# Dell Storage Center Replication and Oracle Data Guard

Dell Engineering
November 2015

A Dell Best Practices Guide

# Revisions

| Date | Description |
|---|---|
| November 2015 | Initial release |

# Acknowledgements

Engineering: Mark Tomczik

# Table of contents

# 1 Introduction

IT leaders are continually reacting to market changes and economic conditions by establishing efficient IT organizations to help transform their businesses. One of the top five priorities of business transformation is business continuity, which ensures that critical business functions are available during a disastrous event. Business continuity is significant because it is estimated that a very small percentage of companies survive a catastrophic data loss.

One important subset of business continuity is disaster recovery (DR). This is a culmination of processes, procedures, and policies that keep critical systems and IT infrastructure available during a disaster or planned outage. DR also includes characteristics like recovery time objective (RTO), which is the duration of time and service level when the business process must be restored, and recovery point objective (RPO), which is the maximum tolerable period that data might be lost. An optimal DR solution not only helps achieve RTO and RPO, but also reduces the costs and management of business continuity environments.

## 1.1 Objective

This paper illustrates how Dell™ Storage Center replication technology and Oracle® Data Guard work together to provide robust and comprehensive disaster avoidance and recovery solutions. This information enables customers to address disaster avoidance, disaster recovery, and non-disaster needs with more certainty and choose an applicable solution for the infrastructure.

This document is not a how-to or comprehensive user guide, but offers sample use cases and solutions that could be deployed to meet specific needs. Keep in mind that the suggestions and solutions presented in this document should be reviewed and adjusted for individual environments. Dell advises customers to understand the type of DR tools and replication types available, application and business processes, and requirements before designing and implementing a data protection strategy.

For detailed information on the Oracle installation and configuration used in this paper, visit the My Oracle Support site. Additional resources are listed in appendix C.

## 1.2 Audience

This document is intended for DBAs, system administrators, and storage administrators seeking DR solutions using Storage Center replication and Oracle Data Guard to protect against local infrastructure and data threats. Readers should be familiar with Dell SC Series arrays and Dell Enterprise Manager, and have prior experience in configuring and operating the following:

- RAID, Fibre Channel (FC), Serial-Attached SCSI (SAS), LUNs, multipathing
- General IP networking concepts
- Familiarity with Oracle Linux® and administration
- Oracle 11g databases and architecture
- Oracle Grid Infrastructure
- Oracle Real Application Clusters (RAC) or single instance database
- Oracle Automatic Storage Management (ASM)
- Oracle Data Guard

# 2 Technology overview

Replication technology shares data in a way that ensures consistency between multiple or redundant storage devices in different locations to improve high availability. It includes fault tolerance, accessibility, and reliability features to help provide business continuity and address local or regional data threats. Sharing data is performed by either hardware or software components. Both sets of components should be evaluated and benefits assessed before deciding on a replication solution. The replication method chosen depends on the platform being replicated, the applications, and business expectations.

## 2.1 Dell Storage Center replication

The Storage Center replication technology, Remote Instant Replay, provides advanced business continuity and DR solutions. Remote Instant Replay creates an exact mirror copy of a local volume onto a remote SC Series array. It can provide minimal data loss (asynchronous replication) to zero data loss (synchronous replication) between arrays as well as the functionality to efficiently simulate replications prior to go-live. This eliminates any impact to the source and target application environments during testing.

Some key benefits of replicating SC Series storage are:

- Minimal storage capacity using thin provisioning for replays
- Optimized management of replication definitions
- Rapid recovery from planned outages

In this scenario, a remote SC Series array is configured to communicate with a local SC Series array. If using FC rather than iSCSI protocol, both arrays must be connected to the same FC fabric, and the zoning must allow the local and remote arrays to communicate. When proper zoning is done, each array appears as remote to the other, and both must be added to Dell Enterprise Manager. Quality of Service (QoS) definitions control how bandwidth is used to send replication data between local and remote SC Series arrays. QoS definitions must be set before creating a replication.

While SCv2000 Series arrays have some replication functionality, the rest of the SC Series portfolio has full replication capabilities. For information on replication with SCv2000, see the *Dell Enterprise Manager 2015 R1 Administrator's Guide* available on Dell.com/support.

## 2.2 Oracle Data Guard

Oracle Data Guard allows from one to as many as nine standby databases to be managed, monitored, and automated to provide disaster protection and high availability for mission critical applications. Oracle Data Guard has the ability to switch roles of the database in the DR environment: The standby database becomes the new production database, and the original production database becomes the new standby database. The two types of standby databases are logical and physical.

## 2.3 Hybrid replication solutions

Both Dell SC Series replication and Oracle Data Guard can be deployed in the same environment, as illustrated in Figure 1.

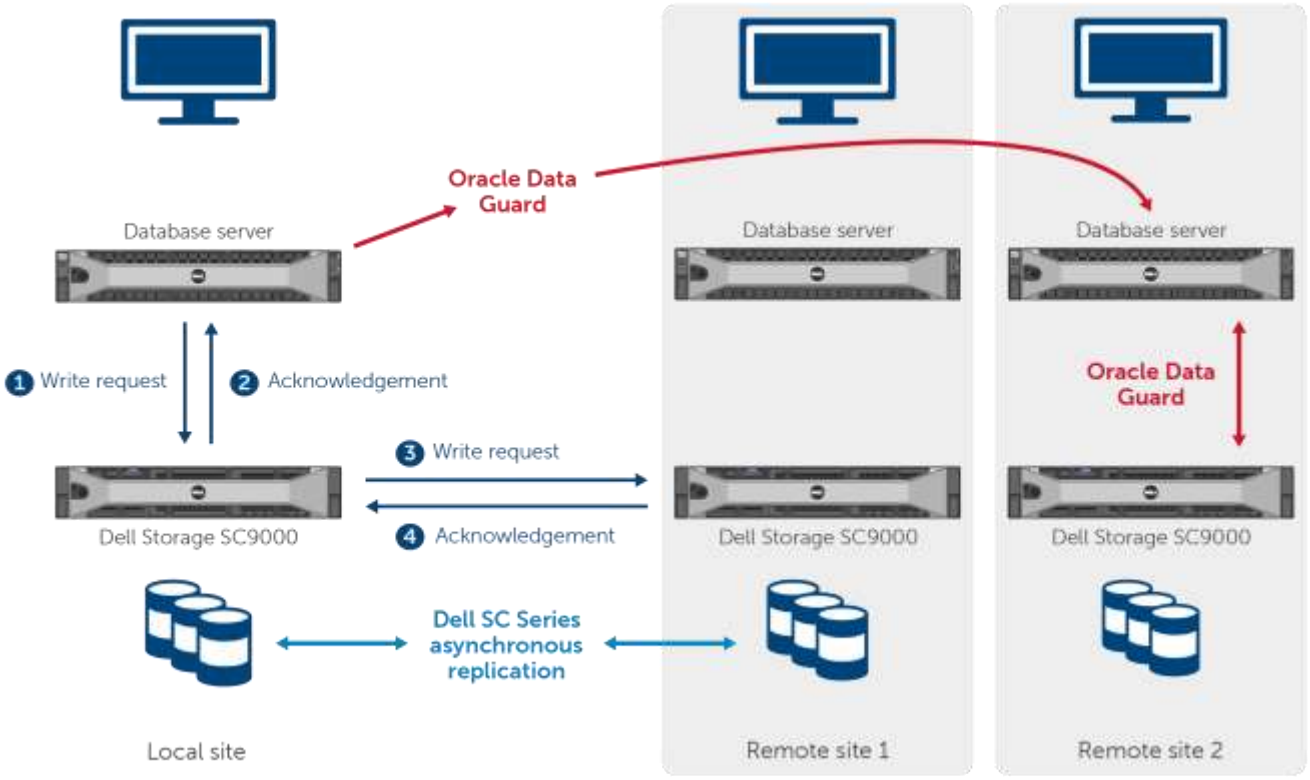


Figure 1 Example of Dell SC Series asynchronous replication and Oracle Data Guard in a hybrid solution

# 3 Dell Storage Center replication

Remote Instant Replay, also known as remote replication, replicates replays between local and remote SC Series arrays to safeguard data against local or regional data threats. If the source SC array becomes unavailable, the replicated volumes on the destination array can be activated to regain access to the data. Remote replication supports several disk storage replication methods over Fibre Channel and cost-effective IP0-based (iSCSI).

The following replication types are supported by SC Series storage:

- Synchronous
- Legacy synchronous
- Asynchronous

The main distinguishing factor between the replication methods is how writes are committed.

When replicating volumes, deduplication should be considered, especially if replication is performed over lower-bandwidth links. Deduplication reduces the amount of data transferred to the destination and can be advantageous when replicating volumes over slower links. If the transfer of the initial replication is too slow, portable volumes (a feature of SC Series) can be used to transport the initial replication data to a remote SC array.

Regarding data, it is replicated from the source volume to the lowest storage tier of the destination volume. If the replication contains an Oracle database, it may make sense to change this default behavior by modifying the settings for a replication and place the replication on a higher tier of storage for the replicated volume. Depending on how the remote database will be used, this change in storage tier may yield better performance metrics.

SC Series storage also supports replicating a single volume to multiple remote SC arrays. SC Series supports three topologies for multiple remote replications:

- **Mixed mode**: A source volume is replicated in parallel to multiple remote arrays.
- **Cascade mode**: A source volume is replicated in series to multiple remote arrays.
- **Hybrid mode**: This combines the mixed and cascade modes.

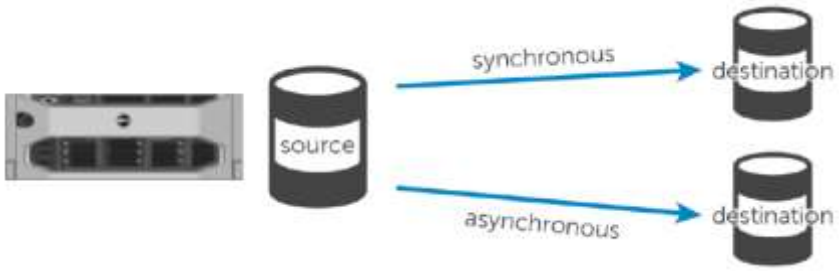


Figure 2 Mixed topology


Figure 3 Cascade topology

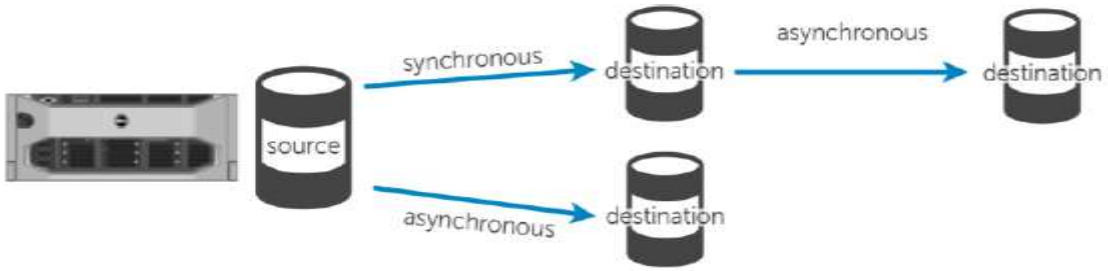


Figure 4 Hybrid mode

Different topologies provide a number of beneficial use cases:

- With mixed topologies, data can be protected in multiple remote or local locations, yielding multiple choices for recovery.
- A cascade topology can quickly provide copies of Oracle database for parallel testing and development activities.
- A hybrid topology combines the features of both mixed and cascade topologies and helps businesses address data protection needs.

It is recommended to simulate replications prior to releasing the replication definitions into any environment (production, test, development). Simulations help to ensure the replication has been configured, performs as expected, and provides estimated requirements for replication.

For further information on replication topology limitations, see the *Dell Enterprise Manager 2015 R1 Administrator's Guide* on Dell.com/support. Multiple topologies are not covered by this document.

## 3.1 Remote Instant Replay and modes of operation

Dell SC Series provides the ability to define and use local replays on a volume, but to take advantage of replicating the replay, a Remote Instant Replay license must be acquired for both the source and destination arrays. Replications can be created manually, or created according to a schedule that exists on the replay profile in Storage Center to which a volume belongs. Replay profiles can be defined as a normal or consistent.


Figure 5 Replay profile types in Dell Storage Center

With consistent replay profiles, the I/O to all volumes belonging to the profile will be suspended during the creation of the replay, thus guaranteeing a consistent point-in-time snapshot of the volumes. A replay created using this type of replay is known as a point-in-time replay. Therefore, if Oracle databases are to be replicated, unless all database files (such as control files, online redo logs, data, and index) reside on the same SC Series volume, all volumes that support the database must reside in the same consistent replay profile and replications created from it.

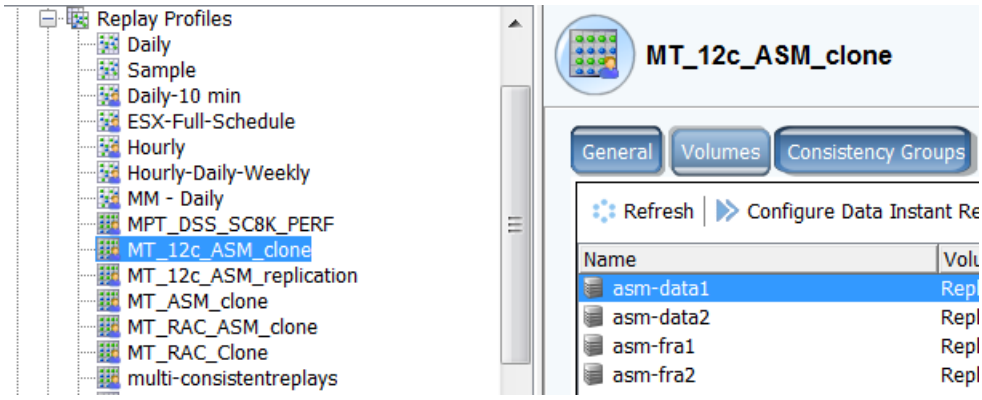


Figure 6 Consistent replay profile for an Oracle database

This guarantees that the volume images are consistent with a point in time, but does not guarantee that the database residing on the replay is consistent. In order to guarantee database consistency, the database must be placed in BEGIN BACKUP mode prior to creating the point-in-time replay. After the point-in-time replay is created, the database can be taken out of BEGIN BACKUP mode.

### 3.1.1 Synchronous replication

Zero data loss is guaranteed when using synchronous replication. With synchronous replication, a write either completes successfully on both local and remote sites, or not at all. An acknowledgement between the local and remote arrays is necessary to determine the outcome of the write. If an acknowledgement between the arrays occurs, then an acknowledgement is sent back to the application. Should an acknowledgement not be received, a failure will be issued to the local array where the application write request originated. At this point, the application can process the error for the requesting transaction. Because applications must wait for this acknowledgement when they are performing write operations, overall application performance decreases. The degree to which performance decreases will be dependent on a number of factors. One of the main factors is the distance between the local and remote

sites. For this reason, it is imperative that performance sizing of the entire infrastructure supporting the synchronous replication and application be performed prior to implementation.

With synchronous replication, the replay and active replay (current, unfrozen source volume data) are copied to the destination volume after they are frozen. There is no way to disable the active replay as there is with asynchronous replication.

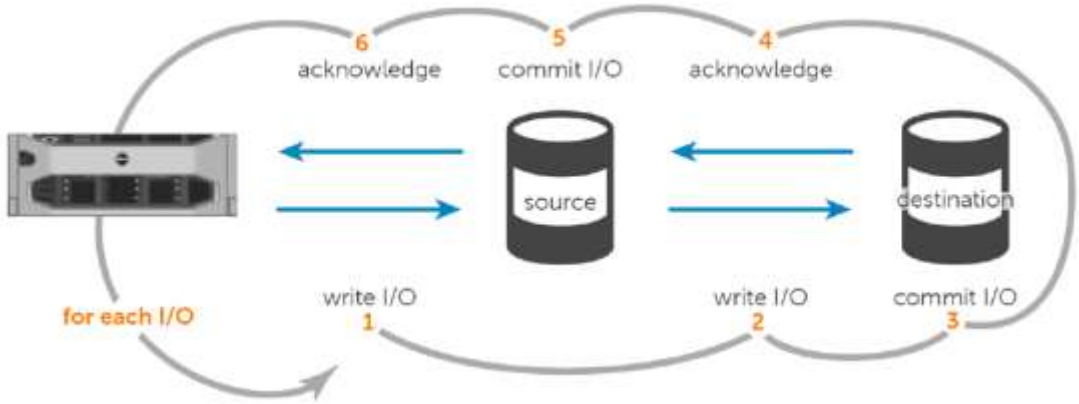


Figure 7 Synchronous replication write I/O sequence

There are two synchronous replication modes: high availability and high consistency. Both of these can have their replications paused.

#### 3.1.1.1 High availability

I/O requests to the source volume are accepted even when the destination volume is unavailable, or when a latency exceeds a predefined Storage Center threshold of 20 seconds. However, increased application latencies will be observed and the destination volume will become stale if it becomes unavailable, but data availability will not be sacrificed for data consistency. This stale state means that the destination volume will be out of date, or inconsistent with the source volume. When Storage Center detects a stale volume on the destination site, I/O requests will be journaled at the source volume. After the destination volume becomes available within a specific latency threshold, the journaled I/O residing at the source volume will be flushed to the destination volume and committed. Any inflight I/O from the application that occurs during this time will be journaled, After all journaled I/O has been processed at the destination, the source and destination will again be consistent and the stale state removed.

Synchronous high availability replication could migrate to asynchronous replication with **Replicate Active Replays** enabled (semi-synchronous). For this reason, synchronous high availability replication for Oracle databases is not recommended, because write order fidelity (WOF) is not guaranteed and may result in a corrupt database at the destination site. To mitigate this, the source database needs to be re-replicated.

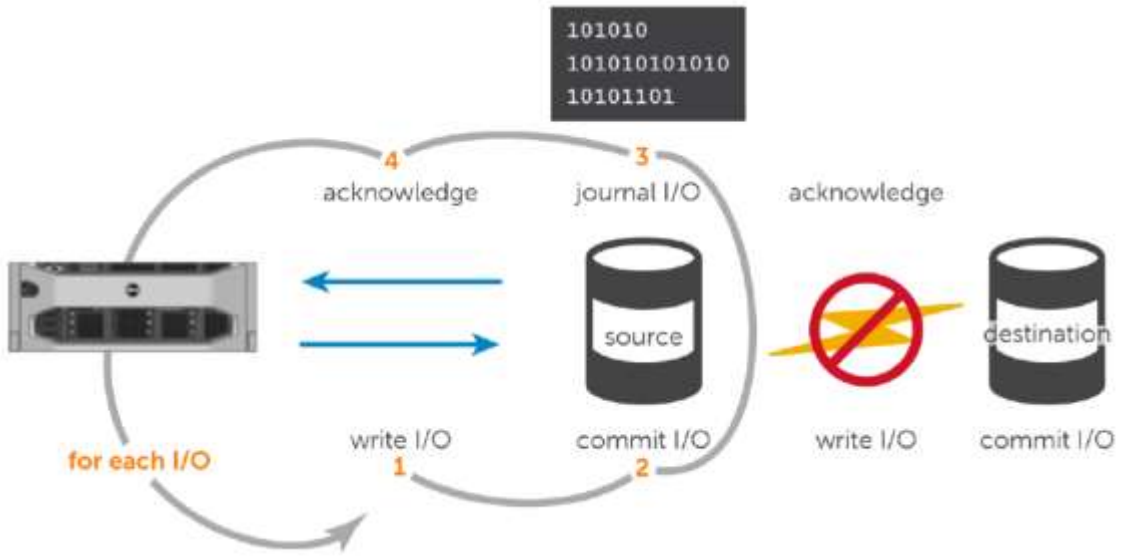


Figure 8 High availability synchronous replication in a stale state

### 3.1.1.2 High consistency

Should the destination volume become unavailable, or if the replication is suspended by the administrator, I/O requests to the source volume will not be permitted, which can result in an application interruption. Latencies can also impact the application should the infrastructure and remote destination volume be unable to absorb the data rate of change.

In either mode, when the destination volume becomes available, the transfer of replays and active replay data resumes.

### 3.1.1.3 Pausing replications

Several use cases can be addressed with pausing replications. One such use case is that it can be used to address high-bandwidth demands of the replication link. In some designs, QoS is shared and other processes can temporarily demand more of the link. Another use would be to pause a replication in anticipation of a planned outage or a different schedule replication.

## 3.1.2 Legacy synchronous replication

Like synchronous replication, legacy synchronous replication writes data to both the source and destination volumes simultaneously to make sure the data is synchronized. Legacy synchronous replication is available in SCOS versions prior to 6.3 and therefore not covered in this document.

## 3.1.3 Asynchronous replication

With asynchronous replication, an application write operation will be considered complete as soon as the local storage acknowledges the write, and the application must wait for that write acknowledgement before it can proceed. But since the write does not have to complete on the remote storage device before the application receives the acknowledgement, there is considerably less performance degradation with

asynchronous writes. Also, since the remote site does not necessarily have the latest data replica, should the local site sustain an outage, the remote storage is not guaranteed to have the current copy of the data.

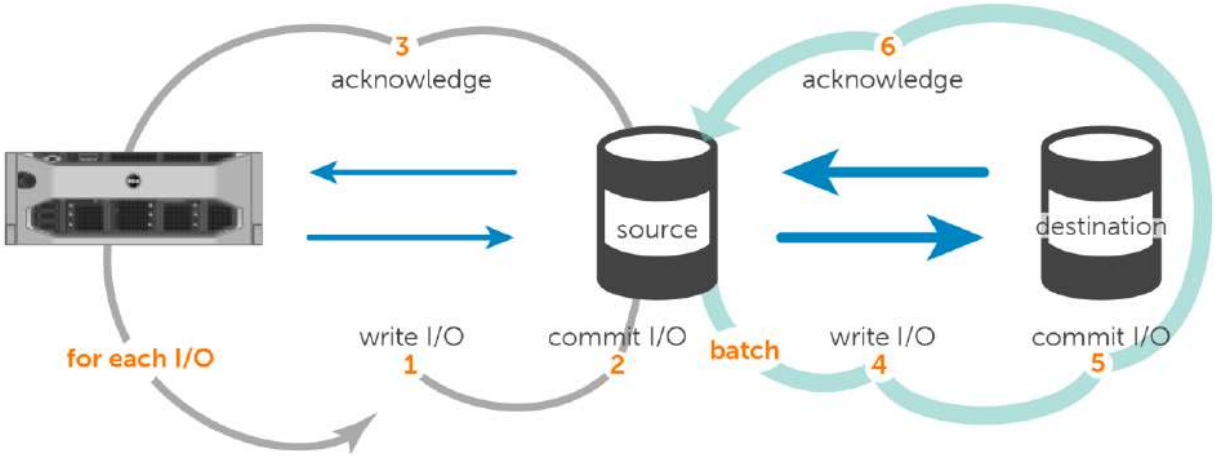


Figure 9 Asynchronous replication write I/O sequence

By default with asynchronous replication, a replay from the source volume is copied to the destination volume after it is frozen. A feature with asynchronous replication called **Replicate Active Replay** (semi-synchronous replication) can change this default behavior.

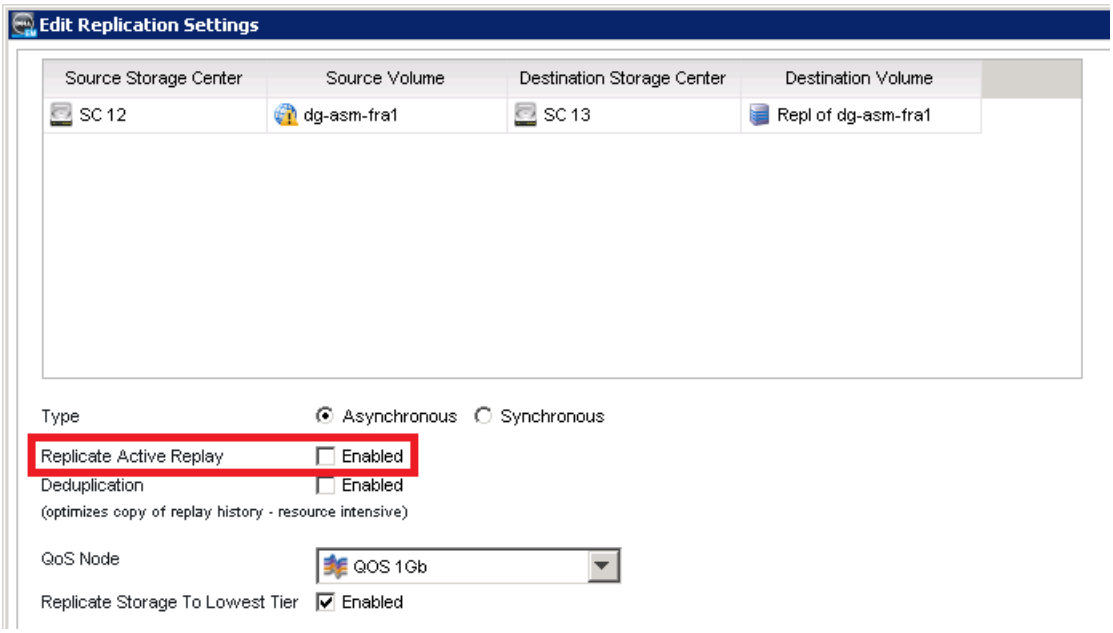


Figure 10 Replicate Active Replay feature in Dell Enterprise Manager

If **Replicate Active Replay** is disabled, only the replay is copied to the destination after it is frozen. There is no attempt to keep the active replay (current, unfrozen source volume data) in sync with the remote destination volume. If **Replicate Active Replay** is enabled and assuming the link (fiber, iSCSI, QoS) and

destination array have the bandwidth to support the current rate of change, the replay and active replay are replicated. Since **Replicate Active Replay** does not support write ordering (a concept better known as WOF), the writes at the remote site are not necessarily applied in the same order as they were applied at the local site. Therefore, asynchronous replication with **Replicate Active Replay** cannot guarantee that the volume at the destination be consistent with the source volume. The net result is that in a DR scenario, RPO and protection from data loss cannot be guaranteed. For this reason, semi-synchronous replication should not be used with Oracle databases because there is no guarantee that the destination database will be in a consistent recoverable state.

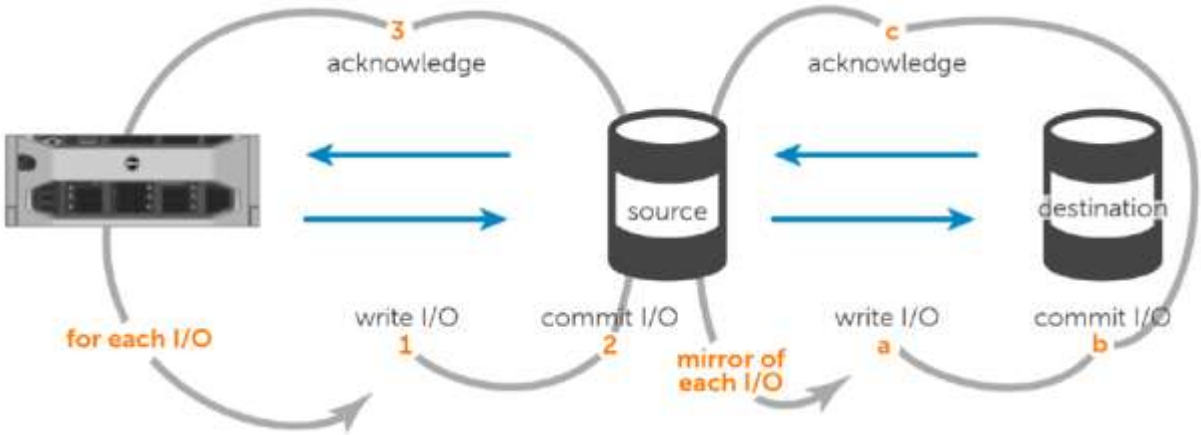


Figure 11 Semi-synchronous replication write I/O sequence

When using asynchronous replication, the default storage tier at the remote site will always be the lowest available storage tier defined for the destination volume and have slower performance metrics. So when using asynchronous replication with Oracle, consider changing the storage tier of the destination volume so the destination database will have the expected performance metrics.

To unmap the replication volume from the target array, first make sure the replication is 100 percent synced and completed, then pause or delete the replication. The delete will unmap the source volume from the target array. If the source volume is unmapped from the target array while the replication is active, it may interfere with the replication and make the target replication unusable.

## 3.2 Replication Quality of Service (QoS)

QoS must be defined for each replication to control how the bandwidth is used to send replicated data between the remote and local array. For mission-critical systems, a QoS with high transfer rates may be more appropriate then lower transfer rates, but this may not always be the case. For example, if a mission-critical system will undergo a planned outage and transactions have been throttled down for business and end-user reasons, a QoS with a lower transfer rate may be applicable.

To create a QoS node, change the view in Enterprise Manager to **Replications & Live Volumes**, select the **QoS Nodes** tab, and select **Create QoS Node**.

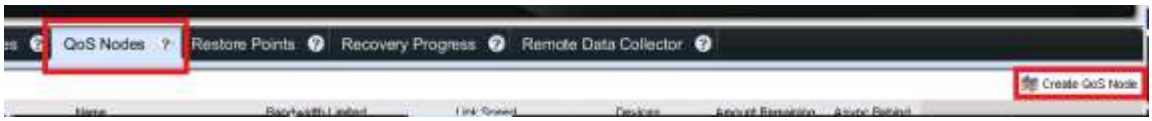


The **Create Replication QoS** wizard appears.

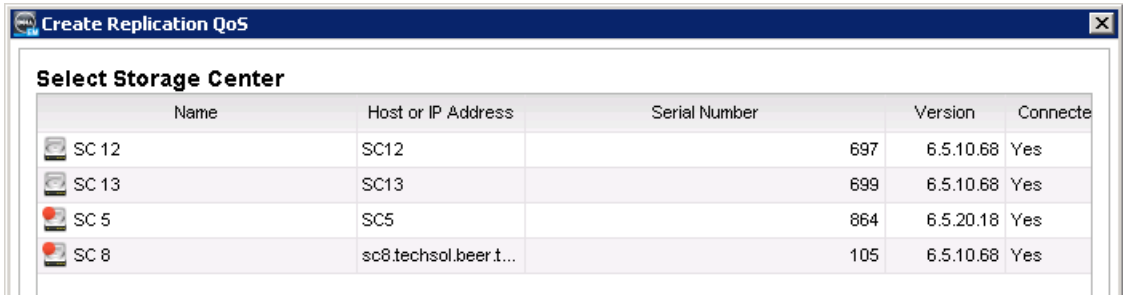


Select the source SC array of the replication for which a QoS node should be created, then select **Next**. Enterprise Manager will display the **Create Replication QoS** page.

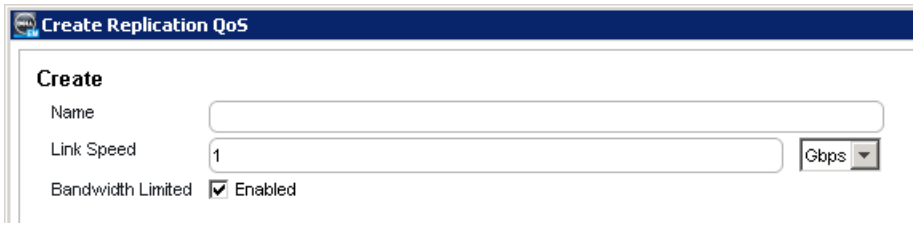


Enter the name and link speed in either **Gbps** or **Mbps**. Unless there are specific reasons to not limit the bandwidth, make sure **Bandwidth Limited** is **Enabled**. The bandwidth chosen will be determined by business requirements. When done, select **Finish**. Enterprise Manager will then display the schedule dialog. By default, the schedule will provide 100 percent of the bandwidth 24x7. Adjustments to this schedule may be necessary to ensure the available bandwidth is not over consumed by the replications during high transfer rates of other activity.

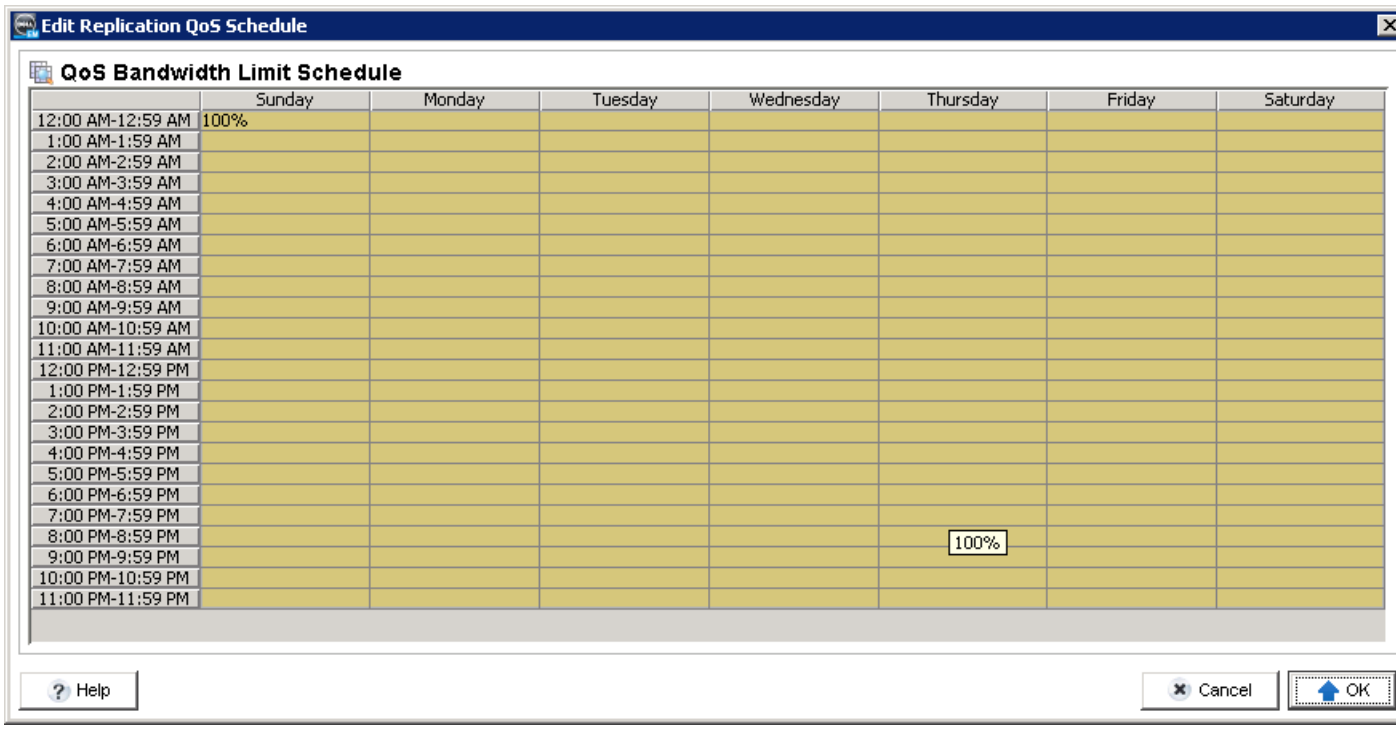


Care should be used if any time is marked as blocked because that will cause any synchronized replications to become unsynchronized during that time. Blocking is not recommended for synchronized replications of Oracle databases.

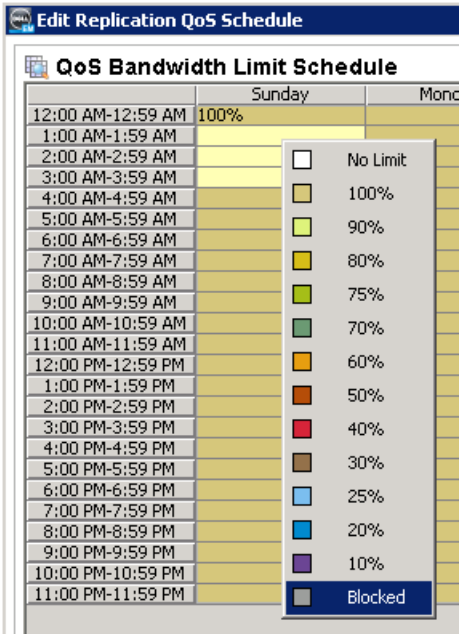

Any attribute of the QoS definition (link speeds, schedules, and bandwidth) can be changed at any time, and QoS definitions should be deleted if they are no longer used or there are no future plans to use them. However, to delete them, they must not be currently in use.

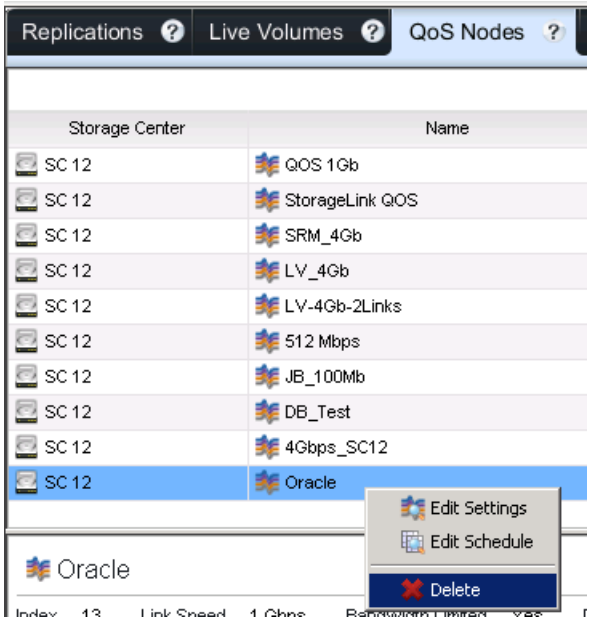

# 4 Dell Storage Center replication use cases

In general, a business continuity plan is only one of the reasons why replication technology is valuable to a business. Replication not only protects an organization's most valuable asset when a planned or unplanned outage occurs, but it also provides methods to resolve additional business requirements, such as:

- Creating image backups of SC Series LUNs at remote locations for backup and recovery
- Efficiently re-provisioning a database from one location to another for read-only production database reporting or analytical processing
- Re-provisioning copies of databases for development, testing, and training activity at remote locations

For many business needs, asynchronous replication provides a noble balance between meeting service level agreements (SLAs) of RPO and RTO, without the exorbitant or prohibitive costs of an infrastructure that includes dark fibre. For this reason, asynchronous replication can be used for most replication solutions, but it is not a replication panacea. For the use cases where business-critical applications require zero data loss, regardless of the application type and transactions, synchronous replication should be chosen. But synchronous replication comes with a high cost, and stakeholders and application owners may be apprehensive of the cost and elect for an asynchronous replication solution or a synchronous replication provided by the application. For Oracle, one possible synchronous solution can be derived from implementing Oracle Data Guard.

For the purposes of this document, Dell asynchronous replication was configured in a lab. Since Oracle Data Guard is a software solution and Dell Storage Center is application agnostic, Oracle Data Guard was not configured.

# 5 Oracle Data Guard

Oracle Data Guard is part of Enterprise and Personal editions of the Oracle RDBMS software, and it can provide high availability and disaster recovery. Data Guard provides methods for managing, creating, maintaining, and monitoring remote images of the local database to protect against disasters. If a planned or unplanned outage strikes the local database, Data Guard can switch to any one of the remote images and convert its role as the production role. Since Dell SC Series Storage is completely software agnostic, the Data Guard environments were not set up in a lab for the purposes of this document.

## 5.1 Available configurations

Every Data Guard configuration has one primary database, and at least one and up to nine standby databases. The primary database is the database that is being protected by Data Guard. A standby database is a synchronized or transactionally consistent copy of the primary database, and is initially created from a copy of the primary database. Once the standby database is created and Data Guard configured, it is automatically maintained by having the redo from the primary database applied to it.

### 5.1.1 Standby databases

Standby databases come in different types, all of which can exist in the same Data-Guard-configured solution, and they can reside on Dell SC Series storage. These database types include:

- Physical standby databases
- Logical standby databases
- Snapshot standby database

This document only discusses using physical standby databases with Dell SC Series storage. For additional information on the different types of standby databases and the benefits and features they provide, refer to Oracle documentation listed in appendix C.

### 5.1.2 Available services

Standby databases are kept synchronized with the primary database by a set of comprehensive services. These services provide the means to create, maintain, manage, and monitor the standby databases, and ensures that the primary database survives disasters and data corruptions. Available services are:

- Redo transport services
- Apply services
- Role transitions (switchover vs failover)

For information on these services, their installation, configuration and usage, refer to Oracle documentation.

## 5.2 Data Guard protection modes

Protection modes guarantee that a standby database is synchronized with the primary database at the minimum required level, and does not have any gaps in the redo logs that are applied to the standby database.

### 5.2.1 Maximum availability

This is the highest level of data protection that is possible without compromising the availability of the primary database. Transactions are allowed to commit as soon as all the redos they generate have been written to online redo logs and to at least one synchronized standby database

If a redo cannot be written to at least one standby database, the primary database operates as if it were in maximum performance mode to preserve the availability of the primary database. Maximum availability guarantees that no data loss will occur if the primary database fails, but only if the complete set of redos from the primary database is sent to at least the standby database.

### 5.2.2 Maximum performance

This is the default protection mode and it offers the highest level of data protection available that does not affect the performance of the primary database. Transactions are allowed to commit as soon as all the redos they generate have been written to online redo logs. The redos are also written asynchronously with respect to the transaction commitment to at least one standby database. This asynchronous write allows the primary database to be unaffected by delays in writing any of the redos to the standby database.

### 5.2.3 Maximum protection

Should the primary database fail, zero data loss is only guaranteed with maximum protection. Oracle provides the protection by writing the redo data to the online log as well as to at least one synchronized standby database before the transactions are allowed to commit on the primary database. To guarantee zero data loss, should the primary database experience issue in writing its redo to the standby database, the primary database will shut down.

## 5.3 Complementary Oracle technologies

Oracle offers other technologies that complement Data Guard to keep mission-critical systems running with greater levels of data protection and availability. When combined, the following technologies create an Oracle maximum availability architecture (MAA):

- Oracle Real Application Clusters (RAC)
- Flashback database
- Recovery Manager (RMAN)

For more information on these technologies, see Oracle documentation.

# 6 Infrastructure testing

This section reviews the architecture, configuration, and diagrams depicting the infrastructure used in lab testing to support the content of this paper. This section should be used as an overview of possible configurations but not as an explicit set of instructions to ensure a successful implementation. For a complete list of implementation requirements, refer to the documentation for all components within the target environment, such as the following:

- Dell Storage Center administration guide
- Dell Storage Linux best practices
- Oracle Linux installation documentation
- Oracle documentation
- Brocade switch documentation
- Cisco switch documentation

More information for supporting documentation can be found in appendix C.

## 6.1 Hardware components

The following hardware components were used in support of this document:

- Dell PowerEdge™ R730xd
- Dell PowerEdge R710
- Dell PowerEdge R620 (Dell Enterprise Manager)
- Dell Storage SC4020
- Cisco® Ethernet switch
- Brocade® Fabre channel switch

## 6.2 Software components

The following software was used in support of this document:

- Oracle Unbreakable Enterprise Linux Kernel 6.5
- Oracle 11g (11.2.0.4) Grid Clusterware
- Oracle 11g (11.2.0.4) RDBMS single instance standalone
- Dell SC4020 Storage Center OS (SCOS) 6.5.10
- Dell Enterprise Manager 2015 R1, build 15.3.1.110
- Microsoft® Windows Server® 2008 R2 Enterprise SP1, 64-bit

## 6.3 Testing performed

After Oracle Grid Infrastructure and database 11.2.0.4 were installed on the local and remote servers, a database with two LUNs was then created and added to a consistent replay profile.

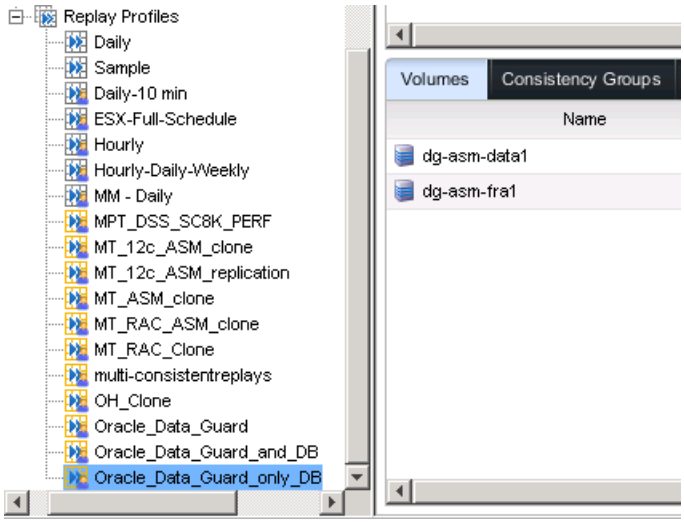


Figure 12 Consistent replay profile for replication

Asynchronous replications were defined for each of these LUNs.

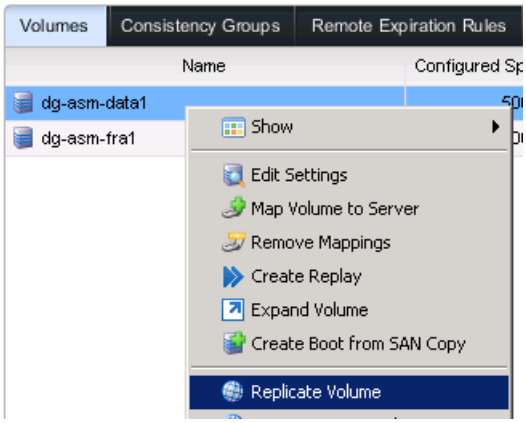


Figure 13 Defining volume replication in Dell Enterprise Manager

A destination SC array was selected.

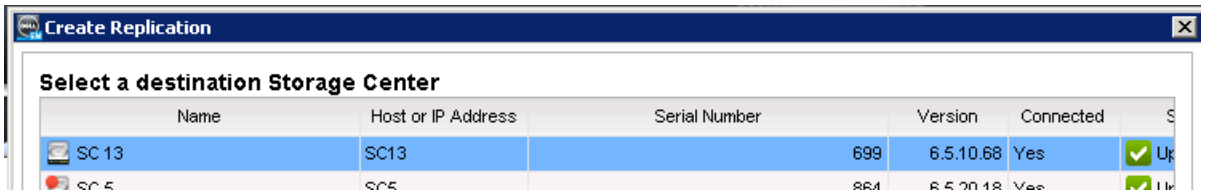


Figure 14 Selecting the remote SC for replications

Then, the create replication wizard was displayed.

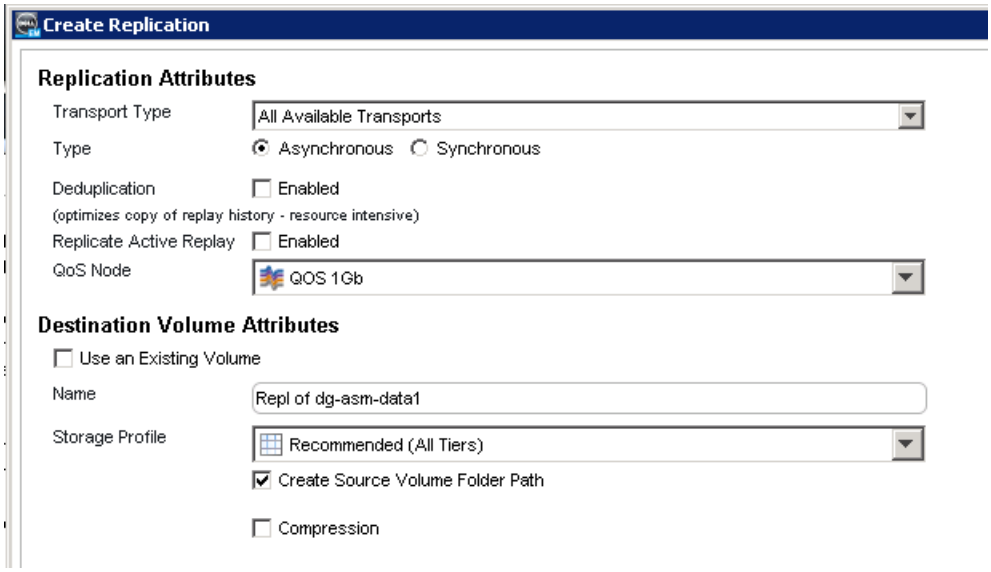


Figure 15 Dell Enterprise Manager wizard for creating replications

The appropriate transport protocol type and deduplication can be selected if needed. The type was set to Asynchronous and QoS Node to the desired QoS bandwidth definition. **Replicate Active Replay** was not selected because WOF is not guaranteed and may result in a corrupt database at the destination site. Then the appropriate destination volume attributes were selected, followed by **Finish**.

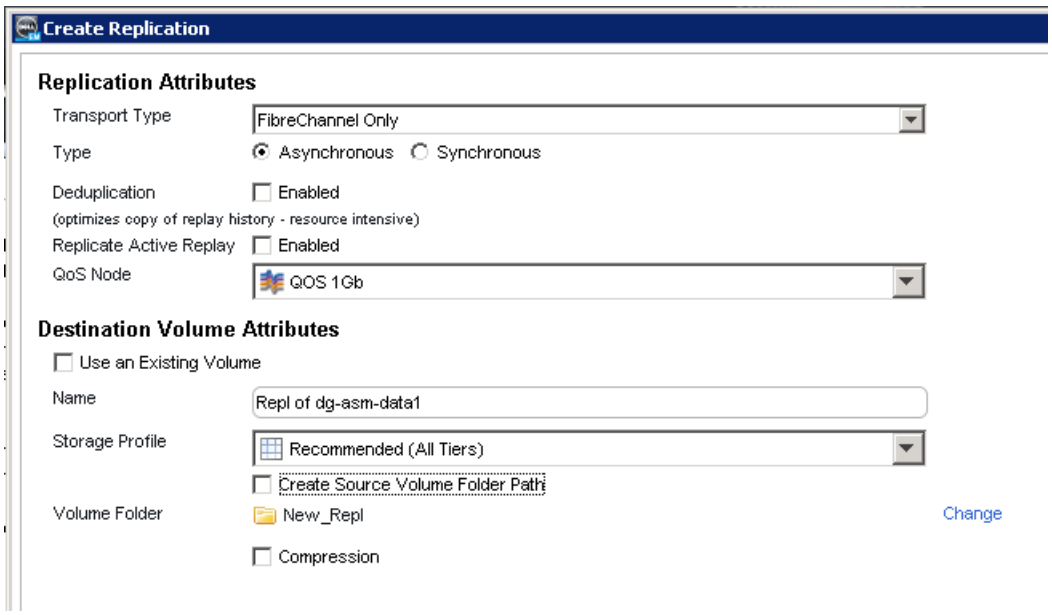


Figure 16 Asynchronous replication options

Enterprise Manager displayed the volumes in the consistency group.

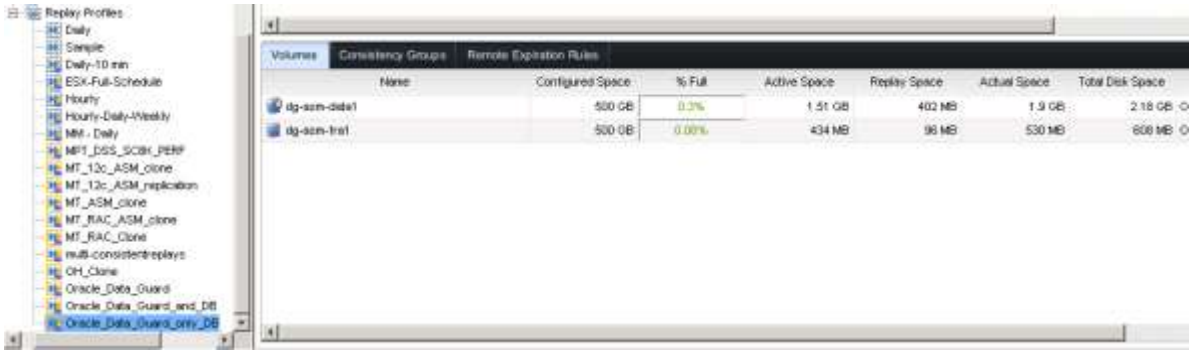


Figure 17 Database volumes in a consistent replay profile

Then, the process was repeated to create a replication definition for all remaining volumes in the consistent replay profile.

After replication definitions were created for all the volumes in the consistency group, the consistency group was right-clicked and **Create Replay** was selected.

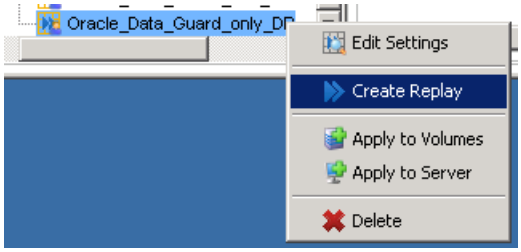


Figure 18 Create Replay selected to create the replication

The Create Replay wizard prompted for expiration parameters, the appropriate values were provided, and **OK** was selected.

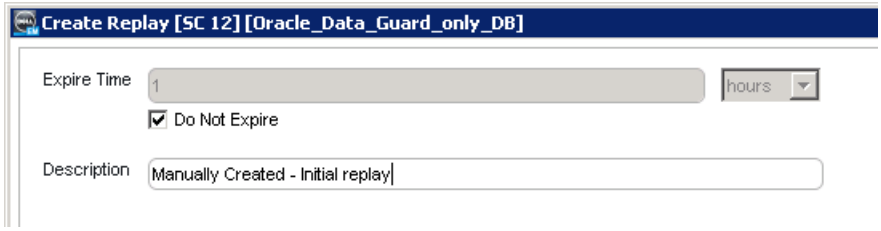


Figure 19 Expiration and description of the replays

A refresh of the destination SC array displayed the newly replicated volumes.

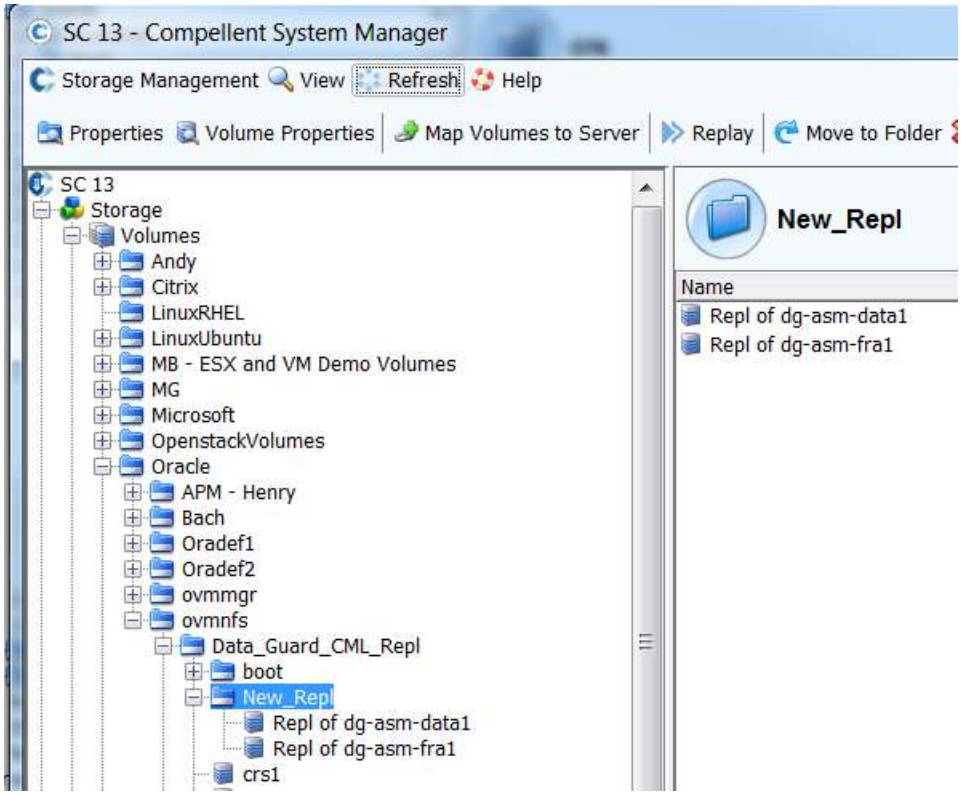


Figure 20 Replicated volumes

The volumes were then replicated to the destination site.

In Enterprise Manager, the display was changed to view **Replications & Live Volumes.**

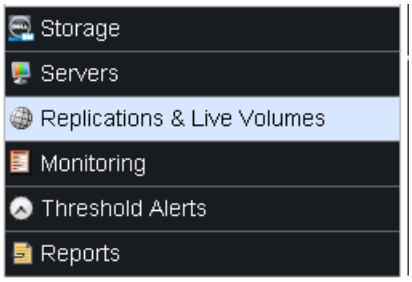


Figure 21 Replication view in Enterprise Manager

The status of the replication was then verified. Before the replicated volumes were used, the replication had to be 100% complete and synced.

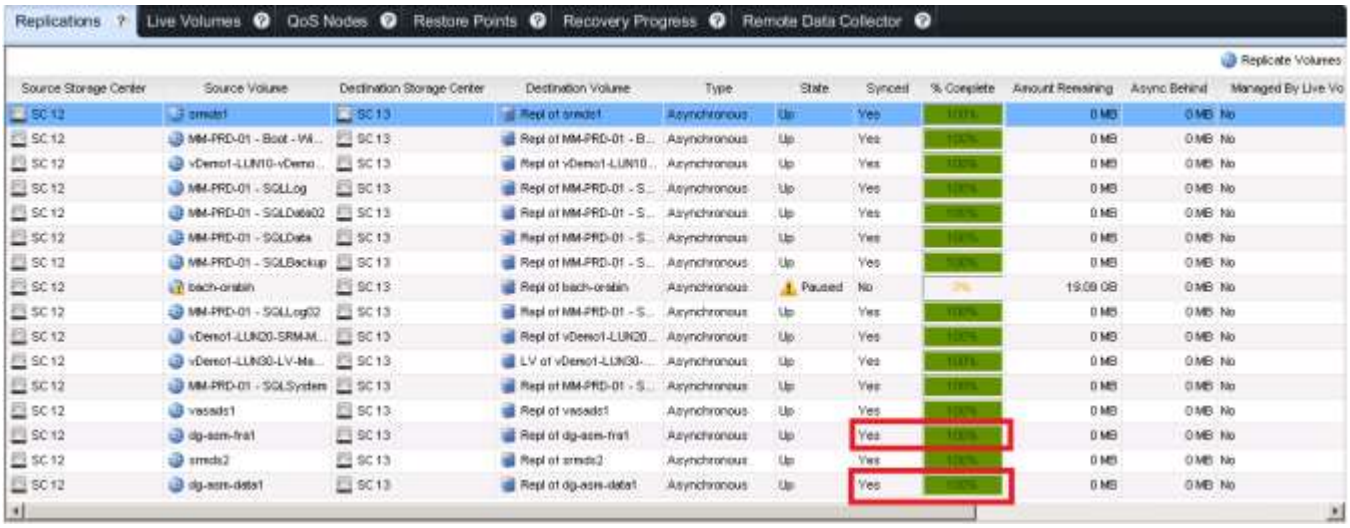


Figure 22 Replication status

Once the replication was synced at 100%, and before the replicated image was used in Oracle at the remote site, the replication definition needed to be broken to ensure that the destination would only receive I/O from Oracle at the remote site, and not from the source or local SC array in the replication.

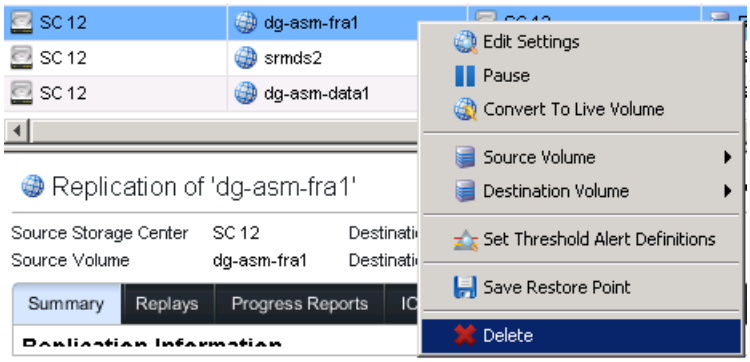


Figure 23 Deleting replications

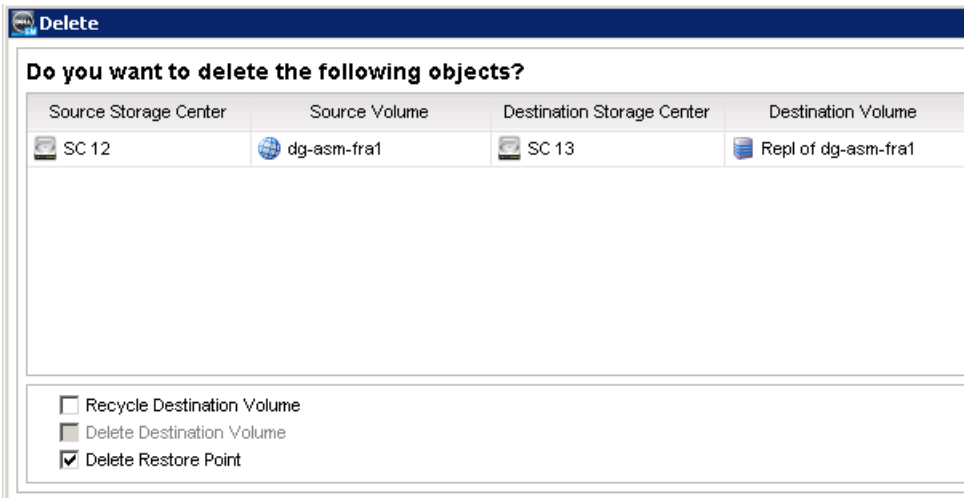


Figure 24 Confirmation to delete the replication definition for a volume

Then, **OK** was selected and the delete process was repeated for the remaining volumes in the consistency group. Once the replication definitions were deleted for the volumes existing in the consistency group, the destination volumes were mapped to a server, presented to Oracle, and then the database was started. Oracle performed crash recovery and placed the database in the consistent state recorded in the replicated replays.

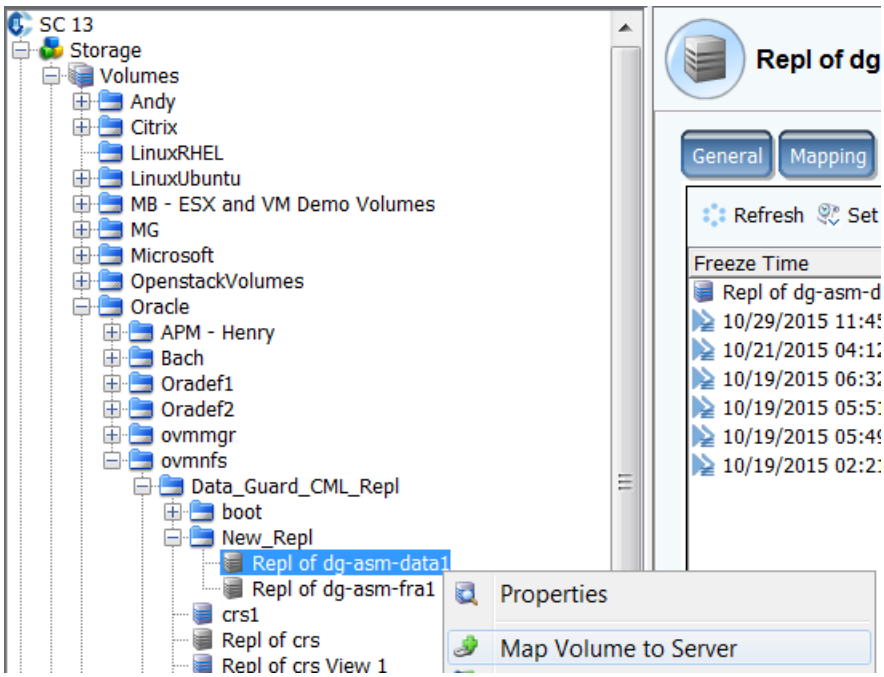


Figure 25 Replicated volumes mapped to a destination server

# 7 Conclusion

This paper illustrates how Dell SC Series storage replication, a hardware replication solution, can provide a quick and effective way to protect a site from local disasters and meet the RTO and RPO SLAs of business applications and databases, or can be used as a provisioning tool for cloning an Oracle database to remote sites. It also covers use cases where integrating SC Series replication with Oracle Data Guard, a software replication solution, provides robust recovery solutions. Finally, the paper presents recommendations and best practices when configuring Dell SC Series storage asynchronous and synchronous replication modes in an Oracle environment.

In summary, implementing Dell SC Series replication can deliver the following benefits:

- Reduced licensing costs
- DR for standard edition Oracle databases
- Database recovery
- Database cloning

# A Configuration details

## A.1 Dell Enterprise Manager and Data Collector configurations

For Enterprise Manager requirements, see the *Dell Enterprise Manager Release Notes* (license required).

Table 1 Enterprise Manager, Data Collector, and application server details

| Component | Description |
|---|---|
| Enterprise Manager version | 2015 R1, build 15.3.1.110 |
| Server model | Dell PowerEdge R620 |
| System revision | 1 |
| Processors | Intel® Xeon® E5-2670 v2 @ 2.50 GHz, 2 processors, 10 cores each |
| Memory | 129 GB RAM (8 x 16 GB DDR-3, 1600 MHz DIMMs) |
| OS | Microsoft Windows Server 2008 R2 Enterprise SP1, 64-bit |
| BIOS version | 2.4.3 |
| Firmware version (iDRAC) | 1.66.65 (Build 07) |
| iDRAC 7 NIC | Dedicated |
| Lifecycle Controller firmware | 1.4.2.12 |
| IDSDM firmware | N/A |
| HBAs | QLogic® QLE2562 8 Gb 2-port FC adapter |
| HBA firmware | 3.21.04 |
| Boot | SAN boot |
| Integrated NIC 1 | Intel 2P X540/2P I350 rNDC |
| IPv6 enabled | No |
| IPv4 enabled | Yes |
| .NET Framework | 4.0.30319 full |
| Mozilla® Firefox® version | 36.0 |
| External database | MySQL version 14.14, distribution 5.5.41 for Win64 (x86) |
| MySQL database storage | 500 GB Dell SC Series volume |
| Microsoft Windows Installer | 5.0.7601.17514 |

# B Linux kernel settings

```
net.ipv4.ip_forward = 0
net.ipv4.conf.default.rp_filter = 1
net.ipv4.conf.default.accept_source_route = 0
kernel.sysrq = 0
kernel.core_uses_pid = 1
net.ipv4.tcp_syncookies = 1
net.bridge.bridge-nf-call-ip6tables = 0
net.bridge.bridge-nf-call-iptables = 0
net.bridge.bridge-nf-call-arptables = 0
kernel.msgmnb = 65536
kernel.msgmax = 65536
fs.aio-max-nr = 1048576
fs.file-max = 6815744
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
kernel.shmall = 32989426
kernel.shmmax = 135124690944
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
kernel.shmall = 1073741824
```

# C Additional resources

For Copilot support of Dell SC Series products:

- Global online support
- Email: support@compellent.com (non-emergency business hours)
- Phone: 866-EZ-STORE (866-397-8673) (United States only)

The Dell SC Series Customer Portal is an online portal for existing customers. A valid portal account is required. Once logged in, click Knowledge center.

Dell TechCenter is an online technical community for IT professionals and is a great resource to discover and learn about a wide range of technologies such as storage, servers, networking, software, and cloud management.

Dell.com/support is focused on meeting your needs with proven services and support.

Referenced or recommended Dell publications:

- *Dell Storage Center with Red Hat Enterprise Linux (RHEL) 6x Best Practices:* http://en.community.dell.com/techcenter/extras/m/white_papers/20437964
- *Dell Storage Center and Oracle Best Practices* http://en.community.dell.com/techcenter/extras/m/white_papers/20438041
- *How to install Oracle 11gR2 with Red Hat Enterprise Linux/Oracle Linux 6*: http://en.community.dell.com/techcenter/enterprise-solutions/m/oracle_db_gallery/20212999.aspx
- *Dell Enterprise Manager 2015 R1 Release Notes* (license required)
- *Dell Storage Center Dell Storage Client 2015 R1 Administrator's Guide* (license required)
- *Dell Enterprise Manager 2015 R1 Administrator's* Guide (license required)
- *Dell Compellent Storage Center, Flash Optimization, Upgrade Guide* (license required)
- *Dell Compellent – Oracle DR with Remote Replication*: http://en.community.dell.com/techcenter/extras/m/white_papers/20438046
- *How to configure Asynchronous Replication between Compellent Storage Centers:* http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/4874.how-do-i-configure-asynchronous-replication-between-compellent-storage-centers
- *How to configure Synchronous Replication between Compellent Storage Centers:* http://en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/4876.how-to-configure-synchronous-replication-between-compellent-storage-centers

Referenced or recommended Oracle publications (some may require Oracle Support license):

- *Oracle Database 11gR2 Database Administration Database Performance Tuning Guide:* https://docs.oracle.com/cd/E11882_01/server.112/e41573/iodesign.htm#PFGRF015
- *Oracle Database Storage Administrators Guide* http://download.oracle.com/docs/cd/B28359_01/server.111/b31107.pdf

- *Oracle 11g Database Concepts:* http://download.oracle.com/docs/cd/B28359_01/server.111/b28318/toc.htm
- *Oracle Database Release Notes 11gR2 for Linux:* https://docs.oracle.com/cd/E11882_01/relnotes.112/e23558/toc.htm
- *Oracle Database Installation Guide 11gR2* https://docs.oracle.com/cd/E11882_01/install.112/e47689/toc.htm
- *Oracle Grid Infrastructure Installation Guide 11gR2* https://docs.oracle.com/cd/E11882_01/install.112/e41961/toc.htm