VMware ESX Server and Storage Architecture Best Practices for Performance, Backup, and Disaster Recovery

Brad Maltz
Principal Solutions Consultant
International Computerware, Inc.

VMWORLD 2006

Agenda

- Overview of the Complete Solution
  - Description
  - What is our Approach
  - How do we break it down
- Storage Requirements
  - Design Best Practices
- Backup
- Disaster Recovery
- Network Requirements
- Missing Pieces

Nirvana / Utopia

- Optimally Performing Virtual Machines
  - Storage and Networking are designed by Best Practices
  - Standardization helps to maintain levels of Performance
- Complete Backup of VMDK Files
- File level backup of Windows VMDKs
- Full Site Replication of Virtual Machines
- DR for Physical Servers as VMs

BUT storage is the central focal point!!!!

Case Study Overview

- During the presentation we will reference the following generalized case study
- The case study is conglomeration of environments I have architected
- The hardware and software that are utilized can be substituted for other vendors
- Customer Requirements
  - RPO of servers is between 10 mins and 24 hours (depending on application)
  - RTO of servers is between 4 hours and 2 weeks (depending on application and tier)
  - Consolidate backups with SAN based backup
    - Low I/O Servers via File Level and All VMDK Files
  - Meet performance requirements of Exchange, SQL and Kronos

Servers in Case Study

Financial SQL Server	ICI-FINSQL
Domain / DHCP / DNS	ICI-DC01
Exchange 2003	ICI-EXCH01
SAS Servers	ICI-SAS01 thru ICI-SAS04
Blackberry	ICI-BB
Citrix	ICI-CITRIX01
Kronos Server	ICI-KRONOS
Web Servers	ICI-WEB01 thru ICI-WEB03
All Development Machines	ICIDEV-Test01 thru ICIDEV-Test06
Image Server	ICI-PSPIN

Case Study / Site Example

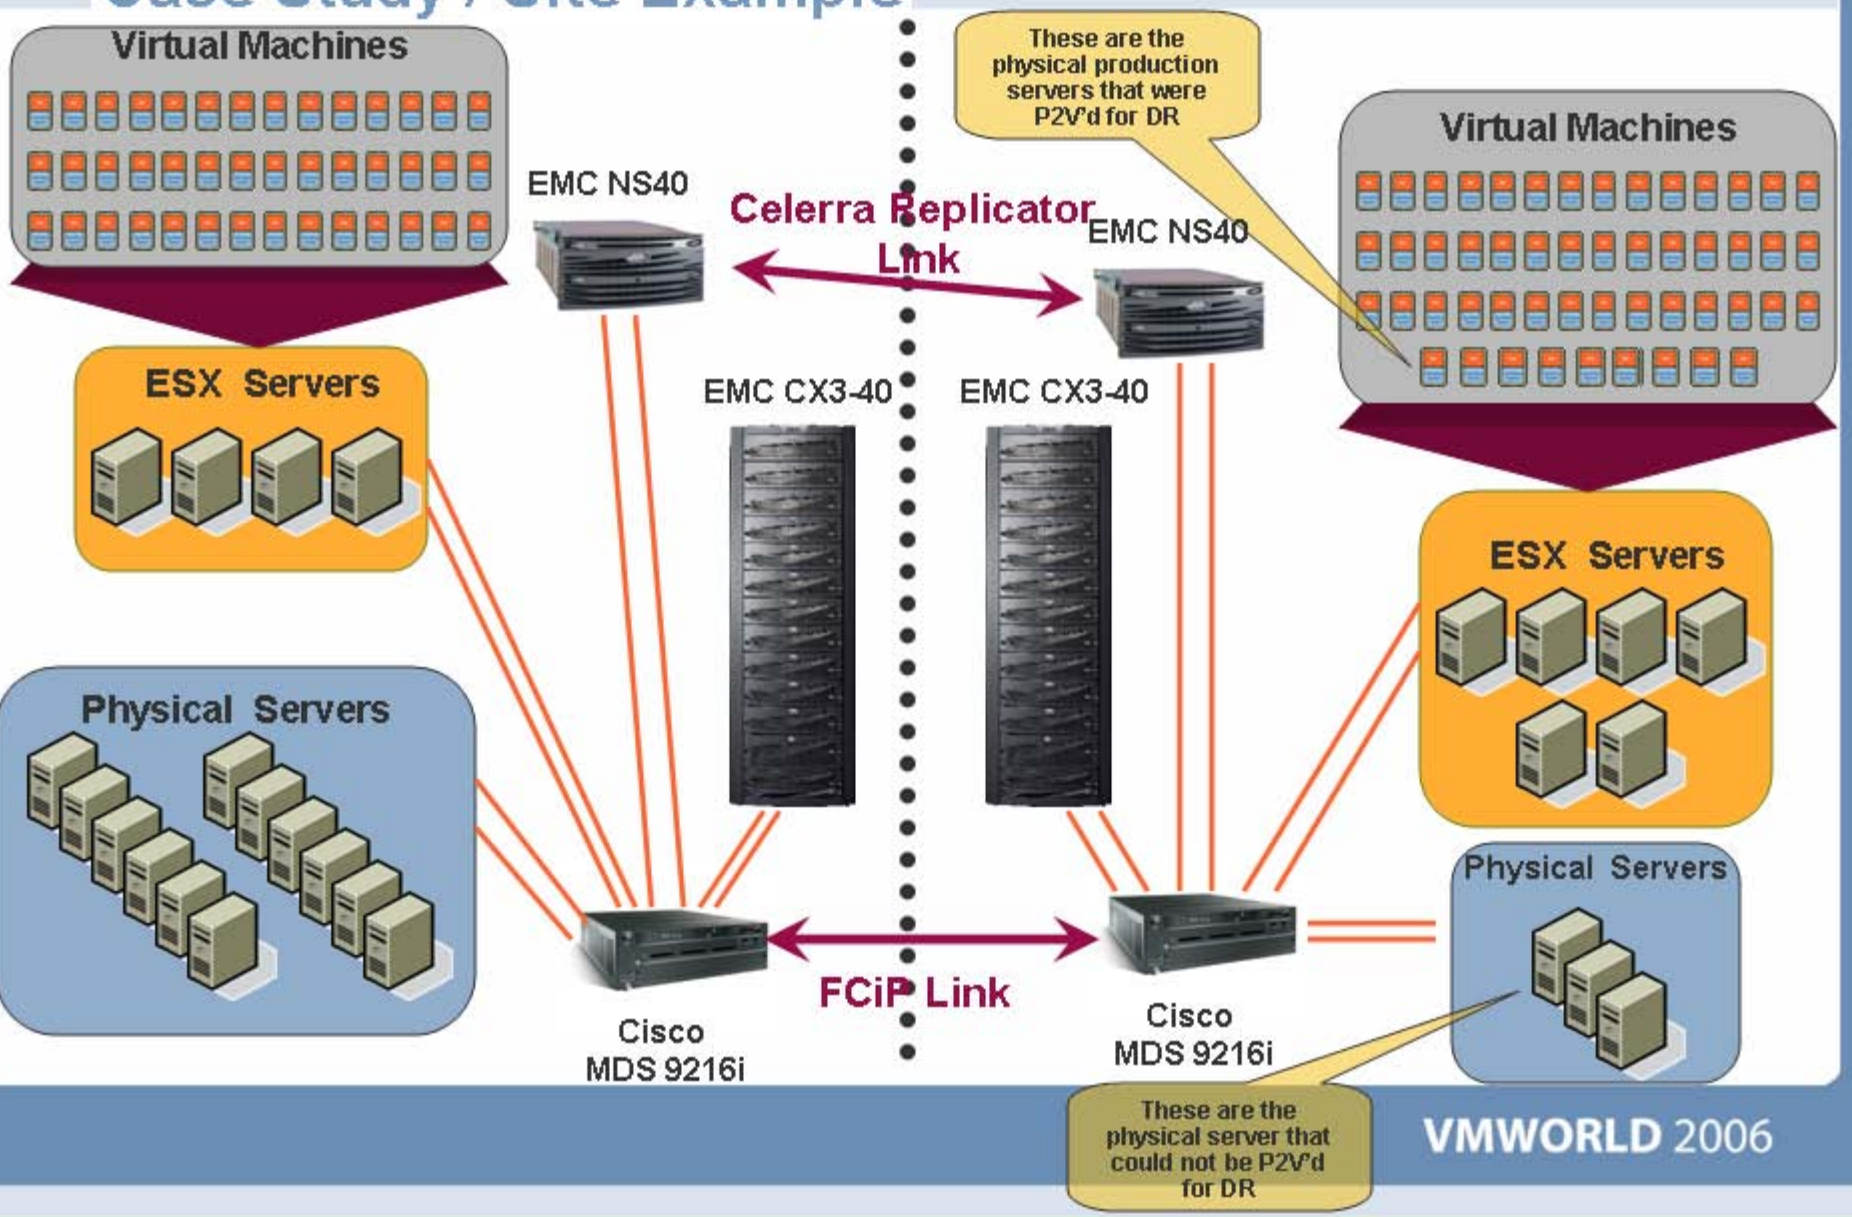

The Approach

- Architect for the present and the future
- Create the complete solution and break into phases
  - Phases are based on requirements, prerequisites and budget
- Tier the environment
  - Servers / VMs / Applications
  - Storage
- Architectural Considerations
  - Performance
  - Backup
  - Disaster Recovery

Architectural Considerations

- Performance
  - Most Virtual Machines will have different performance characteristics
  - Create heterogeneous environment
  - The disk layout should reflect the performance needs
    - What Raid Type
    - What Spindles
    - VMFS vs. RDM / SAN Pass-through
  - How much CPU and memory will be needed per VM
  - How much overhead should be reserved for HA
  - These considerations should also be applied to the DR Site
  - Sizing enough servers, storage/spindles, bandwidth on networks

Architectural Considerations (cont'd)

- Backup and Recovery
  - Host Based, VMDK / VCB, Imaging
  - How will each virtual machine be backed up?
  - Does the backup need to be done at the file level or only VMDK level?
- Disaster Recovery
  - Tiering of Servers, P2V for Physical Servers, Replication Solutions
  - How important is each virtual machine?
  - How will it be replicated?
  - When does it need to be recovered?
  - To what point does it need to be recovered?

Now What?

- **Understand the infrastructure**
  - Storage
    - SAN
    - NAS
    - Switches
  - Backup
    - Standard Host Backup Agents
    - VCB
  - Disaster Recovery
    - Storage Layout
    - Replication Methods
  - Networking
    - Bandwidth
  - Missing Links
    - PlateSpin PowerConvert
    - High IO Servers
    - How to Recover

Underlying Infrastructure:

Case Study: The Storage Area Network

- EMC Storage Arrays
  - CLARiiON CX3-40
  - Celerra NS40
- Cisco MDS SAN Switches
  - MDS 9216i

EMC CLARiiON CX3-40

- Why EMC CX3-40?
- Multiple Disk Types
  - 73, 146, 300, 500
  - 7,200, 10,000 and 15,0000 RPM
  - 2 GB and 4 GB Speeds
- San and iSCSI Connectivity in Single Array
- N+1 Redundancy
- Multiple Raid Types
- Very High Performance
- Supports All Major Operating Systems for Heterogeneous Environment
- Multiple Replication Methods
  - Snapview
  - Mirrorview /A and /S
  - SanCopy

EMC Celerra NS40

- Why NS40
- NFS, CIFS and iSCSI
- Utilizes CLARiiON or Symmetrix for Storage
- Native Replication
- Integrates into Active Directory or NIS
- Simple File Storage
- Consolidate File Servers
- Easy Backup via NDMP
- Simple Web Based GUI

CISCO MDS 92xx and 95xx Series

- Why Cisco?
  - Blade Chassis Architecture
  - Combines iSCSI, Fibre, and FCiP connectivity
  - From 14 Fibre to 526 Fibre Ports
  - From 2 to > 24 iSCSI / FCiP Ports
- This is the SAN Replication Enabler
- Allows multiple SANs to connect via Ethernet Based Connectivity while combining Fibre connectivity in one solution

Storage Design Best Practices

- These Best Practices are derived from many years of storage implementations and VMware Consulting
- Take Advantage of Different Raid Types
- Keep in Mind –
  - Replication Methods
  - Tiers of Servers
  - Backup Methods
  - Performance Impact

CLARiiON Raid Group Layout

- The CX3 can do
  - Raid 5, Raid 3, Raid 1, Raid 1/0, 2 Disk Raid 1/0, Raid 0, Single Disk and Hot Spare
- Do NOT use Raid 0 !!!!!!!!!!
- Use Raid 3 on ATA Drives
- Raid 5 can be anything from a 2+1 to 15+1
  - Sweet spot is anywhere from 4+1 to 8+1
  - BUT 4+1 or 8+1 creates a nice layout and helps to maximize the storage used while maintaining good parity to disk ratio

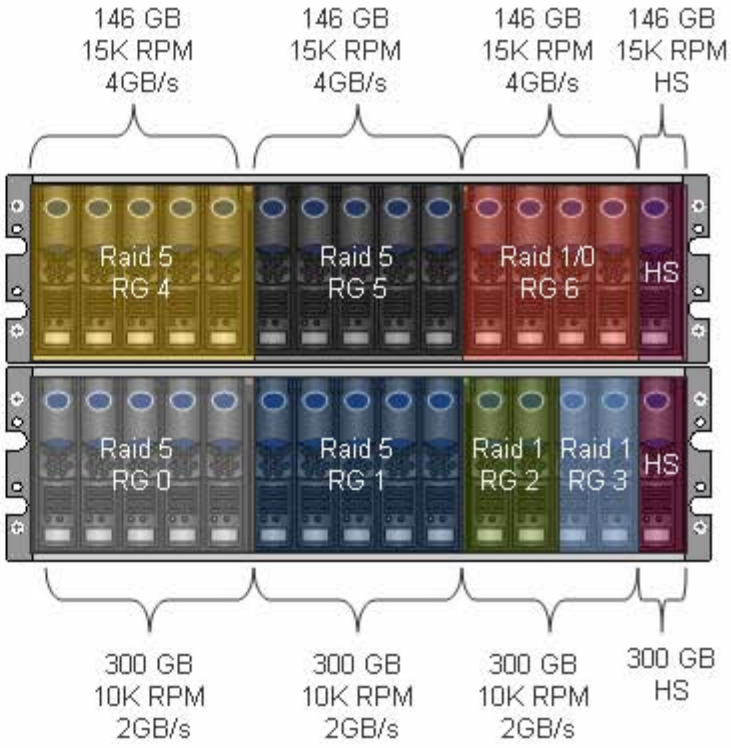

Storage Design Best Practices (continued)

- Create Multiple Types of VMFS Volumes
  - > VMFS_OS_01
    - For all Operating System Drives and Swap Files
      - i.e. C:\ or / (root)
  - > VMFS_PAGE_01
    - For all Page Files or Swap Files
  - > VMFS_DATA_01
    - All Data drives < 50 GB or 150 GB
    - Not High Performance Drives or Cluster Drives
  - > VMFS_OTHER_01
    - Templates, Images, General Storage
- Increment the number when adding a new drive
  - > i.e. VMFS_OS_01 already exists, so add VMFSOS02

VMFS_OS_01 Both Versions

- 20 GB OS Drive
  - 10 VMs per VMFS so
    - 20 GB x 10VMs = 200 GBs
  - Assume 2 GB Virtual Memory per VM so adding Swap space
    - 200 GB + (10 VM * 2 GB) = 220 GB
  - Last but not least Snap Space approx 20 % of total space
    - .20 x 220 GB = 44 GB + 220 GB = 264 GB
    - 264 is uneven, lets make it either **250 or 275**

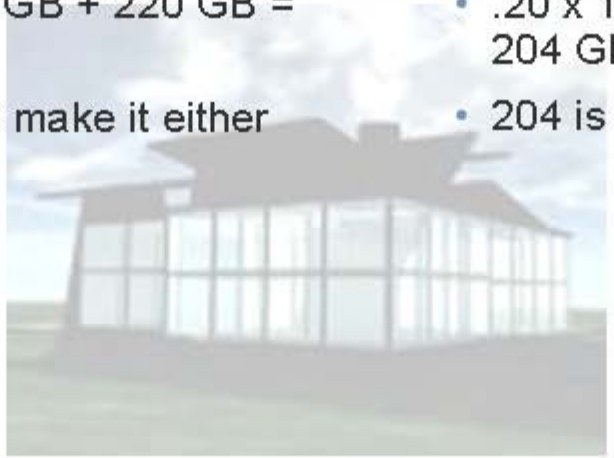

- 15 GB OS Drive
  - 10 VMs per VMFS so
    - 15 GB x 10VMs = 150 GBs
  - Assume 2 GB Virtual Memory per VM so adding Swap space
    - 150 GB + (10 VM * 2 GB) = 170 GB
  - Last but not least Snap Space approx 20 % of total space
    - .20 x 170 GB = 34 GB + 170 GB = 204 GB
    - 204 is uneven, lets make it **200**

VMFS_PAGE_01

- The page LUN is paired with the VMFS_OS LUN
- Make the Page File P:\ and move it there
- So if there are 10 VMs on the VMFS_OS_01 then
  - > 10 VM x 2 GB of Virtual Memory = 20 GB
  - > 20 GB + 10 GB for Page File = 30 GB
- The Page File is the item on the C:\ that fills up the snap space and does not need to be replicated

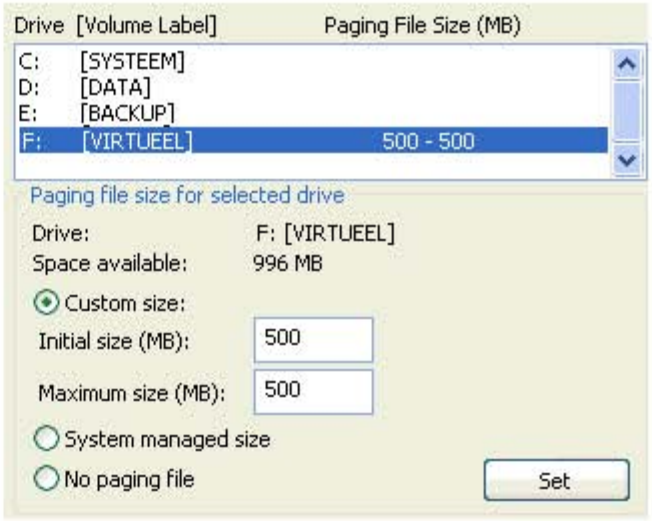

VMFS_DATA_01

- Between 300 GB and 400 GB
- Also for meta files of Raw Disk Mappings aka Mapped SAN LUNs
- For VMDKs larger than ~150 GB, sometimes 50 GB if Very High I/O

- MetaLUNs
  - Should be used ONLY with Extents or RDMs
  - Concatenate not Stripe if expanding with extents
  - Striping can be used with RDMs
  - Different Raid Groups
- LUN Migrate Utility

VMFS_OTHER_01

- For Templates, ISOs and any other type of files
- Can be NAS, iSCSI, or ATA
- Fibre channel might be a waste if the storage environment is tiered

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	[illegible]

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

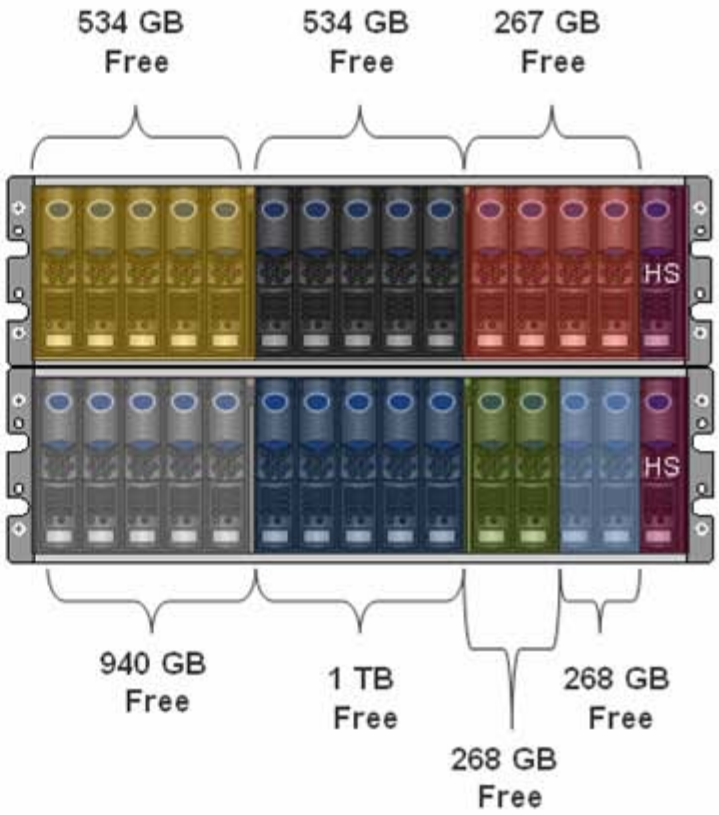

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	[illegible]

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

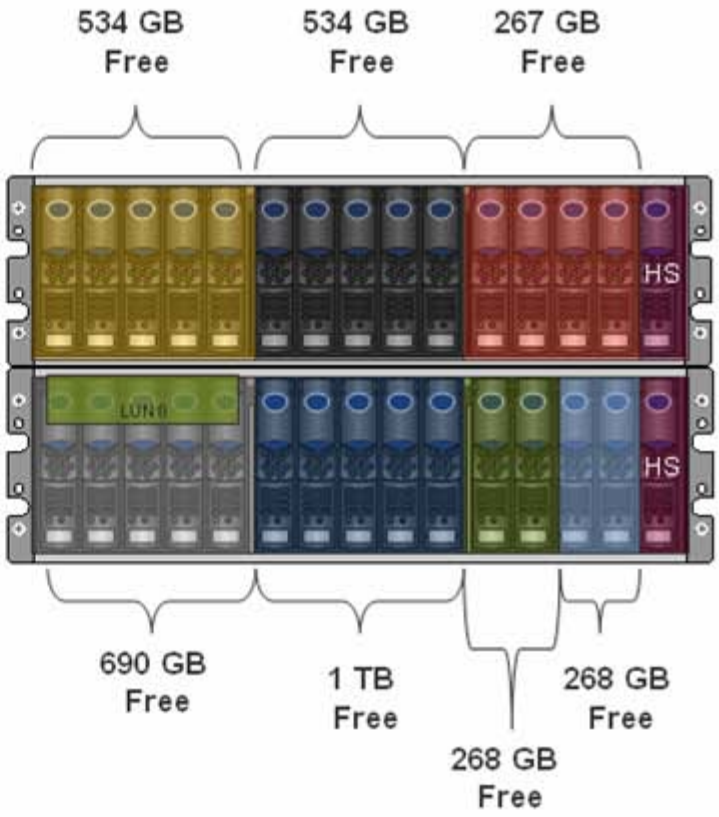

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	VMFSOTHER01 200 GB

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

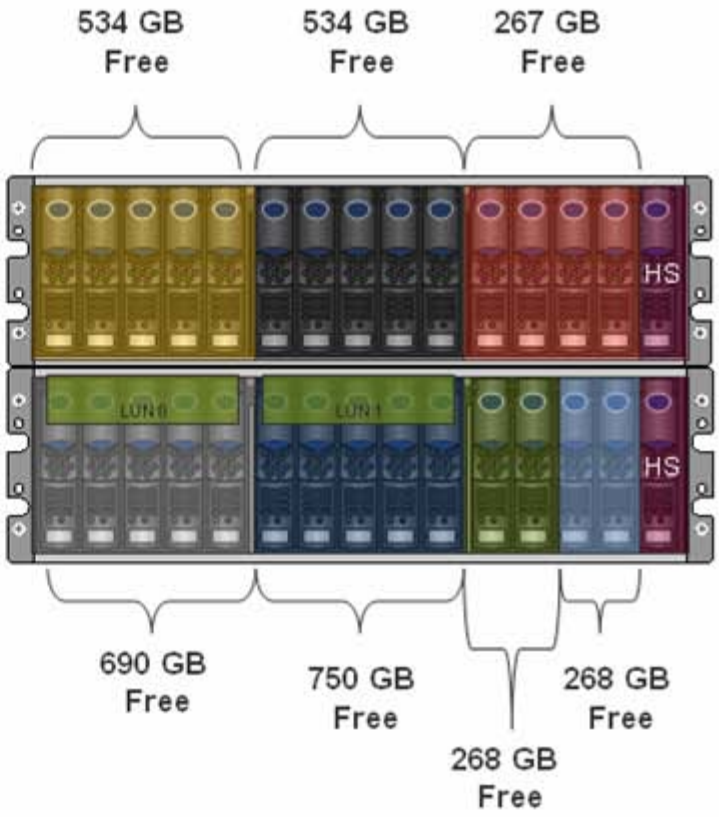

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

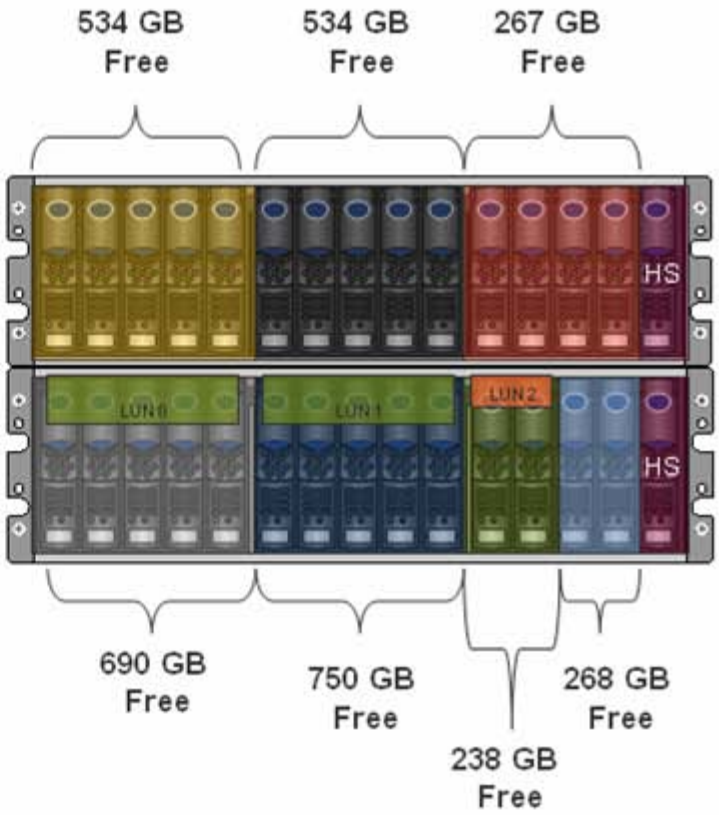

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

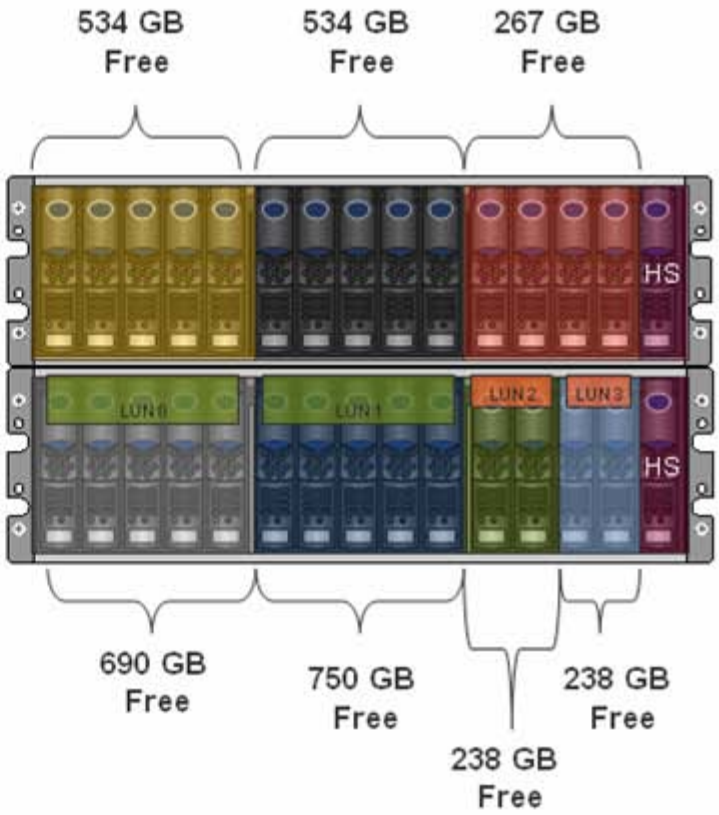

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

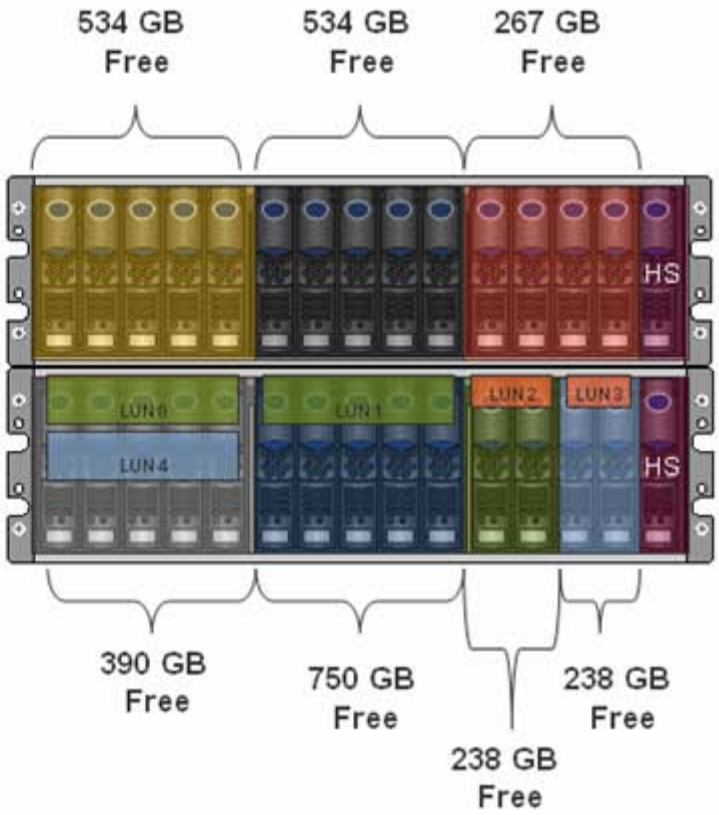

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	[illegible]

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

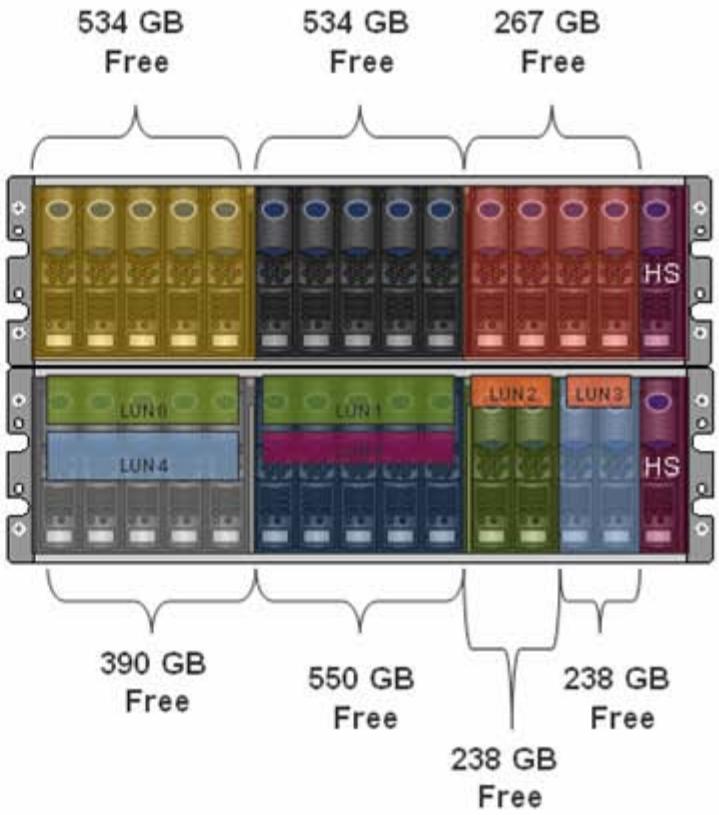

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	VMFSTIER01 200 GB

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

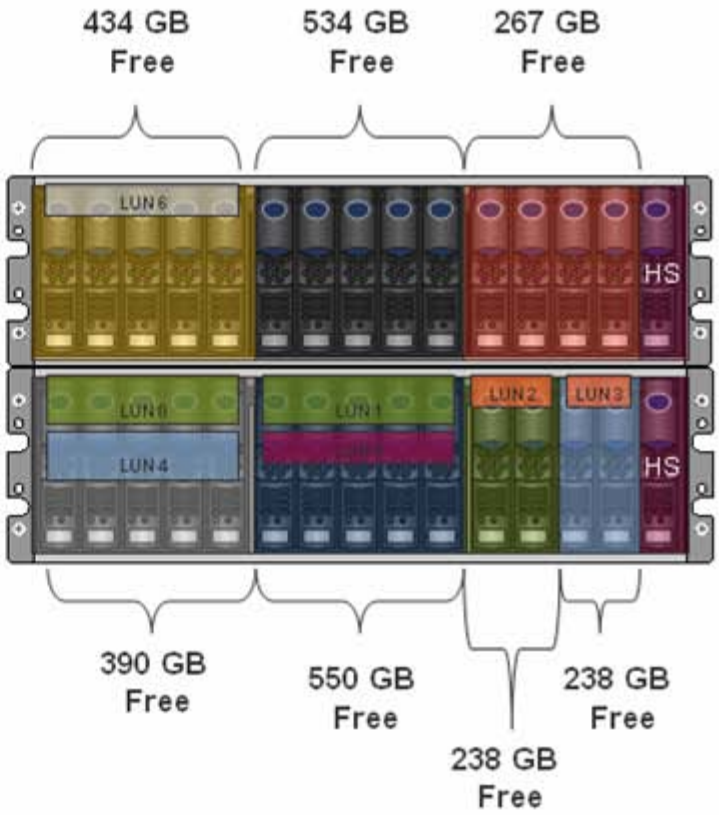

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

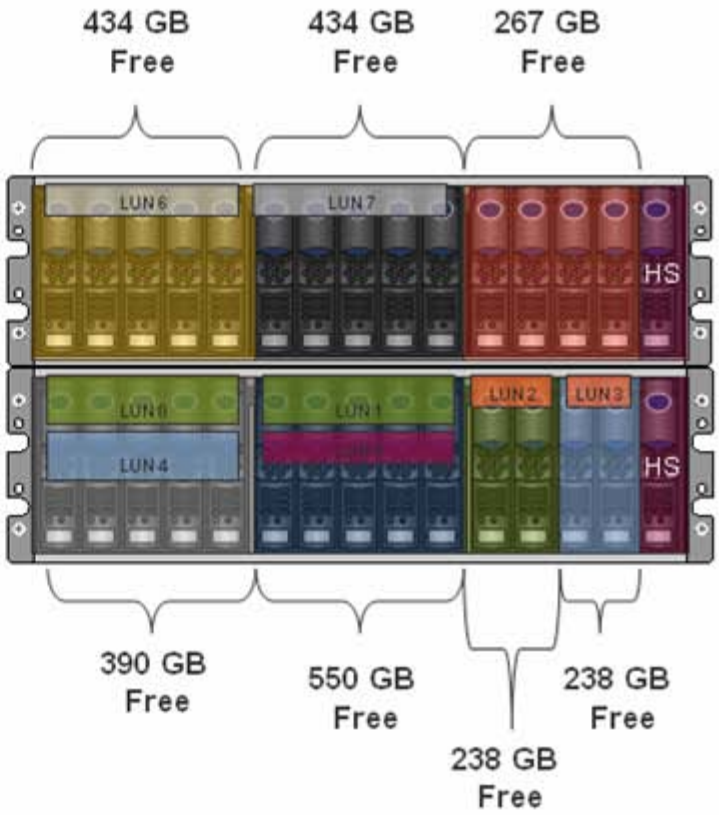

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

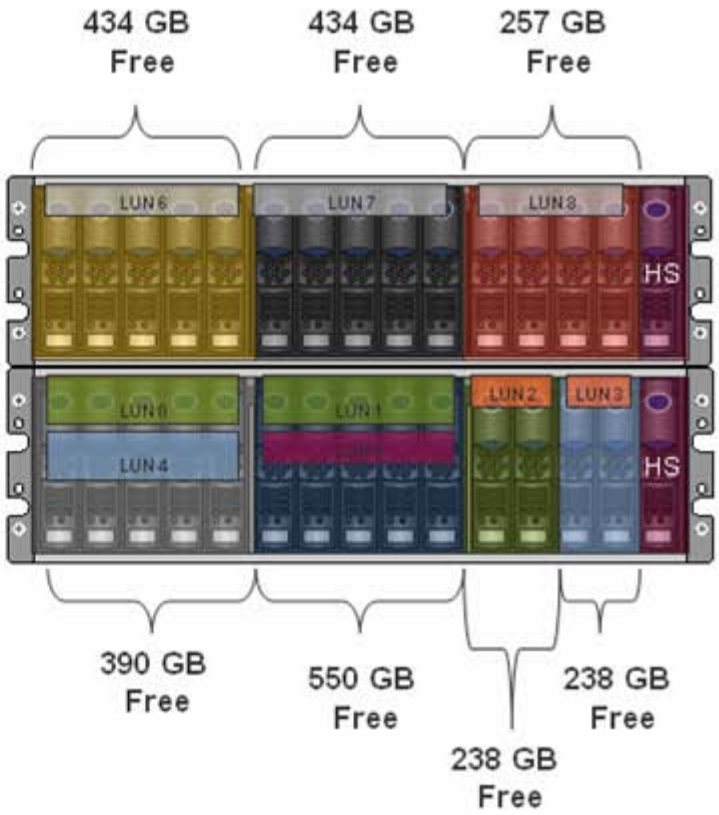

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

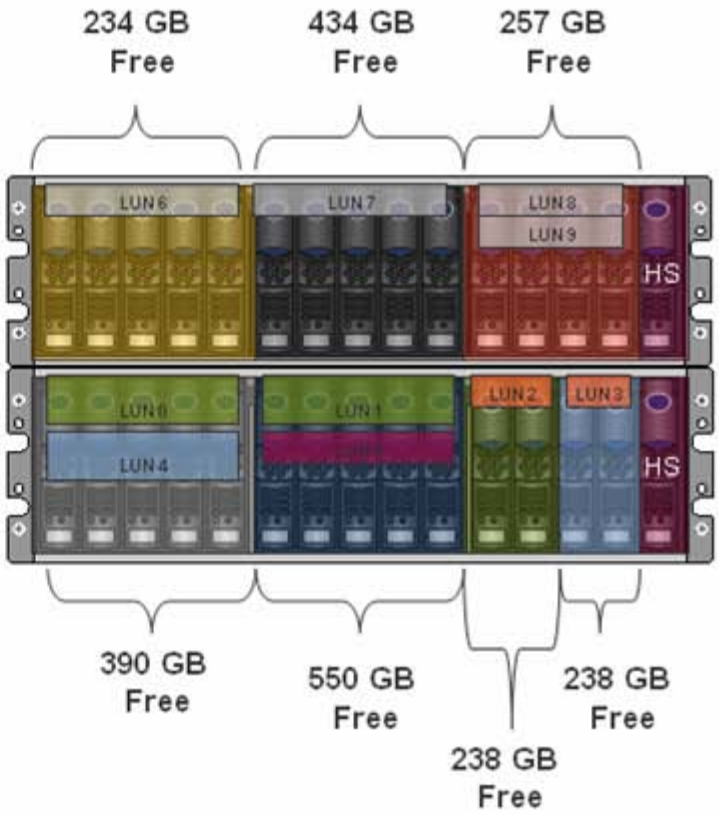

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

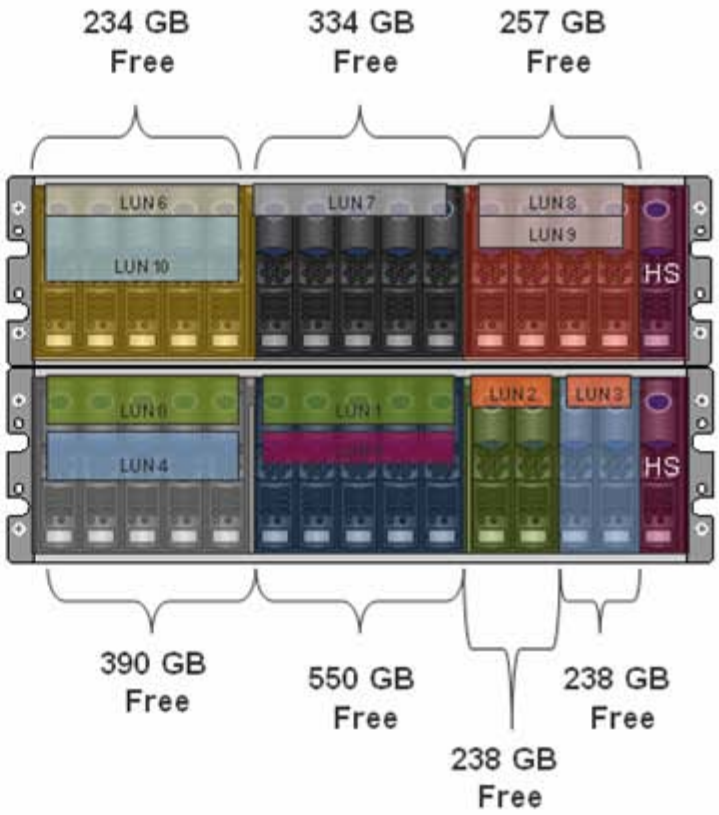

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

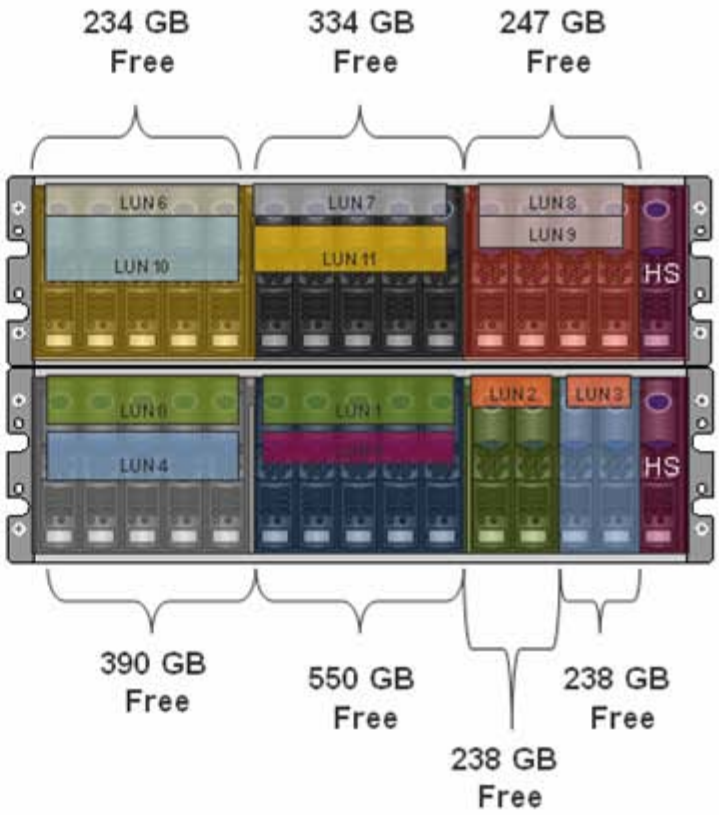

Ok so now what?

- How does the Tiering work with our Storage Best Practices?
  - > Create each VMFS on a different type of disk for different characteristics

VMFSOS01 250 GB	VMFSPAGE01 30 GB	VMFSDATA01 300 GB
VMFSOS02 250 GB	VMFSPAGE02 30 GB	

  - > San Pass Thru LUN / RDM

Exchange DB 1 50 GB	Exchange Log 1 10 GB	SQL DB 1 100 GB
Exchange DB 2 50 GB	Exchange Log 2 10 GB	SQL Log 2 50 GB
	Kronos Data 1 200 GB	

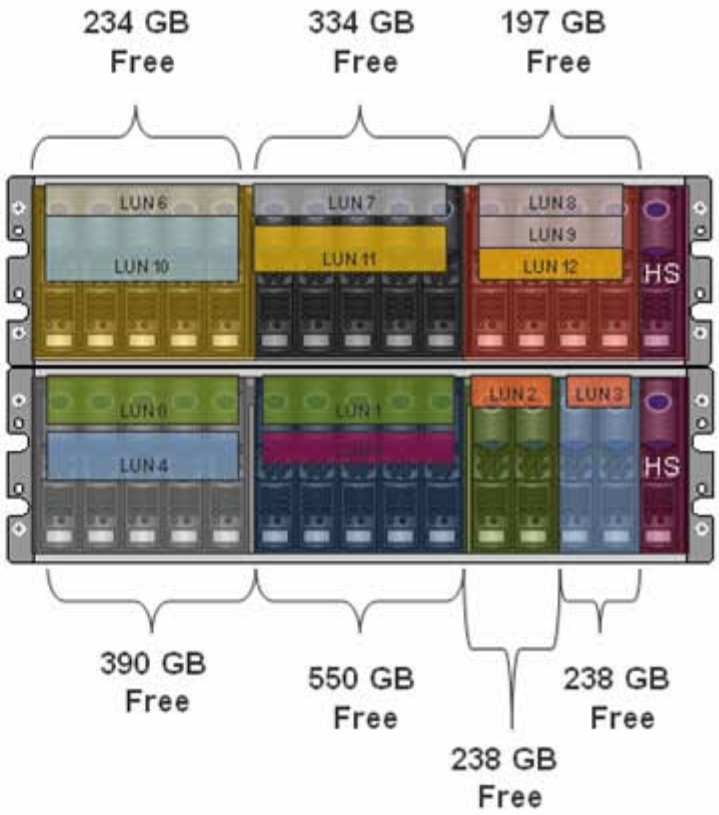

Decisions, Decisions, Decisions

- How do we decide what VMDK goes where?
  - Conform to performance requirements
  - DO NOT be scared of multiple VMFS volumes
    - VI3 makes management VERY simple
  - Always Reserve space for ESX snapshot functionality
- Only use MetaLUNs in extreme cases
- Windows Offset via Diskpart !!!
  - 128k or 64k

Underlying Infrastructure:

Backup Considerations

- What Backup Software is best?
  - Any on the VMware support matrix
- Full vs. Incremental Backups?
  - A mix of full and incremental is ideal
- How often do we do VMDK backups?
  - Depends on the VM
  - Weekly at a minimum
- What's the retention on each VMDK?
  - Based on company policy
- How does VMware Consolidated Backup fit into the backup scheme?
  - VCB will help to perform a SAN backup of the VMDKs
- When do we use VCBs file level backup?
  - If there is a low rate of change of data on the server *i.e. Web Servers*

Standard Network Backups

- Install Host Based Agent in each VM
- Utilize the network to backup the VM
- Perfect for Databases and high I/O applications
  - Use the application specific agents for packages such as SQL, Oracle, Exchange, etc…
- Backup to Virtual Tape Library (i.e. CLARiiON Disk Library) increases backup and restore performance
- DO NOT attach tape drives via Fibre to VMs NOT SUPPORTED
- Restore Process – SIMPLE
  - Standard restore via Backup Agent
  - Good for data restore not VM restore

VMware Consolidated Backup

- Offloads the backups to the SAN
- Mounts VMDK files on a physical windows server
  - Utilizes Snapshots to quiesce the server
- Mounts File Structure of VMDK files of Windows Servers on the physical windows server
  - Allow for Full and Incremental backups
  - To perform Incremental, perform backups by modify date not archive bit
- IMPORTANT: Do Not Present VMFS Volumes Before Installing VCB Agent and Disabling Automount in Windows
- Restore Process – NOT SIMPLE:
  - If small file is to be restored, restore to backup server and perform network copy
  - If whole application restore is needed, install backup agent and restore application
  - If VMDK restore is needed, use vcbRestore to do image based restore

Disaster Recovery

SnapView and SanCopy

- Provides up to eight replicas per source
  - Clones—Full-image copies of source LUN
  - Snapshots—Logical point-in-time views of source LUN
- Useful for
  - Backup and Restore
  - Testing
  - Upgrade Protection
  - Disaster Recovery
- Instant Restore

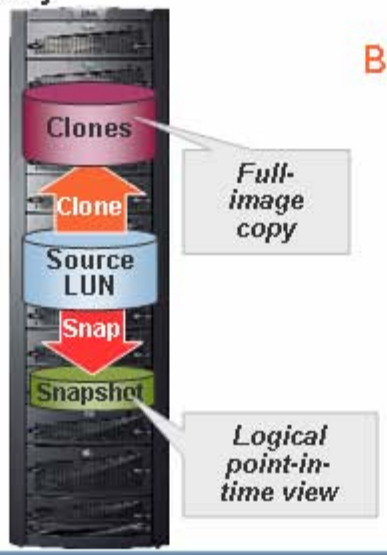

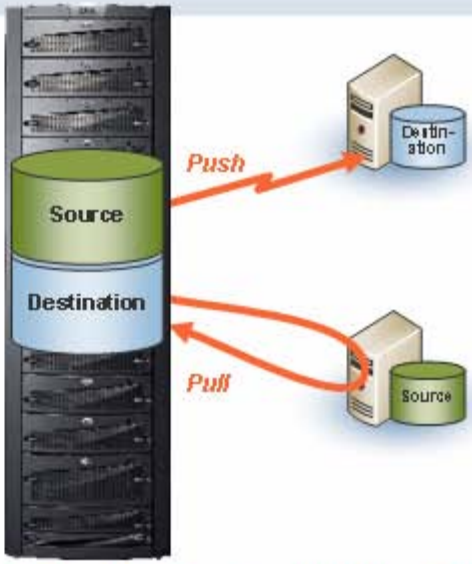

Bi-Directional Copy Capabilities via "Pull" and "Push" Feature

- CLARiiON arrays capable of hosting SAN Copy:
  - CX3-80, CX3-40, CX3-20, CX700, CX500, CX600, and CX400
- Full or incremental copies of data residing on a SAN Copy-hosted array can be "pushed"
  - To any EMC or supported third-party array
- Full copies of data residing on any supported array can be "pulled" to a SAN Copy-hosted array

CLARiiON MirrorView

MirrorView/Synchronous

RPO: Zero seconds

- Both images identical
- Limited distance
- High network bandwidth

MirrorView/Asynchronous

RPO: 30 minutes to hours

- Target updated periodically
- Unlimited distance
- Restartable copy on secondary if session fails
- Optimized for low network bandwidth (consumes 100 Mb/s maximum)

Disaster Recovery Considerations

- VMware ESX and Virtual Center are key components in optimizing a DR Strategy
- Quicker Recovery Time Objective
- Simplify the DR plan
  - You do not know who might need to implement the DR Plan

- Planning for DR, consider:
  - Service Level Agreements
    - With vendors, customers, internal departments
  - Document the Recovery Process for the DR Site
  - Document the Process for returning to the Production Site
  - Perform regular DR Tests
  - Creating DR Tiers based on RTO

Disaster Recovery Tiering

- While Tiering the environment, there are usually 4 tiers:
  - 1 – *aka* Oh @#!%, WE ARE DOWN!!!!!
    - 4 Hour Recovery
    - Email, Blackberry, Point of Sale System, Oracle, etc…
    - Obviously Domain, DHCP, DNS
  - 2 – *aka* Just 20 hours to go….
    - 24 Hour Recovery
    - Web Servers, ERP, CRM, SQL
  - 3 – *aka* Hey Charlie, is there a problem?
    - 48 Hour – 1 Week Recovery
    - IT Software(Cisco Works, ECC, OpenView, etc…)
  - 4 – *aka* Do we really need these applications?!?!?!
    - > 2 Weeks Recovery
    - Test Environment, Lab Servers

Sample Tiering

- **Tier 1**
  - Domain / DHCP / DNS — ICI-DC01
  - Exchange 2003 — ICI-EXCH01
  - Blackberry — ICI-BB
  - Citrix — ICI-CITRIX01
  - Image Server — ICI-PSPIN
- **Tier 2**
  - Financial SQL Server — ICI-FINSQL
  - Kronos Server — ICI-KRONOS
  - SAS Servers — ICI-SAS01 thru ICI-SAS04
- **Tier 3**
  - Web Servers — ICI-WEB01 thru ICI-WEB03
- **Tier 4**
  - All Development Machines — ICIDEV-Test01 thru ICIDEV-Test06

DR's Affect on Storage Layout

- How does our replication scheme effect this layout?
  - Group the VMDKs by replication Tier and schedule
  - Smaller VMFS partitions allows a more granular replication and recovery
  - In Case of a failure, lesser number of VMs are affected
- Create source and target LUNs from same type spindles
  - Fibre for Fibre and ATA for ATA
- When mirroring with MirrorView type software:
  - Remember the ability to fracture the mirror
  - Consistency groups are great for RDM of the same server
- When replicating with SanCopy type software:
  - SanCopy sessions can be performed from Clones for the production volume
  - Incremental sessions can save a significant amount of time
  - These sessions can be scripted with the navicli from within a VM

The Recovery Process via SAN Replication

- Power Up DR ESX Servers
- Place Mirrored LUNs in Storage Groups
- Rescan San for LUNs
- Now the tricky part
  - Option 1: Open the VMFS_OS_01 and register the VMX for each VM
  - Option 2: Create a script that will traverse the directory and register the VMs for you
- The environment is running ☺

Underlying Infrastructure:

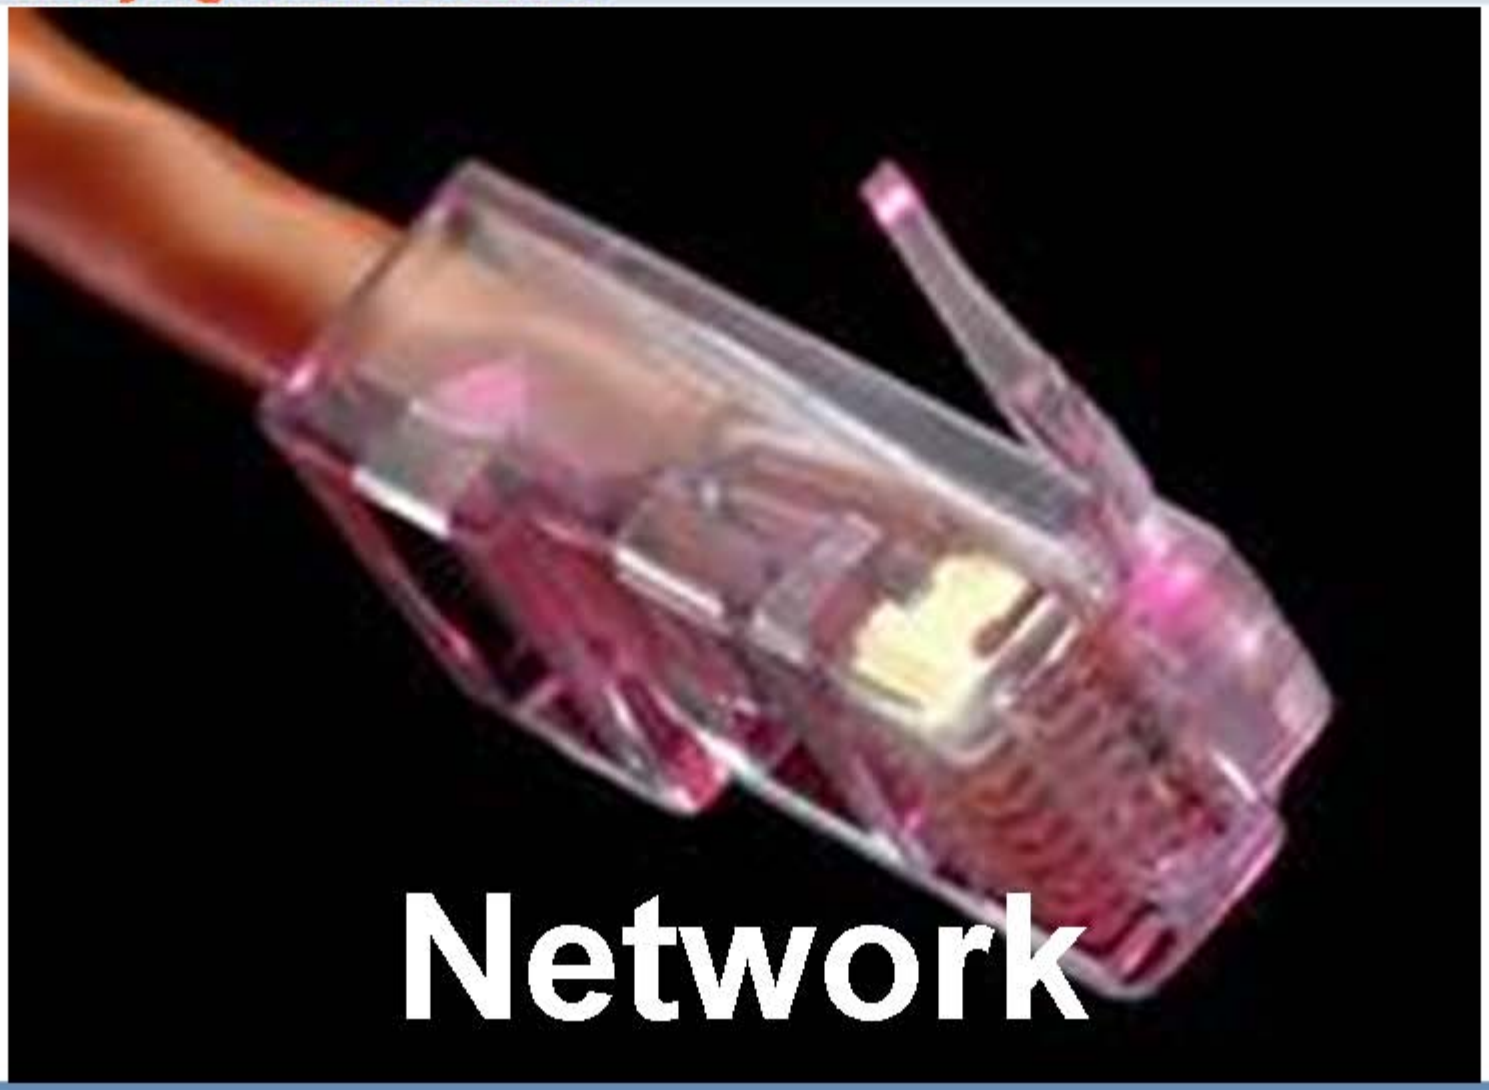

Bandwidth for Backup

- Without VCB
  - Create multi-NIC vSwitches
  - Ensure a Gigabit Network is available
  - Create a separate Port Group for Backups (if possible)
- With VCB
  - Bandwidth is not a factor, SAN does all the work

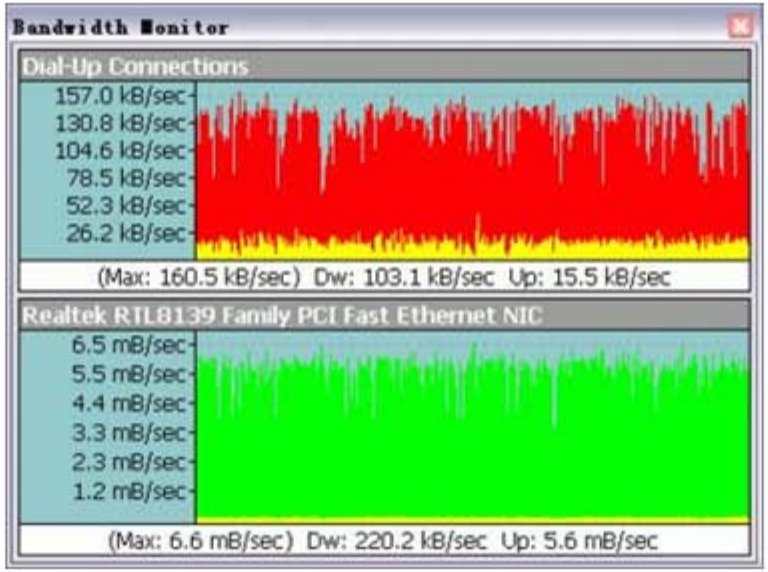

Bandwidth for Replication

- How much data is replicated?
- How often is the data replicated?
- Is the replication data sharing bandwidth?
  - If so is the replication and production data on the same schedule or differing?
- SO

  Sync Rate = 2 Hours

  Rate of Change = 1 Gigabyte per Hour

  Every 2 Hours we need to transfer 2 GBs of data

  2 Gigabytes = 2048 Megabytes = 16384 Megabits

T1 = 16384 Mb / 1.54 Mbps = 10638 seconds / 3600 sph = ~ 2 h 57 min

T3 = 16384 Mb / 44.736 Mbps = 366.24 seconds / 3600 sph = ~ 6 min

Dual T1s = 2 h 57 min / 2 = 1 h 27 min ← Sufficient

Underlying Infrastructure:

What's Missing

Almost There…

- So far we have taken care of a general replication of all the C:\ VMDKs and most of the data volumes, BUT
  - What about all of those Physical servers in production?
  - What about those high I/O servers?
- To simplify our recovery in DR we will want to look towards using ESX as much as possible
  - How about P2V the physical servers for DR so we can recover them in ESX
- And those high I/O servers that have issues with quiescing their I/O
  - Use a host based replication restart solution

Physical Servers – PlateSpin PowerConvert

- PlateSpin PowerConvert solves the physical server issue for Disaster Recovery
- Any physical server that is deemed suitable can be captured nightly to an image or virtual machine
  - These images or VMDKs can then be stored on a LUN on the CLARiiON
  - These LUNs will be Sancopied or Mirrored over to DR on a regular basis
- Using this method we can setup a recovery scenario that looks like this

Physical Servers – PlateSpin PowerConvert

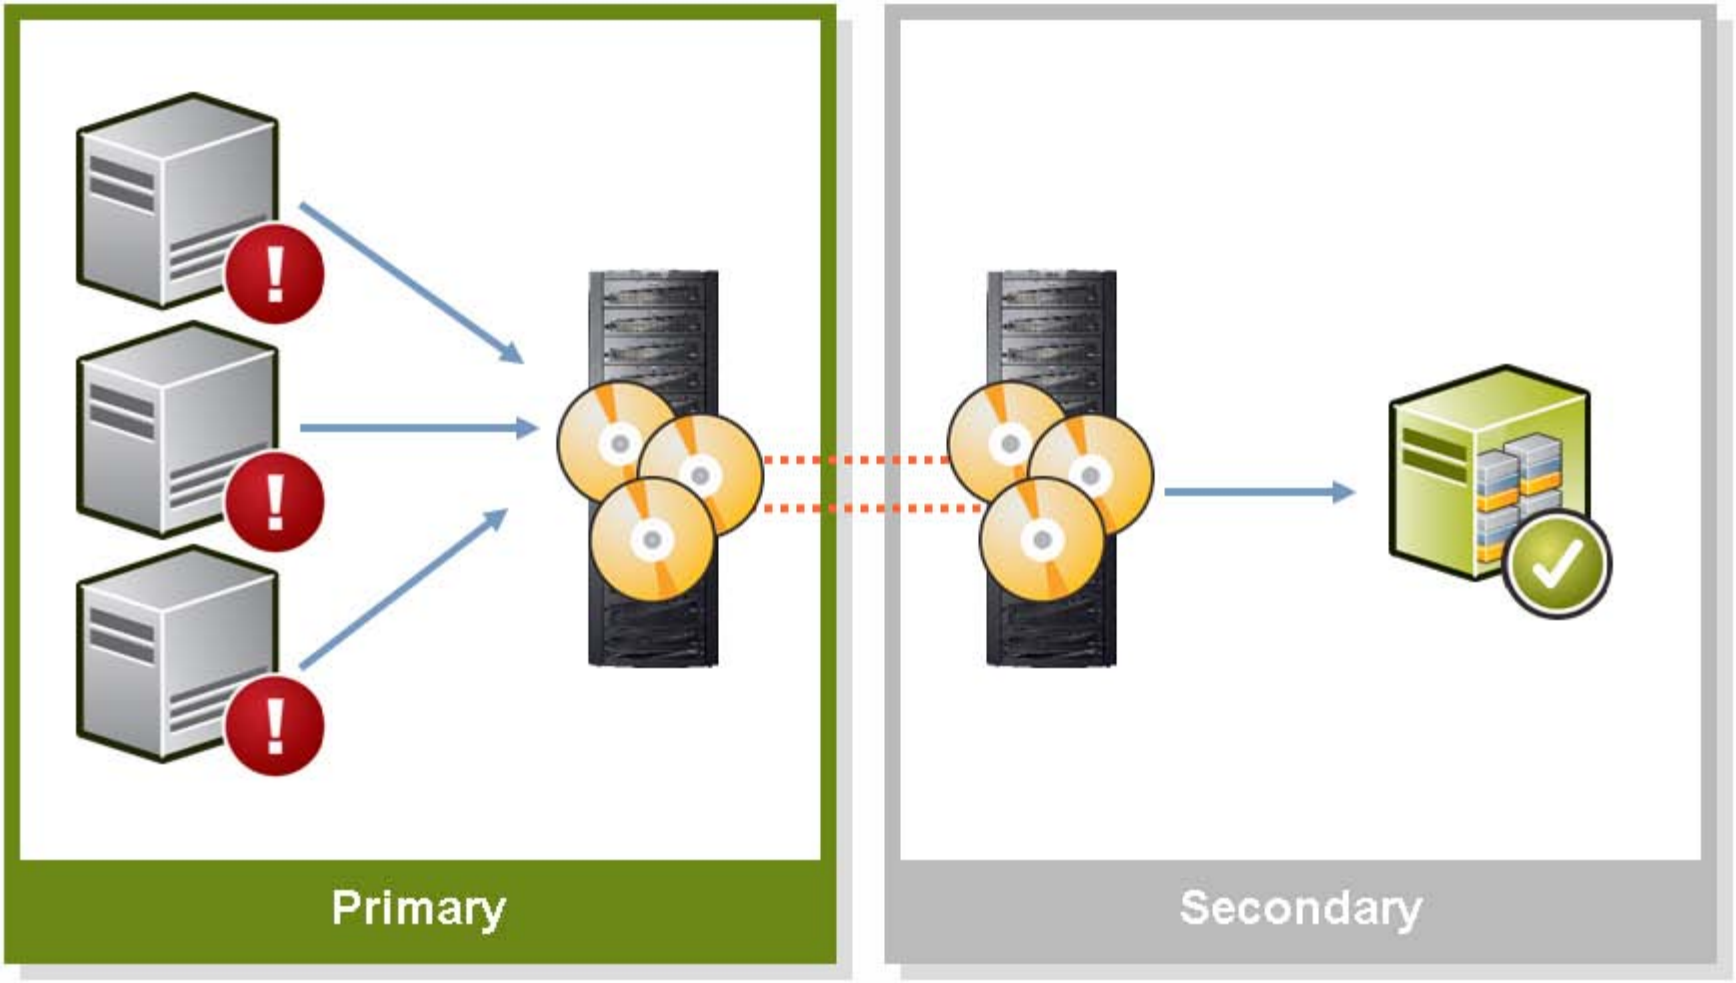

PowerConvert P2V for Recovery

- Images are great, but …
  - Image deployment increase RTO
- So why not …
  - Schedule whole system backups directly into a VM
- PlateSpin PowerConvert 6.5 with Live Incremental P2V
  - Easily test the recovery process
  - Use the same recovery process for physical and virtual
  - V2P is available to restore back to replacement physical hardware

Physical Servers – PlateSpin PowerConvert

Backup

Replicate

Test

Failure

Recover

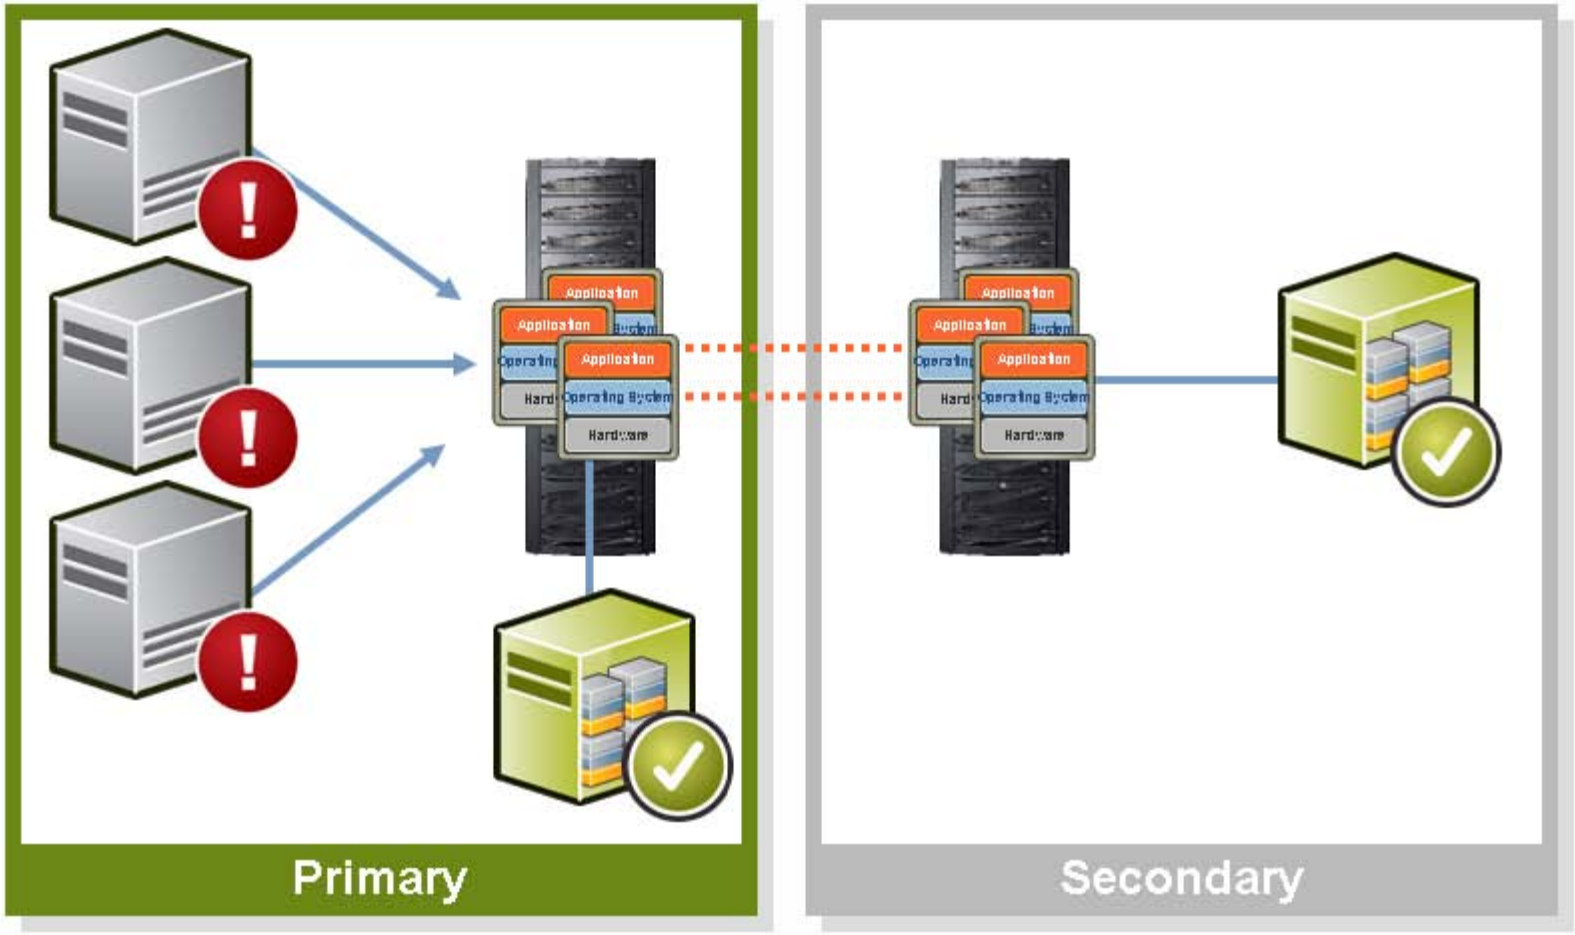

The Recovery Process via SAN and PlateSpin

- Power UP DR ESX Servers
- Place Mirrored LUNs in Storage Groups
- Rescan San for LUNs
- Now the tricky part
  - Option 1: Replicate SQL VC Database
  - Option 2: Open the VMFS_OS_01 and register the VMX for each VM
  - Option 3: Create a script that will traverse the directory and register the VMs for you
- The environment is running ☺

High I/O Quiescable Servers

- Now, last and definitely NOT least… the High I/O quiesceable servers
- These solutions can be used to make Exchange, SQL or other applications geographically clusterable
- Enhances our DR process by alleviating the concern for the quiesce issue with databases
- Fills the gap with the VMs that are affected by a small amount of data loss during failover
- Solutions
  - EMC Replistor
  - NSI DoubleTake
  - NeverFail

The Recovery Process via SAN, Replistor and PlateSpin

- Power UP DR ESX Servers
- Place Mirrored LUNs in Storage Groups
- Rescan San for LUNs
- Now the tricky part
  - Option 1: Replicate SQL VC Database
  - Option 2: Open the VMFS_OS_01 and register the VMX for each VM
  - Option 3: Create a script that will traverse the directory and register the VMs for you
- Replistor failover requires an existing Windows domain infrastructure to be running
- The environment is running ☺

Side Note

For a more comprehensive discussion on PlateSpin PowerConvert please visit the PlateSpin Booth

Questions

My Contact Information

bmaltz@iciamerica.com
http://www.iciamerica.com

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
http://www.vmware.com/vmtn/vmworld/sessions/

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

VMWORLD 2006