

What's the best disk storage for my i5/OS workload?

Sue Baker

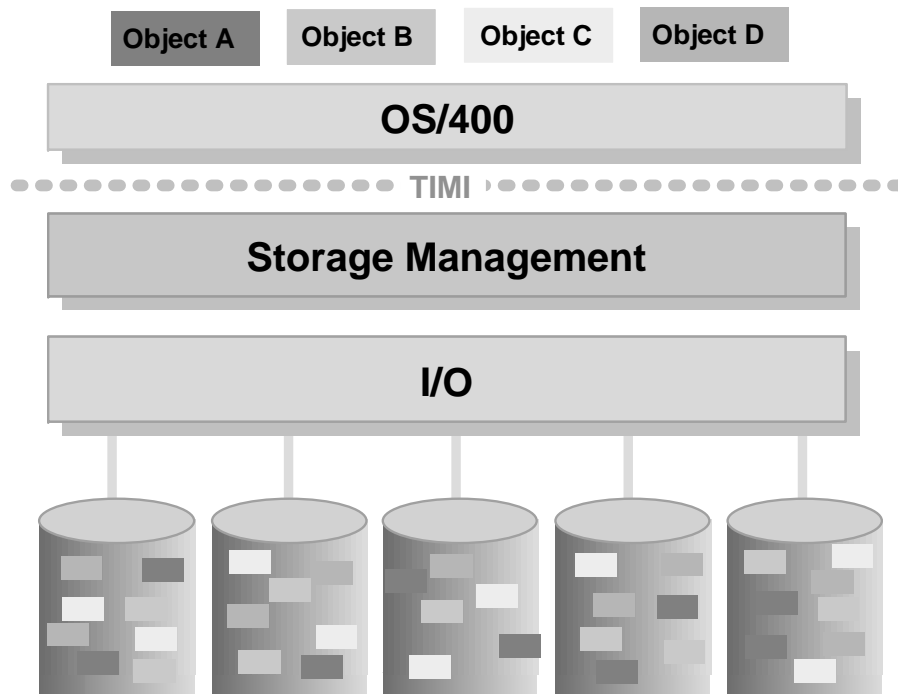
IBM System i Advanced Technical Support

Agenda

- Storage management styles
- Storage technologies for i5/OS
- Considerations for implementing and operating
- Decision scenarios
- Summary

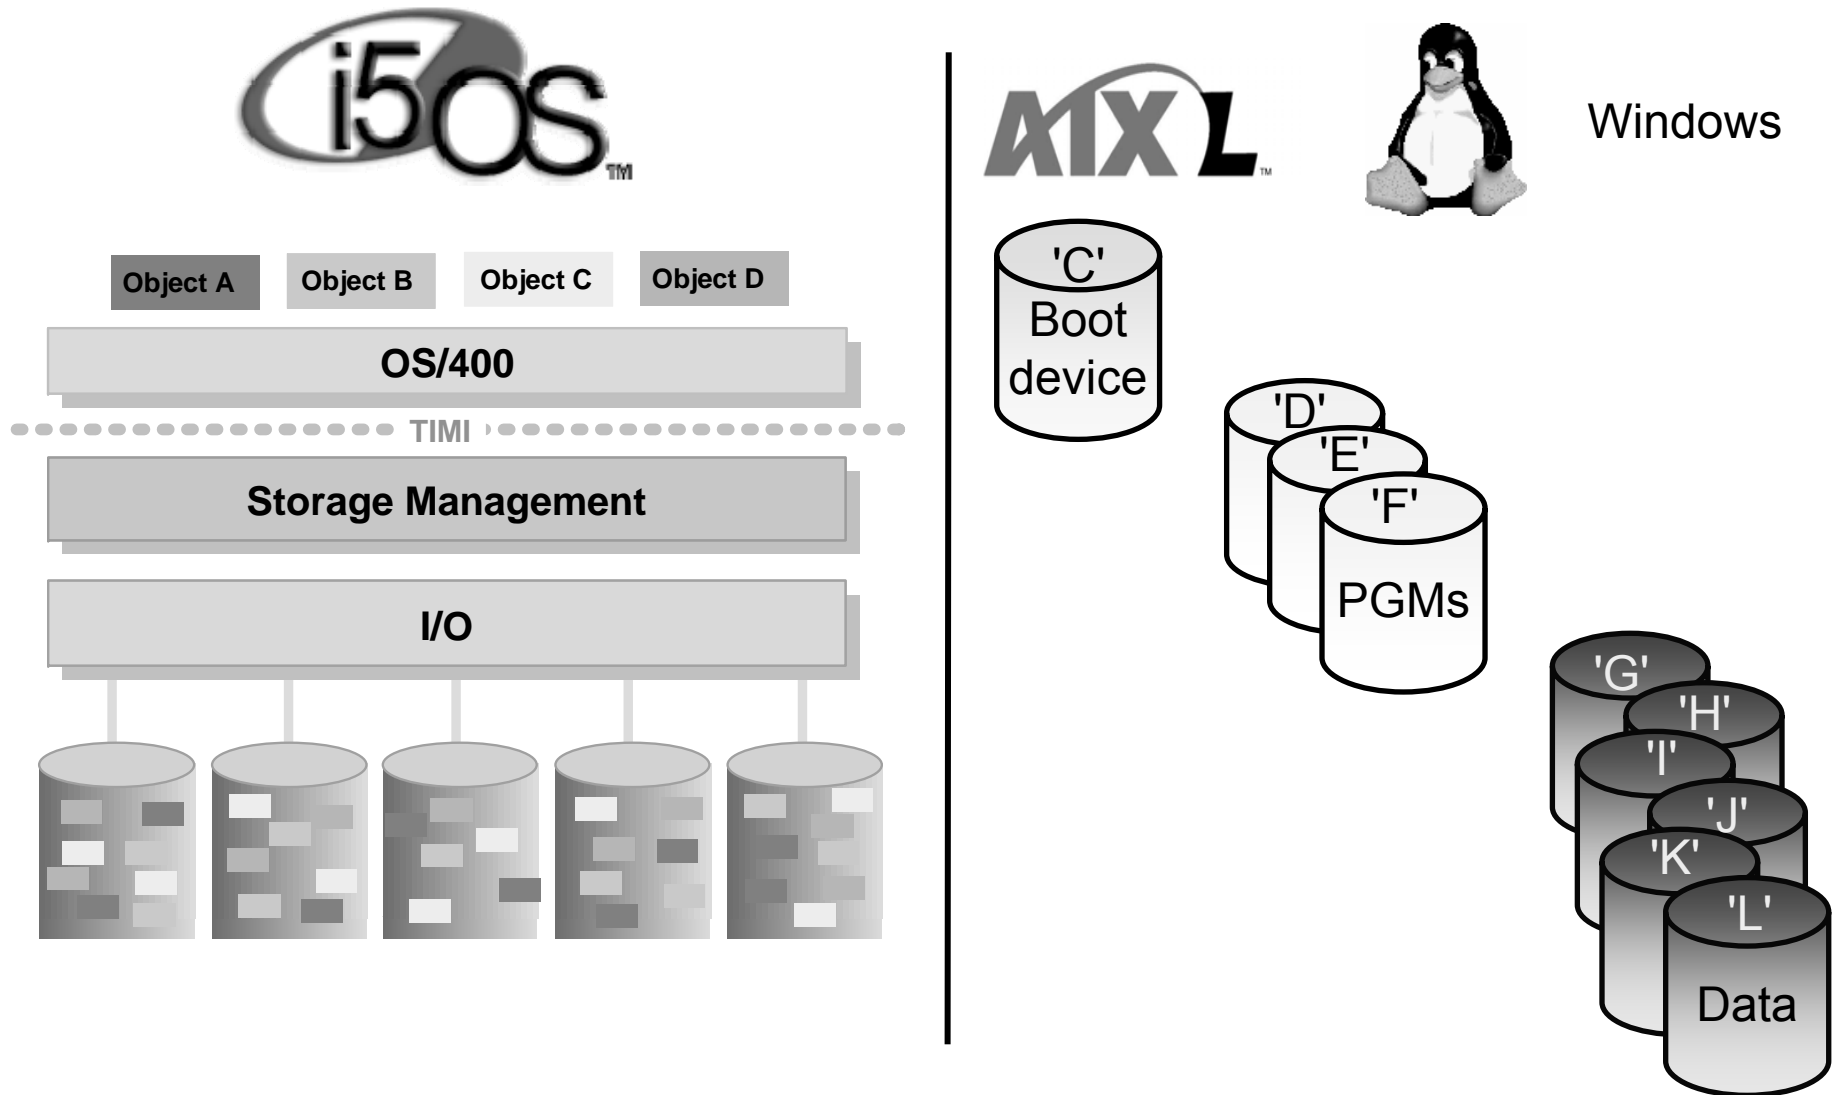


Storage management styles - i5/OS




- Is autonomic
 - i5/OS automatically spreads data across all disks
 - Designed for ease-of-use, self-management
 - Added disk capacity is utilized automatically
 - No continuous disk monitoring needed
- Is easy to operate – single GUI management and wizards
- High performance for all environments

Storage management styles – i5/OS compared to ...




Storage Technologies for i5/OS

Integrated




- Fast, consistent and automatic
- Integration simplicity for i5/OS, xSeries, AIX, and Linux
- Small, medium, large solutions

SAN attached



- Cross server flexibility
- IBM DS4000*, DS600, or DS8000
- EMC only other SAN vendor for i5/OS workloads

Virtualized



- Minimal cost to add partitions
- Requires:
 - V6R1
 - POWER6 processors

- Distributed read and write caches
- Protection via
 - RAID-5 and RAID-6,
 - i5/OS disk mirroring,
 - Hot spare for RAID-5 and RAID-6

- Centralized read and write caches
- Protection via
 - RAID,
 - RAID with hot spare,
 - i5/OS disk mirroring
- SAN = Storage Area Network

- Utilizes main storage as cache
- “Host” is i5/OS or VIOS
- Underlying hardware
 - Provides protection
 - Integrated or SAN attached




Higher performance storage solutions for i5/OS

	Integrated Storage	IBM Systems Storage (SAN, fibre attached)
Maximum capacity	Up to 380 TBs	Up to 308 TBs - multiple subsystems can attach to a each i5/OS partition for up to 380 TBs per System i
Cache	Distributed	Centralized
	<ul style="list-style-type: none"> ■ Write <ul style="list-style-type: none"> ● Up to 1.5GB per controller ● Up to 810 GBs per system ■ Read <ul style="list-style-type: none"> ● Up to 1.6GBs per controller ● Up to 864 GBs per system 	<ul style="list-style-type: none"> ■ Write <ul style="list-style-type: none"> ● Up to 8 GBs ■ Read <ul style="list-style-type: none"> ● Up to 256 GBs
“Volumes” per controller	Up to 36	Up to 32 <i>V6R1 + new IOAs – up to 64</i>
Protection	RAID-5, RAID-6, or i5/OS disk mirroring	RAID-5 or RAID-10 with hot spares - plus i5/OS disk mirroring

System i Integrated Storage “Top to Bottom” Performance

Ultra-High performance integrated storage

- Distributed array level processors and caches provide the most cost effective, scalable, consistent performing solution available
- 70 GBs to 380 TBs scalability
 - Up to 2700 disk units
- i5/OS, Linux, and AIX partitions
- Virtual disks for Windows, Linux, AIX, and i5/OS





Ultra-high performance IOA

- 1.5 GB write cache
- 1.6 GB read cache
- PowerPC 750 Mhz CPU
- Four Ultra-4 SCSI buses
- Up to 540 IOAs on a single System i
 - Up to 810 GB of total write cache
 - Up to 864 GB of total read cache
- Up to 36 disk units per IOA
 - 35, 70, 144, 282 GB
 - 15k disk unit options
- Concurrent maintenance for cache battery

Latest System i disk I/O adapters (IOAs)

- 1.5 GB and EXP24 1.5 GB IOAs – common to both:
 - 1.5GB write cache –
2 times more than #2780 (757 MB)
 - 1.6GB read cache –
1.6 times more than #2780 (1GB)
 - DDR2 memory
 - Faster PowerPC processor
 - Ultra-4 SCSI (up to 320 MB/s transfer speeds)
 - Concurrent battery maintenance
 - Smart IOA – no IOP required when running V5R4 on POWER5 and later systems¹
 - Reduces cost
 - Reduces card slot requirements
 - RAID-5, RAID-6, and i5/OS Disk Mirroring capable
- EXP24 1.5 IOA has built in Redundant Write Cache protection
- Auxiliary write cache is required to support RAID-5 or RAID-6



Latest System i disk enclosures – EXP24

■ IBM TotalStorage EXP24 Disk Enclosure


- More densely packed and more flexibly located in the computer center
 - Saves floor space
 - Up to 240 disk drives in one 2m rack.
(Compared to 90 disk units in previous enclosures)
 - 24 disk slots in just 4U of space
 - Feature codes for use with System i servers – with associated System i terms & conditions
 - Requires i5/OS V5R3 or later
 - Also supported for AIX 5L and Linux partitions
 - Supported on models 800/810/825/870/890 and 520/525/550/570/595
- Uses same 15K RPM disk drives as used in #5094/5294/5095/0595 I/O enclosure & system units



Comparing disk enclosures



#5294 (36U)
 28 PCI slots
 90 disk slots
 15 disk/controller



#0553 (42U)
 24 PCI slots
 192 disk slots
 24 disk/controller

**113% more
 disk in same
 footprint (but
 6U taller)**

i5/OS Load Source Enhancements

- With i5/OS V6R1, larger disk drives are supported as load source devices.

- SCSI – 282.25GB (feature code 4329 or 1269)
- SAS – 283.7GB (feature code 3678)


- V6R1 Load source specify codes

- #0840 for SAS 283GB
- #0841 for SCSI 282GB

- Minimum Load Source size for integrated disk with POWER6 processors is 35GB

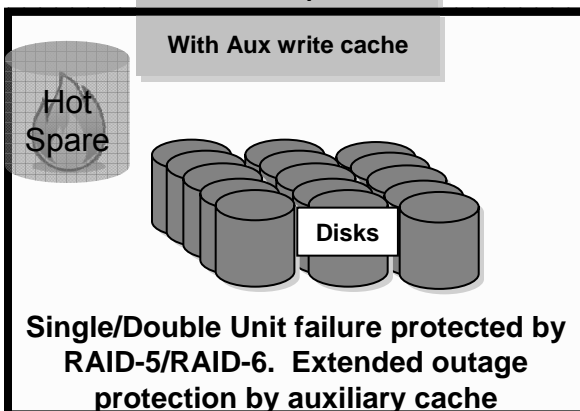

- Load source located on SAN attached storage has a minimum size of 17GB

- Note: as previously announced, IBM Manufacturing no longer requires the i5/OS load source to be in POWER6 570 System Unit

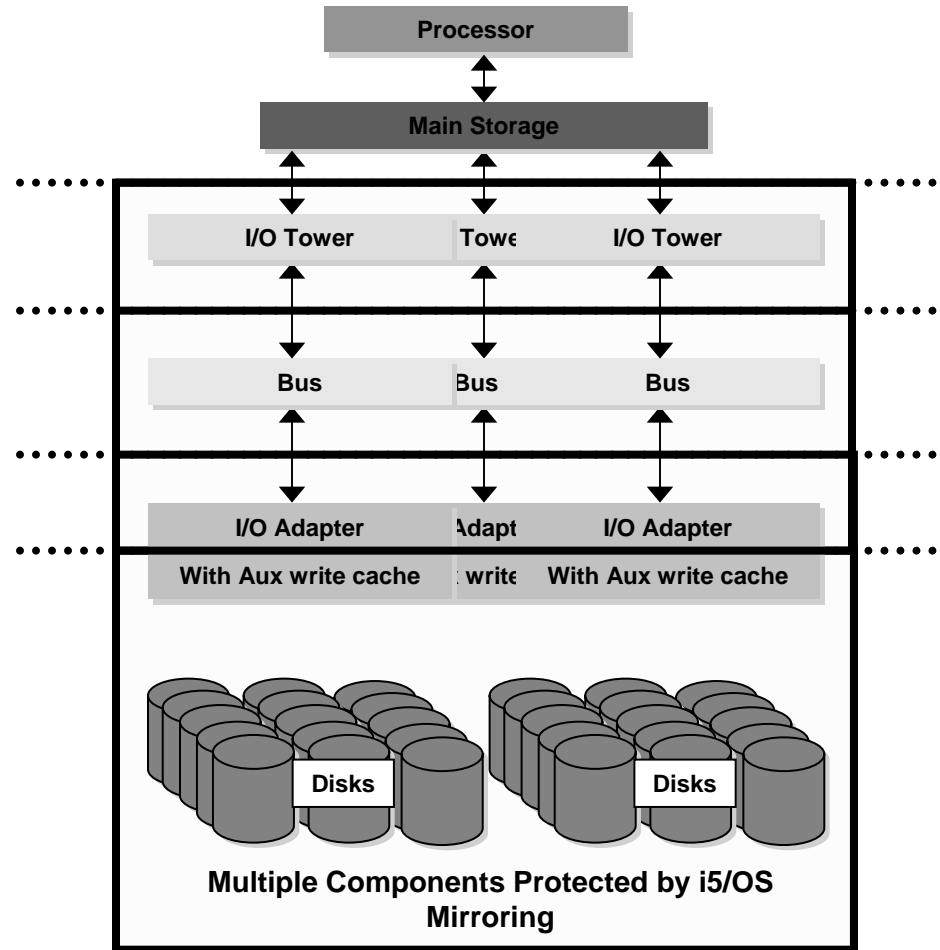


Integrated storage protection options


RAID-5 / RAID-6
plus write cache
protection



i5/OS disk
mirroring



Comparing RAID-5 and i5/OS mirroring hardware requirement




Note: mirroring could require fewer disks than RAID

IBM Systems Storage and System i

Continuous, consistent IBM investment for SAN attached storage on System i

- **4GB Smart IOA Dual Port (Smart IOA)**
- **Improved boot from SAN**
 - *IPL from fibre attached tape*
 - *Multipath for load source on fibre*
- **Tagged Command Queuing**
- **Header Strip Merge**
- **4Gb fiber**
- **Advanced 'DS with i' education**
- **SAN on iSeries Redbook - 3**
- **BRMS support for FlashCopy environment**
- **Boot from San (2847 IOP)**
- **CLI - Command Line Interface**
- **Additional i5/OS LUN sizes for increased flexibility**
- **SAN on iSeries Redbook - 2**
- **i5/OS Multipath fiber**
- **2Gb fiber**
- **iSeries Copy Services for IBM TotalStorage (a.k.a. Toolkit)**
- **Disk Magic and OS/400 Performance Tools coordination**
- **Enhance ESS cache algorithms for iSeries**
- **SAN on iSeries Redbook - 1**
- **1Gb Fiber**

2007



2001

...and more to come....

Flexible SAN Options on POWER6 via Fibre

- **DS8000 Enterprise Storage Solution**
 - Supported on new dual port Smart IOA fibre adapter and earlier IOP required fibre adapters
- **DS6000 Midrange to Enterprise Storage Solution**
 - Not supported on new dual port Smart IOA card, earlier IOP required fibre adapters continue to provide support
- **DS4000 Midrange Storage Solution**
 - Supported via VIOS running on POWER6 with a client i5/OS V6R1 partition



DS8000 Overview

	DS 8100	DS 8300	
Server Processors	2-way POWER5+	4-way POWER5+	
Cache	16 to 128 GB	32 to 256 GB	
Host Ports			
FICON® (4Gb/s) (4 ports per adapter)	8 to 64	8 to 128	
Fibre Channel (4Gb/s) (4 ports per adapter)	8 to 64	8 to 128	
ESCON® (2 ports per adapter)	4 to 32	8 to 64	
Device Ports (2 ports per adapter)	8 to 32	8 to 64	
Drives	16 to 384	16 to 640	
	<u>15K RPM</u>	<u>10K RPM</u>	<u>7.2K RPM</u>
	73 GB 146GB 300GB	146 GB	500 GB
Physical Capacity	1.1 to 192TB	1.1 to 512TB	
Number of Frames	1 to 2	1 to 4	

Note: Achieving maximum TBs requires use of 500 GB 7200 RPM disk units

DS6000 Overview



			DS 6800
Server Processors			PowerPC 750GX 1GHz
Cache			4 GB
Host Ports			
FICON® (2Gb/s) (4 ports per adapter)			8
Fibre Channel (2Gb/s) (4 ports per adapter)			8
ESCON® (2 ports per adapter)			0
Device Ports (2 ports per adapter)			8
Drives			4 to 128
15K RPM	10K RPM	7.2K RPM	
73 GB	146 GB	500 GB	
146 GB	300 GB		
Physical Capacity			.3 to 64TB
Number of drawers			1 to 8

Note: Achieving maximum TBs requires use of 500 GB 7200 RPM disk units

DS400 Overview

	<u>DS4700</u> (follow-on to DS4300)	<u>DS4800</u> (follow-on to DS4500)
Machine/Models	1814-70x,-72x	1814-80x,-82x,-84x,-86x
Drive Types – <i>SATA not supported for i5/OS</i>	SATA or Fibre channel (FC)	SATA or FC
Max Physical Drives	112 drives	224 drives
Max Logical Drives (LUNs)	1024	2048
Total Disk - SATA - Fibre	84-112 TB	168-224 TB
	33-112 TB	67-224 TB
Performance - MB/sec - Disk IO/s - Cache Disk IO/s	990 MB/sec	1240-1600 MB/sec
	44,000 io/sec	62,000-86,000 io/sec
	120,000 io/sec	375,000-575,000 io/sec
Cache	Mdl 70: 2 GB, Mdl 72: 4 GB	Mdl 80/82: 4 GB
		Mdl 84: 8 GB, Mdl 86 – 16 GB
Max DS4K Partitions	128	512
RAID levels	0,1,3,5,7	0,1,3,5,7
	RAID 6 with SM 10.10	(no RAID-6)


i5/OS SAN evolution – where we started



- Minimum SAN configuration requirements
 - One integrated internal disk for load source per system or per LPAR
 - SCSI RAID Adapter for iSeries or IBM eServer™ i5 systems
 - IOP/IOAs for fibre channel SAN connectivity for System Storage
 - Additional LUN on System Storage® protecting integrated internal load source
- Recovery of mirrored load source required during:
 - Primary load source fails
 - Enabling system level (without IASP) functions using System Storage FlashCopy® or Remote Mirror and Copy (PPRC)
- Load source recovery steps requires additional steps, and can take considerable time


Fibre Channel boot from SAN at POWER5

- Internal load source no longer required
- Minimum SAN configuration requires
 - One 2847 IOP for fibre channel i5/OS load source per system or LPAR
 - Two for redundant path protection, with additional LUN defined for mirroring load source inside storage subsystem
 - One fibre IOA for LUN attachment
 - Two for enabling multipath I/O for non-load source LUNs
- Load source recovery back to integrated internal storage not required
- Facilitates easier use of System Storage Copy Services functions

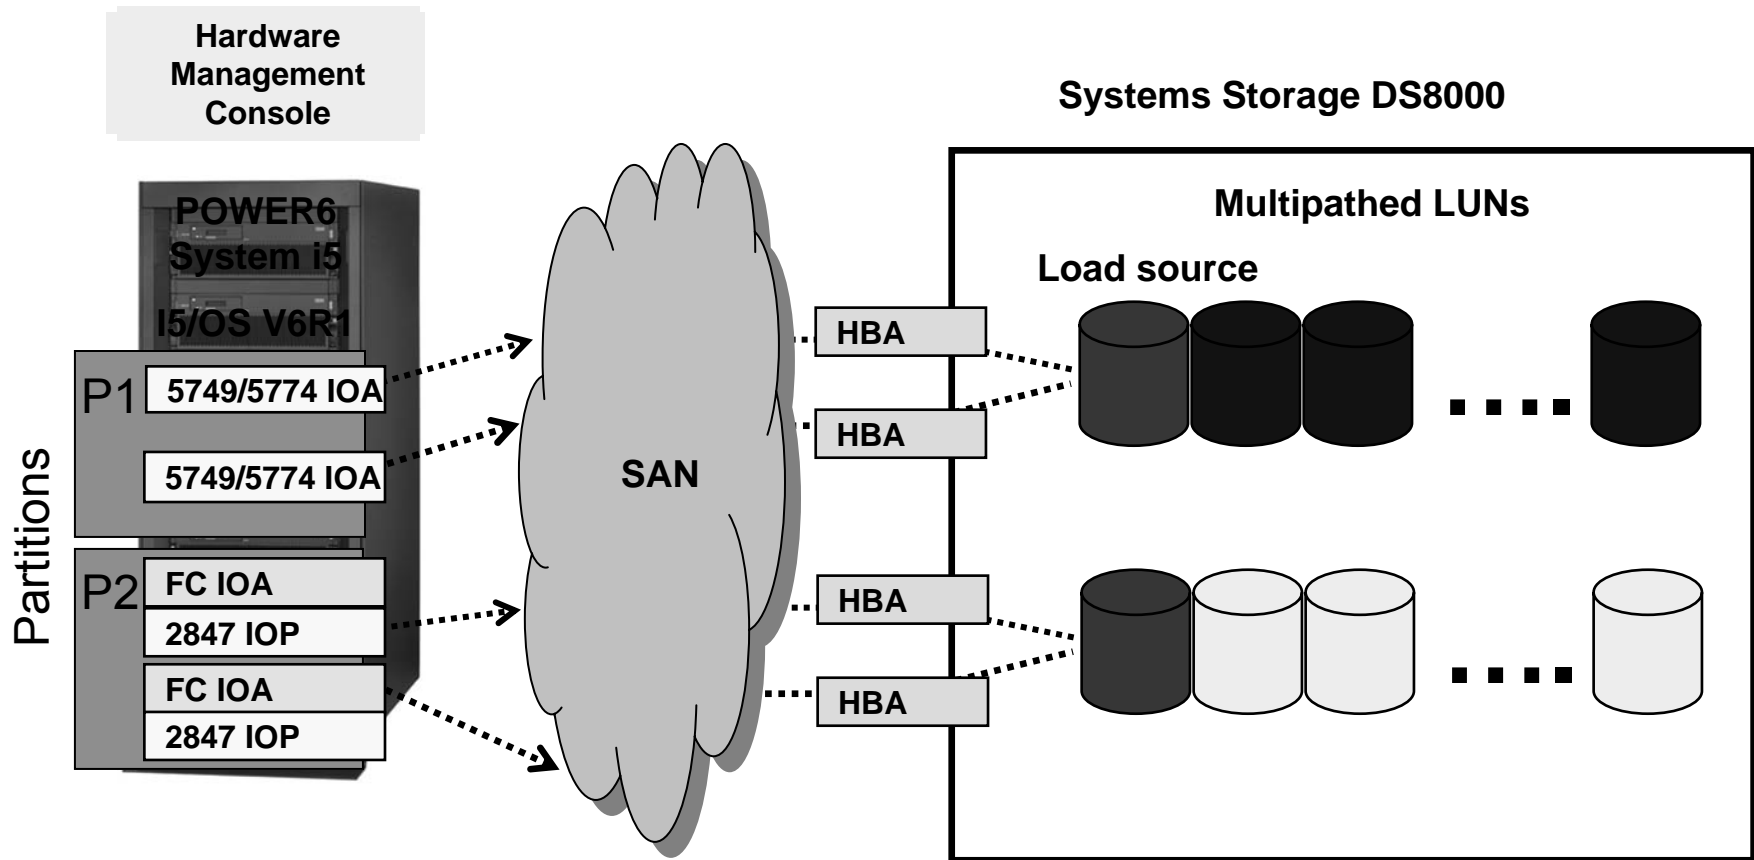


Fibre Channel Load Source for i5/OS V6R1

- Applies to POWER6 and POWER5
- No load source disk required in CEC
- Load Source can be multipathed – remote load source mirroring not required for redundancy
- Minimum SAN configuration requirements
 - One fibre IOA
 - Two IOAs recommended for redundant card protection
- Facilitates easier use of System Storage Copy Services functions



Enabling i5/OS V6R1 Load Source in a SAN



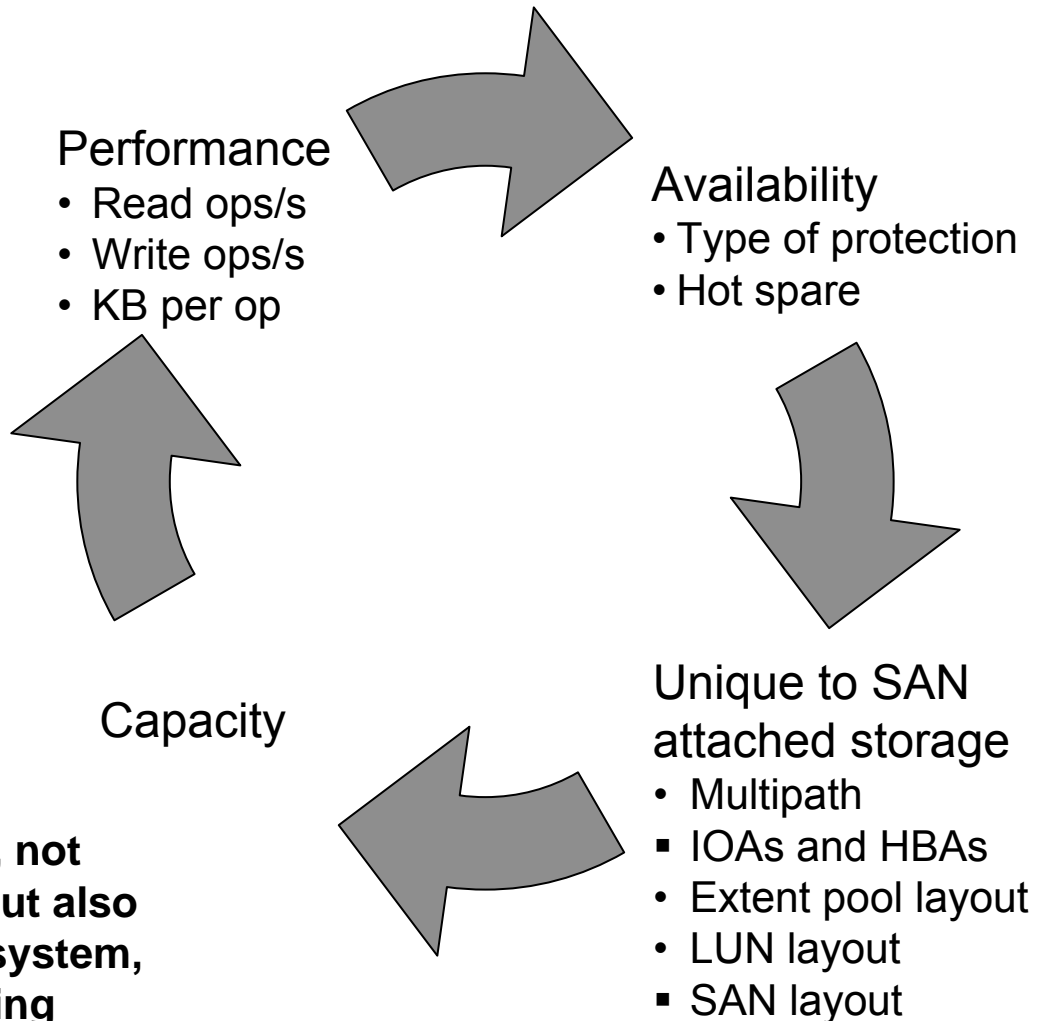
Note:
Smart IOAs do not support ESS 800 and DS6800
IOP based adapters must be in HSL-2 type Expansion

Considerations for implementing and operating

- Sizing
- Planning
- Installation
- Maintaining
- Monitoring
- Troubleshooting

Sizing and Initial Planning Considerations

- For each system or partition
- For each disk pool (ASP)



Note: For SAN attached storage, not only the i5/OS workloads but also each additional operating system, system or partition (including disk pools) must be considered.


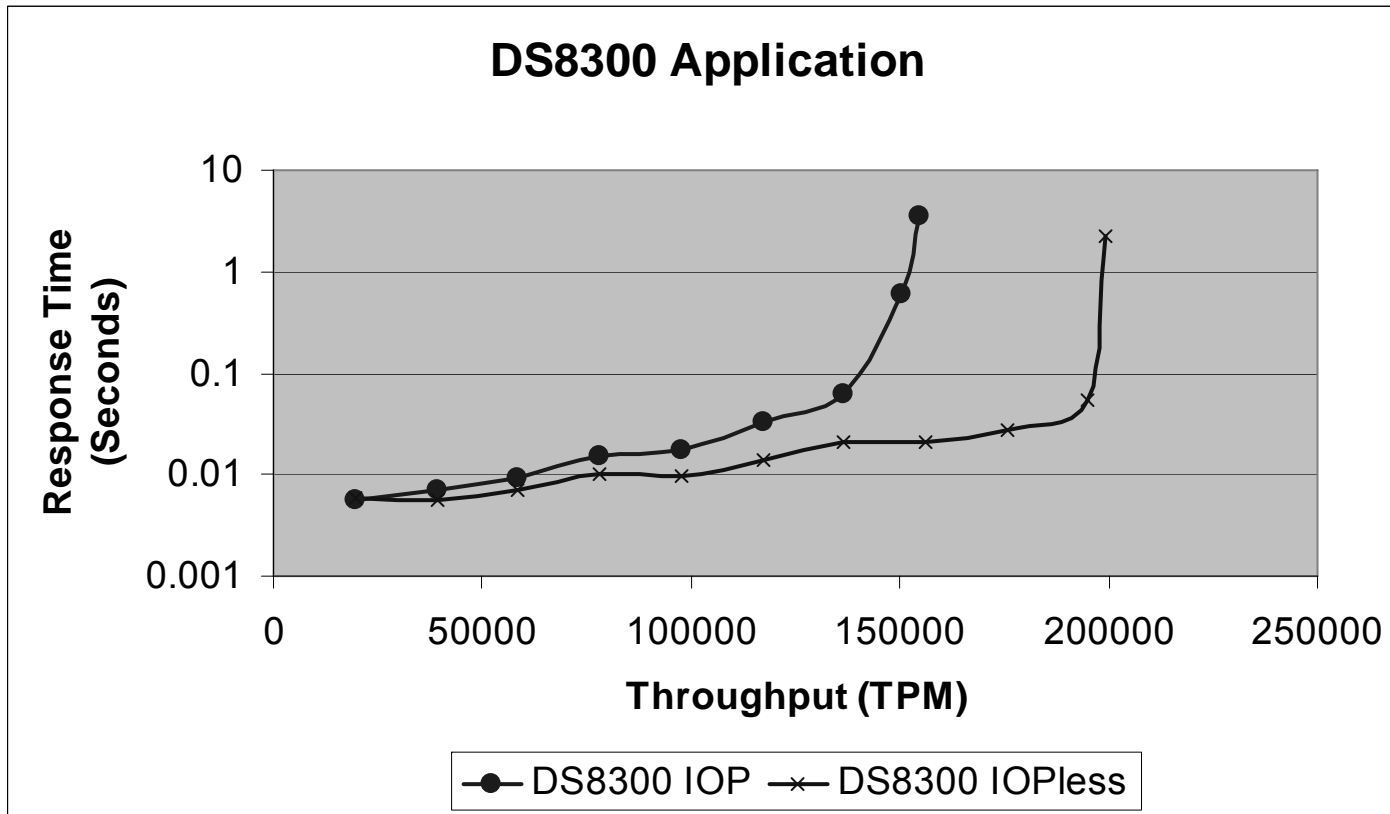
More on Planning plus Installation and Maintenance Considerations

- Installation prerequisites
 - IOA placement
 - PTFs
- Who installs the hardware - IBM service or customer?
 - Integrated storage – either¹
 - SAN attached storage
 - System i IOAs – either¹
 - DS8000 – IBM service
 - DS6800 – IBM service
- Disk protection set up (RAID-5, RAID-6 or Mirroring) – either¹
- Disk protection setup (RAID-5 or RAID-10) and extent pool creation – either¹
- Storage Area Network setup – either¹
- ASP Balancing (when adding to existing hardware) – customer
- Post-install maintenance
 - “Firmware” updates (AKA microcode, LIC)
 - System i – customer
 - DS6800 – IBM service²
 - DS8000 – IBM service (*coordination between IBM service and customer required*)
 - SAN fabric – customer

Monitoring and troubleshooting

- Monitor hardware health with Electronic Service Agent
 - Each system
 - Each partition
 - Each disk subsystem
- Monitor impact of configuration changes for effects on
 - Server
 - Storage response
(Critical in shared storage subsystem environments)
- Monitor performance with
 - i5/OS collection services
 - Performance Tools for i5/OS
 - PM for System i
 - TPC for Disk
 - SAN switches and directors
- Troubleshooting resources
 - HMC logs (System i and DS8000)
 - i5/OS
 - Problem Activity Log (PAL)
 - Performance reports
 - Wait time
 - Service time
 - Utilization
 - Disk subsystem
 - SMC logs (DS6800)
 - TPC reports
 - I/O activity spread
 - NVS full %
 - Rank Read ST
 - Read cache hits
 - Rank activity balance
 - Component overload
 - SAN

Greatly Enhanced SAN Disk Performance for i5/OS



- Configuration the same except the Fibre Channel cards.
- 8 IOP-IOA pairs replaced by 4 smart IOAs.
- Smart IOA performance wins.

Common storage decision scenarios

- You have both storage and high availability decisions to make
- The need for HA, plus the increasing availability of storage based copy solutions may make storage copy services appear a convenient single decision
- HA/DR requirements and their storage requirements should first be defined as separate business problems/goals/strategies to avoid missing business expectations

For HA, this means fully identifying the set of availability problems that you are attempting to address . . .

- Backup window reduction
- Availability for planned outage events (e.g. Hdw/OS/App upgrades)
- Availability for unplanned outage events
- Recovery from disaster related outage events
- Workload balancing

.... Recovery Time Objective (RTO) and Recovery Point Objective (RPO) must be understood

For storage, this means fully identifying the set of storage problems that you are attempting to address . . .

- Near term need to move storage between servers
- Corporate requirement for consolidated storage systems
- High performance vs. 2nd tier, reduced cost storage for archive/compliance
- Disk based copy services and related costs


.... for each application, for each server and partition, and for each operating system

System i Multi-Server Availability

Switched IASP Clusters


- Strategic building block for System i High Availability and Disaster Recovery solutions
- Utilizes i5/OS Cluster Resources for automation and application interface

Switched IASP




- Simple availability
- Single copy of data
- Local only
- Storage agnostic

Cross Site Mirroring



- i5/OS copy of IASP
- HA and DR
- Enhanced with i5/OS V5R4
- Storage agnostic


Advanced SAN Copy Services



- Switched IASPs plus SAN Copy Services
- HA and DR


Storage Replication Clusters

HA-ISV Logical Replication



- Widely deployed
- Hardware flexibility
- HA and DR, with concurrent access
- Supports i5/OS Cluster
- Storage agnostic

Basic SAN Copy Services



- Full System Level SAN copy
- Use SAN boot
- i5 models only
- Data copy or simple DR

Disk storage summary

Disk protection choice	Level of protection	Relative performance	What is protected
RAID-5	Basic	Standard	Single disk unit failure within a parity set.
RAID-6	Basic +	Slight degradation	Multiple disk unit failures within a parity set.
RAID with Auxiliary Write Cache	Intermediate	No effect	Same protection as RAID choice, plus reduced risk of write cache data loss. Outage occurs with disk IOA failure.
RAID with Auxiliary Write Cache and hot spare	Intermediate +	No effect	Same protection as RAID choice, plus reduced risk of write cache data loss. Plus reduced window of multiple drive failure risk. Outage occurs with disk IOA failure.
External storage with multipath	Intermediate to Best	Standard to improved	Redundancy within storage subsystem minimizes outage and data loss due to multiple component failure in subsystem. Multipath minimizes outage due to IOA or path failure.
i5/OS disk mirroring	Best	Improved	Potentially multiple hardware components including HSL, IO tower, IOP, IOA, and disk with no outage required during component failure.

The e nd, Thank You!

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: i5/OS, i5/OS logo, AIX, AIX 5L, BladeCenter, Blue Gene, DB2, e-business logo, eServer, IBM, IBM Logo, Infoprint, IntelliStation, iSeries, pSeries, OpenPower, POWER5, POWER5+, Power Architecture, TotalStorage, Websphere, xSeries, z/OS, zSeries

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation 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.

UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

NOTES:

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. Users of this document should verify the applicable data for their specific environment.

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

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

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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.