

# IBM i and External Storage

MRMUG February, 2012

IBM i can use external storage

- Direct attached
- Using VIO Server
- In combination and with internal storage

All disk presented appear and are configured and utilized the same as internal disk

- A disk is a disk

With external disk, protection is provided by the Storage device

- Don't turn on RAID or mirroring in OS/400 for external disk



# Define Your Base Needs

- ▶ **Reasons for Internal Storage on IBM i**
  - ▶ Simplicity and ease-of-use
  - ▶ Lower cost
  - ▶ High Performance
  - ▶ Entry level IBM i PowerHA
- ▶ **Reasons for external storage on IBM i**
  - ▶ Storage capacity flexibility
  - ▶ Performance flexibility and storage tiers
  - ▶ Meet common corporate storage platform standards
  - ▶ Advance IBM i PowerHA combined with storage copy services for HA, DR and on-line backups for midrange to large IBM i
  - ▶ Advanced RAS, CHARM and hibernation functions



# Define Your End-To-End System Requirements

Strategy and requirements by workload or LPAR:

- ▶ Production
- ▶ Development
- ▶ Test, QA
- ▶ Archive
- ▶ Growth
- ▶ HA/DR

Single solution vs. Mix-and-Match



## **Traditional Influencers**

Storage model line or vendor preference

Capacity, Performance, Flexibility and Cost

## **New, Strategic Influencers**

RAS

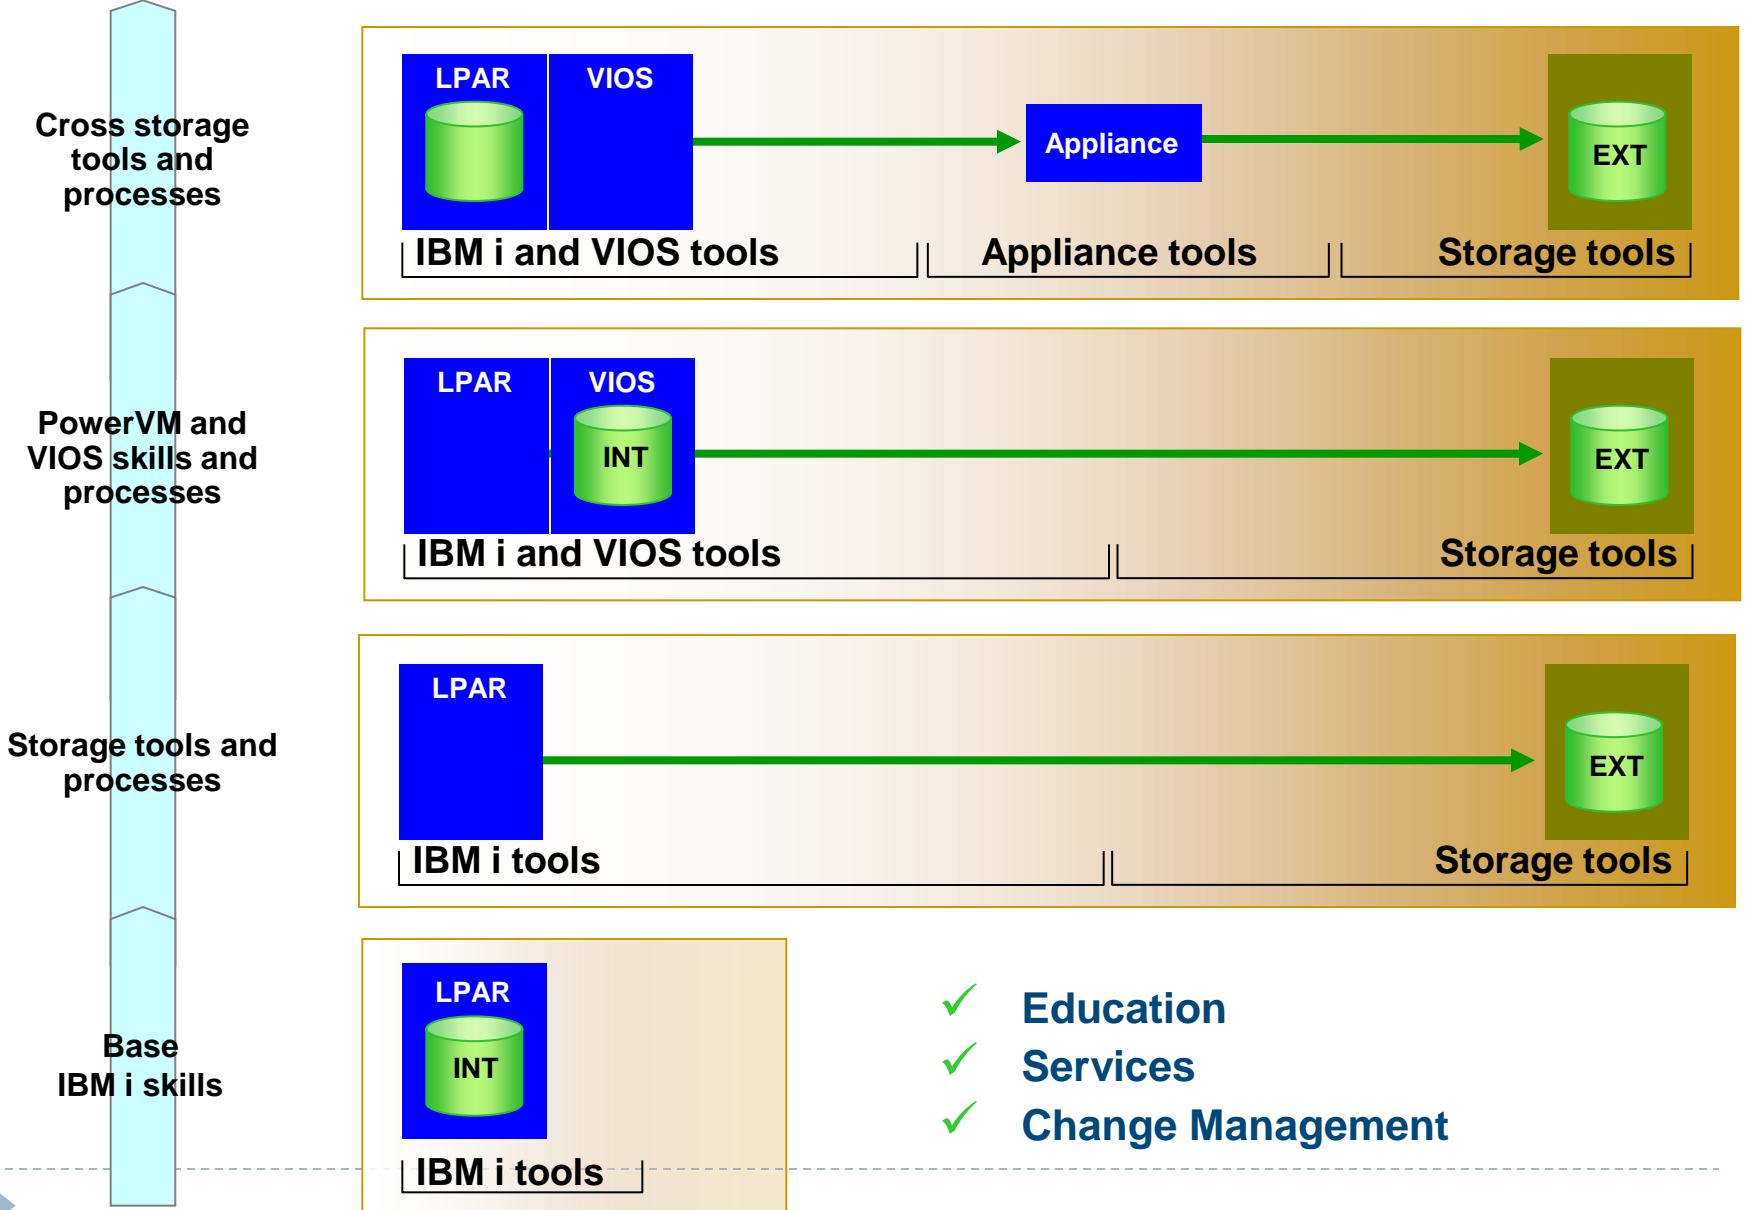
CHARM

PowerHA


Hibernation & Mobility

Data Retention

# Plan for Operations, Systems Management and Skills



# Disk Storage Options for IBM i



# Support for IBM Storage Systems with IBM i

**Table as of  
April 12, 2011**

		N Series @@	DS3200 DS3400 DS3500 DS3950	DS4700 DS4800 DS5020	Storwize V7000	DS5100 DS5300	DS6800	SVC	XIV	DS8100 DS8300	DS8700 DS8800
Rack / Tower Systems	IBM i Version Hardware	<u>5.4 / 6.1 / 7.1</u> POWER5/6/ 7	<u>6.1 / 7.1</u> POWER6/7 Not DS3200#, Yes DS3500##	<u>6.1 / 7.1</u> POWER6/ 7	<u>6.1 / 7.1</u> POWER6/ 7	<u>6.1 / 7.1</u> POWER6/7	<u>5.4 / 6.1</u> POWER5/6/ 7  <u>Not 7.1 ###</u> POWER5/6/ 7	<u>6.1 / 7.1</u> POWER6/ 7	<u>6.1 / 7.1</u> POWER6/7	<u>5.4 / 6.1 / 7.1</u> POWER5/6/7	<u>5.4 / 6.1 / 7.1</u> POWER5/6/7
	IBM i Attach	IFS / NFS (NAS)	VIOS	VIOS	VIOS	Direct* or VIOS – VSCSI and NPIV%	Direct	VIOS	VIOS	Direct or VIOS – VSCSI and NPIV**	Direct or VIOS – VSCSI and NPIV**
Power Blades	IBM i Version Hardware	<u>6.1 / 7.1</u> POWER6/7 IFS / NFS (NAS)	<u>6.1 / 7.1</u> POWER6/7 @, #, ##	<u>6.1 / 7.1</u> POWER6/ 7 (BCH)	<u>6.1 / 7.1</u> POWER6/ 7 (BCH)	<u>6.1 / 7.1</u> POWER6/7 (BCH)	Not supported	<u>6.1 / 7.1</u> POWER6/ 7 (BCH)	<u>6.1 / 7.1</u> POWER6/7 (BCH)	<u>6.1 / 7.1</u> POWER6/7 (BCH)	<u>6.1 / 7.1</u> POWER6/7 (BCH)
	IBM i Attach	IFS (NAS)	VIOS	VIOS	VIOS	VIOS	n/a	VIOS	VIOS	VIOS NPIV**	VIOS NPIV**

## Notes

- This table does not list more detailed considerations, for example required levels of firmware or PTFs required or configuration performance considerations
- POWER7 servers require IBM i 6.1 or later
- This table can change over time as addition hardware/software capabilities/options are added
- # DS3200 only supports SAS connection, not supported on Rack/Tower servers which use only Fibre Channel connections, supported on Blades with SAS
- ## DS3500 has either SAS or Fibre Channel connection. Rack/Tower only uses Fibre Channel. Blades in BCH support either SAS or Fibre Channel. Blades in BCS only uses SAS.
- ### Not supported on IBM i 7.1. But see SCORE System RPQ 846-15284 for exception support
- \* Supported with Smart Fibre Channel adapters – NOT supported with IOP-based Fibre Channel adapters
- \*\* NPIV requires Machine Code Level of 6.1.1 or later and requires NPIV capable HBAs (FC adapters) and switches
- @ BCH supports DS3400, DS3500, DS3950 & BCS supports DS3200, DS3500
- @@ N Series can only be used as file server. No load source/boot support. Support only through IFS. No IBM i data base support
- % NPIV requires IBM i 7.1 TR2 (Technology Refresh 2) and latest firmware released May 2011 or later

**For more details, use the System Storage Interoperability Center: [www.ibm.com/systems/support/storage/config/ssic/](http://www.ibm.com/systems/support/storage/config/ssic/)**

Note there are currently some differences between the above table and the SSIC. The SSIC should be updated to reflect the above information

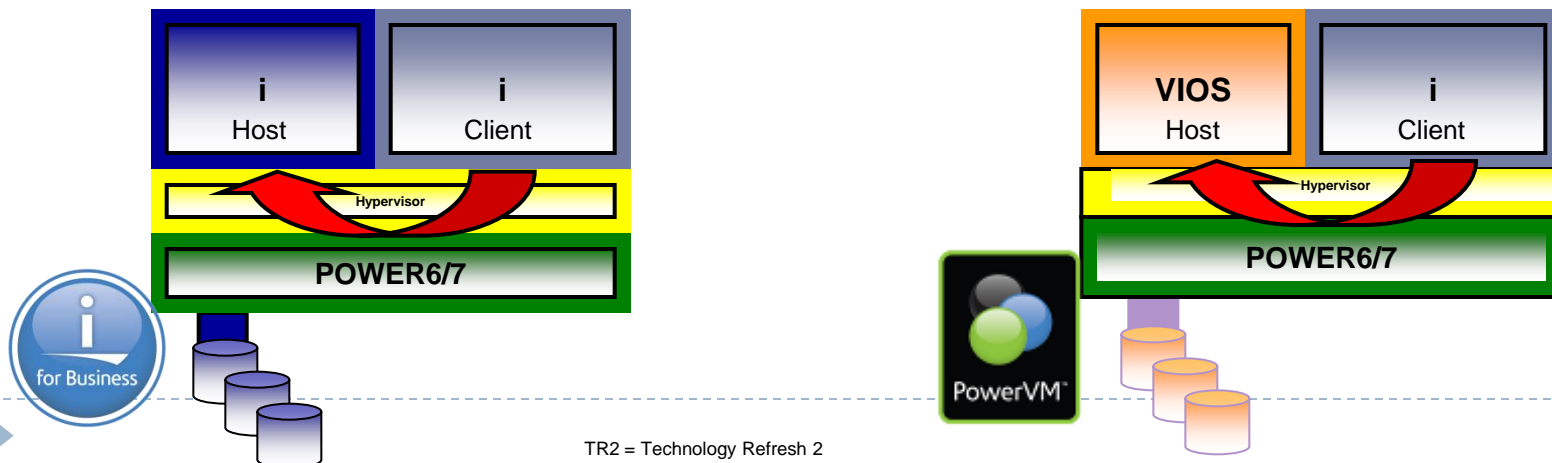
# Virtual Storage for Partitions on Power Servers IBM i

- IBM i hosting

- **IBM i partition uses I/O resources from host IBM i partition**
- Best option for Windows Integration on IBM i
- Familiar IBM i environment
- Limited support of PowerHA Geographic Mirroring (NWSSTG full copy only)
- Supports internal and external storage

- Virtual I/O Server (VIOS) hosting

- **IBM i partition uses I/O resources from a VIOS partition**
- Best environment for IBM i, AIX and Linux
- Typically requires the least amount CPU
- Faster provisioning - fewest setup steps
- No network storage space required
- Simple startup - VIOS is always active
- Provides tape virtualization
- Supports PowerHA-DS8000 Copy Services
- LPAR Hibernation (IBM i 7.1 TR2) new!
- Supports internal and external storage
- NPIV for native attach functionality





# Advanced IBM i Storage Resiliency Design Strategies

Disk protection choice	Level of protection	Relative performance <sup>1</sup>	CHARM <sup>2</sup>
<b>RAID-5 with protected write cache</b>	Basic	Standard	None
<b>RAID-6 with protected write cache</b>	Basic +	Slight degradation	None
<b>RAID-5/6 with protected write cache and hot spare</b>	Better	No effect	None
<b>RAID-5 with dual controllers</b>	Basic	Standard with active-active	Basic
<b>RAID-6 with dual controllers</b>	Better	Improved with active-active	Basic
<b>External (SAN) attached storage with multipath fiber</b>	Basic to Best	Standard to improved	Better to Best
<b>IBM i Disk Mirroring</b>	Best	Improved	None to Basic
<b>VIOS with redundancy</b>	Basic to Best	Standard to improved	None to Best

<sup>1</sup> Assumes the same quantity of physical disk units. R/W ratio of 50/50 to 30/70.

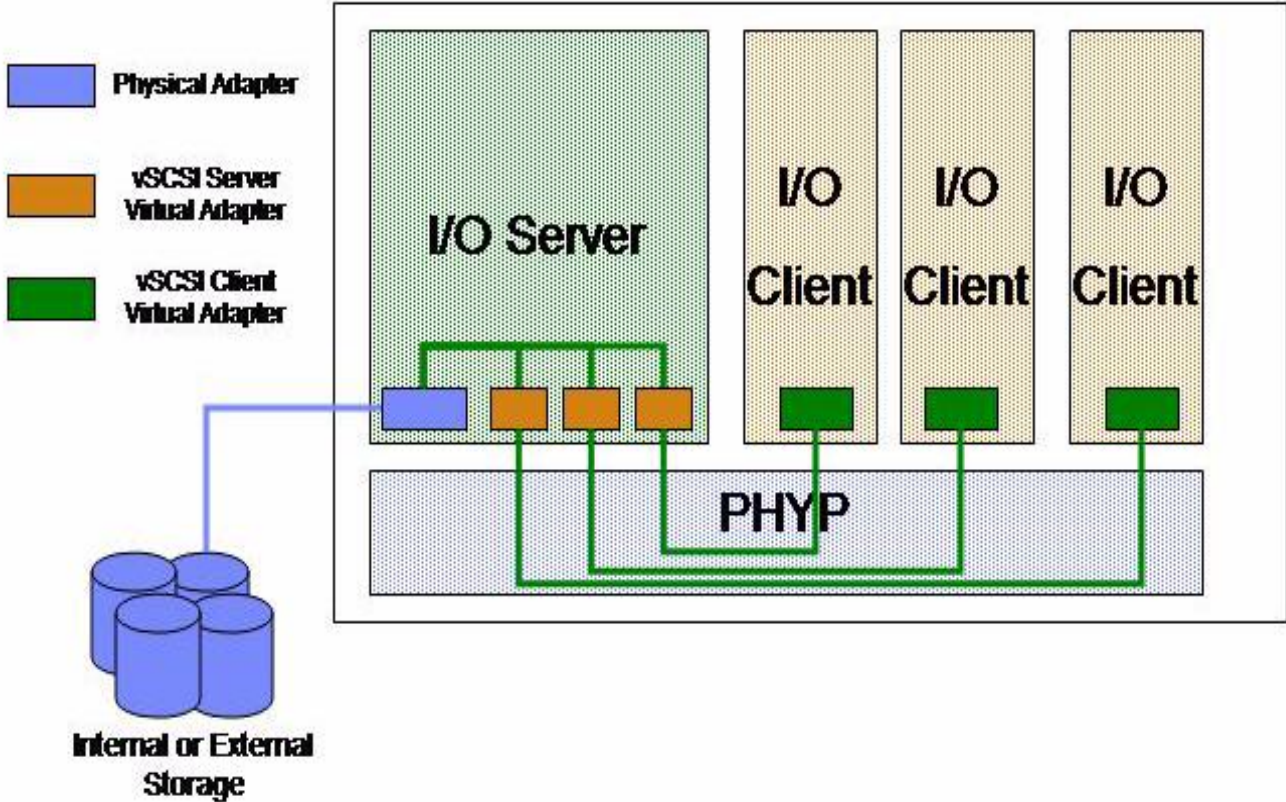
<sup>2</sup> CEC Hot Add and Repair Maintenance. Assumes advance I/O design and operation skills. These are suggested better practices, other options available. Detailed planning required

# 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

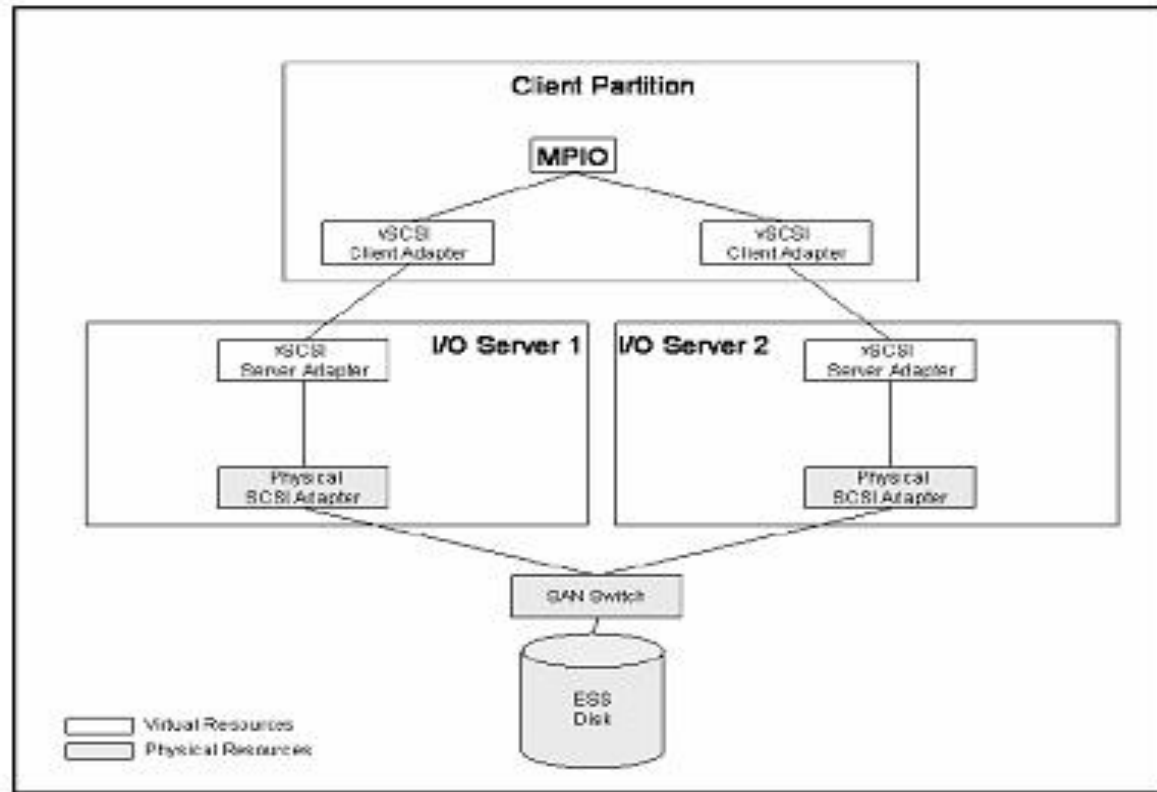


# VIO Server/Client Overview



# VIO Server Configuration with MPIO

- Client sees one hdisk – with two MPIO paths  
lspath -l hdisk0
- Paths are fail-over only. No load balancing in client MPIO
- hdisk1 in each VIO server attached to vscsi server adapter as a raw disk
- Set reserve\_policy attribute on hdisk1 to no\_reserve in each VIO server
- LUN appears in each VIO server as hdisk1
- Single RAID5 LUN carved in ESS, made visible to one Fibre Channel adapter in each of the VIO servers

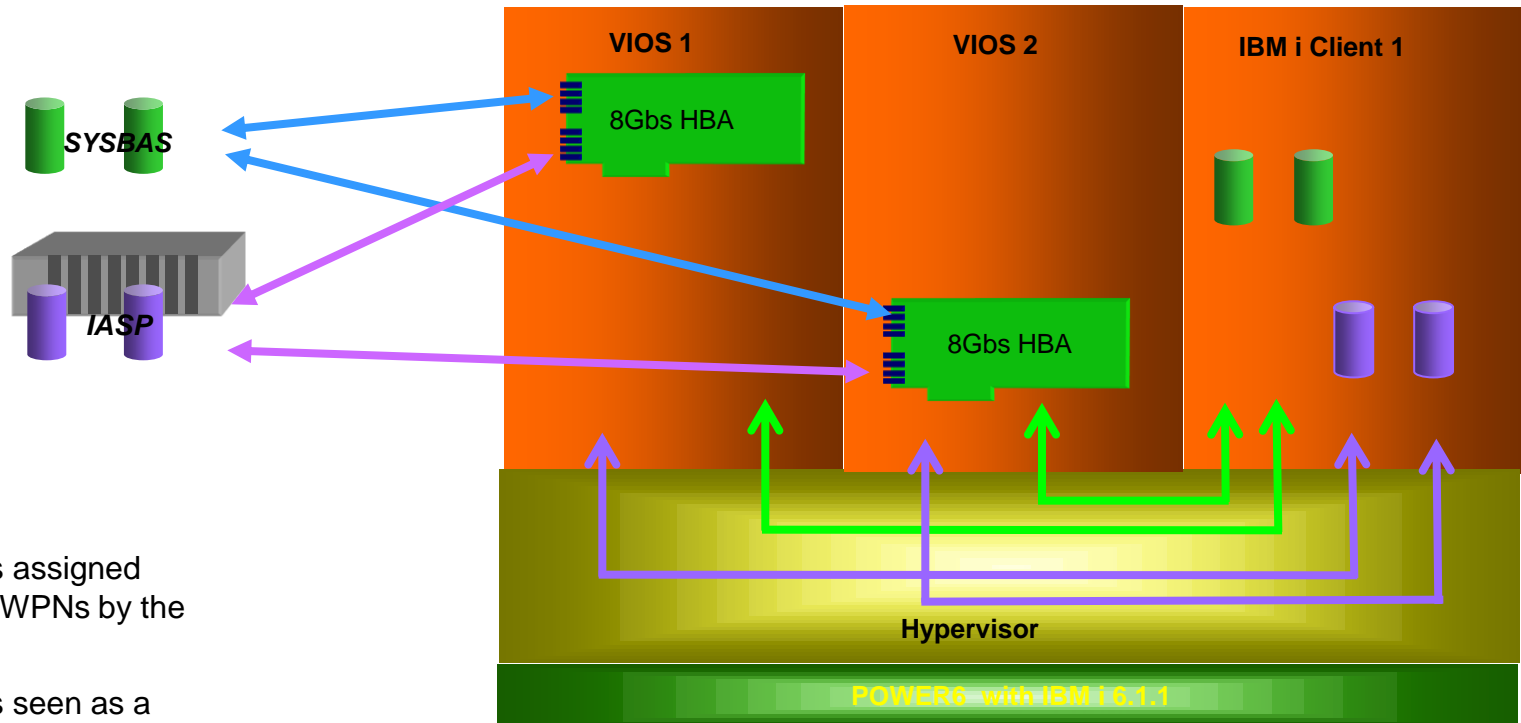


MPIO in client LPAR  
automatically configures



# High Availability Solutions

PowerHA for IBM i plus DS8000 NPV (Virtual Fiber)




- Each port is assigned separate WWPNs by the Hypervisor
- Each port is seen as a separate adapter by IBM i – so PowerHA resets it individually.
- Reduces the hardware for a single partition from 4 to 2 adapters for PowerHA

*Note, This configuration can support up to 64 IBM i partitions without adding any more adapters*

# Scheduled Downtime: Online Backups

## • 6.1 Quiesce Function – reduced disruption backup

- Suspends transactions and operations to ensure as much in-flight data as possible is written to disk before a snapshot is taken
  - More than a memory dump – places transactions at database boundaries if possible
  - Intended for use with applications running commitment control
  - Backup will still be seen as ‘abnormal’ but much friendlier
- Best use is with storage based copy technologies (not direct tape)
  - Use with IASP FlashCopy or Geographical Mirroring solutions
- Command or API support - CHGASPACT \*SUSPEND, or \*RESUME
- Requires assessment and testing, esp. applications without commitment control



Performing a FlashCopy - <http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/index.jsp?topic=/rzaig/rzaigmanageiasp.htm>

# Storwize V7000 Overview

## *Modular Hardware Building Blocks*



## *Software inherited from prior Offerings plus Enhancements*

### *Software inherited from SVC and DS8000 RAID*

- **RAID 0, 1, 5, 6, 10**
- **Storage Virtualization** (Internal and external disks)
- **Non-disruptive Data Migration**
- **Global & Metro Mirror** (Multi-site)
- **FlashCopy** (256 targets, cascaded, incremental...)
- **Thin Provisioning**

### *New and enhanced Software functions*

- **New GUI** (Easy-to-use, web based, XIV like)
- **Easy Tier™ SSD exploitation**
- **RAID & enclosure RAS services and diagnostics**
- **Integration with IBM Systems Director**
- **Enhancements to TPC, FCM and TSM support**

# What is the Storwize V7000

- **Storwize V7000 is a midrange storage system that can also virtualize external storage**
  - Based on proven SVC software technology
- **SAS network for enclosure expansion**
  - SSD, 10K SAS and Nearline SAS
  - Maximum 240 disk drives using 2.5" drives
  - Maximum 120 disk drives using 3.5" drives
  - 120...240 drives depending on combination
- **Dual-active, hot-swappable nodes**
- **Two host interface options**
  - Fibre channel
  - iSCSI
- **Customer installable and maintainable**
- **All primary components are hot-swappable CRUs (Customer Replaceable Units)**





# Storwize V7000 Advanced Functions

- **Thin Provisioning – Included**

- Allows for storage optimization by consuming real capacity only when data is written

- **FlashCopy – Included**

- Allows use of FlashCopy function with internal Storwize V7000 volumes as well as external storage under Storwize V7000 management
- FlashCopy Manager optionally available to help manage

- **Easy Tier – Included**

- Allows for management of hotspots automatically by migrating extents from spinning drives to solid state drives as needed for higher performance
- TPC for Disk Midrange Edition optionally available to help manage

- **Metro and/or Global Mirror – Optional**

- Allows you to replicate data synchronously or asynchronously between Storwize V7000 systems
- TPC for Replication optionally available to help manage

- **External Virtualization – Optional**

- Allows you to bring external fibre channel disk systems under Storwize V7000 control providing access to all the functions of the virtualization software

