PowerVM for IBM i

Kris Whitney
IBM i Base HW, Virtualization and Storage Architect

Agenda

- PowerVM Offering
- SR-IOV
- VIOS Storage
- Live Partition Mobility, PowerVC for IBM i, Cloud

PowerVM Editions are tailored to client needs

PowerVM Editions offer a unified virtualization solution for all Power workloads

- **PowerVM Standard Edition**
  - *Production deployments*
  - *Server consolidation*
- **PowerVM Enterprise Edition**
  - *Multi-server deployments*
  - *Cloud infrastructure*
- **PowerVM for IBM PowerLinux Edition**
  - *Same Function as PowerVM EE*
  - *Restricted to Linux VMs only*

PowerVM Editions	**Standard**	**Enterprise**
Concurrent VMs	**20 per core** (up to 1000)**	**20 per core** (up to 1000)**
Virtual I/O Server	✓ ✓	✓ ✓
NPIV	✓	✓
Suspend/Resume	✓	✓
Shared Processor Pools	✓	✓
Shared Storage Pools	✓	✓
Thin Provisioning	✓	✓
Live Partition Mobility		✓
Active Memory Sharing		✓
PowerVP*		✓

AIX for Business Linux

PowerVM

- ** *Requires eFW7.6 or higher*
- * *Requires eFW7.7 or higher*

PowerVM v2.2.4

Virtualization without Limits

PowerVM

- ✓ ***Reduces IT infrastructure costs***

  Consolidate diverse workloads save operational costs

- ✓ ***Improves service levels***

  Virtualized resources can be applied dynamically to workloads as needed

- ✓ ***Manages risk***

  Unrivaled flexibility enables rapid response to business change minimizing risk

Announce – 10/5/2015 GA – 12/04/2015

- ✓ *Direct OpenStack Enablement*
- ✓ *Shared Storage Pool Enhancements*
- ✓ *Mobility for SRIOV Adapters*
- ✓ *Simplified Management*

- **New!! PowerVM NovaLink Architecture**
  - Allows direct OpenStack connection to PowerVM host
  - Improves Cloud scalability and OpenStack Community adoption of PowerVM drivers
  - Initially for POWER8 servers not managed by a HMC; Tech preview for HMC managed POWER8 systems
- **vNIC Adapters –** Enables VM Mobility for VMs with SRIOV Adapters and improved performance
- **Live Partition Mobility Improvements**
  - Better NPIV Validation, Improved performance, Allow LPM when one VIOS has failed, Improved Resiliency, Target vswitch can now be selected
- **Shared Storage Pools Improvements**
  - New Storage Tiers within a storage pool. Up to 10 Tiers provide better control for performance and segregation of data
  - Ability to Grow a Virtual Disk

PowerVM v2.2.5

Virtualization without Limits

PowerVM

- ✓ *NovaLink Enhancements*
- ✓ *I/O Improvements*
- ✓ *Resiliency Improvements*

✓ *Reduces IT infrastructure costs*

Consolidate diverse workloads save operational costs

✓ *Improves service levels*

Virtualized resources can be applied dynamically to workloads as needed

✓ *Manages risk*

Unrivaled flexibility enables rapid response to business change minimizing risk

Announce – 10/11/16

- **PowerVM GA 10/21/2016**
  - **vNIC Failover –>** improves availability for SR-IOV Configurations
  - **Shared Ethernet Adapter(SEA) Improvements –>** Improves Performance and Resiliency
    Large Send Performance improvements for Linux & IBM i, Failover Health checks, Failback Controls
  - **Live Partition Mobility Improvements –>** Redundant Multiple Data Movers increase LPM Reliability and Performance
  - **Shared Storage Pools Improvements –>**Scalability and Resiliency Improvements Support Larger More reliable SSPs
    32 Nodes Supported in SSP
- **PowerVM NovaLink Enhancements v1.0.0.5 GA 12/16/16**
  - NovaLink partition can run on Red Hat as well as Ubuntu Linux -> Provides more options for Clients
  - SR-IOV Support -> Improves I/O Performance and Quality of Service options
  - Software Defined Networking Tech Preview(Open vSwitch) -> Network Overlays enable Cloud Deployments
- **Firmware v860 GA 10/28/16**
  - Reduction of Hypervisor Memory Usage through large page Usage -> Especially helpful for SAP HANA Workloads
- **HMC v860 GA 10/28/16**
  - Ability to export Performance data to csv format
  - Dynamic Setting of Simplified Remote Restart property
  - Reporting on Energy Consumption via Rest APIs and export facility

SR-IOV

I/O Virtualization on POWER

IO Bus Virtualization with Dedicated Adapters

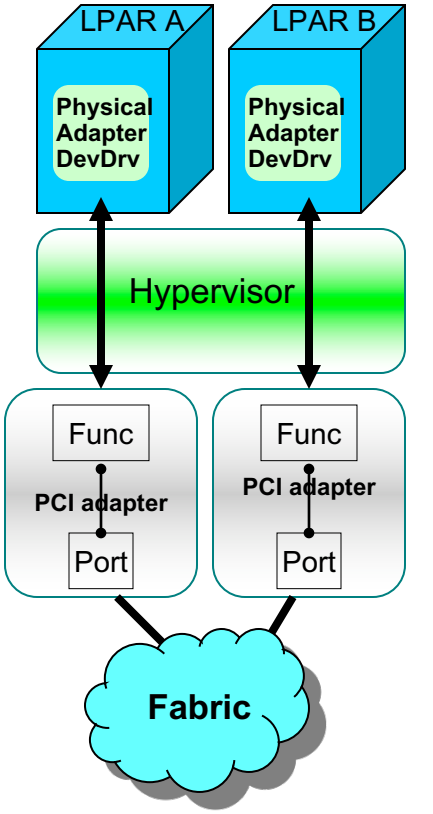

Increasing Adapter BW & LPAR Density per Slot

IO Adapter Virtualization with VIO Server

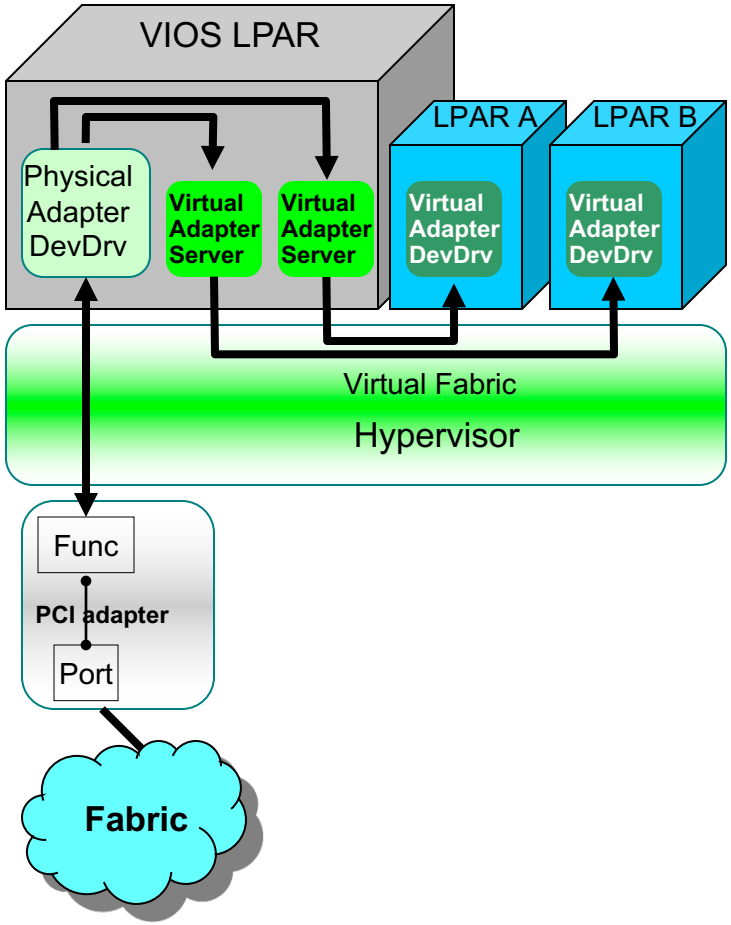

Power Systems SR-IOV Solution

Features

- **Adapter sharing**
  - Improves partition to I/O slot ratio
  - Sharing by up to 48 partitions per adapter. Additional partitions with Virtual I/O Server (VIOS)
- **Direct access I/O**
  - Provides CPU utilization and latency roughly equivalent to dedicated adapters
  - Adapter sharing with advanced features such as Receive Side Scaling (RSS) and adapter offloads.
- **Adapter resource provisioning (QoS)**
  - User designates desired capacity for a logical port.
- **Simple server I/O deployment**
  - Minimal steps to add a logical port to partition or partition profile.
  - Integrated solution (i.e. common UI, tools, etc.)
- **Flexible deployment models**
  - Single partition
  - Multi-partition without VIOS
  - Multi-partition thru VIOS
  - Multi-partition mix of VIOS and nonVIOS

IO Adapter Virtualization with SR-IOV

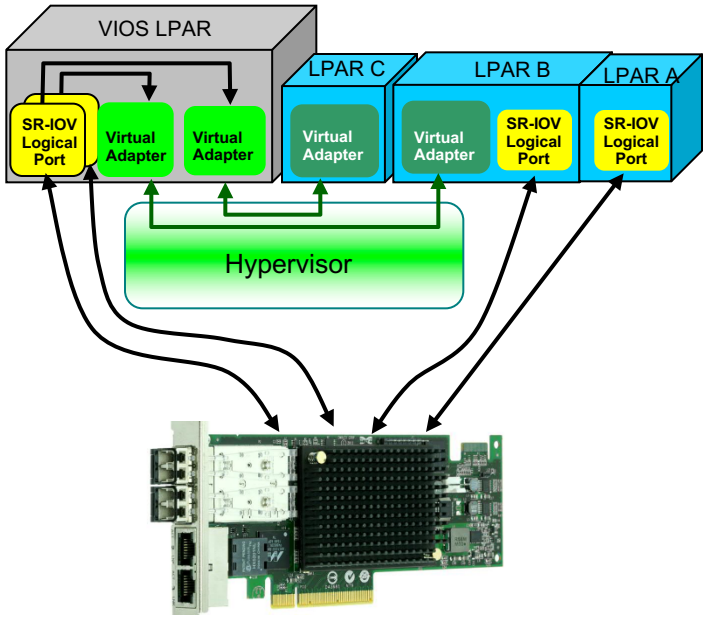

Flexible Deployment

- **Single partition**
  - All adapter resources available to a single partition

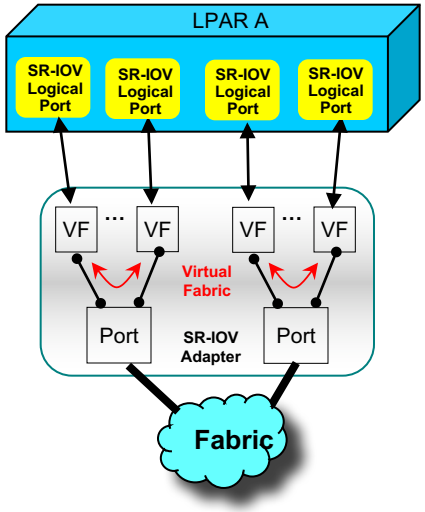

- **Multi-partition without VIOS**
  - Direct access to adapter features
  - Capacity per logical port
  - Fewer adapters for redundant adapter configurations.

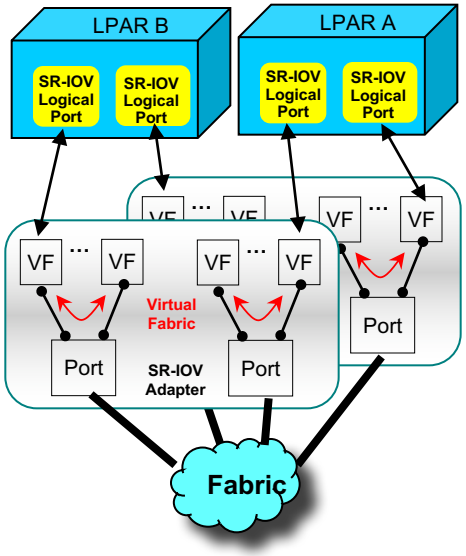

Flexible Deployment

- **Multi-partition thru VIOS**
  - Adapters shared by VIOS partitions
  - Fewer adapters for redundancy
  - VIOS client partitions eligible for Live Partition Mobility
  - Allows class of service between VIOS clients

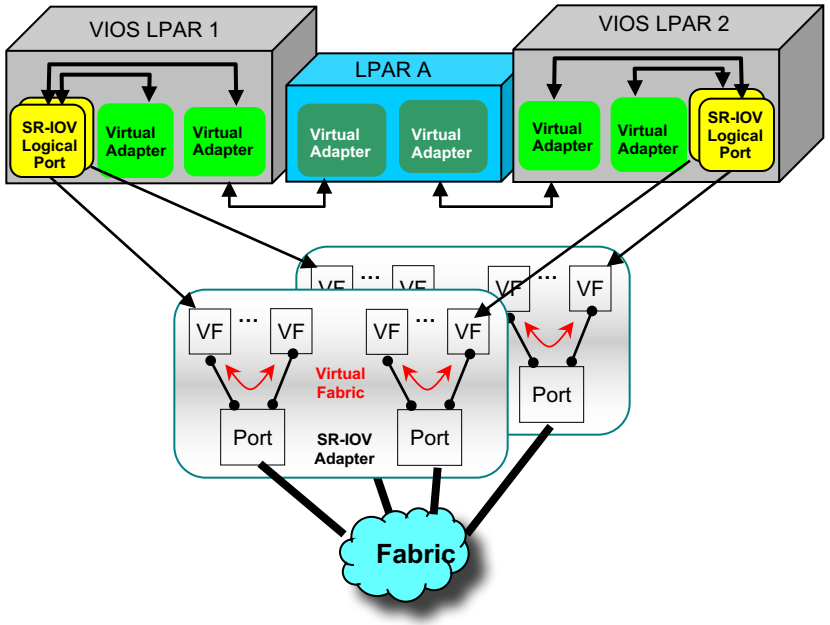

- **Multi-partition mix of VIOS and nonVIOS**
  - For VIOS partitions same as Multi-partition thru VIOS above
  - Direct access partitions
    - Path length & latency comparable to dedicated adapter
    - Direct access to adapter features
    - Entitled capacity per logical port

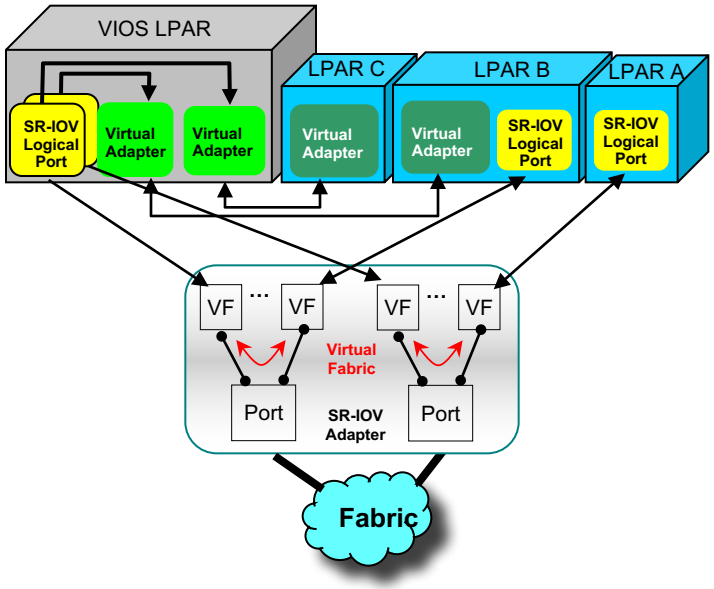

Performance - SR-IOV and VIOS/SEA

- **Shared Ethernet Adapter Bridging**
  - Traffic flows between the physical adapter and client partition through the hypervisor and VIOS partition
  - Within the VIOS the traffic flows between the virtual adapter driver and physical adapter driver through the Shared Ethernet Adapter bridge support.
  - Latency and CPU utilization overhead
    - Hypervisor copies packets between LPAR and VIOS
    - VIOS Shared Ethernet Adapter bridge function and adapter drivers
  - For a 10Gbps link the maximum observed throughput for a single virtual adapter is about 2.8 Gbps.

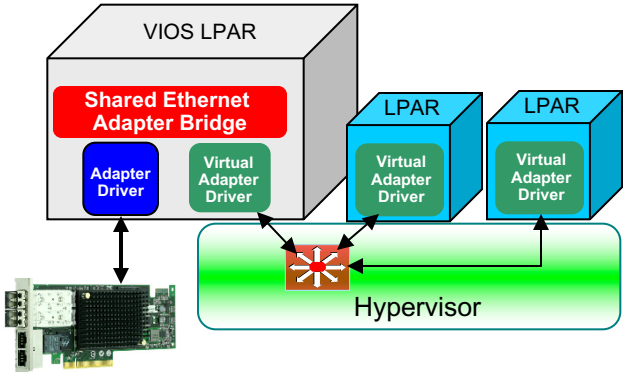

VIOS Shared Ethernet Adapter Bridging

- **SR-IOV Direct Access Adapter Sharing**
  - LPAR has direct access to the adapter
  - Latency and CPU utilization on par with adapter dedicated to an LPAR
  - For a 10Gbps link the throughput for a single logical port (VF) is about 9.1Gbps (line rate)

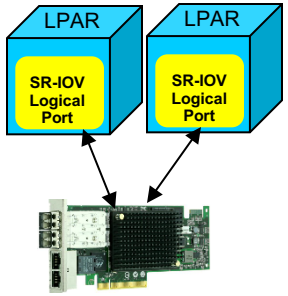

SR-IOV Direct Access Adapter Sharing

Power Systems SR-IOV vNIC Solution

Virtual I/O Enhancements with SR-IOV

- One-to-one relationship between client partition virtual adapter and adapter VF
- Performance Optimized
  - Lower latency and CPU utilization
  - Data flows between client partition memory and adapter (i.e. eliminates data copies)
  - Leverages adapter offload capability
    - Multiplex/demultiplex of I/O operations
    - Adapter switch for partition to partition communication
- Extends VF QoS capability to client partitions
- Client partitions eligible for advanced virtualization features (e.g. LPM, VM Mirror)

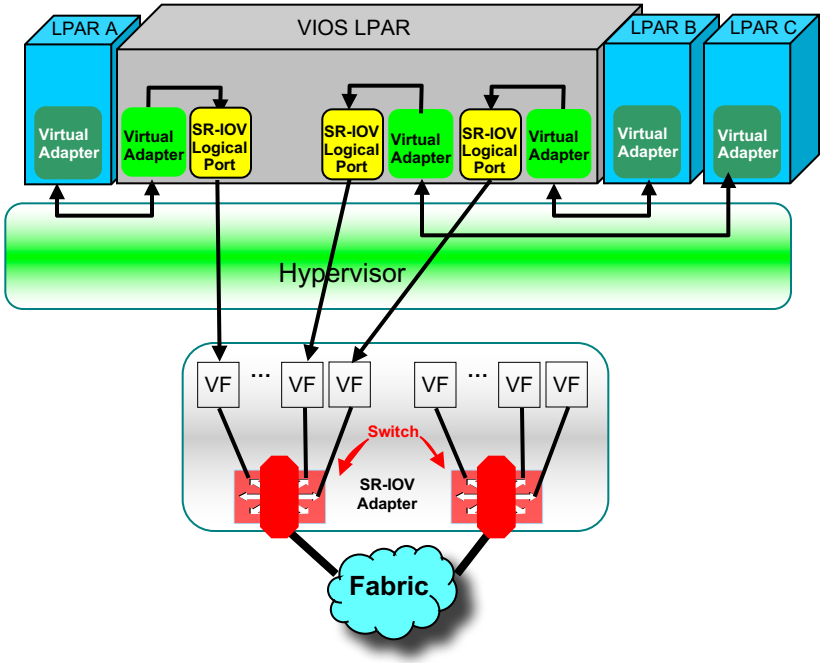

VF = Virtual Function

Power Systems SR-IOV vNIC Failover Solution

Virtual I/O Enhancements with SR-IOV

- One-to-one relationship between client partition virtual adapter and adapter VF
- Performance Optimized
  - Lower latency and CPU utilization
  - Data flows between client partition memory and adapter (i.e. eliminates data copies)
  - Leverages adapter offload capability
    - Multiplex/demultiplex of I/O operations
    - Adapter switch for partition to partition communication
- Extends VF QoS capability to client partitions
- Client partitions eligible for advanced virtualization features (e.g. LPM, VM Mirror)
- Able to failover to up to 5 additional ports.

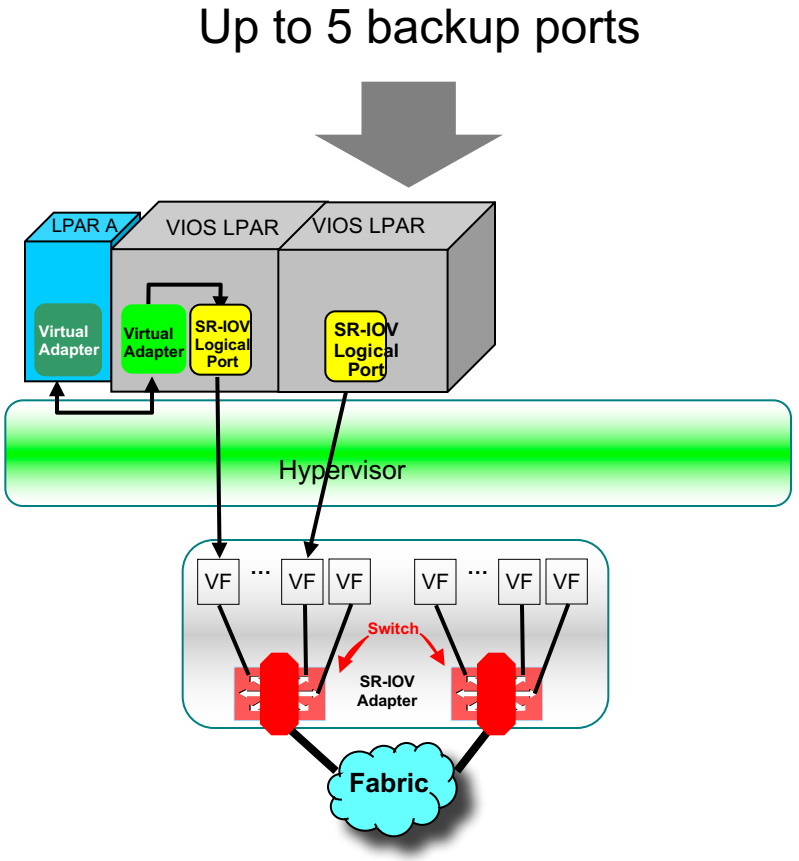

VF = Virtual Function

Gaining Momentum - SR-IOV Announcements

- Feb 2015 - IBM announces that SR-IOV NIC is now supported on the Power E870 and E880 (9119-MME and 9119-MHE) system enclosures when placed in the system unit.
  - PCIe2 LP 4-port (10Gb FCoE and 1GbE) SR&RJ45 Adapter (#EN0L)
  - PCIe2 LP 4-port (10Gb FCoE and 1GbE) SFP+Copper and RJ45 Adapter (#EN0J)
  - These are Low Profile (LP) adapters that are equivalent to EN0H and EN0K
    - LP is required for E870/E880 CEC slots

- May 2015 – IBM fulfills remaining SOD plans for SR-IOV support on POWER8 with firmware fw830 plus additional adapters
  - SR-IOV capability is now available on the entire POWER8 server line in the system unit as well as placement in the PCIe IO drawer
  - New SR-IOV capable adapters
    - EN15 & EN16: PCIe Gen3 - 4 ports 10GBASE-SR (10Gbs OPTICAL-SR)
    - EN17 & EN18: PCIe Gen3 - 4ports 10GSFP+Cu (10Gbs SFP+ TWINAX)
    - Max of 64 partitions per adapter

Adapter	Total # of Ports	Port Physical Interface	Max VFs Per Port	Max VFs Per Adapter
(EN0L and EN0J) SR & Cu – LP	**4**	**2 x 10G FCoE SR or Copper**	**20**	**48**
		2 x 1GbaseT	**4**	
(EN15 and EN16) **Full Height or LP**	**4**	**4 x 10Gbs Optical-SR**	**16**	**64**
(EN17 and EN18) **Full Height or LP**	**4**	**4 x 10Gbs SFP+CU (TWINAX)**	**16**	**64**

SR-IOV Software Support

Software	SR-IOV Support
AIX	AIX 6.1 TL9 SP5 and APAR IV68443 or later AIX 7.1 TL3 SP5 and APAR IV68444 or later AIX 7.1 TL2 SP7 or later (planned availability 3Q 2015) AIX 6.1 TL8 SP7 or later (planned 3Q 2015)
IBM i	IBM i 7.1 TR10 or later IBM i 7.2 TR2 or later
Red Hat	Red Hat Enterprise Linux 6.6 or later Red Hat Enterprise Linux 7.1, big endian, or later Red Hat Enterprise Linux 7.1, little endian, or later
SUSE	SUSE Linux Enterprise Server 12 or later
Ubuntu	Ubuntu 15.04 or later
PowerVM	Firmware 830 available June, 2015 and HMC V8.830

SR-IOV Hardware Support

Model	SR-IOV Mode Support
S814	Slots C6, C7, C10, C12
S822	Slots C2, C3, C5, C6, C7, C10, C12 with both sockets populated
S824	Slots C2, C3, C4, C5, C6, C7, C10, C12 with both sockets populated
S812L	Slots C6, C7, C10, C12
S822L	Slots C2, C3, C5, C6, C7, C10, C12 with both sockets populated
S824L	Slots C2, C3, C4, C5, C6, C7, C10, C12
E850	All internal slots
E870	All internal slots
E880	All internal slots
I/O Drawer	Slots C1 and C4 of the 6-slot fan-out module

Storage Virtualization

What is the VIOS?

- A special purpose appliance partition
  - Provide I/O virtualization
  - Advanced Partition Virtualization enabler
- First GAed 2004
- Built on top of AIX, but not an AIX partition
- IBM i first attached to VIOS in 2008 with the IBM i 6.1
- VIOS is licensed with PowerVM

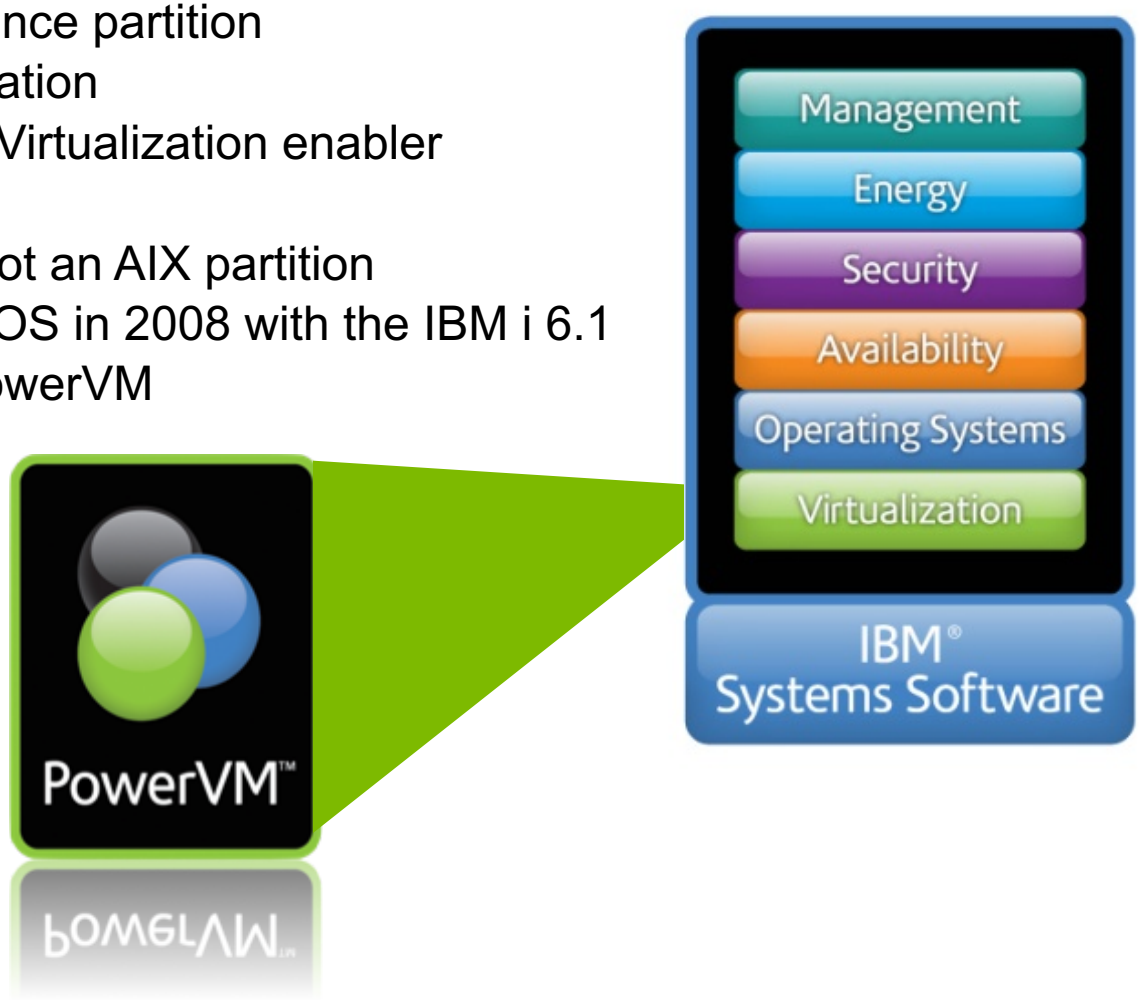

IBM i + VSCSI (Classic)

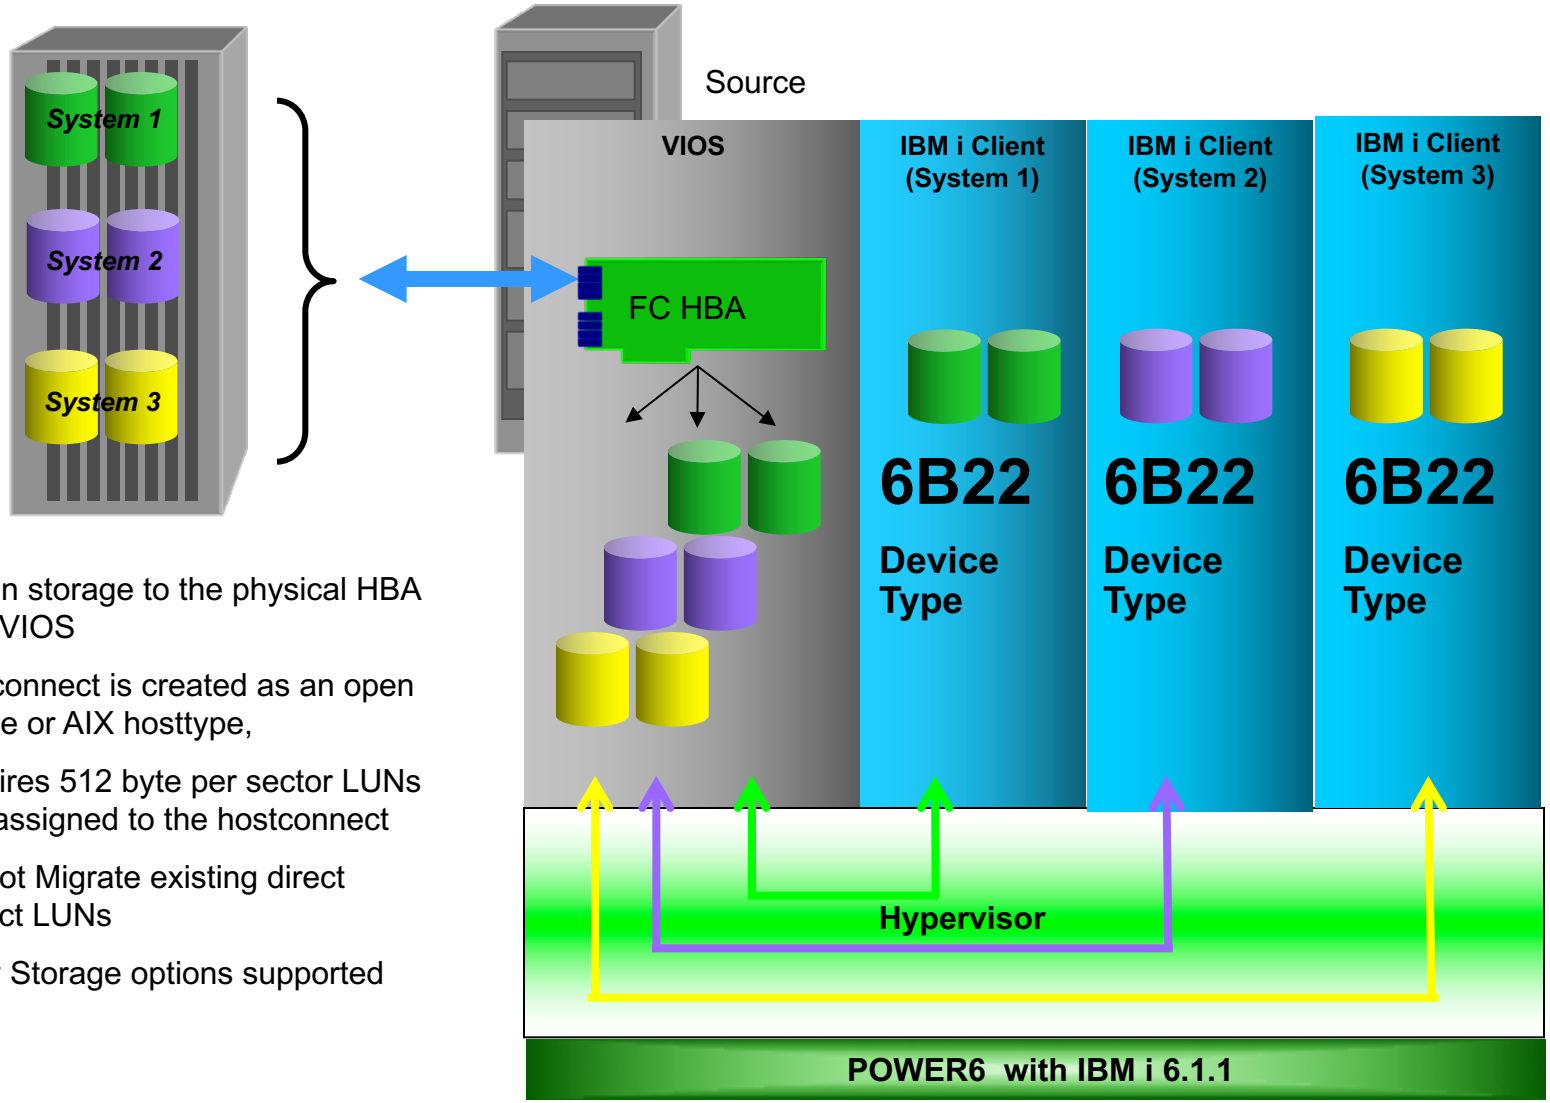

- Assign storage to the physical HBA in the VIOS
- Hostconnect is created as an open storage or AIX hosttype,
- Requires 512 byte per sector LUNs to be assigned to the hostconnect
- Cannot Migrate existing direct connect LUNs
- Many Storage options supported

IBM i + NPIV (Virtual Fiber Chanel)

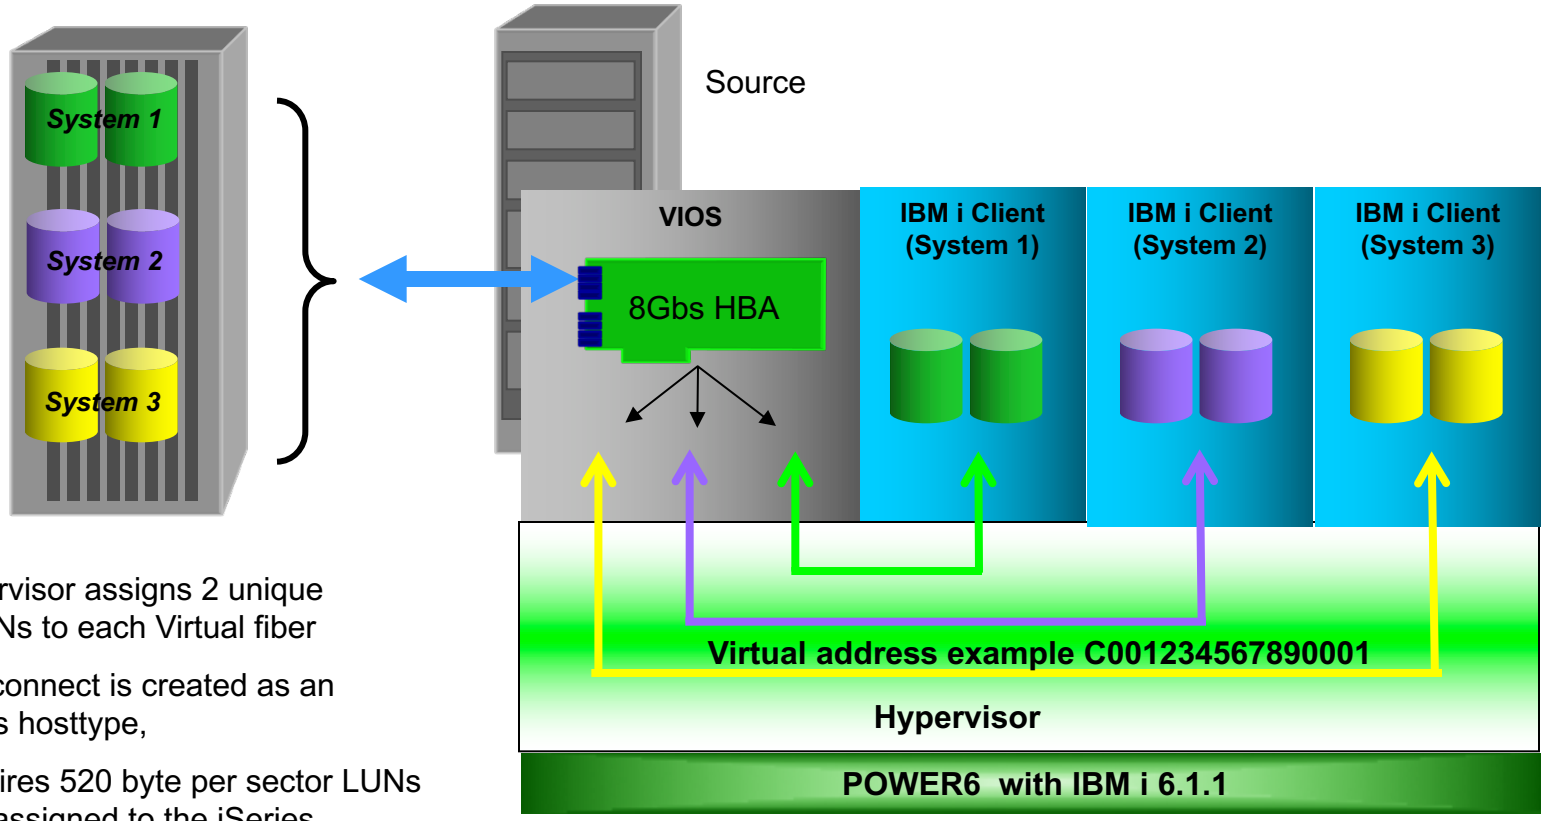

- Hypervisor assigns 2 unique WWPNs to each Virtual fiber
- Hostconnect is created as an iSeries hosttype,
- Requires 520 byte per sector LUNs to be assigned to the iSeries hostconnect on DS8K
- Can Migrate existing direct connect LUNS
- DS8100, DS8300, DS8700, DS8800, DS5100 and DS5300 SVC, V7000, V3700 supported

Note: an NPIV (N_port) capable switch is required to connect the VIOS to the DS8000 to use virtual fiber.

NPIV Configuration - Server Adapter Mappings

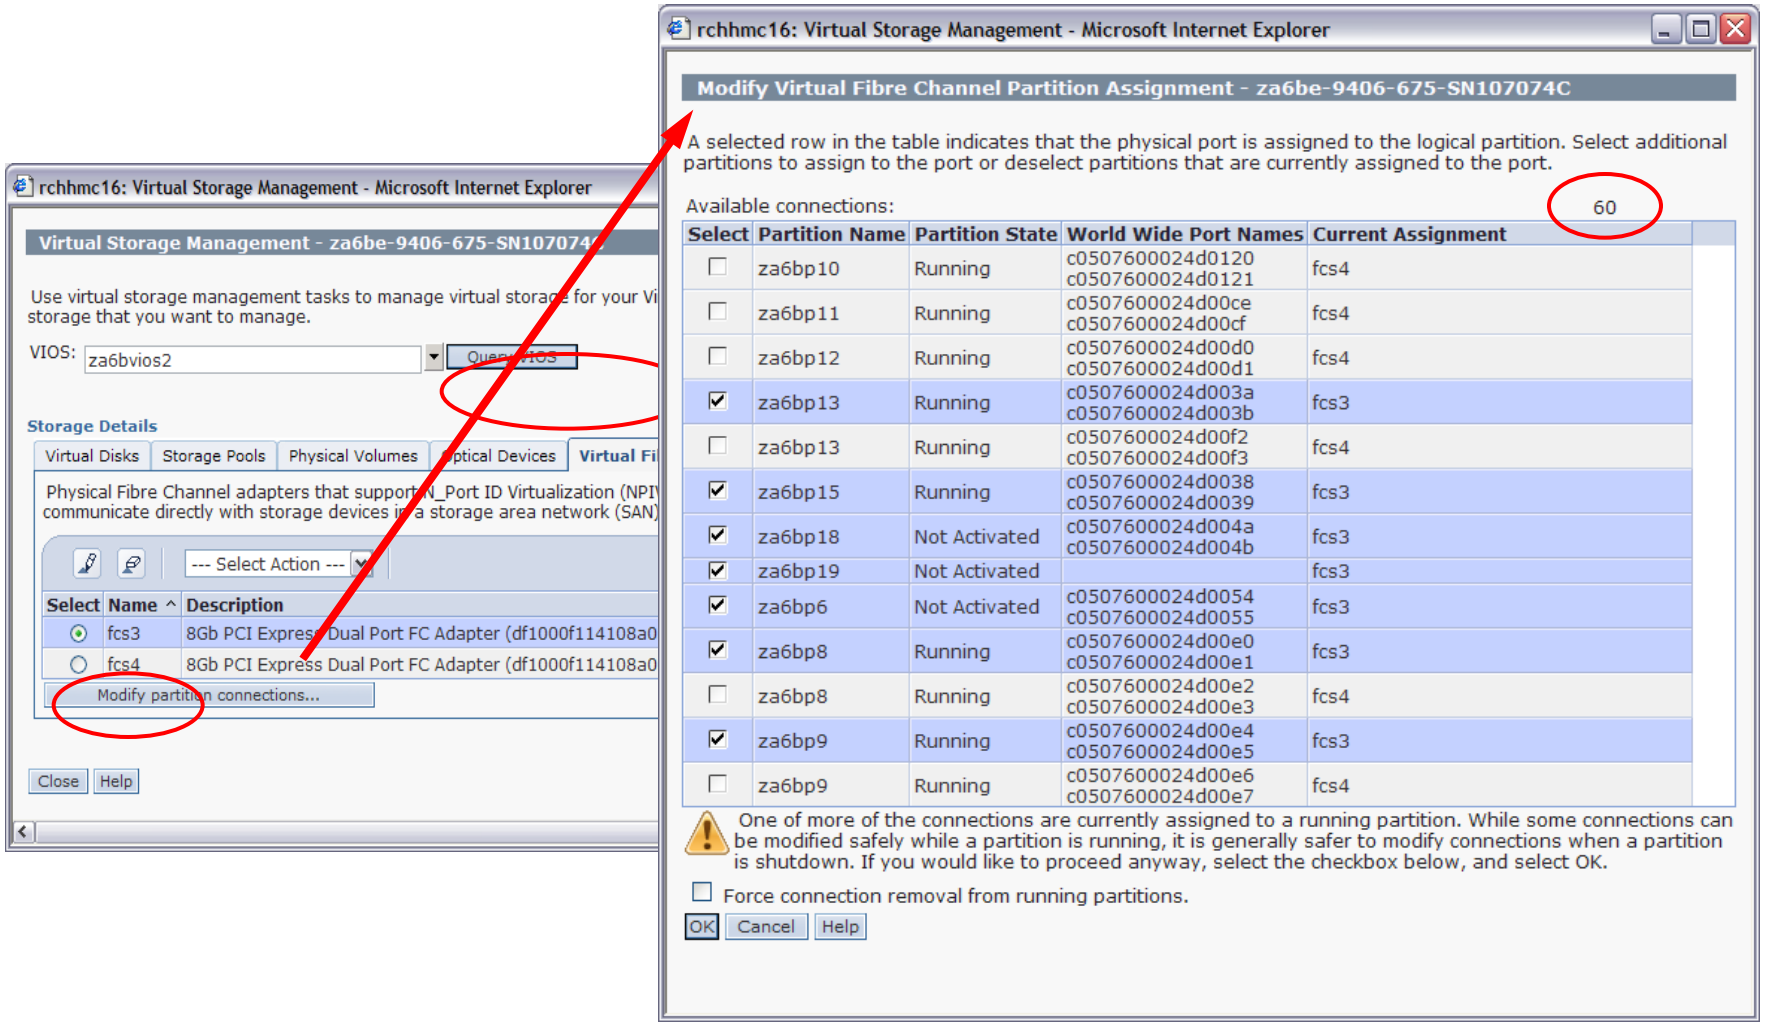

Virtualizing disk storage with IBM i or VIOS

- Single IBM i or VIOS host provides access to SAN or internal storage
  - AIX, IBM i, or Linux client partitions
  - Protect data via RAID-5, RAID-6, or RAID-10

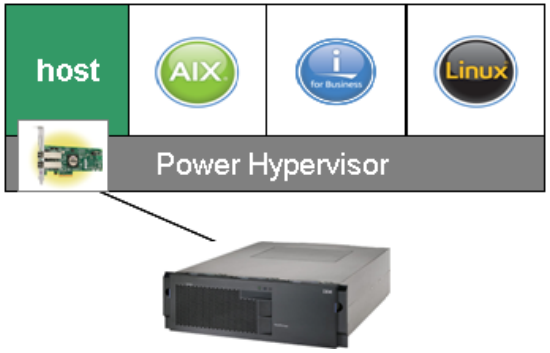

- **Redundant** IBM i or VIOS hosts provide access SAN or internal storage
  - AIX, IBM i, and Linux client partitions
  - Client LPAR protects data via mirroring
  - Two sets of disk and adapters

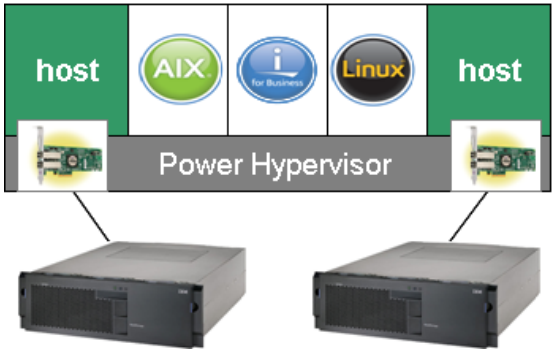

- **Redundant** VIOS hosts multiple paths to attached SAN storage with MPIO
  - AIX, IBM i, and Linux client partitions
  - One set of disk

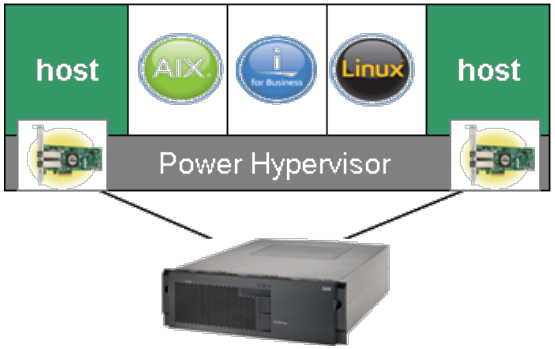

Dual VIOS attaching to EXP24S with up to 24 disks Single IBM i partition

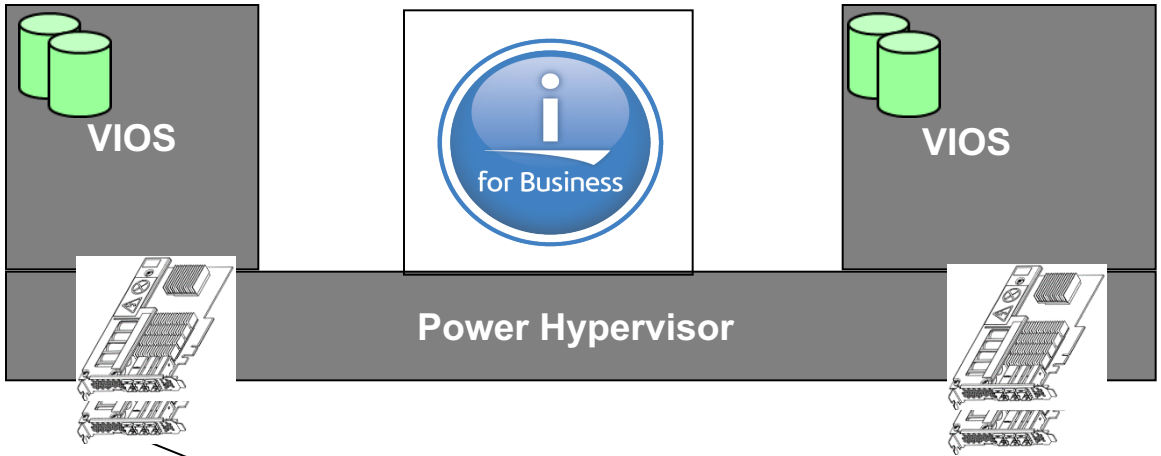

2 x EJ0L SAS controllers

2 x EJ0L SAS controllers

EXP24S I/O drawer
(MODE 2) with up to
24 disks

IBM i Single Page (4K) I/O with different Sector Sizes

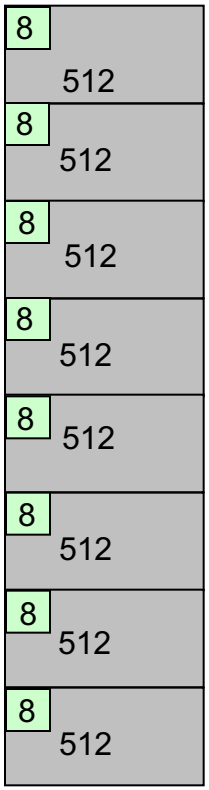

8 x 520 Byte

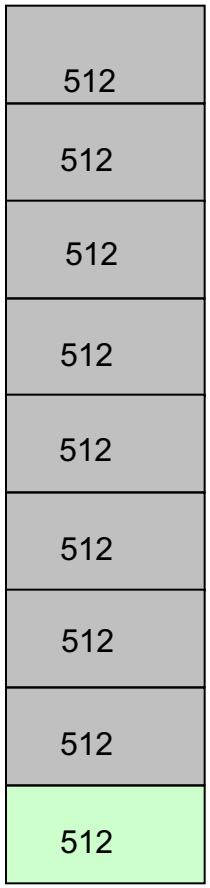

← Header

9 x 512 Byte

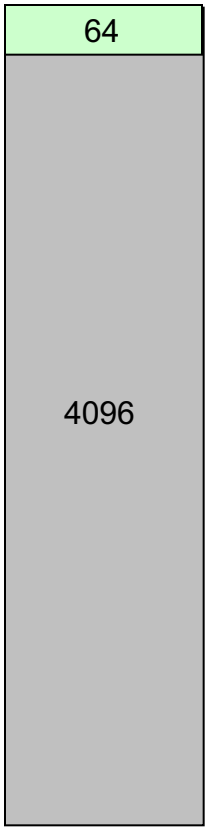

1 x 4160 Byte Sector

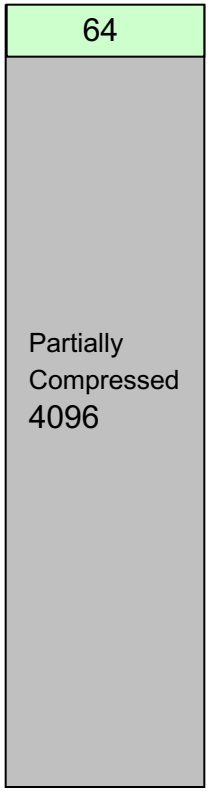

IBM i 7.2 and later

1 x 4096 Byte Sector

VIOS Configuration

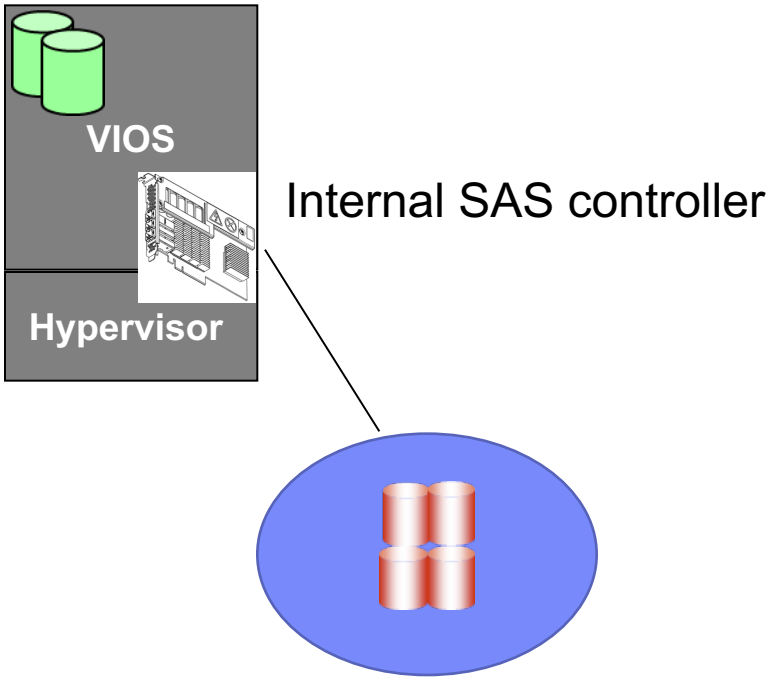

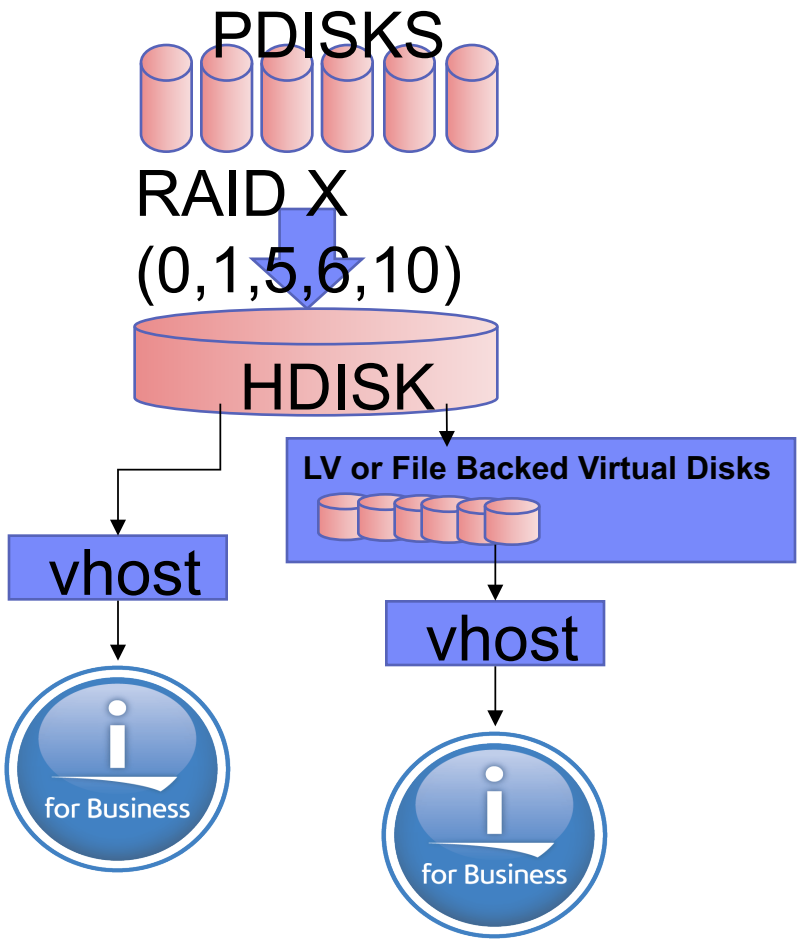

Configuring Raid for Performance with 5XX sector Drives

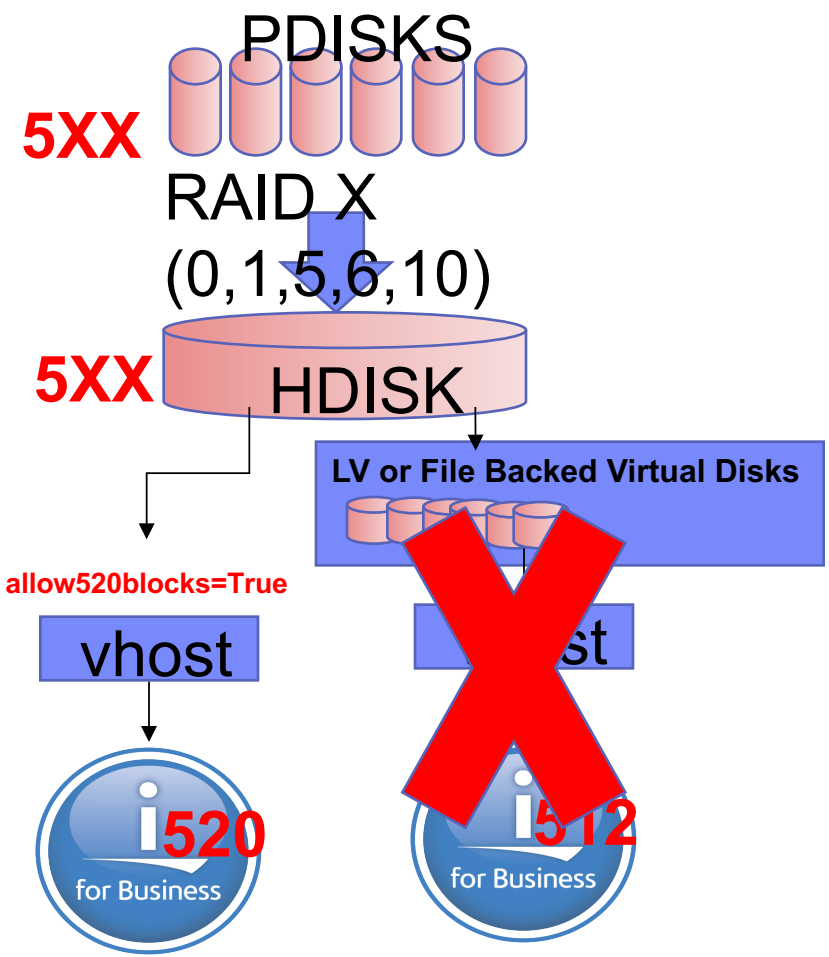

- All PCIe2/PCIe3 SAS adapters (e.g. FC 5913, ESA3, and EJ0L), have been enhanced with hardware acceleration that is optimized on 8 sector I/O boundaries. When IBM i attaches a 512 byte sector device, its I/O is on 9 sector boundaries. Using one of these newer adapters with IBM i using 512 byte sector storage will cause significant write performance degradation verses what the adapter is capable of doing.
- In order to use these adapters in a VIOS environment, you must virtualize the entire hdisk to the IBM i lpar.

Dual VIOS, Raid 0 Configuration

- Due to performance optimizations in the OS and DB2, IBM i still prefers seeing multiple LUNs.
- In order to maximize the number of LUNs seen by IBM i and still have access to 520 byte sectors RAID 0 is the best solution.
- VIOS rootvg could be installed on a RAID10 array or use mirrorios to have redundancy for the VIOS LPAR.

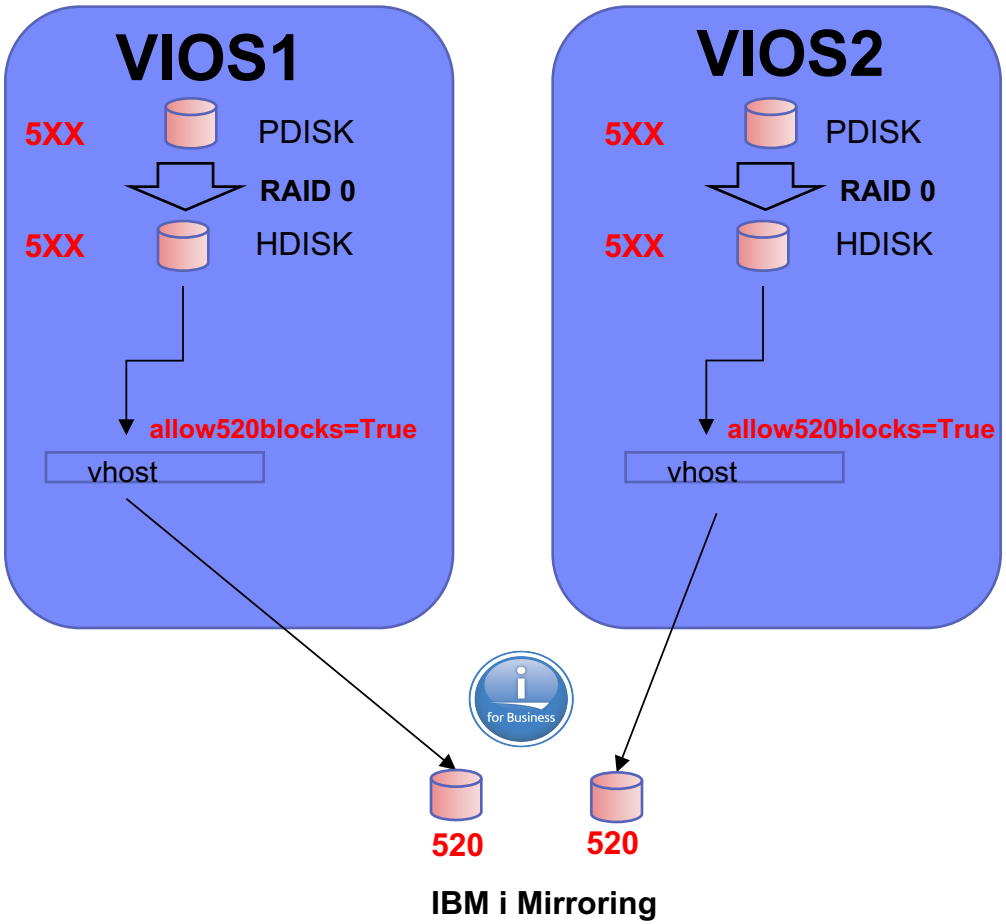

Configuring Raid for Optimal Performance with 4K Drives

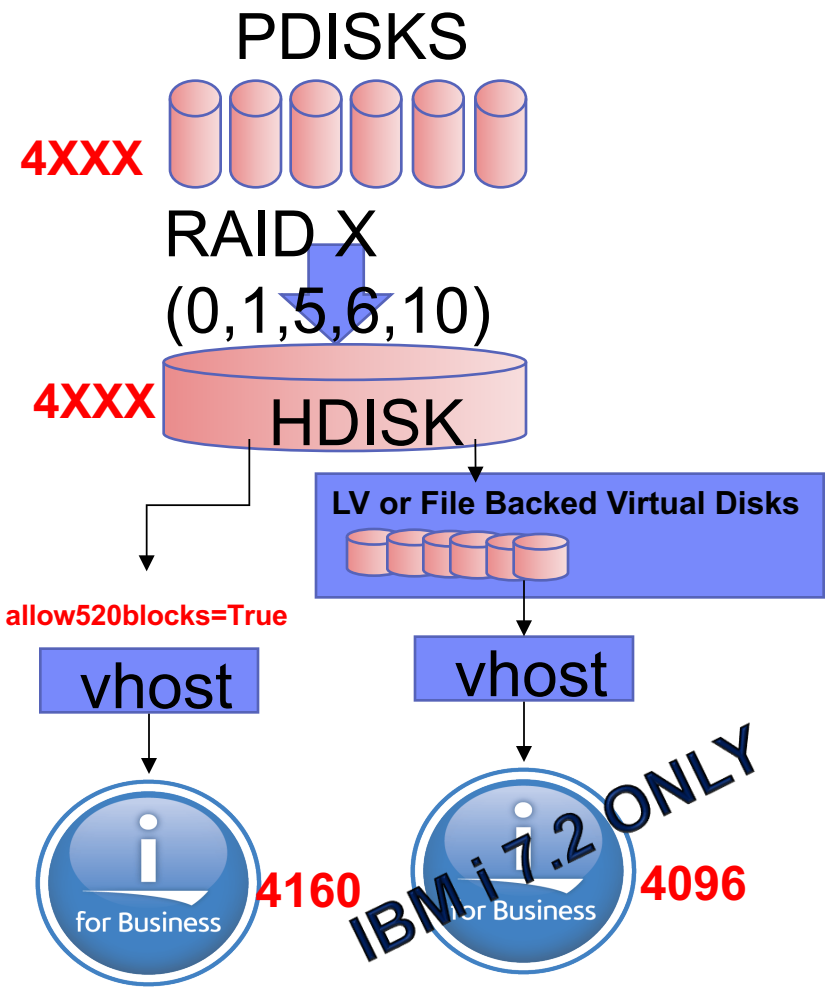

- Power8 now supports the usage of 4K drives.
- For IBM i 7.1, the same recommendation of using Raid0 and directly mapping the hdisk to the vhost still applies. This configuration will give you the best performance.
- If you have all 4K drives in a volume group and use LV or File backed virtual disks with IBM i 7.2, you can utilize new function in IBM i 7.2 to attach LV and File Backed virtual disks and still be aligned on the adapter hardware boundaries.

Dual VIOS Using Split Backplane Summary

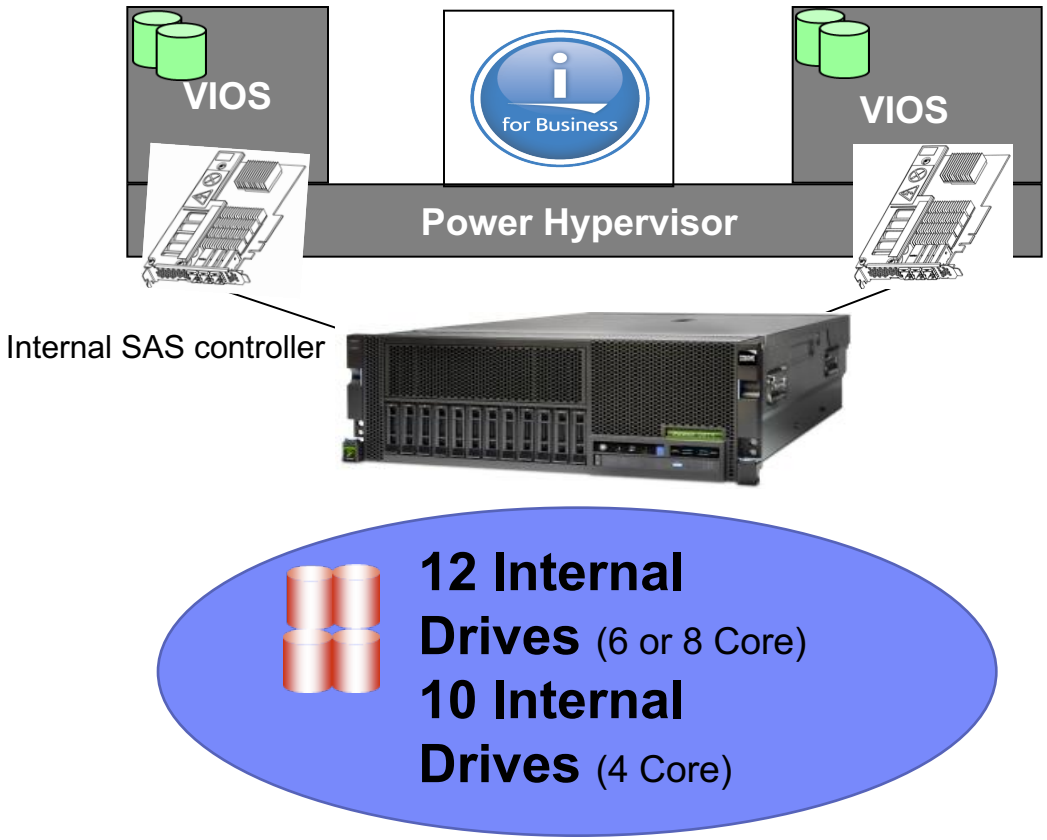

- When using the split backplane on Power8 you do not have any write cache on the adapter which typically severely impacts IBM i workloads. SSDs are highly recommended with these adapters.
- If you use VIOS with these adapters and virtualize the disk to IBM i, you should not use LV or File Backed virtual disks for IBM i*.
- *LV or File Backed may be used with IBM i 7.2 if the physical disks are 4K sector drives.

Dual VIOS Using Split Backplane with an EXP24S

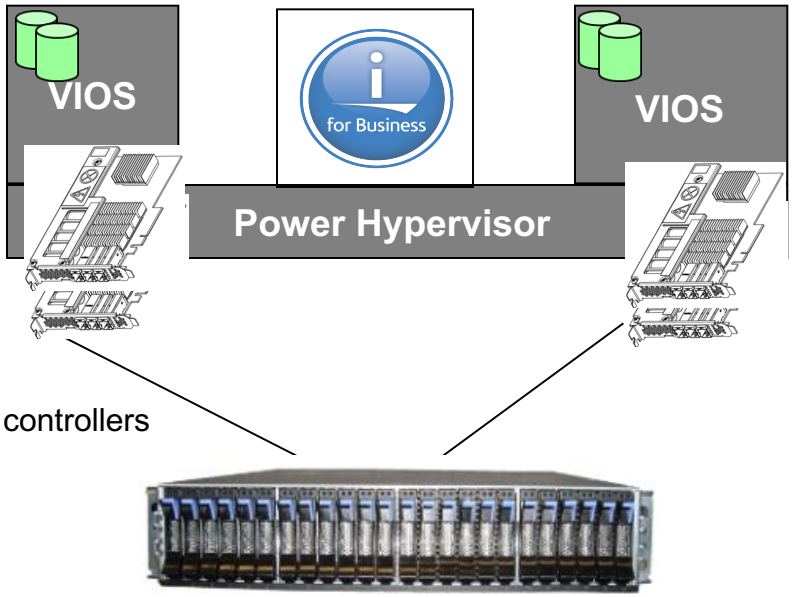

EXP24S I/O drawer (MODE 2) with up to 24 disks

- If you use an expansion drawer you can use the I/O Adapters with cache and split the EXP24S between the VIOSes. This gives much better performance than the split backplane of the system adapter without cache.
- If you use VIOS with these adapters and virtualize the disk to IBM i, you should not use LV or File Backed virtual disks for IBM i*.
- *LV or File Backed may be used with IBM i 7.2 if all the physical disks in the Volume Group are 4K sector drives.

Alternative to Using VIOS

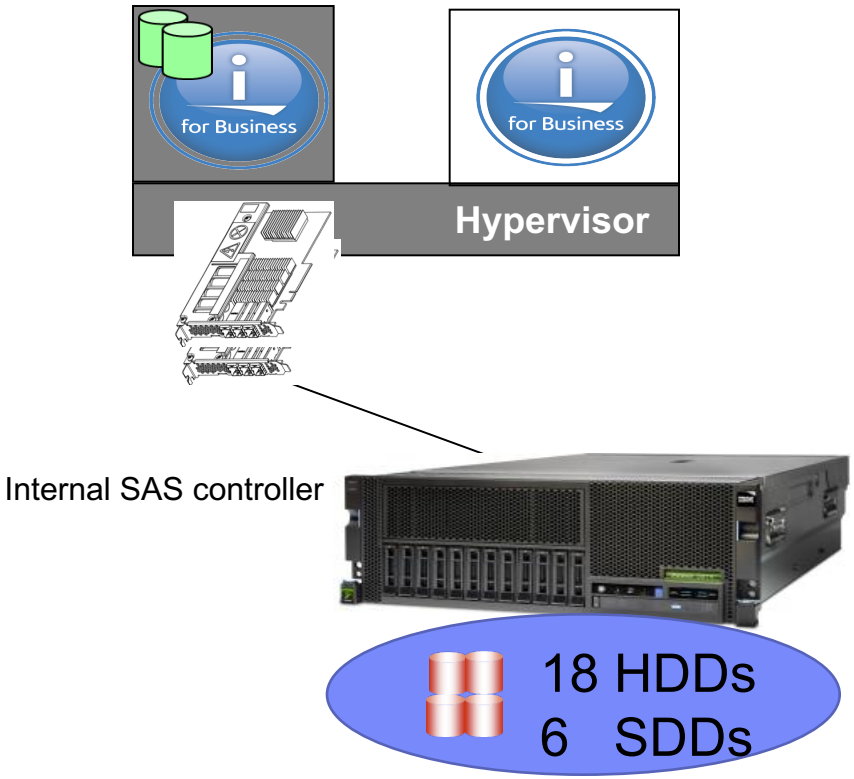

- IBM i can virtualize disk, optical, tape and Ethernet to another IBM i LPAR.
- SAS adapter performance degradation with IBM i virtualization is significantly less than with the wrong configuration on VIOS.
- Ability to create virtual disks of any size
- Ability to set SSD preference
- Ability to use PowerHA to replicate Virtual Disks on the Server
- Ability to use 4096 sector virtual disks
  - New Parameter on CRTNWSSTG in IBM i 7.2
  - Independent of the physical drive format

Redundant IBM i I/O virtualization servers

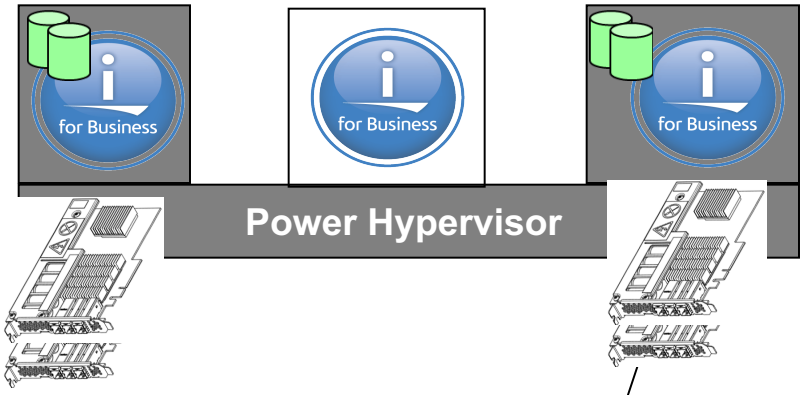

18
Drives
6 SSDs

EXP24S I/O drawer with 24 disks

A combination of the high function raid adapter in the system and expansion can be used.

If you use the split backplane without cache it is highly recommended to use SSDs for performance on IBM i.

12 Internal Drives
6 OR 8 Core
10 Internal Drives
4 core

Virtualization Comparison

	iVirtualization	VIOS with SAS	VIOS with External Storage
512 Byte Performance	+	–	+
520 Byte Performance	+	+	+
4160 Byte Performance	+	+	+
4096 Byte Performance (7.2)	+	+	+
Live Partition Mobility	–	–	+
PowerVC Support	–	–	+
MultiPath I/O	–	–	+

More Information

- Document on how to configure the VIOS
  - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20i%20Technology%20Updates/page/SAS%20Adapter%20Performance%20Boost%20with%20VIOS
- More Information on storage virtualization with IBM I
  - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/beb2d3aa-565f-41f2-b8ed-55a791b93f4f/page/IBM%20i%20Virtualization%20and%20Open%20Storage

It’s not really the size, but the number!

Even though the external storage array has many drive, IBM i still needs multiple LUNs to perform well.

- Database optimizes to the number of LUNs IBM i can see.
- IBM i Storage Management is optimized to scale with the number of LUNs
- Journal can be impacted by the number of LUNs

Don’t mix capacities in the same ASP

Take caution when increasing LUN size and dramatically reducing the quantity of LUNs.

Dynamically increasing the LUN size is not supported on IBM i !

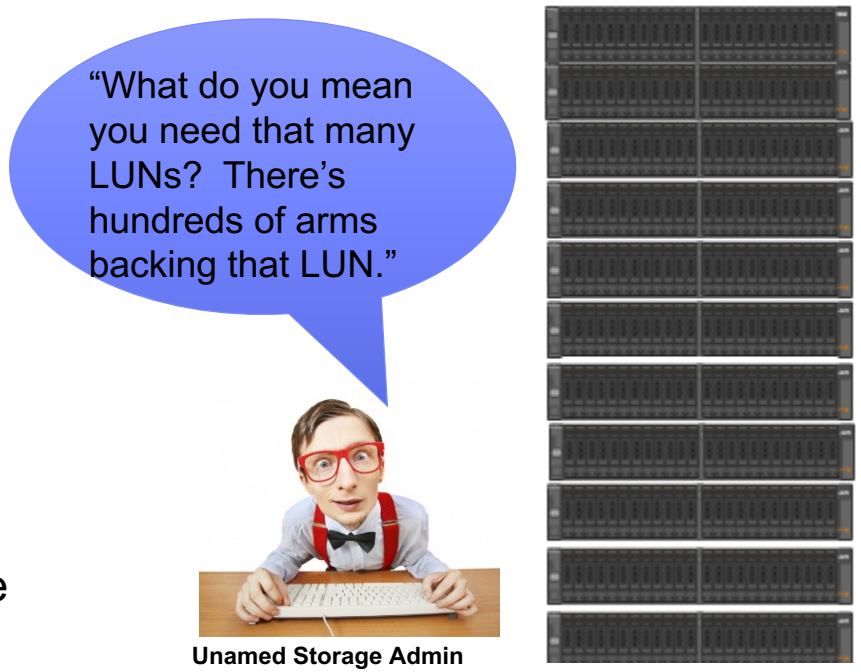

No true multipath support for tape in prior releases

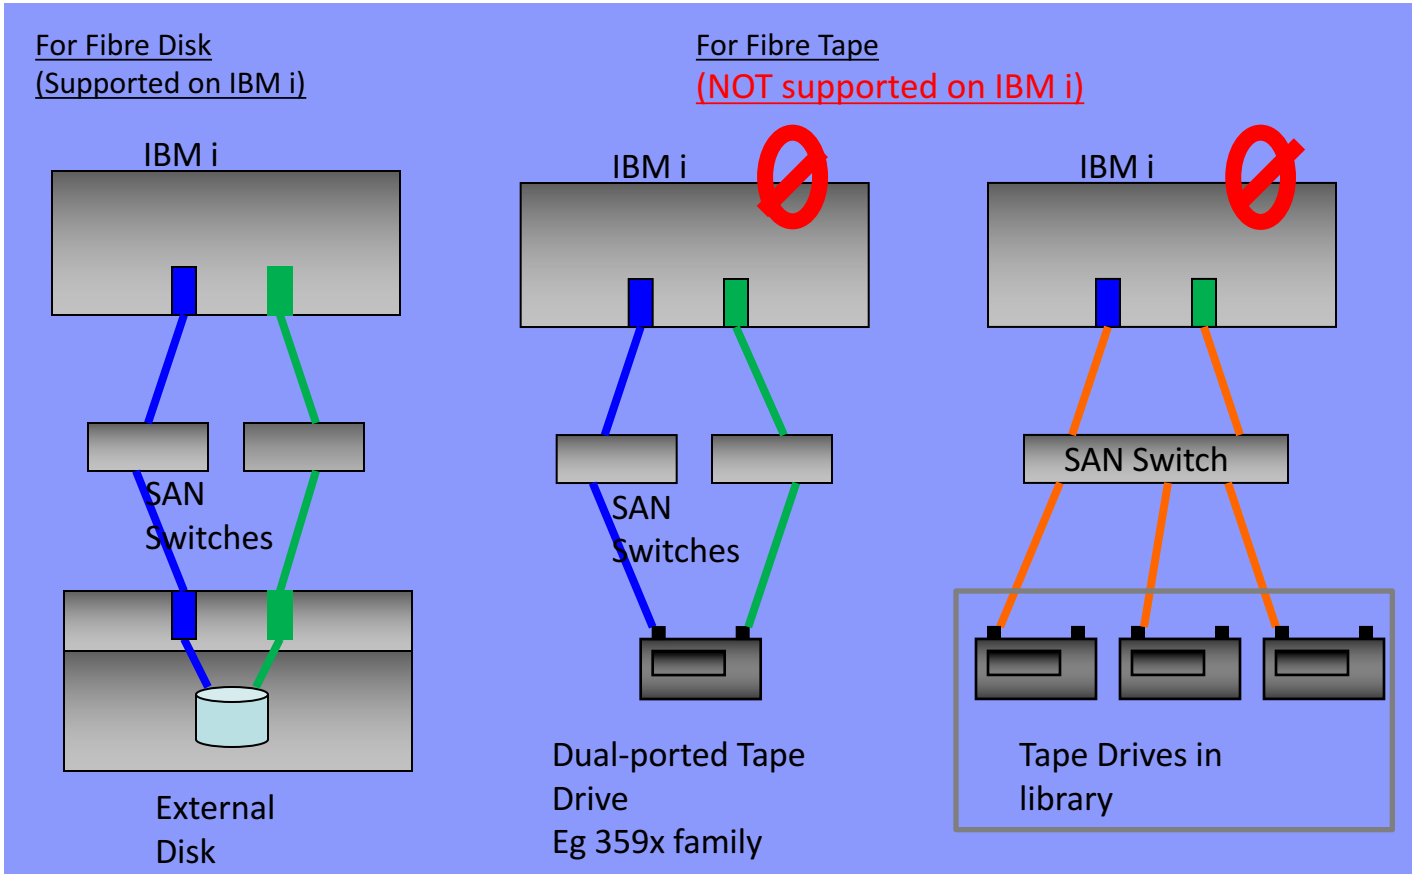

NEW - Tape multipath support with V7R2 TR2

- Will be supported on newer technology Fibre Channel drives
  - LTO5 in TS3100/3200, TS3310, TS3500/4500, 7226 enclosure
  - LTO6 in TS3100/3200, TS3310, TS3500/4500, 7226 enclosure
  - 3592-E07 in TS3500/4500
  - 3592-E08 in TS3500/4500
  - ProtecTIER 3.3.5.1
  - Will also be supported with future LTO and 3592 technology
- Up to 8 paths per device.
- Native attach, VIOS/NPIV attach, or both
- Function is being staged in over time:
  - Only Manual failover is planned to be available at TR2 GA
    - Vary off/on or deallocate/allocate to switch paths
  - Dynamic automatic failover function - post TR2 GA (PTFs).
    - Will not support distance solutions.
    - Will not support WORM media

VIOS – Storage attach

Three categories of storage attachment to IBM i through VIOS

1) Supported (IBM storage)

- tested by IBM; IBM supports the solution and owns resolution
- IBM will deliver the fix

2) Tested / Recognized (3rd party storage including EMC and Hitachi)

- IBM / storage vendor collaboration, solution was tested (by vendor, IBM, or both);
  - CSA in place, states that IBM and storage vendor will work together to resolve the issue

- IBM or storage vendor will deliver the fix

3) Other

- not tested by IBM, maybe not have been tested at all

No commitment / obligation to provide fix

Category #3 (Other) was introduced in the last few years, “other” storage used to invalidate the VIOS warranty. IBM Service has committed to provide some limited level of problem determination for service requests / issues involving "other” storage. To the extent that they will try to isolate it to being a problem within VIOS or IBM i, or external to VIOS or IBM i (ie. a storage problem). No guarantee that a fix will be provided, even if the problem was identified as a VIOS or IBM i issue

FlashSystem 900

Introducing IBM FlashSystem 900, the next generation in our lowest latency offering

- IBM MicroLatency™ with up to 1.1 million IOPS
- 40% greater capacity at a 10% lower cost per capacity
- IBM FlashCore™ technology, our secret sauce

Technical collaboration with Micron Technology, our flash chip supplier

- IBM enhanced flash technology
- MLC NAND flash offering with Flash Wear Guarantee

VAAI UNMAP and VASA support with IBMSIS for improved cloud storage performance and efficiency

Performance at-a-glance

Minimum latency	
Write	90 µs
Read	155 µs
Maximum IOPS 4 KB	
Read (100%, random)	1,100,00
Read/write (70%/30%, random)	800,000
Write (100%, random)	600,000
Maximum bandwidth 256 KB	
Read (100%, sequential)	10 GB/s
Write (100%, sequential)	4.5 GB/s

IBM MicroLatency module type	1.2 TB					2.9 TB				5.7 TB			
Modules quantity	4	6	8	10	12	6	8	10	12	6	8	10	12
RAID 5 capacity (TB)	2.4	4.8	7.2	9.6	12	11.6	17.4	23.2	29.0	22.8	34.2	45.6	57.0
Raw Capacity (TB)	7.1	10.7	14.2	17.8	21.4	26.3	35.1	43.9	52.7	52.7	70.3	87.9	105.5

IBM i Exploitation of Flash Systems

Flash System 840
Flash System 900
Absolute performance*
Up to 1.1 M IOPs
110us latency
MicroLatency™
4-48 TB

Flash Systems

SVC/Storwize – 1Q14
IBM i 7.1, 7.2
VIOS/NPIV – 1H15
IBM i 7.2
Native – 1H15
IBM i 7.2

Flash System Solutions

Flash System v840
Flash System v9000

Built in SVC Functionality

PowerHA support

SVC copy services paired with high performance storage

VIOS/VSCSI – 1Q14
IBM i 6.1, 7.1, 7.2
VIOS/NPIV – 1Q14
IBM i 7.1, 7.2
Native – 1Q14
IBM i 7.1, 7.2

*Performance has not be verified with IBM i

Live Partition Mobility
PowerVC for IBM i
Cloud for IBM i

Live Partition Mobility

Move a running partition from one Power7 (or newer) server to another with no application downtime

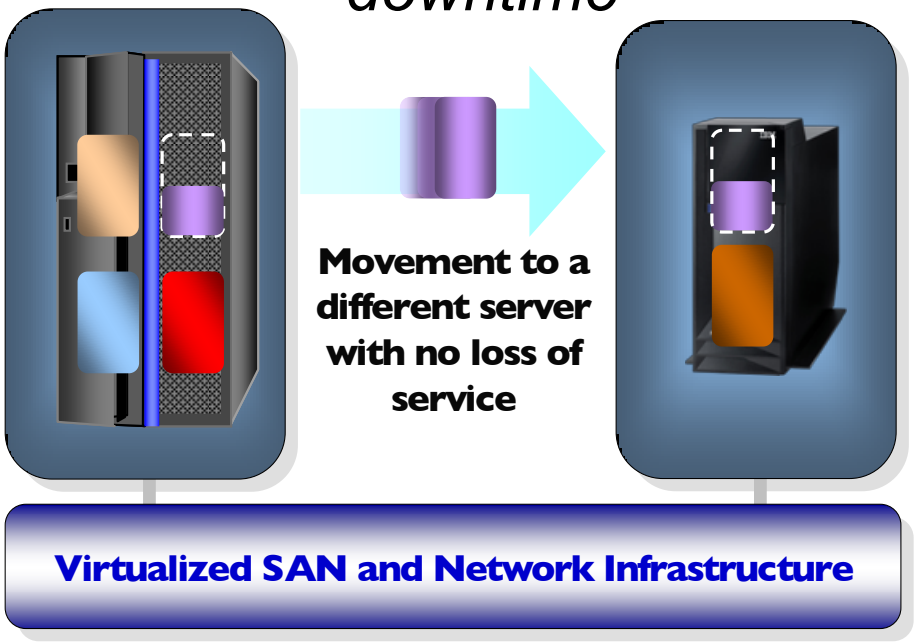

- Reduce planned downtime by moving workloads to another server during system maintenance

- Rebalance processing power across servers when and where you need it

Live Partition Mobility requires the purchase of the optional PowerVM Enterprise Edition

42

Fax support for Virtual Environments

- In 7.2 WAN can run in an Ethernet to Ethernet Device Server with multiple RS232 serial ports
  - Provides true Virtual Serial ports for WAN applications
  - Clients running IBM Facsimile Support for i, 5798-FAX, can use this new support
- Expands advanced virtualization capabilities:
  - Reduces total cost by allowing
  - Fewer PCI slots for applications requiring a modem
    - One Ethernet adapter can provide both TCP/IP connectivity and WAN
  - Allows IBM i client partitions with virtual I/O to use FAX and other WAN applications
    - Support for Flex and Blades - previously had no WAN capability
- Minimal disruption for existing WAN applications
  - No application changes
  - Simple configuration change required for IBM i partition

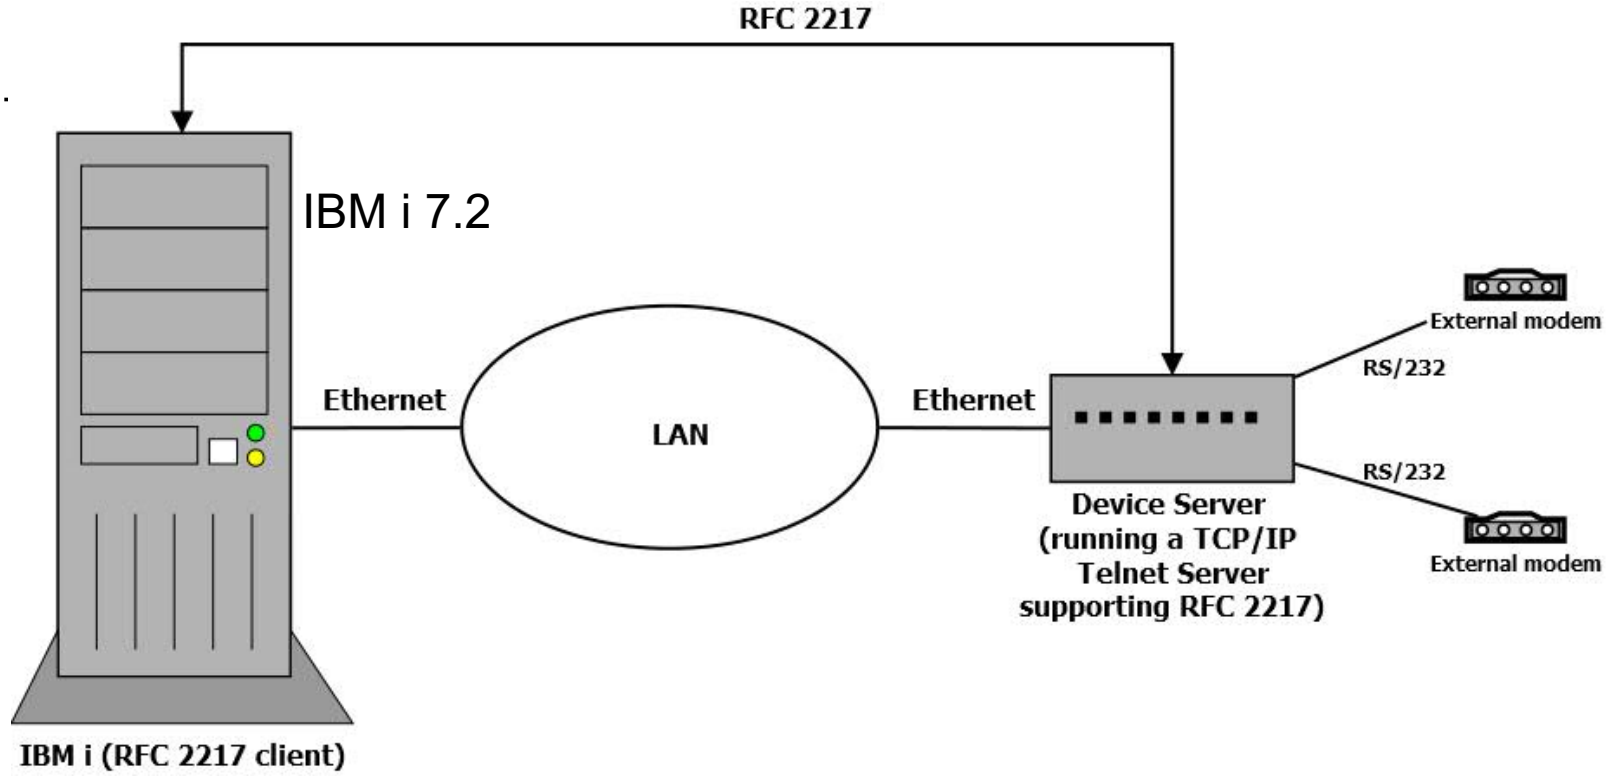

HyperSwap – IBM i 7.2

The IBM i OS stays up during the switch.

The first stage is with the entire OS

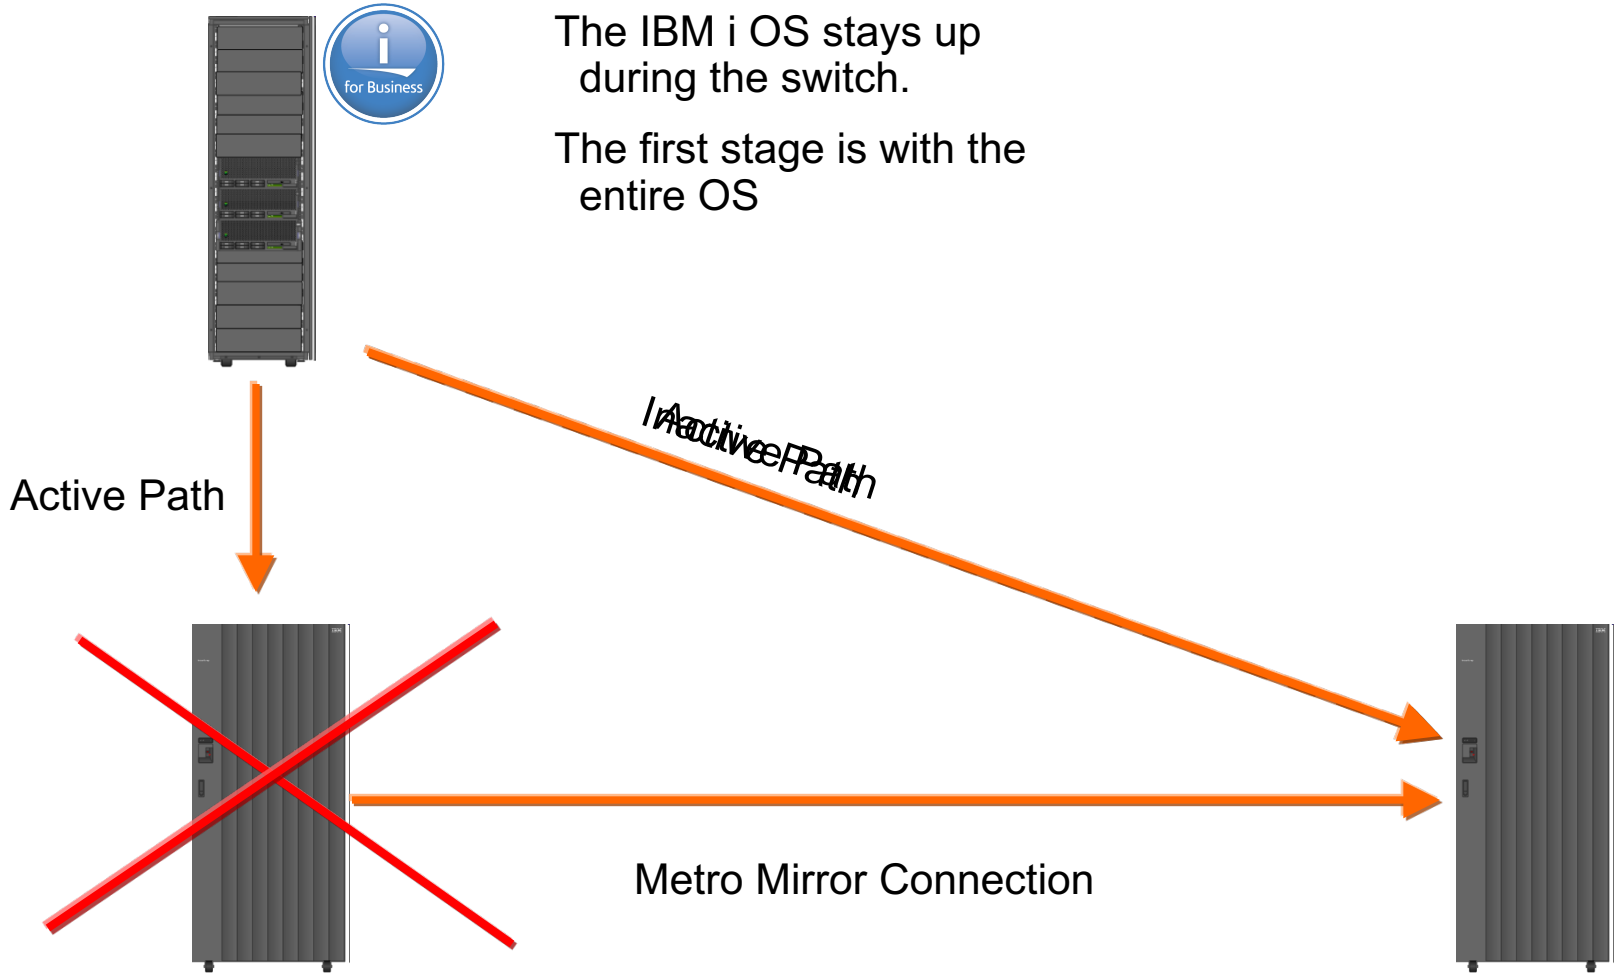

HyperSwap with LPM – IBM i 7.2

After LPM, the affinity to the storage box might be non-optimal

After the HyperSwap the storage affinity can be maintained

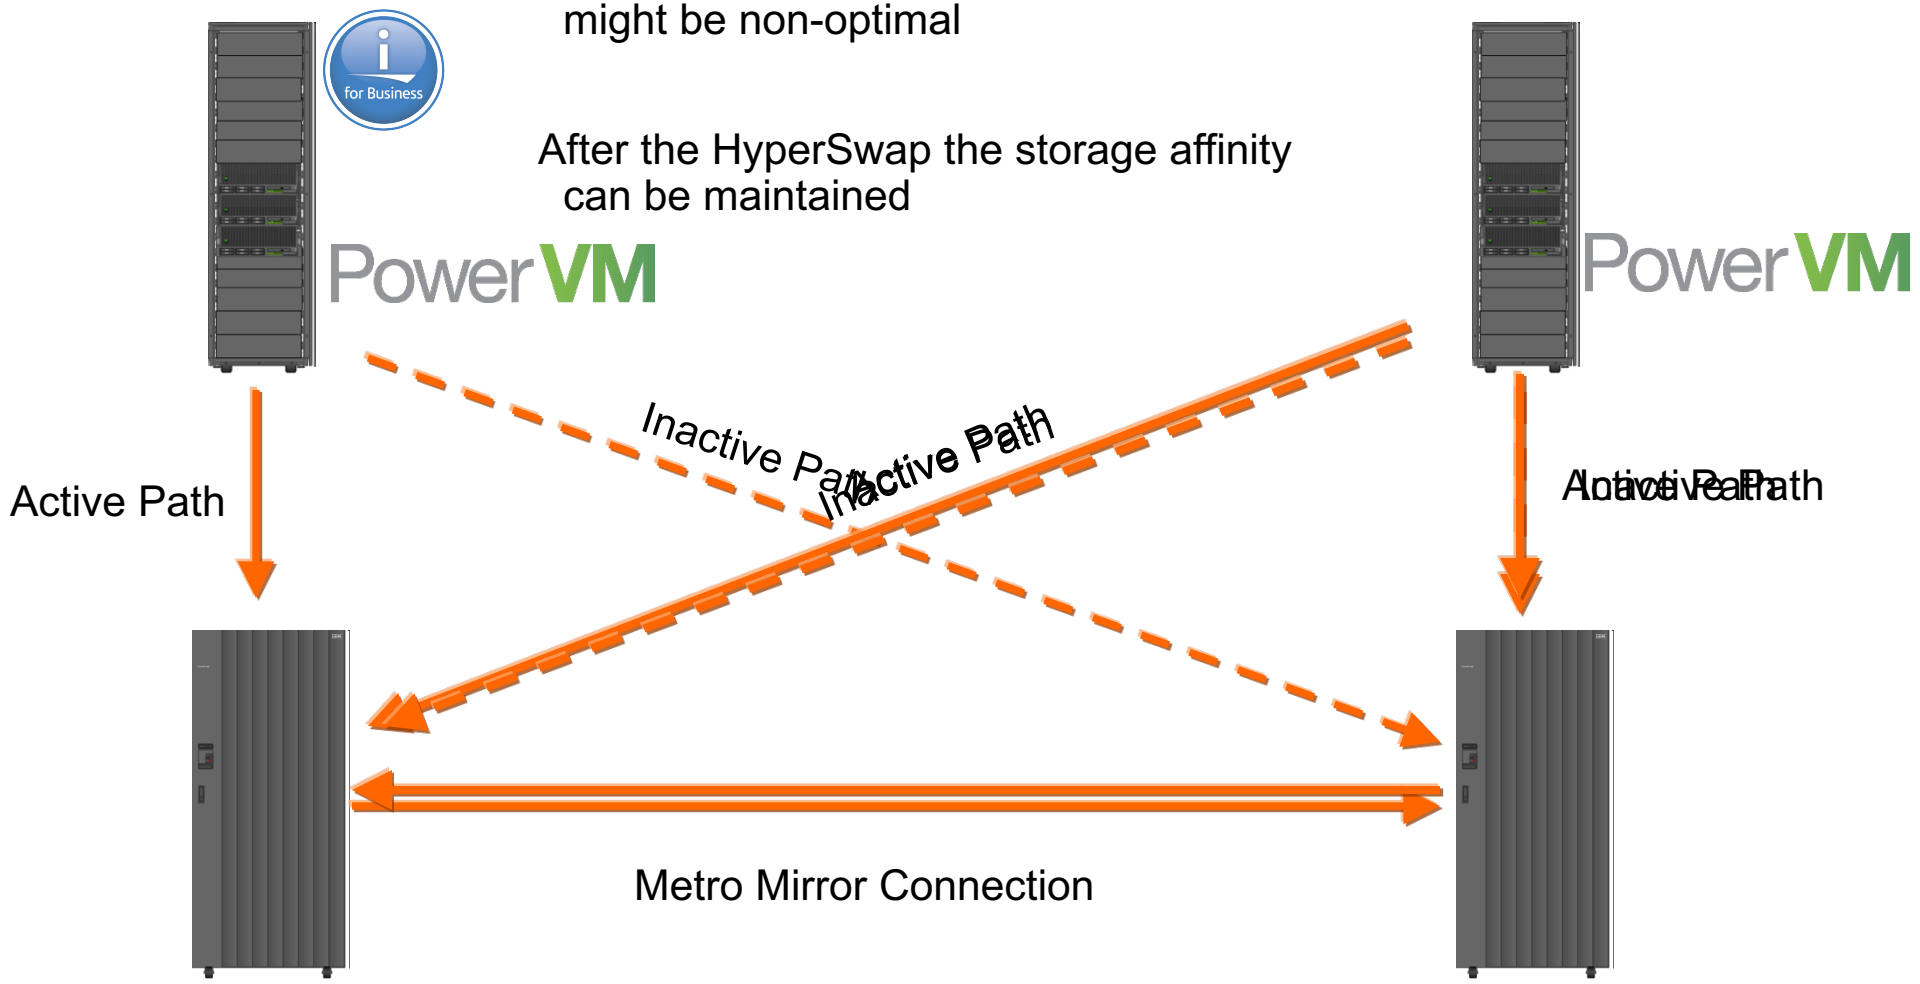

Power Enterprise Pools

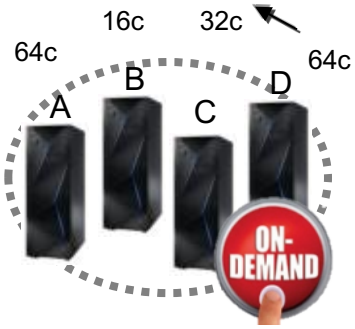

Create activations pool and shared between Server(s)

Migrate CPU/Memory Capacity Between Servers

Instant change, No IBM involvement

Move your Applications with Live Partition Mobility (LPM)

Move your Activations between machines

- Enables **workload balancing**
- Simplifies **systems maintenance**
- **Disaster Recovery**
- Supports your cloud environment

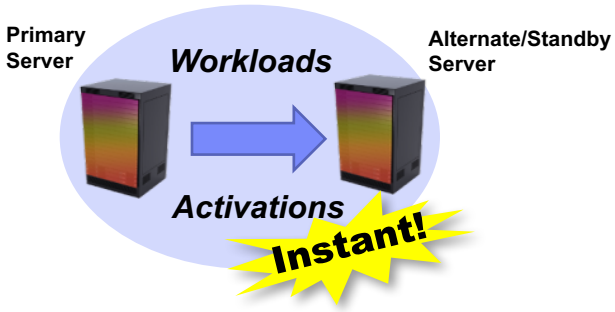

- **Two systems are better than one…but Do not Cost twice as much**
- **User controlled, always available, no IBM involvement**

Systems Management Transition

The market shift to Cloud drove our adoption of Open standards-based Systems Management

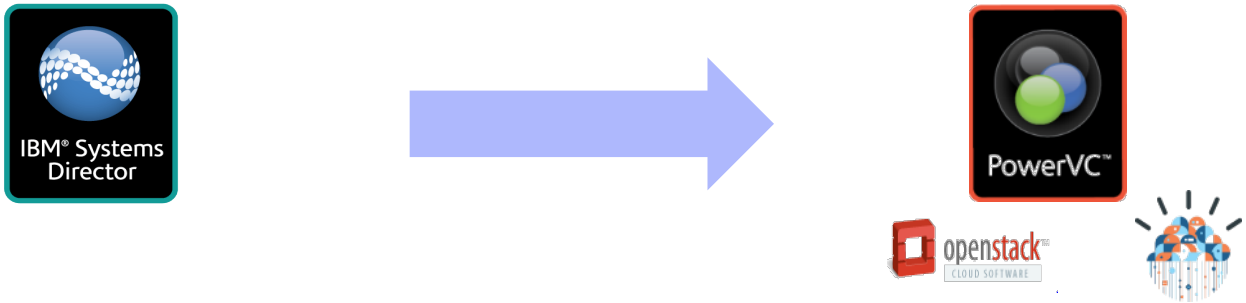

- Cross IBM Platform management
- Proprietary interfaces
- IBM Eco System only
- IBM provided resource support
- Monolithic design
- Broad functionality
- Built for in-house IT management

- Heterogeneous Management at the Cloud level
- Open, standard APIs
- Rich, diverse Eco System
- Hardware vendor provided device support
- Loosely coupled design
- Extendable, Fit-For-Purpose functionality
- Built for Private, Public, and Hybrid Clouds

Key PowerVC Capabilities

- Virtualization management for Power Systems – PowerVM and PowerKVM
- Key Capabilities
  - Advanced virtualization management for Power Systems
  - Virtual machine capture and deployment
  - Virtual machine relocation
  - Policy based VM placement
  - One-click system evacuation
  - Optimization and rebalancing
  - Quick setup and Time to Value
- Based on OpenStack
  - Leverage open community
- Capabilities beyond OpenStack
  - Simplified user interface
  - Platform EGO scheduler
  - Reliability and serviceability

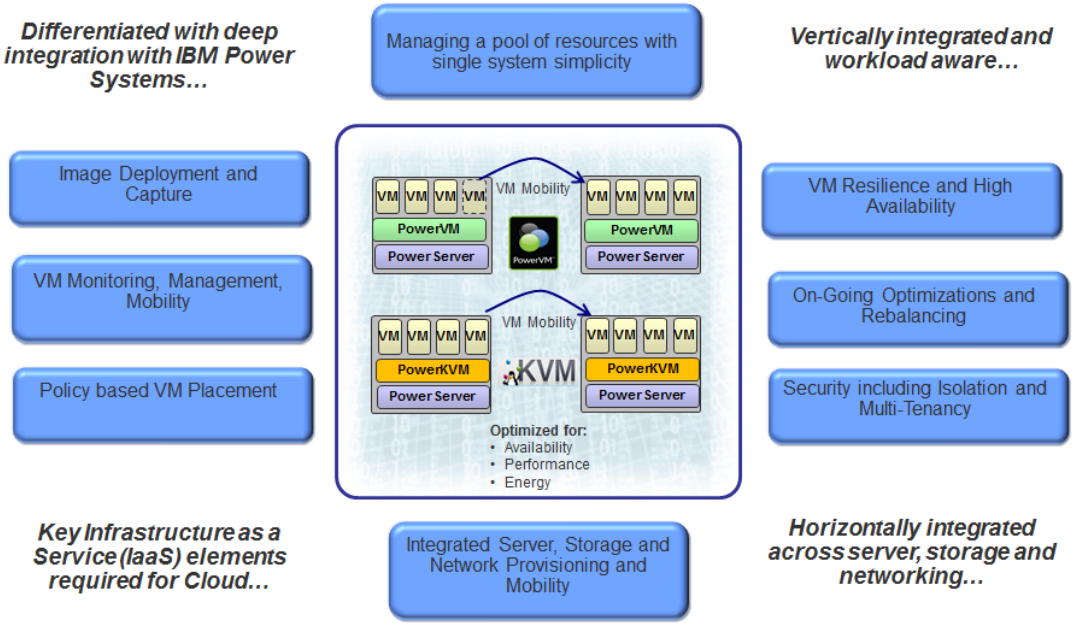

Why OpenStack for Power?

- Community development improves speed of innovation
  - Over 12,000 people in the community
  - Covering 130 countries
  - Rapid growth of community
    - Apr 2012 – 150 orgs, 2600 individuals
    - Jan 2013 – 850 orgs, 6600 individuals
    - Sept 2013 over 12,000 individuals
    - Sept 2015 over 17,000 individuals
- Protects current investment with simple path to new technology
  - Broad industry support and ecosystem for extensive device support and cloud standards
  - Open and extensible architecture to quickly integrate into existing infrastructures
- Open alternative to proprietary cloud stacks
  - Open APIs provide flexibility and agility
  - Foundation for private and public clouds built on best practices of industries leading thinkers

PowerVC Setup and Configuration

Simple, Intuitive with a Focus on Time to Value…

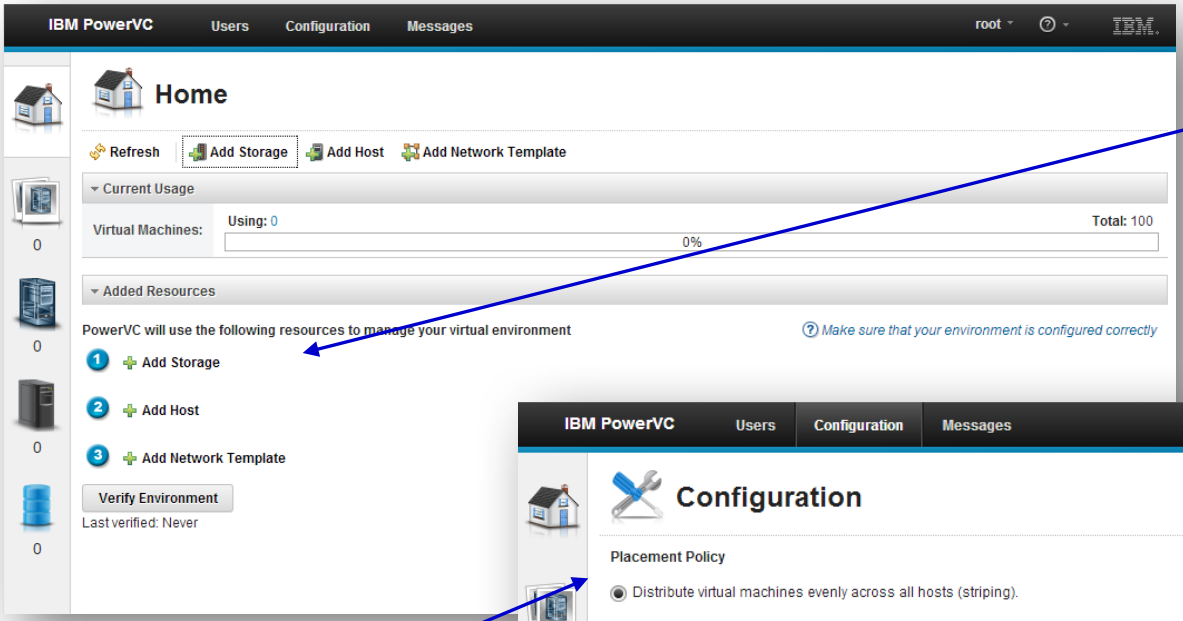

1. Add Storage to be managed…

- Provide IP address
- Provide user-id & password

2.Add Servers to be managed…

- Provide IP address of IVM/HMC
- Provide user-id & password

3. Add Network Template…

- Provide VLAN ID
- Provide IP Configuration

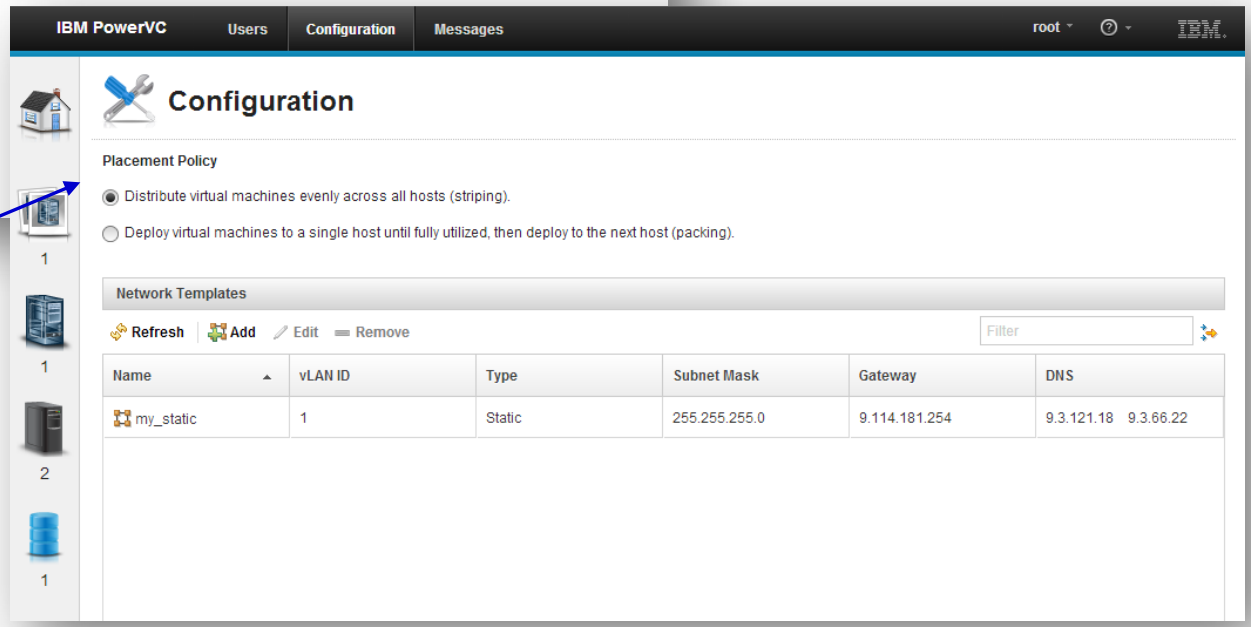

Configuration Placement Policies…

- Stripe Virtual Machines
- Pack Virtual Machines

Virtual Machine Management

Providing the fundamental visibility and management for Power virtual machines…

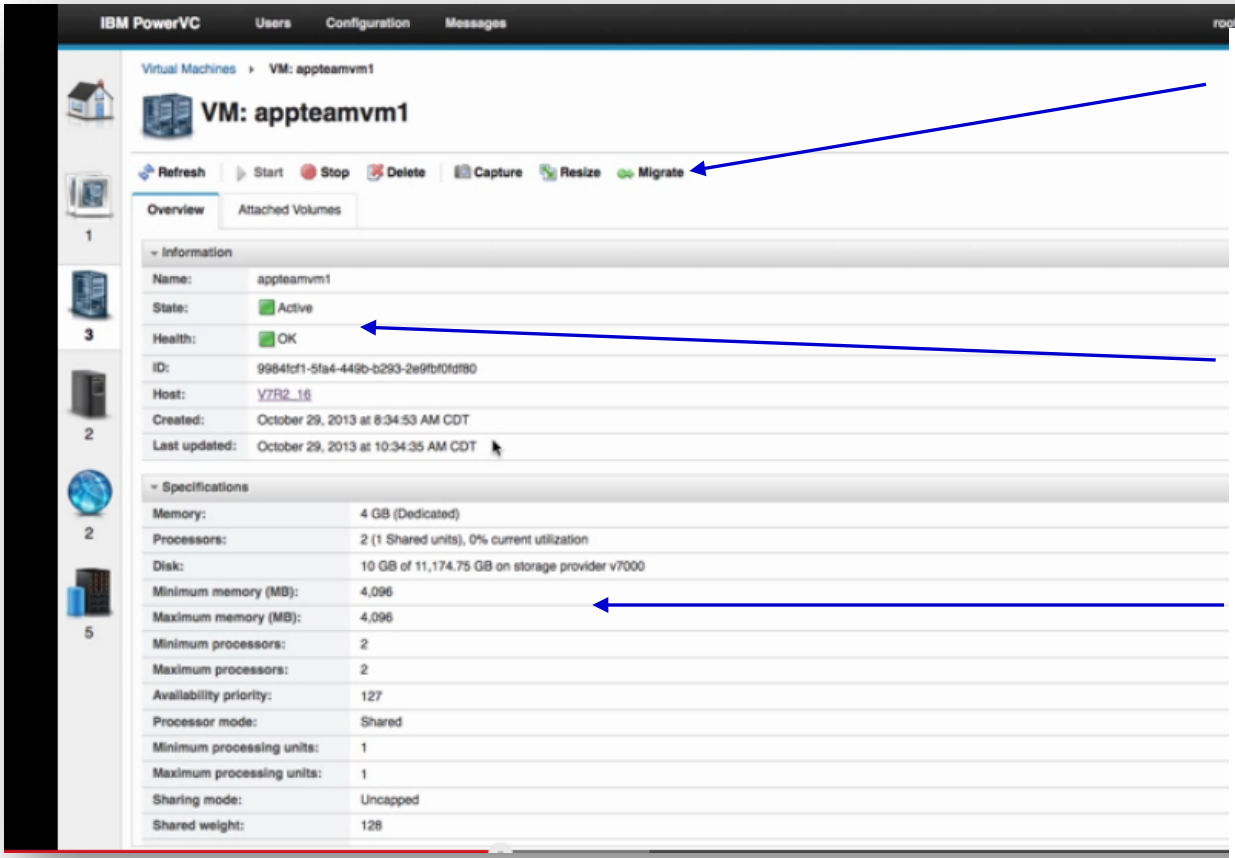

[Virtual Machine Management…]

- Start and stop the virtual machine
- Delete the virtual machine
- Capture VM as Image
- Resize including DLPAR
- Migrate using placement policy

[Virtual Machine Health…]

- Virtual Machine State
- Virtual Machine Status
- Red/Green/Yellow Indicators

[Virtual Machine Properties…]

- Processor, memory, disk
- Related host information
- Network configuration
- Disk volumes (separate tab)

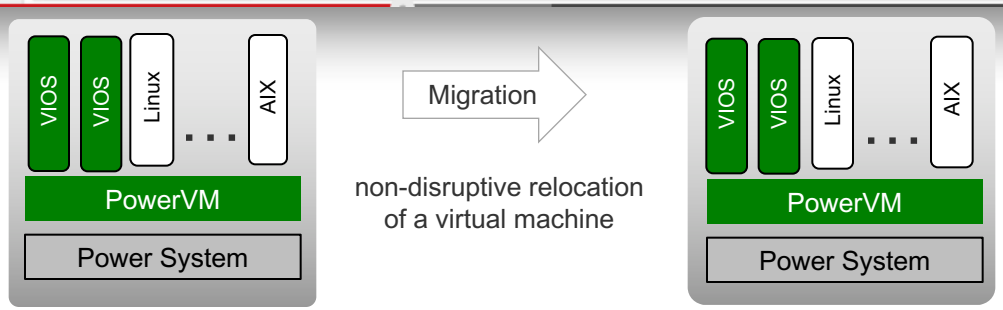

The target host can be selected by the user or selected based on the placement policy

One Click System Evacuation(Q4 2014)

Provides easy, graceful way to prepare for maintenance

- Automatically relocate all virtual machines to other machines
  - Use the PowerVC scheduler to determine the target host or manually select the destination host
  - Clears the system of virtual machines without excessive administrator work
- Alternatively, fence off the physical host to prevent new virtual machines from being deployed or moved to that host
  - Option to allow administrators greater control of relocation operation

PowerVC 1.2.3 Host Groups(June 2015)

Host Groups allow the PowerVC administrator to create a logical boundary around a group of physical servers

- Each server can only be in one host group
- Deployment, mobility and remote restart are only allowed within the group
- Each group has its own placement policy
- Hosts are placed in the default group when added

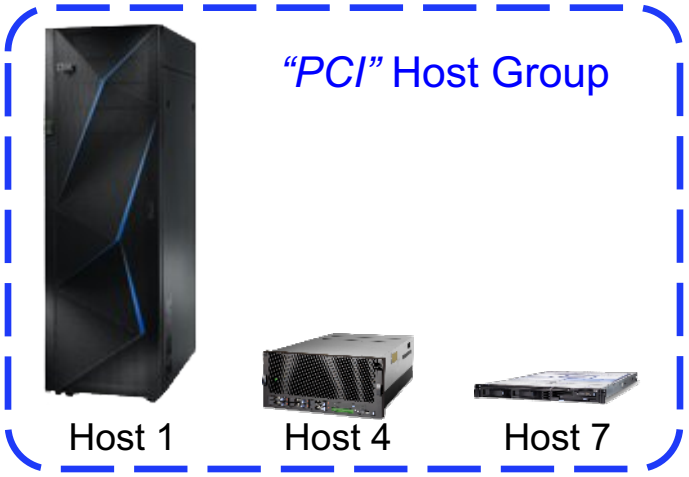

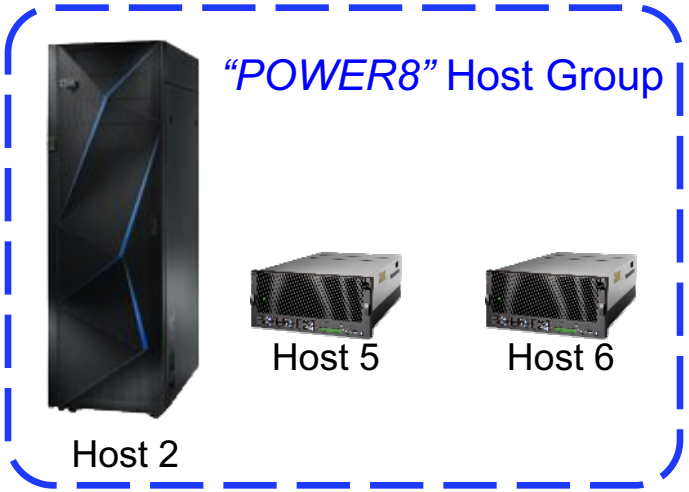

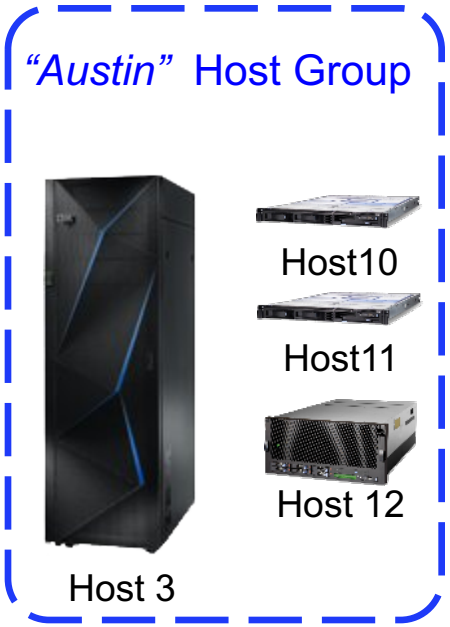

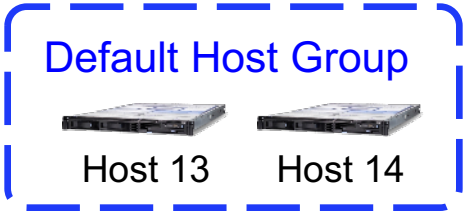

PowerVC 1.2.3 Advanced Placement(June 2015)

Scheduler support VM placement based on CPU & Memory capacity and CPU Utilization

The PowerVC scheduler takes the capacity of servers into account to determine which host to deploy or relocate VMs to. Hosts with the greatest free CPU or memory allocation becomes the target of the next VM.

The scheduler can also take host CPU utilization into account when scheduling VMs

Free CPU / Memory Capacity

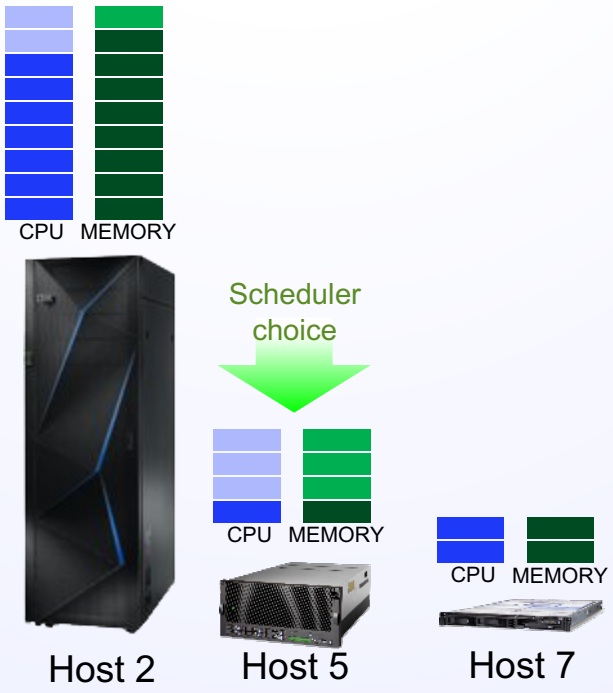

Host CPU Utilization

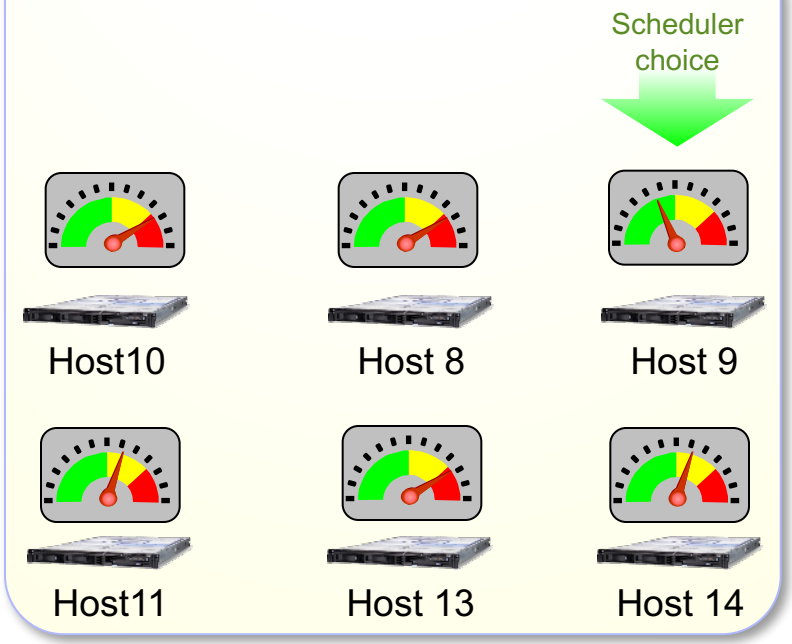

PowerVC Placement Policies

	Policy Description	Initial Placement	
Packing	▪ Pack workload on fewest physical servers ▪ Maximizes usable capacity, reduces fragmentation, reduce energy consumption	✓	
Striping	▪ Spread workload across as many physical servers as possible ▪ Reduce impact of host failures, higher application performance	✓	
CPU Balance	▪ Place VMs on the hosts with the least allocated CPU ▪ Higher application performance	✓	
Memory Balance	▪ Place VMs on the hosts with the most available memory ▪ Improve application performance	✓	
Affinity	▪ Affinity specifies that VMs should be placed on the same host or few hosts ▪ Useful for collocating VMs on the same host(s)	✓	
Anti-Affinity	▪ Do not place VMs on same host ▪ Useful for ensuring VMs are not collocated ▪ Availability cluster support (e.g. PowerHA) ▪ Higher application performance	✓	

PowerVC 1.2.3 VM Collocation Affinity and Anti-affinity (June 2015)

VM with no affinity requirements – can go anywhere within the host group

VMs with affinity – must be placed on the same host

VMs with anti-affinity – cannot be placed on the same host

Host 1 Host 2 Host 3

Affinity and Anti-affinity provide control over which VMs can be placed on the same host
-VMs with Affinity must be deployed to the same host
-VMs with Anti-Affinity must not be placed on the same physical host

PowerVC Remote Restart(June 2015)

Improved recovery from unexpected system failures

PowerVC Remote VM Restart enables restarting VMs from a failed host on another server

- Works with AIX, IBM i or Linux VMs
- Requires a human decision to perform restart using PowerVC
- Host Group policy controls VM placement
- Supports both PowerVM and PowerKVM
- Requires POWER8 with firmware 8.20

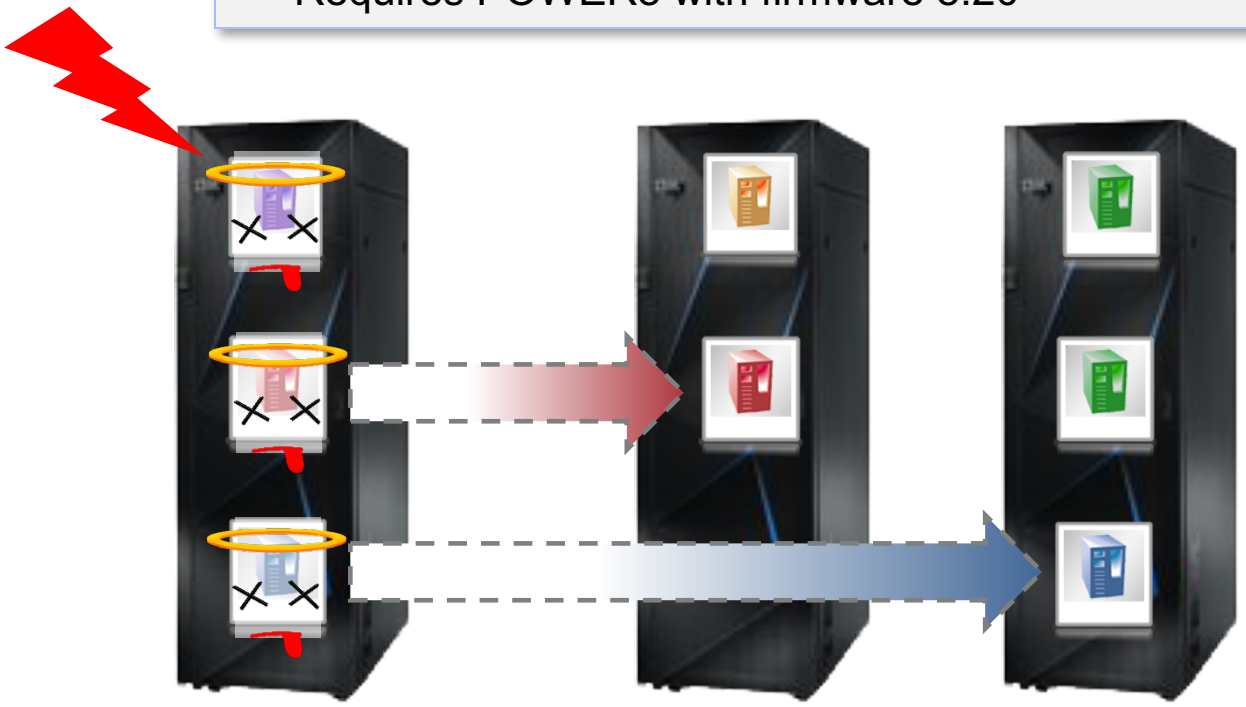

PowerVC v1.3 Dynamic Resource Optimizer(Q4 2015)

Policy-based automation to balance workloads

PowerVC v1.3 Dynamic Resource Optimizer allows for automated rebalancing of workloads between servers

- Server workload can be automatically balanced two ways:
  - Relocating Virtual Machines between servers
  - Moving processor capacity between servers using Enterprise Capacity on Demand
- Works with AIX, IBM i or Linux VMs

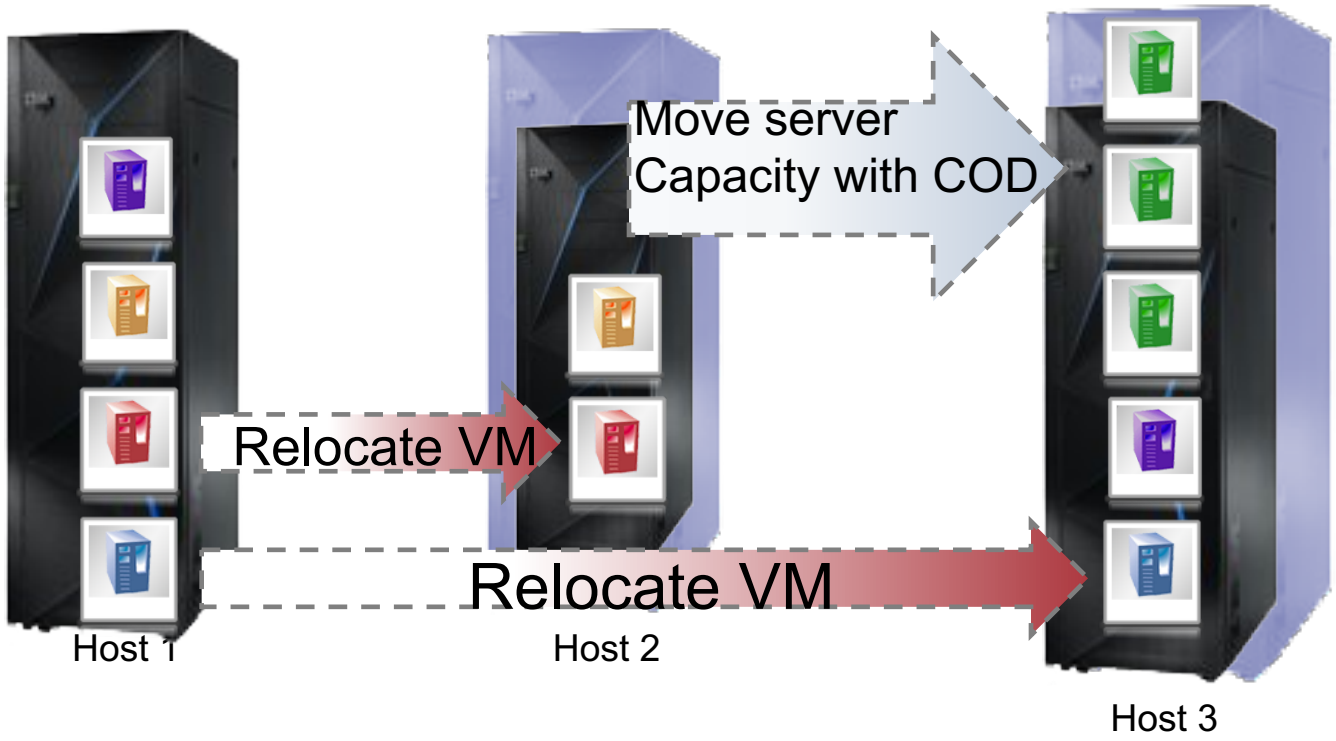

58

PowerVC Multi-disk capture and deployment (June 2015)

Capture

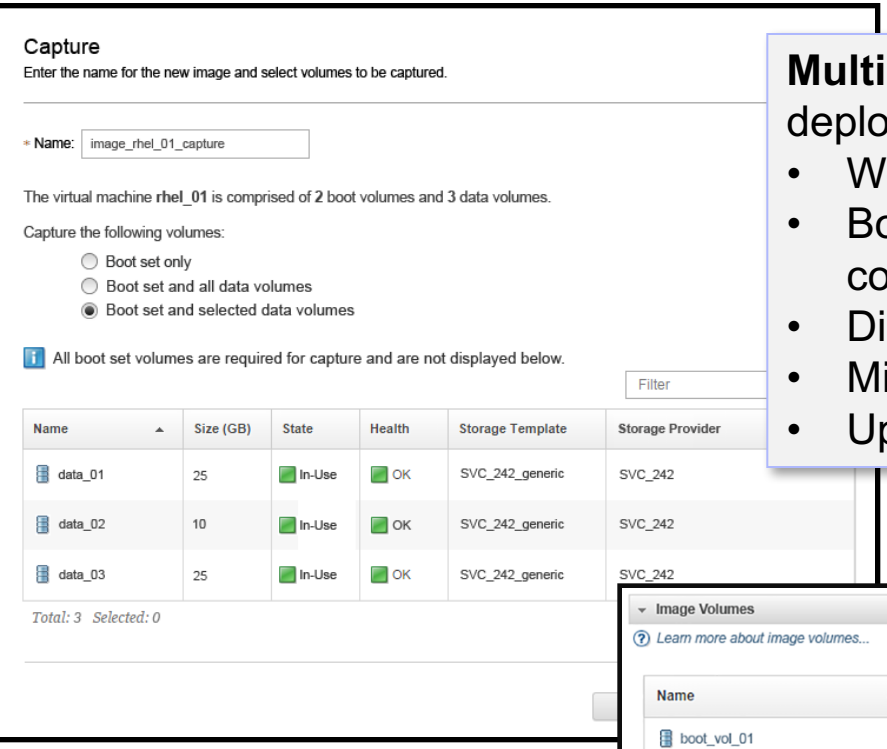

Multi-disk capture and deployment allows capture and deployment of boot and data volumes

- Works with AIX, IBM i or Linux VMs
- Boot and data volumes can be captured separately and combined and deployed together
- Disk volumes do not have to be on the same device
- Mirrored boot volumes are captured and deployed
- Up to 64 volumes supported

Deploy

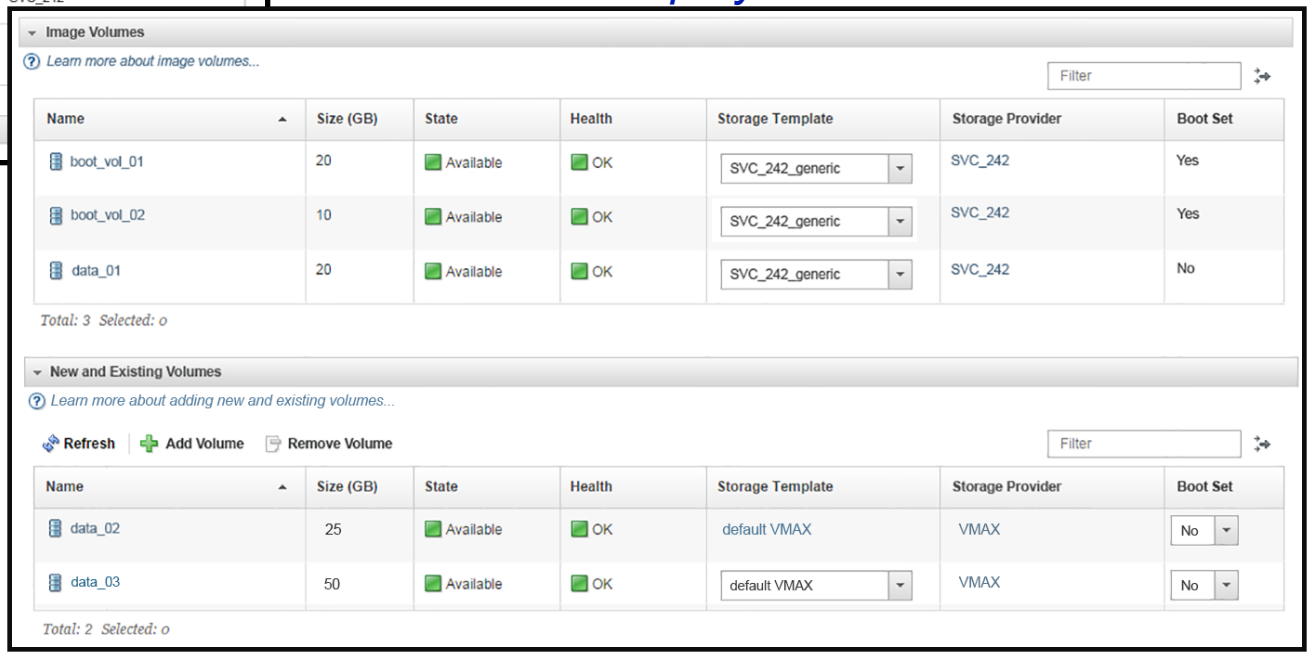

PowerVM NovaLink: Power Systems Platform Management Evolution

Goal: Simplify PowerVM virtualization, accelerate cloud enablement, and improve scale

Key Benefits

- ❑ Improved management scalability – support more virtual machines
- ❑ Aligns PowerVM with the OpenStack community scale model – simplifying future OpenStack exploitation
- ❑ Simplifies management configuration – HMC not needed for virtual machine deployment and configuration
- ❑ Enables flexibility to use any OpenStack based manager to manage PowerVM
- ❑ Uniform management for PowerVM and PowerKVM based systems

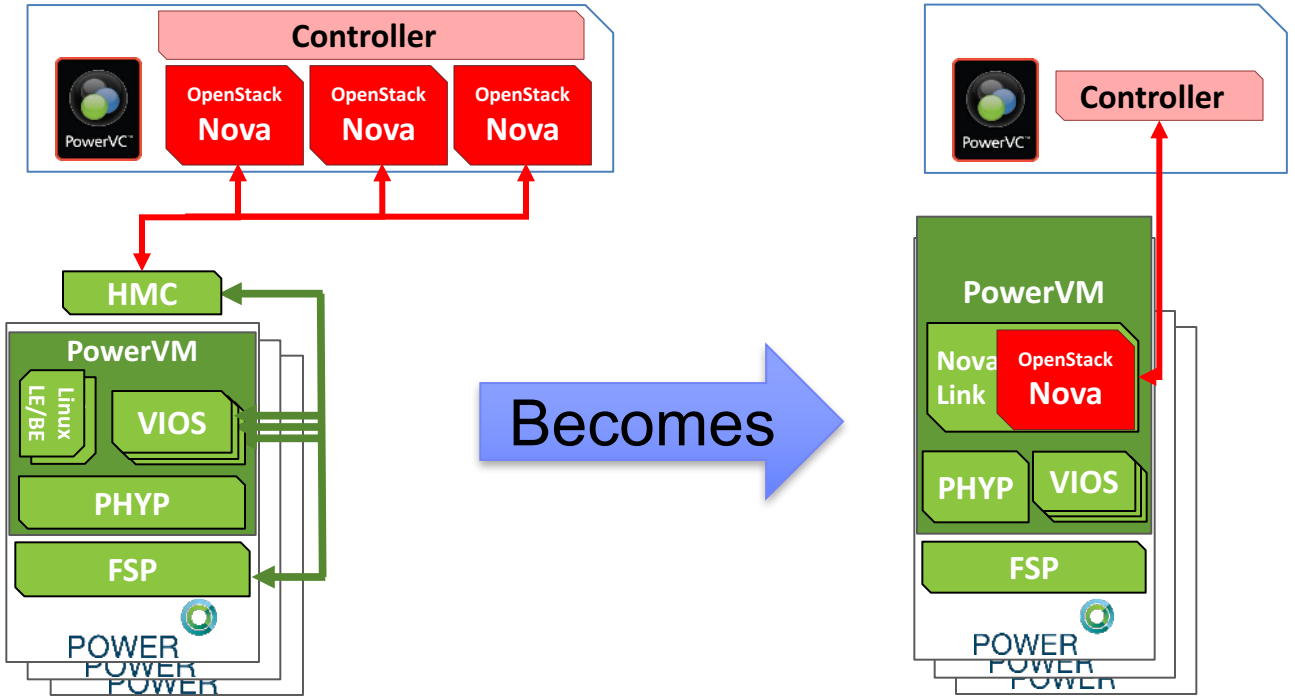

VMware vRealize Virtualization for Power & z Systems

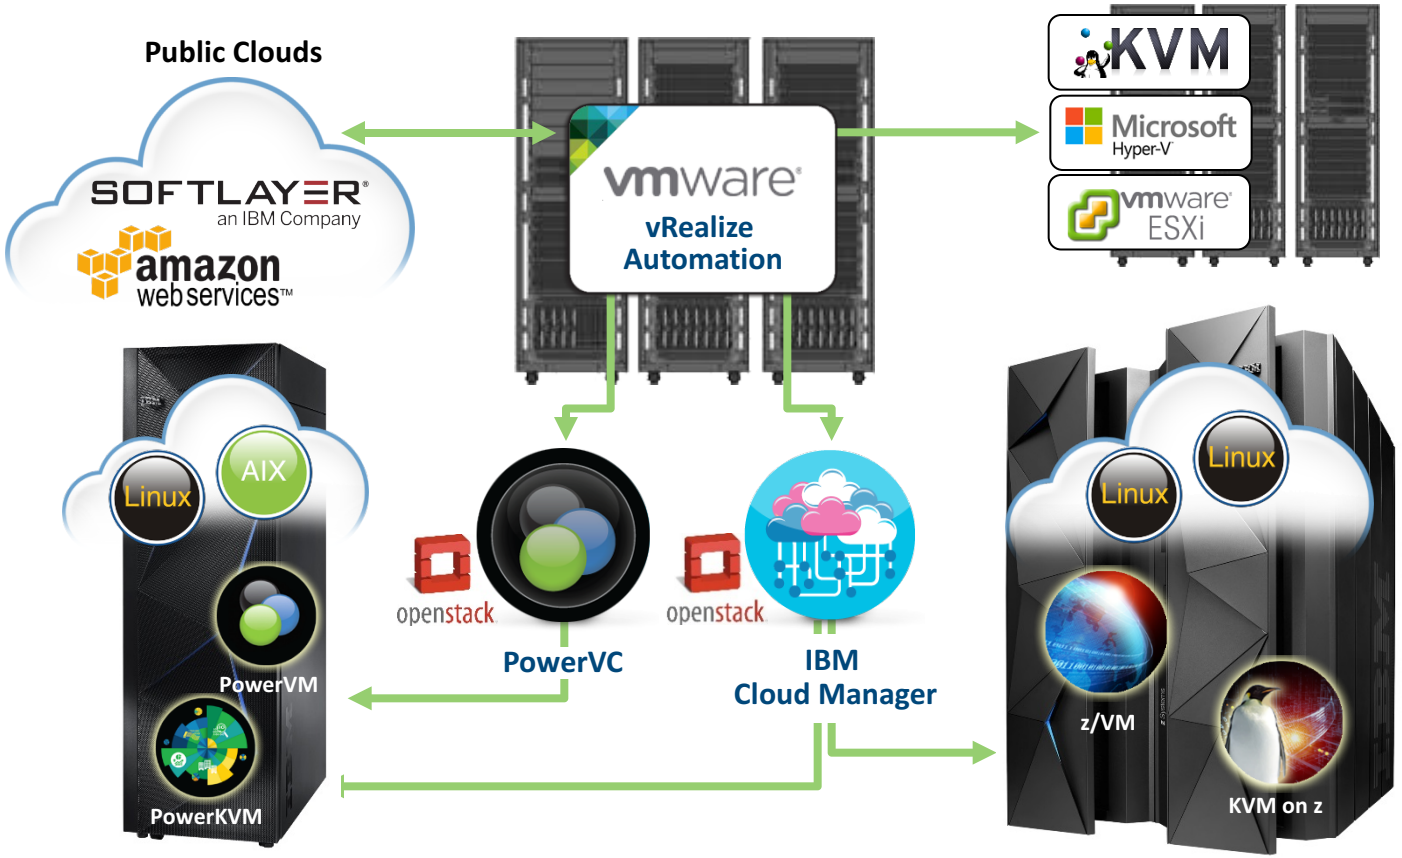

PowerVC v1.3.2

- ✓ ***Automates VM provisioning and best practices***
- ✓ *Improve resource utilization to* ***reduce capital expense and power consumption***
- ✓ ***Increase agility*** *and execution to quickly respond to changing business requirements*
- ✓ ***Increase IT productivity*** *and responsiveness*
- ✓ ***Manage scalability*** *without adding complexity*

Announce – 10/11/2016 GA – 12/16/2016

✓Improvements in Management of High Availability for PowerVM

✓New storage management capabilities

✓Enhanced Policies for Dynamic Resource Optimization

✓Improved Management Support for PowerVM NovaLink

- **New HA Capabilities**
  - Automated Policy-Based VM Restart for NovaLink and HMC Configurations -> Enables faster recovery from server failures
  - VM Restart when System is Powered Off -> Expands the coverage of HA events
- **New Storage Capabilities**
  - Support for NPIV Hitachi VSP & USP Storage(**Won't be in Announcement)**
  - Improved Zoning Control -> Allows clients to have fewer zones
- **Dynamic Resource Optimizer(DRO) Improvements**
  - Balances workload based on Memory usage as well -> Allow memory constrained environments to be automatically balanced optimizing systems and reducing labor costs
- **PowerVM NovaLink Management Improvements**

  Support for
  - **SR-IOV vNIC and vNIC Failover Configurations**
  - NovaLink Partition running Red Hat Linux
  - PowerVM Shared Storage Pools
  - **VM Console Launch**

IBM Cloud Storage Solutions for i

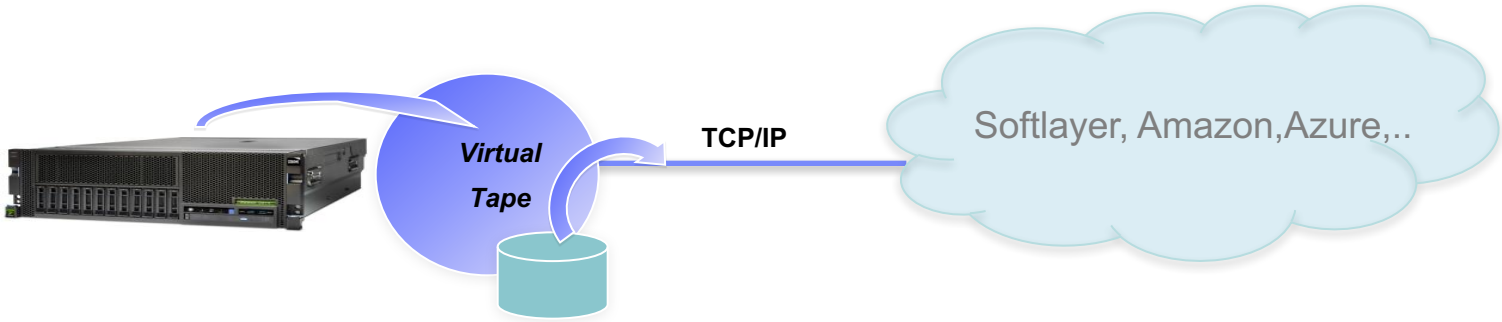

- An API that enables deployment of IBM i data to a public cloud
  - Targeted for customers with under 1 Tbyte of data
- Auto save and synchronize files in the IBM i IFS directory
  - Roll your own backup/recovery (bandwidth considerations)
- Product offering will feature
  - BRMS with virtual tape management
  - Security via VPN

Cloud storage – cashed backup IBM i environment

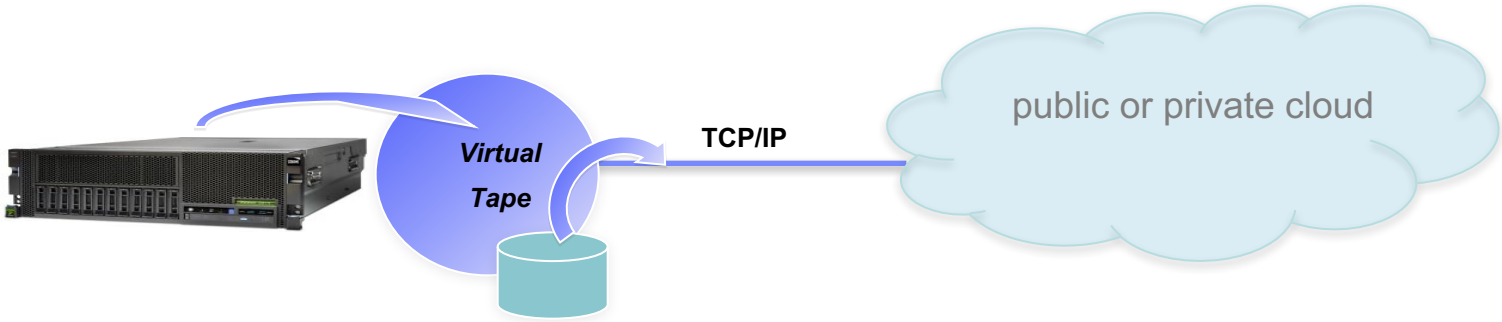

- Foundational topology is enabled via Virtual Tape
  - Physical storage cache is via local disk
  - Data is saved from i as tape objects
- Tape objects are converted to cloud objects
  - Cloud provider has an object format that enables saves to generic disk of any kind
  - To deploy to the cloud, IBM groups the tape objects into cloud objects
- Cloud objects will be transmitted asynchronously to a cloud provider
  - IBM i will leverage BRMS to manage save process from virtual tape to public cloud

Cloud Storage clients concepts

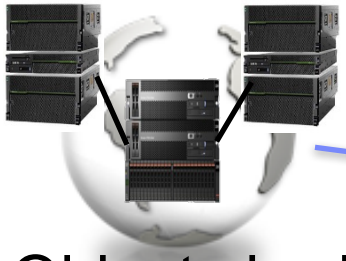

Object sharing among multiple systems

- PTFs
- ISOs
- Files
- others

Backup / Archive

- Move images/Files offsite

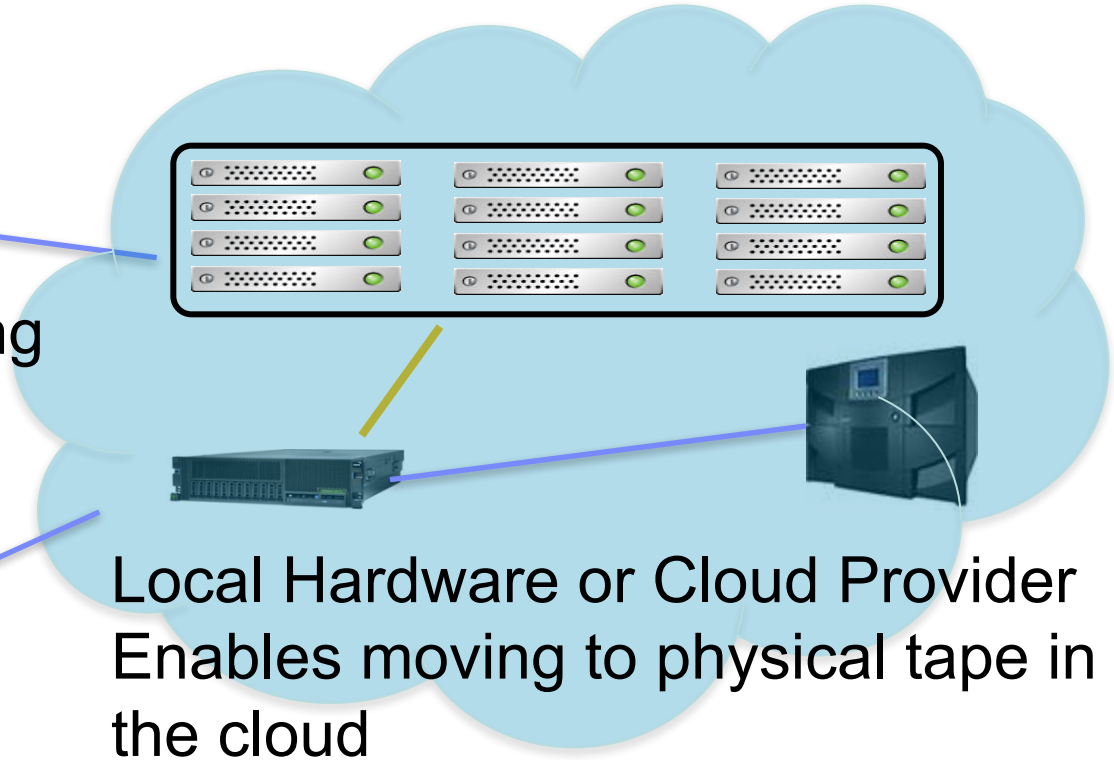

Local Hardware or Cloud Provider
Enables moving to physical tape in the cloud
Enables recovery testing offsite

Learn more about PowerVM on the Web

http://www.ibm.com/systems/power/software/virtualization

PowerVM

(… or Google 'PowerVM' and click I'm Feeling Lucky)

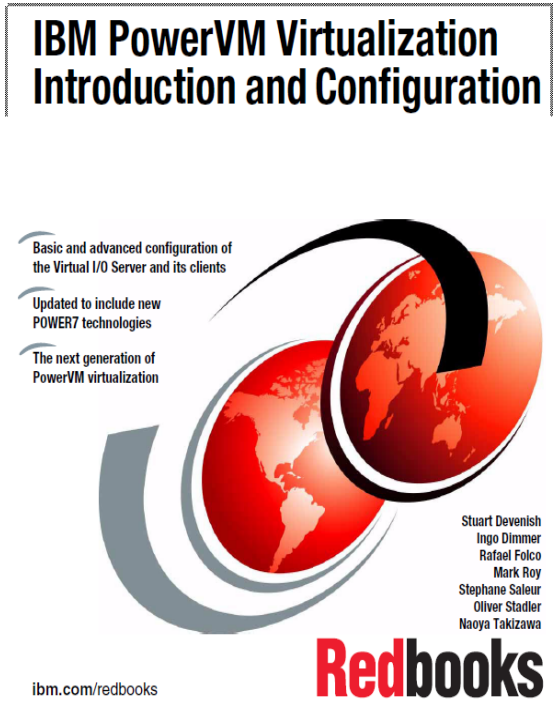

PowerVM resources include white papers, demos, client references and Redbooks

72

Resources and references

- Techdocs – http://www.ibm.com/support/techdocs
  (presentations, tips & techniques, white papers, etc.)

- IBM PowerVM Virtualization Introduction and Configuration - SG24-7940
  http://www.redbooks.ibm.com/abstracts/sg247940.html?Open
- IBM PowerVM Virtualization Managing and Monitoring - SG24-7590
  http://www.redbooks.ibm.com/abstracts/sg247590.html?Open
- IBM PowerVM Virtualization Active Memory Sharing – REDP4470
  http://www.redbooks.ibm.com/abstracts/redp4470.html?Open
- IBM System p Advanced POWER Virtualization (PowerVM) Best Practices - REDP4194
  http://www.redbooks.ibm.com/abstracts/redp4194.html?Open
- Power Systems: Virtual I/O Server and Integrated Virtualization Manager commands (iphcg.pdf)
  http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphcg/iphcg.pdf

Questions?

Trademarks and Disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.
Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at http://www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and are used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

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

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.