

Implementing a cross-site infrastructure with VMware and LeftHand Networks

Jon Hutchings

Oxford University Computing Services



Agenda

- Project Background and Timescale
- Overall Goal
- Systems and Services
- Design Overview
- Final Design Detail
- Implementation
- Summary and Futures




Project Background and Timescale

- Mid 2007 : Business System & Projects approached NSMS for VMware virtualization advice, alongside external consultants.
- Mid 2007 : ICT approach NSMS for help with virtualization of ICT resources
- Late 2007: Agreement on a joint BSP & ICT infrastructure managed by NSMS
- Early 2008 : Detail design decided upon
- March 2008 : Equipment ordered
- Mid July 2008 : Solution delivered to BSP and ICT

Overall Goal

- Consolidate large numbers of ageing physical servers
- Improve system reliability
- Protect against a large scale site failure
- Break the link between a service and a piece of hardware, allowing for better long term maintenance and upgrades.



Systems and Services

- File Serving (Windows / Netware)
- Mail (Notes/ Domino)
- Database (Oracle, MS SQL)
- Domain Controllers
- Etc. etc.
- x86 based, mixture of rack mount and tower legacy servers

Virtualization

- X86 Server Virtualization
 - Extensive experience within the NSMS team with VMware Virtual Infrastructure (VI).
 - VI felt to be the only platform mature enough for business critical systems.
- Storage Virtualization
 - Not key to initial project design, however a nice side effect of using an iSCSI SAN solution such as Lefthand SAN I/Q

Design Overview (1)

- Worked with consultants Aenigma who developed the detail design
- Initial design was for a single site solution for BSP. A cluster of VMware ESX servers and a fibre channel or iSCSI SAN
- With ICT joining the project thoughts moved to a dual site “failover” solution



Design Overview (2)

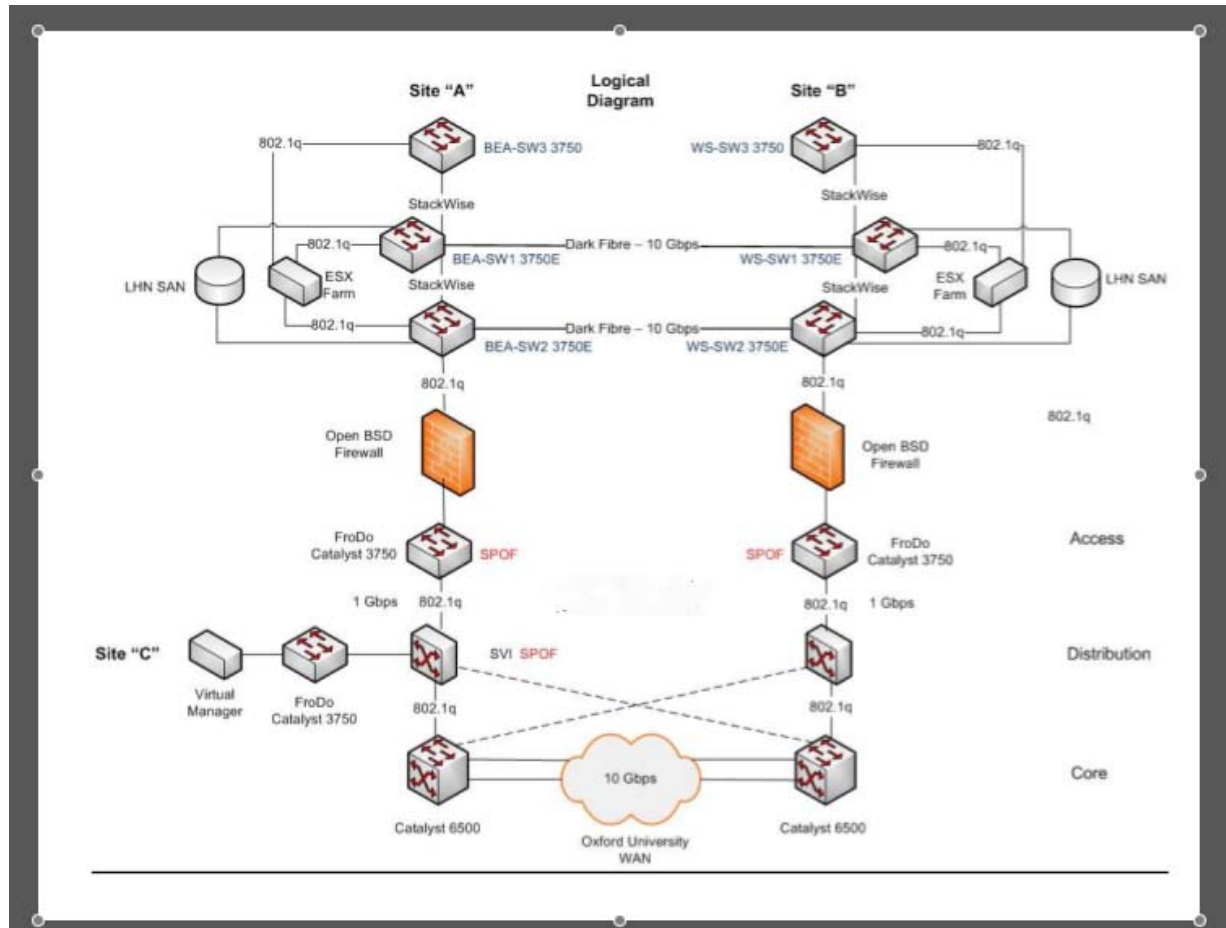
- Initial designs based around VMware VI3.0.1 and open source iSCSI SAN solution provided by Aenigma
- Business Decision to use a “Major Vendor” for SAN solution to ensure long term support.
- Lefthand networks SAN I/Q offered best feature and performance set.



Final Design

- 14 ESX (3.5) servers 7 per data center
- 2 VMware Clusters
 - BSP 4 per site
 - ICT 3 per site
- Single Virtual Center managing both clusters
- NSMS Active Directory used to provide Virtual Center accounts with delegated roles to allow each team access to their share of the infrastructure
- 10 LHN SAN I/Q nodes (HP DL320S)
 - 5 per site, each node 12 drives, 3.6TB raw per node.
 - All LUNS 2 way replicated with data written to both sites
 - approx 16TB replicated
- 2x10GB fiber links for storage traffic
- 2x42U racks, 1 per site
- Looped Network Design! Reliant on spanning tree to prevent loops
- esXpress Backup used to provide VM level backups to large ftp systems at each site

Final Design (Network)



Final Design (VMware)

- Each server cluster uses VMware HA - will auto-restart VMs in case of hardware failure of physical hosts
- DRS (distributed resource scheduling) on each cluster allows the VMs to be balanced across both sites. Rules allow individual machines to be fixed to a physical server(s) except in case of failure.
- VMware Update manager allows automated patching of hosts, with VMs automatically migrated to other available hosts to allow for zero downtime (to the VMs) patching and upgrading.

Final Design (Storage)

- Lefthand Networks SAN I/Q
- 5 nodes per site, presented via private fibre links as a single storage cluster.
- Volumes 2-way replicated, with a copy at each site.
- iSCSI between VMs and storage over private fibre, using static routes on switch-stacks at each site for direct access within VM machines.
- VM providing a casting vote in case of site failure located at 3rd site.

Implementation

- Implemented by a small team (2-3) over 3 months
- ESX servers built with V-Ployment scripted deployment tool.
- Implementation of underlying SAN and Virtual Infrastructure delivered on Schedule in July
- Only significant design change was a switch to Ubuntu Server for firewalls
- Network Configuration key to the success of the project

Summary and Futures

- Successful Implementation, a lot of planning up front enabled implementation to proceed with only minor delays
- Currently 250+ VMs running on the infrastructure
- 22TB allocated storage - 90%+ in use!
- OUCS Networks Team involvement crucial to the project.

References and Thanks!

- OUCS Networks Team!
- <http://h18006.www1.hp.com/storage/highlights/lefthandsans.html>
- www.vmware.com
- www.esxpress.com
- www.nrlcomputers.net (V-Ployment)
- www.aenigmacorp.com

jon.hutchings@oucs.ox.ac.uk