

Combining IBM Storwize V7000 IP Replication and Oracle Data Guard

Reference Guide for Database and Storage Administrators

Shashank Shingornikar



Storage



Overview

This IBM® Redpaper™ publication combines Oracle Data Guard and IBM remote copy technologies and discusses the use of the near disaster and far disaster concepts that are best suited for customer environments.

The Oracle Data Guard technology is native to the Oracle database and does not depend on the server and storage replication. It requires less network bandwidth to transfer the data from a primary site to a standby site, because it transports only log files. But it is a more complex implementation when compared to storage replication and must be implemented for each database independently.

The IBM remote copy feature is a flexible data-mirroring technology that allows replication between volumes on two more disk storage systems. It does not use any server resources. Failover and failback processes are simple with the storage replication technology. It supports disaster recovery capabilities not only for the Oracle database, but for all other applications and databases within the enterprise. One disaster recovery solution can take care of all other applications, providing benefit to the customers.

The intended audience for this paper is any technical lead, system administrator, storage administrator, or Oracle database administrator in a production environment who is experienced with disaster recovery concepts and procedures. After reading this paper, the technical staff can understand how to perform failover and failback procedures on Oracle database with Automatic Storage Management (ASM) using the Data Guard methodology and IBM Metro Mirror and Global Mirror (copy services) technologies. Both technologies are explained in detail in this paper.

Scope of this paper

Although IBM Storwize® V7000 supports both Fibre Channel (FC) and iSCSI protocols, in this paper, the discussion is limited to storage exposed over iSCSI.

This paper provides:

- ▶ Guidelines for implementing Data Guard
- ▶ Guidelines for implementing Metro Mirror and Global Mirror
- ▶ Techniques used for switching the primary site along with Data Guard

This technical paper does not:

- ▶ Discuss any performance impact and analysis from a user perspective.
- ▶ Replace any official manuals and documents from IBM on Storwize V7000.
- ▶ Replace any official manual or support document provided by Oracle.

Introduction to the IBM Storwize V7000 system

The IBM Storwize V7000 system combines hardware and software to control the mapping of storage into volumes in a storage area network (SAN) environment. The system provides many benefits to storage administrators, including simplified storage administration, integrated management of IBM servers and storage, and enterprise-class performance, function, and reliability.

The IBM Storwize V7000 system includes rack-mounted units, called *enclosures*. Each enclosure can be a 12- or 24-drive model and has two canisters and two power supplies located at the back. The system includes the following types of enclosures:

- ▶ Control
- ▶ Expansion

A system can support more than one control enclosure, and a single control enclosure can have several expansion units attached to it.

Figure 1 illustrates the IBM Storwize V7000 system.

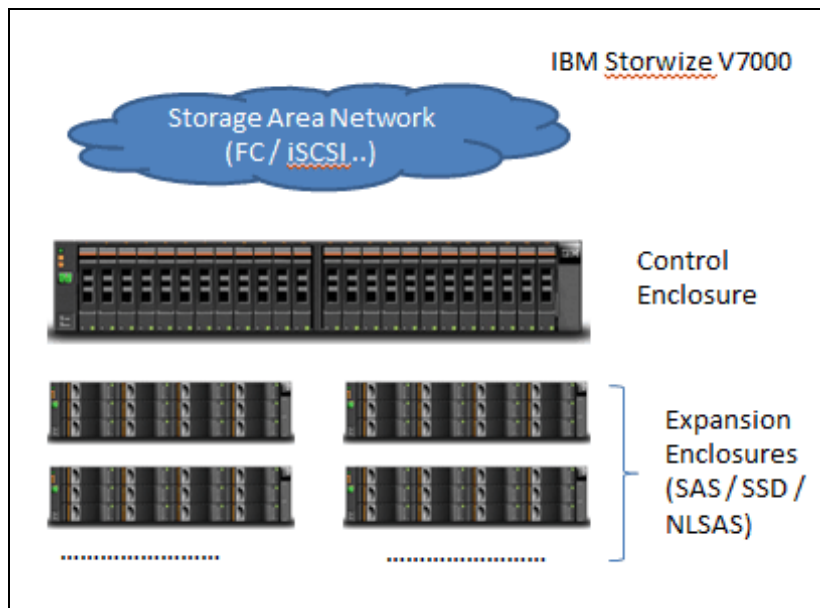


Figure 1 IBM Storwize V7000

Node canisters are always installed in pairs as part of a control enclosure. Each control enclosure represents an I/O group. Any expansion enclosure that is attached to a specific control enclosure also belongs to the same I/O group.

The system also includes an easy-to-use management graphical user interface (GUI), which helps you to configure, troubleshoot, and manage the system.

For more information about the IBM Storwize V7000 system, refer to:

<http://www.ibm.com/support/knowledgecenter/ST3FR7/welcome>

Understanding RTO and RPO

IT systems have become increasingly critical to the smooth operation of companies. Thus, the importance of ensuring their continued operation, including rapid recovery when systems fail, has increased. Before selecting a disaster recovery strategy, a disaster recovery planner must refer to the organization's *business continuity plan*. This plan indicates the key metrics of recovery point objective (RPO) and recovery time objective (RTO) for various business processes.

The metrics specified for the business processes must then be mapped to the underlying IT systems and infrastructure that support those processes. After mapping the RTO and RPO metrics to IT infrastructure, the disaster recovery planner can determine the most suitable recovery strategy for each system.

- ▶ RPO refers to the point in time to which data must be recovered as defined by the organization.

This is generally a definition of what an organization determines as an acceptable loss in a disaster situation.

- ▶ RTO refers to the duration of time and a service level within which a business process must be restored, following a disaster or disruption, in order to avoid unacceptable consequences associated with a break in business continuity.

The ideal solution will have both: a low RPO (in minutes) and RTO (ranges from minutes to hours). It is important to test a disaster recovery solution to find whether it is suitable and efficient for business continuity.

IBM Storwize V7000 copy services

Remote copy services are used to maintain two copies of data separated by distance. The remote copy can be maintained in one of the following modes:

- ▶ Synchronous

Synchronous remote copy ensures that updates are committed at both the primary and the secondary sites before the application considers the updates to be complete; therefore, the secondary site is fully up-to-date if it is needed in a failover. However, the application is fully exposed to the latency and bandwidth limitations of the communication link to the secondary site.

- ▶ Asynchronous

In asynchronous remote copy, the application is provided an acknowledgment that the write is complete before the write is committed at the secondary site. Thus, on a failover, certain updates (data) might be missing at the secondary site.

The following copy services features are available for all supported hosts that are connected to Storwize V7000:

- ▶ *IBM FlashCopy®* makes an instant, point-in-time copy from a source volume to a target volume. The FlashCopy feature copies data instantaneously from a source volume to a target volume. This copy is taken at a particular point in time as hosts continue to access the data. You must create a mapping between the source volume and the target volume. A FlashCopy mapping can be created between any two volumes of the same size in a clustered system. It is also possible for a FlashCopy target volume to be included as a source volume in a Metro Mirror or Global Mirror relationship.

Figure 2 illustrates the FlashCopy feature.




Figure 2 The IBM FlashCopy feature

- *Metro Mirror* provides a consistent copy of a source volume on a target volume. Data is written to the target volume synchronously after it is written to the source volume so that the copy is continuously updated. Metro Mirror is a copy service that provides a continuous, synchronous mirror of one volume to a second volume. The different systems can be up to 300 kilometers apart. So by using Metro Mirror you can make a copy to a location offsite or across town. Because the mirror is updated in real time, no data is lost if a failure occurs. Metro Mirror is generally used for disaster recovery purposes, where it is important to avoid data loss.

Figure 3 illustrates the Metro Mirror relationship.

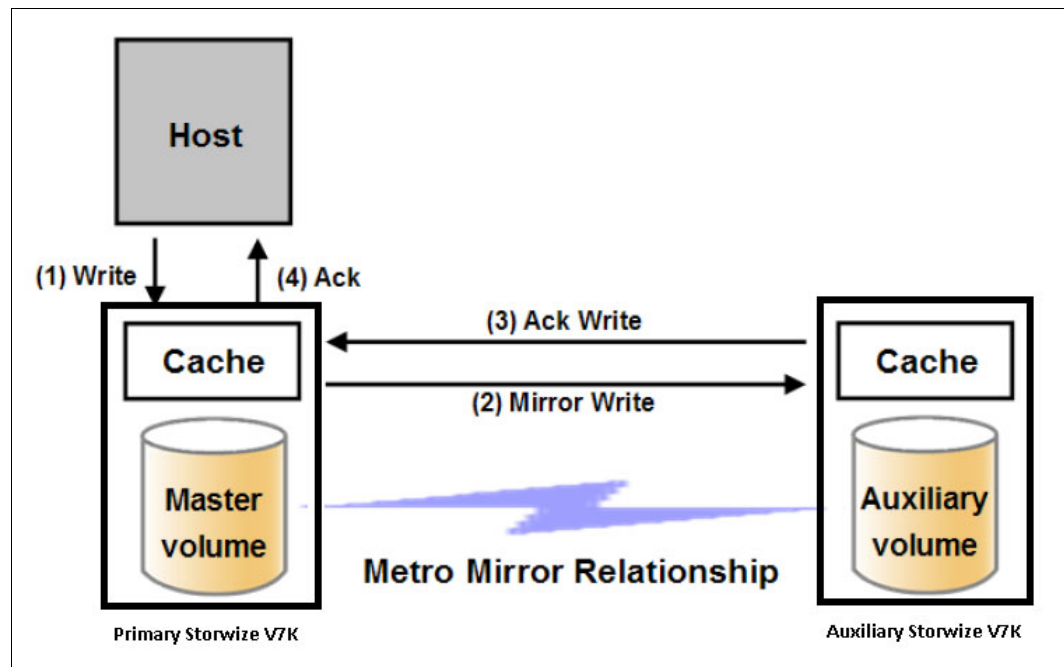


Figure 3 Metro Mirror relationship

- *Global Mirror* and *Global Mirror Change Volumes* provide a consistent copy of a source volume on a target volume. Data is written to the target volume asynchronously so that the copy is continuously updated. However, the copy might not contain the most recent updates in the event that a disaster recovery operation is performed.

Global Mirror is a copy service that is similar to Metro Mirror. Both provide a continuous mirror of one volume to a second volume. But with Global Mirror, the copy is asynchronous. You do not have to wait for the write to the secondary system to complete. For long distances, performance is improved compared to Metro Mirror. However, if a failure occurs, you might lose data. Global Mirror uses one of the two methods to replicate data:

- *Multicycling Global Mirror* is designed to replicate data while adjusting for bandwidth constraints and is appropriate for environments where it is acceptable to lose a few minutes of data if a failure occurs.
- For environments with higher bandwidth, *Noncycling Global Mirror* can be used so that less than a second of data is lost if a failure occurs.

Global Mirror works well for data protection and migration when recovery sites are more than 300 kilometers away.

Figure 4 illustrates the Global Mirror relationship.

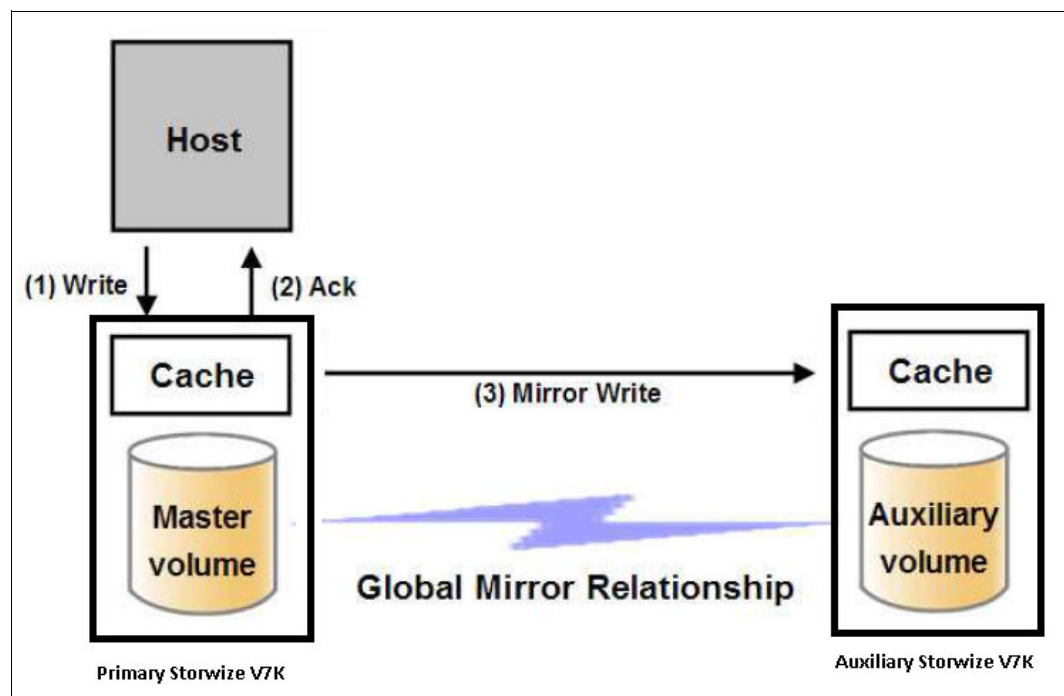


Figure 4 Global Mirror relationship

For this paper, the discussion is limited to Metro Mirror and Global Mirror features only.

Metro Mirror and Global Mirror are two types of remote-copy operations that enable you to set up a relationship between two volumes, where updates made to one volume are mirrored on the other volume. The volumes can either be on the same system (*intrasystem*) or on two different systems (*intersystem*). To use Metro Mirror and Global Mirror functions, you must have the remote copy license installed on each enclosure that you plan to use these functions.

To facilitate remote copy, local and remote Storwize V7000 systems discover each other's remote copy capable IP addresses and then set up IP sessions across those IP addresses to establish a partnership over native IP links.

When creating and configuring IP partnerships, consider the following requirements and constraints:

- ▶ Transmission Control Protocol (TCP) ports 3260 and 3265 are used for IP partnership communications, Therefore, these ports need to be opened in firewalls between the systems.
- ▶ Virtual LAN (VLAN) tagging of the IP addresses configured for remote copy is currently not supported.
- ▶ IP partnerships between the two systems can be either over IPv4 or IPv6 only and not both.
- ▶ iSCSI host access can be configured on node port configured for IP partnership, but this configuration is not recommended.

Note: For detailed information about requirements and constraints, refer to:

http://www.ibm.com/support/knowledgecenter/ST3FR7_7.3.0/com.ibm.storwize.v7000.730.doc/svc_ippartnershipreqs.html?cp=ST3FR7%2F0-0-2-4-2-4&lang=en

Consistency groups

A *consistency group* is a container for FlashCopy mappings, Global Mirror relationships, and Metro Mirror relationships. You can add many mappings or relationships to a consistency group. However, FlashCopy mappings, Global Mirror relationships, and Metro Mirror relationships cannot appear in the same consistency group.

When you use a consistency group, you can perform copy operations on the entire group instead of the individual mappings. Metro Mirror or Global Mirror relationships can only belong to one consistency group; however, they do not have to belong to a consistency group.

Relationships that are not part of a consistency group are called *stand-alone relationships*. A consistency group can contain zero or more relationships. All relationships in a consistency group must have matching primary and secondary systems, which are sometimes referred to as *master* and *auxiliary* systems. All relationships in a consistency group must also have the same copy direction and state.

The system supports the following types of relationships and consistency groups:

- ▶ Metro Mirror
- ▶ Global Mirror without change volumes (cycling mode set to *None*)
- ▶ Global Mirror with change volumes (cycling mode set to *Multiple*)

Oracle Data Guard

Oracle Data Guard provides a solution for high availability, enhanced performance, and automated failover. You can use Oracle Data Guard to create and maintain multiple standby databases for a primary database. The standby databases can be started in the read-only mode to support reporting users and then returned to the standby mode. Changes to the primary database can be relayed automatically from the primary database to the standby

databases with a guarantee of no data lost in the process. The standby database servers can be physically separate from the primary server.

In a Data Guard implementation, a database running in *archive* mode is designated as the primary database for an application. One or more standby databases, accessible through Oracle Net Services, provide for failover capabilities. Data Guard automatically transmits redo information to the standby databases, where it is applied. As a result, the standby database is transitionally consistent.

Depending on how you configure the redo application process, the standby databases might be in sync with the primary database or might lag behind it. The redo log data is transferred to the standby databases through log transport services, as defined through your initialization parameter settings. Log Apply Services apply the redo information to the standby databases.

Figure 5 shows a multi-standby database Data Guard implementation.

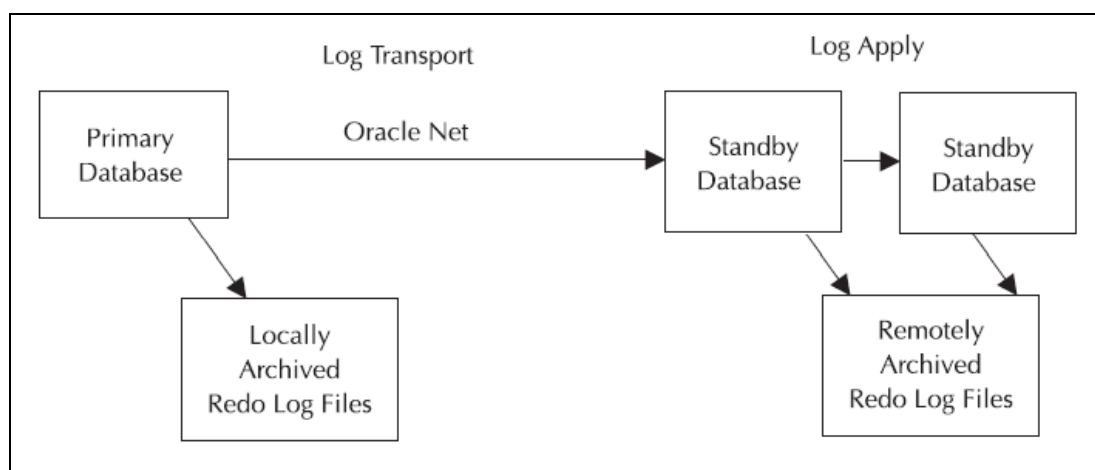


Figure 5 Multi-standby database Oracle Data Guard implementation

In case of a network outage, Data Guard can automatically synchronize the data by applying the redo data to the standby database that was archived at the primary database during the outage period. Data Guard ensures that the data is logically and physically consistent before it is applied to a standby database. For further details, refer to the Oracle Data Guard documentation.

The following types of standby databases are available from Oracle database version 11g onwards:

- ▶ Physical standby database
- ▶ Logical standby database
- ▶ Snapshot standby database

For high availability enhancements available in Oracle 12c, refer to the Maximum Availability Architecture (MAA) Oracle technology network information that is available at:

<http://www.oracle.com/technetwork/database/availability/maximum-availability-wp-12c-1896116.pdf>

In this paper, the standby database discussion is limited to the physical standby database.

Data protection modes

When you configure the primary and standby databases, you need to determine the level of data loss that is acceptable to the business. In the primary database, you define its archive log destination areas, at least one of which refers to the remote site that is used by the standby database. The ASYNC, SYNC, ARCH, LGWR, NOAFFIRM, and AFFIRM attributes of the LOG_ARCHIVE_DEST_n parameter for the standby database directs Oracle Data Guard to select among the following modes of operations:

- ▶ **Maximum protection (or no data loss) mode**

When using this mode, at least one standby location must be written to before a transaction commits in the primary database. The primary database shuts down if the standby database's log location is unavailable.

- ▶ **Maximum availability mode**

When using this mode, at least one standby location must be written to before a transaction commits in the primary database. If the standby location is not available, the primary database does not shut down. The redo that was generated is queued, and when the standby locations are available, it is transported and applied to the standby databases.

- ▶ **Maximum performance mode (default)**

When using this mode, transactions can commit before their redo information is sent to the standby locations. Commits in the primary database occur as soon as writes to the local online redo logs is complete. The ARCH process handles the writes to the standby locations by default.

For concepts and administration of data guard, refer to Oracle's documentation at:

<http://docs.oracle.com/database/121/SBYDB/concepts.htm#SBYDB00010>

IP partnership configuration

This section describes how to configure IP partnership between two Storwize V7000 systems.

To establish an IP partnership between two IBM Storwize V7000 systems, complete the following steps:

1. Set up system IP
2. Set up Challenge Handshake Authentication Protocol (CHAP) authentication (if required)
3. Configure IP addresses for remote ports
4. Establish IP partnerships

Set up system IP

You can set up the IP system by using the Storwize V7000 GUI or the command-line interface (CLI). Administrators need to set up system IPs to access the Storwize V7000 system as the initial step. But, depending on IP partnership deployment, an IP change might be required.

Connect to the Storwize system using a browser. Then click **Settings** → **Network** → **Management IP Address**, and enter the required IP address. Figure 6 shows the configuration of the IP address.



Figure 6 Configure cluster IP address

You can achieve the same configuration using the CLI as follows:

```
svctask chsystemip -clusterip 9.118.47.202 -mask 255.255.254.0 -gw 9.118.46.1
-port 1
```

Set up Challenge Handshake Authentication Protocol (CHAP) authentication (if required)

You can configure the Challenge Handshake Authentication Protocol (CHAP) secret independently for access to iSCSI hosts and remote copy partnerships. The same CHAP secret applies to both, however. Thus, depending on whether you want CHAP authentication for iSCSI hosts, remote copy partnerships, or both, you must configure parameters.

Connect to the Storwize system using a browser. Then click **Settings** → **Network** → **iSCSI** → **Modify CHAP Configuration**. Figure 7 shows the configuration of the CHAP secret.

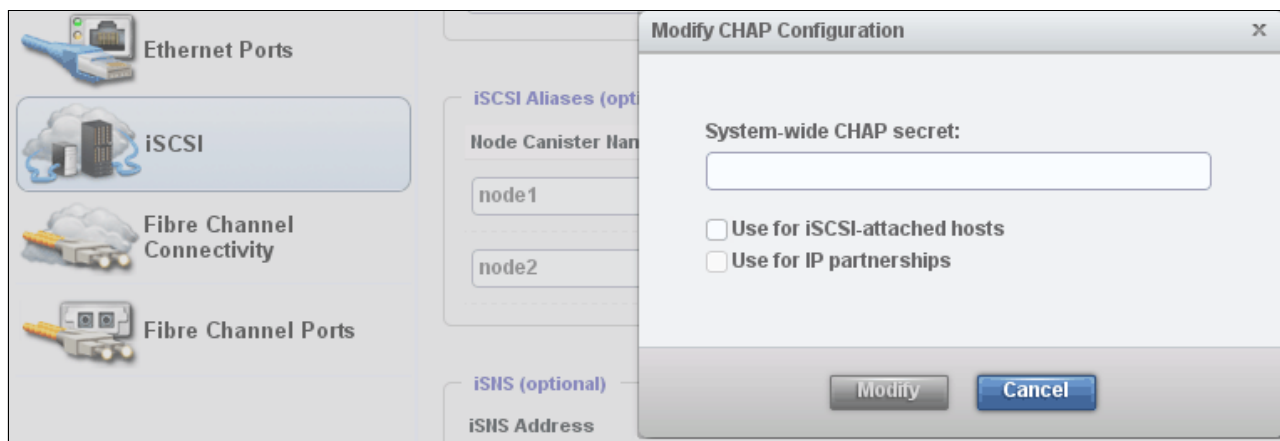


Figure 7 CHAP secret configuration

You can use the CLI to complete the same configuration as follows:

```
svctask chsystem -chapsecret ISVSetup14
```

Configure IP addresses for remote ports

You also need to configure IP addresses for the data path that is used in remote copy. First, connect to the Storwize system using a browser. Then click **Settings** → **Network** → **Ethernet Ports**. Right-click an Ethernet port, and click **Modify**.

Figure 8 shows the configured IP addresses. Remote port configuration must be done on the primary and the auxiliary Storwize systems.

Show IPv6 Actions Filter					
Node Name	Port	State	Speed	IP	
node1	1	✓ Active	1Gb/s		
node1	2	✓ Active	1Gb/s	9.118.47.246	
node1	3	✓ Active	10Gb/s	192.168.1.140	
node1	4	✓ Active	10Gb/s		
node2	1	✓ Active	1Gb/s		
node2	2	✓ Active	1Gb/s	9.118.47.162	
node2	3	✓ Active	10Gb/s	192.168.1.144	
node2	4	✓ Active	10Gb/s		

Figure 8 IP address configuration for remote ports

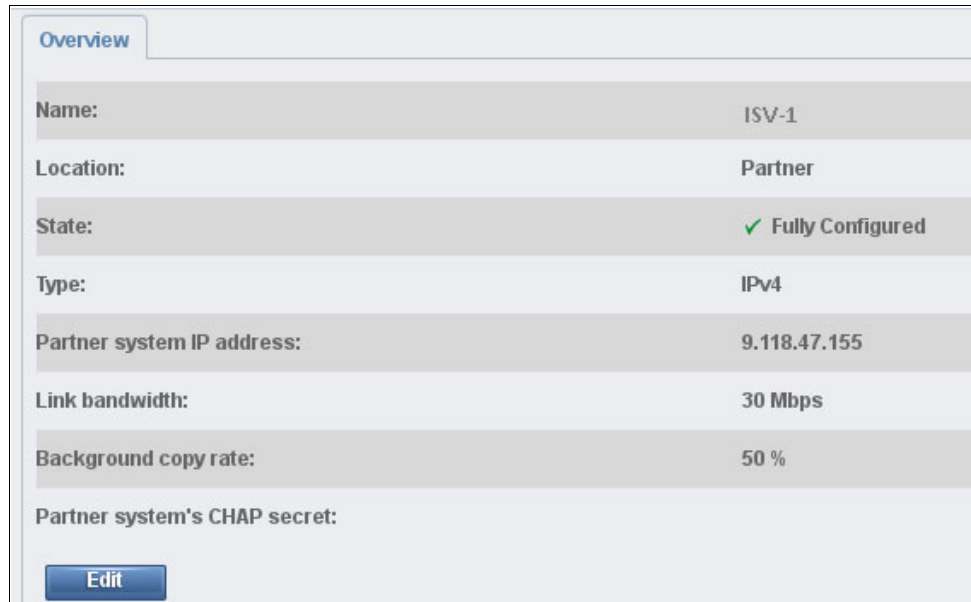
You can achieve the same configuration using the CLI as follows:

```
svctask cfgportip -node node1 -ip 9.118.47.246 -mask 255.255.254.0 -gw 9.118.46.1 -remotecopy 1 1
```

Establish IP partnerships

With the basic configuration in place, you next configure the IP partnership. Connect to the Storwize system using the browser, and then click **Copy Services** → **Partnerships** → **Create Partnership**.

A sample of configured partnership is shown in Figure 9.



The screenshot shows a web interface with a tab labeled 'Overview'. Below the tab is a form displaying the configuration for an IP partnership. The form has several rows, each with a label on the left and a value on the right. The values are: Name: ISV-1, Location: Partner, State: Fully Configured (with a green checkmark icon), Type: IPv4, Partner system IP address: 9.118.47.155, Link bandwidth: 30 Mbps, and Background copy rate: 50 %. At the bottom of the form is a button labeled 'Edit'.

Name:	ISV-1
Location:	Partner
State:	✓ Fully Configured
Type:	IPv4
Partner system IP address:	9.118.47.155
Link bandwidth:	30 Mbps
Background copy rate:	50 %
Partner system's CHAP secret:	
<input type="button" value="Edit"/>	

Figure 9 Configured IP partnership

On the primary system, the following command creates a new IP partnership.

```
svctask mkippartnership -type ipv4 -clusterip 9.118.47.202 -linkbandwidthmbits 30  
-backgroundcopyrate 50
```

On the auxiliary system, the following command creates a new IP partnership.

```
svctask mkippartnership -type ipv4 -clusterip 9.118.47.155 -linkbandwidthmbits 30  
-backgroundcopyrate 50
```

You can view the status of the defined IP partnership using a browser or from the command line using the following command.

```
lspartnership
```

Configuring iSCSI on Storwize V7000

This section describes the iSCSI configuration setup on Storwize V7000 systems created using the GUI.

You can start the iSCSI configuration wizard by clicking **Settings** → **Network**. 10 Gbps ports are configured for iSCSI with a unique IP address per port on each node, thus providing redundant channels of communication per node (refer to Figure 10). The iSCSI name assigned to the system is also displayed on this page and cannot be modified. An iSCSI alias can be given in a free-form text.

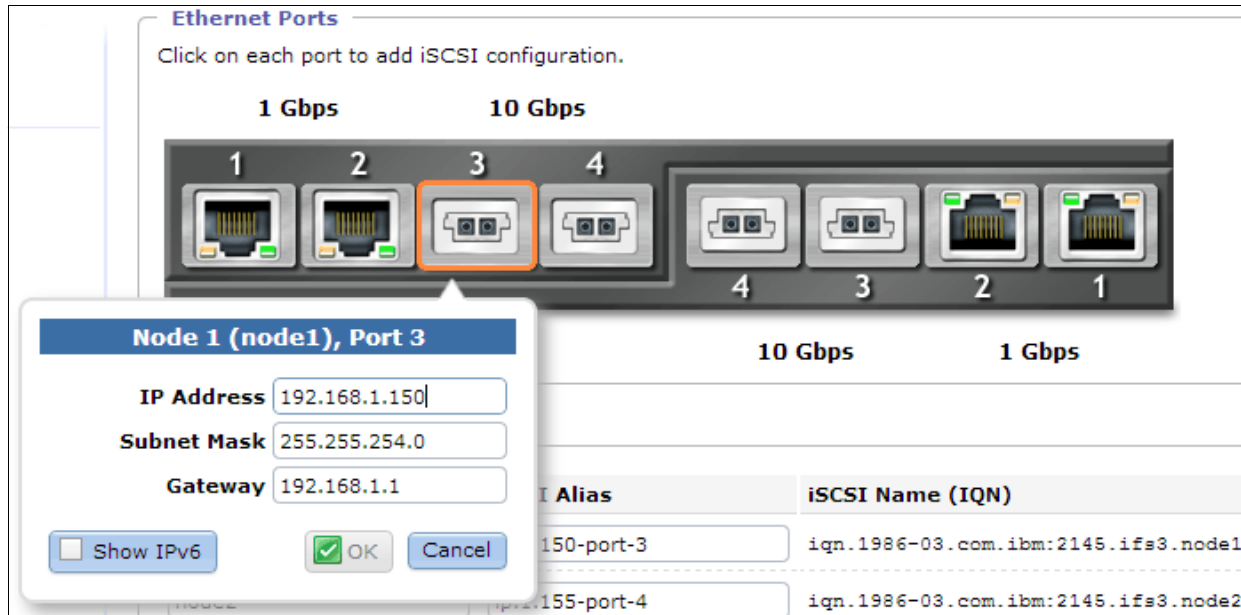


Figure 10 iSCSI configuration

Both 1 GB and 10 GB Ethernet ports can be used for iSCSI traffic, but only the 1 GB Ethernet ports can be used for management traffic.

All IP addresses (service and configuration) associated with a clustered-system Ethernet port must be on the same subnet. However, IP addresses associated with a node Ethernet port used for iSCSI traffic can be configured to belong to different subnets.

For more information about iSCSI configuration, refer to:

<http://www.ibm.com/support/knowledgecenter/ST3FR7/welcome>

Creating and mapping Storwize V7000 volumes using iSCSI

A *volume* is a logical disk that the clustered system presents to a host connected over a Fibre Channel or Ethernet network. Using volumes allows administrators to more efficiently manage storage resources.

For detailed information regarding how to create and map Storwize V7000 volumes to Linux, refer to:

<https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102563>

Lab setup

The lab setup represents two sites disaster recovery environment for both Oracle database and storage. The disaster recovery environment has the following characteristics.

- ▶ *Near disaster recovery* is storage-based replication that is achieved using Metro Mirror or Global Mirror.
- ▶ *Far disaster recovery* is an Oracle Data Guard-based solution in asynchronous mode.
- ▶ Database has Oracle ASM-based storage.
- ▶ ASM disks are virtual disks (VDIs) representing storage logical unit numbers (LUN) are carved from multiple managed disks (MDisks).
- ▶ Each LUN is exposed to database host using iSCSI.

Refer to the following iSCSI networking best practices guide for networking setup:

<http://ibm.co/1MjXJZQ>

The lab setup was created using three IBM System x based servers with Intel Xeon processors running Red Hat Enterprise Linux (RHEL) version 6.2. The two Storwize systems had the micro code level 7.2.

Figure 11 illustrates the lab topology.

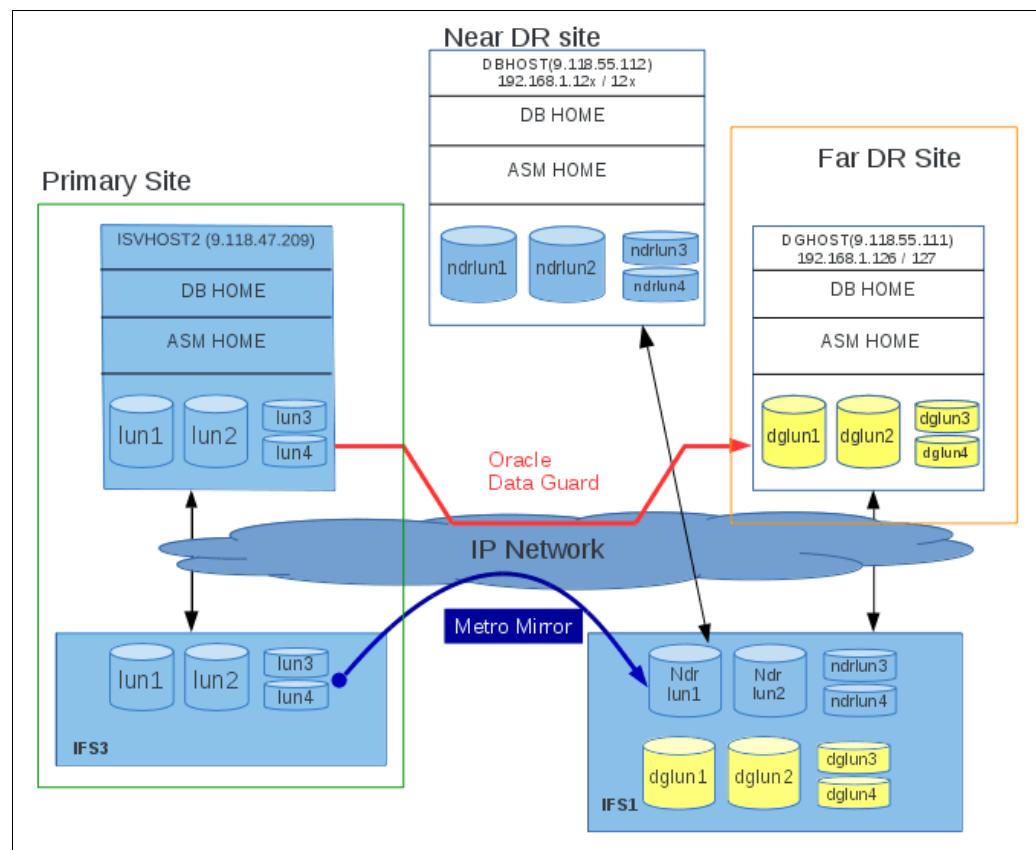


Figure 11 Lab topology

Although, in the lab setup, NDR and FDR sites are hosted on the same Storwize system, in the production environment, use different Storwize systems to host the environment so that the solution is more resilient to site failures.

As part of a failover strategy, storage sync mechanisms, such as Metro Mirror or Global Mirror, do not allow write actions on auxiliary storage. Thus, even if the storage LUNs are mapped to database host, Oracle processes, such as ASM or database, cannot be running.

Also for a disaster recovery strategy, use the following guidelines:

- ▶ Update the NDR site host name and IP address to match the domain name server (DNS) entry for primary site.
- ▶ If a host name change is not part of the failover policy, all the application servers Transparent Network Substrate (TNS) or Java Database Connectivity (JDBC) link name must be updated with the database host name and the IP address of the NDR site.
- ▶ Because the primary site is not available, the managed recovery process on the Data Guard host stalls. This process must be restarted in order to resume recovery from the NDR site.
- ▶ If the host name is used in the TNS entry, before restarting the managed recovery process, `TNSNAMES.ORA` must be modified on the Data Guard host (changed from primary to NDR).

At this point, the test set up has the following properties:

- ▶ Oracle database is up and running at the primary site.
- ▶ Oracle Data Guard setup is available on the FDR site created as per the best practices suggested by Oracle.
- ▶ The NDR site is created (as clone) for the primary site.

For testing, the test team used the combinations of Metro Mirror and Oracle Data Guard as well as Global Mirror and Oracle Data Guard. The following database host names and database names were used:

- ▶ `ISVHOST2` is the primary site database system.
- ▶ `DBHOST` is the NDR site database system.
- ▶ `DGHOST` is the FDR site database system.
- ▶ Database name on the primary site is `ORCL` with the `DB_UNIQUE_NAME` parameter set to the `ORCL` value.
- ▶ Database name on the FDR site is `ORCL` with the `DB_UNIQUE_NAME` parameter set to the `DGORCL` value.

Metro Mirror and Oracle Data Guard

This section describes the normal working scenario of the solution together with Metro Mirror and Oracle Data Guard.

To define the Metro Mirror relationship, you must follow the steps provided in “IP partnership configuration” on page 8.

Creating an Oracle Data Guard setup is independent of the Metro Mirror setup and, thus, can be completed before and after Metro Mirror configuration.

When Metro Mirror is active, the LUNs participating in the relationship are not available for read/write access at NDR (auxiliary site). Although the role can be reversed for copying, the partner is always put in the read-only mode. Oracle Data Guard alternatively works at the application level. So the standby database is available for read-only queries. Alternatively, the database can also be opened temporarily for the read/write mode using the flashback database technique.

Defining and starting the copy services process (Metro Mirror) undergoes various status changes, such as Inconsistent Stopped, Inconsistent Copying, or Consistent Synchronized.

Similar to the Metro Mirror status changes, when the Oracle Data Guard is implemented, the choice is either to create it from running the database (option available from 11gR2 onwards) or from the last full backup. In both cases, you need to create a standby control file from running the primary database. When created, the standby database also lags behind the most recent change on the primary database. Starting the managed recovery on the standby database allows and maintains the standby in the most updated state.

Having set up the protection for the primary site by means of both Metro Mirror and Data Guard enables you to look at the possible disaster recovery scenarios listed in Table 1 that were used in the exercise.

Table 1 Disaster recovery scenarios

Case	Primary	Secondary (NDR)	Data Guard (FDR)
Case 1: Primary site failure	Not available	Database available in RW mode	Database in the read-only mode (Master=secondary)
Case 2: Primary + secondary site failure	Not available	Not available	Database in the read/write mode (Master mode)

Case 1: Primary site failure

This scenario describes the steps taken when the primary site is not available.

Using Metro Mirror, although the data in the Oracle memory structures cannot be accounted for, the last acknowledged write made to storage guarantees the availability of the written data to the NDR site. Thus, in the absence of the primary site, the database on the NDR site can be started by without having to perform any crash recovery.

The following steps illustrate the sequence of database start on the NDR site. The test setup does not contain DNS to resolve the host name to IP addresses. The `/etc/hosts` file is used instead of DNS.

1. Log on to the Oracle Data Guard host (DGHOST) and connect to the standby database using the SQL*Plus utility.
2. Stop the managed recovery using the **ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL** command.
3. Edit the `tnsnames.ora` file, and change the TNS entry, ORCL, to adapt to the new database host (`host=DBHOST`).
4. Log in to the Storwize V7000 system using the web browser. Click **Copy Services** → **Remote Copy**.
5. Right-click the consistency group, and click **Stop**.
6. Log on to DBHOST as an Oracle grid infrastructure owner, and start the Oracle High Availability Services (HAS).
7. Log on to DBHOST as an Oracle database user, and issue the **database startup** command at the SQL*Plus prompt.

8. When the database has started on DBHOST, log on to DGHOST, and start the managed recovery using the **ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT** command.

It is a good practice to run the **tail** command on both the ASM and database alert log file to view any instance startup messages.

Figure 12 illustrates this scenario. In the absence of a primary site, the failover to the NDR site is effective. For Oracle Data Guard, the NDR site becomes the source for log shipping.

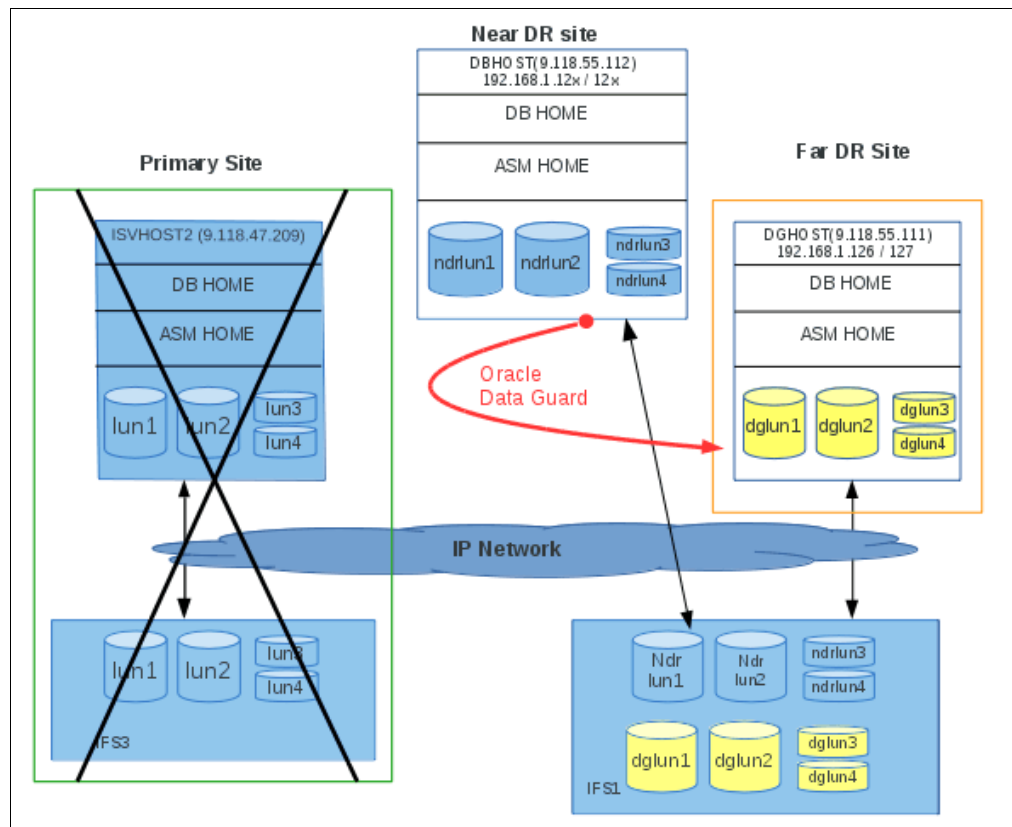


Figure 12 Failover to the NDR site

Metro Mirror role reversal

This section describes the steps taken when the primary site is available and the Metro Mirror role reversal needs to be performed.

In the event of re-availability of the primary site, a failback to the primary site can be performed. The most important step in performing failback is configuring the Metro Mirror relationship. In this case, NDR is the primary site, and the former primary site acts as the secondary or auxiliary site. Thus, the replication can now copy the data from the NDR site to the primary site. The steps to configure the replication using Metro Mirror remain the same, except for the changes that are needed to indicate the primary and auxiliary storages.

Defining and starting the copy services process (Metro Mirror) undergoes various status changes, such as Inconsistent Stopped, Inconsistent Copying, and Consistent Synchronized.

To start the database on the primary site:

1. Log on to Oracle Data Guard host (DGHOST), and connect to the standby database using the SQL*Plus utility.
2. Stop the managed recovery using the **ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL** command.
3. Edit the `tnsnames.ora` file, and change the TNS entry, ORCL, to adapt to the new database host (host=ISVHOST2).
4. Log in to Storwize V7000 using a web browser, and navigate to **Copy Services** → **Remote Copy**.
5. Right-click the consistency group, and click **Stop**.
6. Log on to ISVHOST2 as an Oracle grid infrastructure owner, and start the Oracle HAS service.
7. Log on to ISVHOST2 as an Oracle database user, and issue the **database startup** command at the SQL*Plus prompt.
8. When the database starts on ISVHOST2, log on to DGHOST, and start managed recovery using the **ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT** command.

Figure 13 illustrates the Metro Mirror role reversal. Storage at the NDR site copies the data on the primary site.

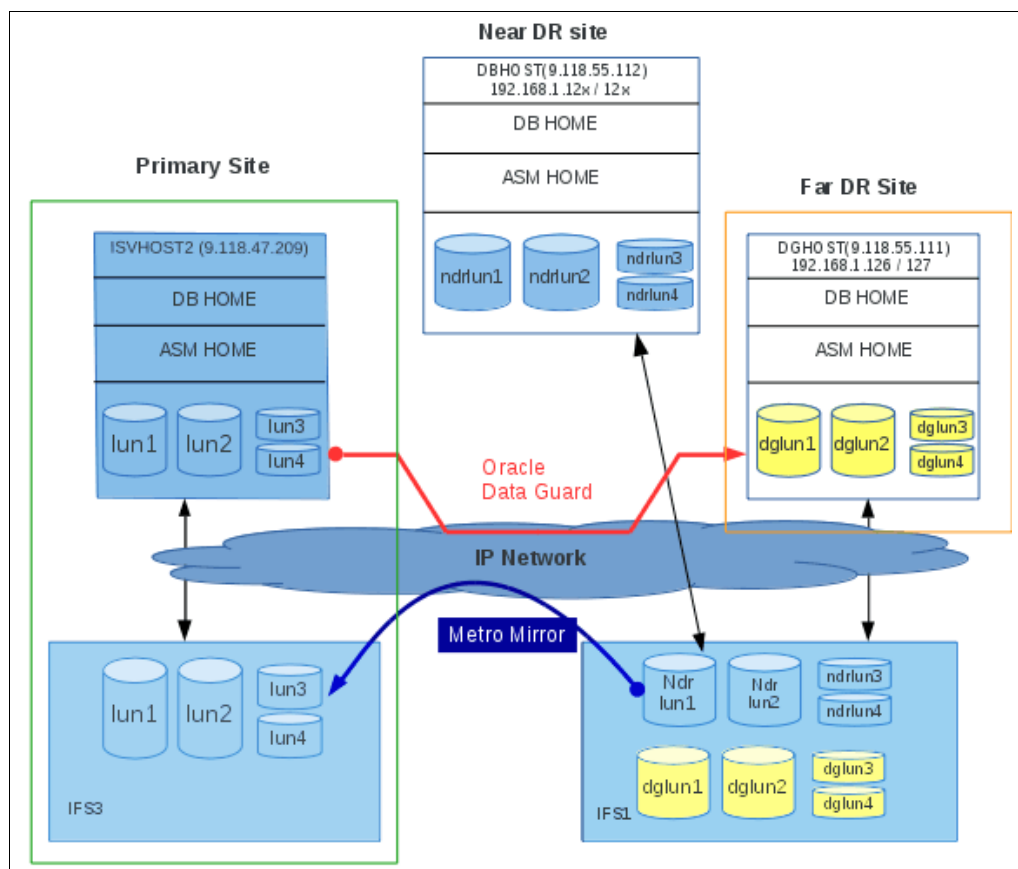


Figure 13 Metro Mirror role reversal

Case 2: Primary and secondary site failure

This scenario describes the steps taken when both the primary and secondary sites are not available. In the event of unavailability of both primary and secondary site failure, the FDR site that is maintained using Oracle Data Guard is made active.

To start the database on the FDR site:

1. Log on to the Oracle Data Guard host (DGHOST), and connect to the standby database using SQL*Plus.
2. At the SQL*Plus prompt, specify to stop the managed recovery using the **ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL** command.
3. Convert the standby database to a primary database using the **ALTER DATABASE ACTIVATE STANDBY DATABASE** command.
4. View the database alert log file, and confirm that the standby database was converted to the primary database.
5. Open the new primary database for read/write operations.
6. Back up the new primary database.

Figure 14 illustrates this scenario.

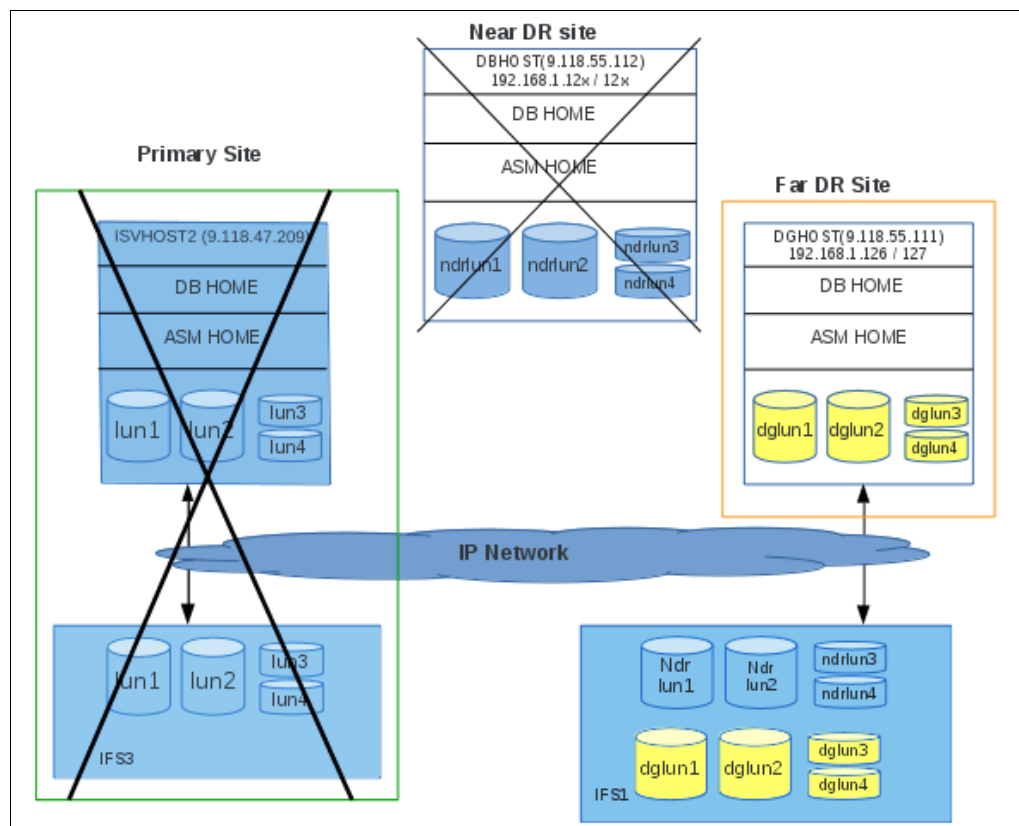


Figure 14 Far disaster recovery site host active as primary

Global Mirror + Oracle data guard

This section describes the normal working scenario of a solution together with Global Mirror and Oracle Data Guard.

To define the Global Mirror relationship, you must follow the steps provided in “IP partnership configuration” on page 8.

Creating an Oracle Data Guard setup is independent of the Global Mirror setup and, thus, can be done before or after the Global Mirror configuration.

In the test setup, a Global Mirror relationship was created between the primary and the NDR sites. When Global Mirror is active, the LUNs participating in the relationship are not available for read/write access at the NDR (auxiliary) site. Although the role can be reversed for copying, the partner is always put in read-only mode. The standby database created using Oracle Data Guard is available for read-only queries. Alternatively, the standby database can also be opened temporarily for the read/write mode using the flashback database technique.

Defining and starting the copy services process (Global Mirror) undergoes various status changes, such as Inconsistent Stopped, Inconsistent Copying, or Consistent Synchronized.

The disaster recovery scenarios tested with Global Mirror and Oracle Data Guard are the same as those tested for Metro Mirror and Oracle Data Guard. The steps listed for database startup on the NDR and FDR sites and role reversal are also same.

Summary

This paper combines the use of the IP replication technique available at storage using copy services, such as Metro Mirror and Global Mirror and Oracle Data Guard, to create a multi-site disaster recovery solution based on the iSCSI protocol offered by Storwize V7000. This paper also provides references that act as guidelines for Ethernet network and OS-specific considerations, understanding and configuring IP partnership, and deploying Oracle 12c using iSCSI.

Resources

The following resources provide useful references to supplement the information included in this paper:

- ▶ IBM Systems on PartnerWorld®
https://www.ibm.com/partnerworld/wps/servlet/ContentHandler/stg_ove_sys-ibm-systems
- ▶ IBM Storwize V7000 Knowledge Center
<http://www.ibm.com/support/knowledgecenter/ST3FR7/welcome>
- ▶ iSCSI networking best practices
<http://ibm.co/1MjXJZQ>
- ▶ Configure IBM System Storage® SVC using IP Partnership
<http://ibm.co/108qFpe>
- ▶ Oracle 12c database administrators guide
<http://docs.oracle.com/database/121/ADMIN/toc.htm>
- ▶ Oracle 12c grid infrastructure information
http://docs.oracle.com/database/121/nav/porta1_16.htm
- ▶ Concept and administration of Oracle Data Guard
<http://docs.oracle.com/database/121/SBYDB/concepts.htm#SBYDB00010>

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Shashank Shingornikar is currently working as an ISV Storage Solutions Architect with IBM India Software Development Lab in Pune, India. He has more than 18 years of experience in Oracle technology products and has expertise on Oracle data migration and disaster recovery solutions. Shashank is leading the solutions for IBM Storage products such as DS8K, XIV®, and Storwize V7000 with Oracle. Shashank holds a B.Sc degree from Mumbai University.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks publications

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5286-00 was created or updated on October 13, 2015.

Send us your comments in one of the following ways:


- Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- Send your comments in an email to:
redbooks@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

FlashCopy®
IBM®
PartnerWorld®

Redbooks®
Redpaper™
Redbooks (logo) ®

Storwize®
System Storage®
XIV®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.



REDP-5286-00

ISBN 0738454516

Printed in U.S.A.

Get connected

