



IBM i Virtualization and Open Storage Read-me First

Vess Natchev

Power Blades and Virtualization
IBM Rochester, MN
vess@us.ibm.com

December 5th, 2008



This “read-me first” document provides detailed instructions on using IBM i 6.1 virtualization and connecting open storage to IBM i. It covers prerequisites, supported hardware and software, planning considerations, install and post-install tasks such as backups. The document also contains links to many additional information sources.

Table of Contents

1. IBM i virtualization solutions

- 1.1. IBM i logical partition (LPAR) hosting another IBM i partition
- 1.2. IBM i using open storage as a client of the Virtual I/O Server (VIOS)
- 1.3. IBM i on a Power blade

2. IBM i hosting IBM i supported configurations

- 2.1. Hardware
- 2.2. Software and firmware

3. IBM i hosting IBM i concepts

- 3.1. Virtual SCSI and Ethernet adapters
- 3.2. Storage virtualization
- 3.3. Optical virtualization
- 3.4. Network virtualization

4. Prerequisites for implementing IBM i hosted LPARs

- 4.1. Storage planning
- 4.2. Performance
- 4.3. Dual hosting

5. Implementing IBM i client LPARs with an IBM i host

6. Post-install tasks and considerations

- 6.1. Configure IBM i networking
- 6.2. How to perform IBM i operator panel functions
- 6.3. How to display the IBM i partition System Reference Code (SRC) history
- 6.4. Client IBM i LPARs considerations and limitations
- 6.5. Configuring Electronic Customer Support (ECS) over LAN
- 6.6. Copying storage spaces
- 6.7. Backups

7. IBM i using open storage through VIOS supported configurations

- 7.1. Hardware
- 7.2. Software and firmware

8. IBM i using DS4000 storage through VIOS concepts

- 8.1. Virtual SCSI and Ethernet adapters
- 8.2. Storage virtualization
- 8.3. Optical virtualization
- 8.4. Network virtualization

9. Prerequisites for attaching DS4000 storage to IBM i through VIOS

- 9.1. Storage planning
- 9.2. Performance
- 9.3. Dual hosting and multi-path I/O (MPIO)

10. Attaching DS4000 storage to IBM i through VIOS

- 10.1. DS4000 configuration
- 10.2. VIOS installation and configuration
- 10.3. IBM i installation and configuration

11. Post-install tasks and considerations

- 11.1. Configure IBM i networking
- 11.2. How to perform IBM i operator panel functions
- 11.3. How to display the IBM i partition System Reference Code (SRC) history
- 11.4. Client IBM i LPARs considerations and limitations
- 11.5. Configuring Electronic Customer Support (ECS) over LAN
- 11.6. Backups

12. DS4000 Copy Services and IBM i

- 12.1. FlashCopy and VolumeCopy
- 12.2. Enhanced Remote Mirroring (ERM)

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

- 13.1. IBM i and SVC concepts
- 13.2. Attaching SVC storage to IBM i

14. SVC Copy Services and IBM i

- 14.1. FlashCopy
 - 14.1.1. Test scenario
 - 14.1.2. FlashCopy statements
- 14.2. Metro and Global Mirror
 - 14.2.1. Test scenario
 - 14.2.2. Metro and Global Mirror support statements

15. Additional resources

16. Trademarks and disclaimers

1. IBM i virtualization solutions

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

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

An IBM i 6.1 LPAR can host one or more additional IBM i LPARs, known as virtual client LPARs. Virtual client partitions can 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, 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 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 document.

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 whose purpose is to provide virtual storage, optical and networking resources to one or more client partitions. The most immediate benefit 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 logical units (LUNs) are physically attached to VIOS via a Fibre Channel connection and then made available to IBM i. While IBM i does not directly attach to the SAN in this case, once open storage LUNs become available through VIOS, they are managed the same way as integrated disks or LUNs from a DS8000 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 JS22 or JS12. Running IBM i on a Power blade is beyond the scope of this document. See 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 supported configurations

2.1. 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 a POWER6-based system can be virtualized to the client IBM i partition. **Virtualization of tape devices is not supported.** The following table lists the supported hardware:

Hardware type	Supported for IBM i hosting IBM i	Notes
IBM Power servers	Yes	Includes IBM Power 520 Express, IBM Power 550 Express, IBM Power 560 Express, IBM Power 570 and IBM Power 595 Does not include IBM POWER6 processor-based blade servers, such as IBM BladeCenter JS12 and JS22
IBM Power 575	No	
POWER5-based systems or earlier	No	
Storage adapters (Fibre Channel, SAS, SCSI)	Yes	Must be supported by IBM i 6.1 and supported on POWER6-based IBM Power server
Storage devices and subsystems	Yes	Must be supported by IBM i 6.1 and supported on POWER6-based IBM Power server
Network adapters	Yes	Must be supported by IBM i 6.1 and supported on POWER6-based IBM Power server
Optical devices	Yes	Must be supported by IBM i 6.1 and supported on POWER6-based IBM Power server
Tape devices	No	

To determine the storage, network and optical devices supported on each IBM Power server model, refer to the Sales Manual for each model: <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 Web site: <https://www-304.ibm.com/systems/support/i/planning/upgrade/futurehdwr.html>.

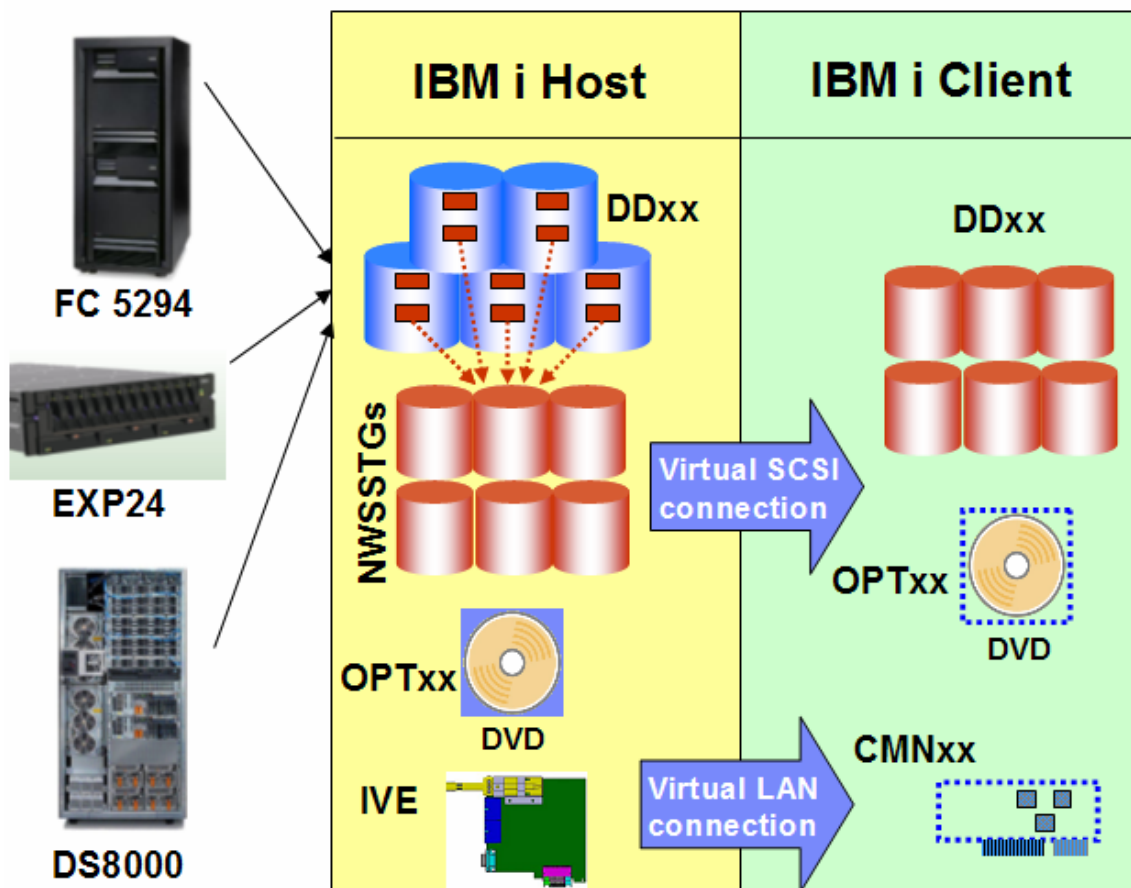
2.2. Software and firmware

Software or firmware type	Supported for IBM i hosting IBM i	Notes
IBM i 6.1 or later	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 firmware 320_040_031 or later	Yes	This is the minimum system firmware level required

HMC firmware HMC V7 R3.2.0 or later	Yes	This is the minimum HMC firmware level required
-------------------------------------	-----	---

3. 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. The following diagram shows the full solution and its components:



3.1. Virtual SCSI and Ethernet adapters

IBM i hosting IBM i uses an existing function of the system firmware, or Power Hypervisor: the capability to create virtual SCSI and Ethernet adapters in a partition. Virtual adapters are created for each LPAR in the Hardware Management Console (HMC). Virtual SCSI 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 types of adapters to the partition in the HMC.

3.2. Storage virtualization

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

- One virtual SCSI server adapter in the host partition
- One virtual SCSI client adapter in the client partition

This virtual SCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one virtual SCSI pair can exist for the same client partition in this environment. To minimize performance overhead on the host partition, the virtual SCSI 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 are 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:

- One Network Server Description (NWSA) object
- One Network Server Storage Space (NWSSTG) object

The NWSA object associates a virtual SCSI server adapter in IBM i (which in turn is connected to a virtual SCSI client adapter in the HMC) with one or more NWSSTG objects. At least one NWSA 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). The following screenshot 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

  Opt  Name      Server  Seq  Link  Access  Stg
  ---  ---      ---     ---  ---   ---     ---
  ---  CP10D01  CP10A   2    *DYN *UPDATE
  ---  CP10D02  CP10A   3    *DYN *UPDATE
  ---  CP10D03  CP10A   4    *DYN *UPDATE
  ---  CP10D04  CP10A   5    *DYN *UPDATE
  
```

The next screenshot 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										
Unit	Type	Size (M)	% Used	I/O Rqs	Request Size (K)	Read Rqs	Write Rqs	Read (K)	Write (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

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 virtual SCSI 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.

3.3. Optical virtualization

Any optical drive supported in the host IBM i LPAR can be virtualized to an IBM i client LPAR. An existing virtual SCSI connection can be used, or a new connection can be created explicitly for optical I/O traffic. By default, if a virtual SCSI connection exists between host and client, all physical 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 install of the client partition, as well as for installing Program Temporary Fixes (PTFs) or applications later. If the optical drive is writeable, the client partition will be able to write to the physical media in the drive.

3.4. 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. 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. In the host partition, the virtual Ethernet adapter is then associated with the physical network adapter via a routing configuration – either Proxy ARP or Network Address Translation (NAT). This allows the client partition to send network packets via 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.

4. Prerequisites for implementing IBM i hosted LPARs

4.1. 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 what percent of the virtual storage is used in the client. For example, if a 500-GB storage space is created, it will occupy that amount of physical storage in the host IBM i LPAR, even if the disk capacity is only 50% utilized in the client LPAR.

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 non-configured 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 virtual SCSI connection for an IBM i client partition, which is 16. If more than 16 NWSSTGs are needed for a client LPAR, additional virtual SCSI connections can be created dynamically in the HMC.

4.2. Performance

As described in section 2.2, disk I/O operations in an IBM i virtual client partition result in I/O requests to the physical disk adapter(s) and drives 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.* Because the host partition is a standard IBM i partition, all the recommendations in the *Performance Capabilities Reference* manual (<http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>) will apply to it. Use the manual's suggestions 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 is 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.

4.3. 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 section 2.2. In addition, a second virtual SCSI client adapter exists in the client LPAR, connected to a virtual SCSI 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.

5. 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 virtual SCSI configuration in the HMC
- 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*: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.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 in section 4.2 of this document.

6. Post-install tasks and considerations

6.1. 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 PCI slot

- A logical port on a Host Ethernet Adapter (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, see the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Power Systems Logical Partitioning Guide* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.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, see the topic **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** in the *Power Systems Logical Partitioning Guide*: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

Lastly, a virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, consult the topic **Configuring a virtual Ethernet adapter using the HMC** in the *Power Systems Logical Partitioning Guide*: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.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 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: <http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/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 Network Address Translation (NAT). To configure Proxy ARP or NAT in the IBM i host partition, follow the instructions in section 5.2 of the Redbook *Implementing POWER Linux on IBM System i Platform* (<http://www.redbooks.ibm.com/redbooks/pdfs/sq246388.pdf>).

6.2. How to perform IBM i operator panel functions

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

- Sign onto 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, then the desired function.

6.3. How to display the IBM i partition System Reference Code (SRC) history

- Sign onto 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 → Reference Code History
- To display words 2 through 9 of a reference code, click the radio button for that code.

6.4. Client IBM i LPARs considerations and limitations

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

6.5. 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. Consult the topic **Setting up a connection to IBM** in the Information Center: http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaji/rzaji_setup.htm.

6.6. 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 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, then press Enter
- Enter a name of up to 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 (up to 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
- Press Enter

To deploy the new client partition, follow the instructions in section 5 to create the necessary virtual SCSI configuration in the HMC and the NWSD object in the host IBM i partition.

6.7. 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. Therefore, the simplest backup and restore 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, consult 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/iphathat/iphathat.pdf>).

Once a tape adapter connected to a tape drive or library is available to the client partition, use the Backup and Recovery topic in the Information Center to manage backups: http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/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. See the topic **Saving IBM i server objects in IBM i** in the *Logical Partitioning Guide* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphathat/iphathat.pdf>).

7. IBM i using open storage through VIOS supported configurations

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

- 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 a SAN Volume Controller (SVC). SVC in turn connects to one or more storage subsystems through the SAN fabric

The support statements about individual storage subsystems below apply to the first connection method: IBM i → VIOS → storage subsystem. Furthermore, **the support statements in the tables below apply to the end-to-end solution of IBM i using open storage as a client of VIOS.** This document will 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 statement, including other clients, see 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, see the data sheet on the SVC overview Web site:

<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 processor-based blade servers, such as IBM BladeCenter JS12 or JS22. **For the full supported configurations statement for the IBM i on Power blade solution**, see the Supported Environments document at:

<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>.

7.1. Hardware

Hardware type	Supported by IBM i as a client of VIOS	Notes
IBM Power servers	Yes	Includes IBM Power 520 Express, IBM Power 550 Express, IBM Power 560 Express, IBM Power 570 and IBM Power 595
IBM Power 575	No	
POWER5-based systems or earlier	No	
DS4800 using Fibre Channel drives	Yes	Supported by IBM i as a client of VIOS on both IBM Power servers and IBM BladeCenter JS22 and JS12
DS4800 using SATA drives	Yes	It is strongly recommended that SATA drives are used only for test or archival IBM i applications for performance reasons
DS4700 using Fibre Channel drives	Yes	Supported by IBM i as a client of VIOS on both IBM Power servers and IBM BladeCenter JS22 and JS12
DS4700 using SATA drives	Yes	It is strongly recommended that SATA drives are used only for test or archival 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 IBM i applications for performance reasons
DS4200 and earlier DS4000	No	

and FAStT models		
EXP810 expansion unit attached to DS4800 or DS47000	Yes	See supported drives comments above
EXP710 expansion unit attached to DS4800 or DS4700	Yes	See 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
Fibre Channel adapters	Yes	Must be supported by VIOS*

*For a list of Fibre Channel adapters supported by VIOS, see the VIOS Datasheet at: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>.

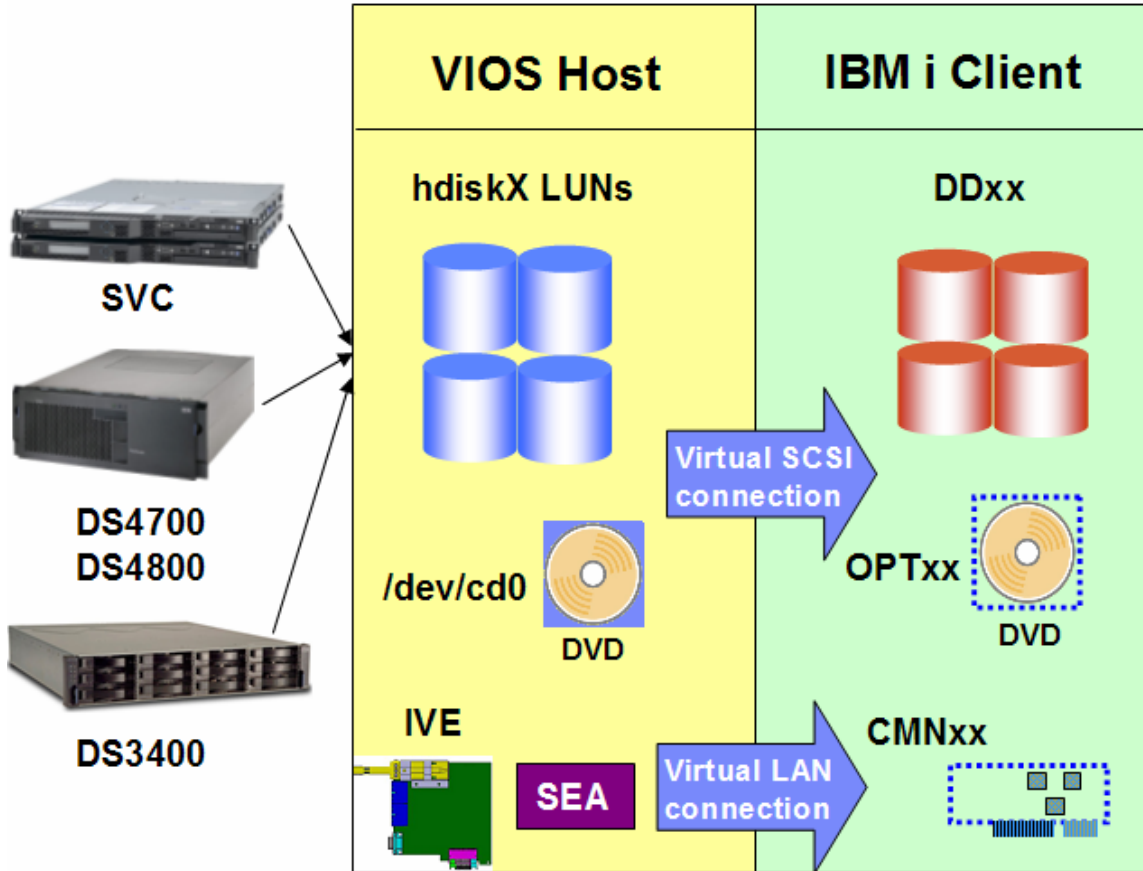
7.2. Software and firmware

Software or firmware type	Supported by IBM i as a client of VIOS	Notes
IBM i 6.1 or later	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 later	Yes	This is the minimum level of VIOS required
IBM Power server system firmware 320_040_031 or later	Yes	This is the minimum system firmware level required
HMC firmware HMC V7 R3.2.0 or later	Yes	This is the minimum HMC firmware level required
DS4000 controller firmware 6.60.08 or later	Yes	This is the minimum DS4000 controller firmware level required. Controller firmware 7.10.23 or later is supported.
DS4000 Storage Manager 9.x or later	Yes	Storage Manager 10.10.x or later is supported
DS4000 Copy Services	Yes	See Section 12 of this document for details
AIX/VIOS host kit license for DS4000	Yes	This license is required in order to attach DS4000 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

8. IBM i using DS4000 storage through VIOS concepts

The capability to attach to DS4000 through VIOS extends the IBM i storage portfolio to include 512-byte-per-sector open storage. The existing IBM i storage portfolio includes integrated SCSI or SAS disk, as well as Fibre Channel-attached storage subsystems that support 520 bytes per sector, such as the IBM DS8000 product line. IBM i cannot currently attach directly to DS4000, because its Fibre Channel adapter code requires 520 bytes per sector. Therefore, logical units

(LUNs) created on the DS4000 are physically connected through the Fibre Channel fabric to VIOS. Once recognized by VIOS, the LUNs are virtualized to one or more IBM i virtual client partitions. The following diagram illustrates both the 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 DS4000 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 to 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 DS4000 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.

As illustrated in the diagram above, the concepts and components for using SVC and DS3400 storage with IBM i as a client of VIOS are the same as for DS4000.

8.1. Virtual SCSI 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 virtual SCSI and Ethernet adapters in a partition. Virtual adapters are created for each LPAR in the Hardware Management Console (HMC).

Virtual SCSI 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. 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.

8.2. Storage virtualization

8.2.1. HMC configuration

For VIOS to virtualize LUNs created on a DS4000 to an IBM i client partition, both HMC and VIOS objects must be created. In the HMC, the minimum required configuration is:

- One virtual SCSI server adapter in the host partition
- One virtual SCSI client adapter in the client partition

This virtual SCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one virtual SCSI pair can exist for the same client partition in this environment. To minimize performance overhead in VIOS, the virtual SCSI 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 are transferred directly from the Fibre Channel adapter in VIOS to a buffer in memory of the IBM i client partition.

In an IBM i client partition, a virtual SCSI client adapter is recognized as a type **290A DCxx** storage controller device. The screenshot below depicts the virtual SCSI client adapter, as well as several DS4000 LUNs and an optical drive virtualized by VIOS:

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

Opt Resource          Type-model  Status          Text
--  -
  -  DC01              290A-001   Operational    Storage Controller
  -  DD001              6B22-050   Operational
  -  DD002              6B22-050   Operational
  -  DD003              6B22-050   Operational
  -  DD004              6B22-050   Operational
  -  DD005              6B22-050   Operational
  -  OPT01              632C-002   Operational    Optical Storage Unit
  -  DD006              6B22-050   Operational
  -  DD007              6B22-050   Operational
```

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

```
$ lsdev | grep vhost
vhost0          Available   Virtual SCSI Server Adapter
$
```

8.2.2. VIOS configuration

No additional configuration in the IBM i client partition is necessary. In VIOS, however, a new object must be created **for each DS4000 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 DS4000 LUNs in this case
- Logical volumes (*hdX and other*)
- Files in a directory

For both simplicity and performance reasons, **it is recommended to virtualize DS4000 LUNs to IBM i directly as physical devices (hdiskX)**, and not through the use of logical volumes or files. (See section 9.2 for a detailed performance discussion.) A vtscsiX device links a LUN available in VIOS (*hdiskX*) to a specific virtual SCSI adapter (*vhostX*). In turn, the virtual SCSI 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 DS4000 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 DS4000 LUNs, can be linked to a single vhostX virtual SCSI server adapter and made available to IBM i. Up to 16 LUNs can be virtualized to IBM i through a single virtual SCSI connection. Each LUN typically uses multiple physical disk arms in the DS4000. If more than 16 LUNs are required in an IBM i client partition, an additional pair of virtual SCSI server (VIOS) and client (IBM i) adapters must be created in the HMC. Additional LUNs available in VIOS can then be linked to the new vhostX device through vtscsiX devices, making them available to IBM i.

8.2.3. Adding devices to VIOS

VIOS does not autoscans 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 with the following command:

- **cfgdev**

Once it has detected the device, VIOS will autoconfigure it and make it available for use. It can then be virtualized to the IBM i client partition.

8.3. Optical virtualization

8.3.1. 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 virtual SCSI server adapter and made available to IBM i through a *virtual target optical device*, *vtoptX*. The same optical drive can be made 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.

8.3.2. 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.

8.4. Network virtualization

Virtualizing a network adapter and using a virtual LAN (VLAN) for 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 partitions in the HMC. 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
```

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 (1-lan)
```

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.

9. Prerequisites for attaching DS4000 storage to IBM i through VIOS

9.1. Storage planning

The first storage planning consideration is having enough available capacity in the DS4000 to create the AIX/VIOS LUNs that will be virtualized to IBM i by VIOS. As mentioned in the supported hardware section (7.1), it is **strongly recommended that only Fibre Channel 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 drives may also be used. This recommendation is not meant to preclude the use of SATA drives for other clients of VIOS or other host servers; it applies only to production IBM i workloads.

Creating and attaching LUNs to VIOS requires that the DS4000 have the AIX/VIOS host kit software installed. Because 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. Actual LUN size will be based on the capacity and performance requirements of each IBM i virtual client partition.

9.2. Performance

When creating a DS4000 LUN configuration for IBM i as a client of VIOS, it is crucial to plan for both capacity and performance. It may seem that because LUNs are virtualized for IBM i by VIOS instead of being directly connected, the virtualization layer will necessarily add a significant performance overhead. However, internal IBM performance tests clearly show that the VIOS layer adds a very small amount of overhead to each I/O operation. Instead, our tests demonstrate that **when IBM i uses DS4000 LUNs virtualized by VIOS, performance is almost entirely determined by the physical and logical configuration of the DS4000**.

The IBM Rochester, MN, performance team has run a significant number of tests with IBM i using DS4000 storage as a client of VIOS. The resulting recommendations on configuring both the DS4000 and VIOS are available in the latest Performance Capabilities Reference manual (PCRM) at <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>. Chapter 14.5 focuses on virtualized storage for IBM i, with sections 14.5.1.1 – 14.5.1.3 specifically addressing the IBM Power server (non-blade) environment. In most cases, an existing IBM i partition using physical storage will be migrated to DS4000 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 DS4000 on which LUNs are created, while following the suggestions in the PCRM sections above.

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 the DS4000. Work with IBM Techline or your IBM Business Partner for a Disk Magic analysis. The latest version of Disk Magic includes support for DS4000 and IBM i as a virtual client of VIOS.

9.3. 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 DS4000 LUNs. The LUNs would also become unavailable if VIOS is brought down for scheduled maintenance or a release upgrade. To remove this dependency, two VIOS partitions can be used to simultaneously provide virtual storage to one or more IBM i client partitions.

The configuration for two VIOS partitions for the same client partition uses the same concepts as that for a single VIOS described in section 8.2. In addition, a second virtual SCSI client adapter exists in the client LPAR, connected to a virtual SCSI 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 DS4000, 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 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.

Note that the dual-VIOS solution above provides a level of redundancy by attaching **two separate sets of DS4000 LUNs** to the same IBM i client through separate VIOS partitions. It is not an MPIO solution that provides redundant paths to the **same set of LUNs**. There are two MPIO scenarios possible with VIOS:

- Two VIOS partitions providing redundant paths to the same set of LUNs on a single DS4000
- A single VIOS partition using two Fibre Channel adapters to connect to the same set of LUNs

Unlike an AIX client of VIOS, an IBM i client partition does not support the first MPIO scenario. In other words, **it is not possible to create a single set of LUNs on the DS4000 and make them available to IBM i through separate VIOS partitions**. The IBM i virtual I/O driver does not currently support this type of MPIO.

On the other hand, if a VIOS has two Fibre Channel adapters assigned in the HMC and the correct host configuration is created in the DS4000, VIOS will have redundant paths to the LUNs connected to the Fibre Channel adapters. The driver in VIOS that provides this MPIO support is RDAC (Redundant Disk Array Controller), which is included in a standard VIOS install. The necessary configuration on the DS4000 is described in Chapter 5.1.5 of the Redbook *IBM System Storage DS4000 and Storage Manager V10.10*, available at: <http://www.redbooks.ibm.com/redpieces/abstracts/sg247010.html?Open>. Once VIOS has redundant connections to the DS4000 LUNs created for IBM i, the IBM i client partition will implicitly take advantage of that redundancy without any additional configuration on the client side. If one of the physical paths to the LUNs fails in VIOS, the IBM i client partition will not be notified; recovery steps would be performed in VIOS.

10. Attaching DS4000 storage to IBM i through VIOS

As described in section 8, IBM i joins the VIOS virtualization environment, allowing it to use DS4000 storage. The setup process involves three main steps:

- DS4000 configuration
- VIOS installation and configuration
- IBM i installation and configuration

The VIOS and DS4000 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 DS4000 configuration requirements in order to achieve good performance, as referenced in section 9.2.

10.1. DS4000 configuration

Because VIOS is based on AIX and uses AIX Fibre Channel drivers, the configuration instructions for an AIX host are used. Consult Chapter 5 of the Redbook *IBM System Storage DS4000 and Storage Manager V10.10*, available at:

<http://www.redbooks.ibm.com/redpieces/abstracts/sg247010.html?Open>. Remember to refer back to the configuration suggestions in Chapter 14.5 of the latest PCRM at <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>.

10.2. VIOS installation and configuration

Consult 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>.

10.3. 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. Consult the topic **Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC** in the *Logical Partitioning Guide*:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphata.pdf>.

11. Post-install tasks and considerations

11.1. 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 PCI slot
- A logical port on a Host Ethernet Adapter (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, see 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/iphata.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, see the topic **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** in the *Logical Partitioning Guide*:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphata.pdf>.

Lastly, a virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, consult 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/iphath/iphath.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 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: <http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/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. See the topic **Configuring virtual Ethernet on the Virtual I/O Server** in the *PowerVM Editions Guide*, available at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>.

11.2. How to perform IBM i operator panel functions

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

- Sign onto 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, then the desired function.

11.3. How to display the IBM i partition System Reference Code (SRC) history

- Sign onto 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 → Reference Code History
- To display words 2 through 9 of a reference code, click the radio button for that code.

11.4. Client IBM i LPARs considerations and limitations

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

11.5. 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. Consult the topic **Setting up a connection to IBM** in the Information Center:

http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaji/rzaji_setup.htm.

11.6. Backups

As mentioned above, IBM i as a client of VIOS can use a mix of virtual and physical I/O resources. Therefore, the simplest backup and restore 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, consult 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/iphathat/iphathat.pdf>).

Once a tape adapter connected to a tape drive or library is available to the client partition, use the Backup and Recovery topic in the Information Center to manage backups: http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int_215989.

12. DS4000 Copy Services and IBM i

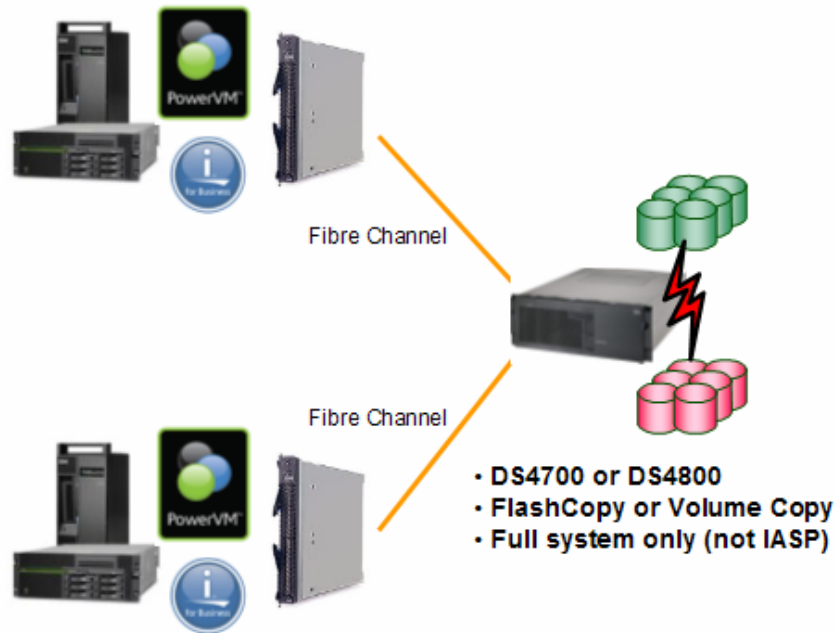
IBM has conducted some basic functional testing of DS4000 Copy Services with IBM i as client of VIOS. Below, you will find information on the scenarios tested and the resulting statements of support for using DS4000 Copy Services with IBM i.

12.1. FlashCopy and VolumeCopy

12.1.1. Test scenario

The following diagram shows the test environment used for FlashCopy and VolumeCopy:

- Production LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS

12.1.2. FlashCopy and VolumeCopy support statements

The use of DS4000 FlashCopy and VolumeCopy 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 on the DS4000 storage subsystem, in VIOS and in IBM i are required**. Currently, no toolkit exists that automates this solution and it is not part of IBM PowerHA for IBM i. The components of the solution – DS4000 FlashCopy/VolumeCopy, VIOS and IBM i – must be managed separately and require the corresponding skill set. Note also 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 FlashCopy and VolumeCopy are supported by IBM 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 '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

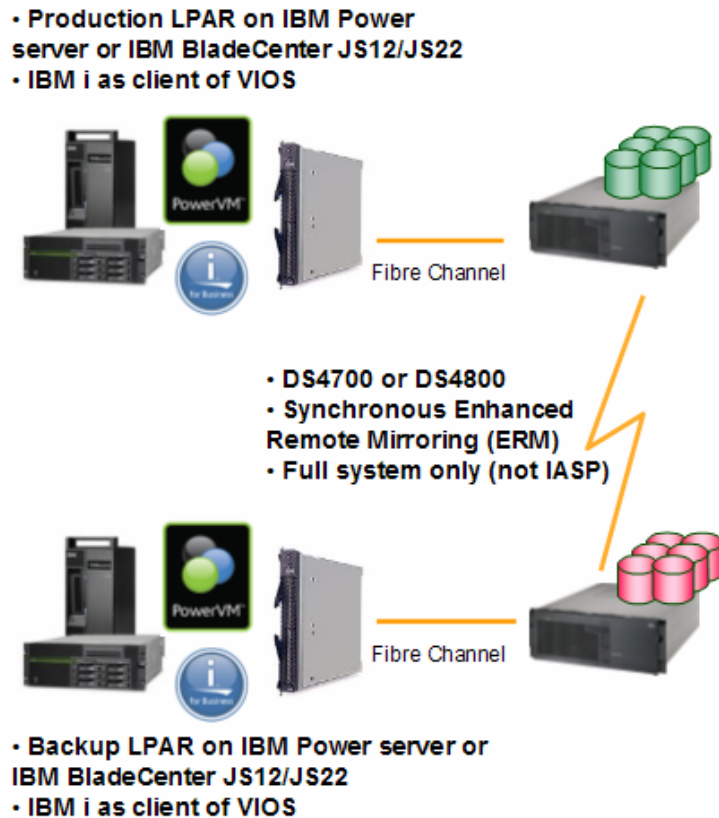
- Full-system FlashCopy and VolumeCopy of the production IBM i logical partition (LPAR) after only using the IBM i 6.1 memory flush to disk (quiesce) function are **not supported**
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is running are **not supported**
- FlashCopy and VolumeCopy 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 FlashCopy and VolumeCopy with IBM i, contact IBM Lab Services using this Web site: <http://www.ibm.com/systems/services/labservices/contact.html>.

12.2 Enhanced Remote Mirroring (ERM)

12.2.1. Test scenario

The following diagram shows the test environment used for ERM:



12.2.2 ERM support statements

The use of DS4000 Enhanced Remote Mirroring 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 on the DS4000 storage subsystem, in VIOS and in IBM i are required**. Currently, no toolkit exists that automates this solution and it is not part of IBM PowerHA for IBM i. The components of the solution – DS4000 ERM, VIOS and IBM i – must be managed separately and require the corresponding skill set. Note also 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, consult: <http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html>.

Support statements for ERM:

- DS4000 ERM is supported by IBM as a client of VIOS on both IBM Power servers and IBM Power blades
- Only Synchronous ERM (DS4000 Metro Mirror) is supported
- Asynchronous ERM (DS4000 Global Copy) or Asynchronous ERM with Write Consistency Group (DS4000 Global Mirror) are **not supported**
- Full-system Synchronous ERM for a planned switchover (IBM i production LPAR is powered off) is supported
- Full-system Synchronous ERM 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 Enhanced Remote Mirroring with IBM i, contact IBM Lab Services using this Web site: <http://www.ibm.com/systems/services/labservices/contact.html>.

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

13.1. 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 the SVC supports.

The concepts of using SVC storage for IBM i with VIOS are the same as those for connecting to DS4000 with VIOS. 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 a DS4000. Available *hdisks* are then assigned to an IBM i partition by mapping them to a Virtual SCSI server adapter (*vhost*) in VIOS, which is in turn connected to a Virtual SCSI client adapter in IBM i. Once 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, see section 8 of this document.

13.2. Attaching SVC storage to IBM i

The general process of attaching SVC storage to IBM i follows the one for using DS4000 storage:

- SVC 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 Redbook *Implementing the IBM System Storage SAN Volume Controller V4.3* (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, consult chapters 8, and 9 or 10 of the same Redbook, using the instructions for AIX hosts.

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

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. Consult the topic **Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC** in the *Logical Partitioning Guide*: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphata.pdf>.

14. SVC Copy Services and IBM i


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

14.1. FlashCopy

14.1.1. Test scenario

The following diagram shows the test environment used for FlashCopy:

- Production LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS

14.1.2. 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 exists that automates this solution 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. Note also 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 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

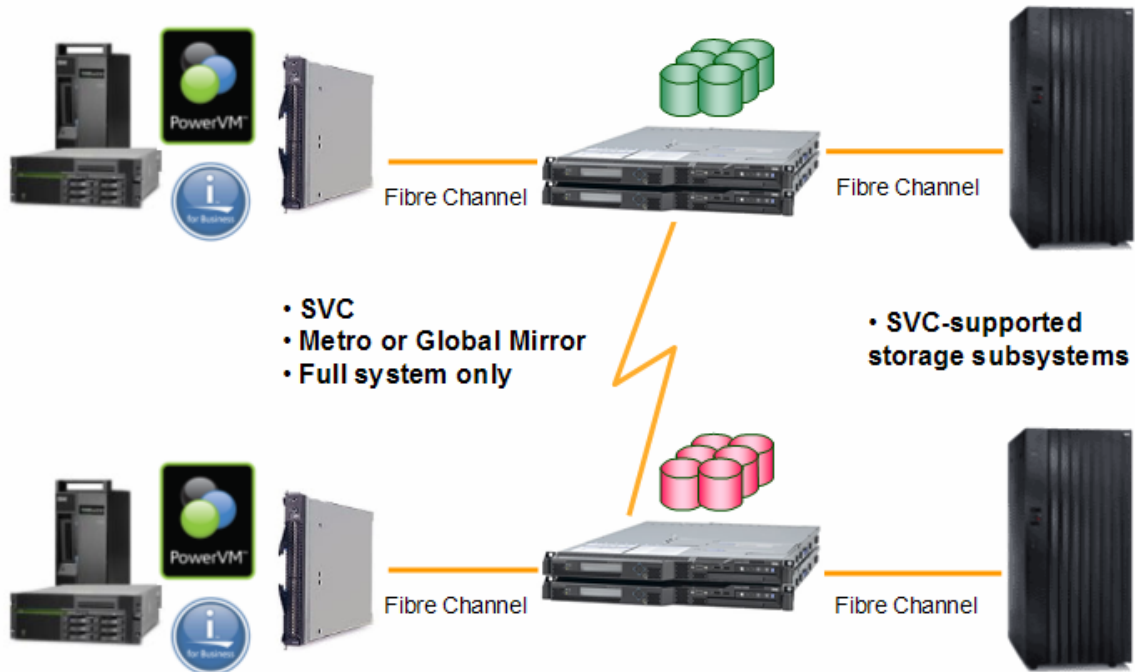
For assistance with using SVC FlashCopy with IBM i, contact IBM Lab Services using this Web site: <http://www.ibm.com/systems/services/labservices/contact.html>.

14.2. Metro and Global Mirror

14.2.1. Test scenario

The following diagram shows the test environment used for Metro and Global Mirror:

- Production LPAR on IBM Power server or IBM BladeCenter JS12/JS22*
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS

14.2.2. Metro and Global Mirror support statements

The use of SVC Metro and Global Mirror 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 exists that automates this solution 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. Note also 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, consult: <http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html>.

Support statements for SVC Metro and Global Mirror:

- Both SVC Metro and Global Mirror are supported by IBM 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 or Global Mirror is **not supported**
- Both Metro and Global Mirror are supported for a planned switchover (IBM i production partition is powered off)
- Both Metro 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 and Global Mirror

For assistance with using SVC Metro and Global Mirror with IBM i, contact IBM Lab Services using this Web site: <http://www.ibm.com/systems/services/labservices/contact.html>.

15. Additional resources

15.1. 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/iphath/iphath.pdf>.
- IBM i installation:
http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzahc/rzahc1.htm&tocNode=int_216451
- IBM i PTF installation:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/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/systems/scope/i5os/topic/rzaji/rzajioverview.htm>
- ECS over LAN:
http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaji/rzaji_setup.htm
- IBM i virtual client partitions topics in the IBM i Information Center:
<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahc/rzahcbladei5limits.htm>
- IBM PowerHA for IBM i: <http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html>.

15.2. Storage

- *IBM System Storage DS4000 and Storage Manager V10.10* (Redbook):
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247010.html?Open>
- DS4000 Web site: <http://www.ibm.com/systems/storage/disk/ds4000/index.html>
- *Performance Capabilities Reference* manual (chapter 14.5 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

15.3. VIOS

- VIOS Web site: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html>
- *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>

15.4. SVC

- SVC overview Web site: <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 V4.3* (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>.

16. 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. Consult 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 inquiries, 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 PVM, 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, SPECcapc, SPECchpc, 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.