HITACHI
Inspire the Next

Best Practice Library

Guidelines for Installing Hitachi HiCommand® Dynamic Link Manager and Multi-Path I/O in a Virtual I/O Environment

By Catherine Anderson, Advanced Technical Consultant

November 2006

Hitachi Data Systems

Executive Summary

This document provides best practices for the design and implementation of Hitachi HiCommand® Dynamic Link Manager and Hitachi storage systems with IBM's Advanced Virtualization feature for IBM p5® servers and IBM AIX® 5.3.

Although the recommendations documented here may generally represent good practices, configurations may vary.

No guarantees expressed or implied.

For further information on Hitachi Data Systems products and services, contact your Hitachi Data Systems representative, or visit Hitachi Data Systems online at http://www.hds.com.

Trademarks

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
1.0	November 2006	Initial Release

Source Documents for this Revision

IBM Virtualization Redbooks:

- http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf
- http://www.redbooks.ibm.com/redbooks/pdfs/sg245768.pdf

VIO FAQ's:

- http://techsupport.services.ibm.com/server/vios/documentation/faq.html

VIO Support Page:

- http://techsupport.services.ibm.com/server/vios

VIO Sizing:

- http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?topic=/iphb1/iphb1_vios_planning_vscsi_sizing.htm

Audience

This document is intended for AIX System Administrators responsible for designing and implementing the IBM Advanced virtualization feature of the p5 servers with AIX 5.3 and Hitachi Data Systems storage systems.

Authors

The information included in this document represents the expertise of a number of skilled practitioners within Hitachi Data Systems in the successful delivery of storage solutions to our customers. Contributors to this document include:

Catherine Anderson	Advanced Technical Consultant, Solution Sales and Services
Rick Butcher	IBM Alliance Manager, Global Solutions Strategy and Development
Melvyn Aguilar	Technical Response Center, Global Support Operations

Delivery of this guide was managed by Technical Marketing, Global Solutions Strategy and Development.

A broad range of additional skills and talents working across many disciplines within Hitachi Data Systems provided many hours of effort on quality assurance of this material. Thanks go to:

Russ Davies	Senior Program Manager, Technical Operations Partner Programs
Mike Sarem	Open Systems Consultant, Customer Support and Services
Delroy Blake	Solution Support Consultant, Solutions Sales and Services

Hitachi, Ltd. Engineers

Contact Information

All documents in this series can be found on www.hds.com.

Please send us your comments on this document. Make sure to include the document title, and revision. Please refer to specific section(s) and paragraph(s) whenever possible.

E-mail: TM@hds.com

Fax: 408-562-5477

Mail:
Technical Marketing, M/S 34-88
Hitachi Data Systems
750 Central Expressway
Santa Clara, CA 95050-2627

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

Contents

Best Practice Library

Guidelines for Installing Hitachi HiCommand® Dynamic Link Manager and Multi-Path I/O (MPIO) in a Virtual I/O Environment

By Catherine Anderson, Advanced Technical Consultant

Introduction to Virtual I/O

Problem Description

The introduction of IBM p5™ servers and the Advanced POWER™ Virtualization feature allows processors and I/O devices to be virtualized and shared. Virtualizing I/O adapters, called Virtual I/O or VIO, allows the sharing of these resources by multiple logical partitions. There are cost and consolidation advantages to utilizing this technology. But there are also challenges in creating and managing a virtual environment. This paper documents the requirements for installing Hitachi HiCommand® Dynamic Link Manager path manager software and IBM AIX® Multi-path I/O (MPIO) in a Virtual I/O server partition to enable sharing of Fibre Channel Host Bus Adapters between client logical partitions.

Technical Overview

POWER 5 and Advanced POWER Virtualization (APV)

The POWER Hypervisor is the foundation for virtualization on a System p5 server. It enables the hardware to be divided into multiple partitions, and ensures isolation between the partitions. It is always active on POWER5-based servers, and provides inter-partition communication that enables the Virtual I/O Server's virtual SCSI and virtual Ethernet function.

IBM's Advanced POWER Virtualization (or APV) feature is a set of hardware (firmware) and software that enable Virtual I/O (VIO) on IBM System p5 and IBM eServer® p5 servers to take advantage of special features in the IBM POWER5 or POWER5+ processor. The main components of the APV feature are the Virtual I/O Server (VIOS) that allows virtualization of SCSI storage devices and the Shared Ethernet Adapter (SEA). The APV also includes firmware enabling of the IBM Micro-Partitioning technology, which allows the establishment of sub-partitions to as little as 10 percent of a physical processor and a Partition Load Manager to balance workloads across physical processors.

The following diagram is from the IBM redbook on Virtualization and shows AIX 5.3, Virtual I/O, and micro-partitioning on the p5 Server:

Figure 1: Virtualization on the IBM p5 Server

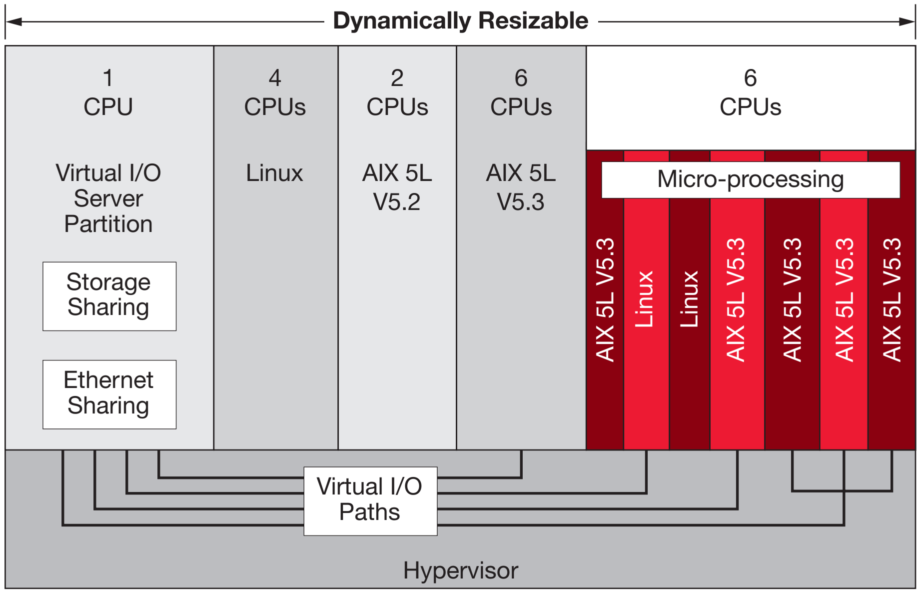

Virtualization is illustrated on the IBM p5 Server showing AIX 5.3, Virtual I/O and Micro-partitioning.

Two of the features of the p5 server Advanced Virtualization are micro-partitioning and Virtual I/O. Micro-partitioning is the ability to create a partition with as little as 10% of a CPU. This allows customers to run multiple applications and or multiple operating systems on a single processor. This can create a problem if there are not enough I/O slots available for all the partitions that can be created. The Advanced POWER Virtualization feature is a package that enables and manages the virtual environment on POWER5 systems. Advance Power Virtualization addresses the issue of not enough I/O slots with the Virtual I/O feature. The main technologies include:

- Virtual I/O Server
  - Virtual SCSI Server
  - Shared Ethernet Adapter
- Micro-Partitioning technology
  - Partition Load Manager

The primary benefit of Advanced POWER Virtualization is to increase overall utilization of system resources by allowing only the required amount of processor and I/O resource needed by each partition to be used.

- The optional Advanced POWER Virtualization feature allows partitions to be created that are in units of less than one CPU (sub-CPU LPARs) and allows the same system I/O to be virtually allocated to these partitions.
- With Advanced POWER Virtualization, the processors on the system can be partitioned into as many as 10 LPARs per processor.

AIX 5.3

AIX 5.3 is required to take advantage of the virtualization features of the p5 system. It supports the Advanced POWER Virtualization hardware feature, which includes:

- Micro-partitioning support for a single processor being shared by up to 10 logical partitions
- Virtual SCSI disks that allow partitions to access storage without requiring a physical storage adapter
- Virtual networking: Virtual Ethernet provides high-speed connections between partitions; Shared Ethernet Adapter provides connectivity between internal and external VLANs.

Virtual I/O (VIO) Overview

Virtual I/O was introduced with AIX 5.3 and the IBM p5 systems. Virtual I/O is the ability to share the physical resources (processors and I/O devices) in a Virtual I/O Server (VIOS) between multiple client logical partitions (VIOC) on the same p5 server. The Virtual I/O Server (see Figure 1 above) is a component of the IBM System p5 Advanced POWER Virtualization hardware feature which allows the sharing of physical resources between logical partitions including virtual SCSI and virtual networking (see note below). This allows more efficient utilization of physical resources by sharing physical resources between logical partitions.

The virtual VIOS and VIOC adapters are presented as SCSI adapters and all mandatory SCSI commands are supported. The devices (internal or SAN disks) exported from the VIO Server appear as hdisks and can be used as such in the VIO Client. The logical partitions, Virtual I/O server and clients, are created using the Hardware Management Console (HMC).

The enabling technologies behind virtual I/O are the advanced capabilities of POWER5 and IBM AIX 5L Version 5.3.

Note: Virtual networking enables inter-partition communication without the need for physical network adapters assigned to each partition and is beyond the scope of this discussion.

Virtual I/O Restrictions[1]

- Logical volumes used as virtual disks must be less than 1TB in size.
- Exported logical volumes used as virtual disks cannot be mirrored, striped or have bad block relocation enabled (Not a concern when using RAID disks).
- Virtual SCSI supports certain Fibre Channel, parallel SCSI, and SCSI RAID devices as backing devices (see the VIO support page for a list of supported configurations which includes Hitachi Data Systems/Hitachi HiCommand Dynamic Link Manager/MPIO at http://techsupport.services.ibm.com/server/vios).
- Virtual SCSI does not impose any software limitations on the number of supported adapters. A maximum of 256 virtual slots can be assigned to a single partition. Every virtual slot that is created requires resources in order to be instantiated. Therefore, the resources allocated to the Virtual I/O Server limits the number of virtual adapters that can be configured.
- The SCSI protocol defines mandatory and optional commands. While virtual SCSI supports all of the mandatory commands, some optional commands may not be supported at this time. This implies that exported Hitachi HiCommand Dynamic Link Manager devices can be used as client boot devices. However, this has not been tested and is not supported at this time. Please check with your Hitachi Data Systems representative for current Dynamic Link Manager software boot support.

1 For clarity and specificity, the restrictions are taken verbatim from the IBM VIO FAQs page which can be found at: http://techsupport.services.ibm.com/server/vios/documentation/faq.html

- Exported Hitachi HiCommand Dynamic Link Manager software devices cannot be used as boot devices for client partitions. Boot devices must be excluded from Dynamic Link Manager software's control by specifying them in the /usr/DynamicLinkManager/drv/dlmfdrv.unconf file. A single LU may be exported from two VIO servers to use as a client boot device if the reserve policy is disabled.
- The Virtual I/O Server is a dedicated partition to be used only for VIOS operations. No other applications can be run in the Virtual I/O Server partition. Exceptions are path managers and agents. Please refer to "Additional Software in VIOS" below for more information on installing agents in the Virtual I/O Server.
- VIO is currently supported for new disk installations only. Please check with IBM and Hitachi Data Systems for current support. This is required as Path Managers may not use the same disk identification method. As Path Managers implement the UDID method of identification (see below) this restriction could change. A current backup should always be available when making any changes in a VIO environment.
- Considerations for VSCSI devices:

  The VIOS uses several methods to uniquely identify a disk for use as a virtual SCSI disk, they are:

  – Unique device identifier(UDID)

  – IEEE volume identifier

  – Physical volume identifier(PVID)

  Each of these methods may result in different data formats on the disk. The preferred disk identification method for virtual disks is the use of UDIDs.

  MPIO uses the UDID method. Most non-MPIO disk storage multi-pathing software products use the PVID method (including Hitachi HiCommand Dynamic Link Manager version 5.6 and below) instead of the UDID method. Because of the different data format associated with the PVID method, customers with non-MPIO environments should be aware that certain future actions performed in the VIOS LPAR may require data migration, that is, some type of backup and restore of the attached disks. These actions may include, but are not limited to the following:

  – Conversion from a non-MPIO environment to MPIO

  – Conversion from the PVID to the UDID method of disk identification

  – Removal and rediscovery of the Disk Storage ODM entries

  – Updating non-MPIO multi-pathing software under certain circumstances

  – Possible future enhancements to VIO
- Because of the methods VIOS uses to identify the disks, Hitachi HiCommand Dynamic Link Manager software version 5.6.1 or above with the unique_id enabled is the recommended path manager level in a VIO Server. MPIO by default uses the unique_id method of identification. Hitachi HiCommand Dynamic Link Manager by default does not use the unique_id. By enabling the unique_id attribute with Dynamic Link Manager software version 5.6.1, Hitachi HiCommand Dynamic Link Manager will not require the conversion listed above for converting from MPIO to Hitachi HiCommand Dynamic Link Manager or vice versa. Hitachi HiCommand Dynamic Link Manager software version 5.6.1 with the unique_id enabled allows more flexibility in changing a Virtual I/O configuration.

AIX 5.3 and p5 Advanced POWER Virtualization

Business Impact

The Power 5 architecture and AIX 5.3 are very attractive to customers. This technology allows system administrators to create logical partitions using only the resources necessary to efficiently support applications. Processing power can be allocated as needed and I/O can be virtualized to support many virtual client logical partitions. The cost savings of sharing physical resources is beneficial to customers both large and small. An additional benefit is the ability to move resources dynamically between partitions. However, I/O devices can be dedicated to partitions if the workload requires non-shared resources. A client partition has the flexibility to use both virtual and dedicated I/O resources. Today only AIX 5.3 is supported with Virtual I/O. Future supported platforms may include Linux.

Evaluation Criteria

Hitachi Data Systems should be contacted prior to implementing Virtual I/O to ensure the latest configuration information is available. Also see “Sizing and Design” below for other considerations.

The hardware requirements for Virtual I/O are a p5 server with the Advanced POWER Virtualization feature and a Hitachi storage system. Software required includes AIX 5.3, Hitachi HiCommand Dynamic Link Manager or MPIO, and the VIO server code. A single VIO server is supported but not recommended for a highly available system. Two VIO servers with access to the same disk devices are recommended for high availability.

Virtual I/O Configurations

Although a single VIO Server is supported, multi-pathing software is recommended for high availability. Multi-pathing software enhances RAID system availability and performance by providing for path failover and load balancing in the event of channel-connection failures. Both Hitachi HiCommand Dynamic Link Manager and MPIO support Virtual I/O Server and multi-pathing. Consider these three important configuration guidelines:

Either Hitachi HiCommand Dynamic Link Manager or MPIO can be used with Hitachi enterprise storage (that is to say the TagmaStore Universal Storage Platform or the Hitachi Lightning 9900 or 9900V Series systems).

Only Dynamic Link Manager can be used with Hitachi midrange TagmaStore Advanced Modular Storage and Workgroup Modular Storage systems.

MPIO can be installed in the VIOS, but is required in the VIOC because virtual SCSI adapters are used. To enable multi-pathing with MPIO in the VIOC, it is highly recommended to use two VIO Servers for high availability.

Basic Configuration

The following VIO configurations are tested and qualified:

- Single VIOS, single path to storage, multiple VIO clients
- Single VIOS, multi-path to storage using MPIO, multiple VIO clients
- Single VIOS, multi-path to storage using Hitachi HiCommand Dynamic Link Manager software, multiple VIO clients
- VIO coexistence on client with Dynamic Link Manager software for direct (natively) FC-attached disks

Figure 2: Single VIOS with Multi-path to Storage

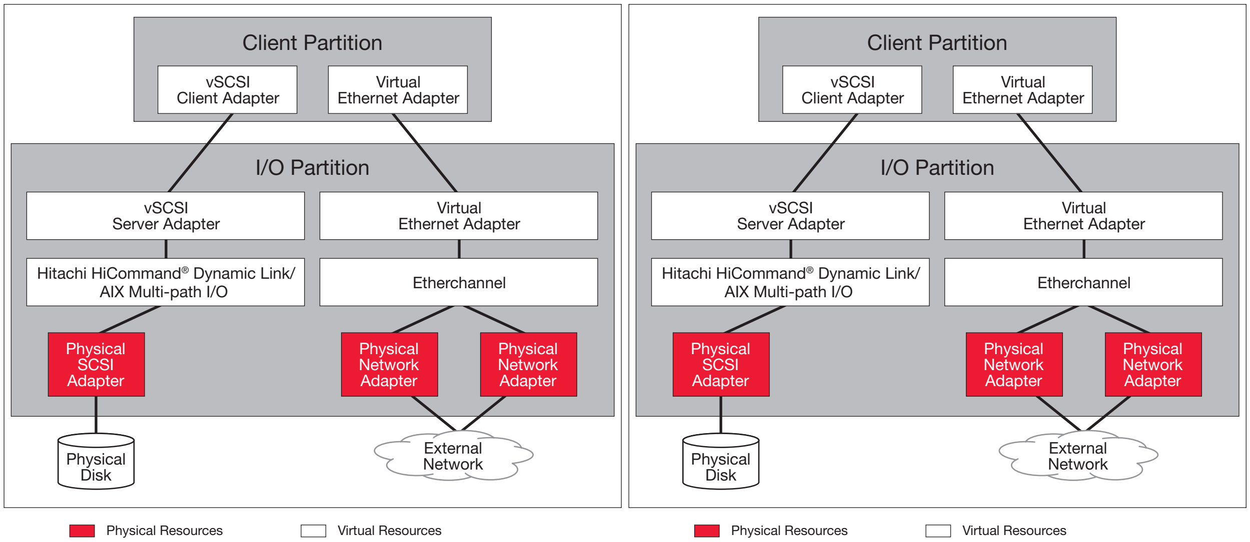

MPIO or Hitachi HiCommand Dynamic Link Manager software support multi-pathing in a single Virtual I/O Server with a single path or dual paths to disks.

Recommendations:

- Hitachi HiCommand Dynamic Link Manager software version 5.6.1 with unique_id enabled (see Hitachi HiCommand Dynamic Link Manager software documentation).
- The exported virtual devices should be physical volumes. Exporting logical volumes is supported but not recommended. There are restrictions on exporting logical volumes as virtual SCSI devices.
- If using two (or more) paths and a switch, change the HBAs attribute to fast_fail.
- Hitachi HiCommand Dynamic Link Manager devices should not be exported for use as client boot. This is not a supported configuration.

The configurations above are supported, but have a single point of failure in the VIOS if a reboot or maintenance is needed in the Virtual I/O server.

Hitachi HiCommand Dynamic Link Manager software in this environment may not fail over a path correctly when an open (varyonvg) is issued from the VIO client at VIOS level 1.1.2.62 (see Hitachi HiCommand Dynamic Link Manager release notes at time of installation).

Advanced Configuration

The following VIO configurations are tested and qualified:

- Dual VIOS, single path to storage from each VIOS, multi-path to VIO client, multiple VIO clients
- Dual VIOS, multi-path to storage from each VIOS, multi-path to VIO client, multiple VIO clients (Dynamic Link Manager software or MPIO)

Figure 3: A Dual VIOS Configuration

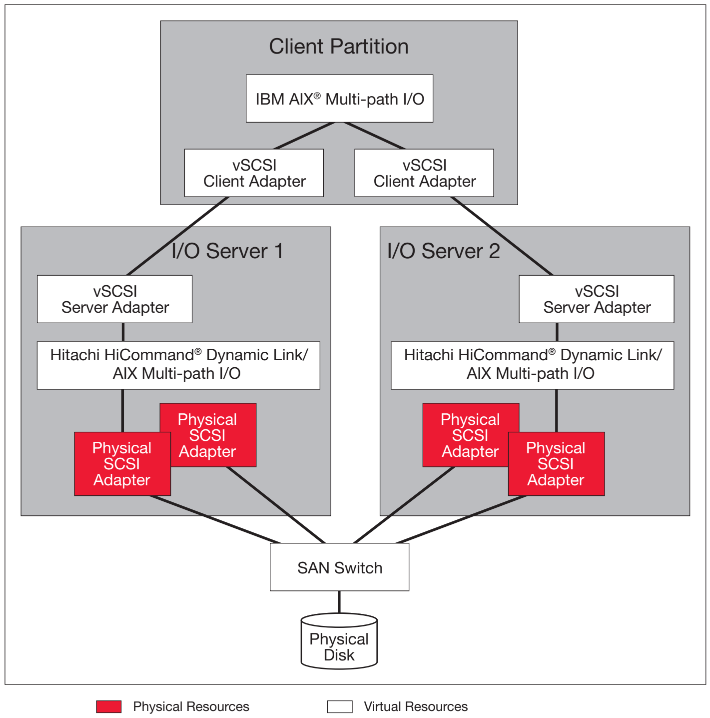

As in Figure 2, MPIO or Hitachi HiCommand Dynamic Link Manager software support of multi-pathing is illustrated here in a dual VIOS configuration.

Requirements for this configuration:

- Reservation policy must be disabled:
  - Hitachi HiCommand Dynamic Link Manager software:
    /usr/DynamicLinkManager/bin/dlnkmgr set rsv on (0) –s
  - MPIO
    chdev –l hdiskx –a reserve_policy=no_reserve

    chdev –l hdiskx –a algorithm=round_robin

Recommendations:

- Hitachi HiCommand Dynamic Link Manager software version 5.6.1 with unique_id enabled (see Hitachi HiCommand Dynamic Link Manager software documentation).
- The exported virtual devices must be physical volumes. Exporting logical volumes is not supported in a dual VIO server configuration.
- Change the HBAs to fast_fail when using more than one path and a switch.
- This is the recommended configuration to reduce single points of failure.
- MPIO in the client is failover only (active/passive). However, MPIO in the client can be changed so both VIO Servers are utilized. This can be accomplished by changing the priority of the path for each device.

- Hitachi HiCommand Dynamic Link Manager devices should not be exported for use as client boot. This is not a supported configuration.

Clustered Configuration

IBM's High Availability Cluster Multi-Processing (HACMP) software is recommended for high availability of the LPAR resources—within a server for software failures or between servers for software and hardware failures as illustrated in Figures 4 and 5 below. It improves computing availability in a manner similar in concept to multi-pathing software which performs the same failover function in the event of a path failure.

Figure 4: Active/Active Cluster Between Two Logical Partitions on One p5 Server

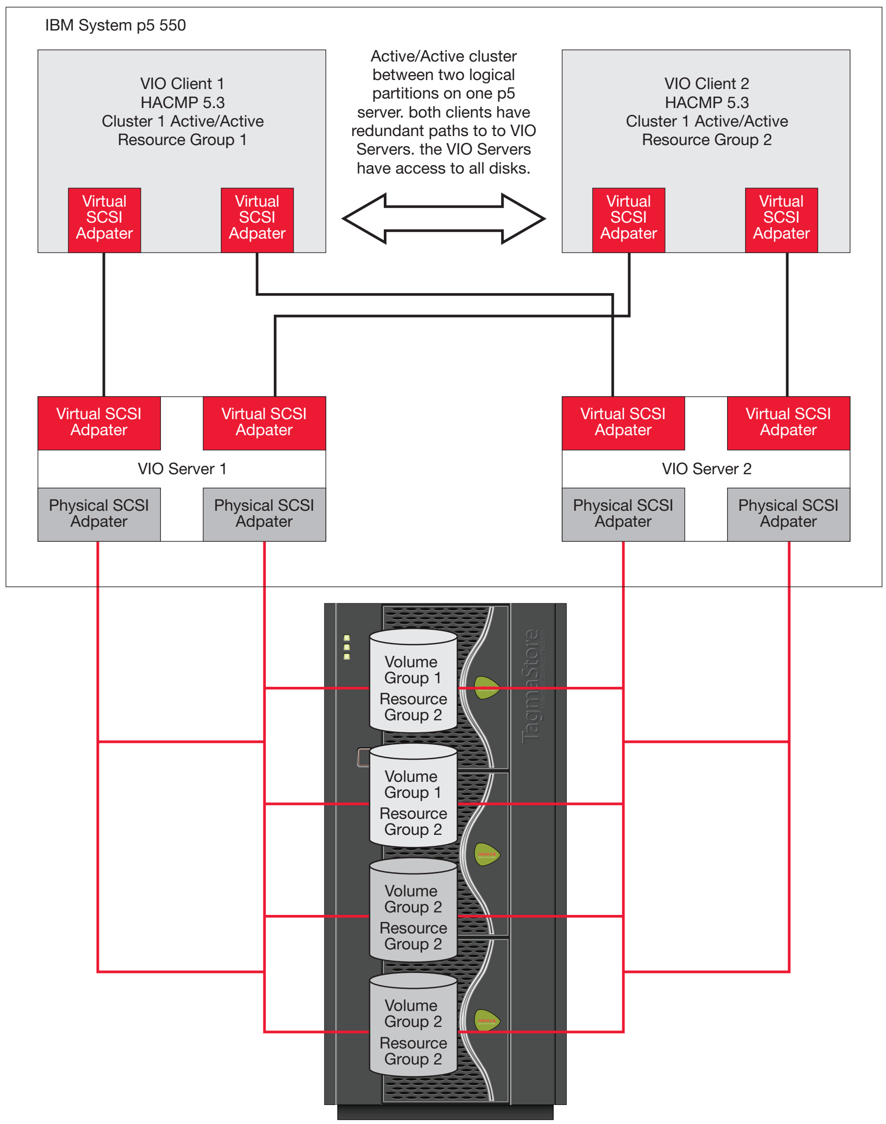

This illustration depicts an Active/Active Cluster between Two Logical Partitions within a single p5 Server.

Note: HACMP runs on all servers in a clustered computing environment but is implemented only in VIOC partitions and not in VIOS partitions.

Figure 5: Active/Active Cluster Between Two p5 Servers

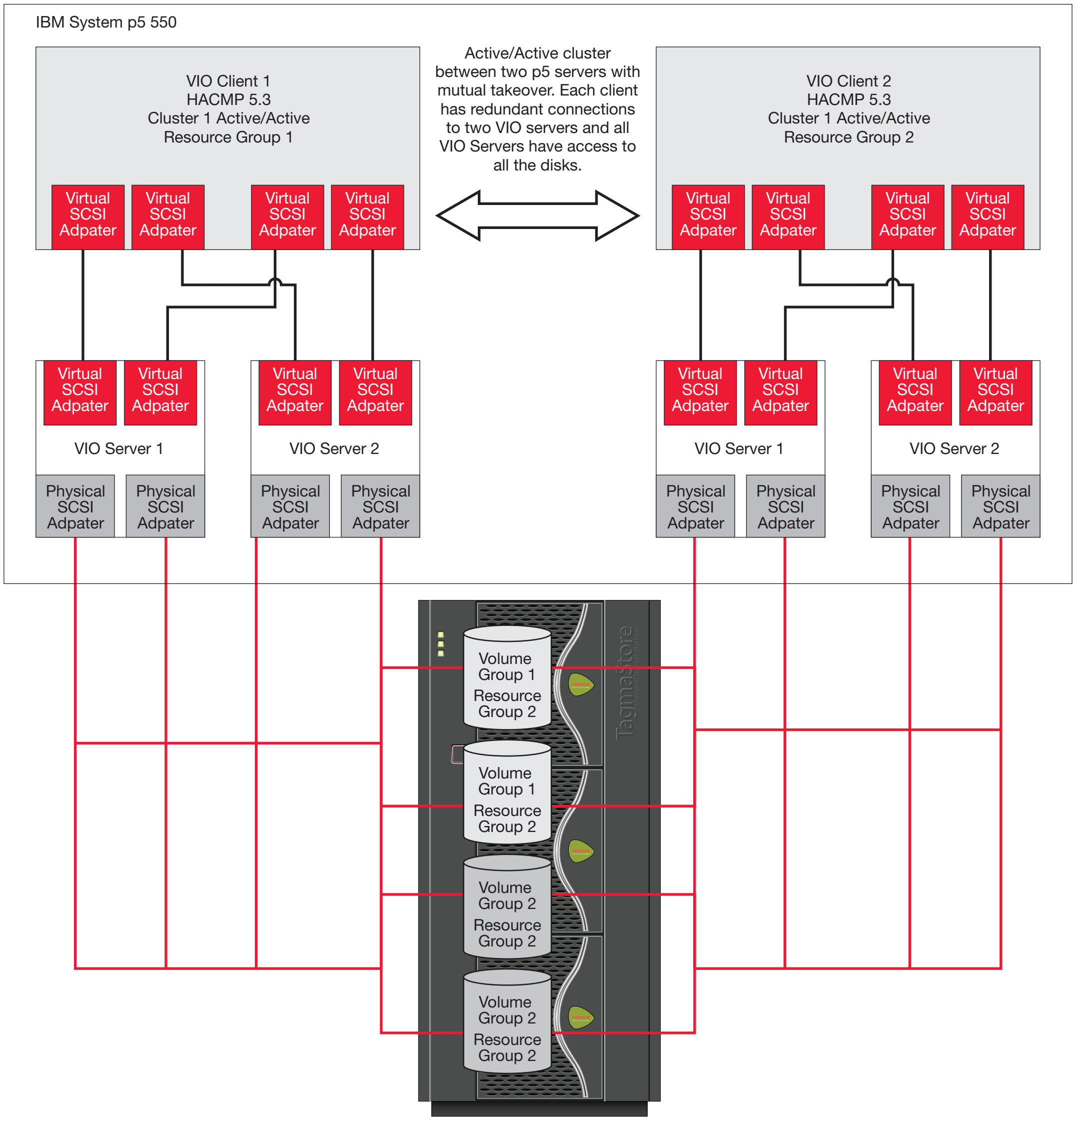

In contrast to Figure 4, this illustration also depicts an Active/Active Cluster, but between two p5 Servers.

It is important to realize that HACMP implementation is not trivial and it is best to consult both IBM and Hitachi Data Systems for the latest information prior to implementation. Technical issues to evaluate when considering HACMP are described below in "Virtual I/O in HACMP Cluster Environments". And the reader may refer to the following link: http://www.redbooks.ibm.com/redbooks/pdfs/sg246769.pdf

SAN Boot

SAN boot is supported for qualified environments using a single path or MPIO. Please consult Hitachi Data Systems for qualified environments. The information required to determine support is the OS level, HBA bios level, switch, and subsystem.

SAN boot is not supported today with Hitachi HiCommand Dynamic Link Manager software; Hitachi HiCommand Dynamic Link Manager cannot manage SAN boot devices even in a Virtual I/O environment. To implement SAN boot in an environment using Dynamic Link Manager software, the disks must be excluded from Dynamic Link Manager software's control by putting the hdisks in the /usr/DynamicLinkManager/drv/dlmfdrv.unconf file.

To provide high availability for the SAN boot devices in a single path environment, AIX Logical Volume Manager may be used to mirror the *rootvg* to another disk on another path.

Hitachi HiCommand Dynamic Link Manager Software with SAN Boot Using LVM Mirroring

Figure 6: SAN Boot with LVM Mirroring and Hitachi HiCommand Dynamic Link Manager

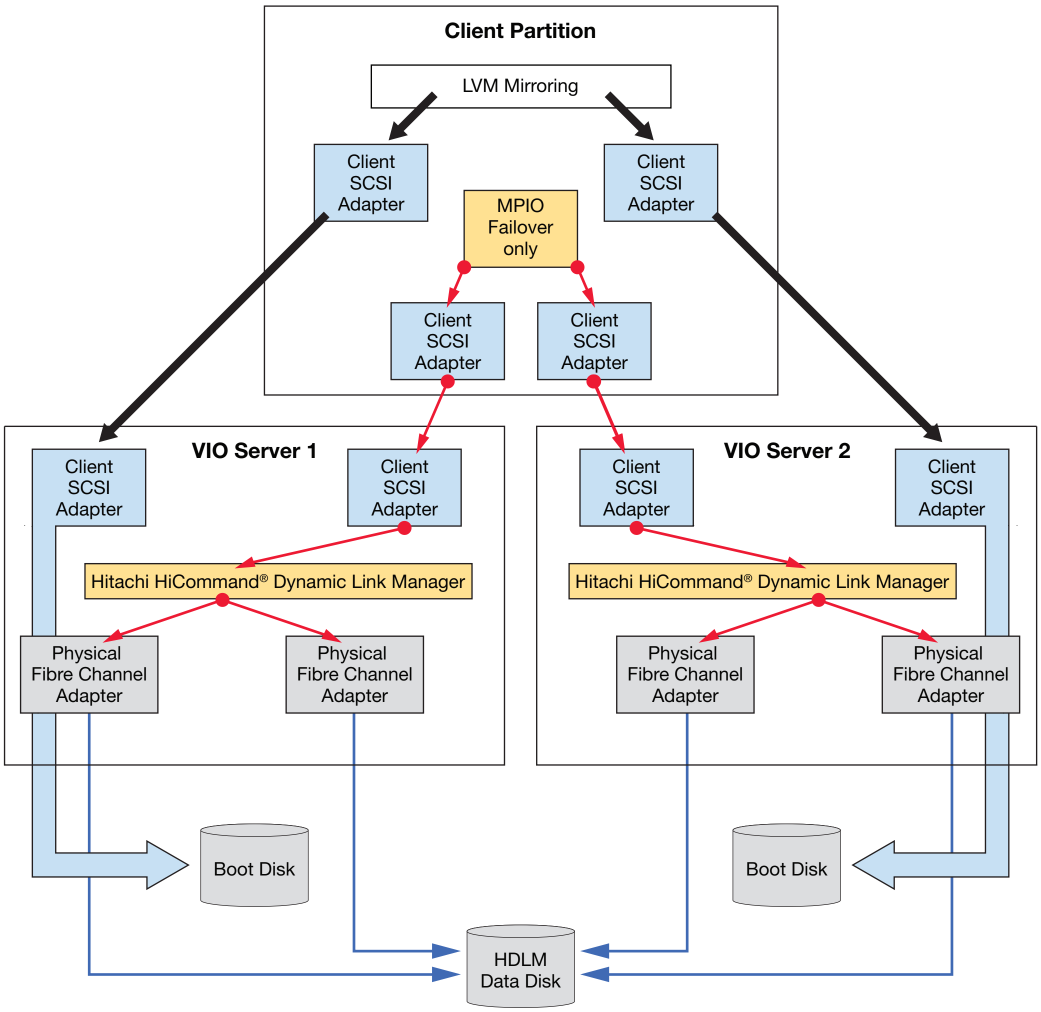

LVM Mirroring is used here for redundancy of the SAN boot LUNs that have been excluded from Hitachi Dynamic Link Manager.

Two LUNs on separate paths must be excluded from Hitachi HiCommand Dynamic Link Manager software's control and exported to the client. In the client, AIX LVM mirroring is used to provide redundancy for the two LUNs exported to be used as boot devices. Separate client adapters are not required, but are recommended.

SAN Boot Using MPIO

Figure 7: SAN Boot using MPIO

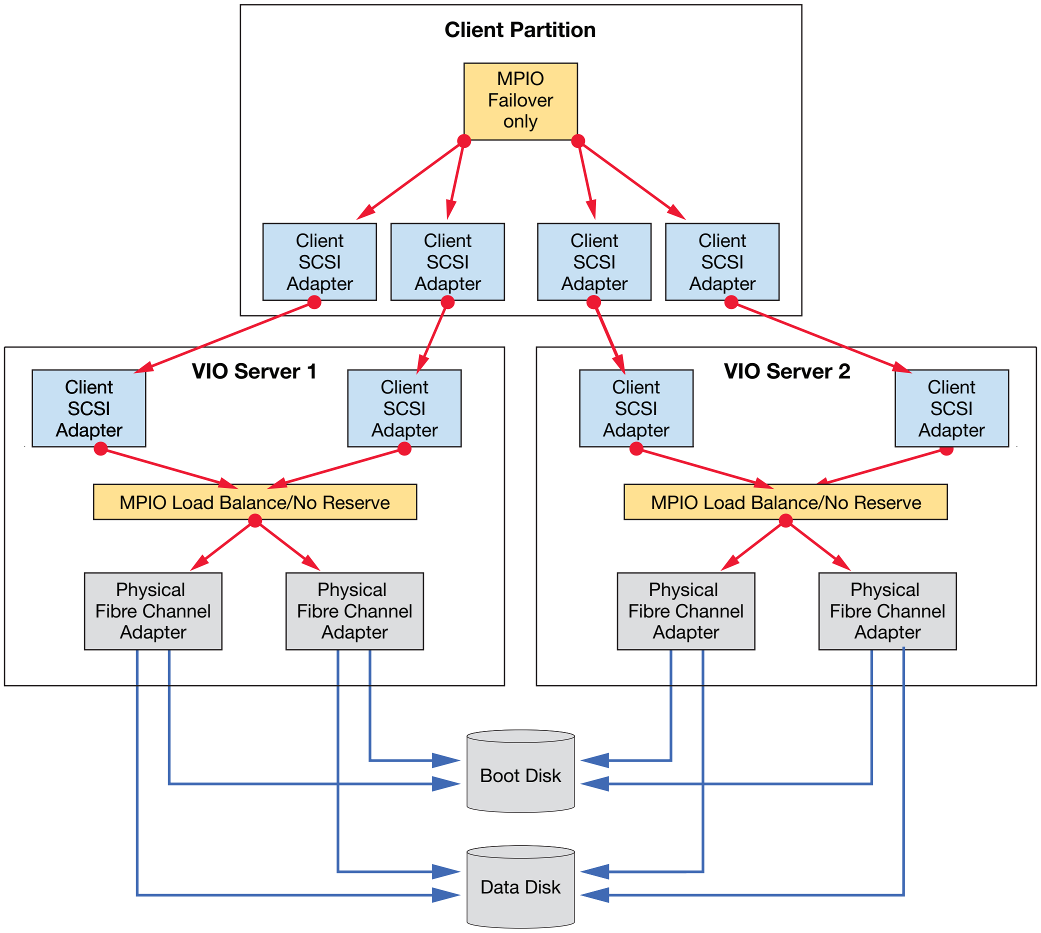

Figure 7 shows MPIO providing path failover to support SAN Boot.

MPIO manages multiple paths to the boot device. The reserves must be disabled and load balancing may be used in the VIO servers. MPIO in the client is failover only. Separate virtual client adapters are not required, but are recommended.

SAN Boot with Hitachi HiCommand Dynamic Link Manager Software Using a Single Boot Device

Figure 8: SAN Boot Using a Single Boot Device with Hitachi HiCommand Dynamic Link Manager Software

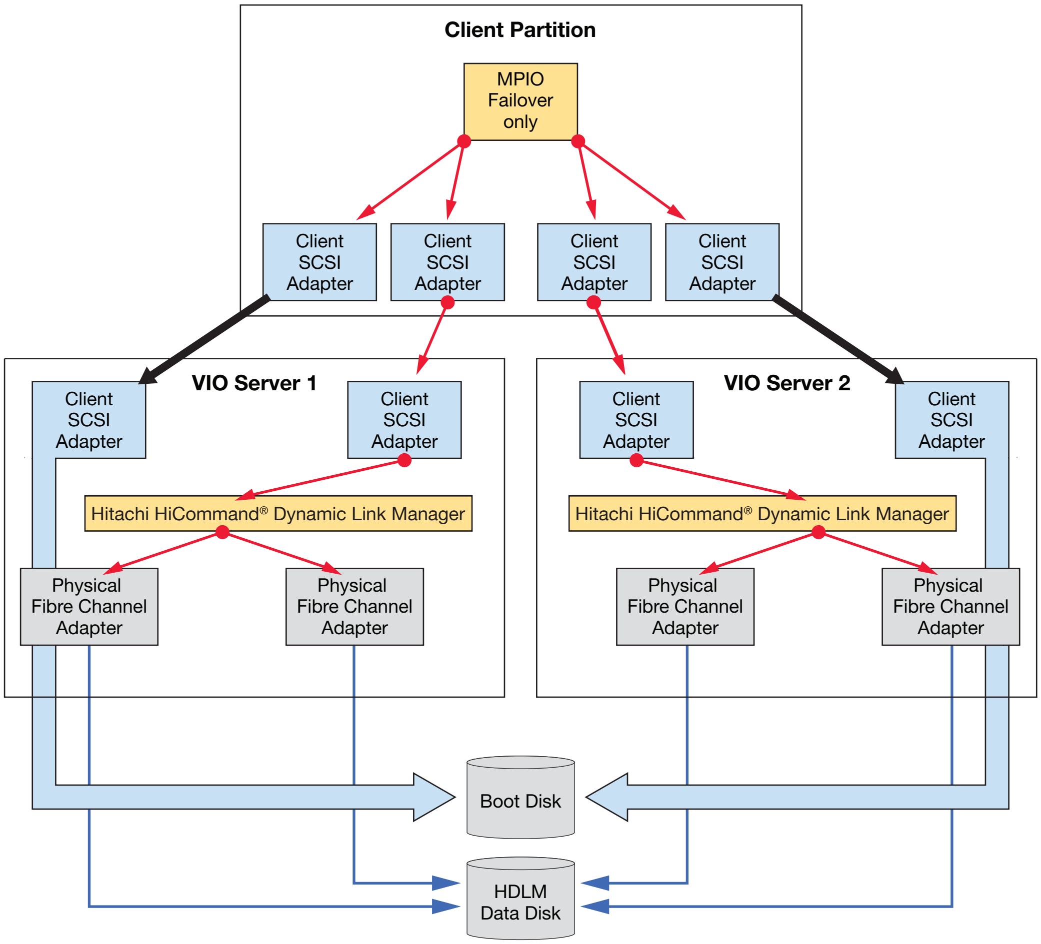

As in Figure 7, MPIO provides path failover to the boot disk while Hitachi HiCommand Dynamic Link Manager software balances access to data disks.

The boot device must be excluded from Hitachi HiCommand Dynamic Link Manager software control by entering the hdisk name in /usr/DynamicLinkManager/drv/dlmfdrv.unconf on both VIO servers. For example, if the boot disk is 'hdisk1', the contents of /usr/DynamicLinkManager/drv/dlmfdrv.unconf on each VIO server will be the single line entry: "hdisk1".

The 5.0.52.1 ODM update is required so the reserve on the disk can be disabled. There is a single path to the boot device from both VIO servers. MPIO is used in the client to provide failover only for the exported boot device. Separate virtual client adapters are not required, but are recommended.

Implementation Planning

Resource Requirements

- P5 server with the Advanced Virtualization Feature
- VIOS must have FixPack 6.2 applied (ioslevel 1.1.2.62). Please check the Hitachi HiCommand Dynamic Link Manager release notes for VIO Server levels supported.
- VIOC LPARs must be AIX 5.3 ML 1, with APARs IY70148, IY70336, and IY70082, and higher
- HBA 6239, 5716, 5758/5759
- Hitachi HiCommand Dynamic Link Manager software version 5.4.2 or above (Hitachi HiCommand Dynamic Link Manager software version 5.6.1 with unique_id is recommended)
- Subsystem requirements:
  - Port option “0F” (AIX) is required
  - Hitachi HiCommand Dynamic Link Manager software or MPIO can be used for Hitachi Lightning 9900 and 9900 V Series systems, and TagmaStore Universal Storage Platforms and Network Storage Controllers
  - Hitachi HiCommand Dynamic Link Manager software only can be used for Hitachi Thunder 9500 V Series systems, and TagmaStore Adaptable Modular Storage and Workgroup Modular Storage
  - Microcode: Contact your Hitachi Data Systems representative for microcode requirements
- All critical fixes, and interim fix IY72974, found at http://techsupport.services.ibm.com/server/aix.efixmgmt/fixes.html, must be applied
- Virtual devices must have a PVID prior to being exported to a VIO client when using Hitachi HiCommand Dynamic Link Manager without unique_id. It is recommended that all exported devices have a PVID.
- Devices with existing data (volume groups) cannot be used as virtual devices
- Storage must be configured for AIX. Select “AIX” as platform in the host groups
- An Hitachi ODM update (disk pre-defines) is required for disk definition

  For Hitachi HiCommand Dynamic Link Manager software, the ODM updates are:

  - 5.0.0.1 Base level
  - 5.0.0.4 Update for device support
  - 5.0.52.1 Update for AIX 5.2
  - 5.0.52.2 Update for HP/HDS compatibility

  For MPIO, the ODM updates are:

  - 5.4.0.0 Base level
  - 5.4.0.1 Update for MPIO health check

Sizing & Design

Performance should be considered when implementing virtual SCSI devices. With the use of POWER Hypervisor calls, virtual SCSI will use additional CPU cycles when processing I/O requests. More CPU cycles on the Virtual I/O Server are used when there is a load on the virtual SCSI devices. The virtual SCSI requires about 20 percent of a CPU for large block streaming or as much as 70 percent for small block, high transaction rate workloads for a dual disk adapter environment. Provided that there is sufficient CPU processing capacity available, the performance of virtual SCSI should be comparable to dedicated I/O devices. Suitable applications for virtual SCSI include boot disks for the operating system or Web servers, which will typically cache a lot of data.

Performance on a single physical port will not exceed the normal I/O performance characteristics of the port. Therefore the sum of all of the VIOS partitions attaching to a single physical port will be limited to that port's capabilities. For this reason, it is important to spread critical high I/O requirements over multiple ports.

When designing a virtual I/O configuration, performance is an important aspect that should be given careful consideration. Virtual SCSI runs at low priority interrupt levels, while virtual Ethernet runs on high priority interrupts due to the latency differences between disks and LAN adapters. A client that produces extensive network activity has the potential to impact performance of a client that requires extensive disk activity as the virtual Ethernet has higher priority interrupts. Either a larger Virtual I/O Server with more CPU resources will be required or a second Virtual I/O Server can be considered to split the two high throughput clients (see Advanced Configuration). The Virtual I/O client can support both virtual and physical adapters for applications requiring high I/O bandwidth.

The IBM Systems Hardware Information Center contains information on Virtual I/O sizing and should be reviewed prior to implementing Virtual I/O.

IBM has documents on VIO sizing and design that should be consulted prior to implementation:
http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?topic=/iphb1/iphb1_vios_planning_vscsi_sizing.htm

Installation & Configuration

VIO Server and VIO Client Implementation

Use the IBM documentation at:
http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm
to install and configure VIO server and client partitions. Search on "VIO install" for detailed instructions. The following instructions assume the VIOS and VIOC partitions have been created, but the virtual Fibre Channel adapters have not yet been created.

Create the VIO hosting partitions:

- Select memory, CPU units (minimum 10%), and physical adapters.

Create the VIO client:

- Select memory, CPU units (minimum 10%)
- Usually no physical adapters (if there are two physical Fibre Channel adapters then Hitachi HiCommand Dynamic Link Manager software can be installed in the client partition).

Figure 9: Example of VIO Servers and Client from the Hardware Management Console

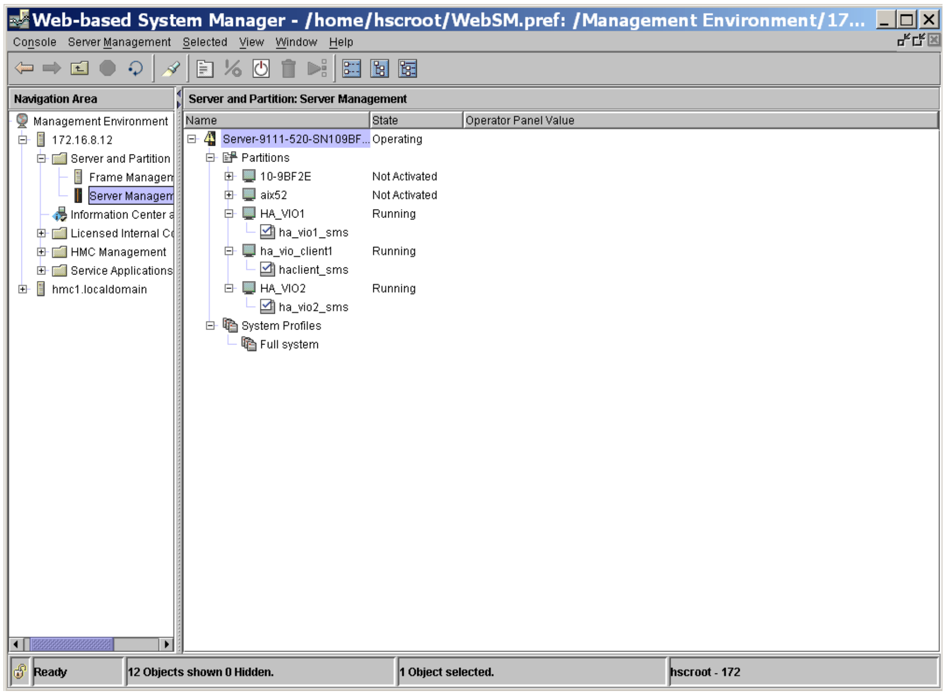

Figure 9 illustrates the view of VIO servers and client from the HMC (installed and active).

Figure 10: Create a Virtual Adapter in the First VIO Host Partition

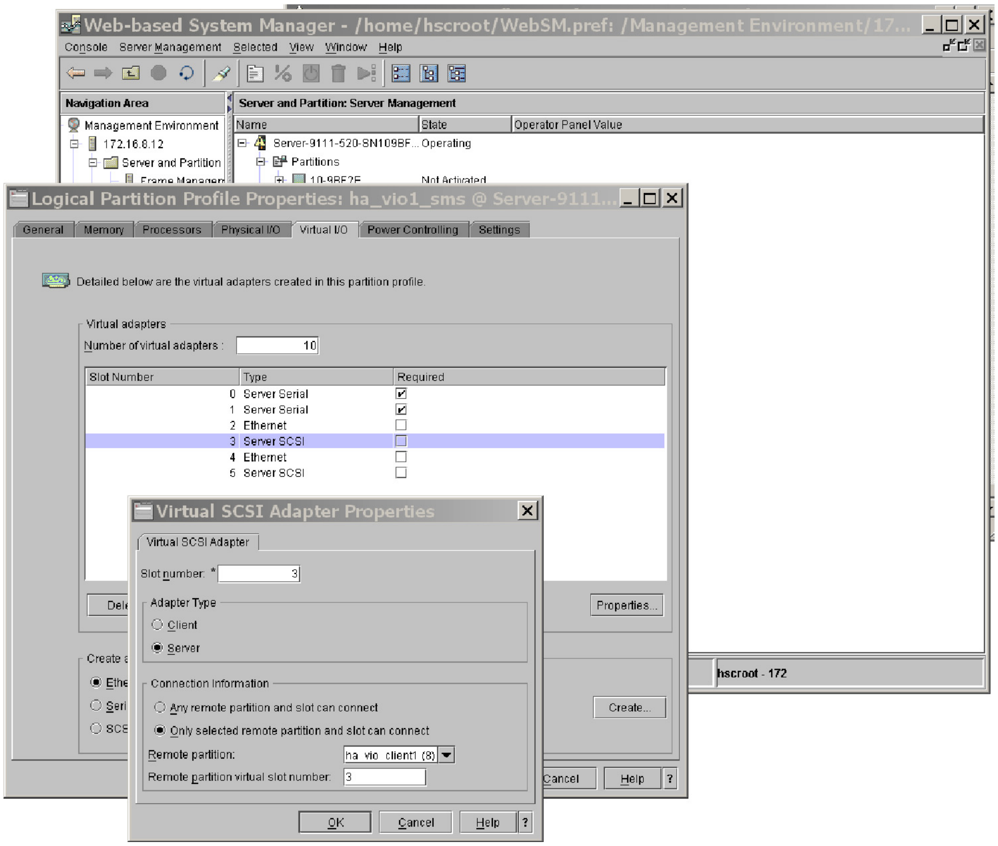

Virtual slot number three is used in creating a virtual adapter in the first VIO host partition.

Create a virtual adapter in the first VIO host partition.

Virtual slot number three will be used in the first VIO server for the Hitachi HiCommand Dynamic Link Manager software devices in VIO client.

Choose server and the remote client partition and the remote partition virtual slot number. It is easier to keep track of if they are the same.

Rules for virtual slot numbers:

This slot number does not refer to any physical hardware location on your system. You can, therefore, assign slot numbers to virtual adapters in any way that makes sense to you, provided that you follow these guidelines:

- You can use any slot number from 2 up to (but not including) the maximum number of virtual adapters. Slots 0 and 1 are reserved for system-created virtual adapters. By default, the system displays the lowest unused slot number for this logical partition.
- You cannot use a slot number that was used for any other virtual adapter on the same logical partition.

Figure 11: Create Virtual Adapter in the Second VIO Host Partition

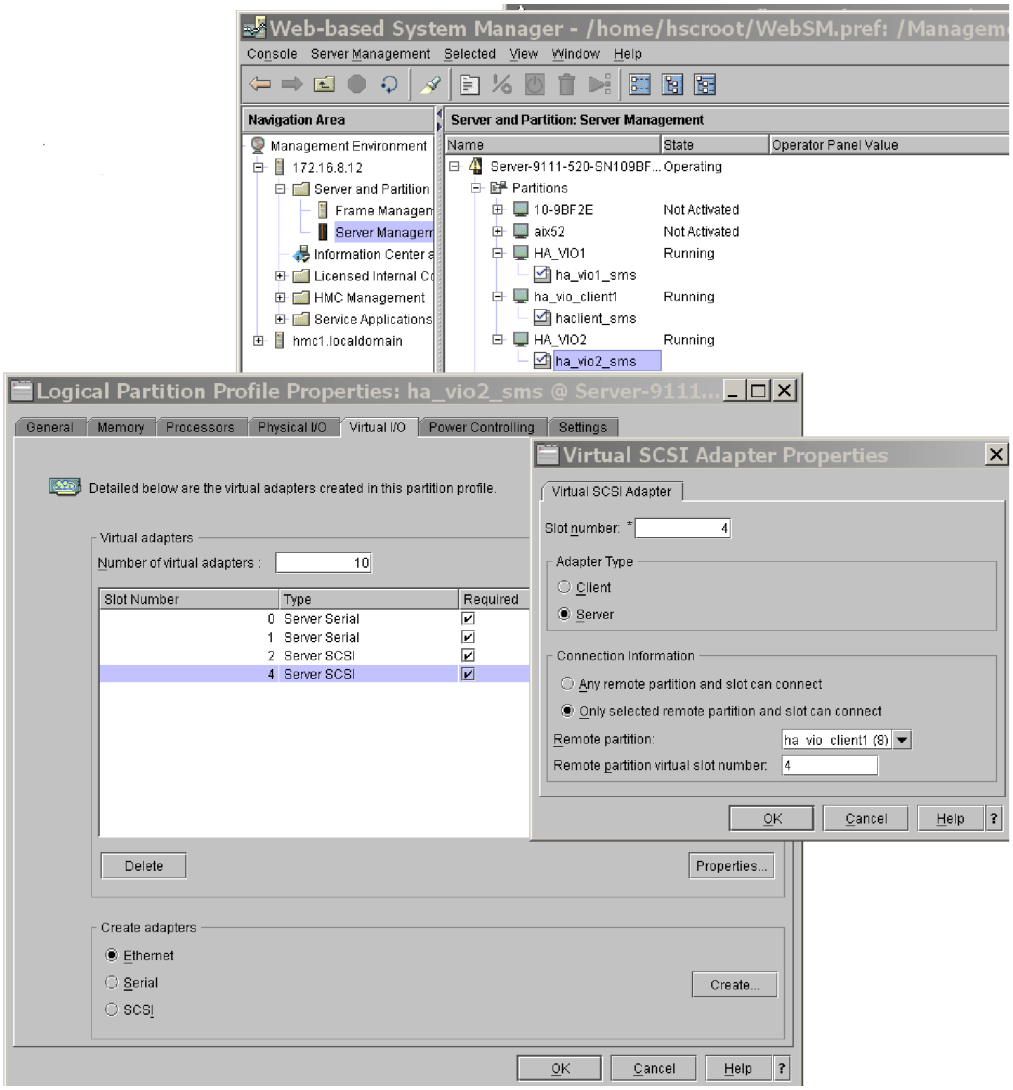

Virtual adapter four is defined in the second VIO host partition.

Virtual adapter four is defined in the second VIO host partition.

Note: The partitions must be (re)activated to recognize the changes in the partition profiles.

Figure 12: Example of Matching Client Adapters and Physical Host Adapter Numbers

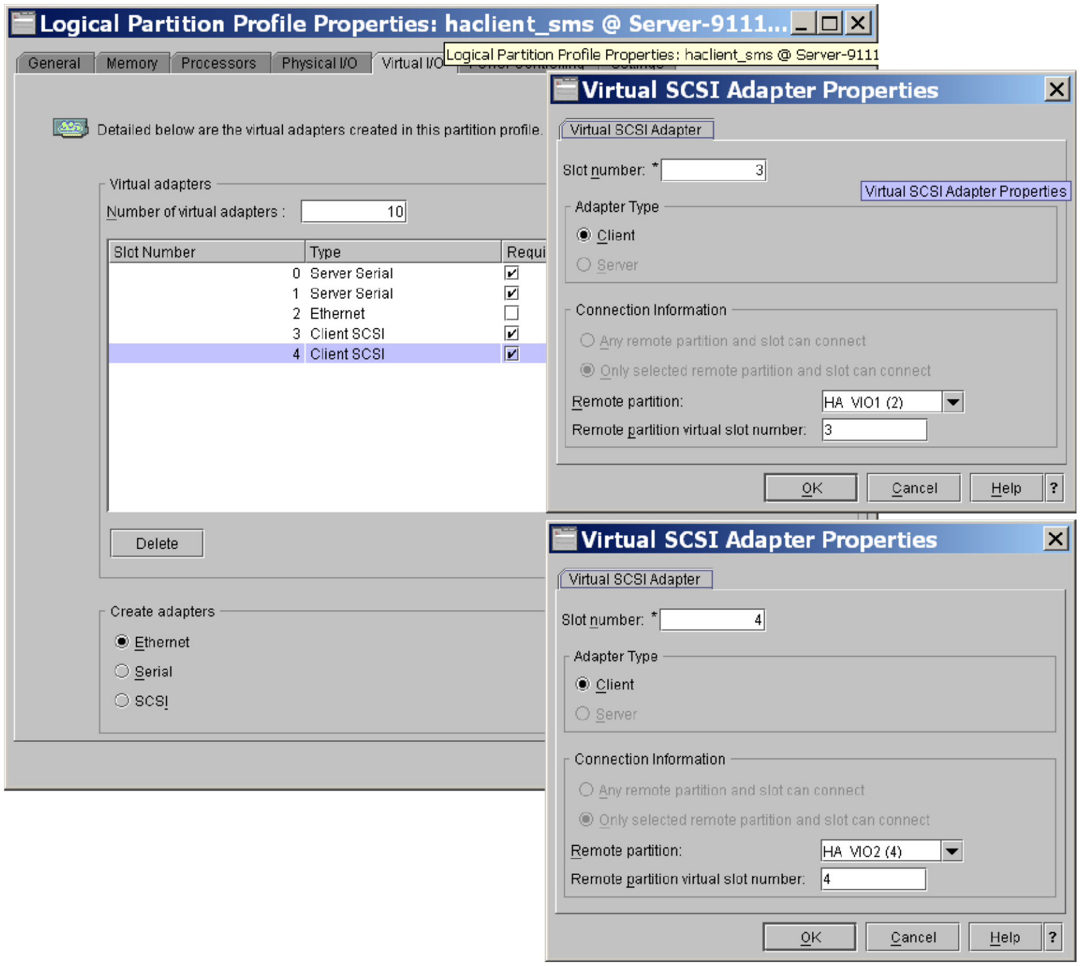

Figure 12 illustrates use of matching client adapter and physical host adapter numbers. Virtual SCSI adapter slot 3 is VIOS 1 slot 3 and virtual SCSI adapter slot 4 is VIOS 2 slot 4.

Define the client partition virtual adapters.

Select profile to be modified and click on properties. Select the Virtual I/O tab. Add the virtual I/O adapters for the client. There will be one from each VIO host partition (server).

In this example, the client adapters match the host adapter numbers. Virtual adapter three is “VIOS 1” and virtual adapter four is “VIOS 2”.

Install Hitachi HiCommand Dynamic Link Manager Software in the VIOS

The following steps are detailed installation instructions for Hitachi HiCommand Dynamic Link Manager software in the VIO server partition(s). These steps can be done before or after creating the virtual adapters. However, the virtual adapters must be created prior to making the virtual devices.

1. Login using the padmin command

2. Enter the following command to install the ODM update and Hitachi HiCommand Dynamic Link Manager software

   $ oem_setup_env

3. Remove any existing devices that will be used as virtual devices:

   # rmdev –dl hdisk(x)

4. Change the Fibre Channel adapters to fast fail:

   # chdev –l fscsix –a fc_err_recov=fast_fail (-P)

Use the "-P" option if the adapters are in use and the hdisks cannot be removed. A reboot after installing Hitachi HiCommand Dynamic Link Manager software will change the adapter attribute.

Note: The –P option cannot be used with Dynamic Link Manager devices or with hdisks that are part of Dynamic Link Manager devices.

5. Install the ODM updates (5001+5004+50521)

   # cfgmgr –vl fcs(x)

   Run once for each Fibre Channel adapter. Use the "lsdev –Cc disk" command to ensure the disks are defined as "Hitachi".

6. Install Dynamic Link Manager software (Use the Hitachi HiCommand Dynamic Link Manager software installation procedures in the User's Guide)

   6.1. Change the "dlmodmset" parameters

      a. *dlmodmset –e on*
         Default value for Hitachi HiCommand Dynamic Link Manager software version 5.6

      b. *dlmodmset –b xxxxx*
         Use the Hitachi HiCommand Dynamic Link Manager software pbuf calculation for AIX 5.3 to determine the correct number to set. The formula is:
         *pv_min_pbuf> * <# of LU> + 2560*
         where*: <pv_min_pbuf> is the pv_min_pbuf value displayed in the "ioo" command*

      c. *dlmodmset –u on*
         Required for Hitachi HiCommand Dynamic Link Manager software to use the unique_id attribute

   6.2. Create a /usr/DynamicLinkManager/drv/dlmfdrv.unconf file for hdisks Dynamic Link Manager software should not manage (i.e. boot disks; see SAN Boot with Hitachi HiCommand Dynamic Link Manager Software Using a Single Boot Device above).

   6.3. Run dlmcfgmgr
   Note: hdisks must be discovered prior to running the dlmcfgmgr command

   6.4. Review Hitachi HiCommand Dynamic Link Manager software system parameters (auto failback, intermittent error monitor, and the reservation level)

      a. *# dlnkmgr view –sys*

      b. Ensure the reserve is disabled when using more than one VIO server

- *# dlnkmgr set –rsv on 0 –s*

6.5. Assign PVIDs for the virtual devices

a. *# chdev –l dlmfdrv(x) –a pv=yes*
Note: The dlmfdrv devices must be assigned PVIDs to be used as virtual devices unless using the unique_id attribute. If using the unique_id attribute, assigning a PVID is still recommended.

6.6. Exit to return to the VIO host partition and reactivate if necessary.

6.7. Create the virtual devices

- *$ mkvdev –vdev hdisk1 –vadapter vhost1 - for internal boot device*
- *$ mkvdev –vdev dlmfdrv6 –vadapter vhost0*
- *$ mkvdev –vdev dlmfdrv7 –vadapter vhost0*

6.8. List the virtual hosts:

- *$ lsdev –virtual* (See Figure 13)
- *$ lsmap –all* (See Figure 14)

Figure 13: List Virtual Devices

```
$ lspv | grep dlmfdrv
dlmfdrv0        none                                  None
dlmfdrv1        none                                  None
dlmfdrv2        00c9bf2edc14ab25                      None
dlmfdrv3        00c9bf2e0c90f881                      None
dlmfdrv4        00c9bf2e3809ee33                      havg
dlmfdrv5        00c9bf2e3809f35b                      havg
dlmfdrv6        00c9bf2ead8571cc                      None
dlmfdrv7        00c9bf2e3809fb3e                      None
dlmfdrvio       none                                  None
dlmfdrv8        none                                  None
$ lsdev -virtual
name            status     description

ent3            Available  Virtual I/O Ethernet Adapter (l-lan)
ent4            Available  Virtual I/O Ethernet Adapter (l-lan)
vhost0          Available  Virtual SCSI Server Adapter
vhost1          Available  Virtual SCSI Server Adapter
vsa0            Available  LPAR Virtual Serial Adapter
vtscsi0         Available  Virtual Target Device - Disk
vtscsi1         Available  Virtual Target Device - Disk
vtscsi2         Available  Virtual Target Device - Disk
ent2            Defined    Shared Ethernet Adapter
$
$ _
MA*   a
```

The 'lsdev –virtual' command in Figure 13 shows that dlmfdrv4 and 5 have a volume group and therefore cannot be used as virtual devices.

Note: In Figure 13 above, dlmfdrv4 and 5 have a volume group and cannot be used as virtual devices. The “mkvdev” command will fail if dlmfdrv4/5 is used. Dlmfdrv6 and 7 are created as ‘vtscsi0’ and ‘vtscsi1’ on vhost0. Vhost1 is the virtual adapter used for hdisk0, which is the internal boot device for the client partition.

This is shown in Figure 14 below. The client partition id, "8", is shown because the client partition is active when this was captured. Normally, in an implementation, this would be zeros until the client is activated.

Figure 14: Hdisk0 Virtual Adapter Vhost1 Is the Client Partition Internal Boot Device

```
$ lsmap -all
SVSA            Physloc                                      Client Partition ID
--------------- -------------------------------------------- ------------------
vhost0          U9111.520.109BF2E-V2-C3                      0x00000008

VTD                   vtscsi0
LUN                   0x8100000000000000
Backing device        dlmfdrv6
Physloc               U787A.001.DNZ0C5Z-P1-C1-T1-L9

VTD                   vtscsi1
LUN                   0x8200000000000000
Backing device        dlmfdrv7
Physloc               U787A.001.DNZ0C5Z-P1-C1-T1-L10

SVSA            Physloc                                      Client Partition ID
--------------- -------------------------------------------- ------------------
vhost1          U9111.520.109BF2E-V2-C5                      0x00000008

VTD                   vtscsi2
LUN                   0x8100000000000000
Backing device        hdisk0
Physloc               U787A.001.DNZ0C5Z-P1-T10-L5-L0

$ _
MA*   a                                                               p1  25/003
```

Figure 14 provides an example where Vhost1 is the virtual adapter used for hdisk0 which is the internal boot device for the client partition.

Configuration on VIO Server 1 is now complete. Activate VIO Server 2 and make the same changes (i.e. install the ODM update, change the adapter attributes, install Hitachi HiCommand Dynamic Link Manager software, create the virtual devices).

After configuration on VIO Server 2 is complete, run the "*lsmap –all*" and "*lsdev –virtual*" commands to review the configuration.

Figure 15: Review Configuration with 'lsdev –virtual' and 'lsmap –all' Commands

```
Last login: Wed Jul 13 14:05:35 CDT 2005 on /dev/vty0 from viomike

$ lsdev -virtual
name            status     description

vhost1          Available  Virtual SCSI Server Adapter
vsa0            Available  LPAR Virtual Serial Adapter
vtscsi0         Available  Virtual Target Device - Disk
vtscsi1         Available  Virtual Target Device - Disk
$ lsmap -all
SVSA            Physloc                                      Client Partition ID
--------------- -------------------------------------------- ------------------
vhost1          U9111.520.109BF2E-V4-C4                      0x00000008

VTD                   vtscsi0
LUN                   0x8100000000000000
Backing device        dlmfdrv6
Physloc               U787A.001.DNZ0C5Z-P1-C2-T1-L8

VTD                   vtscsi1
LUN                   0x8200000000000000
Backing device        dlmfdrv7
Physloc               U787A.001.DNZ0C5Z-P1-C2-T1-L9

$ _
MA*   a                                                                p1  25/003
```

After configuring VIO Server 2, review the configuration with the "lsdev –virtual" and "lsmap –all" commands.

Run "*lspv*" to display the PVIDs. The PVID is the identifier that MPIO on the client uses to configure the devices when using Hitachi HiCommand Dynamic Link Manager without the unique_id attribute. Hitachi HiCommand Dynamic Link Manager and MPIO using the unique_id attribute will display the same PVID on the client and on the server. Also, note the location codes for the virtual devices. In Figure 16 below, Vtscsi1 is at location U9111.520.109BF2E-V4-C4-L2. C4-L2 is the virtual adapter address and LUN number.

Figure 16: Display the PVIDs and Location Codes for Virtual Devices

```
$ lspv
dlmfdrv0        none                                 None
dlmfdrv1        none                                 None
dlmfdrv2        00c9bf2edc14ab25                     None
dlmfdrv3        00c9bf2e0c90f881                     None
dlmfdrv4        00c9bf2e3809ee33                     havg
dlmfdrv5        00c9bf2e3809f35b                     havg
dlmfdrv6        00c9bf2ead8571cc                     None
dlmfdrv7        00c9bf2e3809fb3e                     None
dlmfdrvio       none                                 None
$ lsdev -virtual
name            status      description

vhost1          Available   Virtual SCSI Server Adapter
vsa0            Available   LPAR Virtual Serial Adapter
vtscsi0         Available   Virtual Target Device - Disk
vtscsi1         Available   Virtual Target Device - Disk
$ lsdev -vpd | grep vhost
  vhost1          U9111.520.109BF2E-V4-C4
Virtual SCSI Server Adapter
$ lsdev -vpd | grep vtscsi
  vtscsi1         U9111.520.109BF2E-V4-C4-L2
Virtual Target Device - Disk
  vtscsi0         U9111.520.109BF2E-V4-C4-L1
Virtual Target Device - Disk
MA*   a
```

Running "lspv" displays the PVIDs, "lsdev –virtual displays location codes for virtual devices.

While it is not necessary, it is recommended that the dlmfdrv device numbers match on both VIO servers. It is also recommended that the LUN numbers are consistent. This makes management and problem determination much easier.

Configuration on both VIO servers is now complete. Activate the client to display the devices. This assumes the client configuration in the Hardware Management Console is complete. Login to the VIO client and run "*lspv*", "*lspath*", "*lscfg –vl hdisk(x)*", and "*lsattr –El hdisk(x)*" as shown below in Figure 17.

Figure 17: Activate Client to Begin Connecting Virtual Client Devices to Virtual Server Devices

```
********************************************************************************
Last unsuccessful login: Thu May  5 14:36:40 MDT 2005 on /dev/vty0
Last login: Fri Jul 15 15:19:10 MDT 2005 on /dev/vty0

# lspv
hdisk30         00c9bf2e178949c4                    rootvg          active
hdisk0          none                                None
hdisk1          none                                None
# lspath
Enabled hdisk0  vscsi3
Enabled hdisk1  vscsi3
Enabled hdisk0  vscsi0
Enabled hdisk1  vscsi0
Enabled hdisk30 vscsi1
# lsattr -El hdisk0
PCM            PCM/friend/vscsi Path Control Module        False
algorithm      fail_over        Algorithm                  False
max_transfer   0x20000          Maximum TRANSFER Size      True
pvid           none             Physical volume identifier False
queue_depth    3                Queue DEPTH                False
reserve_policy no_reserve       Reserve Policy             False
# lscfg -vl hdisk0
  hdisk0           U9111.520.109BF2E-V8-C3-T1-L810000000000  Virtual SCSI Disk D
rive
# _
MA*   a                                                               p1  25/003
```

Using the "lscfg –vl hdisk(x)" command in the client is the first of two steps to connect the virtual client devices to the virtual server devices.

How do you connect virtual client devices to virtual server devices? Use the "*lscfg –vl hdisk(x)*" command in the client (see Figure 17 above).

 # lscfg –vl hdisk0
hdisk0 U9111.520.109BF2E-V8-C3-T1-L810000000000 Virtual SCSI Disk Drive

And use the "*lsdev –attr –dev vtscsi(x)*" command in the Server (see Figure 18 below):

Figure 18: Complete Connection of Virtual Client Devices to Virtual Server Devices from Virtual Server

HA_VIO1

```
$ lsdev -attr -dev vtscsi0
attribute        value              description          user_settable

LogicalUnitAddr  0x8100000000000000 Logical Unit Address False
aix_tdev         dlmfdrv6           Target Device Name   False
$ lsdev -attr -dev vtscsi1
attribute        value              description          user_settable

LogicalUnitAddr  0x8200000000000000 Logical Unit Address False
aix_tdev         dlmfdrv7           Target Device Name   False
$ _
MA*    a                                                          p1  25/003
```

HA_VIO2

```
$ lsdev -attr -dev vtscsi0
attribute        value              description          user_settable

LogicalUnitAddr  0x8100000000000000 Logical Unit Address False
aix_tdev         dlmfdrv6           Target Device Name   False
$ set -o vi
$ lsdev -attr -dev vtscsi1
attribute        value              description          user_settable

LogicalUnitAddr  0x8200000000000000 Logical Unit Address False
aix_tdev         dlmfdrv7           Target Device Name   False
$ _
MA*    a                                                          p1  25/003
```

Using the "lsdev –attr –dev vtscsi(x)" command in the Server, as shown in Figure 18 above, is the second of two steps to connect the virtual client devices to the virtual server devices causing the LogicalUnitAddr in the server matches the location code in the client.

Note the LogicalUnitAddr in the server matches the location code in the client.

The virtual disks are defined in the client. Create a volume group on the virtual disks in the client partition using standard LVM commands

 # mkvg –y'vioc_vg1' –s '32' hdisk0 hdisk1

Figure 19: Create a Volume Group

```
# lspv
hdisk30         00c9bf2e178949c4                    rootvg          active
hdisk0          00c9bf2e4ee7a89d                    vioc_vg1        active
hdisk1          00c9bf2e4ee7a96f                    vioc_vg1        active
# lscfg -vl hdisk0
  hdisk0           U9111.520.109BF2E-V8-C7-T1-L810000000000  Virtual SCSI Disk
rive
# lscfg -vl hdisk1
  hdisk1           U9111.520.109BF2E-V8-C7-T1-L820000000000  Virtual SCSI Disk
rive
# _
```

Create a volume group on the virtual disks in the client partition using the standard LVM commands such as "mkvg –y'vioc_vg1' –s '32' hdisk0 hdisk".

New PVIDs are assigned if not using the unique_id attribute (dlmodmset –u on). This is because with Hitachi Dynamic Link Manager, the PVID method is used to identify the disks instead of a unique id. Please see the **VIO restrictions** concerning unique ids and PVIDs.

Install MPIO in the VIOS

The following steps are detailed installation instructions for installing MPIO in the VIO server partition(s). These steps can be done before or after creating the virtual adapters. However, the virtual adapters must be created prior to making the virtual devices.

1. Login using the padmin command
2. Enter the following command to install the ODM update

   $ oem_setup_env
3. Remove any existing devices that will be used as virtual devices:

   # rmdev –dl hdisk(x)
4. Change the Fibre Channel adapters to fast fail:

   # chdev –l fscsix –a fc_err_recov=fast_fail (-P)

   Use the "-P" option if the adapters are in use and the hdisks can not be removed. A reboot will change the adapter attribute.
5. Install the MPIO ODM updates (5.4.0.0 and 5.4.0.1)
6. *# cfgmgr –vl fcs(x)*

   Run once for each Fibre Channel adapter. Use the "*lsdev –Cc disk*" command to ensure the disks are defined as "Hitachi".
7. Ensure the reserve is disabled when using more than 1 VIO server

   # chdev –l hdiskx –a reserve_policy=no_reserve

8. Change the load balancing algorithm if the reserve is disabled:

 # chdev –l hdiskx –a algorithm=round_robin

9. Assign PVIDs for the virtual devices:

 # chdev –l hdisk(x) –a pv=yes

The devices should be assigned PVIDs to be used as virtual devices.

10. Exit to return to the VIO host partition and reactivate if necessary.

11. Create the virtual devices:

$ mkvdev –vdev hdisk1 –vadapter vhost1 - for internal boot device

$ mkvdev –vdev hdisk2 –vadapter vhost0 –dev client_disk1

$ mkvdev –vdev hdisk3 –vadapter vhost0 –dev client_disk2

12. List the virtual hosts:

$ lsdev –virtual

$ lsmap –all

13. In the VIO client partition, create a volume group:

 # mkvg –y'mpiovg' –s '32' hdisk0 hdisk1

lspv

hdisk30	00c9bf2e178949c4	rootvg	active
hdisk0	00c9bf2e83a7e4b9	mpiovg	active
hdisk1	00c9bf2e05198ed4	mpiovg	active

From our previous example of the VIO Servers, there was not a PVID on the second disk that was exported.

$ lspv

hdisk0	none	None
hdisk1	none	None
hdisk2	00c9bf2e83a7e4b9	None
hdisk3	**none**	**None**
hdisk4	00c9bf2e0c90f881	None
hdisk5	00c9bf2e3809ee33	havg
hdisk6	00c9bf2e3809f35b	havg
hdisk7	00c9bf2ead8571cc	None
hdisk8	00c9bf2e40c68507	None
hdisk9	00c9bf2e4ee360de	None
hdisk10	00c9bf2e4ee36fb4	None
hdisk11	none	None
hdisk12	none	None
hdisk13	00c9bf2e552c4f42	rootvg active

$ chdev -dev hdisk3 -attr pv=yes

hdisk3 changed

$ lspv

hdisk0	none	None
hdisk1	none	None
hdisk2	00c9bf2e83a7e4b9	None
hdisk3	**00c9bf2e05198ed4**	**None**
hdisk4	00c9bf2e0c90f881	None
hdisk5	00c9bf2e3809ee33	havg
hdisk6	00c9bf2e3809f35b	havg
hdisk7	00c9bf2ead8571cc	None
hdisk8	00c9bf2e40c68507	None
hdisk9	00c9bf2e4ee360de	None
hdisk10	00c9bf2e4ee36fb4	None
hdisk11	none	None
hdisk12	none	None
hdisk13	00c9bf2e552c4f42	rootvg active

By running the ***chdev*** command, the PVID is now also on the VIO Server. It is recommended to assign PVIDs

Additional Software in VIOS

The VIOS is a critical resource and is not intended to run applications. The VIOS is based on the AIX operating system. Logging in with the padmin userid uses a restricted shell and a generic command line interface is provided for management.. The VIOS command, oem_setup_env, will put the user into an AIX-like environment. This environment is only supported for the installing and setup of certain software packages like Hitachi HiCommand Dynamic Link Manager software for path management, and Hitachi Storage Services Manager, Hitachi Tuning Manager and Hitachi Device Manager software agents. Any other tasks performed in this environment are not supported. The VIOS will continue to support the current user interfaces, however, the underlying operating system may change at any time.

Hitachi HiCommand Dynamic Link Manager software and the Hitachi Storage Services Manager, Hitachi Tuning Manager and Hitachi Device Manager software agents are installed as in a standalone AIX server. The Hitachi Storage Services Manager agent in the VIOS will report the physical resources, the HBAs and disks. The Hitachi Storage Services Manager agent in the VIOC will report the virtual SCSI disks and file systems. There is no correlation today between the physical and virtual environment in Hitachi Storage Services Manager software.

Virtual I/O in HACMP Cluster Environments

The Service Implementation Plan between Hitachi Data Systems and IBM applies to configurations involving HACMP in VIO environments.

Requirements:

- AIX 5L v5.3 ML 3, or above
- MPIO in the VIO Server and Client, or Hitachi HiCommand Dynamic Link Manager in the VIO Server and MPIO in the client
- VIO Server 1.2 or above
- HACMP

HACMP v5.2 with APAR IY68370 (or higher) and APAR IY68387

HACMP 5.3—An ISR (Hitachi Support Request) is required.

- RSCT (Reliable Scalable Cluster Technology)

  The following RSCT filesets are required:

  – rsct.basic.hacmp 2.4.2.1

  – rsct.basic.rte.2.4.2.2

  – rsct.compat.basic.hacmp.2.4.2.0

Additional HACMP and Virtual SCSI (vSCSI) Guidance

The following is additional guidance for HACMP/VIO configurations.

- Reservation policies must be disabled for the shared disks on all VIO servers that are managing the disks.
- The volume group must be defined as “Enhanced Concurrent”.
- All nodes must access shared volumes through virtual SCSI. Disks cannot be shared between an LPAR using virtual SCSI and a node directly accessing those disks.
- From the VIO server, physical disks (hdisks) are shared, not logical volumes or volume groups.
- All volume group construction and maintenance on shared disks is done from the HACMP nodes, not from the VIO server (CSPOC is highly recommended).

Appendix A References

VIO Sizing:

http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp?topic=/iphb1/iphb1_vios_planning_vscsi_sizing.htm

IBM Virtualization Redbooks:

http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf

http://www.redbooks.ibm.com/redbooks/pdfs/sg245768.pdf

VIO FAQs:

http://techsupport.services.ibm.com/server/vios/documentation/faq.html

VIO Support Page:

http://techsupport.services.ibm.com/server/vios

Appendix B Glossary

Hardware Management Console (HMC) is the IBM term for the workstation and software that manages partitioned System p5 severs.

High Availability Cluster Multi-Processing (HACMP) is an IBM trademarked product designed for AIX platforms which can run on multiple computers or nodes in the cluster. At any one point in time, a given node is either active (actively running an application) or passive (waiting to take over if another node fails). File system data can be shared between the nodes in the cluster.

Hitachi Dynamic Link Manager™ Path Manager Software is server-based software which enhances RAID systems by providing automatic failover and load balancing from server-to-RAID channel-connection failures. It provides increased reliability and performance.

Logical Partition (LPAR) refers to a virtualized collection of computing resources that are securely separated from each other and abstracted from the actual physical resource for purposes of resource sharing and increased resource utilization.

Micro-Partitioning, the trademarked IBM feature, is enabled by the POWER Hypervisor which is included in IBM Advanced POWER Virtualization feature technology. Micro-Partitioning enables users to allocate less than a full physical processor to a logical partition and to define logical partitions as small as 1/10th of a processor, and change logical partitions in increments as small as 1/100th of a processor.

Multi-path I/O (MPIO) refers to software that allows the configuring of one or more physical paths between a host and its storage devices to provide fault tolerance. IBM AIX 5L V5.2 and later supports an MPIO driver which is functionally equivalent to the Hitachi Dynamic Link Manager. Multi-Path I/O software routines establish separate controllers, ports, switches, etc. so that should one controller, port or switch fail, the operating system can route I/O through the remaining controller.

POWER Hypervisor is Power5 hardware that supports partitioning and dynamic resource movement across multiple operating system environments.

Virtual I/O (VIO) refers to the sharing of single physical I/O adapters and disks to support multiple logical partitions of the same server, enabling the on demand allocation and management of I/O devices. Virtual I/O allows consolidation of storage resources and minimizes the number of I/O adapters and disks required to service multiple workloads.

Virtual I/O Client (VIOC) is any LPAR on the host that is using virtualized devices provided by other partitions.

Virtual I/O Server (VIOS) is any virtual server or I/O partition that provides virtual storage and shared Ethernet adapter capability to client logical partitions on POWER5 systems. It allows a physical adapter with attached disks on the Virtual I/O Server partition to be shared by one or more partitions, enabling clients to consolidate and potentially minimize the number of physical adapters required.

Virtual LAN provides network virtualization capabilities that enable you to prioritize traffic on shared networks.

Virtualization is the software and/or hardware abstraction of system resources—processors, memory and I/O devices—into a common pool for simplified allocation, access and management.

Hitachi Data Systems

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
www.hds.com
info@hds.com

Asia Pacific and Americas
750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
Phone: 1 408 970 1000
info@hds.com

Europe Headquarters
Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
Phone: + 44 (0)1753 618000
info.eu@hds.com

TM-001-00 November 2006