Using Dell EqualLogic™ Auto Replication to Synchronize Remote Offices and Headquarters

Database Solutions Engineering

By

Jisha J

Dell Product Group
October 2009

Executive Summary

Several organizations operate in a branch model, where a number of remote sites are linked to a single headquarter office. This whitepaper presents Dell™ Reference branch model architecture with Dell™ PowerEdge™ T710 servers, Dell EqualLogic™ PS6000 and PS4000 arrays, Microsoft® Windows™ 2008 R2 and Microsoft® SQL Server™ 2008 SP1. This whitepaper discusses the use of Dell EqualLogic Auto Replication to synchronize the branch offices with the headquarter site. The effect of implementing replication at the branch office and the headquarter databases is also discussed in detail. This whitepaper also emphasizes the capability of PS arrays to perform the headquarter-branch reverse synchronization. The scalability of Dell EqualLogic PS4000 arrays at the branch sites was analyzed based on a standard OLTP workload.

Contents

Introduction

Many companies operate in a *Branch office* model, where a number of remote offices are connected to the company headquarters. Most of the day-to-day database transactions occur at the branch sites and the data at the headquarters is collected and used for analysis and reporting purposes. The databases at the branch offices are updated throughout the day and, periodically, the headquarter database needs to be synchronized with all the branch office databases.

Consider a retail store having a number of branches spread over geographically separated locations. All the sales transactions occurring at the retail sites need to be propagated periodically to the store headquarters located remotely. The consolidated data at the headquarters is analyzed and queried at a later point of time for reporting purposes or for data mining.

Figure 1 illustrates a typical branch office model with a headquarters site and number of branches, namely A, B, C and D.

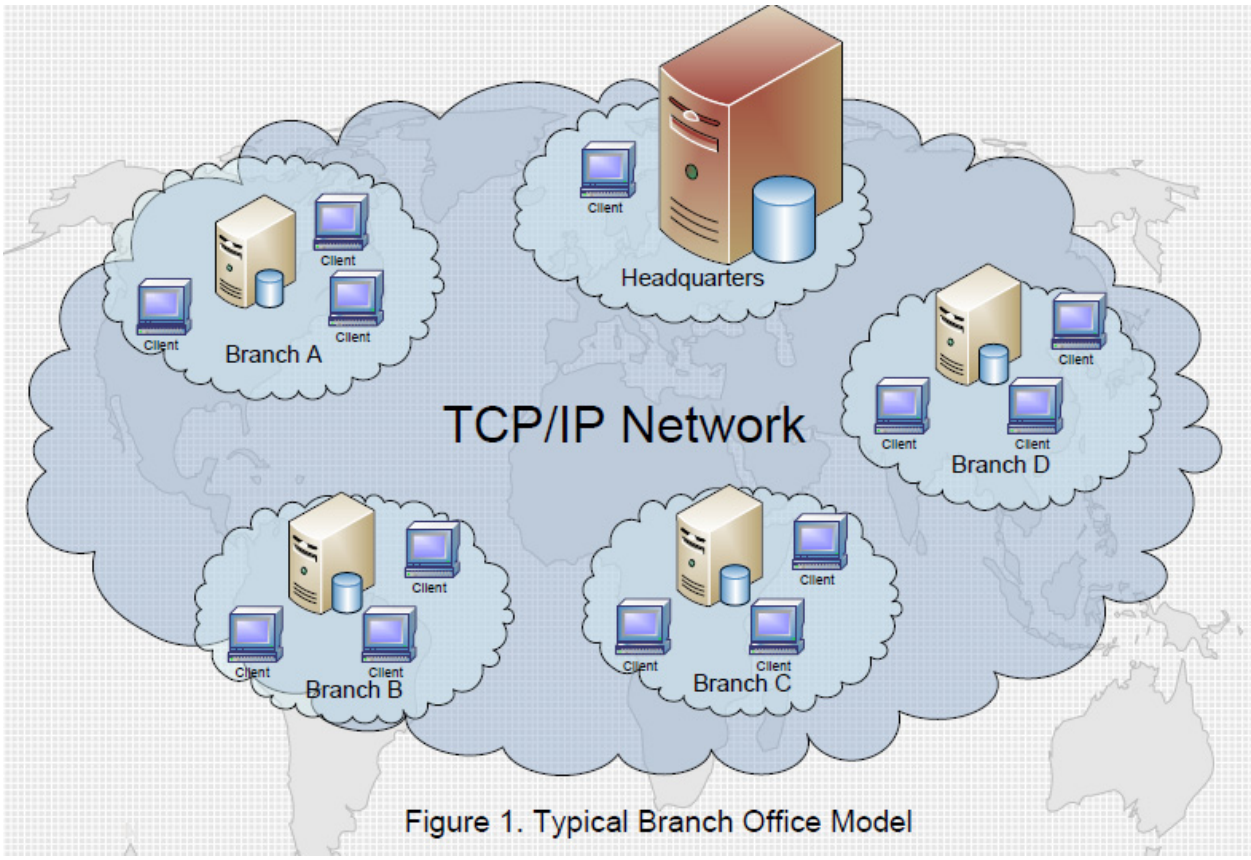

Figure 1. Typical Branch Office Model

One of the major considerations for implementing a branch office model is the data synchronization between the branch offices and the headquarters. There are several ways to accomplish this at the application/database level. The latest Dell EqualLogic™ storage arrays incorporate a number of techniques to achieve this at the storage block level.

Another important factor that is worth considering is the database and storage performance during the data synchronization activity. This factor will be dependent on the hardware implementation and application workloads.

Audience and Scope

This whitepaper is intended for customers, solution architects, storage administrators and database administrators who are interested in evaluating, planning or implementing a branch office model with Dell EqualLogic PS arrays.

This whitepaper provides reference architecture with Dell PowerEdge™ servers, Dell EqualLogic PS storage arrays, Microsoft® Windows operating system and Microsoft® SQL Server™ database. The paper talks about different methods of synchronizing the headquarter database with the branch office database, using the tools available from Dell EqualLogic. This paper will also present the experimental results helpful in evaluating the performance and scalability aspects of Dell EqualLogic PS4000 at any branch office implementation.

Dell Reference Architecture Overview

Dell provides various architectural building blocks to set up a complex branch office model. Some of the Dell components which may be implemented as part of a branch office model are discussed in the following sections.

Dell PowerEdge servers are high performing servers from Dell which may be used to host headquarter and branch databases. Dell PowerEdge T710 is a dual socket server which supports both dual-core and quad-core Intel Nehalem processors. It allows a maximum memory configuration of 144 GB. It supports SAS, SATA and SSD drives with 2.5" or 3.5" hard disk form factors.

Dell EqualLogic PS arrays are built on virtualized peer storage architecture that simplifies the deployment and administration of storage environments. Dell's EqualLogic PS6000 may be used to accommodate the datacenter and its critical applications. The Dell EqualLogic PS4000 may be configured for remote office and Small and Medium Business (SMB) storage deployments.

Microsoft Windows Server 2008 R2 comes with a number of features for the SMB segment. High fault tolerance and improved branch office security are some of the highlights of Microsoft Windows 2008 R2.

Microsoft SQL Server 2008 databases serve can be used as a good backend for the branch and headquarter sites. The close integration of SQL Server databases with the Dell EqualLogic arrays makes for a user-friendly and high-performing package.

Dell recommends the branch model reference architecture comprising of Dell PowerEdge T710 servers, Dell EqualLogic PS6000XV at headquarters and PS4000XV arrays at the branch sites, Microsoft Windows 2008 and Microsoft SQL Server 2008.

Figure 2 provides a high-level view of the Dell Reference Architecture of a branch model.

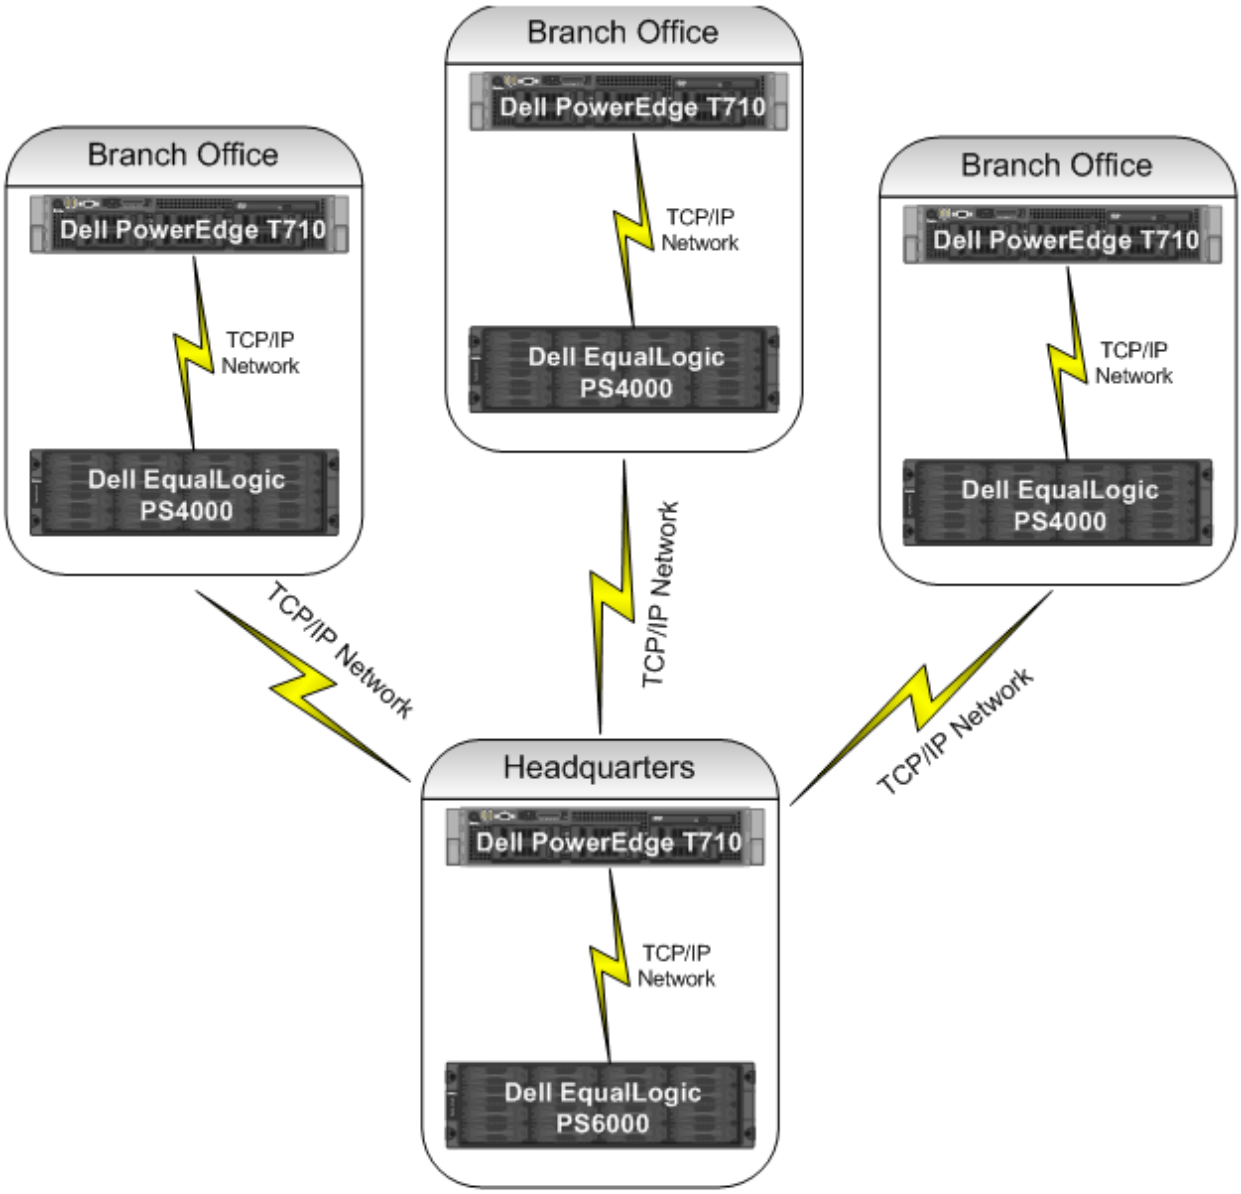

Figure 2. Dell Branch Model Reference Architecture

SAN Headquarters (SANHQ) is another tool from Dell which provides centralized performance and event monitoring for multiple Dell EqualLogic PS Series storage groups. This section discusses the use of Dell EqualLogic SAN Head Quarters in managing a branch office model.

SANHQ is capable of monitoring multiple PS storage groups. It is useful in the detailed monitoring of storage performance. It installs on an application server and logs information received via SNMP from the Dell EqualLogic storage groups. For monitoring from your workstation, a client-only installation

allows you to view the logged information without accessing the console of the server where you have installed SAN HeadQuarters(SANHQ).

SANHQ provides the detailed statistics related to the storage group, pools, members, volumes, inbound replicas, network ports and disks. This allows the administrators to get an idea of the overall storage performance and storage or network bottlenecks.

In this whitepaper, we have used the SANHQ tool to gather the storage performance details.

Synchronizing the Branch Office Database with the Headquarters

The branch databases need to be synchronized with the head office site at specific intervals. At the head office, the copy of the branch office data may be used for various operations like analysis and reporting, disaster recovery, etc. This section will discuss the usage of the Dell EqualLogic replication feature to accomplish the synchronization. This section will also discuss some headquarter- branch office synchronization considerations of the organization.

Dell EqualLogic Auto Replication

Dell EqualLogic Auto replication provides a powerful method for achieving branch office-headquarter synchronization. This section discusses the procedure for deploying Dell EqualLogic replication to synchronize the branch office databases with the headquarter site.

Auto replication uses the standard TCP/IP and ISCSI protocols for the asynchronous replication, which makes you capable of replicating data across any network. This feature requires no special equipment and, in addition, the distance barrier is not very significant in this case. The network considerations are discussed briefly in the later sections of this paper.

The replication is setup between the Dell EqualLogic storage groups (or partners), and not the individual storage arrays. This, in turn, helps in having variable storage capacities between the replication partners. An individual peer storage group may be partnered with up to 16 peer storage groups, which makes it a good solution for branch model, where data from a number of branches may have to be replicated to a centralized location.

Dell EqualLogic Auto replication may be deployed to meet any of the following objectives.

1. Replication
   Replication of the SQL Server branch database creates an exact copy of the branch database for use at the head office. This copy may be accessed from the headquarter site for updates, analysis and reporting purposes. During the replication configuration you may choose to keep the failback snapshot on the local array to enable the disaster recovery of the database volume on the same site.

2. Change Propagation
   The Dell EqualLogic auto replication allows having only the changes be propagated to the headquarter storage. The subsequent replicas contain only the differences from the last replica. Each of these replicas may be restored as a separate database drive to verify or analyze as separate point-in-time copies of the databases.

3. Archiving
   The replication may also be used as a means of archiving the branch data at a remote site, for example the headquarter site or at any other branch offices.

Centralized Replication

The ability of the PS storage array group to be partnered with any other PS storage group makes it capable of having a *Centralized Replication* model, depicted in Figure 3. Here the branches A, B and C are linked to a central headquarter site, HQ. At headquarter site HQ, space is delegated for each of the branch databases with which it is connected. In this model, the branches A, B and C will be able to replicate PS array volumes to the HQ site. During PS auto replication, the replicas transferred from each branch consume the space reserved for each of them. The local reserve space at the branch sites ensure the consistency of the volume changes during replication.

Note: Local reserve space should be allocated on all the branch offices and replica reserve space for all the branch office volumes should be created on the headquarter array.

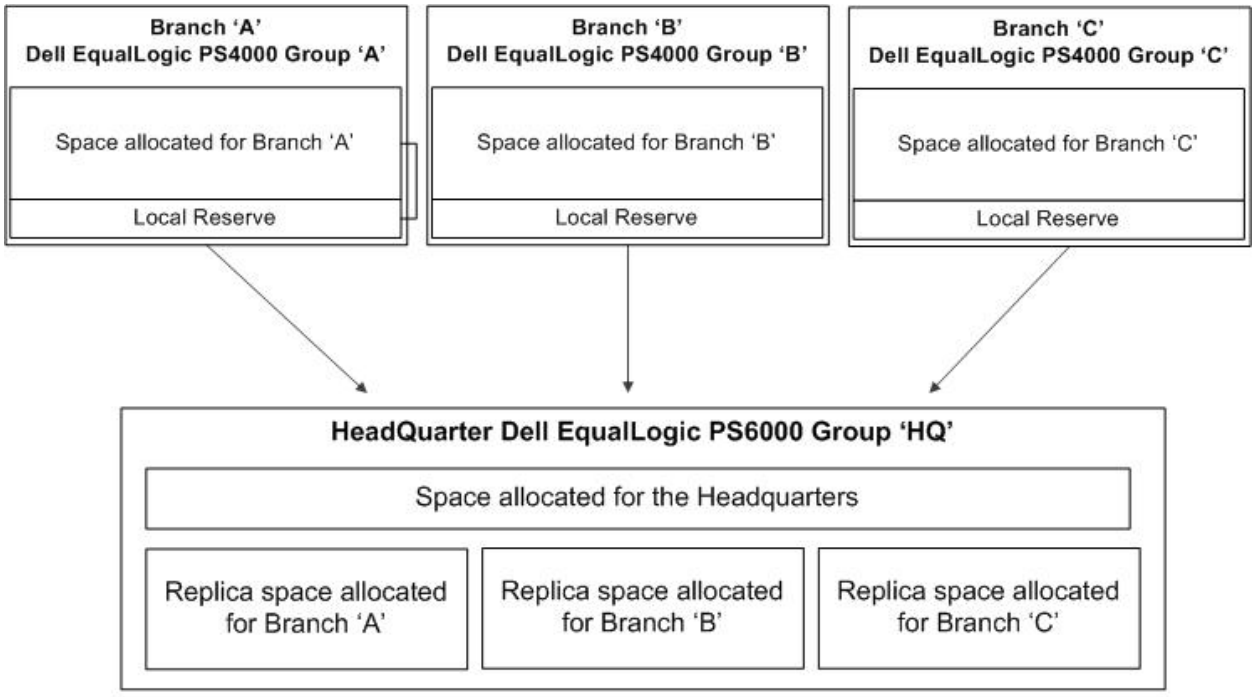

Figure 3. Centralized Replication

Procedure for Replicating a Branch Database to the Headquarters

Replication may be configured either using the group manager interface or the Dell EqualLogic Host Integration Tool kit (HIT). One of the major components of HIT is Auto Snapshot Manager (ASM) VSS provider and VSS requestor. The EqualLogic ASM integrated with the Microsoft Volume Shadow Copy Service(VSS) provides the ability to create data and application consistent Smart Copies of NTFS volumes and VSS aware applications including SQL Server Databases. If the group manager is used to replicate databases spanning multiple volumes you will have to manually create the volume collection set for all the volumes spanned by the database. The replication using either collection set or smart copy set ensures consistency of the database being replicated.

At a high level, the procedure for setting up Auto Replication between two Dell EqualLogic storage groups consists of two steps:

1. Configure Replication Partner both at the source side and destination side:
   The two storage groups between which the replication is configured are called *Replication Partners*. Each PS storage array group may be configured to have a maximum of 16 replication partners. In the branch model, this restricts the number of branches configured with the headquarters to 16.
   Using the Dell EqualLogic group manager, each of the PS array groups should be mutually configured as partners. For example, each branch array group should be configured as Replication Partner from the headquarter array group. Also, the headquarter array group should be configured as Replication Partner from each branch array group.
   The following details should be provided for configuring an array as a Replication Partner.
   a. Exact group name and ip address of the partner
   b. Passwords for the mutual authentications between the replication partner storage groups
   c. Delegated space for storing the replicas from partner
      Delegated space is the total space allocated on a PS group that will be receiving incoming replicas. All the Remote Replica Reserves for all incoming volumes should fit inside the delegated space.
      At the branch office side, the delegated space may be configured as 0MB if the headquarters is not expected to replicate back to the branch. At the headquarter storage group, enough delegated space should be set to accommodate all the replicas from the particular branch.
2. Configure Volume Replication:
   After configuring replication partnership, all the database volumes at the branch office need to be individually configured to be able to replicate to the headquarter array. The following details need to be provided for configuring the volume for replication.
   a. The Replication Partner to which the volume is to be replicated
      Any volume may be configured to be replicated to any of the configured replication partners at a time.

b. Remote Replica reserve
   Replica reserve is the space reserved on the group to where the replicas are stored (secondary group); to store the replica set (The set of the subsequent replicas created on a specific volume).
c. Local reserve
   Local reserve is the space reserved on the group where the volume resides (primary group) to track the changes on the volume when replication is in progress.
d. Failback Snapshot
   This feature allows the volume replica to be temporarily hosted in the local reserve for failback operations.

Please refer to the PS Groups Sizing Replication Space at the Dell EqualLogic support site (www.equallogic.com/support) for detailed guidelines on sizing the delegated space, remote replica reserve and local reserve.

After the volume replication configuration, you are given an option to create the replica immediately or to schedule it later. Once all the database volumes are configured to enable replication, you may create a collection set if the database spans over multiple volumes. These operations can be done from the GUI, CLI or the ASM/ME window. (There is also an option to create the replicas manually, where replica will be created using a transportable media instead of having to replicate through the network.)

After the volume replication is completed, there are two options to make the database volume accessible to the headquarter site:

1. Clone replica:
   This option clones the replica to an independent volume on the headquarter array. The original replicas are kept intact in the replica reserve and may be used at a later point in time.
2. Promote replica set:
   Replica set refers to the set of replicas of a specific volume or collection. The replica set may be promoted to an independent volume in the array. In that case, the replica reserve and the volume (and the snapshot) space in the headquarter array will get freed up. Once a replica is promoted, you will need to reconfigure the volume replication to be able to replicate the same volume again to the partner. The promotion of a replica set is a permanent operation and may be used for disk archival.

You will have to add the headquarter server/initiator to the new volume (cloned replica or promoted replica) access list. After setting the right permissions, the iscsi initiator at the server should be able to connect to the newly added volume.

Since auto replication is an asynchronous replication, you will be required to detach the existing branch database and reattach the database files in the new cloned replica volume or the promoted replica volume. This requires a minimal downtime of a few seconds for the branch data at the headquarters. The database downtime should match the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) requirements of the organization.

RTO refers to the maximum time that a system can be down beyond which the organization cannot afford downtime. RTO requirements are driven by the mission critical nature of the business. RPO indicates the data loss tolerance of an organization. RPO refers to the maximum amount of data that an organization can lose before causing heavy damage to the organization. These factors may play a major role in the availability of database services hosted from the headquarters.

Figure 4 is a high-level, pictorial representation of the replication procedure for a centralized model between the branch array, A, and the Headquarter array, HQ. The replication partner configuration and volume replication configuration are one time activities (steps highlighted in light red). For periodic data synchronization using cloned replicas, the activities highlighted in light blue may need to be performed iteratively.

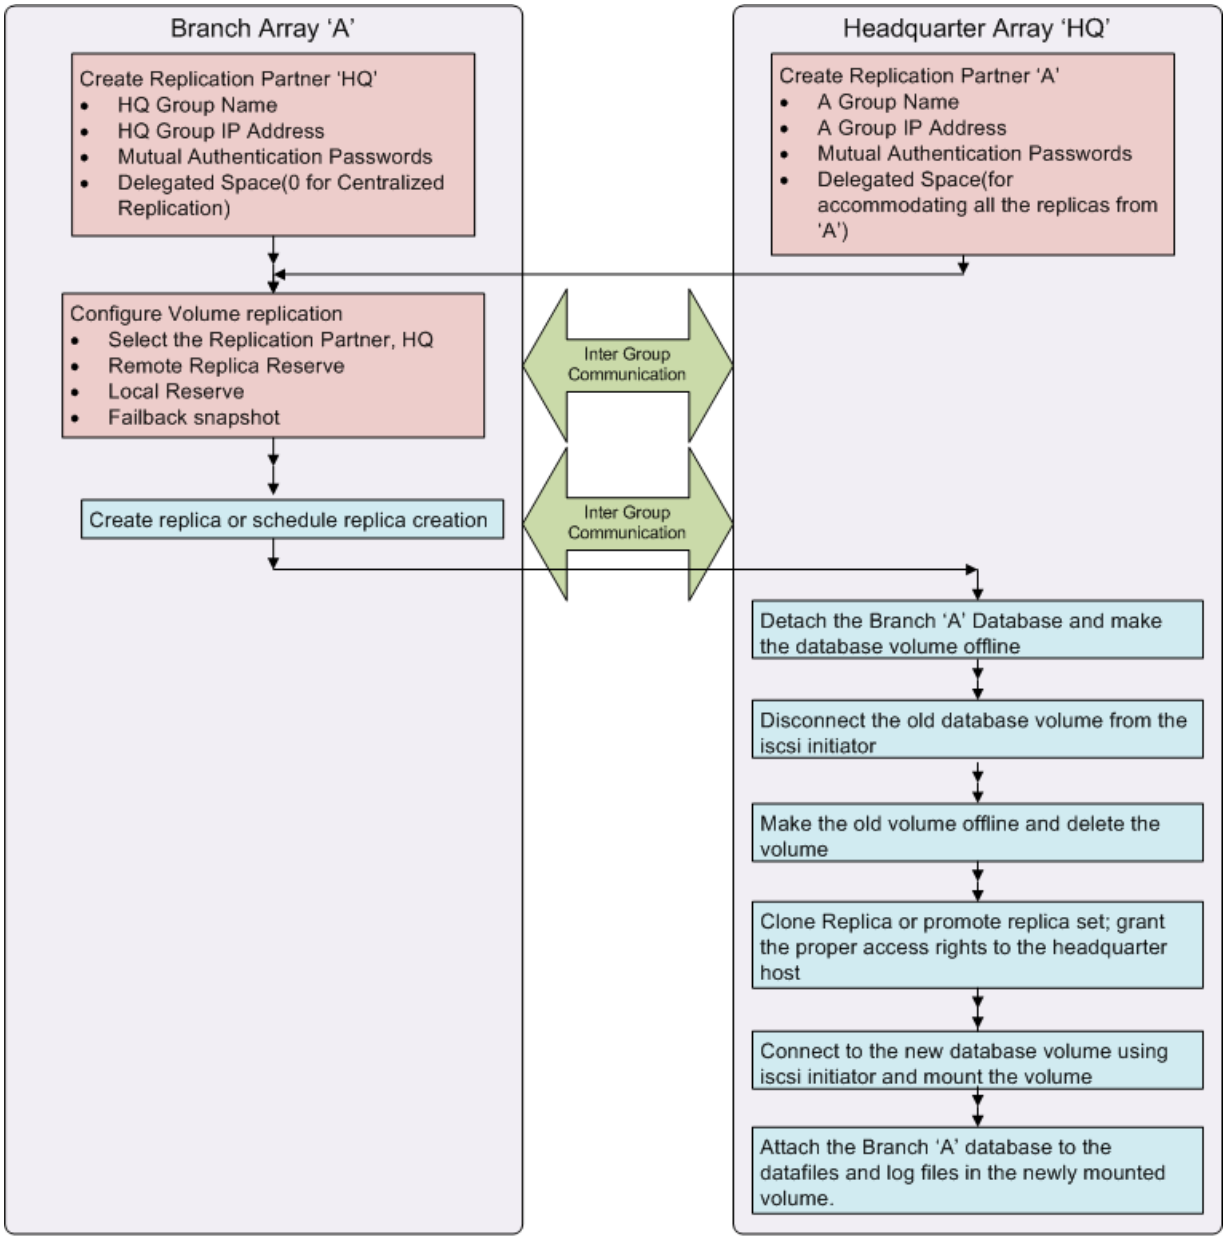

Figure 4. Procedure for replicating branch database to the headquarter site

Auto replication enables you to have multiple point-in-time copies of the same database volumes at the headquarter site. This makes it possible to have independent branch database copies of different time frames for data analysis and reporting.

For more information and detailed administration steps, please refer the PS Group Administration guide available at the Dell EqualLogic support site (www.equallogic.com/support).

Branch Database Change Rates and Network Characteristics

The initial replica of the database volume may always be almost equal to the original database size. The subsequent replicas will be dependent on the volume changes since the last replica creation. Therefore, once a volume is fully replicated to the destination array, the following replicas created tend to be of very small size.

The EqualLogic Replication replica size (in turn, the replication time) is not always directly proportional to the database volume change rates. The type of the database load is a major factor in determining the replica size. If the database workload is random, then you may find that the replica size is much different than the size expected. This is the same case with the local reserve and the replica reserve. For example, a 10GB database with random workload with a 2GB change may generate a replica size of 3GB or may generate a replica size of 1GB. The replica size is highly dependent on the rate of changes as well as the location of the changes.

The main network considerations on implementing Dell EqualLogic Auto Replication are the bandwidth and latency. The higher the latency, the lesser the bandwidth used for replication. For example, in the case of high latency connections, the replication will still work, but it may take more time to complete.

In high bandwidth-low latency scenarios, where the database change rate is high, it may be beneficial to schedule replications frequently. This ensures that the replicas are small in size and the replication process gets completed at a fast pace. In high bandwidth scenarios, where the database change rate is low, it is recommended to schedule replications at specific periodic intervals. This avoids unnecessary replication schedules.

In low bandwidth-high latency scenarios, the frequent replication schedules may overtax the network. In such situations, the replications may take a longer time to complete and there are chances of the network getting tied up by the replication process. In addition, if the scheduled replications are very far apart, the replica size may also increase to a greater extend, which increases the replication time. Care has to be taken in designing and scheduling replications under these conditions. For example, suppose a volume replication is scheduled every 30 minutes for a database volume with high change rate. In this case, there is the chance that the first replication will not complete before the next replication is scheduled to begin. If that happens the subsequent replications may be ignored.

The impact of network bandwidth and latency on replication may be summarized as follows:

	Low Bandwidth	High Bandwidth
High Latency	Very slow replication, as the replication speed is limited by both the low bandwidth and high latency constraints.	The replication speed may be slow, as it is limited by the high latency. Lesser bandwidth may be used for replication.
Low Latency	A higher percentage of the available bandwidth will be used for replication, depending on the latency.	Very efficient and high speed replication as the network parameters pose the least hindrance for the replication transfer.

Effect of Replication on the Branch Database

Auto Replication is executed between any two Dell EqualLogic storage groups. When replication is initiated, a snapshot of the volume is created on the source array. Using this data as a reference, data is sent to the destination array. When replication is initiated, a snapshot of that volume is created. That is used as a reference and data is sent to the replication partner. Replication I/O's are delivered at a lower priority than any new incoming I/O requests from hosts. The replication scheduling may have a minimal effect on the client I/O, but the frequent client requests may delay the replication time.

Tests were executed to analyze the effect of replication at the host or client. The following host scenarios were analyzed for full (initial) database volume replication in order to make sure that there was sufficient replication load at the storage level.

1. Replication on an idle host
2. Replication on a host under load

The test setup may be summarized as follows. The storage array replication partners were connected to each other using TCP/IP network spanning across a few routers. A PS volume of 250GB was mounted to the host machine (branch server) and a SQL Server database of around 106GB was installed on the volume. In the first scenario, replication was triggered with the database server being idle. The host statistics were captured from the host machine during the particular time frame. The CPU, memory and network statistics were captured using the Windows Performance Monitor (perfmon). The collected data showed that there was no measurable effect on the host while performing the replication. The performance parameters showed no significant change during the replication time. The details of the replication are as shown in Table 1.

Replication data size	106.48 GB
Replication Duration	20 min 6 sec
Replication Speed	5424 MB/min

Table 1. Replication on idle host

Since replication was carried out on the database volume for the first time, the replication data size is equal to the database size. If you schedule subsequent replications, the replica size will decrease and be proportional to the database/volume change rate. As replication duration and speed are largely dependent on the network bandwidth and latency, variations may happen to the replication statistics.

The second scenario was designed to capture the effect of database volume replication when the database was under load. To understand the effect of replication on the database writes, data of approximately 10GB was written to the database in approximately 20 minutes of time. Server and storage statistics were captured during the data load. In the second phase of the scenario, replication of the volume under load was initiated when the system was under load. Server and storage performance statistics were again captured during replication. The performance statistics may be depicted as follows, in figure 5a.

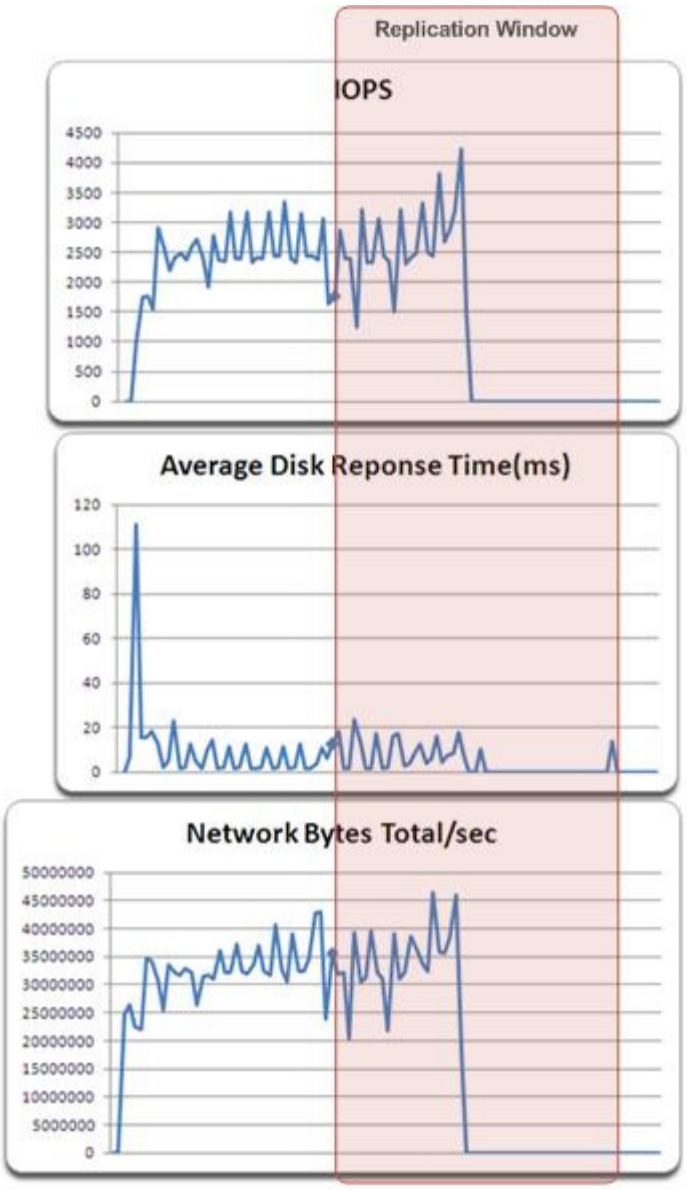

Figure 5a. Effect of Replication on Database Server I/O

From figure 5a, you can see that the IOPS and the network traffic seem to be following almost the same pattern regardless of the replica creation. We could see an increase in average disk response time from 5.7ms to 9ms. This is particularly because of the replica creation going on at the backend storage side. The colored window depicts the time period for the replica creation. The replica requests are at a very low priority compared to the database or client side I/Os. Therefore, the effect of replication on the database performance is very negligible. As the replica transfer is completely done through the TCP/ IP connection established between the replication partner arrays, the database server network performance was not affected by the replica transfer.

When replication is initiated, a snapshot of the specific volume is created and data is sent to the partner with this snapshot as the reference. Since the replication is executed at a lower priority, the replication

duration tends to increase when the volume is under load. Figure 5b shows the storage network details, from SANHQ monitoring tool, for the particular time period. The related workload and replication details are given in Table 2.

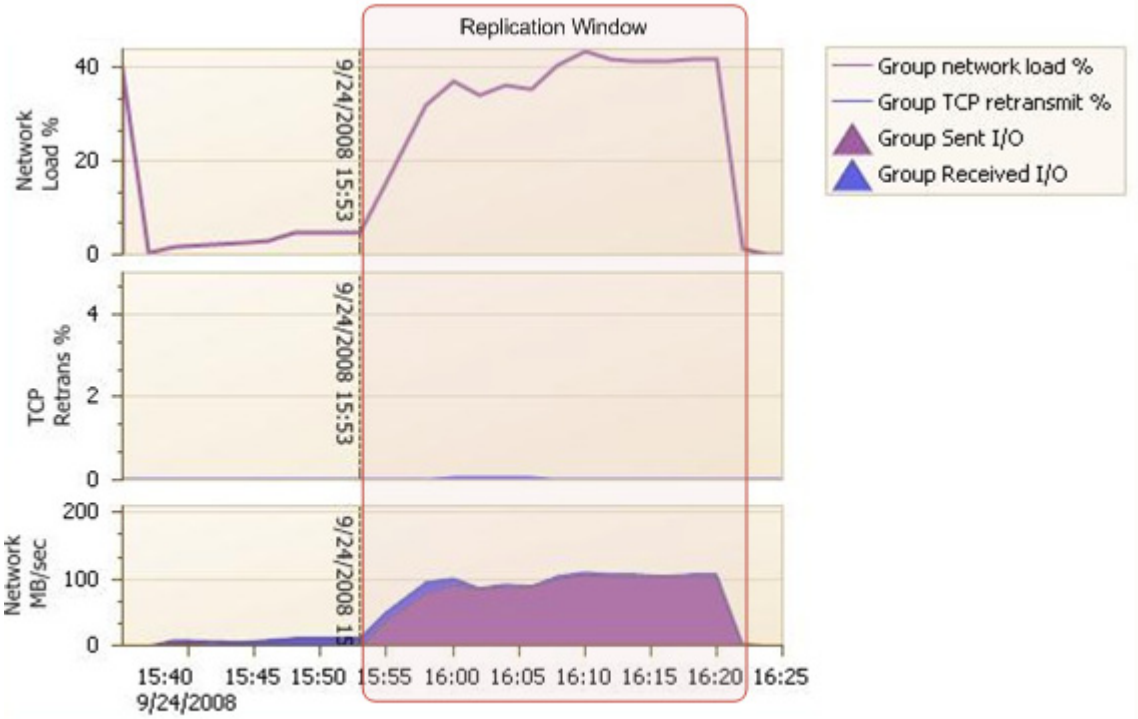

Figure 5b. Storage Network Details

Workload Details	Type	100% Write(10GB data write in 20 minutes)
	User load	50
	Start Time	15:40
	End Time	16:01
Replication Details	Replication Data Size	136.48 GB
	Replication Duration	25 min 36 sec
	Replication Speed	5459 MB/min
	Replication Start Time	15:53

Table 2. Workload and Replication Details

The *group received IO* is the amount of writes at the database. The *group sent I/O* represents the replica transfer load to the partner. After the workload was completed (in figure at 16:01), the network of the storage was utilized for the replica creation exclusively. We could see that the storage network load rose slightly above 40% during the replica transfer. The network throughput was maintained to be around 100MB/sec throughout the replica creation period. We observed that both the network interface ports of PS4000 were actively used during the activity.

It may also be noted that the 10GB writes contributed an additional 30GB to the overall replica size (The initial database size was around 106GB). This is because the replica size depends on the location of the change, in addition to the database change rate.

Based on these tests, we may conclude that the storage processes the database I/O with high priority compared to the simultaneous replica I/Os. The storage network bandwidth is shared between the database workload and the replica workload.

Effect of Concurrent Replications on the Headquarter Site

The headquarter site may be linked to a number of branch offices. Usually after off peak hours, all the branch office data may need to be synchronized with the headquarters. This necessitates concurrent replications to be carried out.

Multiple, as well as single, replications were initiated to the headquarter array (Dell EqualLogic PS6000) to understand the effect of concurrent replications on the headquarter site and headquarter array. The host and storage performance statistics were captured to study the behavior.

Single Replication

A database volume of 200GB was replicated to the headquarter array. The total size of the SQL Server database was approximately 136GB. There was no load running on the headquarter database server. The details of the replication are as given in Table 3.

Replication data size	136.48 GB
Replication Duration	30 min 36 sec
Replication Speed	4567 MB/min

Table 3. Single Replication

There was no noticeable impact on the headquarter server during the sampling time. As the replication is executed between the peer storage partners, it is not expected to impact the headquarter server. Figure 6a represents the storage network details captured by the SANHQ tool for the single replication.

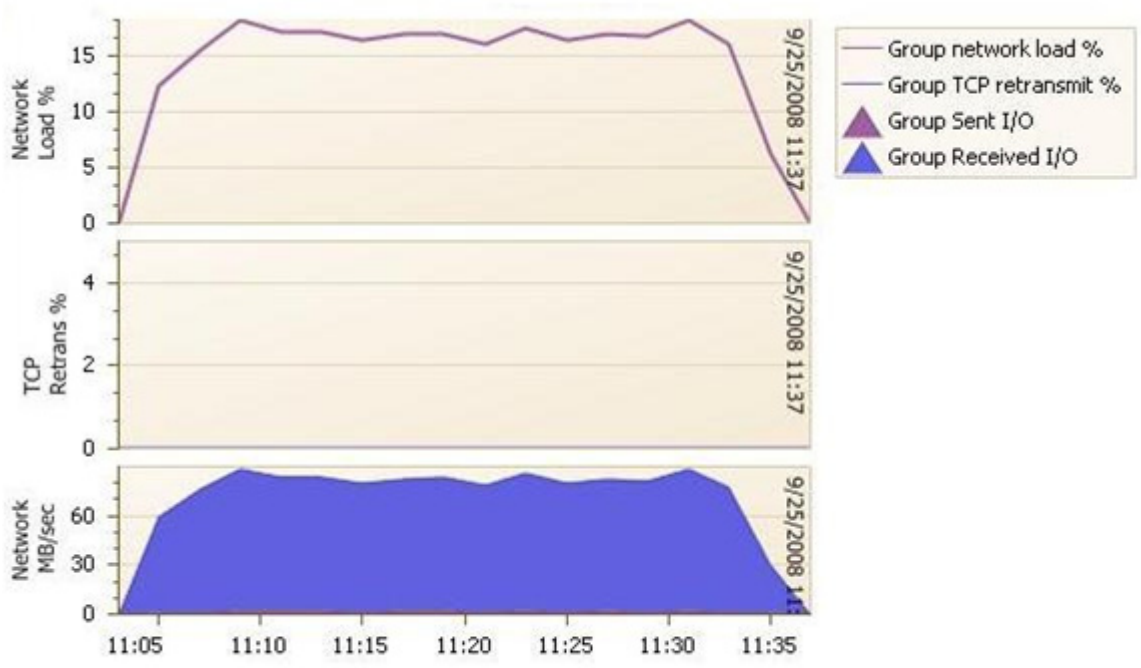

Figure 6a. Single Replication : Storage Network Details

The storage network load during the replication was approximately 15% and the storage array received the replicas at approximately 80MB/sec during the specified time interval.

Multiple Concurrent Replications

Three of the other PS6000 arrays were partnered with the headquarter PS6000 array. Four replications were manually triggered to the headquarter array simultaneously. The details of the replications are given in Table 4. Replications 1 and 2 were initiated from the same source array. Replications 3 and 4 were triggered from separate source arrays.

Replication	Replication Start Time(Approximate)	Replication Data Size(GB)	Replication Duration	Replication Speed(MB/min)
1	15:55	38.61	15 min 32 sec	2545
2	15:55	199.98	47 min 37 sec	4300
3	15:55	136.48	37 min 37 sec	3715
4	16:22	136.48	33 min 51 sec	4128

Table 4. Multiple Replications

From Table 4 we can see that even though the replications were triggered at the same time, replication 4 was executed at a later time. Only three of the four triggered replications were active at a single point of time. Note that the inactive replication may still show the *In Progress* status, but three of the replica transactions may be actively processed at any point of time, for a specific array group.

Figure 6b depicts the storage network details captured by the SANHQ tool for multiple replications.

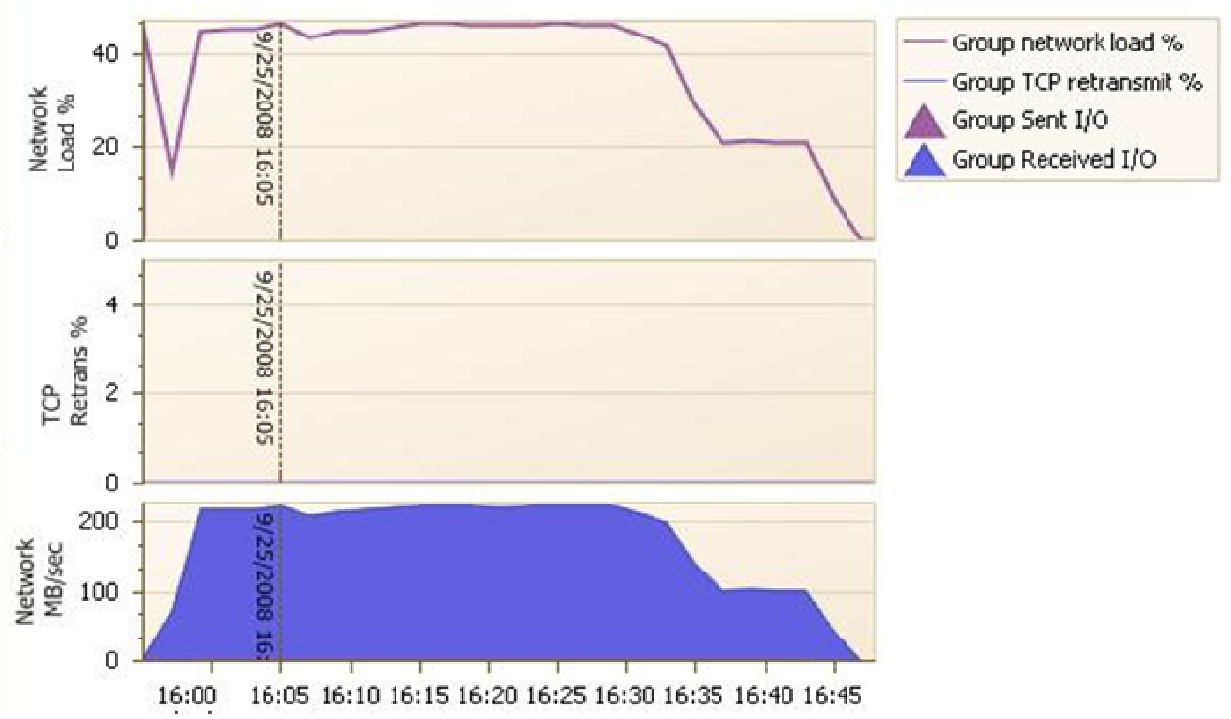

Figure 6b. Multiple Replications : Storage Network Details

The storage network load hit a peak of approximately 45%. The network throughput was 200MB/sec during the transfer. It may be observed that towards the end of the sampling period, the network load and throughput is reduced. This reduction is because the storage network utilization is dependent on the number of active replications and the replica transfer speed.

The vertical line in figure 6b denotes the point in time where the network utilization was at its peak value for the entire replication.

Headquarter-Branch Reverse Synchronization

There are some scenarios where it is required that data be push from the headquarters back to the branch offices. For example, suppose an organization planned to update the price details for a particular product for a specific branch. In that case, the data may be updated in the headquarter database and propagated to the branch database. Another example use case may be updating sales summary data back to the branches. These updates may take place less frequently and this reverse synchronization may need to be triggered on demand (i.e. as soon as the updates are completed at the headquarter site).

Branch reverse synchronization may be achieved using *Bi-directional replication* (or *Reciprocal Replication*), supported by Dell EqualLogic arrays. In this case, the replica space should be delegated at each of the branch sites to receive the replicated volumes from the headquarters. Once the specific

storage groups are configured as Replication Partners, each of the groups can replicate to each other freely. Each site administrator has complete control over their replicas but does not have access to the partner replica or volumes. At a high level, a simple bi-directional replication may be depicted as in figure 7.

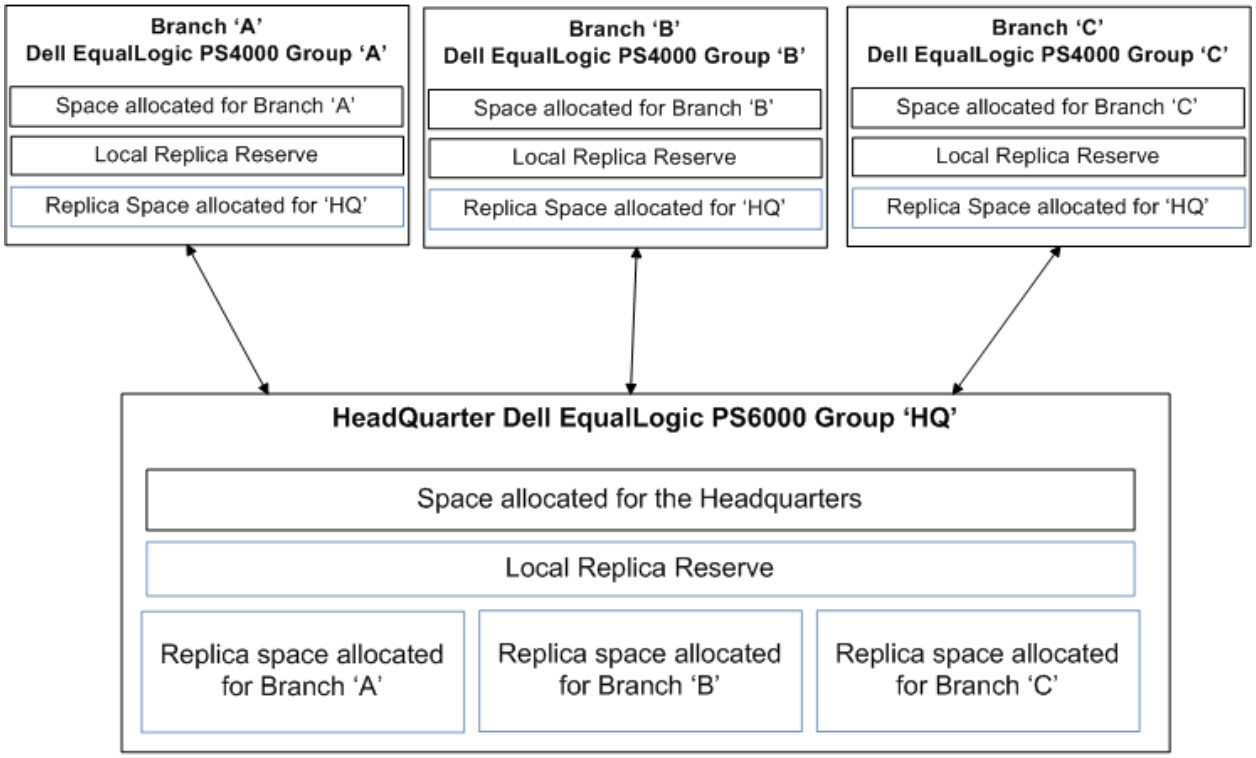

Figure 7. Bi-Directional Replication

One important factor to be noted is that frequent replications may overtax the storage/network bandwidth. The frequency of the replication should be determined based on actual business requirement and the network capability.

Headquarters SQL Server Database Instance Configuration

The headquarters database should have a consolidated view of all the branch office data. This data may be managed using a single SQL Server instance or multiple SQL Server instances. In other words, separate SQL Server instances may be deployed to manage individual branch office data or a single instance may be installed for managing all the branch office data. Both these scenarios have their own advantages and disadvantages.

A separate SQL Server instance consumes an extra set of system resources to load the SQL binaries to the memory. A separate SQL Server instance may be beneficial to segregate the sysadmin rights on one branch database from the other. The multiple instance memory management may be taken care by the SQL Server Enterprise edition Resource Governor. The tempdb database is another point of consideration. When you deploy separate instances to manage individual branch databases, each

instance will have its own tempdb space. In that case, the tempdb utilization of a database will not affect the performance of the other databases. Each of the SQL Server instances will be independent of each other. Using multiple instances also allows you to manage the availability and backup of the individual databases. If you are using a SQL server edition other than the enterprise edition you are required to pay for the multiple instances.

However, since a single instance may have less resource over head than multiple instances, having a single SQL Server instance may be faster than having multiple instances.

Database Workload Analysis

This section discusses the expected workload patterns at branch offices and headquarters. Experiments were carried out to understand the scalability of Dell EqualLogic PS4000XV arrays in the branch office context.

Workload Analysis of Headquarter Database

At the headquarter site, the branch data is most frequently used for analysis or reporting purposes. Therefore, the head quarter databases are usually characterized by large, long running queries. There may be a few updates or inserts to the databases, mainly the summary data changes, based on the business requirement. So, in general, we can expect more of an OLAP (Online Analytical Processing) workload at the headquarter site. The headquarter I/Os may be mostly be sequential in nature with large I/O size of approximately 1MB. The database and the operating system at the headquarters may be optimized for performance considering these factors.

Workload Analysis of Branch Office Database

Branch office sites are usually characterized by a huge number of inserts, updates and deletes. At the end of the business hours, data is propagated to the headquarters site. Very little querying happens at the branch sites. Therefore, the workload at the branch databases is mostly of OLTP (Online Transaction Processing) nature. The access pattern is more of a random nature, with a very small I/O size of approximately 8kb.

Scalability Study of Dell EqualLogic PS4000 Array

Dell EqualLogic PS4000 storage arrays serve as a good backend for the branch office SQL Server databases. OLTP workloads were executed on Dell EqualLogic PS4000 arrays to understand the scalability of PS4000 arrays at the branch office sites. The test configuration is as given in Table 5.

Components		Details
Hardware	Server	Model : Dell PowerEdge T710
		Processor : 2 *Quad core Intel Xeon Processors E5530 @ 2.40GHz
		Memory : 12 GB(3*4GB RDIMM 1067MHz)
	Storage	Model : Dell EqualLogic PS4000XV
		Firmware : 4.1.5
		Hard drives : 300GB 15k SAS drives
	Network Interface card	Model : Broadcom BCM5709C NetXtreme II GigE
		Driver version : 4.8.4.1
	Network Switch	Dell PowerConnect 6248
Software	Operating System	Microsoft Windows Server 2008 R2 Enterprise Edition x64
	Database	Microsoft SQL Server 2008 SP1 x64
Benchmarking Tools	Quest Benchmark Factory	Version : 6.0
	Windows Reliability and Performance Monitor(perfmon)	Version : 6.1.7600

Table 5. Test Configuration

The workload details are as given in Table 6.

Database Size	~100GB
Type of Workload	OLTP
User Load	5-50 users(at 5 user increments)
SQL Memory target	512 MB(to force the load to the storage)

Table 6. PS4000XV Scalability test -Workload Details

The Benchmark factory for Databases, a benchmarking tool from Quest, was used to simulate the real-time user load. The Windows Performance Monitor (perfmon) was used to capture the performance variables. The parameters Disk Transfer/sec, % Processor Utilization, Avg. disk sec/ Transfer and Avg. disk queuelength were collected on the database server to estimate the IOPS, CPU utilization, volume response time and disk queue length respectively.

The results of the experiment may be plotted as in Figures 8a, 8b, 8c and 8d.

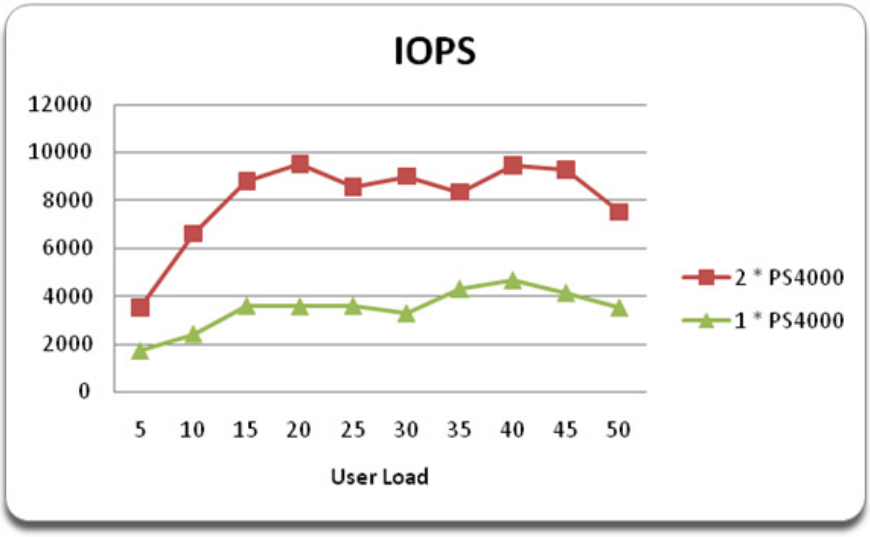

Figure 8a. PS4000XV Scalability Test - IOPS

As in figure 8a, the experiments showed that the SQL Server Database volume IOPS scale almost doubled on adding another PS array. The IOPS increased mainly because of the increased number of disks at the backend to carry out the requests.

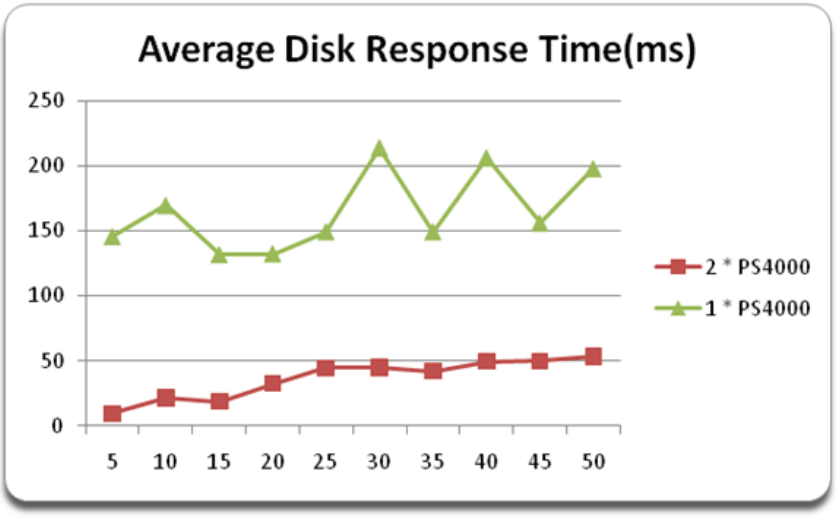

Figure 8b. PS4000XV Scalability Test - Average Disk Response Time

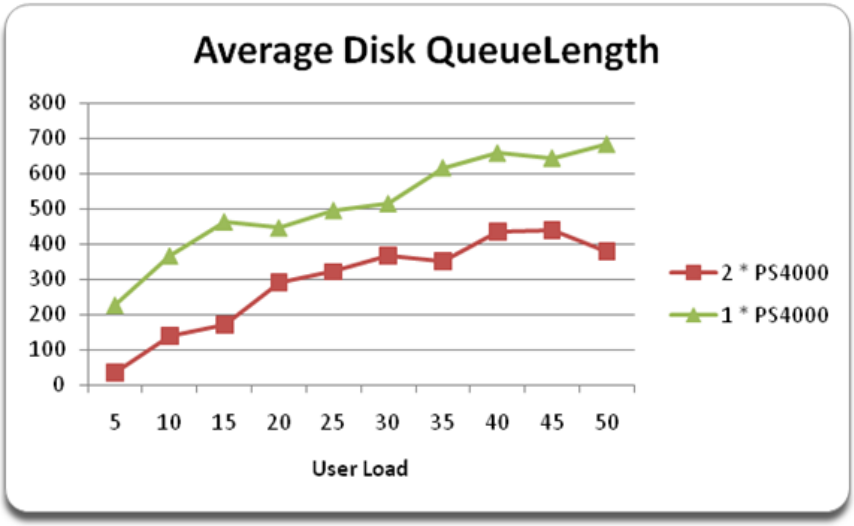

Figure 8c. PS4000XV Scalability Test - Average Disk Queuelength

Figure 8b depicts the improvement in disk response time. The queue length at the volume, given in figure 8c, also showed a significant reduction in numbers. Note that the heavy queues are mainly due to the database stress generated. In real-time scenarios, it is not desirable to have heavy queues to the database disk. The SQL Memory target is one of the parameters that may lead to heavy database disk queues. Here, the SQL Memory target was forced to 512MB to stress the underlying storage.

Conclusion

Dell recommends a branch model reference architecture comprised of Dell PowerEdge Servers, Dell EqualLogic PS arrays, Microsoft Windows and Microsoft SQL Server. One of the concerns in implementing a branch model is the synchronization of the sites. The Dell EqualLogic Centralized replication model provides an easy solution for updating the headquarter database with the branch office changes. The Dell EqualLogic Bi-directional replication model also helps in pushing the updates from the headquarters to the branch office. Auto Replication is an asynchronous operation which spans over long distances using TCP/IP protocols. It may be implemented to accomplish data replication, change propagation and archiving purposes.

Dell EqualLogic PS arrays process the replication requests at a lower priority so that the database workloads are minimally impacted during a replication operation. Also, since the replication is implemented between the PS storage groups, the network traffic on the host is not affected.

The nature of the expected workloads at the remote sites and the headquarters also differ. At the branch sites the workload is more of an OLTP nature, whereas, at the headquarters, most of the queries are executed for analysis and reporting purposes.

At the branch site, the Dell EqualLogic PS4000 arrays provide a cost effective, reliable and high performing database backend. It can be scaled up to two arrays per group to offer more capacity and performance. To summarize, Dell EqualLogic Auto Replication facilitates the easy and reliable synchronization between the headquarter office and the remote sites with minimal disruption of the host performance.

Glossary

Snapshot: Snapshot creates a point in time copy of the volume, which allows the volume (and in turn database) to be restored to that particular point. This enables to achieve business continuity in case of any logical errors. Snapshots are stored in snapshot reserve space allocated during the volume creation.

Clones: Clones are exact copies of the volumes, which may be used to replace the volume in case of any volume (drive) failures. Clones are created as a separate entity. Therefore, the clone and the base volume may be mounted as different volumes at any point of time.

Replica: smart copy set may be created for the volumes, if the volume replication has already been configured from the Dell EqualLogic Group Manager. Auto replication enables the volumes to be transferred among different storage array groups.

References

1. Dell PowerEdge T710

   http://www.dell.com/poweredge

2. The Dell EqualLogic PS Series

   http://www.dell.com/Equallogic

3. TPC Benchmarks

   http://tpc.org/information/benchmarks.asp

4. How to choose the right branch-office IT architecture

   http://www.microsoft.com/southafrica/midsizebusiness/businessvalue/branchofficeIT.mspx

5. Managing Branch Offices with Windows Server 2008 R2

   http://www.microsoft.com/windowsserver2008/en/us/branch-office.aspx#2008

6. SQL SERVER ADVANCEDPROTECTION AND FASTRECOVERY WITHDELL EQUALLOGIC AUTOSNAPSHOT MANAGER

   www.equallogic.com/.../WP911_SQL-Server_Advanced-Protection.pdf

7. PS groups: Sizing Replication Space

   www.equallogic.com/support

8. PS Administration Guide

   www.equallogic.com/support