



XIV Nextra™

Storage Reinvented



Information Explosion Creates Storage Challenges

How much data does mankind store?

- IDC* says about 281 exabytes in 2007
- By 2011, we'll reach 1,773 exabytes
- That's 600% growth in 4 years

281,000 PB

Clients need a simple solution to address modern information storage challenges

1,773,000 PB

*IDC White Paper Sponsored by EMC, "The Diverse and Exploding Digital Universe," March 2008

What if...

You had a disk array for today's challenging world?

- You no longer needed to manage spindles?
- You no longer needed tiers of storage to meet SLA's, cost, and Green goals?
- Performance scaled with capacity
- Your storage was self balancing and healing?
- Your storage was designed to accommodate new technology?
- SNAPshots were fast, easy, & affordable
- You could recover from a failed 1TB drive in < 30 minutes!
- You could provision storage in under 1 minute!

Your choice: disks or systems?

Usable TB	#XIV sys. Systems	Traditional arrays	
		#Disks	#RAID grps
40	1	176	22
80	1	344	43
160	2	680	85
240	3	1,024	128
480	6	2,040	255
960	12	4,080	510

IBM XIV Storage Profile

- XIV was acquired by IBM on December 31, 2007
 - Moshe Yanai XIV CEO (father of EMC's Symmetrix) was appointed an IBM Fellow
- XIV has a unique virtualized grid technology that changes the disk system paradigm: Storage Reinvented
- XIV innovation has been tested and accepted
 - More than 100 systems in production to date
 - 100% of evaluation systems have gone into production at end of testing
 - Product in development for more than 6 years
 - More than 3 years in production
 - More than 50 patents filed
- XIV now part of IBM. For our customers, this means:
 - Next-generation storage product
 - IBM integration, support and service



Key Attributes for Enterprise Information Infrastructure

- **Performance** - Consistent performance under all conditions by eliminating hot spots
- **Reliability** - Business data more critical than ever, with no tolerance for downtime
- **Functionality** - Functions (e.g. replication, thin provisioning) that scale without performance penalties and are inherently built-in to the architecture
- **Power and Space Savings**- “Green”, Minimize power usage, cooling and floor-space
- **Manageability** - Simplicity and ease of use so the data explosion does not require an explosion of work to manage
- **Cost** - Reasonable capital cost and minimal ongoing cost - so business can concentrate its efforts on its core and not on IT

Information explosion can compound challenges for the enterprise

Current Enterprise Storage Solutions

Building blocks:


- Disks
- Cache
- Controllers
- Interface
- Interconnect

PERFORMANCE
RELIABILITY
SCALABILITY

→
→
→


= \$\$\$

With this legacy architecture, scalability is achieved by using more powerful (and more expensive) components



Available Solutions Add Cost and Complexity: Creating the Need for ILM

- ILM tries to cope with storage pains by using multi-tiered storage
 - Tiered storage management and data classification is costly and complex
 - Excessive data movements create reliability and performance issues
 - Utilization rates remain low (50% or less), with limited ability to execute thin provisioning



Imagine prioritizing electricity at home...



Laundry Power?



Lamp Power?



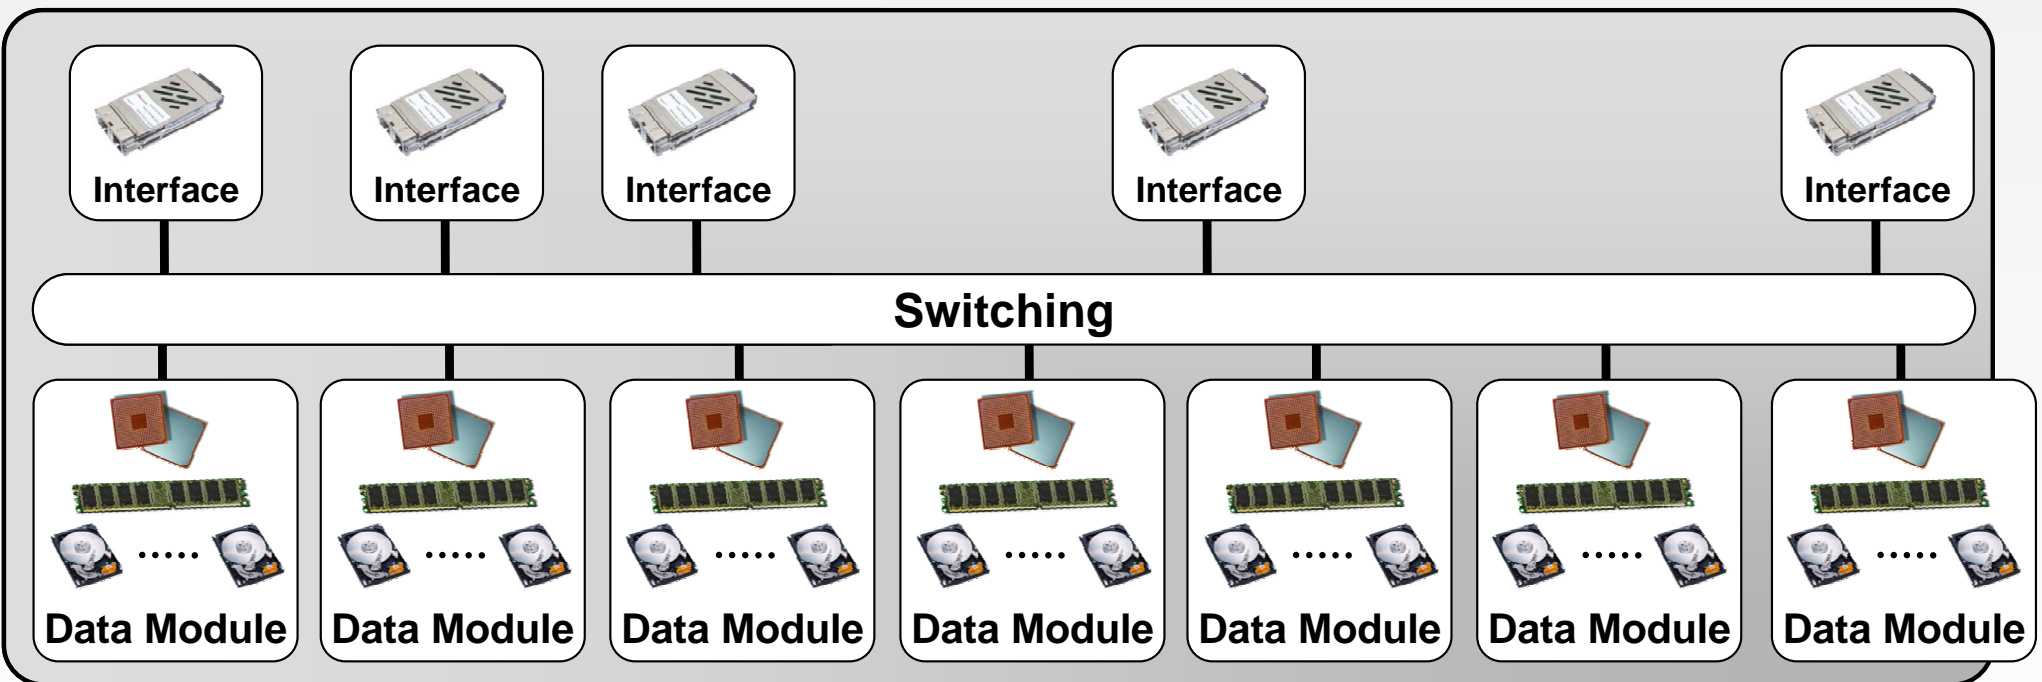
TV Power?

XIV Nextra Architecture

Design principles:

- Massive parallelism
- Granular distribution
- Off-the-shelf components
- Coupled disk, RAM and CPU
- User simplicity

Scale Out



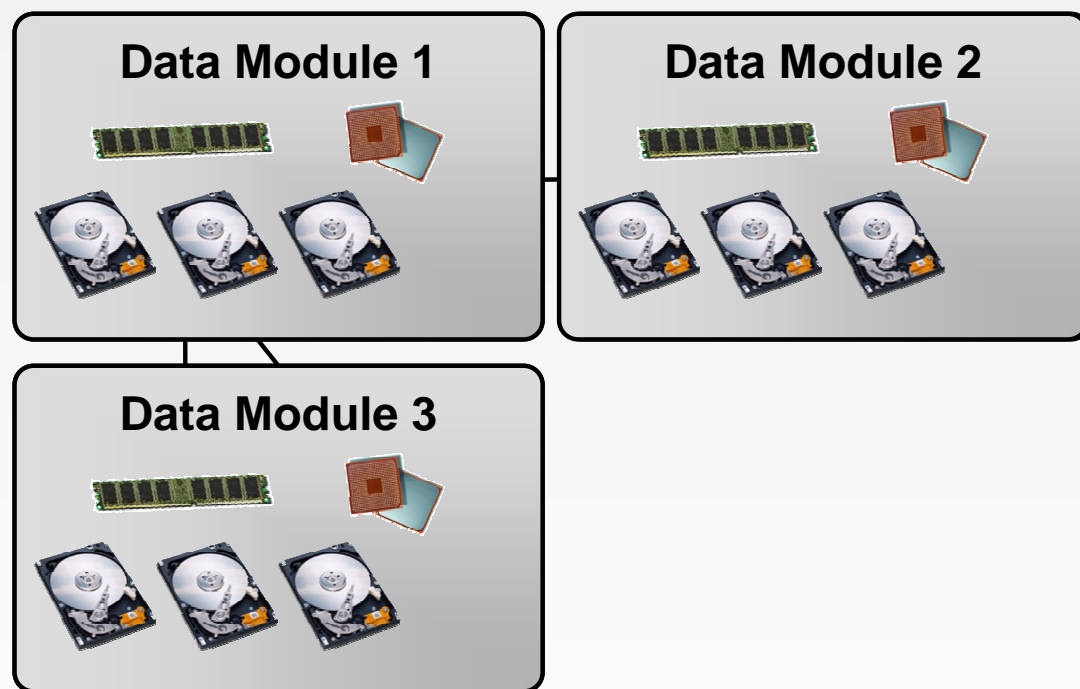

System Description

- 180 disks per rack
 - 12 disks in 2U modules
 - Unified data and interface modules
 - 15 modules per rack
- 1TB disks
- Single rack provides 80TB net capacity
- 120GB of memory per rack (8GB per module)
- Up to 24 4GB FC ports
- Up to 12 iSCSI ports
- 3 UPS modules



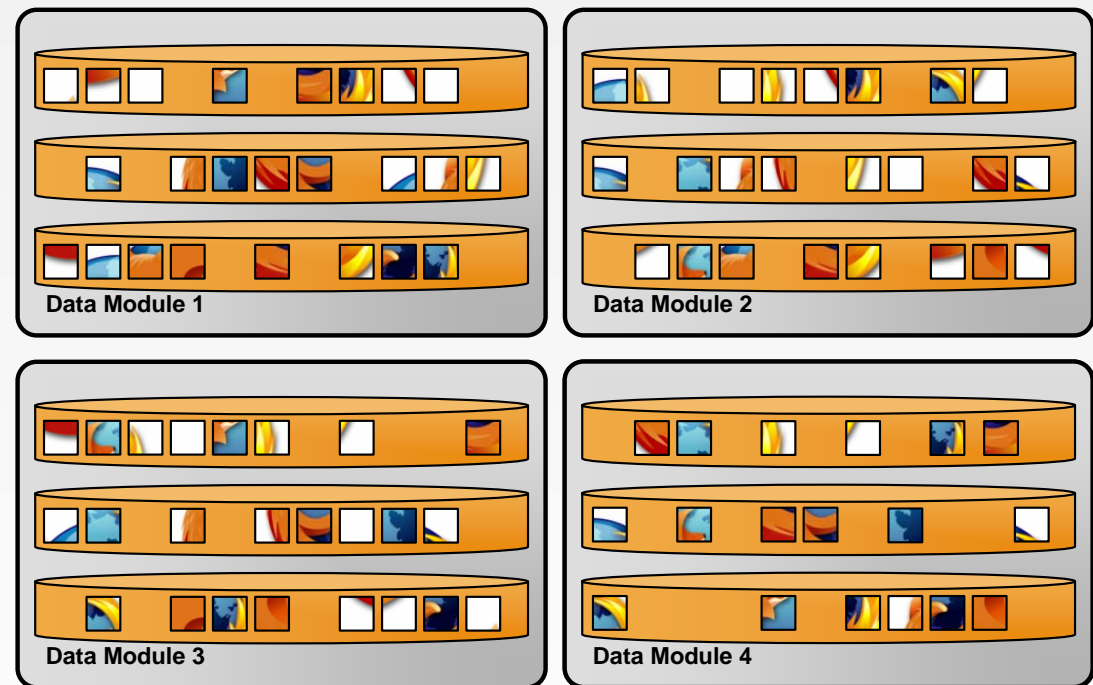
Nextra Distribution Algorithm on System Changes

- Data distribution only changes when the system changes
 - Equilibrium is kept when new hardware is added
 - Equilibrium is kept when old hardware is removed
 - Equilibrium is kept after a hardware failure



Nextra Distribution Algorithm on System Changes

- Data distribution only changes when the system changes
 - Equilibrium is kept when new hardware is added
 - Equilibrium is kept when old hardware is removed
 - Equilibrium is kept after a hardware failure






[hardware upgrade]

Nextra Distribution Algorithm on System Changes

- Data distribution only changes when the system changes
 - Equilibrium
 - Equilibrium
 - Equilibrium

The fact that distribution is full and automatic ensures that all spindles join the effort of data re-distribution after configuration change.

Tremendous performance gains are seen in recovery/optimization times thanks to this fact.




Nextra Distribution Algorithm


- Each volume is spread across all drives
- Data is "cut" into 1MB "partitions" and stored on the disks
- Nextra's disks behave like connected vessels, as the distribution algorithm aims for constant disk equilibrium.

Nextra disks behave like connected vessels, as the distribution algorithm aims for constant disk equilibrium.


Thus, Nextra's overall disk usage approaches 100% in all usage scenarios.




Data Module



Data Module




Data Module




Performance - Does not degrade over time!

Legacy Architectures



- COMPLEXITY GROWS WITH CAPACITY
- REQUIRES MORE Pre-Planning
- INCREASED RISK to existing production
- REQUIRES MORE monitoring & tuning
- LOWER capacity utilization

XIV Nextra



- COMPLETELY Automated Process
- NO Pre-Planning
- AUTOMATIC load balancing
- HIGHEST capacity utilization

Replacing Outdated Hardware

- New hardware can be added to the system
 - Better performance, less power, more density
- Outdated hardware can be phased out and removed
- All system components are replaced, with:
 - No down time
 - No host configuration
 - No administration effort


SNAPs with No Limitations

- SNAPs on each Server.
- High Performance
- Unlimited

High Performance, unlimited SNAPs provide:


- Easier Physical Backup to Tape
- Instant recovery from Logical Backup
- Easy creation of Test Environment
- Boot-from-SAN with easy rollback
- Easy Data-Mining on Production data

Distributed SNAP on each Server.
Memory operations
as fast as
on volumes



Restore Volume from SNAP copy
Each Server's intermediate is placed to
On a SNAP, each Server simply points to
randomly as it has its own data blocks
to original volume. Memory only

Operation



Thin Provisioning

- Defining logical volumes bigger than physical capacity
- Installing physical capacity only if and when needed
- Results:
 - Reduced overall direct storage cost
 - Storage expenses spread over time, exploiting price reductions
 - Reduced power, cooling, and space consumption
 - Easier management

Management: Creating a Volume

Create Volumes [X]

Total Storage Capacity: 5052 GB

515 944 3592

☒ **Allocated** ☐ **Added** ☐ **Free**

Number of Volumes: 5

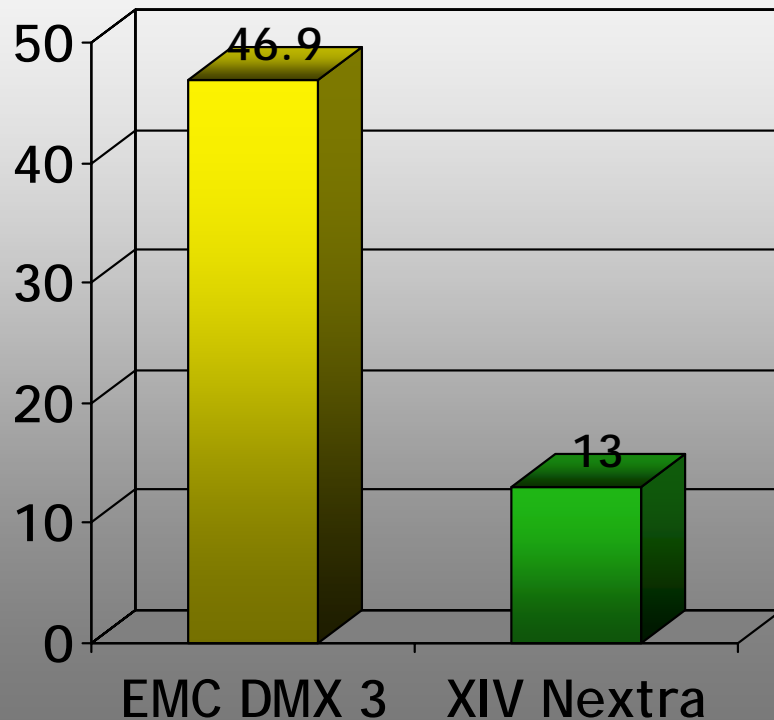
Volume size: 188 GB

Volume name: * Sys01 med Vol 1 ... 5

Create **Cancel**

Green IT: Dramatically Lower Power Consumption


Power Consumption in a 240TB System



- Approximately 1/5 of the power consumption per TB, when compared to an equivalent tier-1 system
- XIV Nextra with 1TB drives;
EMC DMX-3 with 146GB drives
- Source:
www.emc.com/products/symmetrix/DMX_series/pdf/C1304_Symmetrix_DM3_SS_Idv.pdf

Power Consumption per Real Requirements

- Thin provisioning
 - Save 20-50% of storage capacity
- No orphaned space due to virtualization
 - Save 10-20% of storage capacity
- Differential copies for backup instead of full copies
 - Save 15-30% of storage capacity
- Overall, the same requirements are met with storage capacity that is on an order of magnitude less
- Total result: huge power savings



XIV Nextra TCO

- Less storage needed, thanks to:
 - Thin provisioning
 - Management efficiency (no orphaned space)
 - Differential copies
- Reduced capital costs, power, cooling, and space
- Manage the same systems with less staff, due to ease of management
- Provide your organization with better response times
- Easily replace outdated modules with new ones
 - Get more capacity, performance, and power efficiency
 - No administrative effort

Scalability

2 Racks

4 Racks

8 Racks

2TB drives

Remote mirroring

Asynchronous mirroring

3-way mirroring

Availability

CCL to any version

Host interfaces

Infiniband

FCoE

10GBE

Performance

Infiniband Interconnect

SSD as caching device

Security

LDAP

Disk encryption

2Q 2009

4Q 2009

2010

The Bottom Line: Real-World Benefits

- **Performance**
 - Massive parallelism, spindle utilization, and cache effectiveness boost performance dramatically in all conditions
- **Reliability**
 - Revolutionary self healing takes minutes, not hours
- **Functionality**
 - Thin provisioning and replication built into the architecture
- **Manageability**
 - Simple, easy management; a logical volume has only two parameters: name and size
- **Cost**
 - Off-the-shelf components
 - Power saving



Thank You

John Sheehy
jes@e-techservices.com

