

# VIOS “best practices”

*Thomas Prokop*

*Version 1.0*

# This document has content and ideas from many

Most notable is John Banchy.

Other include:

- ▶ Bob G Kovacs - Austin
- ▶ Jorge R Noguerras - Austin
- ▶ Dan Braden - ATS
- ▶ John Tesch - ATS
- ▶ Ron Barker – ATS
- ▶ Viraf Patel – ATS
- ▶ Steve Knudson – ATS
- ▶ Stephen Nasypany - ATS
- ▶ Tenley Jackson - Southfield
- ▶ Greg R Lee - Markham
- ▶ Armin M. Warda - Bonn
- ▶ Nigel Griffiths - UK
- ▶ Indulis Bernsteins – Australia
- ▶ Federico Vagnini – Italy
- ▶ Ray Williams – Calgary
- ▶ Timothy Marchini – Poughkeepsie
- ▶ Jaya Srikrishnan – Poughkeepsie
- ▶ Laurent Agarini – France
- ▶ Giuliano Anselmi - Italy
- ▶ Bill Armstrong – Rochester
- ▶ Ron Young – Rochester
- ▶ Kyle Lucke – Rochester
- ▶ Fred Ziecina – Rochester
- ▶ Bret Olszewski – Austin
- ▶ Rakesh Sharma – Austin
- ▶ Satya Sharma – Austin
- ▶ Erin C Burke – Austin
- ▶ Vani Ramagiri – Austin
- ▶ Anil Kalavakolanu - Austin
- ▶ Charlie Cler – St Louis
- ▶ Thomas Prokop – Minneapolis
- ▶ Jerry Petru – Minneapolis
- ▶ Ralf Schmidt-Dannert – Minneapolis
- ▶ Dan Marshall - Minneapolis
- ▶ George Potter – Milwaukee

# ***Disclaimers***

IBM has not formally reviewed this document. While effort has been made to verify the information, this document may contain errors. IBM makes no warranties or representations with respect to the content hereof and specifically disclaim any implied warranties of merchantability or fitness for any particular purpose. IBM assumes no responsibility for any errors that may appear in this document. The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

If you have any suggestions or corrections please send comments to:  
[tlproko@us.ibm.com](mailto:tlproko@us.ibm.com)

# Agenda

# Agenda

- **Opening remarks**
- **VIOS**
  - ▶ Implementation choices
  - ▶ Performance monitoring
  - ▶ Managing VIOS
  - ▶ Other general remarks
  
- **Questions will be entertained at the end, or may be sent to our moderator**



# Best practices??

## Is there really such a thing as best?

Definitions of Best Practice on the Web: google "define:Best Practice"

- A technique or methodology that, through experience and research, has proven to reliably lead to a desired result.  
[www.tjiss.net/glossary\\_b.html](http://www.tjiss.net/glossary_b.html)
- In medicine, treatment that experts agree is appropriate, accepted and widely used. Health care providers are obligated to provide patients with the best practice. Also called standard therapy or standard of care.  
[www.mdanderson.org/patients\\_public/about\\_cancer/display.cfm](http://www.mdanderson.org/patients_public/about_cancer/display.cfm)
- A way or method of accomplishing a business function or process that is considered to be superior to all other known methods.  
[www.qaproject.org/methods/resglossary.html](http://www.qaproject.org/methods/resglossary.html)
- A superior method or innovative practice that contributes to the improved performance of an organization, usually recognized as "best" by other peer organizations.  
[www.asq.org/glossary/b.html](http://www.asq.org/glossary/b.html)
- Recommendations regarding processes or techniques for the use or implementation of products or services.  
[longwoods.com/website/eletter/glossary.HTML](http://longwoods.com/website/eletter/glossary.HTML)
- It generally identifies behaviour commensurate with state-of-the-art skills and techniques in a given technical/professional area.  
[www.intesasanpaolo.com/script/sir0/isInvestor/eng/glossario/eng\\_glossario.jsp](http://www.intesasanpaolo.com/script/sir0/isInvestor/eng/glossario/eng_glossario.jsp)
- An outcome of benchmarking whereby measured best practice in any given area is actively promoted across organisations or industry sectors. ...  
[www.european-quality.co.uk/index.php](http://www.european-quality.co.uk/index.php)
- Guide and documentation to describe and standardize the use of metadata elements that best support a community's needs. Source: previous University of Melbourne Metadata Glossary  
[www.infodiv.unimelb.edu.au/metadata/glossary.html](http://www.infodiv.unimelb.edu.au/metadata/glossary.html)
- the methods and achievements of the recognised leader(s) in a particular field  
[www.et.teiath.gr/tempus/glossary.asp](http://www.et.teiath.gr/tempus/glossary.asp)
- Etc...

# Assumptions

- **You should have a good grasp of VIOS capabilities**
  - ▶ This is not a VIOS features or abilities discussion
  - ▶ Not an implementation session
- **The content is mostly based on VIOS 1.4/1.5**
  - ▶ Strongly encourage migrating up from VIOS 1.3 or earlier
- **Based on production level requirements – assume dual VIOs**
- **All Power, Power/VM, AIX, VIOS related questions will be entertained**
- **There is no single perfect answer.**
- **Best practices are like R.O.T. They are simply guidelines that often lead to good results for most situations.**

# VIO review



# Virtual I/O Server

## ■ Economical I/O Model

- ▶ Efficient utilization of shared physical resources

## ■ Reduced Infrastructure Costs

- ▶ Less SAN → HBAs, cables, switches, ...
- ▶ Less network → adapters, cables, switches, ...
- ▶ Reduce data center foot-print (i.e. power, space)

## ■ Quick deployment

- ▶ Allocate only virtual resources

## ■ Facilitates Server Consolidation

- ▶ Breaks the max. physical number of adapters issue
- ▶ Reduce SAN management costs - VIOS distributes disk space


## ■ Nice(-ish) Standard interface

- ▶ Regardless of backend storage, client sees same generic interface



# Virtual I/O Server - In Practice

- **Delivered as an Appliance**
  - ▶ One current version → 1.5.2.1 (Fix Pack 11.1)
  - ▶ SWMA=2 years support. If a fix is required, its on the latest VIOS
  - ▶ Only install approved disk device drivers software
  - ▶ Many things BANNED:
    - NFS Server, NIM Server, Web Server, Backup Server ...
  - ▶ Discourage “root” user work or even logging in at all! (after setup)
- **Has meant Very Low VIOS Problem Rate**
  - ▶ Version 1.2 fixed some initial issues
  - ▶ Proves → people fiddling causes problems
- **Problems can be self inflicted – but simple to fix**
  1. High throughput VIOS & below 1GB memory
  2. Low CPU entitlement, capped or silly VP settings




# Virtual I/O Server - In Practice

## ■ CPU

- ▶ Typical entitlement is .25
- ▶ Virtual CPU of 2
- ▶ Always run uncapped
- ▶ Run at high(er) priority

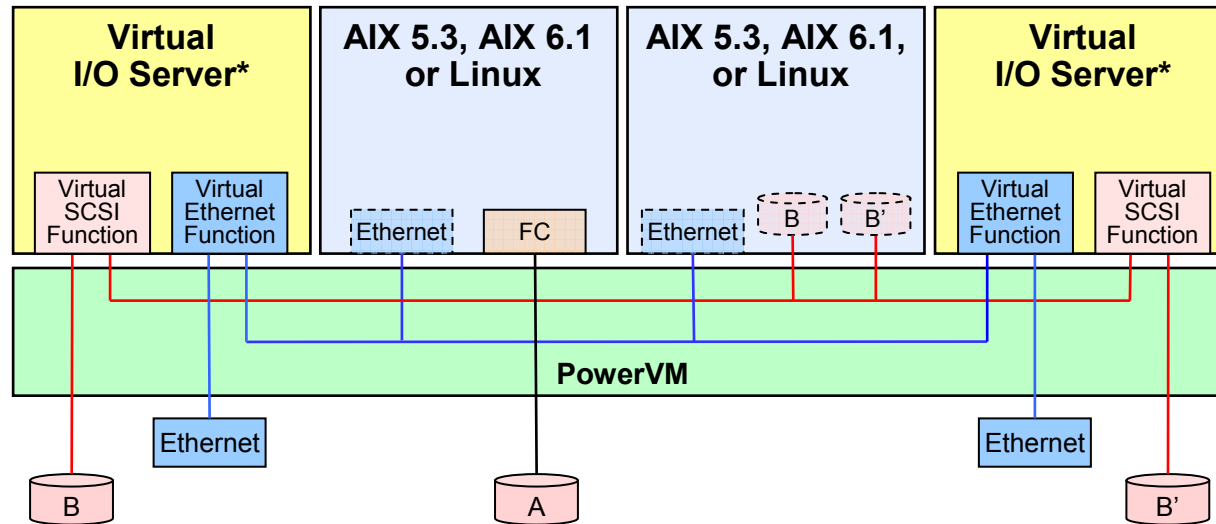
## ■ Memory

- ▶ > 1G, typically 2-4G
- ▶ More if very high device ( vscsi and hdisk ) counts
  - \* Small LUNs drive memory requirements up



# Virtual I/O Overview

# Virtual I/O



## Virtual I/O Architecture

- ▶ Mix of virtualized and/or physical devices
- ▶ Multiple VIO Servers\* supported

## Virtual SCSI

- ▶ Virtual SCSI, Fibre Channel, and DVD
- ▶ Logical and physical volume virtual disks
- ▶ Multi-path and redundancy options

## Benefits

- ▶ Reduces adapters, I/O drawers, and ports
- ▶ Improves speed to deployment

## Virtual Ethernet

- ▶ VLAN and link aggregation support
- ▶ LPAR to LPAR virtual LANs
- ▶ High availability options

\* Available on System p via the Advanced POWER virtualization features. IVM supports a single Virtual I/O Server.

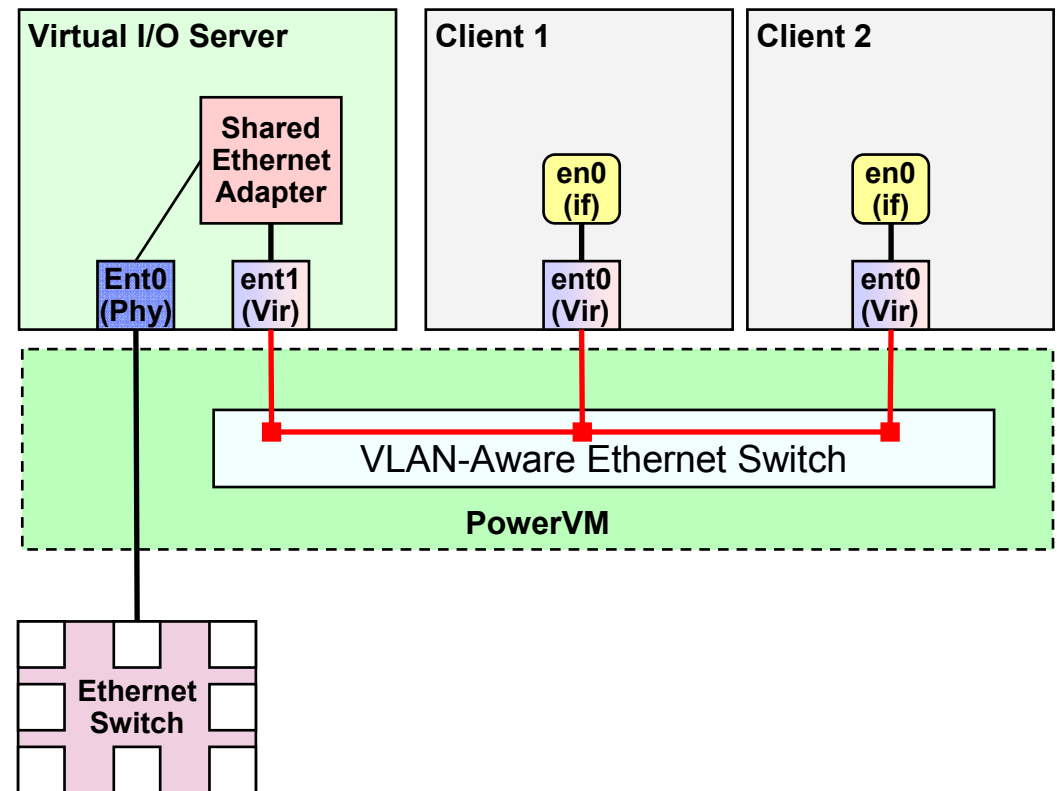
# Virtual Ethernet Options



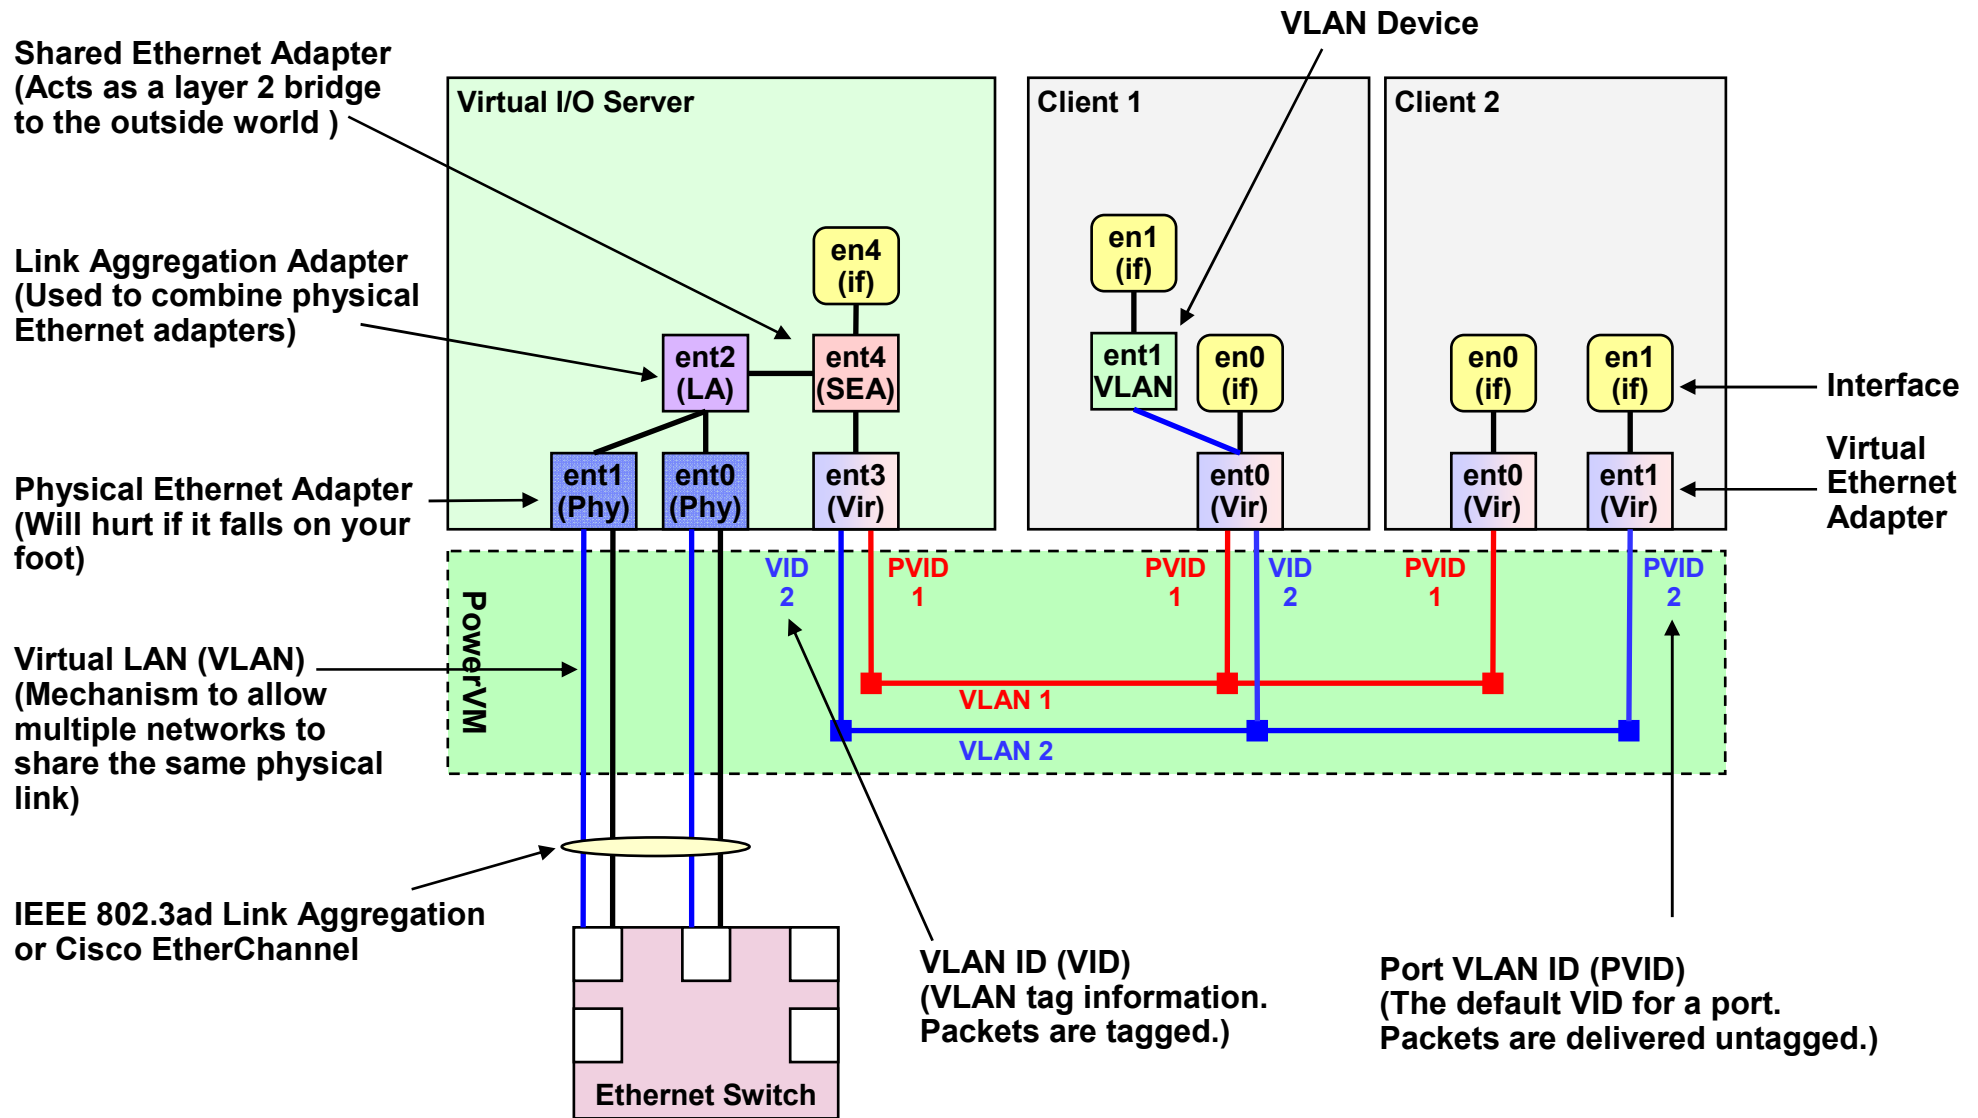
# IBM PowerVM Virtual Ethernet

## Two basic components

- ▶ VLAN-aware Ethernet switch in the Hypervisor
  - Comes standard with a POWER5/6 server.
- ▶ Shared Ethernet Adapter
  - Part of the VIO Server
  - Acts as a bridge allowing access to and from an external networks.
  - Available via PowerVM

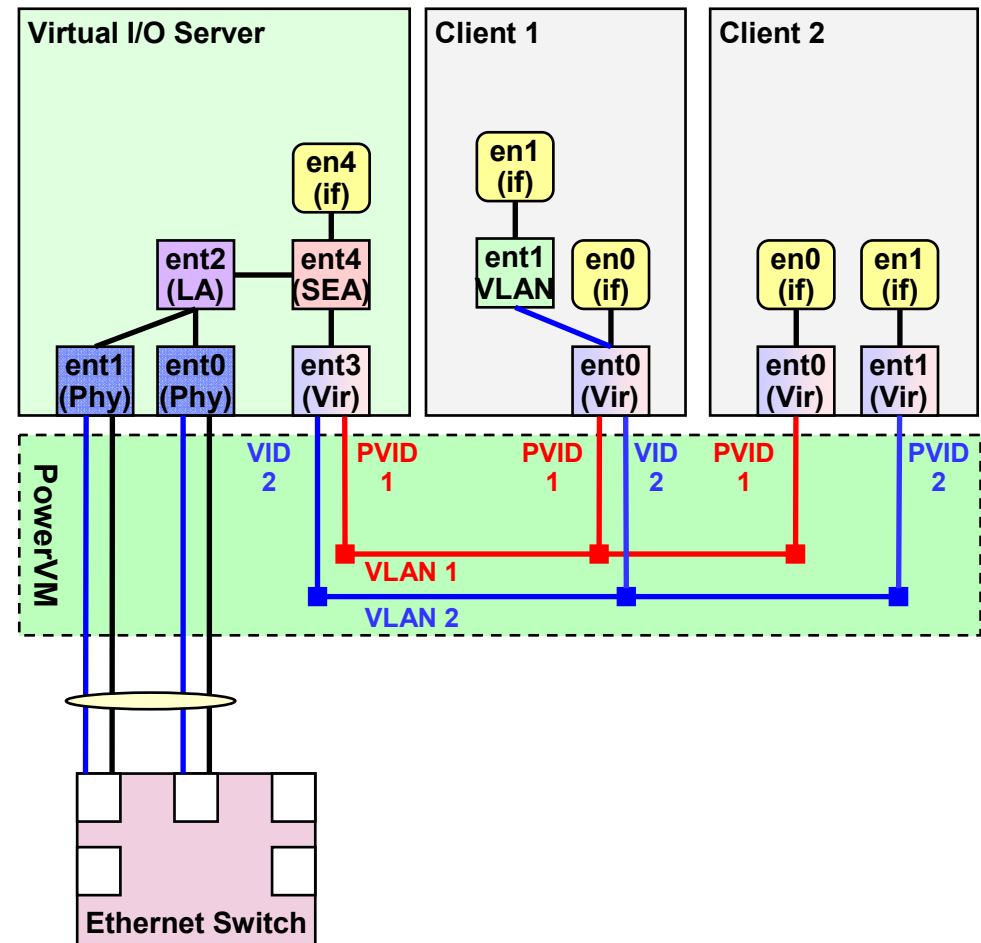


# Virtual I/O Network Terms

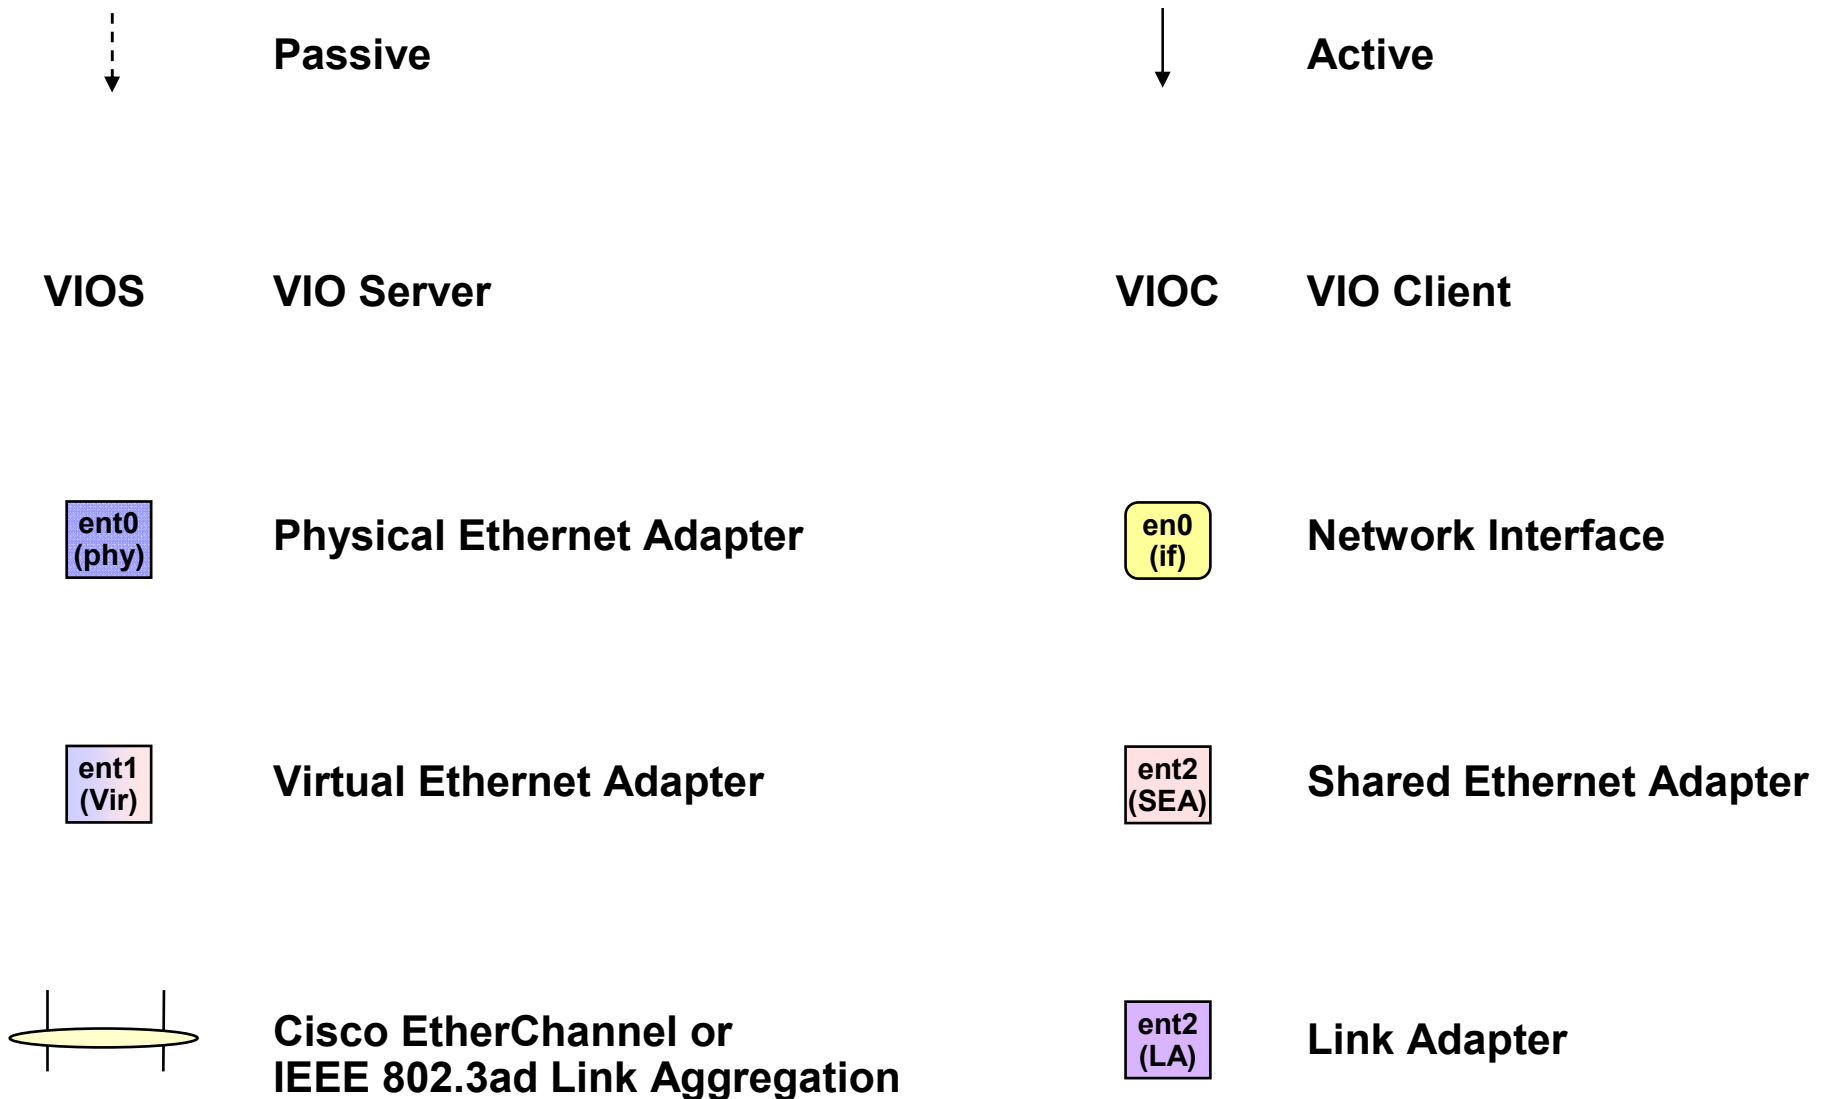


# Limits

- ▶ Maximum 256 virtual Ethernet adapters per LPAR
- ▶ Each virtual adapter can have 21 VLANs (20 VIDs, 1 PVID)
- ▶ Maximum of 16 virtual adapters can be associated with a single SEA sharing a single physical network adapter.
- ▶ No limit to the number of LPARs that can attach to a single VLAN.
- ▶ Works on OSI-Layer 2 and supports up to 4094 VLAN IDs.
- ▶ The POWER Hypervisor can support virtual Ethernet frames of up to 65408 bytes in size.
- ▶ The maximum supported number of physical adapters in a link aggregation or EtherChannel is 8 primary and 1 backup.



# Virtual Ethernet Legend



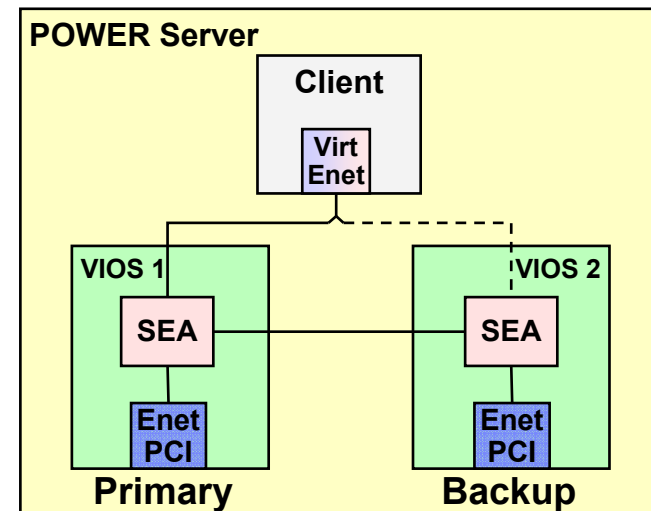
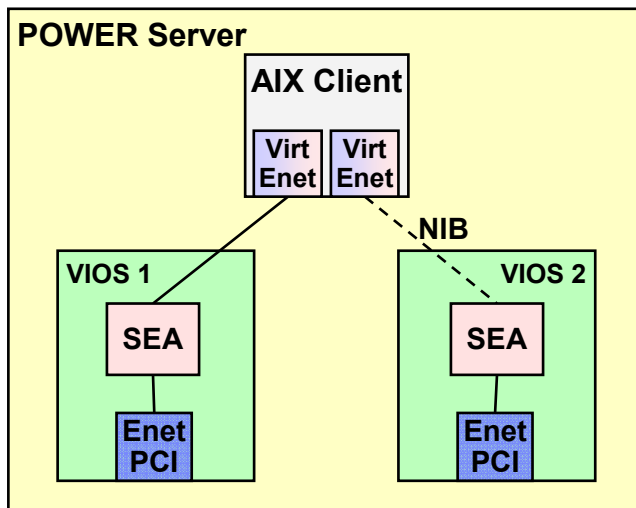
# Most common High Availability VIOS Options

## Network Interface Backup

- ▶ Set up at client level.
- ▶ Needs to ping outside host from each client to initiate NIB failover.
- ▶ Can load share clients across SEAs but LPAR to LPAR communications will happen through external switches
- ▶ VLAN-tagged traffic is not supported.
- ▶ LA is supported in VIOS
- ▶ Supported only on AIX.

## Shared Ethernet Adapter Failover

- ▶ Set up in the VIOS's only.
- ▶ Optional ping is done in VIOS on behalf of all clients.
- ▶ Cannot load-share clients between the primary and backup SEA.
- ▶ VLAN-tagged traffic is supported.
- ▶ LA is supported in VIOS
- ▶ Supported on AIX and Linux on POWER.



# Virtual Ethernet Options

## AIX Network Interface Backup (NIB), Dual VIOS

### ■ Complexity

- ▶ Requires specialized setup on client (NIB)
- ▶ Needs to ping outside host from the client to initiate NIB failover

### ■ Resilience


- ▶ Protects against single VIOS / switch port / switch / Ethernet adapter failures

### ■ Throughput / Scalability

- ▶ Allows load-sharing between VIOS's

### ■ Notes

- ▶ NIB does not support tagged VLANs on physical LAN
- ▶ Must use external switches not hubs





# Virtual Ethernet Options

## Shared Ethernet Adapter Failover, Dual VIOS

### Complexity

- ▶ Specialized setup confined to VIOS

### Resilience


- ▶ Protection against single VIOS / switch port / switch / Ethernet adapter failure

### Throughput / Scalability

- ▶ Cannot do load-sharing between primary and backup SEA (backup SEA is idle until needed).
- ▶ SEA failure initiated by:
  - Backup SEA detects the active SEA has failed.
  - Active SEA detects a loss of the physical link
  - Manual failover by putting SEA in standby mode
  - Active SEA cannot ping a given IP address.

### Notes

- ▶ Requires VIOS V1.2 and SF235 platform firmware
- ▶ Can be used on any type of client (AIX, Linux)
- ▶ Outside traffic may be tagged



# Virtual Ethernet Options

## Shared Ethernet Adapter Failover, Dual VIOS with VLANs

### ■ Complexity

- ▶ Specialized setup confined to VIOS's
- ▶ Requires VLAN setup of appropriate switch ports

### ■ Resilience

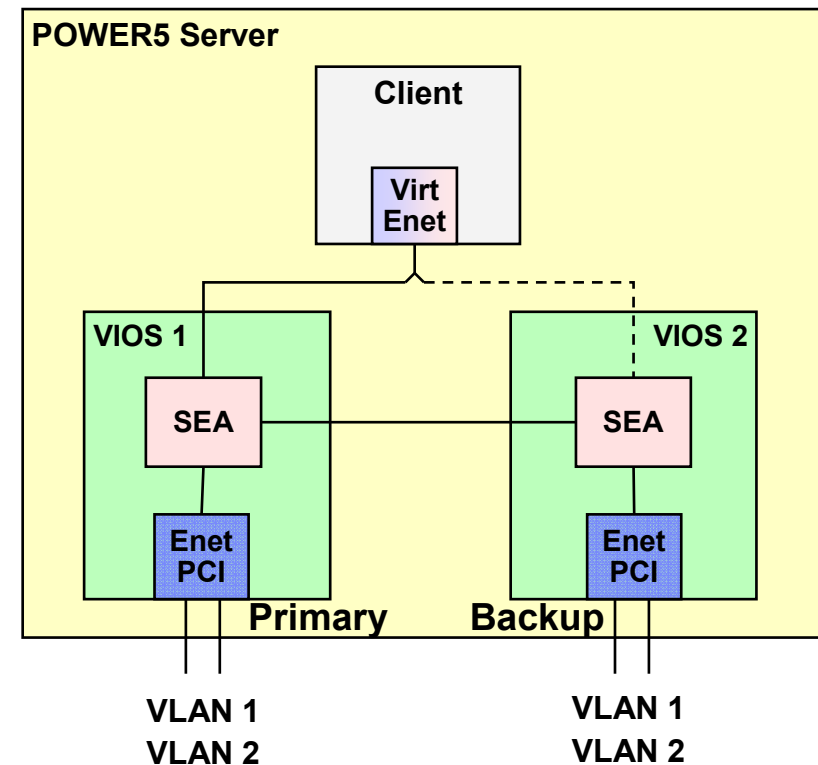
- ▶ Protection against single VIOS /switch port / switch / Ethernet adapter failure

### ■ Throughput / Scalability

- ▶ Cannot do load-sharing between VIOS's (backup SEA is idle until needed)

### ■ Notes

- ▶ Tagging allows different VLANs to coexist within same CEC
- ▶ Requires VIOS V1.2 and SF235 platform firmware



# Virtual Ethernet Rules of Thumb and Other Topics “best practices”

# Virtual Ethernet

## ■ General Best Practices

- ▶ Keep things simple
- ▶ Use PVIDs and separate virtual adapters for clients rather than stacking interfaces and using VIDs.
- ▶ Use hot-pluggable network adapters for the VIOS instead of the built-in integrated network adapters. They are easier to service.
- ▶ Use dual VIO Servers to allow concurrent online software updates to the VIOS.
- ▶ Configure an IP address on the SEA itself. This ensures that network connectivity to the VIOS is independent of the internal virtual network configuration. It also allows the ping feature of the SEA failover.
- ▶ For the most demanding network traffic use dedicated network adapters.

# Virtual Ethernet

## ■ Link Aggregation

- ▶ All network adapters that form the link aggregation (not including a backup adapter) must be connected to the same network switch.

## ■ Virtual I/O Server

- ▶ Performance scales with entitlement, not the number of virtual processors
- ▶ Keep the attribute `tcp_pmtu_discover` set to “active discovery”
- ▶ Use SMT unless your application requires it to be turned off.
- ▶ If the VIOS server partition will be dedicated to running virtual Ethernet only, it should be configured with threading disabled (Note: this does not refer to SMT).
  - `chdev -dev entX thread=0` (disable)
- ▶ Define all VIOS physical adapters (other than those required for booting) as desired rather than required so they can be removed or moved.
- ▶ Define all VIOS virtual adapters as desired not required.

# Virtual Ethernet Performance

## ■ Performance - Rules of Thumb

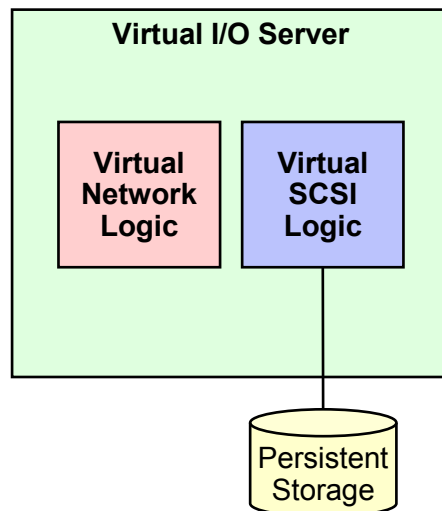
- ▶ Choose the largest MTU size that makes sense for the traffic on the virtual network.
- ▶ In round numbers, the CPU utilization for large packet workloads on jumbo frames is about half the CPU required for MTU 1500.
- ▶ Simplex, full and half-duplex jobs have different performance characteristics
  - Full duplex will perform better, if the media supports it
  - Full duplex will NOT be 2 times simplex, though, because of the ACK packets that are sent; about 1.5x simplex (Gigabit)
  - Some (rare) workloads require simplex or half-duplex
- ▶ Consider the use of TCP Large Send Offload
  - Large send allows a client partition send 64kB of packet data through a Virtual Ethernet connection irrespective of the actual MTU size
  - This results in less trips through the network stacks on both the sending and receiving side and a reduction in CPU usage in both the client and server partitions



# Shared Ethernet Threading


## Virtual I/O Server Priority

- ▶ Ethernet has no persistent storage so it runs at a higher priority level than virtual SCSI.
- ▶ If there is not enough CPU, virtual SCSI performance could suffer.
- ▶ To solve this issue, IBM made an enhancement in VIOS 1.3.



## Virtual I/O Server 1.3

- ▶ SEA can use kernel threads for pack processing.
- ▶ More evenly distributes processing power between virtual Ethernet and virtual SCSI.
- ▶ Threading is now the default.
- ▶ Can be turned off to provide slightly better Ethernet performance. (Recommended for Ethernet only VIOS).
- ▶ `chdev -dev entX thread=1` (Enable)



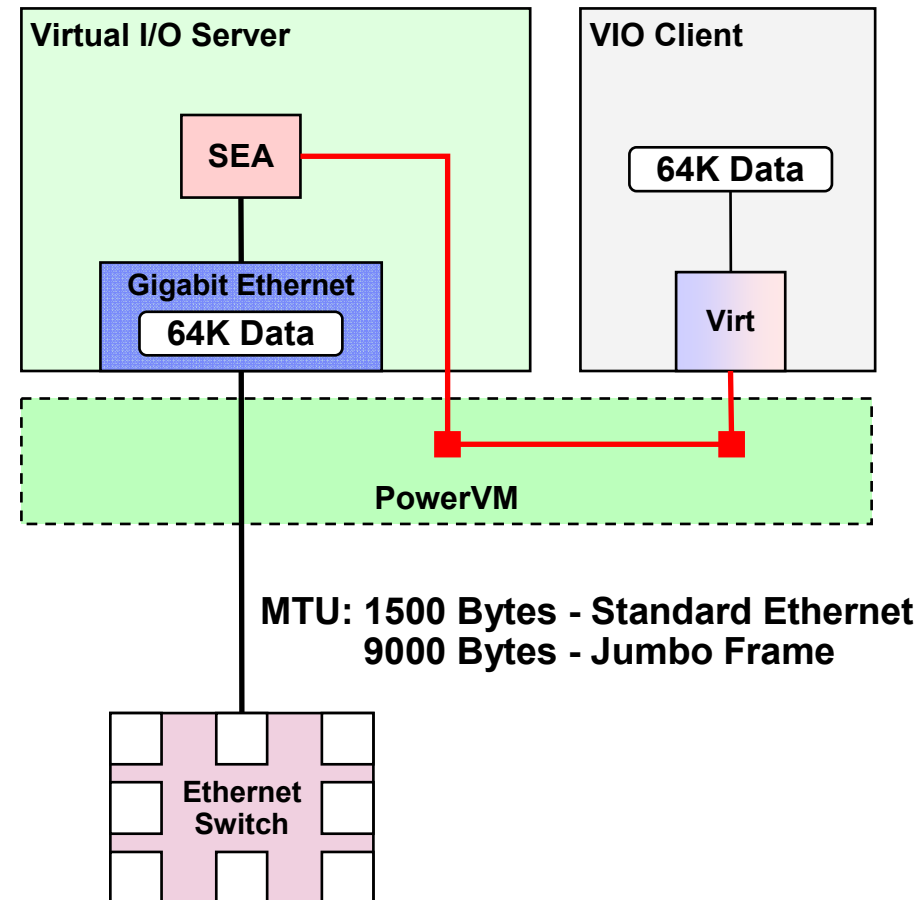
# TCP Large Send Offload

## Large Send Offload

- ▶ Client IP stack sends up to 64K of data to the Gigabit Ethernet adapter
- ▶ Gigabit Ethernet adapter breaks the packet into smaller TCP packets to match the outgoing MTU size.
- ▶ Available on VIOS 1.3 when SEA is bridging traffic (not for SEA generated traffic)
- ▶ Use with PCI-X GigE adapters lowers CPU utilization and provides best throughput
- ▶ Also available for AIX client-to-client traffic not involving the VIOS.

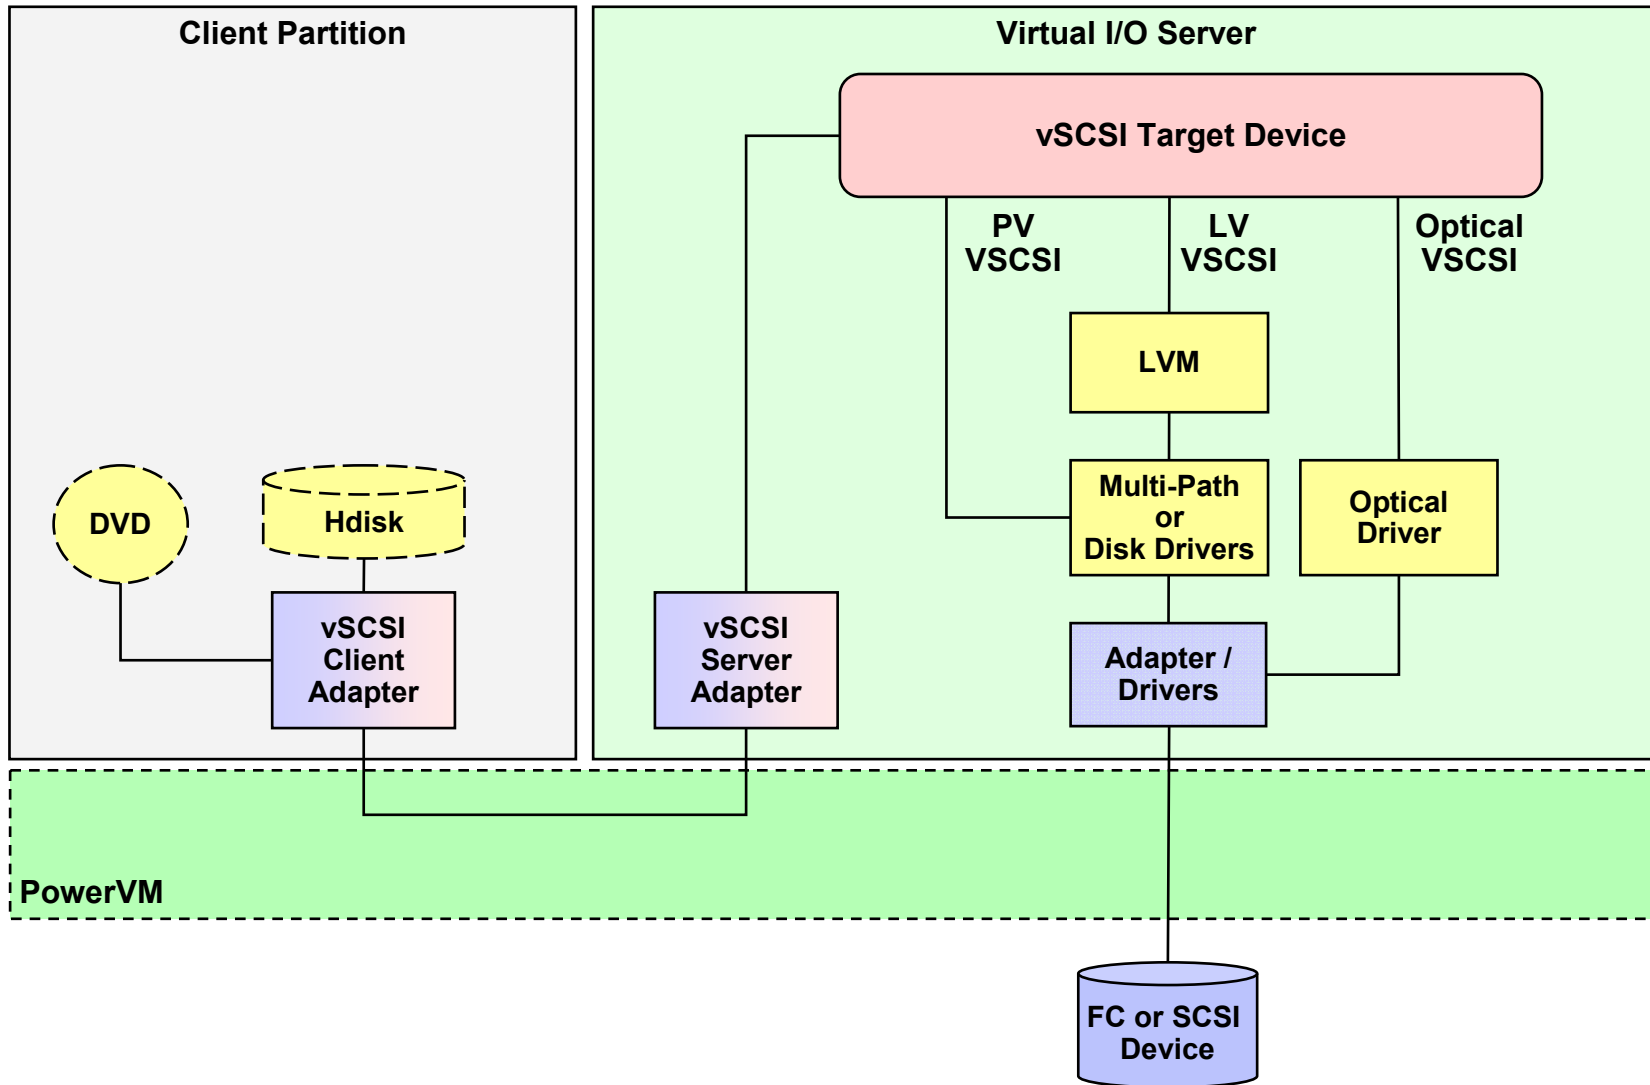
## Setup

- ▶ Run the following command on all virtual Ethernet adapters
  - `chdev -dev entX -attr large_send=yes`
- ▶ Run the following command on the SEA
  - `chdev -dev entX -attr largesend=1`



# Virtual SCSI Options

# Virtual SCSI Basic Architecture



# VIOS Multi-Path Options

VIOS attached storage <sup>1</sup>	Storage software running on the VIOS <sup>3</sup>	MPIO supported on the AIX client
IBM System Storage DS8000 / DS6000 and ESS	SDDPCM	YES
	SDD <sup>2</sup>	NO
IBM TotalStorage DS4000 / FASTT	RDAC	YES
SAN Volume Controller	SDD <sup>2</sup>	NO
	SDDPCM	YES
IBM System Storage N3700, N5000, N7000	MPIO (default PCM)	YES
NetApp	MPIO (default PCM)	YES
EMC	PowerPath <sup>2</sup>	YES
	MPIO (default PCM)	YES
HDS	HDLM <sup>2</sup>	YES
	MPIO (default PCM)	YES
HP	AutoPath	YES
	MPIO (default PCM)	YES

1. See vendor documentation for specific supported models, microcode requirements, AIX levels, etc.
2. Virtual disk devices created in certain non-MPIO VIOS environments may require a migration effort in the future. This migration may include a complete back-up and restore of all data on affected devices. See the VIOS FAQ on the VIOS website for additional detail.
3. Not all multi-path codes are compatible with one another. MPIO *compliant* code is not supported with non MPIO *compliant* code for similar disk subsystems (e.g., one cannot use MPIO for one EMC disk subsystem and PowerPath for another on the same VIOS, *nor* can one use SDD and SDDPCM on the same VIOS).

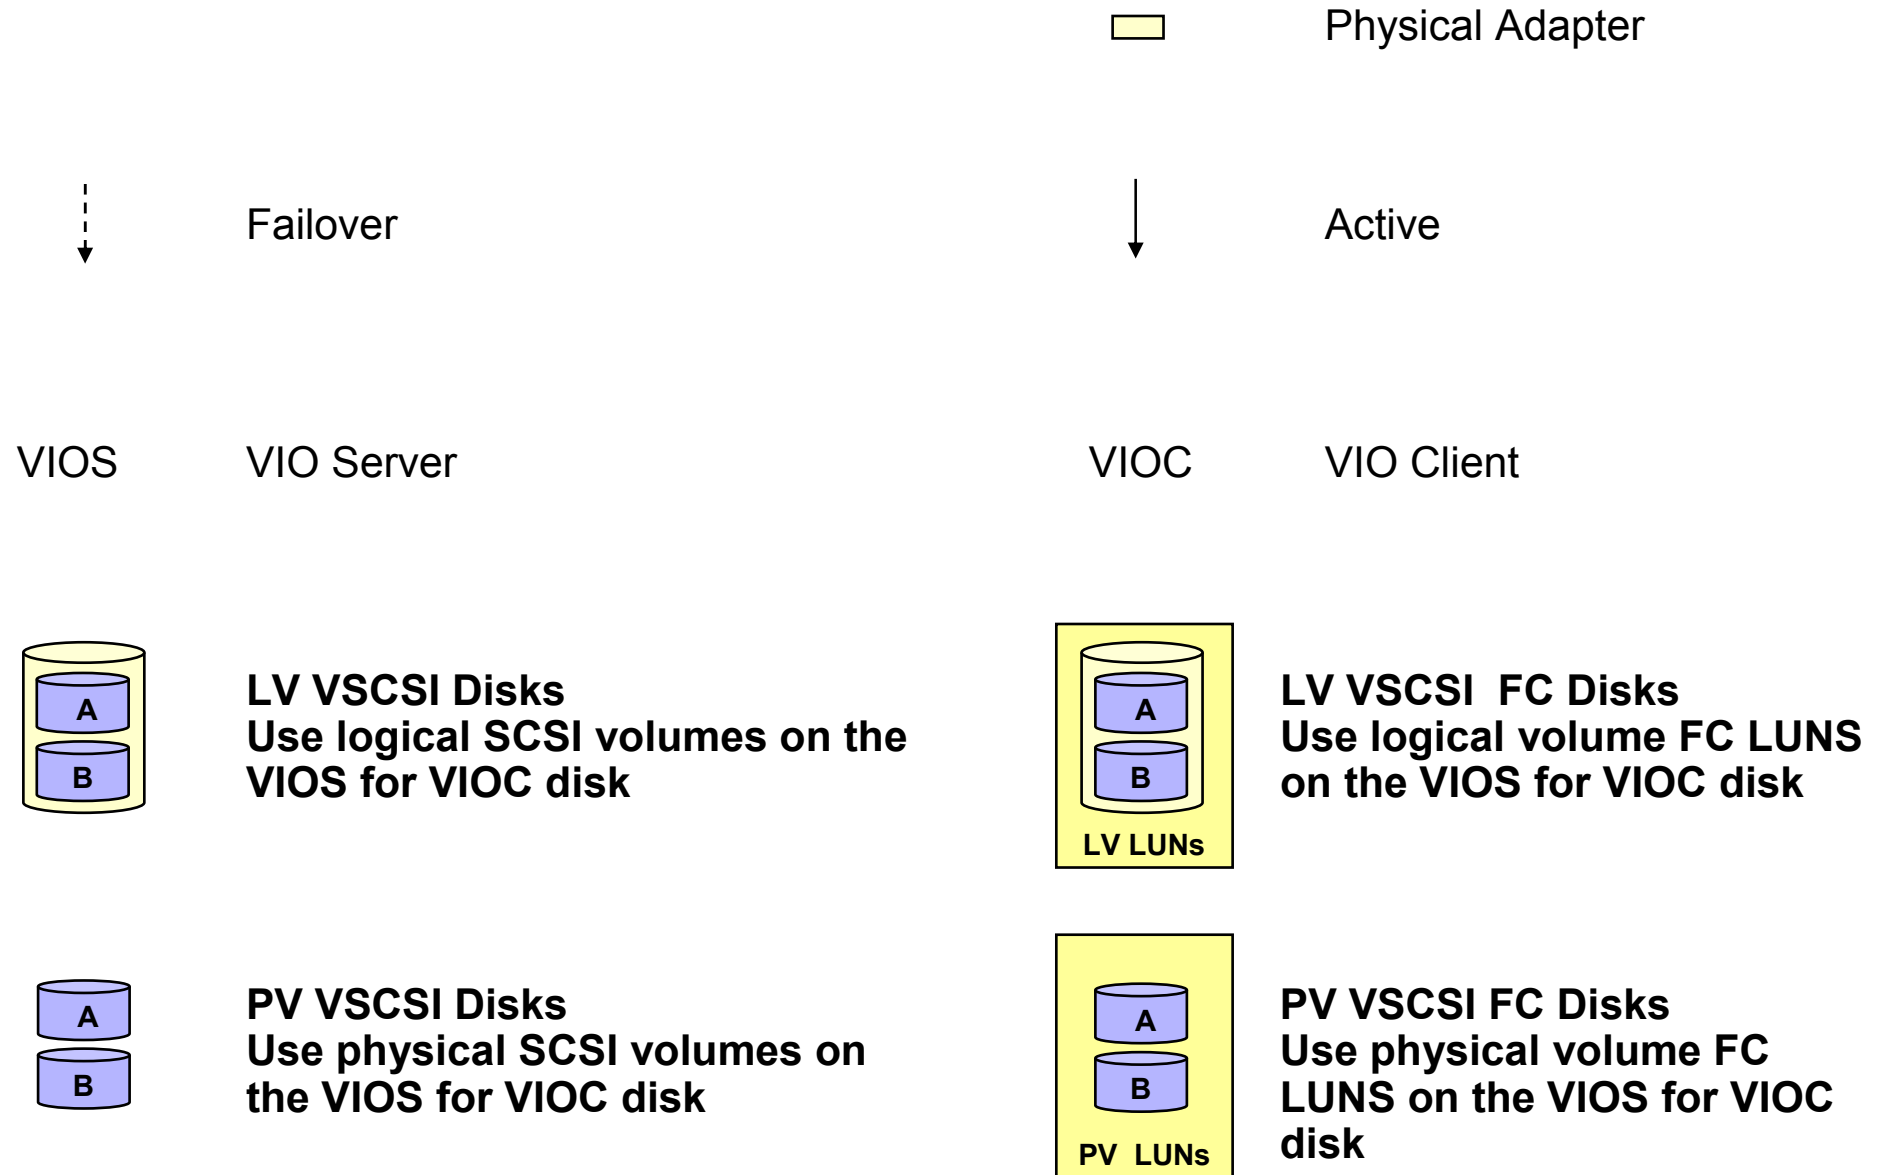
Separate sets of FC adapters are required when using different multi-path codes on a VIO. If incompatible multi-path codes are required, then one should use separate VIOSs for the incompatible multi-path codes. In general, multi-path codes that adhere to the MPIO architecture are compatible.

Sources:

<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>

[http://w3-1.ibm.com/sales/systems/portal/\\_s.155/254?navID=f220s240&geoID=AM&prodID=pSeries&docID=rsfcsk.skit&docType=SalesKit&skCat=DocumentType](http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geoID=AM&prodID=pSeries&docID=rsfcsk.skit&docType=SalesKit&skCat=DocumentType)

# Virtual SCSI Legend



# Virtual SCSI commonly used options

## AIX Client Mirroring, Single Path in VIOS, LV VSCSI Disks

### ■ Complexity

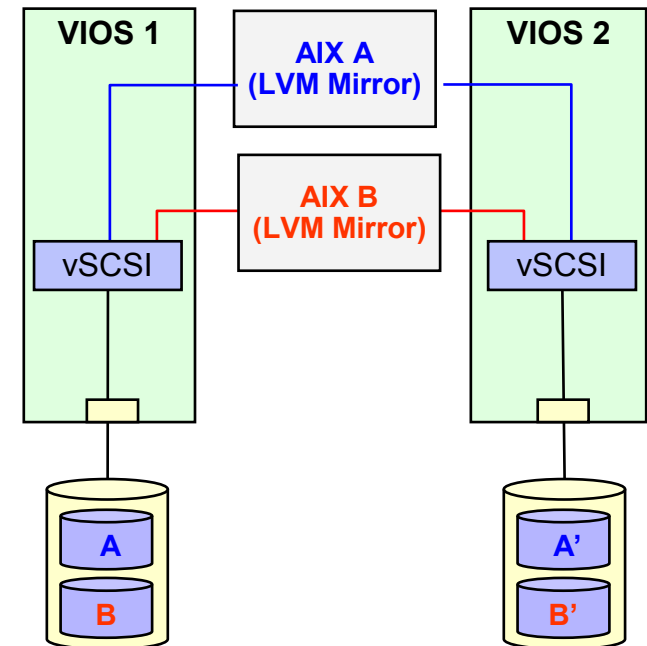
- ▶ More complicated than single VIO server but does not require SAN ports or setup
- ▶ Requires LVM mirroring to be setup on the client
- ▶ If a VIOS is rebooted, the mirrored disks will need to be resynchronized via a varyonvg on the VIOC.

### ■ Resilience

- ▶ Protection against failure of single VIOS / SCSI disk / SCSI controller.

### ■ Throughput / Scalability

- ▶ VIOS performance limited by single SCSI adapter and internal SCSI disks.





# Virtual SCSI Commonly used options

## AIX MPIO Default PCM Driver in Client, Multi-Path I/O in VIOS

### Complexity

- Requires MPIO to be setup on the client
- Requires Multi-Path I/O setup on the VIOS

### Resilience

- Protection against failure of a single VIOS, FC adapter, or path.
- Protection against FC adapter failures within VIOS

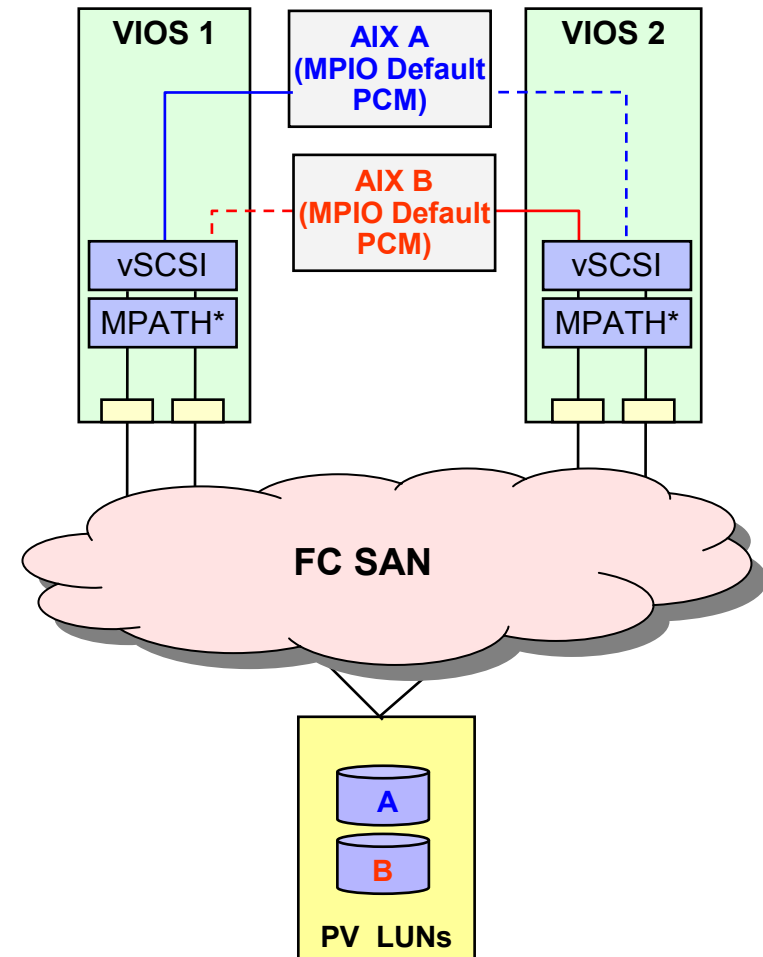
### Throughput / Scalability

- Potential for increased bandwidth due to Multi-Path I/O
- Primary LUNs can be split across multiple VIOS to help balance the I/O load.

### Notes

- Must be PV VSCSI disks.
- \* may facilitate HA and DR efforts
- VIOS Adapters can be serviced w/o VIOC awareness

Default MPIO PCM in the Client  
Supports Failover Only



\* Note: See the slide labeled VIOS Multi-Path Options for a high level overview of MPATH options.

# Virtual SCSI commonly used options

## AIX MPIO Default PCM Driver in Client, Multi-Path iSCSI in VIOS

### ■ Complexity

- ▶ Requires MPIO to be setup on the client
- ▶ Requires Multi-Path I/O setup on the VIOS

### ■ Resilience

- ▶ Protection against failure of a single VIOS, Ethernet adapter, or path.
- ▶ Protection against Ethernet adapter failures within VIOS

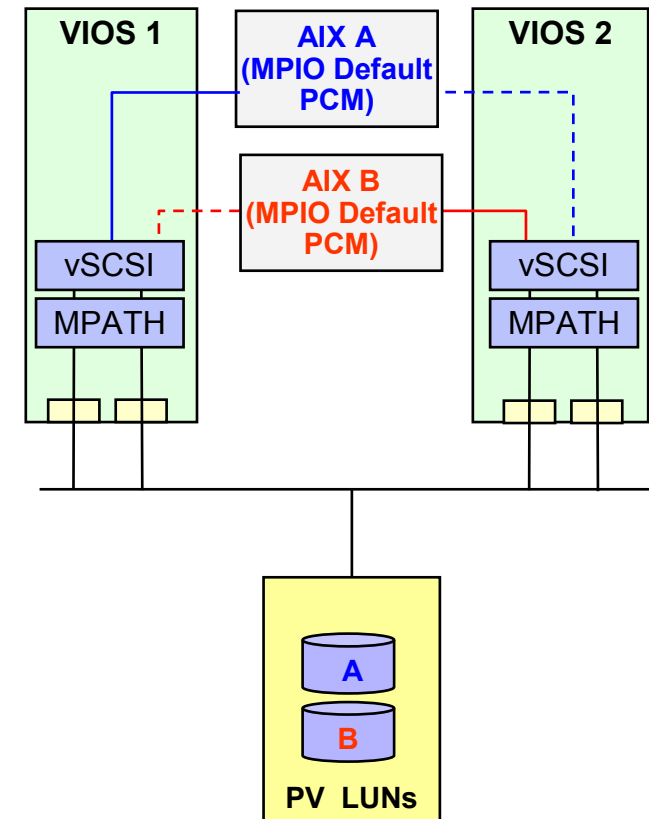
### ■ Throughput / Scalability

- ▶ Potential for increased bandwidth due to Multi-Path I/O
- ▶ Primary LUNs can be split across multiple VIOS to help balance the I/O load.

### ■ Notes

- ▶ Must be PV VSCSI disks.
- ▶ \* may facilitate HA and DR efforts
- ▶ VIOS Adapters can be serviced w/o VIOC awareness

Default MPIO PCM in the Client  
Supports Failover Only



\* Note: See the slide labeled VIOS Multi-Path Options for a high level overview of MPATH options.

# Virtual SCSI “best practices”

# Virtual SCSI General Notes

## ■ Notes

- ▶ Make sure you size the VIOS to handle the capacity for normal production and peak times such as backup.
- ▶ Consider separating VIO servers that contain disk and network if extreme VIOS throughput is desired, as the tuning issues are different
- ▶ LVM mirroring is supported for the VIOS's own boot disk
- ▶ A RAID card can be used by either (or both) the VIOS and VIOC disk
- ▶ For performance reasons, logical volumes within the VIOS that are exported as virtual SCSI devices should not be striped, mirrored, span multiple physical drives, or have bad block relocation enabled..
- ▶ SCSI reserves have to be turned off whenever we share disks across 2 VIOS. This is done by running the following command on each VIOS:  
  

```
# chdev -l <hdisk#> -a reserve_policy=no_reserve
```
- ▶ Set VSCSI Queue depth to match underlying real devices, don't exceed aggregate queue depth
- ▶ You don't need a large number of VSCSI adapters per client, 4 is typically sufficient

# Virtual SCSI General Notes....

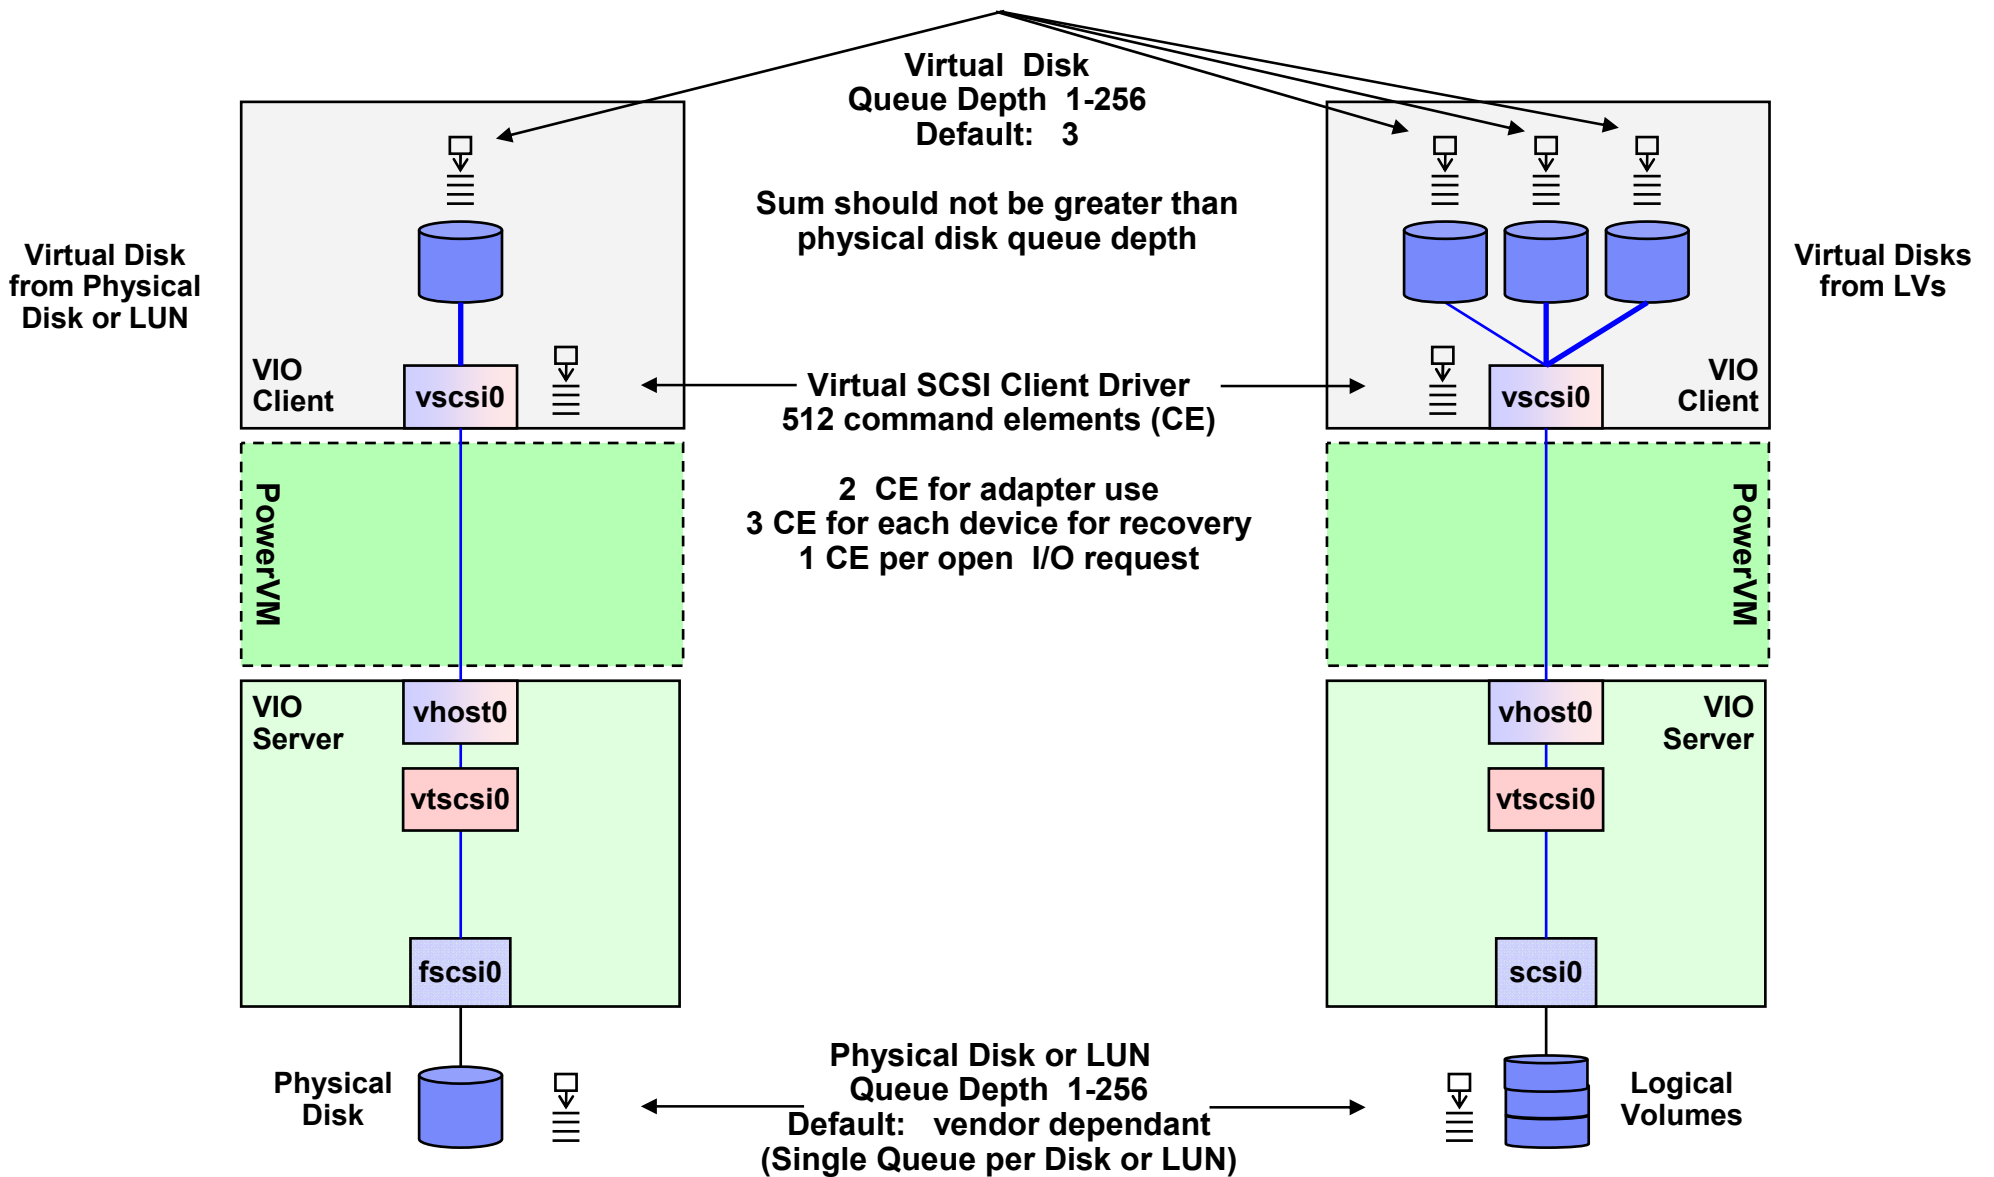
## ■ Notes

- ▶ If you are using FC Multi-Path I/O on the VIOS, set the following fscsi device values (requires switch attach):
  - dyntrk=yes (Dynamic Tracking of FC Devices)
  - fc\_err\_recov= fast\_fail (FC Fabric Event Error Recovery Policy)  
(must be supported by switch)
  
- ▶ If you are using MPIO on the VIOC, set the following hdisk device values:
  - hcheck\_interval=60 (Health Check Interval)
  
- ▶ If you are using MPIO on the VIOC set the following hdisk device values on the VIOS:
  - reserve\_policy=no reserve (Reserve Policy)

# Fibre Channel Disk Identification

- **The VIOS uses several methods to tag devices that back virtual SCSI disks:**
  - ▶ Unique device identifier (UDID)
  - ▶ IEEE volume identifier
  - ▶ VIOS Identifier
  
- **Comments**
  - ▶ The tagging method may have a bearing on the block format of the disk. The preferred disk identification method for virtual disks is the use of UDIDs. This is the method used by AIX MPIO support (Default PCM and vendor provided PCMs)
  - ▶ Because of the different data format associated with the PVID method, customers with non-MPIO environments should be aware that certain actions performed in the VIOS LPAR, using the PVID method, may require data migration, that is, some type of backup and restore of the attached disks.
  - ▶ Hitachi Dynamic Link Manager (HDLM) version 5.6.1 or above with the unique\_id capability enables the use of the UDID method and is the recommended HDLM path manager level in a VIO Server. Note: By default, unique\_id is not enabled
  - ▶ PowerPath version 4.4.2.2 has also added UDID support.
  
- **Note: care must be taken when upgrading multipathing software in the VIOS particularly if going from pre-UDID version to a UDID version. Legacy devices (those vSCSI devices created prior to UDID support) need special handling. See vendor instructions for upgrading multipathing software on the VIOS.**
  
- **See vendor documentation for details. Some information can be found at:**
  - ▶ EMC: [www.emc.com/interoperability/matrices/EMCSupportMatrix.pdf](http://www.emc.com/interoperability/matrices/EMCSupportMatrix.pdf)
  - ▶ Hitachi: [www.hds.com/assets/pdf/wp\\_hdlm\\_mpio\\_in\\_vio\\_TM-001-00.pdf](http://www.hds.com/assets/pdf/wp_hdlm_mpio_in_vio_TM-001-00.pdf)

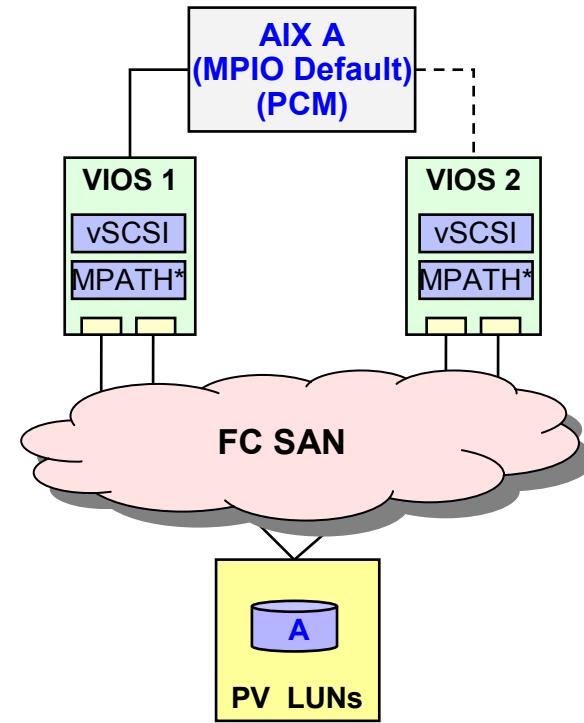
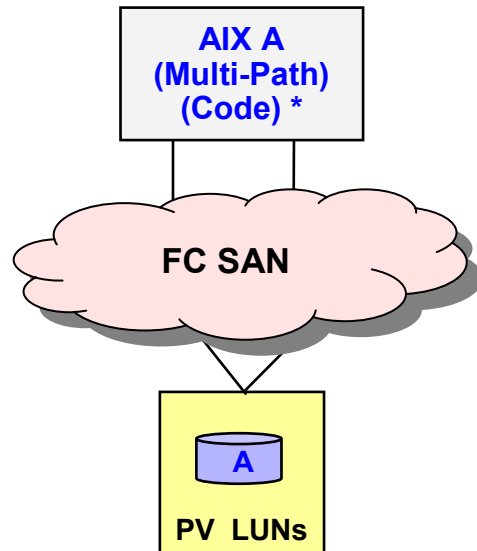
# SCSI Queue Depth





# Boot From SAN

# Boot From SAN



## ▪ Boot Directly from SAN

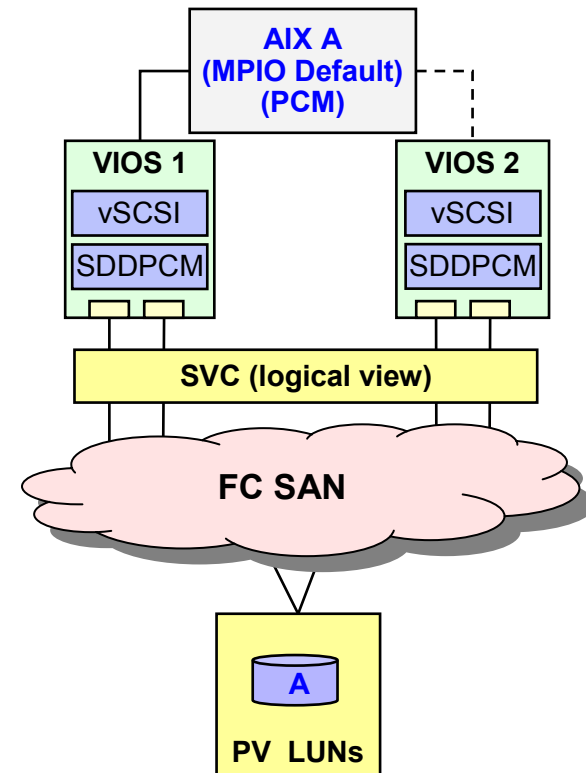
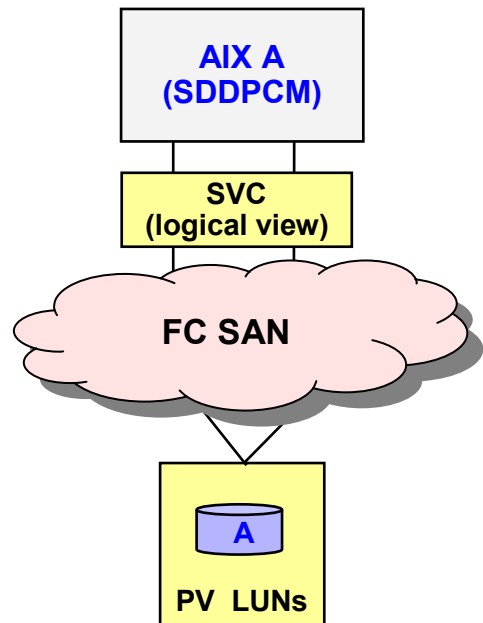
- ▶ Storage is zoned directly to the client
- ▶ HBAs used for boot and/or data access
- ▶ Multi-path code of choice runs in client

## ▪ SAN Sourced Boot Disks

- ▶ Affected LUNs are zoned to VIOS(s) and assigned to clients via VIOS definitions
- ▶ HBAs in VIOS are independent of any HBAs in client
- ▶ Multi-path code in the client will be the MPIO default PCM for disks seen through the VIOS.

\* Note: There are a variety of Multi-Path options. Consult with your storage vendor for support.

# Boot From SAN using a SAN Volume Controller



## ▪ Boot from an SVC

- ▶ Storage is zoned directly to the client
- ▶ HBAs used for boot and/or data access
- ▶ SDDPCM runs in client (to support boot)

## ▪ Boot from SVC via VIO Server

- ▶ Affected LUNs are zoned to VIO(s) and assigned to clients via VIOS definitions
- ▶ HBAs in VIOS are independent of any HBAs in client
- ▶ Multi-path code in the client will be the MPIO default PCM for disks seen through the VIOS.

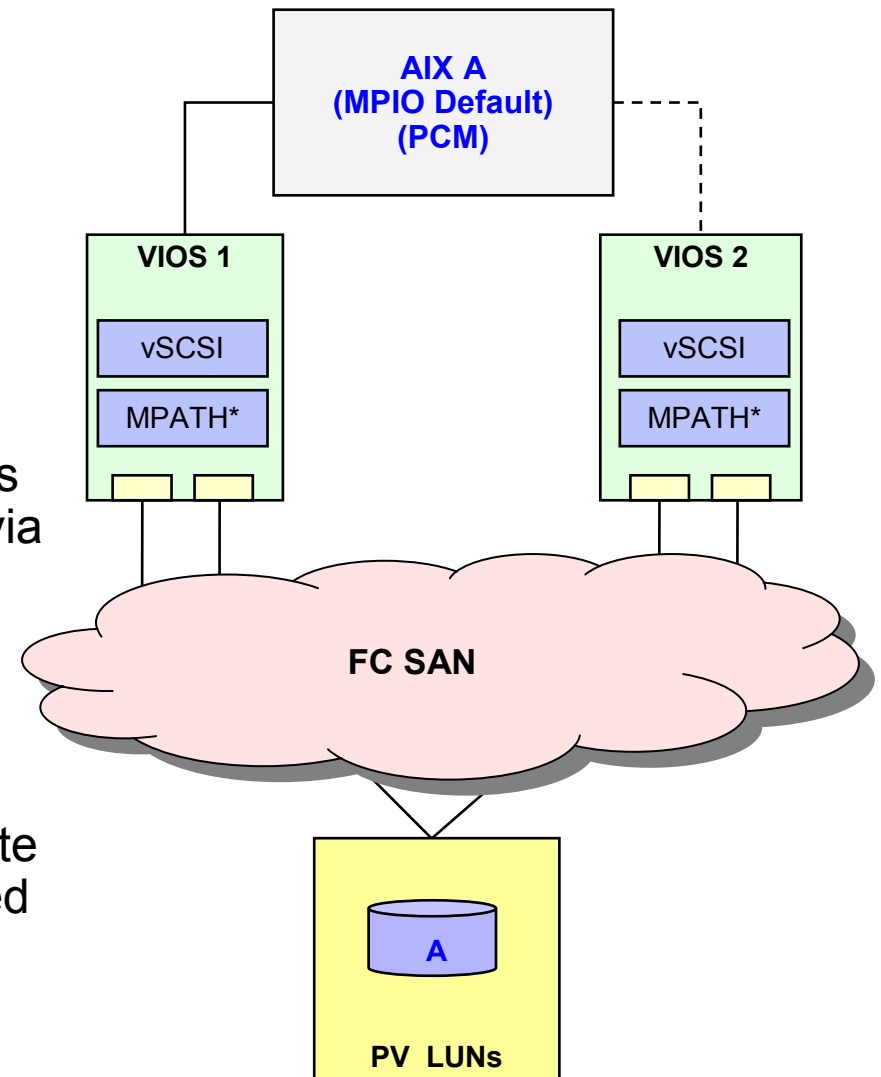
# Boot from SAN via VIO Server

## ■ Client

- ▶ Uses the MPIO default PCM multi-path code.
- ▶ Active to one VIOS at a time.
- ▶ The client is unaware of the type of disk the VIOS is presenting (SAN or local)
- ▶ The client will see a single LUN with two paths regardless of the number of paths available via the VIOS

## ■ VIOS

- ▶ Multi-path code is installed in the VIOS.
- ▶ A single VIOS can be brought off-line to update VIOS or multi-path code allowing uninterrupted access to storage.



# Boot from SAN vs. Boot from Internal Disk

## ■ Advantages

- ▶ Boot from SAN can provide a significant performance boost due to cache on disk subsystems.
  - Typical SCSI access: 5-20 ms
  - Typical SAN write: 2 ms
  - Typical SAN read: 5-10 ms
  - Typical Single disk : 150 IOPS
- ▶ Can mirror (O/S), use RAID (SAN), and/or provide redundant adapters
- ▶ Easily able to redeploy disk capacity
- ▶ Able to use copy services (e.g. FlashCopy)
- ▶ Fewer I/O drawers for internal boot are required
- ▶ Generally easier to find space for a new image on the SAN
- ▶ Booting through the VIOS could allow pre-cabling and faster deployment of AIX

## ■ Disadvantages

- ▶ Will loose access (and crash) if SAN access is lost.
  - \* SAN design and management becomes critical
- ▶ If dump device is on the SAN the loss of the SAN will prevent a dump.
- ▶ It may be difficult to change (or upgrade) multi-path codes as they are in use by AIX for its own need.
  - You may need to move the disks off of SAN, unconfigure and remove the multi-path software, add the new version, and move the disk back to the SAN.
  - **This issue can be eliminated with boot through dual VIOS.**

# System DUMP Considerations

## ■ System DUMP Alternative

- ▶ If dump device is on the SAN the loss of the SAN will prevent a dump
- ▶ The primary and secondary dump spaces can be on rootvg or non-rootvg disks.
  - The dump space can be outside of rootvg as long as it is not designated as permanent.
  - Dump space that is designated as permanent (better term - persistent) remains the dump space after a reboot. Permanent dump space must reside in rootvg. If you want to have your primary or secondary dump space outside of rootvg you will have to have to define the dump space after every reboot using a start up script.
  - The primary and secondary dump spaces can be on different storage areas. For example, the primary space could be on a SAN disk while the secondary space is on an internal VIOS disks (LV-VSCSI or PV-VSCSI).
- ▶ If the primary dump space is unavailable at the time if a system crash the entire dump will be placed in the secondary space.

# Boot from VIOS Additional Notes

## ■ Notes

- ▶ The decision of where to place boot devices (internal, direct FC, VIOS), is independent of where to place data disks (internal, direct FC, or VIOS).
- ▶ Boot VIOS off of internal disk.
  - LVM mirroring or RAID is supported for the VIOS's own boot disk.
  - VIOS may be able to boot from the SAN. Consult your storage vendor for multi-path boot support. This may increase complexity for updating multi-path codes
- ▶ Consider mirroring one NIM SPOT on internal disk to allow booting in DIAG mode without SAN connectivity
  - `nim -o diag -a spot=<spotname> clientname`
- ▶ PV-VSCSI disks are required with dual VIOS access to the same set of disks



# Monitoring

# topas - CEC monitoring screen

- **Split screen accessible with -C or the "C" command**
  - Upper section shows CEC-level metrics
  - Lower sections shows sorted list of shared and dedicated partitions
- **Configuration info retrieved from HMC or specified from command line**
  - c means capped, C - capped with SMT
  - u means shared, U - uncapped with SMT
  - S means SMT
- **Use topas -R to record to /etc/perf/topas\_cec.YYMMDD and topasout to display recording**

```

Sample Full-Screen Cross-Partition Output
Topas CEC Monitor                Interval: 10                Wed Mar 6 14:30:10 2005
Partitions Memory (GB)           Processors
Shr: 4      Mon: 24 InUse: 14    Mon: 8 PSz: 4 Shr_PhysB: 1.7
Ded: 4      Avl: 24              Avl: 8 APP: 4 Ded_PhysB: 4.1

Host      OS    M    Mem  InU Lp Us Sy Wa Id  PhysB  Ent %EntC Vcsw PhI
-----
-----shared-----
ptools1  A53  u    1.1  0.4  4 15  3  0  82  1.30  0.50 22.0 200  5
ptools5  A53  U    12   10   1 12  3  0  85  0.20  0.25  0.3 121  3
ptools3  A53  C    5.0  2.6   1 10  1  0  89  0.15  0.25  0.3  52  2
ptools7  A53  c    2.0  0.4   1  0  1  0  99  0.05  0.10  0.3 112  2
-----
-----dedicated-----
ptools4  A53  S    0.6  0.3   2 12  3  0  85  0.60
ptools6  A52   1.1  0.1   1 11  7  0  82  0.50
ptools8  A52   1.1  0.1   1 11  7  0  82  0.50
ptools2  A52   1.1  0.1   1 11  7  0  82  0.50

```

# VIOS Monitoring



- Fully Supported

1. Topas
2. Topasout
3. PTX
4. xmp perf/rmi
5. SNMP
6. lpar2rrd
7. IVM
8. ITMSESP
9. System Director
10. Tivoli Toolset
11. IBM ME for AIX

- Online only but has CEC screen
- Needs AIX 5.3TL5+fix bizarre formats, improving
  - Works but long term suspect (has been for years)
- Remote extraction from PTX Daemon, code sample available
  - Works OK VIOS 1.5 supports directly
- Takes data from HMC and graphs via rrdtool, 1 hour level
  - Integrated Virtualisation Manager has 5 minute numbers
  - Tivoli Monitor SESP new release is better
  - Included with AIX/VIOS, optional modules available
- Tivoli Monitor and ITMSESP plus other products
- Bundled ITMSE, TADDM, ITUAM offering

- Not Supported tools **but workable**

- ▶ Simple to remove + do not invalidate support but confirm with UNIX Support first
- 10. nmon - freely available, run fine as it is AIX
- 11. Ganglia - open source, popular cluster monitor – monitors anything
- 12. lparmon - demo tool from Dallas Demo Center on Alpha-works

- Pick one and retain data for future analysis!

- [http://www.ibm.com/collaboration/wiki/display/WikiPtype/VIOS\\_Monitoring](http://www.ibm.com/collaboration/wiki/display/WikiPtype/VIOS_Monitoring)


# The IBM ME for AIX Solution Components

## A Brief Overview

- **Systems management offering that provides discovery, monitoring, performance tracking and usage accounting capabilities that are integrated with the AIX® operating system, providing customers with the tools to effectively and efficiently manage their enterprise IT infrastructure.**
- **The following are the underlying products that are included:**
  - ▶ IBM Tivoli Monitoring (ITM), version 6.2
  - ▶ IBM Tivoli Application Dependency Discovery Manager (TADDM) 7.1
  - ▶ IBM Usage and Accounting Manager (IUAM) Virtualization Edition 7.1
- **A common prerequisite for all the above products is a limited license of DB2® 9.1, which is also included in the offering**

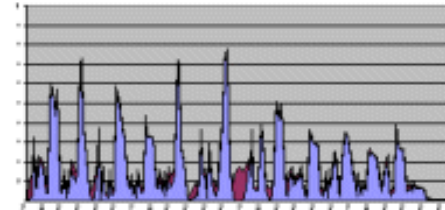


# Performance Tools



**nmon**

- AIX & Linux Performance Monitor




**nmon analyser**

- Excel spreadsheet to graph nmon output



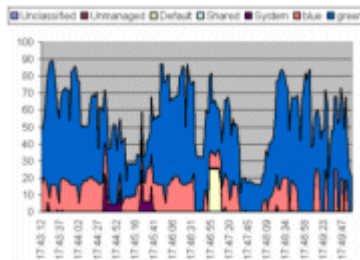
**Ganglia**

- Performance Monitoring with POWER5 additions

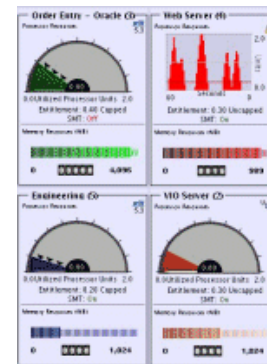


**SEA Monitoring**

- For the VIOS



**Workload Manager (WLM) for AIX**



**lparmon**

- System p LPAR Monitor

<http://www-941.ibm.com/collaboration/wiki/display/WikiPtype/Performance+Other+Tools>

# VIOS management



# Virtual I/O Server – regular version needs regular updates

- VIOS product stream:
  - Typically two updates per year
  - Supplies new function
  - New install media on request or VIOS purchase
    - Can be hard to locate old versions within IBM
      - Keep your copy!, especially for DR!
  - Fix Packs –
    - Latest downloadable from Website
    - Older via CD
- VIOS service stream
  - Fix-packs migrate VIOS installations to
    - Latest level
    - All inclusive
    - Deliver fixes → customers should expect pressure to upgrade to the latest version as a default - if they have problems
  - All VIOS fixes and updates are published to the VIOS website





# VIOS management

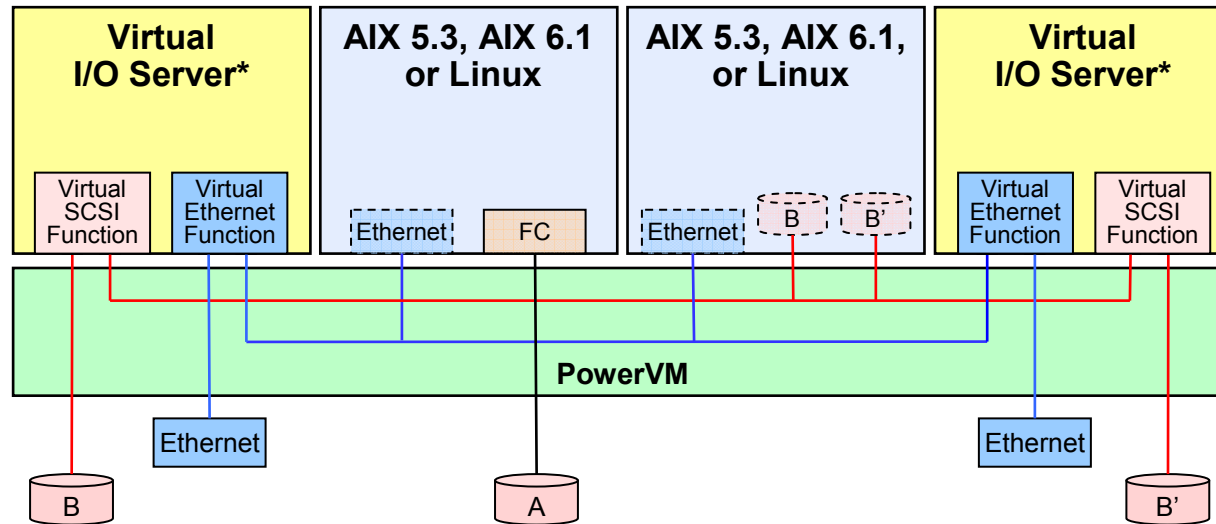
- **Plan for yearly upgrades**
  - ▶ Think about SLAs
- **Separate VIOS for**
  - ▶ Production
  - ▶ Non-production
- **Separate virtual ethernet switches**
  - ▶ Allows complete v-enet separation
  - ▶ Can be used for firewall, DMZ...
  - ▶ Very low memory requirement
- **Configure mutual 'failover' for maximum flexibility**
  - ▶ Ensures all interfaces work
  - ▶ Reduces surprises during planned or unplanned outages
  - ▶ Make sure external switches are HA
- **It's an appliance**
  - ▶ Don't change 'AIX' tunables
  - ▶ Don't open ports
  - ▶ Apply only supported VIOS fixes
  - ▶ Install only supported tools
  - ▶ Use supported commands
- **Setup for high availability**
  - ▶ 2 x VIOS ( or more)
  - ▶ mirror VIOS rootvg
  - ▶ different VIOS for different class of work
    - have a test VIOS pair
    - impact on HA ( hacmp ) decisions esp networking
    - LA and multipath for single interfaces
  - ▶ Don't mix multipath drivers / HBAs

# VIOS and DR

- **UDID device signature**
  - ▶ Use for portability between real and VSCSI
  - ▶ Multipath device and Disk subsystem configuration options
  - ▶ AIX will use PVID internally at VSCSI/VIOC level
- **Virtual / Real devices and slots must be tracked and re-matched if target HW is different**
- **Use NIM for VIOS iosbackup**
- **Consider a CMDB ( Such as Tivoli TADDM)**
- **Standards help, documentation is great**
  - ▶ Use SPT export from HMC
- **Etc...**

# Other general remarks

# Virtual I/O



- **Virtual devices require virtual slots**

- ▶ Pick a sensible naming convention
- ▶ Don't skimp on slots, lpar reboot to add more
- ▶ I typically assign 10 slots per client LPAR
  - 0,1 = Serial
  - 2->5 = ethernet
  - 7->9 = vscsi
  - Prefix in VIOS with VIOC lpar ID

- **Make all VIOS ethernet LA**

- ▶ Allows hot-add of real adapter to LA

- **Configure for performance**

- ▶ Use highest speed adapters possible
- ▶ Alternate odd/even lpar to VIOS

- **Configure for attachment density**

- ▶ Use multiport adapters, even if you only use one of the ports

# References

## Must Read List:

- **VIO Server Home**
  - <https://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/home.html>
- **Virtual I/O Server FAQ Webpage**
  - <https://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/faq.html>
- **Virtual I/O Server Datasheet**
  - <https://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
- **Using the Virtual I/O Server Manual**
  - <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1kickoff.htm>
- **IVM User Guide**
  - [http://p505ivm.aixncc.uk.ibm.com/help.faces?page=guided\\_setup.jspETHERNET](http://p505ivm.aixncc.uk.ibm.com/help.faces?page=guided_setup.jspETHERNET)
- **IVM Help button on the User Interface**
  - <http://p5904.aixncc.uk.ibm.com/help.faces>
- **VIOS Server 1.5 Read me**
  - <https://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/installreadme.html>
- **Installing VIOS from HMC**
  - <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1installmanp5.htm>
- **VIOS Command Manuals**
  - [http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1\\_vios\\_commandslist.htm](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm)
  - [https://www-01.ibm.com/servers/resourcelink/lib03030.nsf/pages/VirtualIIOServerCommandReference/\\$file/sa76-0101.pdf](https://www-01.ibm.com/servers/resourcelink/lib03030.nsf/pages/VirtualIIOServerCommandReference/$file/sa76-0101.pdf)
- **VIO Operations Guide**
  - [https://www-01.ibm.com/servers/resourcelink/lib03030.nsf/pages/AdvancedPOWERVirtualizationOperationsGuide/\\$file/sa76-0100.pdf](https://www-01.ibm.com/servers/resourcelink/lib03030.nsf/pages/AdvancedPOWERVirtualizationOperationsGuide/$file/sa76-0100.pdf)

### Other good stuff

- **AIX6**
  - [http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=877&letternum=ENUSZP07-0385#Header\\_15](http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=877&letternum=ENUSZP07-0385#Header_15)
- **WPAR Manual**
  - [http://publib.boulder.ibm.com/infocenter/pseries/v6r1/index.jsp?topic=/com.ibm.aix.wpar/wpar-kickoff.htm&tocNode=int\\_20313](http://publib.boulder.ibm.com/infocenter/pseries/v6r1/index.jsp?topic=/com.ibm.aix.wpar/wpar-kickoff.htm&tocNode=int_20313)
- **WPAR Mgr**
  - <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS207-272&appname=USN>
- **Firefox 1.7 for AIX**
  - [https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=mozilla&S\\_PKG=dl\\_17\\_ww&S\\_TACT=104CBW71&lang=en\\_US&cp=UTF-8](https://www14.software.ibm.com/webapp/iwm/web/reg/download.do?source=mozilla&S_PKG=dl_17_ww&S_TACT=104CBW71&lang=en_US&cp=UTF-8)

# More good stuff

- **Fibre Channel marketing bulletins: IBM eServer pSeries and RS/6000**
  - ▶ [http://w3-1.ibm.com/sales/systems/portal/\\_s.155/254?navID=f220s240&geoID=AM&prodID=IBM%20eServer%20And%20TotalStorage%20Products&docID=rsrcsk.skit&docType=SalesKit&skCat=DocumentType](http://w3-1.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geoID=AM&prodID=IBM%20eServer%20And%20TotalStorage%20Products&docID=rsrcsk.skit&docType=SalesKit&skCat=DocumentType)
- **Workload Estimator**
  - ▶ <http://www-912.ibm.com/wle/EstimatorServlet>
- **Systems Agenda**
  - ▶ [http://www-03.ibm.com/systems/unix/agenda\\_unix\\_wp.pdf](http://www-03.ibm.com/systems/unix/agenda_unix_wp.pdf)
- **Virtual I/O Server**
  - ▶ <http://techsupport.services.ibm.com/server/vios/documentation/home.html>
- **Hitachi Best Practice Library: Guidelines for Installing Hitachi HiCommand dynamic Link Manager and Multi-Path I/O in a Virtual I/O Environment**
  - ▶ [http://www.hds.com/pdf/wp\\_hdlm\\_mpio\\_in\\_vio\\_TM-001-00.pdf](http://www.hds.com/pdf/wp_hdlm_mpio_in_vio_TM-001-00.pdf)
- **Introduction to the VIO Server**
  - ▶ [http://www-03.ibm.com/servers/eserver/pseries/hardware/whitepapers/virtual\\_io.pdf](http://www-03.ibm.com/servers/eserver/pseries/hardware/whitepapers/virtual_io.pdf)
- **Configuring Shared Ethernet Adapter Failover**
  - ▶ [http://techsupport.services.ibm.com/server/vios/documentation/SEA\\_final.pdf](http://techsupport.services.ibm.com/server/vios/documentation/SEA_final.pdf)



**End of  
Presentation**

# Special notices

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 inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Many of the pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: [http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux\\_pseries.html](http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html)

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.

Revised February 6, 2004

# Special notices (cont.)

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), alphaWorks, AS/400, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DFDSM, DirectTalk, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM, IBM(logo), ibm.com, IBM TotalStorage Proven, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, POWERparallel, PowerPC, PowerPC(logo), Predictive Failure Analysis, PS/2, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SP1, SP2, SpaceBall, System/390, The Engines of e-business, THINK, ThinkPad, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TrackPoint, TURBOWAYS, UltraNav, 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/L(logo), AIX 5L, AIX PVMe, AS/400e, BladeCenter, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, Domino, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, GigaProcessor, HACMP, HACMP/6000, Hypervisor, i5/OS, IBMLink, IMS, Intelligent Miner, Micro-Partitioning, iSeries, NUMACenter, ON DEMAND BUSINESS logo, OpenPower, POWER, Power Architecture, Power Everywhere, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, S/390 Parallel Enterprise Server, ThinkVision, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, X-Architecture, z/Architecture.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark 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 and/or other countries.

Intel, Itanium and Pentium are registered trademarks and Xeon and MMX are trademarks of Intel Corporation in the United States and/or other countries.

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 and/or other countries.

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.

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

Revised February 8, 2005

# Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM eServer p5, pSeries, OpenPower and IBM RS/6000 Performance Report at [http://www-1.ibm.com/servers/eserver/pseries/hardware/system\\_perf.html](http://www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.html)

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX V4.3 or AIX 5L. IBM C Set++ for AIX and IBM XL FORTRAN for AIX with optimization were the compilers used in the benchmark tests. The preprocessors used in some benchmark tests include KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX and MASS for AIX were also used in some benchmarks.

For a definition and explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	<a href="http://www.tpc.org">http://www.tpc.org</a>
SPEC	<a href="http://www.spec.org">http://www.spec.org</a>
Linpack	<a href="http://www.netlib.org/benchmark/performance.pdf">http://www.netlib.org/benchmark/performance.pdf</a>
Pro/E	<a href="http://www.proe.com">http://www.proe.com</a>
GPC	<a href="http://www.spec.org/gpc">http://www.spec.org/gpc</a>
NotesBench	<a href="http://www.notesbench.org">http://www.notesbench.org</a>
VolanoMark	<a href="http://www.volano.com">http://www.volano.com</a>
STREAM	<a href="http://www.cs.virginia.edu/stream/">http://www.cs.virginia.edu/stream/</a>
SAP	<a href="http://www.sap.com/benchmark/">http://www.sap.com/benchmark/</a>
Oracle Applications	<a href="http://www.oracle.com/apps_benchmark/">http://www.oracle.com/apps_benchmark/</a>
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	<a href="http://www.siebel.com/crm/performance_benchmark/index.shtml">http://www.siebel.com/crm/performance_benchmark/index.shtml</a>
Baan	<a href="http://www.ssaglobal.com">http://www.ssaglobal.com</a>
Microsoft Exchange	<a href="http://www.microsoft.com/exchange/evaluation/performance/default.asp">http://www.microsoft.com/exchange/evaluation/performance/default.asp</a>
Veritest	<a href="http://www.veritest.com/clients/reports">http://www.veritest.com/clients/reports</a>
Fluent	<a href="http://www.fluent.com/software/fluent/fl5bench/fullres.htm">http://www.fluent.com/software/fluent/fl5bench/fullres.htm</a>
TOP500 Supercomputers	<a href="http://www.top500.org/">http://www.top500.org/</a>
Ideas International	<a href="http://www.idesinternational.com/benchmark/bench.html">http://www.idesinternational.com/benchmark/bench.html</a>
Storage Performance Council	<a href="http://www.storageperformance.org/results">http://www.storageperformance.org/results</a>

Revised February 8, 2005

# Notes on performance estimates

## rPerf

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other pSeries systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

rPerf estimates are calculated based on systems with the latest levels of AIX 5L and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM @server pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

Revised June 28, 2004