



Dell EqualLogic Best Practices Series

Enhancing SQL Server Protection using Dell EqualLogic Smart Copy Snapshots

A Dell Technical Whitepaper

Storage Infrastructure and Solutions Engineering

Dell Product Group

July 2011

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2011 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerConnect™*, *EqualLogic™*, *PowerEdge™* and *PowerVault™* are trademarks of Dell Inc. *Broadcom™* is a registered trademark of Broadcom Corporation. *Intel®* is a registered trademark of Intel Corporation in the U.S. and other countries. *Microsoft®*, *Windows®*, *Windows Server®*, and *Active Directory™* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Table of Contents

1	Introduction	4
1.1	Audience	4
2	SQL Server backup considerations	5
2.1	Using Snapshots	6
3	Test system configuration	8
3.1	Storage volume layout	10
3.2	Testing SQL Server backup processing impact	11
3.2.1	Analysis and conclusion	12
3.3	Testing SQL Server snapshot processing time and impact	12
3.3.1	Time to create a SQL Server snapshot as a function of user load	13
3.3.2	Time to create SQL Server snapshots while retaining previous snapshots	13
3.3.3	Impact on system and database performance while creating SQL Server snapshots	14
3.4	Testing EqualLogic Smart Copies	15
3.4.1	Time to create EqualLogic Smart Copy snapshots	16
3.4.2	Impact on system and database performance while creating Smart Copy snapshots	17
3.5	Testing time to complete point-in-time recovery	18
3.5.1	Test 1 results: database recovery from SQL Server full backup plus log replay	19
3.5.2	Test 2 results: In-place restore from Smart Copy plus log replay	20
3.5.3	Analysis and conclusion	20
4	Best practice recommendations	24
4.1	Including EqualLogic Smart Copy snapshots as part of your SQL Server database protection strategy	24
4.2	Best Practices: Auto-Snapshot Manager/Microsoft Edition	25
4.3	General backup strategy best practices	26
4.4	Best practices when using VMware ESX Server	27
Appendix A	Auto Snapshot Manager / Microsoft Edition	29
Appendix B	Test system component details	31
	Related publications	33

Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions team between January 2011 and April 2011 at the Dell Labs facility in Round Rock, Texas.

The team that created this whitepaper:

Arun Dendukuri, Puneet Dhawan and Chris Almond

We would like to thank the following Dell team members for providing significant support during development and review:

Anthony Fernandez, Darren Miller and Suresh Jasrasaria

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information. You can submit feedback as follows:

- Use the "[Post a new thread](#)" link here:

[http://www.delltechcenter.com/page/Enhancing+SQL+Server+Protection+using+EqualLogic+Smart+C
opy+Snapshots](http://www.delltechcenter.com/page/Enhancing+SQL+Server+Protection+using+EqualLogic+Smart+C
opy+Snapshots)

1 Introduction

Backup and recovery operations are the focus of business continuity and data protection plans and often the main source of anxiety for IT departments. Few businesses are fully satisfied with their backup and recovery solutions. Not only must data be protected from complete site failures, such as those resulting from natural disasters, data must also be protected from corruption or data loss, such as that resulting from a computer virus or human error.

Properties of an ideal backup solution include:

- Data integrity is maintained during backup operations to ensure that restored data is reliable.
- Multiple copies of data are retained in safe locations, either local (for example, in the same building) or remote (in a different geographic location).
- Minimal impact on application performance and other IT operations occurs during backup processing.
- Restoration of data from backup can be accomplished quickly and effectively, with minimal impact on other IT operations or end user activities.
- The ability to support the business's RPO (Recovery Point Objective) and RTO (Recovery Time Objective).

SQL Server® full database backup is the foundation for every DBA's data protection plan. In most cases full backup is performed daily, and backup of transaction logs is completed multiple times through the day. SQL Server® supports online backups, allowing end users and SQL Server® jobs to be active while the backup operation occurs. Even so, large databases can take long time to backup. Incremental and differential backups provide a mechanism to reduce backup processing time and impact on system performance, however the restore process can still be time consuming and complex.

The goal of this paper is to present capabilities of the EqualLogic Auto-Snapshot Manager/Microsoft® Edition (ASM/ME) that in conjunction with SQL backup can help to improve RPO and RTO goals without any disruption to the database applications. ASM/ME can be used to create transactionally consistent smart copies (snapshots, clones or replicas) of SQL Server® databases. The scope of this paper is limited to EqualLogic Smart Copy snapshots only.

Note: If you are using SQL Server® database mirroring for high availability, then using EqualLogic Smart Copy snapshots in conjunction with mirrored databases can give you even more flexibility in designing efficient backup and restore strategies. We did not use database mirroring during the tests conducted for this paper.

1.1 Audience

This white paper is primarily targeted to database administrators, Storage Administrators, ESX/VMware administrators and Database Managers who are interested how they can use Dell™ EqualLogic™ storage to optimize their backup and restore strategies in Microsoft® SQL Server® environments. We assume the reader already has operational knowledge of SQL Server® backup and restore strategies, configuration and management of EqualLogic SANs and iSCSI SAN network design, and familiarity with VMware® ESX Server environments.

2 SQL Server backup considerations

A copy of data that can be used to restore and recover the data is called a backup. Backups let you restore data after a failure. Microsoft® SQL Server® enables you to back up and restore your databases. The SQL Server® native backup and restore utility provides an important safeguard for protecting critical data stored in SQL Server® databases. A well-planned backup and restore strategy helps protect databases against data loss or corruption caused by a variety of failures and meets organizations Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO). With good backups, you can recover from many different kinds of failures, such as:

- User/administrator errors such as dropping a table by mistake.
- Hardware failures such as a storage controller failure, damaged disk drives or permanent loss of a server.
- Software problems such as driver or firmware bugs that cause SQL data corruption.
- Natural disasters, etc.

As a best practice, you should regularly test your backups and check if your backup/restore strategy meets your RPO and RTO. It is also important to periodically reevaluate your business SLAs to validate your backup strategy.

The following factors can impact achievable RPO and RTO goals in a SQL Server® backup strategy:

Backup Window The time during which the backup is created. SQL Server® databases are generally backed up on a daily basis. Backups are typically scheduled during lean periods to minimize impact of backup on database performance and time to complete the backup. As we'll see later in this paper, backup processing increases CPU utilization and storage I/O on production servers. By keeping the length of the backup window (time it takes to create the backup) as short as possible you minimize the processing impact on production servers.

Backup Type and Frequency The frequency and type of backup will impact how far or close to the failure point you can restore to (RPO) and the time it'll take to bring the database online (RTO). To minimize performance impact of database backups on the production system and increase the frequency or granularity of backups, database administrators typically use either differential or more frequent Transaction Log backups (incremental) to recover SQL databases to specific points in time. To restore using transaction log backups, you need to restore the last full and differential database backup (if any) and then apply all transaction logs up to desired recovery point.

Backup Media The type of media, such as disk or tape that used for storing backup data sets. The type of media used to backup data will impact how fast you can create and restore from backups. Restoring from a disk based backup target is typically much faster than from a tape based target.

Restoration Tools and Processes The tools and processes your enterprise has in place to restore data ultimately determine what kind RTOs you are able to support. Some tools may require multiple manual steps in order to complete the restore process. A long and complicated recovery process can increase the chances for human error. The number different teams involved in data recovery can also impact restoration times.

2.1 Using Snapshots

In general, a snapshot preserves a point-in-time copy at the time the snapshot function is invoked. Snapshots can complement your backup strategy by helping you improve your achievable RPO and RTO goals. The snapshot can be used to facilitate point-in-time roll-back of a database. It can also be placed on-line so that you can recover data from it. A backup application can also utilize the snapshot as a source from which you create backups.

SQL Server® 2008 includes a native snapshot feature that database administrators can use to copy and retain the state of database at a particular time. Snapshots are also available at the storage layer. Using the storage vendor provided snapshot tools database administrators can take SQL Server® snapshots and keep point in time database copies on the SAN.

Though snapshots can be very helpful to improve overall SQL protection strategy, careful consideration must be paid when selecting the snapshot functionality to protect your SQL Server® environment. In SQL Server® environments, snapshots should have the following characteristics:

Application Consistency In a SQL Server® environment snapshots need to be able to create transactionally consistent point-in-time copies. If the snapshots are not consistent then you will not be able to perform a successful or complete recovery of the database from the snapshot. To meet this requirement, the system creating the snapshot must work in concert with SQL Server® to quiesce database activity when creating the snapshot.

Execution Speed To protect critical applications without compromising on performance and uptime, it is important that the protection tools take minimal time to create a consistent point in time copy of the data. Database snapshots creation time should be as short as possible.

Performance Impact Not all snapshots are created and maintained in the same way. For instance, due to the way that native snapshots are implemented in SQL Server®, increasing the number of snapshots retained over a period of time can have a negative impact on system performance. Ideally, snapshot functionality should have minimal impact on the application performance even if multiple snapshots are created and retained.

Easy Restoration You should be able to easily recover data from the snapshot, or restore the system to the point when the snapshot was created. The process should not be time consuming, and should allow you to accomplish a complete or partial recovery faster than if you used just backups and recovery logs.

Eliminate Backup Window using off-host backups

A good snapshot implementation should provide near instantaneous and non-disruptive consistent copy of the database. If the snapshots are transportable or can be attached to a host different from the production host, backup applications running on a proxy backup server can use them as a backup source. This allows you to off-load backup processing impact from the product host to a backup proxy host. Using this method you can schedule backups to occur at any time, without concern for production system impact.

In the following sections, we present results of a set of tests conducted by Dell Labs to do the following:

- Measure the impact of backups and SQL Server® native snapshots on SQL performance
- Test efficiency of SQL native snapshots and snapshots provided by EqualLogic SAN to keep point in time copies of the database
- Show how using EqualLogic snapshots in conjunction with existing SQL Server® backup strategies can improve overall data protection.

3 Test system configuration

To conduct the system testing detailed in this paper we created the SQL Server® test system as shown in Figure 1.

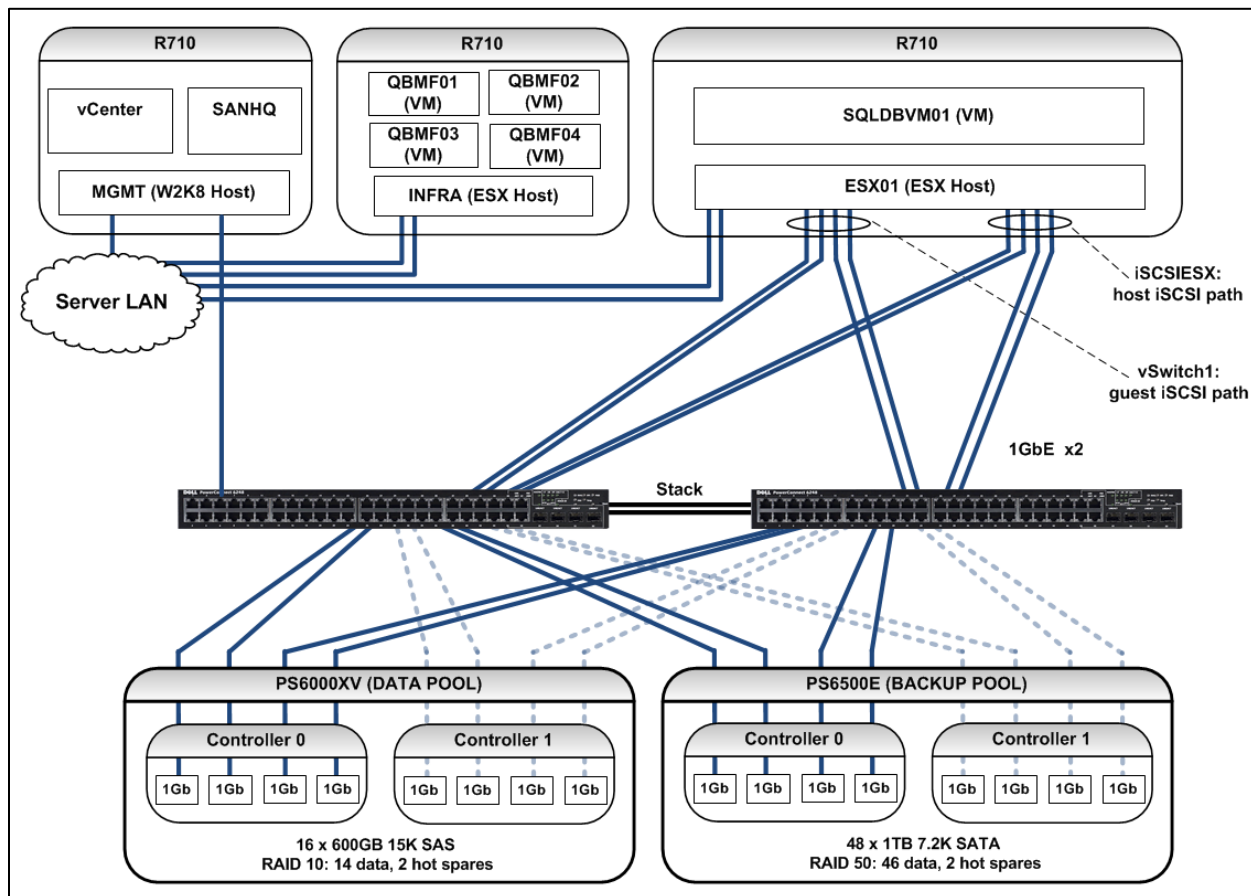


Figure 1 – Test System Configuration

Some key design details in the test system configuration:

- We Installed and configured SQL Server® 2008 R2 in a Windows Server 2008 R2 virtual machine (SQLDBVM01) hosted on the VMware ESXi 4.1 server. A Dell PowerEdge R710 server was used to run VMware ESXi server software.
 - The SQL Server® virtual machine was configured to use two virtual CPUs and 24GB of reserved memory.
 - The two local disks installed in the R710 server were configured as a RAID1 set. ESX was installed onto these disks, and the guest virtual machine OS disk partitions were also hosted within the VMFS file system on these disks.
- A second VMware ESX 4.1 server (INFRA) was used to host a group of four Windows 2008 R2 workload simulation virtual machines, with each running an instance of Quest Benchmark Factory.

- The MGMT server was a Dell PowerEdge R710 server running Windows 2008 R2 natively. It was used to host the following management and monitoring tools: EqualLogic SAN Headquarters¹ (SANHQ), VMware vCenter, remote desktop access to the workload VMs, and SQL Server® Management Studio².
- The SAN switches consisted of two Dell PowerConnect 6248 switches, configured as a single stack. As shown in Figure 1, we used redundant connection paths from the database server virtual machine via each switch in the stack. We also created redundant connection paths from each array controller to each switch in the stack.
- Network configuration details for ESX01:
 - The on-board 4-port LOM (“LAN on motherboard”) Broadcom 5709 network controller was used for the Server LAN connection paths.
 - An additional Intel Gigabit VT Quad Port network adapter was installed and used for the connection paths between the database server (SQLDBVM01) and the volumes in the Data Pool on the PS6000XV array. Via this path, the Windows 2008 iSCSI initiator within the SQLDBVM01 virtual machine was used to connect to the Data Pool volumes. This path is labeled “**vSwitch1: guest iSCSI path**” in Figure 1.
 - Another additional Intel Gigabit VT Quad Port network adapter was installed and used for the connection path between the database server and the Backup Pool on the PS6500E array. The ESX based iSCSI initiator was used in this path. A single pool was created on the PS6500E and a volume in this pool was presented to SQLDBVM01 as the backup target. This path is labeled “**iSCSI ESX: host iSCSI path**” in Figure 1.
- We used one (1) EqualLogic PS6000XV³ consisting of 16 x 600GB 15K RPM SAS disk drives in a RAID 10 configuration to host SQL Server® volumes in the Data Pool and its snapshots.
- We used one (1) EqualLogic PS6500E⁴ consisting of 48 x 1TB SATA drives in a RAID 50 configuration to host the Backup Pool data volume and other SQL Server® components.

Detailed configuration specifications for each test system component are provided in Appendix B

¹ <http://www.equallogic.com/products/default.aspx?id=5829>

² <http://msdn.microsoft.com/en-us/library/ms174173.aspx>

³ EqualLogic PS 6000XV details: <http://www.equallogic.com/products/default.aspx?id=7897>

⁴ EqualLogic PS 6500E details: <http://www.equallogic.com/products/default.aspx?id=7905>

3.1 Storage volume layout

A single pool was created from each array. Table 1 lists the volumes and Figure 2 shows the volume layout.

Table 1 – Storage volumes

Volume	Size	Purpose
Data Pool (PS6000XV RAID10)		
Datavol1	100GB	Masterdata, msdb and model database
Sqlldb2-vol1	100GB	Database files
Sqlldb2-vol2	100GB	Database files
Sqlldb2-vol3	100GB	Database files
Sqlldb2-logvol	100GB	Transaction Log files
Backup Pool (PS6500E RAID 50)		
Bkpvol1	1024GB	Store Database backup files and backup of transaction logs
sqlsvr	100GB	SQL Server® software binaries
tempvol1	100GB	SQL Server® tempdata
templates	1024GB	Virtual machine templates




Figure 2 – Storage Volumes

3.2 Testing SQL Server backup processing impact

In this test we measured the impact on server and database performance when running the SQL Server® native backup utility to create a full database backup across the SAN fabric to a volume hosted on a second storage array.

Test details:

- Database size at the Full Backup point: 130GB
- Database Workload: During the test timeline Quest Benchmark Factory was used to simulate a TPC-C style workload. A constant 4000 user workload was simulated, generating an average I/O load of 220 transactions per second against the database.
- During this test 24GB of RAM and two CPUs were allocated to the database server VM, providing aggregate computing power of 4.52 GHz.

Figure 3 illustrates the data flow across the SAN during creation of the database backup.




Figure 3 – Database backup data flow

In Figure 4 we show the impact on CPU utilization during the backup process.




Figure 4 – CPU utilization impact when creating SQL Server backup using native backup utility

3.2.1 Analysis and conclusion

As expected, we saw a significant increase in CPU load on the system during the time that the backup utility was creating the backup. We also measured a 30% average increase in application response time during the same period.

Based on the test results it is evident that running backups during peak periods will have direct impact on:

- The SQL Server® host performance: backup processing CPU cycles may impact other product workloads.
- Database performance: increased query response times.

As illustrated above, this extra load on the production host is what database administrators want to avoid, hence organizations try to find backup windows or periods of low activity during which backups can be completed. Using off-host backups as described earlier in this paper the CPU cycles consumed by backup processing can now be used by the SQL Server® workload.

3.3 Testing SQL Server snapshot processing time and impact

In this series of tests we measured the impact on server and database performance when running the SQL Server® native snapshot utility to create point-in-time copies of a database.

The snapshot utility in SQL Server® 2008 allows you to create point-in-time copies of a database. SQL Server® snapshots are read-only. Multiple snapshots can exist for a single source database. The snapshots must always reside on the same server instance as the source database. Each database

snapshot is transactionally consistent with the source database at the time the snapshot is created. A snapshot will persist until it is dropped.

3.3.1 Time to create a SQL Server snapshot as a function of user load

In the first test we measured how long it took to complete snapshot creation at different simulated user loads. Test details:

- The same system configuration was used as in the previous section.
- The TPC-C style workload was simulated while varying user load in increments from 100 to 4000 users. (The time to create the snapshot for the 100 user workload was approximately 1 second. This data point is not included in Figure 5.)
- No prior snapshots existed when each of the snapshots in this test was created. For each data point, we dropped the previous snapshot before creating the next one.

The results of this test are presented in Figure 5.




Figure 5 – Time to create SQL Server snapshots at different user loads

The result of this test clearly indicate that time window required to complete creation of a SQL Server® database snapshot depends on the system workload at the time the snapshot is created.

3.3.2 Time to create SQL Server snapshots while retaining previous snapshots

In this test we measured the time to create SQL Server® snapshots while the test system was under load. The goal of this test was to measure the cumulative performance impact on the system during

snapshot creation, while retaining multiple snapshots. Differences between this test and the previous one:

- During creation of all snapshots in the series we ran a constant TPC-C style workload on the system, simulating 4000 users.
- We did not drop any previous snapshots during the test sequence.

The results of this test are presented in Figure 6.

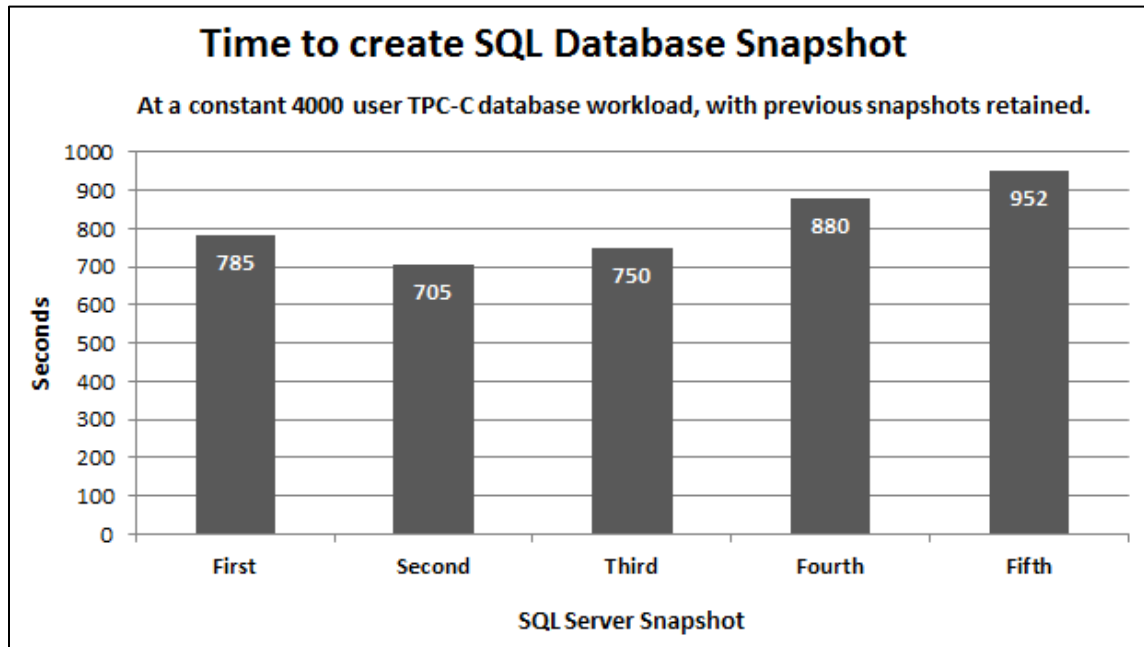


Figure 6 – Time to create SQL Server snapshots (previous snapshots retained)

Based on the results of these tests, it is clear that with varying load and number of snapshots kept, the time to take a snapshot is variable and hence difficult to predict. Minimizing number of snapshots retained would decrease the time it takes to complete the snapshot operation.

3.3.3 Impact on system and database performance while creating SQL Server snapshots

During the cumulative snapshot creation test in the previous section we also measured system CPU utilization, database transactions per second (TPS) and average database response time. The results we gathered are shown in Figure 7. The top chart in Figure 7 shows CPU utilization during the time when we created the five SQL Server® snapshots discussed in the previous section. The bottom part of Figure 7 shows the transient impact on database transactions per second (TPS) and average response time during the creation of the snapshot. From the results we can conclude that:

- Retaining multiple SQL Server® database snapshots will incrementally increase the workload on the system. By the time the fifth snapshot was created, average CPU utilization had increased by ~30% over what it was before the first snapshot was created.
- Creating SQL Server® snapshots causes significant short term impact on database performance. As shown in Figure 7, we observed a large increase in average response time and a large decrease in TPS during creation of the snapshot.




Figure 7 – System and database performance impact when creating SQL Server snapshot

In the next section, we now look at how snapshots at EqualLogic SAN layer perform under similar test conditions.

3.4 Testing EqualLogic Smart Copies

EqualLogic Auto-Snapshot Manager/Microsoft® Edition (ASM/ME) gives a SQL Database administrator the ability to create a “Smart Copy Snapshot” of the database. A “Smart Copy Snapshot” is a transactionally consistent point in time copy of the SQL Server® database. Please see **Appendix A** for more details about EqualLogic Auto-Snapshot Manager/ Microsoft® Edition.

In this section we present the results of the following system tests:

- A. Time to create an EqualLogic “Smart Copy” snapshot of a SQL Server® database
- B. System impact during Smart Copy creation

3.4.1 Time to create EqualLogic Smart Copy snapshots

Test details:

- As in the previous tests, we used Quest Benchmark Factory to simulate a TPC-C style workload. A constant 4000 user workload was simulated, generating an average I/O load of 220 transactions per second against the database.
- The Smart Copy snapshots consumed snapshot reserve assigned to the same volume hosting the database.
- Previous snapshots were retained before taking new snapshots.

The results of this test are shown in Figure 8.




Figure 8 – Time to create ASM Smart Copy Snapshots

The time reported in Figure 8 is the total time to complete snapshot operation as reported in the SQL error log.

Comparing the results of this test to the SQL Server® snapshot test results shown in Figure 6, we see that the time required to create EqualLogic Smart Copy snapshots is much less than the time required to create native SQL Server® database snapshots:

Average time to create a snapshot at a 4000 user workload while retaining previous snapshots	
SQL Server® database snapshots:	814.4 seconds
ASM/ME Smart Copy snapshots:	4.8 seconds

3.4.2 Impact on system and database performance while creating Smart Copy snapshots

We also measured system and database performance impact while creating a series of Smart Copy snapshots. Figure 9 shows database server CPU utilization and database transactions per second (TPS) while creating a series of ASM/ME Smart Copy snapshots. Figure 10 shows the difference in average application response time, as measured from Quest Benchmark Factory, during creation of SQL Server® snapshots vs. ASM/ME Smart Copies.

We can draw the following conclusions from the results shown in Figure 9 and Figure 10:

- We measured no significant increase in CPU utilization during creation of ASM/ME Smart Copy snapshots of SQL Server® databases under a heavy transactional workload.
- Transactions per second (TPS) was also not affected during creation of Smart Copy snapshots.
- We did not measure a noticeable increase in average response time when creating ASM Smart Copy snapshots. We measured a much higher transient increase in average database response time during creation of native SQL Server® snapshots than during creation of ASM/ME Smart Copy snapshots.




Figure 9 – Performance impact while creating a series of ASM/ME Smart Copy snapshots




Figure 10 – Comparing response times during creation of SQL Server snapshots and ASM/ME Smart Copies

In section 3.4, we illustrated the non-disruptive nature of EqualLogic Smart Copy Snapshots on SQL Server® performance and time efficiency of smart copies to quickly create consistent protection points. Backup applications such as Symantec Backup Exec or CommVault Simpana running on Dell DL2200 backup appliance take advantage of EqualLogic snapshots to backup SQL data from a snapshot on EqualLogic SAN rather than from data volumes attached to production host. By leveraging EqualLogic snapshots, database administrators can create more recovery points without waiting for the backup window, thereby improve RPO. As we show in subsequent sections, EqualLogic snapshots can also improve RTO of SQL Server® applications.

3.5 Testing time to complete point-in-time recovery

In this test we measured and compared the time required to complete point-in-time recovery or achievable RTO of the database using two different recovery methods:

- A. Recovery using a SQL Server® full backup and SQL Server® log replay
- B. Recovery using an EqualLogic Auto-Snapshot Manager/Microsoft® Edition Smart Copy Snapshot and SQL Server® log replay

Test details:

- Database size at the Full Backup point: 130GB
- Database Workload: During the test timeline Quest Benchmark Factory was used to simulate a TPC-C style workload. A constant 4000 user workload was simulated.
- Transaction log or incremental backups were created at 15 minute intervals
- Test 1:
 - We simulated a database corruption event at approximately 3 hours and 20 minutes after the last full backup was created.
 - We used the SQL Server® native backup utility to perform the restore and recovery process using the last full backup plus necessary log backups.

- Test 2
 - We performed a system recovery to the same point in time as in test #1.
 - In this test case we restored using the most recent Auto-Snapshot Manager snapshot Smart Copy (using Auto-Snapshot Manager in-place restore with the apply logs option).
 - After completing the in-place restore we replayed necessary incremental or transaction log backups from that point forward using the SQL Server® native backup utility.

Figure 11 below illustrates the recovery timeline components involved in each test. Note in Figure 11 the difference in recovery set components that are needed to complete database recovery (highlighted for each method).




Figure 11 – Recovery point timeline comparison

3.5.1 Test 1 results: database recovery from SQL Server full backup plus log replay

The total time required to complete database restore, recovery and place it back on-line was **41.5 minutes**. The time required to complete each component of the process is shown in Table 2.

Table 2 – Database recovery from SQL Server full backup plus log replay

Recovery Task	Time to complete (seconds)
Restore full database (130 GB) to T0	900
Replay of 13 transaction Logs	1560
Recover database	29
Total	2489 (41:29)

3.5.2 Test 2 results: In-place restore from Smart Copy plus log replay

The total time required to complete database restore, recovery and place it back on-line was **5.5 minutes**. The time required to complete each component of the process is shown in Table 3.

Table 3 – Database recovery using Smart Copy in-place restore plus log replay

Recovery Task	Time to complete (seconds)
Smart Copy restore	168
Replay of 1 transaction Log	135
Recover database	25
Total	328 (5:28)

3.5.3 Analysis and conclusion

In the case of database corruption, data loss events or other recovery scenarios that require restoring from a known good state, the results of this test clearly show the potential for achieving significantly faster system recovery times when using EqualLogic ASM/ME Smart Copy snapshots vs. using just the SQL Server® native backup utility.




Figure 12 – Results of recovery time comparison

Conclusions from the results presented in Figure 12:

- The ASM Smart copy based recovery process was approximately 7.5 times faster than the recovery process using the native SQL Server® backup utility for recovering the database to the same point (which involves restore plus log replay).
- Instead of a typical restore process that involves first restoring a full backup (and then any differential or incremental log backups) in a recovery scenario, restoring to a point in time using EqualLogic smart copy snapshot is much more efficient.
- We did not use SQL Server® differential backups in this test. Differential backups are commonly used, and they can help to accelerate recovery times just as we are showing with use of Smart Copies. Even if SQL differential backups were created at the same times that ASM Smart Copies were, the Smart Copy based method would still have been at least 3 times faster, based on comparison to just the full backup restore and log backup recovery time components.

To explain why the EqualLogic based ASM Smart Copy method provides much faster recovery times it is important to understand the difference in data flow that occurred during each of the recovery processes we tested. For Test 1, the copy-restore processes for the full backup and for the transaction log backups both read from files stored in the backup data volumes in the backup pool. In Test 2, the in-place restore of the Smart Copy snapshot does not cause data to be copied across the SAN from the backup pool to the SQL Server® data pool. Instead, the Smart Copy in-place restore just restores the data volume hosting the database to the most recent volume snapshot. No read/write I/O operations are required to do this. This is why an in-place restore from a Smart copy will typically take much less time to complete than a full database restore when using the SQL Server® native backup utility. This is illustrated in the difference between data flow patterns between Test 1 and Test 2, as shown in Figure 13 and Figure 14 below.

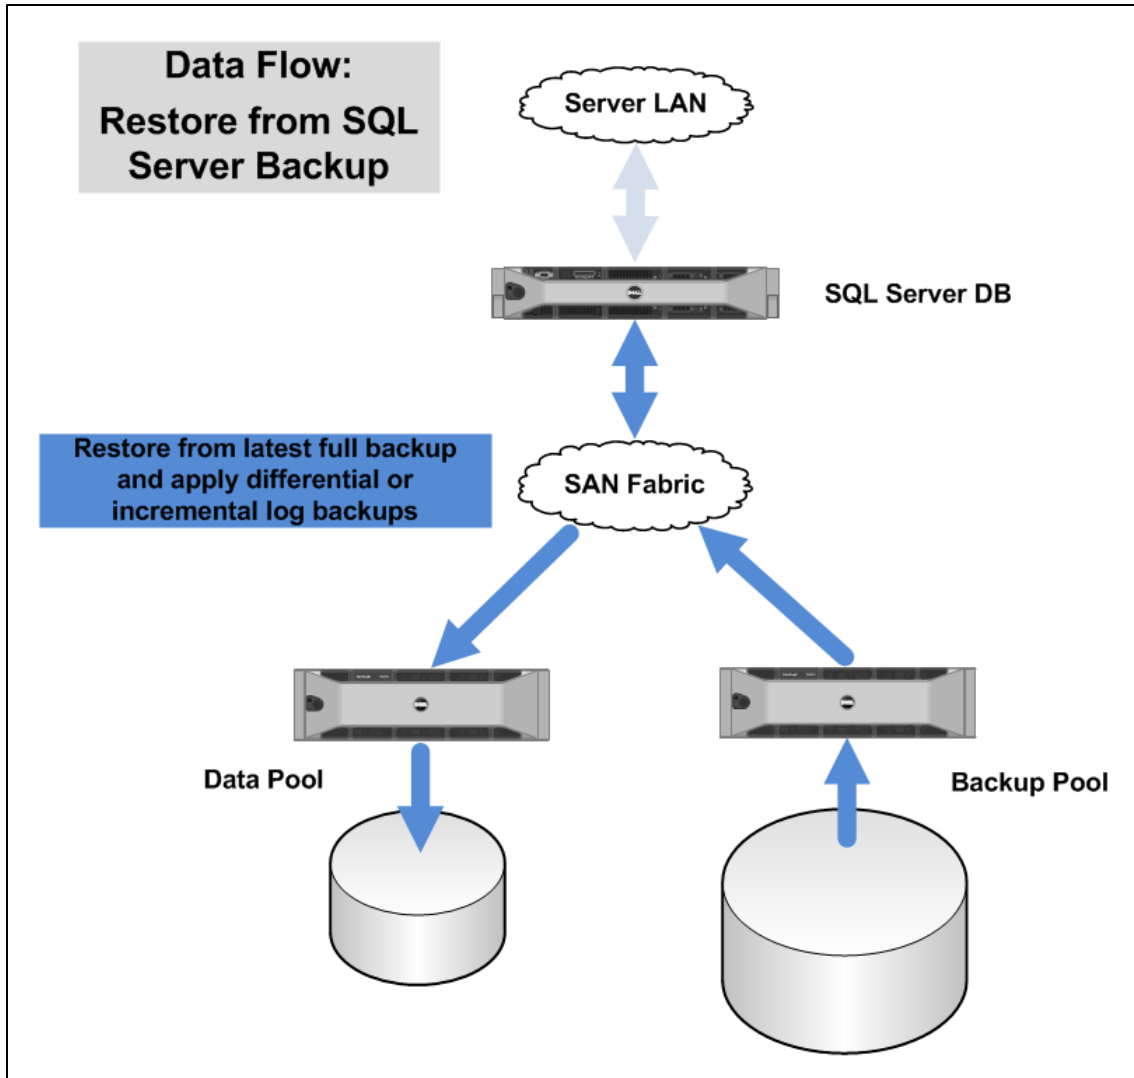


Figure 13 – Restore and recovery data flow path using native SQL Server backup utility

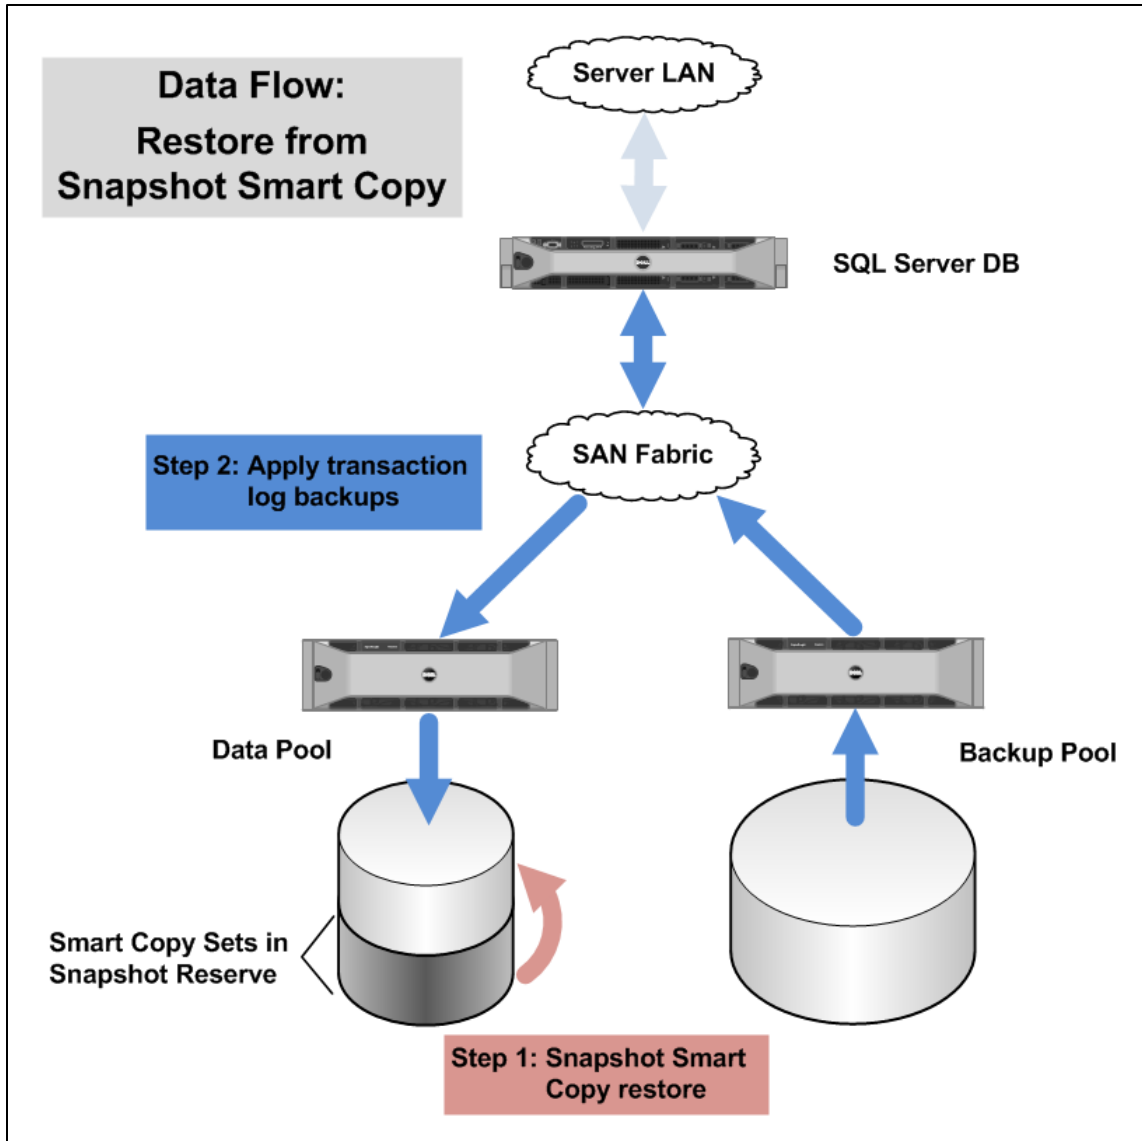


Figure 14– Restore and recovery data flow path using EqualLogic Smart Copy in-place restore

4 Best practice recommendations

In the following sections we summarize best practices to follow for using Auto-Snapshot Manager and for SQL backup strategies in general.

4.1 Including EqualLogic Smart Copy snapshots as part of your SQL Server database protection strategy

In order to include EqualLogic smart copies or snapshots in your overall SQL protection strategy and measure the impact on achievable RPO and RTO, a few considerations are necessary: You should pay attention to the following considerations when including EqualLogic smart copies or snapshots in your overall SQL protection strategy.

Know your workload:

It is important that you understand the performance characteristics of your database workload. If you plan to use snapshots to backup SQL data by attaching the snapshots to a backup server (off-host backup), be aware that the backup process will place an additional IOPS and throughput load on your SAN. You should make sure that the SAN can handle the additional load (largely sequential reads) that is typically created by backup processing. Test your backup process and measure the additional load on SAN by using the EqualLogic SAN Headquarters (SANHQ) monitoring tool.

Know your recovery times from Smart Copies:

You should measure recovery times when restoring from a backup set vs. restoring from an Auto-Snapshot Manager Smart Copy. This information will allow you to plan for optimal Recovery Time Objectives (RTOs). We measured and compared recovery times for each method. The results from that test are provided in *Section 3.5, Testing time to complete point-in-time recovery*, on page 18.

Monitor your snapshot reserve utilization:

ASM Smart Copies create array based snapshots. Therefore, Smart Copies require additional space allocation (snapshot reserve) in the EqualLogic volume to store any data changed after a snapshot was created. The default value of volume snapshot reserve space is 100% of volume size. The actual reserve setting may be higher or smaller than that, depending on your workload profile and the number of snapshots you want to keep. It is important to carefully monitor reserve utilization and tune the reserve setting for optimal storage utilization. The number and size of Smart Copies you can retain is limited by the volume snapshot reserve setting. If you reach the reserve limit you will need to increase the reserve setting or delete an older Smart Copy (or array based snapshot) before you can create another Smart Copy. A good way to measure reserve utilization is to start with the default (%100) reserve allocation and create a series of array based snapshots every hour during a period of normal production workload. While creating the snapshots use SANHQ to monitor how reserve usage increases. Note that when you perform a point-in-time recovery from a Smart Copy, an additional snapshot of the volume hosting the database is first created and then the database is restored to the point when the Smart Copy was created. Therefore, it is important to include one additional snapshot in your snapshot reserve sizing calculation.

Carefully assess your RPO goals:

Depending on your target RPO goals, you should ascertain whether you can fulfill them with just Smart Copies alone or would need a combination of backup sets in conjunction with Smart Copies. For example, consider the following case:

- You cannot tolerate a data loss period of more than 15 minutes
- The size of your SAN volumes combined with the data change rate limits you to creating one snapshot per hour with a maximum snapshot history of four before you have to delete the oldest Smart Copy snapshot.
- You run multiple backup schedules against the same database: regular full backups once per day, Smart Copy snapshots once per hour, transaction log backups every 15 minutes, and off-host differential backups once every four hours.

In a recovery scenario, you could revert back to the closest Smart Copy (if the data loss occurred within 4 hours of starting the recovery process and then apply the transaction log backups to roll the database forward to the point of data loss. This way you will restore much quicker than if you used backup sets only.

Consider other methods to meet aggressive RTO/RPO goals:

If you have very aggressive RTO and RPO goals that cannot be met with recovery from a backup set or a combination of backup set recovery and ASM/ME Smart Copies, then you will need to consider other options: failover clustering, database mirroring, log shipping, or data replication.

Understand the difference between regular EqualLogic volume snapshots and ASM/ME Smart Copies:

You can also create a snapshot of a volume using EqualLogic group management functions directly. If you create a snapshot of a volume hosting an active SQL Server® database instance using this method, then the resulting snapshot will not provide the same level of data consistency that an ASM/ME Smart Copy snapshot does. Group level snapshots can still be considered “crash consistent”. Database activity is not quiesced before the snapshot is created, therefore there is no guarantee that the snapshot will support a clean recovery the way ASM/ME Smart Copy snapshots can.

4.2 Best Practices: Auto-Snapshot Manager/Microsoft Edition

Auto-Snapshot Manager utilizes EqualLogic volume snapshot features. You should follow the best practices best practices in this section when using ASM/ME.

Planning snapshot reserve for ASM Smart Copies

- Monitor snapshot reserve space utilization. If the snapshot reserve fills up, then you will need to either increase the snapshot reserve allocation for the SQL data volumes or delete older snapshots to free up reserve space.
- If you are not sure how to set the size of snapshot reserve for SQL Server® data volumes, then you could start by using the default reserve allocation for the volume. To determine if the reserve size is set optimally you can create a series of snapshots while the system is under load and monitor how quickly the reserve space is consumed by write activity.

Database Layout

- As a best practice, make sure that SQL Server® databases and associated transactions logs should each reside in dedicated volumes. This will prevent creation of “torn” Smart Copies⁵ during selective restore scenarios. This will also enable faster restore and recovery times for individual SQL Server® databases.
- The Volume Shadow Copy Service (VSS) imposes a limit on the time it takes to create a Smart Copy. Smart Copies involving many volumes might exceed this time limit (30 seconds). If this limit is exceeded then ASM/ME displays an error message stating that VSS can no longer hold I/O writes. To resolve this problem, the best practice is to reduce the number of volumes included in the Smart Copy. You might also have to reduce the number of volumes assigned to collections.

When creating ASM Smart Copies

- If your intention is to restore smart copy on a temporary basis for data validation or to restore lost data then Side-by-Side restore of Snapshot Smart Copies is the best option.
- If you are planning to use a Smart Copy to do an “in place” restore of the database and also need to roll forward using transaction logs, then you should make sure that all required log backups are available before proceeding.
- Use the latest EqualLogic controller firmware. If using EqualLogic firmware 5.0 or later, selective restore of database leverages the data copy offload feature on the EqualLogic SAN. This feature bypasses the file system’s copy operation while restoring data, Therefore, it completes much faster than a traditional copy operation that involves data copy from SAN to host and back to SAN.
- By default, backup documents created by Auto-Snapshot Manager are saved as files with a “.bcd” extension. After a smart copy set is imported, the backup document extension is changed to “.imported”, which indicates that the backup document cannot be used again. The location for backup documents should be managed on a central share that can be backed up regularly. We recommend that you store backup documents on a file share accessible to both computers.
- To promote a snapshot/clone Smart Copy on a different computer than the one that created it, the computer must be running Auto-Snapshot Manager and must have access to the secondary group storing the snapshot/clone set. It is recommended to give access on those volumes to the host where the database is transported.

4.3 General backup strategy best practices

- Use of snapshots, whether they are software based or storage based, should not be considered as a substitute for a traditional backup and restore strategy. You should first design and implement a backup and restore strategy that is not snapshot based. After that, you should consider how to use snapshot based tools (such as the Auto-Snapshot Manager Smart Copy feature highlighted in this report) to compliment your backup/restore strategy to improve flexibility in meeting specific aspects of your RPO and RTO requirements.

⁵ See the “Multiple Databases on Volumes” section in the *Auto-Snapshot Manager Microsoft Edition User Guide* for an explanation of torn Smart Copies.

- Careful scheduling is necessary when you have different backup and snapshot schedules running against the same database.
- It is very important that backup destinations reside on storage devices that are separate from where the databases are stored. I/O bandwidth can be very critical to the performance of the database. You want to separate the sequential I/O load that is typically generated by backup and restore processing from the random I/O load that is typically generated by transaction processing databases.
- Use the CHECKSUM option of the Backup command. With this option enabled, the backup will verify the page checksums (if they are present), and generate a separate backup checksum for the backup stream that is stored on the backup media. This option will also cause both the backup processing workload and the amount of time to create the backup to increase.

Example:

```
BACKUP DATABASE sqlldb2
TO DISK = N'G:\Backup\sqlldb2_bak
WITH CHECKSUM
```

- You should always schedule backup operations when database activity is low.
- Back up first to disk whenever possible. Backing up to disk will greatly increase the performance of the backup process and free the resources of SQL Server®. Using file backups also simplifies the restoration process.
- To back up a database that has the database files damaged, use the NO_TRUNCATE or the COPY_ONLY and CONTINUE_AFTER_ERROR options of the BACKUP command.
- Test Backup files periodically using the RESTORE VERIFYONLY command. This command will verify the backup, but not restore it. It checks that the backup set is complete and readable. This command does not verify the structure of the data contained in the backup volume.

```
RESTORE VERIFYONLY
FROM DISK = N'CG:\Backup\SQLDB2-032211.bak '
WITH CHECKSUM
```

4.4 Best practices when using VMware ESX Server

In our lab test environment we used VMware ESX server to host SQL Server® database virtual machines as well as the Quest Benchmark factory work load simulation virtual machines. We share the following best practice recommendations below for running VMware ESX based virtual machines in conjunction with EqualLogic Storage and/or Microsoft® SQL Server® environments.

ESX host configuration

- We recommend in any configuration where you are using the ESX host based iSCSI initiator that you evaluate and take advantage of EqualLogic aware connection and path management by installing and using the EqualLogic Multipathing Extension Module⁶ (MEM) for vSphere 4.1. We used MEM to optimize I/O performance for the connection path between the “iSCSI ESX”

⁶ See the *EqualLogic Multipathing Extension Module Installation and User Guide* for vSphere version 4.1, available here: https://www.equallogic.com/support/download_file.aspx?id=947

host iSCSI connection path between the SQL database server and the PS6500E backup pool array.

- Per Microsoft® SQL Server® storage requirements, SQL Server® guest operating system images must be deployed on physical disk drives separate from physical drives hosting SQL Server® data. The VMDK files containing the database server OS file systems were stored on the local disks installed on the R710 ESX host.
- You should configure separate virtual switches for VM network traffic and iSCSI storage traffic on the ESX hosts.
- Jumbo frames should be enabled on vSwitches handling iSCSI SAN traffic. At least two server NICs dedicated for iSCSI traffic need to be configured as uplink NICs to the iSCSI vSwitch.

Virtual machine and guest OS configuration

- In order to use the Auto-Snapshot Manager/Microsoft® Edition features, you must setup iSCSI SAN storage access for Windows based virtual machines to use a direct access path and the guest OS (Windows) iSCSI initiator, as illustrated in Figure 1 and Figure 2.
- When using the Windows 2008 Server iSCSI initiator within a virtual machine (guest iSCSI), the following recommendations apply:
 - Create virtual NICs of type vmxnet3 within the guest VM for connection to iSCSI virtual switches.
 - Enable TSO (TCP Segmentation Offload) and LRO (Large Receive Offload) in the guest VM NICs for iSCSI traffic.
 - We recommend you use the EqualLogic MPIO DSM installed as part of the EqualLogic Host Integration Toolkit (HIT Kit) in the guest OS.

Appendix A Auto Snapshot Manager / Microsoft Edition

This appendix provides a brief introduction to EqualLogic Auto-Snapshot Manager/Microsoft® Edition. For more in depth information, please see the following references:

- **Auto-Snapshot Manager/Microsoft Edition User Guide** (v3.5.1):
<https://www.equallogic.com/support/download.aspx?id=10243>
(registered support.equallogic.com ID required for access)
- **SQL Server Database Protection Using Auto-Snapshot Manager / Microsoft Edition: Advanced Operations:**
<http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=1143>
(registered support.equallogic.com ID required for access)

EqualLogic ASM/ME is a Microsoft® management Console snap-in tool that enables you to create and manage Smart Copies. Using ASM/ME you can create three types of Smart copies: snapshots, clones or replicas.

ASM/ME uses the Microsoft® Volume Shadow Copy Service (VSS). VSS provides a framework for backing up and restoring data in the Windows server environment. Using ASM/ME, you can quickly create fast, coordinated and consistent copies of SQL database volumes residing in an EqualLogic PS series group.

The relationship between ASM/ME and the Microsoft® VSS copy service is in the following figure.

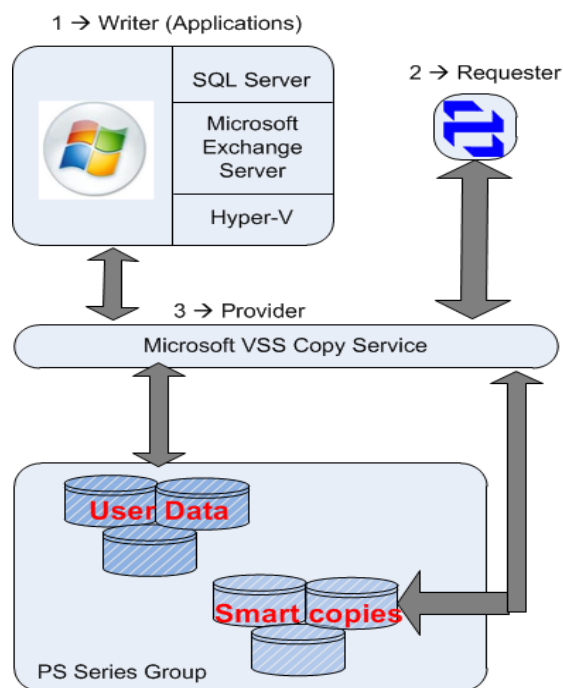


Figure 15 - ASM Integration

Auto-Snapshot Manager manages the interaction with SQL Server® to prepare the database for the Smart Copy operation. When you create an ASM Smart Copy, SQL Server® first places the database in a consistent state, and then Auto-Snapshot Manager creates the Smart Copy. The result is a data-consistent point-in-time copy (snapshot, clone, or replica) of the SQL Server® database. Auto-Snapshot Manager also manages recovery of SQL Server® databases. Since ASM is application aware, it automatically recognizes all EqualLogic volumes that are used by one or more SQL Server® instances. When you create a smart copy using ASM/ME, all volumes that are part of a SQL Server® instance are included in the smart copy operation.

Important: You can also create snapshots, clones, or replicas of volumes using EqualLogic group management functions directly. If you create a snapshot of a volume hosting an active SQL Server database instance using this method, then the resulting snapshot will not provide the same level of data consistency that an ASM/ME Smart Copy snapshot does. Group level snapshots can still be considered "crash consistent". But, database activity is not quiesced before the snapshot is created, thus they cannot provide the clean recovery capability that ASM/ME Smart Copy snapshots can.

Appendix B Test system component details

This section contains an overview of both the hardware and software configurations used throughout the testing described in this document.

Table 4 – Test Configuration Hardware Components

Test Configuration – Hardware Components	
SQL Server® (ESX01)	<p>One (1) Dell PowerEdge R710 Server running VMware ESX v4.1, hosting a single SQL Server® Database virtual machine:</p> <ul style="list-style-type: none"> • BIOS Version: 2.1.15 • 2 x Quad Core Intel® Xeon® E5520 Processors • 96 GB RAM, 2.26 GHz • 2 x 146GB 10K SAS internal disk drives • Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 5.2.7, driver version 5.2.14 • Two (2) Intel Quad Port VT network adapters (Intel 8257 1Gb). Firmware level 1.3.19.12.
I/O Workload Generators (INFRA)	<p>One (1) Dell PowerEdge R710 Server running VMware ESX v4.1, hosting a four (4) Windows Server 2008 R2 virtual machines:</p> <ul style="list-style-type: none"> • BIOS Version: 2.1.15 • Quad Core Intel® Xeon® X5570 Processor • 96 GB RAM, 2.26 GHz • 2 x 146GB 10K SAS internal disk drives • Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 5.2.7, driver version 5.2.14
Network	<p>2 x Dell PowerConnect 6248 1Gb Ethernet Switch</p> <ul style="list-style-type: none"> • Firmware: 3.2.0.9
Storage	<p>1 x Dell EqualLogic PS6000XV:</p> <ul style="list-style-type: none"> • 14 x 600GB 15K RPM SAS disk drives as RAID 10, with two hot spare disks • Dual quad-port 1GbE controllers running firmware version 5.0.2 <p>1 x Dell EqualLogic PS6500E:</p> <ul style="list-style-type: none"> • 148x 1TB SATA drives as RAID 50, with two hot spare disks • Dual quad-port 1GbE controllers running firmware version 5.0.2

Table 5 – Test Configuration Software Components

Test Configuration – Software Components	
Database Server VM (SQLDBVM01)	Windows Server 2008 R2 Enterprise Edition <ul style="list-style-type: none">• EqualLogic Host Integration Toolkit (HIT) v3.4.2• EqualLogic Auto-Snapshot Manager/Microsoft® Edition v3.4.2• SQL Server® edition / version details2008 R2
Workload Servers (QBMF01-04)	8 x Windows Server 2008 R2 Enterprise Edition Workload generators (running within VMs): <ul style="list-style-type: none">• Quest Benchmark Factory version 6.1.1
Monitoring and Management	EqualLogic SAN Headquarters version 2.1 VMWare vCenter version 4.0 Microsoft® SQL Server® Management Studio

Related publications

The following Dell publications are referenced in this document or are recommended sources for additional information.

- ***Microsoft SQL Server Database Protection Using EqualLogic Auto-Snapshot Manager / Microsoft Edition***
<http://www.equallogic.com/WorkArea/DownloadAsset.aspx?id=5247>
- ***EqualLogic Configuration Guide***
<http://www.delltechcenter.com/page/EqualLogic+Configuration+Guide>
- ***Auto-Snapshot Manager/Microsoft Edition User Guide*** (v3.5.1):
<https://www.equallogic.com/support/download.aspx?id=10243>
(registered support.equallogic.com ID required for access)
- ***SQL Server Database Protection Using Auto-Snapshot Manager / Microsoft Edition: Advanced Operations***:
<http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=1143>
(registered support.equallogic.com ID required for access)



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.