Dell EqualLogic Best Practices Series

# Microsoft SQL Server 2008 Backup and Restore using Dell EqualLogic and PowerVault DL2200 with CommVault Simpana

A Dell Technical Whitepaper

**Storage Infrastructure and Solutions Engineering**
**Dell Product Group**
**March 2012**

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.

# Table of Contents

# Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions of Dell Inc.

The team that created this whitepaper:

**Lakshmi Devi Subramanian, Puneet Dhawan,** and **Camille Daily**

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

**Ananda Sankaran, Bob Ganley,** and **Rob Young**

# Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information by sending an email to SISfeedback@Dell.com.

SISfeedback@Dell.com

# Executive summary

Finding backup windows of low activity to reduce the extra load on a production server during backup operations has been a concern for many organizations. One way of achieving this is by off-loading backup operations to a proxy server that can directly read source data through the SAN and write to backup target, a process termed as off-host backup. This paper provides best practices and reference architecture for off-host server backup operations of virtualized Microsoft® SQL Server® 2008 R2 databases, running on Dell™ EqualLogic™ SAN, and using the Dell PowerVault™ DL2200 with CommVault® Simpana®.

CommVault Simpana's SnapProtect™ feature is a snapshot backup solution, which enables backup of application consistent storage snapshots of source volumes to target devices. With EqualLogic storage it can be configured to read volume snapshots from source PS series SAN and backup to storage attached to PowerVault DL2200 or any other disk or tape based backup target. The source volume snapshots on PS series SAN are created in an application consistent manner by leveraging the Microsoft's Volume Shadowcopy Service (VSS) framework on the production host SQL Server.

From tests and data analysis, we conclude in this paper that SnapProtect greatly reduces the impact on the resources at the production SQL server during backup when compared to those performed using native SQL Server Backup utility. During full backup, the average CPU utilization percent at production SQL server:

- Remained the same at 1.5 % before and during full backup when using SnapProtect, where as
- Increased from 1.7% to 13% during full backup when using native SQL Server Management Studio for backup.

Considering performance, the true value of SnapProtect was seen when multiple jobs were run at the same time. Running four database backup jobs in parallel compared to four database backups performed in sequence:

- Increased the average backup throughput by 94%, and
- Improved the backup duration by 45%.

Both source and target Deduplication features on DL2200 appliance offered significant capacity savings on subsequent full backups at the backup target array. The capacity savings during backups with either source or target deduplication

- Improved by 96% from the original backup data size to the third full deduplicated backup after data change, and
- Improved by 99% at the fourth full deduplicated backup when there was no data change.

# 1 Introduction

The PowerVault DL2200, backup to disk appliance powered by CommVault is designed to address data management, protection, and recovery pain points by delivering a disk-based backup and archive solution with deduplication and built-in array snapshot support. The EqualLogic PS Series arrays are designed to provide scalable storage, allowing organizations to grow and match their retention requirements without performance degradation. EqualLogic arrays combined with the DL2200 provide a complete data protection solution and the result is an ideal consolidated storage solution that provides an excellent return on investment.

CommVault's SnapProtect feature uses the volume snapshot capabilities of the EqualLogic SAN. This allows the system to backup large amounts of application data from SAN snapshots without using production server resources. With the option of off-load SnapProtect backups to a proxy server, the production server sees little to no resource utilization from the SnapProtect operation. Deduplication and the use of multiple backup streams, supported by CommVault provide a smart and efficient method to backup and store data. These immensely reduce the storage capacity and optimize the backup and restore process.

A series of tests were conducted at Dell storage labs to characterize data backup operations on Microsoft SQL Server running on a VMware® virtualized environment. The objectives of these are to:

- Test efficiency of off-host backup of SQL Server data volumes through the SnapProtect feature, using EqualLogic PS series arrays and PowerVault DL2200, and
- Test efficiency of source and target deduplication using the EqualLogic PS series arrays and PowerVault DL2200.

## 1.1 Audience

This white paper is targeted for database administrators, storage administrators, ESXi/VMware administrators, and database managers who have EqualLogic and PowerVault DL2200 storage and want to use CommVault Simpana to protect Microsoft SQL Server environments. It is assumed that 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 ESXi Server environments.

## 1.2 Terminology

The following terms are used throughout this document.

**Group**: A PS Series group consists of one or more PS Series arrays connected to an IP network. A group may contain up to 16 arrays.

**Pool**: A container where each member (array) is assigned when added to a group and data volumes assigned to hosts span across a pool. A pool can have up to 8 members.

**CommServe™**: The master server used in CommVault Simpana. It communicates with all clients and media agents coordinating all operations such as backups, restores, copies, media management, etc. within a CommCell. This server uses a Microsoft SQL Server database and therefore must be implemented on a Microsoft Windows system.

**Media agent**: Manages the transmission of data between clients and backup media when using CommVault Simpana. Media agents have broad operating system support including Windows, Linux, and UNIX® options.

**CommCell™ console**: A graphical user interface that helps to run backup/restore operations in CommVault Simpana. In addition the CommCell console also provides a number of other features to help control and manage data.

**iDataAgent (iDA)**: Software modules used for backup and restore operations for specific operating systems or applications.

**Clients**: Hosts running iDataAgents to protect data.

**Subclient**: A client can have multiple subclients, each of which can be associated with different source data.

**Disk library**: A storage resource with an associated mount path that is used in backup to store backup data.

**Storage policy**: A logical object through which a subclient is protected. The storage policy defines how data is backed up and replicated as well as retention requirements.

**Simpana SnapProtect**: Allows creation of point-in-time snapshots of the data, used for various data protection operations. SnapProtect backup works in conjunction with software and hardware storage arrays to provide snapshot functionality for data protection operations.

**Note:** For additional information see the *Dell EqualLogic Configuration Guide*, available at http://www.equallogic.com/resourcecenter/assetview.aspx?id=9831.
In addition, CommVault provides *Online Documentation,* available at http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/default.htm.

# 2 Test system configuration

To conduct the system testing detailed in this paper, the SQL Server test system shown in Figure 1 and Figure 2 were created.

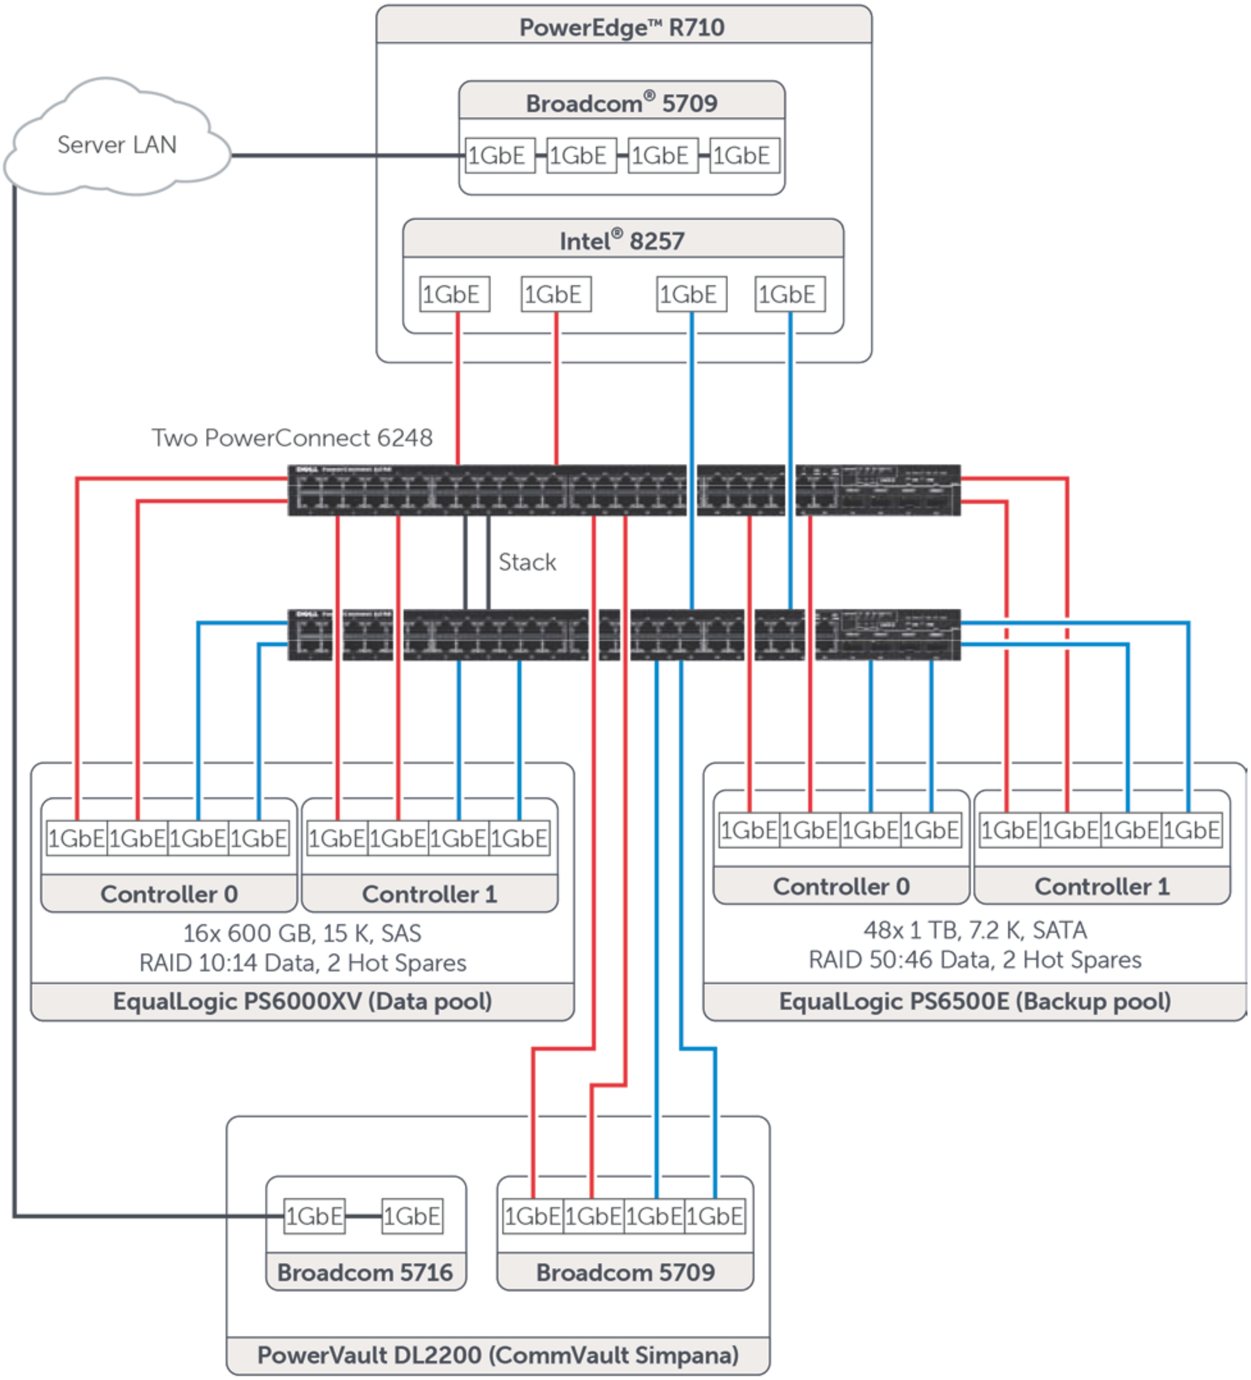

Figure 1 SQL Server LAN and iSCSI SAN Connectivity

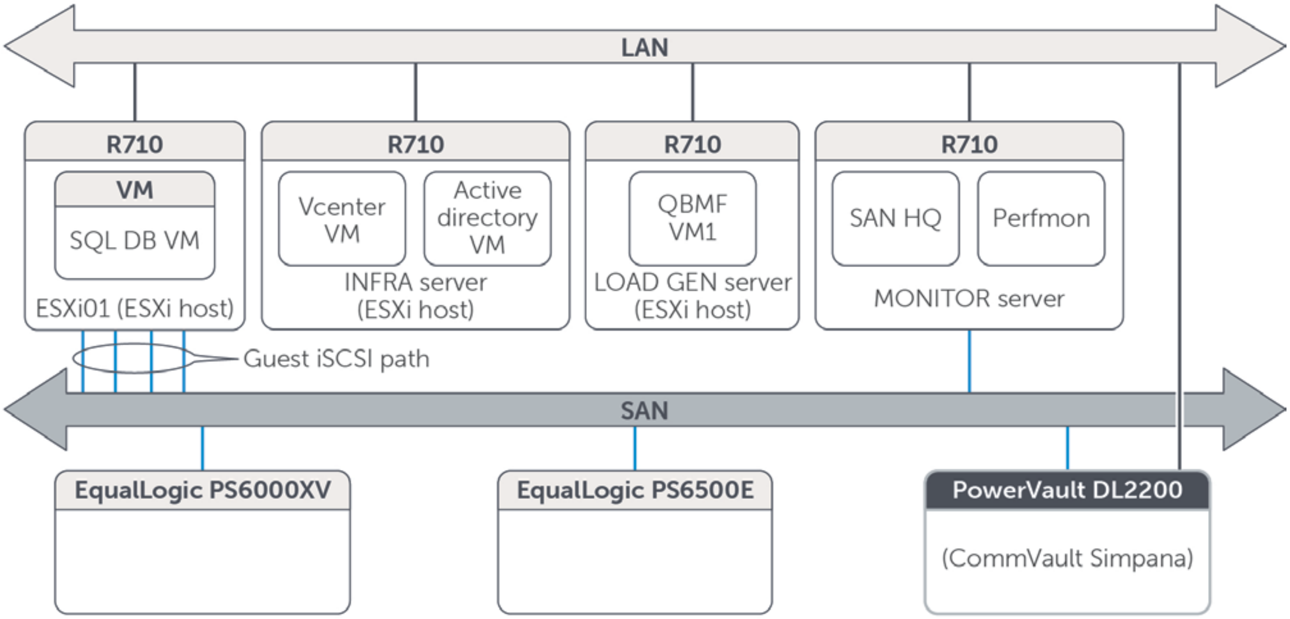

Figure 2 High Level view of Test System Components

Key design details in the test system configuration include:

- Installed and configured SQL Server 2008 R2 Enterprise Edition in a Windows Server 2008 R2 virtual machine (SQL DB VM) hosted on the VMware vSphere® ESXi4.1 on a Dell PowerEdge R710 server.
- The virtual machine that hosted the SQL Server was configured to use eight virtual CPUs and 16GB of reserved memory.

Network configuration details for ESXi 01 (Host):

- The on-board 4-port LOM (LAN on motherboard) Broadcom 5709 network controller was used for the Server LAN connection paths. Refer to Figure 3.
- An additional Intel® Gigabit VT Quad Port network adapter was installed in the server and used for the connection paths between the database server (SQL DB VM) and the volumes in the data pool on the PS6000XV array and backup pool on the PS6500E array. These four NIC ports were assigned on the physical server to be used as uplinks to vSwitch1 configured for iSCSI SAN access.
- The on-board 4-port Broadcom 5709 network controller on the DL2200 was used for the connection paths to the volumes in the data pool on the PS6000XV array and backup pool on the PS6500E array using the iSCSI initiator in DL2200.

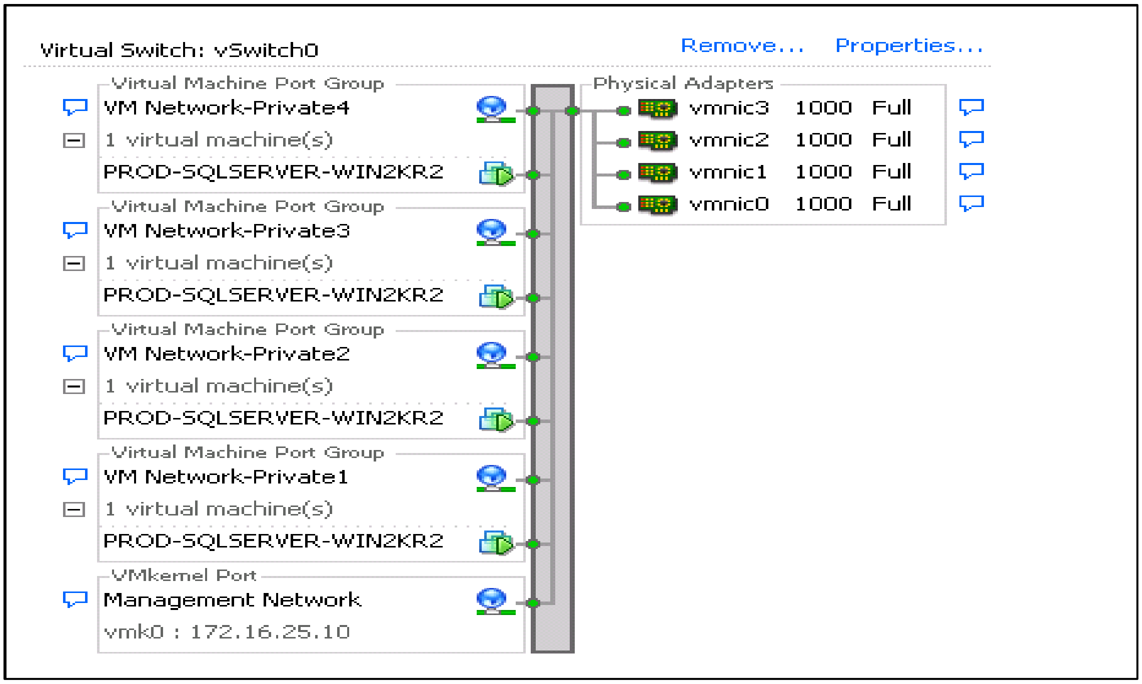

Virtual Switch: vSwitch0

Virtual Machine Port Group
VM Network-Private4
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-Private3
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-Private2
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-Private1
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

VMkernel Port
Management Network
vmk0 : 172.16.25.10

Physical Adapters

| | | |
|---|---|---|
| vmnic3 | 1000 | Full |
| vmnic2 | 1000 | Full |
| vmnic1 | 1000 | Full |
| vmnic0 | 1000 | Full |

**Figure 3 vSwitch0 Configuration**

- Created a separate vSwitch ("vSwitch0") for server side LAN network and a separate vSwitch ("vSwitch1") for iSCSI SAN access.
- Created virtual network adapters (type VMXNET 3) within the VM and assigned them to the vSwitch1 (Refer to Figure 4) on the vSphere host.
- Used EQL MPIO DSM via Host Integration Tools (HIT) kit to setup multiple paths from the guest VM to the storage volumes. These paths are labeled as "Guest iSCSI Path" in Figure 2.

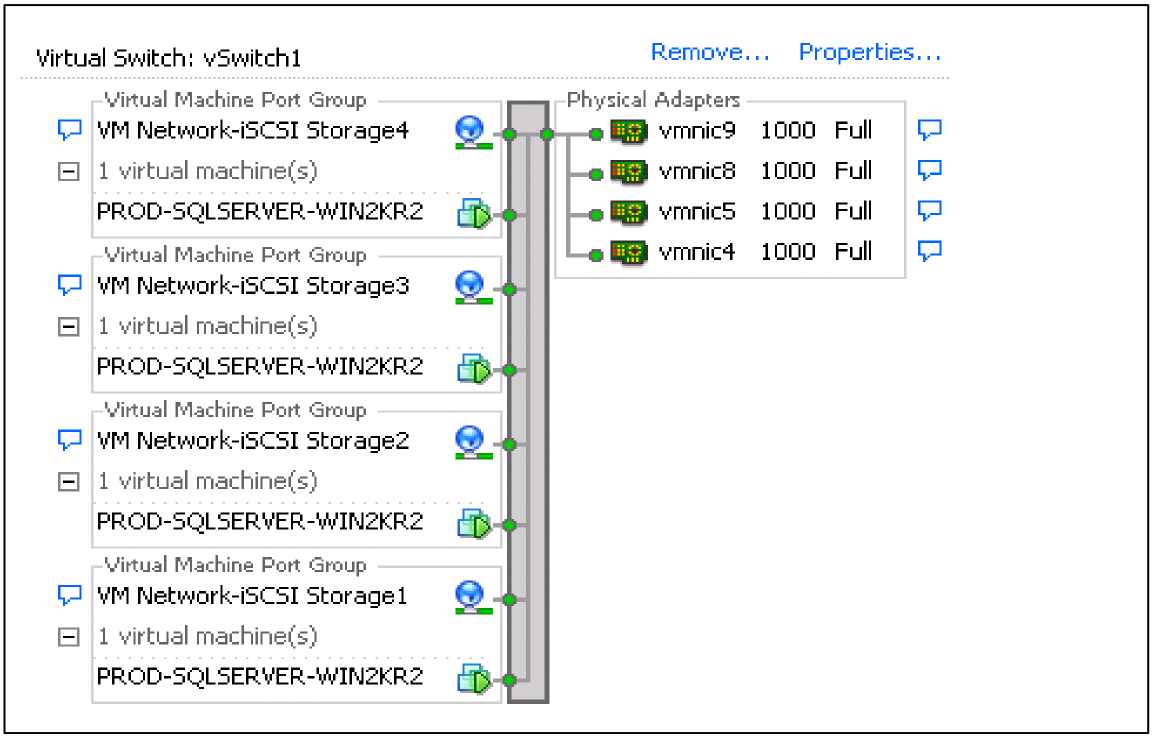

Virtual Switch: vSwitch1

Virtual Machine Port Group
VM Network-iSCSI Storage4
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-iSCSI Storage3
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-iSCSI Storage2
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Virtual Machine Port Group
VM Network-iSCSI Storage1
1 virtual machine(s)
PROD-SQLSERVER-WIN2KR2

Physical Adapters

| | | |
|---|---|---|
| vmnic9 | 1000 | Full |
| vmnic8 | 1000 | Full |
| vmnic5 | 1000 | Full |
| vmnic4 | 1000 | Full |

**Figure 4 vSwitch1 Configuration**

- The two local disks installed in the R710 server were configured as a RAID 1. ESXi 4.1 was installed on these disks, and the guest virtual machine OS disk partitions were also hosted within the VMFS file system on these disks.
- A second VMware ESXi 4.1 server (INFRA) was used to host virtual machines for vCenter and Active Directory.
- A third VMware ESXi 4.1 server (LOAD GEN) was used to host a Windows 2008 R2 workload simulation virtual machine, running an instance of Benchmark Factory® by Quest Software® (QBMF).
- The MONITOR server was a PowerEdge R710 server running Windows 2008 R2 natively. It was used to host EqualLogic SAN Headquarters (SANHQ) and Perfmon for monitoring.
- The SAN switches consisted of two Dell™ PowerConnect™ 6248 switches, configured as a single stack. Created redundant connection paths from each array controller to each switch in the stack.
- The DL2200 with CommVault Simpana was connected to the SAN switches and also to the server LAN. Refer to Figure 1 for connectivity details. This is the backup server that was used to offload backup tasks.
- One  EqualLogic PS6000XV consisting of 16 x 600GB 15K RPM SAS disk drives in a RAID 10 configuration was used  to host SQL Server volumes and snapshots (data pool).
- One  EqualLogic PS6500E consisting of 48 x 1TB SATA drives in a RAID 50 configuration was used to host the backup data volume (backup pool)  and other SQL Server components.
- Detailed configuration specifications for each test system component are provided in Appendix A.

# 3 Test studies

The following test scenarios study the performance impact on the production server when using CommVault Simpana's off-host SnapProtect feature for database backup and restore. The impact of deduplication during the off-host SnapProtect backup was also analyzed. Figure 5 illustrates the data flow across the SAN during the off-host SnapProtect database backup operations. Using SnapProtect, the data flow happens directly from production EqualLogic PS Series SAN (Data Pool on PS6000XV) to the backup target (Backup Pool on PS6500E array) via the DL2200 appliance, thus bypassing the SQL server database production host.

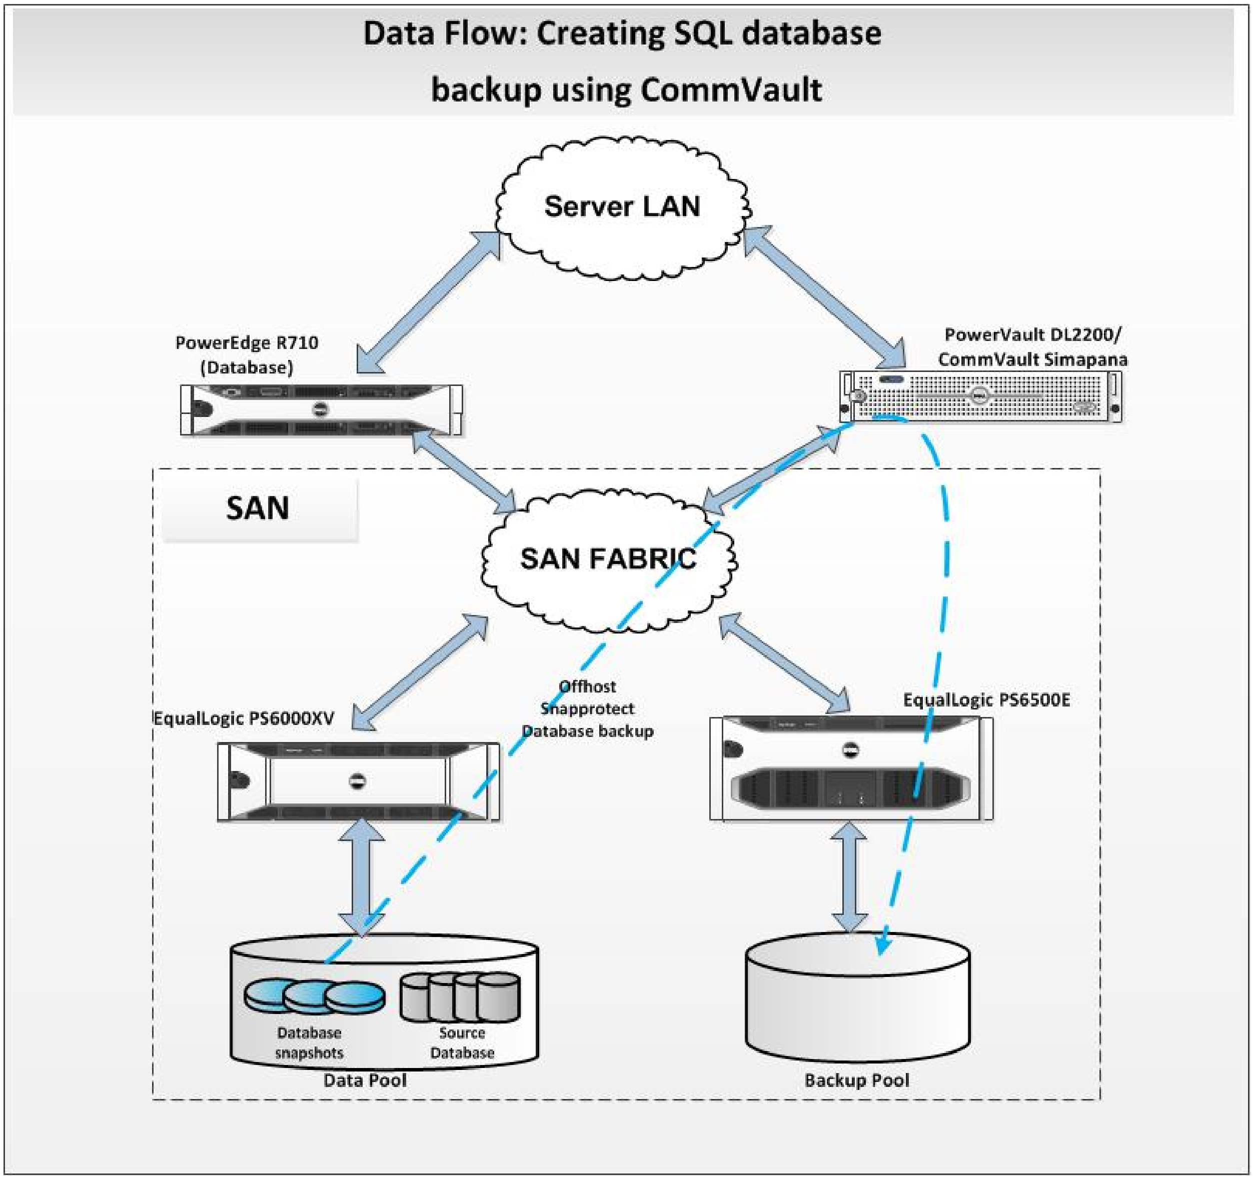
Data Flow: Creating SQL database backup using CommVault
Server LAN
PowerEdge R710 (Database)
PowerVault DL2200/ CommVault Simapana
SAN
SAN FABRIC
EqualLogic PS6000XV
Offhost Snapprotect Database backup
EqualLogic PS6500E
Database snapshots
Source Database
Data Pool
Backup Pool

Figure 5 Database Snap Backup flow using CommVault SnapProtect

## 3.1 SnapProtect Operations

In SnapProtect, a backup job is scheduled using the CommCell console at the DL2200. When the backup job is started:

**First**, CommVault's SQL iData Agent in the client SQL Server host machine communicates with the SQL Server application to put it in a consistent mode, flushing logs and ensuring consistent copies of data are created using Windows Volume Shadow Copy Service (VSS).

**Second**, the SQL iData Agent in the client SQL Server host machine will use the VSS framework to communicate with the PS Series SAN to create the snapshots. VSS will create the snapshots with the help of EqualLogic VSS provider (installed via HIT/ME) on the same machine.

**Third**, these snapshots are then mounted to the specified proxy server, which is the DL2200 here. The media agent in the server indexes and stores all info related to the snapshot operations in the database present in CommServe. The snapshots are mounted like a regular OS drive/file system and the required contents are read. The snapshots are then backed up to disk by the media agent present in CommServe.

- The CommVault software performs efficient indexing for all operations performed. Index data, gathers information regarding the location of files for recovery processing, and enables rapid browsing of backed up data.
- The Backup and Recovery software uses a two-part synchronized indexing scheme, consisting of a centralized Meta database catalog, residing within the CommServe Storage Manager, and an index co-located with the Media Agent software. This approach provides efficient scalability to accommodate data growth, support redeployment of storage resources, and increase reliability of the entire system. A permanent copy of the index is stored on the backup media.
- The Backup and Recovery software also maintains an active copy of the index on the client computer disk (i.e. the SQL Server host), where the media agent is installed. This provides an index cache on the same machine disk where media agent is installed, in addition to the index stored on the backup media.
- As new data is written to media, new indices are created. Configurable parameters let administrators set the size of the disk cache and duration of the local index. If the index exceeds the pre-configured capacity, older indices are overwritten using the least recently used cache algorithm.
- Requests for information found in the indices are satisfied from the cached copy of the index. If the index is no longer on disk, the permanent index on the backup media is accessed.

An array considers each snapshot as a full backup, but the movement of backup data to the media depends on the type of backup specified at CommVault console. For instance, if the backup type specified is a full backup, then CommVault data movement counts a snapshot as full. Similarly if the schedule calls for an incremental backup, the snapshot will be indexed as an incremental job and only index the changed items. Applications such as Exchange and SQL will only transfer the changed files from snap copies to backup media if incremental/differential schedules are executing the snap and are selected for protection.

## 3.2 Test #1: SQL Server performance impact studies using SnapProtect

For this test, four ~150GB databases were created using Quest Benchmark factory. The goal of this test was to study the performance impact on the production SQL Server during off-host SnapBackup and also study the reduction in backup window when multiple databases were backed up at the same time. To achieve this, four CommVault backup subclients were created and each was assigned a different database to perform a SnapBackup.

In CommVault, a backup copy operation provides the capability to copy snapshots of the data to any media. It is useful for creating additional standby copies of data on disk and these can be performed during the SnapProtect backup or at a later time. There are two types of backup copy movement operations to the media.

- *Inline Backup Copy:* Backup copy operations performed during the SnapProtect backup job , and
- *Offline Backup Copy:* Backup copy operations performed independent of the SnapProtect backup job at a later time.

For this test, a customer environment was simulated where the snapshots are taken throughout the day at regular intervals and the most recent snapshot gets backed up to disk at the end of the day. These recent snapshots need to be backed up before they age. For example, aging occurs when the snapshot reserve area runs out of space and older snapshots are deleted or the specified data aging policy expires. If they are not backed up before the snapshots expire, they are lost. In this test, to copy the snapshots of multiple databases, offline backup copy was performed to back-end disks hosted as EqualLogic volumes in the backup pool.

### 3.2.1 Storage volume layout

For the first test scenario, storage volumes were created as shown in Table 1.

Table 1 Storage volumes layout

| Volume | Size | Purpose |
|---|---|---|
| **Data Pool (PS6000XV RAID10)** | | |
| SQL-SERVER-DB-VOL1 | 150GB | Database#1 files |
| SQL-SERVER-DB-VOL2 | 150GB | Database#2 files |
| SQL-SERVER-DB-VOL3 | 150GB | Database#3 files |
| SQL-SERVER-DB-VOL4 | 150GB | Database#3 files |
| SQL-SERVER-DB-LOG1 | 50GB | Transaction Log#1 files |
| SQL-SERVER-DB-LOG2 | 50GB | Transaction Log#2 files |
| SQL-SERVER-DB-LOG3 | 50GB | Transaction Log#3 files |
| SQL-SERVER-DB-LOG4 | 50GB | Transaction Log#4 files |
| SQL-SERVER-DB-TEMPDB | 100GB | TempDB Files |
| **Backup Pool (PS6500E RAID 50**) | | |
| DB-BACKUP4 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP5 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP6 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP7 | 1TB | SQL Server DB Backup Files |

### 3.2.2 Backup duration

The backup duration depends on I/O throughput and decreases with increased throughput. So for this test, backup throughput was measured at the DL2200, by performing a full SnapProtect backup of one database by one client and linearly increasing the number of subclients to perform a SnapBackup of as many as four databases simultaneously. The snapshots for the databases were taken from the CommCell console and then an offline backup was run for the databases. The number of simultaneous backup copy operations can be specified under *specific storage policy > All Tasks > Run Backup copy* in the CommCell browser. For instance, if two subclients were run at the same time, then set the *No of simultaneous jobs to run* field to *2* (Refer to Figure 6).

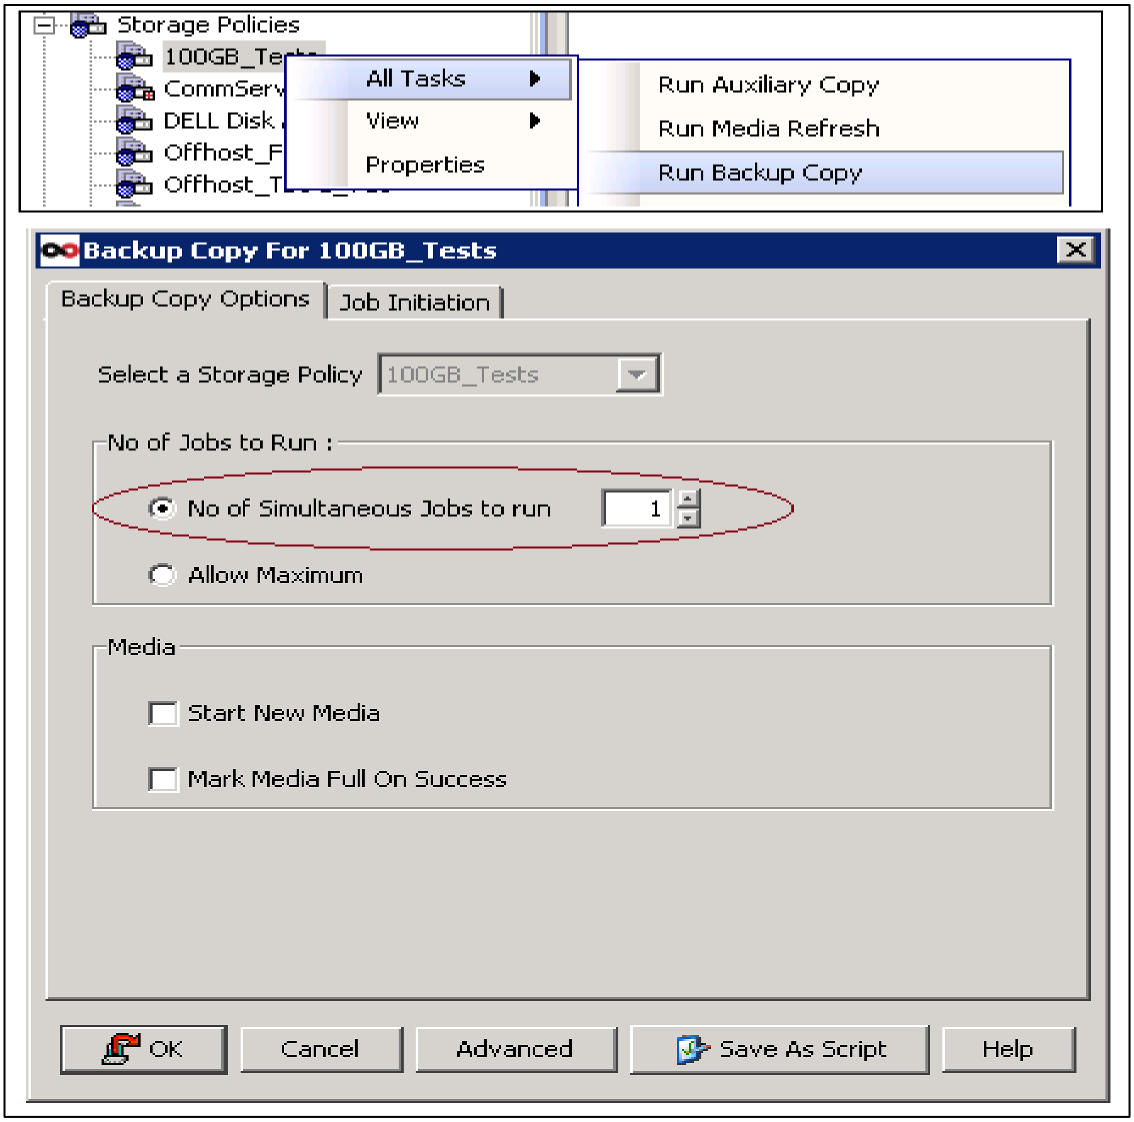

Select a Storage Policy: 100GB_Tests

● No of Simultaneous Jobs to run: 1

○ Allow Maximum

□ Start New Media

□ Mark Media Full On Success

Figure 6 Running an Offline backup copy

The read and write throughput was measured at the DL2200 by running Windows Perfmon utility. To get better performance, shared target disk libraries were created within CommVault and the Spill & Fill option across the mounted backup disk volumes was enabled for the library. This option allows the backup-write operations to happen in parallel across backup disk volumes and this improves performance. Figure 7 shows a high-level view of the multiple backup copies performed by multiple subclients simultaneously. The backup copy operation occurs after the snapshots are mounted in the DL2200.

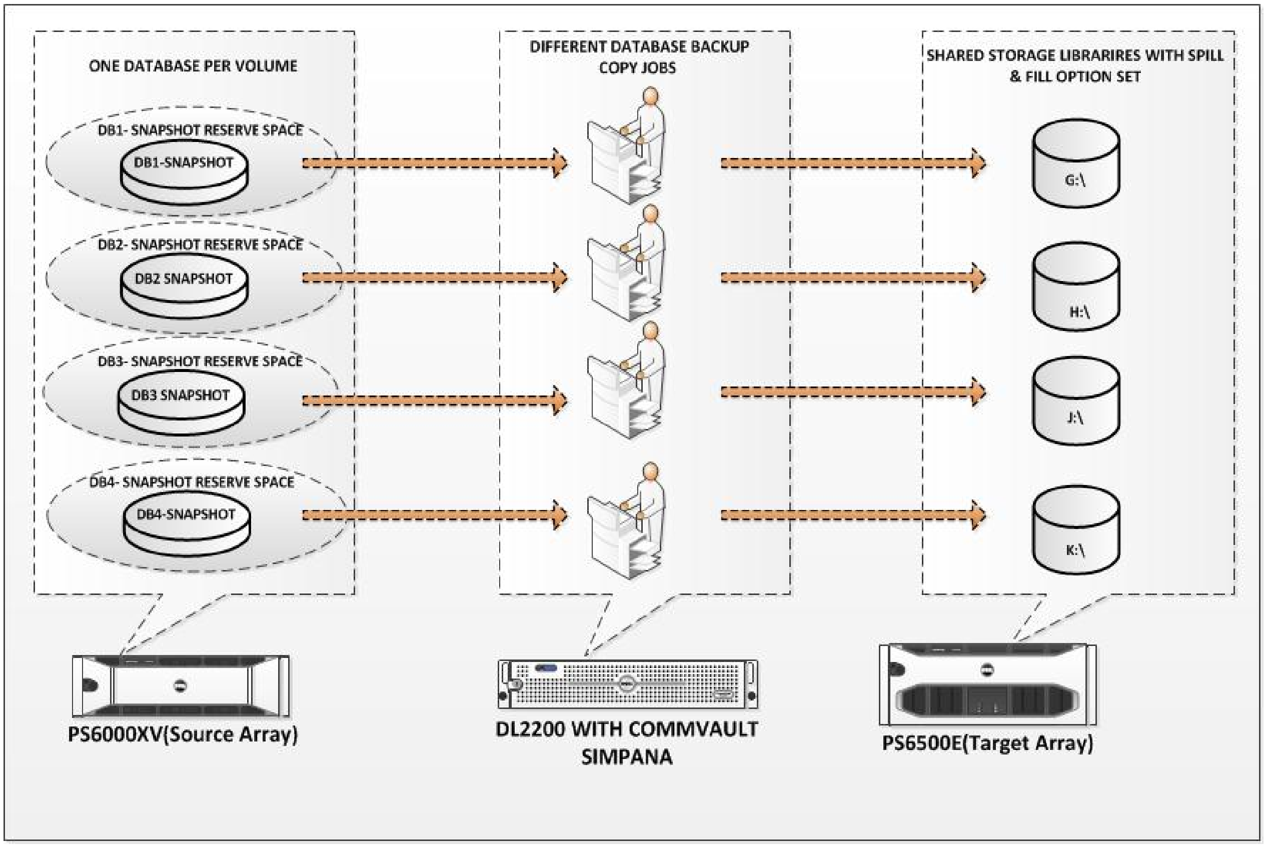
ONE DATABASE PER VOLUME
DB1- SNAPSHOT RESERVE SPACE
DB1-SNAPSHOT
DB2- SNAPSHOT RESERVE SPACE
DB2 SNAPSHOT
DB3- SNAPSHOT RESERVE SPACE
DB3 SNAPSHOT
DB4- SNAPSHOT RESERVE SPACE
DB4-SNAPSHOT
DIFFERENT DATABASE BACKUP COPY JOBS
SHARED STORAGE LIBRARIRES WITH SPILL & FILL OPTION SET
G:\
H:\
J:\
K:\
PS6000XV(Source Array)
DL2200 WITH COMMVAULT SIMPANA
PS6500E(Target Array)

Figure 7 High-level view of Backup copy of multiple DBs

Figure 8 shows the average and max write backup throughput and backup duration in each job with one to four database backups running simultaneously. As the number of database backup copies increased from one to four, the backup throughput increased due to multiple readers and writers involved during multiple database backups. This in turn reduced the backup duration for the multiple backup jobs. Backup throughput was better when multiple jobs were run at the same time rather than running each job sequentially.

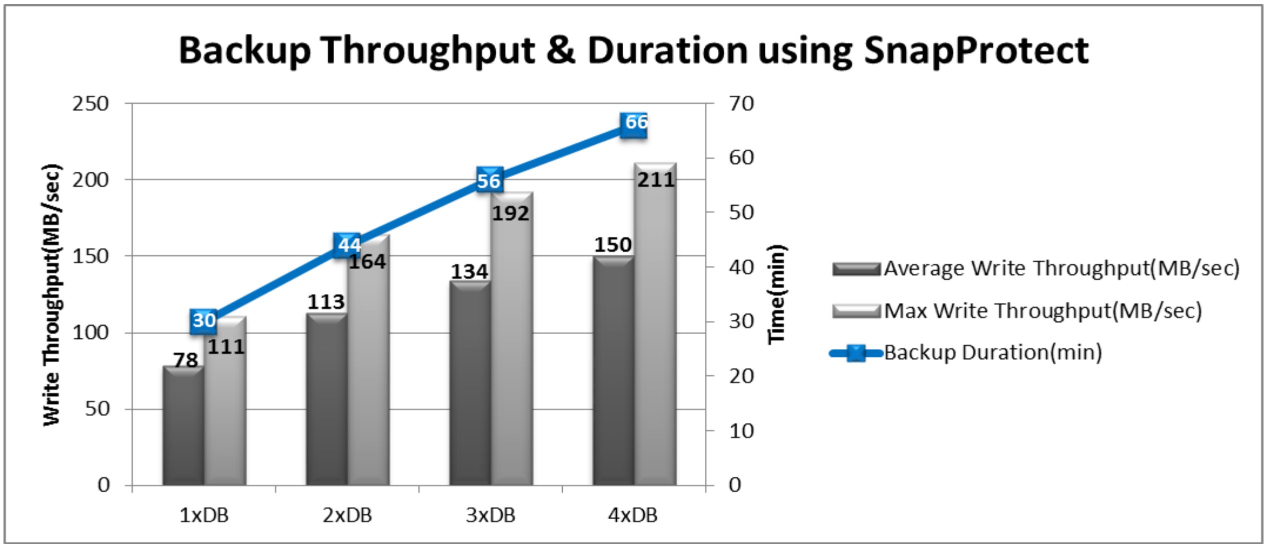

Figure 8 Backup Throughput and duration comparison

The backup duration was less for the simultaneous backups when compared to backups that were taken sequentially (Refer to Figure 9).The sequential duration was plotted theoretically by multiplying the duration of single database backup with the number of databases backed up.

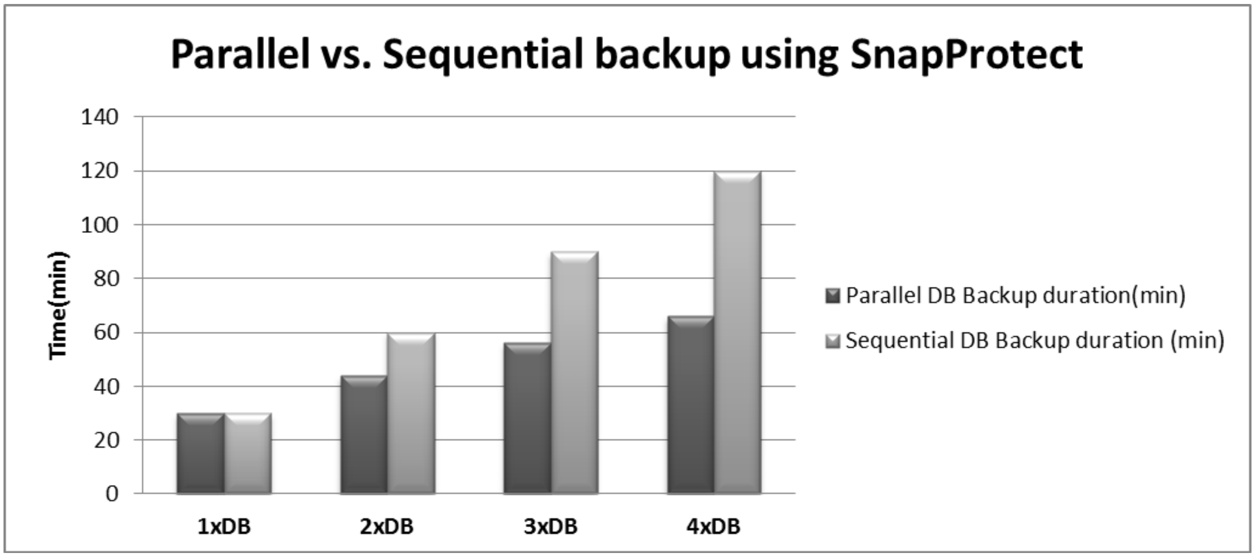

Figure 9 Backup Duration Comparison

### 3.2.3 CPU utilization at production SQL Server

The SnapBackup occurs by mounting the snapshots on the DL2200 appliance, thus taking the backup process off the production host. Minimal impact was seen at the production SQL Server in all of the four backup jobs. Figure 10 shows the impact of multiple database backup on CPU utilization at the production server. In all of the backups, the CPU utilization was minimal and did not impact the production server resources for operations.

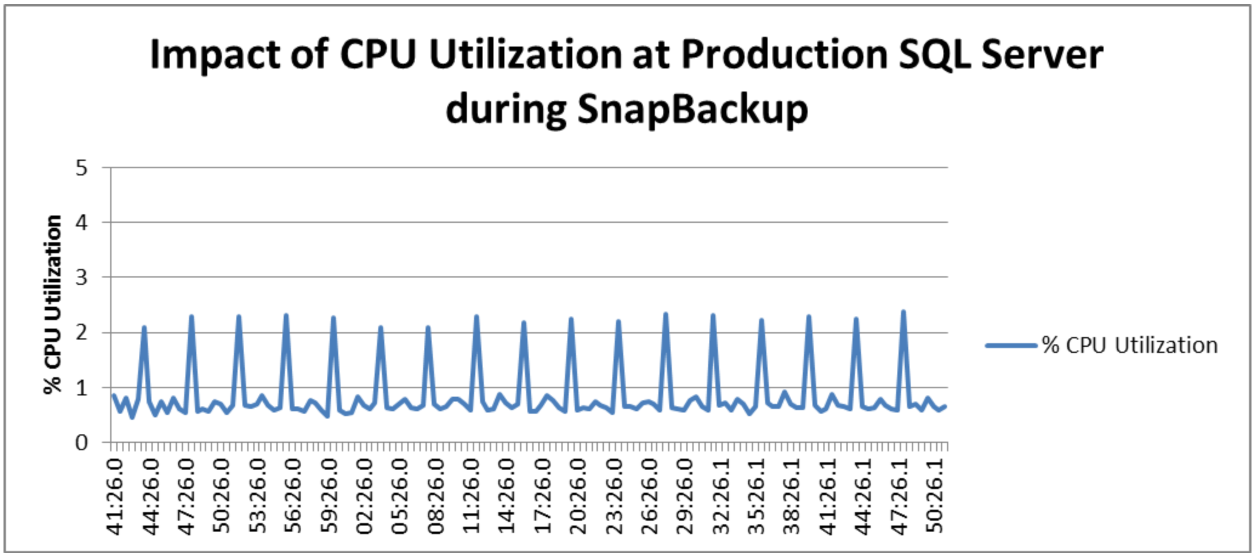

Figure 10 Impact on CPU Utilization at Production SQL Server during SnapBackups

## 3.3 Test #2: SQL Server performance impact studies during full & Tlog SnapProtect backup

This test simulated an environment where the typical production duration is an eight hour workday, full backups are taken daily, and transaction log backups are scheduled to run every 30 minutes during the eight hour work day.

- A single database with total size at the full backup point: ~1TB  was used
- To simulate this test, the user load from Benchmark Factory was run for 15 hours; a full backup was scheduled after 1.5 hours of user load run and Tlog (Transaction log) backups after 1.5 hours of full backup, for every 30 minutes for seven hours.
- Restore to a new database was performed from the full and transaction log backups.

SnapProtect full and log backups were performed as shown in Figure 11. Since SQL iDataAgent supports SnapProtect for only full and differential/incremental backups; transaction log backups follow the regular approach and get backed up through the LAN instead of off-host through the SAN. The performance counters measured for this test were CPU utilization, backup throughput, application response times and disk latencies.

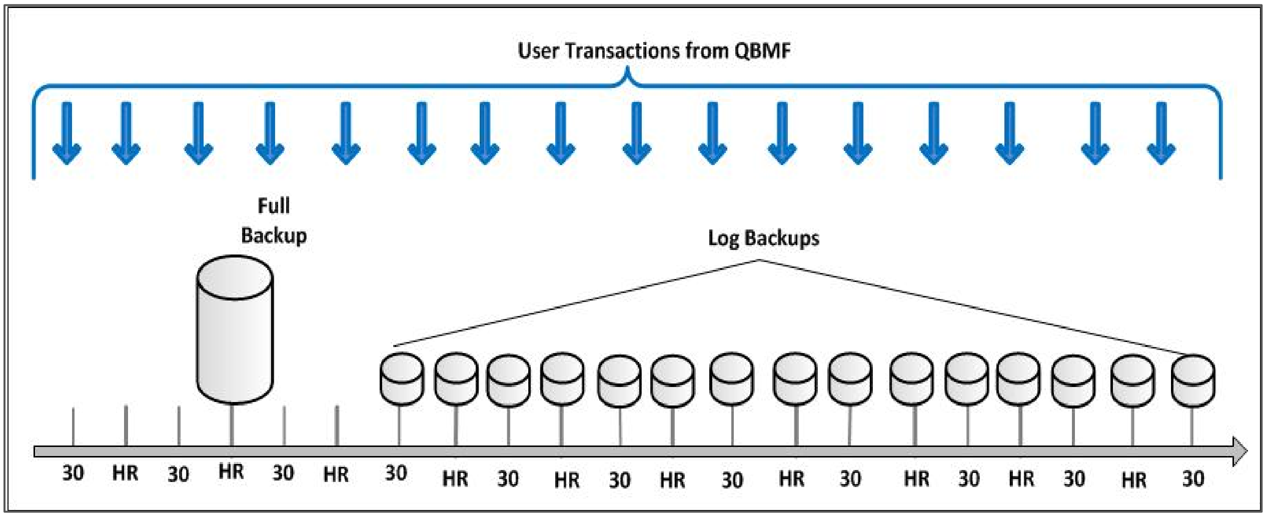

Figure 11 Backup Operations in CommVault Test#2

### 3.3.1 Storage volume layout

The storage volume layout specified in Table 2 was used during testing and source database files spread across four volumes. The backup volumes in the target array were exposed to the DL2200 through the software iSCSI initiator on the Windows Server 2008 R2 operating system running within the DL2200 appliance. For the CommVault media agent to access these volumes, the target volumes in backup pool were added to the disk libraries as mount paths.

Table 2 Storage volumes layout

| Volume | Size | Purpose |
|---|---|---|
| **Data Pool (PS6000XV RAID10)** | | |
| SQL-SERVER-DB-DATA-VOL1 | 500GB | Database files |
| SQL-SERVER-DB-DATA-VOL2 | 500GB | Database files |
| SQL-SERVER-DB-DATA-VOL3 | 500GB | Database files |
| SQL-SERVER-DB-DATA-VOL4 | 500GB | Database files |
| SQL-SERVER-DB-LOG | 600GB | Transaction Log files |
| SQL-SERVER-DB-TEMPDB | 100GB | TempDB Files |
| **Backup Pool (PS6500E RAID 50)** | | |

| Volume | Size | Purpose |
|---|---|---|
| DB-BACKUP4 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP5 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP6 | 1TB | SQL Server DB Backup Files |
| DB-BACKUP7 | 1TB | SQL Server DB Backup Files |

### 3.3.2 Full and Tlog backup

Figure 12 compares the CPU utilization at the Production SQL Server during backup using CommVault's SnapProtect feature and native SQL backup utility (SQL Server Management Studio here). When the backup was performed using the native SQL Server Management Studio on the same setup, the average CPU utilization increased from 1.7% to 13% during full backup, because the backups were in-host using the host server resources ( i.e.  Through the SQL server host) and hence the impact on CPU was high. During the off-host full backup, the impact on the CPU utilization at production SQL server was very low using SnapProtect. The reason for the minimal impact at the production server using SnapProtect was due to the fact that the backups were performed off-host through the SAN by the DL2200 proxy server without using the resources at the production SQL Server.

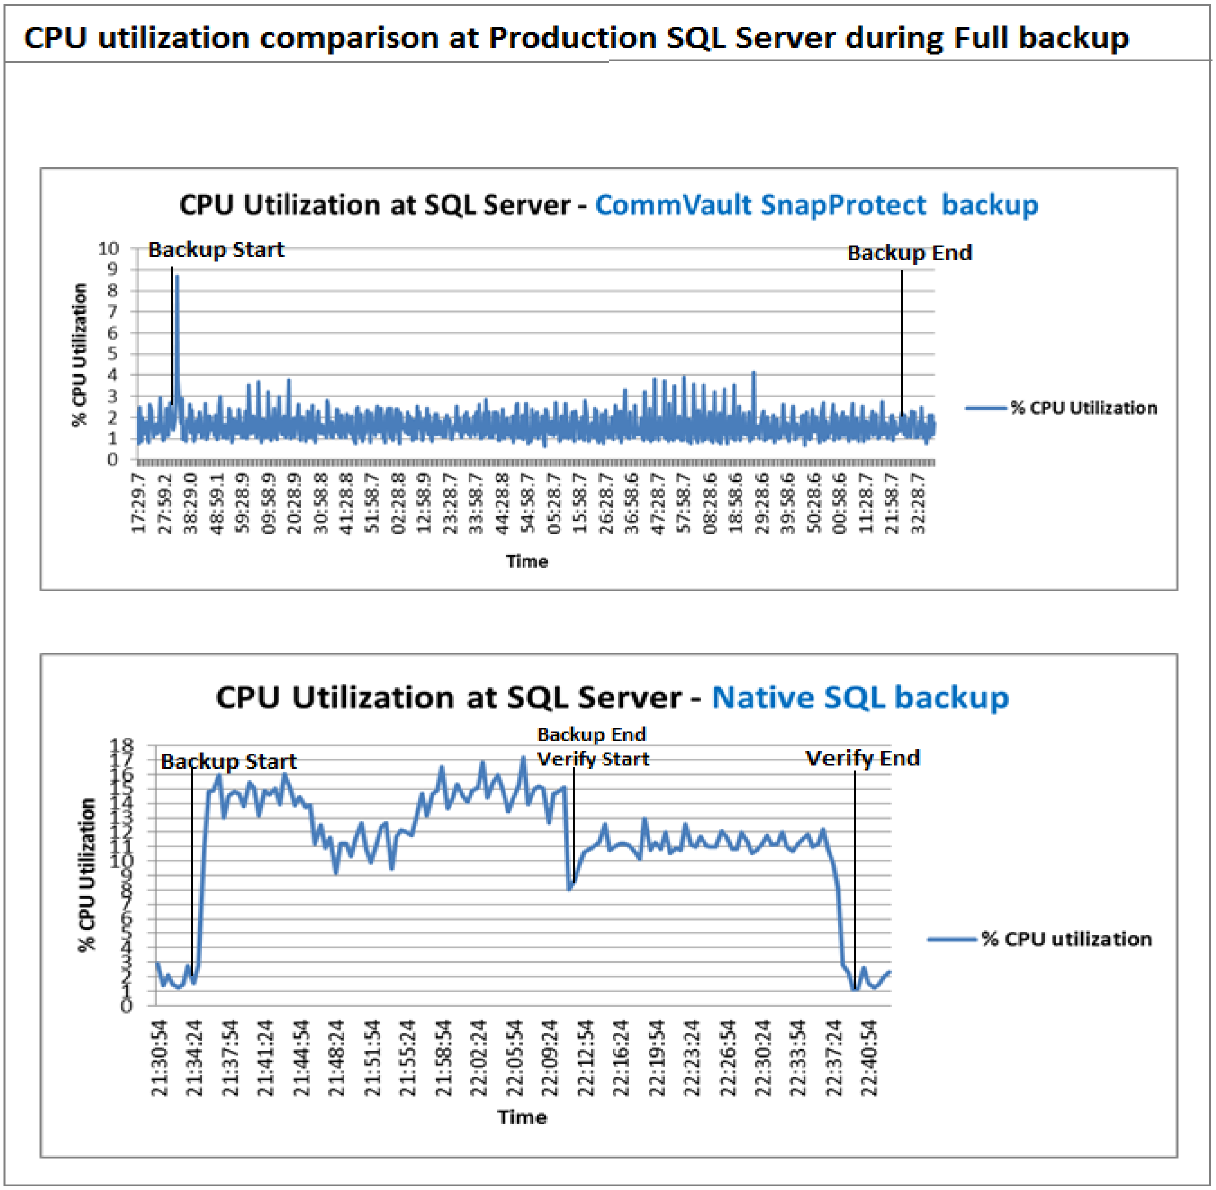

Figure 12 CPU Utilization Comparison during Full backup

Figure 13 shows the CPU utilization during Tlog backups using CommVault and native SQL Server Management studio. Since CommVault SnapProtect is supported only on full, differential backups, and incremental backups for SQL iData Agent, the CPU impact on the production SQL Server during Tlog backups was similar to the CPU utilization seen using native SQL server tools. The Tlog backups work like regular backups through the LAN in SnapProtect and consume CPU resources at lower utilization levels.

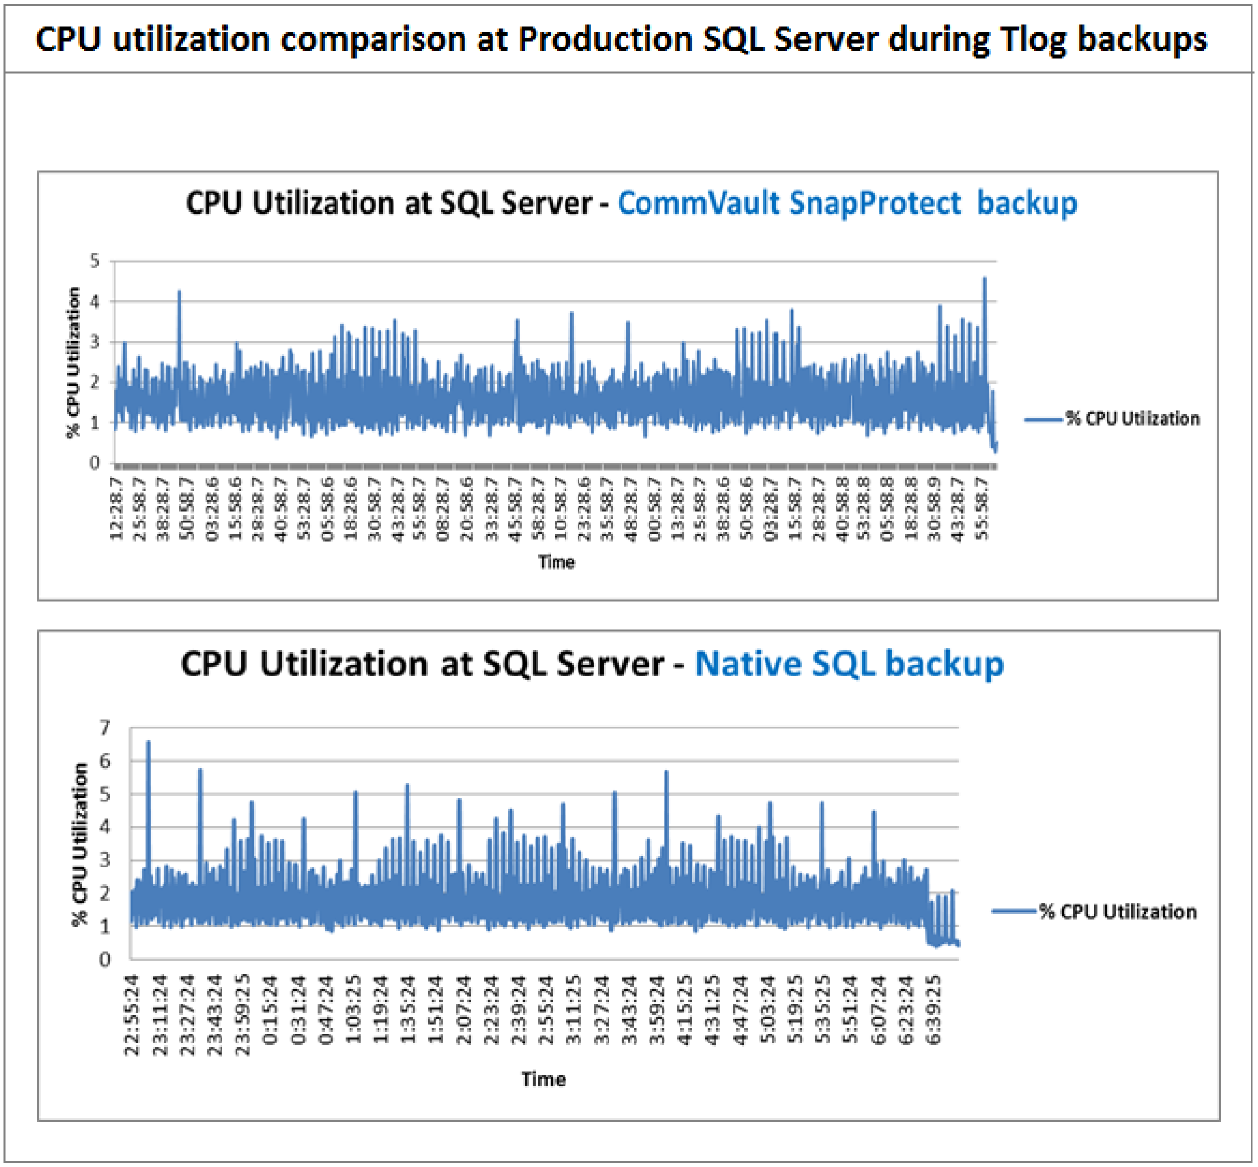

Figure 13 CPU utilization comparison during Tlog backups

The total time required to complete full database and transactional log backups was measured along with backup throughput during each operation. The baseline full backup in Figure 14 is the full backup taken when no user transactions were running on the production server. Figure 14 compares the increase in backup time for the backups taken during the user load vs. the backup taken with no user load. The backup throughput decreased by 7% and the backup duration increased by 18% during the full backup performed while the user transactions were running due to both backup and user transactions competing for storage resources.

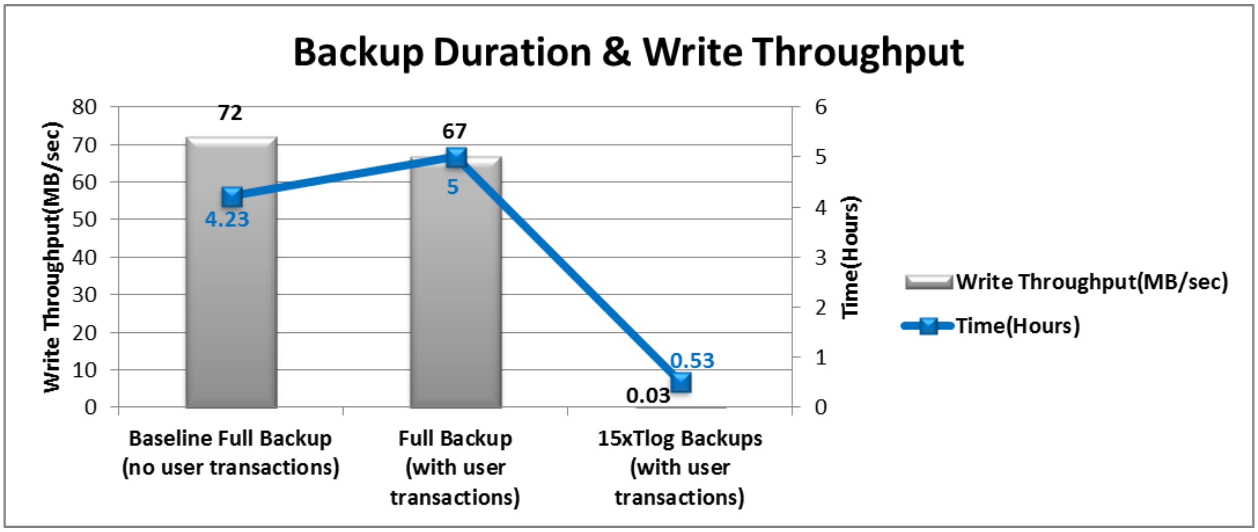

Figure 14 Backup throughput and duration comparison

Figure 15 compares the application response times during the full backup that was performed using CommVault's SnapProtect and native SQL Server Management Studio. The application response times are measures of the SQL transaction query completion times for the queries submitted by the simulation tool (QBMF).Since the backups were done off-host using CommVault; there was no impact on the application response times on the user transactions running during backup. However, for the same user load, there was a 77% increase in the application response time on the user transactions when backup was performed using native SQL Server Management Studio. This is because with the native SQL backup, host server's resources were utilized for backup operations and impacted the resources available for transaction processing.

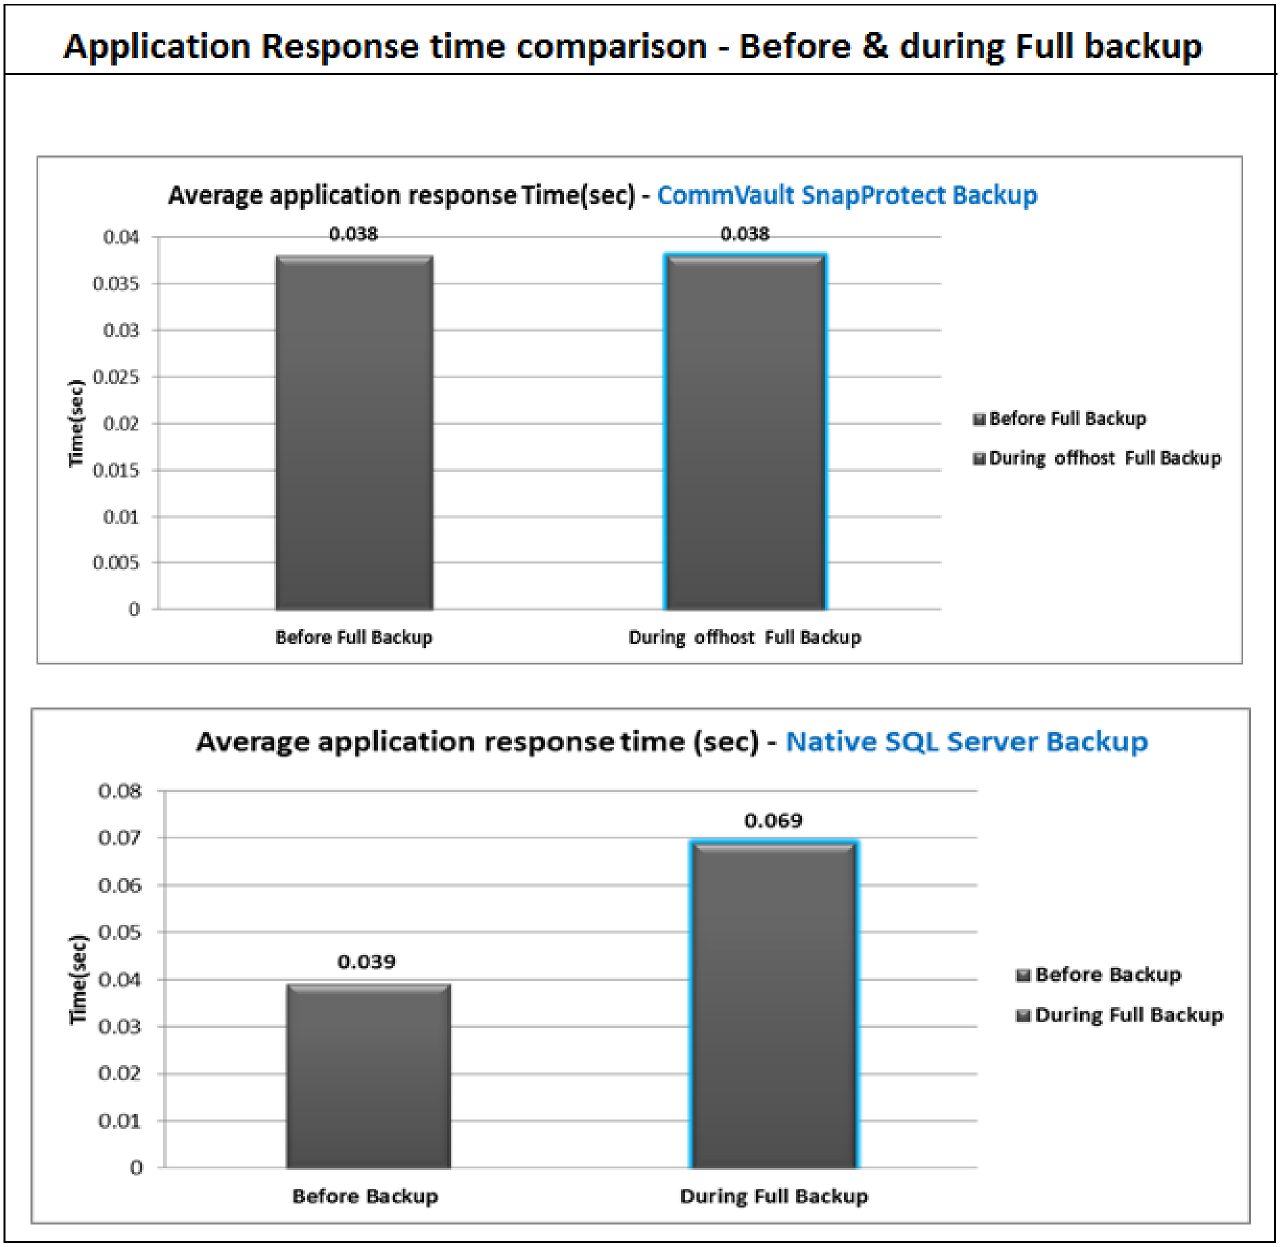

Figure 15 Application response time comparison during full backup

There was minimal impact on the disk read latencies at the source array (SQL Server production DB data) during off-host snap backup. The production server and the array were able to absorb this latency, resulting in minimal impact on the application response times as shown in Figure 14.

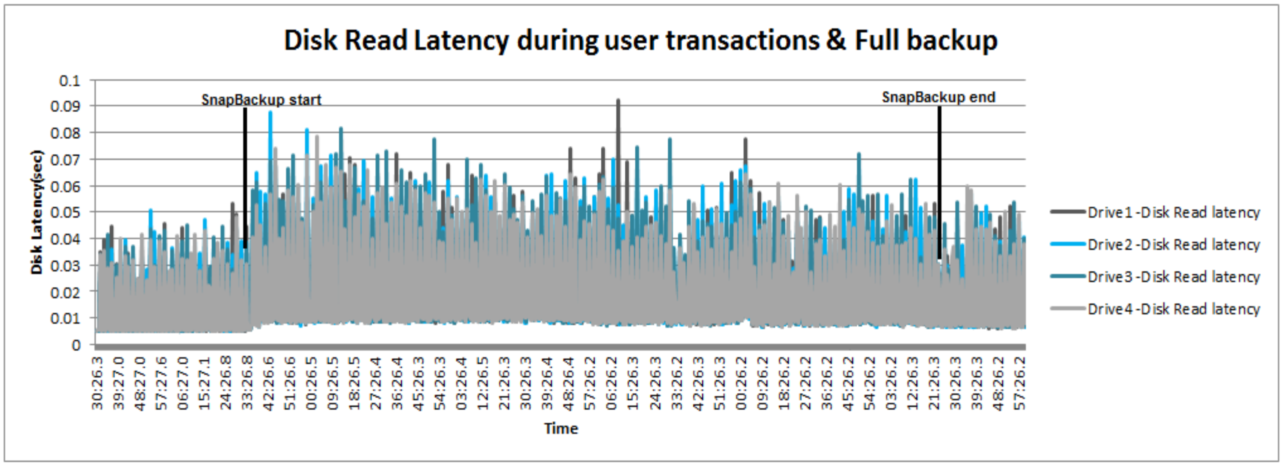

Figure 16 Impact on disk read latencies during off-host full backup

### 3.3.3 Full and Tlog restore

As verifying backup data integrity is very crucial, it is recommended to perform a restore operation immediately after the first full backup to understand the process. In CommVault, restores can be done in two ways.

**Restore from snapcopy:** These restores are performed by mounting the SnapProtect backup snapshots and restoring from the same. For this to happen, the storage array snapshots need to be present for the data volumes of the databases that were backed up.

**Restore from primary/diskcopy:** These restores are performed from the primary copy that was backed up to the disk from the mounted snapshots during backup. This restore method is used when the snapshots are not present for restoring. Snapshots may be deleted from the array due to factors such as low reserve disk space on the array; number of snapshots exceeds the threshold etc. The jobs corresponding to these deleted snapshots can no longer be used for any data recovery or backup copy operations. Restore from the diskcopy can be performed by setting the appropriate copy precedence number. By default in SnapProtect, snapcopy has copy precedence 1, diskcopy has precedence 2, and default restore is snap restore if no precedence is specified. Copy precedence can be specified using the following steps.

1. From the CommCell Browser, navigate to **Client Computers** > *<Client>* > *<Agent>*.
2. Right-click the entity that contains the snapshots you want to restore, and point to **All Tasks** > **Browse Backup Data**.
3. Click **OK**.
4. From the Browse window, select the data you want to restore in the right pane and click **Recover All Selected**.
5. From the **Restore Options** in the **All Selected Items** window, click **Advanced**.
6. Click the **Copy Precedence** tab and select the **Restore from Copy Precedence** checkbox.
7. In the **Copy Precedence** field, type the copy precedence number for the backup copy. It would be **2** for disk copy restores. (Refer to Figure 17)

The type of restore i.e. either from snapcopy or diskcopy would depend on Recovery Point Objective and Recovery Time Objective defined for a specific business continuity strategy.

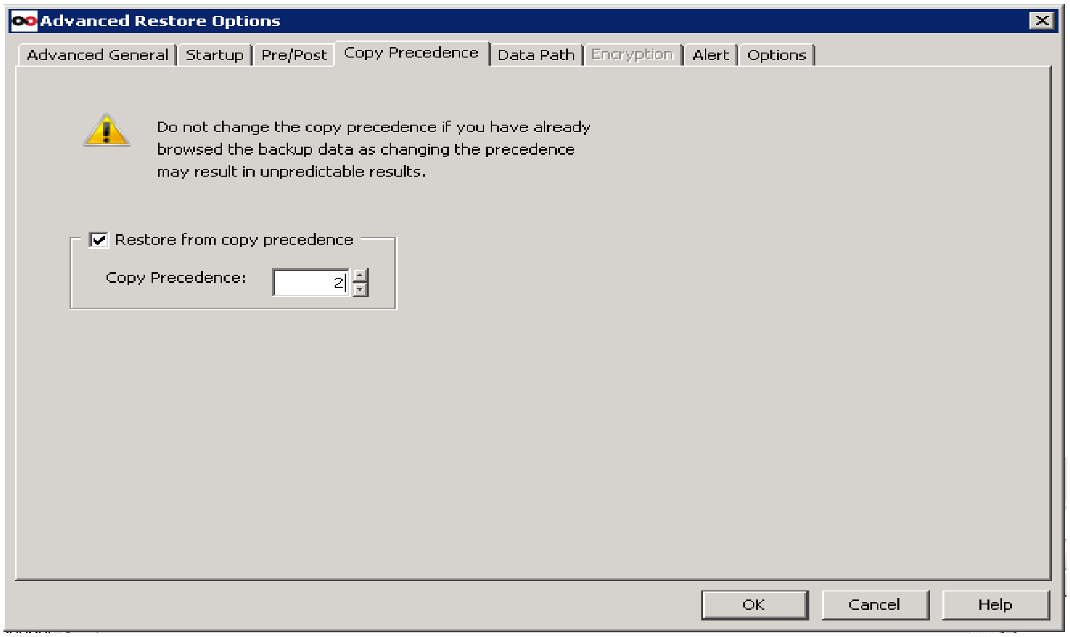

Do not change the copy precedence if you have already browsed the backup data as changing the precedence may result in unpredictable results.

☑ Restore from copy precedence

Copy Precedence: 2

Figure 17 Specifying copy precedence to restore from disk copy backup

To restore a database to a point in time, the last full backup needs to be restored first followed by the transaction log backups until the required point in time. For this test, the full backup was restored initially, followed by 15 Tlog backup restores. Here, full backup was restored from the diskcopy, simulating a customer scenario where the snapshots are not available and were backed up to disk before they expired.

Figure 18 shows the breakdown of restore throughput for the full backups taken when no user load was running on the production server and shown as "Baseline-Restore" and Full, Tlog backups that were taken when user load was running on the production server. The additional time taken for full backup restore was due to the restore of additional data created by user transactions running during backup. Restoring the transaction logs also added time for the restore process along with full backup restore. Both the full restore scenarios incurred similar throughput.

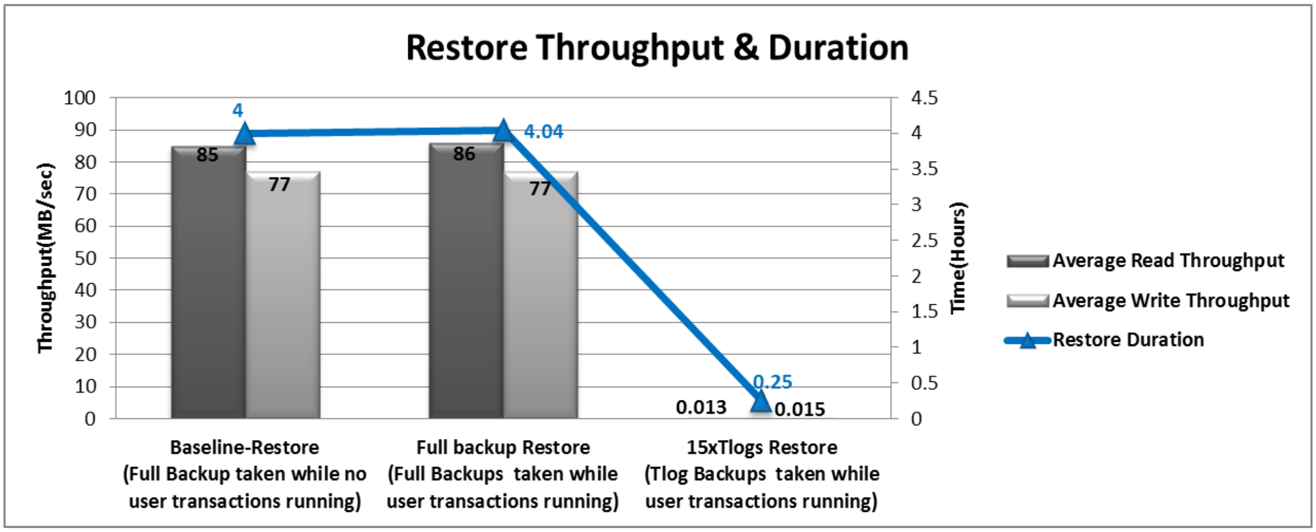

Figure 18 Restore throughput and duration comparison

## 3.4 Test #3: SQL Server performance impact studies during deduplication

### 3.4.1 Deduplication terminology

Deduplication focuses on data within a given object and provides an efficient method for storing data by identifying and eliminating duplicate items in backups. Block level deduplication is performed at the data block level by comparing blocks of data against each other.

Deduplication uses a hashing algorithm to compare data. A signature generation module computes the hashed signature for each block and then compares it with the existing signatures maintained in a deduplication store to determine whether it is identical. Based on the comparison, the CommVault media agent performs one of the following operations:

- If the signature is unique, the data is stored and an entry added to the deduplication store for subsequent comparisons.
- If the signature is identical to an existing signature, additional entries are created in the deduplication store with pointers to the existing storage.

The deduplication datastore (or the deduplication database) serves as the repository for signatures associated with all blocks that are backed up. It also has the reference counts to copies of the blocks that are backed up using the storage policy copy. Deduplication datastores are maintained for each storage policy copy that has deduplication enabled.

Data deduplication can be source or target based on where it occurs. Deduplication that occurs close to where data is created is often referred to as *Source Deduplication.* When it occurs where the data is stored, it is commonly called *Target Deduplication*.

- Deduplication on the source (client) side identifies and eliminates redundant data from the client. This reduces the data transfer rate over the network from the client to the media agent.
- The signature is generated and stored in the Deduplication datastore on the client and in the media agent. For subsequent backups, the signatures are generated in the client and compared with the existing signatures in the deduplication datastore in the media agent.
- Target deduplication is the process of removing duplicates of data in the secondary store. This reduces the amount of disk needed to store your data, but it does not change the amount of bandwidth needed to get the backups to the backup server. The signature is generated and stored in the deduplication datastore on the target which would be DL2200 here. The local hard drives on the DL2200 are configured to store the signatures.

For this test, several deduplication scenarios were performed to analyze the efficiency in each. The tests performed were,

- Source deduplication
- Impact of deduplication block sizes
- Target side deduplication

Below steps were followed in each of the above tests.

1. Used a Database of size: ~1TB
2. Enabled deduplication feature supported by CommVault and performed an initial full backup with no user load running from QBMF.
3. Ran the QBMF transaction mix for 30 minutes to generate data change in the data set backup and took subsequent full backups to understand the data set capacity reduction.

### 3.4.2 Source side deduplication

In this section, tests were performed for source side deduplication to study the impact on production SQL Server and overall dedupe and backup performance. For this test,

- A storage policy with source deduplication was created. Client side compression and signature generation were enabled at the subclient in the CommCell console. The software compression and the signature generation for this deduplication happen at the source (production SQL Server) by default because of the source side deduplication storage policy. However, when SnapProtect is enabled, there is an option to specify a proxy server (either the DL2200 or the production server which have Media Agent in them) to mount the snapshots for the backup. The software compression and signature generation happen at the selected proxy server and it is a way to choose where to run the deduplication workload. The tests in this section compare the impact of choosing,
  - Production SQL server as proxy (Refer to Figure 19) and
  - PowerVault DL2200 as proxy (Refer to Figure 21)
- Production SQL server was specified as proxy server for performing the SnapProtect backup operations for this test. It was specified in the subclient SnapProtect operations tab under properties. The media agent that performs the backup of the data to the target array is at the DL2200 (Refer to Table 3).
- DL2200 features the drives other than the operating system drives (The operating system drives reside in slots 12 and 13 within the DL2200 system) mounted as folders located in the tree structure C:\DiskStorage\<number>. As a best practice, the deduplication datastore was placed on C:\DiskStorage\1\1 within DL2200.

**Production SQL server as Proxy**

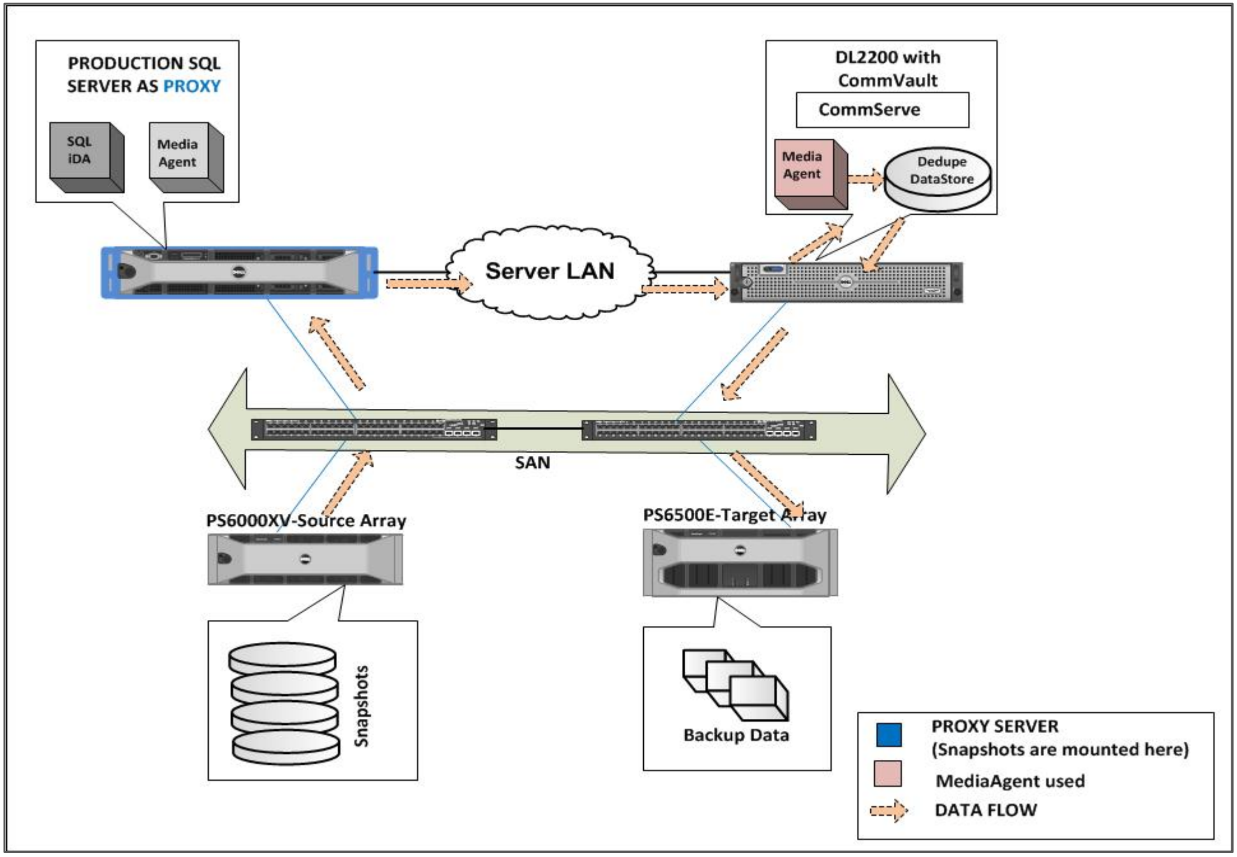

Figure 19 Data flow in source dedupe with production server as proxy

Figure 19 represents the high-level dataflow of source dedupe with production SQL server as proxy. Here the snapshots get mounted in the production SQL server, software compression and signature generation/hashing happen in the same server. The data is first split into 128KB blocks (this is configurable) and signatures are generated for these blocks. These initial signatures are stored in the dedupe datastore located at the DL2200. For the subsequent full backups, the data is split and similar signatures get generated and compared against the signatures in the dedupe datastore. Only the signatures of the changed blocks get added to the datastore and reference is made in the datastore for matching signatures.

Table 3 Configuration details

| | |
|---|---|
| STORAGE POLICY | Client side Deduplication |
| PROXY | Production Server |
| MEDIA AGENT USED | DL2200 |
| DEDUPE DATABASE LOCATION | DL2200 |
| DEDUPE BLOCK SIZE | 128KB |

The production server's resources were used for the same and high CPU utilization was observed (Refer to Figure 20). The dedupe/signature generation happens at the source before the data is sent to the target. Since the production server cannot access the target libraries, the data was sent through LAN to the DL2200, for the media agent at the DL2200 to do the data backup copy process. Since all the dedupe processing happen at the production server, there was very little resource impact seen in the DL2200.

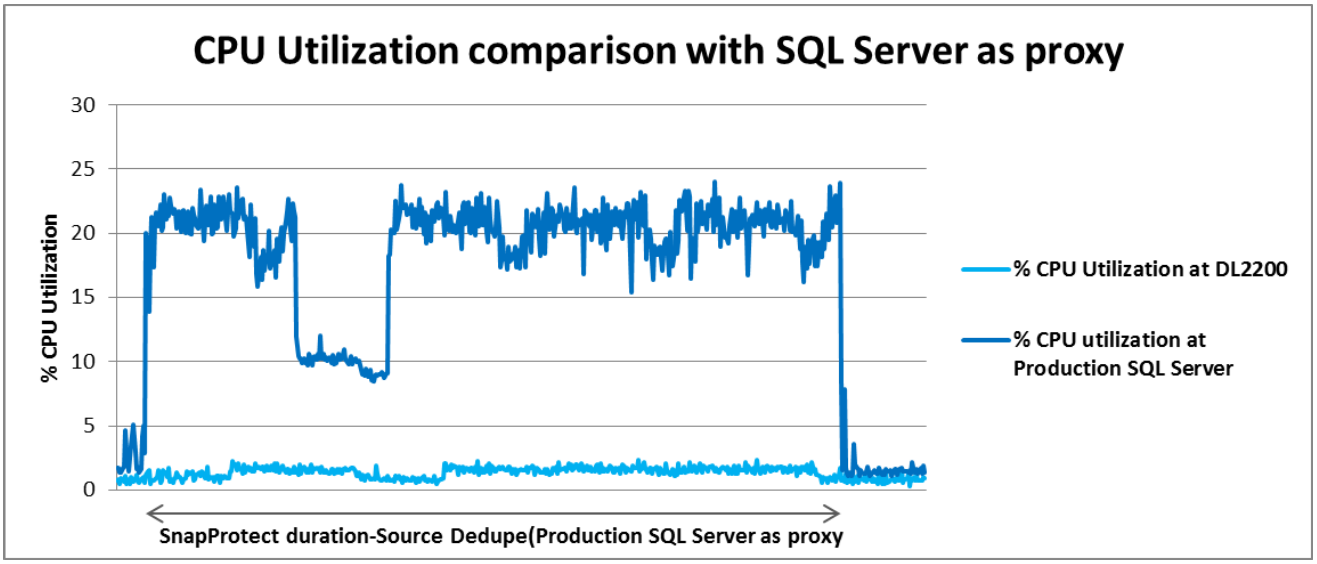

Figure 20 Impact on CPU utilization at production server vs. DL2200 with production server as proxy

### PowerVault DL2200 as Proxy Server

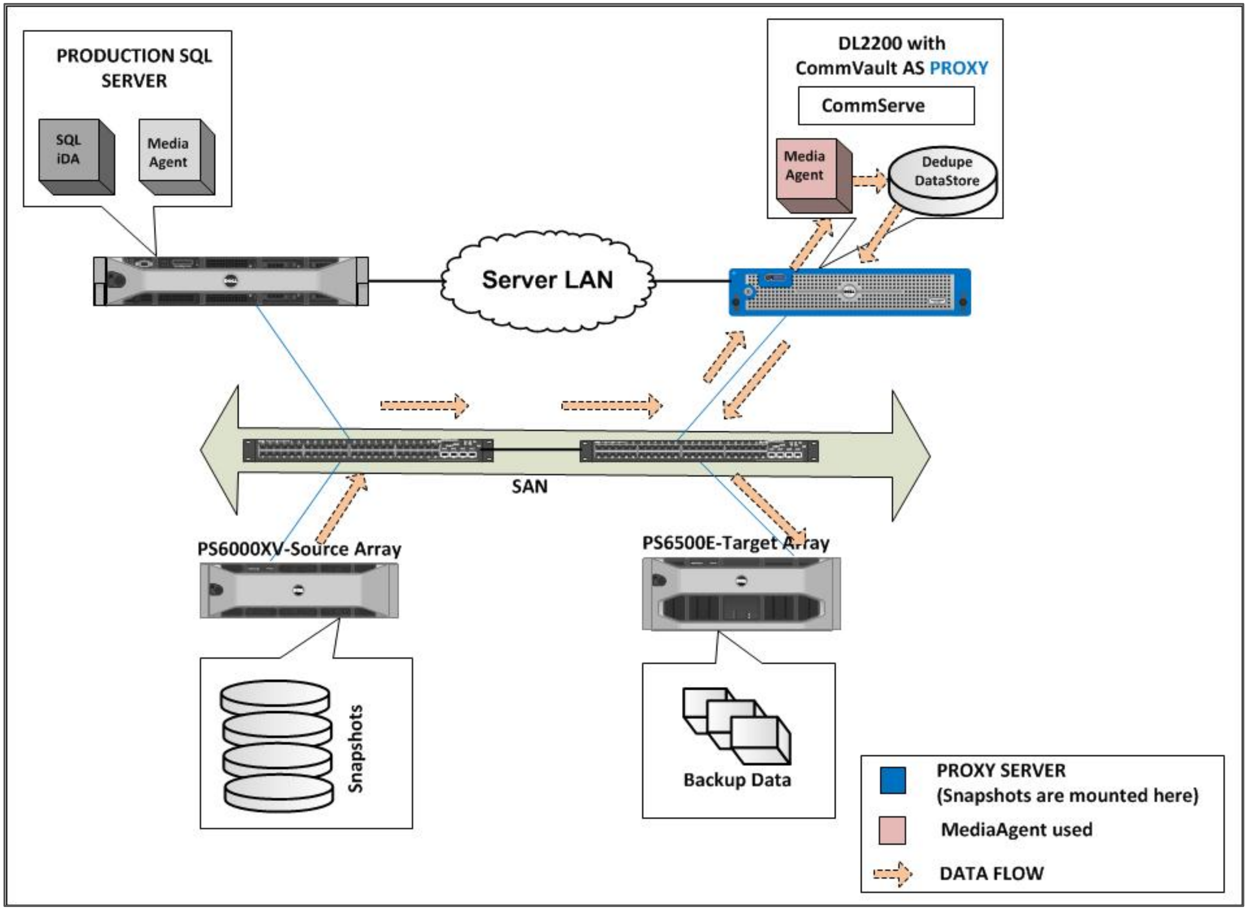

Figure 21 Data flow in source dedupe with the DL2200 as proxy

Figure 21 shows the data flow in this scenario. The DL2200 server was specified as the proxy server for performing SnapProtect backup operations for this test. It was specified in the subclient SnapProtect operations tab under properties (Refer to Table 4). The snapshots get mounted at the DL2200 server, software compression and signature generation/hashing occur in the DL2200.The Media Agent in the DL2200 is responsible for dedupe.

Table 4 Configuration details

| | |
|---|---|
| STORAGE POLICY | Client side Deduplication |
| PROXY | DL2200 |
| MEDIA AGENT USED | DL2200 |
| DEDUPE DATABASE LOCATION | DL2200 |
| DEDUPE BLOCK SIZE | 128KB |

Since the compression and hashing happen at the DL2200, there was very little resource impact seen in the production server (Refer to Figure 22) and DL2200 server resources were used and high CPU utilization was observed at the DL2200 (Refer to Figure 22). Also the DL2200 has visibility to both source and target volumes, and the datacopy occurs off-host through the SAN. Therefore, the production server resources have minimal impact.

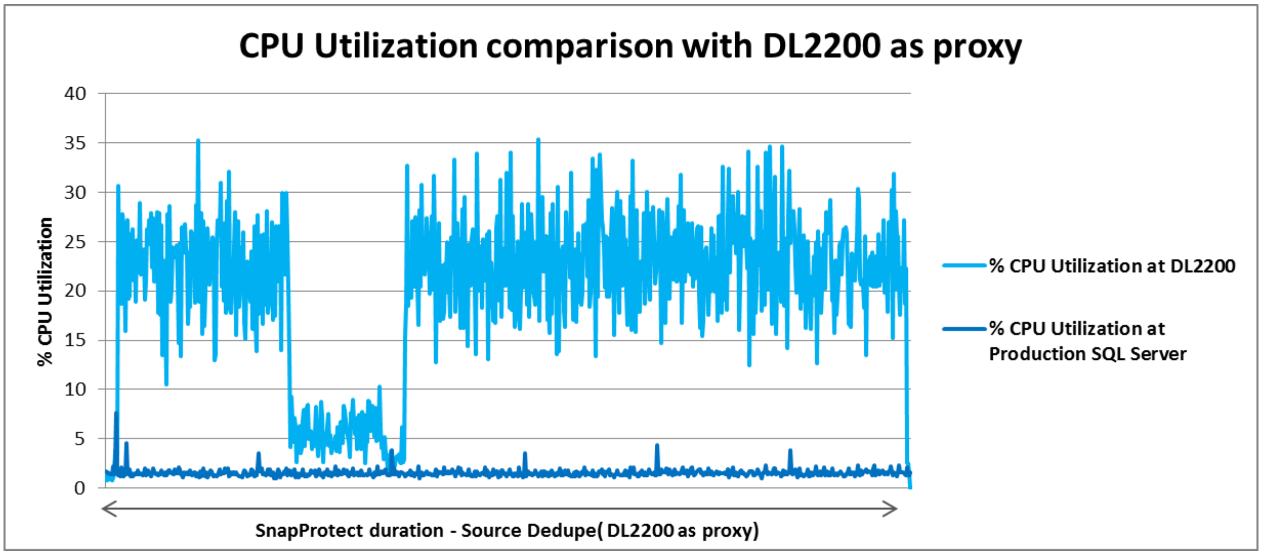

Figure 22 Impact on CPU utilization at production SQL Server vs. DL2200 with DL2200 as proxy

### Performance comparison

The source dedupe test steps below, were performed to measure backup time and backup capacity savings. The backup duration measured at the DL2200 includes the time it takes for snapshot creation, movement to proxy, complete indexing and then copy to the disk target location.

Table 5 Deduplication steps performed for this test

| Serial. No | Steps | Description |
|---|---|---|
| **Source Dedupe with Prod SQL Server as Proxy. Media Agent and Dedupe datastore at DL2200 Server.** | | |
| 1 | Dedupe#1- Full Backup | Performed a full backup with source side Deduplication (Production server as proxy) using SnapProtect backup. |
| 2 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate data change. |
| 3 | Dedupe#2- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 1 |
| 4 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate data change. |
| 5 | Dedupe#3- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 3 |
| 6 | Dedupe#4- Full Backup | Performed a full backup. This is to verify the capacity savings when there is no data change. Since no user transactions were run after step 5, there is no change in data; hence the capacity savings is expected to be ~ 100%. |
| **Source Dedupe with the DL2200 Server as Proxy. Media Agent and Dedupe datastore at the DL2200 Server.** | | |
| 1 | Dedupe#1- Full Backup | Performed a full backup with source side Deduplication (DL2200 as proxy) on an off-host SnapProtect backup. |

| | | |
|---|---|---|
| 2 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate differential. |
| 3 | Dedupe#2- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 1 |
| 4 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate data change. |
| 5 | Dedupe#3- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 3 |
| 6 | Dedupe#4- Full Backup | Performed a full backup. This is to verify the capacity savings when there is no data change. Since no user transactions were run after step 5, there is no change in data; hence the capacity savings is expected to be ~ 100%. |

Using the production server or the DL2200 as proxy during off-host source side deduplication has its pros and cons. Snap Protect with deduplication involves compression, signature generation, indexing and backup to the media. So when production SQL server was used as proxy, only the smaller data set that was already deduped on the client was sent across the LAN and hence the backup duration was less. However, this incurs additional CPU utilization in the production server due to the dedupe processing and compression happening in the same server.

On the other hand, when the DL2200 was used as proxy with source dedupe the hashing and deduplication occurs at the DL2200 before the data was sent to the backup target.  Also the DL2200 has visibility to both source and target volumes and the data copy occurs off-host through the SAN and hence there was minimal impact at the production server, but incurs additional CPU utilization at the DL2200. When DL2200 was used as proxy, all the deduplication processing, indexing, and backup to the media are taken care by the Media Agent at the DL2200. All the Deduplication and the additional data movement tasks pose higher resource utilization at the DL2200 and hence the backups were slightly longer. The backup durations in both cases are shown in Figure 23.

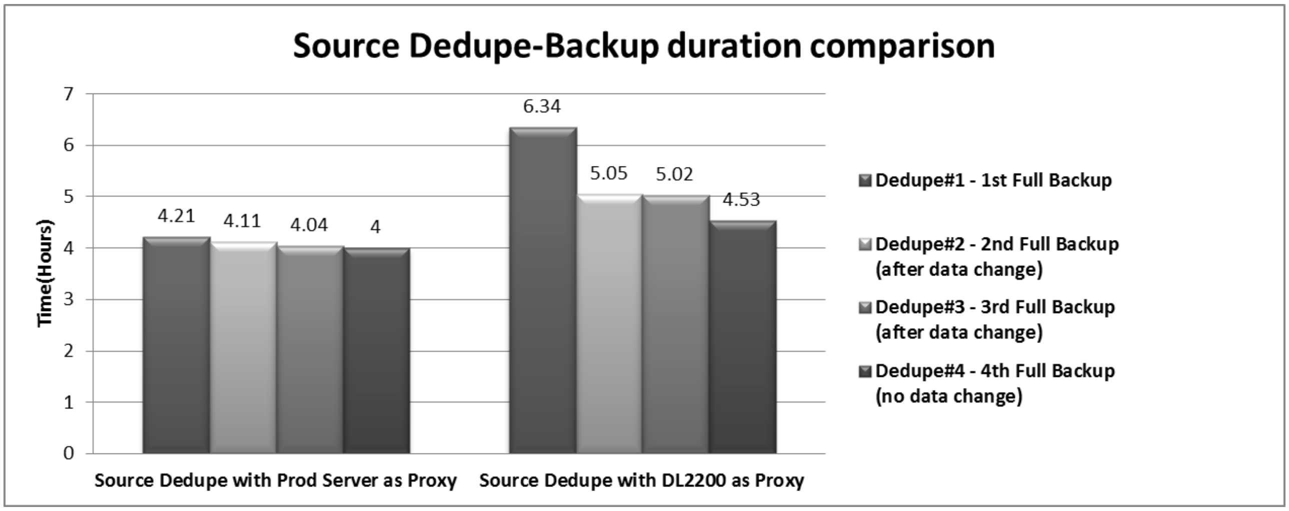

Figure 23 Backup time comparisons

Figure 24 shows the capacity reduction after each deduplication. As expected, when there was no data change, the capacity savings percentage was 99%.

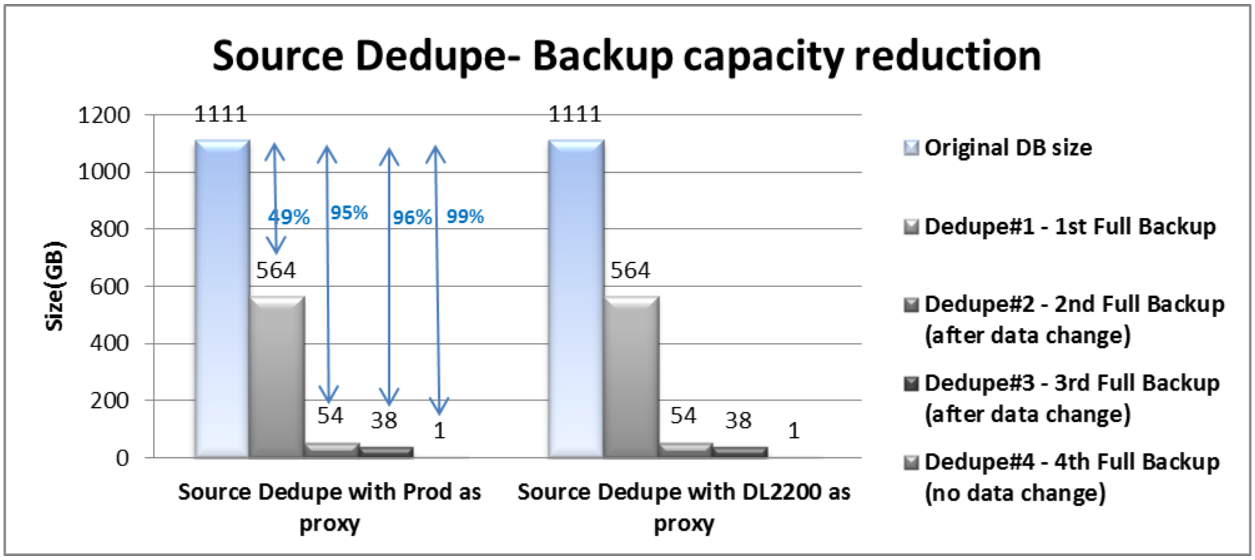

Figure 24 Backup capacity size reduction

### 3.4.3 Deduplication block size studies on source side deduplication

For the source side deduplication test, several full backups were run with 128KB and 256KB block sizes to see the performance impact on the production server and DL2200. The DL2200 was used as a proxy for the off-host SnapProtect source deduplication. To set the dedupe block size: select the desired block size in the *Block Level Deduplication Factor* field on the *Storage Policy Properties (Advanced)* tab and click *OK* to save the changes. For this test, the below steps were performed for 128KB and 256KB Block sizes.

Table 6 Deduplication steps performed for this test

| Serial. No | Steps | Description |
|---|---|---|
| 1 | Dedupe#1- Full Backup | Performed a full backup with source side Deduplication (DL2200 as proxy) on an off-host SnapProtect backup. |
| 2 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate data change. |
| 3 | Dedupe#2- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 1 |
| 4 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate differential. |
| 5 | Dedupe#3- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 3 |
| 6 | Dedupe#4- Full Backup | Performed a full backup of the same data in step 5. This is to verify the capacity savings when there is no data change. Since there is no change in data from step 3, the capacity savings is expected to be ~ 100%. |

The backup duration measured at the DL2200 includes the time it takes for snapshot creation, movement to proxy, complete indexing, and then copy to the disk target location. By default the dedupe block size for databases is set to 128KB. With 128KB block size, the data set reduction percentage was higher compared to the data set reduction percentage when using 256KB block size (Refer to Figure 25).The reduction percentage decreased when using larger block sizes because it makes the deduplication engine compare more data in a single chunk. However, the capacity savings were ~100% for the "Dedupe#4" (Full backup with no change in dataset from previous full backup) in both the block sizes. Hence, bigger block sizes means lesser deduplication, hence lesser capacity reduction.

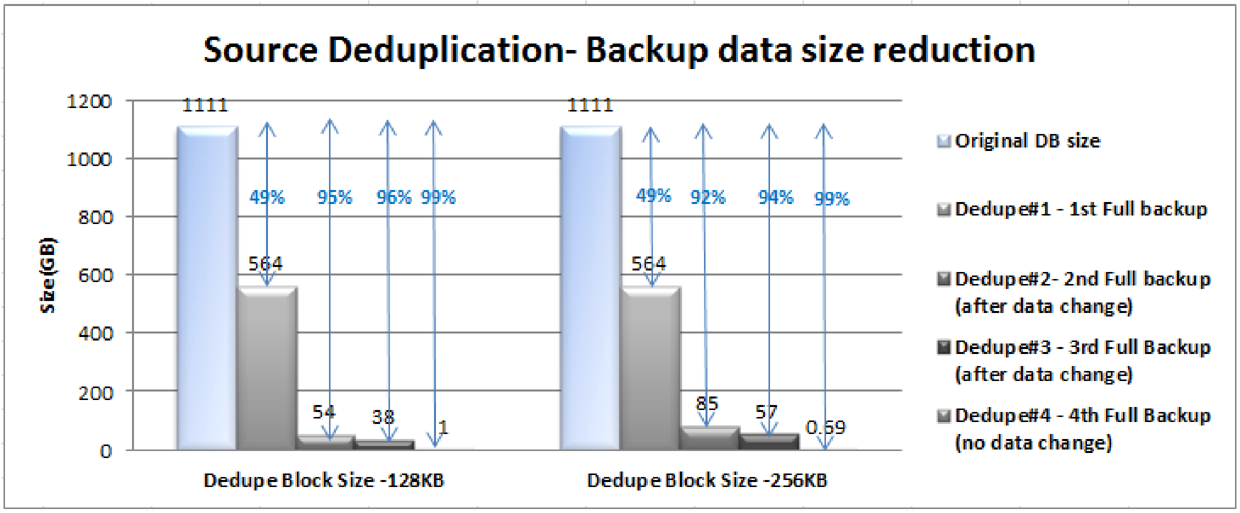

Figure 25 Impact of dedupe block sizes on backup data size

Figure 26 compares the backup duration when using 128KB and 256KB block sizes. With 256KB block size, the backup duration was faster because, bigger the block sizes, lesser number of blocks to compare and hence faster backups. Using larger block sizes compares more data in a single chunk and is recommended for structured data (databases) or large file types. This is because larger block sizes can result in lower deduplication rates when used in smaller unstructured data (files) environments. So using smaller block sizes here would allow for more granular comparisons to be made and deliver a greater disk savings benefit.

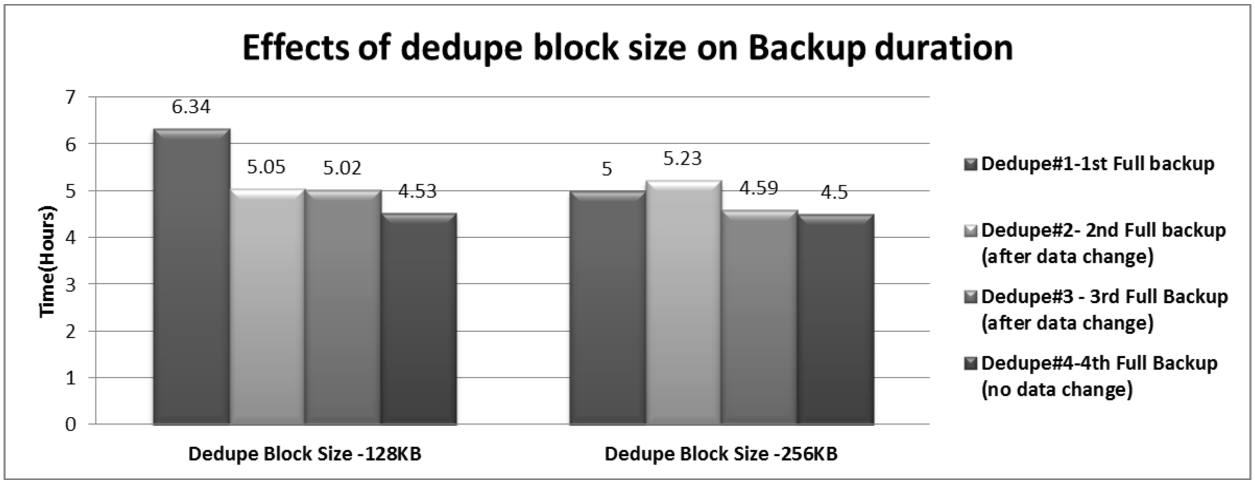

Figure 26 Impact of dedupe block sizes on backup duration

### 3.4.4 Target deduplication

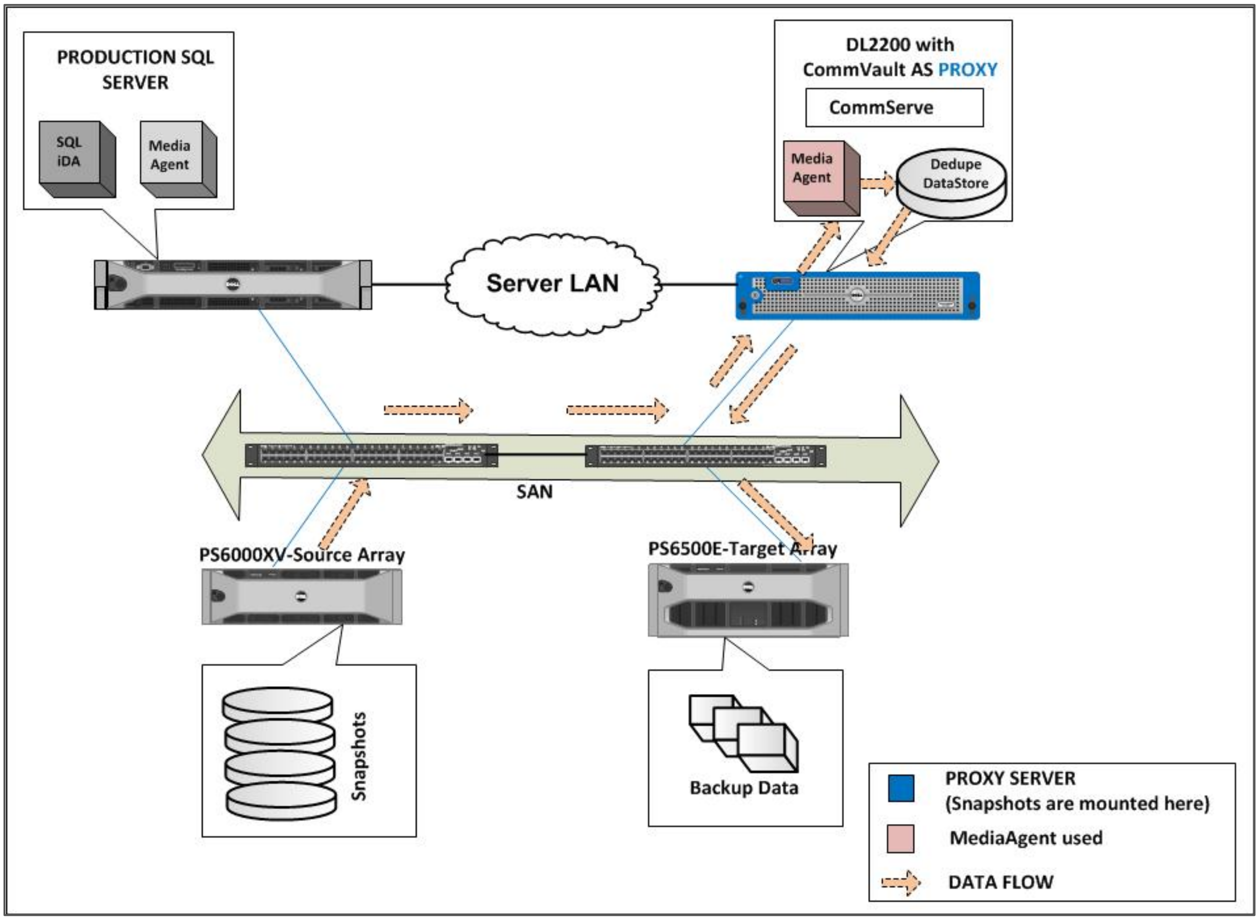

Figure 27 Data flow in target dedupe

A storage policy with target deduplication was created enabling compression and signature generation at the media agent in the DL2200. The difference between target dedupe and source dedupe with DL2200 as proxy is in the way they are created using storage policy and the way the agents perform dedupe and backup, however the data flow looks very similar in both the cases. Figure 27 shows the data flow in target deduplication .The dedupe datastore was created at the DL2200 as the media agent within DL2200 was used for dedupe. DL2200 access the data to be backed up through the SAN. Compression and signature generation are performed at the DL2200 by the media agent, signatures get stored in the dedupe datastore and the deduped backup data gets written to the target array. Table 7 shows the configuration details set for this test.

Table 7 Configuration Details

| | |
|---|---|
| STORAGE POLICY | Target side Deduplication |
| PROXY | DL2200 Server |
| MEDIA AGENT USED | DL2200 Server |
| DEDUPE DATABASE LOCATION | DL2200 Server |
| DEDUPE BLOCK SIZE | 128KB |

Since the compression and signature generation take place at the DL2200, there was minimal impact seen on the production server (Refer to Figure 28). The initial spike for the CPU utilization at the production SQL Server was caused by the communication establishment between the DL2200 and SQL Server, and also when snapshots get initiated at the SQL Server. The CPU utilization at the DL2200 was high due to the dedupe processing, indexing and backup to the media. (Refer to 28).

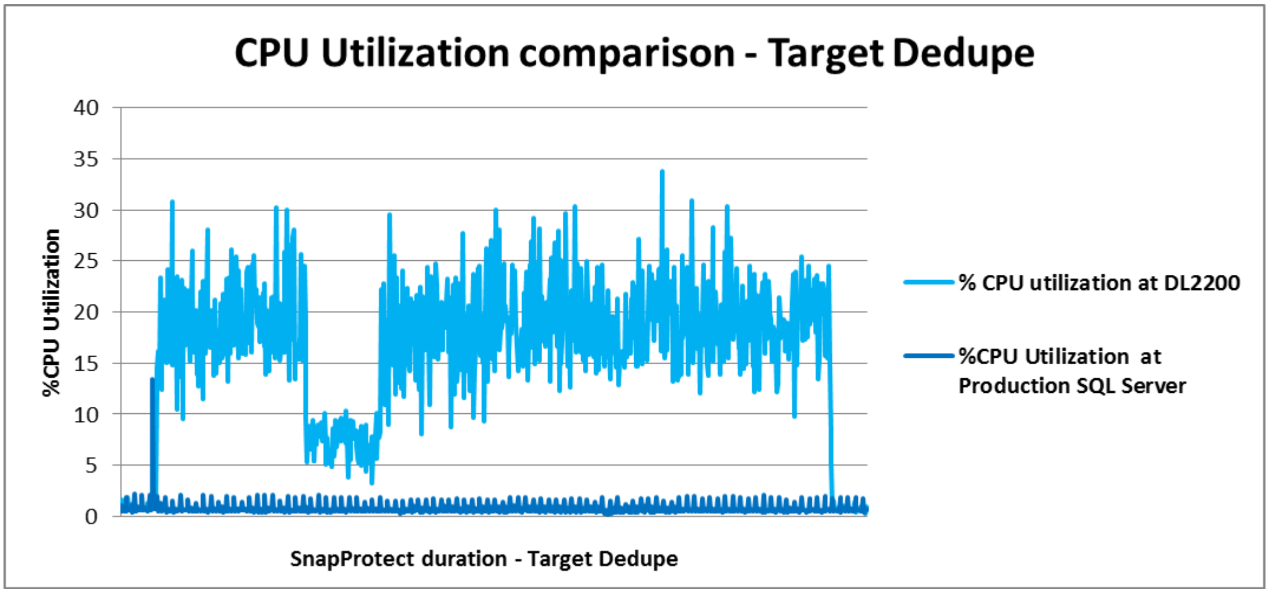

Figure 28 Impact on CPU Utilization at production server vs. DL2200

Next, in the target deduplication impact studies, the steps specified in Table 8 were performed for source and target dedupes. The backup duration measured at the DL2200 includes the time it takes for snapshot creation, movement to proxy, complete indexing, and then copy to the disk target location.

Table 8 Deduplication steps performed or this test

| Serial. No | Steps | Description |
|---|---|---|
| 1 | Dedupe#1- Full Backup | Performed a full backup with target Deduplication (DL2200 as proxy) on an off-host SnapProtect backup. |
| 2 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate differential. |
| 3 | Dedupe#2- Full Backup | Performed a full backup which has some differential compared to the full backup performed in step 1 |
| 4 | User Transactions on Database | Ran user transactions from QBMF on the database for 30 minutes to generate data change. |
| 5 | Dedupe#3- Full Backup | Performed a full backup which has some differential compared to the full backup performed step 2 |
| 6 | Dedupe#4- Full Backup | Performed a full backup. This is to verify the capacity savings when there is no data change. Since no user transactions were run after step 5, there is no change in data; hence the capacity savings is expected to be ~ 100%. |

Target Deduplication greatly reduces the impact on CPU utilization at the production SQL server and provides backup capacity savings for the backup data. Figures 29 and 30 compare the backup duration and backup data size reduction in source dedupe with production SQL server as proxy and target dedupe with DL2200 server as proxy. Target dedupe took longer to complete since resource utilization at DL2200 was high due to dedupe and backup processing at the DL2200. This process is similar to source deduplication with the DL2200 as proxy. Both target dedupe and source dedupe with DL2200 yielded similar results and either one can be chosen for backup capacity savings.

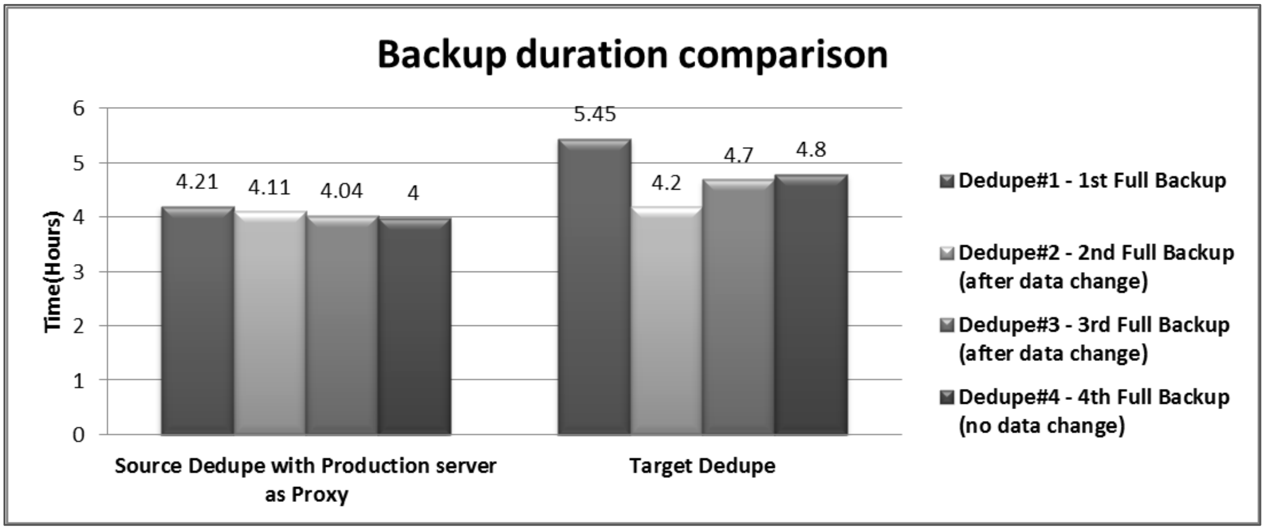

Figure 29 Backup duration comparisons

The capacity reduction was almost the same in both test scenarios (Refer to 30).

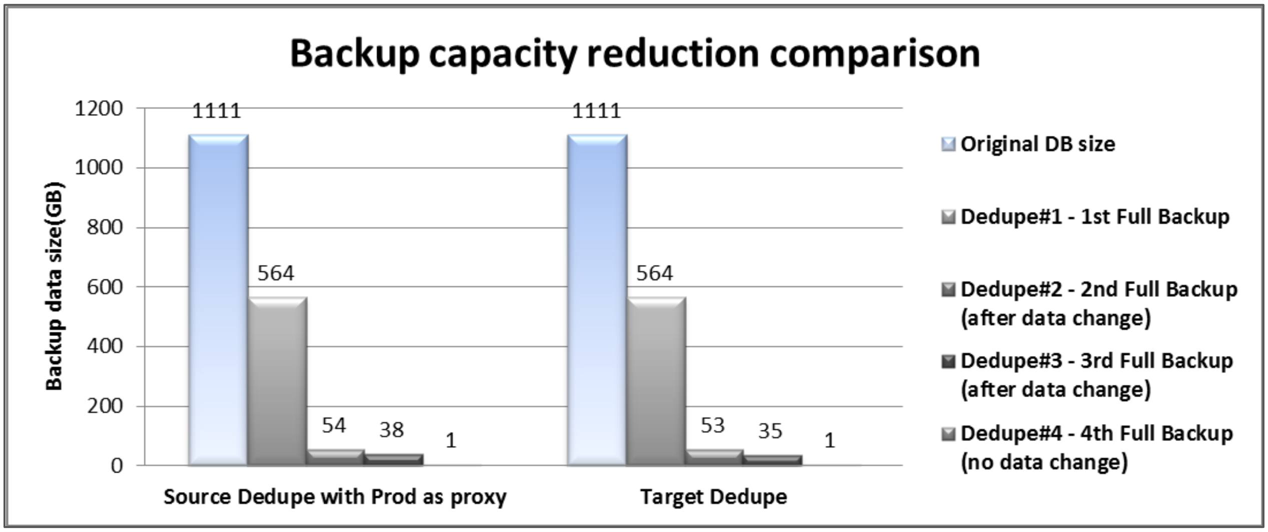

Figure 30 Backup data size reduction comparison

## 3.5 Analysis and conclusion for Dedupe studies

SnapProtect backup greatly reduces the impact on the resources at the production SQL server during backup when compared to the backups performed using native SQL backup utility.

Source and target Deduplication have their pros and cons. The usage depends on the goal for the backup operation. For instance, if the goal is to reduce the impact on the production SQL server during backup, then target deduplication would be a better choice since all dedupe and backup processes are off-loaded to the proxy server specified in target dedupe. Source Deduplication without using off-host proxy would be ideal for leveraging the existing network infrastructure or even for slower infrastructure since only unique segments are transferred across the network. However as long as additional CPU usage is acceptable at the production server, then source deduplication without using off-host proxy would be valuable. Benefit can be achieved from both of the dedupe scenarios for backup capacity reduction at the target array.

In SnapProtect while using the source dedupe, the proxy feature can be used to move where the deduplication processing takes place. Refer to section 3.4.2 for the details. When a backup server such as the DL2200 is used as the proxy in source dedupe, the backup takes place off-host resulting in minimal impact at the production server. Even though the dataflow looks similar to target deduplication, there are minor differences in the way the agents perform dedupe and backup, hence the difference in backup duration between target dedupe and source dedupe with DL2200 as proxy.

Both source and target Deduplication offer significant capacity savings on subsequent full backups at the backup target array. In the tests performed, the capacity savings during backups with either source or target deduplication, improved by 49% from the original backup data size to the first full dedupe backup, 95% to second full backup after data change, 96% to third full backup after data change, and 99% at the fourth full backup with no data change.

**Note:** These results are specific to our test environment and would vary depending upon the nature of data, number of storage arrays, type of arrays, RAID configurations and CommVault Service Pack version used. However, the analysis presented provides insight on the impact on Production SQL Server using SnapProtect and Deduplication.

Since the publishing of this paper, CommVault has released further enhancements on performance by providing multi-threaded support for individual and multiple databases spread across multiple volumes. These enhancements may provide better performance scalability by leveraging the multiple volumes for multithreaded backups during backup copy execution from Dell EqualLogic SAN storage.

# 4 Best practice recommendations

## 4.1 Network infrastructure

The following are network infrastructure design best practices:

- Design redundant SAN component architectures. This includes the NICs on the servers and switches for the storage network (including server blade chassis switches and external switches).
- Make sure that the server NIC ports and storage NIC ports are connected so that any single component failure in the SAN will not disable access to any storage array volumes.
- Enable flow control on both the server NICs and switch ports connecting to the server and storage ports.
- Enable jumbo frames on the server ports and switch ports.
- Disable spanning tree on switch ports connecting to end devices such as server ports and storage ports. Enable PortFast for these ports instead.

**Note:** General recommendations for EqualLogic PS Series array network configuration and performance is provided in the document titled *PS Series Array Network Performance Guidelines*.
http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229

## 4.2 Storage

Use the following best practices when configuring Dell EqualLogic storage arrays for a data protection solution.

- When possible, use high performance drives for the arrays that host the SQL Server database volumes. For the test configuration in this paper, 15K RPM SAS drives were used on the PS6000XV hosting the database volumes.
- Dedicate separate storage pools for production database data volumes and backup volumes for data protection reasons.
- Choose the appropriate RAID type for the PS Series arrays hosting the source database volumes and backup target volumes.
    - Microsoft best practices call for deploying SQL Server log files on RAID 10 volumes when planning for IO intensive workloads such as OLTP where possible for best performance and protection from failures. Since log is write-intensive mostly, RAID 10 would provide better throughput for write-intensive operations .For SQL implementations, RAID 50 can be used to provide maximum storage capacity, in addition to the performance benefits of striping.
    - The test configuration in this paper used RAID 10 for the PS6000XV source array that hosts data and log volumes and used RAID 50 on the PS6500E target array that hosts backup volumes.

**Note:** General recommendations for deploying SQL server in EqualLogic PS Series array is provided in the document titled *PS Series Grouped deploying Microsoft® SQL Server in an ISCSI SAN*.
http://www.equallogic.com/uploadedfiles/Resources/Tech_Reports/tr1013-sql-server.pdf

**Note**: The EqualLogic Auto-Snapshot Manager/Microsoft Edition application in conjunction with SQL backup can help to improve RPO and RTO goals without any disruption to the database applications. ASM/ME helps to create transactionally consistent smart copies (snapshots, clones or replicas) of SQL Server® databases. An in-depth discussion on SQL Server data protection using EqualLogic smart copy snapshots can be found at http://en.community.dell.com/techcenter/storage/w/wiki/enhancing-sql-server-protection-using-equallogic-smart-copy-snapshots.aspx.

## 4.3 VMware vSphere ESXi Server/VM

The lab test environment for this paper was comprised of a VMware ESXi server to host SQL Server database virtual machine as well as the Quest Benchmark factory work load simulation, vCenter, and Active Directory virtual machines. The following best practices are applicable for running VMware ESXi based virtual machines in conjunction with EqualLogic storage and/or Microsoft SQL Server environments.

Virtual machine and guest OS configuration

- For these tests, iSCSI SAN storage access was setup for Windows based virtual machines to use a direct access path and the guest OS (Windows) iSCSI initiator, as illustrated in Figures 1 and 2.
- When using the Windows 2008 Server iSCSI initiator within a virtual machine (guest iSCSI), the following best practices 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.
    - Install the Dell EqualLogic Host Integration Toolkit (HIT) for Windows within the guest OS. This installs the EqualLogic DSM for the Windows Server MPIO framework. The DSM provides multi-path optimizations tailored to the EqualLogic storage arrays.

**Note:** The iSCSI volumes were natively formatted as NTFS and directly accessed within the Windows 2008 Server VM.

## 4.4 Data backup using CommVault Simpana and EqualLogic

The following section contains best practices for data backup using CommVault Simpana SnapProtect backup from and to EqualLogic arrays.

### 4.4.1 General SnapProtect best practices

**Data retention/aging**

Data Aging is the process of removing old data from secondary storage to allow the associated media to be reused for future backups. By default, all backup data is retained infinitely. However, the data retention should be changed based on the needs. More information on Data aging can be found in *CommVault Online Documentation.* http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/data_aging.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/features/data_aging/data_aging.htm

In SnapProtect, while snapshots provide a fast and efficient manner to create copies of SAN volumes, the snapshots are still stored with the SAN volumes and may share data with the base volume. This means that both primary application data and its snapshots are vulnerable to catastrophic loss scenarios. So it is recommended to design backup environments that ensure copies of data are regularly created and backed up to an external storage (Example: The DL2200 used in this paper) before the snapshot's retention policy expires.

- The jobs for the snapshot are pruned based on the retention policy of the snapshot copy.
- The snapshots related to the pruned jobs are deleted from the array periodically.

**Snapshot Reserve**

Snapshot reserve space is needed to be allocated before creating snapshots. For each volume where the data needs to be protected, ensure there is sufficient Snapshot reserve space. It is configured as a percentage of the volume reserve (allocated space) for the volume.

Snapshot reserve defaults to 100% of the base volume capacity. This is to ensure that a single snapshot can be held even if 100% of the base volume is modified, but it can be changed at any time without taking the volume offline. Snapshot reserve can be increased to values greater than 100% (up to 10,000%) or it can be completely disabled by setting the value to 0%. To conserve free pool space, snapshot reserve can always be disabled for a volume if there are no plans to use snapshots.

Although snapshot reserve is not used until volume or snapshot writes occur, it is immediately consumed from free pool space and storage wouldn't report as "free" space and this helps in space planning thus avoiding other methods to track space commitments. For example, if a fully-allocated (not thin provisioned) 200GB volume is created and 50% snapshot reserve is specified, free pool space is reduced immediately by 300 GB (200 GB for the volume reserve and 100 GB for snapshot reserve), even though there are no pages in use. Therefore, before creating a volume, consider how much snapshot reserve, if any, to allocate to the volume.

**Configuring magnetic libraries**

Writers and mount paths are used in either *Fill and spill* or *Spill and Fill* allocation manner. Fill and Spill means the first mount path disk space or writers must be filled before the resources of the next mount path are used.

Spill and Fill assigns one writer from each mount path in turn until all mount paths are being used before adding additional writers to each mount path. This distributes the data transfer operations across more disk resources, typically yielding better performance.

**Backup data verification**

CommVault Simpana software typically protects/archives data on various types of media. Once these operations have been performed, there is no way to establish if the data is valid for recovering until a data recovery operation has been performed on that data.

It is important to confirm the validity of the backups to avoid future data loss with data verification operation. Data is checked to ensure it is valid for recovery and for a successful auxiliary copy operation. Data can be verified on all copies or on a specific copy. More information on Data Verification can be found in *CommVault Online Documentation.*

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/features.htm?var1=http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/features/data_verification/data_verification.htm

**Performance**

- If the main goal of database backup is to reduce the impact on the production server and optimize application response times then SnapProtect would be an optimal choice. If in addition, backup capacity savings is another goal then SnapProtect with deduplication would be the proper choice.
- Instead of running a single backup job at the CommCell console, running multiple offline backup copy jobs enable multiple reader and writer streams to execute and this in turn help achieve better performance compared to a single job in SnapProtect.

**Things to remember in SnapProtect**

- SQL iData Agent supports the SnapProtect backup only for full and differential backups. Transactional Log backups always use the traditional backup method. Log backups are stored in the Primary (classic) copy.
- In SnapProtect, differential snap backups behave like full backups if the differential data is more than 64k.
- It is a good practice to create the destination restore folders manually before restoring a snap backup to avoid a restore failure.

### 4.4.2 Deduplication

Deduplication provides a smart and proficient method to store data by identifying and eliminating the duplicate items in backups. It is essential to evaluate the performance and capacity needed on the disks used to store the dedupe datastore. CommVault recommends the below evaluation methods for the same.

Table 9 Dedupe datastore size calculation

| | |
|---|---|
| Running the Tool | 1. Create a folder on the disk where you wish to estimate the performance.<br>2. Copy the following file to this folder:<br>`C:/Program Files/CommVault/Simpana/Base/SIDB2.exe` |
| Usage | `SIDB2 -simulateddb -p <SidbLocation> -in <Instance#> [-datasize] [-dratio] [-blocksize] [tlimit] [-diskperf -tpath] [-user] [-password] [-domain]`<br>Where:<br>Options in [] denotes optional arguments<br>`-simulateddb` is the keyword to simulate the deduplication database to evaluate the disk compatibility for hosting the deduplication store.<br>`-p` is the location (an empty directory) where the deduplication store will be located.<br>`-in` is the instance of the software using the tool.<br>`-datasize` is the application data size in GB. Number.<br>`-dratio` is the expected deduplicaiton ratio. Number (default value is 5.)<br>`-blocksize` is the deduplication data block size in KB. Number (default is 128.)<br>`-tlimit` is the value in microsecond. Number (default value is 1000.) -`tlimit` and `-datasize` arguments cannot be used together.<br>`-samplesize` is the size of the sample. Number (default values is 10000.)<br>`-diskperf` and `-tpath`. Diskperf is the keyword to measure disk performance and tpath is the path of the disk. If you use `-diskperf`, `-tpath` is mandatory.<br>`-keepddb` is the option to keep the deduplication database files. The files are removed by default.<br>-`stopCounter` signifies how much additional iteration to process after reaching the threshold time. This is to limit spikes caused by caching. (Default value is 50.) |

| | |
|---|---|
| Example 1 | For the details on the projected average transaction time for an insert/query in the Deduplication database based on the size of the application data that is backed up, use the tool with the `-simulateddb` and `-datasize` options.<br>COMMAND<br>SIDB2 -simulateddb -in instance001 -p d:\dedup_store -datasize 500<br>SAMPLE OUTPUT<br>The disk is capable of hosting a deduplication DB for: 0.500 TB of Application Data Size 0.100 TB of data on disk 146.0 microseconds average Q&I overhead perblock Throughput for DDb server 3156 GB per Hour |
| Example 2 | For recommendations on the maximum application data size that can be backed up using the store based on the average access time for each record, use the tool with the `-simulateddb` and `-tlimit` options.<br>EXAMPLE<br>`SIDB2 -simulateddb -in instance001 -p d:\dedup_store -tlimit 1000` |
| Example 3 | For recommendations on disk performance, use the tool with the `-simulateddb` and `-diskperf` options.<br>EXAMPLE<br>`SIDB2 -simulateddb -in instance001 -p d:\dedup_store -datasize 100 -diskperf -tpath d:\disktest` |

### Dedupe Datastore disk performance calculation

Use the following steps to measure the disk throughput for the disk in which you plan to create the Deduplication Store.

Table 10 Dedupe Datastore disk performance calculation

| | |
|---|---|
| Running the Tool | 1. Create a folder on the disk where you wish to estimate the performance.<br>2. Copy the following file to this folder:<br>`C:/Program Files/CommVault/Simpana/Base/CvDiskPerf.exe` |
| Usage | `CvDiskPerf -READWRITE -PATH <SIDB path> -RANDOM -FILECOUNT <filecount> -USER <username> -PASSWORD <password> -DOMAIN <domain> -OUTFILE <outputfile>`<br>Where:<br>`-READWRITE` is the option to measure read/write performance.<br>`-PATH` is the deduplication store mount path to be tested for performance.<br>`-RANDOM` is the keyword to measure random read/write operations (Optional). By default, sequential read/write operations are measured.<br>`-FILECOUNT` is the number of files used in the read and write operations (Optional). Default value is 512.<br>`-USER, -PASSWORD,` and `-DOMAIN` are options to provide specific user credentials to impersonate access to the path provided in the –PATH option (Optional). By default, the application user-credential will be used. If domain name is not provided, then the default domain will be used.<br>`-OUTFILE` is the location of the output file to store the disk performance results (Optional). Default value is `'.\CvDiskPerf.txt'` |
| Sample Commands | `CvDiskPerf -READWRITE -PATH c:\SIDB01 -OUTFILE c:\temp\perf.txt`<br>`CvDiskPerf -READWRITE -RANDOM -PATH c:\SIDB01 -OUTFILE c:\temp\perf.txt`<br>`CvDiskPerf -READWRITE -RANDOM -PATH c:\SIDB01 -USER commuser –PASSWORD commpw -OUTFILE c:\temp\perf.txt` |
| Output | The details of the disk performance are stored in the output file provided in the `-OUTFILE` option. The contents of a sample output file are given below:<br>`DiskPerf Version        : 1.0` |

```
Path Used                 : f:\
Read-Write type           : RANDOM
File Count                : 500
Total Bytes Written       : 1048576000
Time Taken to Write(S)    : 7.113515
Throughput Write(GB/H)    : 494.217709
Total Bytes Read          : 1048576000
Time Taken to Read(S)     : 7.581667
Throughput Read(GB/H)     : 463.700792
```

Ensure that the average read throughput of the disk is around 500 GB per hour, and the average write throughput of the disk is around 400GB per hour. More information on dedupe datastore disk performance can be found in *CommVault Online Documentation*.

http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/english_us/prod_info/dedup_disk.htm

# Appendix A 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 11 Test Configuration Hardware Components

| Test Configuration – Hardware Components | |
|---|---|
| SQL Server®<br>(ESXi01) | One (1) Dell PowerEdge R710 Server running VMware ESXi v4.1, hosting a single SQL Server® Database virtual machine:<br>BIOS Version: 3.0.0<br>2 x Quad Core Intel® Xeon® E5520 Processors 2.26 GHz<br>48 GB RAM,<br>2 x 146GB 15K SAS internal disk drives<br>Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 6.2.14, driver version 12.6.0<br>Two (2) Intel Quad Port VT network adapters (Intel 8257 1Gb)-firmware level 1.3.19.12 |
| INFRA SERVER | One (1) Dell PowerEdge R710 Server running VMware ESXi v4.1, hosting a two (2) Windows Server 2008 R2 virtual machines for Active Directory and vCenter:<br>BIOS Version: 3.0.0<br>Quad Core Intel® Xeon® X5570 Processor 2.26 GHz<br>48 GB RAM<br>2 x 146GB 15K SAS internal disk drives<br>Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 6.2.14, driver version 12.6.0 |
| LOAD GEN SERVER | One (1) Dell PowerEdge R710 Server running VMware ESXi v4.1, hosting 1 Windows Server 2008 R2 virtual machine for Quest Bench Mark Factory:<br>BIOS Version: 3.0.0<br>Quad Core Intel® Xeon® X5650 Processor 2.26 GHz<br>48 GB RAM<br>2 x 146GB 10K SAS internal disk drives<br>Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 6.2.14, driver version 12.6.0 |
| MONITOR SERVER | One (1) Dell PowerEdge R710 Server with Windows Server 2008 R2 for SANHQ:<br>BIOS Version: 3.0.0<br>Intel® Xeon® X5650 Processor 2.26 GHz<br>48 GB RAM<br>2 x 146GB 10K SAS internal disk drives<br>Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 6.2.14, driver version 12.6.0 |
| DL2200 | One (1) Dell PowerVault DL2200 Server with CommVault Simpana.<br>BIOS Version: 1.6.3<br>Quad Core Intel® Xeon® X5620 Processor 2.26 GHz<br>24 GB RAM<br>6 x 500GB 7.2K SAS internal disk drives<br>PERC H700 Integrated Firmware version – 12.10.1-0001/Driver-4.31.1.64<br>PERC H800 Integrated Firmware version – 12.10.1-0001/Driver-4.31.1.64<br>Broadcom 5709c 1GbE quad-port NIC (LAN on motherboard) – firmware version 6.4.4, driver version 16.2.1 |
| Network | 2 x Dell PowerConnect 6248 1Gb Ethernet Switch<br>Firmware: 3.2.1.3 |
| Storage | 1 x Dell EqualLogic PS6000XV:<br>14 x 600GB 15K RPM SAS disk drives as RAID 10, with two hot spare disks<br>Dual quad-port 1GbE controllers running firmware version 5.1.1.(H1) |

| | |
|---|---|
| | 1 x Dell EqualLogic PS6500E:<br>148x 1TB SATA drives as RAID 50, with two hot spare disks<br>Dual quad-port 1GbE controllers running firmware version 5.1.1(H1) |
| **Test Configuration – Software Components** | |
| Operating systems | Host: VMware vSphere ESXi Server v4.1<br>Guest: Microsoft® Windows Server 2008 R2 Enterprise Edition (virtual machine):<br>o MPIO enabled using EqualLogic DSM for Windows when using guest iSCSI initiator<br>o EqualLogic Host Integration Toolkit(HIT) v3.5.1 installed |
| Applications | SQL Server® 2008 R2 Enterprise Edition<br>CommVault Simpana 9.0/Service Pack-5 on DL2200 |
| Monitoring Tools | EqualLogic SAN Headquarters version 2.1<br>Windows Perfmon |
| Simulation Tools | Quest Benchmark Factory version 6.5.1 |

# Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and your installations.

Referenced or recommended Dell publications:

- *PS Series Array Network Performance Guidelines*
  http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229

- *Protecting Microsoft® SQL Server® with an Integrated Dell / CommVault Solution*
  http://www.dell.com/downloads/global/solutions/Prot_%20micr%20sql%20ser%20int%20dell%20com%20sol.pdf

- *Off Host Data Protection: Dell PowerVault Backup to Disk Appliance powered by CommVault SnapProtect and Dell EqualLogic PS arrays*
  http://dell.commvault.com/files/WP_Off_Host_Data_Protection_SnapProtect_Dell_Equallogic.pdf

- *PS Series Groups Deploying Microsoft® SQL Server in an ISCSI SAN*
  http://www.equallogic.com/uploadedfiles/Resources/Tech_Reports/tr1013-sql-server.pdf

- *Microsoft SQL Server 2008 Backup and Restore using Dell EqualLogic*
  http://en.community.dell.com/techcenter/storage/w/wiki/3499.microsoft-sql-server-2008-backup-and-restore-using-dell-equallogic-by-sis.aspx

The following resource was referred for all CommVault SnapProtect functions and is recommended for additional information.

- *CommVault Online Documentation*
  http://documentation.commvault.com/commvault/release_9_0_0/books_online_1/default.htm

For EqualLogic best practices white papers, reference architectures, and sizing guidelines for enterprise applications and SANs, refer to Storage Infrastructure and Solutions Team Publications at:

- http://dell.to/sM4hJT

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.