IBM

DS8000 Copy Services for IBM i with VIOS

- Establish multiple paths to multiple VIOS partitions
- Learn about VIOS with Node Port ID Virtualization
- Understand Disaster Recovery scenarios

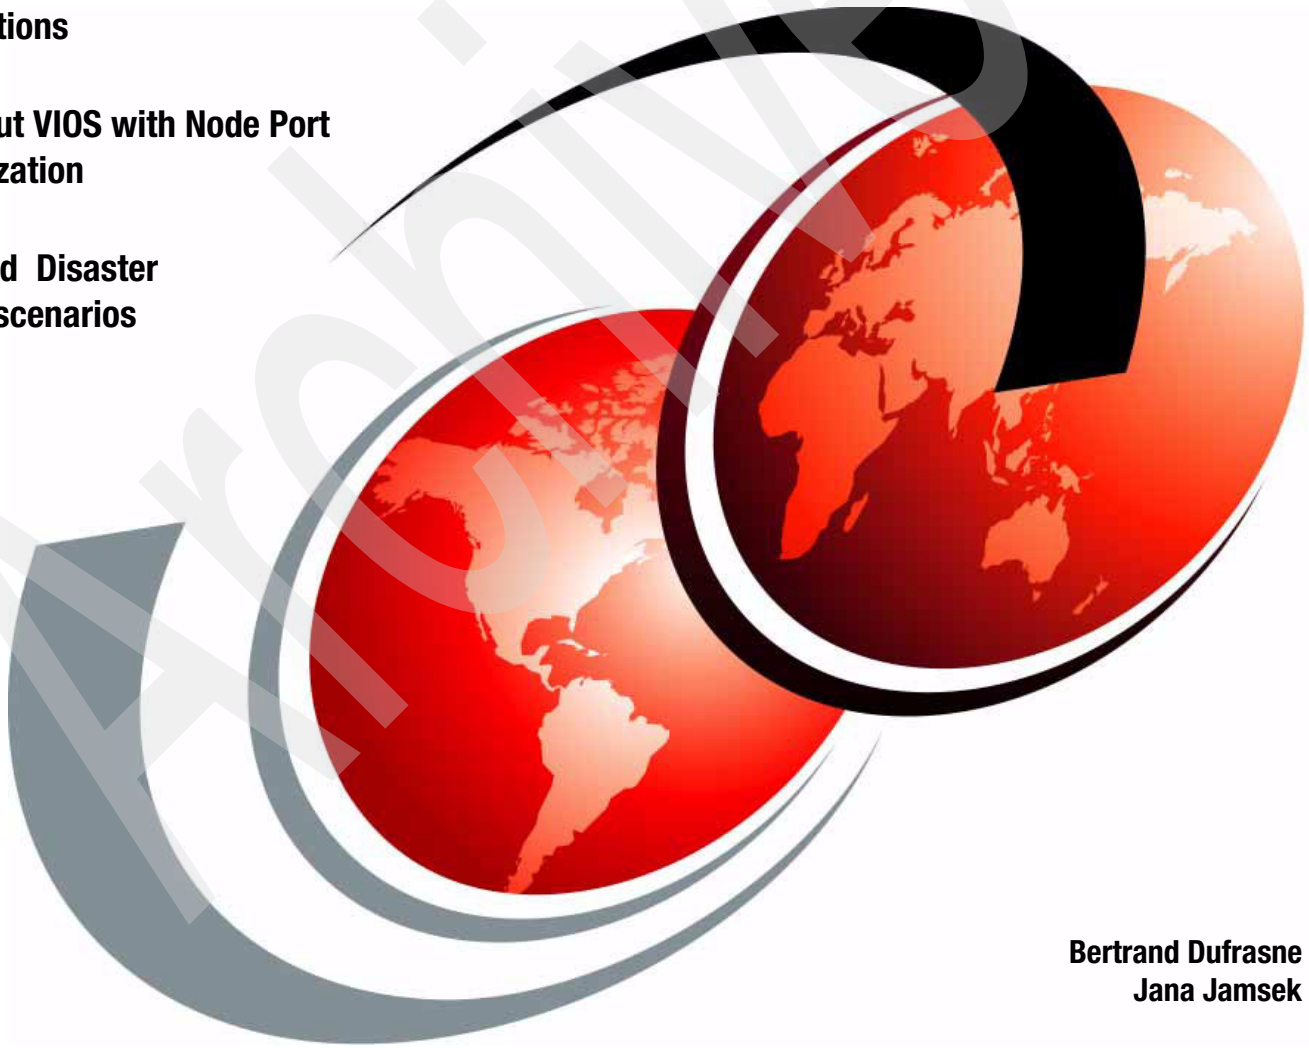

Bertrand Dufrasne
Jana Jamsek

Redpaper

ibm.com/redbooks

International Technical Support Organization

DS8000 Copy Services for IBM i with VIOS

November 2009

REDP-4584-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (November 2009)

This edition applies to the IBM System Storage DS8000 License Machine Code level 5.4.30.xxx (bundle version 64.30.xx.x) with VIOS 2.1.2.0.

Contents

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

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.

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:

AIX®
AS/400®
DS8000®
FlashCopy®
i5/OS®
IBM®
Micro-Partitioning™
Power Systems™
POWER6®
PowerHA™
PowerVM™
POWER®
Redbooks®
Redpaper™
Redbooks (logo)  ®
System i®
System p®
System Storage™
System z®
Tivoli®
XIV®
z/OS®

The following terms are trademarks of other companies:

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

Preface

IBM® System Storage™ DS8000® Copy Services provide functions for Disaster Recovery and Business Continuity solutions for an IBM i server. Many IBM i customers with natively connected DS8000 use Metro Mirror or Global Mirror functions to provide a set of consistent data at the remote site. This approach allows fail over or restoring data in case of disaster at the production site. You may also use IBM FlashCopy® technology to provide point-in-time copy of disk, from which you can back up data to minimize the necessary downtime.

This IBM Redpaper™ publication is intended for those who want to plan or implement DS8000 Copy Services with IBM i. It discusses the use of DS8000 Copy Services for an IBM i system connected with VIOS Node Port ID Virtualization (NPIV). For this paper, we created test scenarios with FlashCopy, Metro Mirror, and Global Mirror of Sysbas. We used DS8000 DS command-line interface (DS CLI) commands to manage the scenarios.

The team who wrote this paper

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

Bertrand Dufrasne, is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage disk products at the International Technical Support Organization, San Jose Center.

Jana Jamsek is an IT Specialist in IBM Slovenia. She works in Storage Advanced Technical Support for Europe as a specialist for IBM Storage Systems and i5/OS® systems. Jana has eight years of experience in the System i® and AS/400® areas, and eight years of experience in storage. She has a Master's degree in computer science and a degree in mathematics from the University of Ljubljana, Slovenia. She has authored several IBM Redbooks® publications.

Thanks to the following people at IBM for their contributions to this project:

Ingo Dimmer
Peter Kimmel
James Lembe
Walter Lipp
Vess Natchev
Lamar Reavis
Russ Vanduine

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail 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

1

Introduction

Although the DS8000 can be attached to a native IBM i partition, it can also be connected to IBM i through the Virtual I/O Server (VIOS).

A connection through VIOS is especially interesting for the IT centers that have many small IBM i partitions. With a native DS8000 connection, one Fibre Channel (FC) adapter, or preferably two FC adapters (for availability and multipathing), are required for every partition, which might be rather expensive. If you use VIOS, however, the FC adapters can be installed in the Virtual I/O Server only; IBM i clients use virtual connectivity to the VIOS.

The IT centers, whose strategy is virtualization of all their Power partitions, might also want to consolidate the attachment of DS8000 disk for all the partitions through the Virtual I/O Server.

DS8000 Copy Services provide functionality for Disaster Recovery and Business Continuity solutions for an IBM i server. Many IBM i customers with natively connected DS8000 use Metro Mirror or Global Mirror to provide a set of consistent data at the remote site. This allows them to fail over or restore data in case of disaster at the production site. They also use FlashCopy to provide point in time copy of disk, from which they backup their data to minimize the needed downtime.

Copy Services for natively connected DS8000 can be implemented with IBM i system disk pool and user disk pools (Sysbas) or with Independent storage pool (IASP). Customers have a choice of different management tools for their solutions, such as PowerHA™, System i Copy Services Toolkit, and Tivoli® Storage Productivity Center for Replication (TPC-R).

This paper relates scenarios using DS8000 Copy Services for an IBM i system connected with VIOS Node Port ID Virtualization (NPIV). The tests were done with FlashCopy, Metro Mirror, and Global Mirror of Sysbas. We used DS8000 DS CLI commands to manage the scenarios.

When connecting DS8000 to IBM i with VIOS NPIV, the VIOS is completely transparent to the user. You define DS8000 LUNs as for an IBM i partition, and you do not need to perform device discovery in VIOS before using Copy Services target volumes in IBM i.

Although not used in developing this paper, we expect that PowerHA with IASP, and TPC-R management will also be used with implementations in this environment.

2

Virtualization with PowerVM

This chapter introduces PowerVM™ and the VIOS virtualization software.

Virtualization capabilities are becoming a very important consideration for server design and planning. With commercial and environmental concerns growing, a strong pressure exists for reducing the footprint of servers and their overall power consumption.

Virtualization in Power systems provides rapid, cost-effective response to these business needs. Virtualization in Power systems enables you to increase the use of servers, which reduces:

- Server management and administration costs, because of fewer physical servers
- Power and cooling costs with increased use of existing servers
- Time to market, virtual resources can be deployed immediately

Virtualization technology of Power systems is provided by IBM PowerVM brand.

2.1 IBM PowerVM overview

IBM PowerVM is a special software appliance associated with IBM POWER® Systems, that is, the converged IBM i and IBM p server platforms. It is licensed on a POWER system processor basis.

IBM PowerVM is a virtualization technology for AIX®, IBM i, and Linux® environments on IBM POWER processor-based systems. IBM Power Systems™ servers coupled with PowerVM technology are designed to help clients build a dynamic infrastructure, reducing costs, managing risk and improving service levels.

PowerVM offers a secure virtualization environment, which has the following major features and benefits:

- Consolidates diverse sets of applications built for multiple operating systems on a single server: AIX, IBM i, and Linux.
- Virtualizes processor, memory, and I/O resources to increase asset utilization and reduce infrastructure costs.
- Dynamically adjusts server capability to meet changing workload demands.
- Moves running workloads between servers to maximize availability and avoid planned downtime.

Virtualization technology is offered in three editions on Power Systems:

- PowerVM Express Edition
- PowerVM Standard Edition
- PowerVM Enterprise Edition

They provide logical partitioning (LPAR) technology by using either the Hardware Management Console (HMC) or the Integrated Virtualization Manager (IVM), Dynamic LPAR operations, Micro-Partitioning™ and Virtual I/O Server capabilities, and Node Port ID Virtualization (NPIV).

PowerVM Express Edition

PowerVM Express Edition is offered only on the IBM Power 520 and Power 550 servers. It is designed for clients who want an introduction to advanced virtualization features, at a highly affordable price.

With PowerVM Express Edition, clients can create up to three partitions on a server (two client partitions and one for the Virtual I/O Server and Integrated Virtualization Manager). They can leverage virtualized disk and optical devices, and try out the shared processor pool. All virtualization features, such as Micro-Partitioning, Shared Processor Pool, Virtual I/O Server, PowerVM LX86, Shared Dedicated Capacity, N Port ID Virtualization, and Virtual Tape, can be managed by using the Integrated Virtualization Manager.

PowerVM Standard Edition

For clients ready to get the full value from their server, IBM offers PowerVM Standard Edition. This provides the most complete virtualization functionality for UNIX® and Linux in the industry. This option is available for all IBM Power Systems servers.

With PowerVM Standard Edition, clients can create up to 254 partitions on a server. They can leverage virtualized disk and optical devices and try out the shared processor pool. All virtualization features, such as Micro-Partitioning, Shared Processor Pool, Virtual I/O Server, PowerVM Lx86, Shared Dedicated Capacity, N Port ID Virtualization, and Virtual Tape, can be managed by using an HMC or the IVM.

PowerVM Enterprise Edition

PowerVM Enterprise Edition is offered exclusively on POWER6® servers. It includes all the features of PowerVM Standard Edition, plus a capability named PowerVM Live Partition Mobility.

PowerVM Live Partition Mobility allows for the movement of a running partition from one POWER6 technology-based server to another with no application downtime. This results in better system utilization, improved application availability, and energy savings. With PowerVM Live Partition Mobility, planned application downtime, because of regular server maintenance, can be a thing of the past.

2.2 Virtual I/O Server

Virtual I/O Server (VIOS) is virtualization software that runs in a separate partition of your POWER system. Its purpose is to provide virtual storage and networking resources to one or more client partitions.

The Virtual I/O Server owns the physical I/O resources such Ethernet and SCSI/FC adapters. It virtualizes those resources for its client LPARs to share them remotely using the built-in hypervisor services. These client LPARs can be quickly created, typically owning only real memory and shares of processors without any physical disks or physical Ethernet adapters.

Virtual SCSI support allows VIOS client partitions to share disk storage that is physically assigned to the Virtual I/O Server logical partition. This virtual SCSI support of VIOS is used to make storage devices such as XIV® that do not support the IBM i proprietary 520-byte per sectors format available to IBM i clients of VIOS.

VIOS owns the physical adapters, such as the Fibre Channel storage adapters connected to the IBM XIV Storage System. The LUNs of the physical storage devices seen by VIOS are mapped to VIOS virtual SCSI (VSCSI) server adapters created as part of its partition profile.

The client partition with its corresponding virtual SCSI client adapters defined in its partition profile connects to the VIOS virtual SCSI server adapters through the hypervisor with VIOS performing SCSI emulation and acting as the SCSI target for IBM i.

Figure 2-1 shows an example of the Virtual I/O Server owning the physical disk devices and its virtual SCSI connections to two client partitions.

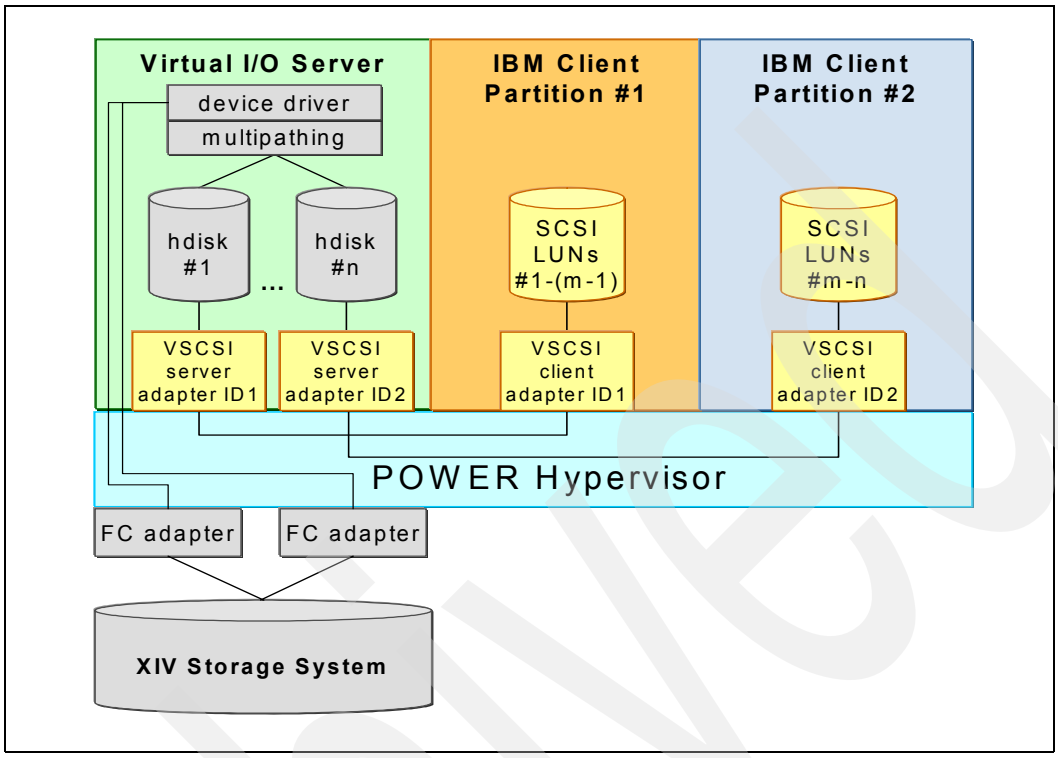

Figure 2-1 VIOS Virtual SCSI support

2.3 Node Port ID Virtualization

The VIOS technology has been enhanced to boost the flexibility of Power Systems servers with support for Node Port ID Virtualization (NPIV). NPIV simplifies the management and improves performance of Fibre Channel SAN environments by standardizing a method for Fibre channel ports to virtualize a physical node port ID into multiple virtual node port IDs. The VIOS takes advantage of this feature and can export the virtual node port ids to multiple virtual clients. The virtual clients see this node port ID and can discover devices just as though the physical port was attached to the virtual client. The VIOS does not do any device discovery on ports that use NPIV. Therefore, no devices are shown on the VIOS that is connected to NPIV adapters. The discovery is performed by the virtual client; all the devices found during discovery are seen only by the virtual client. This approach allows the virtual client to use Fibre Channel SAN (FC SAN) storage specific multipathing software on the client to discover and manage devices.

Figure 2-2 on page 7 shows a managed system configured to use NPIV, running two Virtual I/O Server partitions each with one physical Fibre Channel card. Each Virtual I/O Server partition provides virtual Fibre Channel adapters to the virtual I/O client. For increased serviceability you can use multipath I/O (MPIO) in the AIX virtual I/O client.

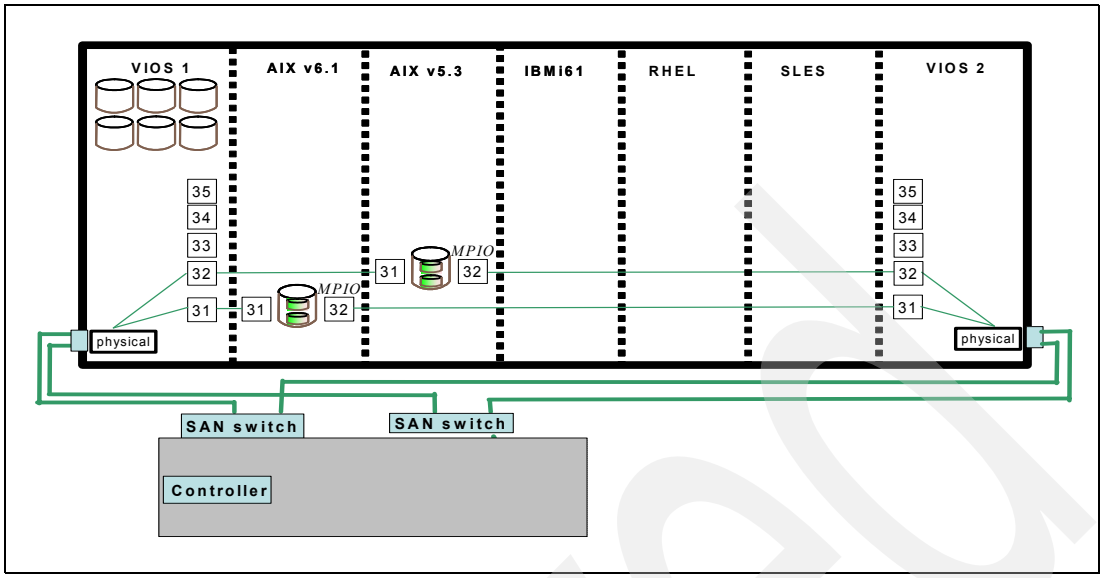

Figure 2-2 Virtual I/O Server partitions with NPIV

Further information regarding PowerVM virtualization management can be found in *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590.

3

Connecting PowerVM clients to DS8000

This chapter reviews the requirements and general setup procedures for connecting a PowerVM client, such as an IBM i client, to the DS8000 Storage System.

3.1 Requirements for connecting DS8000 to IBM i with VIOS

In this section, we describe hardware and software requirements for connecting a DS8000 to an IBM i partition for the following configurations:

- Connecting IBM i with virtual SCSI adapters in single path through VIOS

  The hardware and software requirements are:

  - VIOS and IBM i partitions reside on IBM POWER Systems POWER6 server model
  - System firmware 320_040_031 or later
  - PowerVM Standard (PID 5765-PVS) or Enterprise Edition (PID 5765-PVE)
  - VIOS release 1.5.0 with fix pack 10.1 or later
  - HMC V7 R3.2.0 or later
  - IBM i version 6.1.0 or later

- IBM i partition with virtual SCSI adapters and multipath through two VIOS partitions

  The hardware and software requirements are:

  - VIOS and IBM i partitions reside on the IBM POWER Systems POWER6 server model.
  - System firmware 320_040_031 or later
  - VIOS release 2.1.2.0 or later
  - HMC V7 R3.2.0 or later
  - IBM i version 6.1.1 or later

- IBM i partition in multipath through two VIOS partitions with NPIV

  The hardware and software requirements are:

  - VIOS and IBM i partitions reside on the IBM POWER Systems POWER6 server model
  - 8 Gb FC adapter in VIOS; the latest firmware is recommended
  - NPIV enabled SAN switches
  - System firmware 01EM350_014 or later
  - VIOS release 2.1.2.0 or later
  - HMC V7 R3.5.0 build level 20090707 or later

For information about how to create the VIOS and IBM i client partitions in the Power6 server, refer to *IBM i and Midrange External Storage*, SG24-7668. That book also explains how to install VIOS in a partition on the Power6 server.

3.1.1 Multipath capability for IBM i through two VIOS partitions

IBM i provides a multipath capability to access a DS8000 LUN by using multiple connections and establishing one path through each connection. Up to 16 paths to the same LUN or set of LUNs are supported. Multipathing provides redundancy against a connection failure, and also increases performance by using all available paths to the LUNs for I/O operations.

With VIOS release 2.1.2.0 or later, and IBM i release 6.1.1 or later, establishing multiple paths to a set of LUNs, with each path using a connection through a different VIOS partition, is possible. Using multipathing with different VIOS partitions provides redundancy against connection failure and also against VIOS partition failure. Up to 16 multipath connections can be implemented to the same set of LUNs, each through a different VIOS. However, we expect that most IT centers will establish no more than two such connections.

3.2 Virtual SCSI adapters in multipath with two VIOS partitions

In this section, we discuss how to connect and implement an IBM i partition with multipathing through two VIOS partitions over virtual SCSI adapters.

3.2.1 Connecting with virtual SCSI adapters in multipath, two VIOS partitions

In our setup, we use two VIOS partitions and two virtual SCSI adapters in the IBM i partition. Each SCSI adapter is assigned to a virtual adapter in one VIOS. We connect the same set of DS8000 LUNs to each VIOS, and map them to virtual SCSI adapters serving the IBM i partition. This way, IBM i sees the LUNs through two paths, each path through one VIOS.

For more information about configuring the DS8000, refer to *IBM System Storage DS8000: Architecture and Implementation,* SG24-6786.

To realize this setup, perform the following steps:

1. Assuming that the arrays, ranks, and extent pools are configured in the DS8000, connect the ports of the DS8000 HBAs to FC adapters in VIOS; the connected ports should be defined as SCSI-FCP.
2. In the DS8000, create the volume group as for an IBM System p® sever. Use the following DS CLI command:

   ```
   mkvolgrp -hosttype pSeries prod_flash
   ```

3. Create the volumes for the volume group that was previously defined. Use the following DS CLI command:

   ```
   mkfbvol -extpool P4 -name prod_#h -cap 70 -volgrp V17 4800-4807
   ```

4. For each FC adapter in VIOS, create a host connection for the defined volume group. For example, we used the following command:

   ```
   mkhostconnect -wwname 10000000c98319f6 -hosttype pSeries -volgrp v17
   VIOS_1_adapter
   ```

5. For *each* VIOS partition, perform the following steps to continue with the setup:
   a. Log in to VIOS as administrator. In our example, we use PUTTY to log in, as described in *IBM i and Midrange External Storage*, SG24-7668, (refer to the information about configuring VIOS virtual devices).

      Issue the command **`cfgdev`** to have the VIOS recognize newly attached LUNs.

      In our example, the DS8000 LUNs are listed as `hdisk2` - `hdisk17`, as shown in Figure 3-1 on page 12.

```
 $ lsdev -type disk
name             status      description
hdisk0           Available   SAS RAID 0 Disk Array
hdisk1           Available   SAS RAID 0 Disk Array
hdisk2           Available   MPIO Other FC SCSI Disk Drive
hdisk3           Available   MPIO Other FC SCSI Disk Drive
hdisk4           Available   MPIO Other FC SCSI Disk Drive
hdisk5           Available   MPIO Other FC SCSI Disk Drive
hdisk6           Available   MPIO Other FC SCSI Disk Drive
hdisk7           Available   MPIO Other FC SCSI Disk Drive
hdisk8           Available   MPIO Other FC SCSI Disk Drive
hdisk9           Available   MPIO Other FC SCSI Disk Drive
hdisk10          Available   MPIO Other FC SCSI Disk Drive
hdisk11          Available   MPIO Other FC SCSI Disk Drive
hdisk12          Available   MPIO Other FC SCSI Disk Drive
hdisk13          Available   MPIO Other FC SCSI Disk Drive
hdisk14          Available   MPIO Other FC SCSI Disk Drive
hdisk15          Available   MPIO Other FC SCSI Disk Drive
hdisk16          Available   MPIO Other FC SCSI Disk Drive
hdisk17          Available   MPIO Other FC SCSI Disk Drive
```

Figure 3-1 DS8000 LUNs listed as disk devices in VIOS

b. In the multipath setup (IBM i with two VIOS partitions), each DS8000 LUN is seen by each of the two VIOS partitions. Before connecting the LUN (in each VIOS) to the IBM i client, you must ensure that the LUN is not SCSI-reserved. SCSI reservation is the default in VIOS, and you have to specifically change the LUNs to non-reserved.

   You may first want to check the current reserve policy, with the **`lsdev -dev hdiskx -attr reserve_policy`** command, for a disk device representing the DS8000 LUN, as shown in Figure 3-2.

```
$ lsdev -dev hdisk2 -attr reserve_policy
value

single_path
$
```

Figure 3-2 Current SCSI Reserve policy

   If the current reserve policy is other than `no_reserve`, change it to `no_reserve` by executing the **`chdev -dev hdiskx -attr reserve_policy=no_reserve`** command for each disk device corresponding to a DS8000 LUN. See Figure 3-3.

```
$ chdev -dev hdisk2 -attr reserve_policy=no_reserve -perm
hdisk2 changed
$
```

Figure 3-3 Change SCSI reserve policy to non_reserve

c. Change the attributes for all FC ports in VIOS as is shown in Figure 3-4.

```
*$ chdev -dev fscsi0 -attr fc_err_recov=fast_fail dyntrk=yes -perm
fscsi0 changed
$
```

Figure 3-4 Changing the attributes in FC ports

d. Map the disks (that correspond to DS8000 LUNs) to the virtual SCSI adapters assigned to the IBM i client. First, check the IDs of assigned virtual adapters by following these steps:

   i. Open the partition profile of the IBM i partition (LPAR), select the **Virtual Adapters** tab, and note the corresponding virtual SCSI adapters in VIOS. In our setup, the virtual adapters IDs 15 and 16 in the IBM i client, are connected to the virtual adapter IDs 16 and 19 in corresponding VIOS, as shown in Figure 3-5.

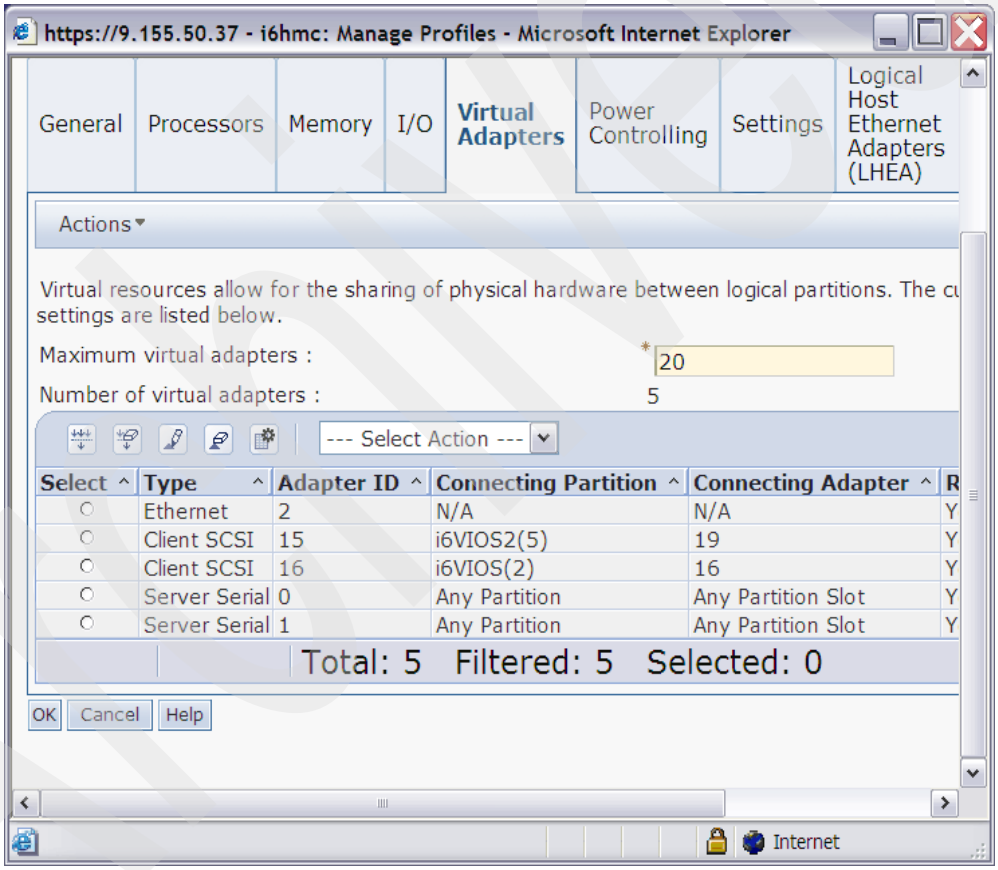

Figure 3-5 Virtual adapters in IBM i

   ii. In VIOS, look for the device name of the virtual adapter connected to the IBM i client. In our example, we use the **`lsmap -all`** command to display the virtual adapters. We are looking for the adapter ID 16, which is the one in virtual slot 16. As shown in Figure 3-6 on page 14, the virtual adapter we are looking for has the device name `vhost1`.

```
$ lsmap -all
SVSA            Physloc                                     Client Partition
ID
--------------- ------------------------------------------- ------------------
vhost0          U9406.MMA.655A620-V2-C15                    0x00000000

VTD                   NO VIRTUAL TARGET DEVICE FOUND

SVSA            Physloc                                     Client Partition
ID
--------------- ------------------------------------------- ------------------
vhost1          U9406.MMA.655A620-V2-C16                    0x00000000

VTD                   NO VIRTUAL TARGET DEVICE FOUND

$
```

Figure 3-6 Virtual SCSI adapters

e. Then, map the disk devices to the SCSI virtual adapter assigned to the SCSI virtual adapter in the IBM i partition, by using the commad **mkvdev -vdev hdiskxx -vadapter vhostx**. An illustration of the command is shown in Figure 3-7.

```
$ mkvdev -vdev hdisk3 -vadapter vhost1
vtscsi1 Available
$ mkvdev -vdev hdisk4 -vadapter vhost1
vtscsi2 Available
$
```

Figure 3-7 Mapping virtual disks to the virtual SCSI adapter

After the DS8000 LUNs have been mapped to the corresponding virtual SCSI adapters in each VIOS, the LUNs are available to the IBM i client in multipath mode.

3.2.2 Implementing IBM i with virtual SCSI adapters in multipath, two VIOS partitions

Install IBM i in the client partition connected to both VIOS partitions, as is described in *IBM i and Midrange External Storage*, SG24-7668 (refer to the information about adding DS storage to IBM i configuration).

Note: Before the installation, be sure that:

- The Load source tagged I/O is pointing to the virtual SCSI adapter assigned to the VIOS adapter with disks.
- The Alternate restart device tagged I/O is pointing to the virtual SCSI adapter that is assigned to the VIOS adapter with optical drive.

In our example, we tagged the virtual adapter 16 as load source and the virtual SCSI adapter 3 as Alternate Device from which to start the installation. See Figure 3-8.

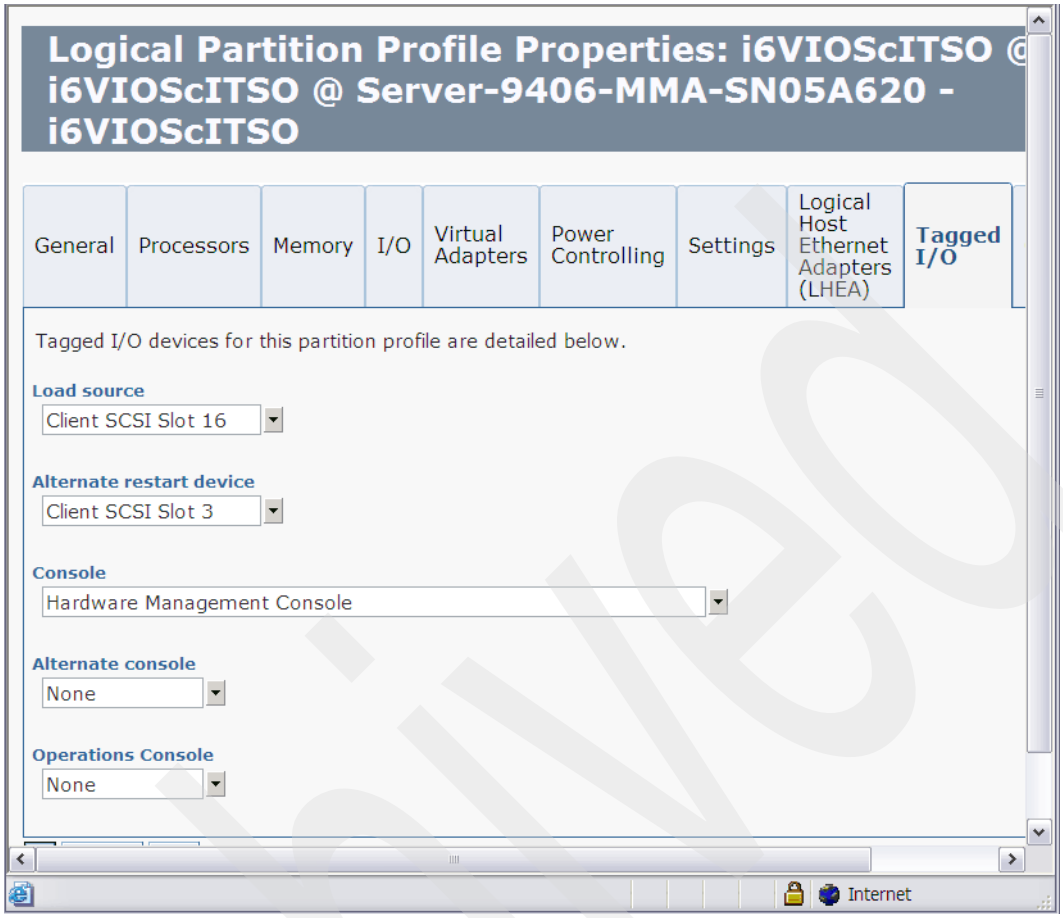

Figure 3-8 Tagged I/O adapters

IBM i multipath is automatically started on the DS8000 LUNs that are connected through two VIOS partitions. Therefore, the LUNs are shown as DMPxxx devices in the System Service Tools (SST), as shown in Figure 3-9.

```
 Display Disk Configuration Status

            Serial                        Resource                    Hot Spare
 ASP Unit   Number          Type Model    Name       Status           Protection
   1                                                 Unprotected
         1  YQYN8QD87B9Z    6B22  050     DMP002     Configured           N
         2  YGGGN3CTXXLM    6B22  050     DMP004     Configured           N
         3  YGS2CQNHUD9D    6B22  050     DMP014     Configured           N
         4  Y8HXRX2T29RZ    6B22  050     DMP008     Configured           N
         5  Y3NUTB5NPNTZ    6B22  050     DMP010     Configured           N
         6  YGW9MWEUWAFV    6B22  050     DMP006     Configured           N
         7  YVGUBXW98SU3    6B22  050     DMP015     Configured           N
         8  Y7Y59T8GRTTJ    6B22  050     DMP012     Configured           N

Press Enter to continue.

 F3=Exit        F5=Refresh         F9=Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 3-9 DS8000 LUNs connected through two VIOS partitions

3.3 Multipath through two VIOS partitions with NPIV

In this section, we discuss the connection and implementation of IBM i in multipath through two VIOS partitions with NPIV partitions.

3.3.1 Connecting with two VIOS partitions for multipath using NPIV

To connect DS8000 LUNs to an IBM i client through two VIOS partitions with NPIV, perform the steps in the following sections:

- Step 1: Log in and enable NPIV capability
- Step 2: Perform steps in each VIOS
- Step 3: Create DS8000 volumes, volume groups, host connections

Step 1: Log in and enable NPIV capability

Follow these steps:

1. Log in to the switch (connecting the FC adapters in VIOS and the DS8000 HBA ports), by using Telnet. For each port issue the **portCfgNPIVPort <PortNumber> <Mode>** command, specifying the relevant port number and mode **1**. This specification enables the NPIV capability on the port.
2. After you enable the NPIV capability, you might want to check that the ports are NPIV-enabled by executing the **portcfgshow** command. The NPIV capability on each port should indicate ON, as shown in Figure 3-10.

```
IBM_B64_SLE04_1:admin> portcfgshow
Ports of Slot 0    0  1  2  3    4  5  6  7    8  9 10 11   12 13 14 15
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed            AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port       .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Long Distance    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port  .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed  .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent Disable.. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
NPIV capability  ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Mirror Port      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
```

Figure 3-10 NPIV capability on switch ports

Step 2: Perform steps in each VIOS

Perform the following steps in *each* VIOS to which the client IBM i partition will be connected:

1. Create a server Fibre Channel adapter in VIOS. If VIOS is already installed in a partition of the Power system, add the virtual FC adapter as follows:
   a. On the HMC, select the managed server to be configured by **Systems Management** → **Servers** → **<servername>**.
   b. Select the VIOS partition for which the virtual FC server adapter is to be configured. From popup menu for the selected VIOS partition, select **Configuration** → **Manage Profiles**, as is shown on Figure 3-11 on page 17.

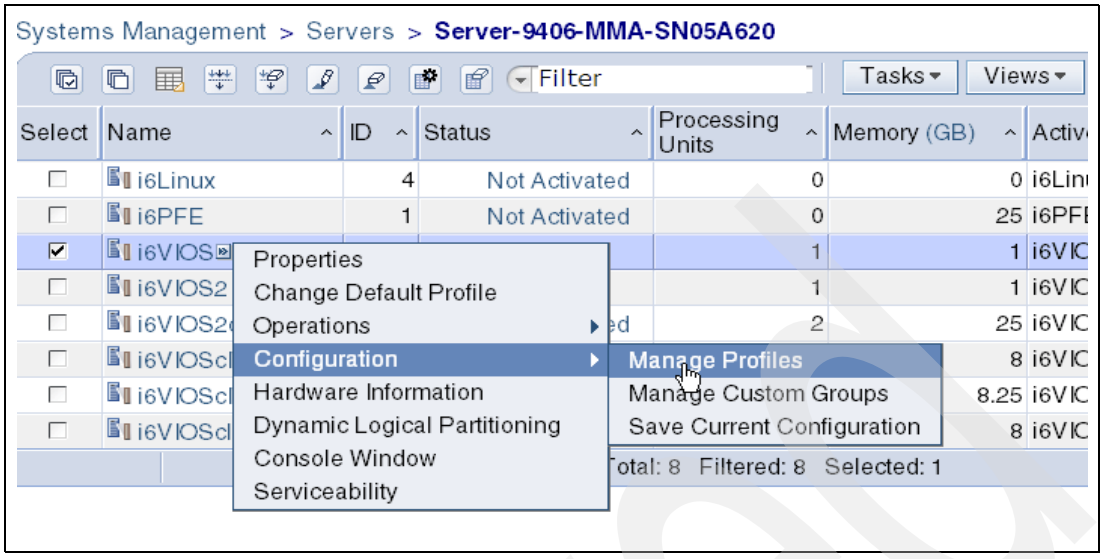

Figure 3-11 HMC: manage server profile

c. In the Managed Profiles window, select the VIOS partition, expand the Actions pull-down menu and select **Edit**, as shown in Figure 3-12.

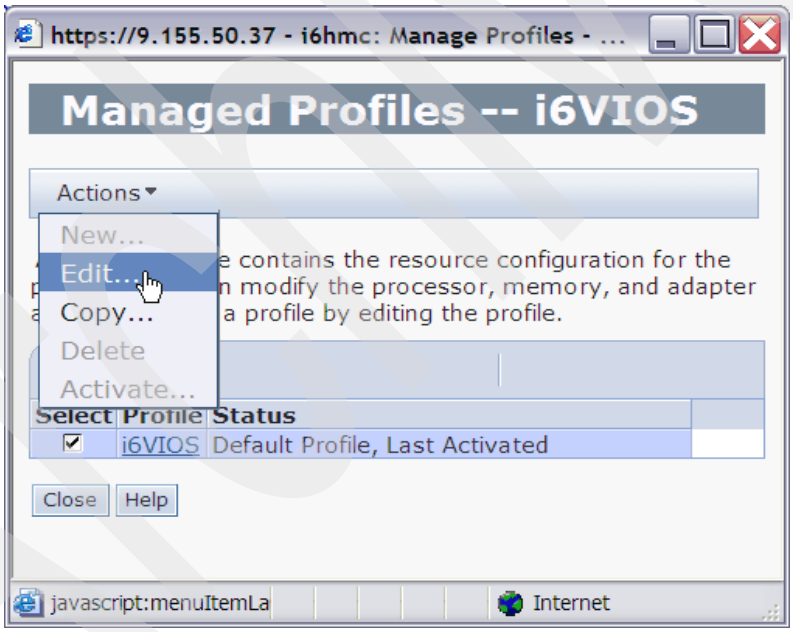

Figure 3-12 Managed Profiles: Edit

d. In the Logical Partition Profile Properties window, select the **Virtual Adapters** tab, expand the Actions pull-down, and select **Create** → **Fibre Channel Adapter**, as shown in Figure 3-13 on page 18.

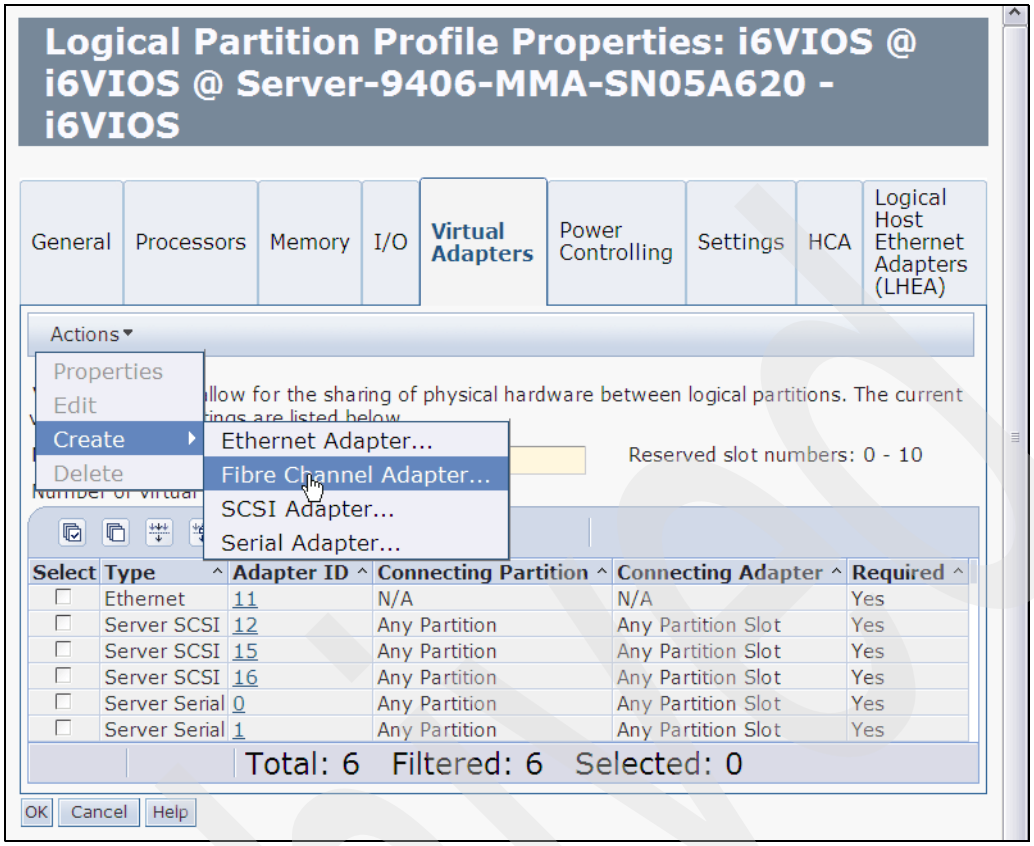

Figure 3-13 Create virtual Server FC adapter

e. In the Create Virtual Fibre Channel Adapter dialog (shown in Figure 3-14), specify the virtual slot of the FC adapter (you may keep the default value). Check **This adapter is required for partition activation**, select the client partition to which the adapter will connect, and select the virtual slot of the connecting client adapter (you may keep the default value). Then, click **OK**.

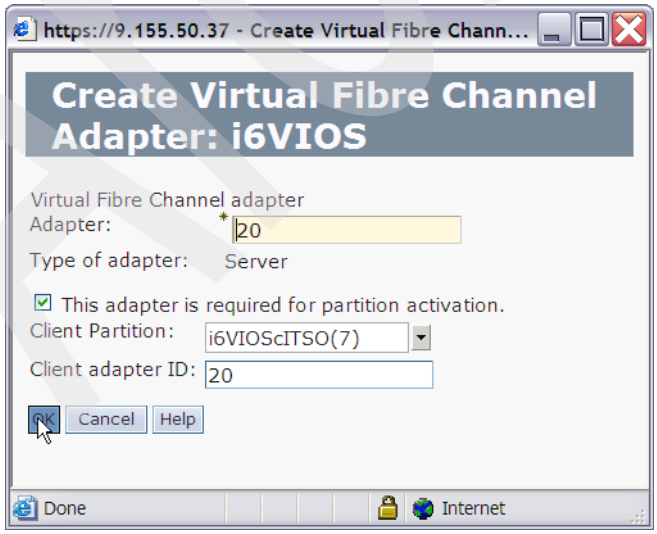

Figure 3-14 Specify values for virtual Server FC adapter

f. In the Logical Partition Profile Properties window, click **OK**, to save the profile.

Note: The changed partition profile will be activated after you power down and restart the partition.

If you are configuring a new VIOS partition, create the virtual FC adapter while creating the partition. For more information about creating a VIOS partition in a Power6 server, refer to *IBM i and Midrange External Storage*, SG24-7668.

2. Create a client FC adapter in an IBM i client partition.

   If the IBM i client is already installed in a partition of the Power system, add the virtual FC adapter by performing the following actions:

   a. In the HMC, select the managed server to be configured with **Systems Management** → **Servers** → **<servername>**.

   b. Select the IBM i partition on which the virtual FC client adapter is to be configured. From pop-up menu of the selected VIOS partition, select **Configuration** → **Manage Profiles**, as shown in Figure 3-15.

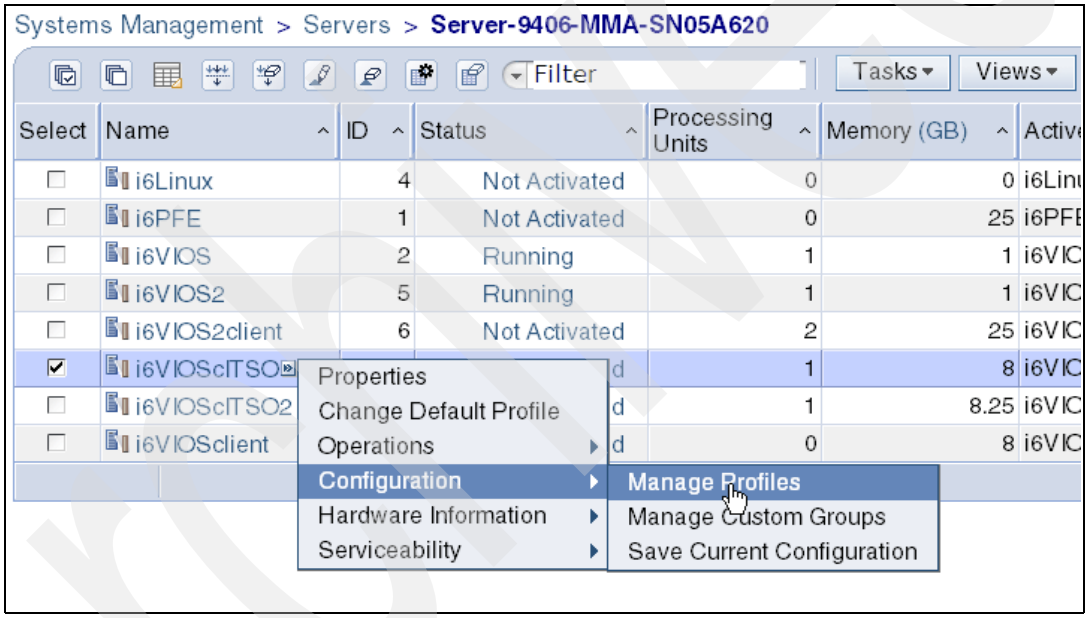

Figure 3-15 .HMC - manage client profile

   c. In the Managed Profiles window, select the VIOS partition, expand the Actions menu and click **Edit**, as shown Figure 3-12.

   d. In the Logical Partition Profile Properties window, navigate to the Virtual Adapters tab and expand the Actions menu. Select **Create** → **Fibre Channel Adapter**, as shown in Figure 3-13.

   e. In the Create Virtual Fibre Channel Adapter window, specify the virtual slot of the FC adapter. Check **This adapter is required for partition activation**, select the server partition to which the adapter will connect, and the virtual slot of the connecting client adapter. Click **OK**. See figure Figure 3-16 on page 20.

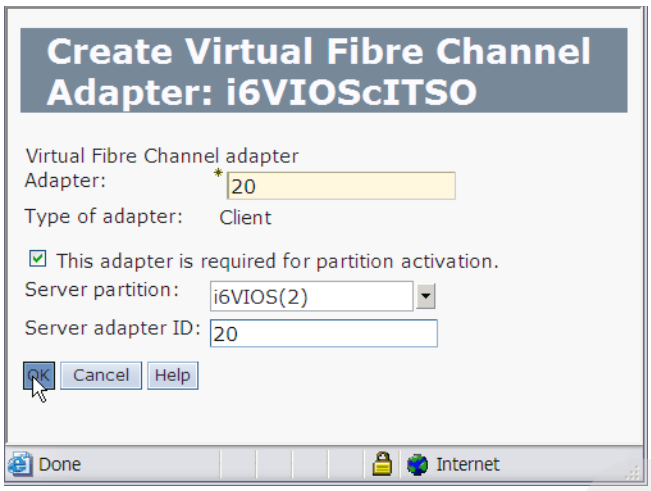

Figure 3-16 Specify values for virtual client FC adapter

f. In the Logical Partition Profile Properties window, click **OK** to save the profile.

> **Note:** The changed partition profile will be activated after you power down and restart the partition.

If you are configuring a new IBM i client partition, create the virtual FC adapter while creating the partition. For more information about creating an IBM i client partition in a Power6 server refer to *IBM i and Midrange External Storage*, SG24-7668.

3. Log on to the Virtual I/O Server partition as user *padmin*.
4. Run the **cfgdev** command to get the virtual Fibre Channel server adapter (or adapters) configured. Check the virtual FC adapter by using the **lsdev -dev vfchost*** command, which lists all available virtual Fibre Channel server adapters in the VIOS. See Figure 3-17.

```
$ lsdev -dev vfchost*
name             status      description
vfchost0         Available   Virtual FC Server Adapter
$
```

Figure 3-17 Virtual FC adapters in VIOS

Use the **lsdev -dev fcs*** command to list all available physical Fibre Channel server adapters in the Virtual I/O Server partition, as shown on Figure 3-18.

```
$ lsdev -dev fcs*
name             status      description
fcs0             Available   4Gb FC PCI Express Adapter (df1000fe)
fcs1             Available   4Gb FC PCI Express Adapter (df1000fe)
fcs2             Available   8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs3             Available   8Gb PCI Express Dual Port FC Adapter
(df1000f114108a03)
fcs4             Available   4Gb FC PCI Express Adapter (df1000fe)
fcs5             Available   4Gb FC PCI Express Adapter (df1000fe)
$
```

Figure 3-18 Physical FC adapters in VIOS

5. Run the **lsnports** command to check the NPIV readiness of the Fibre Channel adapter and the SAN switch. Figure 3-19 shows that the `fabric` attribute for the physical Fibre Channel adapter in slot C2 is set to `1`. This setting means that the adapter and the SAN switch are NPIV-ready. If the value is equal to `0`, then the adapter or SAN switch is not NPIV-ready and you should check the SAN switch configuration.

```
$ lsnports
name             physloc                        fabric tports aports swwpns  aww
pns
fcs2             U789D.001.DQDWXNY-P1-C2-T1          1     64     64    2048    2
048
fcs3             U789D.001.DQDWXNY-P1-C2-T2          1     64     64    2048    2
048
$
```

Figure 3-19 NPIV readiness in VIOS

6. Before mapping the virtual FC adapter to a physical adapter, get the `vfchost` name of the virtual adapter you created and the `fcs` name for the FC adapter from the previous **lsdev** commands output (see Figure 3-18).
7. Use the **vfcmap-vadapter vfchostx -fcp fcsx** command to map the virtual FC adapter to a port of the physical FC adapter. In our example, we map the virtual FC adapter `vfchost0` to the port `fcs2` of the physical FC adapter, by using **vfcmap -vadapter vfchost0 -fcp fcs2** command.

   To list the mappings, use the **lsmap -vadapter vfchost0 -npiv** command, as shown in Figure 3-20 on page 22. In our example, the virtual adapter `vfchost0` is mapped to the physical port `fcs2`.

```
$ lsmap -npiv -all
Name          Physloc                          ClntID ClntName       ClntOS
------------- -------------------------------- ------ -------------- -------
vfchost0      U9406.MMA.655A620-V2-C20               7

Status:NOT_LOGGED_IN
FC name:fcs2                    FC loc code:U789D.001.DQDWXNY-P1-C2-T1
Ports logged in:0
Flags:4<NOT_LOGGED>
VFC client name:                VFC client DRC:

$
```

Figure 3-20 List virtual to physical FC adapter mapping

If needed, create additional virtual server and client FC adapters.

8. Obtain the information about the worldwide port name (WWPN) of the virtual Fibre Channel client adapter created in the virtual I/O client partition. That information will be used for creating zones in the switch. To obtain the WWPN, select the appropriate virtual I/O client partition, then select **Task** → **Properties**. Expand the Virtual Adapters tab, select the Client Fibre Channel client adapter, and then select **Actions** → **Properties** to list the properties of the virtual Fibre Channel client adapter, as shown on Figure 3-21.

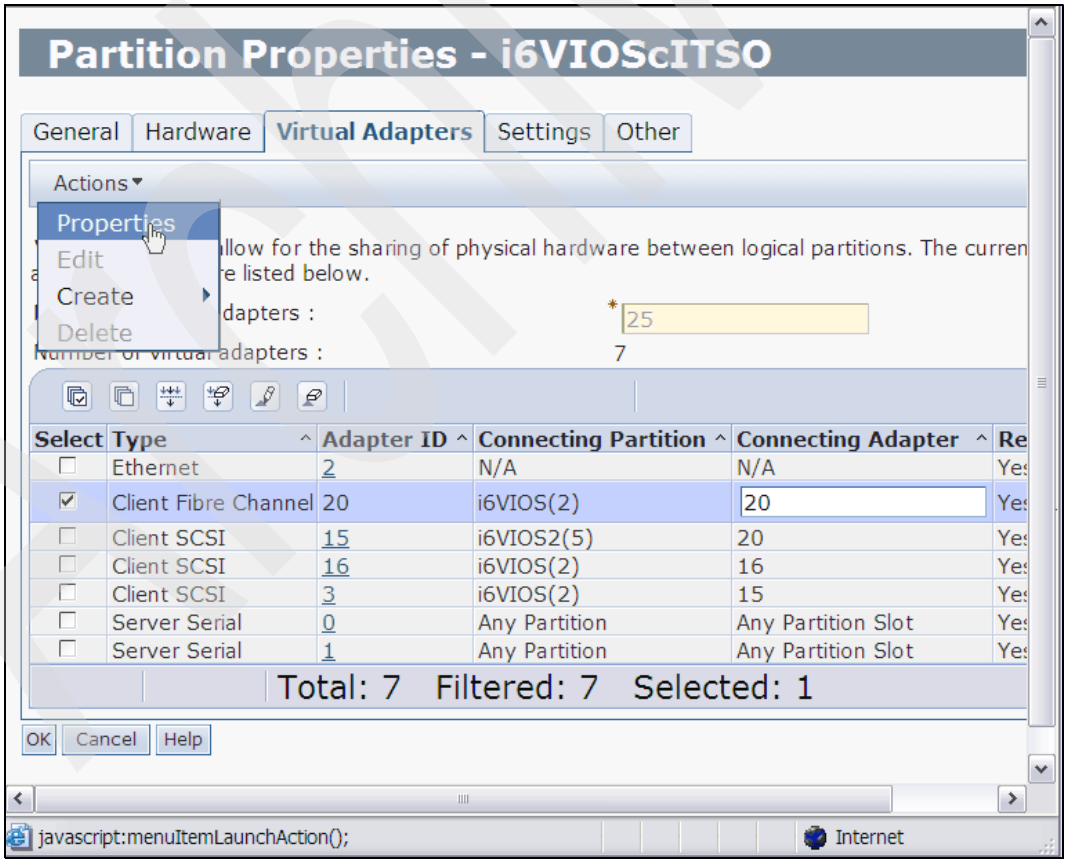

Figure 3-21 Virtual client FC adapter: getting properties

Figure 3-22 shows the properties of the virtual Fibre Channel client adapter. Here you can get the WWPN that is required for the zoning.

> **Note:** To view the newly created adapter, you might have to activate the IBM i client partition. Even if IBM i is not yet installed in the partition, you can start it until you get the message System Reference Code (SRC) B2004158, which indicates missing load source.

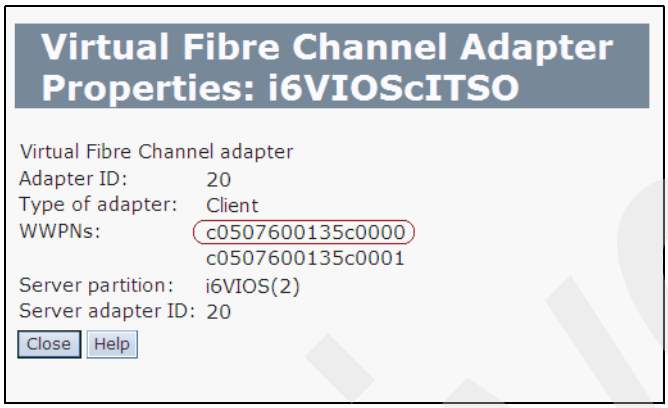

Figure 3-22 WWPNs of virtual client FC adapter

9. Log on to your SAN switch and create a new zoning, or customize an existing one. Create a zone containing the first WWPN of the virtual client FC adapter (WWPN is highlighted in Figure 3-22), and the WWPN of the DS8000 port. An example of such zoning for an IBM SAN64B-2 switch, is shown in Figure 3-23. After the zone is created, save and enable the new zone configuration.

> **Note:** The WWPN of the port for the physical FC adapter in VIOS should not be in the zone. The WWPN of the virtual client FC adapter might not log in to the switch until you IPL the IBM i partition in the Dedicated Service Tools (DST) menu.

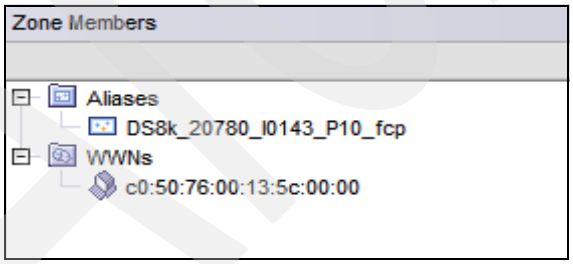

Figure 3-23 Zone with virtual client FC adapter

Step 3: Create DS8000 volumes, volume groups, host connections

Create the DS8000 volumes, volume groups, and host connections as if this were a DS8000 attached natively to the IBM i partition. For the host connection, specify the WWPN of the virtual FC adapter. For more information about configuring the DS8000 for System i attachment, refer to *IBM i and IBM System Storage: A Guide to Implementing External Disk on IBM i*, SG24-7120. An example of creating the configuration with DS CLI is shown in Figure 3-24 on page 24.

Note: When connecting DS8000 LUNs to an IBM i client through two VIOS partitions, you create one set of LUNs, and one set of volume groups for the LUNs. Then, you create two host connections, one for each volume group.

```
dscli> mkfbvol -extpool p4 -os400 A04 -name jj_ITSO_pri_#h 4800-4803
Date/Time: August 18, 2009 8:37:49 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM .2107-7520781
CMUC00025I mkfbvol: FB volume 4800 successfully created.
CMUC00025I mkfbvol: FB volume 4801 successfully created.
CMUC00025I mkfbvol: FB volume 4802 successfully created.
CMUC00025I mkfbvol: FB volume 4803 successfully created.
dscli> mkfbvol -extpool p5 -os400 A04 -name jj_ITSO_pri_#h 4900-4902
Date/Time: August 18, 2009 8:40:29 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM .2107-7520781
CMUC00025I mkfbvol: FB volume 4900 successfully created.
CMUC00025I mkfbvol: FB volume 4901 successfully created.
CMUC00025I mkfbvol: FB volume 4902 successfully created.
dscli> mkvolgrp -type os400mask -volume 4800-4802,4900-4902 jj_ITSO-pri
Date/Time: August 18, 2009 8:42:35 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM .2107-7520781
CMUC00030I mkvolgrp: Volume group V17 successfully created.
dscli> mkhostconnect -wwname C05076001350000 -hosttype iSeries -volgrp V17
Vrtual_FC_pri
Date/Time: August 18, 2009 8:45:58 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM .2107-7520781
CMUC00012I mkhostconnect: Host connection 0022 successfully created.
```

Figure 3-24 Configuring DS8000 LUns for IBM i with VIOS NPIV

3.3.2 Implementing IBM i in multipath through two VIOS partitions with NPIV

Before starting the installation, make sure that the virtual client FC adapter assigned to one of the VIOS, is tagged as the Load source in the HMC. The virtual adapter to which the virtual DVD drive is assigned, should be tagged as Alternate IPL device. The IPL mode should be D-mode and Manual. Proceed with the installation as described in *IBM i and Midrange External Storage*, SG24-7668. An illustration of tagged devices in HMC is shown in Figure 3-25 on page 25.

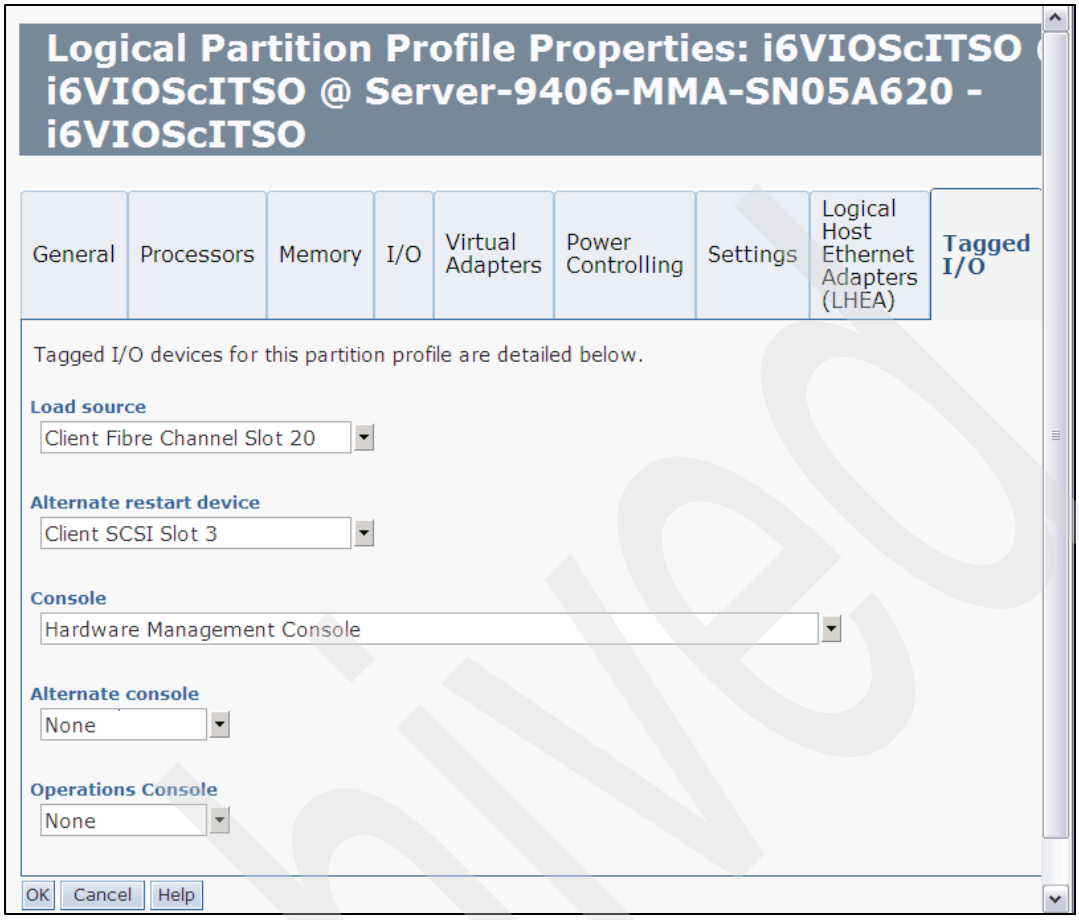

Figure 3-25 Tagged virtual adapters in HMC

Figure 3-26 shows the DST Display Disk Configuration Status window for DS8000 LUNs that are connected in multipath through two VIOS partitions in NPIV.

```
Display Disk Configuration Status

                Serial                          Resource                        Hot Spare
ASP Unit        Number          Type Model      Name        Status              Protection
  1                                             Unprotected
      1         50-4902781      2107  A04       DMP011      RAID 5/Active           N
      2         50-4800781      2107  A04       DMP001      RAID 5/Active           N
      3         50-4801781      2107  A04       DMP003      RAID 5/Active           N
      4         50-4802781      2107  A04       DMP005      RAID 5/Active           N
      5         50-4900781      2107  A04       DMP007      RAID 5/Active           N
      6         50-4901781      2107  A04       DMP009      RAID 5/Active           N

Press Enter to continue.

 F3=Exit        F5=Refresh         F9=Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 3-26 NPIV connected LUNs in Multipath

4

FlashCopy scenarios

This chapter includes a brief overview of FlashCopy operations and describes several usage scenarios in the context of IBM i as a VIOS client.

4.1 FlashCopy overview

FlashCopy creates a point-in-time copy of the data. When a FlashCopy operation is invoked, the process of establishing a FlashCopy source and target volume pair, and creating the necessary control bitmaps is completed in only a few seconds. Thereafter, you have access to a point-in-time copy of the source volume, as though all the data had been copied. As soon as the pair has been established, you can read and write to both the source and target volumes.

Two variations of FlashCopy are available:

- Standard FlashCopy uses a normal volume as target volume. This target volume must be the same size (or larger) as the source volume and that space specified for the target volume is allocated in the storage subsystem.
- FlashCopy SE function uses track space-efficient volumes as FlashCopy target volumes. A track space-efficient target volume has a virtual size that is equal to or greater than the source volume size. However, space is not allocated for this volume at the time the volume is created and the FlashCopy is initiated. Original tracks are copied to the TSE volume when eventually modified on the source volumes. Space in the repository is allocated for only these tracks or for any write operation to the target itself.

Note: Both FlashCopy and FlashCopy SE can coexist on a DS8000.

Figure 4-1 shows these FlashCopy basic concepts.

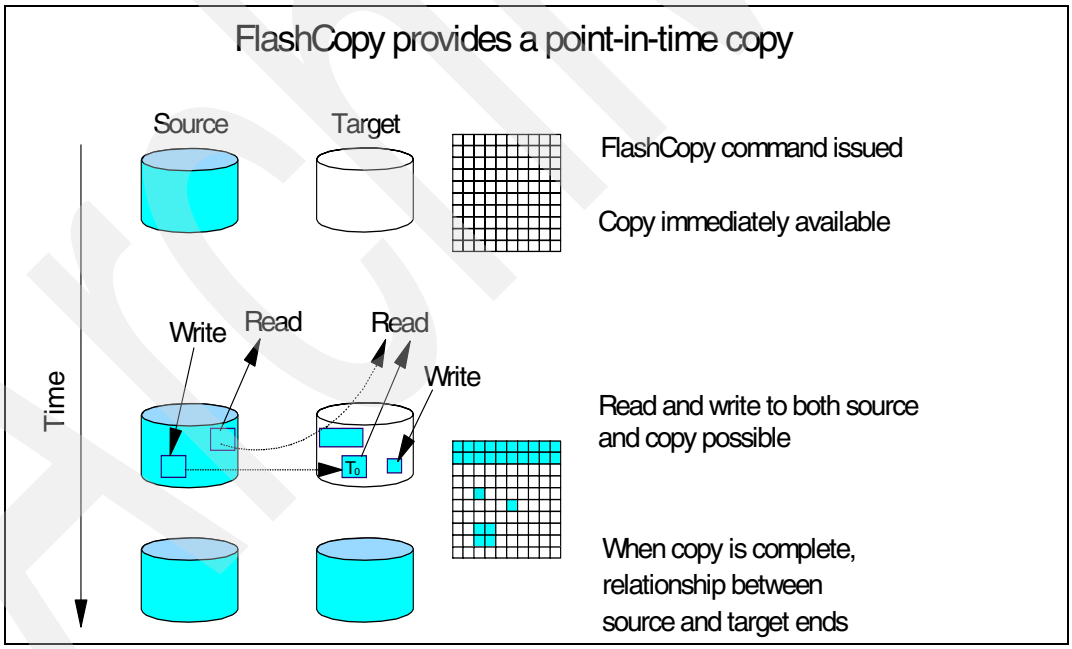

Figure 4-1 FlashCopy concepts

If you access the source or the target volumes during the background copy, standard FlashCopy manages these I/O requests as follows:

- Read from the source volume

  When a read request goes to the source volume, the request is read from the source volume.

- Read from the target volume

  When a read request goes to the target volume, FlashCopy checks the bitmap, and:

  - If the point-in-time data was already copied to the target volume, it is read from the target volume.
  - If the point-in-time data has not been copied yet, it is read from the source volume.

- Write to the source volume

  When a write request goes to the source volume, first the data is written to the cache and persistent memory (write cache). Then, when the update is destaged to the source volume, FlashCopy checks the bitmap and:

  - If the point-in-time data was already copied to the target, then the update is written to the source volume.
  - If the point-in-time data has not been copied yet to the target, then first it is copied to the target volume, and after that the update is written to the source volume.

- Write to the target volume

  When data is written to the target volume while the FlashCopy relationship exists, the storage subsystem ensures that the bitmap is updated. This way, the point-in-time data from the source volume never overwrites updates that were done directly to the target.

The background copy can have a slight impact to your application performance because the physical copy needs some storage resources, but the impact is minimal because the host I/O has priority over the background copy. And if you want, you can issue standard FlashCopy with the *no background copy* option.

No background copy option

If you invoke standard FlashCopy with the no background copy option or FlashCopy SE, the FlashCopy relationship is established without initiating a background copy. Therefore, you can minimize the impact of the background copy. When the DS8000 receives an update to a source track in a FlashCopy relationship, a copy of the point-in-time data is copied to the target volume so that it is available when the data from the target volume is accessed. This option is useful when you do not need to issue FlashCopy in the opposite direction.

FlashCopy consistency group

FlashCopy consistency group allows you to freeze (temporarily queue) I/O activity to a LUN or volume. FlashCopy consistency group helps you to create a consistent point-in-time copy across multiple LUNs or volumes, and even across multiple storage units.

Important: FlashCopy consistency group can create host-based consistent copies; they are not application-based consistent copies. The copies have *power-fail* or *crash* level consistency. This means that if you suddenly power off your server without stopping your applications and without destaging the data in the file cache, the data in the file cache can be lost and you might need recovery procedures to restart your applications. To start your system with FlashCopy consistency group target volumes, you might need the same operations as the crash recovery.

For example, if the consistency group source volumes are used with a journaled file system (such as AIX JFS) and the source LUNs are not unmounted before running FlashCopy, the `fsck` command will likely have to be run on the target volumes.

4.2 IBM i and FlashCopy

When considering to use DS8000 FlashCopy for with BM i, you should be aware of the System i architecture and the way it affects DS8000 Copy Services based solutions. In this section, we describe two important considerations: quiescing data to disk, and cloning an IBM i partition.

4.2.1 Quiescing IBM i data to disk

Because of its single-level storage architecture, IBM i keeps data in main memory until data is written to disk as a result of page-swaps. Before initiating a FlashCopy in IBM i, you must ensure that the data in main memory has been flushed to disk. This step is required to guarantee a consistent and current copy.

Prior to IBM i release V6R1, forcing an IBM i *restricted condition* was the only way to quiesce data to disk. Restricted condition is caused by ending all IBM i subsystems, which implies ending all jobs in the subsystems. When ending a particular job, the data belonging to that job is flushed from memory to disk. In other words, ending all the jobs guarantees that data is quieseced to disk.

IBM i release V6R1 provides a function to write all pending changes to disk and suspend database activity within an auxiliary storage pool (ASP). The database activity remains suspended until a Resume is issued. This is known as quiescing the ASP. The function is initiated by the CHGASPACT command. When running the command, specify the following values:

- The ASP device on which the quiescing will be done. The device can be Sysbas or an independent ASP.
- Option of whether to suspend or resume the ASP device
- Suspend time-out. This value specifies how many seconds the system waits for committed transactions to end their commit-cycle.
- Time-out action. Specifies whether to end or to continue the flushing of data to disk, after the time-out is finished.

An illustration of the CHGASPACT command is shown in Figure 4-2.

```
Change ASP Activity (CHGASPACT)

Type choices, press Enter.

ASP device . . . . . > *SYSBAS       Name, *SYSBAS
Option . . . . . > *SUSPEND     *SUSPEND, *RESUME, *FRCWRT
Suspend timeout  . . . . . > 30           Number
Suspend timeout action . . . . . > *END         *CONT, *END




                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 4-2 CHGASPACT command to quiesce data

When executing the command CHGASPACT *SUSPEND, the following events occur:

1. As much as possible, data is flushed to disk without distinguishing among database and non-database files, transaction, and non-transaction activity.
2. The system waits for the specified time-out to get all current transactions to their next commit-boundary and does not let them continue past that commit-boundary. When CHGASPACT *SUSPEND command is initiated, Commitment Control is informed not to allow any new transactions to start. It maintains a count of open transactions, which are being monitored. As transactions get committed, that count is decremented. When the count reaches zero, no open transactions exist, and the suspend processing continues as described in step 5, even if the time-out has not yet expired.
3. If a time-out occurs (the amount of time to wait has passed and the transaction count is not zero) and the time-out action *END is specified, the transactions are released. A message is sent indicating that the CHGASPACT *SUSPEND failed, and the CHGASPACT is marked as *Failed*.
4. If a time-out occurs and the time-out action *CONT is specified, a diagnostic message is sent indicating that not all transactions were successfully suspended, and processing continues as is described in step 5.
5. If all transactions were successfully suspended within the time limit, or time-out occurred and the time-out action *CONT is specified, the system proceeds with following actions:
   a. Non-transaction- based operations are suspended
   b. Data that is non-pinned in the memory is flushed to disk, in order to write any changes that occurred while the suspend operation was occurring.

When we want the transactions to resume, for example after performing FlashCopy, we issue the following command:

```
CHGASPACT *RESUME
```

4.2.2 Cloning an IBM i partition

Because cloning creates a copy of the whole source system, consider the following information when you create a clone:

- A clone is an exact copy of the original source system in every aspect.
- The system name and network attributes are identical.
- The TCP/IP settings are identical.
- The BRMS network information is identical.
- The NetServer settings are identical.
- User profiles and passwords are identical.
- The job schedule entries are identical.
- Relational database entries are identical.

Attention: Be very careful when you activate a clone. In particular, you have to ensure that it does not connect to the network automatically, because doing so can cause substantial problems within both the clone and its parent system.

Imagine that you are in the process of creating a clone and that your network has a problem with a router. Your network is effectively split in two. You finish your clone and connect it to a new partition ready for use. When you IPL your clone, it might see itself plugged into the network and working correctly. The job scheduler starts updating the external systems that it can see. While this is happening, your live production system is updating those other systems that it can see. The result can be catastrophic.

Important: Do not attach a clone to your network until you have resolved any potential conflicts that the clone has with the parent system.

Ensure that your clone system is customized properly before you attach it to a network.

Important: Although cloning is a highly effective means of backing up a system for disaster recovery, always remember that backing up all objects on the clone does not make sense unless the backup is as part of a full backup for disaster recovery. In particular, if you bring journals and associated receivers or system logs back from the clone to the production system, the data content will not be relevant, because the systems would have a different *data history* reflected in the journals. This inconsistency leads to unpredictable results.

Restriction: You must not attach any clone LUNs to the original parent system unless they have been used in another partition first or deleted and re-created, that is re-formatted, within the DS8000. Failure to observe this restriction can have unpredictable results and could lead to loss of data.

This restriction is because the clone LUNs are perfect copies of LUNs that are on the parent system, and as such, the system would not be able to tell the difference between the original and the clone if they were attached to the same system.

As soon as you use LUNS from the clone in a separate partition, they become owned by that partition, and can then be safely reused in the original partition.

4.3 Testing environment for FlashCopy

On the Power 6 system in our scenario, we set two VIOS partitions, an IBM i production partition, and IBM i backup partition in which an IPL from FlashCopy targets is performed.

In each VIOS, we create two virtual FC adapters, both of them are mapped to one port of the physical FC adapter. The ports connect to the DS8000 through an NPIV enabled switch. Each production and backup partition is connected to each VIOS over one virtual FC adapter. We assign the set of DS8000 production LUNs to both virtual FC adapters in the IBM i production partition. This means that the IBM i establishes two paths to the set of LUNs. Similarly, FlashCopy LUNs are assigned to the backup partition through two virtual FC adapters, so that we also have multipathing enabled in the backup partition.

Our scenario is depicted in Figure 4-3 on page 33. The solid black lines indicate physical connections, the dashed black lines indicate virtual adapters connections, and the blue lines indicate LUN assignments to virtual FC adapters. The LUNs assignment is shown for the production IBM i only; The backup IBM i uses a similar assignment scheme (two virtual FC adapters through the two VIOS partitions).

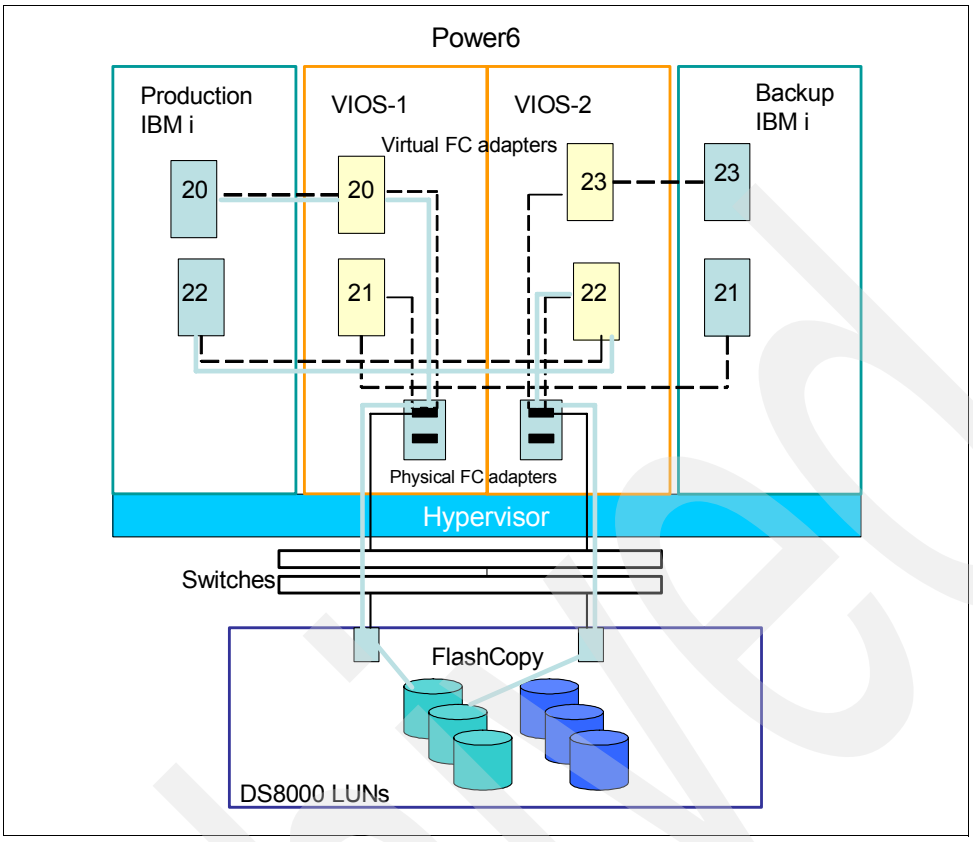

Figure 4-3 Testing environment for FlashCopy with VIOS NPIV

4.4 FlashCopy with production IBM i in restricted condition

In this scenario, we end all subsystems in the production IBM i partition to make sure that all data in main memory is written to disk. Then, we establish a FlashCopy relationship for the DS8000 LUNs in the production partition. As soon as the FlashCopy is done, we restart subsystems in the IBM i production partition. Next, we IPL the backup IBM i partition from the FlashCopy target LUNs. After successful IPL, the backup partition contains a clone of the production IBM i partition, which we use to save the production libraries and other objects to tape, or for testing, or queries and so on.

Usually, we establish a FlashCopy relationship without background copying (*nocopy*) for the backups to tape; for queries and development, we establish FlashCopy with background copying. In our example, we use FlashCopy with the nocopy option.

4.4.1 Steps to perform FlashCopy with restricted condition

Set up each production IBM i partition and backup partition to connect to the DS8000 LUNs with VIOS NPIV in multipathing mode, as is described in 3.3.2, “Implementing IBM i in multipath through two VIOS partitions with NPIV” on page 24.

In our example, we use the connection scheme shown in Figure 4-3.

Follow these steps:

1. In the production IBM i partition, execute the ENDSBS SBS(*ALL) command to end the IBM i subsystems. You may specify *IMMED option to force the subsystems termination, or use the *CNTRLD option with a specified time delay to allow subsystems to complete jobs. After the subsystems are ended the IBM i issues the message: *System ended to restricted condition.*
2. Perform a FlashCopy of the production IBM i LUNs. In our example, we use the DS CLI command **mkflash**. As shown in Figure 4-4, we use FlashCopy consistency groups (option **-freeze**), and we do not use background copying (option **-nocp**)

```
dscli> mkflash -nocp -freeze 4800-4802:4803-4805 4900-4902:4903-4905
Date/Time: August 19, 2009 9:09:06 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM
.2107-7520781
CMUC00137I mkflash: FlashCopy pair 4800:4803 successfully created.
CMUC00137I mkflash: FlashCopy pair 4801:4804 successfully created.
CMUC00137I mkflash: FlashCopy pair 4802:4805 successfully created.
CMUC00137I mkflash: FlashCopy pair 4900:4903 successfully created.
CMUC00137I mkflash: FlashCopy pair 4901:4904 successfully created.
CMUC00137I mkflash: FlashCopy pair 4902:4905 successfully created.
dscli>
```

Figure 4-4 Establish FlashCopy of Sysbas

3. In the production IBM i partition, start the controlling subsystem, which will then start the other subsystems. In our example, we start the controlling subsystem QCLT by using the STRSBS QCTL command.
4. Ensure that the backup IBM i partition tags, for the Load source, a virtual FC adapter connected through VIOS NPIV to the FlashCopy target LUNs. A virtual adapter to which a DVD drive (in VIOS) is assigned, should be tagged as alternate IPL device. In our example, shown in Figure 4-5 on page 35, we tag the virtual FC adapter in virtual slot 23 for the Load source, and the virtual SCSi adapter in virtual slot 5 for the alternate IPL device.

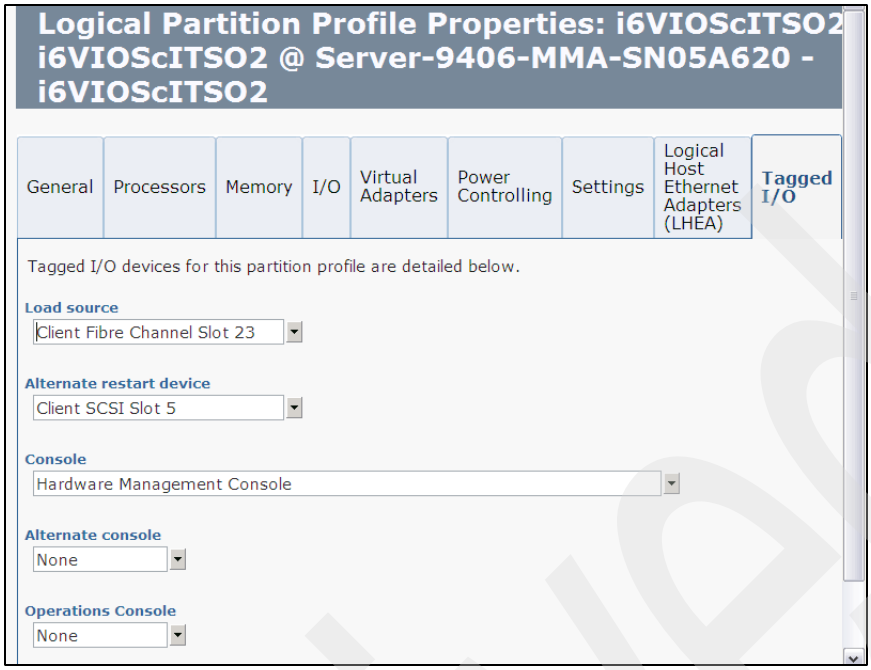

Figure 4-5 Tagged virtual adapters

Also be sure that the IPL source is set to A or B. You may want to use the keylock position *Manual* the first time you IPL form the FlashCopy target LUNs to be able to check the disk units in DST. For subsequent IPLs, use the keylock position *Normal*.

To check the IPL source and keylock position expand the drop down menu at the IBM i client partition in the HMC, and click **Properties**. In the Partition Properties window, select the **Settings** tab. If necessary change the settings to **A**, **B**, and **Normal**. Figure 4-6 on page 36 shows the IPL settings in our example. i

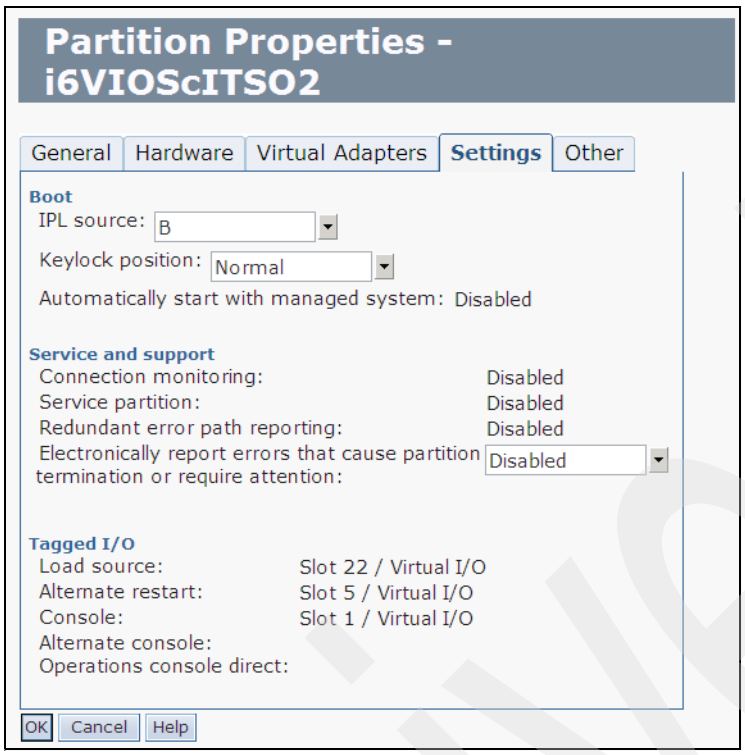

Figure 4-6 Settings for IPL from FlashCopy targets

5. To IPL the backup IBM i client, expand the pop-up menu for the client partition at the HMC, select **Operations** → **Activate**. See Figure 4-7.

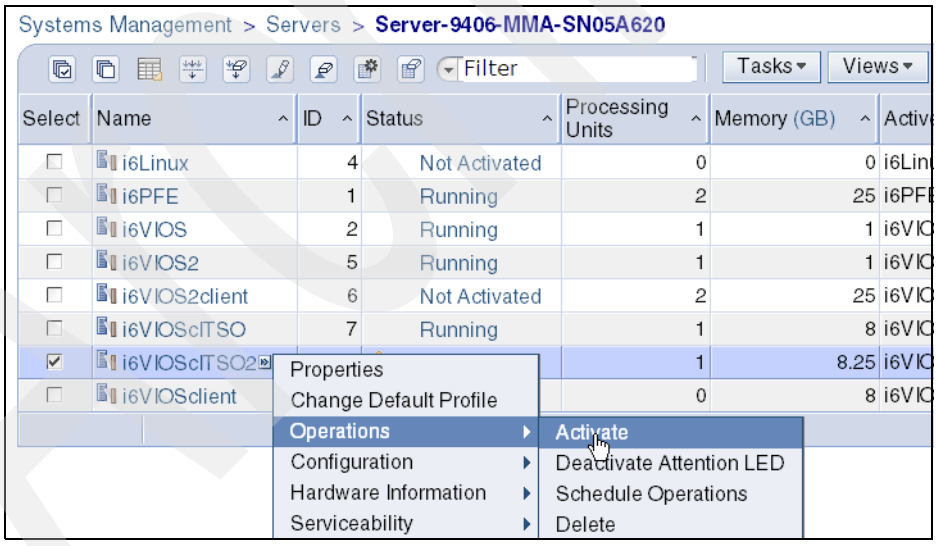

Figure 4-7 Activate backup partition

As expected, the IPL from FlashCopy targets is abnormal because the production IBM i was not shut down, but in a restricted state when the FlashCopy was established.

The backup partition now contains the clone of the production IBM i. Refer also to the considerations described in 4.2.2, “Cloning an IBM i partition” on page 31.

4.5 FlashCopy with quiescing of production Sysbas

To ensure that the data in the IBM i main memory is flushed to disk prior to the FlashCopy, you may use the IBM i Quiesce function, which is activated with the CHGASPACT command.

4.5.1 Steps to perform FlashCopy with Quiesce

Set up each production IBM i partition and back up IBM i partition to connect to the DS8000 LUNs with VIOS NPIV in multipathing mode, as described in 3.3.2, “Implementing IBM i in multipath through two VIOS partitions with NPIV” on page 24.

Follow these steps to perform the Quiesce and FlashCopy:

1. Quiesce the data in the production IBM i by using the CHGASPACT *SUSPEND command. In our example, we use a time-out value of 30 seconds, and the END time-out action. See Figure 4-8.

```
Change ASP Activity (CHGASPACT)

Type choices, press Enter.

ASP device . . . . . > *SYSBAS       Name, *SYSBAS
Option . . . . . > *SUSPEND     *SUSPEND, *RESUME, *FRCWRT
Suspend timeout  . . . . . .   30           Number
Suspend timeout action . . . . .   *end         *CONT, *END




                                                                      Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 4-8 Quiesce Sysbas in production System i

When the Quiesce is achieved, Sysbas enters the suspended state, which is indicated by the message: Access to ASP *SYSBAS is suspended.

2. Start the FlashCopy relationship for the production IBM i LUNs. Figure 4-4 on page 34 shows an example of the DS CLI command used for that purpose.

3. Resume activities in Sysbas by using the CHGASPACT *RESUME command, as is shown in Figure 4-9. The resume operation is successful is indicated by the message: Access to ASP *SYSBAS successfully resumed.

```
Change ASP Activity (CHGASPACT)

Type choices, press Enter.

ASP device . . . . . > *SYSBAS       Name, *SYSBAS
Option . . . . . > *RESUME       *SUSPEND, *RESUME, *FRCWRT




                                                                  Bottom
F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 4-9 Resume activation in Sysbas

4. IPL the backup IBM i partition from the FlashCopy targets, as described in 4.4.1, “Steps to perform FlashCopy with restricted condition” on page 33 (steps 4 and 5).

   Because the production partition was only quiesced, but not shut down prior to the FlashCopy, the IPL is abnormal.

The backup partition now contains the clone of the production IBM i. Refer also to the considerations addressed in 4.5, “FlashCopy with quiescing of production Sysbas” on page 37.

5

Metro Mirror scenarios

This chapter explains Metro Mirror scenarios for IBM i as a VIOS client. The chapter starts with a brief overview of the Metro Mirror operations.

5.1 Metro Mirror overview

Metro Mirror (previously known as synchronous Peer-to-Peer Remote Copy, or PPRC) provides real-time mirroring of logical volumes between two DS8000s that can be located up to 300 km from each other. It is a synchronous copy solution where write operations are completed on both copies (local and remote site) before they are considered to be complete. It is typically used for applications that cannot suffer any data loss in the event of a failure.

As data is synchronously transferred, the distance between the local and the remote disk subsystems will determine the effect on application response time. Figure 5-1 illustrates the sequence of a write update with Metro Mirror.

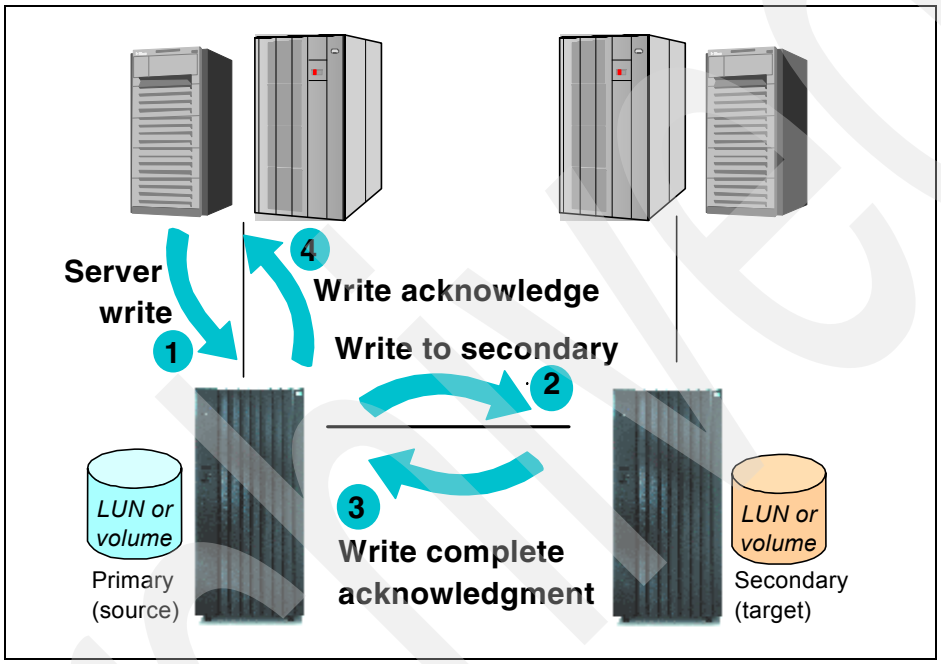

Figure 5-1 Metro Mirror

When the application performs a write update operation to a source volume, the following events occur:

1. Write to source volume (DS8000 cache and NVS)
2. Write to target volume (DS8000 cache and NVS)
3. Signal write complete from the remote target DS8000
4. Post I/O complete to host server

5.1.1 Metro Mirror volume state

Volumes participating in a Metro Mirror session can be found in any of the following states:

- Copy pending

  Volumes are in a *copy pending* state after the Metro Mirror relationship is established, but the source and target volumes are still out-of-sync. In that case, data must continue to be copied from the source to the target volume of the Metro Mirror pair. This might be the case immediately after a relationship is initially established, or re-established after being suspended. The Metro Mirror target volume is not accessible when the pair is in a copy pending state.

- Full duplex

  The *full duplex* state is the state of a volume copy pair whose members are in sync, that is, both source and target volumes contain exactly the same data. The target volume is not accessible when the pair is in full duplex.

- Suspended

  Volumes are in the *suspended* state when the source and target storage subsystems cannot communicate anymore, or when the Metro Mirror pair is suspended manually. In this state, write operations to the source volume are not mirrored onto the target volume. The target volume becomes out-of-sync. During this time, Metro Mirror keeps a bitmap record of the changed tracks in the source volume. Later, when the volumes are re-synchronized, only the tracks that were updated will be copied.

- Target copy pending

  This state indicates that the source volume is unknown or cannot be queried and the target state is copy pending.

- Target full-duplex

  This state indicates that the source volume is unknown or cannot be queried and the target state is full duplex.

- Target suspended

  This state indicates that the source volume is unknown or cannot be queried and the target state is suspended.

- Not remote copy pair

  This state indicates that the relationship is not a Metro Mirror pair.

- Invalid-state

  This state indicates that the relationship state is invalid.

5.1.2 Data consistency

To restart applications at the remote site successfully, the remote site volumes must have consistent data. In normal operation, Metro Mirror keeps data consistency at the remote site. However, in a rolling disaster type of situation, a certain mechanism is necessary to keep data consistency at the remote site.

For Metro Mirror, consistency requirements are managed through use of the *Consistency Group* option. You can specify this option when you are defining Metro Mirror paths between pairs of logical subsystems (LSS) or when you change the default LSS settings. Volumes or LUNs that are paired between two logical subsystems whose paths are defined with the Consistency Group option can be considered part of a consistency group.

Consistency is provided by means of the *extended long busy* (for IBM z/OS®) or *queue full* (for open systems) conditions. These are triggered when the DS8000 detects a condition where it cannot update the Metro Mirror target volume. The volume-pair that first detects the error goes into the queue full condition, so that it will not do any write operations, and an SNMP trap message is issued. These messages can be used as a trigger for automation purposes that will provide data consistency.

Data consistency and dependent write operations are discussed in detail in *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788 (refer to the information about Metro Mirror options and configuration).

In our tests with IBM and VIOS, we did not set up Metro Mirror consistency groups. Nevertheless, we expect that many IT centers using Metro Mirror for an IBM i system, will implement them in conjunction with TPC-R, to provide consistent recovery data in case of disaster.

5.2 Planned and unplanned outages introduction

Scenarios with Metro Mirror include planned and unplanned outages.

Planned outages

The planned outage procedures rely on two facts:

- Metro Mirror source and target volumes are in a consistent and current state.
- Both DS8000s are functional and reachable.

Unplanned outages

In contrast to the assumptions for planned outages, the situation in a disaster is more difficult:

- In an unplanned outage situation, only the DS8000 at the recovery site is functioning. The production site DS8000 might be lost or unreachable.
- In an unplanned situation, volumes at the production and recovery site might be in different states.
- In a planned situation, you can stop all I/Os at the production site to make all volumes across the recovery site reach a consistent status. This cannot be done in an unplanned situation. If not using consistency groups, for example in a power failure, you can only assume consistency at the level of a single volume-pair, not at the application level.

5.3 Environment for Metro Mirror scenarios

When implementing a Disaster Recovery solution with DS8000 Metro Mirror, the Disaster Recovery (DR) IBM i partition typically resides on a remote Power6 system. Each production IBM i and DR IBM i is connected to the DS8000 through two VIOS partitions in NPIV: the relevant IBM i partition and the two VIOS partitions residing in the same Power6.

In our test scenario, we used the same Power6 system for both the production and DR IBM i partition (in a real environment, you would need a separate Power6 system at each site). The setup is as follows:

- Two VIOS partitions
- IBM i production partition
- IBM i Disaster Recovery partition

In each VIOS partition, we create two virtual FC adapters, both of them mapped to one specific port of the physical FC adapter. These ports are connected to both the primary and secondary DS8000 through NPIV enabled switches. Each of the IBM i production and backup partitions is connected to each VIOS through one virtual FC adapter. We assign the set of LUNs on the primary DS8000 to both virtual FC adapters in the production IBM i partition. In doing so, the IBM i has two paths (multipath) to the set of LUNs. Similarly, we assign the Metro Mirror target LUNs in the secondary DS8000 to the DR partition through two virtual FC adapters, each through one VIOS, so that multipath is enabled in the DR partition too.

Our environment for this scenario is shown in Figure 5-2. The solid black lines indicate physical connections, the dashed black lines indicate connection to virtual adapters, and the blue lines indicate LUNs assignment to virtual FC adapters. The assignment of LUNs is shown for the production IBM i only; the DR IBM i uses a similar assignment scheme to two virtual FC adapters through the two VIOS partitions.

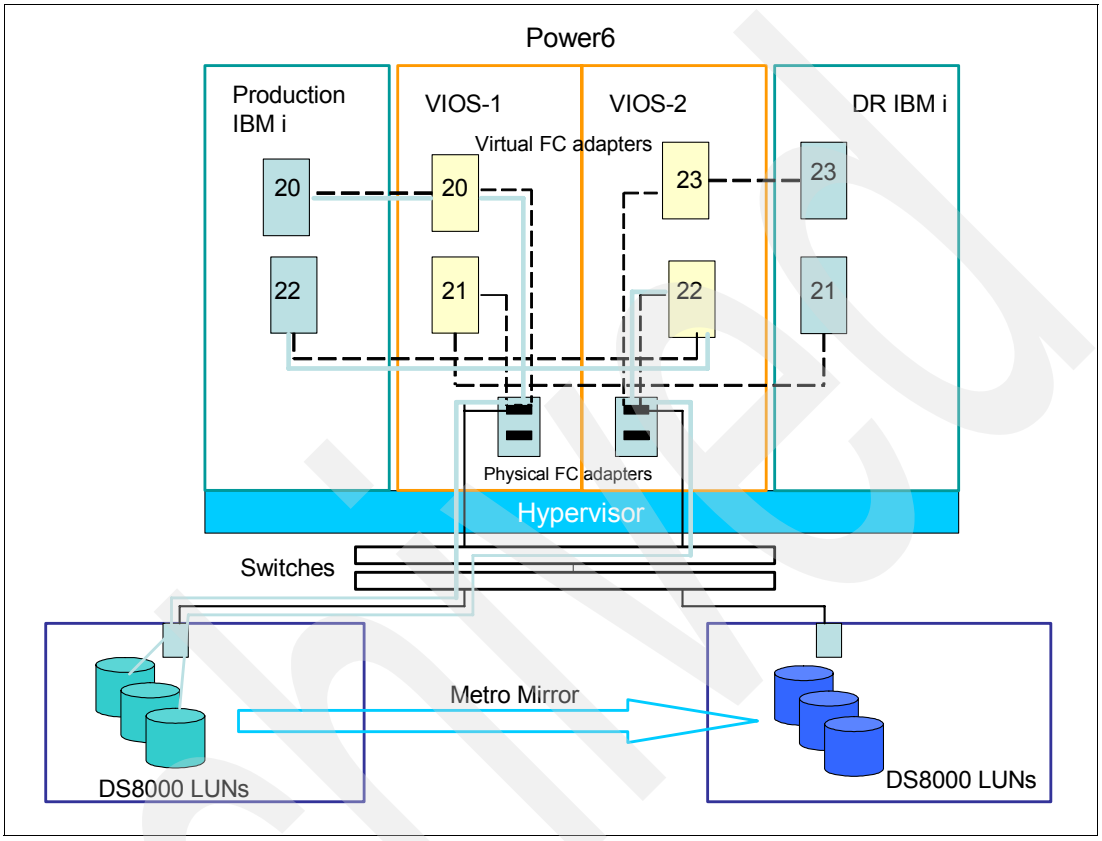

Figure 5-2 Testing environment for Metro Mirror with VIOS NPIV

5.4 Establishing Metro Mirror

Perform the following steps to establish Metro Mirror relationships:

1. At the primary DS8000, look for available ports to establish a remote mirror path. The `lsavail` command displays the port pairs through which the links for remote mirror and copy are established (that is, each link between a port in primary DS8000 and the corresponding port in secondary DS8000).

   In our example, we use DS CLI command **`lsavailpprcport`** for LSS 48, the command and part of resulting available port pairs, shown in Figure 5-3 on page 44.

   Note: In our illustrations of DS CLI commands, we do not specify the device ID of the primary and secondary DS8000s because they are already specified in our DS CLI profiles.

```
dscli> lsavailpprcport -remotewwnn 5005076303FFC663 48:48
Date/Time: August 22, 2009 2:00:38 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
Local Port Attached Port Type
============================
I0043      I0000         FCP
I0043      I0040         FCP
I0043      I0140         FCP
I0143      I0000         FCP
I0243      I0040         FCP
I0243      I0140         FCP
I0343      I0130         FCP
I0400      I0130         FCP
I0401      I0030         FCP
I0401      I0031         FCP
I0401      I0032         FCP
I0401      I0033         FCP
I0401      I0100         FCP
I0401      I0101         FCP
I0401      I0102         FCP
I0401      I0103         FCP
I0403      I0030         FCP
I0403      I0031         FCP
I0403      I0032         FCP
I0403      I0033         FCP
```

Figure 5-3 DS CLI lsavailpprcport

2. Create remote mirror and copy paths for each LSS in the production IBM i partition. A good practice is to create at least two paths for each LSS, each path using a different link. Figure 5-4 illustrates the creation of remote mirror and copy paths for our scenario. We do not use Metro Mirror consistency groups so we do not specify parameter **-consistgrp**.

> **Note:** Create remote mirror and copy path in both directions: from primary to secondary DS8000 and from secondary to primary DS8000. The paths in the reverse direction will be used for Metro Mirror failback.

```
dscli> mkpprcpath -remotewwnn 5005076303FFC663 -srclss 48 -tgtlss 48
I0043:I0040 I0243:I0140
Date/Time: August 22, 2009 2:41:28 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00149I mkpprcpath: Remote Mirror and Copy path 48:48 successfully
established.
dscli> lspprcpath 48
Date/Time: August 22, 2009 2:42:04 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
Src Tgt State   SS   Port  Attached Port Tgt WWNN
=========================================================
48  48  Success FF48 I0043 I0040         5005076303FFC663
48  48  Success FF48 I0243 I0140         5005076303FFC663
dscli>
```

Figure 5-4 DS CLI mkpprcpath command

3. Create the Metro Mirror relationships between the production IBM i LUNs (primary) and DR IBM i LUNs (secondary). Use the **mkpprc** command as shown in Figure 5-5.

```
dscli> mkpprc -type mmir 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 6:32:33 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4800:4800
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4801:4801
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4802:4802
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4900:4900
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4901:4901
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4902:4902
successfully created.
dscli>
```

Figure 5-5 DS CLI mkpprc

4. During the initial Metro Mirror synchronization, the Metro Mirror volume pairs are in *copy pending* state. To verify the state, use an **lspprc** command. When using the command with parameter **-l**, you can also see the number of tracks not yet synchronized. See Figure 5-6.

```
dscli> lspprc -l 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 6:42:03 CEST AM IBM DSCLI Version: 5.4.30.210 DS: IBM.2
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src Cascade
==========================================================================================
4800:4800 Copy Pending -      Metro Mirror 737260             Disabled Disabled
4801:4801 Copy Pending -      Metro Mirror 744239             Disabled Disabled
4802:4802 Copy Pending -      Metro Mirror 742402             Disabled Disabled
4900:4900 Copy Pending -      Metro Mirror 711155             Disabled Disabled
4901:4901 Copy Pending -      Metro Mirror 711636             Disabled Disabled
4902:4902 Copy Pending -      Metro Mirror 712518             Disabled Disabled
dscli>
```

Figure 5-6 DS CLI lspprc during Metro Mirror synchronization

When Metro Mirror pairs are synchronized, they are in a *full duplex* state (Figure 5-7).

```
dscli> lspprc -l 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 7:04:56 CEST AM IBM DSCLI Version:
ID        State       Reason Type         Out Of Sync Tracks
=============================================================
4800:4800 Full Duplex -      Metro Mirror 0
4801:4801 Full Duplex -      Metro Mirror 0
4802:4802 Full Duplex -      Metro Mirror 0
4900:4900 Full Duplex -      Metro Mirror 0
4901:4901 Full Duplex -      Metro Mirror 0
4902:4902 Full Duplex -      Metro Mirror 0
dscli>
```

Figure 5-7 DS CLI lspprc after Metro Mirror is synchronized

The Metro Mirror relationships are now established.

Before switching over to the DR site, be sure to verify that the DR IBM i partition is correctly connected to the DS8000 and that the correct adapter is tagged as Load source. In our example, we check that the WWPN of the tagged adapter is the same as the WWPN on one of the host connections. Relevant screen captures from the HMC and DS CLI are shown in Figure 5-8, Figure 5-9, and Figure 5-10.

Also, check that the IPL source of the DR partition is set to `A` or `B`, and that the IPL mode is set to normal (set the IPL mode to manual only when you want to check for any issue in the DST before proceeding with IPL).

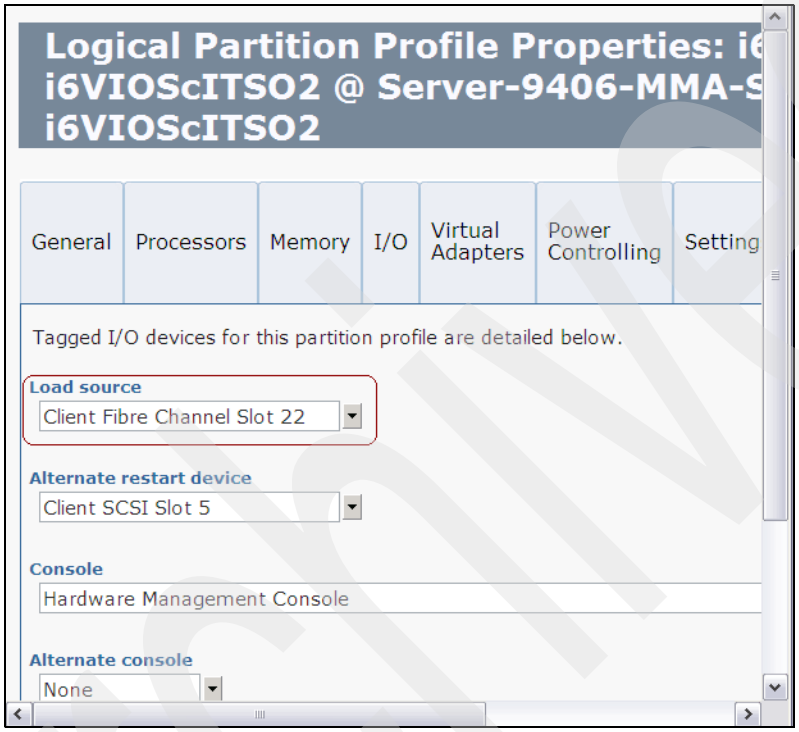

Figure 5-8 Tagged adapter for Load source

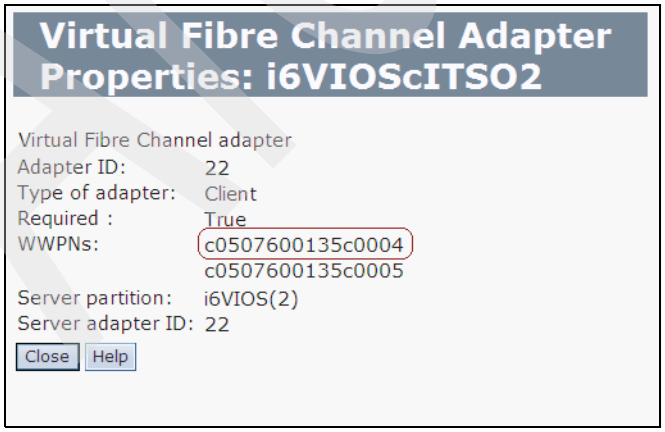

Figure 5-9 WWPN of the tagged adapter

```
dscli> lshostconnect 0009 0013
Date/Time: August 24, 2009 1:37:07 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
Name           ID   WWPN             HostType Profile              portgrp
volgrpID ESSIOport
================================================================================
==============
Vrtual_FC_MM_1 0009 C0507600135C0004 iSeries  IBM iSeries - OS/400      0 V12
all
Vrtual_FC_MM_2 0013 C0507600135C000A iSeries  IBM iSeries - OS/400      0 V12
all
dscli>
```

Figure 5-10 WWPN of tagged adapter in Host connection

5.5 Planned outage with Metro Mirror

A planned outage of the production IBM i partition occurs for situations such as disruptive hardware upgrades, or software installations. The IT centers may choose to temporarily run their production environment on the DR IBM i partition, which is then IPLed from the Metro Mirror target volumes.

After the production IBM i partition is back and running again, you have the option to either replicate the changes made to the DR system during the outage back to the production system, or to replicate changes in the production system to the DR site. The decision depends on the type of planned outage.

To switch from the production system to the DR site in planned outage situations, follow these steps:

1. Power down the production IBM i system by issuing the **pwrdwnsys** command, shown in Figure 5-11.

```
 Power Down System (PWRDWNSYS)

 Type choices, press Enter.

 How to end . . . . . .   *CNTRLD       *CNTRLD, *IMMED
 Controlled end delay time  . . .   10            Seconds, *NOLIMIT
 Restart options:
   Restart after power down . . .   *NO           *NO, *YES
   Restart type . . . . . .   *IPLA         *IPLA, *SYS, *FULL
 IPL source . . . . . .   *PANEL        *PANEL, A, B, D, *IMGCLG









                                                                      Bottom
 F3=Exit   F4=Prompt   F5=Refresh   F10=Additional parameters   F12=Cancel
 F13=How to use this display        F24=More keys
```

Figure 5-11 Power Down of production IBM i

2. Suspend the Metro Mirror at the primary DS8000. Use the **pausepprc** command as shown in Figure 5-12.

```
dscli> pausepprc 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 9:55:18 CEST PM IBM DSCLI Version: 5.4.30.210 DS: IB
.2107-7520781
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4800:4800 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4801:4801 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4802:4802 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4900:4900 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4901:4901 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4902:4902 relationship
successfully paused.
dscli>
```

Figure 5-12 Suspend Metro Mirror

As a result, the source volumes are now in *host source suspended* state (see Figure 5-13), which means that they are enabled for read and write I/O operations. By suspending the Metro Mirror, we ensure that source volumes will be enabled for read and write operations when the Metro Mirror failback takes place.

```
dscli> lspprc 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 10:07:08 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
ID        State     Reason      Type         SourceLSS Timeout (secs) Critical
M
ode First Pass Status
================================================================================
=
====================
4800:4800 Suspended Host Source Metro Mirror 48        60             Disabled
    Invalid
4801:4801 Suspended Host Source Metro Mirror 48        60             Disabled
    Invalid
4802:4802 Suspended Host Source Metro Mirror 48        60             Disabled
    Invalid
4900:4900 Suspended Host Source Metro Mirror 49        60             Disabled
    Invalid
4901:4901 Suspended Host Source Metro Mirror 49        60             Disabled
    Invalid
4902:4902 Suspended Host Source Metro Mirror 49        60             Disabled
    Invalid
```

Figure 5-13 Source LUNs in source suspended state

3. At the secondary DS8000, fail over the Metro Mirror for the relevant volume pairs. Use the **failoverpprc** command, as shown in Figure 5-14.

```
dscli> failoverpprc -type mmir 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 23, 2009 10:21:43 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4800:4800 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4801:4801 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4802:4802 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4900:4900 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4901:4901 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4902:4902 successfully
reversed.
dscli>
```

Figure 5-14 Failover Metro Mirror

After Metro Mirror failover, the Metro Mirror relationship is reversed and the volumes at the DR site are in *host source suspended* state, thus enabled for read and write operations.

4. IPL the IBM i DR partition that is connected to the Metro Mirror target LUNs of the secondary DS8000. Use the HMC to activate the DR partition.

   Because the production IBM i was powered-down before the Metro Mirror suspend and failover, the IPL in DR partition is normal, as shown in Figure 5-15.

```
 Licensed Internal Code IPL in Progress
                                                          08/24/09  01:47:22
 IPL:
   Type  . . . . . . . . . . . . . . . :   Attended
   Start date and time . . . . . . . . :   08/24/09  01:47:03
   Previous system end . . . . . . . . :   Normal
   Current step / total  . . . . . . . :      16     16
   Reference code detail . . . . . . . :   C6004065

 IPL step                                  Time Elapsed   Time Remaining
  Commit Recovery                             00:00:01        00:00:00
  Data Base Initialization                    00:00:01        00:00:00
  Journal IPL Clean up                        00:00:01        00:00:00
  Commit Initialization                       00:00:01        00:00:00
 >Start the operating system

 Item:
   Current / Total . . . . . . . :

 Sub Item:
   Identifier  . . . . . . . . . :
   Current / Total . . . . . . . :
```

Figure 5-15 Normal IPL in DR partition

The DR partition now contains a clone of the production IBM i. Consider the information in 4.2.2, “Cloning an IBM i partition” on page 31.

Two paths are established from the DR IBM i partition to the Metro Mirror target LUNs, each path through one VIOS. Therefore, the volumes are in multipath after the DR partition is IPled, as can be seen in the DST display of the LUNs in Figure 5-16.

```
 Display Disk Configuration Status

           Serial                        Resource                      Hot Spare
 ASP Unit  Number          Type Model    Name       Status             Protection
   1                                               Unprotected
         1  50-4902001      2107  A04   DMP012     RAID 5/Active          N
         2  50-4800001      2107  A04   DMP013     RAID 5/Active          N
         3  50-4801001      2107  A04   DMP014     RAID 5/Active          N
         4  50-4802001      2107  A04   DMP015     RAID 5/Active          N
         5  50-4900001      2107  A04   DMP016     RAID 5/Active          N
         6  50-4901001      2107  A04   DMP017     RAID 5/Active          N

Press Enter to continue.

 F3=Exit        F5=Refresh         F9=Display disk unit details
 F11=Disk configuration capacity   F12=Cancel
```

Figure 5-16 Multipath in DR partition

5. After the production IBM i has been restored, perform the remaining steps to switch back to the production site. First, power down the DR IBM i partition.
6. Depending on the outage, re-establish the Metro Mirror in the direction secondary to primary, or primary to secondary by using a **failbackpprc** command. The re-synchronization of the Metro Mirror pairs are incremental in each case (failback from secondary to primary, or failback from primary to secondary).

   In our example, we want to propagate changes from the DR site to the production site, so we perform the **failbackpprc** at the secondary DS8000, as shown in Figure 5-17.

```
dscli> failbackpprc -type mmir 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 24, 2009 4:15:44 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4800:4800 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4801:4801 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4802:4802 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4900:4900 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4901:4901 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4902:4902 successfully
failed back.
dscli>
```

Figure 5-17 Failback of Metro Mirror

After incremental re-synchronization, the Metro Mirror pairs are now in *full duplex* state.

7. Suspend the Metro Mirror LUNs at the secondary DS8000. Use the **pausepprc** command, as shown in Figure 5-12.
8. Fail over the Metro Mirror LUNs at the primary DS8000. Use the **failoverpprc** command, as shown in Figure 5-18.

```
dscli> failoverpprc -type mmir 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 24, 2009 8:09:10 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4800:4800 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4801:4801 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4802:4802 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4900:4900 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4901:4901 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4902:4902 successfully
reversed.
```

Figure 5-18 Failover Metro Mirror on primary DS8000

9. Fail back Metro Mirror from the primary to the secondary DS8000. Issue a **failbackpprc** command, as shown in Figure 5-19.

```
dscli> failbackpprc -type mmir 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 24, 2009 8:16:34 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4800:4800 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4801:4801 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4802:4802 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4900:4900 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4901:4901 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4902:4902 successfully
failed back.
dscli>
```

Figure 5-19 Failback Metro Mirror on primary DS8000

10. IPL the production IBM i partition.

Note: The production IBM i partition is now a clone of the DR partition at the time the DR partition was powered down (before the failover to the primary site).

5.6 Unplanned outage with Metro Mirror

Unplanned outages result from incidents such as hardware failures (IBM i or DS8000), or even failure of the entire production site. In our tests, we simulated an unplanned outage by immediately shutting down the production partition without first properly powering down the IBM i partition, as follows:

1. If the primary DS8000 is working, suspend Metro Mirror at the primary DS8000.

   Note: If you suspend the Metro Mirror while it is still in working condition (both DS8000 and Metro Mirror links are working), whether you suspend it at the primary or secondary DS8000 will not matter.

   In true disaster situations, when Metro Mirror links are broken and you can only suspend the Metro Mirror at the secondary DS8000, the primary LUNs will only enter a suspended state if there is still communication to the secondary, or on a write operation from the host to the primary DS8000 (a write operation occurs from primary to secondary).

2. Fail over Metro Mirror at secondary LUNs, as is described in step 3 on page 49 (in 5.5, “Planned outage with Metro Mirror” on page 47).
3. If possible, fail back the Metro Mirror, as described in step 6 on page 50 (in 5.5, “Planned outage with Metro Mirror” on page 47).
4. IPL System i partition at the DR site. IBM i was broken when we shut-down the partition; the IO activity on disk space was transferred by Metro Mirror to the secondary site, so the IPL of DR partition is abnormal.

After the production system is back and running again, perform the following steps to switch back to the production site:

1. Failback the Metro Mirror from secondary to primary DS8000. Use **`failbackpprc`**, as shown in Figure 5-17.
2. After Metro Mirror is re-synchronized, power down the DR IBM i.
3. Suspend Metro Mirror LUNs at the secondary DS8000. Use **`pausepprc`**, as shown in Figure 5-12.
4. Failover the Metro Mirror at the primary DS8000. See Figure 5-18.
5. Failback Metro Mirror from the primary to the secondary DS8000. See Figure 5-19.
6. IPL the production IBM i partition. The IPL is normal because the DR IBM i was powered-down before suspending Metro Mirror.

6

Global Mirror scenarios

In this chapter, we describe the environment we used for testing Global Mirror with IBM i when connected to the DS8000 through VIOS with NIPV. We explain the procedures to recover from Global Mirror target disks in case of disaster at the production site, or to use Global Mirror for business continuity during planned outages.

6.1 Global Mirror overview

The section has an overview of Global Mirror. Because Global Copy is a component of Global Mirror, the concepts of Global Copy are also provided.

6.1.1 Global Copy

Global Copy, previously known as Peer-to-Peer Remote Copy Extended Distance (PPRC-XD), copies data non-synchronously and over longer distances than is possible with Metro Mirror. When operating in Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume, instead of sending a constant stream of updates. This causes less impact to application write operations for source volumes and less demand for bandwidth resources, while allowing a more flexible use of the available bandwidth.

Global Copy does not keep the sequence of write operations. Therefore, the copy is a fuzzy copy, but you can make a consistent copy through synchronization (called a go-to-sync operation). After the synchronization, you can issue FlashCopy at the secondary site to make a consistent backup copy. After establishing FlashCopy, you can change the mode back to the non-synchronous mode. See Figure 6-1.

An alternative method to acquire a consistent copy is to pause the applications until all changed data at the local site has drained to the remote site. When all consistent data is replicated to the remote site, suspend Global Copy, restart the applications, issue the FlashCopy, and then resume the non-synchronous (Global Copy) operation.

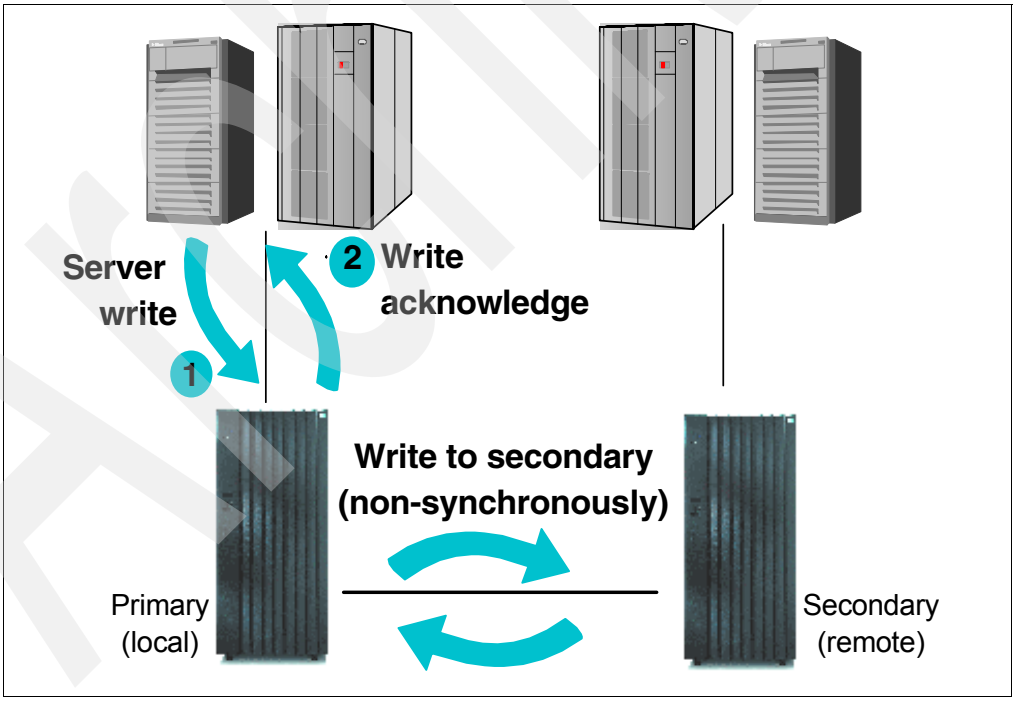

Figure 6-1 Global Copy basic operation

6.1.2 Global Mirror

Global Mirror, previously known as Asynchronous PPRC, is a two-site, long distance, asynchronous, remote copy technology for both System z® and Open Systems data. This solution integrates the Global Copy and FlashCopy technologies. With Global Mirror, the data that the host writes to the storage unit at the local site is asynchronously mirrored to the storage unit at the remote site. With special management steps, under control of the local master storage unit (see the Note in “How Global Mirror works” on page 56), a consistent copy of the data is automatically maintained on the storage unit at the remote site.

Global Mirror operations (shown in Figure 6-2) provide the following benefits:

- Supports virtually unlimited distances between the local and remote sites, with the distance typically limited only by the capabilities of the network and the channel extension technology. This *unlimited* distance enables you to choose your remote site location based on business needs, and enables site separation to add protection from localized disasters.
- Provides a consistent and restartable copy of the data at the remote site, created with minimal impact to applications at the local site.
- Supports data currency where, for many environments, the remote site lags behind the local site typically 3 - 5 seconds, minimizing the amount of data exposure in the event of an unplanned outage. The actual lag in data currency that you experience will depend upon a number of factors, including specific workload characteristics and bandwidth between the local and remote sites.
- Supports dynamic selection of the desired recovery point objective (RPO), based upon business requirements and optimization of available bandwidth.
- Provides efficient synchronization of the local and remote sites with support for failover and failback operations, helping to reduce the time that is required to switch back to the local site after a planned or unplanned outage.

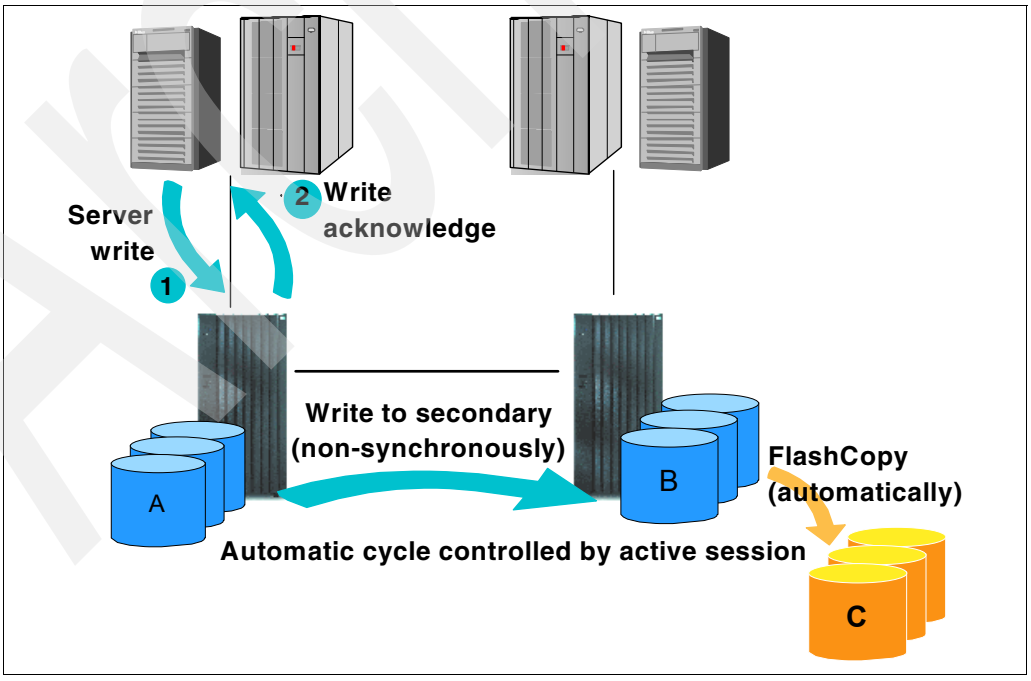

Figure 6-2 Global Mirror basic operation

How Global Mirror works

Figure 6-3 illustrates the basics of how Global Mirror works; everything is in an automatic fashion under the control of the DS8000 microcode and the Global Mirror session.

Important: Global Mirror works like a distributed application, which is usually built on a server-to-client relationship. In Global Mirror, the server and client can be replaced by terms such as master and subordinate. A master coordinates all efforts within a Global Mirror environment. After the master is started and manages a Global Mirror environment, the master issues all related commands over inband communication to its attached subordinates at the local site.

When the master and subordinate are in a single storage disk subsystem, the subordinate is internally managed by the master. However, with two or more storage disk subsystems at the local site, which participate in a Global Mirror session, the subordinate is external and requires separate attention when creating and managing a Global Mirror session or environment.

Figure 6-3 shows that the A volumes at the local site are the production volumes and are used as Global Copy primary volumes. The data from the A volumes is replicated to the B volumes, which are the Global Copy secondary volumes. At a certain point in time, a consistency group is created using all of the A volumes, even if they are located in different storage units. This approach does not affect applications, because the creation of the consistency group is very quick (on the order of milliseconds).

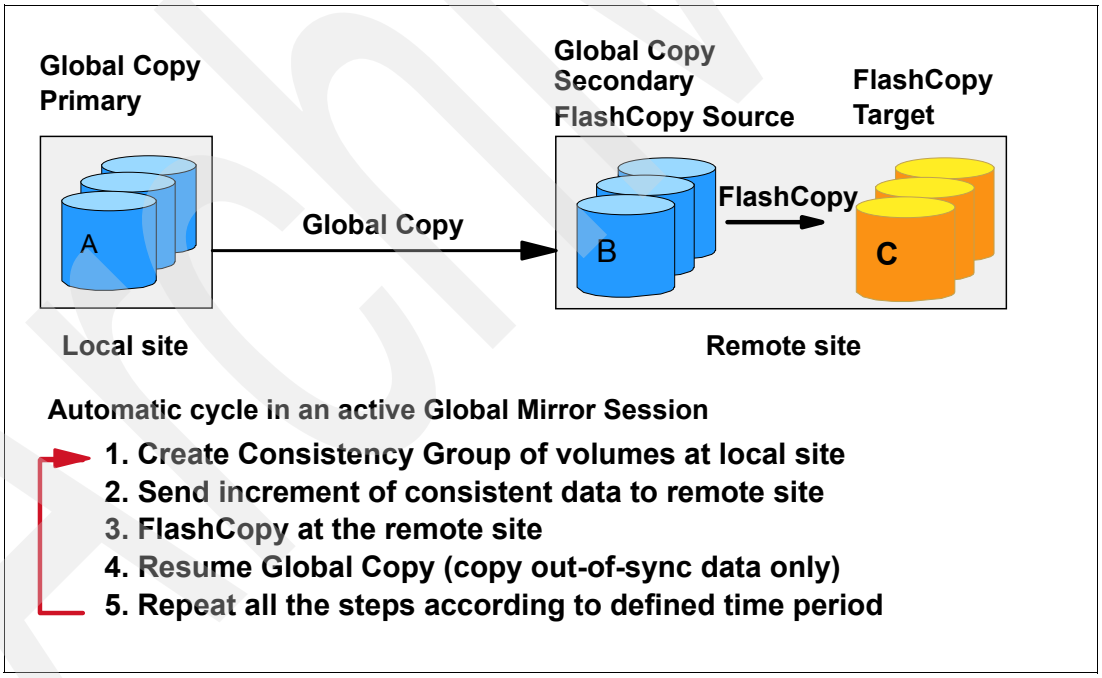

Figure 6-3 How Global Mirror works

After the consistency group is created, the application write operations can continue updating the A volumes. The increment of the consistent data is sent to the B volumes using the existing Global Copy relationships. When the data reaches the B volumes, a FlashCopy is taken to the C volumes.

After this step is complete, a consistent set of volumes have been created at the secondary site. For this very brief moment only, the B volumes and the C volumes are equal in their

content. Because the B volumes, except at the moment of doing the FlashCopy, usually contain a *fuzzy* copy of the data, the C volumes are used to hold the last consistent point-in-time copy of the data while the B volumes are being updated by Global Copy. Therefore, you need the B and the C volume to build consistent data which are as little as possible behind the production data.

The data at the remote site is current within typically 3 - 5 seconds, but this recovery point depends on the workload and bandwidth available to the remote site.

Global Mirror recovery

The two stages to the recovery of a Global Mirror environment are:

1. The first stage is to check the exact status of Global Mirror at the time of the outage. Depending on where Global Mirror was in the consistency group formation process at the time of the failure event, certain actions might be required to ensure that the C copy is consistent.

   If the Global Mirror environment was part way through the FlashCopy process when the failure occurred, this is similar to an in-flight transaction in a database environment. If the commit process has not yet started, then we must revert the FlashCopy to back out the consistency group; if the commit process has started, we must complete this process to commit the consistency group.

2. The second stage is to recover the environment and enable production systems to be restarted on the B devices and prepare for a potential return to the primary site. This is performed with the following process:

   a. Fail over the B devices. This step places the B devices in a primary suspended state.

   b. Fast reverse restore the FlashCopy relationship with the C devices. This step restores the latest consistency group to the B devices and starts a background copy for those tracks that have been modified since the latest consistency group.

   c. Use FlashCopy from the B devices to the C devices to save an image of the last consistency group. This step is optional but preserves an image of the production devices at the recovery point in case this might be required.

   d. Restart the production systems.

With its efficient and autonomic implementation, Global Mirror is a solution for disaster recovery implementations where a consistent copy of the data must be kept at all times at a remote location that can be separated by a very long distance from the production site. For more information about the functionality of Copy Services, refer to *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788.

6.2 Environment for Global Mirror scenarios

When implementing a disaster recovery solution with DS8000 Global Mirror, the Disaster Recovery (DR) IBM i partition resides on a remote Power6 system. Each production IBM i and DR IBM i will be connected to a DS8000 through two VIOS partitions in NPIV, with the relevant IBM i partition and the two VIOS partitions residing in the same Power6.

In our test scenario, we used the same Power6 system for both the production and DR IBM i partitions (in a real environment, you would need a separate Power6 system at each site). The setup is as follows: two VIOS partitions, IBM i production partition, and IBM i Disaster Recovery partition. In each VIOS partition, we create two virtual FC adapters, both of them mapped to one specific port of the physical FC adapter. These ports are connected to both the primary and secondary DS8000 through NPIV-enabled switches. Each IBM i production and backup partition is connected to each VIOS through one virtual FC adapter. We assign the set of LUNs on the primary DS8000 to both virtual FC adapters in the production IBM i partition. In doing so, the IBM i has two paths (multipath) to the set of LUNs. Similarly, we assign the Global Mirror B volumes in the secondary DS8000 to the DR partition through two virtual FC adapters, each through one VIOS, so that multipath is enabled in the DR partition too.

Our environment for this scenario is shown in Figure 6-4. The solid black lines indicate physical connection, the dashed black lines indicate connection of virtual adapters, and the blue lines are assignment of LUNs to virtual FC adapters. The assignment of LUNs is shown for the production IBM i only; The DR IBM i uses a similar assignment scheme to two virtual FC adapters through the two VIOS partitions.

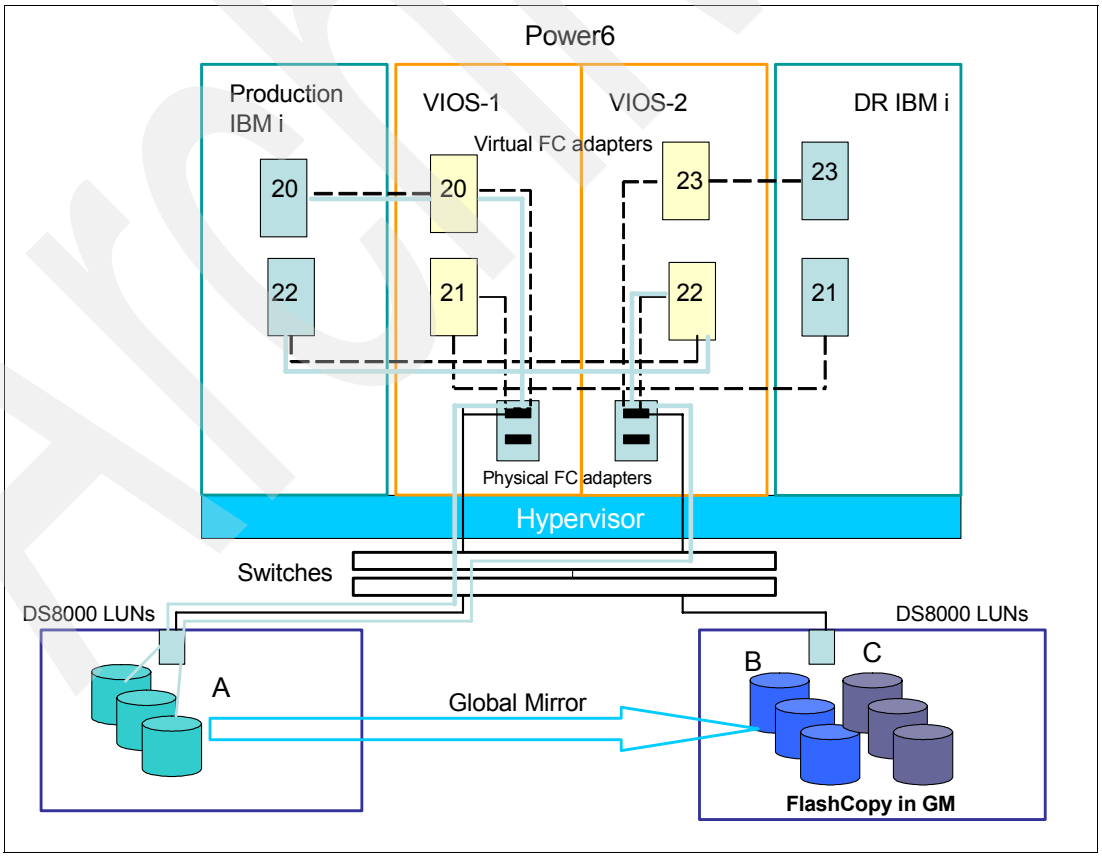

Figure 6-4 Testing environment for Global Mirror with VIOS NPIV

6.3 Establishing Global Mirror

Perform the following steps to establish Global Mirror relationships:

1. Establish a *remote mirror and copy* path, as described in 5.4, “Establishing Metro Mirror” on page 43.
2. At the primary DS8000, create a Global Copy relationship between the IBM i LUNs of the primary and secondary DS8000 (Global Mirror A and B volumes), as shown in Figure 6-5.

```
dscli> mkpprc -type gcp 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 25, 2009 2:31:01 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4800:4800
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4801:4801
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4802:4802
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4900:4900
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4901:4901
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 4902:4902
successfully created.
dscli>
```

Figure 6-5 Create Global Copy relation

3. Wait for Global Copy to synchronize. This is indicated by no *out of sync* tracks being displayed in the output of the DS CLI command **lspprc -l**, part of which is shown in Figure 6-6.

```
dscli> lspprc -l 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 25, 2009 3:12:12 CEST AM IBM DSCLI Version: 5.4.30.210 DS: IBM.2107-7520781
ID        State       Reason Type        Out Of Sync Tracks Tgt Read Src Cascade Tgt Cascade Date
==========================================================================================================
4800:4800 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
4801:4801 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
4802:4802 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
4900:4900 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
4901:4901 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
4902:4902 Copy Pending -      Global Copy 0                  Disabled Disabled    Invalid     -
dscli>
```

Figure 6-6 Global Copy synchronized

4. At the secondary DS8000, establish a FlashCopy relationship between the Global Copy target volumes (Global Mirror volumes B) and the FlashCopy LUNs (Global Mirror volumes C), with the following FlashCopy parameters:

   – `-tgtinhibit`

     Prevents host system write operations to the target while the FlashCopy relationship exists.

   – `-record`

     Keeps record of the tracks that were modified on both volumes within a FlashCopy pair. Select this parameter when you create an initial FlashCopy volume pair that you intend to use with the **`resyncflash`** command.

   – `-nocp`

     Inhibits background copy. Data is copied from the source volume to the target volume only if a track on the source volume is modified.

   – `-persist`

     Keep the FlashCopy relationship until it is explicitly or implicitly terminated.

   The **`-persist`** parameter is automatically selected when the **`-record`** parameter is selected. Therefore, you do not specify the **`-persist`** parameter explicitly.

   See Figure 6-7.

```
dscli> mkflash -tgtinhibit -nocp -record 4800-4802:4803-4805
4900-4902:4903-4905
Date/Time: August 25, 2009 3:41:44 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00137I mkflash: FlashCopy pair 4800:4803 successfully created.
CMUC00137I mkflash: FlashCopy pair 4801:4804 successfully created.
CMUC00137I mkflash: FlashCopy pair 4802:4805 successfully created.
CMUC00137I mkflash: FlashCopy pair 4900:4903 successfully created.
CMUC00137I mkflash: FlashCopy pair 4901:4904 successfully created.
CMUC00137I mkflash: FlashCopy pair 4902:4905 successfully created.
dscli>
```

Figure 6-7 FlashCopy for Global Mirror volumes B to C

5. At the primary DS8000, define the Global Mirror session for the logical subsystems containing the production LUNs. Use the **`mksession`** commands as shown in Figure 6-8.

```
dscli> mksession -lss 48 01
Date/Time: August 25, 2009 7:36:48 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00145I mksession: Session 01 opened successfully.
dscli> mksession -lss 49 01
Date/Time: August 25, 2009 7:37:09 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00145I mksession: Session 01 opened successfully.
dscli>
```

Figure 6-8 Define Global Mirror session

6. Add the Global Mirror A volumes to the session that was defined in the previous step. Use the **chsession -action add -volume** command at the primary DS8000, as shown in Figure 6-9.

```
dscli> chsession -lss 48 -action add -volume 4800-4802 01
Date/Time: August 25, 2009 8:00:34 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00147I chsession: Session 01 successfully modified.
dscli> chsession -lss 49 -action add -volume 4900-4902 01
Date/Time: August 25, 2009 8:01:06 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00147I chsession: Session 01 successfully modified.
```

Figure 6-9 Add volumes to Global Mirror session

7. Verify the results with the **lssession** command, as shown in Figure 6-10.

```
dscli> lssession -l 48
Date/Time: August 25, 2009 8:02:58 CEST PM IBM DSCLI Version: 5.4.30.210 DS: IBM.2107-7520781
LSS ID Session Status Volume VolumeStatus PrimaryStatus       SecondaryStatus   FirstPassComplete AllowCascading
==========================================================================================================================
48     01      Normal 4800   Join Pending Primary Copy Pending Secondary Simplex True              Disable
48     01      Normal 4801   Join Pending Primary Copy Pending Secondary Simplex True              Disable
48     01      Normal 4802   Join Pending Primary Copy Pending Secondary Simplex True              Disable
dscli> lssession -l 49
Date/Time: August 25, 2009 8:03:04 CEST PM IBM DSCLI Version: 5.4.30.210 DS: IBM.2107-7520781
LSS ID Session Status Volume VolumeStatus PrimaryStatus       SecondaryStatus   FirstPassComplete AllowCascading
==========================================================================================================================
49     01      Normal 4900   Join Pending Primary Copy Pending Secondary Simplex True              Disable
49     01      Normal 4901   Join Pending Primary Copy Pending Secondary Simplex True              Disable
49     01      Normal 4902   Join Pending Primary Copy Pending Secondary Simplex True              Disable
dscli>
```

Figure 6-10 List GM session

Note: We do not have to do anything to add the B and C volumes to the Global Mirror session. They are automatically recognized by the Global Mirror mechanism through the Global Copy relationships and the FlashCopy relationships.

8. At the primary DS8000, start the Global Mirror session to initiate the process of consistency group formation. Use the **mkgmir** command. The results can be verified with the **showgmir** command, shown in Figure 6-11 on page 62.

Note: Because of the master subordinate structure of the Global Mirror, you perform the commands for the first logical subsystem only. The other logical subsystems are then processed automatically.

```
dscli> mkgmir -lss 48 -session 01
Date/Time: August 25, 2009 8:21:49 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00162I mkgmir: Global Mirror for session 01 successfully started.
dscli> showgmir 48
Date/Time: August 25, 2009 8:22:35 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
ID                           IBM.2107-7520781/48
Master Count                 1
Master Session ID            0x01
Copy State                   Running
Fatal Reason                 Not Fatal
CG Interval Time (seconds)   0
Coord. Time (milliseconds)   50
Max CG Drain Time (seconds)  30
Current Time                 08/25/2009 20:22:37 CEST
CG Time                      08/25/2009 20:22:37 CEST
Successful CG Percentage     97
FlashCopy Sequence Number    0x4A942BED
Master ID                    IBM.2107-7520781
Subordinate Count            0
Master/Subordinate Assoc     -
dscli>
```

Figure 6-11 Staring Global Mirror

In the **mkgmir** command, the logical subsystem (LSS) specified with the -lss parameter becomes the master. In our example, it is LSS 48. With this command, we also specify the Global Mirror session ID for the session we are starting.

When starting the Global Mirror session, you can change the session tuning parameters, as follows, with the **mkgmir** command:

- `cginterval`

  Specifies how long to wait between the formation of consistency groups. If this number is not specified or is set to zero, consistency groups are formed continuously.

- `coordinate`

  Indicates the maximum time that Global Mirror processing can hold host I/Os in the source disk subsystem to start forming a consistency group.

- `drain`

  Specifies maximum amount of time in seconds allowed for the data to drain to the remote site before failing the current consistency group.

This **showgmir** command shows the Global Mirror current status. If the Copy State field indicates `Running` status, the Global Mirror is satisfactorily operating. If the status is `Fatal`, the Global Mirror failed, and the Fatal Reason field will show the reason for the failure.

The **showgmir** command also shows the current time in the Current Time field, which is the time when the DS8000 received this command. The time corresponding to when the last successful consistency group was formed is shown in the CG Time field. You can obtain the current Recovery Time Objective (RPO) for this Global Mirror session from the difference between the Current Time and the CG Time.

9. Check the status of the volumes in the Global Mirror session by using **lssession** command. In our example we use **lssession** command for LSS 48 and LSS 49. As shown in Figure 6-12, the GM volumes are now in Active status; they were in Join Pending status before starting the Global Mirror.

```
dscli> lssession 48-49
Date/Time: August 25, 2009 8:35:07 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
LSS ID Session Status         Volume VolumeStatus PrimaryStatus
SecondaryStatus   FirstPassComplete
g
================================================================================
==========================
=
48     01      CG In Progress 4800   Active       Primary Copy Pending
Secondary Simplex True
48     01      CG In Progress 4801   Active       Primary Copy Pending
Secondary Simplex True
48     01      CG In Progress 4802   Active       Primary Copy Pending
Secondary Simplex True
49     01      CG In Progress 4900   Active       Primary Copy Pending
Secondary Simplex True
49     01      CG In Progress 4901   Active       Primary Copy Pending
Secondary Simplex True
49     01      CG In Progress 4902   Active       Primary Copy Pending
Secondary Simplex True
dscli>
```

Figure 6-12 Global Mirror volumes after GM is started

6.4 Recovering after failure at the production site

We assume that the production IBM partition is connected through VIOS with NPIV to the production LUNs (Global Mirror A volumes) and that the DR IBM i partition is connected to Global Mirror B volumes. In our test, the partitions are connected to Global Mirror A and B LUNs, as is described in 6.2, “Environment for Global Mirror scenarios” on page 58.

Note: In our illustrations, we show the use of DS CLI commands for recovery with Global Mirror to better explain the mechanism. However, in practice, use a management tool such as TPC-R, PowerHA or the IBM i Copy Services Toolkit when performing a recovery from Global Mirror target LUNs.

6.4.1 Switch over to DR site after the failure on local site

After a production site failure (in our scenario, we shut-down the production partition, with option *immediate*, to simulate the failure at the production site), perform the following tasks to recover from Global Mirror target volumes at the DR site:

1. If the primary DS8000 is available, stop the Global Mirror session at the primary DS8000. Use the **rmgmir** command as shown in Figure 6-13.

```
dscli> rmgmir  -lss 48 -session 01
Date/Time: August 26, 2009 1:15:39 CEST AM IBM DSCLI Version: 5.4.30.210 DS: IBM
.2107-7520781
CMUC00166W rmgmir: Are you sure you want to stop the Global Mirror session 01:?
[y/n]:y
CMUC00165I rmgmir: Global Mirror for session 01 successfully stopped.
dscli>
```

Figure 6-13 Stop Global Mirror session

2. At the secondary DS8000, perform a Global Copy *failover* on Global Mirror B volumes. Use the **failoverpprc** command, as shown in Figure 6-14.

```
dscli> failoverpprc -type gcp 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 1:27:18 CEST AM IBM DSCLI Version: 5.4.30.210 DS: IBM
.2107-75ABTV1
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4800:4800 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4801:4801 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4802:4802 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4900:4900 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4901:4901 successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4902:4902 successfully reversed.
```

Figure 6-14 Failover Global Copy

This step changes the volume state to `Suspended Host Source`. Figure 6-15 shows the status of the Global Mirror B volumes after failover.

```
dscli> lspprc 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 1:27:52 CEST AM IBM DSCLI Version: 5.4.30.210 DS: IBM
.2107-75ABTV1
ID        State     Reason      Type        SourceLSS Timeout (secs) Critical Mo
de First Pass Status
================================================================================
====================
4800:4800 Suspended Host Source Global Copy 48        60             Disabled
   True
4801:4801 Suspended Host Source Global Copy 48        60             Disabled
   True
4802:4802 Suspended Host Source Global Copy 48        60             Disabled
   True
4900:4900 Suspended Host Source Global Copy 49        60             Disabled
   True
4901:4901 Suspended Host Source Global Copy 49        60             Disabled
   True
4902:4902 Suspended Host Source Global Copy 49        60             Disabled
   True
```

Figure 6-15 Status of Global Mirror B volumes after failover

3. Investigate whether all FlashCopy relationships are in a consistent state. To determine the state of the FlashCopy relationship, query all FlashCopy relationships between B and C, which are part of the consistency group. Global Mirror might have been in the middle of forming a consistency group and the FlashCopy operation might have not completed the creation of a full set of consistent C volumes.

   Each FlashCopy pair needs a FlashCopy query to identify its state. Use the **lsflash** command on secondary DS8000 to check the SequenceNum and Revertible field status. Figure 6-16 shows that the Revertible status of all of the FlashCopy pairs is Disabled (that is, non-revertible) and the sequence numbers of all relationships are equal. No special action is required in this case.

   For more information about special action for Global Mirror recovery, refer to *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788.

```
dscli> lsflash 4800-4802:4803-4805 4900-4902:4903-4905
Date/Time: August 26, 2009 1:41:23 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
ID      srclss SequenceNum Timeout ActiveCopy Recording Persistent Revertible
SourceWriteEna
=============================================================================================
4800:4803 48     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
4801:4804 48     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
4802:4805 48     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
4900:4903 49     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
4901:4904 49     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
4902:4905 49     4A9470A2    60      Disabled   Enabled   Enabled    Disabled   Enabled
dscli>
```

Figure 6-16 FlashCopy relation between B and C volumes

4. At this point only the C volumes (logical data) comprise a set of consistent data volumes, although the physical data of the C volumes might be spread over the physical B and C volumes. The B volumes (logical data) do not provide consistent data volumes because Global Copy does not provide data consistency.

   We want to have two good copies of the data at the recovery site. The aim is to have a consistent set of volumes to work with, while still keeping a good copy, which we can resort to if necessary. The next step is to create the same consistency on the B volumes as we have on the C volumes. This can be achieved with the **reverseflash -fast** command. This operation is called *Fast Reverse Restore (FRR)*. You have to use the **-tgtpprc** parameter with the **reverseflash -fast** command because the B volume are also a Global Copy source at this step.

   Note: Although the Fast Reverse Restore operation starts a background copy from the C to the B volumes, in the **reverseflash** command you have to specify the B volumes as the FlashCopy sources and the C volumes as the FlashCopy targets.

An illustration of the **reverseflash** command is shown on Figure 6-17.

```
dscli> reverseflash -fast -tgtpprc 4800-4802:4803-4805 4900-4902:4903-4905
Date/Time: August 26, 2009 2:29:14 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00169I reverseflash: FlashCopy volume pair 4800:4803 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 4801:4804 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 4802:4805 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 4900:4903 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 4901:4904 successfully reversed.
CMUC00169I reverseflash: FlashCopy volume pair 4902:4905 successfully reversed.
dscli>
```

Figure 6-17 Fast Reverse Restore

5. Wait until all Fast Reverse Restore operations and their background copy have completed successfully before proceeding to step 6. Again, when the background copy completes, the FlashCopy relationship ends. Therefore, you should check that no FlashCopy relationships remain to determine that all Fast Reverse Restore operations have completed.

   To check the status of FlashCopy, we use **lsflash**, as is shown in Figure 6-18.

```
dscli> lsflash 4800-4802:4803-4805 4900-4902:4903-4905
Date/Time: August 26, 2009 2:34:48 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00234I lsflash: No Flash Copy found.
dscli>
```

Figure 6-18 FlashCopy relation between B and C ended

6. Re-establish the FlashCopy relationship between B and C volumes by using the **mkflash** command with **-tgtinhibit -nocp -record** parameters, described in 6.3, “Establishing Global Mirror” on page 59. Our example is shown in Figure 6-19.

```
dscli> mkflash -tgtinhibit -nocp -record 4800-4802:4803-4805 4900-4902:4903-4905
Date/Time: August 26, 2009 2:42:03 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00137I mkflash: FlashCopy pair 4800:4803 successfully created.
CMUC00137I mkflash: FlashCopy pair 4801:4804 successfully created.
CMUC00137I mkflash: FlashCopy pair 4802:4805 successfully created.
CMUC00137I mkflash: FlashCopy pair 4900:4903 successfully created.
CMUC00137I mkflash: FlashCopy pair 4901:4904 successfully created.
CMUC00137I mkflash: FlashCopy pair 4902:4905 successfully created.
dscli>
```

Figure 6-19 Re-establish of FlashCopy from B to C volumes

7. The data on B volumes is now consistent and ready for recovery. You can IPL the IBM i DR partition, which is connected to those B volumes through VIOS NPIV. Before doing so, check that the correct virtual FC adapter is tagged for Load source and that IPL source is set to A or B, as is described in 5.4, “Establishing Metro Mirror” on page 43.

   As expected, we experience an abnormal IPL at the DR site. The DR partitions are connected to Global Mirror B volumes through two VIOS NPIVs in multipath; this can be seen in disk resource names starting with DMP. The LUNs in our DR partition after recovery can be seen in Figure 6-20 on page 67.

```
Display Disk Configuration Status

          Serial                        Resource                        Hot Spare
ASP Unit  Number          Type Model    Name       Status               Protection
  1                                               Unprotected
      1   50-4902001      2107  A04     DMP012     RAID 5/Active             N
      2   50-4800001      2107  A04     DMP013     RAID 5/Active             N
      3   50-4801001      2107  A04     DMP014     RAID 5/Active             N
      4   50-4802001      2107  A04     DMP015     RAID 5/Active             N
      5   50-4900001      2107  A04     DMP016     RAID 5/Active             N
      6   50-4901001      2107  A04     DMP017     RAID 5/Active             N



Press Enter to continue.

F3=Exit        F5=Refresh         F9=Display disk unit details
F11=Disk configuration capacity   F12=Cancel
```

Figure 6-20 IBM i disks in multipath in DR partition

Note: The DR IBM i partition now contains a clone of the production IBM i. Refer to the considerations described in 4.2.2, “Cloning an IBM i partition” on page 31.

6.4.2 Returning to the local site

After the production site is back, you can resume normal work at the local site by performing the following steps:

1. If you have not yet done so, create *remote mirror and copy* paths from the secondary to the primary DS8000, as described in 5.4, “Establishing Metro Mirror” on page 43.
2. At the secondary DS8000, re-establish the Global Copy from GM volumes B to A. This step re-synchronizes the data changes that took place at the recovery site during the outage. We use a **failbackpprc** command at secondary DS8000, shown in Figure 6-21.

```
dscli> failbackpprc -type gcp 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 7:05:55 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4800:4800 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4801:4801 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4802:4802 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4900:4900 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4901:4901 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4902:4902 successfully
failed back.
```

Figure 6-21 Failback Global Copy from B to A

3. Wait until the data is synchronized from volumes B to A. To check the synchronization status query the Global Copy out of synch tracks, by using the **lspprc -l** command at the secondary DS8000. See Figure 6-22.

```
dscli> lspprc -l 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 7:15:25 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-75ABTV1
ID        State        Reason Type        Out Of Sync Tracks Tgt Read Src
Cascade Tgt Cascade Date S
==========================================================================
====================
4800:4800 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
4801:4801 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
4802:4802 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
4900:4900 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
4901:4901 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
4902:4902 Copy Pending -     Global Copy 0                  Disabled Disabled
Invalid     -
```

Figure 6-22 Query out of synch. tracks at GC failback

After the data is synchronized, the number of `Out Of Sync Tracks` all show as 0 (zero), and you can proceed with the next step.

4. Power down the IBM i partition at the DR site to ensure that all the data is copied to disk and is consistent. If the downtime caused by the IBM i shutdown is not acceptable, you may quiesce data to disk by using the command CHGASPACT *SYSBAS command, as described in 4.2.1, “Quiescing IBM i data to disk” on page 30. In our example, we power down the IBM i LPAR that is connected to the secondary DS8000.

   Wait until the data from volumes B to A is synchronized: the Global Copy `Out Of Sync Tracks` indicate 0 (zero).

5. At the primary DS8000, fail over Global Copy to the A volumes, as shown in Figure 6-23, so that their status changes to `Suspended Host Source`.

```
dscli> failoverpprc -type gcp 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 7:47:27 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4800:4800 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4801:4801 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4802:4802 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4900:4900 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4901:4901 successfully
reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 4902:4902 successfully
reversed.
dscli>
```

Figure 6-23 Failover Global Copy to volumes A

6. Fail back Global Copy from volumes A to B to re-establish the Global Copy direction from primary to secondary DS8000. See Figure 6-24.

```
dscli> failbackpprc -type gcp 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 26, 2009 7:56:41 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4800:4800 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4801:4801 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4802:4802 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4900:4900 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4901:4901 successfully
failed back.
CMUC00197I failbackpprc: Remote Mirror and Copy pair 4902:4902 successfully
failed back.
dscli>
```

Figure 6-24 Failback Global Copy from volumes A to B

7. The last step before you can IPL the production IBM i, is to start the Global Mirror session again. (If you did not already create the FlashCopy relationships from B to C volumes, do it before starting the Global Mirror.) To start the Global Mirror session, first check the status for the Global Mirror session for each LSS by using the **`lssession`** command. To start the Global Mirror session use the **`mkgmir`** command. After it is started, view the Global Mirror status with the **`showgm`** command.

See Figure 6-25, Figure 6-26 and Figure 6-27.

```
dscli> lssession 48-49
Date/Time: August 26, 2009 8:01:08 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
LSS ID Session Status Volume VolumeStatus PrimaryStatus       SecondaryStatus
FirstPassComplete
================================================================================
====================
48     01      Normal 4800   Active       Primary Copy Pending Secondary
Simplex True
48     01      Normal 4801   Active       Primary Copy Pending Secondary
Simplex True
48     01      Normal 4802   Active       Primary Copy Pending Secondary
Simplex True
49     01      Normal 4900   Active       Primary Copy Pending Secondary
Simplex True
49     01      Normal 4901   Active       Primary Copy Pending Secondary
Simplex True
49     01      Normal 4902   Active       Primary Copy Pending Secondary
Simplex True
dscli>
```

Figure 6-25 Status of GM session before starting GM

```
dscli> mkgmir -lss 48 -session 01
Date/Time: August 26, 2009 8:09:14 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00162I mkgmir: Global Mirror for session 01 successfully started.
dscli>
```

Figure 6-26 Start Global Mirror

```
dscli> showgmir 48
Date/Time: August 26, 2009 8:12:30 CEST PM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
ID                          IBM.2107-7520781/48
Master Count                1
Master Session ID           0x01
Copy State                  Running
Fatal Reason                Not Fatal
CG Interval Time (seconds)  0
Coord. Time (milliseconds)  50
Max CG Drain Time (seconds) 30
Current Time                08/26/2009 20:12:33 CEST
CG Time                     08/26/2009 20:12:33 CEST
Successful CG Percentage    99
FlashCopy Sequence Number   0x4A957B11
Master ID                   IBM.2107-7520781
Subordinate Count           0
Master/Subordinate Assoc    -
```

Figure 6-27 Show status of Global Mirror

8. IPL the production IBM i partition. If you powered down the DR system before the Global Copy failover, the IPL will be normal, shown in Figure 6-28. If you quiesced the DR IBM i without a power-down operation, the production system IPL will be abnormal.

```
Licensed Internal Code IPL in Progress
                                                          08/26/09  20:20:41
IPL:
  Type  . . . . . . . . . . . . . . . :   Attended
  Start date and time . . . . . . . . :   08/26/09  20:20:39
  Previous system end . . . . . . . . :   Normal
  Current step / total  . . . . . . . :      10     16
  Reference code detail . . . . . . . :   C6004057

IPL step                                  Time Elapsed   Time Remaining
 Journal Recovery                            00:00:01        00:00:00
 IFS Initialization                          00:00:01        00:00:00
>Data Base Recovery
 Journal Synchronization
 Commit Recovery

Item:
  Current / Total . . . . . . . :

Sub Item:
  Identifier  . . . . . . . . . :
  Current / Total . . . . . . . :
```

Figure 6-28 IPL of production IBM i after switch back to production site

6.5 Planned outage with Global Mirror

With Global Mirror, the disaster recovery system typically resides at a long distance from the production system. However, some IT centers might use it for business continuity at planned outages.

Perform the following steps to switch to the DR site at planned outages:

1. Power down the production IBM i system.
2. Stop the Global Mirror session as described in 6.4.1, “Switch over to DR site after the failure on local site” on page 64, step 2).
3. Suspend Global Copy at the primary DS8000. Use a **pausepprc** command, as shown in Figure 6-29 on page 72.

```
dscli> pausepprc 4800-4802:4800-4802 4900-4902:4900-4902
Date/Time: August 27, 2009 12:01:47 CEST AM IBM DSCLI Version: 5.4.30.210 DS:
IBM.2107-7520781
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4800:4800 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4801:4801 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4802:4802 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4900:4900 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4901:4901 relationship
successfully paused.
CMUC00157I pausepprc: Remote Mirror and Copy volume pair 4902:4902 relationship
successfully paused.
```

Figure 6-29 Suspend Global Copy before failover to secondary site

4. Fail over the Global Copy as described in step 3 on page 65 (in 6.4.1, “Switch over to DR site after the failure on local site” on page 64).
5. Because the production System i was powered off before the Global Mirror was stopped, we expect consistent data at the Global Mirror C volumes. Therefore you need to observe the status of the FlashCopy at the secondary DS8000, and revert or commit the FlashCopy.

   Perform a Fast Reverse Restore as described in step 5 on page 66 (in 6.4.1, “Switch over to DR site after the failure on local site” on page 64).

   You have to wait until all Fast Reverse Restore operations and their background copy have successfully completed before proceeding to the next step. Again, when the background copy completes, the FlashCopy relation ends. Therefore, you should check whether any FlashCopy relationships remain to determine when all Fast Reverse Restore operations have completed.
6. Re-establish the FlashCopy relation between volumes B and C, as is described in step 6 on page 66 (in 6.4.1, “Switch over to DR site after the failure on local site” on page 64).
7. IPL the IBM i partition that is connected to Global Mirror B volumes, which continues to work during the outage on the production system.

After the production system is ready to take over again, perform the steps described in 6.4.2, “Returning to the local site” on page 67, to continue working at the production site.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 73. Note that some of the documents referenced here might be available in softcopy only.

- *IBM PowerVM Virtualization Managing and Monitoring*, SG24-7590
- *IBM i and Midrange External Storage,* SG24-7668
- *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i,* SG24-7120
- *IBM System Storage DS8000: Architecture and Implementation,* SG24-6786
- *IBM System Storage DS8000: Copy Services in Open Environments*, SG24-6788

Online resources

These Web sites are also relevant as further information sources:

- IBM System Storage DS8000 Web site

  `http://www.ibm.com/servers/storage/disk/ds8000/`
- IBM System Storage DS8000 Information Center

  `http://publib.boulder.ibm.com/infocenter/ds8000ic/index.jsp`

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

`ibm.com/redbooks`

Help from IBM

IBM Support and downloads

`ibm.com/support`

IBM Global Services

`ibm.com/services`

Index

DS8000 Copy Services for IBM i with VIOS

Redpaper™

Establish multiple paths to multiple VIOS partitions

Learn about VIOS with Node Port ID Virtualization

Understand Disaster Recovery scenarios

IBM System Storage DS8000 Copy Services provide functions for Disaster Recovery and Business Continuity solutions for an IBM i server. Many IBM i customers with natively connected DS8000 use Metro Mirror or Global Mirror functions to provide a set of consistent data at the remote site. This approach allows fail over or restoring data in case of disaster at the production site. You may also use IBM FlashCopy technology to provide point-in-time copy of disk, from which you can back up data to minimize the necessary downtime.

This IBM Redpaper publication is intended for those who want to plan or implement DS8000 Copy Services with IBM i. It discusses the use of DS8000 Copy Services for an IBM i system connected with VIOS Node Port ID Virtualization (NPIV). For this paper, we created test scenarios with FlashCopy, Metro Mirror, and Global Mirror of Sysbas. We used DS8000 DS command-line interface (DS CLI) commands to manage the scenarios.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

REDP-4584-00