



Datacenter Storage og virtualiseringsløsninger

*Lars Bo Iversen
Storage Specialist
lars_bo_iversen@dell.com*

*Brian Lauge Pedersen
Technical Solutions Professional - Datacenter
blauge@microsoft.com*



Windows Server® 2008 R2

Microsoft TechNet

Session Objectives And Takeaways

- Session Objective(s):
 - Describe the use of Hyper-V with the Dell iSCSI Target
 - Describe the use of Failover Clustering with Hyper-V and the Dell iSCSI Target
 - Describe the use of Cluster Shared Volumes in Windows Server 2008 R2 with the Dell iSCSI Target

Notes

- There are many Virtualization options out there.
We focus on **Hyper-V** here.
- There are many Clustering options out there.
We focus on the **Failover Clustering** feature included with Windows Server 2008 R2 here.
- There are many iSCSI Initiator options out there.
We focus on the **Microsoft iSCSI Initiator** included with Windows Server 2008 R2 here.
- There are many iSCSI Target options out there.
We focus on the **Dell iSCSI Target**.

Agenda

1. Basic Considerations
2. Hyper-V with the Dell iSCSI Target
3. Failover Clustering with Hyper-V and the Dell iSCSI Target
4. Cluster Shared Volumes with the Dell iSCSI Target
5. Summary
6. Q&A



Windows Server® 2008 R2

1. Basic Considerations

*Before we focus on the Dell iSCSI Target,
a few basics on Hyper-V so you can
understand the scenario*

1.0. Basic Considerations

- All the usual options
 - Direct-Attach, FC SAN or iSCSI SAN?
 - How many spindles?
 - SATA, SAS, FC or SSD?
 - What RAID level?
 - Standalone or Clustered?
- Plus lots to decide specifically for Hyper-V
 - VHD (Virtual Hard Drive) file or pass-through disk?
 - Fixed-size, dynamic or differencing VHD files?
 - Virtual IDE, Virtual SCSI or iSCSI to child?

1.1. Options for Hyper-V (table 1)

How storage is exposed to the parent partition

	VHD file on Parent Partition	Pass-through disk on Parent Partition	Exposed <u>Directly</u> to Child as iSCSI
DAS (Direct-Attach)	X	X	
FC SAN	X	X	
iSCSI SAN	X	X	X

1.2. Options for Hyper-V (table 2)

Different types of VHD files

	Fixed-size	Dynamically Expanding	Differencing
Pre-allocated disk space	X		
Saves disk space		X	X
Files might grow during writes		X	X
Common base VHD			X

What type of virtual hard disk do you want to create?

☒ Dynamically expanding

The .vhd file grows as data is stored to the disk, up to the size you specify in this wizard. The .vhd file does not shrink automatically when data is deleted.

☐ Fixed size

The .vhd file uses the amount of space you specify for the disk size, regardless of how much data is saved to the virtual hard disk.

☐ Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily.

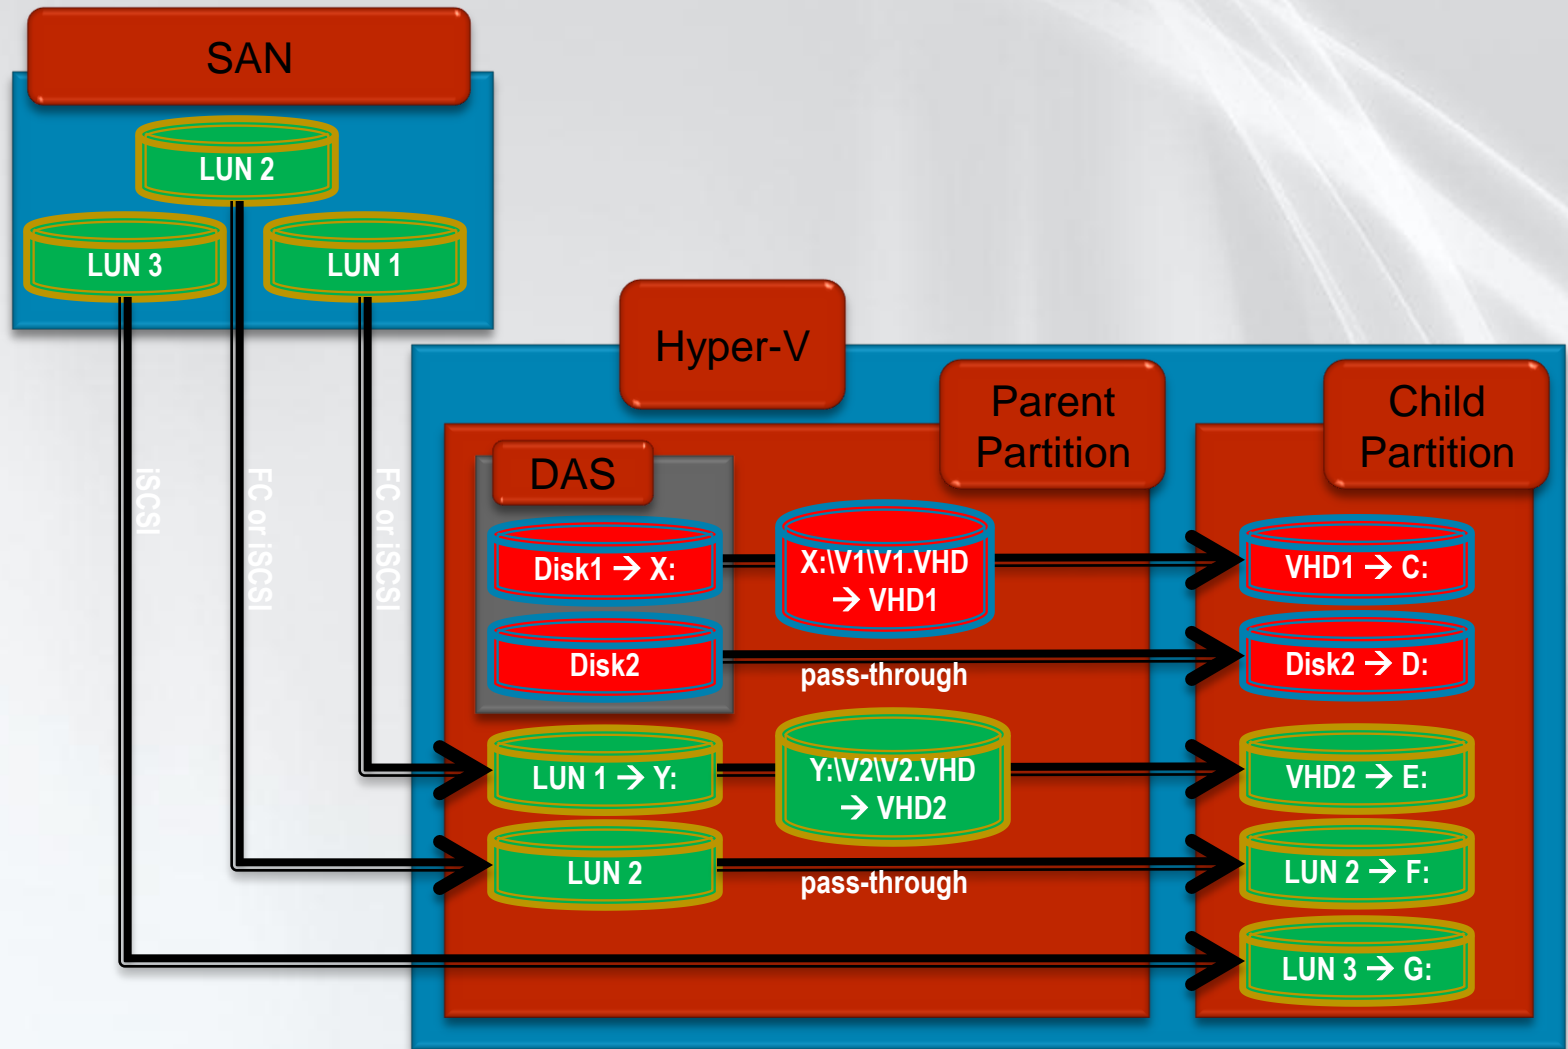
1.3. Options for Hyper-V (table 3)

How storage is exposed to the child partitions

	Exposed to Child as Virtual IDE	Exposed to Child as Virtual SCSI	Exposed <u>directly</u> to Child as iSCSI
Additional Software on Child	Integration Components (optional)	Integration Components	iSCSI initiator
Child sees disk as	Virtual HD ATA Device	Msft Virtual Disk SCSI Disk Device	MSFT Virtual HD SCSI Disk Device
Child max disks	2 x 2 = 4 disks	4 x 64 = 256 disks	Not limited by Hyper-V
Child hot add disk	No	Yes	Yes

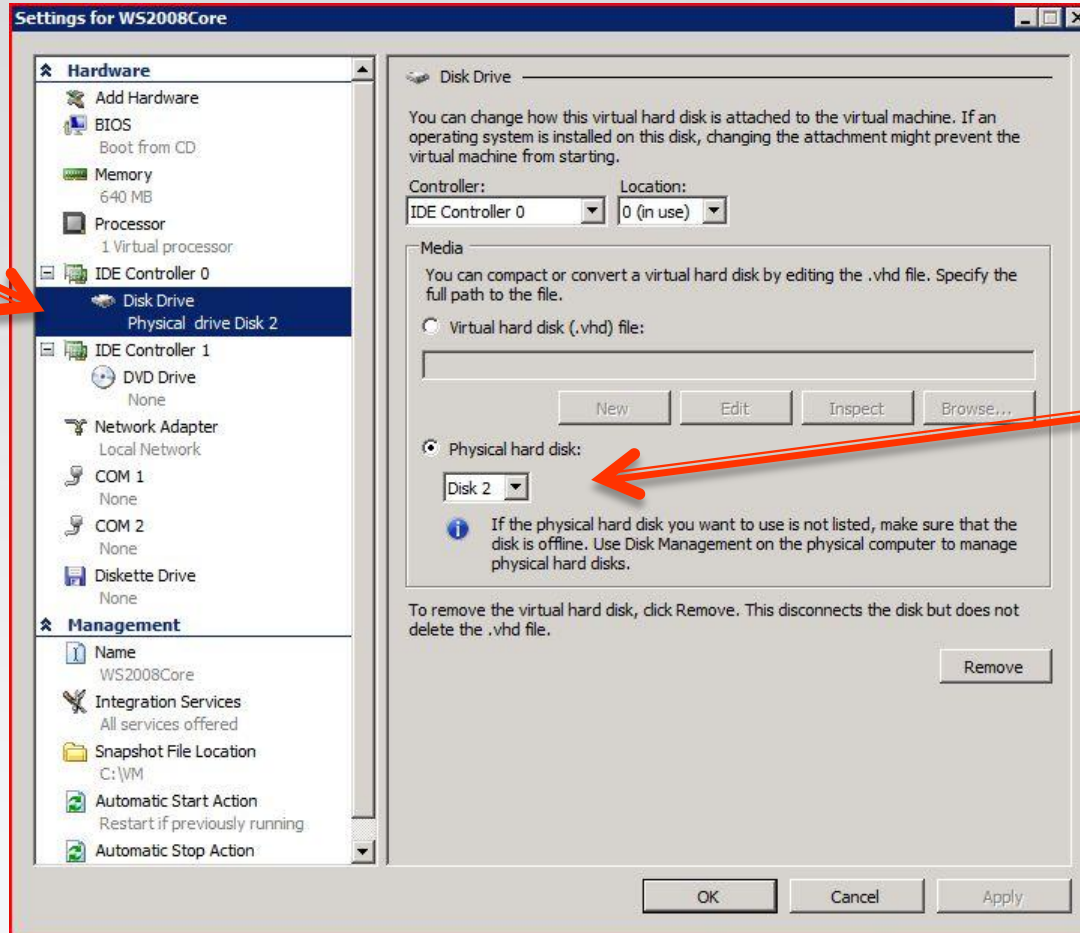
Windows Server 2008 R2

1.4. Options for Hyper-V



1.5. Configuration Screenshots

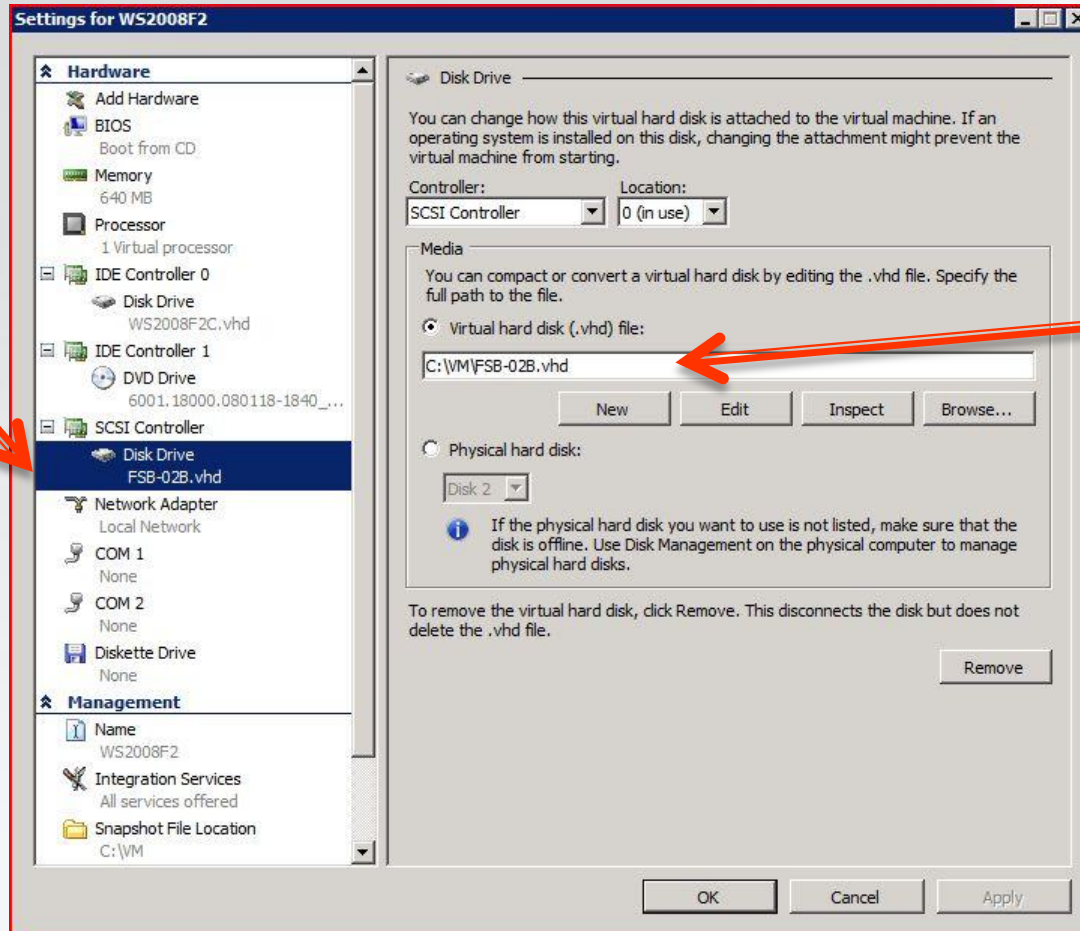
Virtual
IDE



Pass-through
disk

1.5. Configuration Screenshots

Virtual
SCSI



VHD File



Windows Server® 2008 R2

2. Hyper-V with Dell iSCSI Target

Understanding how to use the Dell iSCSI Target in conjunction with the Hyper-V role in Windows Server 2008 R2



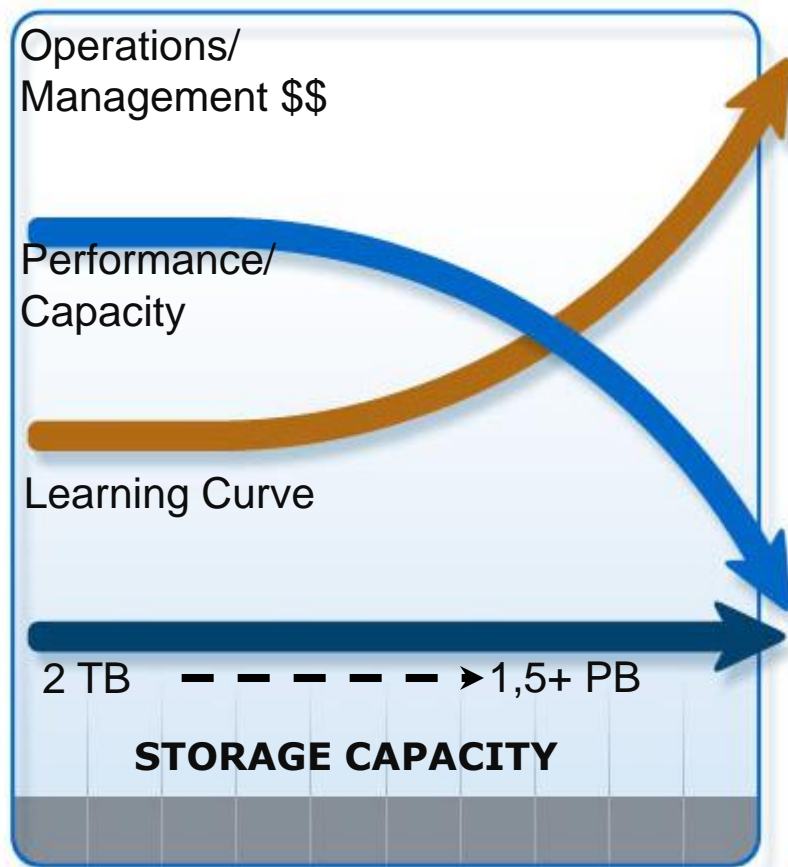
Windows Server® 2008 R2

EqualLogic PS Series Storage

products, Advanced Features, Services & Solutions

Lars Bo Iversen | Storage Specialist

EqualLogic Vision: Storage that Manages Itself



- Non-disruptive growth
 - Add capacity any time
 - Transparent to servers: no downtime
- High performance scaling
 - Linear scaling
- Self-managing
 - System rebalances automatically as it grows
 - System tunes automatically in response to workload
- Integrated, all-inclusive feature set

Equallogic PS Series

Simplify your storage

PS6000 Series

*Powerful.
Virtual.
Enterprise.*



For Large Organizations &
Data Center Deployments

PS4000 Series

*Simple.
Capable.
Affordable.*



For remote/branch
offices & SMB

Equallogic Peer Storage

Drive Types

PSxx00**S**
SSD
50GB, 100GB

PSxx00**XV**
SAS 15K RPM
300GB, 450GB, 600GB

PSxx00**X**
SAS 10K RPM
300GB, 450GB, 600GB

PSxx00**E**
SATA 7.2K RPM
250GB, 500GB, 1TB

PS6000

16 drives
4x1GB or 2x10GB ports
16 arrays in a Group



PS6500

48 drives
4x1GB or 2x10GB ports
16 arrays in a Group



PS4000

16 drives / 2 GB ports
2 arrays in a Group



Benefits of Dell EqualLogic in a 10GbE environment

Investment Protection

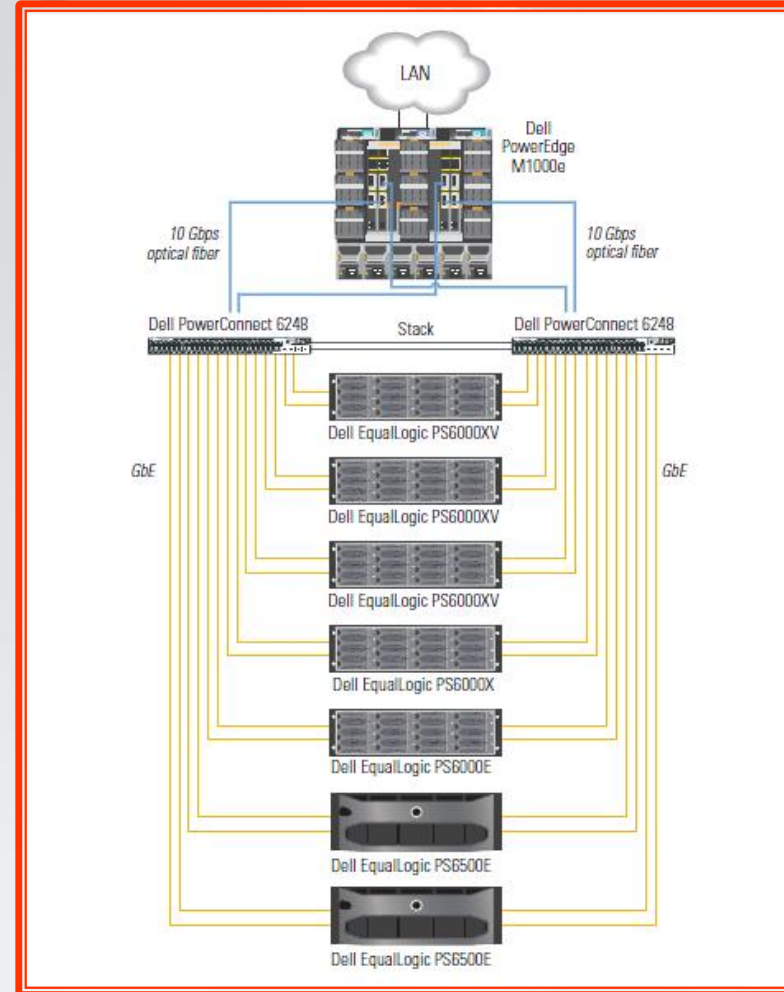
- Utilize Existing 1GbE models with future 10GbE PS Series

Simplify & Scale

- Easily introduce new networking into your virtual data center

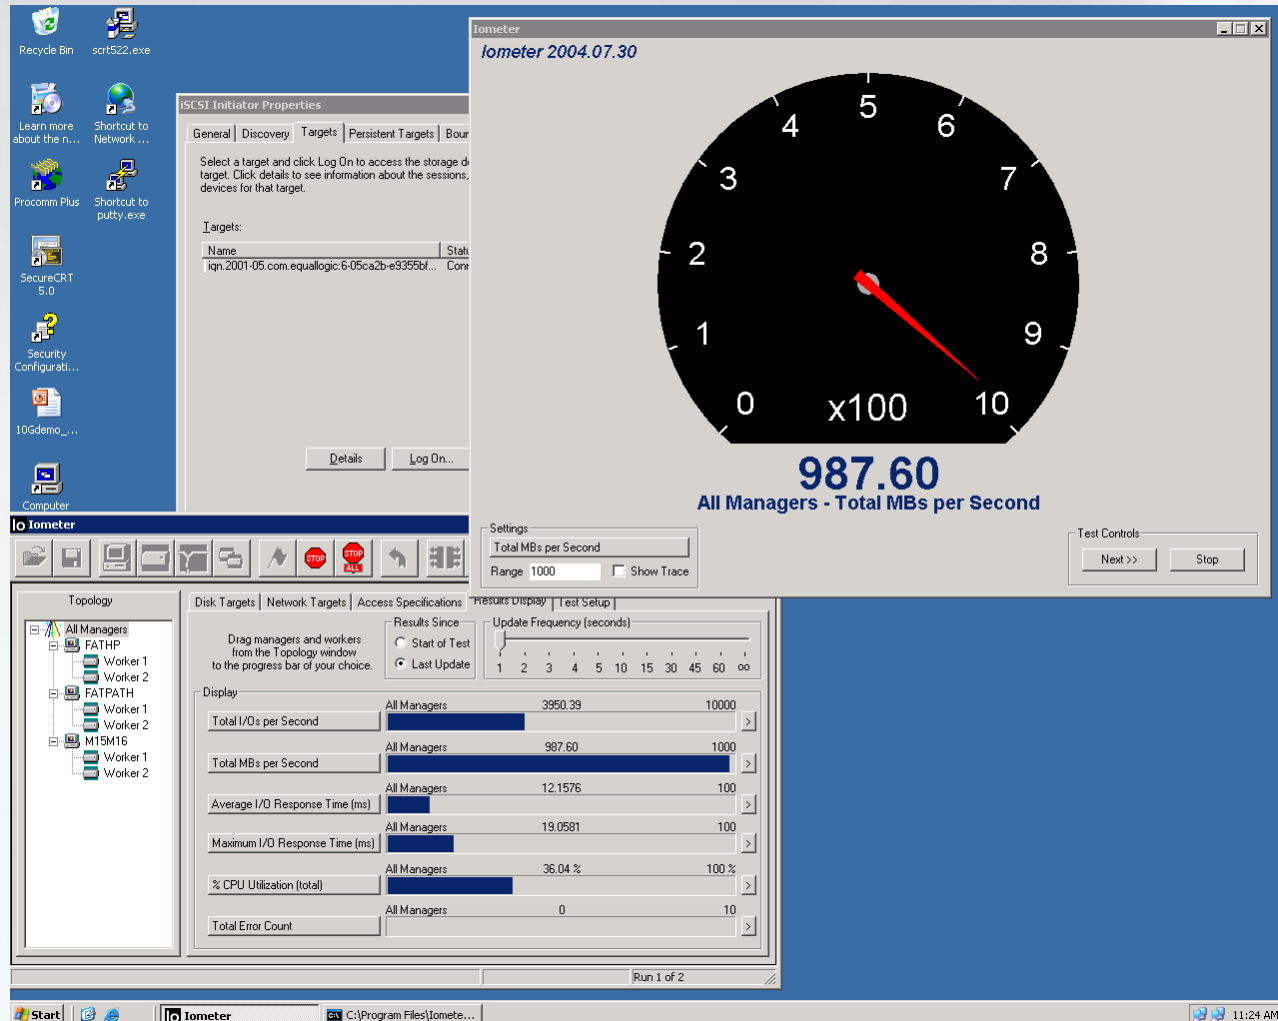
Enhanced Consolidation ROI

- Reduced OPEX with simplified cabling, unified network management and standards based technologies



10Gb Controller

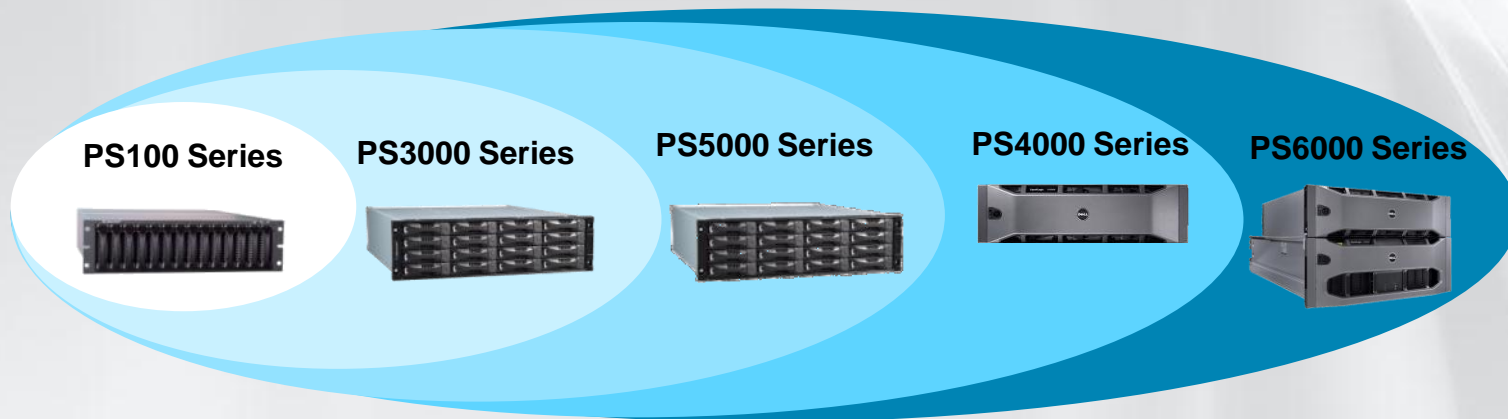
Sequential workload



SEAMLESS SAN EXPANSION

CONTINUOUS SOFTWARE ADVANCEMENTS

Single Multi-Generational SAN



Continuous Advancements of All-inclusive Firmware & Software

-
- Core Data Protection
 - Auto-Snapshot Manager/Microsoft® Edition (ASM/ME)
 - Thin Provisioning
 - VMware® Site Recovery Manager
 - Auto-Snapshot Manager/ VMware® Edition (ASM/VE)
 - SAN Headquarters
 - ASM/ME for Hyper-V™
 - VMware® vStorage

2004



Windows Server® 2008 R2

PS Series Comprehensive Data Management

All-Inclusive with **NO** Additional Costs

BASE SOFTWARE FEATURES

Management

- ✓ Instant on 'set-up' manager
- ✓ Group manager
- ✓ Rapid provisioning
- ✓ Roles-based management

Data Protection and Availability

- ✓ RAID 5, 6, 10, and 50
- ✓ Automatic RAID placement
- ✓ Multi-path / IO support

Maintenance

- ✓ Phone home
- ✓ Enclosure monitoring system
- ✓ Performance monitoring

ADVANCED SOFTWARE FEATURES

Storage Virtualization

- ✓ Complete SAN virtualization
- ✓ Thin provisioning
- ✓ Auto-load balancing
- ✓ Automatic storage pools & tiering
- ✓ Array evacuation

Data Protection and Recovery

- ✓ Writeable snapshots
- ✓ Multi-volume snapshots
- ✓ Instant restore / cloning
- ✓ Multi-way replication for disaster recovery

Multi-Group Monitoring

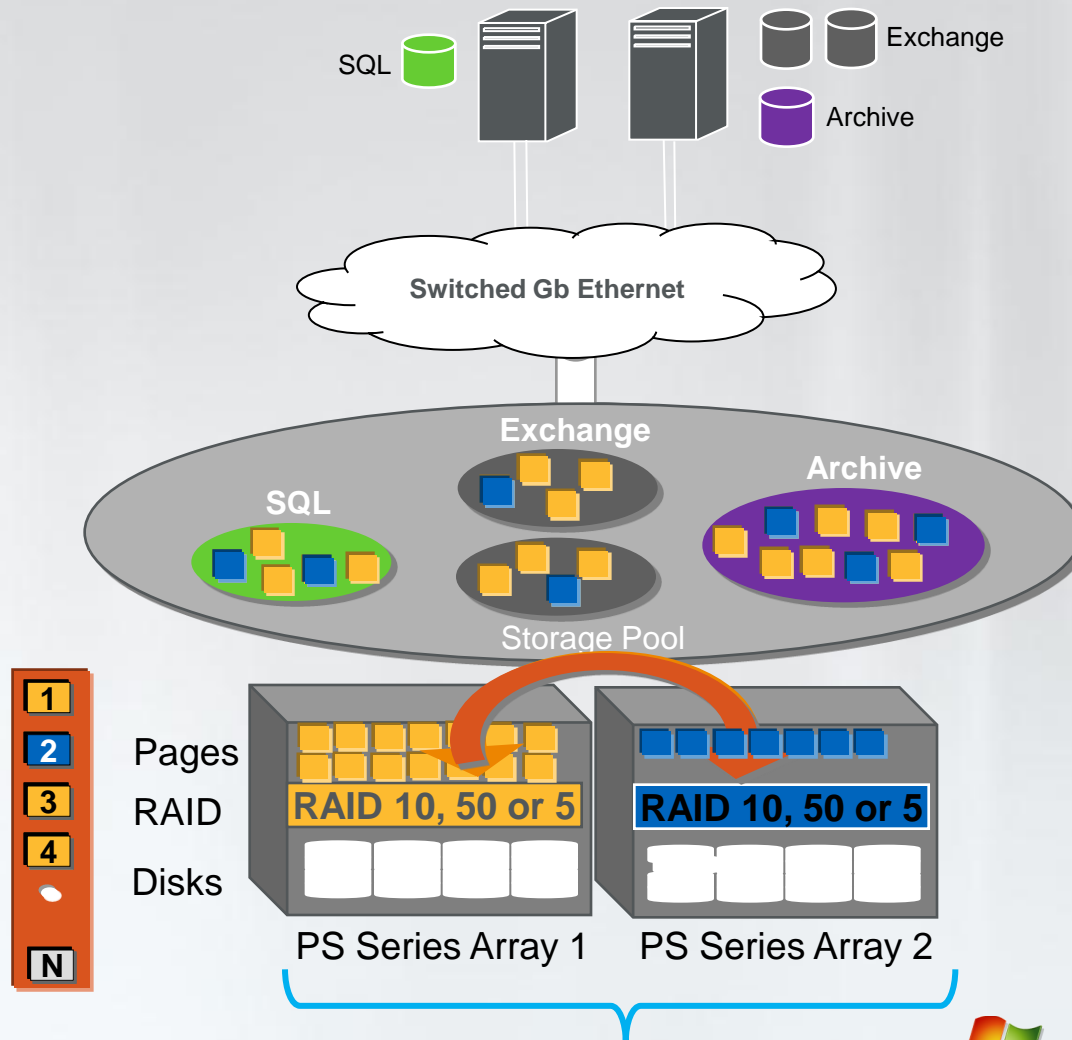
- ✓ SAN HeadQuarters (SAN HQ) event & performance monitoring

Server Management Integration

- ✓ VSS and VDS Providers
- ✓ Automatic MPIO Connection Management
- ✓ Auto-Snapshot Manager/ Microsoft® Edition
 - Hyper-V™, Exchange®, SQL Server and Windows® file systems data, now including MSCS
- ✓ Auto-Snapshot Manager/
 - Hypervisor-aware SAN-based snapshots, clones and replication for rapid recovery
 - Selective VM restore
 - Advanced MPIO integration

PS Series Architecture

Virtualized peer storage



Single-View Management

Single SAN Manager – Any Size, Any Form

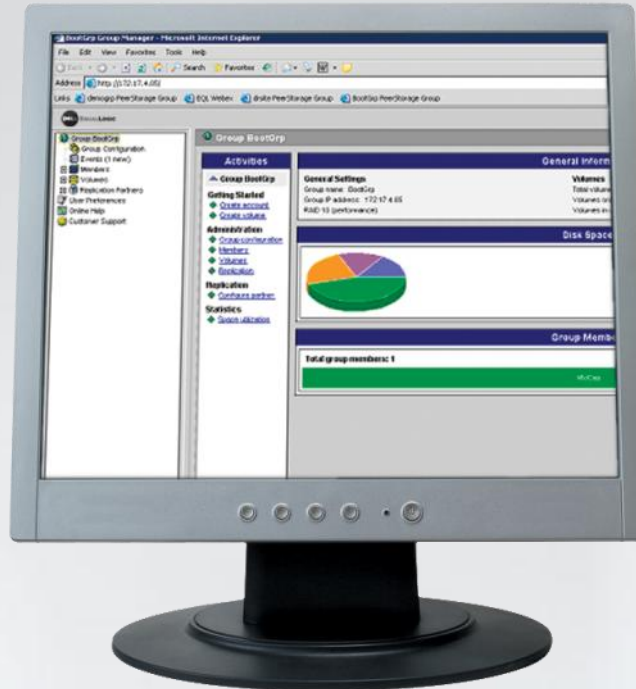
Group Manager — Simple Intuitive GUI

SAN Virtualizer

System Load
Balancer

Volume Manager

Enclosure Monitoring
System — in-depth
SAN reporting



Storage Configuration
Assistant — best
practice hints and clues

Multiple Storage Pool
Creation — online
volume movement
between pools

Data Replication Control

Rapid Provisioning

Reducing time to Storage Deployment

Instant Volume Creation

1. Name
2. Size
3. Snapshot size
4. Security

Online Volume Modification

- Clones
- Snapshot
- Replication
- Volume expansion



Windows Server[®] 2008 R2

EqualLogic SAN HeadQuarters

Rich historical Reporting

Capacity

- Overall Capacity
- Thin Provisioned space

I/O Performance

- IOPS
- I/O in kb/sec
- Latency

Network Data

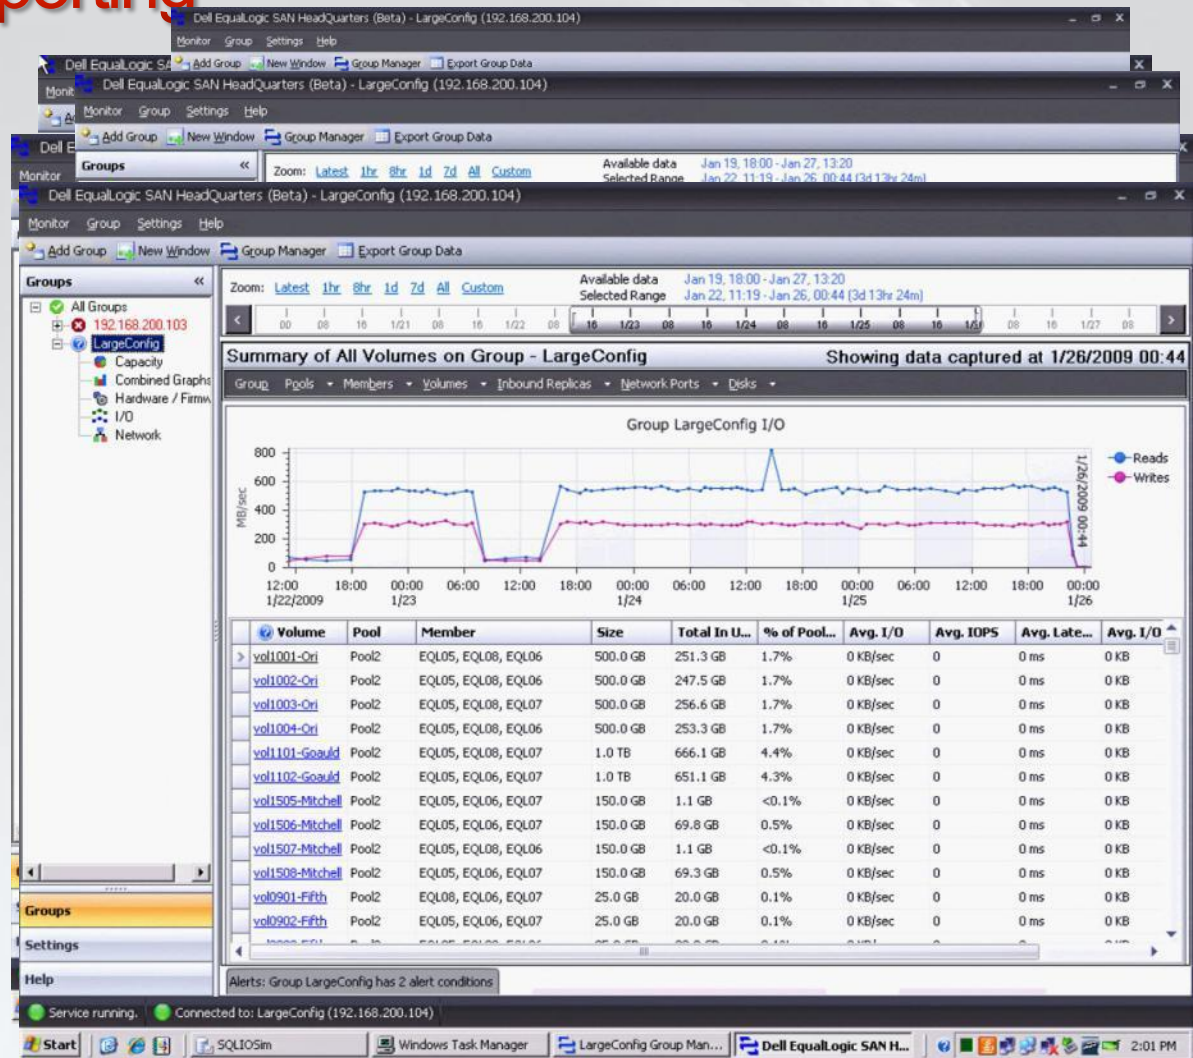
- Network Link Throughput
- Active ports/Slowest port

Member Hardware and Configuration

- Pool
- Status/RAID Policy

Volumes Data

- Pool
- Binding



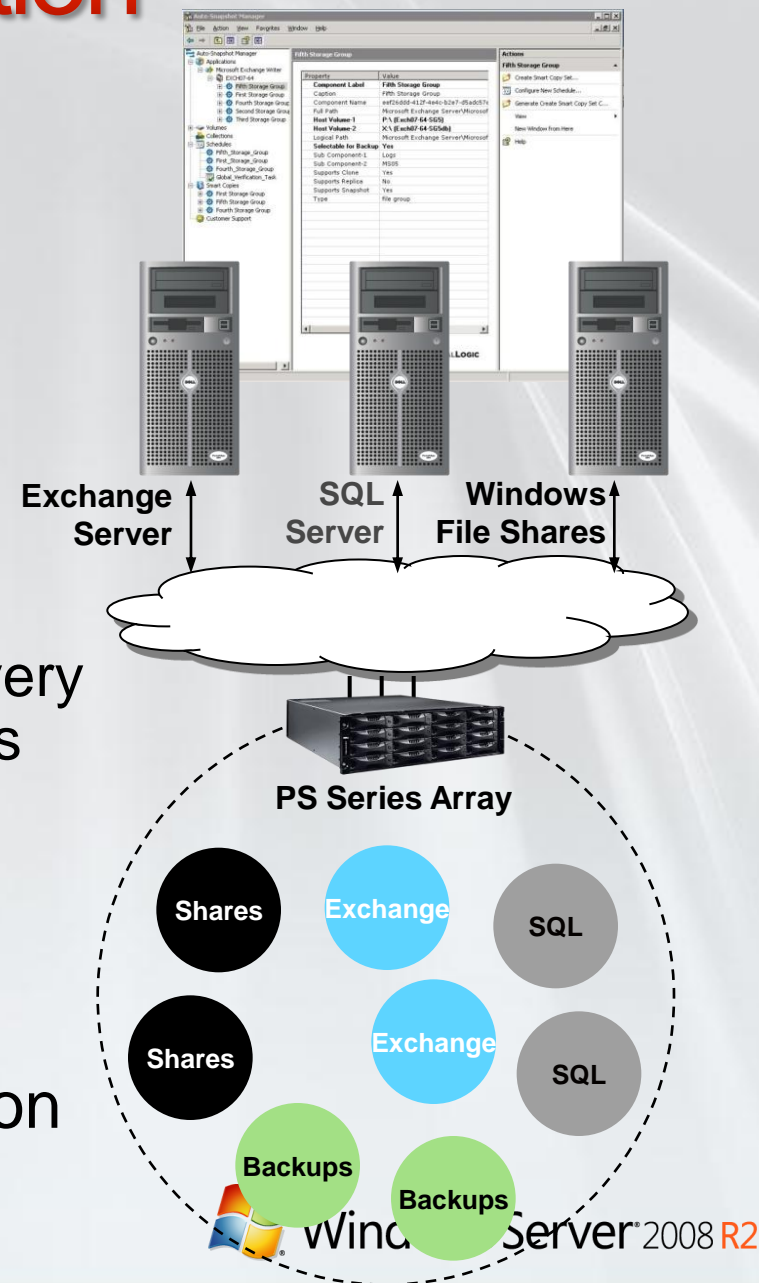
Application Data Protection Auto-Snapshot Manager Microsoft Edition

Auto-Snapshot Manager

Integrated with Microsoft® Volume ShadowCopy Service

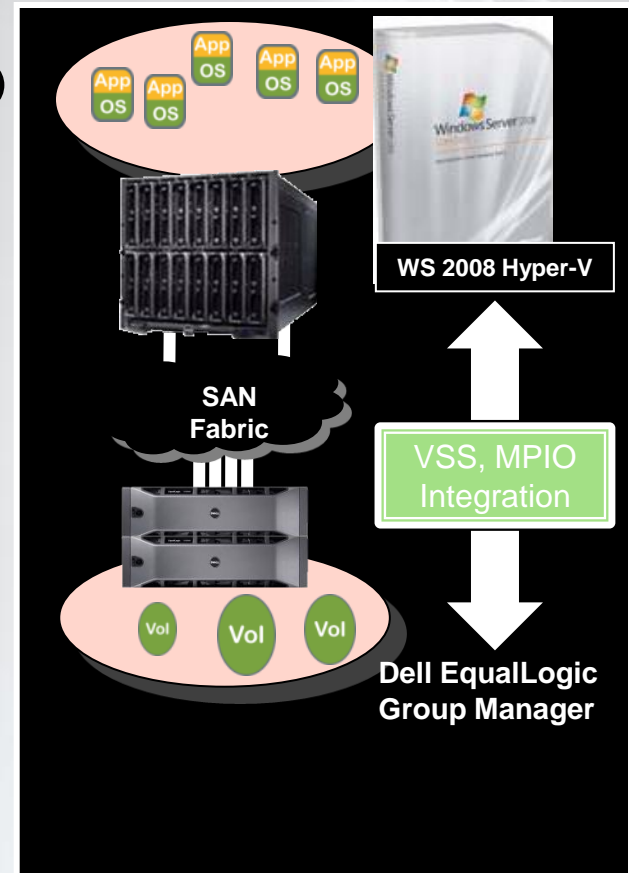
- Application aware protection
 - SQL, Exchange®
 - Hyper-V™
- Delivers powerful and flexible
 - Data protection and rapid recovery
 - Test/Development environments
 - Selective restore of application objects
- Available at no extra charge

Simplify IT: Spend less time managing and protecting application data

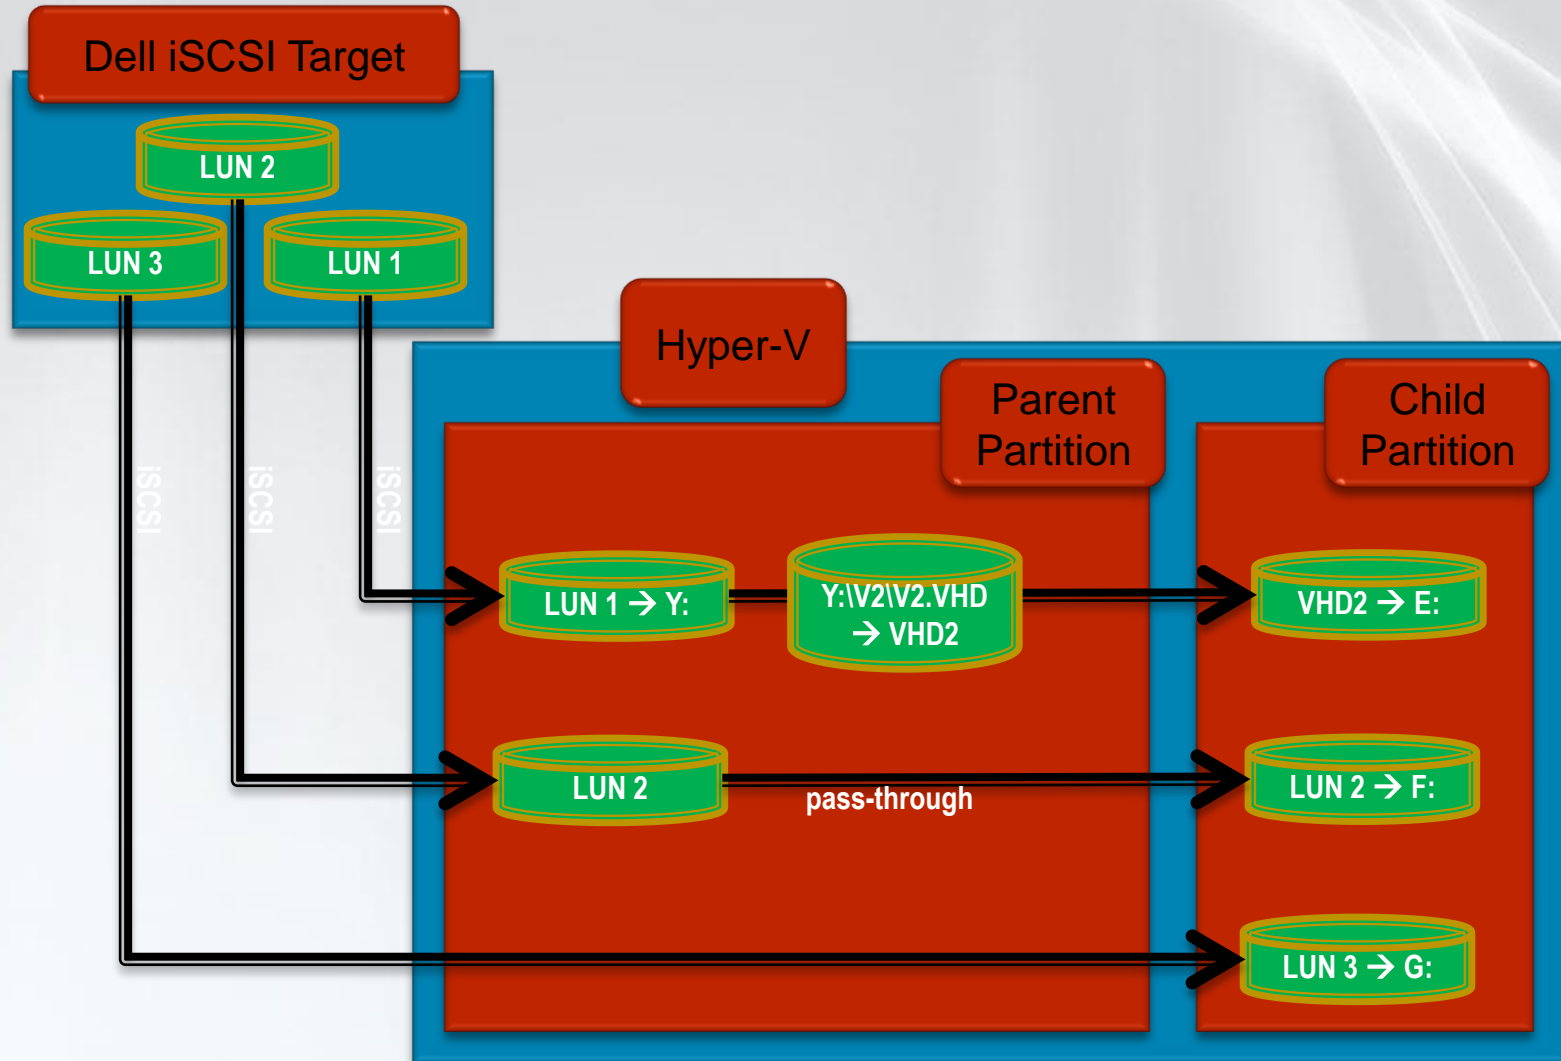


Introducing ASM/ME Smart Copy for Microsoft Hyper-V

- Newest Addition to Auto Snapshot Manager/Microsoft® Edition
- Tightly integrated with Volume Shadow Copy Services (VSS)
 - Uses Hyper-V™ VSS protection and recovery capabilities
- Create and manage Smart Copy Snapshots for point-in-time copies of Hyper-V™ Virtual Machines
- Rapid restoration of Hyper-V™ Virtual Machines



iSCSI directly to child



2.5. Comparing the three options

Using the Microsoft iSCSI Software Target with Hyper-V

Feature	iSCSI to Parent, VHD file	iSCSI to Parent, Pass-through	iSCSI directly to Child
Easier to manage from parent	X		
Easier to manage from child			X
Expanding/Differencing VHDs	X		
Hot add disk with Hyper-V V1			X
Hot add disk with Hyper-V V2	X*	X*	X
Hyper-V VHD snapshots on Parent	X		
iSCSI Target VSS snapshots on Parent	X	X	
iSCSI Target VSS snapshots on Child			X
Backup with Hyper-V VSS Writer			
Child-based Failover Clustering			X
Child boot from disk (no third party tools)	X**	X**	

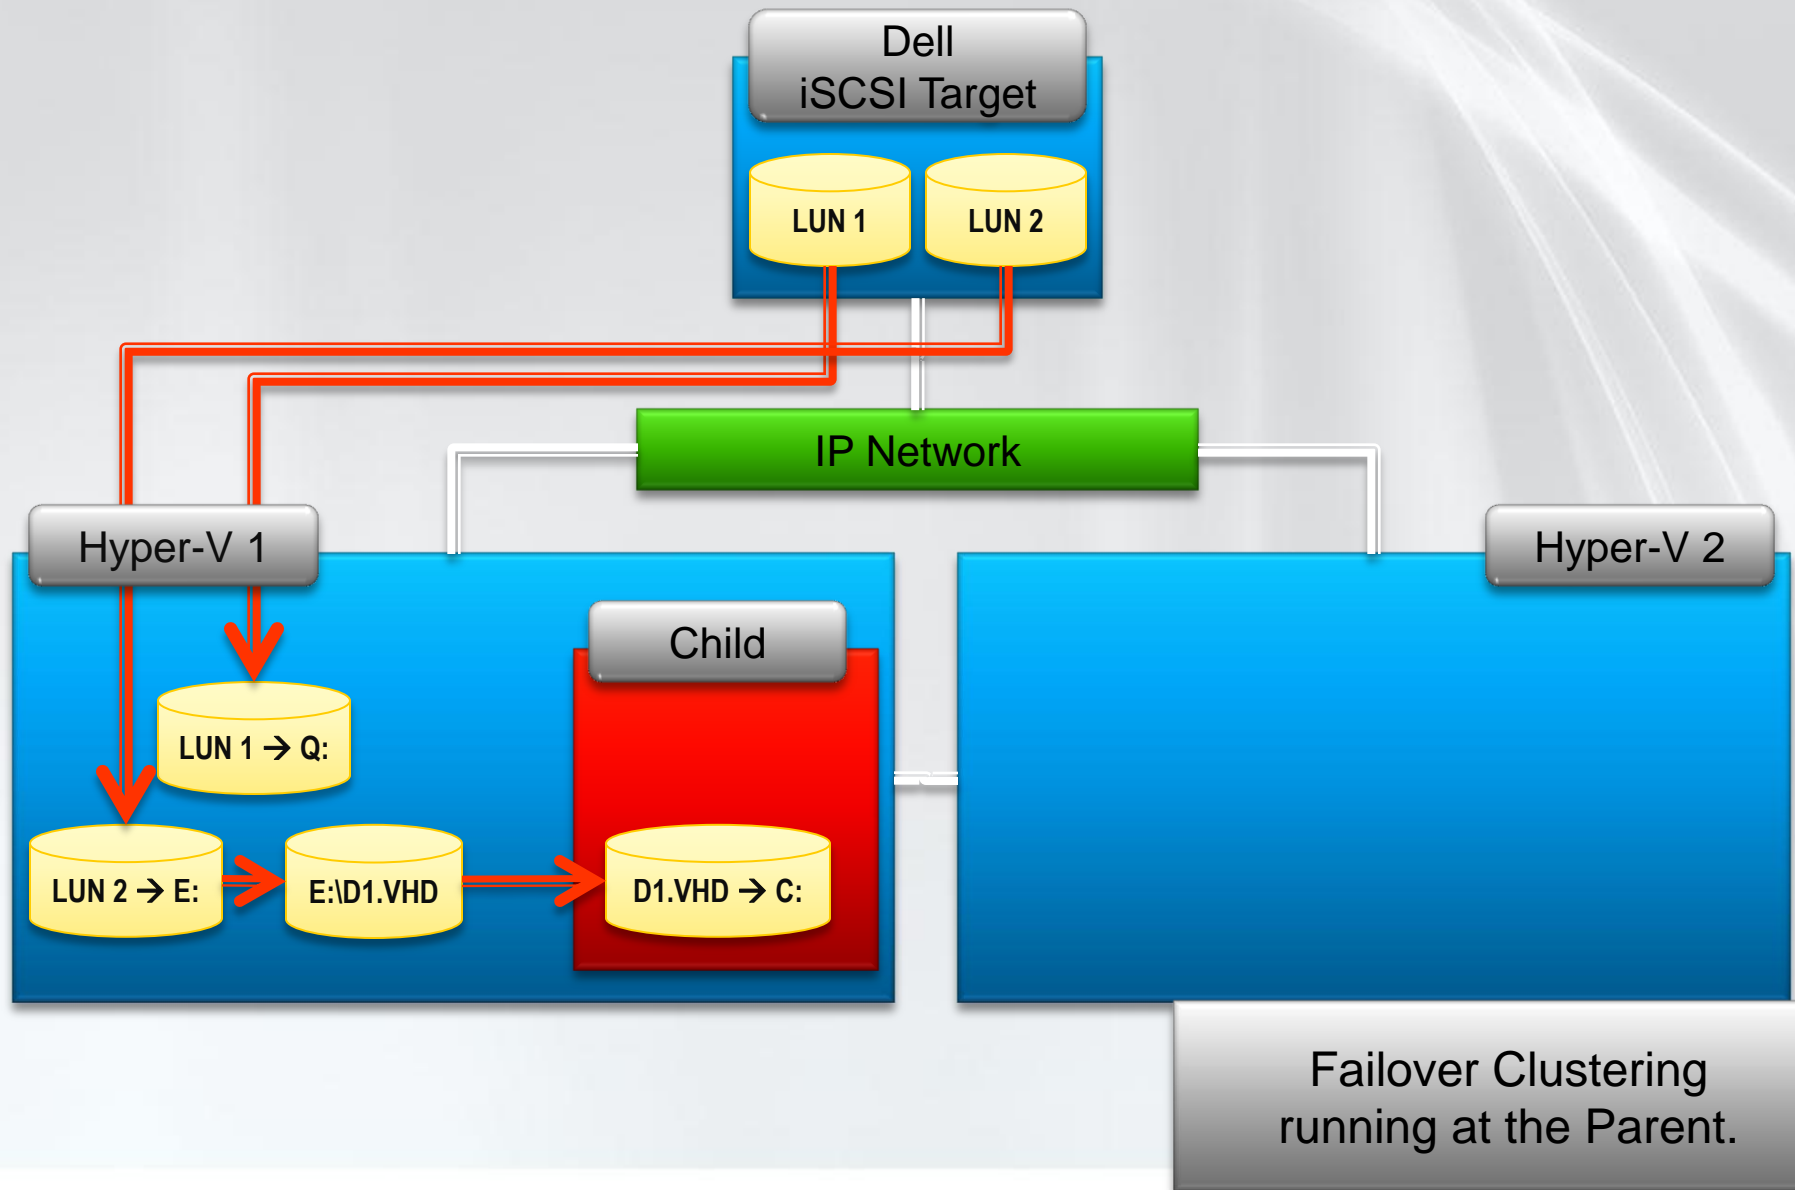


Windows Server® 2008 R2

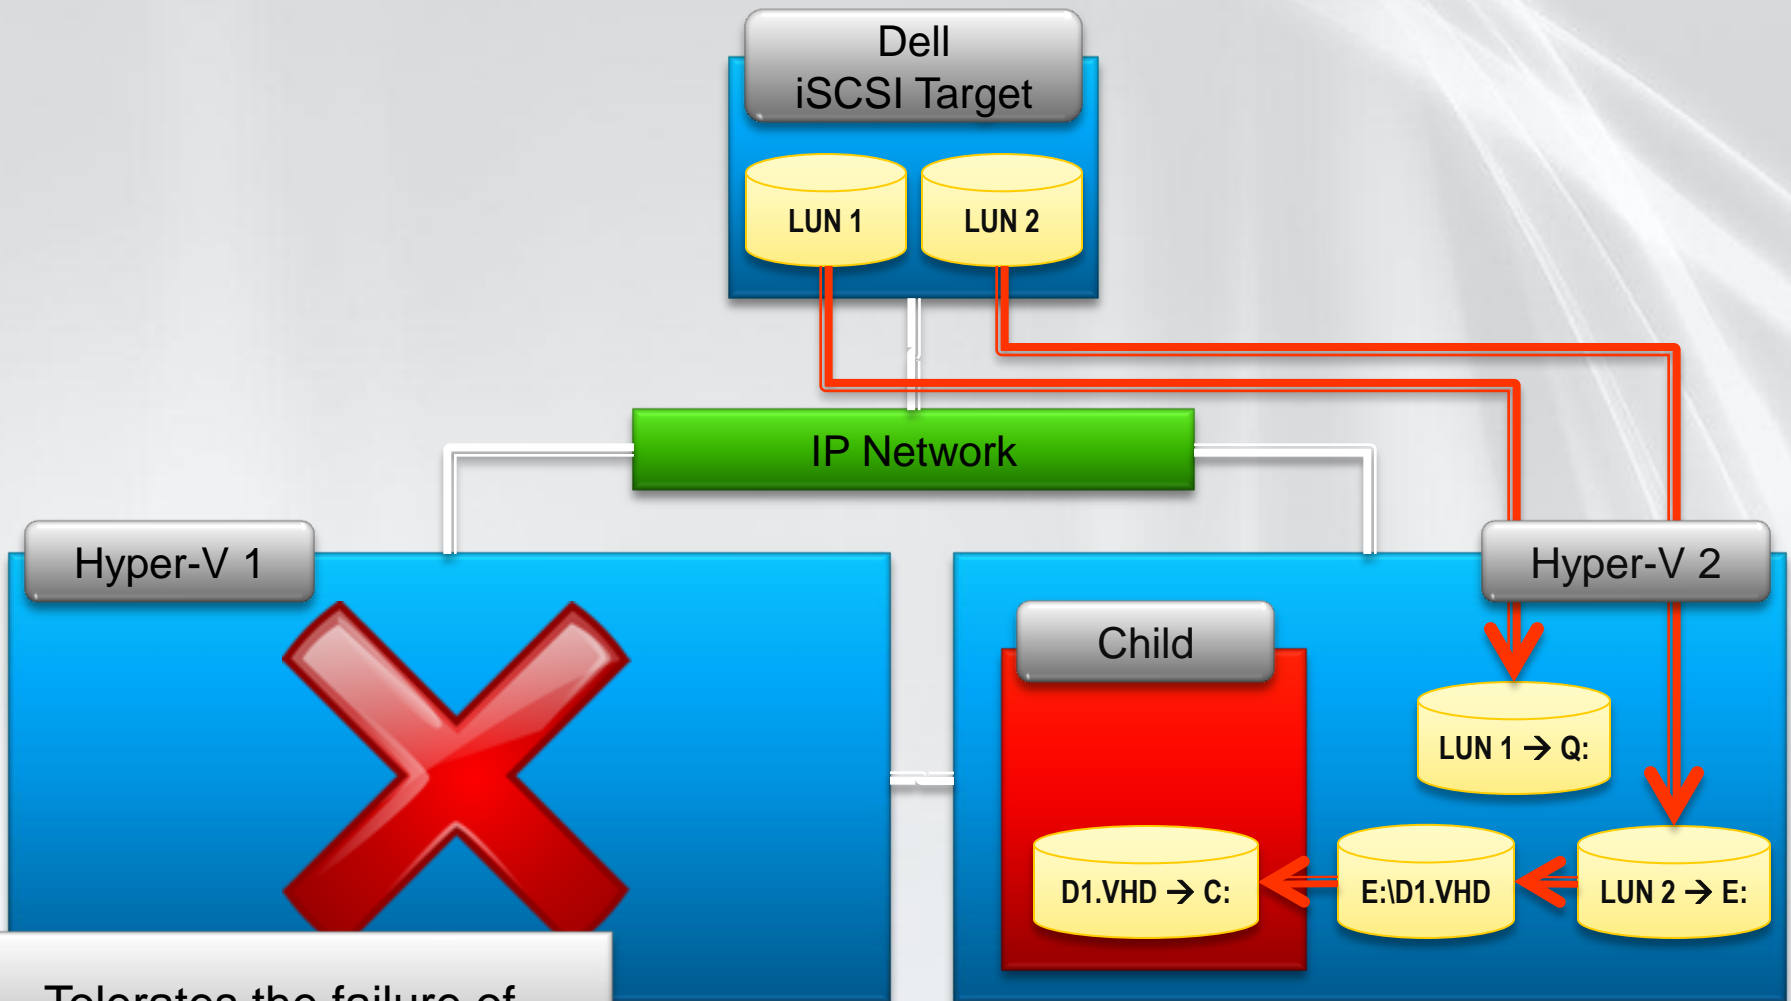
3. Failover Clustering with Hyper-V and iSCSI Target

Understanding how to use the Dell iSCSI Target in conjunction with the Failover Cluster and the Hyper-V role in Windows Server 2008 R2

3.1a. Parent-based, two physical servers (before)

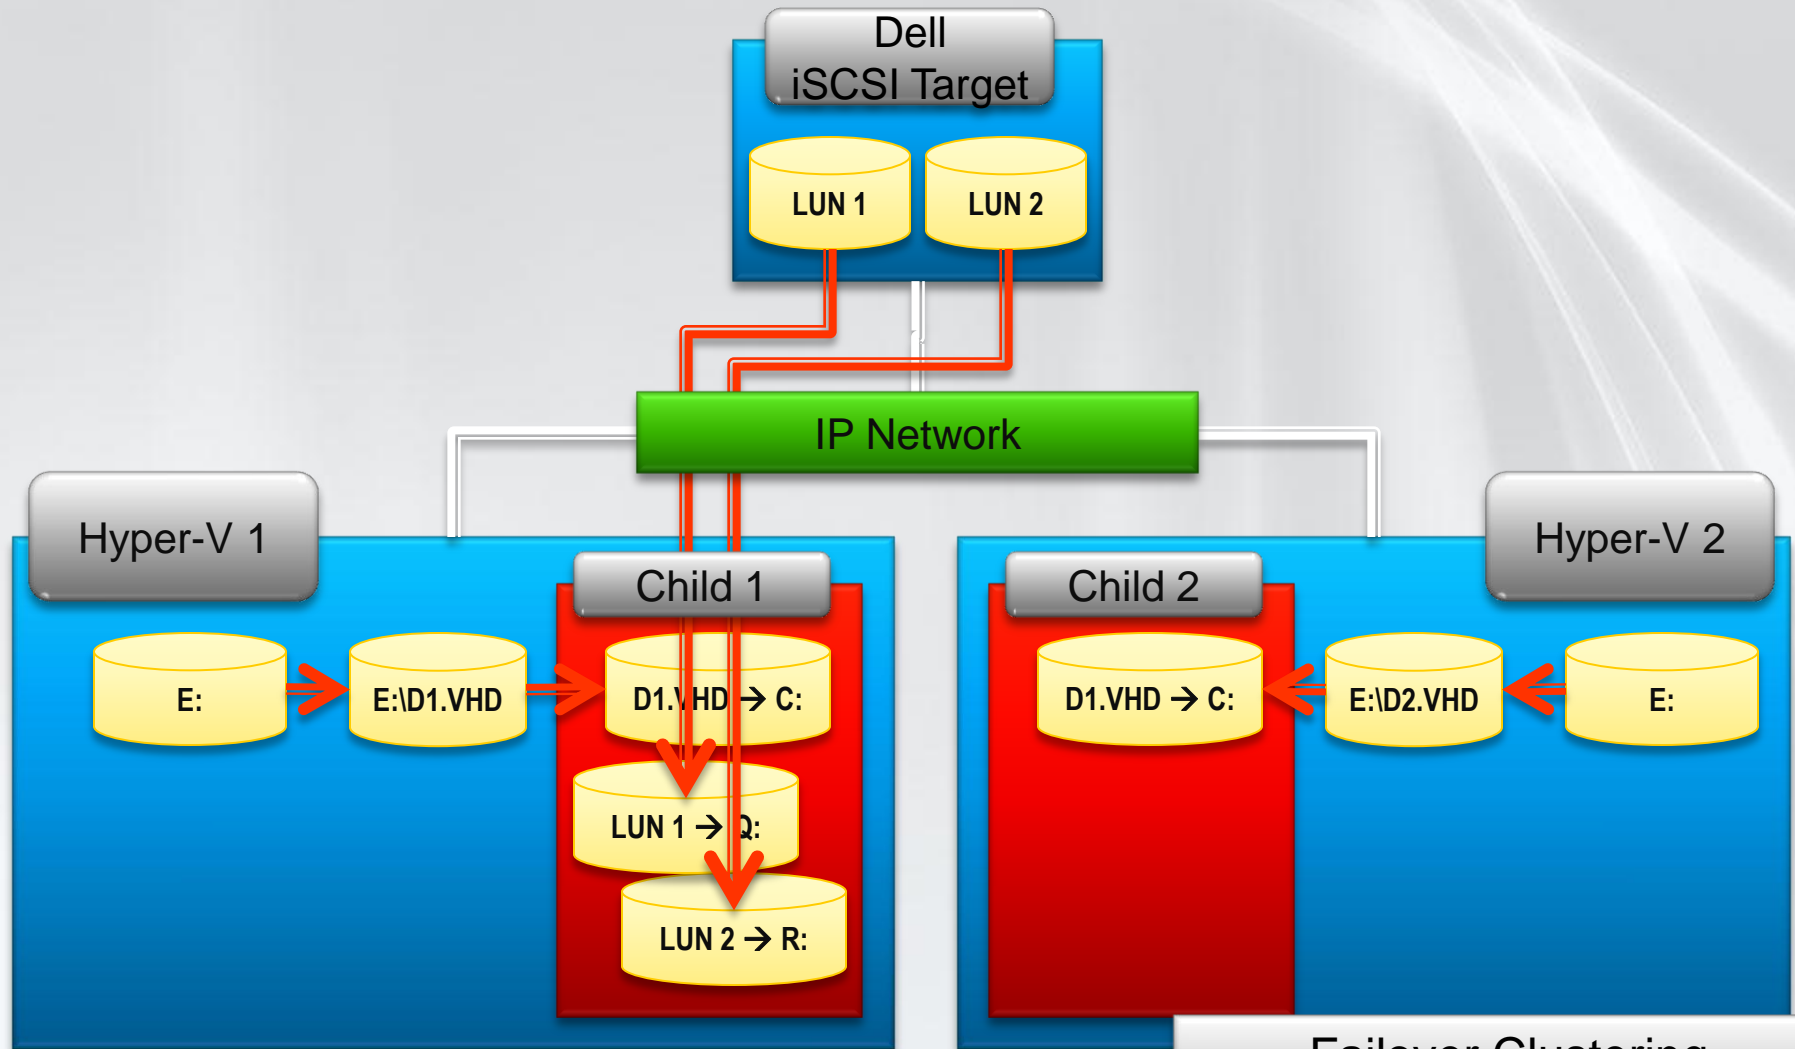


3.1b. Parent-based, two physical servers (after)



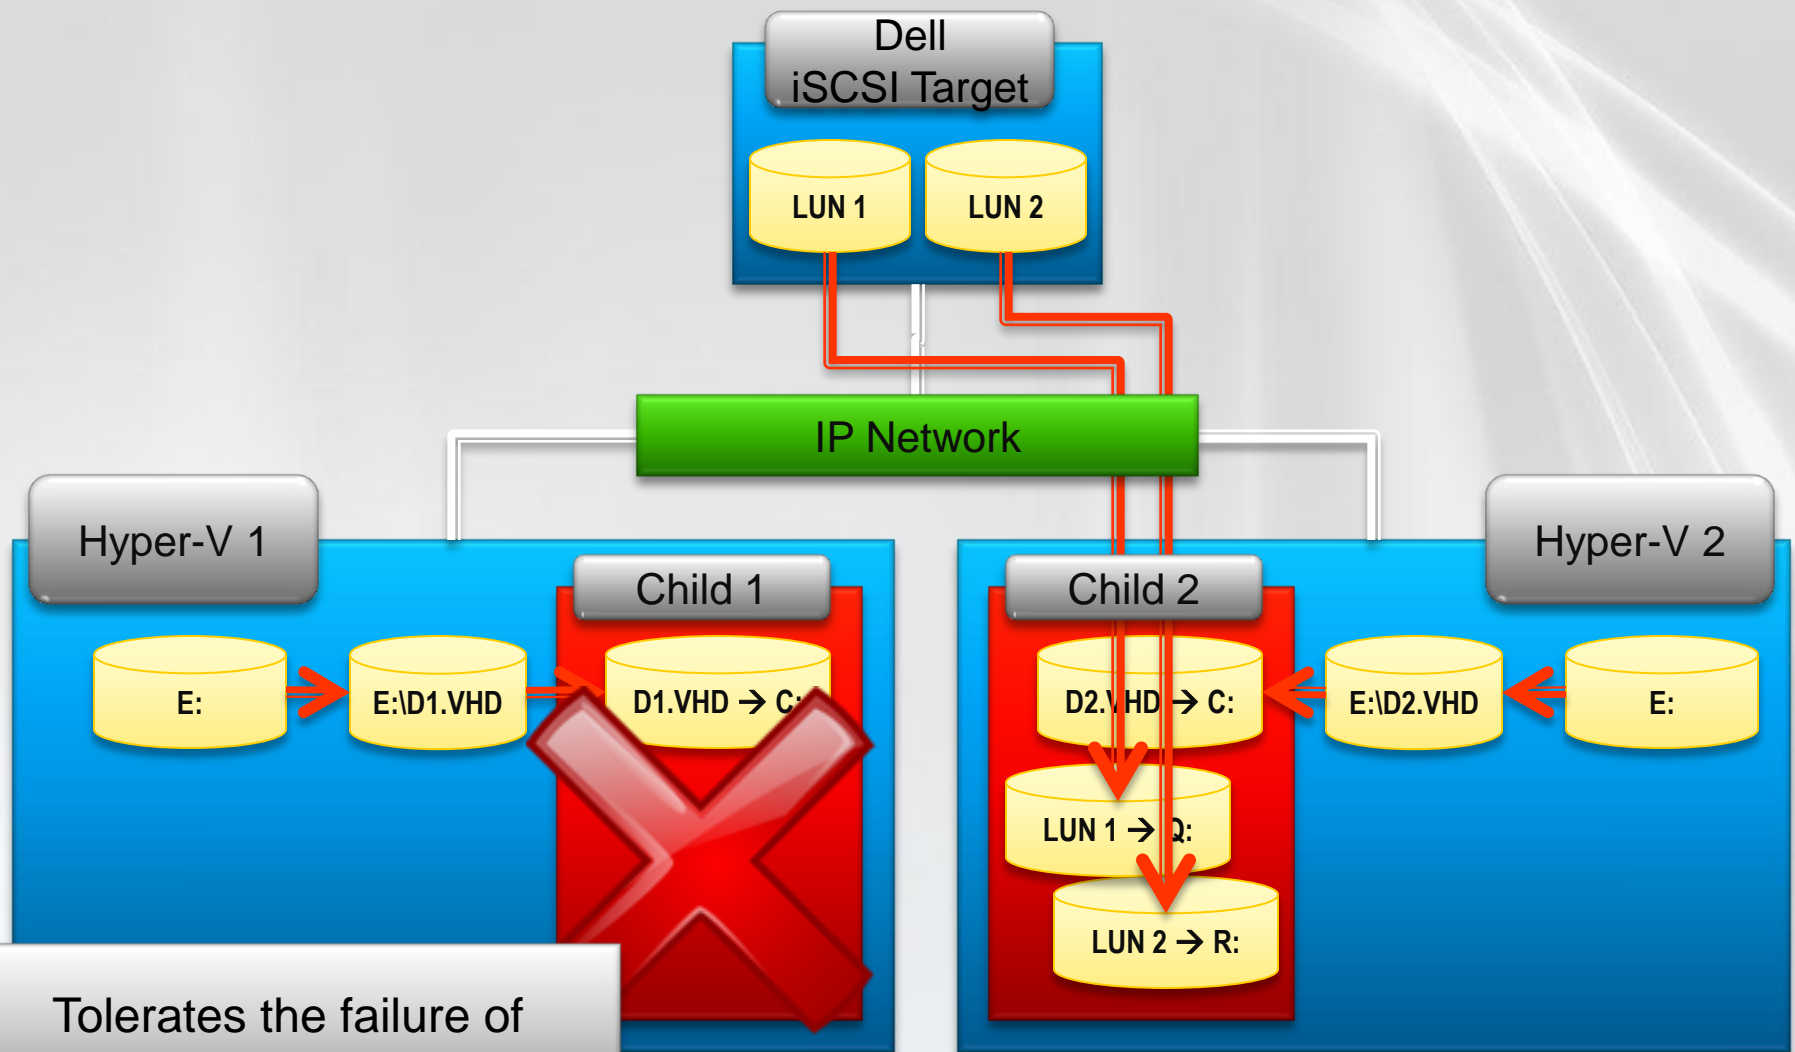
Tolerates the failure of one of the Hyper-V hosts (physical computer)

3.2a. Child-based, two physical servers (before)



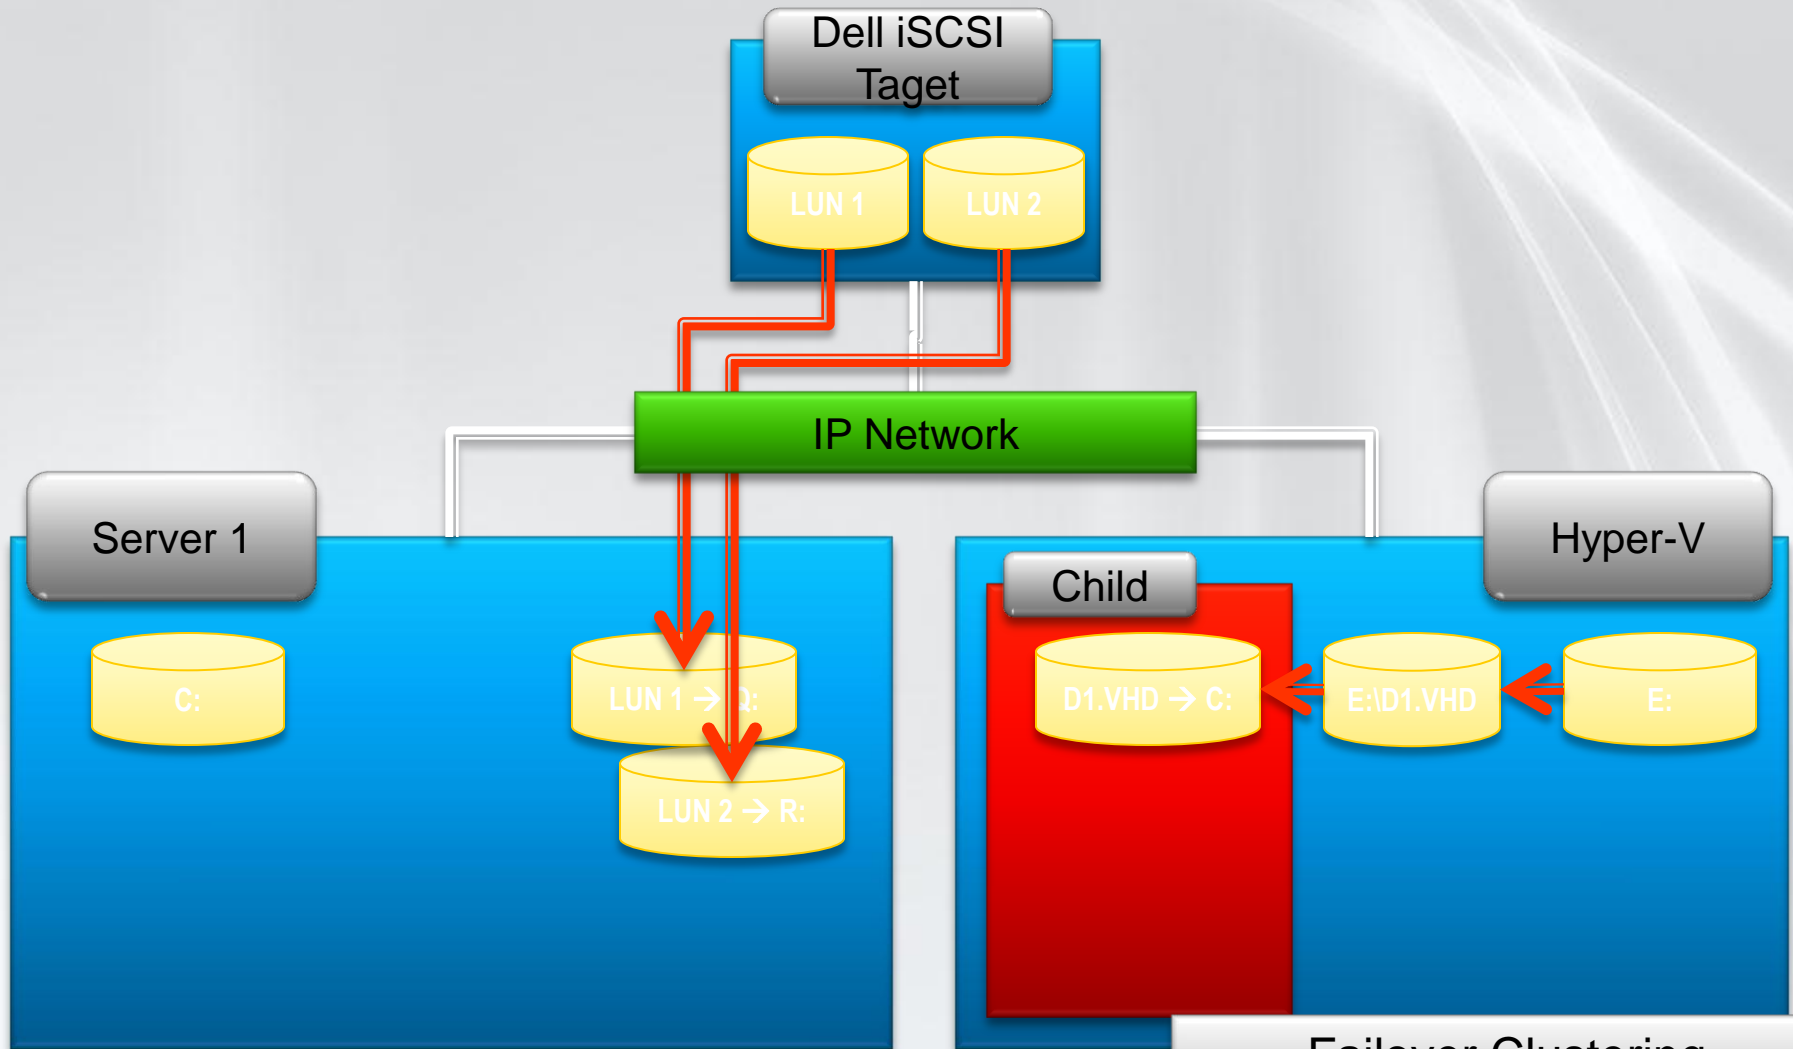
Failover Clustering
running at the Child.
Requires iSCSI.

Dell iSCSI Target



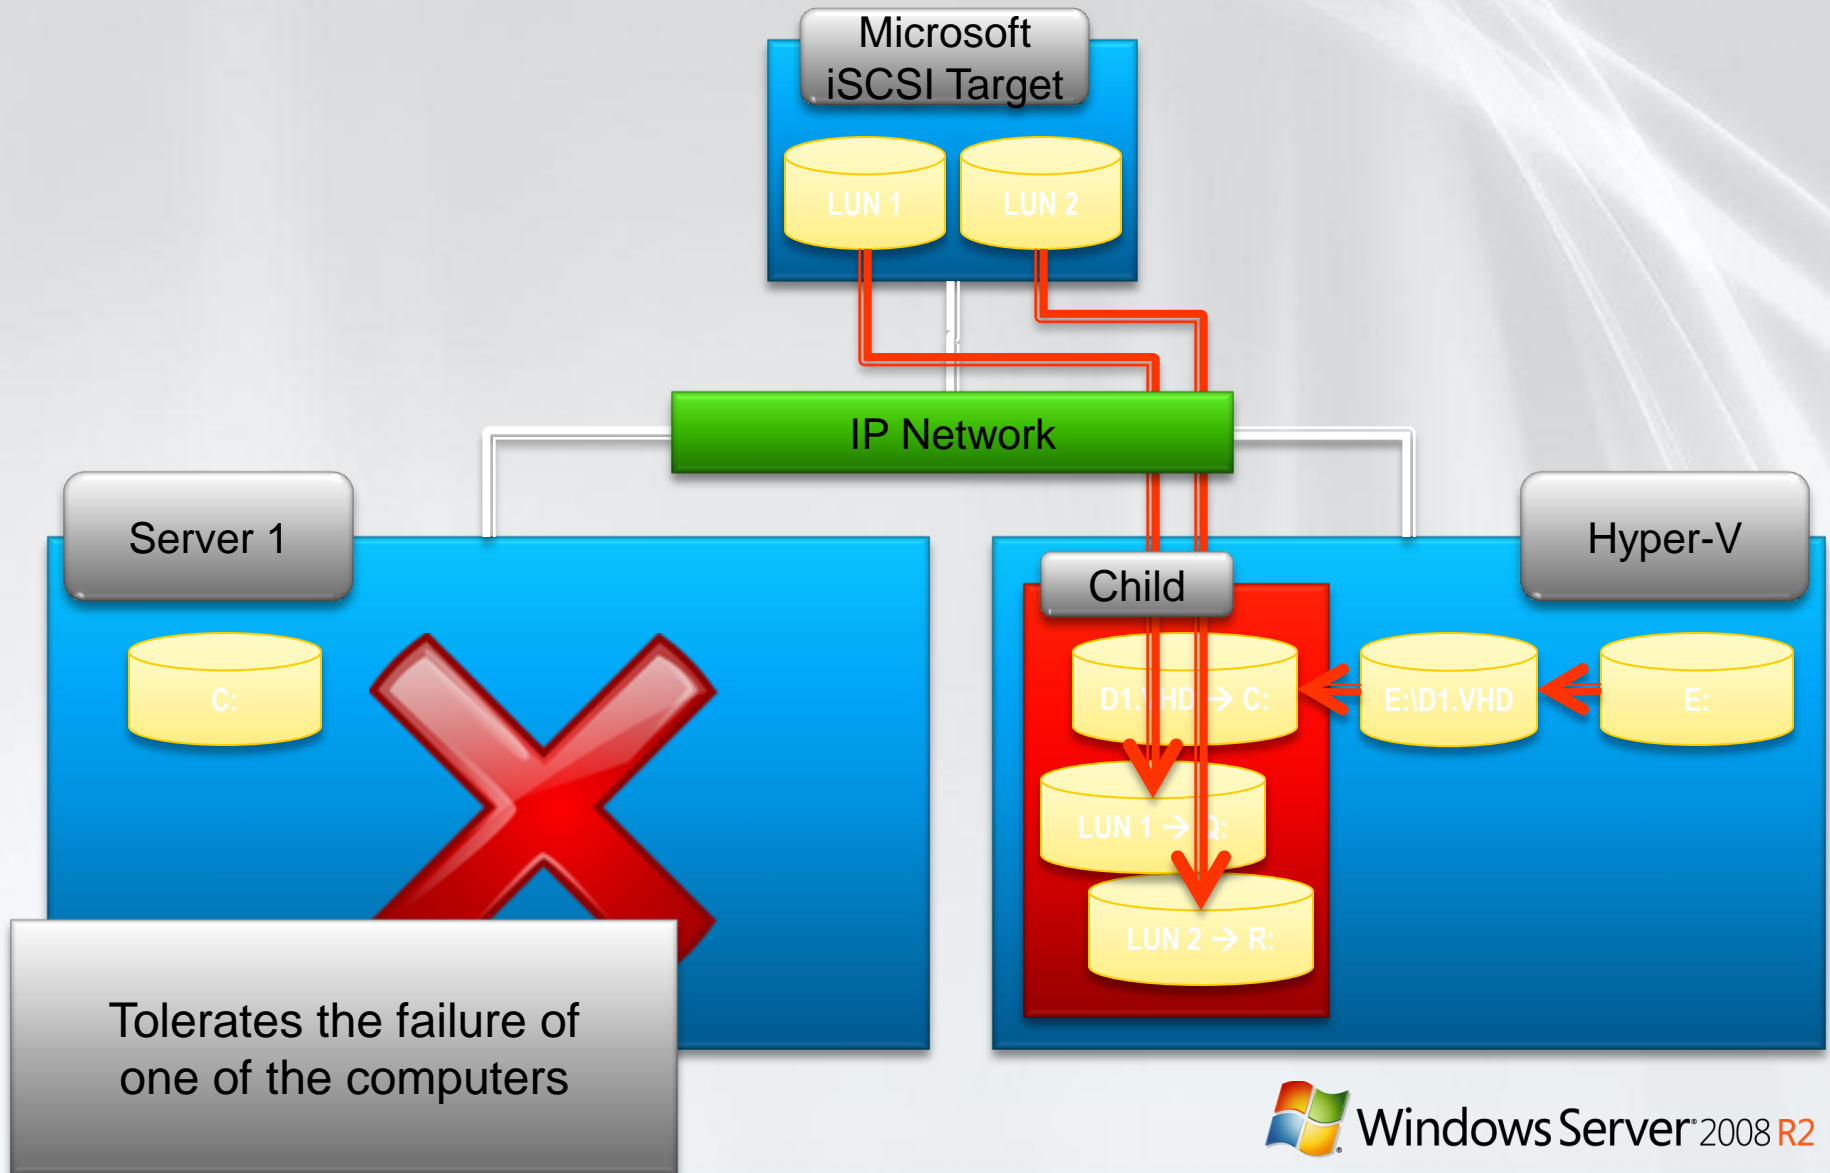
Tolerates the failure of one of the Hyper-V hosts (physical computer)

3.3a. Mixed Physical/Virtual (before)

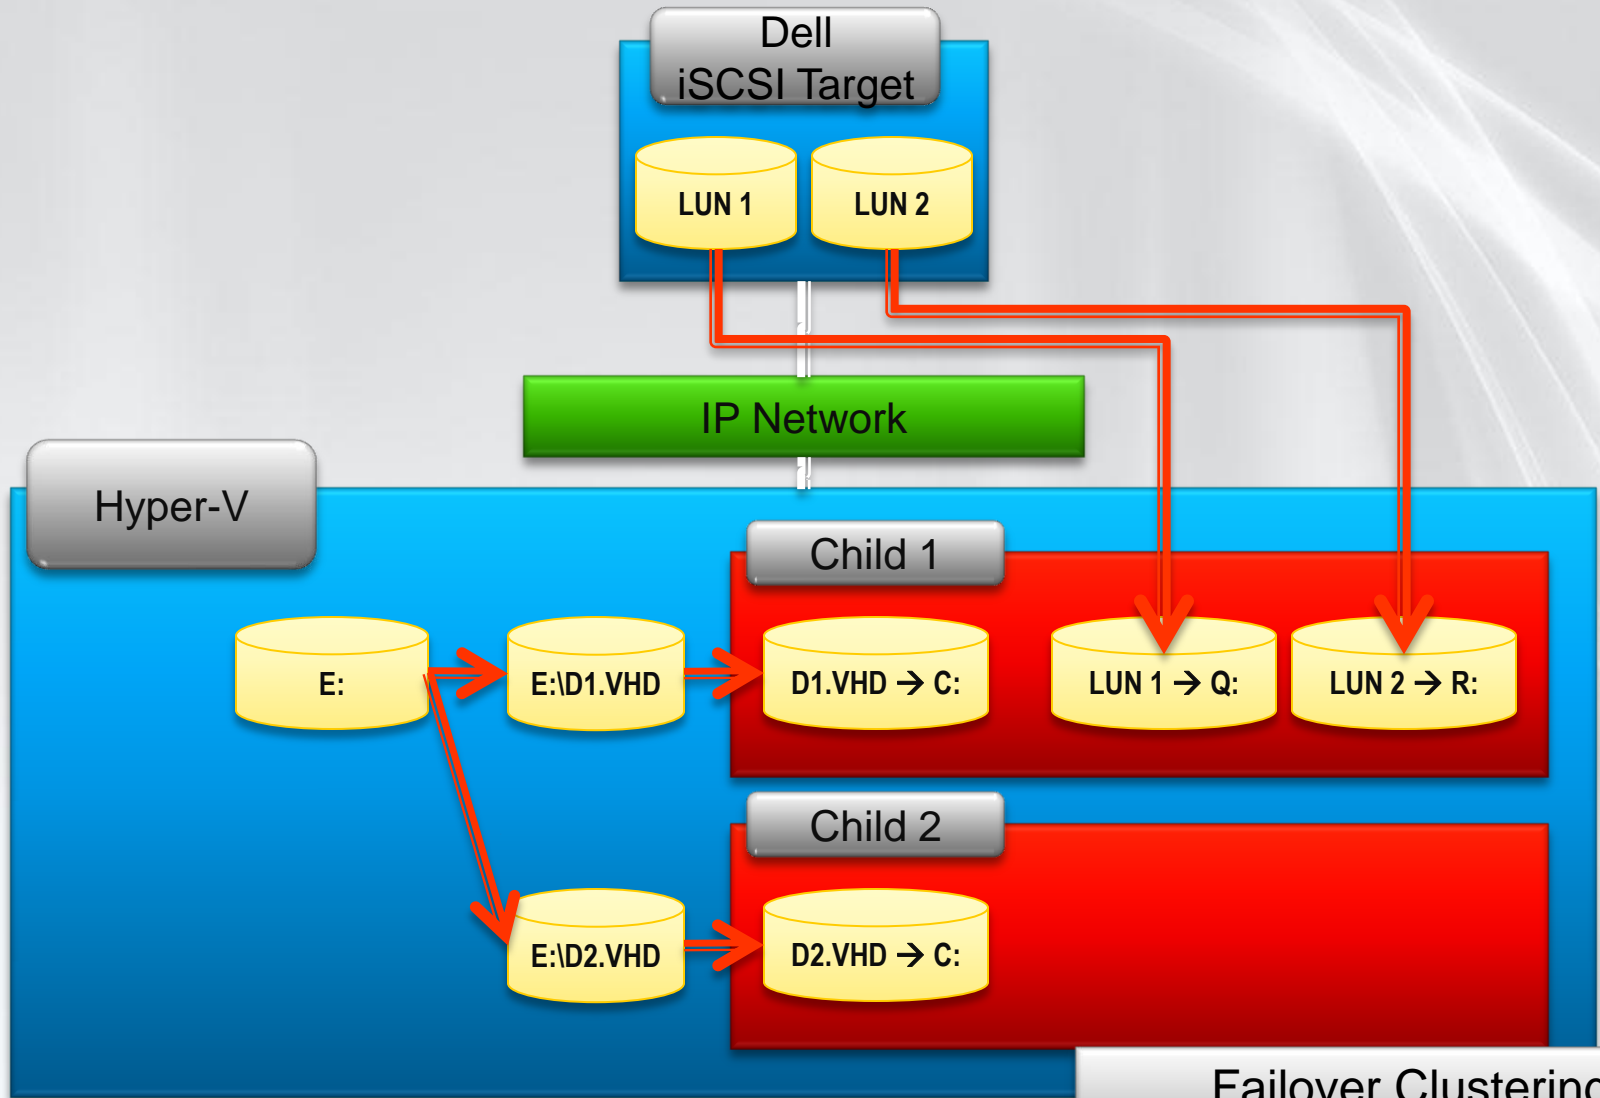


Failover Clustering
running at the Child.
Requires iSCSI.

3.3b. Mixed Physical/Virtual (after)

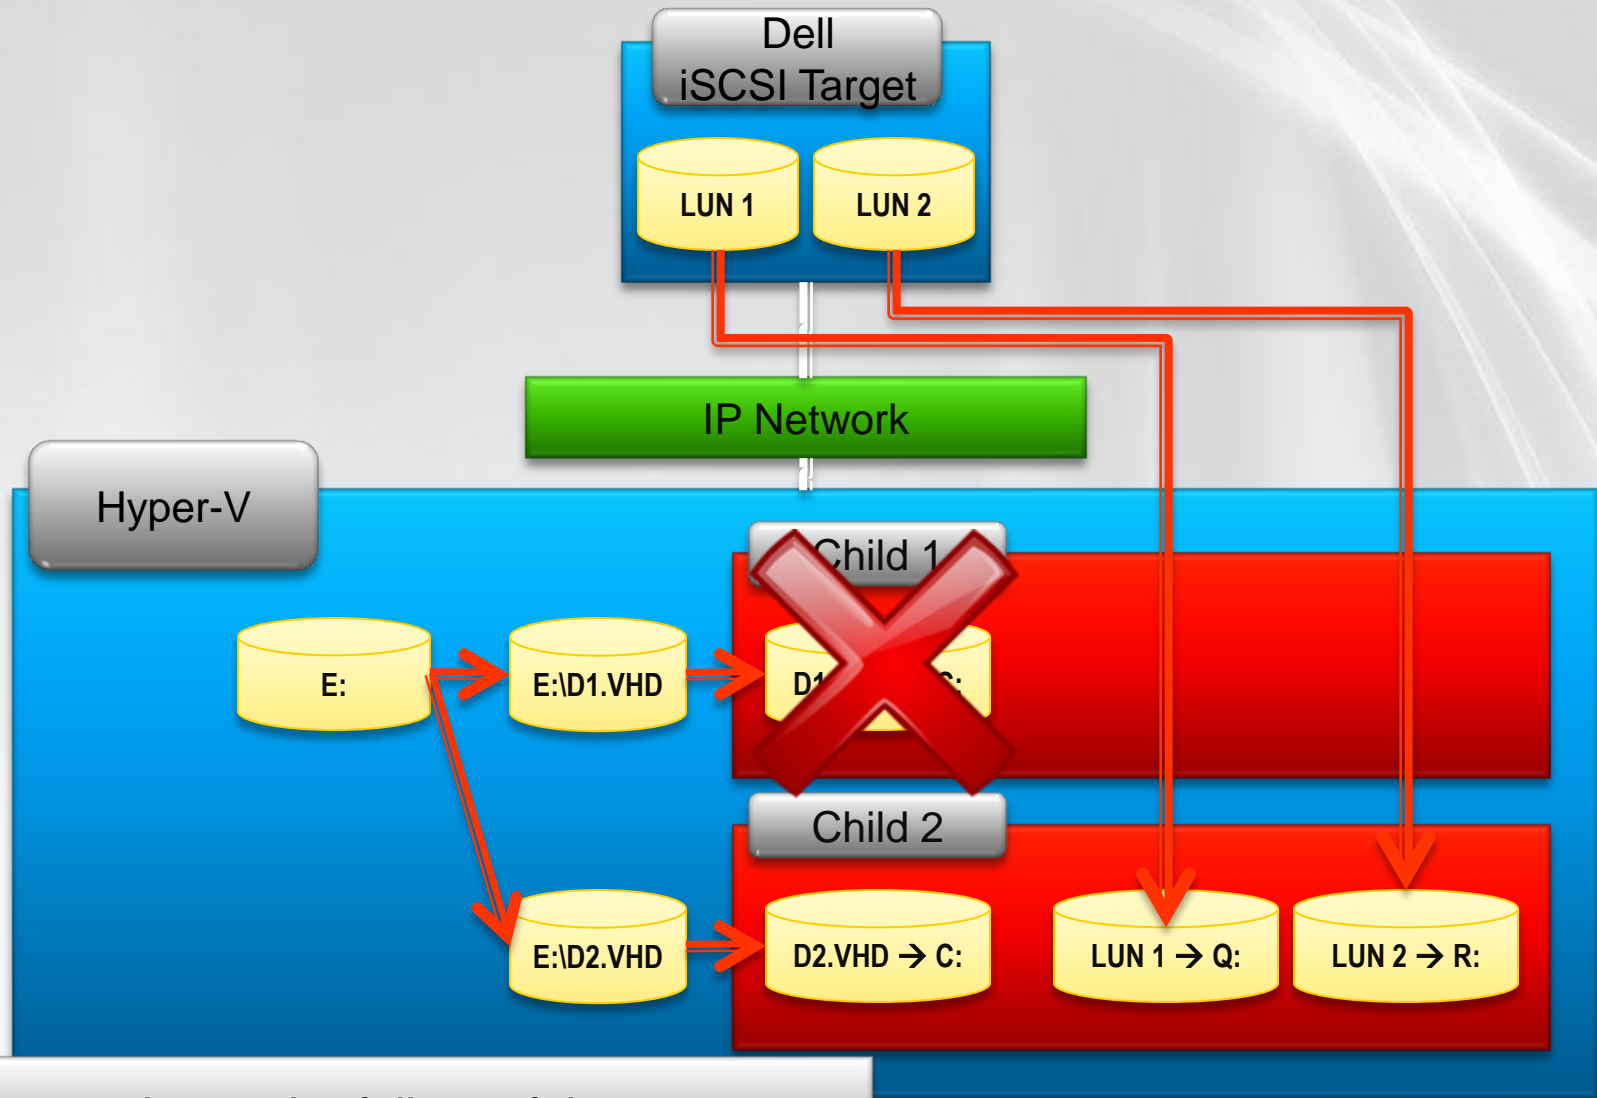


3.4a. Child-based, one physical server (before)



Failover Clustering
running at the Child.
Requires iSCSI.

3.4b. Child-based, one physical server (before)



Does not tolerate the failure of the computer.
Good for testing, learning, demos.



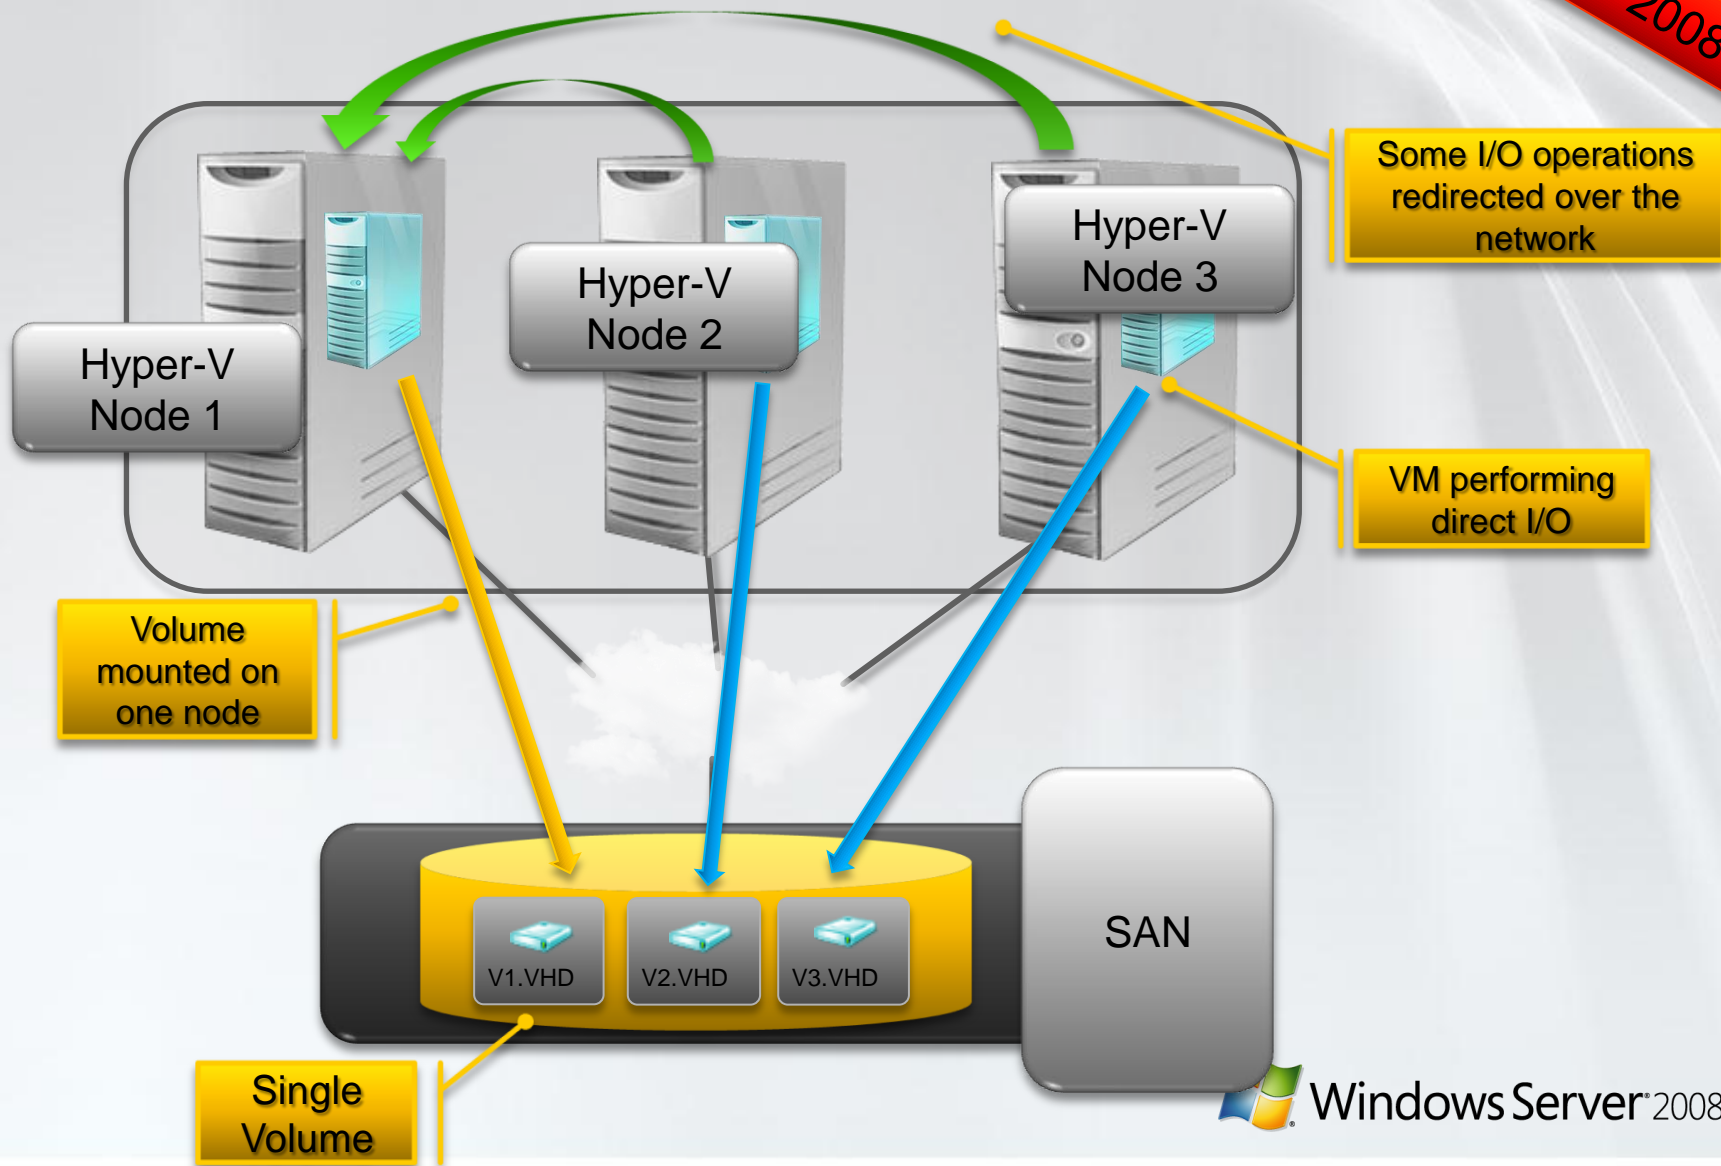
Windows Server® 2008 R2

4. Cluster Shared Volumes with the Dell iSCSI Target

*Understanding the Cluster Shared Volumes
feature of Windows Server 2008 R2, when
combined with Hyper-V and the Dell iSCSI
Target*

4.0. Cluster Shared Volumes

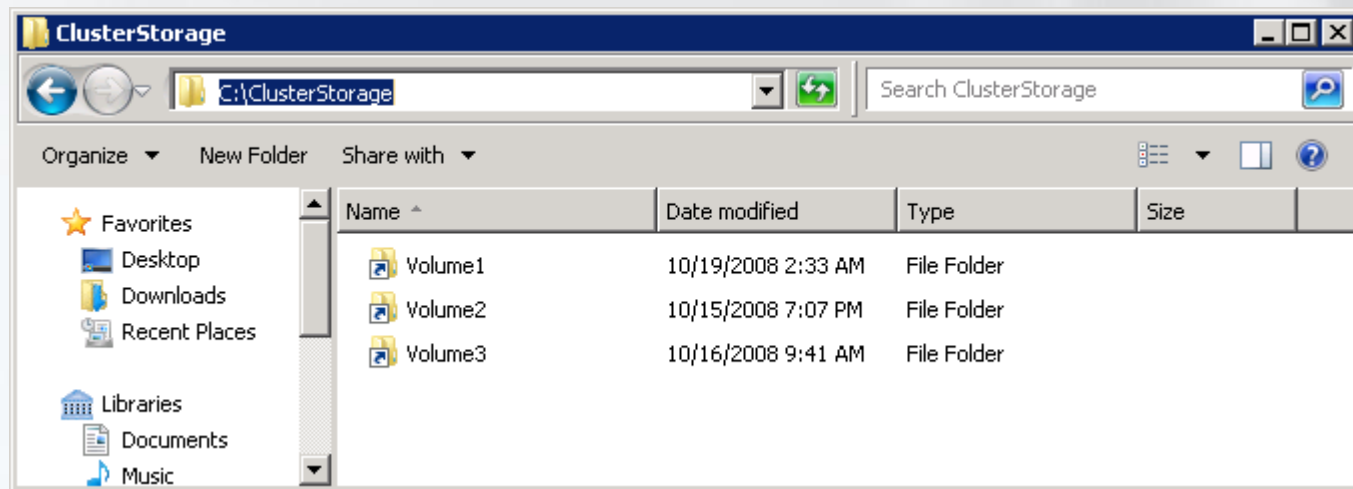
Windows Server 2008 R2



4.2. Single Name Space

Windows Server 2008 R2

- CSV provides a single consistent file name space
 - Files have the same name and path when viewed from **any** node in the cluster
 - CSV volumes are exposed as directories and subdirectories under the “ClusterStorage” root directory
 - C:\ClusterStorage\Volume1\<root>
 - C:\ClusterStorage\Volume2\<root>
 - C:\ClusterStorage\Volume3\<root>



4.3. Compatibility

Windows Server 2008 R2

- No special hardware requirements
- No file type restrictions
- No directory structure or depth limitations
- No special agents or additional installations
- Uses well-established NTFS file system

It just works!

4.4. Coordinator Node

Windows Server 2008 R2

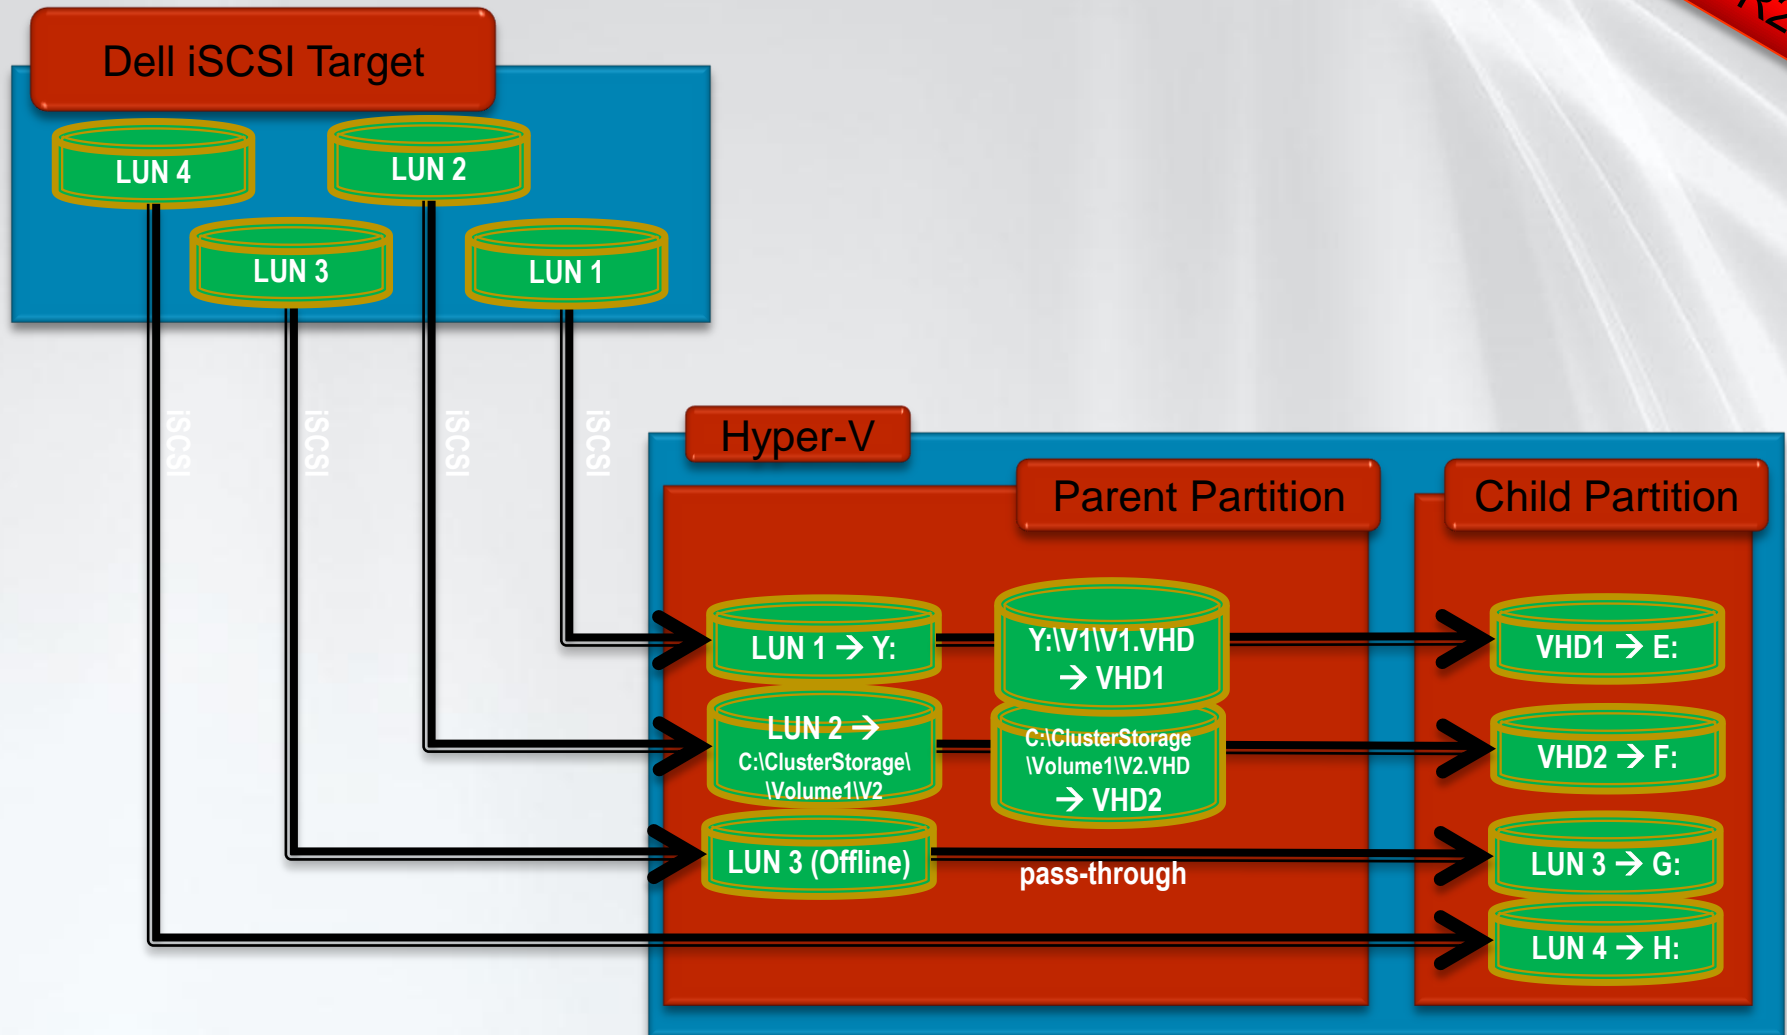
- Windows Server 2008
 - 1 node owns the disk resource
 - 1 node accesses the disk
- Windows Server 2008 R2
 - 1 node owns the disk resource, the Coordinator node
 - Every node accesses the CSV disk
 - The Coordinator manages CSV disk access



4.5. Options for Hyper-V

When using the Dell iSCSI Target
including the additional option with Cluster Shared Volumes

Windows Server 2008 R2



5. Summary

- Dell offers iSCSI Storage (Equallogic) which can be used to provide storage for Hyper-V
- There many Failover Clustering Options when you combine Hyper-V and the Dell iSCSI Target
- Windows Server 2008 R2 introduces a new option: Clustered Shared Volumes, which also can be combined with the Dell iSCSI Target and Hyper-V

Campus Days
14.-16. januar 2010



THANK YOU!



Windows Server® 2008 R2

Microsoft TechNet

Intel & Microsoft Present:

Maximizing Hyper-V iSCSI Performance

Breakthrough iSCSI Performance and Reliability



Jordan Plawner:

Senior Product Planner, Storage Networking, Intel

Suzanne Morgan:

Senior Program Manager, Windows Storage, Microsoft

Jim Schwartz:

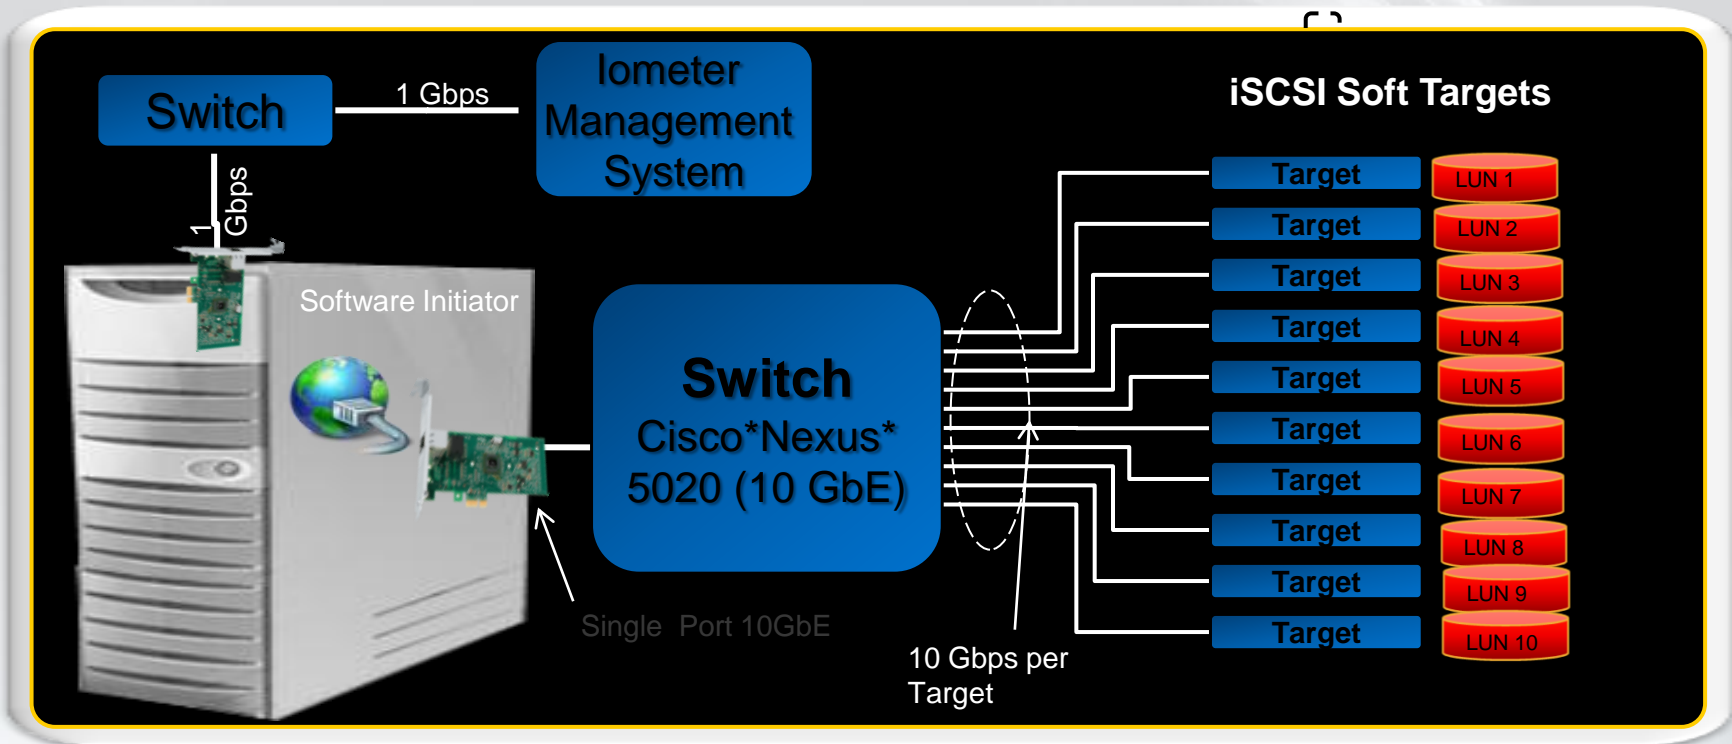
Solutions Marketing Director, Microsoft Virtualization

January 14th, 2010

<http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?EventID=1032432956&EventCategory=4&culture=en-US&CountryCode=US>



iSCSI Test Configuration



Server

- Windows Server 2008 R2
- Microsoft iSCSI Initiator
- Intel®Xeon® Processor 5580, quad core, dual socket, 3.2 Ghz, 24GB DDR3, MTU 1500, Outstanding I/Os =20

Adapter

- Intel® Ethernet Server Adapter X520 based on Intel® 82599 10GbE Controller

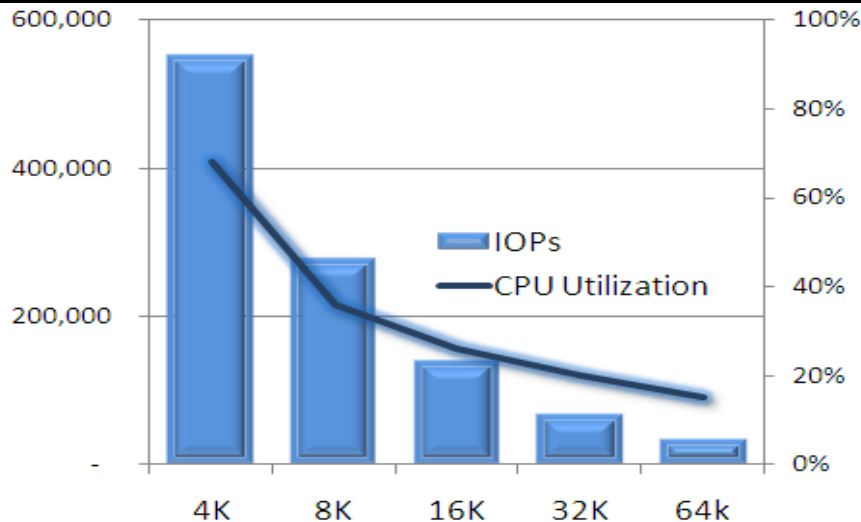
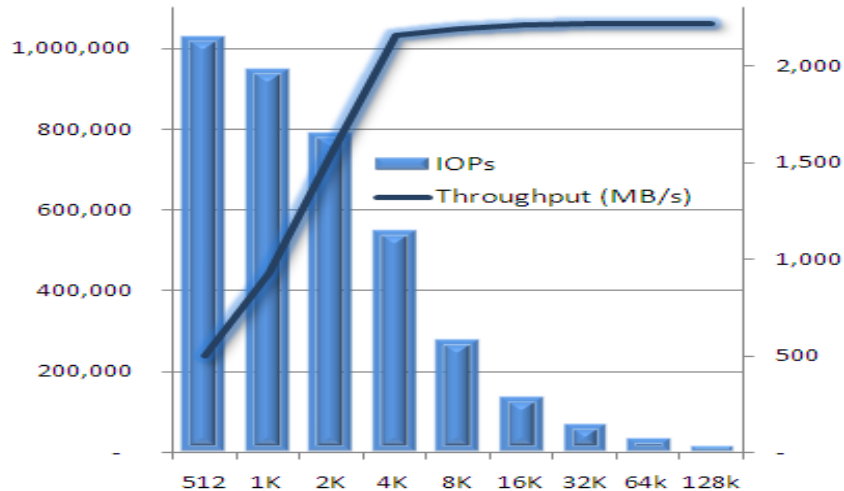
Performance factors

- iSCSI initiator perf optimizations
- Network stack optimizations
- Receive Side Scaling (RSS)
- Intel Xeon 5500 QPI and integrated memory controller
- Intel® 82599: HW Acceleration, multi-core scaling with RSS, MSI-X

Breakthrough Performance at 10GbE

Intel® Xeon® Processor 5580 Platform, Windows Server 2008 R2 and Intel® 82599 10GbE Adapter

Read/Write IOPs and Throughput Test



1,030,000 IOPs

Single Port

10GbE line rate

10k IOPs per CPU point

Performance for real world apps

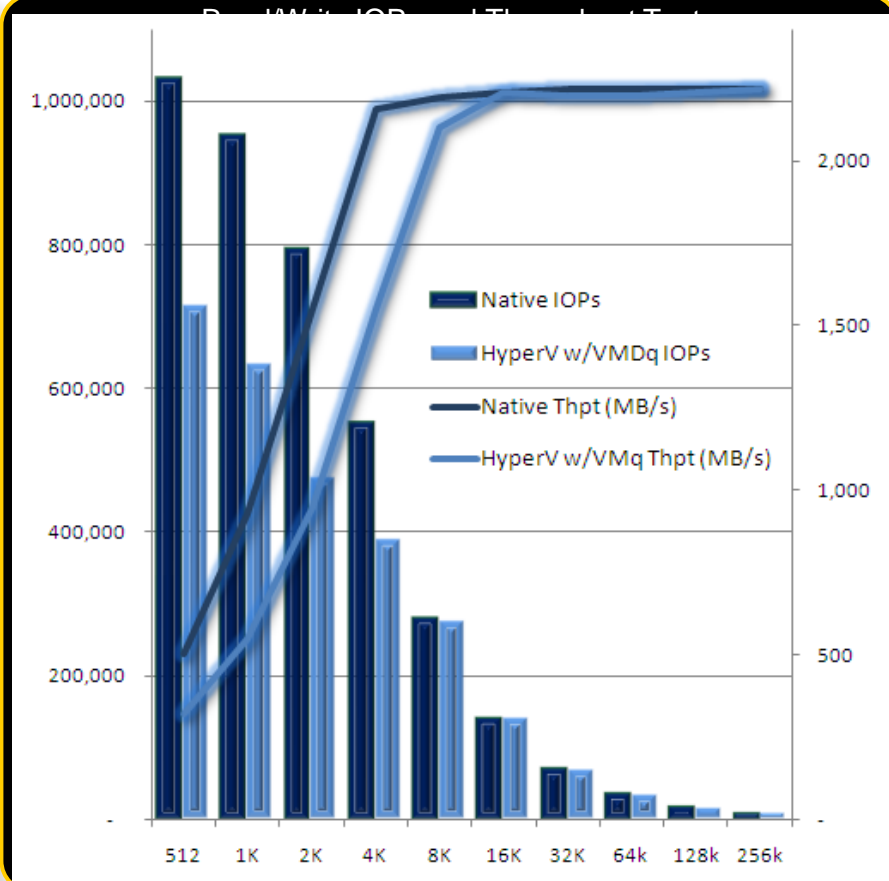
Future ready: Performance Scales

552k IOPs at 4k represents

- 3,100 Hard Disk Drives
- 400x a demanding database workload
- 1.7m Exchange mailboxes
- 9x transactions of large eTailers
- Jumbo frames: >30% CPU decrease is common for larger IO size (jumbo frames not used here)

Breakthrough Performance

iSCSI Performance with Intel® 82599 10G NIC with VMDq, Intel® Xeon 5580 Platform, Windows Server 2008 R2 and R2 Hyper V



715k IOPs -- 10GbE line rate
Intel VMDq and Microsoft VMQ
accelerates iSCSI to the guest
Hyper-V achieves native
throughput at 8k and above
Future ready: Scales with new
platforms, OS and Ethernet
adapters

***Near native iSCSI
performance***