




IBM yazılım zirvesi '09



IBM XIV System Storage İşinizde Verim,Hız ve Güven

Serkan Acar

Senior XIV SR

IBM Systems & Technology Group



Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

What's the pain with Storage today



- Web 2.0 causes massive **data explosion**
- **High administration effort**
 - Highly dynamic environment
 - Need to manually avoid bottlenecks
- **Price/performance** with traditional FC based systems
- Traditional storage systems **don't scale horizontally**
 - Specialized (=expensive) design
 - Dual controller strategy

Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

- Central Cache, CPU, Backplane
- Custom HW design & high cost
- Complex manual tuning
- Cost, performance functionality trade off
- Long time to market

- Parallel Grid Architecture
- Commodity components/cost
- Tier 1 performance & functionality
- Self Healing/Autonomic tuning
- Rapid Time to Market

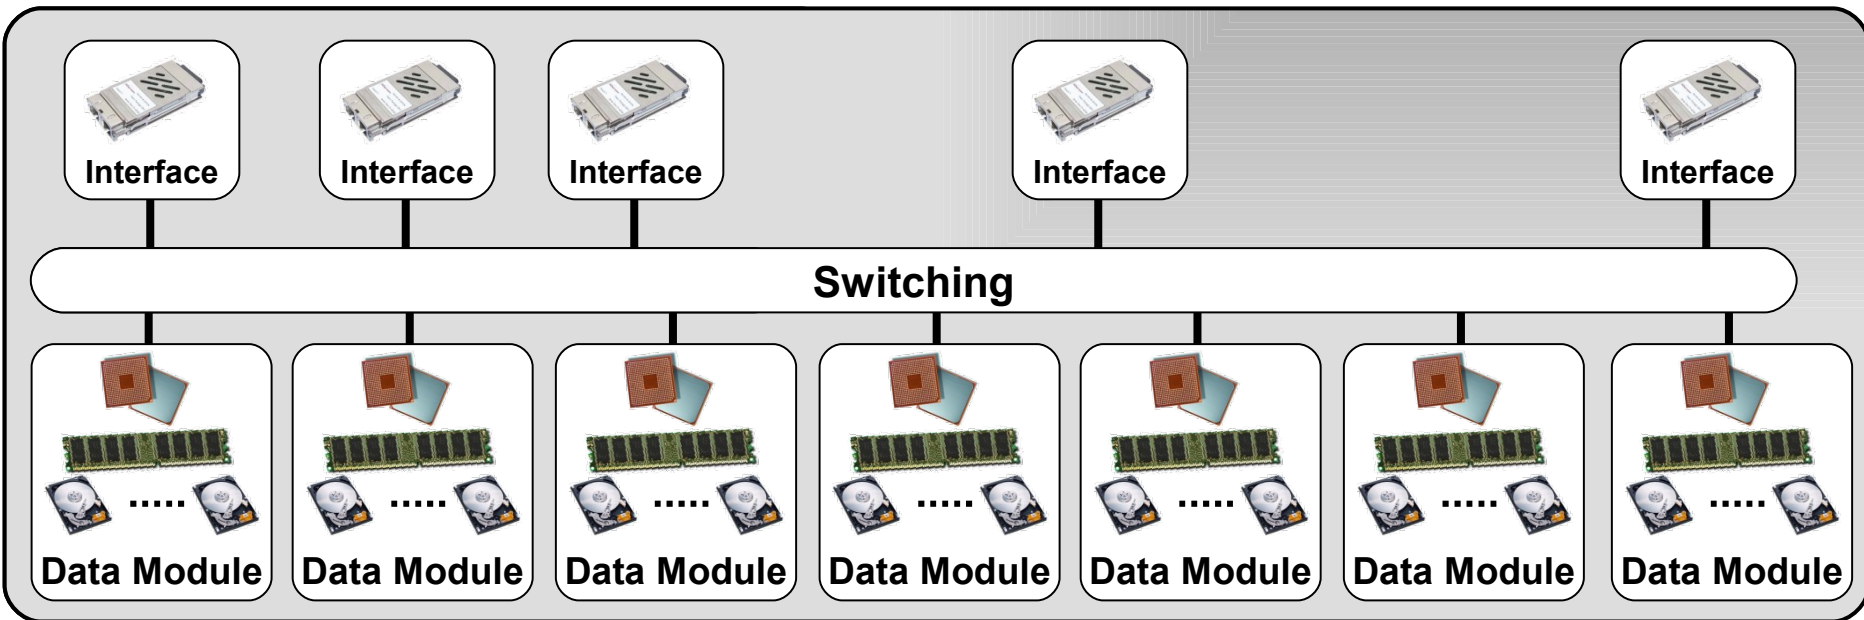


IBM XIV Storage System Architecture

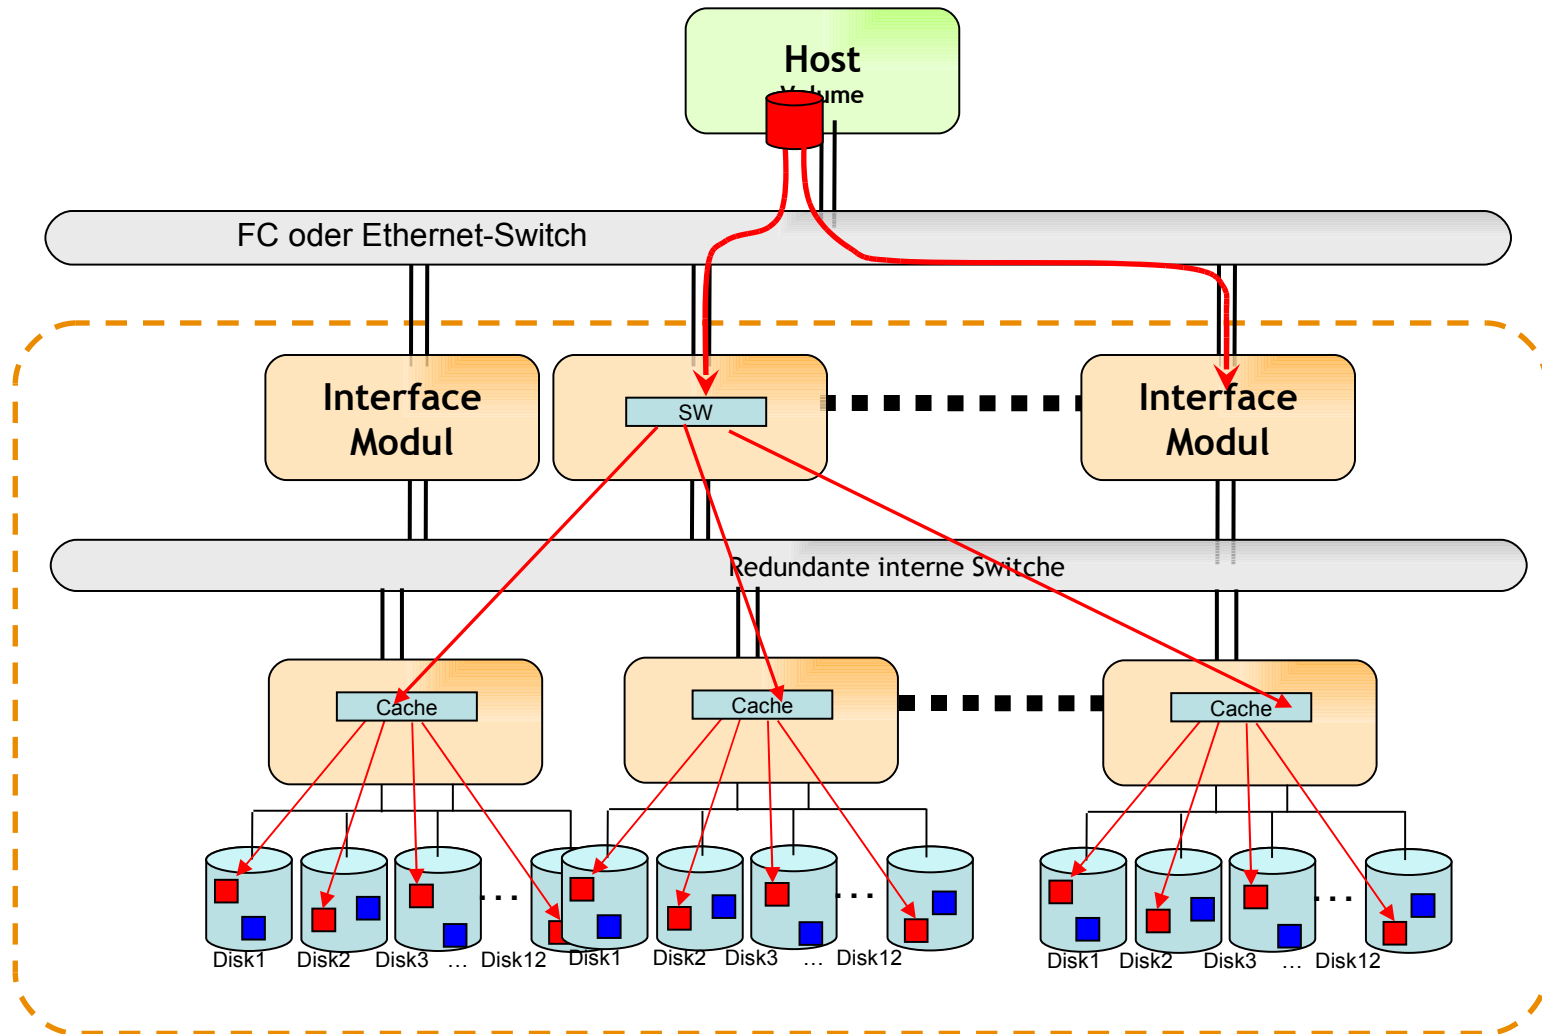
Design principles:

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

Scale Out



IBM XIV Storage is a Grid Architecture Solution



IBM XIV Storage System HW Details

- Available Models (12/2008)

	Full Rack	Half Rack
Modules	15	6
Disks	180	72
Cache	120 GB	48 GB
Net Capacity	79,1 TB	27,2 TB
FC-Ports	24	8
SCSI-Ports	6	0
Redundant UPS	yes	yes



- Module by module upgrade planned for 2009

Agenda

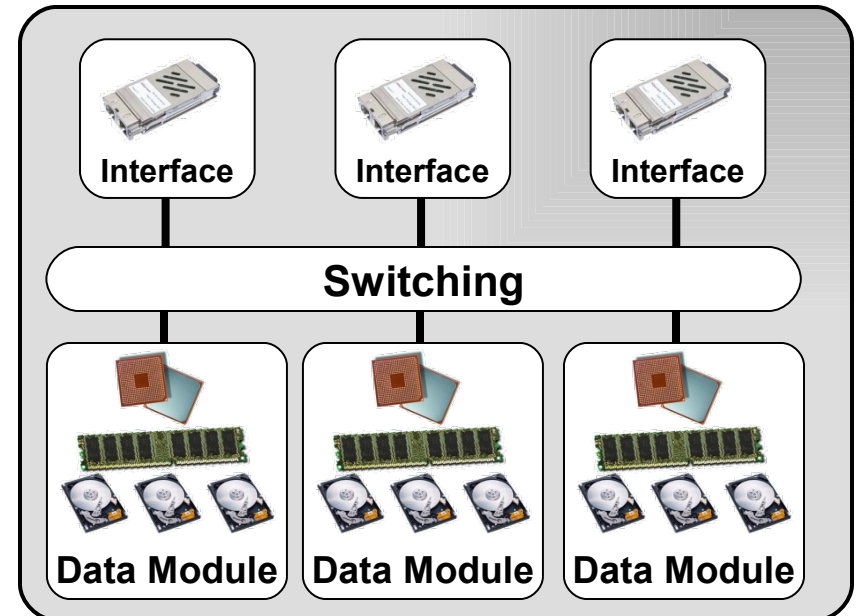

- Why IBM acquired XIV
- A new Architectural Concept: XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

Overview of Functions and Features

- **Virtualization**
- Thin Provisioning
- Snapshots
- Remote Mirroring
- Management

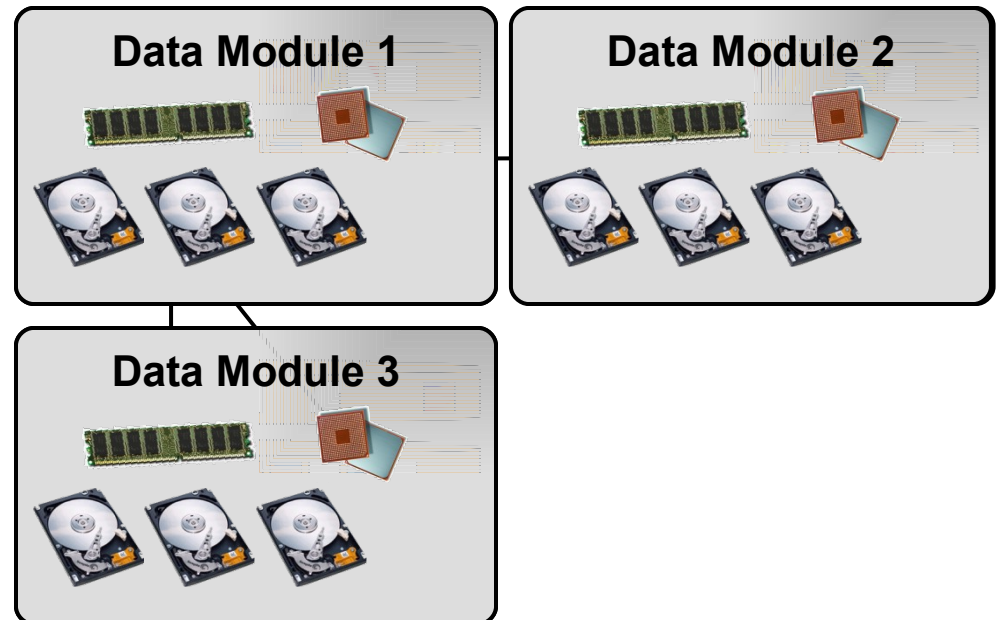

Virtualization

- IBM XIV Storage System has a unique data distribution technique
- Each volume is spread across all drives
- Data is “cut” into 1MB “partitions” and stored on the disks
- XIV algorithm **automatically** distributes partitions across **all** disks in the system pseudo-randomly



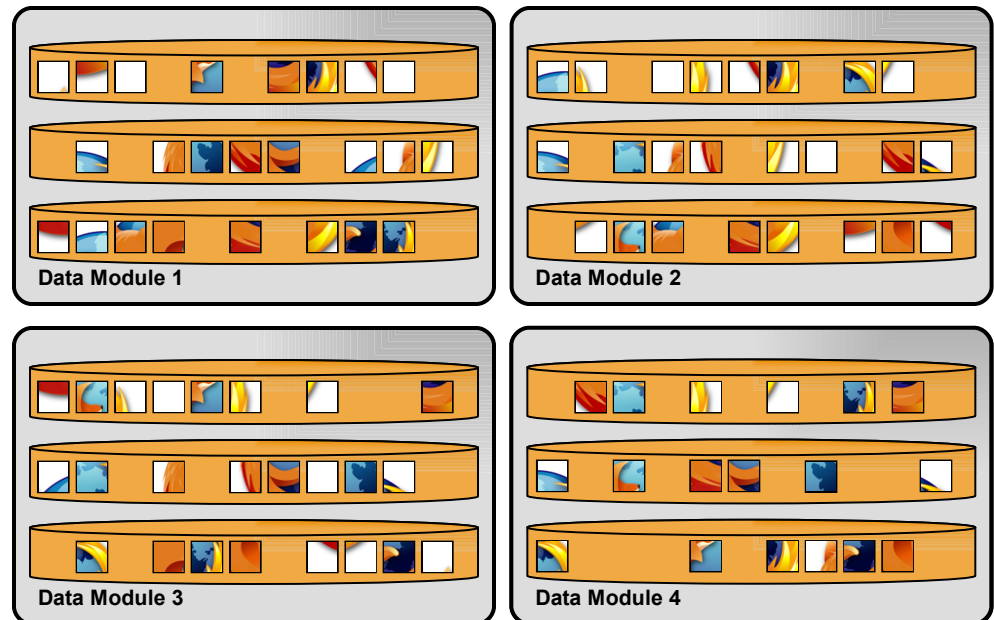

IBM XIV Storage Distribution Technique - System Change

- 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



IBM XIV Storage Distribution Technique - System Change

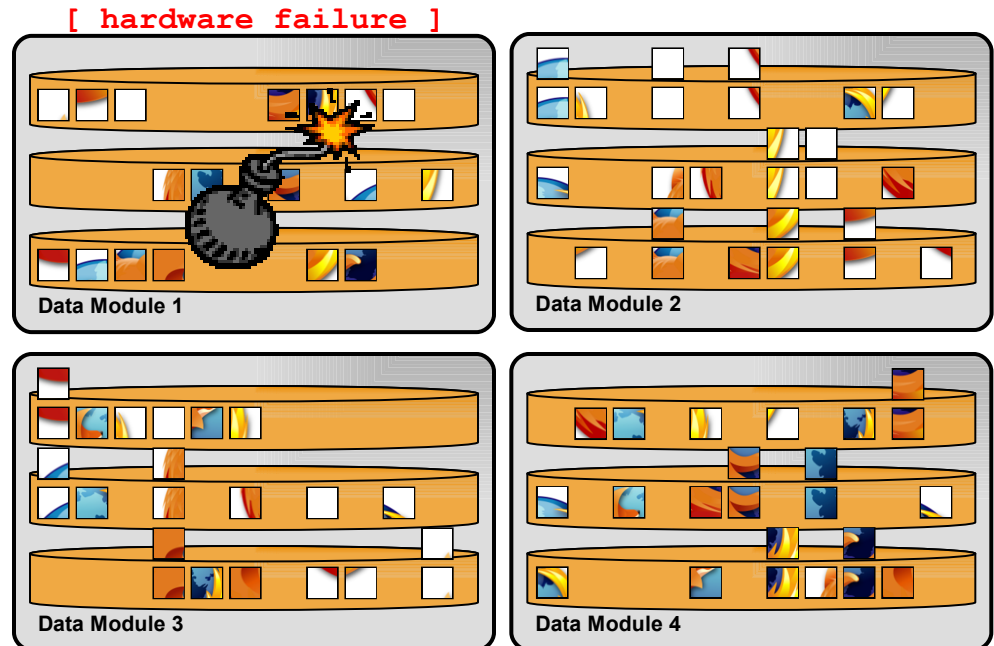

- Data distribution only changes when the system changes
 - Equilibrium is kept when new hardware is added
 - Equilibrium is kept when old hardware is removed
 - Equilibrium is kept after a hardware failure



[hardware upgrade]

IBM XIV Storage Distribution Techniques - System Change

- 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

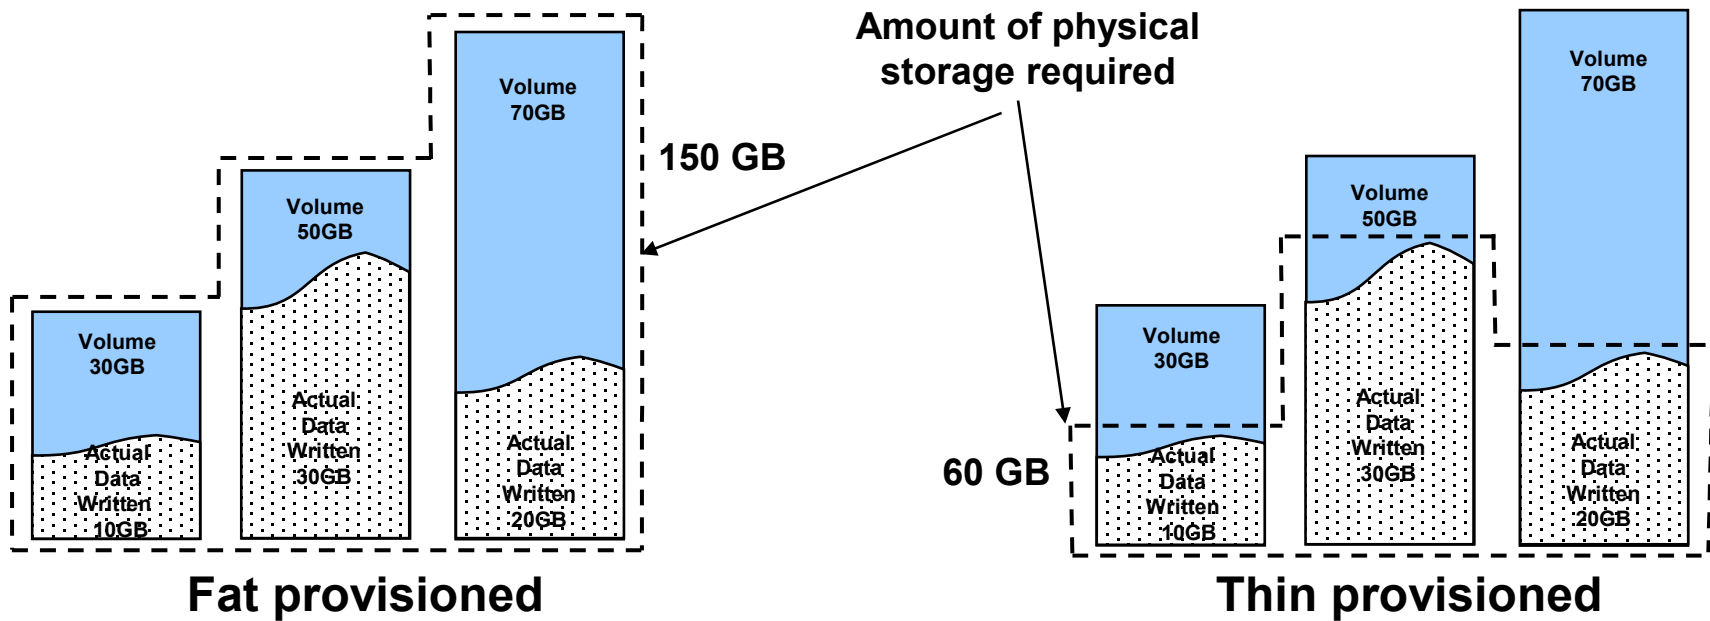


Overview of Functions and Features

- Virtualization
- **Thin Provisioning**
- Snapshots
- Remote Mirroring
- Management

Thin Provisioning – What is it?

- Users define volumes with any logical size
- Users acquire only the physical capacity of XIV Storage needed for data that is actually written
 - The part of the volume that contains no data **does not consume any physical space**



Overview of Functions and Features

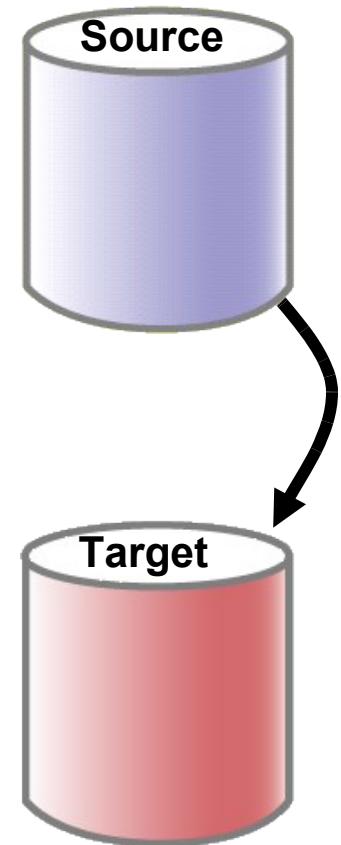
- Virtualization
- Thin Provisioning
- **Snapshots**
- Remote Mirroring
- Management

IBM XIV Snapshots

- Snapshot creation/deletion is instantaneous
 - Takes 150 ms...for any size of system, any capacity
- High performance WITH snapshots
- Unlimited number of snapshots
- Differential snapshots save 15-30% of storage capacity
- Snapshots on snapshots (with clones)

Snapshot Features

- Snapshot support
 - Differential
 - Full Copy
 - Multiple Targets
 - Snapshots of Snapshots
- Up to 16,000 Snapshots supported
- Writable Snapshots
- Consistency Groups

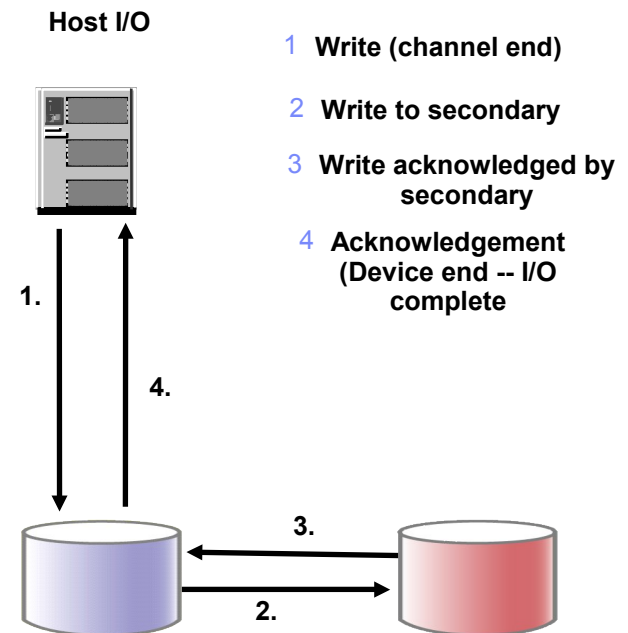


Overview of Functions and Features

- Virtualization
- Thin Provisioning
- Snapshots
- **Remote Mirroring**
- Management

IBM XIV Storage - Remote Mirroring

- Remote Mirroring for Disaster Recovery
- Low granularity – any to any volume replication
- Every I/O is committed to local and remote copies before completion
- Various policies upon link failure
 - Re-sync when link is resumed
 - Full completion or Fail



Overview of Functions and Features

- Virtualization
- Thin Provisioning
- Snapshots
- Remote Mirroring
- **Management**



If he says it about XIV
too . . .

IBM XIV Storage - Management Tools

```

alex@alex-laptop: ~/XIVGUI - Shell No. 3 - Konsole
alex@alex-laptop:~/XIVGUI$ ./xcli -m 9.155.56.100 -u xclitest -p passw0rd mapping_list host=x342_alex
LUN Volume Size Master Serial Number Locked
1 x342_alex_full_1 103 7 no
2 x342_alex_full_2 103 8 no
3 x342_alex_full_3 103 9 no
4 x342_alex_full_4 103 10 no
5 x342_big_01 412 296 no
6 x342_big_02 412 297 no
7 x342_big_03 412 298 no
8 x342_big_04 412 299 no
9 x342_big_05 412 300 no
10 x342_big_06 412 301 no
11 x342_big_07 412 302 no
12 x342_big_08 412 303 no
13 x342_big_09 412 304 no
14 x342_big_10 412 305 no
15 x342_500_01 515 306 no
16 x342_500_02 515 307 no
17 x342_500_03 515 308 no
18 x342_500_04 515 309 no
19 x342_500_05 515 310 no
20 x342_500_06 515 311 no
21 x342_500_07 515 312 no
22 x342_500_08 515 313 no
23 x342_500_09 515 314 no
24 x342_500_10 515 315 no
alex@alex-laptop:~/XIVGUI$
    
```

The screenshot shows the XIV Storage Management GUI. At the top, it displays 'XIV V10.0 MI00050'. The main interface features a 3D rendering of a server rack with 15 bays. To the left, there is a 'Volumes' menu with options: 'Volumes & Snapshots', 'Snapshots Tree', 'Volumes by Pools', 'Consistency Groups', and 'Snapshots Group Tree'. At the bottom, there are three status indicators: 'Hard: 64613 of 78718 GB (82%)', 'IOPS: 1952', and 'Full Redundancy'.

IBM XIV Storage - Management Tools

```
alex@alex-laptop: ~/XIVGUI - Shell No. 3 - Konsole
alex@alex-laptop:~$
LUN Volume Size Master Serial Number LUN Key
1 x342_alex_full_1 103 no
2 x342_alex_full_2 103 no
3 x342_alex_full_3 103 no
4 x342_alex_full_4 103 no
5 x342_big_01 412 296
6 x342_big_02 412 296
7 x342_big_03 412 296
8 x342_big_04 412 299
9 x342_big_05 412 300
10 x342_big_06 412 300
11 x342_big_07 412 302
12 x342_big_08 412 303
13 x342_big_09 412 305
14 x342_big_10 412 305
15 x342_500_01 515 306
16 x342_500_02 515 306
17 x342_500_03 515 306
18 x342_500_04 515 309
19 x342_500_05 515 310
20 x342_500_06 515 310
21 x342_500_07 515 312
22 x342_500_08 515 313
23 x342_500_09 515 313
24 x342_500_10 515 315
alex@alex-laptop:~/XIVGUI$
```

- Intuitive GUI (Java based) with Script Generator

- No dedicated management station

- Command Line Interface (CLI)

- XML over SSL

- Event management (SNMP)

- Complete Event Logging

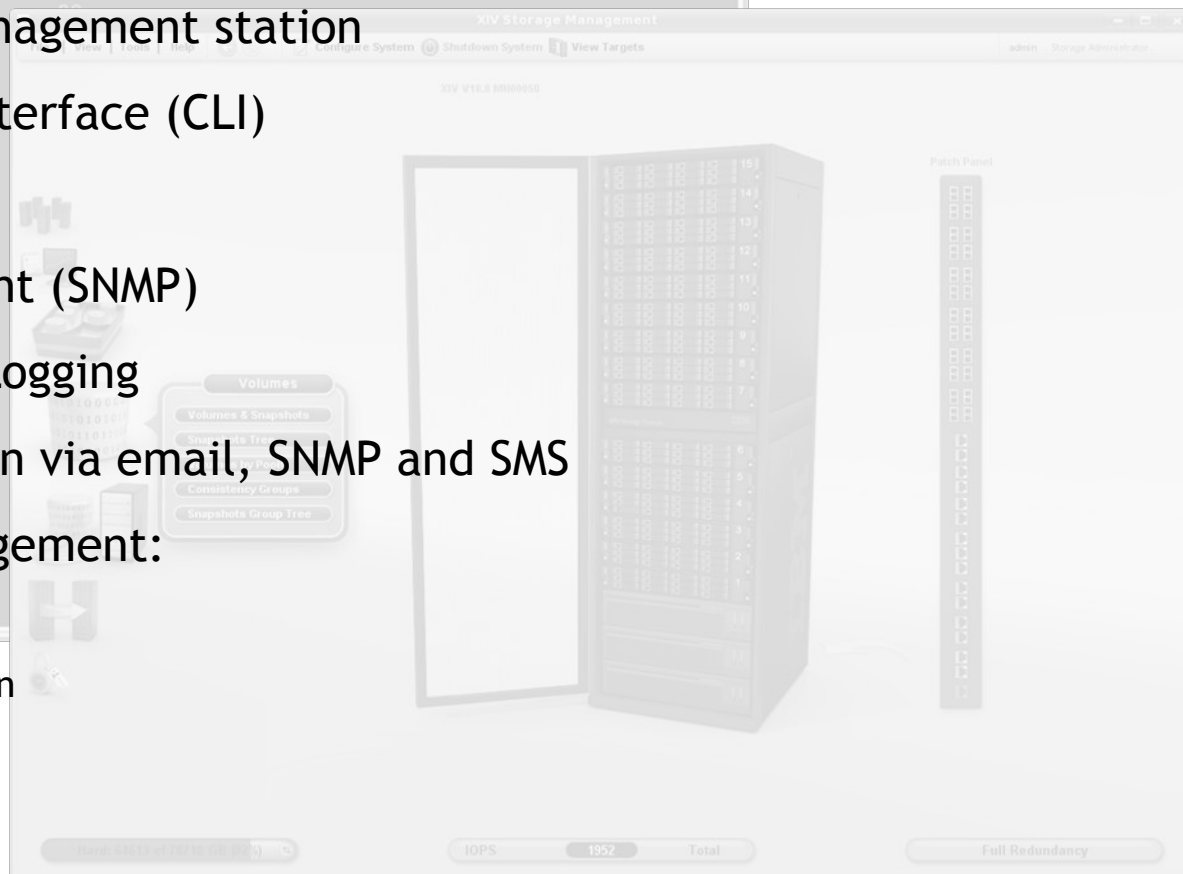
- Events notification via email, SNMP and SMS

- Role based management:

- Storage Admin

- Application Admin

- Operator



IBM XIV Storage Simple Intuitive Management example: Creating a Volume

Create Volumes

Select Pool: PriorityApps_0

Total Capacity: 13400 GB of Pool: PriorityApps_0

3075 3367 6957

Allocated Total Volume(s) Size Free

Number of Volumes: 1

Volume Size: 3367 GB

Volume Name: *Email_Vol_1

Create Cancel

- Used capacity is always known !

Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

The “Allegheney” 2-6-6-6



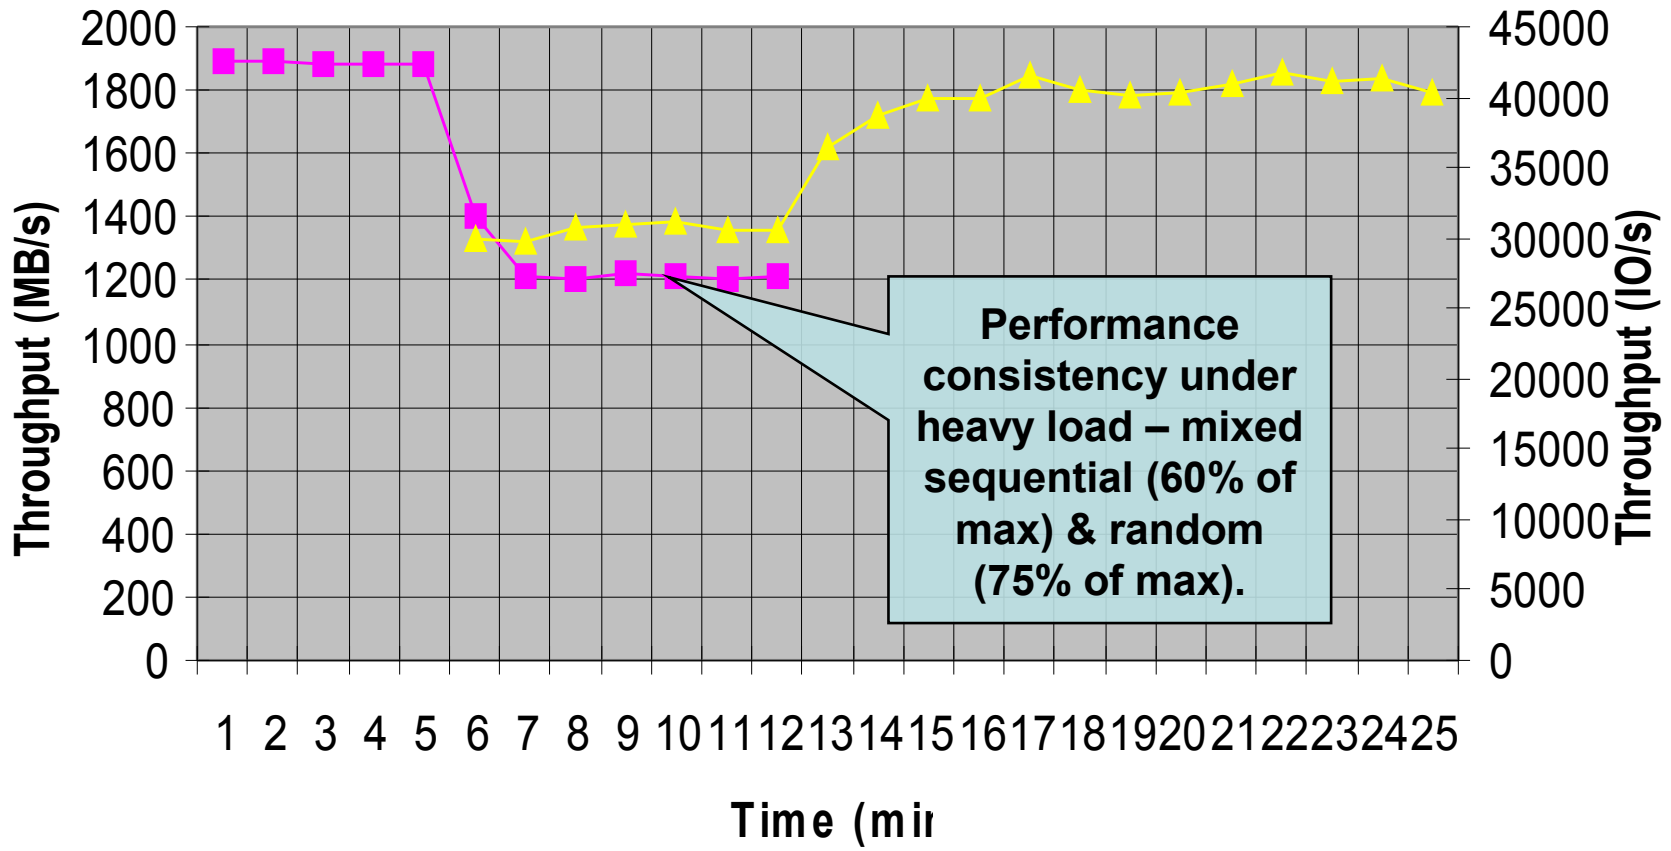
110,200 lbs of work power

Alco FA-2



61,000 lbs x 5 = 305,000 lbs of work power

Mix Workload- 256K Seq. Read/I



■ Throughput - 256K Seq. Read
 ▲ Throughput - DB0

Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

Self-healing


- Self-healing uses spare capacity, not spare hardware
- Self-healing at the module level
 - After a rebuild, the system can survive a second module failure
- Maintains redundancy even after one module and three disks fail
- No maintenance on system during rebuild (while system is non-redundant)

30-minute rebuild (1TB drives)

- Every disk drive is backed by all the other disk drives
 - Upon a failure, all disks participate in the rebuild process
- Only allocated volumes are rebuilt
- Only written data is rebuilt
- In practice, rebuild is over within minutes
- Almost no performance hit
- **Near-zero exposure to double-failure**


Down - 25% DBO (1 already failed)
rebuild 40GB, 34min

Response Time (ns)



Down - 75% DBO
rebuild 50GB (?), 35min

Response Time (ns)



Other reliability features

- Redundant UPS systems (2 out of 3 are enough)
- Dual power source
- Redundant Ethernet switches
- Disk monitoring for early fault detection
- Disk load balancing reduces failure rate
- Double-conversion UPS eliminates power spikes

Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

Power consumption: KW per TB

- Today: 7.5KW per 180TB raw, 79TB net
- 2TB disk drives will double power efficiency
- High-end architecture with SATA drives
 - Provide revolutionary power consumption rate without compromising performance

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, same requirements are met with storage capacity that is on an order of magnitude less

Agenda

- Why IBM acquired XIV
- A new Architectural Concept:
XIV
- Functions
- Performance
- Reliability
- Power Consumption
- Positioning
- Summary

The Bottom Line: Real-World Benefits

- **Reliability**
 - Revolutionary self healing takes minutes, not hours
- **Functionality**
 - Thin provisioning and replication built into the architecture
- **Power and Space**
 - Minimize power, cooling and floor-space with SATA drives
- **Performance**
 - Massive parallelism, spindle utilization, self-healing and cache effectiveness boost performance dramatically
- **Manageability**
 - Simple, easy management; a logical volume has only two parameters: name and size
- **Cost**
 - Off-the-shelf components
 - No charge for software features (Snap, DR, Management)



My Story . . .

To respond more quickly to market fluctuations and to cut costs, we needed to increase the efficiency of our IT infrastructure. We needed to gain more storage capacity and flexibility, and were also in need of a scalable solution that would meet our application requirements. The solution we chose had to be affordable. To accomplish this, we needed a solution that could be implemented robustly, highly available, and with multiple tiers of storage.

My XIV Solution . . .

IBM Business Partner helped us explore various alternatives and select a solution based on XIV without manual tuning. The solution was deployed without any IT staff intervention. The system equilibrium is kept automatically. As storage is added, changed or removed, the system was rapidly deployed and we became autonomous in a very short period of time, realizing direct benefits from the new technology.

Defined volumes and snapshots very easily and without requiring a plan for performance optimization using the system's built-in virtualization capabilities. Automatically and perfectly balanced loads across all applications and physical resources, handling volume provisioning in less than a minute and without any configuration process.

My Benefits . . .

With the XIV solution, we have been able to reduce administrator workload due to XIV's advancements in data placement and data loss prevention. We have rapidly deployed and migrated data from our existing DS 8100 device. Improved data management strategy and staff productivity have been achieved via the ease-of-use and rich feature set. We have been able to provide Tier 1 storage benefits at Tier 3 costs. We have defined volumes and snapshots very easily and without requiring a plan for performance optimization using the system's built-in virtualization capabilities. We now boast automatically and perfectly balanced workloads across all applications and physical resources, handling volume provisioning in less than a minute and without any configuration process.



My Story

Like most of our infrastructure business companies, Systems solution issues present a solution that provides performance, in terms of during cut-off window

The WOW aspects... Consistent Performance, the GUI and Reliability

s IT services to the When the Hitachi Data investigate new performance form by finding a could deliver reliable



My XIV Solution . . .

The company conducted committing to the XIV platform. After XIV training to JoongAng instruction, the company and maintain the new system one day," recalls SK Lee. and half without a single

“With XIV, we’ve improved system management and we’ve enabled automatic load-balancing,” reports JinSoo Lee. “We’ve also removed the bottlenecks during peak production, which had long been a thorn in our side. Now we have new levels of business efficiency, which translates to improved customer satisfaction.”



My Benefits . . .

JoongAng Ilbo has realized XIV storage, including cost administration of the storage infrastructure, it has helped the company reduce IT labor costs significantly. The contributor to that is the innovative XIV GUI, which SK Lee describes as “Fantastic— a groundbreaking tool.”

re, it has helped the company reduce IT labor costs significantly. The contributor to that is the innovative XIV GUI, which SK Lee describes as “Fantastic— a groundbreaking tool.”





My Story . . .

We were experiencing lots of growth and has a relatively small IT department that was being overburdened. We do not have a specific storage group, our UNIX sysadmins are also responsible for storage. So ease of management was a key criterion in our selection process .

My XIV Solution . . .

We've implemented two XIV systems - one for production use (not in production yet), and one for our development environment. We are attaching servers running VMware, AIX, Solaris, Windows to the XIV frames. We are also actively in testing with a new Exchange on VMware environment . This is currently in limit testing but is expected to grow.

My XIV Benefits . . .

Our primary storage admin has been extremely impressed with the ease of use offered by XIV when compared to our existing EMC CLARiiON environment. Our CIO and our IBM partner (VSS) have a very strong relationship, and VSS in conjunction with the XIV sales team were able to position XIV and how it would fit specifically to our existing and planned applications, and demonstrated how much better it was than what they currently had. We now consider the XIV team to a trusted advisor to Carquest, and we look forward to expanding our XIV footprint as we grow our business.





My Story . . .

We outsource our IT to IBM Global Services.. IGS was looking in general to use as few people as possible to manage our environments in order to keep costs under control. Fast, efficient deployment of storage is critical to meeting our growth demands and full redundancy was required to meet client SLA's.

My XIV Solution . . .

Basic implementation and knowledge transfer were provided as part of our IGS-managed XIV solution. Minimal help and knowledge transfer were all we needed to take over and begin using the XIV system. In just a few minutes after power-up, we had storage presented to SVC and began migrating data from our DS 4000 to the XIV.

My XIV Benefits . . .

The XIV's ease of use and management - especially in our heterogeneous environment with SVC as the virtualization layer - have proven to be a major benefit for us. The XIV's incredible (and simple) scalability will help us handle future growth and new projects as well as help us more efficiently manage older systems which are in the process of being retired. The strong relationship between IBM's local storage FTSS - who recommended the customer look at XIV for new storage needs - played a key role in our selection of XIV. The XIV team jointly presented the solution to us along with IGS, and the combination was a real win/win for us.

Serkan Acar

Serkan.acar@tr.ibm.com

+90 530 317 1783
serkanac@msn.com

Bu sunum 22 Ekim 2009 tarihinde İstanbul Swisotel the Bosphorus'da yapılan Yazılım Zirvesi 2009 için hazırlanmıştır.

<http://www.ibm.com/software/tr>

© Copyright IBM Corporation 2009. All Rights Reserved. IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information at www.ibm.com/legal/copytrade.shtml. Other company, product, or service names may be trademarks or service marks of others.