

Live Migration


Aidan Finn

About Aidan Finn




- Technical Sales Lead at MicroWarehouse (Dublin)
- Working in IT since 1996
- MVP (Virtual Machine)
- Experienced with Windows Server/Desktop, System Center, virtualisation, and IT infrastructure
- @joe_elway
- <http://www.aidanfinn.com>
- <http://www.petri.co.il/author/aidan-finn>
- Published author/contributor of several books

Books



System Center
2012 VMM



Windows Server
2012 Hyper-V

Live Migration Basics

W2008 Quick Migration

- Quick Migration in W2008 to move an offline (saved state) virtual machine
- Move:
 - 1 virtual machine at a time between two hosts
 - Service downtime depends on RAM size and disk speed
- Failover Clustering:
 - Did not support queuing
 - Was required and constrained Quick Migration

Quick Migration (W2008)

Host
A



Host
B




W2008 R2 Live Migration

- Live Migration introduced in W2008 R2 Hyper-V clusters to compete with vMotion
- Move:
 - 1 virtual machine at a time between two hosts
 - With no impact on service uptime
- Failover Clustering:
 - Did not support queuing
 - Was required and constrained Live Migration

Live Migration (W2008 R2)

Host
A




Host
B



How Live Migration Worked

Memory



Host 1



Host 2

Live Migration is not High Availability

- And High Availability (HA) is not Live Migration (LM)
- HA:
 - Clustered host dies
 - VM fails over to another host and boots up (minimal downtime)
 - Reactive
- LM:
 - Load balancing, VM relocation, host maintenance
 - VM is live migrated with no downtime to services
 - Proactive

Primary Reason For Virtualisation?

- The Great Big Hyper-V Survey (<http://greatbighypervsurvey.com/>)
- 2011:
 1. Reduce costs – 29.41%
 2. Increase flexibility/agility – 28.19%
- 2012 (unpublished):
 1. Increase flexibility/agility – 25.7%
 2. Reduce costs – 23.7%
- The role of Live Migration is flexibility/agility

Live Migration Desires

- Live Migration of VHDs
- Perform as many simultaneous Live Migrations as hardware and network will allow
- Queue up Live Migrations
- Break down barriers to live migration for more flexibility
- Enter Windows Server 2012

Improving Live Migration

Redefining Live Migration

- Used to mean one thing, but now it's something BIGGER
- *“The ability to move virtual machines and/or their storage from one location to another without impacting service uptime”*

2 Promises

- If a problem occurs during LM, e.g. network break, then the VMs will remain operational
- New features will not prevent LM:
 - SR-IOV
 - Virtual Fibre Channel
 - Can VMware say the same thing?

Live Migration Requirements

1. Windows Server 2012 (Std/DC) Hyper-V or Hyper-V Server 2012
 - (NOT legacy versions)
 - Or later
 - Matching versions *
2. Common AD forest membership
3. At least 1 GbE connection between source and destination

WS2012 Live Migration Improvements

- Memory synchronisation process is improved to reduce iterations
- Recently inactive memory is copied first
- Recently active memory is synchronised after this
- Results in fewer iterations:
 - Makes each individual VM LM quicker
- Supports VMs on SMB 3.0 storage

More Speed!

WS2012 – Faster Live Migration


- Compression
 - Makes the most of existing network/hardware investment
 - Uses spare CPU cycles – CPU monitored by Hyper-V
 - Over 2x improvement in live migration time
 - Enabled by default
- SMB
 - For 10 GbE or faster networks
 - Make the most of network investment for SMB 3.0
 - Converged networks
 - Enables high-speed & low impact live migrations
 - Uses SMB Multichannel to leverage multiple interfaces
 - Uses SMB Direct to minimise impact on hosts' CPUs

Demo: Live Migration Options

Types of Live Migration

Live Migration & SMB Storage

Storage handle moved



Demo: Live Migration & SMB Storage

PowerShell

- Everything you can do in the GUI, you can do in PowerShell
 - Get-Help Move-VM
- LM a VM without moving storage:
 - `Move-VM VM01 -DestinationHost Demo-Host2`
 - `Get-VM -Name SMB -ComputerName Demo-Host2 | Move-VM -DestinationHost Demo-Host1`

Storage Live Migration

- Relocate VHDs (and VHDX) to a new location
 - Folder, LUN, CSV, server, SAN
- Zero downtime to services
- It does increase IOPS
- Be careful when relocating to different folder on same LUN:
 - Roughly double the IOPS created
 - Will temporarily require double the space – more if using Dynamic VHD and it grows during LM

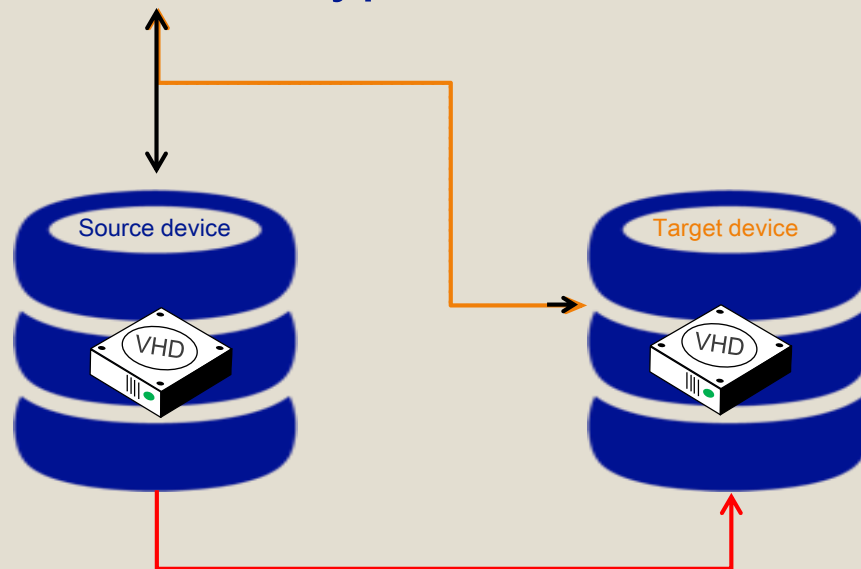
Storage Live Migration

Reads and writes go to new destination VHD



Virtual machine

Computer
running
Hyper-V



Benefits of Live Storage Migration

- Abstracted storage (VHD/VHDX) can be relocated with zero impact to service:
 - More space
 - Change storage tier of VM
 - Move to newer storage
- Enables flexibility and self-service (cloud)
- Does not support Passthrough (raw) Disk

Demo: Storage Live Migration

Demo: Building a Hyper- V Cluster ... The Better Way

PowerShell

- Move a VM from one location to another:
 - Move-VMStorage “*VM01*” –DestinationStoragePath “*E:\Virtual Machines*”

Thinking Outside The ~~Box~~ Cluster

- Not every host is clustered
 - 66% have non-clustered hosts
 - 66% have clustered hosts
 - Meaning 33% have both kinds of host
 - *From The Great Big Hyper-V Survey of 2011*
- For example:
 - Do you need clustered hosts for low cost VDI?
 - Hosting companies have to pass on the cost of idle cluster capacity

Live Migration is not Clustering

- You do not need a Hyper-V cluster to do Live Migration
- Now we can LM VMs around the datacenter:
 - Host to host in the same cluster (as before)
 - Non-clustered host to non-clustered host
 - Cluster to cluster
 - Non-clustered host to cluster, and vice versa

Alternative Host Architecture

- If you do not need HA:
 - Store VMs on SMB 3.0 file shares
 - Use non-clustered hosts
- Can still do Live Migration
 - Load balancing
 - Easy host maintenance
 - Host aging/replacement
- It is not a cluster, so no HA

Taking it one Step Further

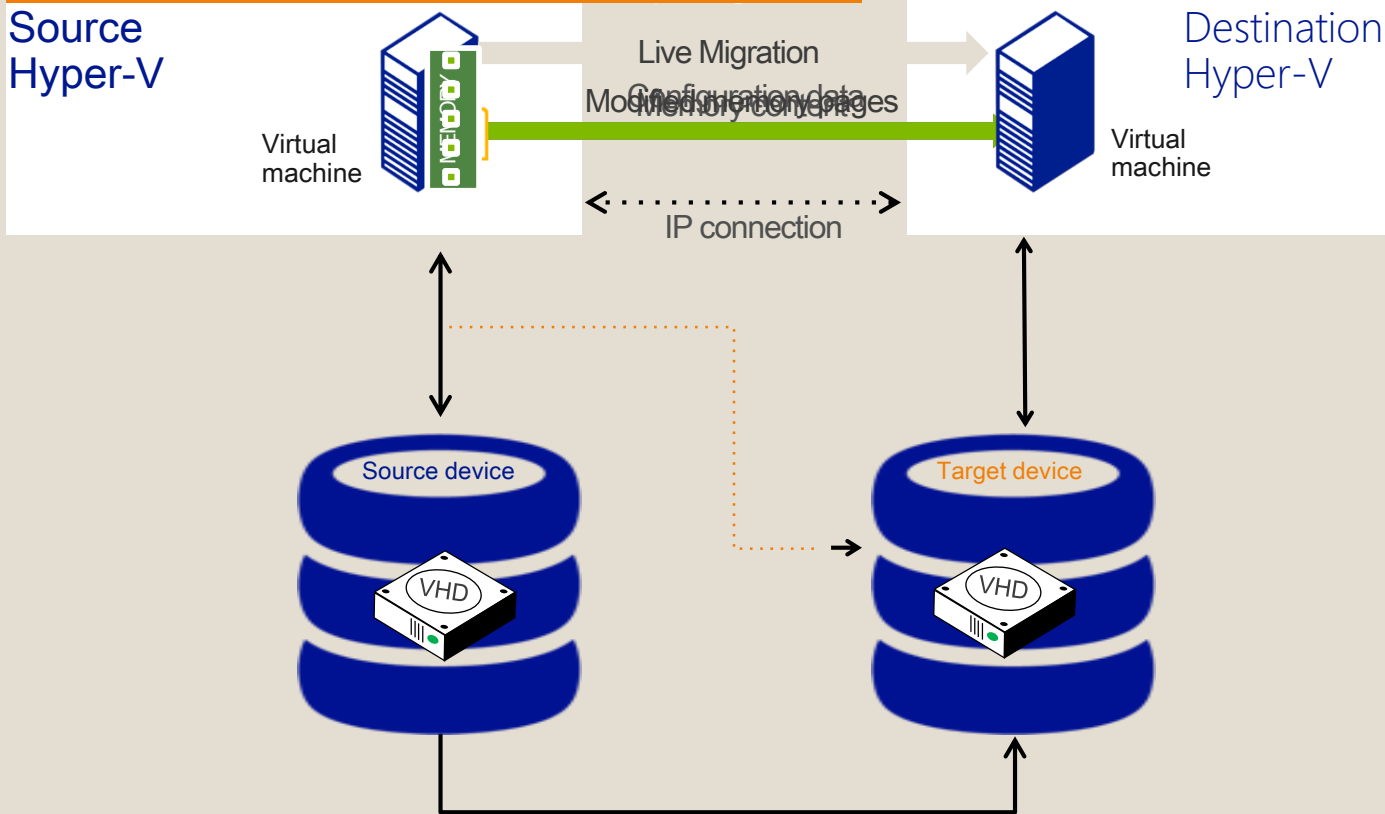
- Source host and destination host might not have shared storage
 - A 2U host can have over 1 TB+ RAM and lots of internal disk
 - Migrating VMs between two different clusters
- WS2012 Hyper-V is breaking down lots of barriers
 - We don't even need shared storage for Live Migration
 - Added flexibility
 - More cost reductions and options

Shared-Nothing Live Migration

- Move storage and VM between two hosts not in the same cluster
 - Therefore the VM is not on common SAN or SMB 3.0 share
- Requires storage relocation
- Process:
 - Storage Live Migration to relocate storage (first 50% of progress)
 - Copy/synchronise memory and switch VM execution (second 50% of progress)

Shared-Nothing Live Migration

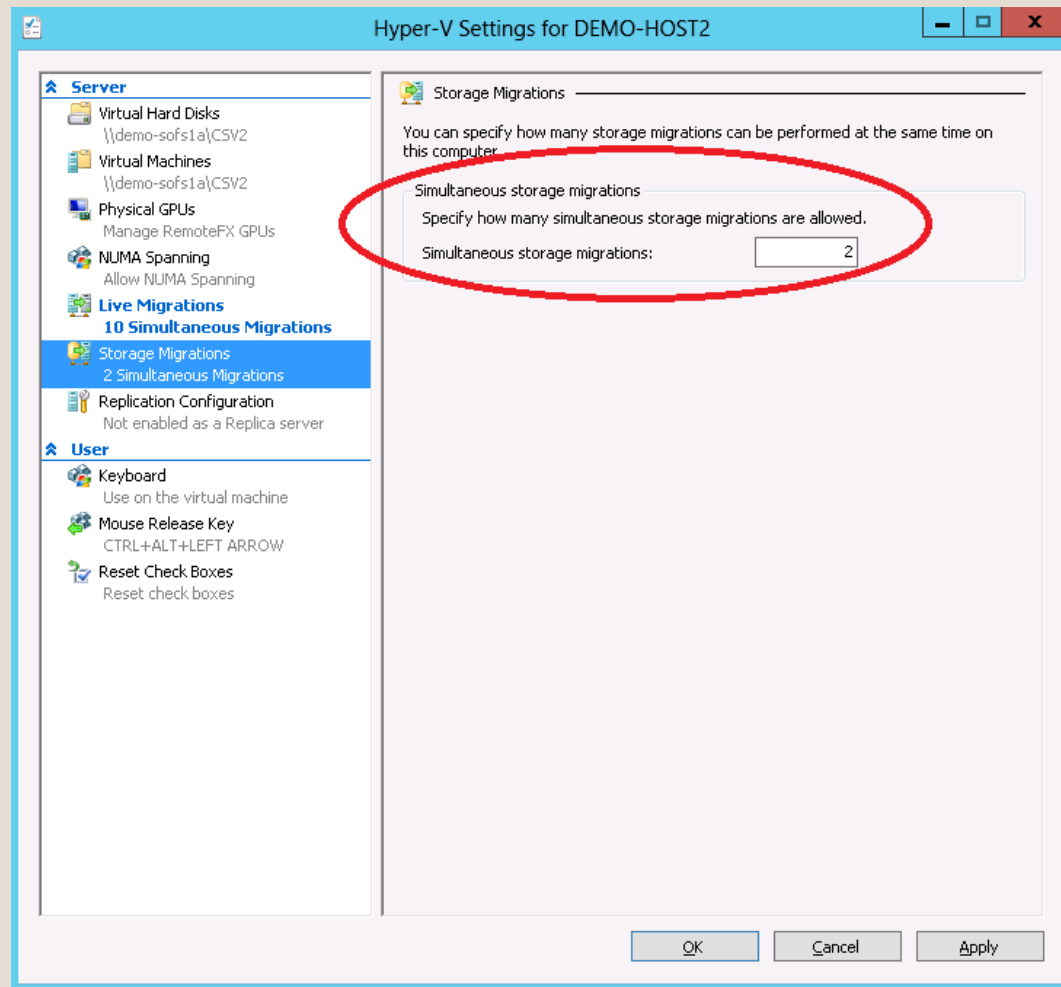
Live Migration Completes



Simultaneous Live Migration

- “Unlimited” simultaneous Live Migrations
- Set on a per host level (Host Settings)
- Unlimited?
- In theory: yes
- In practice: no
 - VMWP.EXE will consume host CPU
 - There is only ever so much network bandwidth

Configuring Live Migration



Configuring Live Migration PowerShell

- Enable-VMMigration
- Set-VMMigrationNetwork 192.168.2.0
- Set-VMHost –
VirtualMachineMigrationAuthenticationType
Kerberos

Cluster Live Migration

The screenshot displays the Failover Cluster Manager interface with the 'Live Migration Settings' dialog box open. The dialog box is titled 'Live Migration Settings' and has a tab labeled 'Networks for Live Migration'. The instructions inside the dialog state: 'Select one or more networks for virtual machines to use for live migration. Use the buttons to list them in order from most preferred at the top to least preferred at the bottom.'

The 'Networks for Live Migration' list contains the following entries:

Name	Selected
<input checked="" type="checkbox"/> vEthernet (Host-LiveMigration)	Yes
<input type="checkbox"/> vEthernet (Host-Cluster)	No
<input type="checkbox"/> iSCSI A	No
<input type="checkbox"/> vEthernet (Host-Parent)	No
<input type="checkbox"/> iSCSI B	No

Buttons for 'Up' and 'Down' are located to the right of the list. At the bottom of the dialog box are 'OK', 'Cancel', and 'Apply' buttons. In the background, the Failover Cluster Manager tree view shows the 'Networks' folder under 'demo-host1' circled in red.

What If?

- Can I LM 1000 VMs at one?
 - In my experience, max simultaneous is capped by Hyper-V based on hardware
 - VMWP.EXE eating host CPU & bandwidth
 - You can set it, but you might not reach it
- What if 2 hosts have different limits set?
 - The lowest of the 2 values will be used at the maximum number of simultaneous LMs

Failover Cluster Improvements

- LM to “Best Possible Host”
- Now supports queuing LMs
- For example:
 - Running 100 VMs on a clustered host
 - Configured to do 10 simultaneous LMs
 - Can select all VMs and start LM
 - 10 start to LM
 - Other 90 VMs wait their turn to LM

Simultaneous Live Migration FUD

- VMware used to say “More LMs is good”
- Now they say: “Are more LMs of any use?”
- Reality: Yes
 - We have bigger hosts (up to 4 TB RAM), typically up to 256 GB
 - It takes time to copy/synchronize 256 GB of changing RAM
 - Do you want to wait a weekend to do host maintenance?
- Enterprise
 - Put in bigger (10 GbE, Infiniband, etc) converged fabrics
 - Live migrate more VMs at once for quicker maintenance

Demo: Simultaneous Live Migration

Live Migration & Kerberos

Kerberos Complications

- Kerberos is used by AD for authentication (who are you?) and authorization (what can you do?)
- Prevents man-in-the-middle, such as:
 - You log into PC-A and use Hyper-V Manager
 - You start to manage a VM on Host-B
 - You want to relocate it to Host-C or FileServer-D
 - Host-B becomes the man-in-the-middle & move is blocked
- Causes issues for non-clustered hosts
- Constrained delegation required to fix this

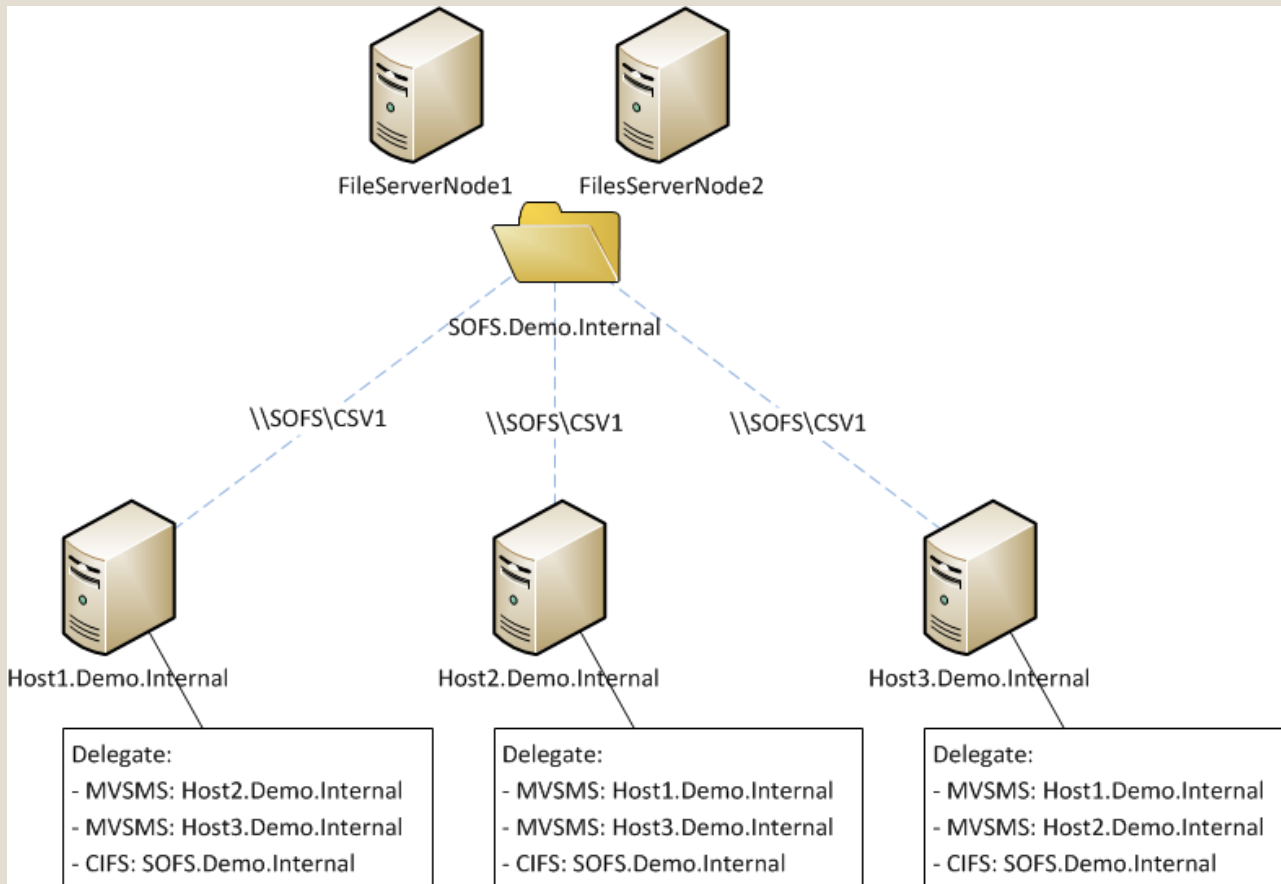
Constrained Delegation Scenarios

For Non-Clustered Hosts & SMB 3.0 Storage:

On non-clustered hosts

1. Delegate all other possible hosts for *Microsoft Virtual System Migration Service*
2. Delegate the file server for *CIFS*

Constrained Delegation Scenarios



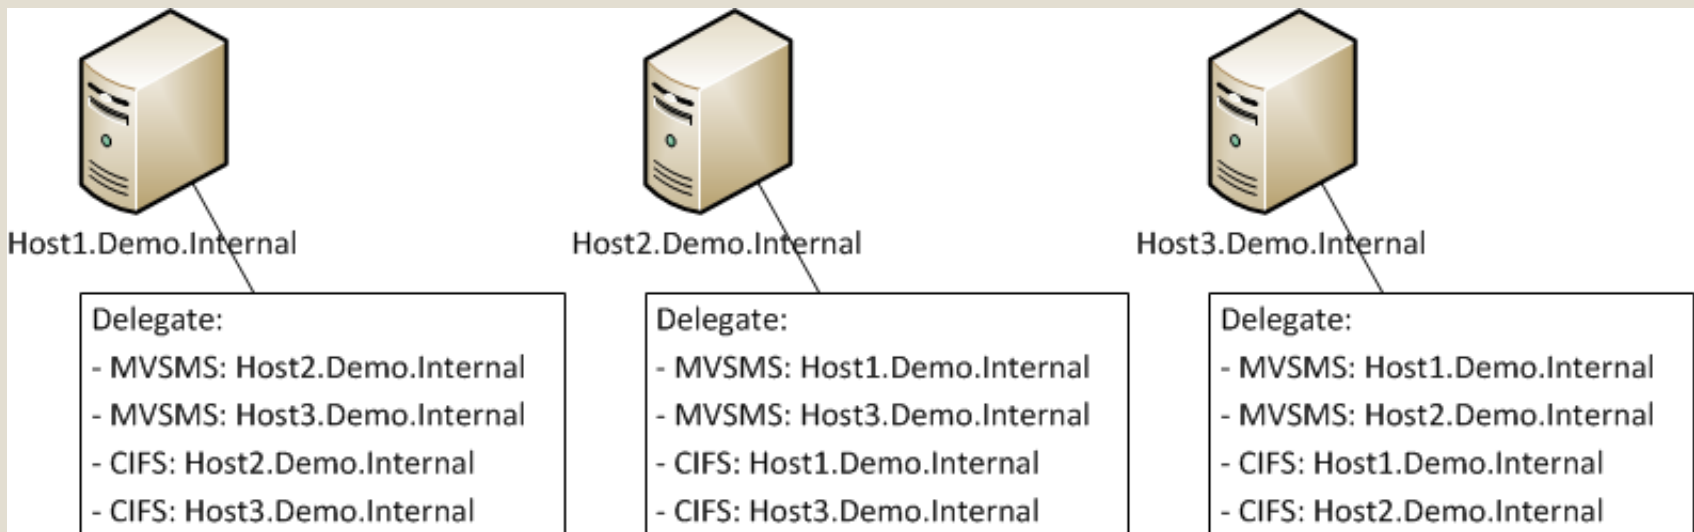
Constrained Delegation Scenarios

For Shared-Nothing Live Migration:

On each host, delegate all other possible hosts:

1. *CIFS*
2. *Microsoft Virtual System Migration Service*

Constrained Delegation Scenarios



Demo: Kerberos Constrained Delegation

Constrained Delegation

1. Enable Advanced view in ADUC
2. Open properties of each host/file server
3. Browse to Delegation tab
4. Add hosts and delegate them for one or both:
 1. CIFS
 2. Microsoft Virtual System Migration Service
5. Reboot the host you modified

DEMO-HOST1 Properties

Location | Managed By | Object | Security | Dial-in | Attribute Editor

General | Operating System | Member Of | Delegation | Password Replication

Delegation is a security-sensitive operation, which allows services to act on behalf of another user.

Do not trust this computer for delegation to any service (Kerberos only)
 Trust this computer for delegation to any service (Kerberos only)
 Trust this computer for delegation to specified services only

Use Kerberos only
 Use any authentication protocol

Services to which this account can present delegated credentials:

Service Type	User or Computer	Port	Service Name
cifs	DEMO-HOST2		
Microsoft Virt...	DEMO-HOST2		

Expanded

OK Cancel Apply Help

Add Services

To allow services to act on behalf of users or computers...

To select one or more services to which this account can present delegated credentials...

Available services:

- alerter
- appmgmt
- browser
- cifs**
- cisvc
- clipsrv
- dcom
- dhcp

Expanded

Add... Remove

DEMO-HOST1 Properties

Location | Managed By | Object | Security | Dial-in | Attribute Editor

General | Operating System | Member Of | Delegation | Password Replication

Delegation is a security-sensitive operation, which allows services to act on behalf of another user.

Do not trust this computer for delegation to any service (Kerberos only)
 Trust this computer for delegation to any service (Kerberos only)
 Trust this computer for delegation to specified services only

Use Kerberos only
 Use any authentication protocol

Services to which this account can present delegated credentials:

Service Type	User or Computer	Port	Service Name
cifs	DEMO-HOST2		
Microsoft Virt...	DEMO-HOST2		

Expanded

OK

Constrained Delegation Script

- That's all a bit mousey
- You can script all that:
<http://tinyurl.com/HyperVKCD>

SMB Delegation Got Much Easier

- Delegation for SMB can be done using WS2012 R2 AD PowerShell cmdlets
- Enable-SmbDelegation –SmbServer "FileServer01" –SmbClient "HVSVR01"
- **Does not** require a reboot of the Hyper-V host

Live Migration Across Network Footprints

Live Migrate Between Networks

- Large data centres can have many independent networks
 - For example, a hosting company replacing network footprint
- This changes the networks that VMs are on
 - Moving VM to different network
 - Requires automated (DHCP!) or manual change
 - DNS takes time to propagate
 - Sadly, some services are IP dependent & require change

Network Virtualization

- Just like VMs virtualise hardware, Network Virtualization abstracts IP subnets
 - Also known as Software Defined Networking (SDN)
- Engineers create policies to abstract VM IP from actual physical IP
 - Best done using System Center 2012 SP1
 - Generic Routing Encapsulation (GRE) or IP Address Rewrite
- No boundaries for LM left in the data centre

Demo: Comparing Live Migration Speeds

Using Live Migration for Zero Downtime Upgrades

Cross-Version Live Migration

- Make it easier to migrate VMs to newer version of Hyper-V
 - New zero service downtime migration option
- Live migrate virtual machines from Windows Server 2012 to Windows Server 2012 R2
 - One-way operation
 - Only from WS2012 to WS2012 R2
- Includes support for:
 - Live Migration of VMs on SMB 3.0 storage
 - Shared-nothing live migration

Cross-Version Compatibility

- Increased compatibility for VMs on Windows Server 2012 R2 from Windows Server 2012 R2
 - Saved states
 - Checkpoints (formerly known as Hyper-V Snapshots)
- Less clean-up work (loss for devs/testers) when:
 - Upgrading hosts
 - Using Cross-Version Live Migration