized Storage Management solution for SQL Server 2000
- Installs and configures the *SANworks* Virtual Replicator software for successful snapshot creation on the customer's existing storage environment
- Assistance with the startup and operation of *SANworks* Virtual Replicator software
- Provides a sample restore procedure of the SQL database from the snapshot
- An orientation session provides an overview of the configuration and utilization of *SANworks* Virtual Replicator software as it applies to this solution

5.0 Operational Guidelines

Designing The System

For most databases, including SQL Server 2000, some general considerations can make a significant difference in performance.

- Use RAID fault-tolerant settings (RAID 1+0) in the disk subsystem design for database storage.
- In order to facilitate database recovery, locate the SQL databases on a separate volume from the transaction logs. This also improves database performance by separating random reads/writes (database files) from sequential reads/writes (log files).
- When creating logical drives using the Compaq Array Controller Utility, create separate arrays for data files and log files. When creating virtual pools, use different arrays for each pool to ensure that the two pools don't share physical disk resources. Using this strategy, your log files will always be located on different volumes than your database.
- Monitor and maintain pool size throughout the operation of the solution. By spreading the database across multiple high performance disks, database performance will be greatly improved. For more information, refer to Section 5.2, Solution Management and Maintenance.

Best Practice: Microsoft recommends that the log files and data files be kept on separate volumes.

- When building database volumes, Microsoft recommends preallocating space for potential database growth, allowing the database to grow into the size of the volume. The Virtualized Storage Management for SQL Server 2000 solution allows on-the-fly growth for database volumes, therefore allowing for more efficient disk space utilization.

Using Online Volume Growth

- Set the array controller to medium for both the Rebuild and Expand priority options. The medium setting allows for faster logical drive extension when the server is heavily utilized.
- Delete all snapshots before growing virtual disks.

Note: The *StorageWorks* MSA 1000 offers a capacity extension feature through the Array Configuration Utility (ACU) that allows you to increase the size of existing logical drives without disturbing the data on those logical drives. If an existing logical drive is full of data, you can extend the logical drive when there is free space on the array. If there is no free space on the array, you can add drives to the array and proceed to extend the logical drive. However, this feature is limited to the physical restraints of the storage array, whereas, *SANworks* Virtual Replicator allows for the growth of a virtual pool leveraging all disk resources across the SAN.

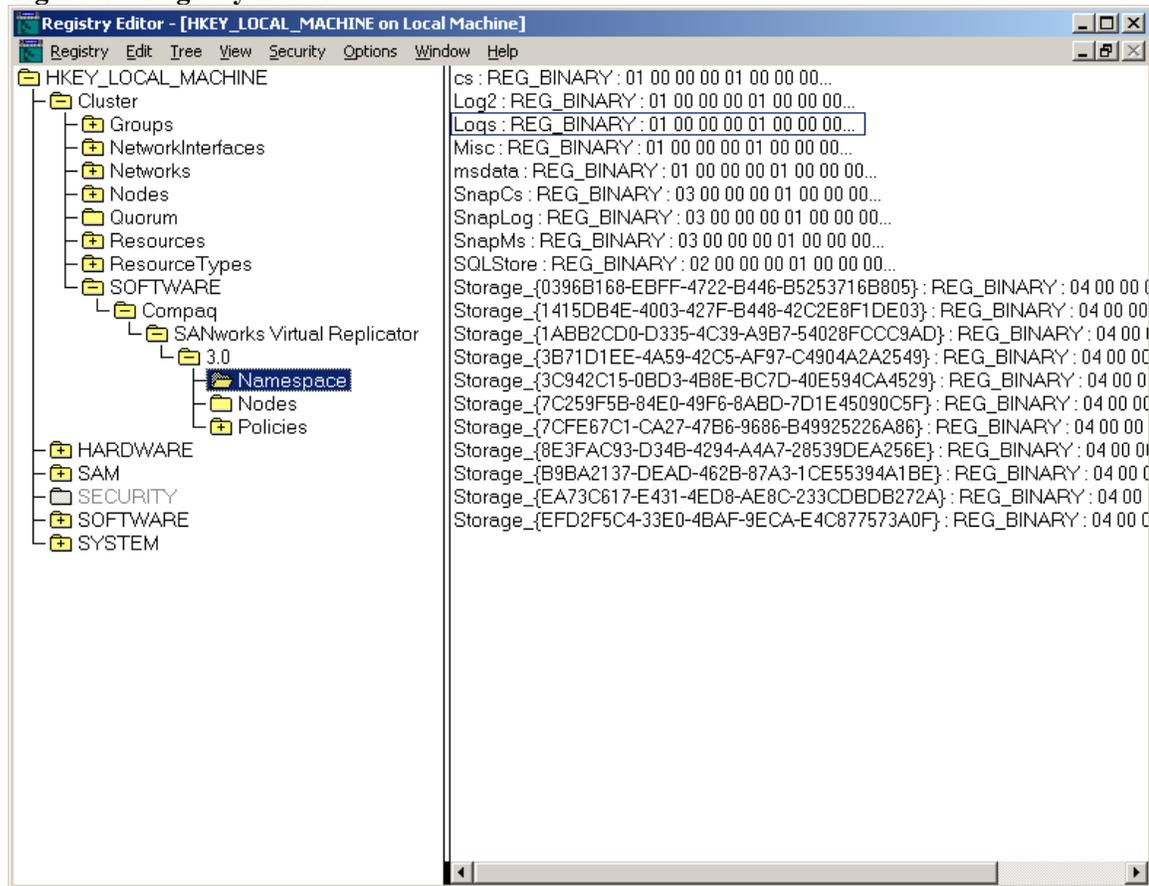
Caution: SANworks Virtual Replicator will not allow use of the Online Volume Growth function unless all open snapshots have been deleted. If the Online Volume Growth function is not responding, verify that all snapshots have been deleted.

- For minimum user impact, perform online volume growth during times of the lowest server and application traffic.
- Close all Microsoft Management Console (MMC) windows before beginning Online Volume Growth. MMC must not be in operation when using the Online Volume Growth feature.

Consideration: The segment size set for the virtual pool will define the maximum size of the virtual disk. Once a specific segment size has been set, Online Volume Growth cannot expand beyond this limit.

Other Related Issues

- When running in a clustered environment, you may encounter an error when trying to add a unit to a virtual pool. You should reboot the cluster nodes to clear the error.
- If the storage pool or virtual disks are lost due to operator error (other than from deletion using the Snapshot Manager), the names are retained in the registry. You need to manually delete the names from the registry in order to use them again. Follow the path outlined in the screenshot below to delete names from the registry. To modify the registry information, click on Start / Run / Regedt32.

Figure 4 – Registry Editor

In the Namespace folder, you can view the names of your old snapshots, virtual disks, and storage pools that existed prior to a loss in configuration. These names must be deleted in order to use them again.

5.1 Operation of Solution

This solution uses snapshots as a component in your overall backup and recovery strategy. This section describes how to use and manage snapshots in a SQL Server 2000 environment.

Using the Compaq DButil Utility

Database files and logs must be in a consistent state before you create a snapshot. Disk writes to the SQL Server database must be queued, flushing the cache and suspending I/O activity to the disks at the time of snapshot creation.

During the brief period when the SQL database is suspended, read access is completely unaffected and read requests will be processed as they would normally. Writes may continue to be submitted but will be queued. The user may see a brief delay on write requests, waiting for the database to resume. Note that suspending the database does not affect the normal transaction isolation level. Users are not logged out of the database and active transactions can continue.

All held transactions are automatically rolled forward once the database is resumed. Snapshot creation is therefore transparent to the user.

Once the snapshot process has been performed (approximately 1 minute) and the SQL database resumed, all queued transactions are committed and I/O activity to the database continues normally.

-Pre and -Post Scripts

DButil uses a simple command line interface that is easily integrated into batch scripts. From the folder where DButil is located, run the following command to write suspend your database:

```
DButil -pre <server\instance name> <database name> <user name>  
<password>
```

Once you have completed your snapshots, resume the database to continue normal operation. DButil uses the following post command to resume the database.

```
DButil -post <server\instance name> <database name> <user name>  
<password>
```

Example scripts are provided as batch files that you can use to automate these processes. You can schedule batch programs to further automate the snapshot process.

Creating Snapshots

It's important to monitor the size of the snapshots because testing has shown that snapshots that are attached to the parent virtual disks for extended periods of time take longer to restore. See Section 3.3, Performance Results for more information on snapshot sizes and Snapback times.

Considerations: Log files generate high I/O and therefore discretion should be used when attaching snapshots to log volumes. Multiple log snapshots will directly impact application performance.

Follow these steps to create snapshots of the database:

Step 1: Perform an online full backup of SQL Server

Best Practice: Always perform a periodic (daily) full online backup using the SQL native backup engine or a backup program of your choice.

Step 2: Prepare the SQL databases

Use DButil to prepare the database before you take a snapshot.

```
DButil -pre <server\instance name> <database name> <user name>  
<password>
```

Be sure to pause the SQL databases from any I/O activity while the snapshot is taken.

Step 3: Create the new snapshot of the database files

With the SQL database suspended, use the Virtual Replicator Snapshot Manager interface to create a new snapshot from the parent virtual disk hosting the database.

Step 4: Resume SQL databases

Resume the SQL database with the DButil command line interface.

```
DButil -post <server\instance name> <database name> <user name>  
<password>
```

Testing has shown these steps can be accomplished in less than two minutes.

Step 5: Unmap the drive letter from the old snapshot

To access the snapshot, you must assign a drive letter to the snapshot volume.

Step 6: Delete the old snapshot

You can have up to 12 snapshots of a virtual disk, but multiple snapshots consume capacity within the storage pool. It is not recommended to keep snapshots if they have grown to more than 25 percent of the parent virtual disk. Snapshots of this size take much longer to restore. It is recommended that you retain the snapshot from the prior period so that one current snapshot is available during the above process. Similar to tape retention schemes, you should maintain a father-son snapshot relationship at all times.

Optional: Backup snapshot to tape

Although not documented in this guide, snapshots can be used as a tape backup image. Once you have the snapshots mapped to the backup server, you can back up the snapshot to tape. Backing up the snapshot to tape provides added protection from a catastrophic failure. Tape backup tools are available to allow backups while the database is online without the risk of data corruption.

Considerations: Since a snapshot shares block level data with the original database, using a snapshot as a backup source does not resolve performance contention issues, as would be addressed with clone based backups.

Sample Script

Much of the process above can be automated using scripts. You can also schedule the scripts to launch at the appropriate times using a scheduler of your choice. Windows 2000 Task scheduler was used in this solution. If you use the downloaded files, you will need to modify the scripts to accommodate your specific environment. No scripts were developed for the restore process, as restores are unplanned and should be performed manually.

The sample script provided below was used in this solution to automate the creation of snapshots. This script is also provided as a .BAT file that can be downloaded from the Virtualized Storage Management for SQL Server 2000 web site. This script should be run from the cluster node that hosts the SQL resources. This script uses DButil and the *SANworks* Virtual Replicator SnapManager command line utility.

```

@ echo off

REM Pause SQL Server to take Snapshot
REM \<folder> is the folder containing DButil
REM DButil is called to prepare the database for backup.
REM usage: DButil <function> <SQL Server Name> <database name> <user>
<password>
REM -pre is the function to prepare the database for snapshot backup

cd \<folder>
DButil -pre VSMSQL MyDatabase dba password

REM Create new snapshot of virtual disks
REM This MUST be done for each volume that has database files

cd\
SNAPMGR SNAPSHOT "SnapD1" /PARENT:"Data1"
SNAPMGR SNAPSHOT "SnapD2" /PARENT:"Data2"

REM Resume SQL Server
REM -post option to DButil ends the SQL backup mode started by -pre

cd \<folder>
DButil -post VSMSQL MyDatabase dba password

REM Map drive letter for the new snapshot

cd\
SNAPMGR SNAPSHOT "SnapD1" /MAP:W
SNAPMGR SNAPSHOT "SnapD2" /MAP:X

REM ** Unmap and Delete old snapshots if exist **

SNAPMGR SNAPSHOT "SnapD1B" /UNMAP
SNAPMGR SNAPSHOT "SnapD2B" /UNMAP

SNAPMGR SNAPSHOT "SnapD1B" /DELETE
SNAPMGR SNAPSHOT "SnapD2B" /DELETE

Exit

```

The first command is the DButil –pre command. This command queues the disk writes to the database and allows the snapshot to be taken without any torn SQL pages. Once the database writes have been queued, you can execute the SnapManager command to create your snapshots. After you have created the snapshots, you resume the SQL database using the DButil –post command. This will commit all queued transactions to the database. Once the database is running, you can map the snapshot and delete the old snapshot.

Sample script download at:

<http://www.compaq.com/products/storageworks/solutions/vsm-sql2000/scripts.html>

License Statement: The DButil software is copyrighted by Compaq Computer Corporation. Download and use of the DButil software is governed by the terms of Compaq's license agreement

as provided. Any reproduction or redistribution of the DButil software not in accordance with Compaq's license agreement terms is expressly prohibited.

As long as the combination of "Server, Instance, and DatabaseName" is unique within the DButil command line, you can run multiple copies of DButil without any interference. Two examples are provided where multiple databases are used:

Example 1: Snapping several SQL databases

- Create a single batch job
- Perform the -pre command for each database to suspend each database
- Perform the snapshot process for each database. Do not perform a snapshot until all databases have been suspended.
- Perform the -post step for each database to resume database operation

Example 2: Snapping multiple SQL databases independently

- Create multiple batch jobs that are self contained with the -pre, snap, and -post commands associated to each unique database.
- Use a batch scheduler to automate the operation, or run the jobs manually.

Database Recovery with Snapshots

Although you can recover databases in different ways, the database administrator must decide which recovery method best suits a particular failure incident. This section describes how to use snapshots to restore a database (shown in order of best practice).

The *SANworks* Virtual Replicator Snapback feature performs restoration directly from a snapshot volume. Snapback automatically replaces your snapshot files by placing them back onto the parent virtual disk. For more information on Snapback, and how it is used, refer to Section 3.3, Performance Results – Test 4: Restore from a Snapshot using Snapback and the *SANworks* Virtual Replicator System Administrators Guide.

With snapshots, you can recover a database that has become corrupt. This means both replacing the data files and leaving the log files intact. With this process, the T-logs will automatically roll forward any transactions that occurred after the snapshot was taken. Keep in mind that if the log files contain certain activity that should not be reapplied to the database, such as inadvertent user errors, you must recover the appropriate log files, either from the log snapshot or from tape, prior to restoring the database.

Database recovery options discussed in this paper:

1. Restore database files only (without T-Logs)

This procedure is recommended in cases when you want to restore to the point in when the snapshot was created.

2. Restore database and apply transactional logs.

This procedure is recommended in cases when you want to restore up to the latest incremental T-log backup.

Procedure Option 1: Recover to a point in time using the snapshot

Restore Scenario: Data and log files need to be restored.

Step 1: Backup the transaction log files

Step 2: Detach the failed SQL Database

Use SQL Enterprise Manager or Query Analyzer to detach the database

Step 3: Snapback the snapshot to the virtual disk

Copy the data files (*.mdf files) and log files (*.ldf, if available) from the snapshot to the virtual disk using the *SANworks* Virtual Replicator Snapback function.

Consideration: Although it is recommended to create a snapshot of the *.Log files, special care should be taken with respect to log files during the Snapback process. During restores, log files need to be treated separately from data files and therefore should reside within separate snapshot volumes. A Snapback of the log file snapshot volume will replace the current log files, preventing all transactions that occurred after the log file snapshot was created from being applied. Therefore, a Snapback of the log files should only be used in instances where the administrator wishes to recover to the point in time of the snapshot.

Step 4: Attach the SQL Database

Use the SQL Enterprise Manager to attach the database once all the copy operations have completed. If you are attaching the data files only then Enterprise Manager will prompt you to create a log file for the database. In this step you will have to direct Enterprise Manager to a pre-allocated virtual disk to host the logs. Once the database is attached, the logs are automatically created from the database files.

Step 5: Create the new snapshot of the database files

Use DButil and *SANworks* Virtual Replicator to create a new snapshot from the new virtual disk and then delete the old snapshot.

Procedure Option 2: Restore the database and apply transactional log backups

Restore Scenario: Data and log files need to be restored. Multiple transaction log backups need to be rolled forward.

Step 1: Backup transaction log files

Step 2: Detach the SQL database

Use SQL Enterprise Manager to detach the database.

Step 3: Run the DButil –restoresns command using the script shown below to put the database into a loading state. In the loading state, the database is ready to accept a restore of the transactional logs. When you previously ran the DButil –pre command, a meta file was created.

The DButil -restore command uses this meta file to create the loading database. If you are restoring a SQL database to a new SQL server with a different SQL server name, you need to edit the name of the meta file to match the new name of the SQL server.

```
dbutil -restoresns <server\instance name> <database> <username>  
<password>
```

Note: Compaq recommends that you run DButil on the original LUN that is being cloned. This creates the meta file on the original LUN. Then the meta file exists on the clone when it is split off. When you run DButil, information is passed to the SQL server from the meta file. The SQL server uses this information to create a Virtual Device (the database in a loading state). You can then mount the clone database disks under the assigned drive letters. After you run the restore command and the clone is mounted to the correct drive letter, you can optionally restore sequential transactional logs in order, starting with the earliest from clone creation. For each transactional log restore, select the No Recovery option for all but the very last transactional log.

Step 4: Copy the snapshot to the virtual disk

The DButil Restoresns command will prompt the user to replace all backup files necessary to start the restore process. This is when you copy all data and log files (*.mdf and *.ldf files) from the snapshots over to the virtual disk using Virtual Replicator's Snapback feature.

Step 5: Apply Incremental T-Log backups

After you have replaced the snapshots and verified the database is in a loading state you will need to apply your transactional logs. This can be done from Enterprise Manager, Query Analyzer, or a third party backup application.

Step 6: Create the new snapshot of the database files

Use DButil and *SANworks* Virtual Replicator management interface to create a new snapshot from the new virtual disk and delete the old snapshot. Be sure to specify the 'with norecovery' option when restoring multiple log files. Only use the 'with recovery' option when restoring the final log file. This will bring back the database to a fully operational state.

Procedure Option 3: Copy / Remount

The copy / remount procedure requires a higher level of administrator intervention. Before you start this procedure, be sure that you have allocated additional capacity within the storage pool to support two virtual disks.

Step 1: Create a new virtual disk

Step 2: Copy the associated files from the previously created snapshot to the new virtual disk

Step 3: Delete the old virtual disk and related snapshots

Step 4: Remount the new virtual disk to the original database drive letter

Procedure Option 4: Operating from the Snapshot

This method uses the snapshot as the database volume. After performing this procedure, your application is operating from the snapshot only. The virtual disk is in a state of vulnerability. *It is critical that you recreate the original virtual disks and perform a complete database recovery as soon as possible and return to a standard operating condition.*

Step 1: Detach the SQL Database

Step 2: Unmap the drive letter of the virtual disk (if still exist)

Step 3: Map the same drive letter to the snapshot volume using SnapManager

Step 4: Attach the snapshots with the drive letters of the virtual disks.

This method is a temporary quick fix to have your database back up in the minimal amount of time. It is also the most risky. Should you loose the snapshot disks for some reason you will have to restore to the last full database backup which could potentially loose a lot of data.

5.2 Solution Management and Maintenance

Mandatory Update Requirement for SANworks Virtual Replicator for V3.0 installations:

This update addresses the registry growth issues that result from products utilizing Microsoft's plug and play architecture that report storage devices to the operating system. The *SANworks* Virtual Replicator V3.0A software update implements a Registry Cleanup Tool to eliminate unused *SANworks* Virtual Replicator entries that are created in the Windows Registry. This update kit also includes the functionality required to import an existing storage unit with production data into a storage pool created with *SANworks* Virtual Replicator in a Windows 2000 cluster.

This tool is available at:

<http://www.compaq.com/products/sanworks/softwaredrivers/>

Maintenance of the storage pool

It is important to monitor the pool space available and delete all unnecessary snapshots. You can use the Snapshot Manager snap-in within Microsoft Management Console or the SnapMgr command line to display information about the existing pool. When free space is less than 30 percent of pool capacity, *SANworks* Virtual Replicator logs warnings every five minutes in the system event log. When the free space drops below 5 percent of pool capacity, *SANworks* Virtual Replicator logs critical events. If you run into a free space limitation situation, you should immediately add more storage to your virtual pool.

It is also important to monitor the size of your snapshots. Snapback time is directly related to the amount of copy-outs (Delspace) on your snapshots. When recovering a database from a snapshot as described in the previous section, the amount of copy-outs to a snapshot also affects

the time it takes for the database to bring the log files and the data files to a consistent state. An administrator should monitor the amount of database transactions during the initial implementation period in order to better plan the snapshot frequency. Refer to Section 3.3, Performance Results, for information on Snapback times vs. copy-outs.

Log Management

It is important to backup the log files apart from the database files. It is recommended that you perform frequent log backups throughout the day, using the SQL Server built in backup utility. Multiple snapshots attached to log files create a performance impact on the database. In general, attach one snapshot to a log file and frequently update the snapshot to maintain minimal growth.

5.3 On-Going Service and Support

Post sales telephone support, providing problem identification and resolution for *SANworks* Virtual Replicator related issues, is available by purchasing part number: QT-66MAB-Z*. Refer to the *SANworks* Virtual Replicator Quick Spec for more information.

6.0 Related Information

6.1 Web / Documentation Links

Compaq ActiveAnswers

[Compaq ActiveAnswers for Microsoft SQL Server for Business Operations – Quick Specs](#)

Compaq ActiveAnswers Microsoft SQL Server for Business Operations is a set of information and methodologies for planning, deploying, and operating a business operations environment based on the Microsoft SQL Server database. Whether it's an order entry, human resources or accounting solution, this ActiveAnswers area provides proven guidance for the most time consuming and complex issues that impact a customers' ability to rapidly and effectively implement a business operations environment, including support for web-based transactions.

SANworks Virtual Replicator by Compaq

<http://www.compaq.com/products/sanworks/vr/index.html>

Q&A on *SANworks* Virtual Replicator

<http://www.compaq.com/products/sanworks/vr/qa.html>

MSA 1000 Array Controller

<http://www.compaq.com/products/storageworks/msa1000/index.html>

Compaq ENSA2 Vision

<http://www.compaq.com/storage/ensa2vision.html>

Compaq Microsoft SQL Server Solutions

<http://www.compaq.com/solutions/showroom/sqlserver.html>

SANworks Secure Path by Compaq

<http://www.compaq.com/products/sanworks/softwaredrivers/securepath/index.html>

6.2 Related Offerings

For positioning details, see Section 2.8 - Positioning with Other Compaq Products or Solutions.

Rapid Restore for SQL Server 2000

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

6.3 Feedback

Forward all comments, questions, and related feedback to:

Microsoft.Storage.Solutions@compaq.com.

Appendix A: Overview of Compaq SANworks Virtual Replicator

SANworks Virtual Replicator V3.0 is a key component of the Compaq Enterprise Network Storage Architecture strategy (ENSA). Virtual Replicator provides advanced, centralized storage management capabilities in Windows 2000 computing environments. Using the virtual replication snapshot feature, you can decrease the amount of time required to restore a database. Virtual Replicator also provides an additional tool that is complementary to regular online backup techniques. Using snapshot technology, you can perform integrity checks and offline copies of the snapshot without stopping the core services of the SQL Application Server. Snapshots should not be viewed as a replacement for tape backups or as a justification for unchecked growth of a SQL database.

Core Capabilities

Storage Pooling/Virtualization – enables the grouping of hardware array storage or physical disks into a logically concatenated pool of disk space. Multiple "virtual disks" can be created from the pool that behave and perform exactly like physical disks. Disk virtualization allows you to optimally tailor disk space to the size required by users and their applications. Storage devices can be added to a pool as needed, increasing the size of the pool. In addition, storage from multiple hardware array controllers can be bound into a storage pool and a virtual disk can be created that spans these storage devices. Storage units with existing production data can be seamlessly imported into a virtual storage pool. The production data is automatically preserved to a virtual disk with any free storage capacity being allocated to the pool. Individual virtual disks up to a maximum size of 2 TBs can be created.

Snapshots – enable the instantaneous creation of multipurpose virtual replicas of SQL database volumes. Snapshots function identically to ordinary physical disks with read and write capability. Snapshots are used without disrupting running applications. For example, when a SQL database is online, the snapshot can be checked for integrity or corruption. The snapshot then becomes a source for the last-known good copy if a restoration is required. Using Windows 2000 Task Scheduler, you can automate the processes for unattended creation, deletion, and integrity checking of snapshots. Virtual Replicator is also designed to integrate with scripts and schedulers for a wide range of products that are designed for Microsoft Windows 2000.

Online Volume Growth – lets you increase storage capacity without disrupting the Windows 2000 operating system. Typically when you grow a RAID set, the Microsoft Windows 2000 operating system does not recognize the size change until you reboot. The Virtual Replicator Online Volume Growth feature directs the operating system to update the size of a physical or virtual disk without a reboot.

This feature was tested in the Virtualized Storage Management for SQL Server 2000 solution, resulting in no disruption and no measurable system performance impacts.

SANworks Virtual Replicator Complements Your Backup Strategy

A snapshot is a copy of your SQL Server database files that share the same disks in the storage pool as the virtual disk. The disks in the pool should be protected with fault-tolerant RAID arrays and a battery backed up controller that protects the volatile cache memory. The snapshot should never be a replacement for a tape backup. You can combine snapshots with a daily backup to tape as an effective strategy for availability.

SANworks Virtual Replicator General Requirements

Operating System: Windows 2000 Server or Advanced Server, SP2 (minimum)

For clustered environments: Windows 2000 Advanced Server, SP2

Memory: 64MB min required

Disk space: 20MB for installation of software

SANworks Virtual Replicator can be deployed on Windows NT or Windows 2000 systems (server or workstation). It is complementary to industry backup solutions (such as VERITAS Backup Exec for Windows NT, Legato NetWorker for Windows NT, CommVault Galaxy v3.7, and Computer Associates ARCserve/IT for Windows NT), as well as NT's basic built-in backup capability. When layered on top of Microsoft Cluster Server (MSCS), Virtual Replicator takes full advantage of MSCS high availability features.

Note: Microsoft SQL Server requirements may exceed the minimum requirements for *SANworks* Virtual Replicator. Consult your Compaq sales or support representatives for assistance in optimizing your specific configuration.

SANworks Virtual Replicator Storage Management

Microsoft Management Console (MMC)

Microsoft Management Console (MMC) is a framework for running the management tools provided by Microsoft and other vendors. MMC allows management of local and remote computers. All *SANworks* Virtual Replicator functions can be managed locally or remotely using the MMC. *SANworks* Virtual Replicator provides two snap-ins for the MMC.

Snapshot Manager (SnapMgr) – Manages and presents information about pools, the virtual disks and their snapshots. For additional scripting capabilities, the SnapMgr provides command line interface (CLI) tools.

The MMC Snapshot Manager provides additional statistics to help manage your environment. Table B-1 provides descriptions for the POOL Properties Statistics Tab.

Table B-1. MMC POOL Properties Statistics Tab

Statistics Tab - MMC POOL Properties	
Read Received	The number of read I/O requests received by the pool software
Reads issued to disk	The number of read I/Os issued to disk by the pool software. This includes reads caused by copy-out operations and split reads
Split Reads	The number of read I/O requests that had to be split into two or more I/Os to disk because the read crossed a segment boundary, and the next segment was not contiguous with the current segment
Write Received	The number of write I/O requests received by the pool software
Write Issued to Disk	The number of write I/Os issued to disk by the pool software. This includes writes caused by copy-out operations and split writes
Split Writes	The number of write I/O requests that had to be split into two or more I/Os to disk because the write crossed a segment boundary, and the next segment was not contiguous with the current segment
Copy-outs	The number of segments that were copied to preserve data for snapshots

Snapshot Planner

Use the Snapshot Planner, prior to installing *SANworks* Virtual Replicator, in order to plan your snapshot allocations. You can only use Snapshot Planner with drives that are not used in a pool and are formatted with NT Disk Administrator.

Snapshot Planner helps predict the cost of snapshots, in terms of disk space and extra I/Os. It tracks actual reads and writes to your existing drives and determines how much space a snapshot of each drive needs. Snapshot Planner also determines how many extra I/Os the snapshot would cause due to copy-out operations. Install Snapshot Planner on a server that simulates production load. The data presented by Snapshot Planner will help in planning the pool space required. Plan to run the Snapshot Planner for at least 24 hours.

Copy-out is the physical copying of a 32 KB (default) segment of disk blocks in order to maintain snapshot integrity. Copy-out occurs only when a segment on the original virtual disk is modified for *the first time*. Subsequent changes to the same segment on the original disk do not require copy-outs. Disk space is allocated from the storage pool as needed for copy-outs. If parent and snapshot diverge, the parent keeps the original frames, while the snapshot is allocated new frames.

Table B-2. Windows 2000 System Monitor Counters Provided with the Snapshot Planner

Counter	Description
Writes to disk	The total number of writes issued by your applications and users.
Extra copy-out I/Os	The extra I/Os caused by the snapshot. This is the number of additional reads and writes incurred by copy-out operations.
Snapshot space (MB)	The disk space used by the simulated snapshot.
Snapshot space (%)	The disk space used by the simulated snapshot, expressed as a percentage of the space you would have used if you had made a normal copy of the disk. It is initially 0%, when you first create the snapshot. At worst it could rise to 100%.

Filter Driver: *SANworks* Virtual Replicator implements a filter driver (sddriver.sys) that intercepts and processes I/O requests bound for the underlying devices. Filter drivers are implemented under the NT I/O manager. Windows 2000 implements a filter driver for software RAID (ftdriver.sys). Other applications such as virus scanners, disk quota managers also install filter drivers. Filter drivers are integrated to use existing function calls within the Windows 2000 I/O subsystem.

SnapMgr Command Line Interface (CLI)

SANworks Virtual Replicator provides a command line interface (CLI) that enables batch jobs to automate tasks. The SnapMgr and NDMgr commands provide the same level of functionality as the two snap-ins for the MMC.

Note: SnapMgr CLI allows administrators to manage all the *SANworks* Virtual Replicator devices. The same functionality that exists locally is also available remotely.

Use these commands in scripts or from the Microsoft Windows NT Command Line.

Table B-3. SnapMgr Command Line Interface

Counter	Description
Writes to disk	The total number of writes issued by your applications and users.
Extra copy-out I/Os	The extra I/Os caused by the snapshot. This is the number of additional reads and writes incurred by copy-out operations.
Snapshot space (MB)	The disk space used by the simulated snapshot.
Snapshot space (%)	The disk space used by the simulated snapshot, expressed as a percentage of the space you would have used if you had made a normal copy of the disk. It is initially 0%, when you first create the snapshot. At worst it could rise to 100%.

Note: Remote administration from a Windows NT Workstation or Server, enables administrators or a service running with administrative privilege, to manage remote computers from the SnapMgr or NDMgr prompt.

Automated Scheduling

SANworks Virtual Replicator software includes automation and scripting capabilities that automate repetitive storage management tasks to simplify managing data and storage. A series of wizards provide a quick way to create a schedule for various snapshot activities that manage recurring tasks. Once created, these schedules automatically launch at your predefined times.

The following are some examples of tasks that can be scheduled:

- Creating and deleting snapshots
- Performing unattended backup operations
- Monitoring resource usage to ensure pool capacity is not exceeded

The schedules can be automated using Windows NT Task Scheduler or a suitable alternative.

Descriptive Terms

Storage Units

SANworks Virtual Replicator storage units are the basic building blocks for physical storage. Storage Units for a SQL Server should follow similar capacity planning exercises as a traditional disk subsystem for your primary storage. The storage unit requires physical storage that is available to the NT system. Storage units are either individual physical disks or logical disk devices (RAID sets). The disk must be unformatted and non-partitioned.

Storage Pool

A storage pool is a collection of one or more physical disks or controller-based fault-tolerant disk sets that can store virtual disks and snapshots. It is important to use controllers that protect the volatile cache memory. A maximum of 8 storage units can be pooled.

Use Storage Units with similar redundancy, read-write, and failure characteristics. Additional Storage Units can be added to a Storage Pool.

Virtual Disk

A virtual disk is a disk in a pool that was created as an empty disk of a specified capacity. To applications and users, a virtual disk looks just like a normal disk. You map a drive letter to it, partition and format it like a normal disk, and then read and write it like a normal disk. But a

virtual disk has one special feature – you can make virtual copies of it in a matter of seconds. The copies are called snapshots.

Virtual Disks are created from available disk blocks in the storage pool. The virtual disk has the same properties as a physical disk device. Multiple virtual disks (8 max) can be created from the same storage pool. Virtual disks may also use blocks obtained from different storage units. For the SQL Server database files, use blocks from a local primary Storage Unit.

When a Virtual Disk is deleted, the disk blocks are returned to the storage pool and available for reallocation.

Snapshot

The term snapshot may have different definitions for applications other than *SANworks* Virtual Replicator. A snapshot in Virtual Replicator is a disk in a pool that was created by making a virtual copy of the parent disk.

When the snapshot is first created, it is an exact virtual copy of the parent disk, made at an instant in time. It has the same capacity and label. It contains exactly the same data. It is as if you had a camera and took a picture of every single byte of data stored on the original disk at a single instant in time.

To applications and users, a snapshot looks like a normal disk. You can read and write to it like a normal disk. The snapshot can have an available drive letter mapped to it and be shared over the network. The snapshot drive appears as any other drive letter within the NT file system (NTFS).

When you first create a snapshot, it does not occupy any disk space. The snapshot only begins to take up disk space when you modify the data stored on either the snapshot or the parent disk. Subsequent snapshots can be created from the parent disk or the point-in-time snapshot replica.