IBM i Virtualization and Open Storage (read-me first)

Mike Schambureck and Keith Zblewski
(schambur@us.ibm.com and zblewski@us.ibm.com)

IBM Lab Services – IBM Power Systems, Rochester, MN

August, 2011

This “read-me first” paper provides detailed instructions on using IBM i 6.1/7.1 virtualization and connecting open storage to IBM i. It provides information on the prerequisites, supported

hardware and software, planning considerations, install and post-install tasks such as backups. The paper also contains links to many additional information sources.

1 IBM i virtualization solutions

IBM i 6.1 introduced three significant virtualization capabilities that allow faster deployment of IBM i workloads within a larger heterogeneous IT environment. This section introduces and differentiates these new technologies.

1.1 IBM i logical partition (LPAR) hosting another IBM i partition

An IBM i 6.1/7.1 LPAR can host one or more additional IBM i LPARs, known as virtual client LPARs. Virtual client partitions typically have no physical I/O hardware assigned and instead leverage virtual I/O resources from the host IBM i partition. The types of hardware resources that can be virtualized by the host LPAR are disk, tape, optical and networking. The capability of IBM i to provide virtual I/O resources has been used successfully for several years to integrate AIX®, Linux® and Microsoft® Windows® workloads on the same platform. The same virtualization technology, which is part of the IBM i operating system, can now be used to host IBM i LPARs. IBM i hosting IBM i is the focus of the first half of this paper.

1.2 IBM i using open storage as a client of the Virtual I/O Server (VIOS)

IBM i virtual client partitions can also be hosted by VIOS. VIOS is virtualization software that runs in a separate partition with the purpose to provide virtual storage, optical, tape and networking resources to one or more client partitions. The most immediate benefit that VIOS brings to an IBM i client partition is the ability to expand its storage portfolio to use 512-byte/sector open storage. Open storage volumes (or *logical units*, LUNs) are physically attached to VIOS through a FC or a Serial-attached SCSI (SAS) connection and then made available to IBM i. While IBM i does not directly attach to the storage area network (SAN) in this case, as soon as open storage LUNs become available through VIOS, they are managed the same way as integrated disks or LUNs from a directly attached storage system. IBM i using open storage through VIOS is the focus of the second half of this read-me first guide.

1.3 IBM i on a Power blade

The third major virtualization enhancement with IBM i 6.1 is the ability to run an IBM i LPAR and its applications on a Power blade server, such as IBM BladeCenter®JS12 or JS22. Running IBM i on a Power blade is beyond the scope of this paper. Refer to the *IBM i on a Power Blade Read-me First* for a complete technical overview and implementation instructions: http://www.ibm.com/systems/power/hardware/blades/ibmi.html.

2 IBM i hosting IBM i

2.1 Supported configurations

2.1.1 Hardware

One of the most significant benefits of this solution is the broad hardware support. Any storage, network and optical adapters and devices supported by the host IBM i partition on an IBM POWER6® or IBM POWER7® processor-based server can be virtualized to the client IBM i partition. Virtualization of tape devices from an IBM i host to an IBM i client is supported as of June 2011. Table 1 lists the hardware supported.

Hardware type	Supported for IBM i hosting IBM i	Notes
IBM Power servers	Yes	**Includes** IBM POWER6 AND POWER7 servers **Does not include** IBM POWER6/7 processor-based blade servers.
IBM Power 575/775 servers	No	
IBM POWER5-based systems or earlier	No	
Storage adapters (FC, SAS, SCSI)	Yes	Must be supported by IBM i 6.1 or higher **and** supported on POWER6/7-based IBM Power server
Storage devices and subsystems	Yes	Must be supported by IBM i 6.1 or higher **and** supported on POWER6/7-based IBM Power server
Network adapters	Yes	Must be supported by IBM i 6.1 or higher **and** supported on POWER6/7-based IBM Power server
Optical devices	Yes	Must be supported by IBM i 6.1 or higher **and** supported on POWER6/7-based IBM Power server
Tape devices	No	

Table 1: Hardware supported

To determine the storage, network and optical devices supported on each IBM Power server model, refer to the sales manual for each model at: http://www.ibm.com/common/ssi/index.wss. To determine the storage, network and optical devices supported only by IBM i 6.1, refer to the upgrade planning website at: https://www-304.ibm.com/systems/support/i/planning/upgrade/futurehdwr.html.

2.1.2 Software and firmware

Software or firmware type	Supported for IBM i hosting IBM i	Notes
IBM i 6.1 or higher	Yes	**Required** on both host and client IBM i partition
IBM i 5.4 or earlier	No	Not supported on host or client partition
IBM Power server system	Yes	This is the minimum system firmware level

firmware 320_040_031 or higher		required
HMC firmware V7 R3.2.0 or higher	Yes	This is the minimum HMC firmware level required
SDMC firmware V6 R7.3.0 or higher	Yes	This is the minimum HMC firmware level required

Table 2: Software and firmware requirements

2.2 IBM i hosting IBM i concepts

The capability of an IBM i partition to host another IBM i partition involves hardware and virtualization components. The hardware components are the storage, optical and network adapters and devices physically assigned to the host IBM i LPAR. The virtualization components are the system firmware and IBM i operating system objects necessary to virtualize the physical I/O resources to client partitions. Figure 1 shows the full solution and its components:

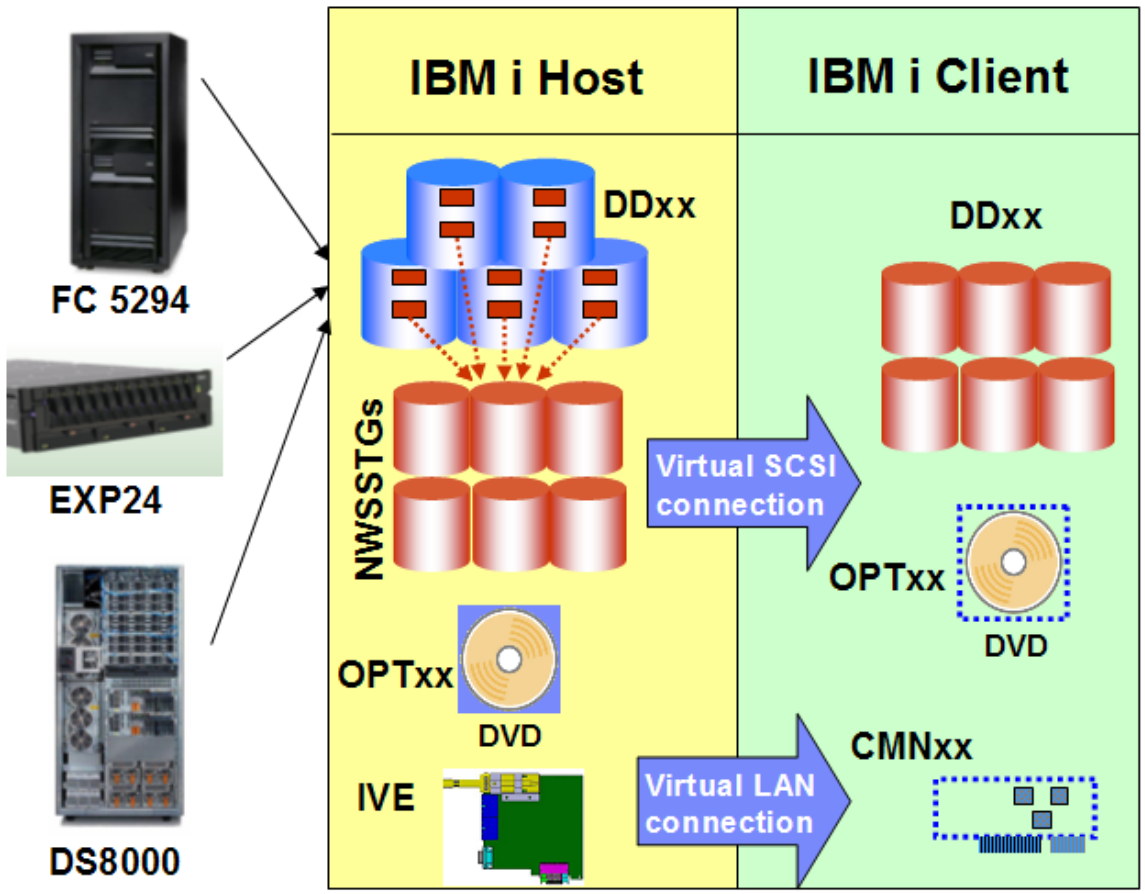

Figure 1: IBM i hosting IBM i components

2.2.1 VSCSI and Ethernet adapters

IBM i hosting IBM i uses an existing function of the system firmware, or IBM Power Hypervisor™: the capability to create VSCSI (VSCSI) and Ethernet adapters in a partition. Virtual adapters are

created for each LPAR in the Hardware Management Console (HMC) or virtual server in the Systems Director Management Console (SDMC). VSCSI adapters are used for storage and optical virtualization; virtual Ethernet adapters are used for network virtualization.

Note that using virtual I/O resources from a host partition does not preclude an IBM i client partition from owning physical hardware. A mix of virtual and physical hardware in the same partition is supported for IBM i in this environment, by assigning both the types of adapters to the partition in the HMC or SDMC.

2.2.2 Storage virtualization

2.2.2.1 Disk virtualization

To virtualize integrated disk (SCSI, SAS or SSD) or LUNs from a SAN system to an IBM i client partition or virtual server, both HMC/SDMC and IBM i objects must be created. In the HMC/SDMC, the minimum required configuration is:

- One VSCSI server adapter in the host partition
- One VSCSI client adapter in the client partition

Note: for the remainder of this paper, the term LPAR (when using an HMC) will also mean virtual server (when using SDMC).
This VSCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one VSCSI pair can exist for the same client partition in this environment. To minimize performance overhead on the host partition, the VSCSI connection is used to send I/O requests, but not for the actual transfer of data. Using the capability of the POWER Hypervisor for Logical Remote Direct Memory Access (LRDMA), data is transferred directly from the physical adapter assigned to the host partition to a buffer in memory of the client partition.

There is no additional configuration required in IBM i in the virtual client partition. In the host partition, the minimum required IBM i setup consists of the following requirements:

- One network server description (NWSD) object
- One network server storage space (NWSSTG) object

The NWSD object associates a VSCSI server adapter in IBM i (which in turn is connected to a VSCSI client adapter in the HMC/SDMC) with one or more NWSSTG objects. At least one NWSD object must be created in the host for each client, though more are supported.

The NWSSTG objects are the virtual disks provided to the client IBM i partition. They are created from available physical storage in the host partition. In the client, they are recognized and managed as standard DDxx disk devices (with a different type and model). Figure 2 shows several storage spaces for a client partition in an IBM i 6.1 host partition:

```
                  Work with Network Server Storage Spaces
                                                              System:   ZA6BP3
Type options, press Enter.
  1=Create   2=Change   3=Copy   4=Delete   5=Display   6=Print   10=Add link
  11=Remove link

                                    Link          Stg
Opt  Name        Server    Seq    Type  Access   Path
_    __________
__   CP10D01     CP10A       2   *DYN  *UPDATE
__   CP10D02     CP10A       3   *DYN  *UPDATE
__   CP10D03     CP10A       4   *DYN  *UPDATE
__   CP10D04     CP10A       5   *DYN  *UPDATE
```

Figure 2: example storage spaces

Figure 3 shows several storage spaces in an IBM i 6.1 client partition:

```
                          Work with Disk Status                      Y213B2P2
                                                              03/26/08  16:43:00
Elapsed time:   00:00:00

                Size    %     I/O   Request   Read  Write   Read  Write    %
Unit  Type      (M)  Used     Rqs  Size (K)    Rqs   Rqs     (K)   (K)  Busy
   1  6B22    34359  14.6      .0        .0     .0    .0      .0    .0     0
   2  6B22    34359   5.6      .0        .0     .0    .0      .0    .0     0
   3  6B22    34359   5.6      .0        .0     .0    .0      .0    .0     0
   4  6B22    34359   5.6      .0        .0     .0    .0      .0    .0     0
   5  6B22    34359   5.6      .0        .0     .0    .0      .0    .0     0
```

Figure 3: Example of what the client partition sees as storage.

Storage spaces for an IBM i client partition do not have to match physical disk sizes; they can be created from 160 MB to 1 TB in size, as long as there is available storage in the host. The 160 MB minimum size is a requirement from the storage management Licensed Internal Code (LIC) on the client partition. For an IBM i client partition, up to 16 NWSSTGs can be linked to a single NWSD, and therefore, to a single VSCSI connection. Up to 32 outstanding I/O operations from the client to each storage space are supported for IBM i clients. Storage spaces can be created in any existing auxiliary storage pool (ASP) on the host, including Independent ASPs. Through the use of NWSSTGs, any physical storage supported in the IBM i host partition on a POWER6-based system can be virtualized to a client partition.

For performance reasons, you might consider creating multiple storage spaces associated with multiple NWSDs. The rule of thumb is six to eight storage spaces for each client partition. This implies that you are also creating multiple sets of VSCSI adapter pairs between the hosting partition and the client partition. Associate each hosting partition's server VSCSI adapter with a separate NWSD by referencing the VSCSI adapter's resource name in the NWSD. Then link storage spaces to the NWSDs. This will supply multiple disk arms for the client partition to use.

2.2.2.2 Tape virtualization

Any tape drive supported in the host IBM i LPAR can be virtualized to an IBM i client LPAR. An existing VSCSI connection can be used, or a new connection can be created explicitly for tape I/O traffic. By default, if a VSCSI connection exists between the host and the client, all physical and virtual TAPxx drives in the host will be available to the client, where they will also be recognized as TAPxx devices. The NWSD parameter **Restricted device resources** can be used to specify which optical devices in the host a client partition **cannot** access.

A virtualized tape drive in the host partition can be used for a D-mode initial program load (IPL) and install of the client partition or applications higher. The client partition will be able to write to the physical media in the drive.

2.2.2.3 Optical virtualization

Any optical drive supported in the host IBM i LPAR can be virtualized to an IBM i client LPAR. An existing VSCSI connection can be used, or a new connection can be created explicitly for optical I/O traffic. By default, if a VSCSI connection exists between host and client, all physical and virtual OPTxx optical drives in the host will be available to the client, where they will also be recognized as OPTxx devices. The NWSD parameter **Restricted device resources** can be used to specify which optical devices in the host a client partition **cannot** access.

A virtualized optical drive in the host partition can be used for a D-mode Initial Program Load (IPL) and installation of the client partition, as well as for installing Program Temporary Fixes (PTFs) or applications higher. If the optical drive is writeable, the client partition will be able to write to the physical media in the drive.

2.2.3 Network virtualization

Virtualizing a network adapter and using a virtual LAN (VLAN) for partition-to-partition communication within a system are existing IBM i capabilities. In order for a client to use a host's physical network adapter, a virtual Ethernet adapter must be created in the HMC in both partitions or the SDMC in both virtual servers. To be on the same VLAN, the two virtual Ethernet adapters must have the same Port Virtual LAN ID (PVID). This type of adapter is recognized by IBM i as a communications port (CMNxx) with a different type (268C). In the host partition, the virtual Ethernet adapter is then associated with the physical network adapter via a routing configuration – either Proxy Address Resolution Protocol (ARP) or network address translation (NAT). This allows the client partition to send network packets via the VLAN and through the VLAN and through the physical adapter to the outside LAN. The physical adapter can be any network adapter supported by IBM i 6.1, including Integrated Virtual Ethernet (IVE) ports, also known as Host Ethernet Adapter (HEA) ports.

2.3 Planning for implementing IBM i hosted LPARs

2.3.1 Storage planning

Because virtual disks for the IBM i client LPAR are NWSSTG objects in the host LPAR, the main prerequisite to installing a new client LPAR is having sufficient capacity in the host to create those objects. Note that the host partition is not capable of detecting the percent of the virtual storage that is used by the client. For example, if a 500-GB storage space is created, it occupies that amount of physical storage in the host IBM i LPAR, even if the disk capacity is only 50% utilized in the client LPAR. The number of NWSSTG objects is important for the client IBM i partition to allow for concurrent disk IO.

It is recommended to closely match the total size of the storage spaces for each client partition to its initial disk requirements. As the storage needs of the client partition grow, additional storage spaces can be dynamically created and linked to it on the host partition. On the client, the new virtual disk will automatically be recognized as a nonconfigured drive and can be added to any existing ASP. The only restriction to consider in this case is the maximum number of storage spaces allowed per VSCSI connection for an IBM i client partition, which is 16. If more than 16

NWSSTGs are needed for a client LPAR, additional VSCSI connections can be created dynamically in the HMC or the SDMC.

2.3.2 Performance

As described in the "IBM i logical partition (LPAR) hosting another IBM i partition" section, disk I/O operations in an IBM i virtual client partition result in I/O requests to the physical disk adapters and drives that are assigned to the host partition. Therefore, *the best way to ensure good disk performance in the client LPAR is to create a well-performing disk configuration in the host LPAR*. As the host partition is a standard IBM i partition, all the recommendations in the *Performance Capabilities Reference* manual at: (http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html) will apply to it. Use the suggestions provided in the manual for maximizing IBM i disk performance for the type of physical storage used in the host, whether it is integrated disk or SAN.

Note that if only the system ASP exists on the host partition, NWSSTG objects are created on the same physical disk units as all other objects. If the host partition is running production applications in addition to providing virtual storage to client partitions, there will be disk I/O contention as both client partitions and IBM i workloads in the host send I/O requests to those disk units. To minimize disk I/O contention, create storage space objects in a separate ASP on the host (Independent ASPs are supported). Performance on the client(s) would then depend on the disk adapter and disk configuration used for that ASP. If the host partition is providing virtual storage to more than one client partition, consider using separate ASPs for the storage space objects for each client. This recommendation should be weighed against the concern of ending up with too few physical disk arms in each ASP to provide good performance.

Disk contention from IBM i workloads in the host LPAR and virtual client LPARs can be eliminated if a separate IBM i LPAR is used just for hosting client LPARs. An additional benefit of this configuration is the fact that an application or OS problem stemming from a different workload on the host cannot negatively affect client partitions. These benefits should be weighed against:

- The license cost associated with a separate IBM i partition
- The maintenance time required for another partition, such as applying Program Temporary Fixes (PTFs)
- The ability to create well-performing physical disk configurations in both partitions that meet the requirements of their workloads

If the host partition runs a heavy-I/O workload and the client partitions also have high disk response requirements, it is strongly recommended to consider using a separate hosting partition, unless separate ASPs on the host are used for storage space objects. If the host partition's workload ranges from light to moderate with respect to disk requirements and the client partitions are used mostly for development, test or quality assurance (QA), it is acceptable to use one IBM i partition for both tasks.

2.3.3 Dual hosting

An IBM i client partition has a dependency on its host: if the host partition fails, IBM i on the client will lose contact with its disk units. The virtual disks would also become unavailable if the host partition is brought down to restricted state or shut down for scheduled maintenance or to apply PTFs. To remove this dependency, two host partitions can be used to simultaneously provide virtual storage to one or more client partitions.

The configuration for two hosts for the same client partition uses the same concepts as that for a single host described in the "IBM i logical partition (LPAR) hosting another IBM i partition" section.

In addition, a second VSCSI client adapter exists in the client LPAR, connected to a VSCSI server adapter in the second host LPAR. The IBM i configuration of the second host mimics that of the first host, with the same number of NWSD and NWSSTG objects, and NWSSG objects of the same size. As a result, the client partition recognizes a second set of virtual disks of the same number and size. To achieve redundancy, adapter-level mirroring is used between the two sets of storage spaces from the two hosts. Thus, if a host partition fails or is taken down for maintenance, mirroring will be suspended, but the client partition will continue to operate. When the inactive host is either recovered or restarted, mirroring can be resumed.

2.3.4 Client IBM i LPARs considerations and limitations

Refer to the topic *Considerations and limitations for i5/OS client partitions on systems managed by the Integrated Virtualization Manager (IVM)* in the Information Center at: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/rzahcbladei5limits.htm. While in this case the IBM i client partition is not being managed by IVM, it does virtual I/O resources and the limitations outlined in the topic above apply to it.

2.4 Implementing IBM i client LPARs with an IBM i host

Installing IBM i in a client LPAR with an IBM i host consists of two main phases:

- Creating the VSCSI configuration in the HMC orSDMC
- Creating the NWSSTG and NWSD objects in the IBM i host partition, and activating the new client partition

The implementation steps are described in detail in the topic *Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC* in the *Power Systems Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf. Note that sufficient available capacity is required in the IBM i host partition to create the storage space objects. When following the detailed implementation instructions, keep in mind the performance recommendations explained in the “Performance section of this paper.

2.5 Post-install tasks and considerations

2.5.1 Configure IBM i networking

Once the IBM i client partition is installed and running, the first system management step is to configure networking. There are three types of network adapters that can be assigned to an IBM i client partition:

- A standard physical network adapter in a Peripheral Component Interconnect (PCI) slot
- A logical port on a HEA
- A virtual Ethernet adapter

Note that both physical and virtual I/O resources can be assigned to an IBM i virtual client partition. If a physical network adapter was not assigned to the IBM i client partition when it was first created, refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Power Systems Logical Partitioning Guide* at: (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf) to assign an available adapter.

An IBM i client partition can also use the new HEA capability of POWER6 processor-based servers. To assign a logical port (logical HEA or LHEA) on an HEA to an IBM i client partition,

refer to the topic **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** in the *Power Systems Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

Lastly, a virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, refer to the topic **Configuring a virtual Ethernet adapter using the HMC** in the *Power Systems Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

In all three cases, the assigned network adapter will be recognized as a communications port (CMNxx) in IBM i. The type of communications port will depend on the network adapter: for example, 5706 for a Gigabit Ethernet adapter, 5623 for an logical HEA and 268C for a virtual Ethernet adapter. In the case of a standard PCI network adapter or an logical HEA, networking can be configured following the process described in the IBM i networking topic in the Information Center at: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm.

If the IBM i client partition is using a virtual Ethernet adapter for networking, additional configuration on the IBM i host is required. The virtual Ethernet adapter allows the client partition to communicate only with other partitions whose virtual Ethernet adapters have the same Port Virtual LAN ID (PVID); in other words, partitions on the same virtual LAN within the system. A routing configuration can be created in the IBM i host partition to allow forwarding of network packets from the outside LAN to the client partition on the virtual LAN. That type of virtual network configuration has been used successfully for several years to provide networking to Linux client partitions with an IBM i host. The two methods for routing traffic from the physical LAN to a client partition on a virtual LAN are proxy ARP and NAT. To configure proxy ARP or NAT in the IBM i host partition, follow the instructions in section 5.2 of the IBM Redbook *Implementing POWER Linux on IBM System i Platform* at: (http://www.redbooks.ibm.com/redbooks/pdfs/sg246388.pdf).

2.5.2 How to perform IBM i operator panel functions

Operator panel functions in an IBM i client partition are performed in the HMC:

- Sign onto the HMC with a profile that has sufficient authority to manage the IBM i client partition
- Select the partition
- Use the open-in-context arrow to select **Serviceability → Control Panel Functions**, then the desired function.

2.5.3 Displaying the IBM i partition System Reference Code (SRC) history

- Sign onto the HMC with a profile that has sufficient authority to manage the IBM i client partition.
- Select the partition.
- Use the open-in-context arrow to select **Serviceability → Reference Code History.**
- To display words 2 through 9 of a reference code, select the option specific to that code.

2.5.4 Configuring Electronic Customer Support (ECS) over LAN

A supported WAN adapter can be assigned to the IBM i client partition for ECS. Alternatively, ECS over LAN can be configured. Refer to the topic **Setting up a connection to IBM** in the

Information Center at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm.

2.5.5 Copying storage spaces

Because an IBM i client partition is installed into one or more storage space objects in the IBM i host partition, new client partitions can be deployed rapidly by copying the storage space(s). Note that each IBM i partition, client or host, must have valid OS and Licensed Product Program (LPP) licenses for the number of processors it uses.

To copy one or more storage spaces that contain an installed IBM i client partition, first shut down the partition during an available period of downtime. Next, log into the host IBM i partition with a security officer-level profile and perform the following steps:

- Enter **WRKNWSSTG.**
- Enter **3** next to the storage space you are going to copy and press **Enter.**
- Enter a name that can be a maximum of 10 characters for the new storage space.
- The size of the original storage space will be entered automatically. The new storage space can be as large or larger (a maximum of 1 TB), but not smaller.
- Enter the correct ASP ID. The ASP where the original storage space exists is the default
- Optionally, enter a text description and press **Enter.**

To deploy the new client partition, follow the instructions in the "Implementing IBM i client LPARs with an IBM i host" section to create the necessary VSCSI configuration in the HMC and the NWSD object in the host IBM i partition.

2.5.6 Backups

As mentioned above, an IBM i client partition with an IBM i host can use a mix of virtual and physical I/O resources.

If the hosting partition has the i7.1 TR2 PTFs applied, any physical or virtual tape devices identified by the hosting partition can be identified by the client partition(s) by creating a vSCSI adapter pair between the hosting IBM i and the client IBM i, or by using the one that already exists,

Another approach is to assign an available tape adapter on the system to it and treat it as a standard IBM i partition. The tape adapter can be any adapter supported by IBM i on IBM Power servers and can be shared with other partitions.

To assign an available tape adapter to the IBM i client partition, refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

When a tape adapter connected to a tape drive or library is available to the client partition, refer to the Backup and Recovery topic in the Information Center to manage backups at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int_215989.

The IBM i host partition can also be used for **system-level** backups of the client partition. Refer to the topic **Saving IBM i server objects in IBM i** in the *Logical Partitioning Guide* at:
http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

3 IBM i using open storage supported configurations

3.1 Storage Area Networks 101

There are some basic SAN concepts that need to be understood by a person coming from an integrated disks background:

- The disk drives installed in a SAN are not directly seen by IBM i. They are typically grouped into RAID sets of drives called arrays, extent pools or storage pools. RAID 0, 1,, 5, 6 and 10 are generally supported. The disks are RAIDed in the SAN and should not* be protected on IBM i.
    - * Note, the DS8000 has an option to present unprotected LUNs directly to IBM i, where they would be protected by IBM i.
    - Most SANs support different types of disk drives (ie fibre channel versus SATA) and different disk drive speeds. Typical IBM i workloads work best with faster FC or SAS drives versus SATA drives.
- After the RAID array is created, the SAN interface is used to create logical drives from the space within the array.
    - The user sets the size of the logical drives from a range of a few GBs to almost 2 TBs each. IBM i can handle drives anywhere in this range, except for restrictions on the load source drive (35 GB minimum, 1.9 TB maximum). There is no need to match the size of the supported integrated (internal) disk drives.
    - IBM i performs best with logical drives that are the same size.
    - IBM i has always been architected to perform best with more disk arms. This does not change with SAN disks. You need to create a good number of logical drives, not one large drive. Refer to the Performance section for more details.
- The next step is to map or assign the logical drives to a host server. This server can be VIOS or a direct connection to IBM i (depending upon the SAN being used. Refer to the next two sections for details).
    - Hosts are known by the Fibre Channel World Wide Port Name (WWPN) or the SAS MAC address of the port on the adapter that the host server owns.
    - A separate step, called SAN zoning, has to take place to configure the SAN switches to connect the host adapter ports to the SAN's host ports.
    - As the logical drives are mapped to the host they are assigned a logical unit number or LUN. A LUN is a unique number within an array. Many people refer the logical drive as a LUN.
- On the VIOS host the mapped logical drives are seen as hdisks and then assigned to an IBM i partition. It is recommended to assign whole hdisks to IBM i. VIOS supports logical volumes where you can subset an hdisk into multiple volumes to assign to a partition. This is not recommended for IBM i partitions.
- On an IBM i, the VIOS assigned drves or the direct attached mapped logical drives are seen as DDxxx drives. They can be added to any auxiliary storage pool (ASP) that you choose.

3.2 IBM i open storage methods

There are three general methods by which IBM i connects to open storage:

- Directly through the SAN fabric without VIOS.
- As a client of VIOS, with VIOS connecting to a storage subsystem through the SAN fabric.

- As a client of VIOS, with VIOS using storage from an IBM System Storage SAN Volume Controller (SVC). SVC in turn connects to one or more storage subsystems through the SAN fabric.

The first set of support tables in the hardware and software sections below applies to the first connection method above (IBM i direct attachment), while the second set of tables in both sections applies to the second method (IBM i → VIOS → storage subsystem). Furthermore, the support statements in the second set of tables (in the following sections) apply to the end-to-end solution of IBM i using open storage as a client of VIOS. This paper does not attempt to list the full device support of VIOS, nor of any other clients of VIOS, such as AIX and Linux. For the general VIOS support statements, including other clients, refer to the VIOS datasheet at: http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html.

When the second connection method to open storage is used (IBM i → VIOS → SVC → storage subsystem), the support statement for IBM i follows that of SVC. For a list of environments supported by SVC, refer to the data sheet on the SVC overview website at: http://www.ibm.com/systems/storage/software/virtualization/svc.

Note that IBM i cannot connect directly to SVC; it must do so as a client of VIOS.

IBM i as a client of VIOS is also supported on POWER6 or higher processor-based blade servers. IBM i cannot attach directly to open storage when running on Power blades. For the full supported configurations statement for the IBM i on Power blade solution, refer to the BladeCenter Interoperability Guide (BIG) at: http://www-947.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5073016

3.3 Direct attached storage

Hardware type	Supported by IBM i for direct attachment	Notes
IBM Power servers	Yes	**Includes** IBM POWER6 AND POWER7 servers
IBM Power 575, 775	No	
POWER5-based systems or earlier	No	
IBM System Storage DS5100 using FC or Serial Advanced Technology Attachment (SATA) drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS5300 using FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
EXP810 expansion unit with FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
EXP5000 expansion unit with FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival IBM i applications for performance reasons
EXP5060 expansion unit with FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival IBM i applications for performance reasons
DS6800	Yes	
DS8100 using FC or FATA	Yes	**It is strongly recommended** that FATA

drives		drives are used only for test or archival IBM i applications for performance reasons
DS8300 using FC or FATA drives	Yes	**It is strongly recommended** that FATA drives are used only for test or archival IBM i applications for performance reasons
DS8700 using FC or FATA drives	Yes	**It is strongly recommended** that FATA drives are used only for test or archival IBM i applications for performance reasons

3.4 VIOS Hosted storage

Hardware type	Supported by IBM i as a client of VIOS	Notes
IBM Power servers	Yes	**Includes** IBM POWER6 AND POWER7 servers
IBM Power 575,775	No	
POWER5-based systems or earlier	No	
DS3200 using SAS or SATA drives.	Yes	Supported by IBM i as a client of VIOS on **both** IBM Power servers and IBM Power blade servers
DS3500 using SAS or SATA drives. Connects through FC or SAS.	Yes	Supported by IBM i as a client of VIOS on **both** IBM Power servers and IBM Power blade servers
DS3950 using FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS4800 using FC drives	Yes	Supported by IBM i as a client of VIOS on **both** IBM Power servers and IBM Power blade servers
DS4800 using SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS4700 using FC drives	Yes	Supported by IBM i as a client of VIOS on **both** IBM Power servers and IBM Power blade servers
DS4700 using SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS3400 using SAS or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS4200 and earlier DS4000 and FAStT models	No	
EXP810 expansion unit attached to DS4800 or DS47000	Yes	Refer to the supported drives comments above
EXP710 expansion unit attached to DS4800 or DS4700	Yes	Refer to the supported drives comments above
SAN Volume Controller (SVC)	Yes	SVC is not supported for direct connection to IBM i
SVC Entry Edition (SVC EE)	Yes	SVC EE is not supported for direct

		connection to IBM i
Storage subsystems connected to SVC	Yes	The supported list for IBM i as a client of VIOS follows that for SVC
DS5020 using FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS5100 using FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS5300 using FC or SATA drives	Yes	**It is strongly recommended** that SATA drives are used only for test or archival of IBM i applications for performance reasons
DS8100 using FC or FATA drives	Yes	**It is strongly recommended** that FATA drives are used only for test or archival of IBM i applications for performance reasons
DS8300 using FC or FATA drives	Yes	**It is strongly recommended** that FATA drives are used only for test or archival of IBM i applications for performance reasons
DS8700 using FC or FATA drives	Yes	**It is strongly recommended** that FATA drives are used only for test or archival of IBM i applications for performance reasons
IBM XIV® Storage System	Yes	
IBM Storwise® V7000 using SSD, SAS or near line SAS drives	Yes	Supported by IBM i as a client of VIOS on **both** IBM Power servers and IBM Power blade servers. **It is strongly recommended** that near line SAS drives are used only for test or archival of IBM i applications for performance reasons
FC adapters	Yes	Must be supported by VIOS*

*For a list of FC adapters supported by VIOS, Refer to the VIOS Datasheet at: http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html.

To verify combinations of adapters, servers and SANS use the System Storage Interoperation Center (SSIC) at: http://www.ibm.com/systems/support/storage/ssic/interoperability.wss

3.5 Software and firmware: Direct attached

Software or firmware type (minimums)	Supported by IBM i for direct attachment	Notes
IBM i 6.1 with LIC 6.1.1 or higher	Yes	
IBM i 6.1 without LIC 6.1.1	No	
IBM i 5.4 or earlier	No	
DS5000 Copy Services	Yes	Refer to the “DS4000 and DS5000” section of this paper for details
IBM i host kit for DS5000	Yes	This license is required in order to attach DS5000 LUNs directly to IBM i
Controller firmware 07.60.26.00 or higher	Yes	
DS Storage Manager 10.60.x5.16 or higher	Yes	

3.6 Software and firmware: VIOS hosted

Software or firmware type (minimums)	Supported by IBM i as a client of VIOS	Notes
IBM i 6.1 or higher	Yes	
IBM i 5.4 or earlier	No	
PowerVM Standard Edition or higher	Yes	Required for VIOS media and activation code
VIOS 1.5 with Fix Pack 10.1 or higher	Yes	This is the minimum level of VIOS required
IBM Power server system firmware 320_040_031 or higher	Yes	This is the minimum system firmware level required
HMC firmware HMC V7 R3.2.0 or higher	Yes	This is the minimum HMC firmware level required
DS4000 controller firmware 6.60.08 or higher	Yes	This is the minimum DS4000 controller firmware level required. **Controller firmware 7.10.23 or higher is supported.**
DS4000 Storage Manager 9.x or higher	Yes	Storage Manager 10.10.x or higher is supported
DS4000 Copy Services	Yes	Refer to the "DS4000 and DS5000" section of this paper for details
DS5000 Copy Services	Yes	Refer to the "DS4000 and DS5000" section of this paper for details
AIX/VIOS host kit license for DS3000/4000/5000, Storwise 7000	Yes	This license is required in order to attach DS3000/4000/5000, Storwise 7000 LUNs to VIOS, and therefore virtualize them to IBM i
SVC code level 4.3.1	Yes	This is the minimum SVC code level required
SVC Copy Services	Yes	Refer to the "SAN Volume Controller (SVC) Copy Services and IBM i" section of this paper for details
XIV firmware 10.0.1b	Yes	This is the minimum XIV firmware level required
XIV Copy Services	Yes	Refer to the "XIV Copy Services and IBM i" of this paper for details

4 IBM i using open storage through VIOS

The capability to use open storage through VIOS extends the IBM i storage portfolio to include 512-byte-per-sector storage subsystems. The existing IBM i storage portfolio includes integrated SCSI, SAS or Solid State disk (SSD), as well as FC-attached storage subsystems that support 520 bytes per sector, such as the IBM DS8000 product line. IBM i cannot currently attach directly to 512-B open storage, because its FC adapter code requires 520 bytes per sector. Therefore, logical units (LUNs) created on open storage are physically connected through the FC fabric to VIOS. As soon as recognized by VIOS, the LUNs are virtualized to one or more IBM i virtual client partitions. Figure 4 illustrates both the physical and virtualization components of the solution:

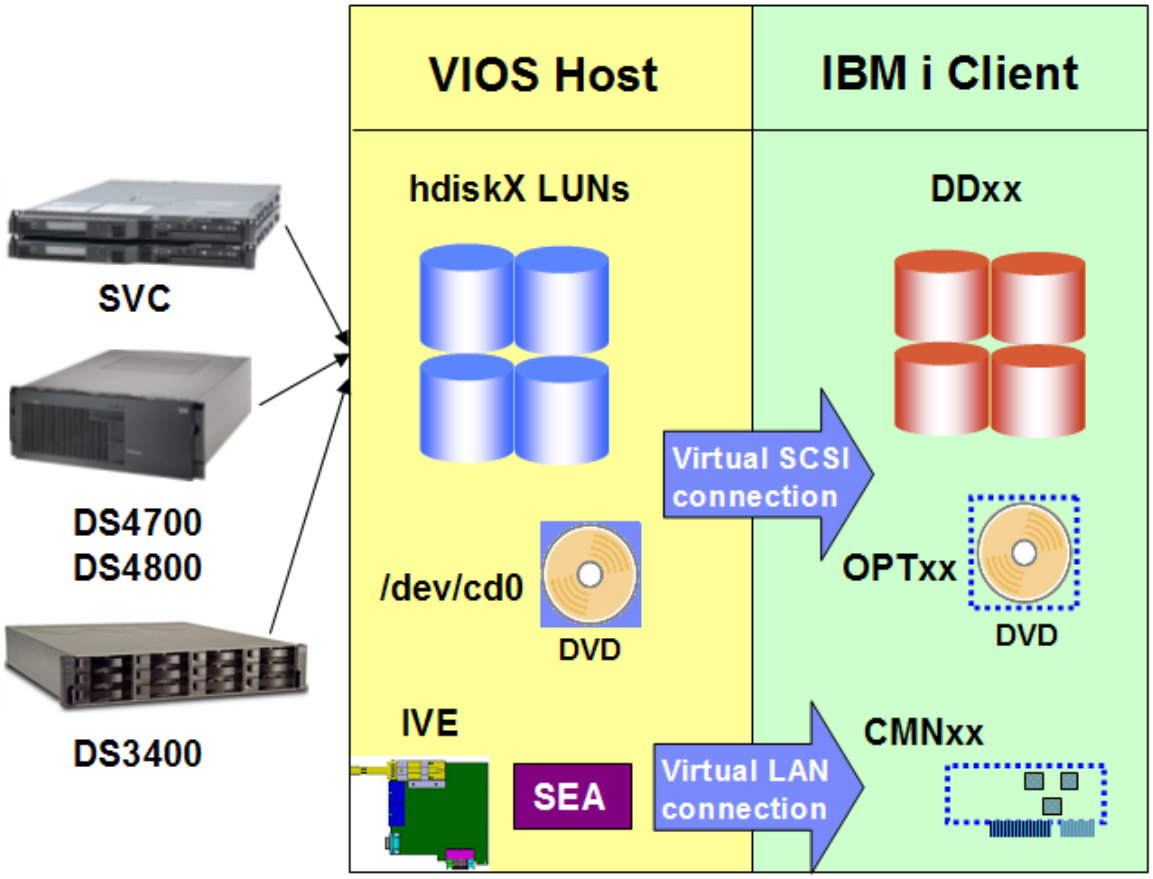

Figure 4: physical and virtualization components of the solution

VIOS has been used successfully for several years to virtualize storage, optical and networking resources to AIX and Linux client partitions. Now IBM i joins this virtualization environment, gaining the ability to use open storage. From a VIOS perspective, IBM i is another client partition; the host-side configuration steps are the same as for AIX and Linux clients. While code changes were made in VIOS to accommodate IBM i client partitions, there is no special version of VIOS in use for IBM i. If you have existing skills in attaching open storage to VIOS and virtualizing I/O resources to client partitions, they will continue to prove useful when creating a configuration for an IBM i client partition.

The hardware and virtualization components for attaching open storage to IBM i illustrated in Figure 4 also apply to using DS5000, XIV and other subsystems supported for this solution, as listed in the "IBM i using open storage supported configurations" section.

Three management interfaces are available for virtualizing I/O resources through VIOS to IBM i: the Hardware Management Console (HMC), the Systesms Director Management Console (SDMC) and the Integrated Virtualization Manager (IVM). All provide logical partitioning and virtualization management functions, however, their support on Power servers and Power blades is different:

- The HMC is supported with all POWER6 and POWER7 servers but not supported with Power blades.
- The SDMC is supported with all POWER6 and POWER7 servers and is supported with Power blades
- IVM is supported with all Power blades and only with Power 520 and Power 550

There is another significant difference between HMC and IVM when managing I/O virtualization for IBM i: when the Power server is IVM-managed, IBM i partitions must be purely virtual. They are not allowed to own any physical I/O adapters. Therefore, when managing Power servers, an HMC or SDMC is used most of the time, with IVM being used occasionally only on Power 520 and Power 550. SDMC is just being introduced as this paper is being updated. When managing Power blades, IVM or SDMC can be used.

4.1 VSCSI and Ethernet adapters

VIOS providing storage to an IBM i client LPAR uses an existing function of the system firmware, or POWER Hypervisor: the capability to create VSCSI and Ethernet adapters in a partition. Virtual adapters are created for each LPAR in the HMC. VSCSI adapters are used for storage and optical virtualization; virtual Ethernet adapters are used for network virtualization.

Note that using virtual I/O resources from VIOS does not preclude an IBM i client partition from owning physical hardware when the server is managed by an HMC or and SDMC. A mix of virtual and physical hardware in the same partition is supported for IBM i in this environment, by assigning both types of adapters to the partition in the HMC/SDMC.

Using IVM for LPAR and virtualization management mandates using only virtual resources for IBM i partitions. However, IVM provides the benefit of automatically creating VSCSI server (in VIOS) and client adapters (in IBM i) as virtual disk, optical and tape resources are assigned to IBM i. For more information on LPAR management with IVM, refer to the IVM topic in the *PowerVM Editions Guide*, available at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm.

Note that IVM will automatically create only the first pair of VSCSI server and client adapters when disk and optical resources are virtualized to IBM i. Additional adapter pairs must be created by using the IVM command line, available through a Telnet or Secure Shell (SSH) session to VIOS. Refer to the VIOS topic in the *PowerVM Editions Guide* for more information on using the IVM/VIOS command line. IVM will automatically create a new VSCSI adapter pair for each tape resource assigned to IBM i.

SDMC offers interfaces to manually create virtual adapters and does not create them by default as IVM does.

4.2 Storage virtualization

4.2.1 SDMC, HMC and IVM configuration

For VIOS to virtualize LUNs created on open storage to an IBM i client partition, configuration takes place both in the HMC, SDMC or IVM and in VIOS. No explicit VSCSI configuration is necessary in IVM the majority of the time.

In the HMC/SDMC, the minimum required configuration is:

- One VSCSI server adapter in the host partition
- One VSCSI client adapter in the client partition

This VSCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one VSCSI pair can exist for the same client partition in this environment. To minimize performance overhead in VIOS, the VSCSI connection is used to send I/O requests, but not for the actual transfer of data. Using the capability of the POWER Hypervisor for Logical Remote Direct Memory Access (LRDMA), data is transferred directly from the FC adapter in VIOS to a buffer in memory of the IBM i client partition.

In an IBM i client partition, a VSCSI client adapter is recognized as a type **290A DCxx** storage controller device. Figure 5 depicts the VSCSI client adapter, as well as several open storage LUNs and an optical drive virtualized by VIOS:

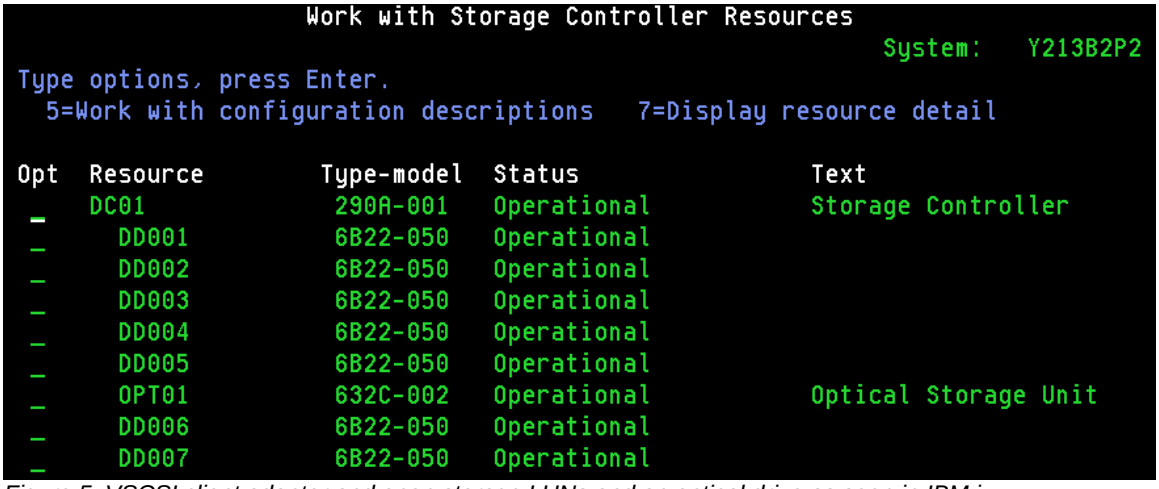

Figure 5: VSCSI client adapter and open storage LUNs and an optical drive as seen in IBM i.

In VIOS, a VSCSI server adapter is recognized as a **vhostX** device:

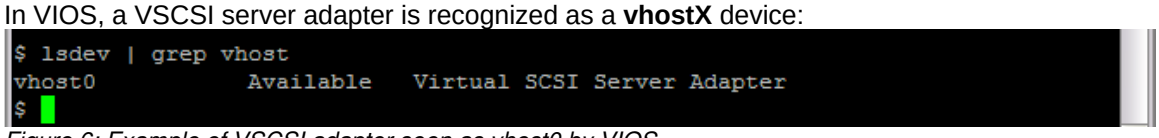

Figure 6: Example of VSCSI adapter seen as vhost0 by VIOS.

4.2.2 VIOS configuration of virtual storage

In VIOS, however, a new object must be created for each open storage LUN that will be virtualized to IBM i: a *virtual target SCSI device*, or *vtscsiX*. A vtscsiX device makes a storage

object in VIOS available to IBM i as a standard DDxxx disk unit. There are three types of VIOS storage objects that can be virtualized to IBM i:

- Physical disk units or volumes (*hdiskX*), which are open storage LUNs in this case
- Logical volumes (*hdX and other*)
- Files in a directory

For both simplicity and performance reasons, **it is recommended to virtualize open storage LUNs to IBM i directly as physical devices (hdiskX)**, and not through the use of logical volumes or files. (Refer to the "Performance" section for a detailed performance discussion.) A vtscsiX device links a LUN available in VIOS (hdiskX) to a specific VSCSI adapter (vhostX). In turn, the VSCSI adapter in VIOS is already connected to a client SCSI adapter in the IBM i client partition. Thus, the hdiskX LUN is made available to IBM i through a vtscsiX device. What IBM i storage management recognizes as a DDxxx disk unit is not the open storage LUN itself, but the corresponding vtscsiX device. The vtscsiX device correctly reports the parameters of the LUN, such as size, to the virtual storage code in IBM i, which in turn passes them on to storage management.

Multiple vtscsiX devices, corresponding to multiple open storage LUNs, can be linked to a single vhostX VSCSI server adapter and made available to IBM i. It is possible to virtualize up to 16 LUNs to IBM i through a single VSCSI connection. This is an IBM i limit, not VIOS. Each LUN typically uses multiple physical disk arms in the open storage subsystem. If more than 16 LUNs are required in an IBM i client partition, an additional pair of VSCSI server (VIOS) and client (IBM i) adapters must be created in the HMC/SDMC or on the IVM/VIOS command line. Additional LUNs available in VIOS can then be linked to the new vhostX device through vtscsiX devices, making them available to IBM i.

Prior to May 2009, creating vtscsiX devices and thus virtualizing open storage LUNs to IBM i was necessarily a task performed only on the IVM/VIOS command line when the server was HMC-managed. When the server is IVM-managed, assignment of virtual resources is performed using the IVM browser-based interface for the first 16 devices for a client partition. However, **the HMC interface has been significantly enhanced, making the use of the IVM/VIOS command line rare** for the first 16 devices for a client partition**.**

For HMC/SDMC provisioning of VIOS storage to work, the VIOS LPAR must have an active Resource Monitoring and Control (RMC) connection to the HMC. The same connection is required for performing dynamic LPAR (DLPAR) changes to the VIOS LPAR and it generally involves successful TCP/IP communication between the LPAR and the HMC. If the HMC displays an RMC error while attempting VIOS storage provisioning, refer to the *DLPAR Checklist* at: http://www.ibm.com/developerworks/systems/articles/DLPARchecklist.html

If the storage configuration is performed while VIOS is down or not yet installed, new storage will be discovered upon boot and the HMC can then be used to assign the correct LUNs to IBM i. If storage changes are made while VIOS and IBM i are running, it will still be necessary to use Telnet or PuTTy to VIOS and run the **cfgdev** command to discover newly added storage prior to using the HMC for assigning it to IBM i.

As soon as they are available in VIOS, open storage LUNs and optical devices can be assigned to IBM i using one of the following options:

4.2.2.1 Assigning storage using IVM

- Click on **View/Modify Virtual Storage.**
- Click the **Physical Volumes** tab.
- Select the hdisks that you want to assign to the IBM i client.

- Click **Modify Partition Assignment**.
- Select the IBM client partition from the list and click **OK.**

4.2.2.2 Assigning storage using the HMC

- Sign onto the HMC as hscroot or another superadministrator-level user ID.
- Expand **Systems Management.**
- Click **Servers.**
- Select the correct managed server.
- In the menu below, expand **Configuration and then** then **Virtual Resources.**
- Click **Virtual Storage Management.**
- Select the correct VIOS from the list and click **Query VIOS.**
- Click the **Physical Volumes** (or **Optical Devices** for optical resources).
- Select the correct LUNs (up to 16) or optical devices (up to 16) and click **Modify assignment…** to assign them to IBM i. Additional LUNs can be selected, as AIX does not limit you to 16, but IBM i will not be able to see them.
- Note: For more than 16 LUNs, you will need to use the make virtual device (mkvdev) command from the VIOS command line interface.

Note that this sequence assumes the existence of the correct VSCSI server and client adapters in VIOS and IBM i, respectively. After the assignment is performed, the new LUNs are immediately available in IBM i as nonsconfigured drives.

4.2.2.3 Assigning storage using SDMC

SDMC can also be used to assign hdisks to the IBM i client partitions (up to 16 per VSCSI adapter pair. Perform the following steps to assign storage using SDMC.

- On the Resources link of the welcome page, expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server for VIOS, right-click on the name and click **System Configuration -> Manage Virtual Server**.
- Click **Storage Devices** on the left-hand navigation panel. The physical volumes (hdiskX) are shown on the right.
- Select the hdisks that you want to assign to the IBM i client and click **Assign**.
- Select the client IBM i partition from the list and click **OK**.

One use of the command line is to add more details to the LUNs to better manage them. The make virtual device (mkvdev) command has a device parameter that can be used to add more descriptive names. Refer to the VIOS command reference for more details at:
http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/iphcg/iphcg.pdf

If the QAUTOCFG system value in IBM i is set to **1** (which is the default), the new virtual resources will become available in IBM i immediately. No action is required on the IVM/VIOS command line.

4.2.3 Adding devices to VIOS

VIOS does not automatically scan for new devices added after the initial boot is complete. When a new device – such as a LUN on a storage subsystem – is added to VIOS, prompt VIOS to scan for new devices by connecting with Telnet or SSH and using the following command:

- **cfgdev**

- If you are using IVM, click **Hardware Inventory** and then click **Configure devices** to run the cfgdev command.

After detecting the device, VIOS will automatically configure it and make it available for use. It can then be virtualized to the IBM i client partition using the HMC, as described in the previous section. Note that if the SAN configuration is performed before the VIOS partition boots, this step is not necessary, as VIOS will recognize all available devices at boot time.

4.3 Optical virtualization

4.3.1 Physical optical devices

CD and DVD drives supported by VIOS can be virtualized to IBM i directly. Optical drives are recognized in VIOS as cdX devices. A cdX device is linked to a vhostX VSCSI server adapter and made available to IBM i through a *virtual target optical device, vtoptX*. The same optical drive can be virtualized to multiple IBM i client partitions by creating separate vtoptX devices linked to different vhostX adapters for the same cdX device. Only one IBM i client partition can use the physical optical drive at a time. If the physical optical drive is writeable, IBM i will be able to write to it. Similar to LUNs, optical devices can be virtualized to IBM i using the enhanced functions of the HMC/SDMC; it is not necessary to use the IVM/VIOS command line to create the vtoptX devices.

4.3.2 VIOS media repository

VIOS provides a capability similar to that of an image catalog (IMGCLG) in IBM i: a repository of media images on disk. Unlike IMGCLGs in IBM i, a single media repository may exist per VIOS. The media repository allows file-backed virtual optical volumes to be made available to the IBM i client partition through a separate vtoptX device. One immediate benefit to IBM i is the ability to import the IBM i install media into VIOS as ISO images, then install the client partition from the ISO images instead of switching real media in a physical DVD drive.

IVM has an interface to create and load the virtual media under the **View/Modify Storage** task, then clicking the **Optical/CD** tab.

SDMC has an interface to create and load the virtual media library too. Use the following steps to do this:

- On the welcome page, click **Resources**,
- Expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server for VIOS, right-click the name and click **System Configuration -> Manage Virtual Server**.
- Click **Media Devices** on the left-hand navigation panel. The media library interface is shown on the right.
- Click **Add** button to create a new library.

The HMC also has an interface to create the virtual media library:

- You need to start at the Managed Server level. Select the managed server and then expand **Configuration->Virtual Resources->Virtual Storage Management**.
- Select a VIOS server and click **Query VIOS**.
- Expand the **Optical devices** tab and look under **Virtual Optical Media** to create your library.

4.4 Network virtualization

Virtualizing a network adapter and using a VLAN for a partition-to-partition communication within a system are existing Power server capabilities. In order for an IBM i client to use a physical network adapter in VIOS, a virtual Ethernet adapter must be created in both of the partitions in the HMC/SDMC. To be on the same VLAN, the two virtual Ethernet adapters must have the same Port Virtual LAN ID (PVID).

A virtual Ethernet adapter is recognized by IBM i as a communications port (CMNxx) of type 268C:

```
                    Work with Communication Resources
                                                          System:   Y213B2P2
Type options, press Enter.
  5=Work with configuration descriptions   7=Display resource detail

Opt  Resource        Type  Status               Text
 _   CMB03           268C  Operational          Comm Processor
 _     LIN02         268C  Operational          LAN Adapter
 _       CMN02       268C  Operational          Ethernet Port
```

Figure 1: Sample communications resources.

In VIOS, the same Ethernet type of device, entX, is used for logical Host Ethernet ports, physical and virtual Ethernet adapters:

```
$ lsdev | grep ent
ent0             Available   Logical Host Ethernet Port (lp-hea)
ent1             Available   Logical Host Ethernet Port (lp-hea)
ent2             Available   Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent3             Available   Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent4             Available   Virtual I/O Ethernet Adapter (l-lan)
```

Figure 8: Example list of Ethernet ports in VIOS.

VIOS provides virtual networking to client partitions, including IBM i, by bridging a physical Ethernet adapter and one or more virtual Ethernet adapters. The virtualization object that provides this Ethernet bridge is called a *Shared Ethernet Adapter (SEA)*. The SEA forwards network packets from any client partitions on a VLAN to the physical LAN through the physical Ethernet adapter. Because the SEA creates a Layer-2 bridge, the original MAC address of the virtual Ethernet adapter in IBM i is used on the physical LAN. The CMNxx communications port that represents the virtual Ethernet adapter in IBM i is configured with an externally routable IP address and a standard network configuration is used. The physical adapter bridge by the SEA can be any network adapter supported by VIOS, including Integrated Virtual Ethernet (IVE) ports, also known as Host Ethernet Adapter (HEA) ports.

The enhanced virtualization management functions of the HMC also allow for network virtualization without using the IVM/VIOS command line. Perform the following steps in the HMC to work with the virtual network management interface.

- Select the correct managed server
- In the menu below, expand **Configuration**, then **Virtual Resources**
- Click **Virtual Network Management**
- Use the drop-down menu to create or modify the settings of a VLAN (referred to as a *VSwitch*), including which SEA it uses to access the external network

SDMC has an interface to create the virtual network too. Perform the following steps to do this:

- On the welcome page click **Resources**,
- Expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server for VIOS, right-click on the name and click **System Configuration -> Manage Virtual Server**.
- Click **Network** on the left-hand navigation panel. The virtual network interface is shown on the right.
- Use the menu to create or modify the settings of a VLAN. Select the **Use this adapter for Ethernet** check box to make the adapter a shared Ethernet adapter (SEA) to allow access the external network.

Additionally, link aggregation is supported for fail over and load balancing across multiple Ethernet ports. The **cfgassist** menu interface can be used from the VIOS CLI to configure link aggregation, SEAs and TCP/IP.

4.4.1 Redundant VIOS Networking

When implementing redundant VIOS partitions as discussed in the Dual hosting section, you also need to consider the networking used there.
Refer to the following steps at:
https://www-304.ibm.com/support/docview.wss?uid=isg3T1011040

Note that step 3 and step 7 are the mkvdev –sea … command to create the shared Ethernet adapters (SEA) has to have all of the parameters, specifically the –attr parameters, included in the one command. Do not try to change an existing SEA to add the attributes.

Note: Step 5 Create a virtual adapter to be used in the SEA adapter on VIOS2. EX: (ent2)

a. Give the virtual adapter the same VLAN ID (PVID) as VIOS1. EX: "1" .
b. Check the box "access external network".
c. Give the virtual adapter a higher trunk priority. EX: "2"

The Trunk Priority of 2 is critical! The default is 1, which will match the SEA already created in the first VIOS partition, With the same trunk priority, or if you try to create the second SEA without all of the attributes shown in step 7 of the referenced document, your VIO server will create a **broadcast storm** on your network! Your network administrator will not be happy with you.

5 Planning for attaching open storage to IBM i through VIOS

5.1 Storage planning

The first storage planning consideration is to have sufficient available capacity in the open storage subsystem to create the AIX/VIOS LUNs that will be virtualized to IBM i by VIOS. As mentioned in the supported hardware the "IBM i using open storage supported configurations section, it is **strongly recommended that only FC or SAS physical drives are used to create LUNs for IBM i as a client of VIOS**. The reason is the performance and reliability requirements of IBM i production workloads. For non-I/O-intensive workloads or nearline storage, SATA or FATA drives can also be used. This recommendation is not meant to preclude the use of SATA or FATA drives for other clients of VIOS or other host servers; it applies only to production IBM i workloads.

If the storage subsystem requires a host kit for AIX/VIOS to be installed before attaching LUNs to these hosts (as is the case with DS5000 or DS4000), that host kit is also required when virtualizing LUNs to IBM i. As the LUNs are virtualized by VIOS, they do not have to match IBM i integrated disk sizes. The technical minimum for any disk unit in IBM i is 160 MB and the maximum is 2 TB, as measured in VIOS. Actual LUN size is based on the capacity and performance requirements of each IBM i virtual client partition and load source disk restrictions.

Direct attached LUNs for IBM i require an "IBM i host kit" on the storage subsystem.

5.2 Performance

When creating an open storage LUN configuration for IBM i as a client of VIOS, it is crucial to plan for both capacity and performance. As LUNs are virtualized for IBM i by VIOS instead of being directly connected it may seem that the virtualization layer will necessarily add a significant performance overhead. However, internal IBM performance tests clearly show that the VIOS layer adds a negligible amount of overhead to each I/O operation. Instead, the tests demonstrate that **when IBM i uses open storage LUNs virtualized by VIOS, performance is almost entirely determined by the physical and logical configuration of the storage subsystem**.

The IBM Rochester, MN, performance team has run a significant number of tests with IBM i as a client of VIOS using open storage. The resulting recommendations on configuring both the open storage and VIOS are available in the latest *Performance Capabilities Reference* manual (PCRM) at: http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html. Chapter 6 focuses on virtualized storage for IBM i. In most cases, an existing IBM i partition using physical storage will be migrated to open storage LUNs virtualized by VIOS. The recommended approach here is to start with the partition's original physical disk configuration; then create a similar setup with the physical drives in the open storage subsystem on which LUNs are created, while following the suggestions in the PCRM sections.

The commonly used SAN disk sizing tool Disk Magic can also be used to model the projected IBM i performance of different physical and logical drive configurations on supported subsystems. You can work with IBM Techline or your IBM Business Partner for a Disk Magic analysis. The latest version of Disk Magic includes support for multiple open storage subsystems and IBM i as a virtual client of VIOS.

5.3 Dual hosting and multi-path I/O (MPIO)

An IBM i client partition in this environment has a dependency on VIOS: if the VIOS partition fails, IBM i on the client will lose contact with the virtualized open storage LUNs. The LUNs would also become unavailable if VIOS is brought down for scheduled maintenance or a release upgrade. To remove this dependency, two or more VIOS partitions can be used to simultaneously provide virtual storage to one or more IBM i client partitions.

5.3.1 Dual VIOS LPARs with IBM i mirroring

Prior to the availability of redundant VIOS LPARs with clients-side MPIO for IBM i, the only method to achieve VIOS redundancy was to use mirroring within IBM i. This configuration uses the same concepts as that for a single VIOS described in the "Storage virtualization" section. In addition, at least one additional VSCSI client adapter exists in the client LPAR, connected to a VSCSI server adapter in the second VIOS on the same Power server. A second set of LUNs of the same number and size is created on the same or a different open storage subsystem, and connected to the second VIOS. The host-side configuration of the second VIOS mimics that of the first host, with the same number of LUNs (hdisks), vtscsiX and vhostX devices. As a result, the client partition recognizes a second set of virtual disks of the same number and size. To achieve redundancy, adapter-level mirroring in IBM i is used between the two sets of virtualized LUNs from the two hosts. Thus, if a VIOS partition fails or is taken down for maintenance, mirroring will be suspended, but the IBM i client will continue to operate. When the inactive VIOS is either recovered or restarted, mirroring can be resumed in IBM i.

5.3.2 Path redundancy to a single set of LUNs

Note that the dual-VIOS solution above provides a level of redundancy by attaching *two separate sets* of open storage LUNs to the same IBM i client through separate VIOS partitions. It is not an MPIO solution that provides redundant paths to a *single set* of LUNs. There are two MPIO scenarios possible with VIOS that remove the requirement for two sets of LUNs:

- A single VIOS partition using two FC adapters to connect to the same set of LUNs
- Two VIOS partitions providing redundant paths to the same set of LUNs on a single open storage subsystem

5.3.2.1 Path redundancy with a single VIOS

If a VIOS LPAR has two or more FC adapters assigned and the correct host configuration is created in the open storage subsystem, VIOS will have redundant paths to the LUNs connected to the FC adapters. VIOS includes a basic MPIO driver, which has been the default instead of the RDAC (Redundant Disk Array Controller) driver since November 2008. The MPIO driver can be used with any storage subsystem which VIOS supports and is included in a default install. In this case, configuration is required only on the storage subsystem in order to connect a single set of LUNs to both ports on a FC adapter owned by VIOS.

This multi-path method can be configured in two ways: round-robin, or failover. For storage systems such as XIV and DS8000 that support active-active multi-path connections, either method is allowed. Path failover will use one path and leave the other idle until the first path fails, while round-robin will always attempt to use both of the paths, if available, to maximize throughput. For systems that only support active-passive connections, such as IBM System Storage DS5000, DS4000, and DS3000 subsystems, failover is the only method allowed.

The access method can be configured on a per-LUN basis from the VIOS command line through a Telnet or SSH session. To show the current multi-path algorithm, run the following command:

- **lsdev –dev hdiskX –attr algorithm**

To change the multi-path algorithm, run the following command:

- **chdev –dev hdiskX –attr algorithm=round_robin**, or
- **chdev –dev hdiskX –attr algorithm=fail_over**

These commands must be repeated for each hdisk.

5.3.2.2 Redundant VIOS LPARs with client-side MPIO (VSCSI)

Beginning with IBM i 6.1 with Licensed Internal Code (LIC) 6.1.1, the IBM i VSCSI (VSCSI) client driver supports MPIO through two or more VIOS partitions to a single set of LUNs (up to a maximum of eight VIOS partitions). This multipath configuration allows a VIOS partition to fail or be brought down for service without IBM i losing access to the disk volumes as the other VIOS partition(s) remain active. Figure 9 illustrates the new capability for IBM i.

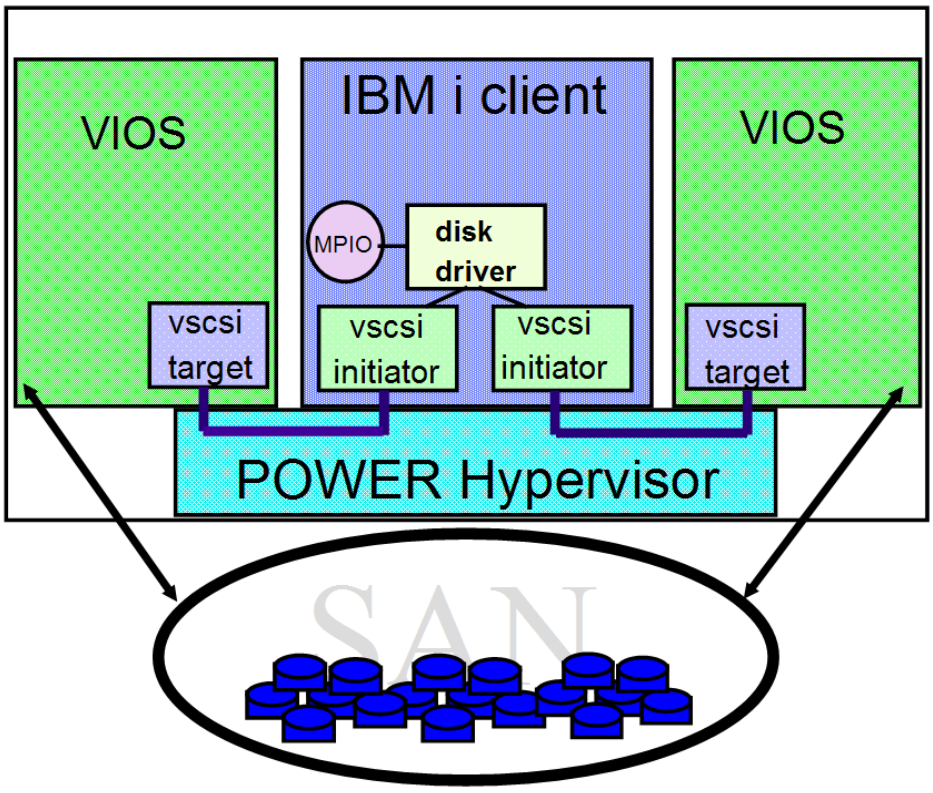

Figure 9: Redundant VIOS server concept.

Note that as with dual VIOS LPARs and IBM i mirroring, a VSCSI connection from IBM i to *both* VIOS partitions is required.On the storage subsystem, it is recommended to create a host for each VIOS containing the worldwide port names (WWPNs) of all FC adapters in that LPAR, and then create a host group comprising those VIOS hosts. A single set of LUNs for IBM i can then be created and mapped to the host group (shared access), giving both of the VIOS LPARs access to the LUNs. Configuring redundant VIOS LPARs with client-side MPIO does not

preclude using MPIO within a VIOS for a higher level of resiliency, as described in the previous section.

For storage systems such as DS8000 and XIV, these connections are typically shared through the round-robin access as described earlier. So it is a good practice to have two or more connections to each VIOS partition.

Note: For D5000, DS4000 and DS3000 storage subsystems with dual controllers, a connection must be made to both of the controllers to allow an active and a failover path. When the volumes are created on these systems, the host OS type should be set to *DEFAULT or AIX* (not *AIX ADT/AVT*, or failover/failback oscillations might occur). For all storage systems, it is recommended that the fabric configuration uses **separate dedicated zones** and FC cables for each connection.

In addition to the VSCSI and storage subsystem configuration, there are several VIOS settings that must be changed. These must be set through a Telnet or SSH session on *both* of the VIOS partitions. The first group of settings is applied to each of the fscsiX devices being used. To show the current fscsiX devices, run the following command:

- **lspath**

Set the fast fail and dynamic tracking attributes on these devices, run the following command:

- **chdev–attr fc_err_recov=fast_fail,dyntrk=yes –perm –dev fscsiX**

Next, restart VIOS. The chdev command must be repeated for each fscsiX device but only one restart is required after all fscsiX attributes have been set.
NOTE: fast_fail is designed for a multipath configuration to not retry failed paths for long periods of time. In a single path configuration, do not set fast_fail since it will not tolerate temporary failures on the single path.

To verify the settings have been changed successfully, run the following command:

- **lsdev –dev fscsiX –attr fc_err_recov**
- **lsdev –dev fscsiX –attr dyntrk**
- Or just **lsdev –dev fscsiX –attr** and review the list of attributes shown.

Importantly, the SCSI reserve policy for *each LUN (or hdisk) on both VIOS LPARs* must be set to *no_reserve* to allow disk sharing. Some storage subsystems such as XIV default to no_reserve and do not require a change, while others such as DS4000 and DS5000 default to *single_path*.
NOTE: The change must be made **prior** to mapping the LUNs to IBM i and it does not require a restart of VIOS.

To show the current reserve policy settings, run the following command:

- **lsdev –dev hdiskX –attr reserve_policy**

To **set** the no_reserve attribute if necessary:

- **chdev –dev hdiskX –attr reserve_policy=no_reserve**
- This is done for every hdiskX

After the VSCSI client driver in IBM i detects a second path to the same set of LUNs through a different VIOS LPAR, the disk names change from DDxxx to DMPxxx. This is identical to multi-

path I/O to a directly attached DS8000 subsystem. If a path is lost, the disk names do not change back to DDxxx, but a path failure message is sent to the QSYSOPR message queue in IBM i. As soon as the path is restored, an informational message is sent to the same queue. To monitor the status of the paths in IBM i:

- Start Services Tools and sign in.
- Work with disk units.
- Display disk configuration.
- Display disk path status.

Note that while the HMC allows a single VSCSI client adapter to be tagged as an IPL device, it is not required to change the tagged resource if that path is lost. As long as the alternate path through the second VIOS is active, the IBM i LPAR will be able to IPL. from the same load source LUN. The alternate path must exist prior to the loss of the original IPL path and prior to powering the IBM i LPAR off.

Another option to VSCSI storage, as just described, is NPIV attached storage facilitated through both VIOS partitions. The sets of commands listed earlier in this section do not come into play because the VIOS partitions become passthru virtual servers.

Also refer to the “Redundant VIOS Networking section for additional recommendations.

5.3.3 Subsystem Device Driver – Path Control Module (SDDPCM)

VIOS also supports the Subsystem Device Driver – Path Control Module (SDDPCM) for certain storage subsystems. Examples of supported subsystems include the SAN Volume Controller (SVC) and IBM DS8000 system. To find out whether a particular storage system supports SDD-PCM for VIOS, refer to its interoperability matrix on the SDDPCM website: https://www-304.ibm.com/support/docview.wss?uid=ssg1S4000201. Note that there are separate support statements for AIX and VIOS. If SDDPCM is supported on your storage subsystem for VIOS, download and install the driver following the instructions in the *Multi-path Subsystem Device Driver User's Guide* at the same location.

6 Attaching open storage to IBM i through VIOS

As described in the “IBM i using open storage through VIOS ” section, IBM i joins the VIOS virtualization environment, allowing it to use open storage. The setup process involves the followingf main steps:

- Open storage configuration
- VIOS installation and configuration
- SAN zoning
- IBM i installation and configuration

The VIOS and open storage configuration steps are the same as for existing clients of VIOS, such as AIX and Linux. These steps are well documented and any existing skills in those areas would apply to IBM i as a client partition, as well. The only significant difference with IBM i is the specific open storage configuration requirements in order to achieve good performance, as referenced in the “Performance” section.

6.1 Open storage configuration

All open storage subsystems follow similar general configuration steps:

- Perform physical disk and Redundant Array of Independent Disks (RAID) configuration
- Create volumes (LUNs)
- Attach those volumes to a host (VIOS in this case)

Naturally, the detailed steps and graphical interfaces used vary by subsystem. Locate the correct documentation for your supported subsystem in IBM Redbooks at: http://www.redbooks.ibm.com. Importantly, follow the instructions for creating and attaching LUNs to an AIX host if you are planning to use VSCSI adapters. For virtual FC the IBM i host type is used. Refer to the associated sections for more information.

6.2 VIOS installation and configuration

Refer to the VIOS planning, installation and configuration topics in the *PowerVM Editions Guide*, available at:
http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm.

As described in the “VIOS configuration of virtual storage” section, as soon as VIOS is installed and physical I/O resources are assigned to it, those resources can be virtualized to IBM i using the HMC/SDMC. It is typically not necessary to use the IVM/VIOS command line to assign virtual disk, optical and network resources to IBM i.

6.3 IBM i using SAN Volume Controller (SVC) storage through VIOS

6.3.1 IBM i and SVC concepts

In October 2008, IBM announced SVC support for IBM i 6.1 as a client of VIOS. The ability to use storage virtualized by SVC further expands the IBM i open storage portfolio and allows clients to leverage their existing SVC infrastructure for IBM i. As mentioned previously, the support statement for IBM i as a client of VIOS includes all storage subsystems that the SVC supports.

The concepts of using SVC storage for IBM i with VIOS are the same as those for other open storage. It is not possible to attach IBM i directly to SVC. Instead, 512-byte-per-sector virtual disks (*VDisks*) are created in SVC and physically connected to VIOS over the SAN fabric. SVC VDisks are recognized in VIOS as hdisks, or physical volumes, similar to LUNs from other open storage subsystems. Available hdisks are then assigned to an IBM i partition by mapping them to a VSCSI server adapter (vhost) in VIOS, which is in turn connected to a VSCSI client adapter in IBM i. As soon as available to the IBM i client partition, SVC VDisks are used and managed no differently from integrated disk units or LUNs. They are assigned a disk type of 6B22, identifying them as virtual disk units provided by VIOS. For a detailed discussion of using virtual storage through VIOS, Refer to the "IBM i using open storage through VIOS " section of this paper.

6.3.2 Attaching SVC storage to IBM i

The general process of attaching SVC storage to IBM i follows the one for using other open storage:

- SVC and zoning configuration
- VIOS installation and configuration
- IBM i installation and configuration

The SVC configurations steps are no different from virtualizing storage to a standard AIX host system. VIOS is installed and configured as for any other client partition, such as AIX or Linux. Existing skills in those areas continue to be applicable when making storage available to IBM i. There are no special versions of SVC and VIOS or additional software necessary for IBM i as a client of VIOS.

Before configuring new storage on SVC, review the planning and performance considerations in chapters 3 and 4 of the IBM Redbook *Implementing the IBM System Storage SAN Volume Controller V5.1* (SG246423), available at: http://www.redbooks.ibm.com/abstracts/sg246423.html?Open. To create VDisks for IBM i on SVC and connect them over the SAN fabric to VIOS, refer to chapters 8, and 9 or 10 of the same Redbook, using the instructions for AIX hosts.

For VIOS installation and configuration, refer to the *PowerVM Editions Guide*, available at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm.

Please Refer to the "Subsystem Device Driver – Path Control Module (SDDPCM)" section when planning for SVC connectivity. This multipath driver is recommended for this connection.

The IBM i client partition configuration as a client of VIOS is the same as that for a client of an IBM i 6.1 host partition. Refer to the **Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC** topic in the *Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

6.4 N_Port ID Virtualization (NPIV) for IBM i

6.4.1 Overview

NPIV is an industry-standard FC protocol that allows VIOS to directly share a single FC adapter among multiple client LPARs. Unlike VSCSI (VSCSI), NPIV does not map a LUN to a virtual target device in VIOS, which the client LPAR can then access as a generic SCSI disk. Instead, a port on the physical FC adapter is mapped to a virtual FC server adapter in VIOS, which in turn is connected to a vitual FC client adapter in IBM i. When the vitual FC client adapter is created, two unique world-wide port names (WWPNs) are generated for it. Through the link to the server vitual

FC adapter and then the physical adapter in VIOS, those WWPNs become available on the SAN and storage can be mapped to them as with any other FC host ports. Note that these WWPNs are unique not just within the Power server, but globally on the SAN. The PowerVM Hypervisor creates two WWPNs per vitual FC client adapter to facilitate Live Partition Mobility for LPARs using NPIV. Live Partition Mobility is currently supported for AIX and Linux LPARs. When a vitual FC client adapter is deleted, the WWPNs are not reused. By default, the PowerVM Hypervisor is capable of creating 32,000 WWPNs for vitual FC client adapters. If additional WWPNs are required, clients can acquire an enablement code from IBM.
Note: If you add vitual FC adapters using DLPAR, use the save configuration to the partition profile option when done. If you instead edit the profile separately, you will generate a new pair of virtual WWPNs.
Figure 10 illustrates using NPIV for IBM i:

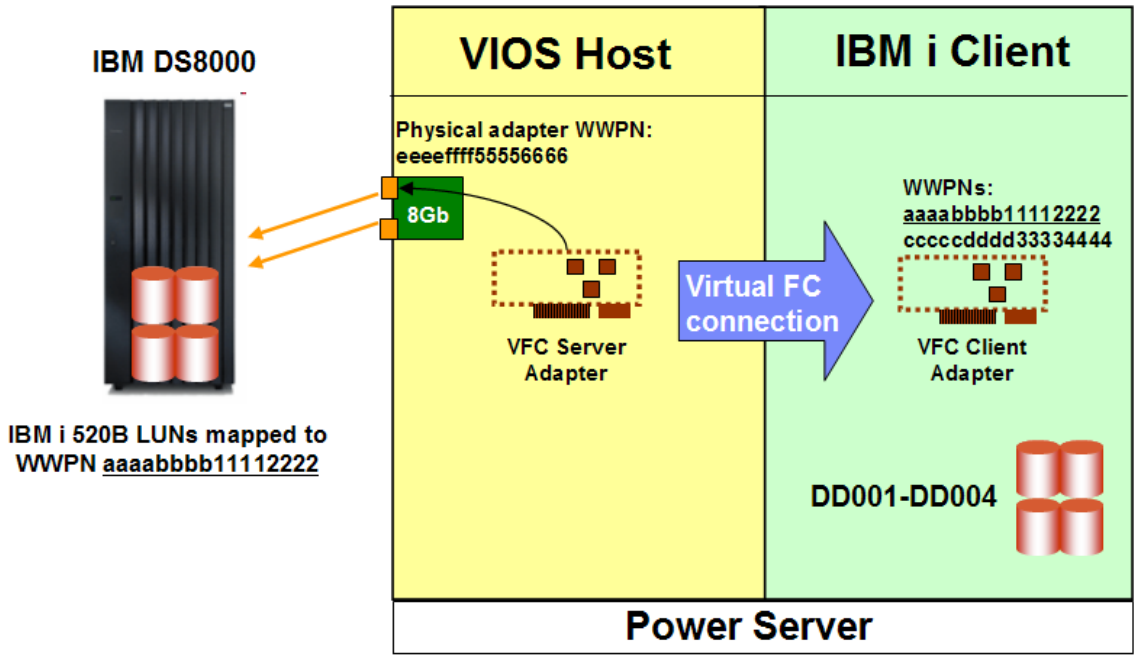

Figure 10: using NPIV for IBM i

From the storage subsystem's perspective, the LUNs, volume group and host connection are created as though the IBM i LPAR is directly connected to the storage through the SAN fabric. While VIOS still plays a role with NPIV, that role is much more of a passthrough one when compared to VSCSI. Note that an 8Gb FC adapter is required; however, the FC switches do not need to be 8Gb. Additionally, only the first switch – the one to which the FC adapter on the Power server is directly connected – must be NPIV-capable. The rest of the SAN fabric does not need to be NPIV-capable.

6.4.2 Supported hardware and software

As of the publication date, NPIV for IBM i supports only the following hardware and software:

- POWER6 processor-based or higher Power servers
- The FC #5735 8Gb dual-port PCI E (PCIe) FC adapter
- The DS8000 storage subsystem series or the DS5100/5300 (Not including POWER blades at this time)
- FC tape libraries:

    - TS3100 (M/T 3573) with LTO3, LTO4 and LTO5 tape drives
    - TS3200 (M/T 3573) with LTO3, LTO4 and LTO5 tape drives
    - TS3310 (M/T 3576) with LTO3, LTO4 and LTO5 tape drives
    - TS3400 (M/T 3577) with TS1120 and TS1130 tape drives
    - TS3500 (M/T 3584) with LTO3, LTO4, LTO5, TS1120, TS1130, TS1140 and 3592-J1A tape drives
    - TS7610 (IBM System Storage TS7610) ProtecTIER® Deduplication Appliance) with software version v2.1
    - TS7650 (IBM System Storage TS7650) ProtecTIER® Deduplication Appliance) with software version v2.4
- IBM i 6.1 with LIC 6.1.1 or higher
- VIOS 2.1.2 or higher

6.4.3 Configuration

There are three general steps in configuring NPIV for IBM i:

- LPAR and VIOS configuration on the Power server
- Storage subsystem or tape library configuration
- SAN zoning

To perform the LPAR and VIOS setup, refer to Chapter 2.9 in the Redbooks *PowerVM Virtualizing Managing and Monitoring* (SG 24-7590) at: http://www.redbooks.ibm.com/abstracts/sg247590.html?Open. While the examples given are for an AIX client of VIOS, the procedure is identical for an IBM i client.

To perform the storage or tape configuration, refer to the Redbooks *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i* (SG24-7120) at: http://www.redbooks.ibm.com/abstracts/sg247120.html?Open or *Implementing IBM Tape in i5/OS* (SG 24-7440) at http://www.redbooks.ibm.com/abstracts/sg247440.html?Open. As mentioned above, from the storage subsystem's or tape library's perspective, the configuration is identical to that for IBM i directly attached through the SAN fabric.
There is a web interface on the tape media library where you need to enable control paths from each device that you want IBM i to be able to work with. Selecting Enable, creates the control paths.
IBM i can not dynamically detect these control paths. To detect the control paths, you need to re-IPL the virtual I/O adapter (IOA).
First determine which virtual IOA has been created for the virtual FC adapters. To do this, enter a WRKHDWRSC *STG command and check for a 6B25 (virtual FC) adapter. Note the IOP/IOA name.
Next, use STRSST command and start a service function (1) -> hardware service manager (7) -> Logical Hardware Resources (2) – System bus resources (1). Enter a 255 in the System bus(es) to work with field and hit enter. This is the virtual adapter bus. Locate the virtual IOA from above and enter a 6 for I/O debug, then option 4 to IPL the IOA. Use F3 to exit SST.
Return to the WRKHDWRSC *STG and use an option 9 to Refer to the tape devices under the VFC IOA. With auto configuration turned on, a new tape device(s) should show up under WRKCFGSTS *DEV TAP*.

Accordingly, DS8000 LUNs might be created as "IBM i protected" or "IBM i unprotected" and will correctly report as such to Storage Management in IBM i. A tape library and a drive within will also report the correct device names and types, such as TAPMLBxx, 3584 and TAPxx, 3580-004 and so on. All tape operations supported with direct FC attachment are supported through NPIV, including hardware tape encryption.

6.5 IBM i installation and configuration

The IBM i client partition configuration as a client of VIOS is the same as that for a client of an IBM i 6.1 host partition. Refer to the **Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC** topic in the *Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

For SDMC use the IBM Systems Director Management Console: Introduction and Overview in IBM Redbooks at: http://www.redbooks.ibm.com/abstracts/sg247860.html?Open

6.6 Post-IBM i install tasks and considerations

6.6.1 Configure IBM i networking

After the IBM i client partition is installed and running, the first system management step is to configure networking. There are three types of network adapters that can be assigned to an IBM i client partition:

- A standard physical network adapter in a PCI slot
- A logical port on a HEA
- A virtual Ethernet adapter

Note that both physical and virtual I/O resources can be assigned to an IBM i virtual client partition. If a physical network adapter was not assigned to the IBM i client partition when it was first created, Refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Logical Partitioning Guide* (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf) to assign an available adapter.

An IBM i client partition can also use the new HEA capability of POWER6 processor-based servers. To assign a logical port (LHEA) on an HEA to an IBM i client partition, Refer to the **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** topic in the *Logical Partitioning Guide* at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

A virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, refer to the topic Configuring a virtual Ethernet adapter using the HMC in the *Logical Partitioning Guide*: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

For the SDMC interfaces for networking refer to the IBM Systems Director Management Console: Introduction and Overview Redbook at: http://www.redbooks.ibm.com/abstracts/sg247860.html?Open

In all cases, the assigned network adapter will be recognized as a communications port (CMNxx) in IBM i. The type of communications port will depends on the network adapter: for example, 5706 for a Gigabit Ethernet adapter, 5623 for an LHEA and 268C for a virtual Ethernet adapter. In the case of a standard PCI network adapter or an LHEA, networking can be configured following the process described in the IBM i networking topic in the Information Center at: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm.

If the IBM i client partition is using a virtual Ethernet adapter for networking, an SEA must be created in VIOS to bridge the internal virtual LAN (VLAN) to the external LAN. Use the

HMC/SDMC and the instructions in section Network virtualization to perform the SEA configuration.

6.6.2 End-to-end LUN device mapping

On October 2009, IBM enhanced both the HMC and VIOS to allow end-to-end device mapping for LUNs assigned to client LPARs, such as IBM i. The new function enables administrators to quickly identify which LUN reporting in VIOS (or, hdisk) is which DDxxx disk device in IBM i. This in turn makes it easier to troubleshoot disk-related problems and safer to change a virtualized disk configuration. In order to correctly perform the mapping, the HMC requires an active RMC connection to VIOS. Such a connection is not required for the IBM i client LPAR, because IBM i does not use RMC for DLPAR operations. However, if the client is an AIX LPAR, an active RMC connection to AIX is also required for a complete end-to-end mapping.
To perform end-to-end LUN device mapping, use the following steps:

- Sign in to the HMC as **hscroot** or another superadministrator-level user ID.
- Expand **Systems Management.**
- Expand **Servers.**
- Click the correct managed system (server).
- Select the correct VIOS by using the checkbox.
- Click **Hardware Information → Virtual I/O Adapters → SCSI.**

6.6.3 How to perform IBM i operator panel functions

If the system is HMC-managed, follow these steps:

- Sign in to the HMC with a profile with sufficient authority to manage the IBM i client partition
- Select the partition.
- Use the open-in-context arrow to select **Serviceability → Control Panel Functions** and then the desired function.

If the system is IVM-managed, follow these steps:

- In IVM, click **View/Modify Partitions** .
- Select the IBM i partition.
- From the **More Tasks** list, select **Operator panel service functions**
- Select the function you wish to perform and click **OK**

If the system is SDMC-managed, perform follow these steps:

- On the SDMC welcome page, select the host you are working with.
- Right-click the host and click **Service and Support-> Control Panel Functions-> (20) Type, Model, Feature**.The function is limited to displaying the hardware information listed.

6.6.4 Displaying the IBM i partition System Reference Code history

If the system is HMC-managed, follow these steps:

- Sign in to the HMC with a profile that has sufficient authority to manage the IBM i client partition.
- Select the partition.
- Use the open-in-context arrow and then select **Serviceability → Reference Code History**.
- To display words 2 through 9 of a reference code, select the option for that code.

If the system is IVM-managed, perform the following steps:

- In IVM, click **View/Modify Partitions**.
- Select the IBM i partition.
- From the **More Tasks** drop-down menu and select **Reference Codes**.
- Click an SRC to display all words.

If the system is SDMC-managed, perform the following steps:

- View the blade's system reference codes from the SDMC welcome page.
- Select the host for the Power server
- Right-click the host and click **Service and Support-> Refernce Code history.**

6.6.5 Client IBM i LPARs considerations and limitations

Refer to the **Considerations and limitations for i5/OS client partitions on systems managed by the Integrated Virtualization Manager (IVM)** topic in the Information Center at: http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahc/rzahcbladei5limits.htm.

6.6.6 Configuring Electronic Customer Support (ECS) over LAN

A supported WAN adapter can be assigned to the IBM i client partition for ECS. Alternatively, ECS over LAN can be configured. Refer to the **Setting up a connection to IBM** topic in the Information Center at:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm.

6.6.7 Backups

IBM i as a client of VIOS can use a mix of virtual and physical I/O resources. One approach is to assign an available tape adapter on the system to the IBM i client partition and treat it as a standard IBM i partition. The tape adapter can be any adapter supported by IBM i on IBM Power servers and can be shared with other partitions. To assign an available tape adapter to the IBM i client partition, refer to the **Managing physical I/O devices and slots dynamically using the HMC** topic in the *Logical Partitioning Guide* at*:*
(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf).
Also refer to the "Tape virtualization" section for another way to assign tapes to the IBM i client.

After a tape adapter connected to a tape drive or library is available to the client partition, refer to the Backup and Recovery topic in the Information Center to manage backups:
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int_215989

7 DS5000 direct attachment to IBM i

7.1 Overview

IBM i has been able to connect to 512-byte-per-sector open storage through VIOS since early 2008 and the solution is implemented in many production environments. On October 2009, IBM further simplified open storage use for IBM i by announcing direct FC attachment to the DS5100 and DS5300 subsystems. With the new enhancement, the IBM i LPAR owns one or more FC adapters, which are connected to the SAN fabric. 51-byte LUNs on the DS5100 or DS5300 are then mapped to the world-wide port names (WWPNs) of the FC adapter(s) in the IBM i LPAR. VIOS is no longer required in this case. Note that one or more VIOS LPARs may still exist on the same Power server to take advantages of other technologies, such as Active Memory Sharing (AMS), NPIV, LPM for AIX and Linux or other I/O virtualization. However, VIOS is no longer required in order for IBM i to use storage from DS5100 and DS5300 system. For example, it is possible to have the same or another IBM i LPAR on the same Power server and still use VIOS for optical virtualization.

A note on terminology: while the new enhancement is referred to as **direct attachment of IBM i to DS5100 and DS5300,** IBM expects that in most cases the attachment will be through a FC switch. While directly attaching FC cables from the adapters on the Power server owned by IBM i to the host channel ports on the storage is supported, this approach decreases the flexibility of the SAN environment. Therefore, **direct attachment** in this case refers to **without VIOS.**

There are two main benefits of directly attaching IBM i to DS5100 and DS5300 systens:

- VIOS skills are not required to implement and use the solution, only IBM i and SAN skills are required.
- The processor, memory, I/O and software expenditure to run a VIOS LPAR is not required

Multipath I/O for IBM i in this environment is supported with a maximum of eight paths. No additional software on the DS5100 or DS5300 systems or IBM i is required. The IBM i FC adapter driver employs the round-robin algorithm by default; it cannot be changed.

7.2 Supported hardware and software

As of the publication date, only the following hardware is supported for IBM i directly attaching to DS5100 and DS5300 systems. This support is also reflected in the “IBM i using open storage supported configurations” section.

- DS5100 (machine-type 1818, model 51A)
- DS5300 (machine-type 1818, model 53A)
- EXP810 expansion unit
- EXP5000 expansion unit
- EXP5060 expansion unit
- FC and SATA HDDs (hard disk drives), with FC drives recommended for production workloads
- POWER6 processor-based or higher Power servers
- FC #5273 4Gb dual-port PCI-e FC adapter
- FC #5276 8Gb dual-port PCI-e FC adapter
- FC #5774 4Gb dual-port PCI-e FC adapter
- FC #5735 8Gb dual-port PCI-e FC adapter

As of the publication date, only the following software or firmware is supported for IBM i directly attached to DS5100 and DS5300 systems. This support is also reflected in "IBM i using open storage supported configurations" section.

- IBM i 6.1 with LIC 6.1.1 or higher
- Controller firmware 07.60.28.00 or higher
- DS Storage Manager 10.60.x5.17 or higher
- IBM i Host Attachment Kit, FC #7735 (required)
- Storage Partitioning (strongly recommended)

7.3 Best practices, limitations and performance

As with other storage subsystems, it recommended to create LUNs for IBM i production workloads on FC drive RAID arrays only in order to meet the performance requirements of the application. While SATA drives are supported for direct attachment to IBM i, they are best suited for test, development and archival of IBM i workloads. Mixing FC and SATA drives within an enclosure is supported along with solid-state drives (SSDs). Hardware encryption of the data on the physical drives within the storage subsystem is supported.

All LUNs mapped directly to IBM i must be protected. Therefore, the supported RAID levels for IBM i direct attachment are 1, 5 and 6. LUNs report as **protected** in IBM i; therefore, IBM i mirroring is not possible. The reliability of the solution is based on the RAID protection within the storage subsystem and on using multiple paths. The maximum LUN size for IBM i in this environment is up to but **not** including 2TB (remember the load source minimum and maximum sizes). A maximum of 64 LUNs per FC adapter port is supported; this statement is identical to that for DS8000 system direct attachment. Dynamic volume (LUN) expansion is not supported for IBM i. After LUNs have been mapped to IBM i, at most 300 from the same subsystem can be added to an ASP with a single operation. More than 300 LUNs can be added to an ASP in general, but they require multiple add operations.

There are no special requirements for the load source LUN, except of course sufficient size to qualify as a load source for IBM i 6.1. The load source does not need to be the very first LUN mapped to IBM i. When performing an Initial Program Load (IPL), an active path to the load source LUN is required. Path failover is not possible during IPL; however, as soon as the IPL is complete, failover functionality resumes.

As with other host types, each port on a FC adapter in IBM i should be zoned separately to controller A or controller B on the storage subsystem. Furthermore, it is strongly recommended to use Storage Partitioning when directly connecting to IBM i. Storage Partitioning is a premium feature on the DS5100 and DS5300 which allows a set of LUNs to be associated only with specified hosts and host ports. Storage Partitioning allows the creation of multiple host groups; only the hosts within a particular host group are allowed access to a group of LUNs. Without that premium feature, all hosts reside in the default host group, which allows IBM i to access LUNs mapped to other hosts. The recommended approach is to use Storage Partitioning and create a host group for each IBM i LPAR using the storage subsystem, placing the hosts associated with that LPAR in the host group. Hosts representing different IBM i LPARs should not be placed in the same host group, as IBM i does not support shared simultaneous access to the same LUNs from multiple LPARs.

Note that migrating from IBM i using DS5100 or DS5300 storage system through VIOS to IBM i using the same subsystem directly requires data migration, typically by save and restore. There is no option to reassign the LUNs from a VIOS host to an IBM i host or to use IBM Tivoli® Storage FlashCopy® Manager or VolumeCopy due to sector size differences..

Performance for IBM i directly attached to DS5100 and DS5300 system is very similar to that of the same adapter and storage configuration attached through VIOS. Direct attachment performance will not be significantly better than that of VIOS attachment for IBM i. Performance results for IBM i directly attached to DS5300 system are in Chapter 5.1.2 of the *Performance Capabilities Reference* manual at: http://www-03.ibm.com/systems/resources/systems_power_software_i_perfmgmt_pcrm_apr2011.pdf.

7.4 Sizing and configuration

There are two main sources of sizing information when planning a configuration involving IBM i directly accessing DS5100 or DS5300 systems: the Disk Magic sizing tool and the *Performance Capabilities Reference* manual mentioned in the previous chapter. It is recommended to use them jointly. While there is not yet a direct attachment model for IBM i with DS5100 or DS5300 system in Disk Magic, the existing VIOS model can still be used to determine the physical storage configuration, such as number of physical drives, number and type of RAID arrays and amount of cache. At the same time, the *Performance Capabilities Reference* manual offers some excellent suggestions on the number of LUNs and paths to use based on the I/Os per second and throughput required. Note that the optimal number of LUNs for IBM i direct attachment is significantly higher than that for VIOS attachment; typically 64 or 128 in the direct case.

Configuration of IBM i direct attachment to DS5100 and DS5300 systems can be divided into two general steps:

- Storage and SAN zoning configuration
- Power server hardware and LPAR configuration

To configure the DS5100 or DS5300 systems, refer to the *IBM Midrange System Storage Hardware Guide* (SG24-7676) from IBM Redbooks available at: http://www.redbooks.ibm.com/abstracts/sg247676.html?Open, while following the best practices and limitations outlined in the previous section of this paper. When creating a host for direct IBM i attachment on the storage subsystem, the host type should be IBM i. To configure the SAN fabric, follow the instructions from your FC switch manufacturer.

To create the IBM i LPAR and assign the FC adapters to it, follow the instructions in the *Logical Partitioning Guide*, available at: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.

For SDMC refer to the IBM Redbooks IBM Systems Director Management Console: Introduction and Overview at: http://www.redbooks.ibm.com/abstracts/sg247860.html?Open

8 Redundant VIOS Virtual Servers hosting IBM i

For a production IBM i virtual server hosted by VIOS, you want redundant VIOS virtual servers in case of failures or scheduled maintenance on the hosting VIOS virtual server. IVM can not be used since it is actually running in VIOS and can only manage that virtual server. For a Power server, an HMC or SDMC can be used. For a POWER blade, only an SDMC can be used. You need to consider the following steps for the configuration of this environment.

- Create two VIOS virtual servers/LPARs using SDMC/HMC. Allocate processor, memory, network adapters and storage adapters for each. For SDMC use the IBM Systems Director Management Console: Introduction and Overview from IBM Redbooks at: http://www.redbooks.ibm.com/abstracts/sg247860.html?Open
- Create an IBM i virtual server/LPAR using SDMC/HMC. Allocate processor and, memory. The networking and storage are typically virtualized, but physical adapters can be assigned. If you are considering moving to IBM i hibernation and eventually to LPM, when it becomes available, you should make the virtual server/LPAR fully virtualized hardware.
- You need shared storage for IBM i that both VIOS virtual servers/LPARs can access. This implies a SAN as integrated disks cannot share access. Review the concepts explained in the "Storage virtualization" section. Select either a VSCSI or virtual FC configuration for the shared storage.
- For networking considerations refer to the "Redundant VIOS Networking" section.

9 Copy Services and IBM i

9.1 DS4000 and DS5000

IBM has conducted some basic functional testing of DS4000 and DS5000 and Copy Services with IBM i as client of VIOS. In this section, you will find information on the scenarios tested and the resulting statements of support for using DS4000 and DS5000 Copy Services with IBM i.

9.1.1 FlashCopy and VolumeCopy

9.1.1.1 Test scenario

Figure 11 shows the test environment used for FlashCopy and VolumeCopy:

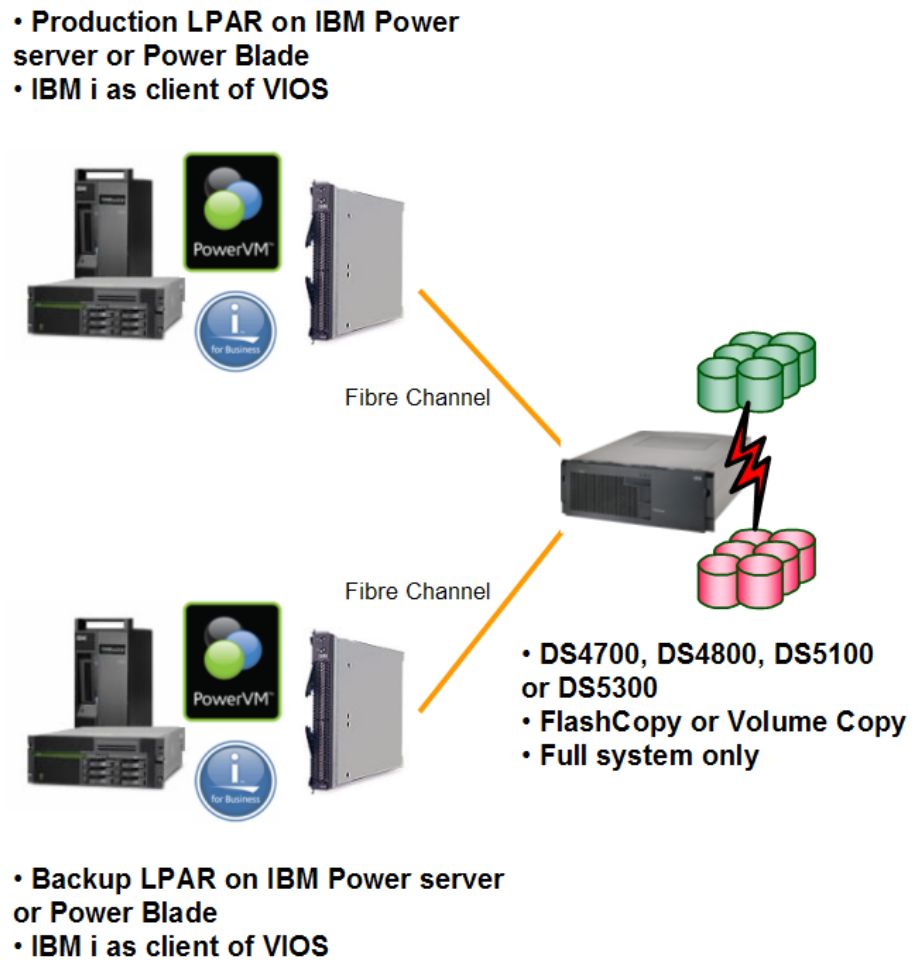

Figure 11Test environment for Flashcopy an d volume copy.

9.1.1.2 FlashCopy and VolumeCopy support statements

The use of DS4000 and DS5000 FlashCopy and VolumeCopy with IBM i as a client of VIOS is supported as outlined in this section. Note that to implement and use this solution, **multiple manual steps on the DS4000 or DS5000 storage subsystem, in VIOS and in IBM i are required**. Currently, no toolkit that automates this solution exists and it is not part of IBM

PowerHA® for IBM i. The components of the solution – DS4000 or DS5000 FlashCopy/VolumeCopy, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organizations and not solely by the IBM i Support Center.

Support statements:

- DS4000 and DS5000 FlashCopy and VolumeCopy are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system FlashCopy and VolumeCopy when the production IBM i logical partition (LPAR) is powered off are supported .
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is in restricted state are supported.
- The DS4000 and DS5000 'disable' and 're-create' functions with full-system FlashCopy and VolumeCopy when the production IBM i LPAR is powered off or is in restricted state are supported.
- Full-system FlashCopy and VolumeCopy of the production IBM i logical partition (LPAR) after only using the IBM i 6.1 memory flush to disk (quiesce) function are **not supported**
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is running are **not supported.**
- FlashCopy and VolumeCopy of independent auxiliary storage pools (IASPs) are **not supported.**
- Having the production and backup IBM i LPAR under the same VIOS is **not supported.**

For assistance with using DS4000 and DS5000 FlashCopy and VolumeCopy with IBM i, contact IBM Systems Lab Services and Training using this website:
http://www.ibm.com/systems/services/labservices/contact.html.

9.1.2 Enhanced Remote Mirroring (ERM)

9.1.2.1 Test scenario

Figure 12 shows the test environment used for Enhanced Remote Mirroring.

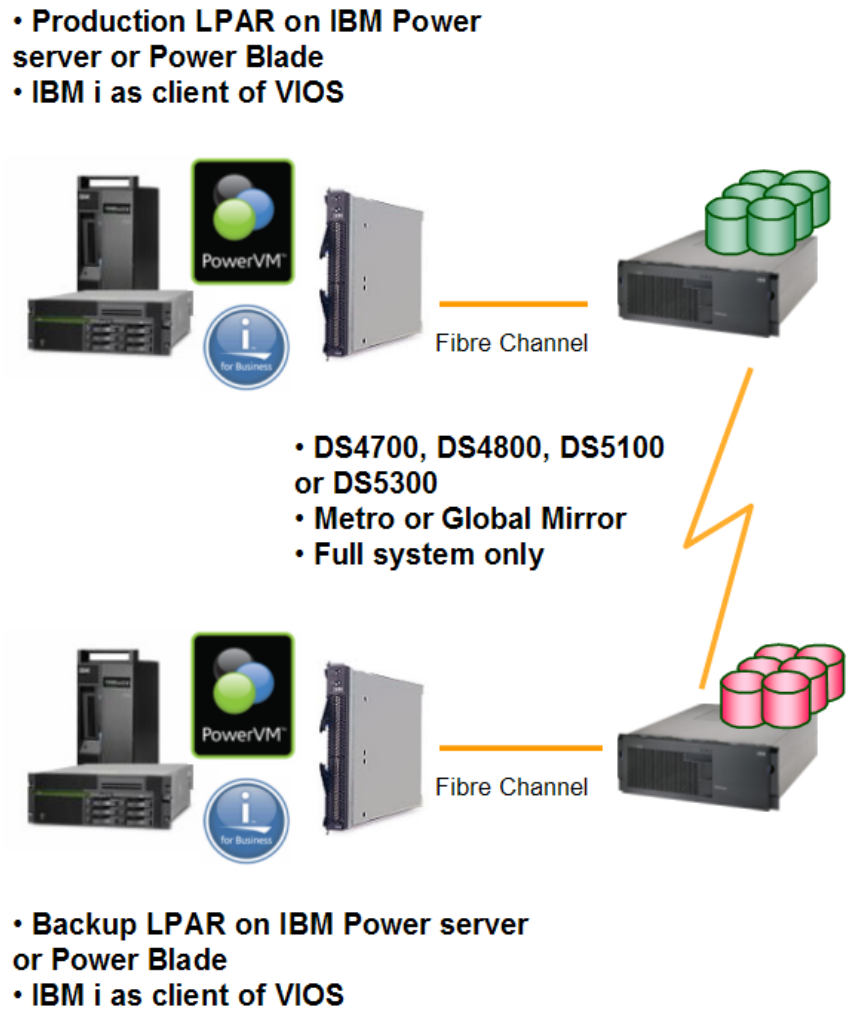

Figure 12: the test environment used for ERM.

9.1.3 ERM support statements

The use of DS4000 and DS5000 Enhanced Remote Mirroring with IBM i as a client of VIOS is supported as outlined in this section. Note that to implement and use this solution, **multiple manual steps on the DS4000 or DS5000 storage subsystem, in VIOS and in IBM i are required**. Currently, no toolkit that automates this solution exists and it is not part of IBM PowerHA for IBM i. The components of the solution – DS4000 or DS5000 ERM, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organization and not solely by the IBM i Support Center.

IBM PowerHA for IBM i is also supported for IBM i as a client of VIOS. PowerHA for IBM i provides an automated, IBM i-driven replication solution that allows clients to leverage their existing IBM i skills. PowerHA for IBM i is supported by the IBM i Support Center. For more information on IBM PowerHA for IBM i, refer to the URL: http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html.

Support statements for ERM:

- DS4000 and DS5000 ERM is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Synchronous ERM (DS4000 and DS5000 **Metro Mirror**) is supported.
- Asynchronous ERM with Write Consistency Groups (DS4000 and DS5000 **Global Mirror**) is supported.
- Asynchronous ERM (DS4000 and DS5000 Global Copy) is **not supported.**
- Full-system ERM (Metro Mirror and Global Mirror) for a planned switchover (IBM i production LPAR is powered off) is supported.
- Full-system ERM (Metro Mirror and Global Mirror) for an unplanned failover (IBM i production LPAR is running) is supported.
- Replication of IASPs is **not supported.**
- Suspending replication and accessing the replicated LUNs on backup site is **not supported.**

For assistance with using DS4000 and DS5000 Enhanced Remote Mirroring with IBM i, contact IBM Lab Services using the follwoing website:
http://www.ibm.com/systems/services/labservices/contact.html.

9.2 SAN Volume Controller (SVC) Copy Services and IBM i

IBM has conducted some basic functional testing of SVC Copy Services with IBM i as client of VIOS. In this section, you will find information on the scenarios tested and the resulting statements of support for using SVC Copy Services with IBM i.

9.2.1 FlashCopy

9.2.1.1 Test scenario

Figure 13 shows the test environment used for FlashCopy.

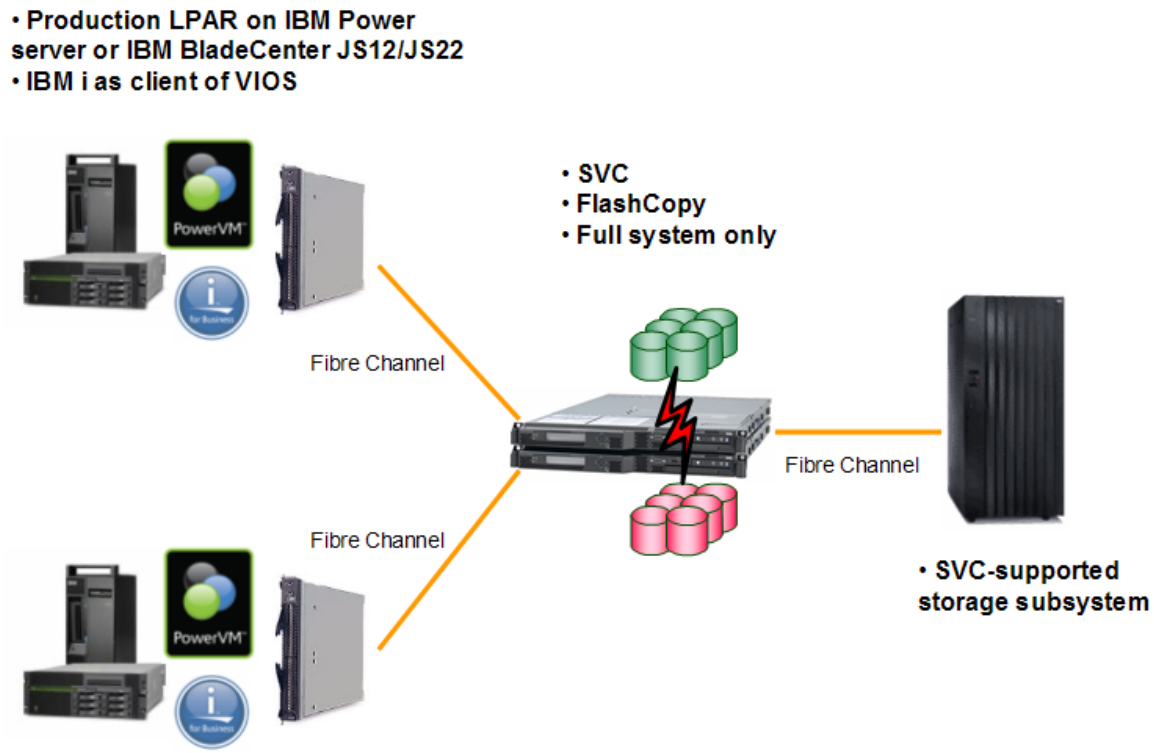

Figure 13:The test environment used for FlashCopy.

9.2.1.2 FlashCopy statements

The use of SVC FlashCopy with IBM i as a client of VIOS is supported as outlined below. **Please note** that to implement and use this solution, **multiple manual steps in SVC, in VIOS and in IBM i are required**. Currently, no toolkit that automates this solution exists and it is not part of IBM PowerHA for IBM i. The components of the solution – SVC FlashCopy, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organizations and not solely by the IBM i Support Center.

Support statements:

- SVC FlashCopy is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system FlashCopy when the production IBM i logical partition (LPAR) is powered off is supported .
- Full-system FlashCopy when the production IBM i LPAR is in restricted state is supported.
- Full-system FlashCopy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function is supported.
- Full-system FlashCopy when the production IBM i LPAR is running is **not supported.**
- FlashCopy of Independent Auxiliary Storage Pools (IASPs) is **not supported.**
- FlashCopy between VDisks using heterogeneous storage subsystems is supported.

For assistance with using SVC FlashCopy with IBM i, contact IBM Lab Services refer to the following website: http://www.ibm.com/systems/services/labservices/contact.html.

9.2.2 Metro and Global Mirror

9.2.2.1 Test scenario

Figure 14 shows the test environment used for Metro and Global Mirror.

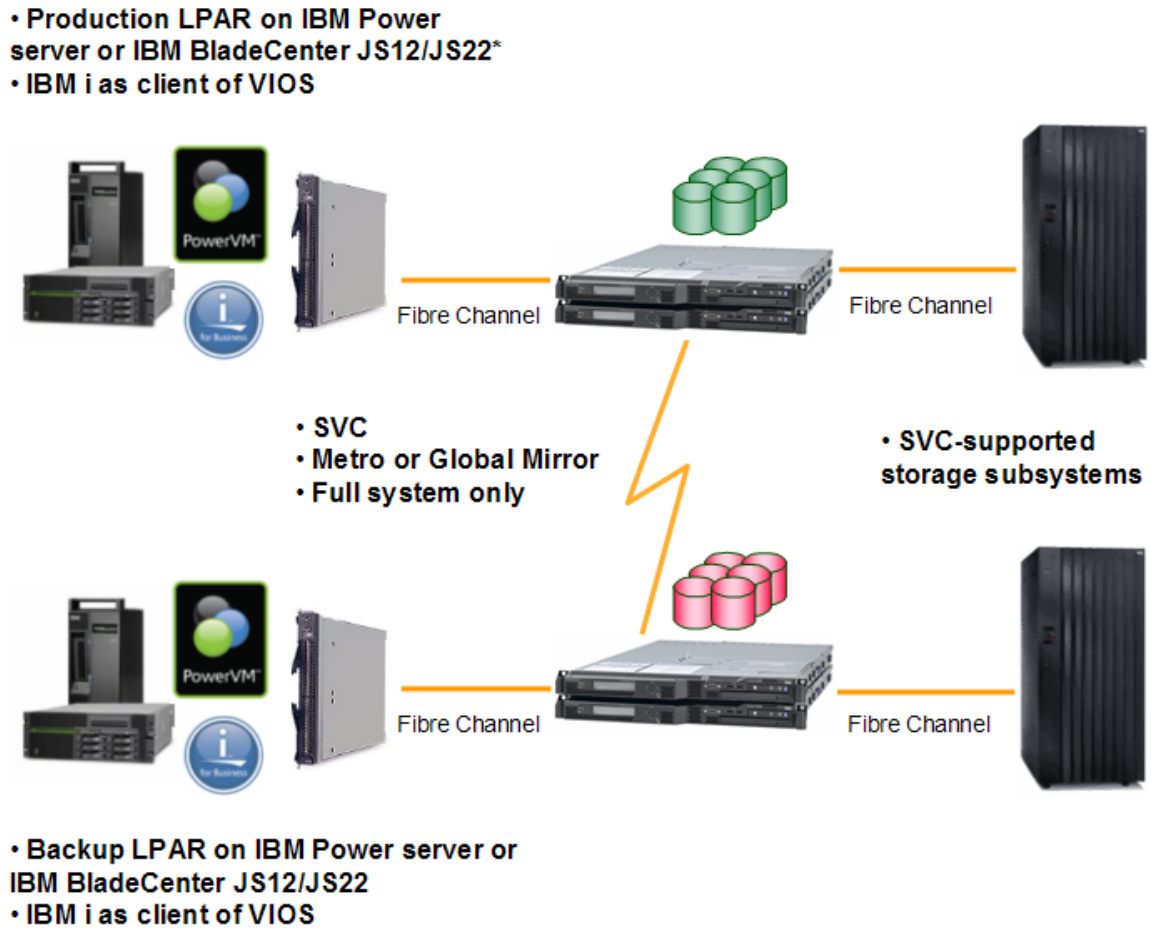

Figure 14: The test environment used for Metro Mirror and Global Mirror.

9.2.2.2 Metro Mirror and Global Mirror support statements

The use of SVC Metro Mirror and Global Mirror with IBM i as a client of VIOS is supported as outlined in the following section. Note that to implement and use this solution, **multiple manual steps in SVC, in VIOS and in IBM i are required**. Currently, no toolkit that automates this solution exists and it is not part of IBM PowerHA for IBM i. The components of the solution – SVC Metro or Global Mirror, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organization and not solely by the IBM i Support Center.

IBM PowerHA for IBM i is also supported for IBM i as a client of VIOS. PowerHA for IBM i provides an automated, IBM i-driven replication solution that allows clients to leverage their existing IBM i skills. PowerHA for IBM i is supported by the IBM i Support Center. For more information on IBM PowerHA for IBM i, refer to the following URL: http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html.

Support statements for SVC Metro and Global Mirror:

- Both SVC Metro Mirror and Global Mirror are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Only full-system replication is supported.
- Replication of IASPs with Metro Mirror or Global Mirror is **not supported.**
- Both Metro Mirror and Global Mirror are supported for a planned switchover (IBM i production partition is powered off).
- Both Metro Mirror and Global Mirror are supported for an unplanned failover (IBM i production partition is running).
- Suspend and resume in case of link failure between SVC clusters is supported for both Metro Mirror and Global Mirror.

For assistance with using SVC Metro and Global Mirror with IBM i, contact IBM Lab Services using the following website: http://www.ibm.com/systems/services/labservices/contact.html.

9.3 XIV Copy Services and IBM i

IBM has conducted some basic functional testing of XIV System Storage Copy Services with IBM i as client of VIOS. The following lists is the resulting statements of support for using XIV Copy Services with IBM i.

Snapshots and Volume Copy:

- XIV snapshots and Volume Copy are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system snapshots and Volume Copy when the production IBM i logical partition (LPAR) is powered off are supported.
- Full-system snapshots and Volume Copy when the production IBM i LPAR is in restricted state are supported.
- Full-system snapshots and Volume Copy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function are supported.
- Full-system snapshots and Volume Copy when the production IBM i LPAR is running are **not supported.**
- Snapshots and Volume Copy of iASPs are **not supported.**

Remote Mirror:

- XIV Remote Mirror (synchronous replication) is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades
- XIV asynchronous replication is also supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades
- Only full-system replication is supported
- Replication of IASPs with Remote Mirror is **not supported**
- Remote Mirror is supported for a planned switchover (IBM i production partition is powered off)
- Remote Mirror is supported for an unplanned failover (IBM i production partition is running)

9.4 DS5000 Direct attach Copy Services

IBM has conducted some basic functional testing of DS5100 and DS5300 Copy Services when directly attached to IBM i. This section summarizes the resulting support statements. Note that using Copy Services when directly attaching DS5100 and DS5300 storage systems involves manual steps in both the DS Storage Manager GUI and IBM i. PowerHA for IBM i does not support DS5100 and DS5300 Copy Services when directly attachied to IBM i (or through VIOS).

FlashCopy and VolumeCopy:

- FlashCopy and VolumeCopy are supported when directly attaching DS5100 and DS5300 to IBM i on POWER6 or higher processor-based Power servers.
- As mentioned earlier, Power blades are **not supported** for direct IBM i attachment to DS5100 and DS5300 in general.
- Full-system FlashCopy and VolumeCopy when the production IBM i logical partition (LPAR) is powered off are supported.
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is in restricted state are supported.
- Full-system FlashCopy and VolumeCopy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function are **not supported.**
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is running are **not supported.**
- FlashCopy and VolumeCopy of Independent Auxiliary Storage Pools (IASPs) are supported.

Enhanced Remote Mirroring (ERM):

- DS5100 and DS5300 ERM with direct attachment to IBM i is supported on IBM Power servers (but **not** on IBM Power blades).
- Synchronous ERM (DS5100 and DS5300 **Metro Mirror**) is supported.
- Asynchronous ERM with Write Consistency Groups (DS5100 and DS5300 **Global Mirror**) is supported.
- Asynchronous ERM (DS5100 and DS5300 **Global Copy**) is **not supported.**
- Full-system ERM (Metro and Global Mirror) for a planned switchover (IBM i production LPAR is powered off) is supported.
- Full-system ERM (Metro and Global Mirror) for an unplanned failover (IBM i production LPAR is running) is supported.
- ERM (Metro and Global Mirror) with IASPs for a planned switchover (IBM i production LPAR is powered off) is supported.
- ERM (Metro and Global Mirror) with IASPs for an unplanned failover (IBM i production LPAR is running) is supported.

9.5 NPIV attached storage Copy Services

The main differences between NPIV and VSCSI I/O virtualization for IBM i are:

- With NPIV, storage and tape devices report into IBM i with their specific characteristics, as opposed to generic SCSI devices.
- IBM i has also gained NPIV access to storage and tape devices to which it could already attached directly, benefiting from transparent FC adapter sharing.

As a result, both the storage subsystem (DS8000) and LUN types involved with NPIV for IBM i are identical to those already supported by PowerHA for IBM i. PowerHA for IBM i continues to manage both Copy Services on DS8000 using the Storage Management Console (SMC) and the

required clustering objects in IBM i from a single interface. Therefore, DS8000 storage system Copy Services are supported by PowerHA for IBM i when using NPIV, including FlashCopy, Metro Mirror and Global Mirror.

10 Appendix:

1 Additional resources

These websites provide useful references to supplement the information contained in this paper

1.1 IBM i

- IBM i on a Power Blade Read-me First: http://www.ibm.com/systems/power/hardware/blades/ibmi.html.
- IBM STG Lab Services: http://www.ibm.com/systems/services/labservices/contact.html.
- Logical Partitioning Guide: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf.
- IBM i installation: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/rzahc1.htm&tocNode=int_216451
- IBM i PTF installation: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzam8/rzam8fixconcepts1.htm
- Latest recommended IBM i PTFs: http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes
- IBM Fix Central: http://www-912.ibm.com/eserver/support/fixes/
- IBM Systems Workload Estimator: http://www-304.ibm.com/systems/support/tools/estimator/index.html
- IBM i networking: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm
- ECS over LAN: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm
- IBM i virtual client partitions topics in the IBM i Information Center: http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=/rzahc/rzahcbladei5limits.htm
- IBM PowerHA for IBM i: http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html.

1.2 Storage

- *IBM System Storage DS4000 and Storage Manager V10.30* (Redbook): http://www.redbooks.ibm.com/redpieces/abstracts/sg247010.html?Open
- DS4000 Website: http://www.ibm.com/systems/storage/disk/ds4000/index.html
- *Performance Capabilities Reference* manual (chapter 6 for best practices for IBM i and VIOS): http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html
- SDD-PCM driver: http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D430&uid=ssg1S4000201&loc=en_US&cs=utf-8&lang=en
- IBM Redbooks site: http://www.redbooks.ibm.com
- *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247120.html?Open.
- *Implementing IBM Tape in i5/OS* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247440.html?Open.

- *IBM Midrange System Storage Hardware Guide*(Redbook): http://www.redbooks.ibm.com/abstracts/sg247676.html?Open.
- *IBM System Storage DS3500: Introduction and Implementation Guide: http://www.redbooks.ibm.com/redpieces/pdfs/sg247914.pdf*

1.3 VIOS

- *PowerVM Editions Guide*: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm
- *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247940.html?Open
- VIOS command reference: http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm
- VIOS Datasheet: http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html
- *PowerVM Managing and Monitoring* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247590.html?Open.

1.4 SVC

- SVC overview Website: http://www.ibm.com/systems/storage/software/virtualization/svc
- Interoperability information: http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html
- *Implementing the IBM System Storage SAN Volume Controller V5.1* (Redbook): http://www.redbooks.ibm.com/abstracts/sg246423.html?Open.
- *SAN Volume Controller Best Practices and Performance Guidelines* (Redbook): http://www.redbooks.ibm.com/abstracts/sg247521.html?Open.
- *SVC V4.3.0 Advanced Copy Services* (Redbook): http://www.redbooks.ibm.com/redpieces/abstracts/sg247574.html?Open.
- IBM SAN Volume Controller 4.2.1 Cache Partitioning (Redpiece): http://www.redbooks.ibm.com/abstracts/redp4426.html?Open.

2 Trademarks and disclaimers

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Refer to your local IBM business contact for information on the IBM offerings available in your area.
Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.
All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.
All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.
IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.
IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.
All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L (logo), AIX 6 (logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business (logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM i, IBM i (logo), IBM, IBM (logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, OpenPower, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC (logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, Quick Place, Rational, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Sametime, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, Tivoli Ready (logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX 5L, AIX PVMe, AS/400e, Calibrated Vectored

Cooling, Chiphopper, Chipkill, Cloudscape, DataPower, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, DS4000, DS6000, DS8000, e-business (logo), e-business on demand, EnergyScale, Enterprise Workload Manager, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, General Purpose File System, GigaProcessor, GPFS, HACMP, HACMP/6000, IBM Systems Director Active Energy Manager, IBM TotalStorage Proven, IBMLink, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, On Demand Business logo, POWER, PowerExecutive, PowerVM, PowerVM (logo), Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power PC, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, pure XML, Quickr, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, Workload Partitions Manager, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
UNIX is a registered trademark of The Open Group in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Itanium, Pentium are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.
Other company, product and service names may be trademarks or service marks of others.