August 2008

Storage Management for Exchange

LeftHand NETWORKS

LeftHand Networks, Inc.

Leader in iSCSI SANs

- Pioneer in the IP SAN market, founded in 1999
- Highly available, simple to manage, and “grow as needed” architecture

Rapid market acceptance and growth

- More than 11,000 installations; over 3,000 customers
- Strategic VARs and resellers in North America and Europe

Market focus

- Mid-tier to enterprise companies
- Exchange, SQL Server, SharePoint
- Virtualization

Quiz

LeftHand Networks got it's name:

1. Because the CEO and founder is left-handed
2. Because it was founded at the Left Hand Brewery
3. Because its HQ is near Left Hand Canyon in Boulder, CO

Storage Requirements for a Dynamic Data Center

Performance	• Support concurrent workloads, no performance degradation
Scalability	• Scale capacity and performance without disruption
Availability	• Data stays online during a site failure, recover from disasters
Manageability	• Simple to administer and change

Enterprise Exchange Server Topology

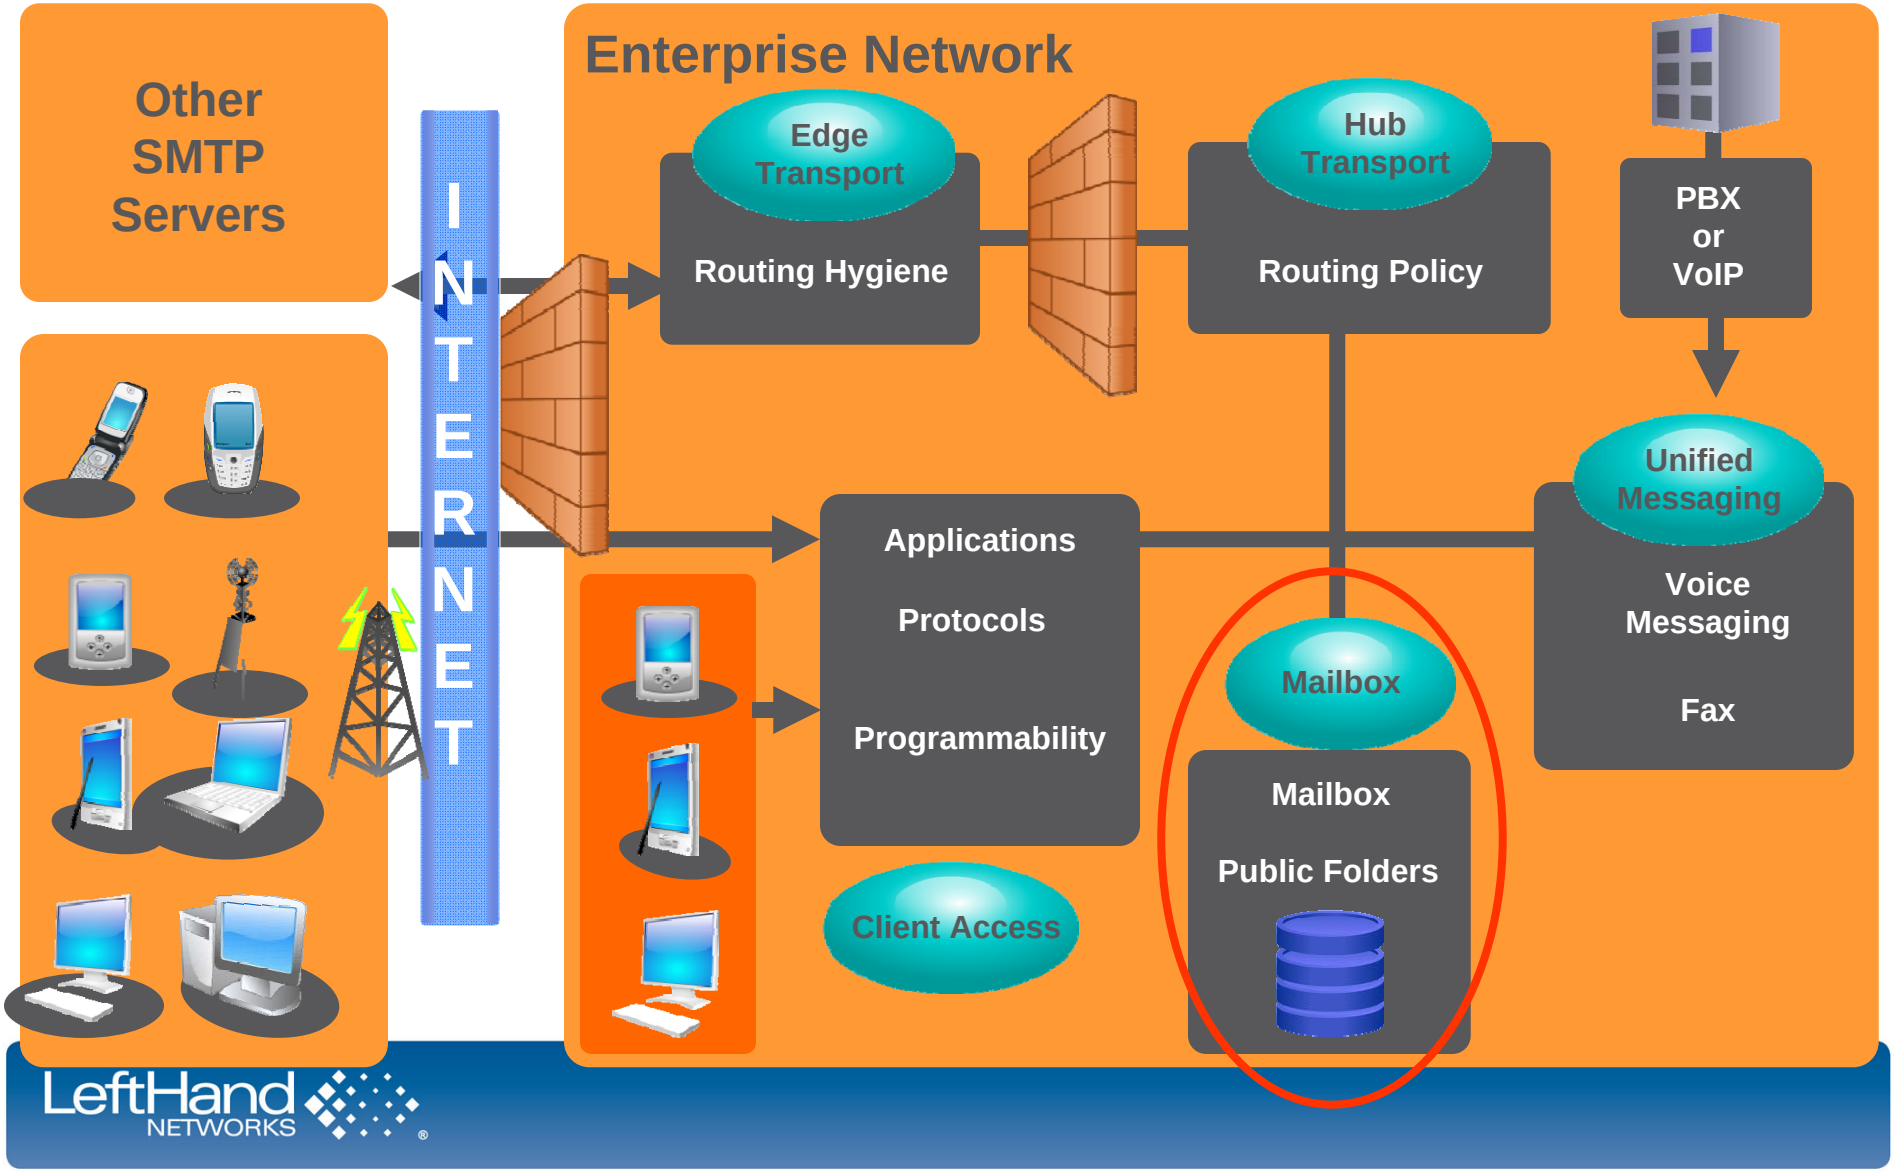

Typical Exchange 2003 Disk Layout

ORG = CompanyA
2200 mailboxes
2 Exchange servers

Storage Group A
1000 Mailboxes

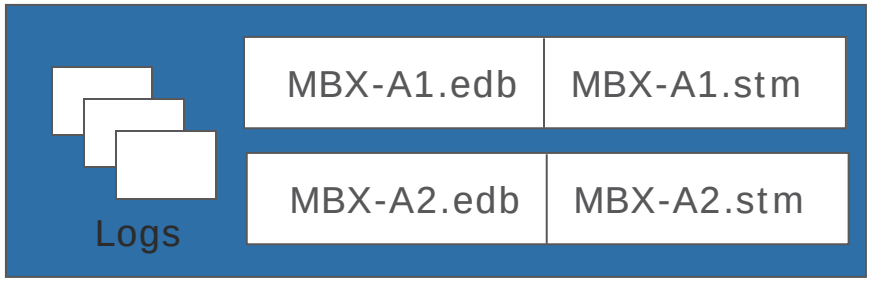

Storage Group B
1200 Mailboxes
Public folders

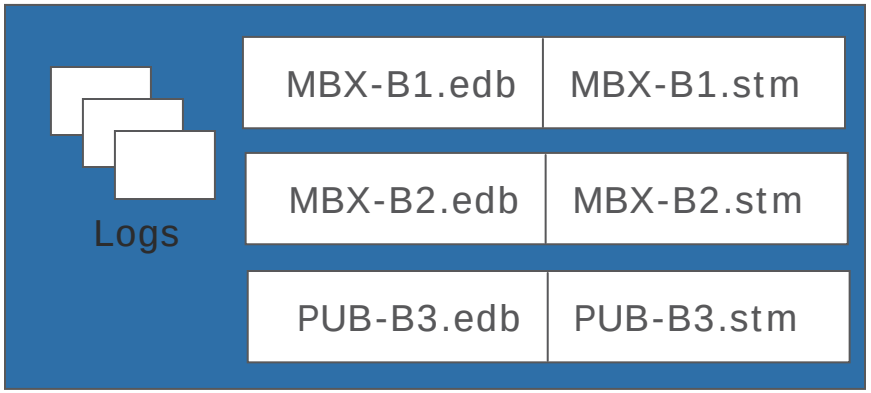

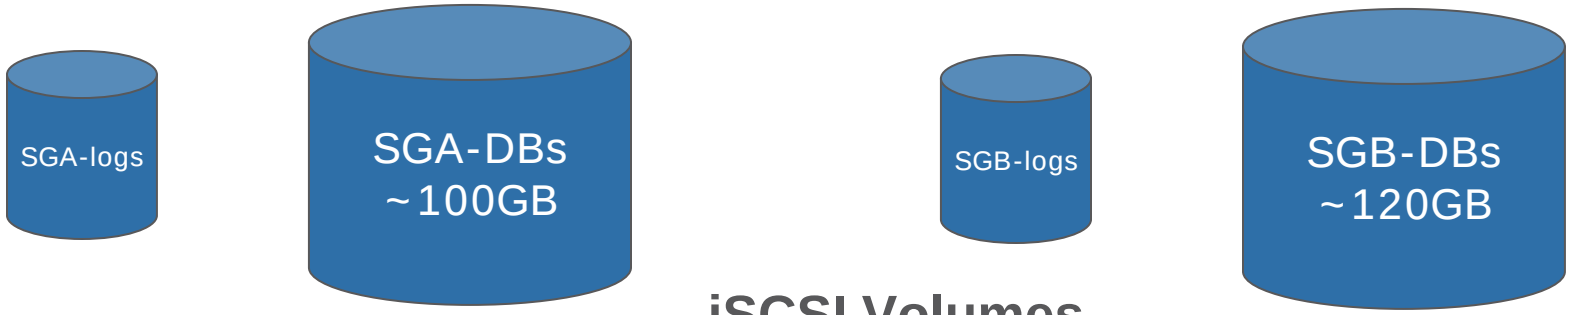

iSCSI Volumes

Typical Exchange 2003 Disk Layout

ORG = CompanyA
2200 mailboxes
2 Exchange servers

Storage Group A
1000 Mailboxes
1000 IOPS

Storage Group B
1200 Mailboxes
Public folders
1200 IOPS

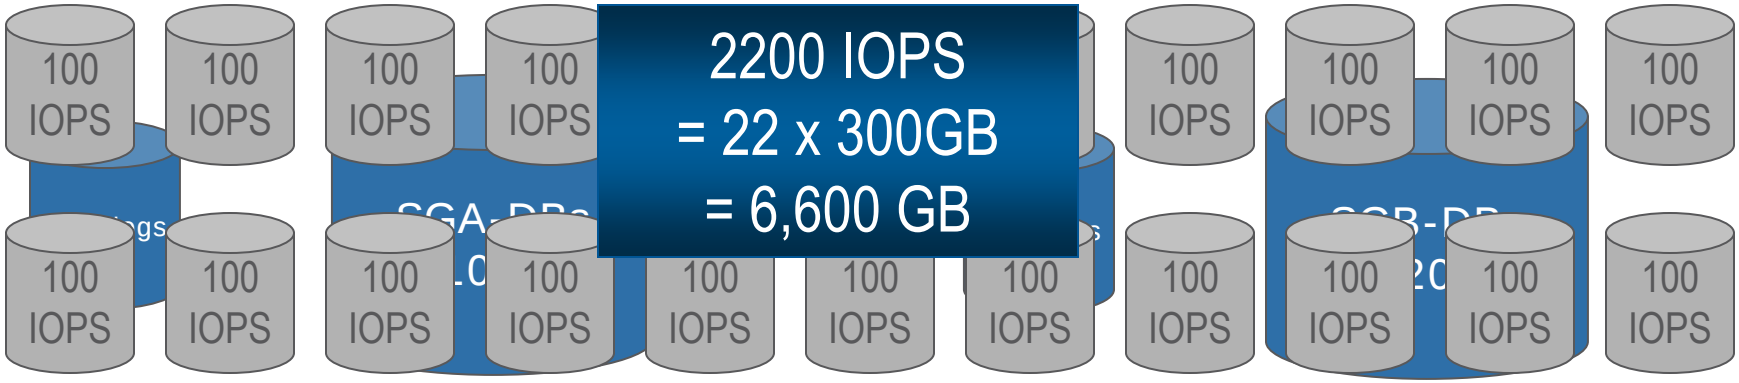

Exchange 2000 / 2003 Performance Sizing

Exchange I/O behavior

- Exchange database I/O ratio: 4KB I/Os, 70/30 r/w split
- Exchange log volume I/O ratio: 64KB I/Os, 100% writes
- Complexities: the IOPS required by RAID, snapshots, and replication

Best practice

- Usually the # of spindles required is driven by IOPS, not capacity
- Best practice (best $ per IOPS) is smaller (146GB) 15K SAS or SCSI drives. 300GB drives can be also used
- SATA is not generally cost efficient except for smaller environments (<500 mailboxes)

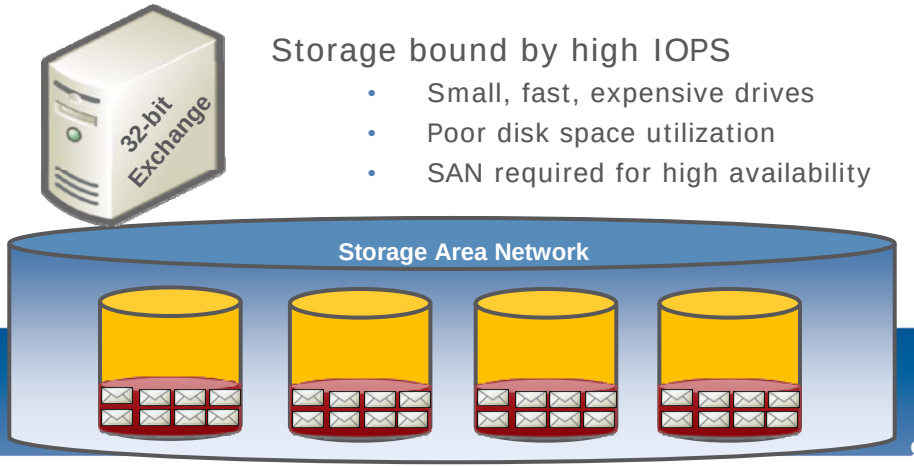

Exchange 2007 Performance Sizing

Exchange I/O behavior

- Exchange database I/O ratio: 8KB I/Os, 50/50 r/w split
  - Many more reads are cached, and many I/O are > 8KB
- Exchange log volume I/O ratio: 128KB I/Os, 100% writes
- Complexities: the IOPS requirements of RAID, snapshots, SAN/iQ replication, Exchange replication, unified storage

- Best practice
  - Spindle requirements are more balanced between IOPS and capacity
  - Best practice (best $ per IOPS) are larger (300GB) 15K SAS or SCSI drives
  - SATA may be cost efficient for up to 1500 mailboxes, or in any environment with very large mailboxes

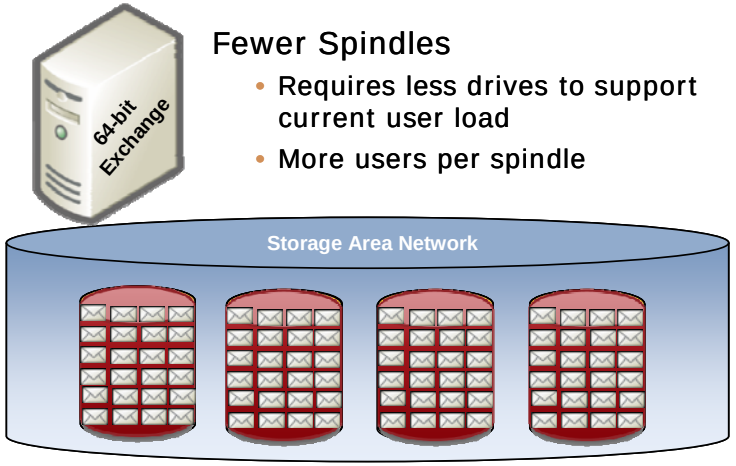

“Rule of Thumb” with SCSI / SAS / SATA Disks

IOPS per Drive	Exch 2000 / 2003	Exch 2007
	Random 4KB R/W with R5	Random 8KB R/W with R5
SCSI	125	115
SAS	130	120
SATA	70	60

IOPS per Mailbox	Exch 2000 / 2003	Exch 2007
Heavy	1.0	0.3
Medium	0.5	0.2
Light	0.3	0.1
POP	0.1	0.1
Mobile	0.5	0.4

How to calculate spindle count:

- Determine number and type of users that will be using system at any one time
- Determine number of spindles
- Multiply answer by 110% for unified messaging (Exch 2007)
- Multiply answer by 150% for replication

Note: these calculations will be for ***optimal*** performance with the described environment (<250ms average response time for Outlook users)

Performance Planning Scenario

You've deployed your brand new Exchange 2007 system

You ended up with more mailboxes than planned, fewer servers, larger database, and your users jumped all over unified messaging and deployed mobile devices for everyone

Performance Censored! (it's bad)

So, how do you solve the problem?

- How much re-planning do you do?
- How much time working with the products / vendors?
- How much downtime to fix the problem?
- How complex is the fix?

Typical Storage Array Architecture

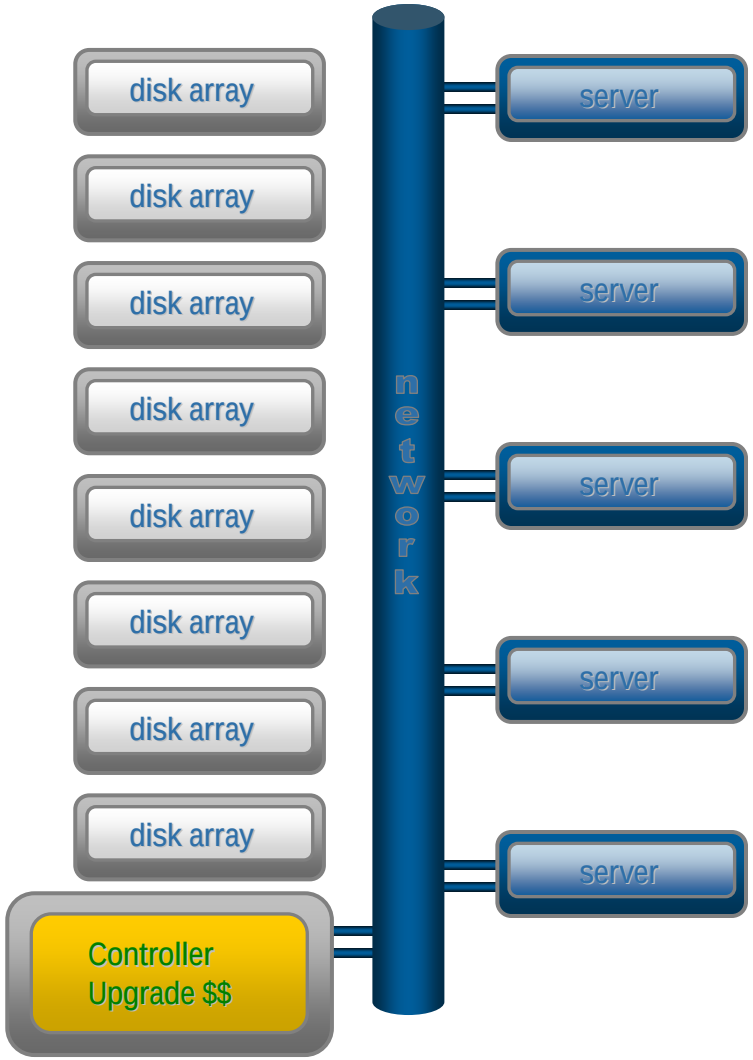

At this point with traditional arrays, the controller typically becomes a bottleneck and a costly upgrade must be implemented to further scale performance and capacity.

Traditional Storage

Example FC Array LUN Configuration

SAN Storage

A SAN unit stored the data for the Internet Platform and Operations group deployment of Exchange Server. The unit comprised switches, two controllers, and three disk shelves, each of which housed fourteen 36.4-GB hard disk drives. The raw capacity of the SAN was 509.6 GB per shelf, or approximately 1.5 terabytes (TB) in total.

Figure 4 and Table 3 show how the SAN drives were allocated. Unused drives 105, 302, 402, 500, 504, 600, 601, and 604 are shown as white rectangles.

Port 1	Disk 100 000	Disk 101 000	Disk 102 000	Disk 103 000	Disk 104 000	Disk 105 000	Disk 108 000	Disk 400 000	Disk 401 000	Disk 402 000	Disk 403 000	Disk 404 000	Disk 405 000	Disk 408 000	**Port 4**
Port 2	Disk 200 000	Disk 201 000	Disk 202 000	Disk 203 000	Disk 204 000	Disk 205 000	Disk 208 000	Disk 500 000	Disk 501 000	Disk 502 000	Disk 503 000	Disk 504 000	Disk 505 000	Disk 508 000	**Port 5**
Port 3	Disk 300 000	Disk 301 000	Disk 302 000	Disk 303 000	Disk 304 000	Disk 305 000	Disk 308 000	Disk 600 000	Disk 601 000	Disk 602 000	Disk 603 000	Disk 604 000	Disk 605 000	Disk 608 000	**Port 6**
Ctrl 1	**Storageworks 2200 Controller Shelf**														**Ctrl 2**

Better by design: SAN/iQ Architecture

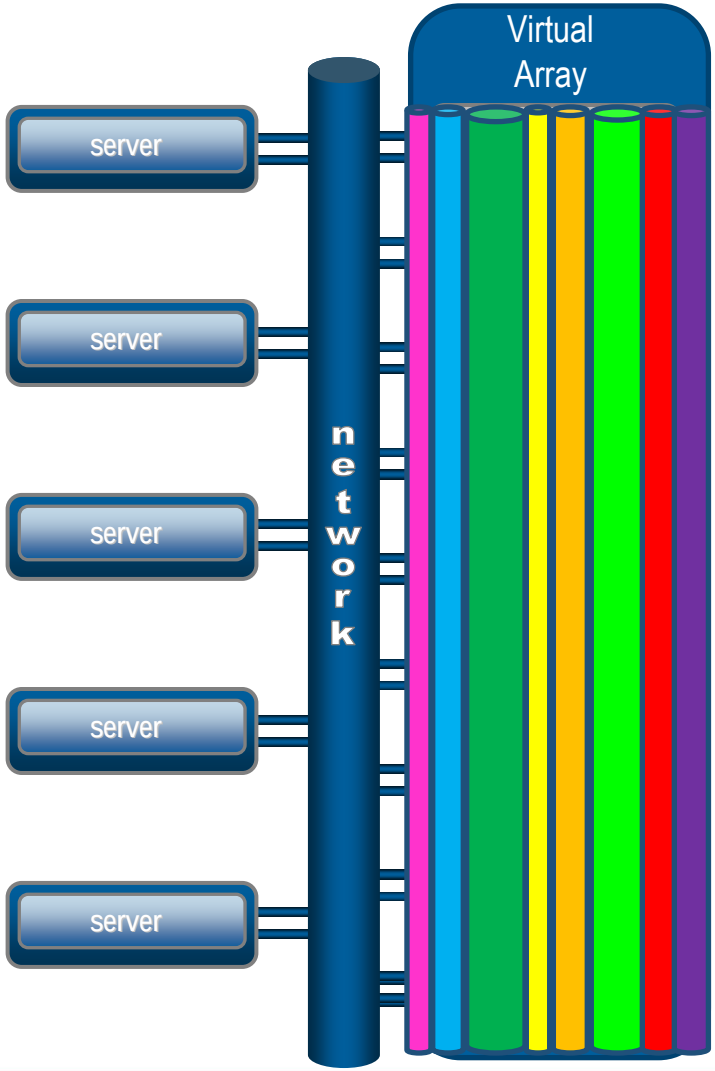

LeftHand differentiators

- True clustered storage
- Highly virtualized for ease of management
- Scalable capacity and performance
- Always on-line storage volumes
- Advanced data availability via network RAID
- Multi-site SAN

Better By Design: Change Management

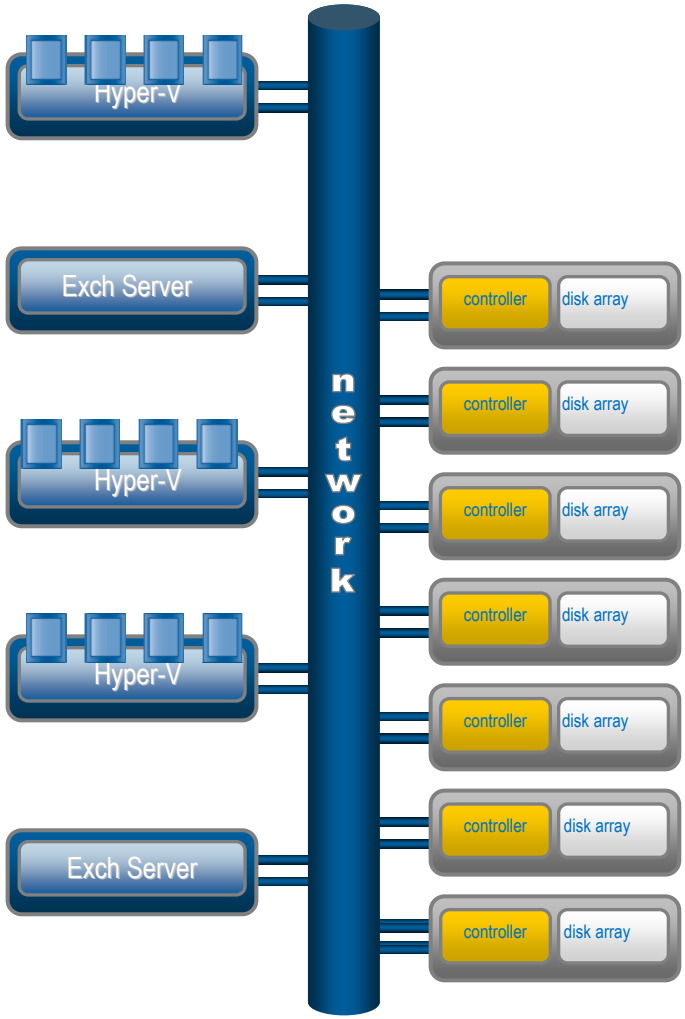

Average Response Time for 2500 mailboxes

250ms

Or, number of mailboxes @ 900ms

5000

iSCSI volumes are automatically restriped as each node is added:

- No downtime
- No re-planning your SAN
- No LUN reconfiguration
- No administrative activity what-so-ever

ESRP Submission

50,000 user Exchange 2007 submission

- Largest iSCSI-based Exchange solution submitted (Exchange 2003 or 2007)
- First iSCSI vendor to submit Exchange 2007 results
- Demonstrates true scalability
- 108TB (30 nodes)

Random I/O Scalability

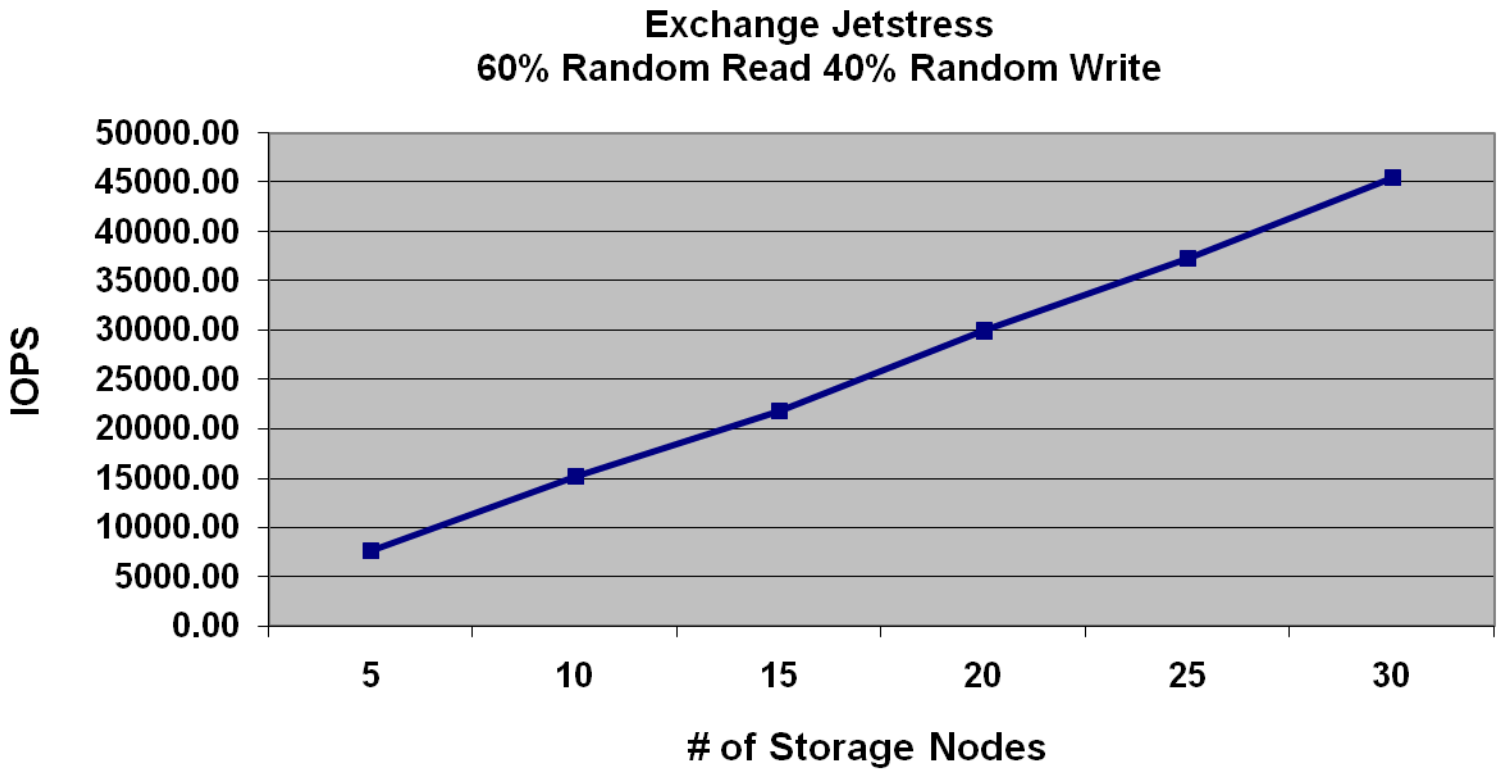

Exchange HA and Data Protection

Many options…

- Traditional backups
- Snapshots (file system, SAN, VSS)
- Proxy host backups, or “in SAN” backups
- Replication (synchronous, asynchronous)
- Continuous data protection
- Exchange LCR/CCR
- RAID, duplicate components

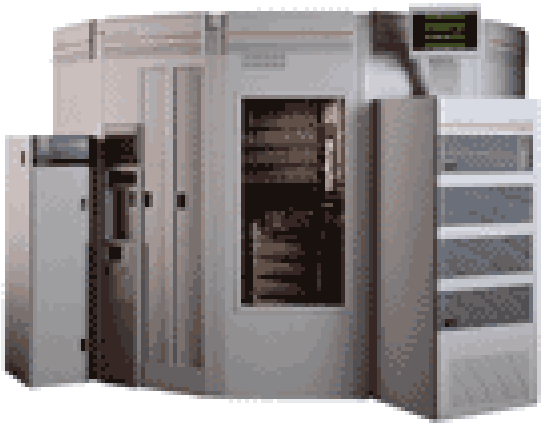

LeftHand Technology Primer

SAN/iQ: storage virtualization

- Clustered for distributed performance
- Network RAID for ultra-high availability

Physical storage nodes

- Preconfigured storage appliances
- Off-the-shelf x86 server platforms
- LeftHand branded or from Dell, IBM, and HP

Full-featured solution at all price points

- Solve any performance and capacity need
- Online scalability
- Beyond-the-box data protection
- Self-managing thin provisioning & snapshots
- Asynchronous replication for DR

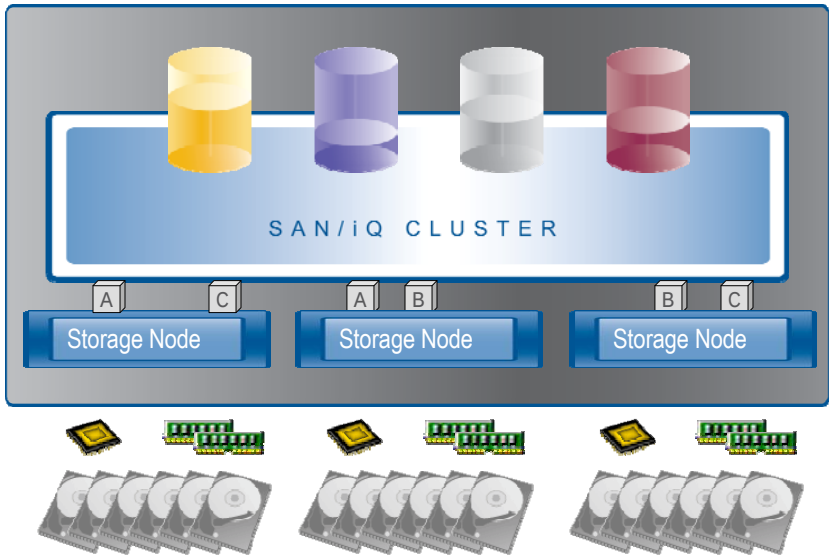

Beyond-the-Box Data Protection

Patented Network RAID

- Volume level granularity
- Replication levels
  - 2-way
  - 3-way
  - 4-way
  - None
- Change level on-the-fly

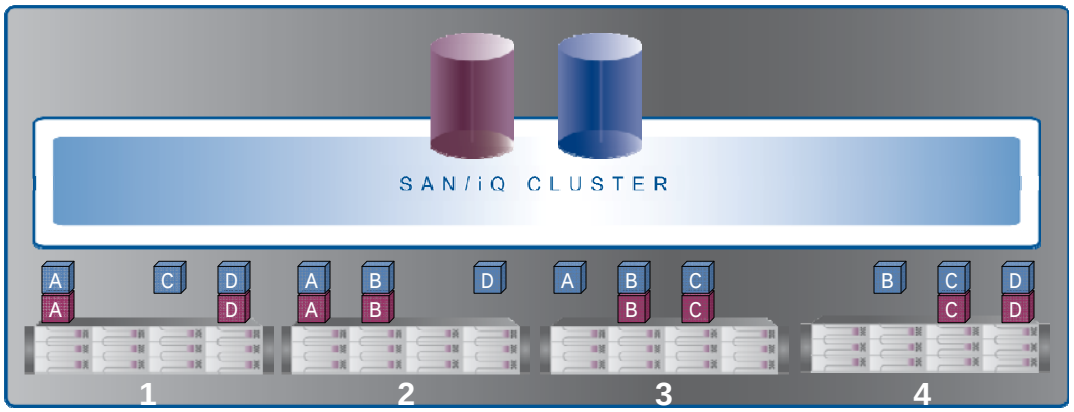

Ensures high data availability

- Survive storage node failure
- No disruption of data access!

Protect Your Data Across Locations

Multisite SAN

- Separate clusters By:
  - Rack
  - Room
  - Floor
  - Building
- Keep data online during:
  - Facility disruption
  - Natural disaster
  - Site maintenance
- Network prerequisites
  - Gigabit connectivity
  - ~4 millisecond latency

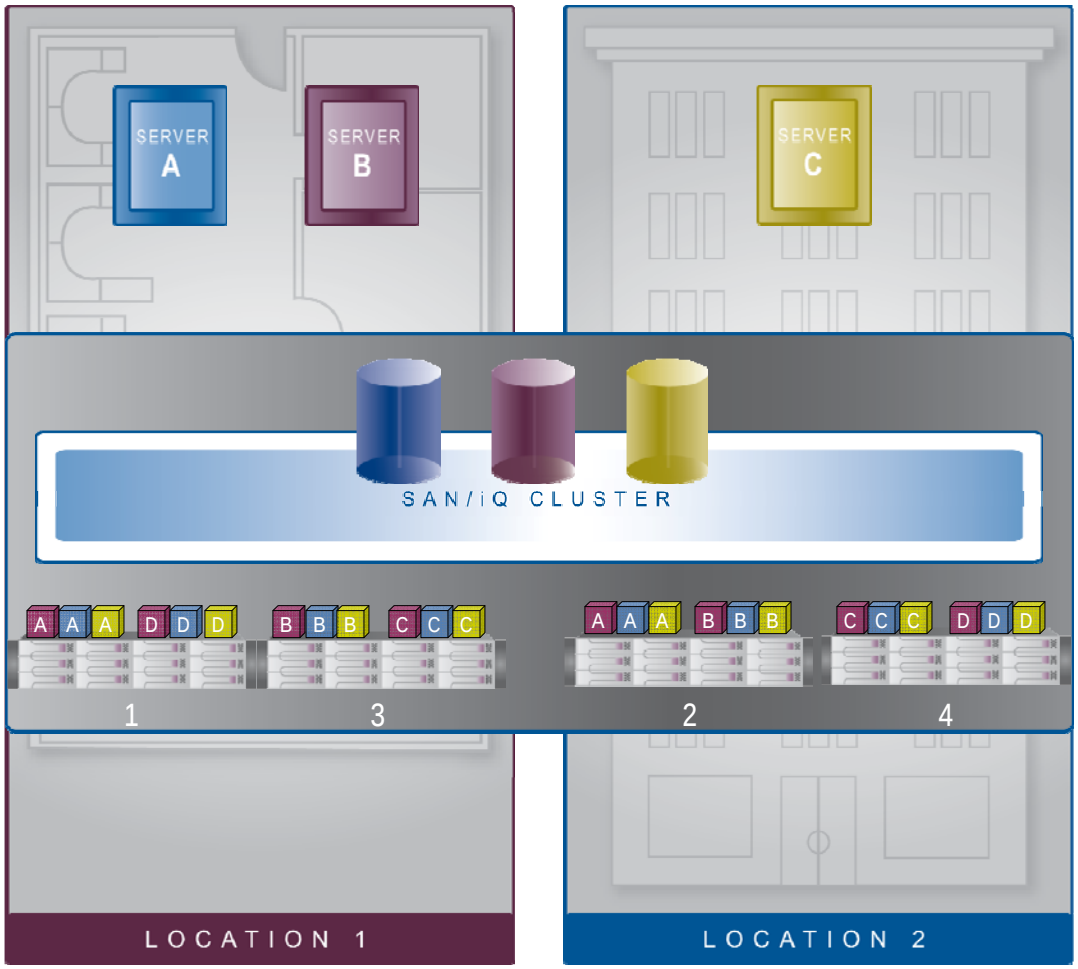

Multisite SAN & Windows Virtualization: Complete HA

Complete high availability

- Survive single and multiple points of failure:
  - Windows server
  - LeftHand storage node
  - Single or multiple locations

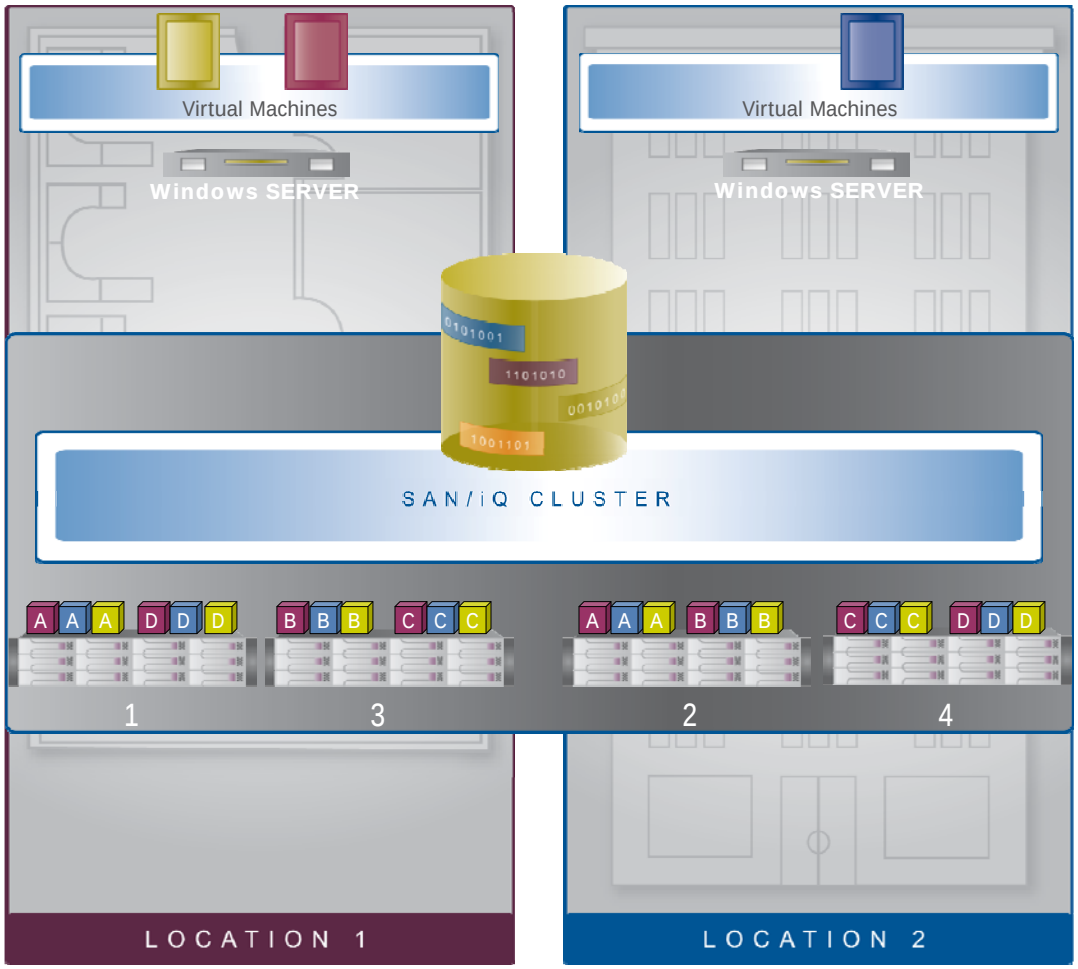

Exchange CCR / LCR

Built in mechanism for replicating Exchange databases

Uses “log forwarding” to replay transaction logs into a replica database on the same server (LCR) or a cluster stand-by server (CCR)

Standby Continuous Replication in SP1

- Allows replication to multiple additional servers
- Many-to-many
- Replay lag time

No point-in-time recovery

- Data Protection Manager 2007

Primarily designed for servers with DAS

Confusion with SAN replication products

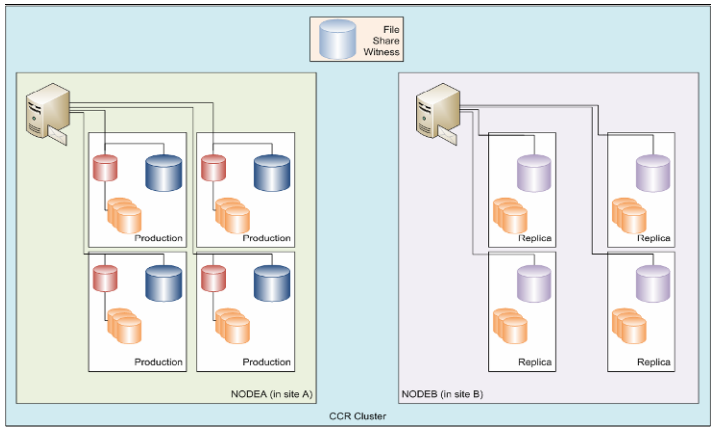

VSS: Synchronizing Exchange

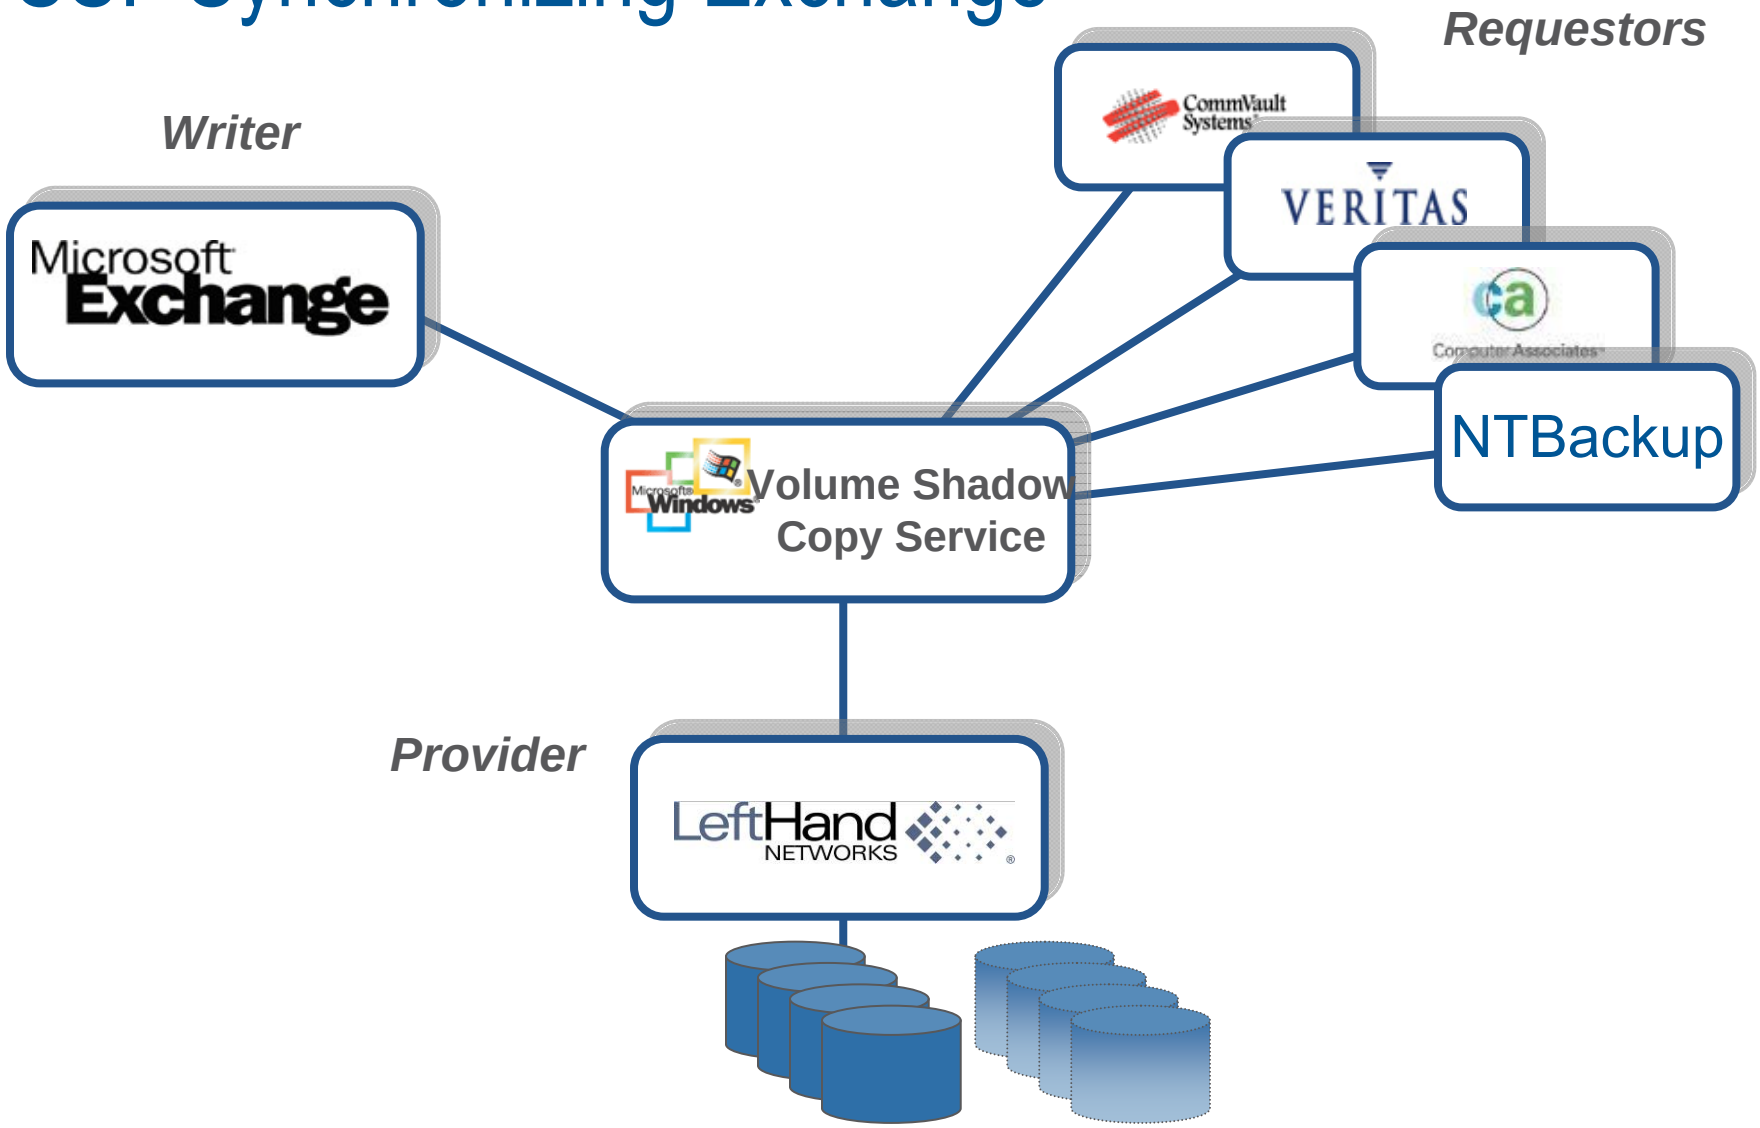

Microsoft System Center
Data Protection Manager 2007

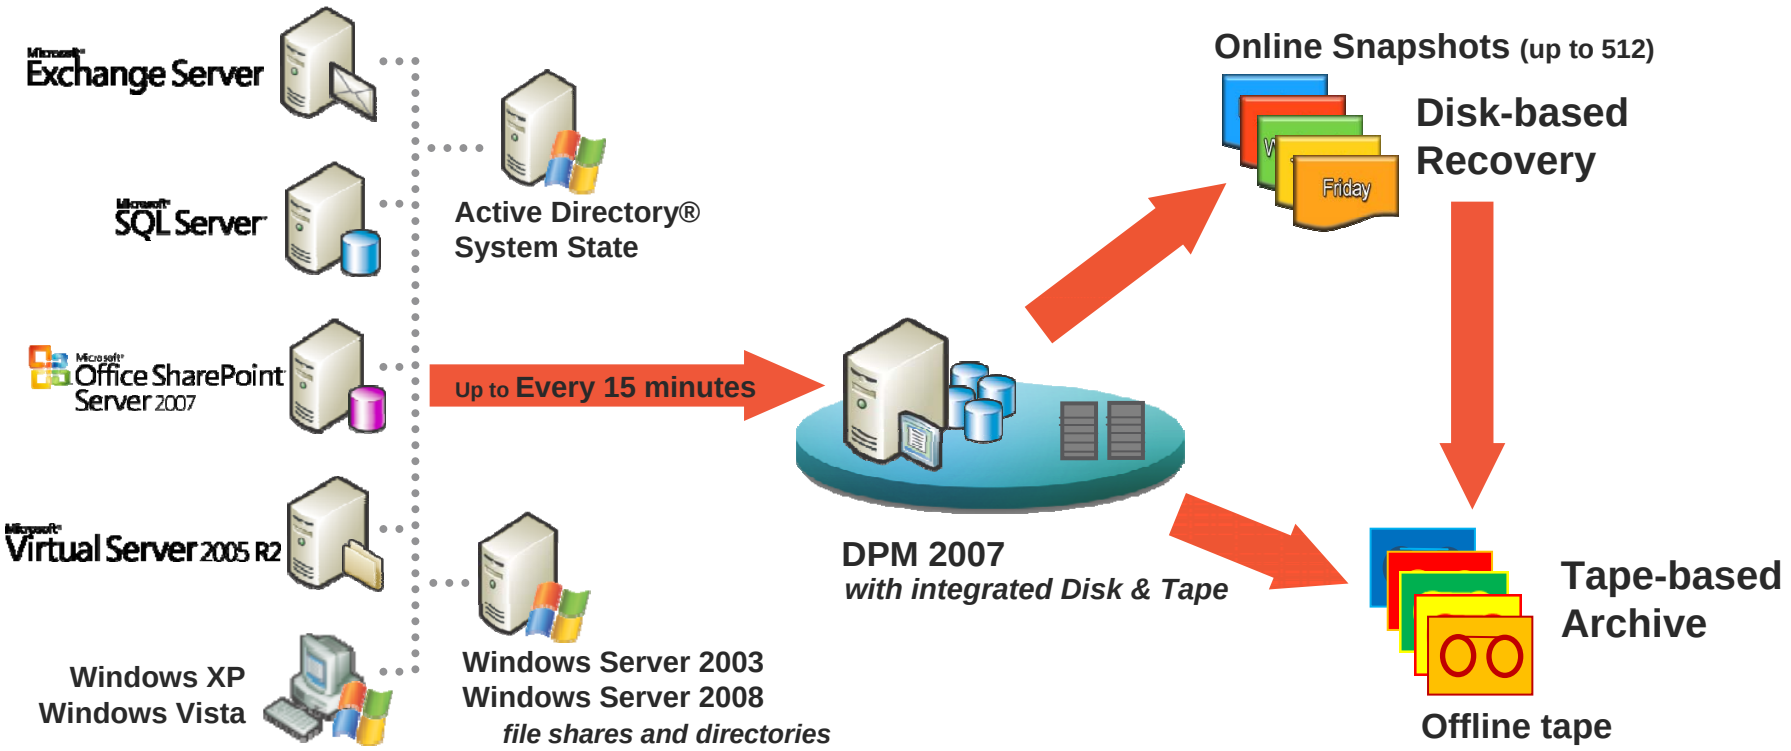

DPM 2007

- Continuous Data Protection for Windows Application and File Servers
- Rapid & Reliable Recovery from disk instead of tape
- Advanced Technology for enterprises of all sizes

Exchange Data Protection Recommendation

For component and site failure protection

- Network RAID (2-way) for SAN protection
- Multi-site SANs w/ Microsoft Clusters for environmental protection
- Remote copy for long distance DR

For logical failures and user error

- VSS Snapshots or DPM 2007 (depending on RPOs)

Bringing It All Together…

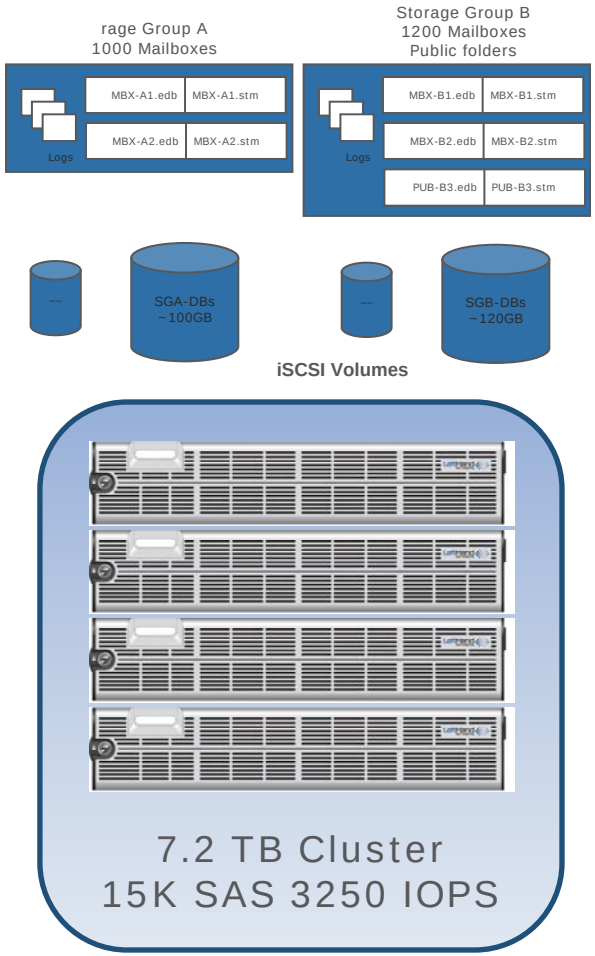

iSCSI SAN from LeftHand Networks

- Provides for unlimited scalability
- Scales both performance and capacity
- Protects against component failures

Bringing It All Together…

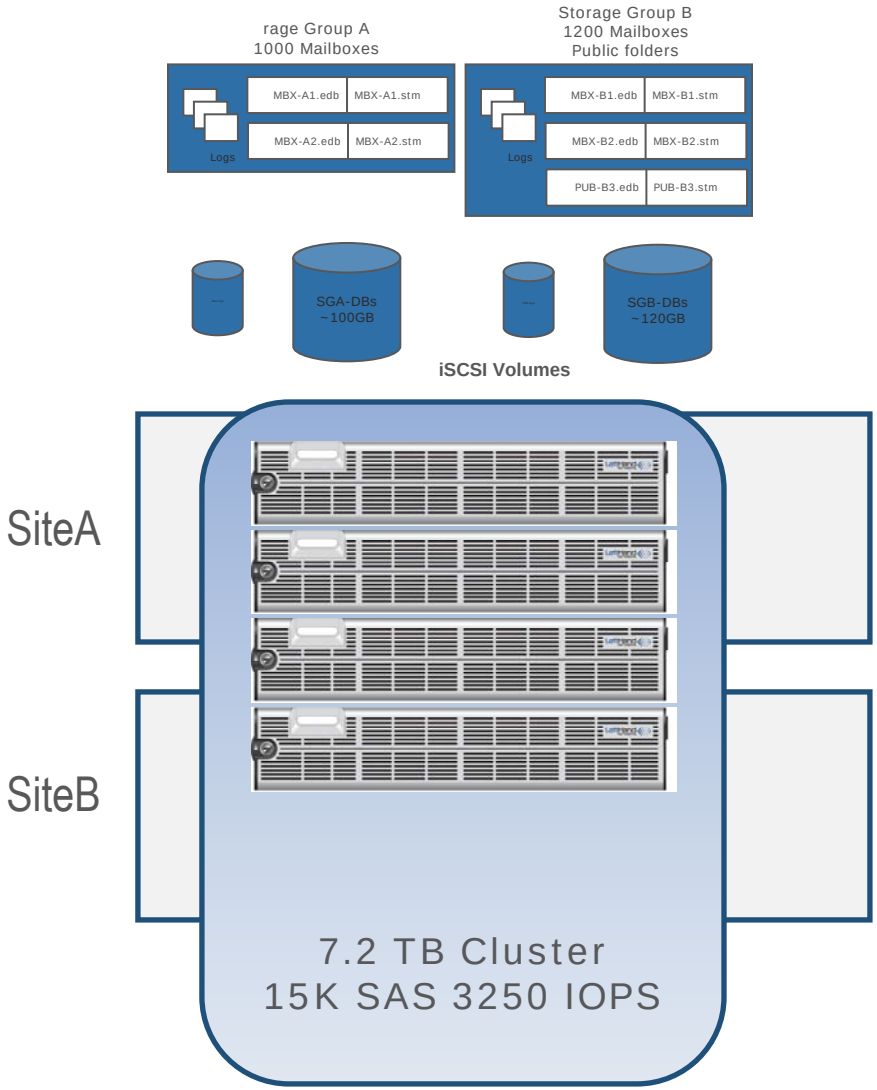

Multi-site SAN

- Site-wide disaster protection

Bringing It All Together…

Data Protection Manager
- Protection against logical failures
- 15 min RPOs
- 2 to 4 weeks retention
- Use SAN for bulk data movement

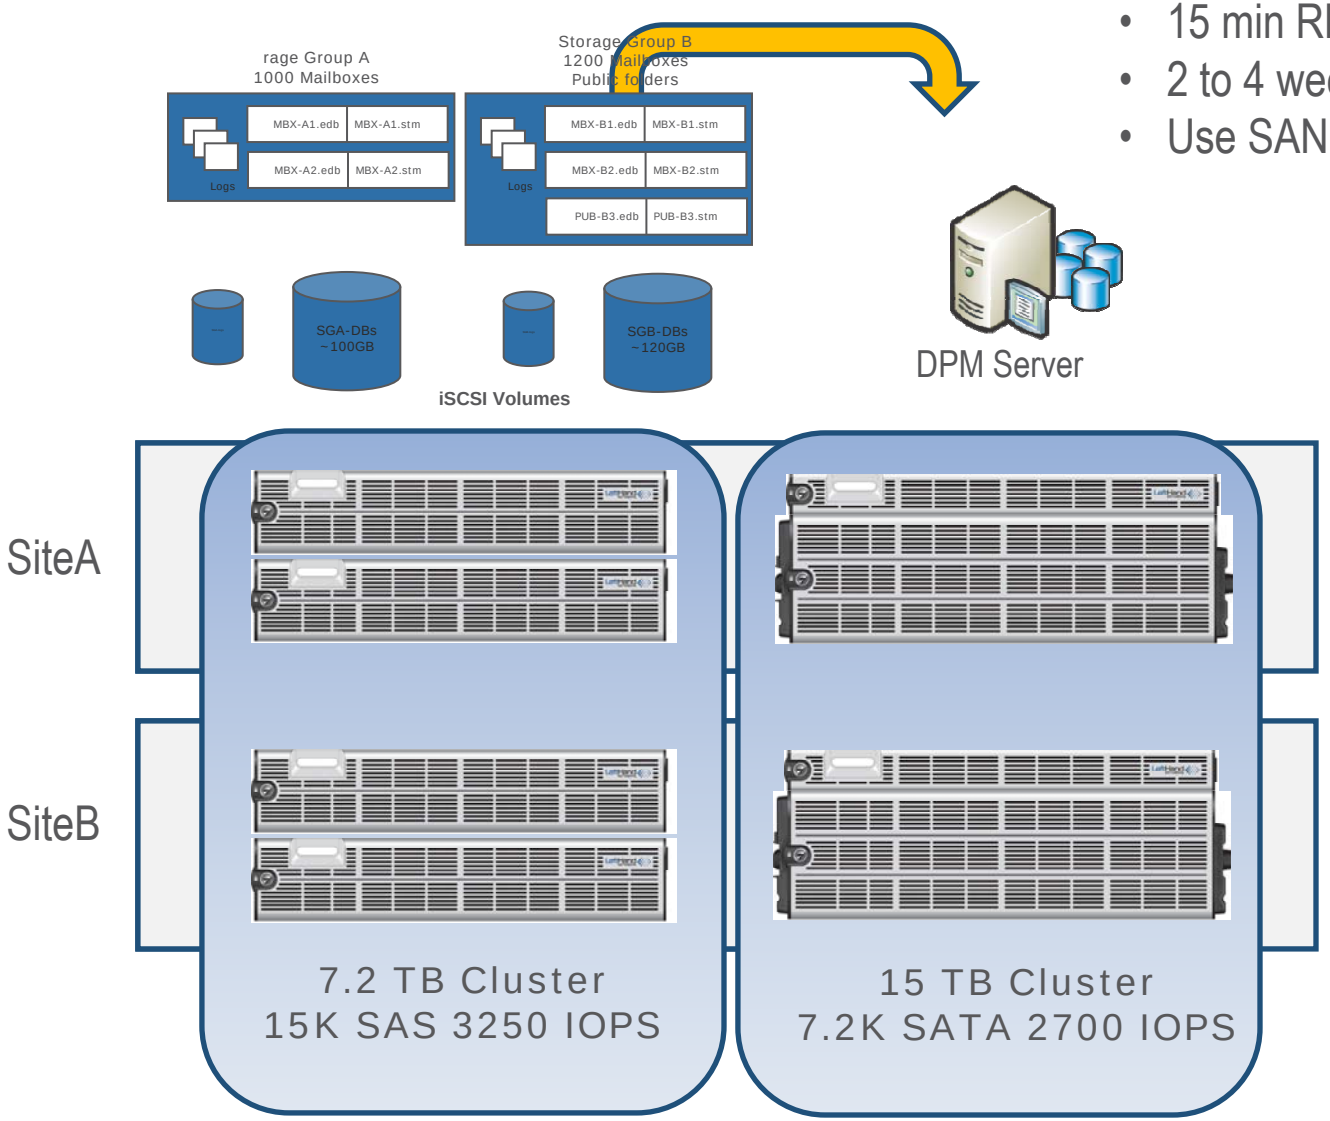

Bringing It All Together…

Remote DPM Server
- Protection against metro-wide disasters
- Can be used for archive site

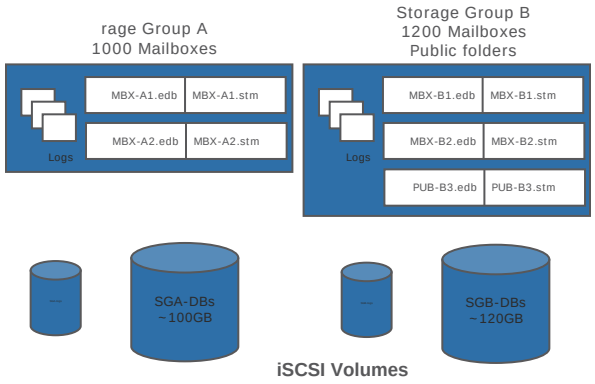

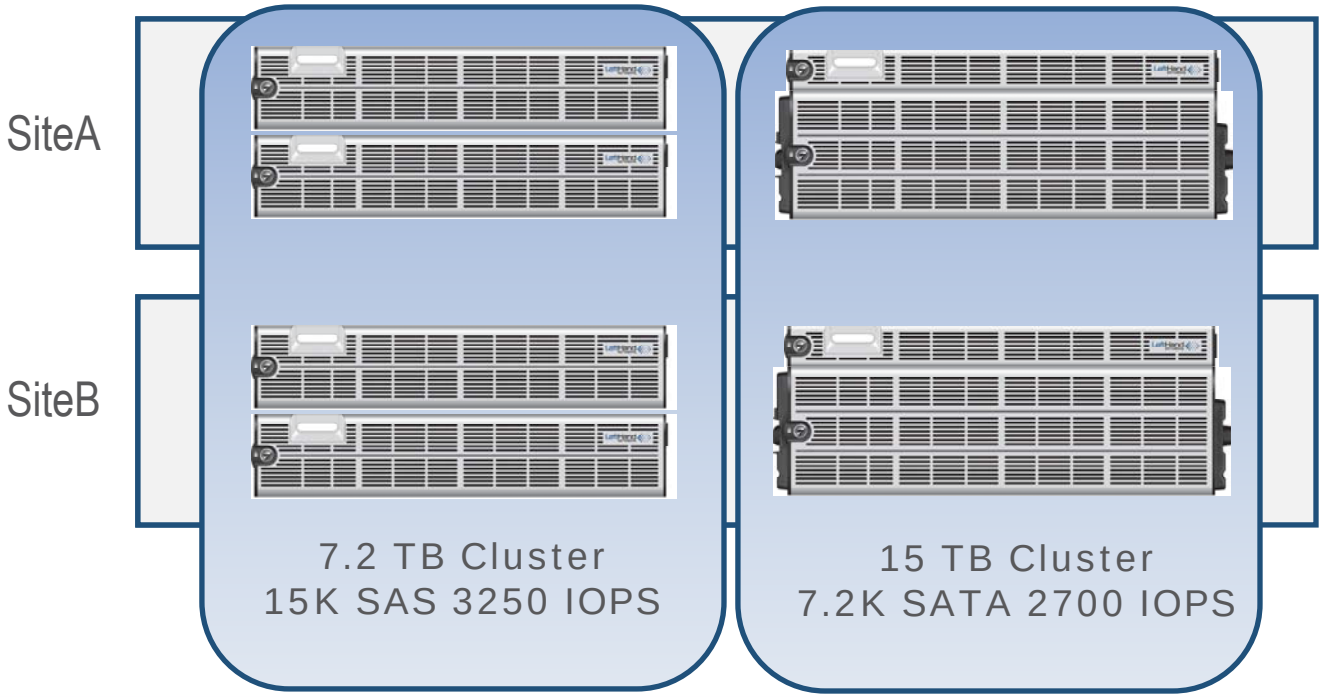

Non-disruptive Management

The following management activities are done without downtime or loss of volume access:

- Adding or removing storage nodes from a cluster
- Restriping (load balancing) iSCSI volumes as new nodes are added
- Recovering a storage node
- Changing volume size
- Changing a volume from thin provisioned to full provisioned, or vice versa
- Changing network RAID settings on a volume
- Migrating a volume from one SAN/iQ cluster to another

These activities occur far more frequently in a virtualized environment – legacy SANs (LUNs) are statically provisioned and far less flexible

Exchange 2007 Resources

Download from www.lefthandnetworks.com

50,000 mailbox Exchange 2007 iSCSI Solution - replicated

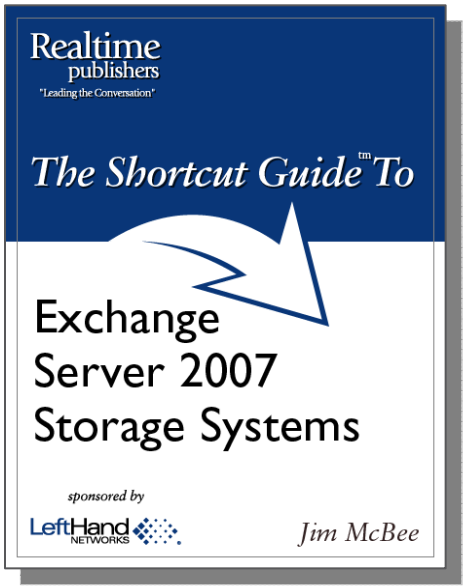

The Shortcut Guide to Exchange Server 2007 Storage Systems eBook
(3 chapters currently available)

Seven Ways an iSCSI SAN Simplifies Management of Exchange 2007 white paper

Summary

Microsoft

LeftHand SANs are easy to manage, highly automated, and do not require special expertise, training, or add-on products

LeftHand SANs address the four key storage requirements:

- Performance
- Scalability
- Availability
- Management

LeftHand delivers all the enterprise-class storage features you need and doesn't break the bank

Demonstration

LeftHand NETWORKS