





WHAT IS NEW IN VMWARE VSPHERE 4?

*Jon Hutchings & Marko Jung
Network Systems Management Service*

*ICT Forum Conference
14. July 2010, Kassam Stadium, Oxford*



NETWORK SYSTEMS MANAGEMENT SERVICE

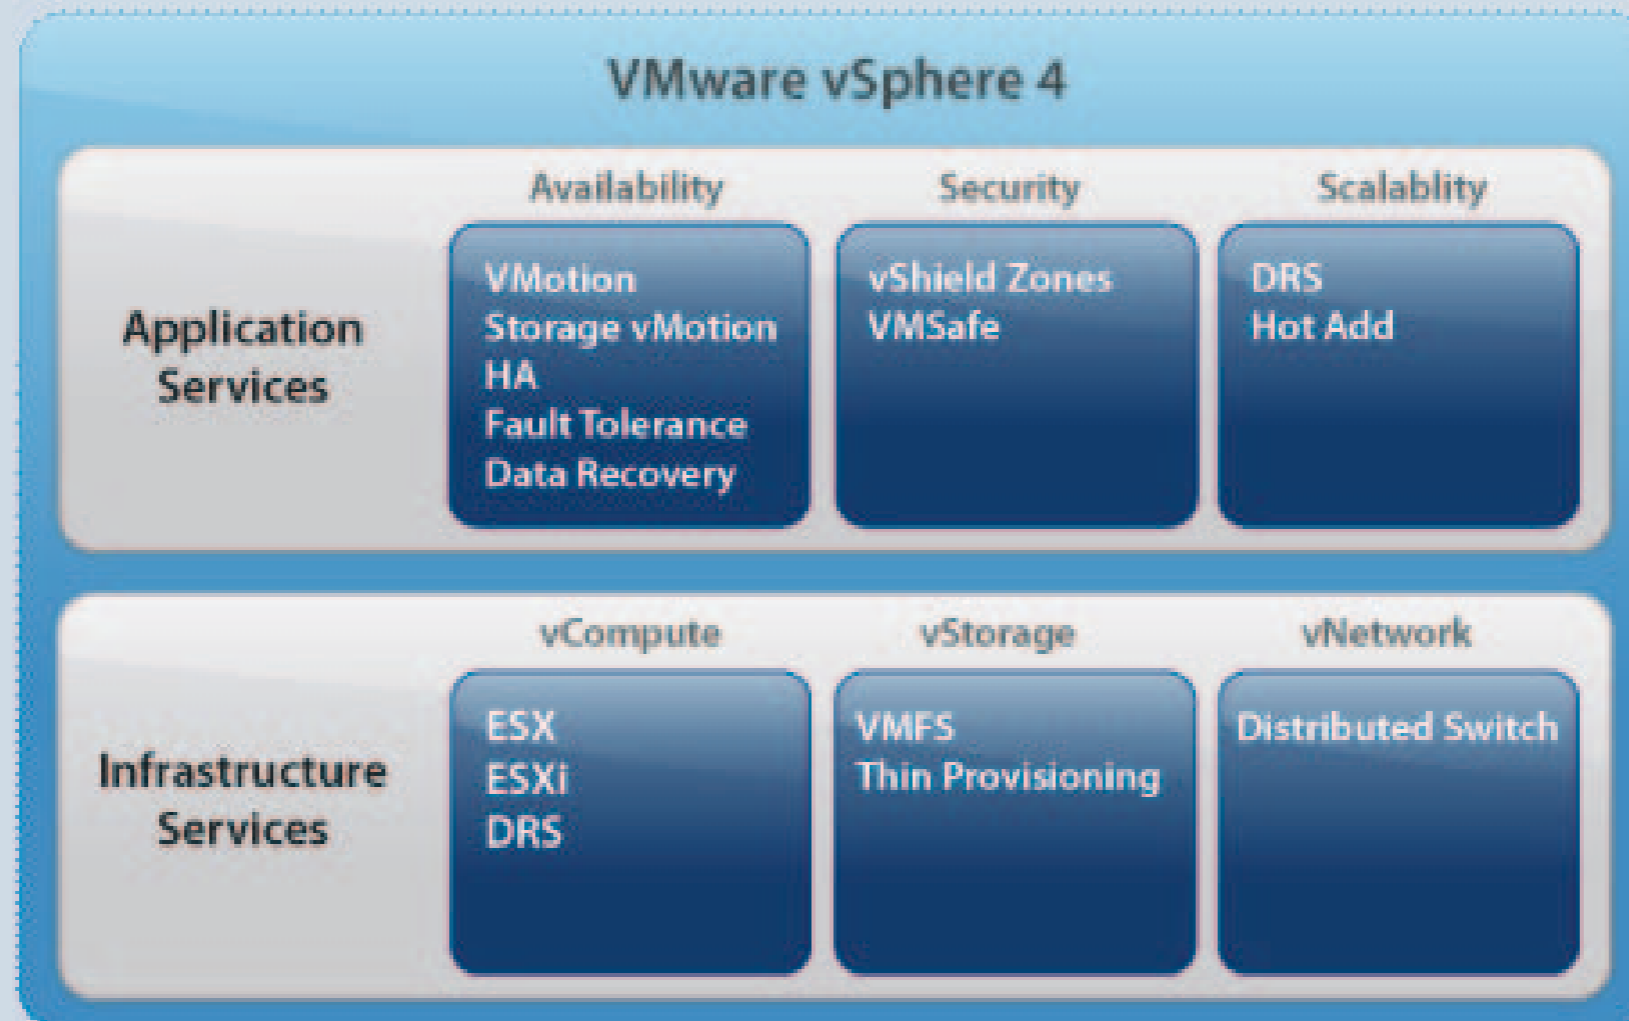
- NSMS includes many VMware trained staff
- Many years of experience from ESX 2.5.x onwards
- Experience of a wide range of VMware environments both small and large
- Able to help at any and all stages from planning to implementation and ongoing management
- NSMS operates 4 vSphere clusters as shared virtual infrastructure
 - 235 Virtual Machines
 - 19 Hosts
 - 443 GHz on 152 CPUs
 - 808 GB Memory

Highly resilient

- Dual homed
- Continuous data replication
- VMware build-in HA



VMware vCenter Suite




VSPHERE 4.0

DISTRIBUTED RESOURCE MANAGEMENT

DISTRIBUTED RESOURCE SCHEDULER (DRS)

Balance computing capacity


- Align resources to meet business needs
- Create rules and policies to prioritise resource allocation to VMs
- DRS balances the total workload of all VMs across all hosts in the cluster
- Can be fully automated or simply provide recommendations for administrators to action



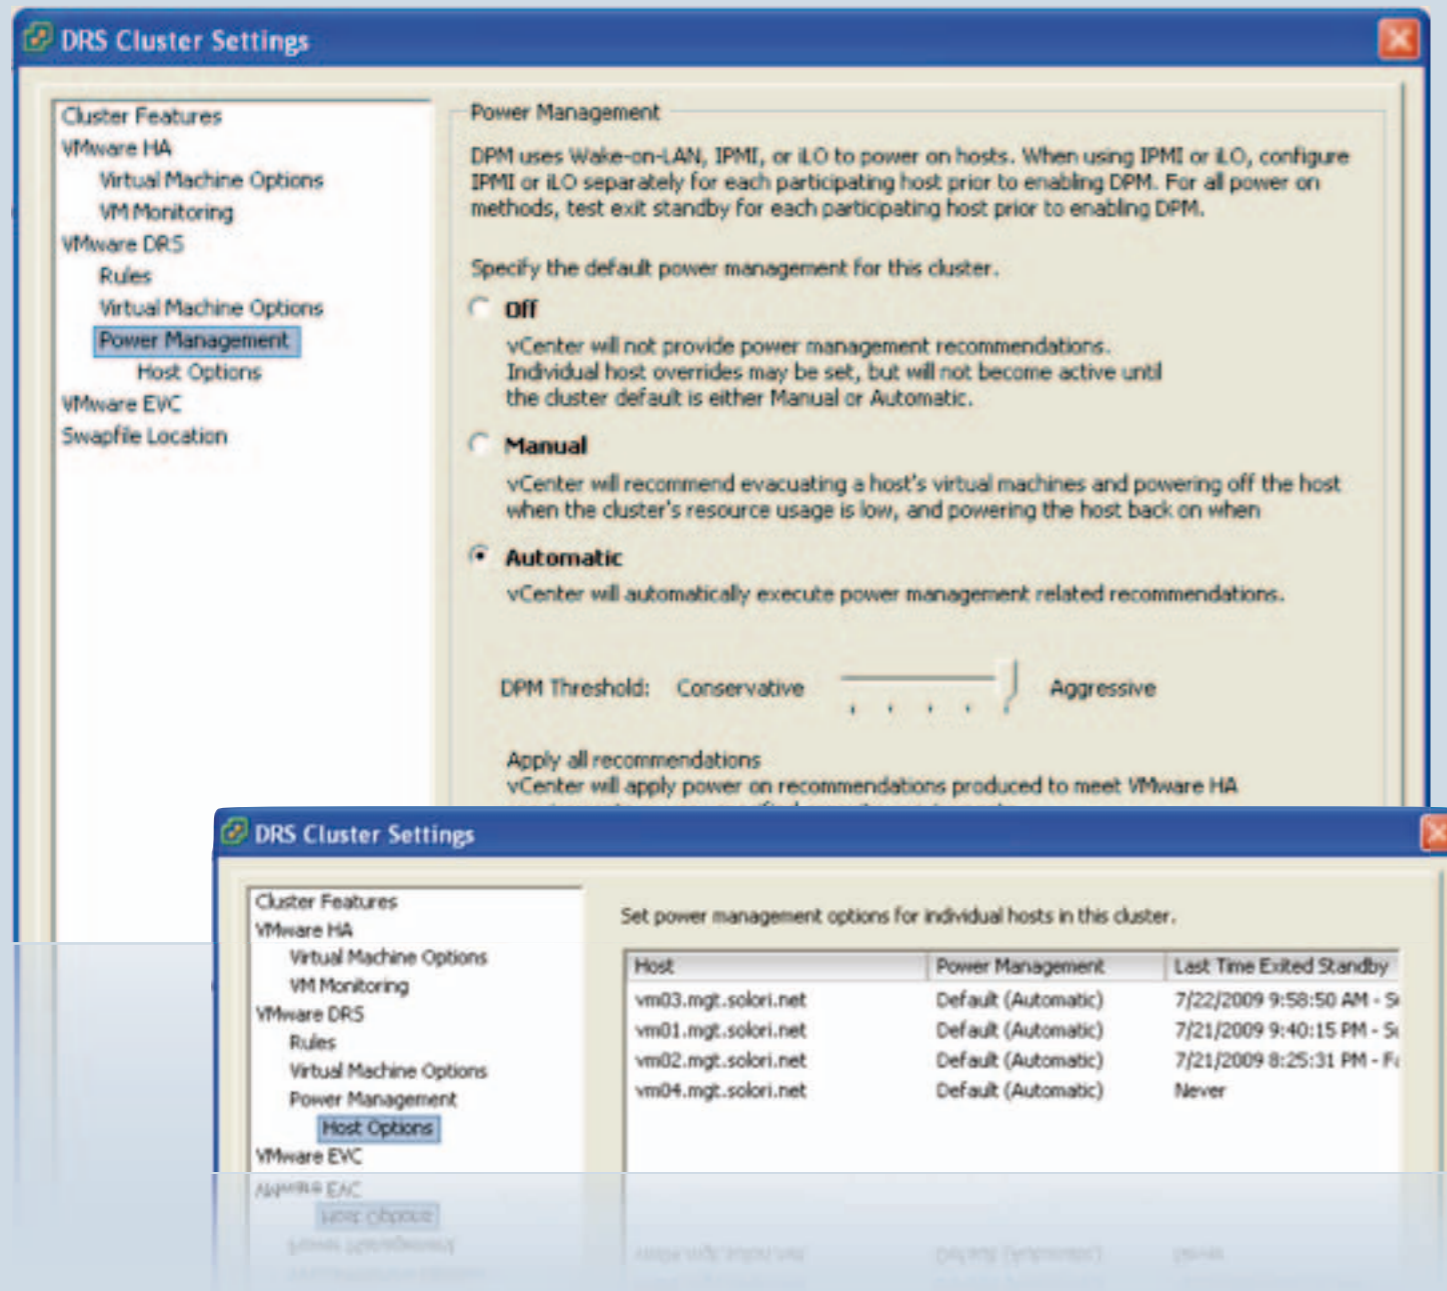
DISTRIBUTED POWER MANAGEMENT (DPM)

DPM monitors resource requirements across a DRS cluster:

- Reduced workload on cluster:
 - Consolidate workload
 - Place unused hosts in standby-mode
 - Resource requirements increase:
 - Bring standby-mode hosts back
 - Distribute workload to ensure service levels
- Cut power and cooling costs during low utilisation times
- Automate energy efficiency management in data centres



DISTRIBUTED POWER MANAGEMENT (DPM)



Levels of operation

- Manual: DPM makes recommendations that appear on the DRS tab
- Automatic: DPM executes host operations if VMs can be migrated based on the priority (aggressive to conservative)

→ You want to disable DPM for individual hosts of a cluster where the 'Last Time Exited Standby' status is 'Never'.

DISTRIBUTED POWER MANAGEMENT (DPM)

Supported technologies

- Intelligent Platform Management Interface (IPMI)
- Integrated Lights-Out (iLO)
- Wake on LAN (WoL)

Set-up

- Test VMotion link of each host
- Configure IPMI, iLO or WoL (on VMotion NIC)
- Test the wake capability of each host by choosing the 'Enter Standby Mode' and 'Power On' commands

The screenshot shows the vSphere Configuration page for a host. The 'Configuration' tab is selected, and the 'IPMI/iLO Settings for Power Management' section is expanded. Below the title bar, there are tabs for 'Allocation', 'Performance', 'Configuration', 'Tasks & Events', 'Alarms', and 'Permissions'. The main content area contains the following information:

IPMI/iLO Settings for Power Management [Properties...](#)

IPMI or iLO settings permit the host to participate in a Distributed Power Management (DPM) enabled cluster. In addition, manual Standby and Power On commands may be issued through vCenter.

User name:	vsphere
BMC IP Address:	192.168.13.131
BMC MAC Address:	00:e0:81:b5:55:...


Below the settings, a tree view shows the vSphere environment structure:

- vsphere.solori.local
 - SOLORI
 - DRS Cluster
 - vm01.mgt.solori.net
 - vm02.mgt.solori.net (standby mode)
 - vm03.mgt.solori.net
 - vm04.mgt.solori.net (standby mode)

VSPHERE 4.0

NETWORKING

vNETWORK DISTRIBUTED SWITCH



ESX Host 1

ESX Host 2

ESXi Host 3



vNETWORK DISTRIBUTED SWITCH

- Provides a central point of control for datacenter-level virtual networking.
- Centralised configuration interface for all virtual machine networking
- Maintains network runtime state for virtual machines as they live migrate from one host to another
- Allows higher security by enabling inline monitoring and centralised firewall services.
- Enable third party development, i.e. Cisco Nexus 1000V

FEATURE	STANDARD SWITCH	DISTRIBUTED SWITCH
Switch Features		
Layer 2 Forwarding	Yes	Yes
IEE 802.1Q VLAN	Yes	Yes
private VLAN		Yes
Multicast Support	Yes	Yes
Network Policy VMotion		Yes
Physical Switch Connectivity		
Ethernet Channel	Yes	Yes
Load Balancing Algorithms	Yes	Yes
VM network port block		Yes
Traffic Management Features		
Tx Rate Limiting	Yes	Yes
Rx Rate Limiting		Yes
Management Features		
Data-centre level management		Yes
vNetwork Switch APIs		Yes

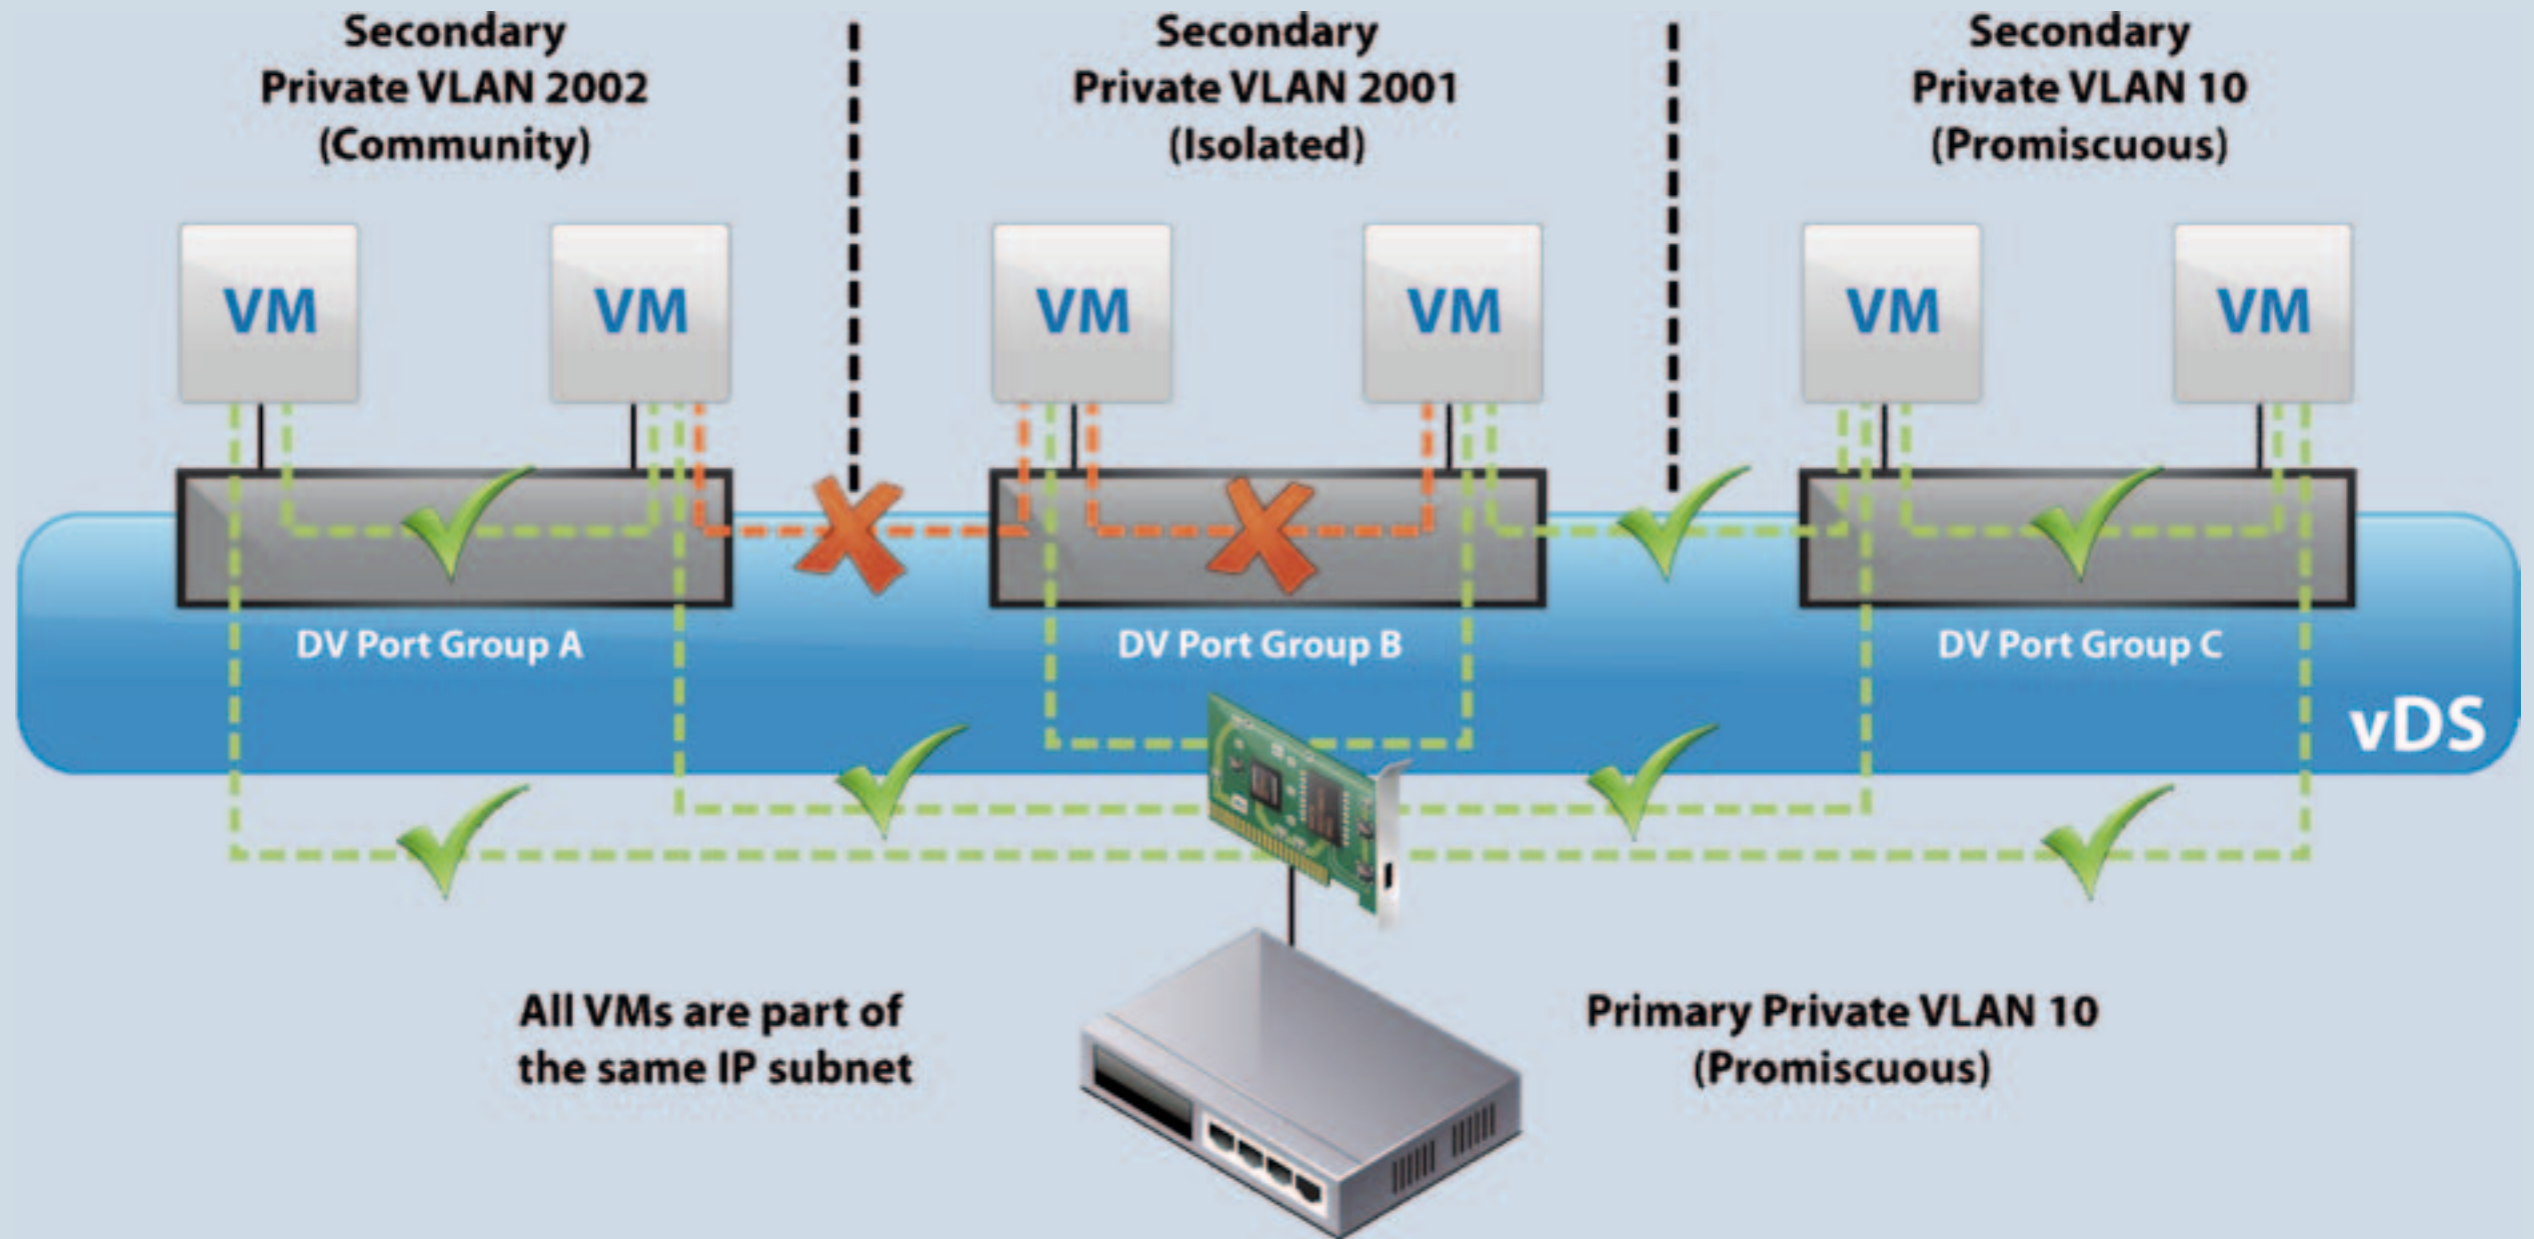
VSPHERE 4.0 NETWORK MAXIMUMS

MAXIMUM	STANDARD SWITCH	DISTRIBUTED SWITCH
switches per VC	4,096	16
switches per ESX host	248	16
port groups per ESX host	512	512
port groups per switch	512	512
ports per host	4,096	4,096
uplinks per host	32	32
ports per switch	4,088	6,000
uplinks per virtual switch	32	32
VLANs/private VLANs	Limited by maximum number of port groups	

PRIVATE VLAN ARCHITECTURE (PVLAN)

- Mechanism to divide a broadcast domain into several logical broadcast domains.
- Private VLAN is an extension to the VLAN standard to add further segmentation.
- Private means that the hosts in the same PVLAN are not able to be seen by the others, except the selected ones in the promiscuous PVLAN.
- The VLAN is further divided using secondary PVLAN types:
 - Promiscuous
 - Isolated
 - Community
- No Double Encapsulation
- Standard 802.1Q Tagging
- Switch software decides which ports to forward the frame, based on the tag and the PVLAN tables.

PRIVATE VLAN ARCHITECTURE (PVLAN)



VSPHERE 4.0

STORAGE

STORAGE CAPACITY MANAGEMENT

Datastores are managed as object in vCenter via datastore view

The screenshot shows the vSphere Client interface for a datacenter. The left-hand navigation pane displays a tree structure under 'vc.demo.local' with 'Datacenter' expanded to show 'Local Datastores' and 'SAN Datastores'. The 'SAN Datastores' section includes 'LUN 1', 'LUN 2', 'Production_Central', 'Efficient', 'Storage1', and 'Storage1 (1)'. The main pane shows the 'Datacenter' view with tabs for Summary, Virtual Machines, Hosts, Datastores, IP Pools, Performance, Tasks & Events, Alarms, Permissions, Maps, and Storage Views. The 'Datastores' tab is active, displaying a table with columns for Identification, Status, Device, Capacity, Free, Type, and Last Update. The table lists several datastores, including Storage1 (1), Storage1 (2), Storage1, Core_infrastructu..., LUN 1, LUN 2, Efficient, and Production_Central, with their respective capacities and free space.

Identification	Status	Device	Capacity	Free	Type	Last Update
Storage1 (1)	Normal	naa.600508b100...	110.50 GB	102.33 GB	vmfs3	10/22/2009 5:11:13 PM
Storage1 (2)	Normal	naa.6001e4f03ba...	134.75 GB	126.58 GB	vmfs3	10/22/2009 5:10:59 PM
Storage1	Normal	naa.600508b100...	110.50 GB	102.33 GB	vmfs3	10/22/2009 5:10:35 PM
Core_infrastructu...	Normal	naa.6006016090...	299.75 GB	186.63 GB	vmfs3	10/22/2009 4:10:56 PM
LUN 1	Normal	naa.6006016090...	299.75 GB	286.20 GB	vmfs3	10/22/2009 4:10:56 PM
LUN 2	Normal	naa.6006016090...	299.75 GB	299.20 GB	vmfs3	10/22/2009 4:10:56 PM
Efficient	Alert	naa.6006016090...	69.75 GB	4.40 GB	vmfs3	10/22/2009 5:10:35 PM
Production_Central	Normal	naa.6006016090...	149.75 GB	124.24 GB	vmfs3	10/22/2009 5:10:59 PM

STORAGE CAPACITY MANAGEMENT

Set Alarms on % FULL and % overcommitted

The screenshot displays the VMware vSphere interface for an Efficient storage object. The left-hand navigation pane shows the hierarchy: vc.demo.local > Datacenter > SAN Datastores > LUN 1 > LUN 2 > Production_Central > Efficient. The main window is titled 'Efficient' and has tabs for Summary, Virtual Machines, Hosts, Performance, Configuration, Tasks & Events, Alarms, Permissions, and Storage Views. The 'Alarms' tab is active, showing a table of triggered alarms:

Name	Defined In	Description
OverCommit Percent	This object	
Datstore usage on d...	vc.demo.lo...	Default alarm to monitor datastore d...

An 'Alarm Settings' dialog box is open, showing the configuration for the 'OverCommit Percent' alarm. The 'Triggers' tab is selected, displaying a table of trigger conditions:

Trigger Type	Condition	Warning	Condition Length	Alert	Condition Length
Datstore Disk Overallocation (%)	Is above	150		200	
Datstore Disk Usage (%)	Is above	75		90	

Below the table, there are two radio button options: 'Trigger if any of the conditions are satisfied' (which is selected) and 'Trigger if all of the conditions are satisfied'. At the bottom of the dialog are buttons for 'Add', 'Remove', 'OK', 'Cancel', and 'Help'.

REPORTING

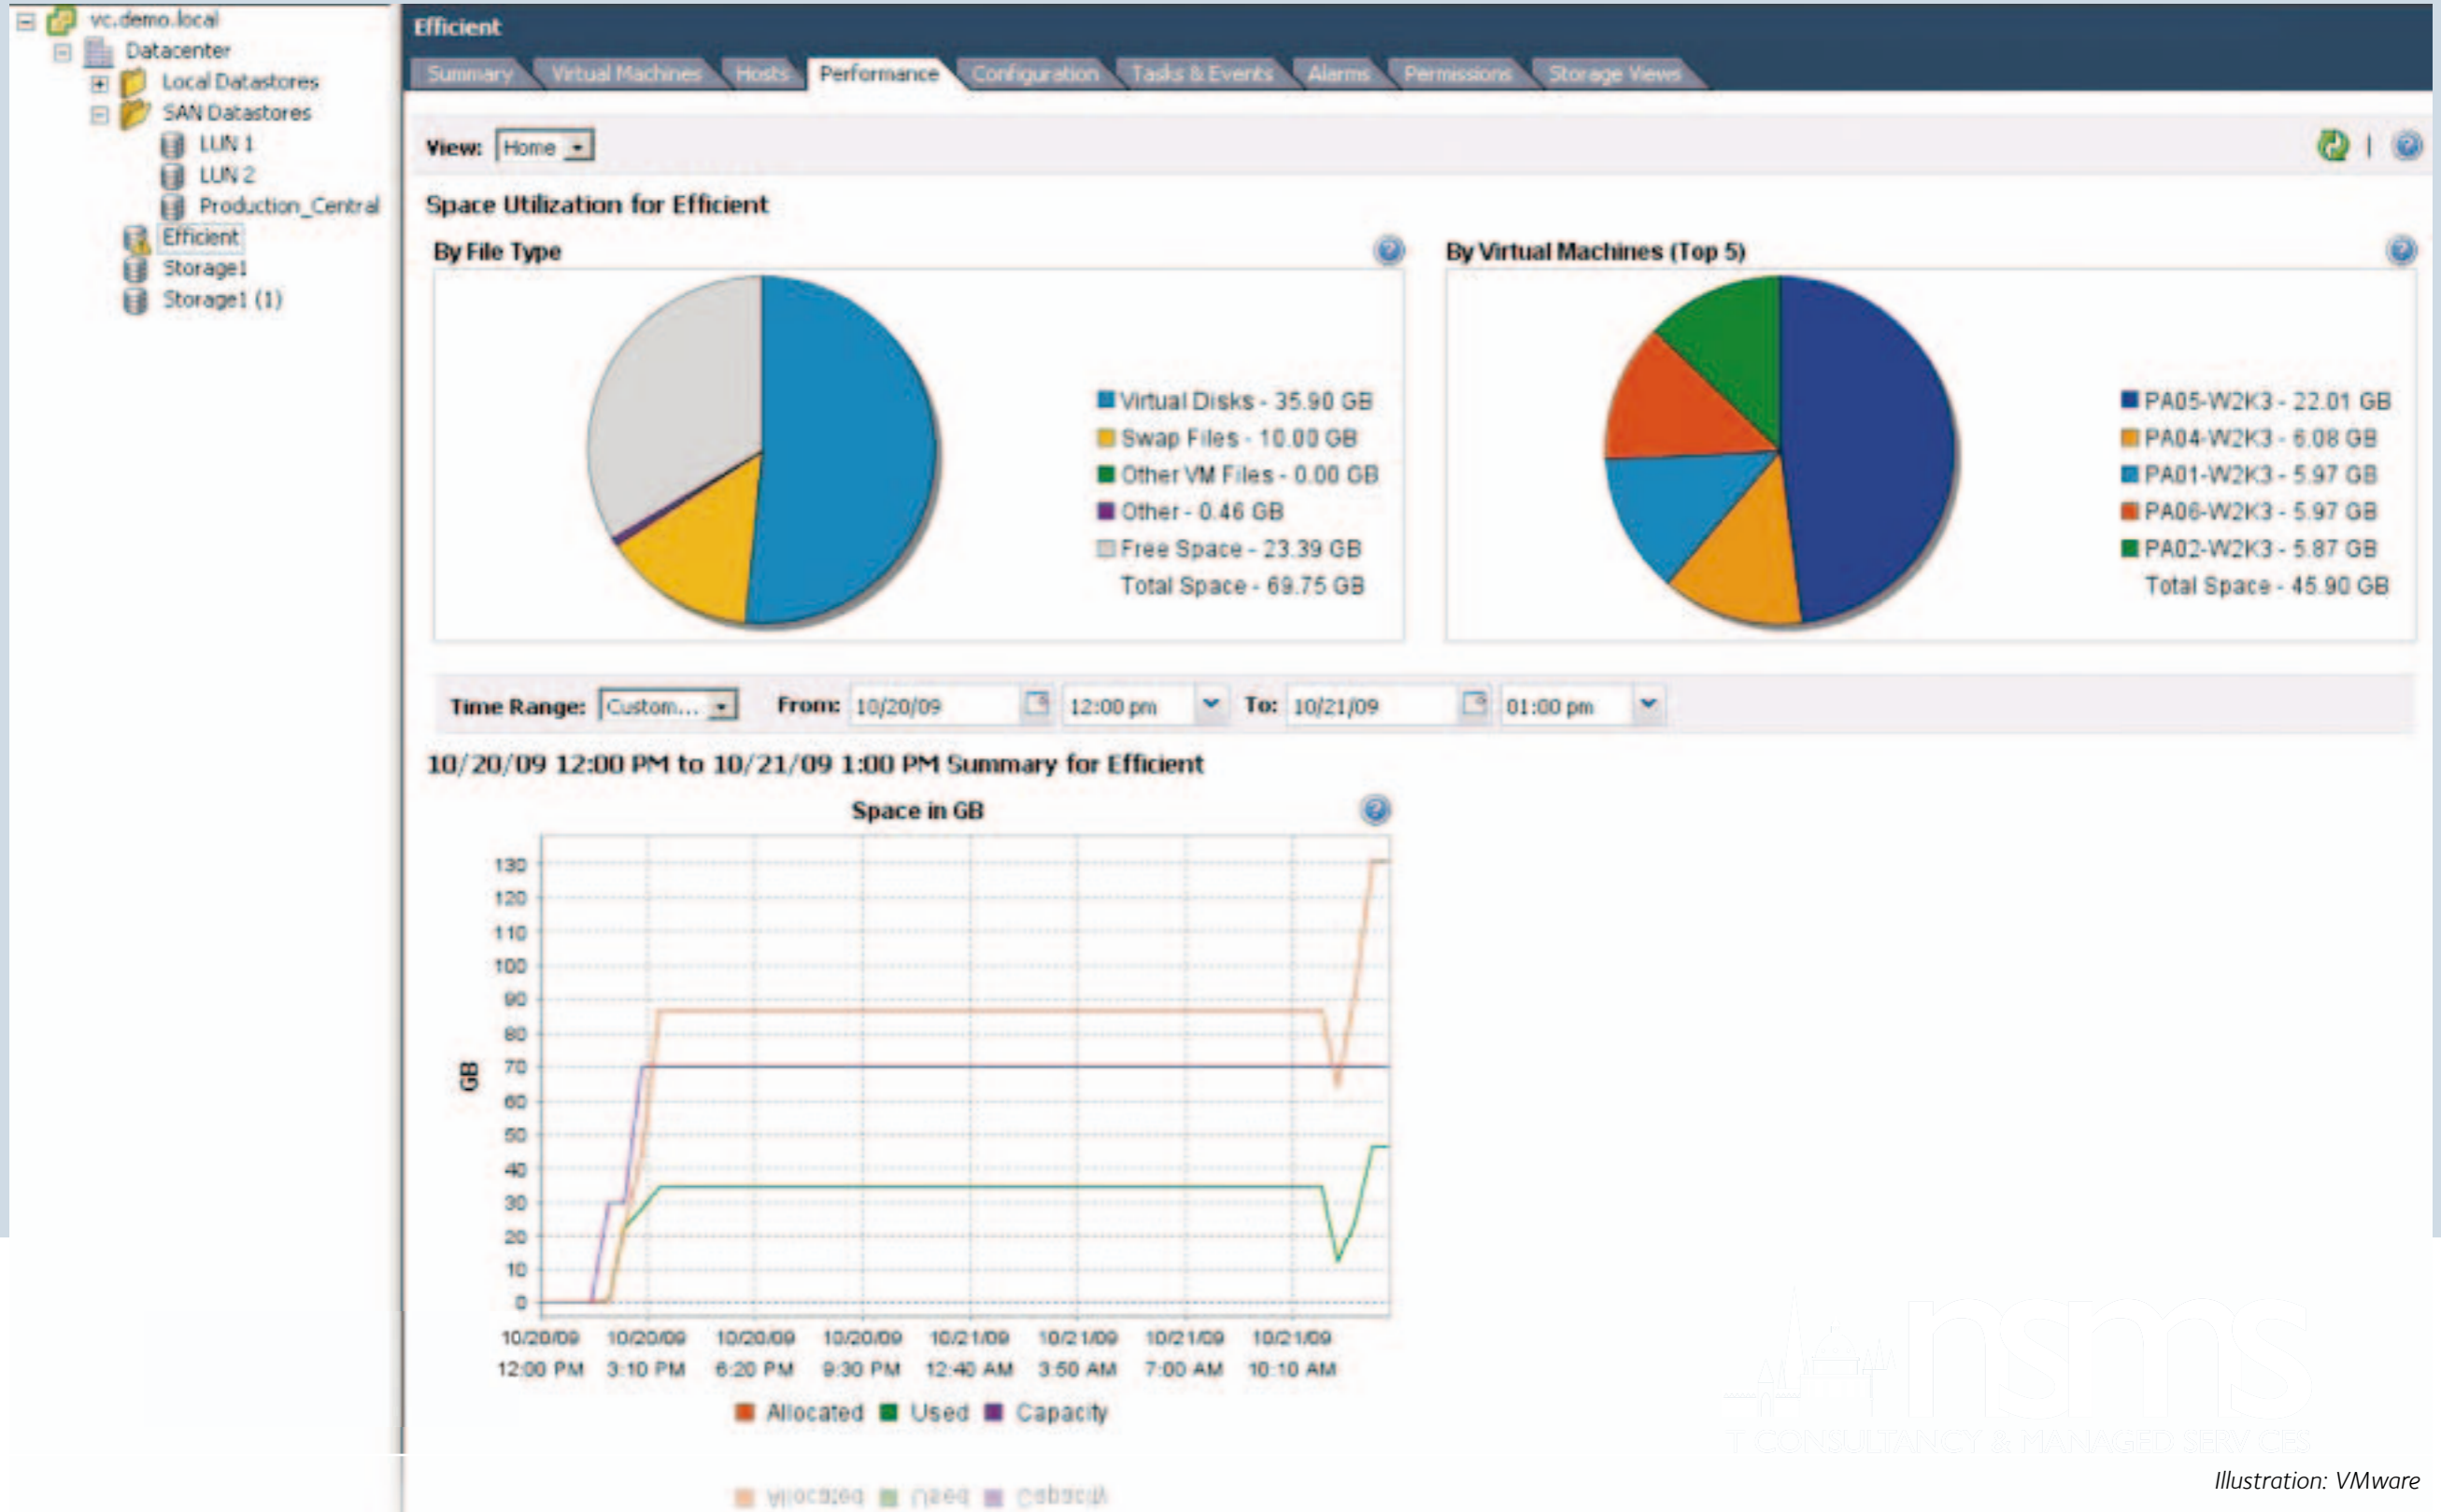


Illustration: VMware

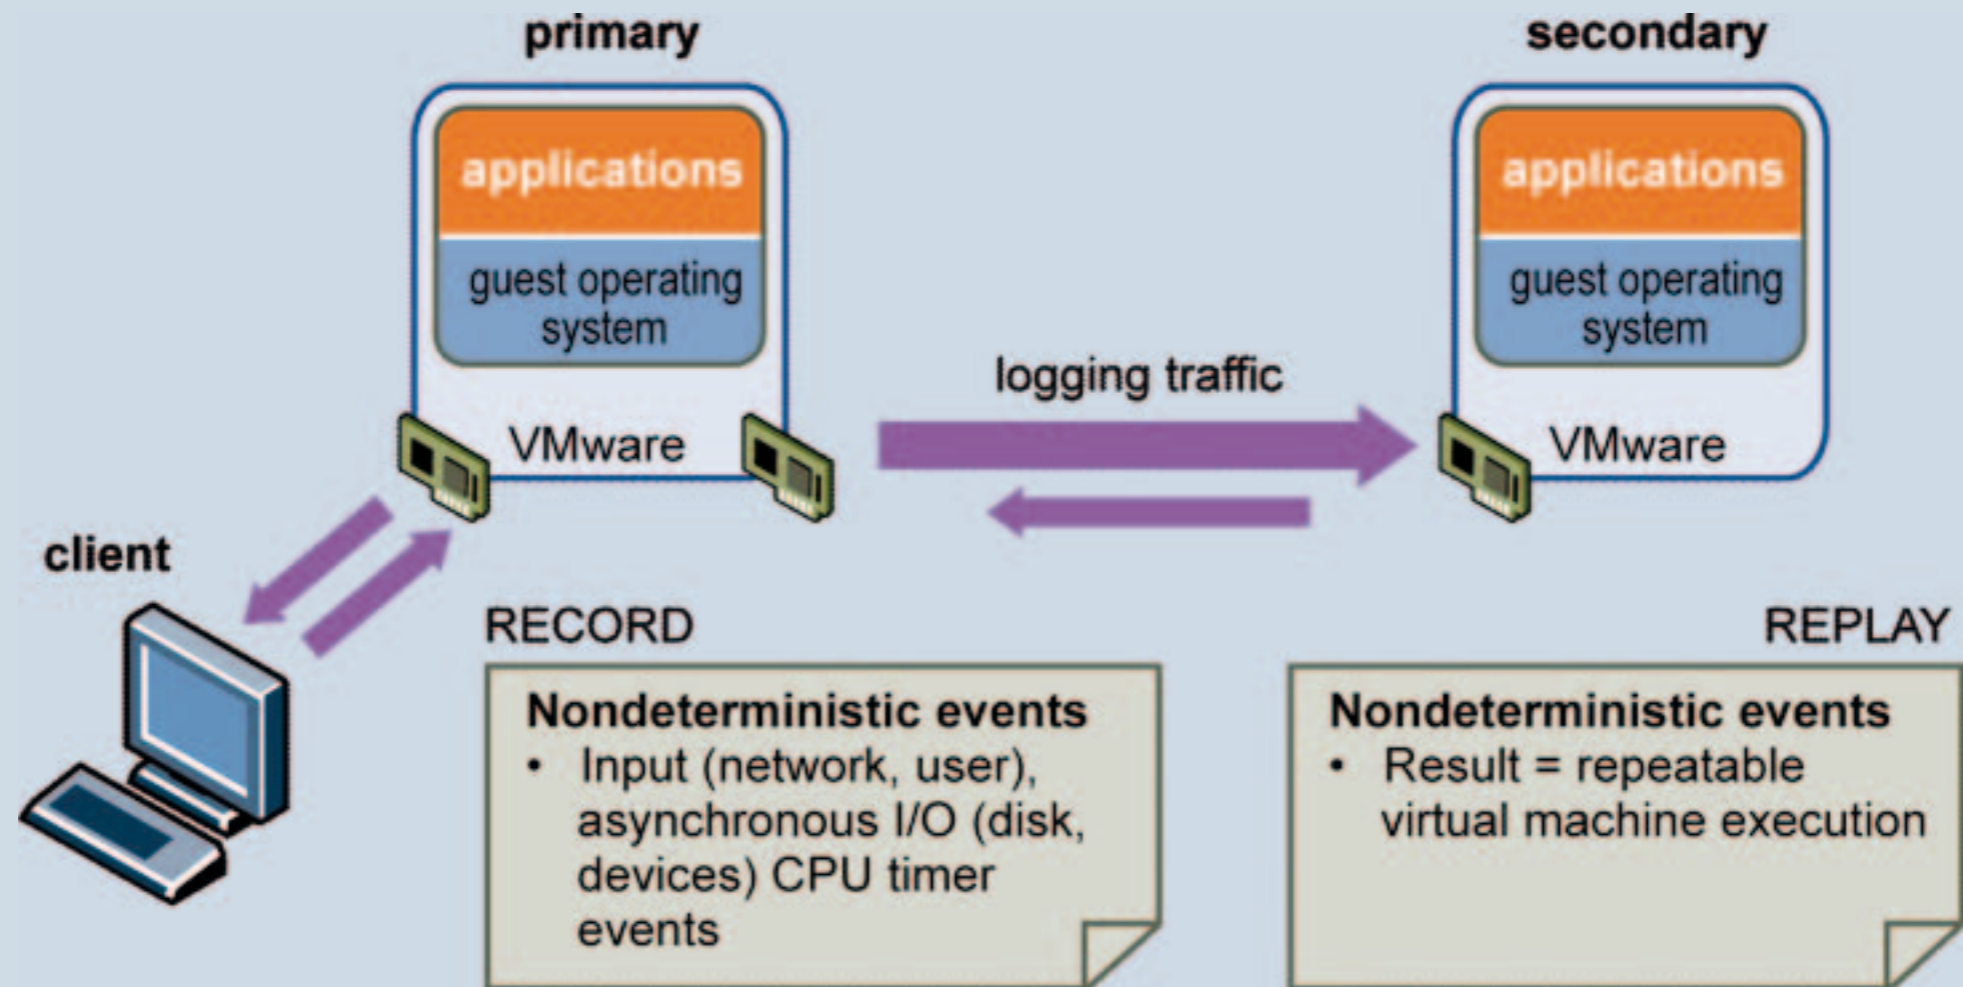
VSPHERE 4.0

HIGH AVAILABILITY & FAULT TOLERANCE

VSPHERE HIGH AVAILABILITY (HA)

- VMware HA protects application availability
- protects against a server failure by automatically restarting the virtual machines on other hosts within the cluster.
- protects against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.
- VMware Tools must be installed for application monitoring to function.
- Test application monitoring thoroughly before making VM live and relying on it.

VSPHERE FAULT TOLERANCE (FT)



VSPHERE FAULT TOLERANCE (FT)

- Provides Continuous availability
- Replicates CPU instructions from a Primary VM to a Secondary copy
- In event of failure Secondary becomes Primary and a new Secondary is created
- Only one VM seen in VC client (unless you look for it)
- Once configured will work even if VC server is down
- Will not help if Primary crashes due to software error
- Large VMs (15GB RAM or more) may not work with FT due to latency of memory copies

VSPHERE FT LIMITATIONS

- No Snapshots
(so no VM level backups)
- No SMP VM support
- No Physical Raw Device Mappings
- No CD/Floppy/USB attached devices
- Limited network device support, no VMXNET3 or legacy devices, only VMXNET2, VMXNET and e1000
- No paravirtualised SCSI devices, only Buslogic and LSI
- No NIC passthrough
- Must use thick provisioned, FT aware discs (converted automatically)
- No hot plug hardware
- May not work with VMs with 15GB or more RAM



VSPHERE 4.X

Photo: <http://www.flickr.com/photos/wiebke-worm/>

VSPHERE 4.X: THE BIGGER THE BETTER

- 3,000 virtual machines per cluster (compared to 1,280)
- 1,000 hosts per VC Server (compared to 300)
- 15,000 registered VMs per VC Server (compared to 4,500)
- 10,000 concurrently powered-on VMs per VC Server (compared to 3,000)
- 500 hosts per virtual Datacenter object (compared to 100)
- 5,000 virtual machines per virtual Datacenter object (compared to 2,500)
- 120 concurrent Virtual Infrastructure Clients per VC Server (compared to 30)
- hostd footprint and memory consumption reduced by 40%



VSPHERE 4.X: NEW FEATURES

Scalable concurrent vMotion

- 4 live migrations over 1GbE
- 8 live migrations over 10GbE (up to 8GB/sec)

vNetwork Distributed Switches can load balance traffic on team physical network interfaces

Multiple Data Recovery virtual appliances

Support for FT in DRS clusters with Enhanced vMotion Compatibility

Refined Distribute Resource Scheduling (DRS) VM affinity rules

ESX/ESXi AD Integration

Enhanced hardware support

- 8 GB Fibre Channel HBAs
- iSCSI TCP Offload Engine (TOE) NICs
- Intel Xeon 7500 / 5600 / 3600 CPU series
- AMD Opteron 4000 / 6000 CPU series

VSPHERE 4.X: PERFORMANCE

Storage I/O Control

- Quality of service prioritisation per host or cluster
- Congestion monitoring and prioritisation
- Limits of I/O operations per VM

Network I/O Control

- Prioritisation now also based on traffic type: virtual machines, vMotion, Fault Tolerance, and IP storage


Transparent Memory Compression (TMC)

- On the fly compression for RAM pages instead of swapping
- Compression cache for small pages (<2KB)
- 15%-25% performance gain depending on memory over-commitment

Improved VM NUMA scheduler



Questions?





WHAT IS NEW IN VMWARE VSPHERE 4?

*Jon Hutchings & Marko Jung
Network Systems Management Service*

*jon.hutchings@oucs.ox.ac.uk
marko.jung@oucs.ox.ac.uk*

