

The VIO cheat sheet

Consolidate your AIX systems

Skill Level: Intermediate

[Christian Pruettt \(pruettc@gmail.com\)](mailto:pruettc@gmail.com)
Senior Systems Administrator
Freelance

20 Jul 2010

Do you want to consolidate multiple AIX® systems under a single System p® server in ways that maximize your investment return on disk and network resources? Have you been confused about technical terms like SEAs, IVE, or NPIV? This article introduces you to the latest AIX technology, virtual I/O, and provides a cheat sheet on how to set up a simple yet robust server and client environment.

In the early days of IBM AIX®, systems administrators were limited to one logical server per physical server. If you wanted to grow your computing environment, you had to purchase a new IBM RS/6000® or pSeries® server. Or, conversely, if you had some free resources you wanted to share to another server, there was no easy way of doing it short of physically moving hardware components.

In 2001, IBM introduced logical partitioning (LPAR) technology, which allowed multiple logical servers to use the same physical server's resources, including processor, memory, disk, and adapters, managed by a special administrative server called the *Hardware Management Console* (HMC). This technology let systems administrators create, modify and remove LPARs, manage resources, and do operating system work on disparate logical AIX and Linux® servers within a single physical server.

With several more iterations of LPAR technology, it became possible to dynamically manage resources and micropartition processor resources, letting multiple LPARs share even the same physical processor simultaneously. Then, IBM introduced virtual I/O (VIO) technology, which allows the same LPARs to use the same storage and network resources, thereby breaking the barrier of disk and adapter isolation.

The basics of VIO

VIO technologies consist of servers, software, and various commands.

VIO servers

VIO technology is based on the idea of having special LPARs that manage the disk and network resources that other LPARs use on the same pSeries or IBM System p servers. Instead of the individual network and disk resources being carved out on an LPAR-by-LPAR basis (especially in cases where there wouldn't be sufficient resources for all the LPARs to possess what they require to function), one or two VIO servers are given control of these resources and share them out to the other LPARs (VIO client LPARs).

VIO software

This VIO server runs a special version of the AIX operating system, with an additional software package called *IOS*. This software comes bundled and is managed independent of the usual operating system commands (for example, `installp`) and versioning structure (technology levels). It is similar to how HMCs have specialized software loaded onto a Linux kernel for a specific purpose.

Note: Installing third-party software or altering the operating system through any means outside of the IOS typically invalidates support from IBM, so it's best to avoid modifying the server in any nonstandard, non-VIO way.

The VIO user ID and commands

Instead of using the root user ID to issue commands, an administrative user ID—`padmin`—is used for all of the VIO controls. This user ID has a shell called the *IOS command-line interface* (`ioscli`) that runs a unique set of commands for managing devices assigned to the VIO servers. Many of these commands are similar to regular AIX commands (for example, `lsdev`) but use different flags and command structures (for example, `lsdev -dev`). But, most of the superuser level commands are new and perform many different operations at once. In fact, when administered properly, systems administrators will rarely have to become root.

Helpful `padmin` commands include:

- **help:** This command lists all of the commands available in the `ioscli`. If you pass a specific command into it, such as `help updateios`, you can see the specific flags and syntax for that command.
- **cfgdev:** This command is the equivalent of the `cfgmgr` command and detects new physical and logical devices added to the VIO server.

- **oem_setup_env**: This command is the equivalent of running `su - root` but without the need to enter a password. Again, you will rarely have to become root on a VIO server.
- **mkvdev**: This command manages the virtual devices that you create and serve up to the VIO client LPARs.
- **mktcpip, lstcpip, and rmtcpip**: These commands manage your networking from the command line and circumvent the need for utilities such as `smitty tcpip`.
- **lsmap**: This command shows the relationships between disk resources and VIO client LPARs.

Planning your environment

As you begin to plan your VIO environment, a successful, fully functioning and highly available VIO environment is directly proportional to the amount of time you invest in checking your hardware, designing some handy spreadsheets, and focusing on the details.

Disk resources

The first area you have to address is how to manage your disk resources to your VIO client LPARs. Within VIO, you use three main methods to serve up disk resources:

- **Method 1: Logical volumes as disks.** This method requires assigning a disk to a VIO server and using the `padmin` user ID to create a volume group and logical volumes (LVs) on that disk. Then, you map those LVs to VIO client LPARs for use, seeing them as individual disks. Typically, each LV is several gigabytes in size to accommodate the need of the VIO client LPARs, such as having the AIX operating system for that LPAR placed on it.
 - *Benefits*: This method can reduce the amount of disks assigned to a VIO server, because multiple VIO client LPARs can access the individual LVs they are assigned on the same physical disk or Logical Unit Number (LUN) from a Storage Area Network (SAN). If there is sufficient free space on the disk, you can create an additional LV on the fly and assign it quickly.
 - *Drawbacks*: This method can cause some resource contention if the VIO client LPARs perform rapid input/output (I/O) on the same physical disk. In some cases, for volume group redundancy on the VIO client LPARs, it can require twice as many physical disks being

assigned to two different VIO servers. And, an errant `rmlv` command can completely knock a VIO client LPAR off of the wire.

- **Method 2: Virtual SCSI disks.** In this method, disks are assigned to the VIO servers and mapped directly to the VIO client LPARs. The VIO servers have no visibility into what is on the disks or how they are being used but simply serve the disks out to the VIO client LPARs.
 - *Benefits:* This method is a quick and easy way of getting disks out to VIO client LPARs; it takes only two short commands to get a disk detected and out the door to a server for use. Plus, the VIO clients do not have to worry about any sort of disk management software (such as SDDPCM) to have redundancy in seeing the disks down two paths when served by two VIO servers.
 - *Drawbacks:* Managing a massive quantity of disks on VIO servers and the clients to which they are mapped can become tricky. Also, if you ever have to boot into System Management Services (SMS), it can take a long time for the VIO server to probe all the disks, and it may identify several as being root volume groups (the trick is to look for the VIO name).
- **Method 3: Virtual Fibre Channel Adapter (NPIV).** In this method, VIO servers become complete pass-throughs in sharing out their Fibre Channel (FC) adapters directly to the VIO client LPARs. Using a new technology called *N-Port ID Virtualization*, a single FC adapter hooked up to a SAN can be used by multiple VIO client LPARs simultaneously. The VIO servers never see any of the disks that are assigned to the VIO client LPARs, because each VIO client LPAR is given its own Worldwide Number (WWN) on each FC adapter, and the LUNs from the SAN are directly mapped to those WWNs.
 - *Benefits:* This method is an extremely elegant way to manage VIO disk resources and simplifies the process of mapping disks. It minimizes the amount of VIO involvement, because after the initial mapping of a FC adapter to a VIO client LPAR is complete, you need run no other commands on the VIO servers—unlike the virtual SCSI disk method, where commands have to be run on each VIO server for each and every disk that is shared out.
 - *Drawbacks:* The main drawback with this method is that some SAN technology is not yet compatible with NPIV technology. For example, I had one tedious experience where I had to manually enter all of the WWNs from my VIO clients into the zone maps, because the SAN could not automatically detect them. And, if you're not careful with your licensing, you can exhaust the range of WWNs that the virtualization technology allocates to the VIO servers.

Network resources

The second area you have to plan out is how to share out your network resources to the VIO client LPARs. Similar to the disk resources, there are two main ways of setting things up:

- **Method 1: Shared Ethernet adapters (SEA).** The main principle behind SEA technology is simple:

```
1 Physical Ethernet Adapter + 1 Virtual  
Ethernet Adapter = 1 Shared Ethernet Adapter
```

When VIO servers are created, they are assigned both physical Ethernet adapters and virtual Ethernet adapters. The VIO client LPARs are told which virtual Ethernet adapters they should use for their communication. The VIO servers then map these virtual adapters to physical Ethernet adapters, and those VIO client LPARs can communicate through the same device.

- *Benefits:* As long as you have a physical entX device available, you can make a new connection for your VIO client LPARs. And, even the VIO servers can have IP addresses configured onto the SEAs for communication, bypassing the need for any sort of specialized administrative network connection.
- *Drawbacks:* Resource contention can occur if you have too many VIO client LPARs going through the same physical Ethernet adapter. If virtual LAN (VLAN) trunking is available, where multiple network subnets can be accessed simultaneously through the same physical adapter, this method does not use that benefit.
- **Method 2: Integrated virtual Ethernet (IVE).** IVE technology is similar to SEA technology but allows access to multiple VLANs through the same physical adapter. Each VLAN is defined both through the HMC and on the VIO server for communication. Then, the VIO client LPARs are told the virtual Ethernet adapters and VLAN numbers they should access through an SEA mapping. The communication to multiple subnets occurs seamlessly.
 - *Benefits:* IVE cuts down on the number of physical Ethernet adapters and connections needed to facilitate communications. It becomes possible to send traffic to production, development, and backup networks all through the same piece of wire.
 - *Drawbacks:* At this time, you cannot spontaneously add new VLANs to an IVE connection. If you need to add a new VLAN to an existing IVE connection, you must first logically destroy and re-create the

underlying SEA device, possibly impeding any VIO client LPARs using that connection. Furthermore, as with older SAN technology and NPIV, older networking equipment may not be able to handle IVE connections.

Redundant VIO servers

The third area you need to plan is to have redundant VIO servers on the same physical pSeries or System P server. If a single VIO server supports a dozen VIO client LPARs and something catastrophic knocks that server offline, everything on top of it will come crashing down.

By having two VIO servers with the same set of resources, VIO client LPARs can continue functioning unimpaired if something takes one of the VIO servers down. The VIO client LPARs will go to the other VIO server for their disk and network resources. Mapping disks to both VIO servers and creating network control channels gives VIO client LPARs two legs to stand on. It also makes it possible to perform IOS upgrades on the VIO servers dynamically without affecting the VIO client LPARs.

The resources that you assign to each VIO server should be as close to identical as possible and designed to maximize availability. Do not mix a slower speed Ethernet adapter on one VIO server with a faster speed one on another. Do not put all of the FC adapters used by both VIO servers in the same physical drawer. Instead, stagger the adapters between multiple drawers and assign them independently. Plan out every possible hardware failure and look for ways to maximize redundancy.

In addition, it is especially important to document how everything is mapped out. Record your environment in a spreadsheet and cross-reference it often with the output of commands like `vfcmmap`. [Figure 1](#) provides an example of a simple sheet that details a System p server with two VIO servers and four VIO client LPARs using a mix of SEA, IVE, virtual SCSI, and virtual FC.

Figure 1. Sample variables spreadsheet

VIO Server	Physical Adapter	Virtual Adapter	SEA	Default VID	Control Channel VID	Trunk Priority	VLANs
vio1	ent0	ent2	ent4	12	13	1	N/A
	ent1	ent3	ent5	14	15	2	100,200
vio2	ent0	ent2	ent4	12	13	2	N/A
	ent1	ent3	ent5	14	15	1	100,200

VIO Server	VIO Client LPAR	Volume Group	HDISK	VHOST	Server Slot	Client Slot
vio1	N/A	rootvg	hdisk0	N/A	N/A	N/A
	client01	rootvg	hdisk1	vhost0	100	100
	client01	datavg	hdisk2	vhost0	100	100
	client02	rootvg	hdisk3	vhost1	200	200
vio2	N/A	rootvg	hdisk0	N/A	N/A	N/A
	client01	rootvg	hdisk1	vhost0	100	101
	client01	datavg	hdisk2	vhost0	100	101
	client02	rootvg	hdisk3	vhost1	200	201

VIO Server	VIO Client LPAR	Fiber Channel Adapter	VFCHOST	Client Fiber Channel
vio1	client03	fcs0	vfchost0	fcs0
	client04	fcs1	vfchost1	fcs0
	client04	fcs2	vfchost2	fcs1
vio2	client03	fcs0	vfchost0	fcs0
	client04	fcs1	vfchost1	fcs0
	client04	fcs2	vfchost2	fcs1

Building the VIO server

Now that you've determined what you need for your environment, the following procedure guides you through building a VIO server. This procedure assumes that you are familiar with the HMC and SMS along with their menu systems.

1. Confirm that advanced power virtualization is available:
 - a. In the HMC, select your managed system.
 - b. Click **Properties**.
 - c. On the **Capabilities** tab, confirm that **Virtual I/O Server Capable** is set to **Available**. If it is not available, contact IBM for an Advanced Power Virtualization code, and install it to make VIO available.

2. Define the VIO LPAR:
 - a. In the HMC, with your managed system selected, click

Configuration > Create Logical Partition > VIO Server.

- b. Name your server, and call this profile `$SERVER.novirtuals`.
 - c. Give it the amount of processors, memory, and I/O resources you desire, but do not create any virtual adapters at this time.
 - d. If you intend to build your VIO server from CD or DVD, assign the drive as needed.
3. Install IOS:
- a. Select the VIO server, and click **Operations > Activate**.
 - b. Click **Advanced** and choose **SMS** for the Boot mode.
 - c. Select the check box to open a terminal screen.
 - d. If you are installing from a CD or DVD, insert the disc and have the server boot from it within SMS.
 - e. If you are using Network Installation Manager (NIM), configure your network adapter settings, and point to your NIM server. Let the server install the IOS on your hard disk.
4. Set up the password, licensing, patching, and mirroring:
- a. When the VIO server is up, log in with the `padmin` user ID and set its password.
 - b. If prompted, run the `license -accept` command to confirm the software licensing.
 - c. If you have an update for the server, use the `updateios` command to install any patches.
 - d. Mirror the root volume group with the `mirrorios` command, if applicable.
 - e. Reboot the VIO server with the `shutdown -restart` command.
5. Clone the server:
- a. Back up the server with the `backupios` command and use that image to build your redundant VIO server (I prefer the ease of NIM for this task).
6. Create the virtual-enabled profile:

- a. In the HMC, make copies of the current VIO servers' profiles, and call them `$SERVER.vio`. These profiles will contain your VIO servers' configurations with virtual devices.
7. Define your virtual Ethernet devices (HMC):
- a. In the HMC, open the virtual-enabled profiles using the **Edit** menu.
 - b. Click the **Virtual Adapters** tab, and change the **Maximum Virtual Adapters** number to something high, like 1000 (So you don't get errors for exceeding the default of 20).
 - c. Click **Actions > Create > Ethernet Adapter**.
 - d. Set the Adapter ID, and enter VLANs if you are using IVE.
 - e. For the main virtual adapter, select the **Access External Network** check box.
 - f. Set different trunk priority numbers between the two VIO servers.
 - g. Repeat the same process for a control channel adapter for redundancy, but do not select the **Access External Network** check box.
 - h. Save your changes, then boot from this profile.
8. Define your virtual Ethernet devices (VIO):
- a. Log in to the VIO servers as `padmin`.
 - b. Check your device list with the `lsdev` command.
 - c. Check the attributes of the virtual Ethernet adapters with the `lsdev -dev entX -attr` command to confirm which adapters are which.
 - d. Run the following command to create an SEA, substituting your `entX` devices and ID number from your spreadsheet:

```
mkvdev -sea $PHYS -vadapter $VIRT -default
$VIRT -defaultid $ID
      -attr ha_mode=auto ctl_chan=$CTRL
```
 - e. If you need to make this SEA available from the VIO Server, use the `mktcpip` command to set an IP address on it. A `ping` test will

quickly confirm whether you have set up everything correctly.

Building the VIO clients

Now that your VIO servers are up, the following procedure guides you through building a VIO client LPAR:

1. Define the VIO client LPAR:
 - a. In the HMC, with your managed system selected, click **Configuration > Create Logical Partition > AIX Server**.
 - b. Name your server, and call this profile `$SERVER.vio` for ease.
 - c. Give it the amount of processors, memory, and I/O resources you desire but do not create any virtual adapters at this time.
2. Create the VIO server disk resources:
 - a. In the HMC, open the VIO servers' virtual-enabled profiles using the **Edit** menu.
 - b. Click the **Virtual Adapters** tab.
 - c. Click **Actions > Create > Fiber Channel Adapter or SCSI Adapter**.
 - d. Enter the slot numbers from your spreadsheet.
 - e. Select the **Only selected client partition can connect** option, and pick your VIO client LPAR.
 - f. Shut down your VIO servers, and activate them from these profiles, or dynamically add the same resources to the LPARs.
Note: You created the VIO client LPAR as a blank slate so that you can define this easily.
3. Edit the VIO client LPAR:
 - a. In the HMC, open the VIO client LPARs' virtual-enabled profiles using the **Edit** menu.
 - b. Click the **Virtual Adapters** tab.

- c. Click **Actions > Create > Fiber Channel Adapter or SCSI Adapter**.
 - d. Enter the slot numbers from your spreadsheet.
 - e. Click **Actions > Create > Ethernet Adapter**, set the Adapter ID, and enter VLANs as needed from your spreadsheet. If you created virtual Fibre Channel adapters, click their properties to obtain their WWNs.
4. Define the virtual SCSI disk maps (VIO):
- a. If you are using virtual SCSI adapters to serve disk resources, map those disks at this time from your SAN (if applicable).
 - b. Log in to the VIO servers with the `padmin` user ID and run `cfgdev` to detect any new disks.
 - c. Examine them with the `lspv` and `lsdev -dev hdiskX -attr` commands.
 - d. Examine the vhosts on the server with the `lsmmap -all` command.
 - e. Run the following command to map the disks to the specified vhosts, giving them virtual target disk (VTD) names to help you track them as you desire:

```
mkvdev -vdev hdiskX -vadapter $VHOST -dev $VTD
```

5. Define the virtual FC maps (VIO):
- a. If you are using virtual FC adapters to serve disk resources, examine the vfchosts on the server with the `lsmmap -all -npiv` command.
 - b. Run the following command to map the FC adapters to the specified vfchosts:

```
vfcmmap -vadapter vfchostX -fcp fcsX
```

- c. Enter in your WWNs into your SAN and carve out and map disks. They will go to the VIO client LPARs.

6. Activate the client LPARs (HMC):
 - a. Select the VIO client LPARs, and click **Operations > Activate**.
 - b. Click **Advanced**, and choose **SMS** for the Boot mode.
 - c. Select the check box to open a terminal screen.

From this point, the installation follows a standard AIX server installation.

Conclusion

VIO technology can help you use resources more effectively, cut down on hardware costs, and consolidate servers in new and powerful ways. This article provided the background and basics on how to make a simple VIO environment work, but the best way to fully understand all the concepts is to put them into practice and set up and configure some servers on your own. The one piece of advice I leave you with is to plan, document, and test everything before putting it into production. It will be worth it.

Resources

Learn

- [VIO and Virtualization wiki](#). Check out the developerWorks VIO wiki.
- [IBM PowerVM Virtualization Managing and Monitoring](#). Check out the IBM Redbook.
- [VIO](#). Check out the Power6 Virtual I/O documentation.
- [AIX and UNIX developerWorks zone](#): The AIX and UNIX zone provides a wealth of information relating to all aspects of AIX systems administration and expanding your UNIX skills.
- [New to AIX and UNIX?](#) Visit the New to AIX and UNIX page to learn more.
- [Technology bookstore](#): Browse the technology bookstore for books on this and other technical topics.

Discuss

- Follow [developerWorks on Twitter](#).
- Get involved in the [My developerWorks community](#).
- Participate in the AIX and UNIX forums:
 - [AIX Forum](#)
 - [AIX Forum for developers](#)
 - [Cluster Systems Management](#)
 - [IBM Support Assistant Forum](#)
 - [Performance Tools Forum](#)
 - [Virtualization Forum](#)
 - More [AIX and UNIX Forums](#)

About the author

Christian Pruett

Christian Pruett is a senior UNIX systems administrator with more than 14 years of experience with AIX, Sun Solaris, Linux, and HP/UX in a wide variety of industries, including computing, agriculture, and telecommunications. He is the co-author of two IBM Redbooks on AIX, has served as a UNIX book review for O'Reilly Publishing, and has

worked on several of the IBM AIX certification exams. He resides in Colorado with his wife and two children. You can reach Christian at pruettc@gmail.com.