



IBM Technical conference

Buenos Aires 30 Mayo 2013

- IBM High End disk DS8870 disk and IBM XIV disk

Ing. Federico Lizarralde – Storage Hardware



AGENDA :


- Introducción a la tecnología IBM High End DS8870
- Descripción de los modelos, arquitectura, características.
- Anuncions y funcionalidades IBM DS8870 High End disk .
- Tecnología IBM XIV High End disk
- Integración, arquitectura, funcionalidades.
- Administración, GUI, monitoreo XIV disk

Familia de subsistemas de discos IBM System Storage



IBM System FLASH Technology

Entry point




IBM V3700

Unified SAN/NAS series disk for open servers



IBM STORWIZE V7000 / SVC

Foundation
Open Enterprise
class



**XIV Enterprise
Open**

Enterprise-class storage . Leading the industry in functionality, performance, TCO



DS8870

Plataforma de administración
unificada

Servicios de copia unificados

Virtualization

Compelling price points

Lider en la industria de servicios y
soporte


Enterprise-class Storage Continuum

IBM System Storage Family ayuda a la innovación:

- Simplifica la infraestructura de storage IT permitiendo la administración a bajos costos y complejidad, mientras se incrementa la habilidad de responder ante las necesidades de los cambios.
- Asegurar continuidad del negocio, seguridad y durabilidad de los datos.
- Administración eficiente de la information a travez del ciclo de vida de los datos (ILM), relativo al valor de los negocios.

Introducing the DS8870

Smarter Storage for enterprise-critical information environments

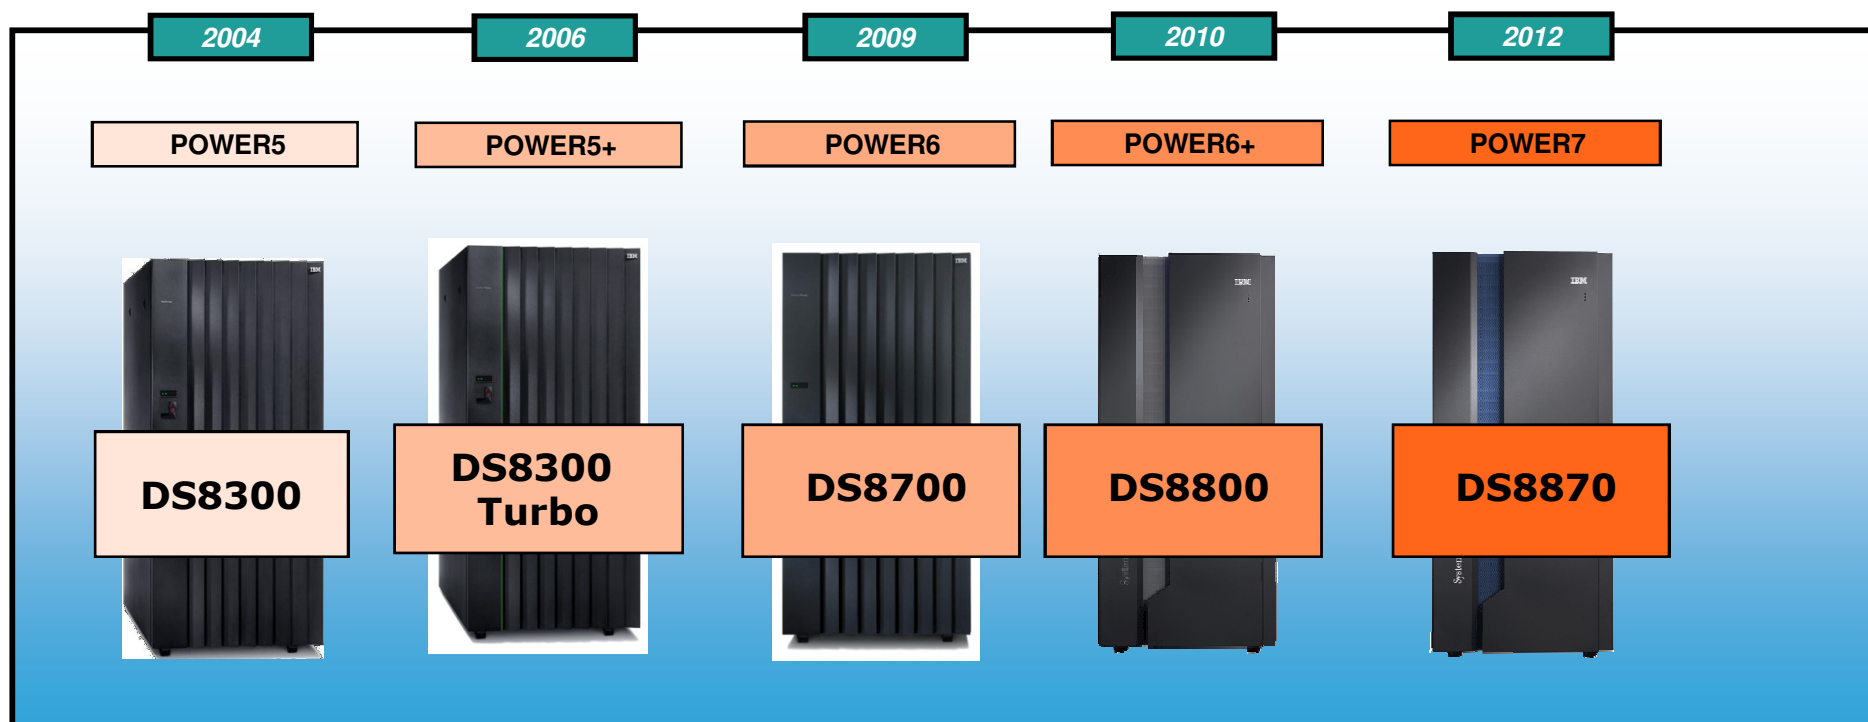


- **Built on the DS8800 base**
 - Exceptionally fast with up to 3× performance increase
 - Proven architecture and code base for optimal reliability with non-disruptive microcode updates
 - Inherited all functionality of DS8800
 - RoHS compliance reduces hazardous material
- **New dual IBM POWER7 controllers – unprecedented scaling**
 - Scalable processor configurations with 2, 4, 8 and 16 cores per controller
 - Scalable cache from 16 GB – 1 TB
 - Everything scales non-disruptively
 - Entry-level Business Class configuration also available
- **New energy-efficient power supply**
 - Improved efficiency, power dissipation, reliability
 - Up to 20% reduction in energy usage
 - Designed to meet upcoming ENERGY STAR standard
- **Full Disk Encryption drives now standard**
 - Client decides when to encrypt or not^{*)}



5th-generation DS8000 enterprise disk system

- Building on a market-proven, reliable code base!
- 94% of the same proven microcode



- Designed for over 5-9's availability natively
- Designed for over 6-9's availability when DS8000 with Metro Mirror is combined with GDPS/PPRC HyperSwap

Business Class and Enterprise Class configuration options

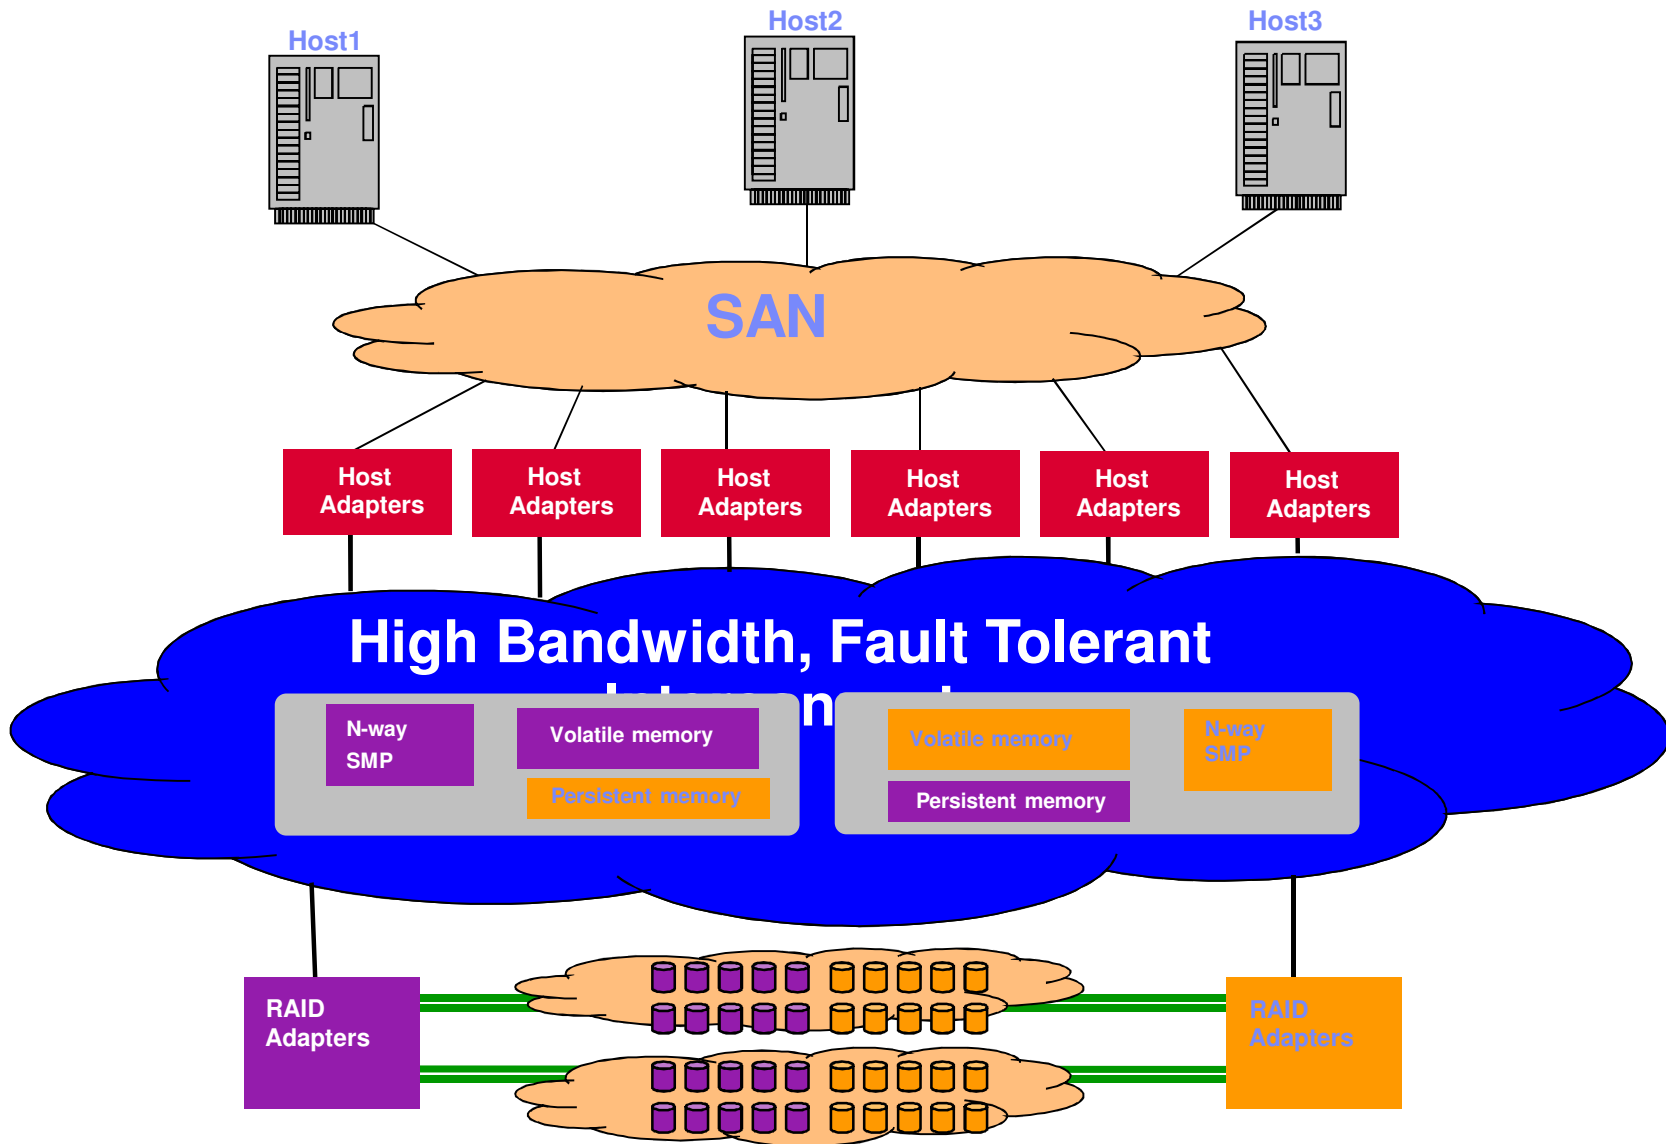


- Non-disruptive upgrade from smallest to largest configuration

Model	Processor	Physical Capacity (max.)	Disk Drives (max.)	Memory (GB)	Host Adapters (max.)	9xE Attach
Business Class						
961	2-core	216 TB	144	16	4	0
961	2-core	216 TB	144	32	4	0
Enterprise Class						
961	4-core	360 TB	240	64	8	0
961	8-core	1 584 TB	1 056	128	16	0–2
961	8-core	2 304 TB	1 536	256	16	0–3
961	16-core	2 304 TB	1 536	512	16	0–3
961	16-core	2 304 TB	1 536	1 024	16	0–3
First Expansion Frame						
96E	N/A	504 TB	336	N/A	8	N/A
Second/Third Expansion Frame						
96E	N/A	720 TB	480	N/A	N/A	N/A

DS8000 Topologia – All Models

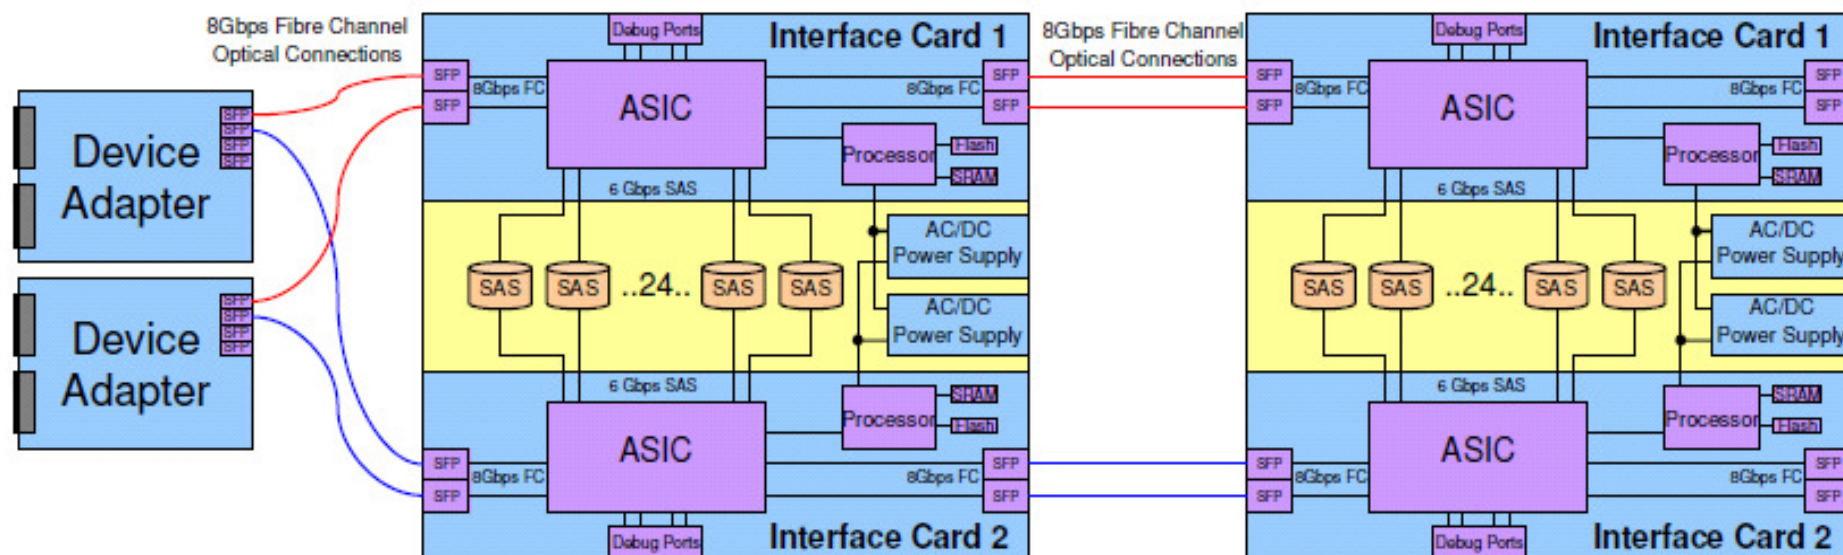
(different power & *Bandwidth*)






DS8870 Storage Enclosure Interconnect

Gigapack Enclosures




- **Device adapter interconnect remains Fibre Channel (Fibre Channel backbone with SAS drives)**
 - Gigapack connects to the DA via Fibre Channel protocol and converts to SAS protocol on the disk side
 - Hardware supports up to 8Gbps FC interconnect and 6 Gbps SAS drives
- **Optical SFP interconnect for adapter-to-drawer and drawer-to-drawer cabling**
 - Allows for increased cabling distance, especially at 8Gbps speeds
- **Support for 24x SFF (2.5") SAS disks in 2U drawer: HDD (SAS 15k, 10k), SSD, NL-SAS**

DS8870 Base Model View: Front and Rear



1. Gigapack enclosures with 240 disk drives
2. HMC
3. POWER7 servers (p740)
4. I/O enclosures
5. DC-UPS power supplies
6. RPC cards (Rack power control)

Schematic View



Base Frame

1st Expansion Frame

2nd Expansion Frame

3rd Expansion Frame

DS8870 Gigapack Storage Enclosure – Overview

- Disk Technology
 - 2.5" Small Form Factor (SFF) SAS 2.0 Disks
 - 2U of vertical rack space
 - Supports 24 disks per enclosure (Gigapack)
 - DDMs only in front
- Disk Capacity (all FDE capable)
 - 146/300 GB 15K rpm
 - 600/900 GB 10K rpm
 - 400 GB Solid-State Disk
 - 3 TB 7200 rpm Nearline-SAS
- No DDM intermix in Storage Enclosure pair
- All disk drives installed within a disk enclosure pair must be of the same type (capacity and speed).
- Rear Components
 - Two Integrated power and cooling units (PSU)
 - Two Integrated FCIC cards
- Cooling is now front to back (entire rack is now front to back)
- Filler plates required to maintain airflow
- Enclosure follows SBB 2.0 standard easing possible future upgrade paths
- Form factor allows for integrated power and cooling
- Enclosures always installed in pairs: top and bottom



Extraordinary performance for enterprise applications



	DS8300 (R4.3)	DS8700 (R5)	DS8800 (R6)	DS8870 (R7)	<i>Increase</i>
	P5+ 4-way	P6 4-way	P6+ 4-way	P7 16-way	<i>vs. DS8800</i>
Seq. Read (GB/s)	3.9	9.7	11.8	21.0	1.8x
Seq. Write (GB/s)	2.2	4.7	6.7	11.0	1.6x
DB z/OS (K IOPS)	165	201	204	640	3.1x
DB Open (K IOPS)	165	191	198	550	2.8x

	Increase vs. DS8300	Increase vs. DS8700
Seq. Read (GB/s)	5.4x	2.2x
Seq. Write (GB/s)	5.0x	2.3x
DB z/OS (K IOPS)	3.9x	3.2x
DB Open (K IOPS)	3.3x	2.9x

Official Storage Performance Council (SPC) results



- **SPC-1 throughput of 451,082 IOPS**
 - #1 result for single, enterprise-class all-HDD system
 - 67% faster than HDS VSP

- **SPC-2 throughput of 15,424 MB/s**
 - #1 result overall
 - 6% faster than prior #1 result
 - 17% faster than HDS VSP
 - 59% faster than DS8800

Note: DS8870 results have been published on 3 October 2012


SPC Benchmark 1 (SPC-1) results page – http://www.storageperformance.org/results/benchmark_results_spc1

SPC Benchmark 2 (SPC-2) results page – http://www.storageperformance.org/results/benchmark_results_spc2

SPC-1 comparison vs. competition^{*)}

Higher throughput with 451K IOPS and lower response time!

SPC-1 Published Results



^{*)} Source: Published SPC-1 results: http://www.storageperformance.org/results/benchmark_results_spc1

- DS8300 Turbo: 480 × 73 GB 15K HDDs, RAID-10, 256 GB Cache
- HDS USP V: 1024 × 146 GB 15K HDDs, RAID-1, 256 GB Cache
- HDS VSP: 1152 × 146 GB 15K HDDs, RAID-1, 512 GB Cache + 512 GB Cache Flash
- 3PAR V800: 1920 × 300 GB 15K HDDs, RAID-1, 8 nodes, 768 GB Cache
- SVC + 2 × DS8700: each DS8700: 1024 × 146 GB 15K HDDs, RAID-10, 384 GB Cache
- DS8870: 1536 × 146 GB 15K HDDs, RAID-10, 32 × 8Gb FCP, 1024 GB Cache

EMC refuses to publish results, so we can assume less-than-impressive performance

Energy consumption comparison



DS8700 with 1024 drives

- Base frame: 6.8 kW
- Exp frame: 7.1 kW
- Exp frame: 6.1 kW
- Exp frame: 6.1 kW
- Exp frame: 3.1 kW

TOTAL: 29.2 kW



DS8800 with 1536 drives

- Base frame: 7.5 kW
- Exp frame: 6.2 kW
- Exp frame: 6.3 kW
- Exp frame: 6.3 kW

TOTAL: 26.3 kW



DS8870 with 1536 drives

- Base frame: 6.0 kW
- Exp frame: 5.6 kW
- Exp frame: 5.8 kW
- Exp frame: 5.8 kW

TOTAL:

23.2 kW

Base frame is 20% more energy efficient



Higher efficiency rating positioned to meet emerging ENERGY STAR requirements

Note: Measurements taken on 100% read miss workload

New VMware VAAI support

- VMware APIs for Array Integration (VAAI) includes 3 “primitives” for offloading virtual machine and storage management operations to storage systems that support them
 - Atomic Test and Set (ATS) enables hardware-assisted locking of files
 - Full Copy – enables the storage array to make full copies of data within the array
 - Block Zeroing enables the array to zero out large numbers of blocks
- ATS and Full Copy support are available with this DS8870 release; Block Zeroing support is planned for Q2/2013
- Code needed for this: R7.0.5/R7.1 JUNE 2013
- For VAAI clients, use administrative indicator feature FC0965
- DS8870 enhancements for VMware planned for 2013 - JUNE
 - VAAI Block Zeroing
 - vCenter Plug-in
 - SRM 5.0/5.1

Full Disk Encryption is now standard

- FDE options across all drive tiers
- Same performance as standard drives
- Standards-based key management software supports key management interoperability protocol (KMIP)
- Key manager supports both disk and tape
- New European Union Privacy Directive makes breaches very costly
- Encryption is the least expensive data disposal technique
- Supports Easy Tier environments
- Over a thousand IBM encryption disk and tape solutions deployed worldwide

"Do not wait for an event-driven reaction to secure your data. Proactively securing your data will help ensure against a worst-case scenario and a financial impact that is likely to far surpass that of the data security purchase itself."

John Monroe
Gartner Hype Cycle for Storage, July 2012



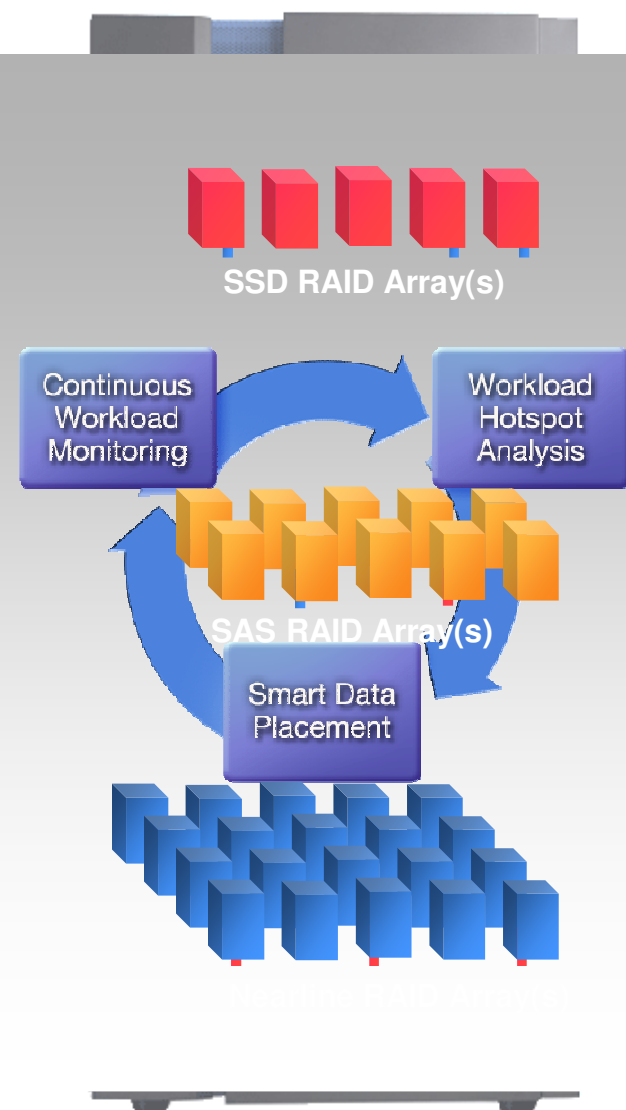
"Within five years, all HDDs and SSDs will be shipped preloaded with some kind of industry-standard FDE technology" – Gartner Hype Cycle for Storage, July 2012

* Requires deployment of Tivoli Key Lifecycle Manager or IBM Security Key Lifecycle Manager


Introducing IBM Easy Tier

SSD Optimisation

- **Monitors performance of each ‘extent’ (1 GiB, sub-volume level) to determine the data ‘temperature’**
- **Creates extent migration plan for optimal data placement every 24 hrs based on performance statistics**
- **Migrates extents within an extent pool according to plan over 24-hour period**
- **A limited number of extents are chosen for migration every 5 minutes**



ET IV: Easy Tier – SSD + SAS Extent Pool






IBM Auto Rebalance Performance


1. Support both **homogeneous pool** and hybrid pool.
2. Automatically maximize storage performance primarily for IOPS. This includes all I/O types so ranks are also secondarily optimized for bandwidth.
3. Auto Performance Rebalance on any storage tier
 - a) After new resource is added to or removed from the storage pool.
 - b) Natural Performance Skew : Caused content of some extent has more access of the other extent.

a




Add Rank

b



Natural Performance Skew

ET IV: Intra-tier Auto-Rebalance – Homogeneous Extent Pool




All performance data contained here were obtained in the specific operating environment and under the conditions and/or configuration described above and is presented as an illustration. Details are available in the whitepaper WP102024, <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102024> "IBM System Storage DS8800 and DS8700 Performance with Easy Tier 3rd Generation". Performance obtained in other operating environments may vary and customers should conduct their own testing.

Tiering data to the server can boost performance – POWER AIX

Now ! JUNE 2013

- Easy Tier moves the hottest data closer to server and overcomes latency of the SAN
- Both server and storage resources remain optimised for performance and cost objectives
- Advanced self-tuning reduces administration costs further and allows IT staff to focus on strategic initiatives




[IBM June Storage SOD RFA link](#)

Applications will guide Easy Tier on increasing performance even more – POWER AIX.


Now! June 2013

- Applications will provide performance tips to Easy Tier
 - Applications will guide what, when, where and how data is placed across the server and storage systems dynamically
- Advanced self-tuning reduces administration costs further and allows IT staff to focus on strategic initiatives
- Integration with existing QoS features enables better multi-tenancy and cloud environments



[IBM June Storage SOD RFA link](#)


Metro Mirror



Metro Mirror

- **Distance**
 - 303 KM
 - 500 KM SAP Adaptive Computing Compliant
- **Performance Improvements**
- Pre-deposit Writes.
- Half The Infrastructure Required
 - SAN Ports
 - Inter Controller Connections


Global Mirror



Global Mirror

- Up To 50% Bandwidth Reduction Over Competitive Architectures
- **3 – 5 Sec RPO**
- Heterogeneous Consistency Groups (zOS, UNIX, Wintel)
- Superior Application Failover Capability Via GDPS Or HACMP
- No External SW Required

Metro Global Mirror



MGM

- Up To 50% Bandwidth Reduction Over Competitive Architectures
- 3 – 5 Sec RPO
- **Incremental Resynch**
- Heterogeneous Consistency Groups (zOS, UNIX, Wintel)
- Superior Application Failover Capability Via GDPS Or HACMP
- No External SW Required

Easy Tier Heat Map Transfer

ET Learning Transfer to PPRC Secondary

■ Primary site:

- Optimise the storage allocation according to the customer workload;
- Save the learning data;
- Transfer the learning data from the primary site to the remote site.

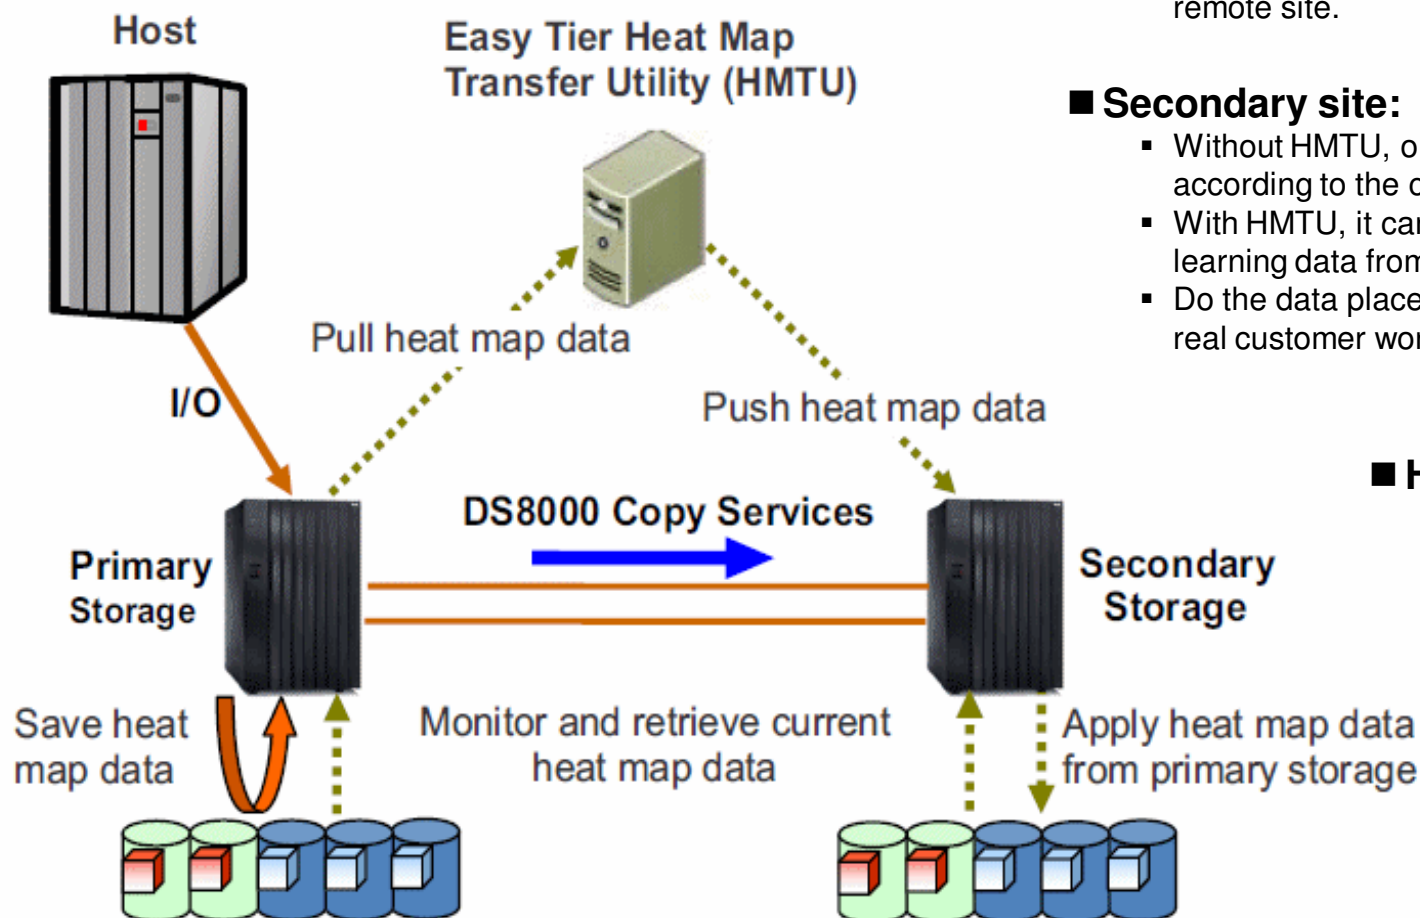
■ Secondary site:

- Without HMTU, only optimise the storage allocation according to the customer **write** workload;
- With HMTU, it can apply the HMTU checkpoint of learning data from the primary site;
- Do the data placement optimisation based on the real customer workload.

■ HMTU available for


- Windows
- Linux (SLES11+, RHEL 4+)

Alternatively: do with
TPC for Replication




Competing Application Workload on DS8000

- By itself, critical App1 succeeds
- Non-critical App2 begins
- Critical App1 slows due to DS8000 contention



Dynamic control of Competing Workloads by IOPM

- I/O Priority Manager detects critical App1 QoS impact
- App2 usage of critical DS8000 resources is controlled
- Controls on App2 restore App1 performance





- **1U drawer ... Up to 30 Flash cards**
 - 400GB Flash provide **Up to 9.5TB Usable (12 TB)**

- **Great performance**
 - Directly connected to DS8000 internal PCIe fabric for performance
 - Up to 400,000 IOPS (100% read) and 270,000 IOPS (100% write)
 - Up to 4.5 GB/s bandwidth

- **Two large-cache SAS adapters especially designed for Flash**
 - TMS common building blocks

- **Enterprise class RAS**
 - RAID 5 and 6 and mirrored/protected write cache

- **Easy Tier Enabled**

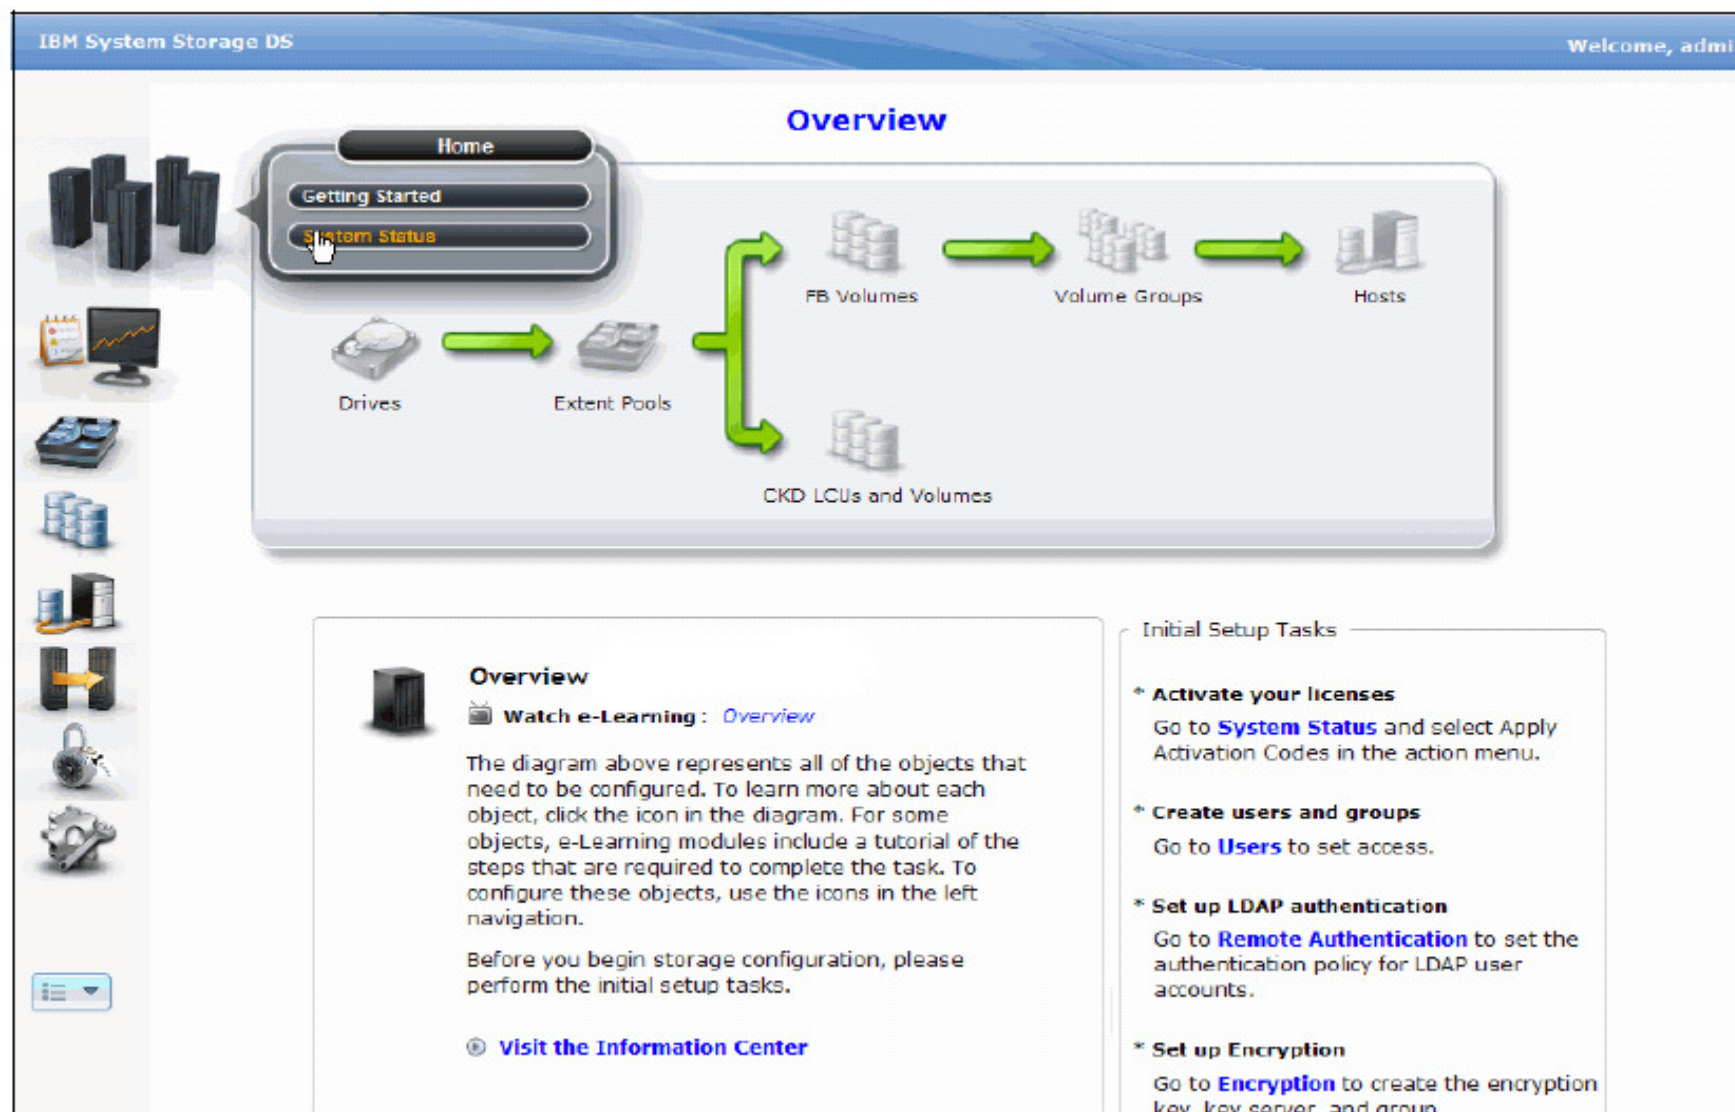



Figure 10-1 DS8000 Storage Manager GUI: Overview window

Familia de subsistemas de discos IBM System Storage



IBM System FLASH Technology

Entry point




IBM V3700

Unified SAN/NAS series disk for open servers



IBM STORWIZE V7000 / SVC

Foundation
Open Enterprise
class



**XIV Enterprise
Open**

Enterprise-class storage . Leading the industry in functionality, performance, TCO



DS8870

Plataforma de administración
unificada

Servicios de copia unificados

Virtualization

Compelling price points

Lider en la industria de servicios y
soporte

Enterprise-class Storage Continuum

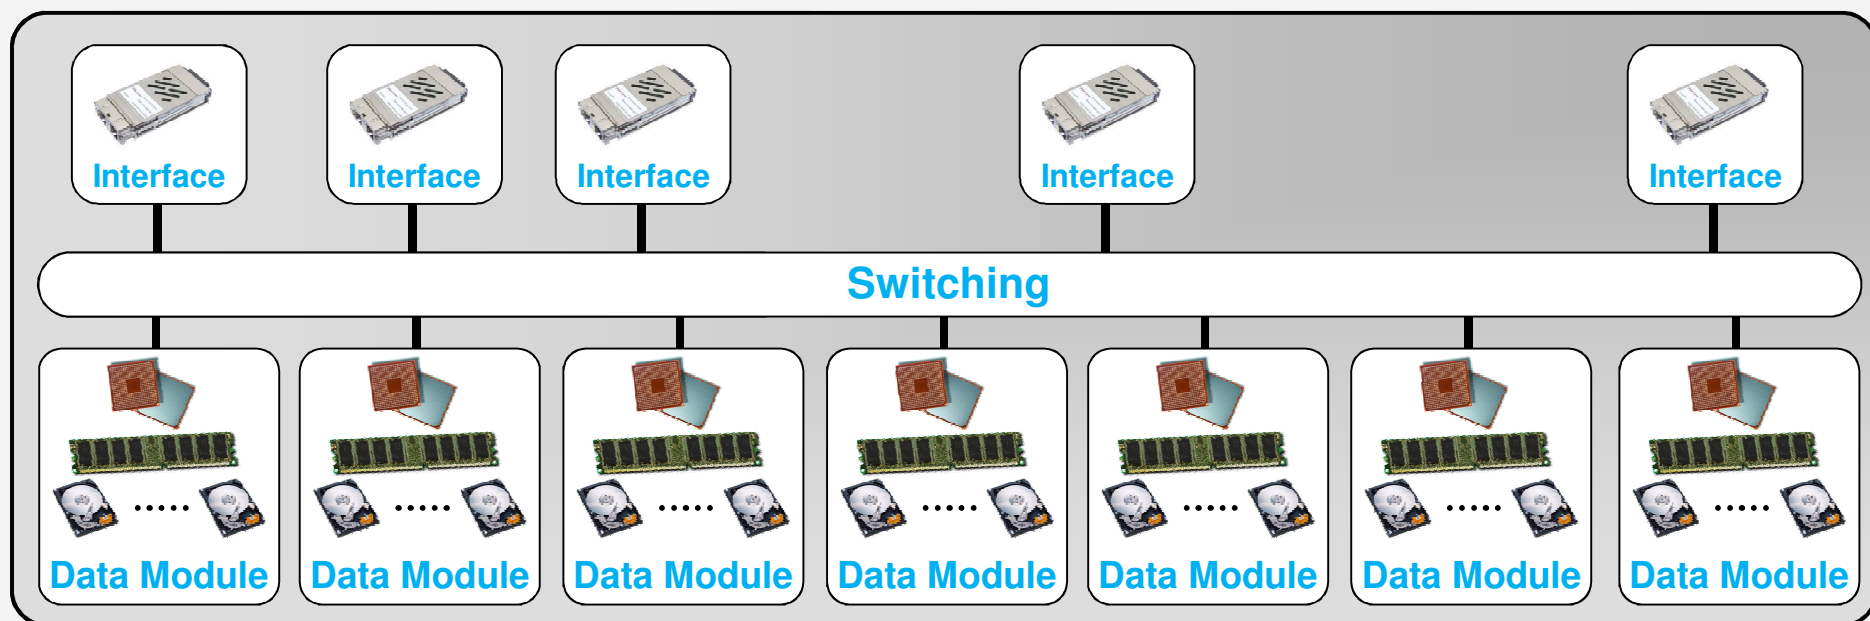
IBM System Storage Family ayuda a la innovación:

- Simplifica la infraestructura de storage IT permitiendo la administración a bajos costos y complejidad, mientras se incrementa la habilidad de responder ante las necesidades de los cambios.
- Asegurar continuidad del negocio, seguridad y durabilidad de los datos.
- Administración eficiente de la information a travez del ciclo de vida de los datos (ILM), relativo al valor de los negocios.

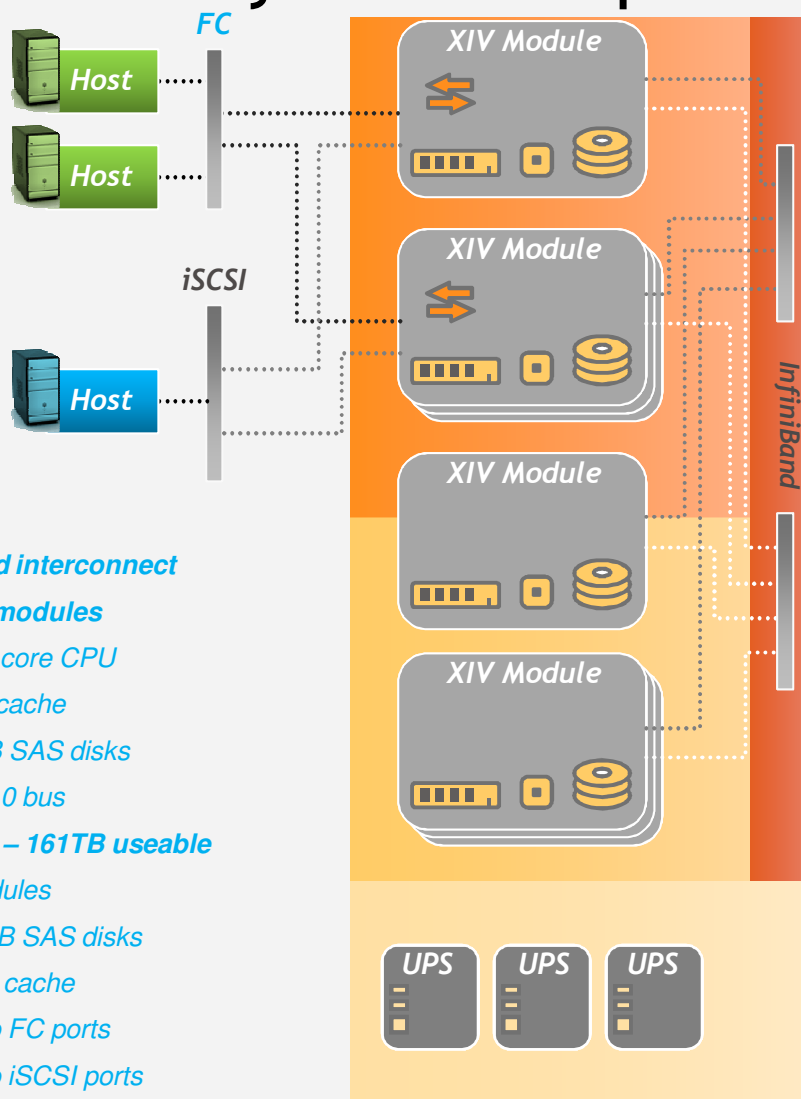
Arquitectura de XIV

Principios de diseño:


- Paralelismo masivo
- Distribución granular
- Componentes “Off-the-shelf”
- Discos, RAM y CPU acoplados
- Simplicidad de uso



XIV Gen3 System Components



- **InfiniBand interconnect**
- **Updated modules**
 - 1 quad core CPU
 - 24 GB cache
 - 12 2GB SAS disks
 - PCIe 2.0 bus
- **Full Rack – 161TB useable**
 - 15 modules
 - 180 2TB SAS disks
 - 360GB cache
 - 24 8Gb FC ports
 - 22 1Gb iSCSI ports
- **Partial Rack – 6,9-14 Modules**
- **IBM T42 Rack**





XIV Gen3 - InfiniBand Module Interconecion

■ InfiniBand

— Switched fabric

- 1 HCAs (host channel adapter) in each node
 - Dual 4x DDR ports
 - 8 lanes of PCIe 2.0
- Redundant switches
 - 4x DDR
 - 2-link switch interconnect



— Fast - see performance section

- High throughput
 - 4x DDR = 32 Gb/sec
- Low latency
 - Chips - 140 nanoseconds
 - **End-to-end between 1-3 microseconds**



— Low CPU/OS overhead

- Hardware based transport
- RDMA technology (module-to-module memory communication)

— Scalable

- 48,000 nodes in subnet
- 2^{128} Nodes in network

— No better choice for XIV Gen3



IBM XIV Gen3 Technology Highlights

Built for Performance

Componentes End-to-end ofrecen extrema performance

- **Nuevos Interface and Data Modules**
 - **More System throughput**
- **24GB cache por module (360GB per System)**
 - **50% more cache capacity**
- **8Gb FC and 1Gb iSCSI en todos los Interface Modules**
 - **Up to 24 8Gb FC ports**
 - **Up to 22 1Gb iSCSI ports**
 - **Over 2X more external bandwidth**
- **InfiniBand for Module Interconnect**
 - **Ofrece un 20X más internal bandwidth (infiniband)**



XIV Gen3 - Rack Configuration Specifications

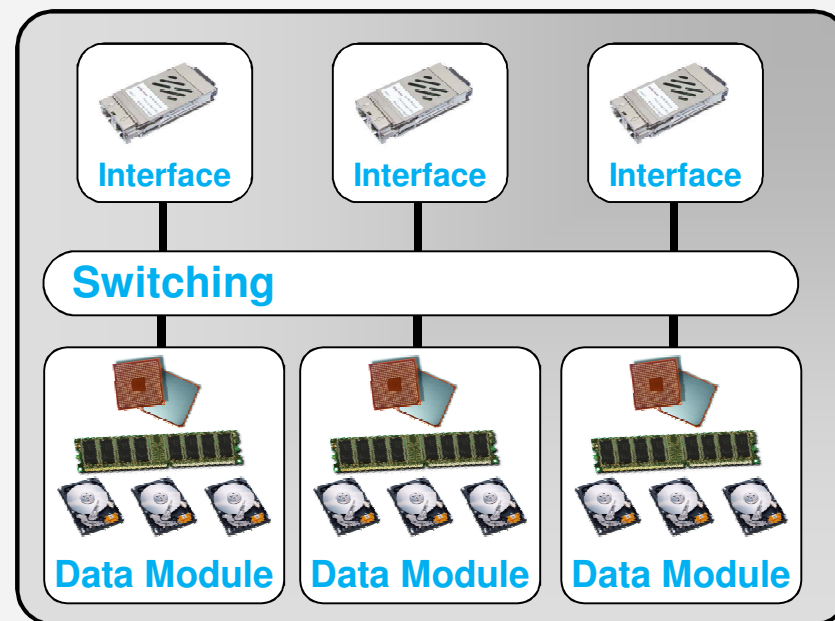

	6 Physical Modules	9 Physical Modules	10 Physical Modules	11 Physical Modules	12 Physical Modules	13 Physical Modules	14 Physical Modules	15 Physical Modules
Maximum Capacity -->	84.1	132.8	154.9	168.3	190.0	203.6	225.3	243.3
Capacity per CoD Activation -->	14.017	14.756	15.490	15.300	15.833	15.662	16.093	16.220
3 CoD Activations -->	42.050							
4 CoD Activations -->	56.067							
5 CoD Activations -->	70.083							
6 CoD Activations -->	84.100	88.533						
7 CoD Activations -->		103.289	108.430					
8 CoD Activations -->		118.044	123.920	122.400				
9 CoD Activations -->		132.800	139.410	137.700	142.500			
10 CoD Activations -->			154.900	153.000	158.333	156.615		
11 CoD Activations -->				168.300	174.167	172.277	177.021	
12 CoD Activations -->					190.000	187.938	193.114	194.640
13 CoD Activations -->						203.600	209.207	210.860
14 CoD Activations -->							225.300	227.080
15 CoD Activations -->								243.300

Figure 3-5 CoD capacities for a XIV Storage System Gen3 with 3-TB drives





Algoritmo de Distribución de XIV

- Todos los volúmenes se reparten en todos los discos
- Los datos se dividen en “particiones” de 1MB y se almacenan en los discos
- XIV distribuye automáticamente las particiones en todos los discos en un esquema pseudo-aleatorio




Distribución del XIV cuando hay cambios

- Se distribuyen los datos sólo cuando hay cambios en el sistema
 - El equilibrio se mantiene cuando se agrega nuevo hardware
 - El equilibrio se mantiene cuando se remueve hardware viejo
 - El equilibrio se mantiene cuando hay una falla de hardware



Distribución del XIV cuando hay cambios



- Se distribuyen los datos sólo cuando hay cambios en el sistema
 - El equilibrio se mantiene cuando se agrega nuevo hardware
 - El equilibrio se mantiene cuando se remueve hardware viejo
 - El equilibrio se mantiene cuando hay una falla de hardware



[hardware upgrade]

Distribución del XIV cuando hay cambios

- Se distribuyen los datos sólo cuando hay cambios en el sistema
 - El equilibrio se mantiene cuando se agrega nuevo hardware
 - El equilibrio se mantiene cuando se remueve hardware viejo
 - El equilibrio se mantiene cuando hay una falla de hardware



Software



Software incluido en el XIV

- Thin provisioning
 - Sin limitaciones de arquitectura
 - Administración de la capacidad y seguimiento simple y directo
 - No se consume espacio si los datos son sólo 0
- Snapshots
 - Creación y borrado de snapshots es instantáneo
 - Cantidad casi ilimitada de SNAPs
 - SNAPs diferenciales ahorran entre un 15% a un 30% de capacidad
 - SNAP de SNAP (con clones)
- Replicación Remota
 - Baja granularidad - replicación de volúmenes “any to any”
 - Cada I/O tiene commit en las copias locales y remotas antes de completarse
 - Se usa el Snap automático para mantener las copias auto-consistentes incluso durante la resincronización después de una falla de comunicaciones
- Automatic data migration
 - XIV se ubica entre los Servidores y el storage actual
 - Migración de volúmenes “Thick” a “Thin”
 - Migración de datos en línea desde cualquier storage



IBM XIV Snapshots - Virtually without Limits

- La creación/eliminación de los Snapshot es instantanea

- Alta performance con ejecución de snapshots



- Ilimitado

- Snapshots

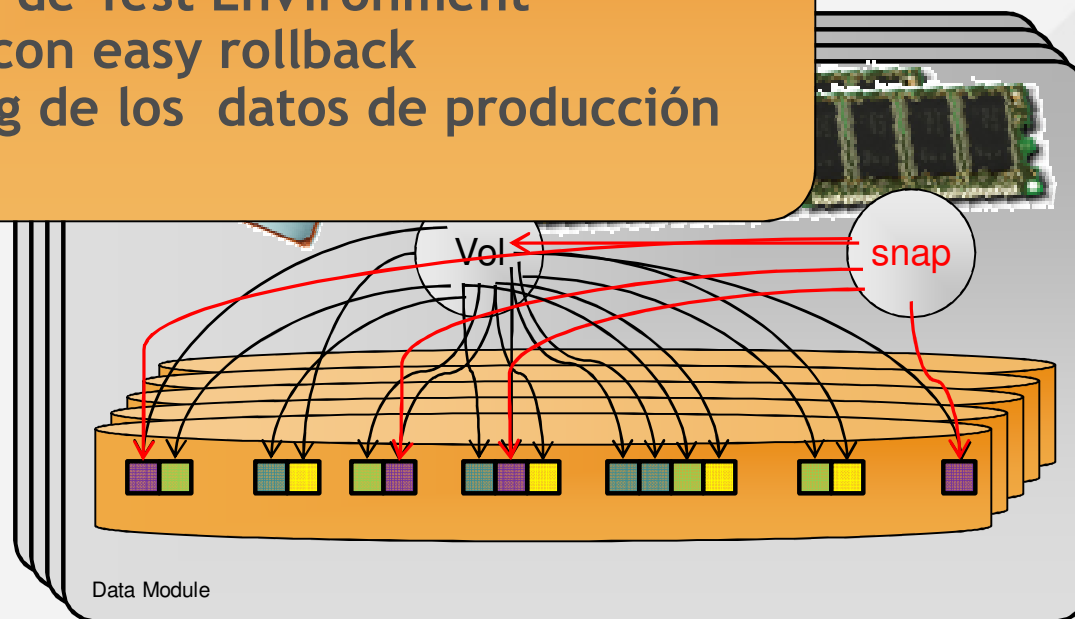
- Snapshots

Alta performance snapshots proporciona:

- Simplicidad para el Backup físico a Tape
- Recupero instantaneo de Logical Backup
- Simple creación de Test Environment
- Boot-from-SAN con easy rollback
- Easy Data-Mining de los datos de producción

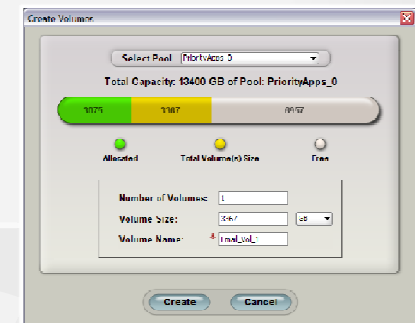


Restore Volume from snapshot copy
Each Server Interpolates its memory to
On a snapshot, each server simply
randomly across system daily local links
points to original volume. memory only
operation



Facilidad de uso

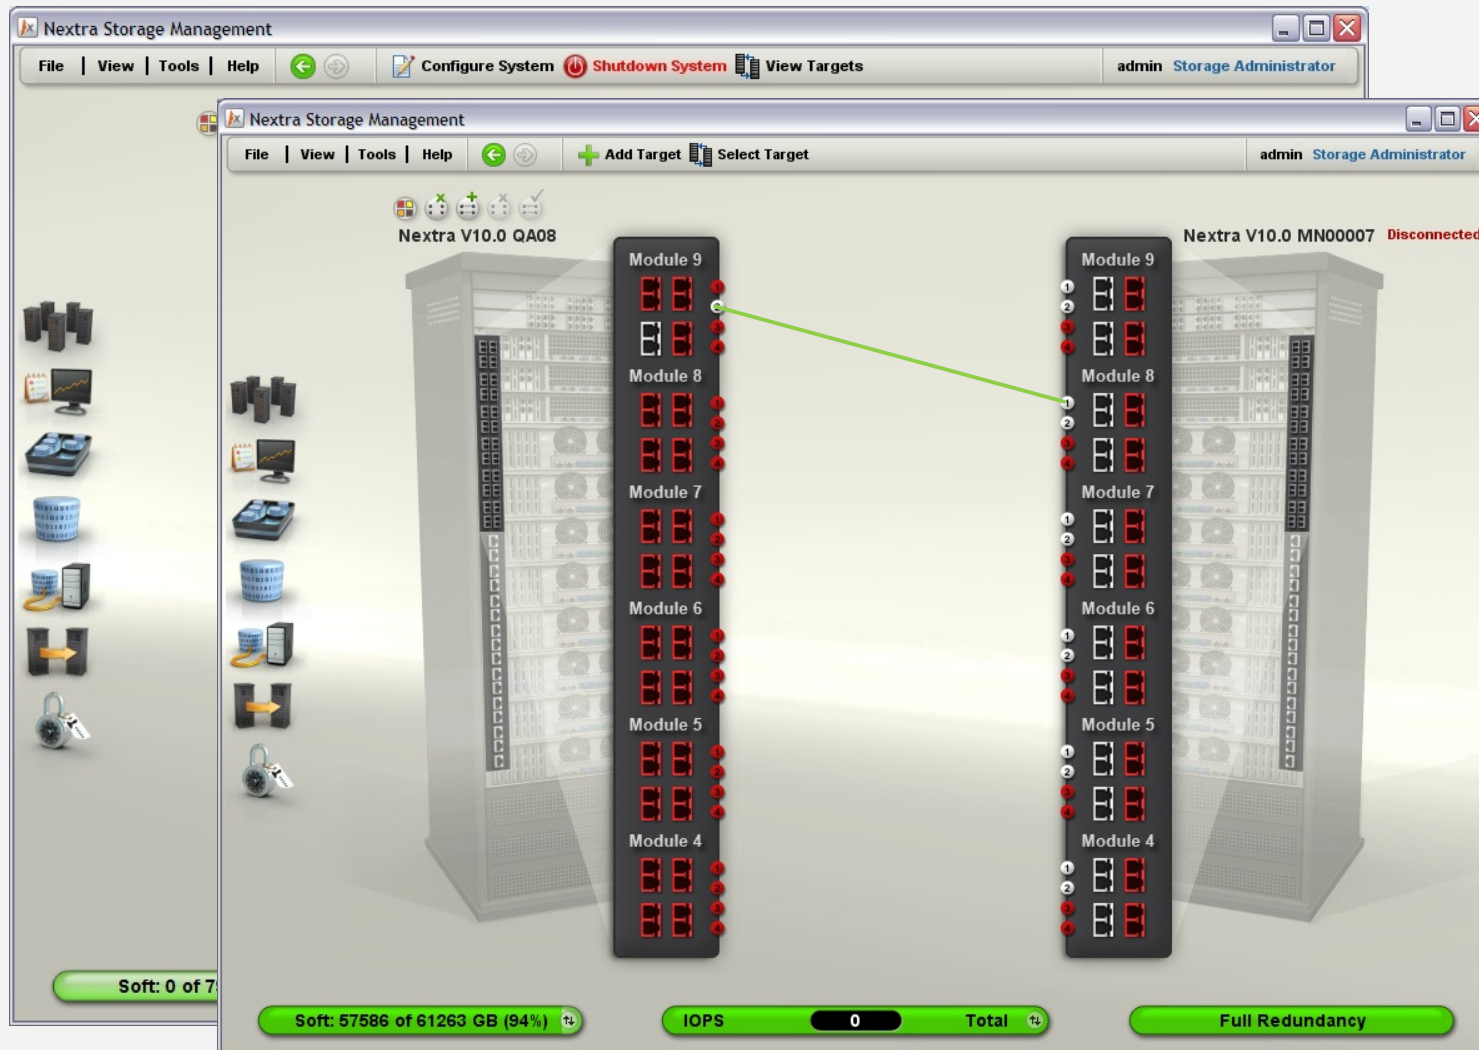
- Use el tiempo para provisionar storage:
 - Definir volúmenes en segundos
 - Redimensionar volúmenes sin dolor
 - Crear snapshots instantáneos
 - Crear ambientes de test con snapshots escribibles
- No pierda tiempo optimizando:
 - No se necesita tuning de rendimiento
 - No se necesitan layouts de disco complejos
- Administración basada en roles:
 - Administradores de aplicaciones manejan sus propios snapshots
 - Integración con LDAP de la organización
- Soporte de TPC a través de SMI-S



☒ **Simplicidad ahorra tiempo y dinero**

☒ **Mejor servicio a la organización**

Facilidad de uso



Administración intuitiva: Creación de un Volumen

The screenshot shows a 'Create Volumes' dialog box with the following elements:

- Select Pool:** A dropdown menu showing 'PriorityApps_0'.
- Total Capacity:** Text indicating '13400 GB of Pool: PriorityApps_0'.
- Progress Bar:** A horizontal bar divided into three segments: green (3075), yellow (3367), and grey (6957).
- Legend:** Three colored circles with labels: a green circle for 'Allocated', a yellow circle for 'Total Volume(s) Size', and a grey circle for 'Free'.
- Input Fields:**
 - Number of Volumes:** A text box containing the value '1'.
 - Volume Size:** A text box containing '3367' and a unit dropdown menu set to 'GB'.
 - Volume Name:** A text box containing '*Email_Vol_1'.
- Buttons:** 'Create' and 'Cancel' buttons at the bottom.

La capacidad utilizada se conoce siempre

Resizing de volumenenes


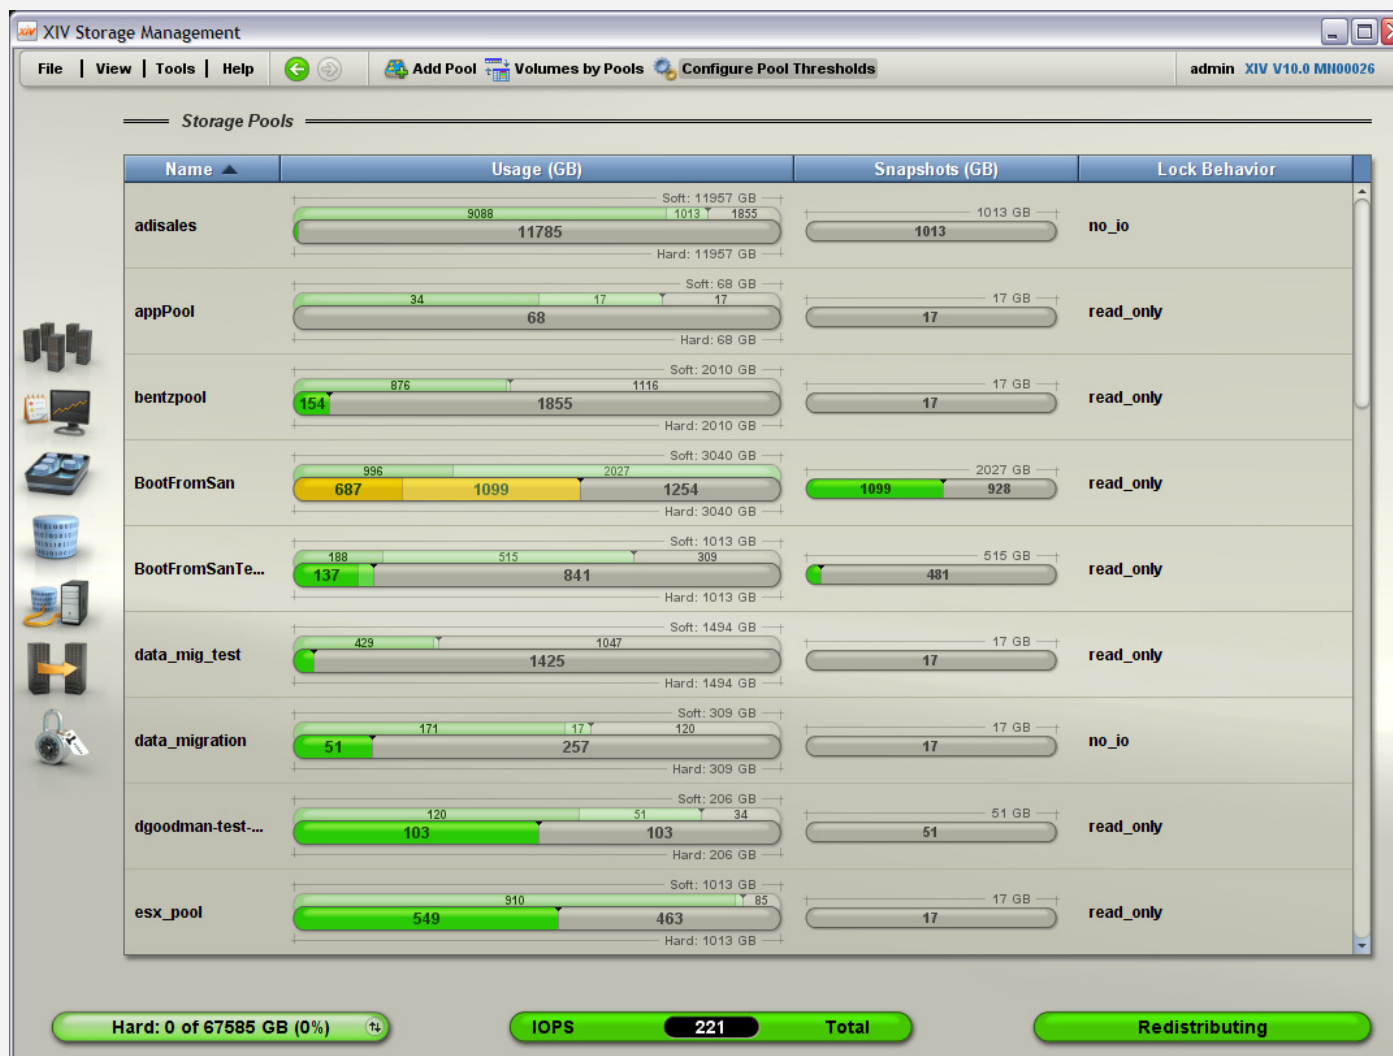


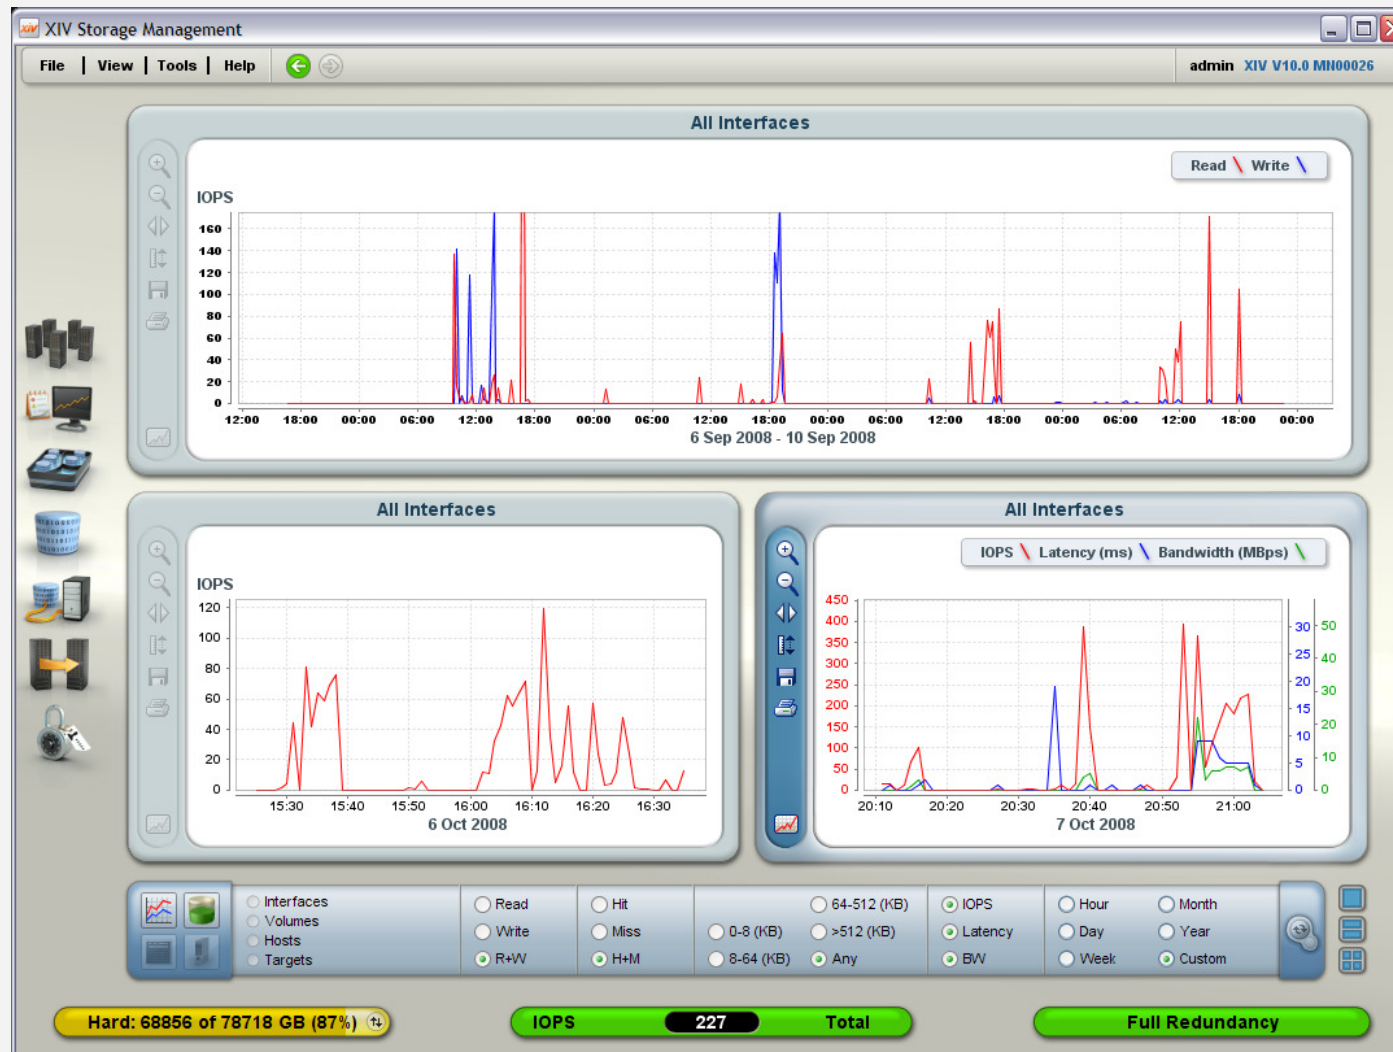
Figure 4-37 Resize an existina volume



Administración intuitiva : Manejo de Storage Pools



Monitoreo en XIV



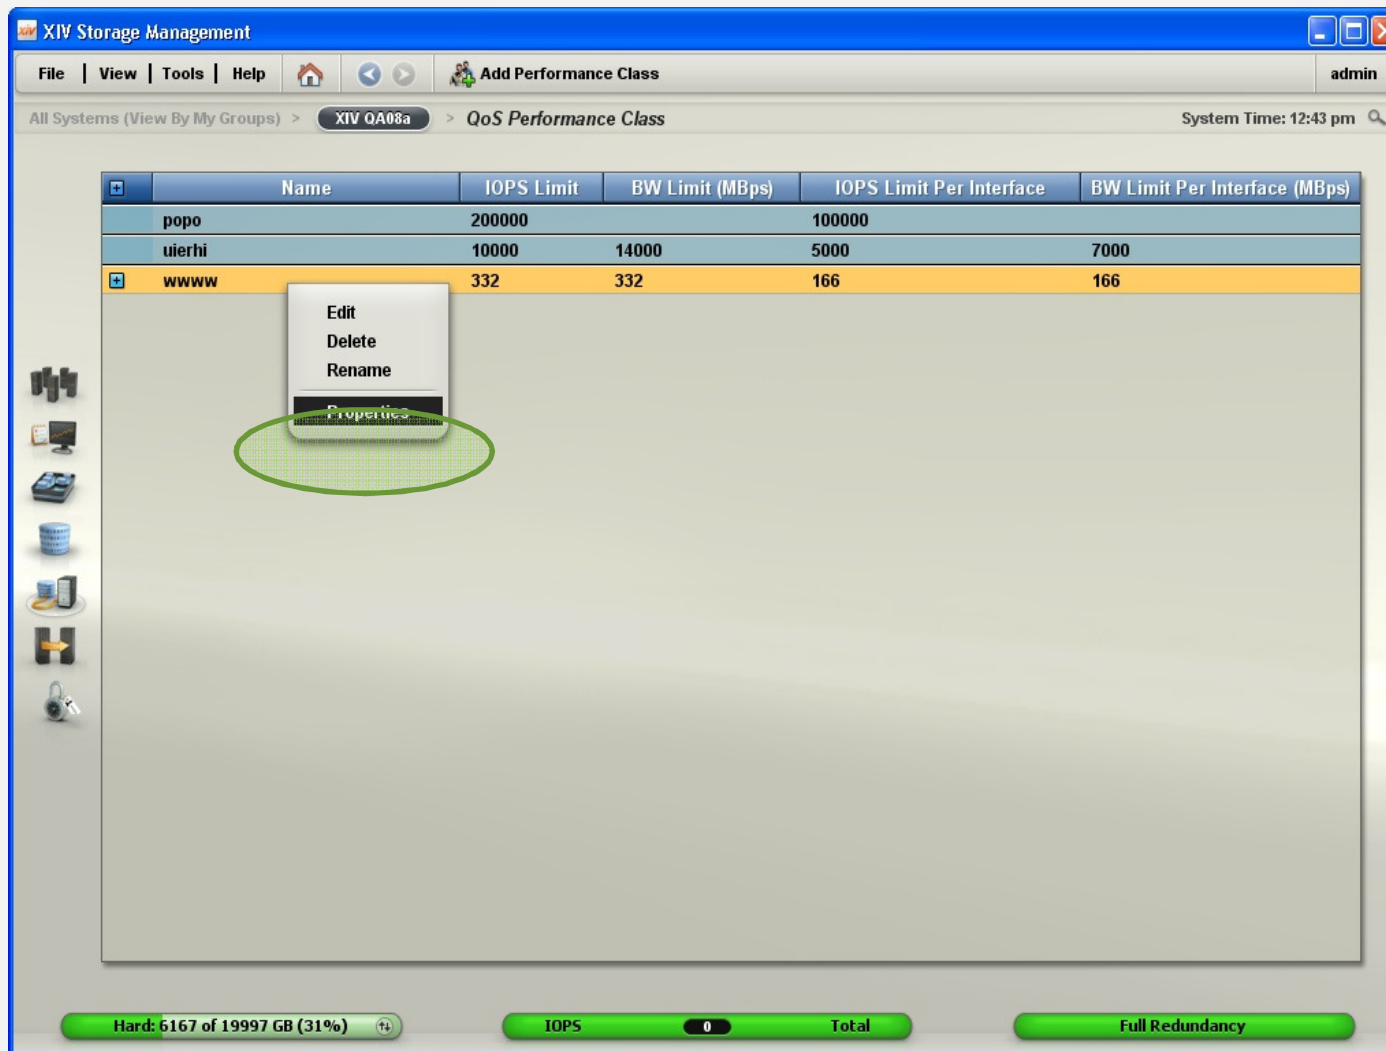
Migración de datos



QoS Working with Performance Classes from the XIV GUI



Viewing Performance Classes Properties



XIV Storage Management

File | View | Tools | Help | Add Performance Class | admin

All Systems (View By My Groups) > XIV QA08a > QoS Performance Class | System Time: 12:43 pm

Name	IOPS Limit	BW Limit (MBps)	IOPS Limit Per Interface	BW Limit Per Interface (MBps)
popo	200000		100000	
uierhi	10000	14000	5000	7000
www	332	332	166	166

Context Menu for 'www':

- Edit
- Delete
- Rename
- Properties**

Bottom Status Bar:

- Hard: 6167 of 19997 GB (31%)
- IOPS: 0
- Total
- Full Redundancy

Performance Classes Properties

XIV Storage Management

File | View | Tools | Help

Home

Back

Forward

Add Performance Class

admin

All Systems (View By My Groups) > XIV QA08a > QoS Performance Class

System Time: 12:43 pm

Name	IOPS Limit	BW Limit (MBps)	IOPS Limit Per Interface	BW Limit Per Interface (MBps)
popo	200000		100000	
uierhi	10000	14000	5000	7000
www	332	332	166	166

Performance Class Properties

Name:

www

IOPS Limit:

332

BW Limit (MBps):

332

IOPS Limit Per Interface:

166

BW Limit Per Interface (MBps):

166

OK

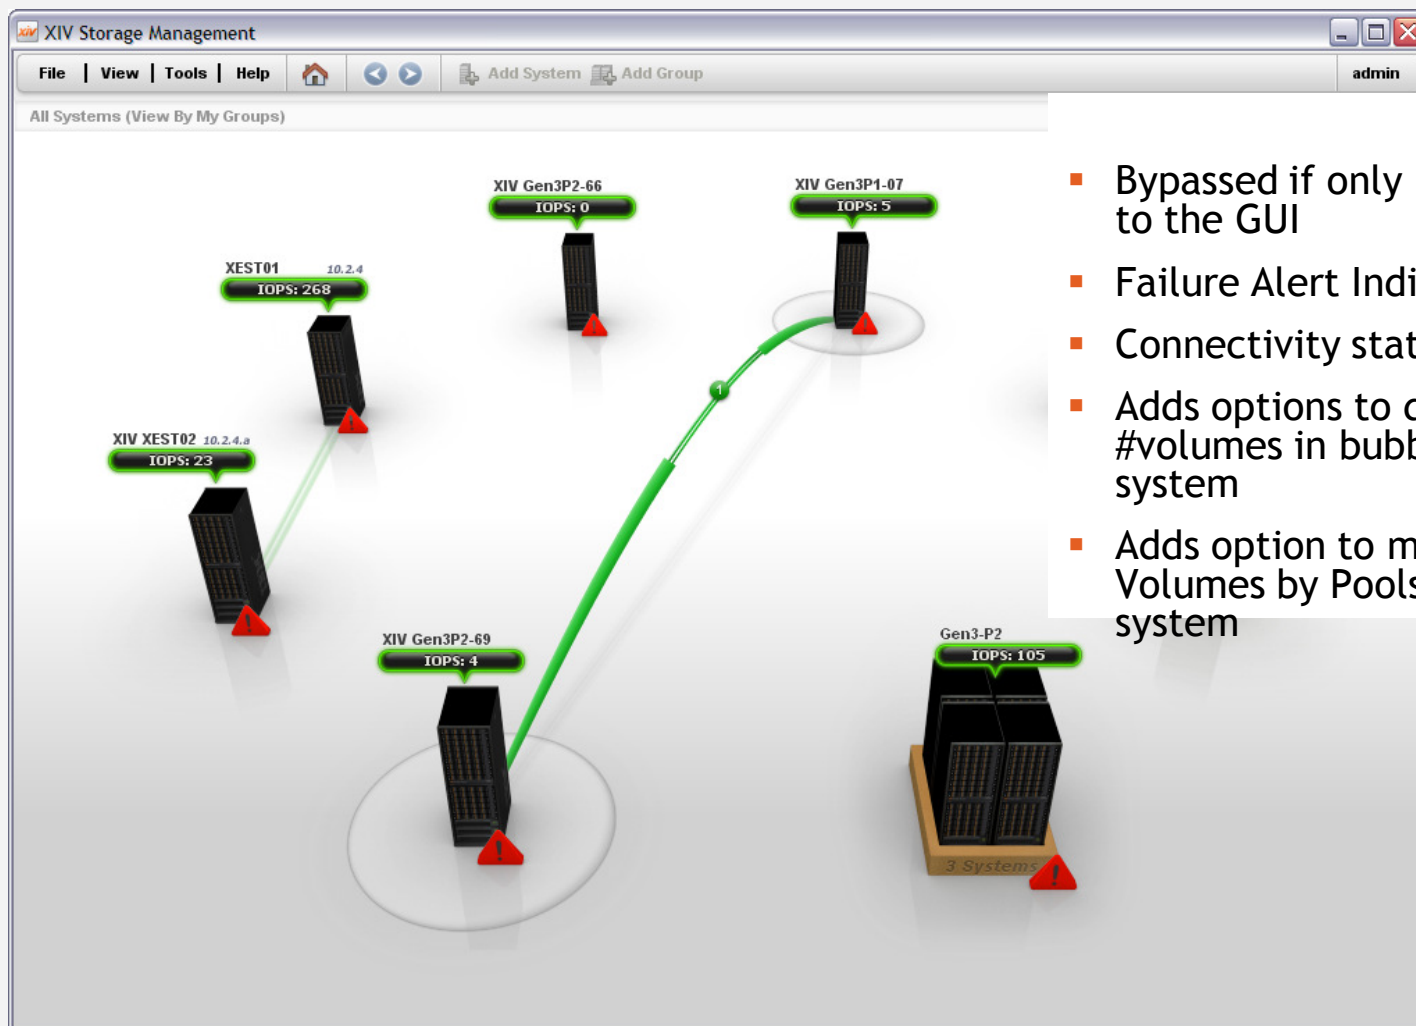
Hard: 6167 of 19997 GB (31%)

IOPS 0 Total

Full Redundancy



All Systems Panel



- Bypassed if only 1 XIV is defined to the GUI
- Failure Alert Indications
- Connectivity status
- Adds options to display #hosts or #volumes in bubble for specific system
- Adds option to move to view Volumes by Pools for specific system



¿Preguntas?

Muchas gracias !

