

vmworld2005

virtualize^{now}

las vegas • october 18-20, 2005

Derived from VMware Education Services course '*Virtual Infrastructure*'
See <http://www.vmware.com/services/education.html> for course schedule

Higher Availability with VMware Software: Implementation and Best Practices SLN349

Craig Williams
VMware Education




**This presentation may contain VMware
confidential information.**

Copyright © 2005 VMware, Inc. All rights reserved. All other
marks and names mentioned herein may be trademarks of their respective
companies.

Networking

Virtual Switches With ≥ 2 Outbound Adapters ('Bonds')

- Each connects virtual machines to an 802.3ad NIC team
 - Up to 32 virtual machines can connect to one bond
 - Zero collisions internal traffic
 - Each virtual NIC will have its own MAC address
 - Improved network performance: network traffic load balancing
 - Redundant NIC operation
 - Bandwidth per virtual machine can be controlled with traffic shaping



Creating a NIC Team

- Up to 10 physical NICs can be placed in one team
- Up to 10 teams can be created
- A physical NIC that is used in a team can no longer be assigned stand-alone

This server has an unassigned adapter

Add it to the virtual switch

The screenshot shows the VMware ESX Server 2.1.0 Network Connections interface. The 'Physical Adapters' tab is selected, showing a table of network adapters. One adapter is listed as 'Unassigned' with a status of 'Not Assigned'. An orange arrow points from the text 'This server has an unassigned adapter' to this row. Another orange arrow points from the text 'Add it to the virtual switch' to the 'Add...' link in the 'Virtual Switches' column of the same row. The interface also shows two virtual switches: 'Engineering LAN' and 'Internal-only Switch', each with its own properties and an 'Add...' link for virtual switches.

Outbound Adapters	1 Assigned (2 Total)
Virtual Switches	2 Add...

Virtual Switch: Engineering LAN	
Properties Edit...	
Outbound Adapter 0	100 Mbps, full duplex
Total Bandwidth	100 Mbps, full duplex
Port Groups	0 Add...

Virtual Switch: Internal-only Switch	
Properties Edit...	
No outbound adapters. Traffic will be routed locally.	
Port Groups	0 Add...

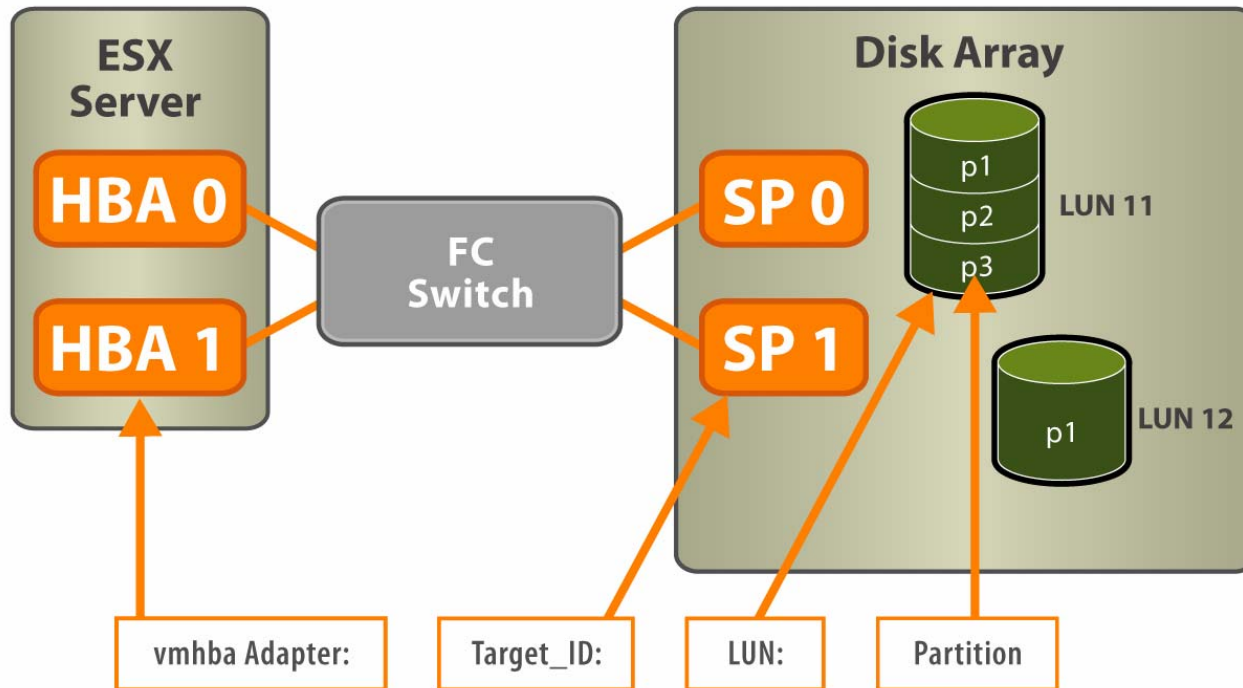
Load Balancing and Switch Failover

- out-mac**
- Default: Each virtual machine's outbound traffic is mapped to a specific physical NIC based on the virtual machine's MAC address
 - Low overhead
 - Compatible with all switches
 - May not spread traffic out evenly
- out-ip**
- A NIC for each outbound packet is chosen based on its source and destination IP addresses
 - Better distribution of traffic
 - Slightly higher CPU overhead
 - Not compatible with all switches; requires 802.3ad link aggregation support
- Standby**
- The bond will use one "home" NIC exclusively until that NIC fails, then fail over to another
 - Useful for falling back on a backup network path

Storage

Addressing SAN LUNs

- VMkernel addresses disk partitions as follows:



Examples:

LUN Addresses

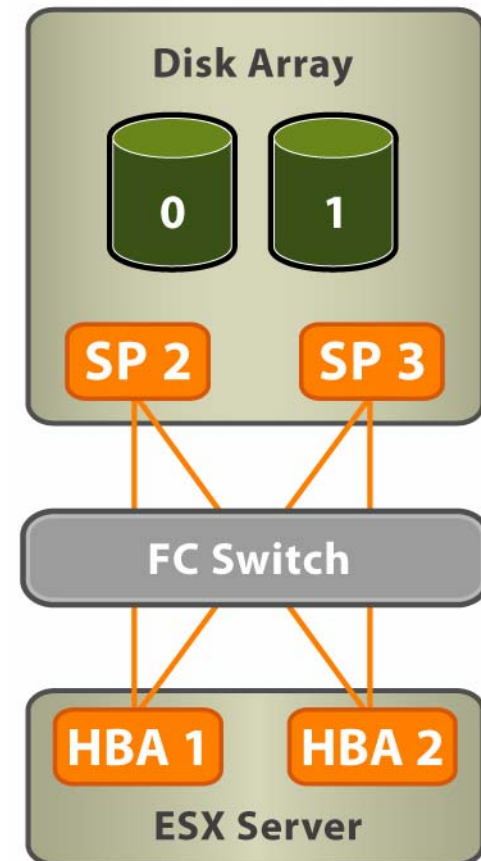
vmhba0:0:11
vmhba1:1:12

Partition Addresses


vmhba0:0:11:3
vmhba1:1:12:1

Multipathing

- Multipathing allows continued access to
- SAN LUN in the event of hardware failure
 - Failover occurs automatically, with a
 - configurable delay
- Exactly one path is active (in use) to any
- LUN at any time
 - Can enable/disable individual failover paths
 - by changing their status
- Two failover policies exist:
 - MRU (Most Recently Used, default)
 - Use last active path, no auto-fail backs
 - Fixed (Preferred path)
 - Revert back to preferred path when available
- Preferred and active paths may be set for each LUN



SAN Failover Path Editing



Active, preferred path to LUN 0

Set new active, preferred path to LUN 1

esx112: Storage Management - Microsoft Internet Explorer

VMware ESX Server 2.1.0 build-7653 | root@esx112.vmware.com

Disks and LUNs | **Failover Paths** | Adapter Bindings | Refresh | Rescan SAN | Help | Close

Storage Management: Failover Paths
Review the current state of paths from your system to SAN LUNs.

SAN LUN vmhba1:2:0 (4 Paths, Policy: mru)

Adapter : Target : LUN	SAN Target	Edit...
◆ vmhba1:2:0	50:06:01:68:10:20:AD:87	
◇ vmhba1:3:0	50:06:01:60:10:20:AD:87	
◇ vmhba2:2:0	50:06:01:68:10:20:AD:87	
◇ vmhba2:3:0	50:06:01:60:10:20:AD:87	

SAN LUN vmhba1:2:1 (4 Paths, Policy: mru)

Adapter : Target : LUN	SAN Target	Edit...
◆ vmhba1:2:1	50:06:01:68:10:20:AD:87	
◇ vmhba1:3:1	50:06:01:60:10:20:AD:87	
◇ vmhba2:2:1	50:06:01:68:10:20:AD:87	
◇ vmhba2:3:1	50:06:01:60:10:20:AD:87	

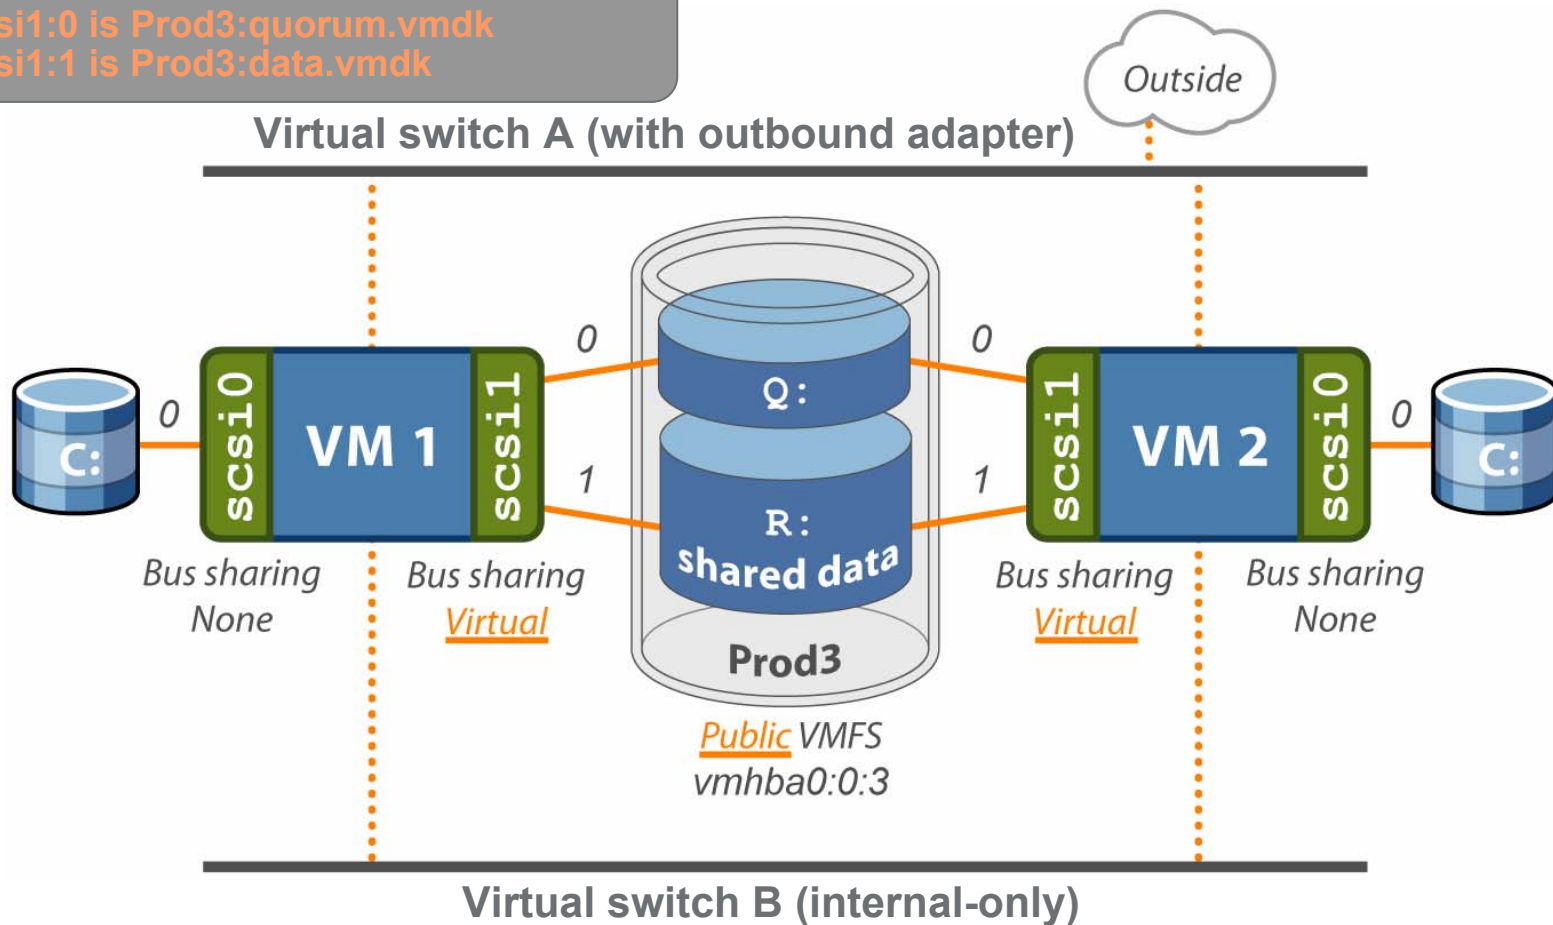
Define paths and failover policy

Four paths to each LUN: two HBAs, two SPs

High Availability Clustering

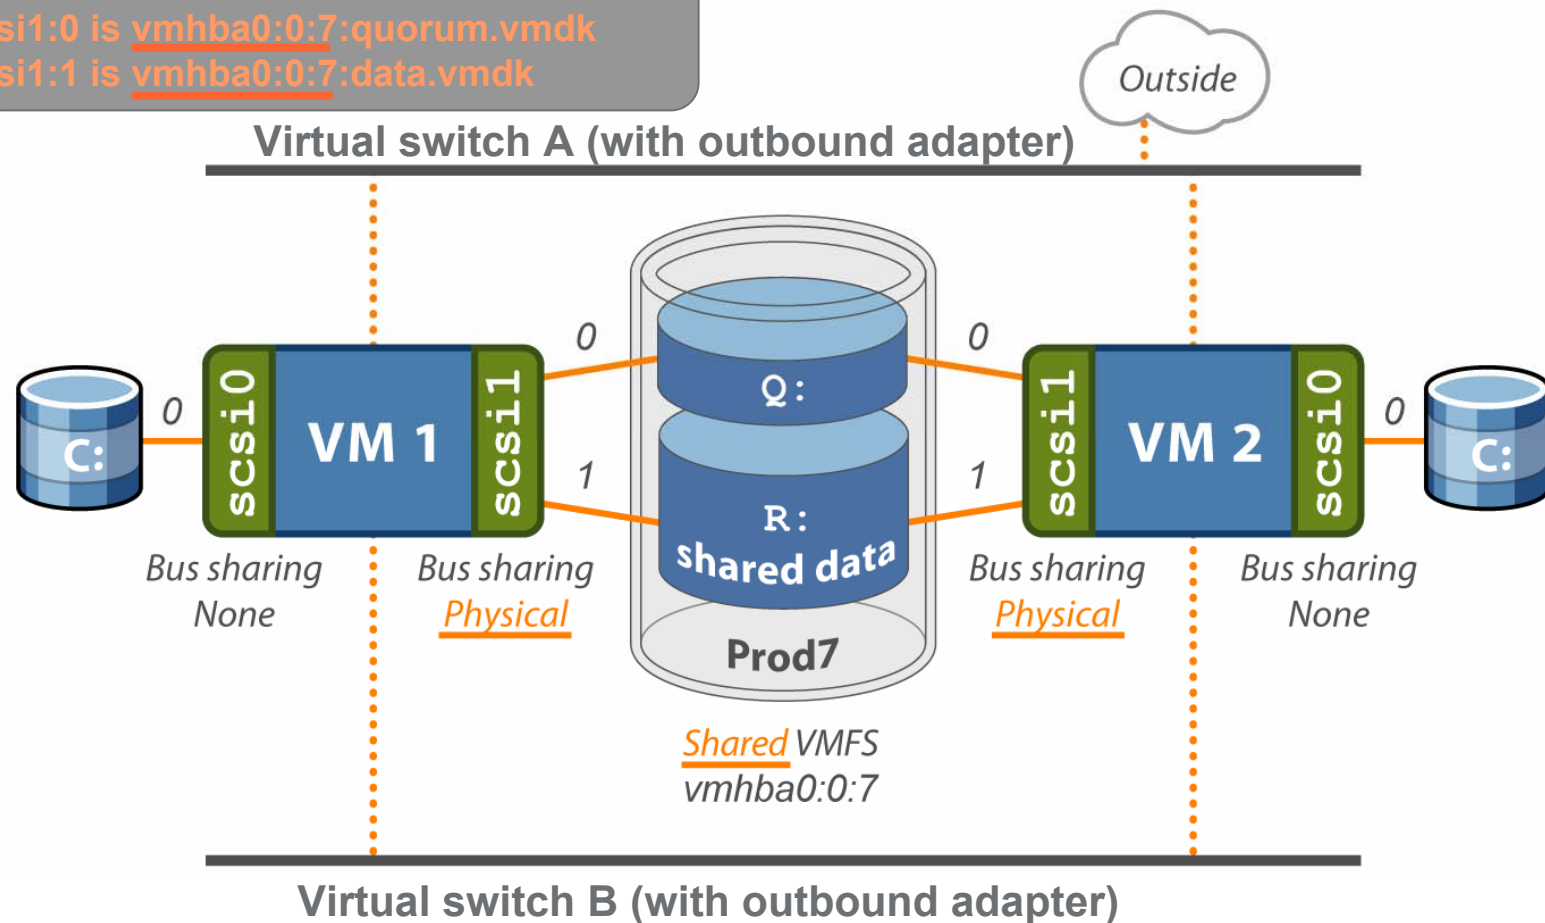
Implementing Cluster-In-A-Box

On both virtual machines:
scsi1:0 is Prod3:quorum.vmdk
scsi1:1 is Prod3:data.vmdk

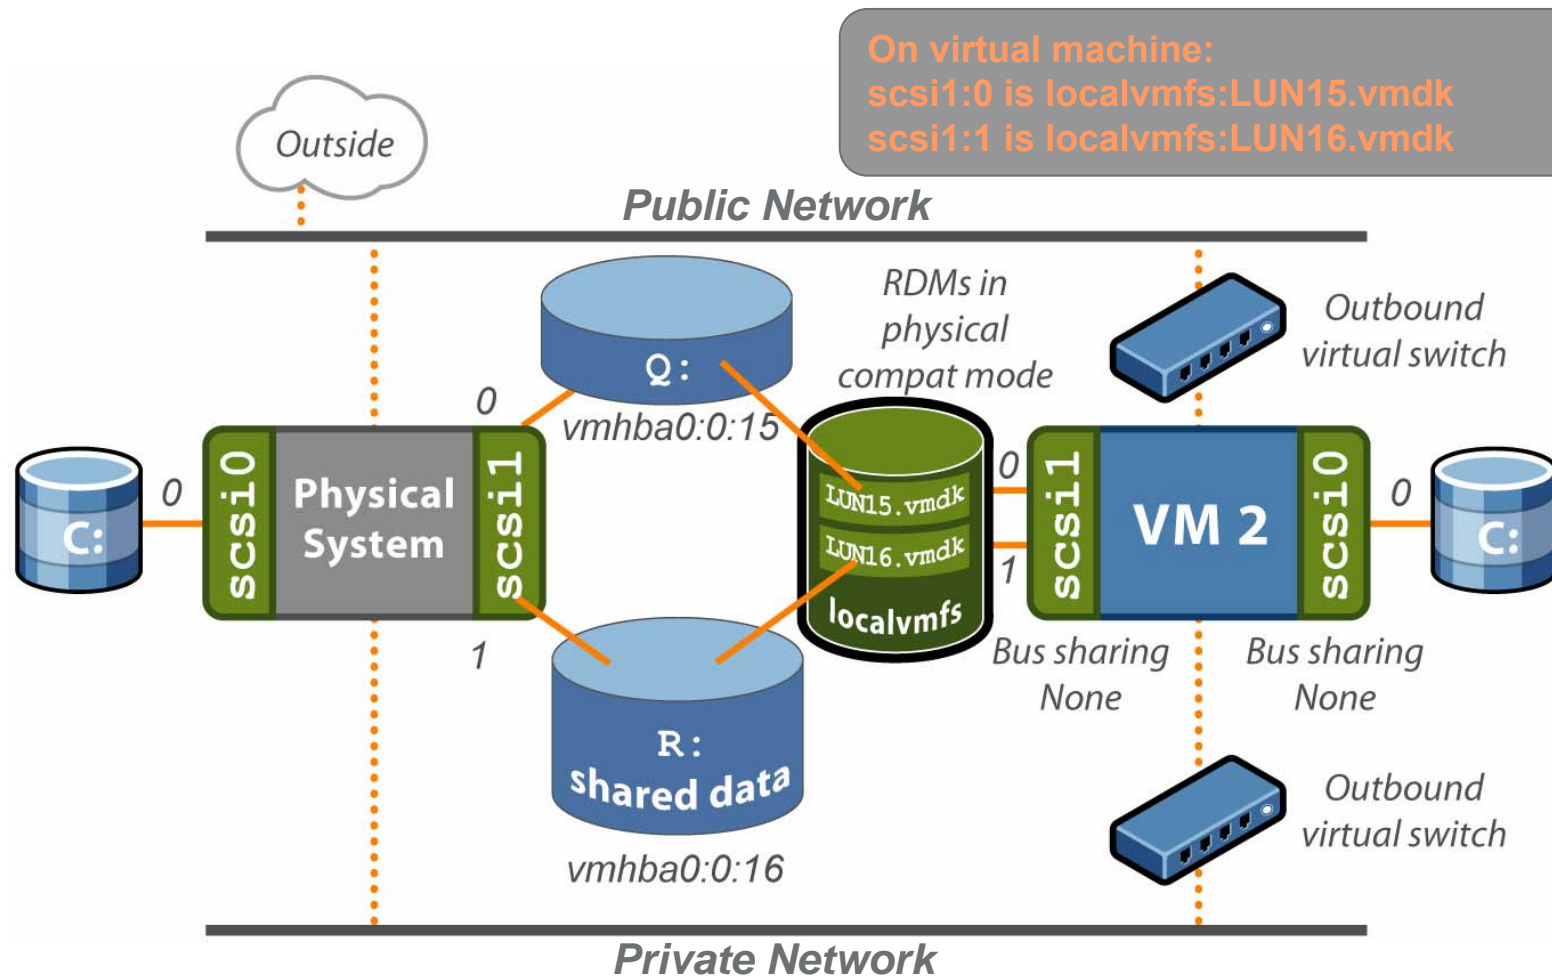


Implementing Cluster-Across-Boxes

On both virtual machines:
scsi1:0 is vmhba0:0:7:quorum.vmdk
scsi1:1 is vmhba0:0:7:data.vmdk



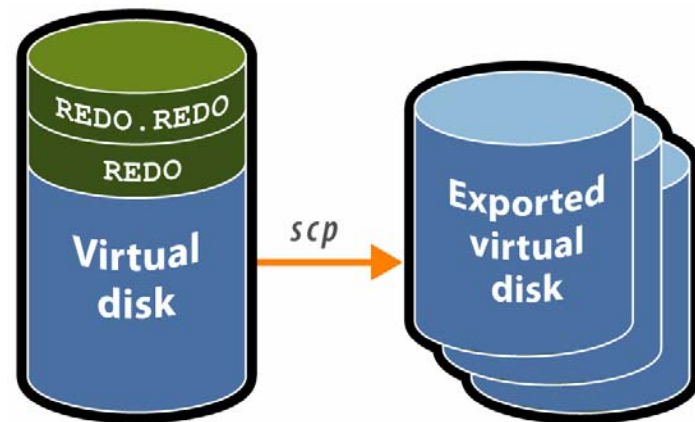
Physical-To-Virtual Cluster



Live Backups

Creating Snapshot Backups

vmsnap.pl	
Pros	<ul style="list-style-type: none"> Allows archiving virtual disks without virtual machine shutdown
Cons	<ul style="list-style-type: none"> Only crash-consistent Requires local and remote storage space
	Best when applications can quiesce themselves, and when a cold backup is also available



To back up:

```
vmsnap.pl -a archiveserver -c /home/vmware/name/name.vmx
```

To restore:

```
vmres.pl -a archiveserver -c /home/vmware/name/name.vmx  
-o owner -g group -v VMFS
```

This presentation covers the current versions of our products. Details about future releases of our products are available in select sessions at VMworld, including:

- PAC879:** The Next Phase of Virtual Infrastructure:
Introducing ESX Server 3.0 and VirtualCenter 2.0
- PAC177:** Distributed Availability Services Architecture
- PAC484:** Consolidated Backup with ESX Server:
In-Depth Review
- PAC485:** Managing Data Center Resources Using the
VirtualCenter Distributed Resource Scheduler
- PAC532:** iSCSI and NAS in ESX Server 3

Higher Availability with VMware Software: Implementation and Best Practices

Craig Williams
VMware Education

