IBM FlashSystems

Simply efficient, simply consistent, simply reliable

Herve Guerin
Architecte @ Tech Data
Herve.guerin@azlan.com

Evolution des composants de stockage

Du disque, au SSD, au Flash

IOPS

Latency

Flash Optimized

Flash Adapters
High IOPs Performance

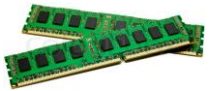

IBM FlashSystem
High capacity optimized flash storage system

100 Micro-Seconds

< 1 Milli-Seconds

Disk Form Factors

Solid State Disk
I/O performance in HDD form factor

‘Hybrid’ HDDs with flash
Improved I/O with spinning media capacities

15k SAS Hard Disks
High performance spinning media

7.2k SAS Hard Disks
Capacity / Cost

5-15 Milli-Seconds

Customer Reference

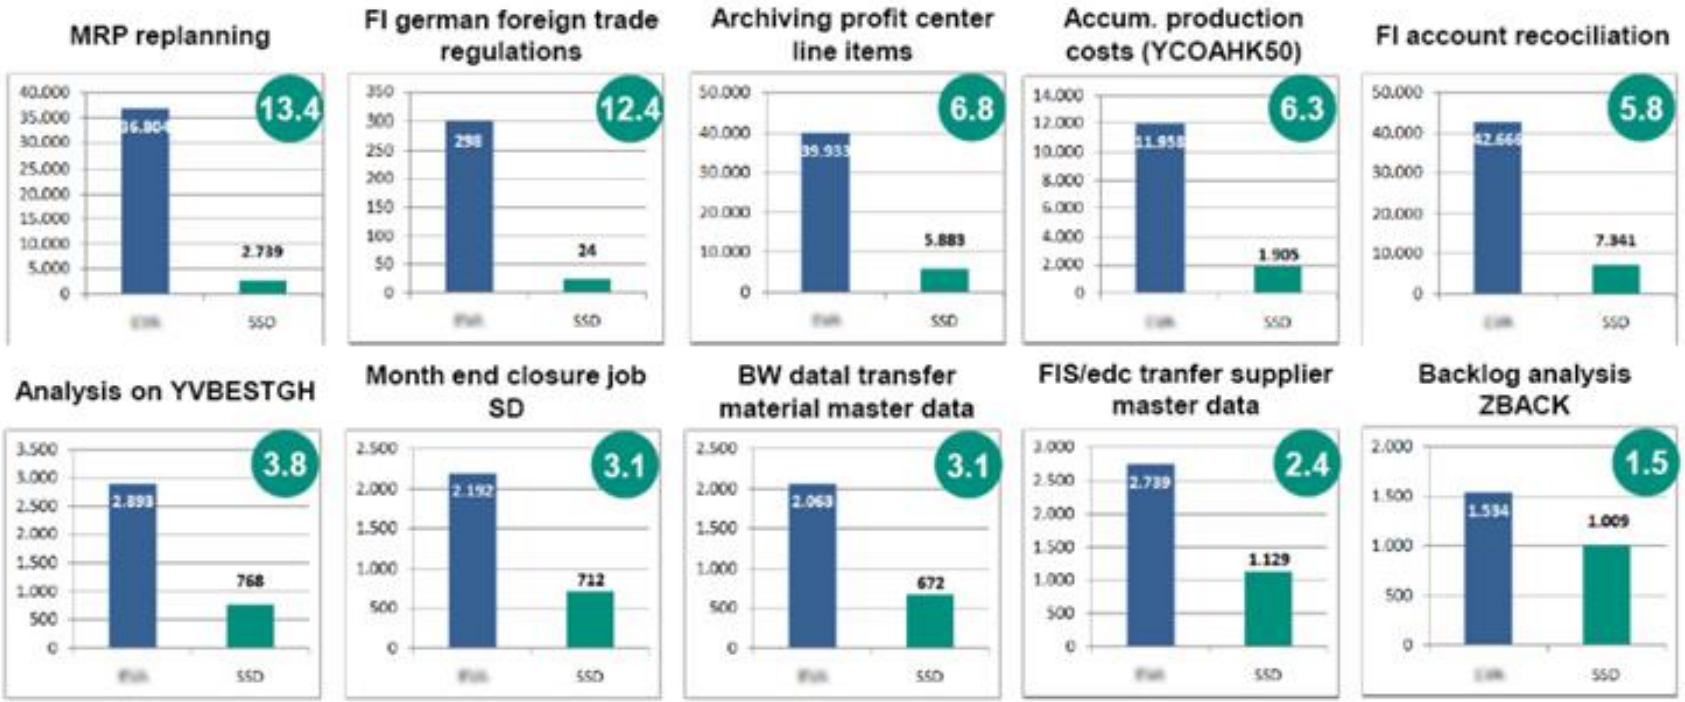

IBM FlashSystem Within The IBM Systems Enterprise Storage Portfolio

IBM FlashCore™ Technology Optimized

FlashSystem A9000

All flash array for cloud service providers

- Best performance with full time data reduction
- Targeting VDI and VMware

FlashSystem A9000R

All flash array for large deployments

- Best performance with full time data reduction
- Targeting mixed workloads

FlashSystem V9000

All flash array for virtualizing the hybrid Data Center

- Best performance with storage services & selectable data reduction
- Targeting database/ analytics workloads

FlashSystem 900

All flash array for application acceleration

- Extreme performance
- Targeting database acceleration & Spectrum Storage booster

IBM Flash Array

Storwize V7000F

Mid-Range

Storwize V5030F

Entry / Mid-Range

SVC

DS8880

High End Server
- **Mainframe**
- **Power**

- Extreme reliability and replication
- Available in All Flash & Hybrid configurations

IBM Spectrum Accelerate

IBM Spectrum Virtualize

IBM FlashSystem Within The IBM Systems Enterprise Storage Portfolio

IBM FlashCore™ Technology Optimized

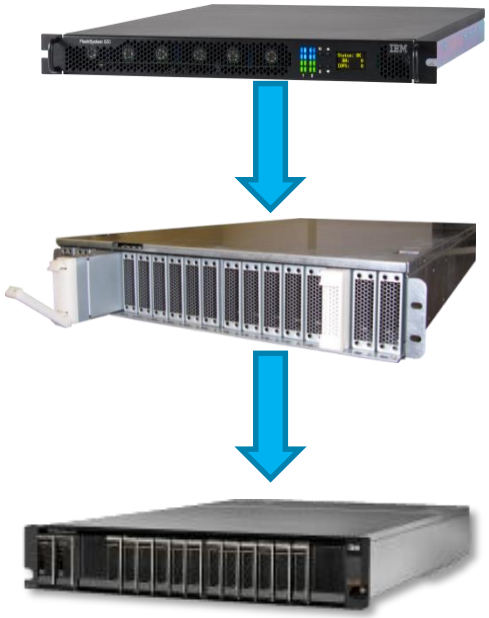

2012 Texas Memory Systems

2014 FlashSystem 840

2015 FlashSystem 900

SW-Based Designs: SW Controlled Data Path

DSK8 XIV V7K | UVSP SVP | 3PAR | VNX VMAX | EF540

IBM | HITACHI Inspire the Next | hp | EMC² | NetApp

What you get:

A consolidation play
Ease of integration
Built-in Tiering
Hot-spots
Another tier to manage
Relative ease of use
Marginal application efficiencies
Proven resiliency

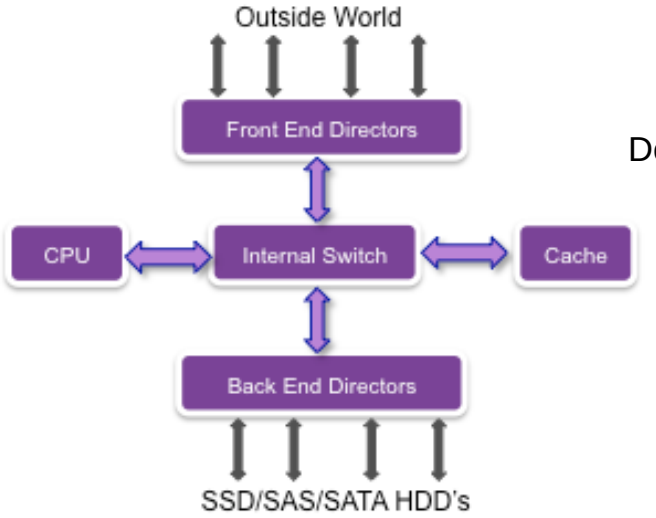

SSD

What you don't get:

Best Latency / Best Performance
Deterministic/Consistent Performance
Architecture design for Flash
Reduced footprint
Maximum application efficiency
Rack space reductions
Power/Cooling savings
Simplicity
Lesser Tiers
Zero Tuning

HW-Based Designs: HW controlled data path

FlashSystem

- Maximum performance and lowest latency by:
  - Optimized FPGA HW data path
  - Custom HW design
  - Massive parallelism handling data
  - 90us write latency into flash media

High Speed track

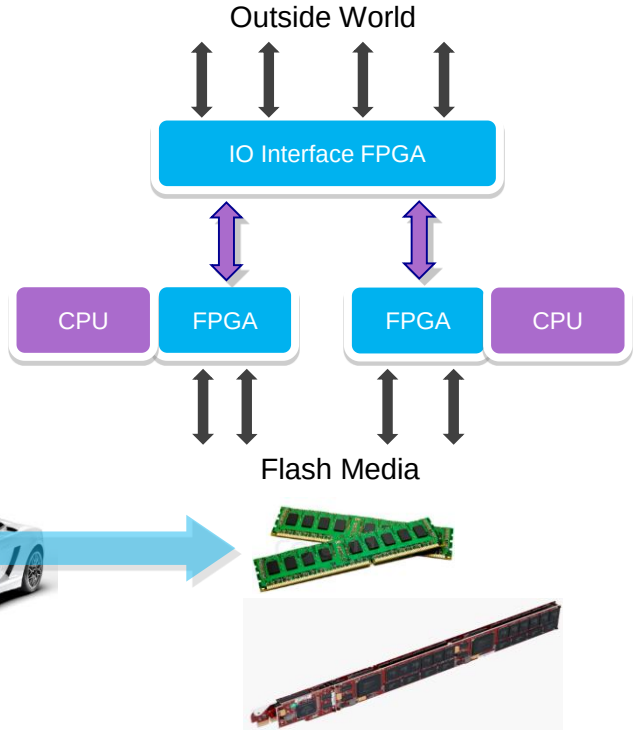

FlashSystem 900

STORAGE ALLIANCE
MICRON VALUED PARTNER | ECOSYSTEM

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

- IBM MicroLatency™ with up to 1.1 million IOPS
- 40% greater capacity (for 2TB & 4TB modules)
- IBM FlashCore™ technology, our secret sauce

Technical collaboration with Micron Technology, our flash chip supplier

- IBM enhanced flash technology
- MLC NAND flash (20nm) offering with Flash Wear Guarantee
- Extra storage capacity: ~50% ➔1.78 TB, 4.39 TB et 8.79TB

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

Module type	1.2 TB					2.9 TB				5.7 TB			
Modules	4	6	8	10	12	6	8	10	12	6	8	10	12
Raw capacity	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
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
RAID 5 capacity (TiB)	2.18	4.37	6.55	8.73	10.91	10.55	15.82	21.1	26.37	20.73	31.1	41.47	51.84

IBM s'engage à remplacer tout module défectueux, sans frais supplémentaire, pendant la durée de maintenance d'une baie FlashSystem, **quelque soit la workload**.

7 ans de garantie

Flash Wear Guarantee: ***"You're covered, no matter what your wear"***

7 ANS GARANTIE

IBM FlashSystem V9000

IBM FlashCore™ Technology Optimized

FlashSystem V9000

All flash array for virtualizing the hybrid Data Center

- Best performance with storage services & selectable data reduction
- Targeting database/ analytics workloads

FlashSystem 900

All flash array for application acceleration

- Extreme performance
- Targeting database acceleration & Spectrum Storage booster

SVC

IBM Spectrum Virtualize

Two-dimensional scaling

Scalable Performance

Maximum Performance Scaled Out (100% Read, 4 Building Blocks) *	Latency (4K)	200µs
	IOPS (4K)	2,520,000
	Bandwidth (128K)	19.2 GB/s
Maximum Performance Scalable Build Block (100% Read, Cache Miss)	Latency (4K)	200µs
	IOPS (4K)	630,000
	Bandwidth (128K)	6.2 GB/s
Maximum Performance Fixed Build Block (100% Read, Cache Miss)	Latency (4K)	200µs
	IOPS (4K)	526,000
	Bandwidth (128K)	6.2 GB/s

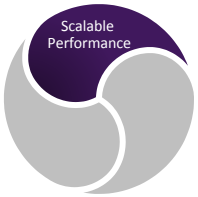

FlashSystem V9000 options for deployment

Scalable Performance

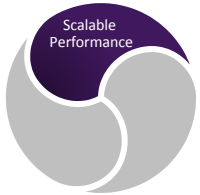

Small Data Center

- Great for database and VDI workloads
- Accelerate, migrate, tier, clone, snapshot, replicate, compress existing storage
- Up to 630K IOPS, 200µs
- Up to 57 TB usable, 285 TB effective

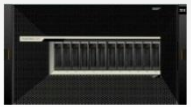

Mixed Workload Accelerator

- Great for multiple mixed workloads that drive huge I/O
- Scale out for more all flash capacity, IOPS and bandwidth
- Up to 2.5M IOPS, 200µs
- Up to 228 TB usable, 1.1 PB Effective

Public or Private Cloud

- Great for Tier 1 disk replacement
- Scale up and out for more processing, more capacity and more I/O
- Up to 2.5M IOPS, 200µs
- Up to 456 TB usable, 2.2 PB Effective

Application Accelerator

- Great for large data sets with big I/O requirements and needing storage services
- Scale up for more all flash capacity
- Up to 630K IOPS, 200µs
- Up to 285 TB usable, 1.4 PB Effective

Virtualized Data Center

- Great for data centers with heterogeneous storage
- Extends core feature set to other storage arrays
- Up to 2.5M IOPS, 200µs
- Up to 3 Exabytes virtualized

IBM FlashSystem Within The IBM Systems Enterprise Storage Portfolio

IBM FlashCore™ Technology Optimized

FlashSystem A9000

All flash array for cloud service providers

- Best performance with full time data reduction
- Targeting VDI and VMware

FlashSystem A9000R

All flash array for large deployments

- Best performance with full time data reduction
- Targeting mixed workloads

FlashSystem 900

All flash array for application acceleration

- Extreme performance
- Targeting database acceleration & Spectrum Storage booster

IBM Spectrum Accelerate

IBM FlashSystem A9000

Interconnect

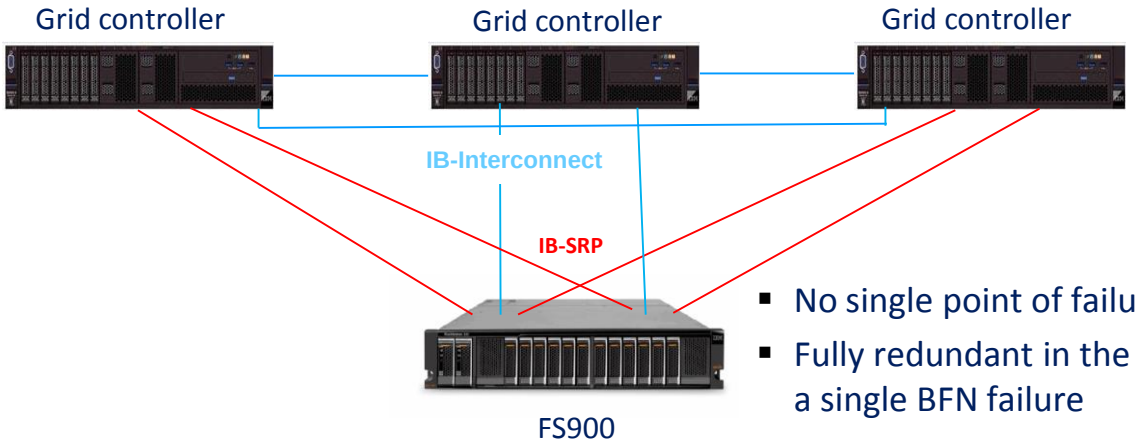

- No single point of failure
- Fully redundant in the event of a single BFN failure

- The Intel based compute server as well as the Flash storage enclosure are interconnected via **IB infrastructure**.
- The connection between the servers are at FDR 56Gb rate using proprietary **RDMA based XIV protocols**.
- The connection to the Flash storage is done at QDR 40Gb rate using **SCSI protocol over IB (SRP**).

Introducing IBM FlashSystem A9000

A highly parallel all-flash platform for the cloud-scale business

FAST

- IBM FlashCore™ technology
- Consistent low latency even with full-time data reduction

CONSOLIDATED

- Hyper-Scale framework
- QoS supports mixed workloads
- Secure multi-tenancy

EFFICIENT

- Flash-optimized pattern removal, deduplication and compression lower TCO
- Thin provisioning and snapshots maximize storage efficiency

SECURE

- IBM FlashCore technology with IBM Variable Stripe RAID™ protection
- Five 9's availability
- Survive controller failures without performance degradation
- Hardware-based data at rest encryption (SKLM / KMIP)

FEATURE-RICH

- Intuitive user interface
- Redirect on Write Snapshots
- Asynchronous Replication
- Synchronous Replication
- VMware integrated
- Linked to OpenStack
- REST API

Introducing IBM FlashSystem A9000

Simply efficient, simply consistent, simply reliable

- 8U modular offering
- Composed of 3 grid controllers & 1 flash enclosure
- Scales via IBM Hyper-Scale Manager

	Flash Enclosure-60	Flash Enclosure-150	Flash Enclosure-300
Effective Capacity [1]	60 TB	150 TB	300 TB
IBM MicroLatency® modules	12 x 1.2 TB	12 x 2.9 TB	12 x 5.7 TB
IOPS[2]	Up to 165,000		
Minimum latency	250 µs		
Host interface	8 or 16 Gb Fibre Channel + 10 Gb iSCSI, or 10 Gb iSCSI		

[1]Based on a 5.26 to 1 data reduction ratio; [2]70/30 read write

Flash-optimized data efficiency suite optimizes economics

Enduring economics

Designed together for comprehensive & complementary reduction

PATTERN REMOVAL	DEDUPLICATION	COMPRESSION	THIN PROVISIONING	SNAPSHOTS

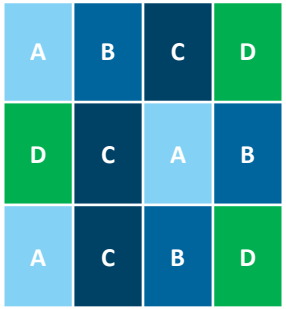

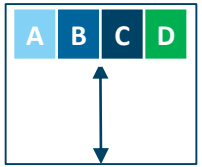

- Supports scalable workloads
- Global, inline
- 8K block size w/4K alignment

- Hardware accelerated
- No performance-degrading background scrubs
- Inline

- Enables high utilization efficiencies

- Redirect-on-write
- Space-efficient
- High-performing

Introducing IBM FlashSystem A9000R

A grid-scale, all-flash storage platform designed to drive your business into the cognitive era

FAST

- IBM FlashCore™ technology
- Consistent low latency even with full-time data reduction

EFFICIENT

- Flash-optimized pattern removal, deduplication and compression lower TCO
- Thin provisioning and snapshots maximize storage efficiency

FEATURE-RICH

- Intuitive user interface
- Redirect on Write Snapshots
- Asynchronous Replication
- Synchronous Replication
- VMware integrated
- Linked to OpenStack

SCALABLE

- Scale out to multiple petabytes to quickly accommodate your growth
- Hyper-Scale framework

AUTOMATED

- Grid-scale automatically redistributes across all resources available
- No hot spots, zero tuning

SECURE

- IBM FlashCore technology with IBM Variable Stripe RAID™ protection
- Five 9's availability
- Survive controller failures without performance degradation
- Hardware-based data at rest encryption (SKLM / KMIP)

Performance & capacity scale together in a single rack

A9000R *Agile integration*

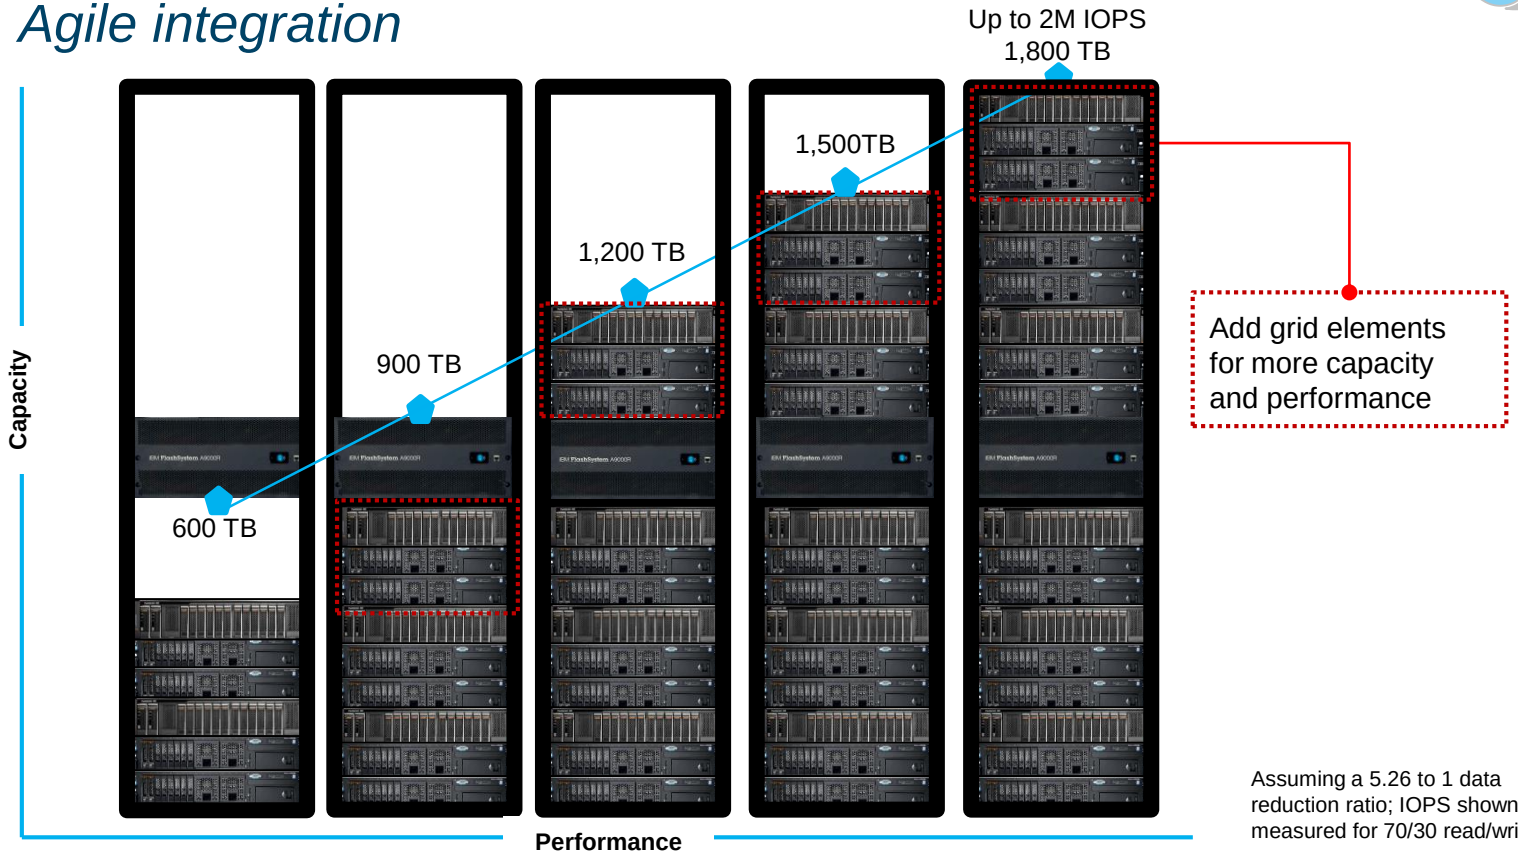

Assuming a 5.26 to 1 data reduction ratio; IOPS shown measured for 70/30 read/write

FlashSystem A9000 & FlashSystem A9000R

Purpose-built for the cloud

IBM FlashCore™ Technology

Differentiate your offering

Secure Multi-Tenancy

Keep your tenants happy

Data Reduction

Optimize capacity for any workload

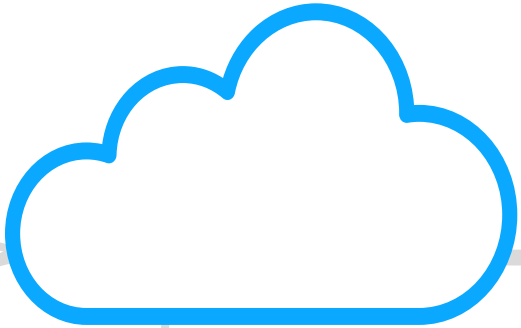

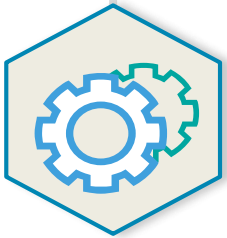

- Microsoft
- VMWare
- OpenStack
- REST API

Cloud Integration

Integrate easily with what you have, and what's to come

Grid-Scale & Hyper-Scale

Scale & grow with your customers

Enterprise Services

Gain peace of mind for you and your tenants

Quality of Service

Eliminate the impact of noisy neighbors

Environment synergies protect your data

Agile integration

Synchronous copy

- Enhance Recovery Time Objective & Recovery Point Objective
- Cross-system consistency groups
- Non disruptive workload mobility

Asynchronous copy

- Replication to FlashSystem A9000 & A9000R or XIV Gen 3*
- Protect more app's from down time
- Transferable software license
- Volumes must be uncompressed

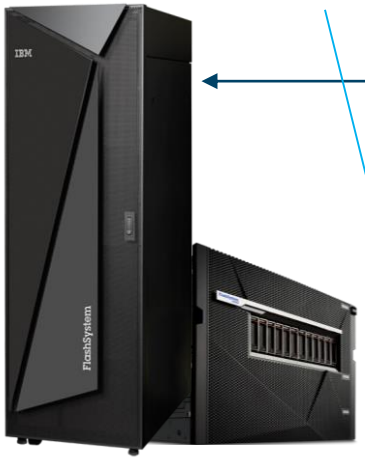

Primary site

- Space efficient snapshots
- VMware/Microsoft/OpenStack unified orchestration for Disaster Recovery, Backup & Restore
- Single pane of storage management

IBM Spectrum Protect

Backup/DR

- Data Protection
- Node Replication

Archive**

- Reporting
- Monitoring
- Analytics

Test/Dev DevOps

- Instant mount*
- Self service portal

Analytics

- Data retention
- Archive

*Replication to XIV Gen 3: Future statement of direction

**For Spectrum Protect server (TSM), not from snapshots.

Tight integration abilities under the VMware umbrella

Agile integration

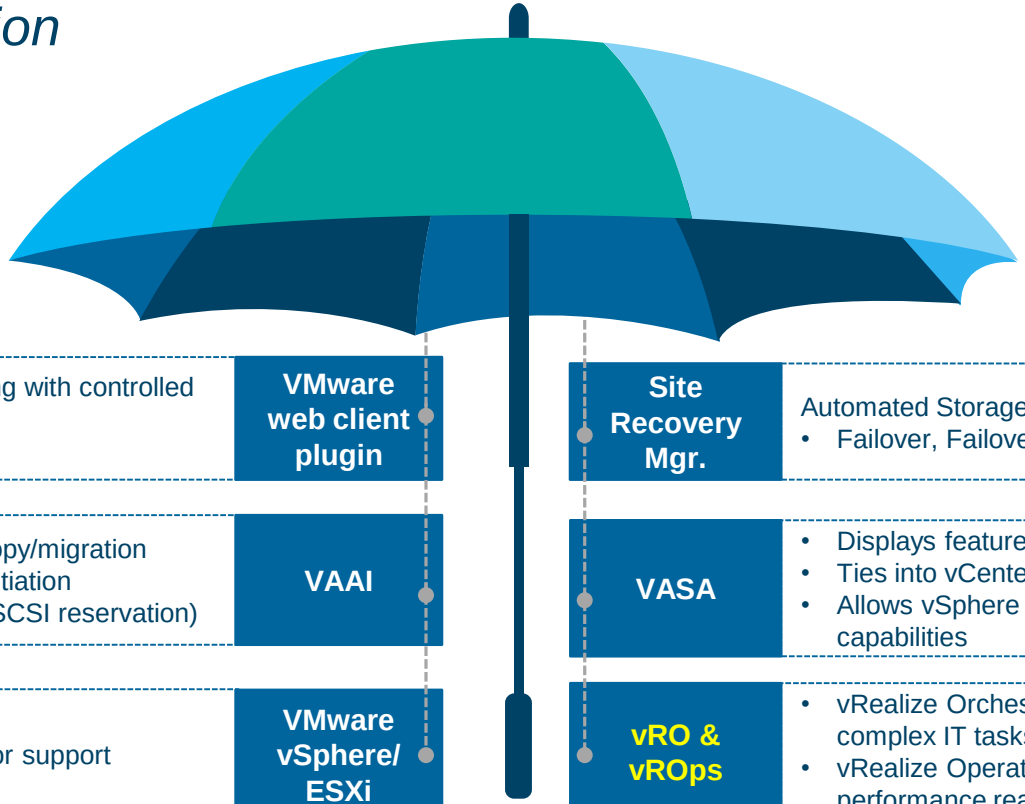

- Self-service provisioning with controlled delegation
- Storage visibility

VMware web client plugin

- HW-accelerated VM copy/migration
- HW-accelerated VM initiation
- Accelerate VMFS (no SCSI reservation)

VAAI

- Certified for Hyper-Visor support

VMware vSphere/ ESXi

Site Recovery Mgr.

Automated Storage & Host

- Failover, Failover testing, Failback

VASA

- Displays features of the physical storage devices
- Ties into vCenter for easy mgmt.
- Allows vSphere admin insight into storage capabilities

vRO & vROps

- vRealize Orchestrator simplifies the automation of complex IT tasks
- vRealize Operations Manager helps to analyze performance real-time in your virtual environment

A brand new interface simplifies management

Enduring economics

Web based application

Holistic View

Powerful Navigation

Instant Analysis

IBM FlashSystem Within The IBM Systems Enterprise Storage Portfolio

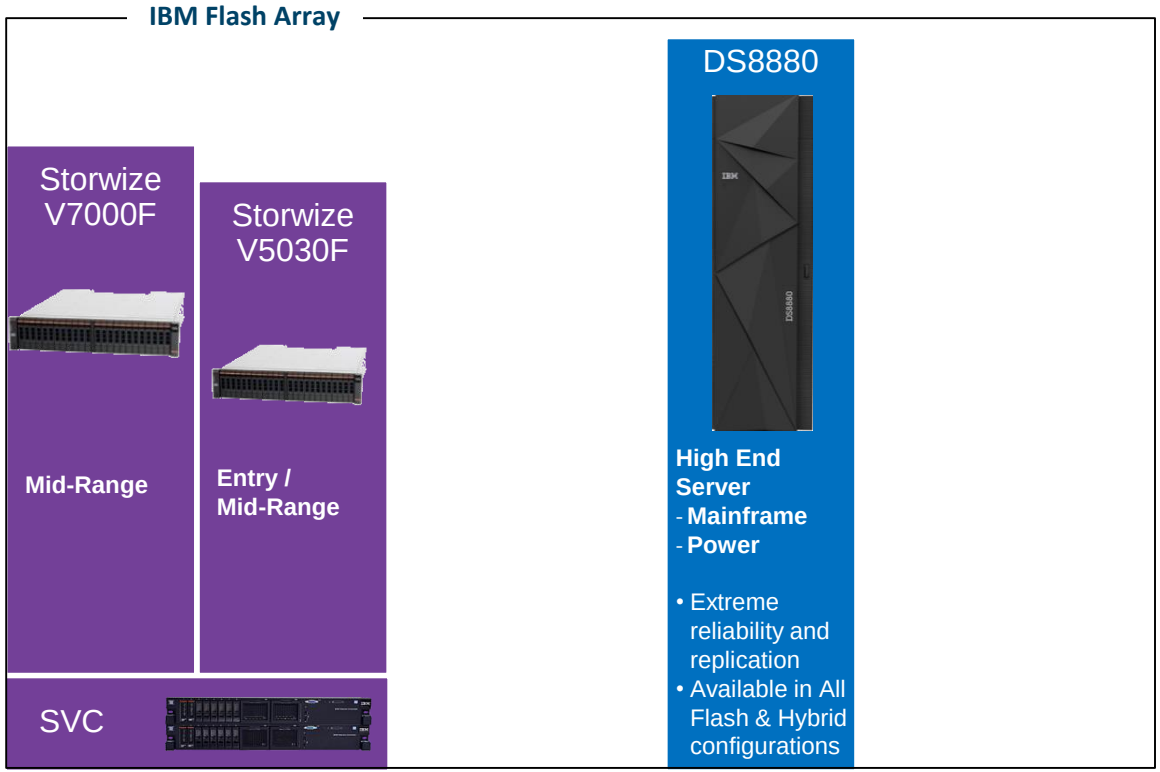

IBM Spectrum Virtualize

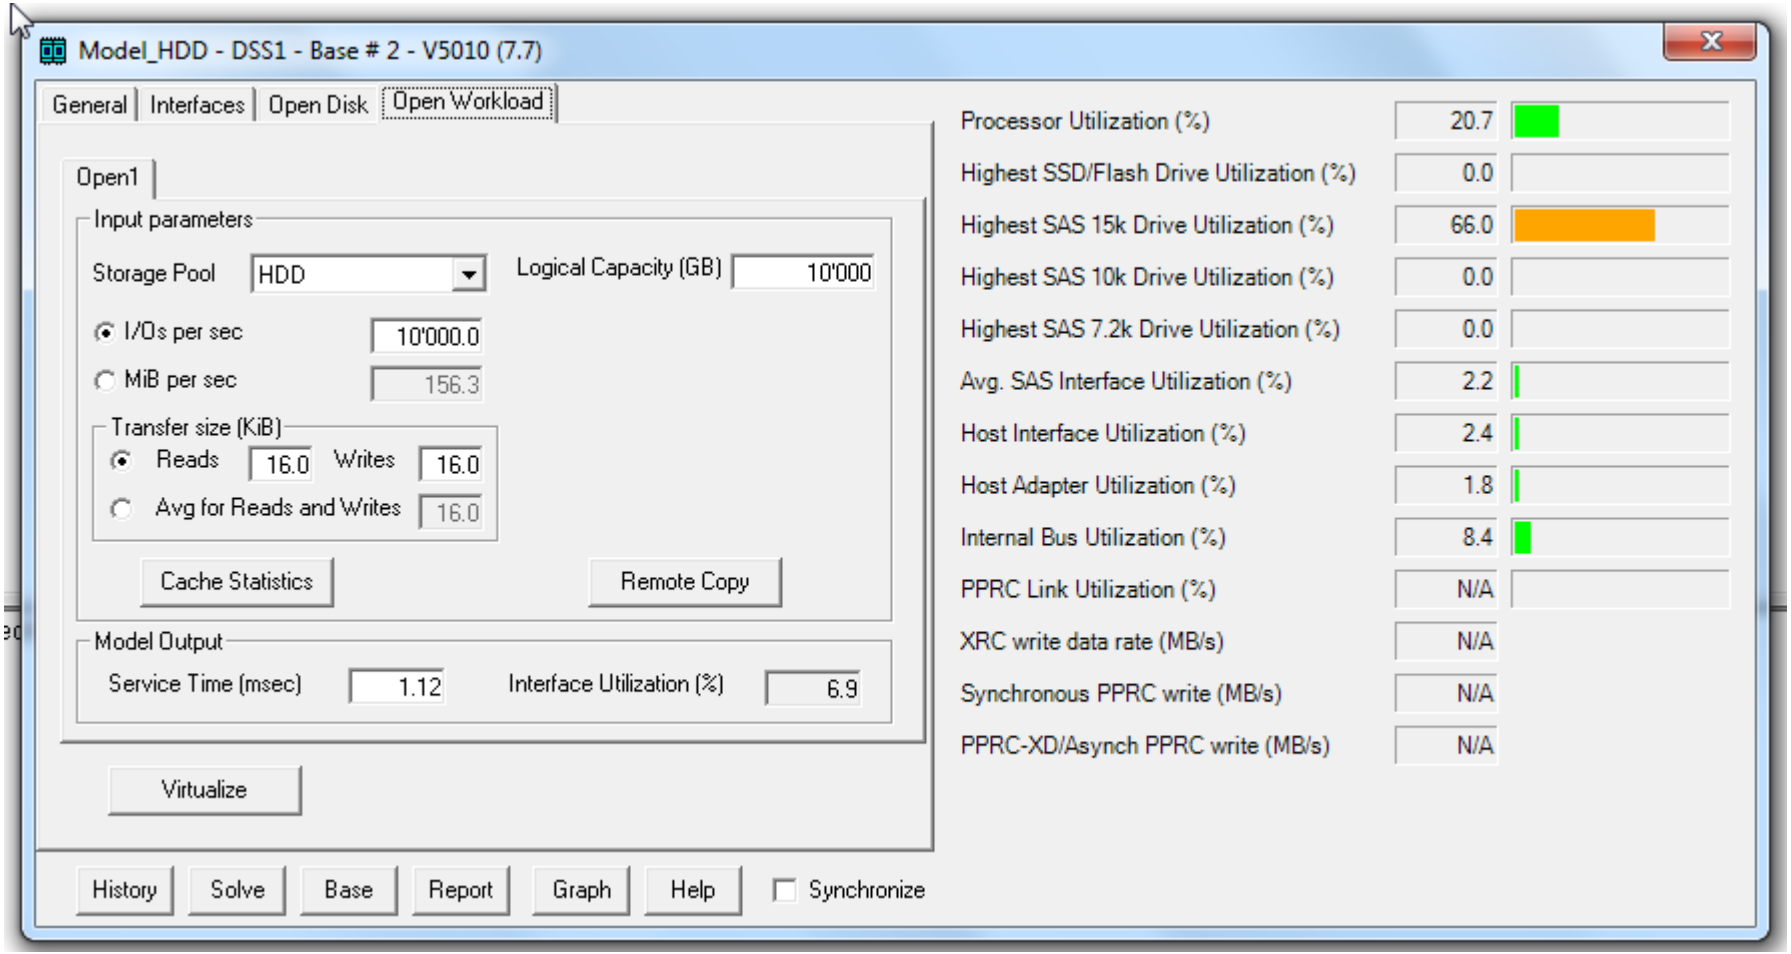
Model_HDD - DSS1 - Base # 2 - V5010 (7.7)
General
Interfaces
Open Disk
Open Workload
Open1
Input parameters
Storage Pool
HDD
Logical Capacity (GB)
10'000
I/Os per sec
10'000.0
MiB per sec
156.3
Transfer size (KiB)
Reads
16.0
Writes
16.0
Avg for Reads and Writes
16.0
Cache Statistics
Remote Copy
Model Output
Service Time (msec)
1.12
Interface Utilization (%)
6.9
Virtualize
History
Solve
Base
Report
Graph
Help
Synchronize
Processor Utilization (%) 20.7
Highest SSD/Flash Drive Utilization (%) 0.0
Highest SAS 15k Drive Utilization (%) 66.0
Highest SAS 10k Drive Utilization (%) 0.0
Highest SAS 7.2k Drive Utilization (%) 0.0
Avg. SAS Interface Utilization (%) 2.2
Host Interface Utilization (%) 2.4
Host Adapter Utilization (%) 1.8
Internal Bus Utilization (%) 8.4
PPRC Link Utilization (%) N/A
XRC write data rate (MB/s) N/A
Synchronous PPRC write (MB/s) N/A
PPRC-XD/Asynch PPRC write (MB/s) N/A

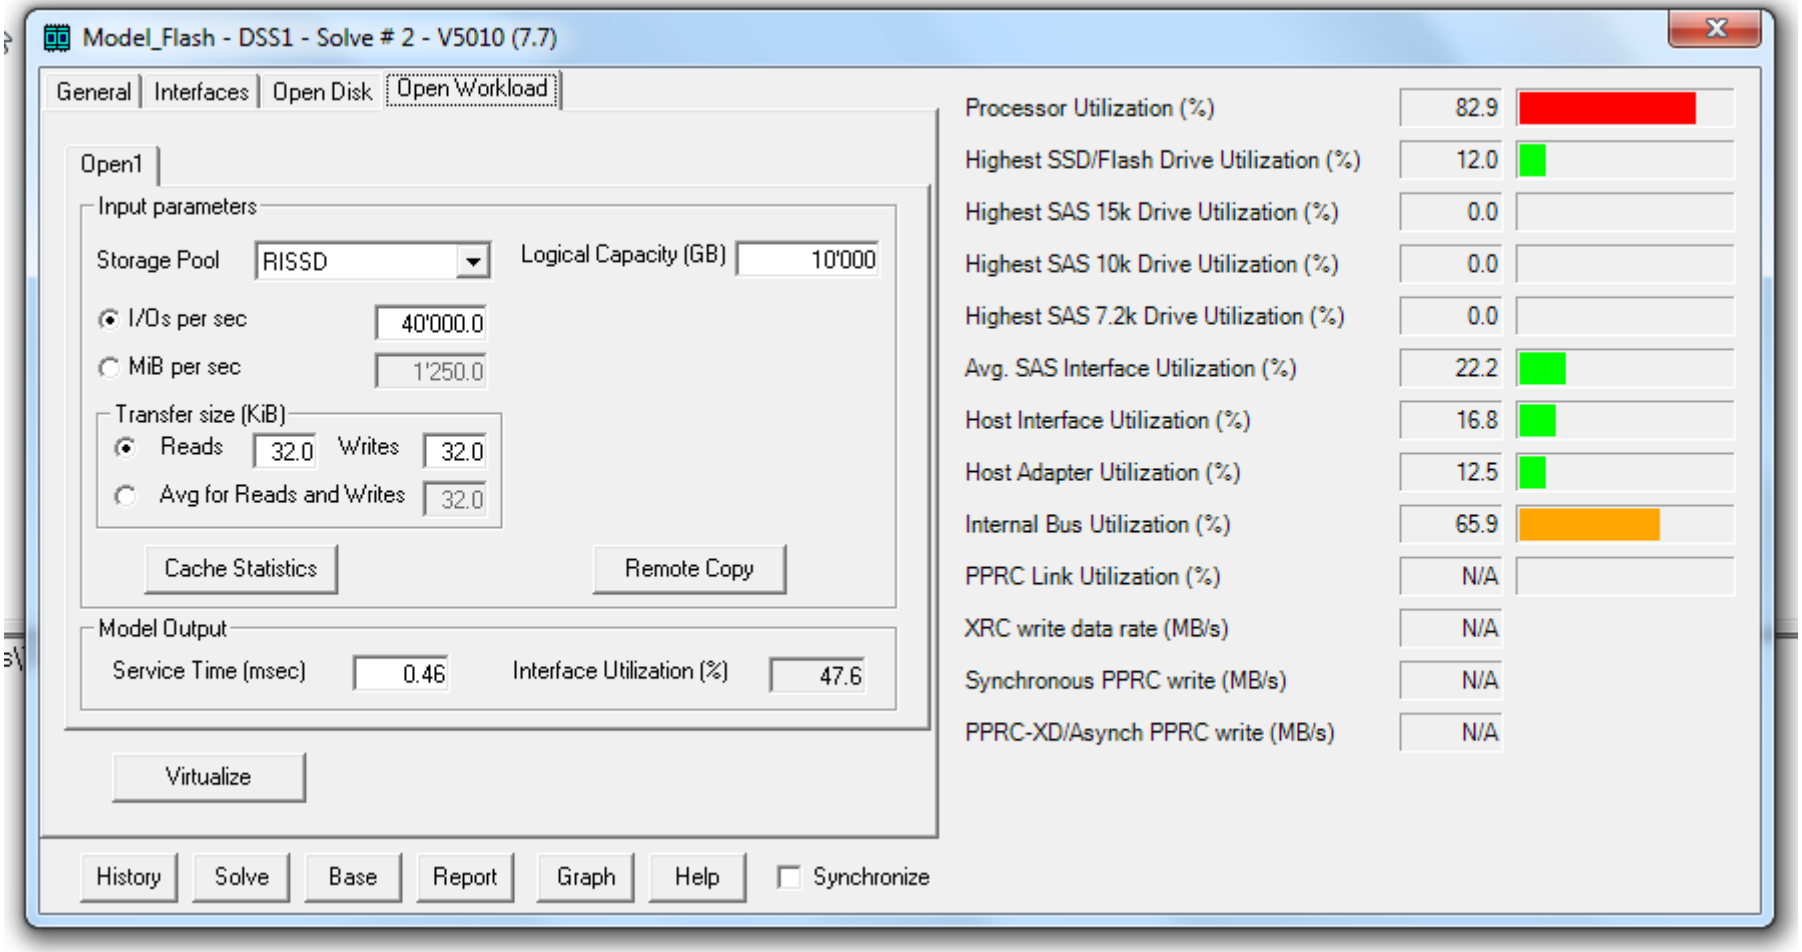
Model_Flash - DSS1 - Solve # 2 - V5010 (7.7)
General
Interfaces
Open Disk
Open Workload
Open1
Input parameters
Storage Pool
RISSD
Logical Capacity (GB)
10'000
I/Os per sec
40'000.0
MiB per sec
1'250.0
Transfer size (KiB)
Reads
32.0
Writes
32.0
Avg for Reads and Writes
32.0
Cache Statistics
Remote Copy
Model Output
Service Time (msec)
0.46
Interface Utilization (%)
47.6
Virtualize
History
Solve
Base
Report
Graph
Help
Synchronize
Processor Utilization (%) 82.9
Highest SSD/Flash Drive Utilization (%) 12.0
Highest SAS 15k Drive Utilization (%) 0.0
Highest SAS 10k Drive Utilization (%) 0.0
Highest SAS 7.2k Drive Utilization (%) 0.0
Avg. SAS Interface Utilization (%) 22.2
Host Interface Utilization (%) 16.8
Host Adapter Utilization (%) 12.5
Internal Bus Utilization (%) 65.9
PPRC Link Utilization (%) N/A
XRC write data rate (MB/s) N/A
Synchronous PPRC write (MB/s) N/A
PPRC-XD/Asynch PPRC write (MB/s) N/A

Disques SSD proposés

IBM propose deux types de SSD:

	Enterprise SSD	SSD «Low Cost»
DWPD	10	1
Positionnement	Workload orientés majoritairement en écriture (exemple typique: IBMi)	Workload orienté majoritairement en lecture (exemple: profil I/O R/W de 70/30)
Technologie SSD	MLC	• MLC pour Toshiba • TLC V-NAND pour Samsung
Offre SSD	• 400/800GB (HGST SSD800MM) • 1,6TB (HGST SSD1600MM) • 3,2TB (Toshiba PX04SMB320)	• 1,92 et 3,84TB (Toshiba SSD-PX04SR) • 7.68 et 15.36TB (Samsung PM1633a)

NB: Les SSDs «low cost» sont également appelés «Read Intensive Flash Drives» du fait de leur positionnement

- Liens utiles
  - https://www.hgst.com/sites/default/files/resources/US_SSD1600MM_ds.pdf
  - http://toshiba.semicon-storage.com/content/dam/toshiba-ss/asia-pacific/docs/product/storage/product-manual/eSSD-PX04SM-product-overview.pdf
  - http://toshiba.semicon-storage.com/content/dam/toshiba-ss/asia-pacific/docs/product/storage/white-paper/WhitePaper_eSSD_e_201404.pdf

  - http://toshiba.semicon-storage.com/content/dam/toshiba-ss/asia-pacific/docs/product/storage/product-manual/eSSD-PX04SR-product-overview.pdf
  - http://www.samsung.com/semiconductor/global/file/insight/2016/06/PM1633a-flyer-0.pdf

- Les disques SSD Enterprise/Low Cost peuvent être intégrés dans des groupes TRAID ou DRAID.
- Depuis Spectrum Virtualize 7.8, une nouvelle dénomination (RI = Read Intensive) a été créée pour différencier les deux types de SSD: SSD (Enterprise SSD) et RI (SSD Low cost). Un nouveau « Tier » a également été créé pour les SSD RI .

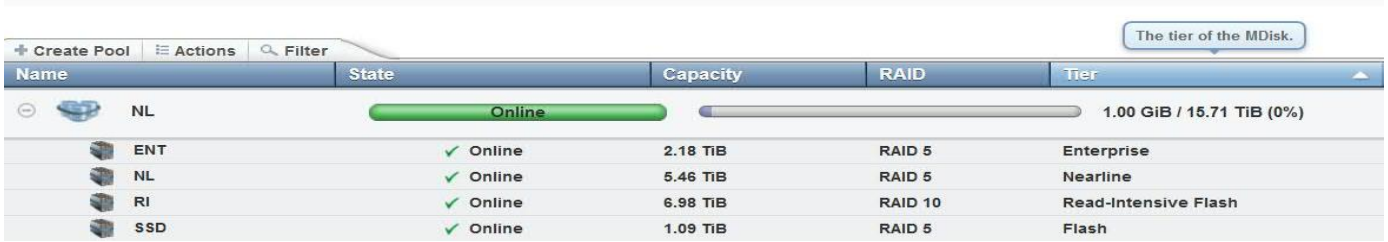

Depuis Spectrum Virtualize 7.7.1, un indicateur du degré d'usure des SSD RI est disponible via GUI: « Write Endurance Limit ». La commande lsdrive du CLI indique « l'Espérance de Vie (EV) » des disques SSD RI via un code couleur: rouge EV < 5 ans, orange ~4,5 < EV < ~5,5 ans et vert EV ~5,5 ans et plus.

- Lorsque le SSD RI atteint 95% d’usure, il est considéré comme défaillant et une procédure de reconstruction du groupe RAID sur un spare est lancé (si l’administrateur en a prévu, bien sur).

Dans la plupart des workloads, l’utilisation de SSD RI ne posera aucun problème lié à la réduction de la durée de vie du support. Par exemple, un ensemble de 20 disques SSD RI 1,92TB (2 groupe DRAID5 de 10 disques) peut supporter 90 000 IOPS de 16KB en continu avec un profil R/W de 70/30, sans usure prématurée.

En fonction du profil I/O des applications, le choix pourra se faire entre les SSD Enterprise et les SSD RI, ces derniers ayant un très net avantage en terme de prix.

Capacity Report

Description	Arrays	Total drives	Data drives	Parity / Mirror drives	Spares	Physical capacity (TB)	Effective capacity (TB)	Effective capacity (TiB)	Eff. util.
TOTAL	3	60	50.7	6.3	3	536.6	452.0	411.1	84 %
By pool									
Pool_1	1	12	9.8	1.2	1	120.0	97.1	88.3	81 %
Pool_2	1	24	20.4	2.6	1	368.6	314.0	285.6	85 %
Pool_3	1	24	20.4	2.6	1	48.0	40.9	37.2	85 %
By preset									
Distributed RAID 5	3	60	50.7	6.3	3	536.6	452.0 *	411.1 *	84 %
By platform									
Open	3	60	50.7	6.3		536.6	452.0	411.1	84 %
By drive type									
AHD6: 10 TB 7.2k 3.5" NL HDD	1	12	9.8	1.2	1	120.0	97.1 *	88.3 *	81 %
AHHD: 15.36 TB 2.5" RI Flash Drive	1	24	20.4	2.6	1	368.6	314.0 *	285.6 *	85 %
AHG2: 2 TB 7.2k 2.5" NL HDD	1	24	20.4	2.6	1	48.0	40.9 *	37.2 *	85 %

* Without subtracting 5 extents per pool


THANK
YOU
GRACIAS
ARIGATO
SHUKURIA
GOZAIMASHITA
EFCHARISTO
DANKSCHEEN
JUSPAXAR
KOMAPSUMNIDA
MAAKE
GRAZIE
MEHRBANI
PALDIES
BOLZÏN
MERCI
BÏYAN
SHUKRIA
TINGKI
YAQHANYELAY
TASHAKKUR ATU
SUKSAMA
EKHMET
SPASSIBO
SNACHALHUYA
NUHUN
CHALTU
WABEEJA
MAITEKA
YUSPAGARÀTAM
HUI
GUI
HÄTUR
SPASIBO
DENKUUJA
NENACHALHYA
UHALCHEESH
EKO-JU
SIKOMO
MAKETA!
MIMMONCHAR
DHANYABAAD
ANNA
ATTO
MERCI
LAH
SANCO
MERASTAWHY
GAEJTHO
AGUYJE
FAKAAUE
TAVTAPUCH
MEDAWAGSE
BAINA