

# What's New in Windows Server 2012 R2 Networking


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

# Agenda

- Administration
- Storage Networking
- Live Migration
- Virtual Machine Scalability
- NIC Teaming
- Networking The Cloud
- A New Data Centre

# Administration

# Diagnostics

- Ping – tried and tested
  - But a pain to script with
- WS2012 R2 adds Test-NetConnection
- Simple PowerShell test with results that you can easily test

```
PS C:\Windows\system32> $Result = Test-netconnection demo-host1
PS C:\Windows\system32> $Result

ComputerName      : demo-host1
RemoteAddress     : 192.168.1.51
InterfaceAlias    : vEthernet (Management)
SourceAddress     : 192.168.1.53
PingSucceeded     : True
PingReplyDetails (RTT) : 1 ms
```

# IP Address Management (IPAM)

- IPAM is a feature of WS2012 and WS2012 R2
- Automated solution to replace Excel and Ping address "management"
- Improvements in WS2012 R2:
  - Role-based access control
  - Integrate with SCVMM 2012 R2 to manage Hyper-V networking
  - Enhanced DHCP management
  - Use SQL Server as an alternative to Windows Internal Database
  - Enhanced PowerShell support

# Storage Networking



# Windows Azure IaaS

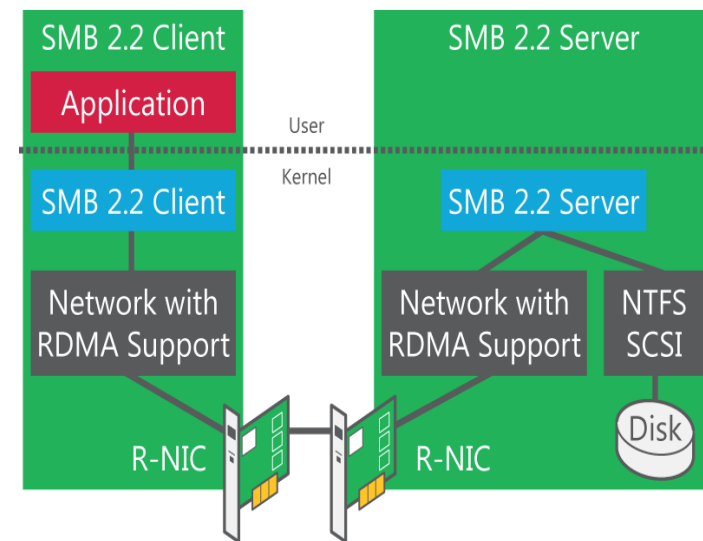
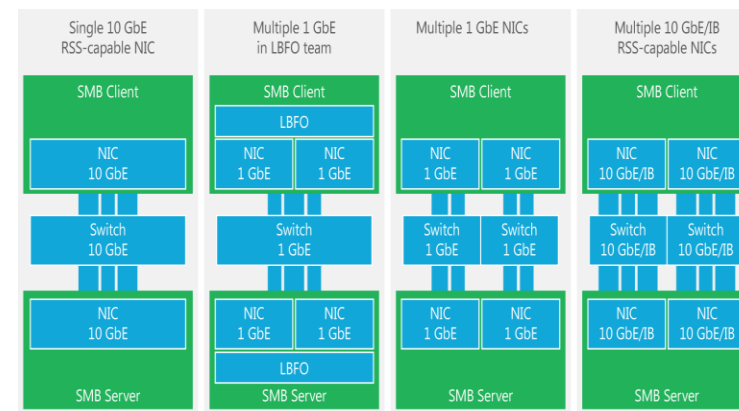
- Windows Azure uses the same Hyper-V virtualization service built-into Windows Server 2012
- Complete virtual machine compatibility between on premise Hyper-V and Azure IaaS

# SMB 3.0

- Server Message Block (SMB):
  - Old protocol
  - Use for client/server file sharing
- Reinvented for WS2012:
  - Called SMB 2.2 in preview & beta
  - Renamed to SMB 3.0 in RC & RTM
- Designed to rival and beat legacy protocols for applications accessing networked storage:
  - iSCSI
  - Fiber Channel
- SMB 3.0 is Microsoft's enterprise *data* protocol

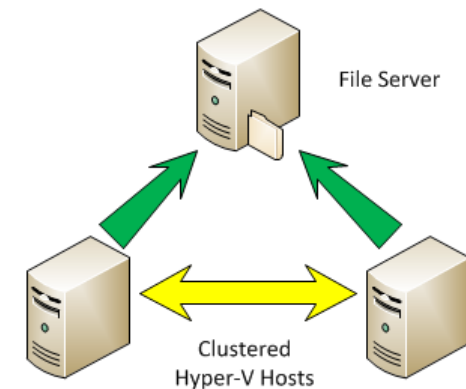

# What Made SMB 3.0 So Good?

- SMB Multichannel
  - Make the most of 1 or more NICs
  - Auto detection, unlike MPIO
  - Used by cluster Redirected IO
  - Use SMB Multichannel Constraints to control NIC selection
- SMB Direct
  - Lots of bandwidth = lots of H/W interrupts = high CPU utilisation
  - Remote Direct Memory Access (RDMA) capable NICs (rNICs)
  - Reduce CPU usage, improve performance
  - Increase scalability of file servers N/W



# Hyper-V Over SMB

- WS2012 (and later) Hyper-V hosts can store VMs on WS2012 (and later) shared folders
  - Shares permissioned for hosts/clusters and administrators
  - VMs stored on a UNC path
- Leverages SMB Multichannel and SMB Direct
- Completely supported for clusters
- MSFT's preferred storage protocol
  - Faster & cheaper
- Enables more flexibility than block storage
  - Live Migration of VMs on shares
  - A share can be used by more than one host or cluster



# Changes to SMB 3.0 in WS2012 R2

- Technically it's SMB 3.02 but still called SMB 3.0
- Improved performance for small I/O workloads
- Improved error messages
- You can uninstall SMB 1.0 (used by Windows XP and W2003)
- New -SmbDelegation PowerShell cmdlets in the AD Module for easier share delegation to hosts (no reboots!)
- There is more to WS2012 R2 SMB 3.0 ...

# Live Migration


# Live Migration

- Technically it's SMB 3.02 but called SMB 3.0
- Improved performance for small I/O workloads
- Improved error messages
- You can uninstall SMB 1.0 (used by Windows XP and W2003)
- New -SmbDelegation PowerShell cmdlets in the AD Module for easier share delegation to hosts (no reboots!)


# Live Migration

- Ability to move a virtual machine's state from one host to another with no perceivable downtime to guest services
- WS2012 added ability to move VMs' file locations
  - Live Storage Migration
  - Shared-Nothing Live Migration

Memory



Host 1



Host 2



# The Status of Live Migration

- Most common type of networking used in virtualization
  - 1 GbE
  - FCoE with some bandwidth used for host networking
- WS2012 optimized the ordering of initial page copies
- More could be done to make the most of existing hardware
- We need faster Live Migration/vMotion:
  - Bigger hosts than ever (up to 4 TB RAM)
  - Bigger VMs than ever (up to 1 TB RAM)
  - Unplanned host maintenance – System Center detects hardware failure
  - Planned host maintenance – perform Cluster Aware Updating quickly

# Compressed Live Migration


- Enabled by default in WS2012 R2
- Uses spare CPU time to compress pages
  - Hyper-V monitors host CPU
  - Automatically switches to/from TCP/IP (legacy) Live Migration as CPU is required by host/VMs
- 2x (or more) improvement of Live Migration time
  - Depends on RAM contents of guest OS
- Make the very most of existing hardware
  - No changes required
- Should be used on < 10 Gbps networking
  - 1 GbE
  - 10 Gbps FCoE: converged storage & host networking

# SMB Live Migration

- Makes the most of an investment in > 10 Gbps networking
  - Example: Converged cluster & Live Migration networking with SMB 3.0 storage
- Uses features of SMB 3.0 for Live Migration
- SMB Multichannel:
  - Aggregate networks for more bandwidth
  - Live Migration at 20 Gbps, 80 Gbps, 128 Gbps or faster!!!
- SMB Direct:
  - More speed!
  - Real benefit is to offload CPU work to RDMA capable NIC

# Converging SMB 3.0

- A few will install dedicated 10 Gbps or faster networks for Live Migration
- Most will converge Live Migration & storage
- Traditional QoS just sees SMB 3.0 so how do you guarantee bandwidth to storage *AND* Live Migration
- Add-WindowsFeature FS-SMBBW
- Set-SmbBandwidthLimit
  - VirtualMachine: Hyper-V over SMB storage traffic
  - LiveMigration: Live Migration over SMB traffic
  - Default: All other types of SMB traffic
  - Set-SmbBandwidthLimit -Category LiveMigration -BytesPerSecond 4GB



# Live Migration Recommendations

- Networks slower than 10 Gbps:
  - Use Compression (default)
  - Remember that Fibre takes bandwidth from FCoE
- 10 Gbps or faster
  - Use SMB

# Demo: Live Migration

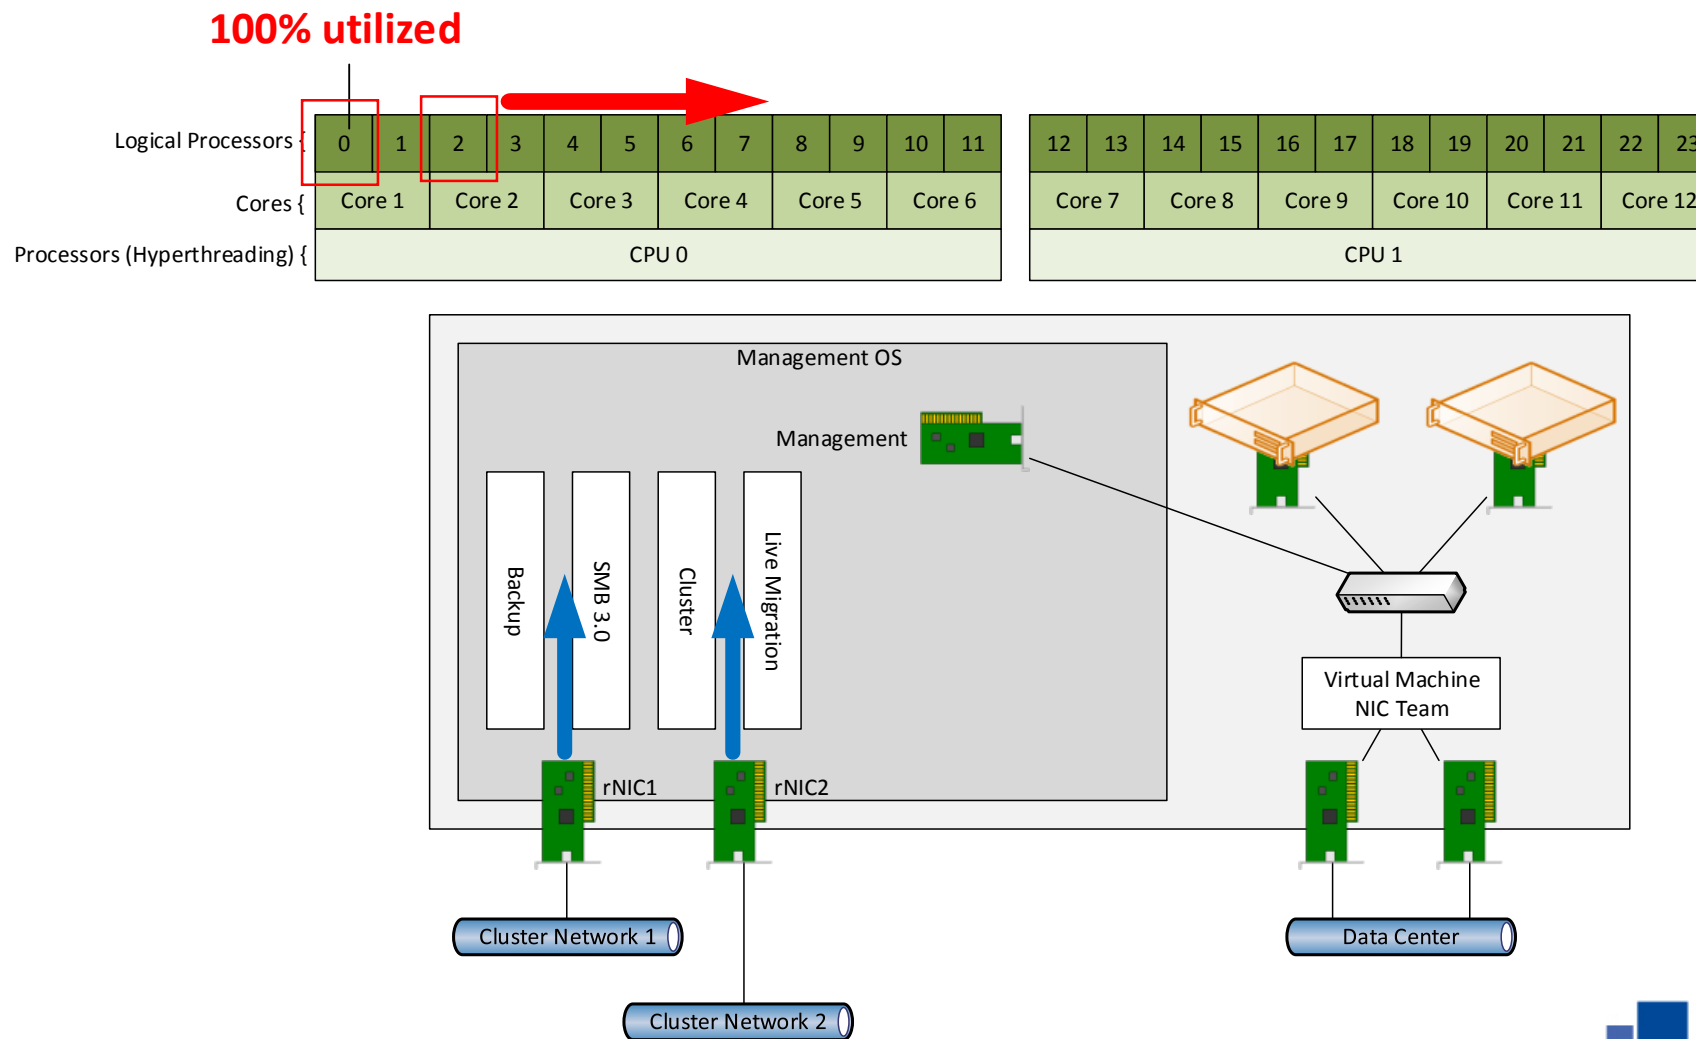
# Virtual Machine Scalability

# Hyper-V VMs Are Enterprise Ready

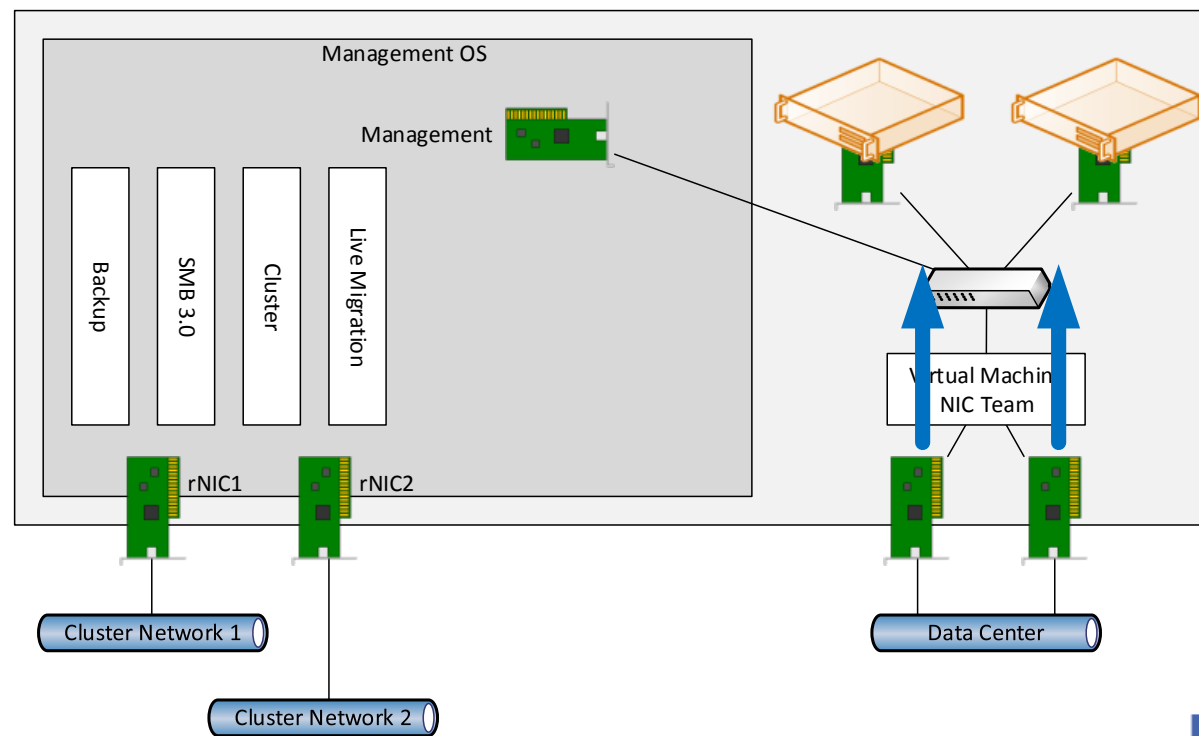

- WS2012 Maximum virtual machine scalability:
  - 64 virtual processors
  - 1 TB RAM
  - 256 \* 64 TB VHDX files
  - VM aware NUMA
- We can truly deploy huge services in Hyper-V VMs
- But what about the networking?
- Can we push huge workloads into those VMs?
- Back to basics ... before we solve the problem 😊



# Receive Side Scaling (RSS)

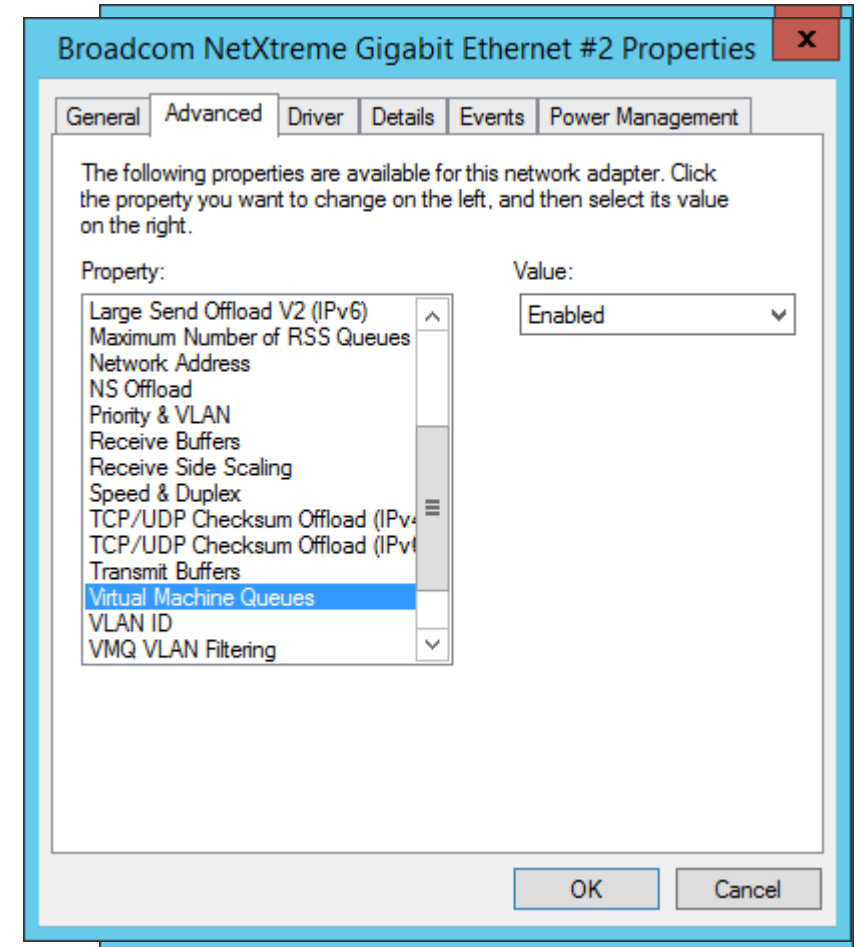


# Dynamic Virtual Machine Queue (dVMMQ)



# RSS & dVMQ

- Consult your network card/server manufacturer
- Upgrade firmware & drivers to avoid issues
- Can use Get- Set-NetAdapterRSS to configure
- RSS and DVMQ are incompatible on the same NIC so design hosts accordingly



# Virtual RSS (vRSS)

- Added in WS2012 R2
- RSS provides extra processing capacity for inbound traffic to a physical server
- Using virtual processors beyond CPU 0
- vRSS does the same thing in the guest OS of a VMM
- Allows inbound networking to VMM to scale out
- Requires VMs with > 1 virtual processors
- DVMQ must be enabled on physical NIC(s) used by virtual switch
- Enable RSS in the advanced NIC properties in the VM's guest OS

# Demo: vRSS

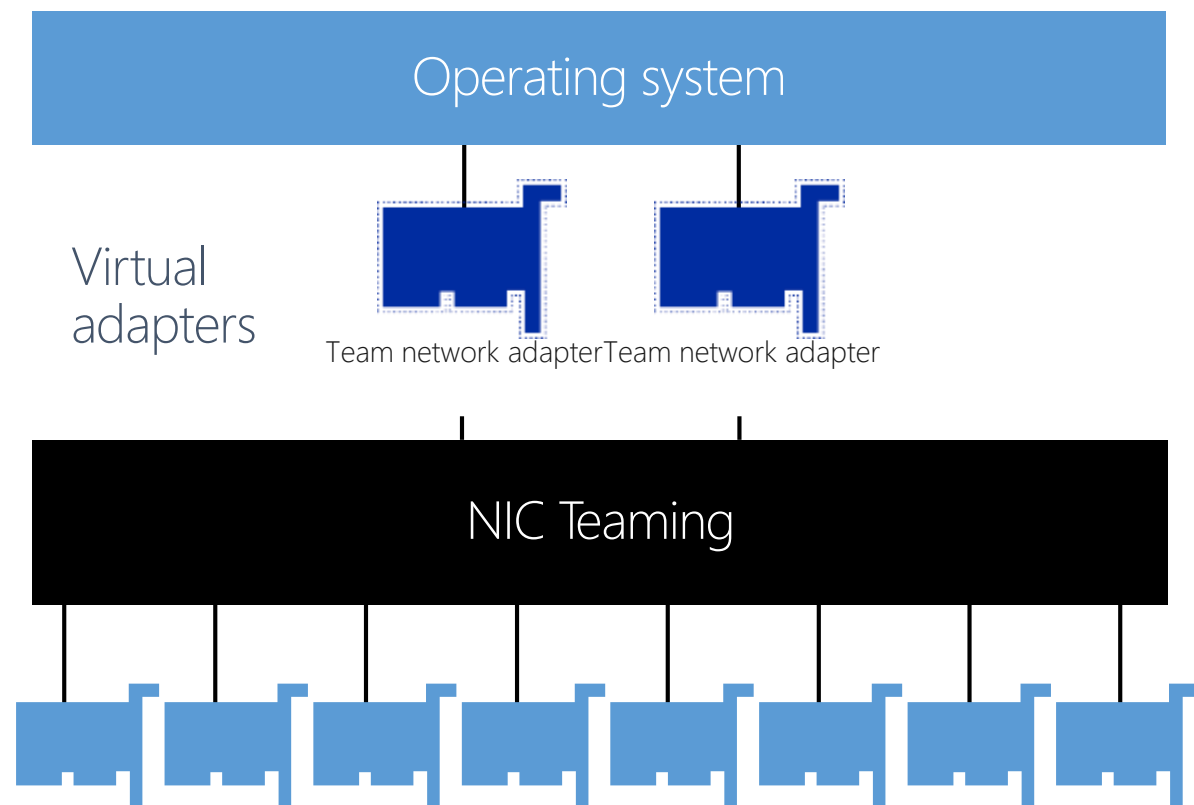
# NIC Teaming

# NIC Teaming in WS2012

- Built-in and supported NIC teaming added in WS2012
  - Up to 32 NIC (physical) of different manufacturers in a single team
  - 2 vNICs in a single team in a guest OS
- LBFO
  - Load balancing: bandwidth aggregation
  - Failover: automatic NIC fault tolerance
- 2 load distribution modes
- Address Hashing
  - Normally used for non-VM networking
  - Spreads data across physical NICs based on destination address/port hashing
- Hyper-V Port
  - Normally used when connecting a virtual switch
  - Limits a VM's NIC to bandwidth of a single physical NIC

# Dynamic Load Balancing

- New load balancing algorithm in WS2012 R2
- Think of it as the best of:
  - Hyper-V Port
  - Address Hashing
- Uses “flowlets” to hash outbound data streams across physical NICs in the host’s team
- Inbound traffic “affinitized” to a single physical NIC (still can failover)
- Default load balancing algorithm

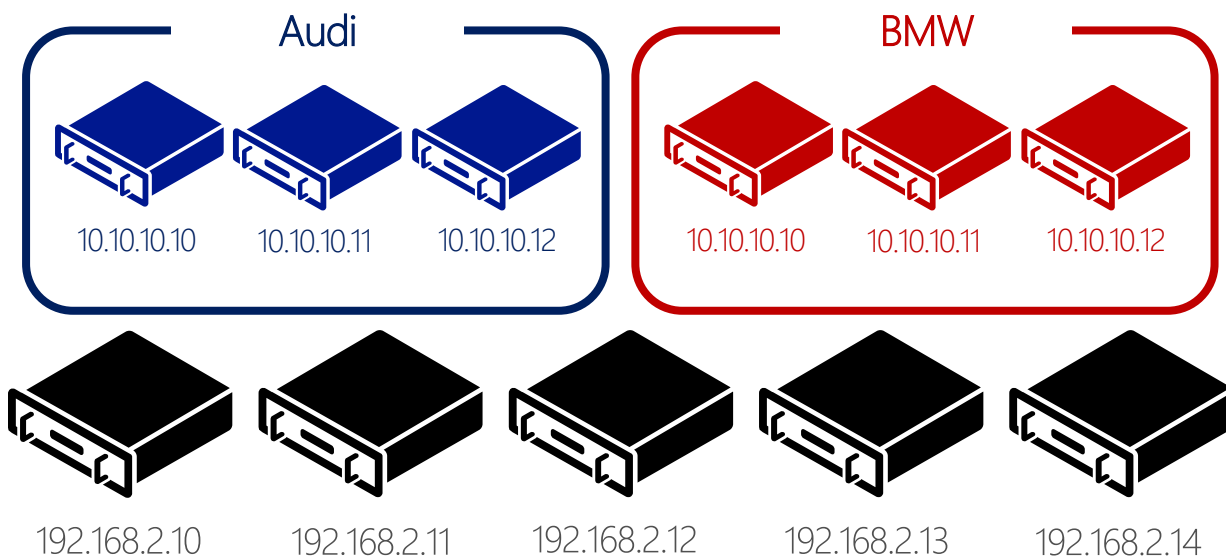




# Networking The Cloud

# Software Defined Networking (SDN)

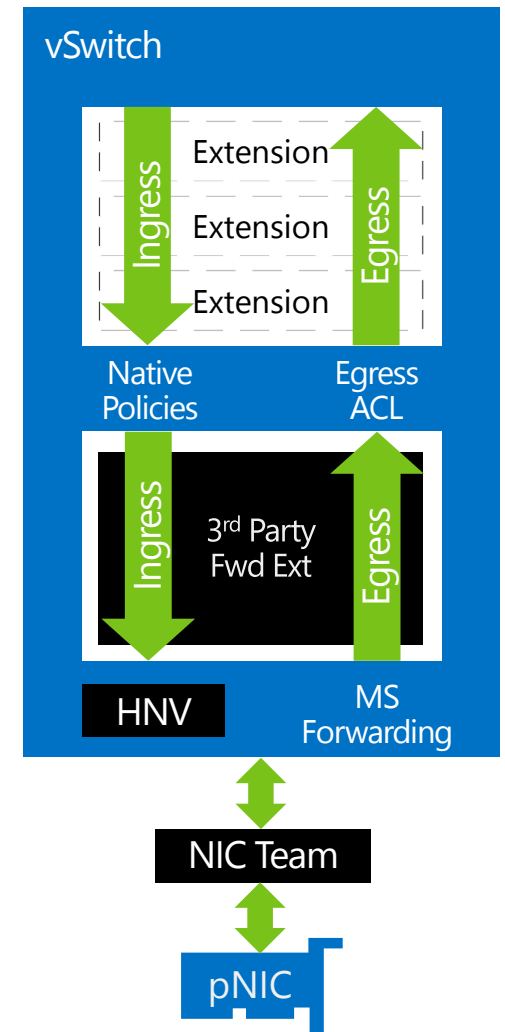
- Secure Isolation for traffic segregation, without VLANs
- VM migration flexibility
- Scalable beyond 4096 networks
- Key Concepts
  - Provider Address – Unique IP addresses on hosts used for encapsulation
  - VM Networks – Boundary of isolation between tenants
  - VM Subnets – routed subnets within VM Networks
  - Customer Address – VM Guest OS IP addresses within the VM Networks
  - Policy Table – maintains relationship between different addresses & networks



Network/VSID	Provider Address	Customer Address
Blue (5001)	192.168.2.10	10.10.10.10
Blue (5001)	192.168.2.10	10.10.10.11
Blue (5001)	192.168.2.12	10.10.10.12
Red (6001)	192.168.2.13	10.10.10.10
Red (6001)	192.168.2.14	10.10.10.11
Red (6001)	192.168.2.12	10.10.10.12

# Improvements to HNV

- NVGRE Gateway is built into Windows Server
  - Deploy a host cluster with VMs (guest cluster) into edge network
  - VMM 2012 R2 service template for the guest cluster
- HNV moved into the virtual switch
  - Third party extensions (e.g. Cisco Nexus 1000V) can manage in/out traffic
- Port ACLs compatible with HNV
- HNV can learn guest assigned IP addresses (tenant DHCP)
- Support for NVGRE offload
  - Requires special NICs



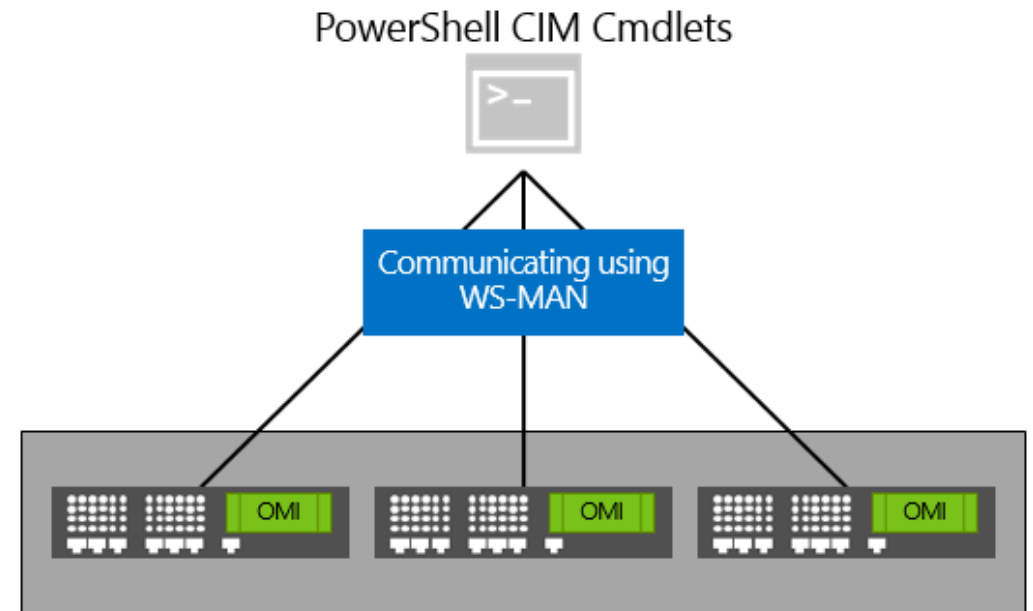
# Port ACLs

- A basic form of firewall rule that is applied by the virtual switch
- Added in WS2012 to control traffic between virtual NICs
  - Specify source/destination address (IPv4/IPv6/MAC) of virtual NIC or network
  - Allow/block/measure actions
- Extended Port ACLs added in WS2012 R2:
  - Can specify a port number for granular control
  - Stateful uni-conditional rules
- Compatible with Hyper-V Network Virtualization

# A New Data Centre

# Data Center Abstraction Layer

- A software-defined data centre:
  - Bare metal computer deployment
  - SMI-S management of block storage
  - Software-defined storage
  - Top-of-rack (TOR) switch management
- Standards-based CIM model
- Switches running Open Management Infrastructure (OMI)
- PowerShell cmdlets
- Common management across vendors
- Automate common tasks
- Networking “just works”
- Not just marketing to Microsoft



# Summary

- Easier administration
- SMB 3.0 expanded to improve infrastructure flexibility
- Hyper-V networking can scale to enterprise levels
- The Cloud OS has been improved based on Microsoft experience in Azure and customer feedback
- Software-defined data centre is more than marketing for Microsoft

# Dankeschön!

**Aidan Finn**

@joe\_elway

[www.aidanfinn.com](http://www.aidanfinn.com)

Petri IT Knowledgebase