

# Implementing VMware Virtual Infrastructure: Experiences and Best Practices

Shane M. Blackwell

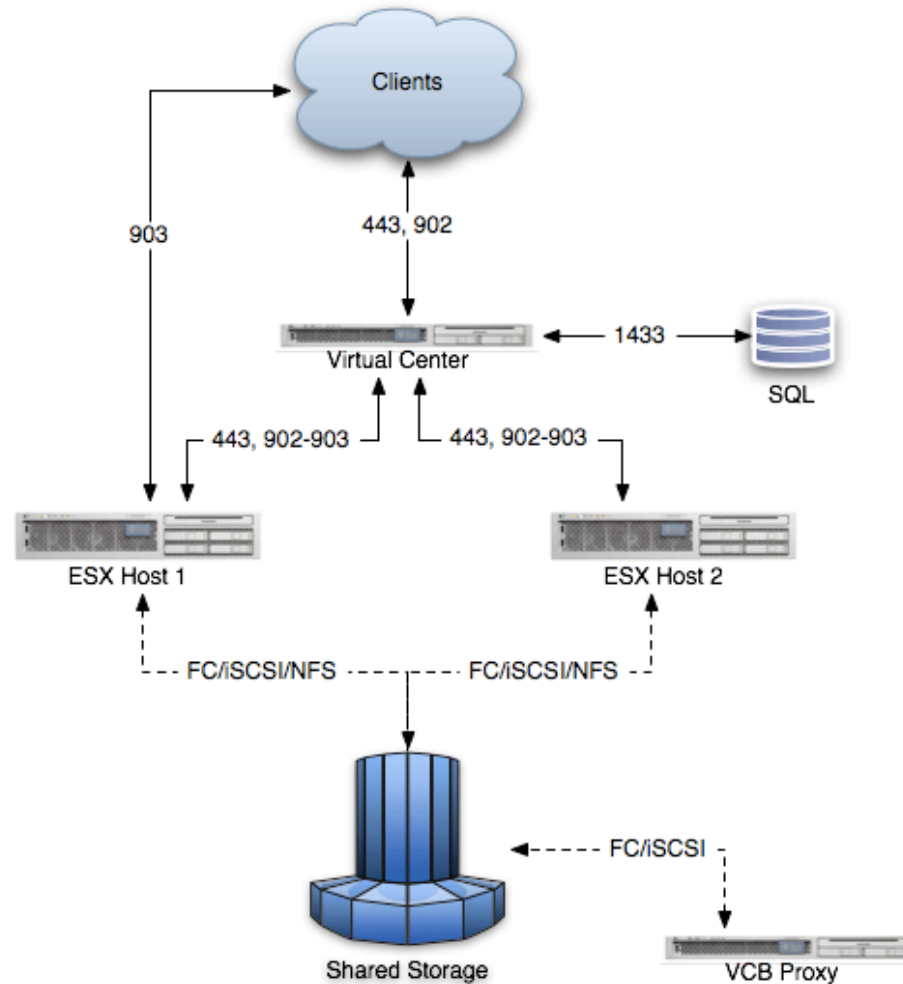
37<sup>th</sup> USG Computing Conference

24 October 2008

Office of Information Technology  
<http://www.oit.gatech.edu>

- Allows for unlimited use of:
  - VMware Virtual Infrastructure Enterprise
    - ESX Server
      - Includes HA, DRS, EVC
    - Virtual Center Server
    - VMware Consolidated Backup
    - VMware Converter Enterprise
- 50% discount on VMware products
- 50% discount on VMware services
- Platinum support contract
- VDI prohibited by the ELA

# Typical VI Architecture



- At least 2 ESX hosts
- Common storage
- 1 Virtual Infrastructure Management (VIM) Server
- 1 Database Server (if one does not exist)
- 1 Consolidated Backup Proxy (optional)

- Provides centralized licensing and management for ESX hosts and their virtual machines
- Requires a MS SQL or Oracle backend database
- Home to several key components
  - VMware HA
  - VMware DRS
  - VMware EVC

- Recommendations
  - Run Virtual Center and its database on separate machines if possible
  - MS SQL preferred to Oracle
    - Set up account in SQL Server to access database
    - Grant SQL account db\_owner permissions on VC database
  - Place on same subnet as ESX hosts
    - Simplifies firewall configurations
  - Install custom SSL certificates

- Virtual Center can be integrated into Active Directory
  - Allows the leveraging of AD security groups
- 9 Default Roles
- Custom Roles can be created based on the default roles

- Ability to reboot virtual machines in the event of a host failure
- Admission control to ensure failover capacity
- Can set default cluster settings for VM restart priority and host isolation

- Common issues
  - HA agents randomly have errors on hosts that appear to clear up in a few seconds
    - Normal behavior according to VMware
  - Persistent HA errors may require reinstallation of the HA agent on the affected host

- Balances workload across hosts in the cluster
- Can set automation levels and migration thresholds
- Allows for affinity and anti-affinity rules

- Recommended Settings
  - Migration Threshold → 2 stars or higher
  - Use affinity rules to ensure that VMs that perform similar functions do not end up on the same host

- Protects against potential CPU mismatch issues between hosts
  - Prevents the addition of incompatible hosts into the cluster
- Use only when starting from scratch
  - VMs from existing clusters may not be able to migrate to the EVC-enabled cluster

- Recommended hardware specifications
  - CPU → Dual-Socket Dual-Core or better
    - AMD → Barcelona
    - Intel → Harpertown
  - Memory → 16 GB or more
  - Local HDD → Minimal
- HBAs for SAN (FC/iSCSI) access
- Make sure server is on the VMware HCL

- Manual configuration process through Virtual Center
- Can leverage bash scripts in the Service Console to automate
  - Console OS (COS) slowly being phased out of ESX

- VMware Update Manager
  - Use Virtual Center to remediate patches against ESX hosts
  - Can also be used to patch Windows based guests
  - Requires an additional database on SQL Server to store patches

- Patch Depots
  - HTTP or FTP based
  - Can be a virtual machine
  - Use esxupdate command on ESX hosts
    - esxupdate -d <DepotURL> scan
    - esxupdate -d <DepotURL> --test update
    - esxupdate -n -d <DepotURL> update

- ESX utilizes VMFS partitions to store virtual machines and their associated files
  - Clustered file system
  - Shared storage (SAN) required (FC or iSCSI)
- Can also use NFS exports as datastores

- **Advantages**
  - Doesn't require expensive SAN hardware
- **Disadvantages**
  - Performance not as good as SAN
  - Limited to 8 NFS datastores per ESX host
- **Recommendations**
  - Use in non-production environments

- **Advantages**
  - Better performance than NFS
  - No additional hardware needed when software initiators are used
  - Less expensive than fibre channel
- **Disadvantages**
  - Easier to saturate bandwidth than fibre channel

- Recommendations
  - Can use for raw device mappings to VMs, test/devl, or mid-tier production datastores
  - Run on separate switch infrastructure
  - Bind multiple NICs when possible
    - Increases aggregate I/O bandwidth

- **Advantages**
  - Best performance
  - Support for multiple paths to SAN LUNs (multipathing)
- **Disadvantages**
  - Additional hardware required (HBAs)
  - More expensive than iSCSI
- **Recommendations**
  - Use for top-tier production environments

- At least 4 NICs needed for proper setup
  - 1 for Service Console (Backup for VMKernel)
  - 1 for VMKernel (Backup for Service Console)
  - 2+ for virtual machine network traffic

- 2 vSwitches recommended
  - Service Console/VMkernel
  - VM network traffic
    - Portgroups for each VLAN
- Recommended settings
  - Set number of ports on VM traffic switch to 256
  - Ensure that SC and VMkernel ports act as failover for one another
  - Make all NICs that are used for VM traffic active

- Start with low CPU, Low I/O services
  - Departmental file servers
  - Web servers
- Test/Dev Application Servers
- Production Application Servers

# What NOT to Virtualize

- Resource intensive applications (at first)
  - Require careful planning and resource allocation

- Windows Server 2003/2008
  - Create templates to provision new VMs from
  - Ensure that appropriate Sysprep tools are located on Virtual Center server
    - C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\syprep
  - Run monthly MS updates on templates
    - Convert template to VM, then revert to template once patching is complete

- Linux
  - Can use existing kickstart infrastructure if available (RedHat)
  - Can also use templates to provision (RHEL4 and RHEL5)
  - Install from ISO

- Allows for quiescing of virtual disks
- NOT appropriate as a backup solution
  - Increased disk usage in datastore
  - Performance hit
- Suitable for use during patching or installation of software in guest OS
  - Snapshot → Perform Update → Delete Snapshot

- Allows for the offloading of backups from ESX hosts to a proxy server
- Integration modules for major backup software vendors
- Script needed to perform backups from proxy
  - VBScript or PowerShell

- Recommendations
  - Make sure to have adequate drive space if backing up to spinning media (i.e. Disk-Disk-Tape)
  - Install correct backup module if going directly to tape

- Moving more towards cloud-based computing
- Fault tolerant virtual machines
- Distributed virtual switches
- Host profiles
- Ability to roll up multiple vCenter installations into one large vCenter

- **Systems Compatibility Guide**
  - [http://www.vmware.com/pdf/vi35\\_systems\\_guide.pdf](http://www.vmware.com/pdf/vi35_systems_guide.pdf)
- **Storage Compatibility Guide**
  - [http://www.vmware.com/pdf/vi35\\_san\\_guide.pdf](http://www.vmware.com/pdf/vi35_san_guide.pdf)
- **ESX/VC Installation Guide**
  - [http://www.vmware.com/pdf/vi3\\_35/esx\\_3/r35u2/vi3\\_35\\_25\\_u2\\_installation\\_guide.pdf](http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_installation_guide.pdf)
- **Other ESX/VC Documentation**
  - [http://www.vmware.com/support/pubs/vi\\_pages/vi\\_pubs\\_35u2.html](http://www.vmware.com/support/pubs/vi_pages/vi_pubs_35u2.html)

# Questions?