

Microsoft®

tech·days

15 From the Desktop to the Cloud
Years of Turning Vision into Value

Hong Kong|2012



Windows 8 Hyper-V Scalability

Bryon Surace, Senior Program Manager
Microsoft Corporation



About this Presentation

- Statements about capabilities or benefits are subject to change
- Packaging and licensing have not yet been determined
- Any concepts shown are for illustration purposes only

Disclaimer:

This presentation contains preliminary information that may be changed substantially prior to final commercial release of the software described herein.

The information contained in this presentation represents the current view of Microsoft Corporation on the issues discussed as of the date of the presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of the information presented after the date of the presentation. This presentation is for informational purposes only.

MICROSOFT MAKES NO WARRANTIES, EXPRESSED, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.

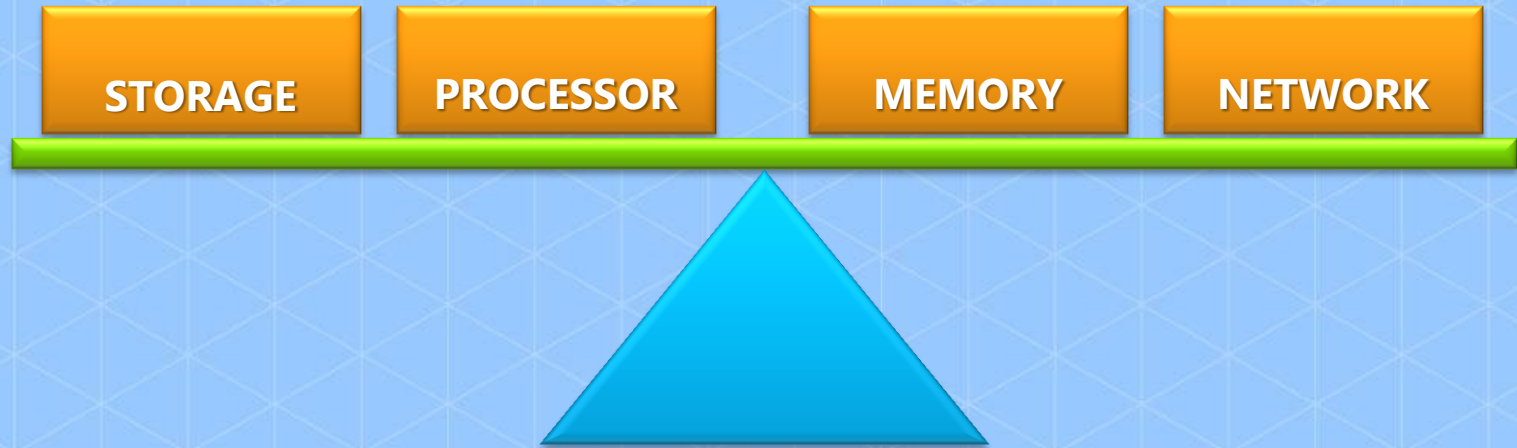
Microsoft may have patents, patent applications, trademarks, copyrights or other intellectual property rights, covering subject matter in this presentation. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this information does not give you any license to these patents, trademarks, copyrights or other intellectual property rights.

Design Principals

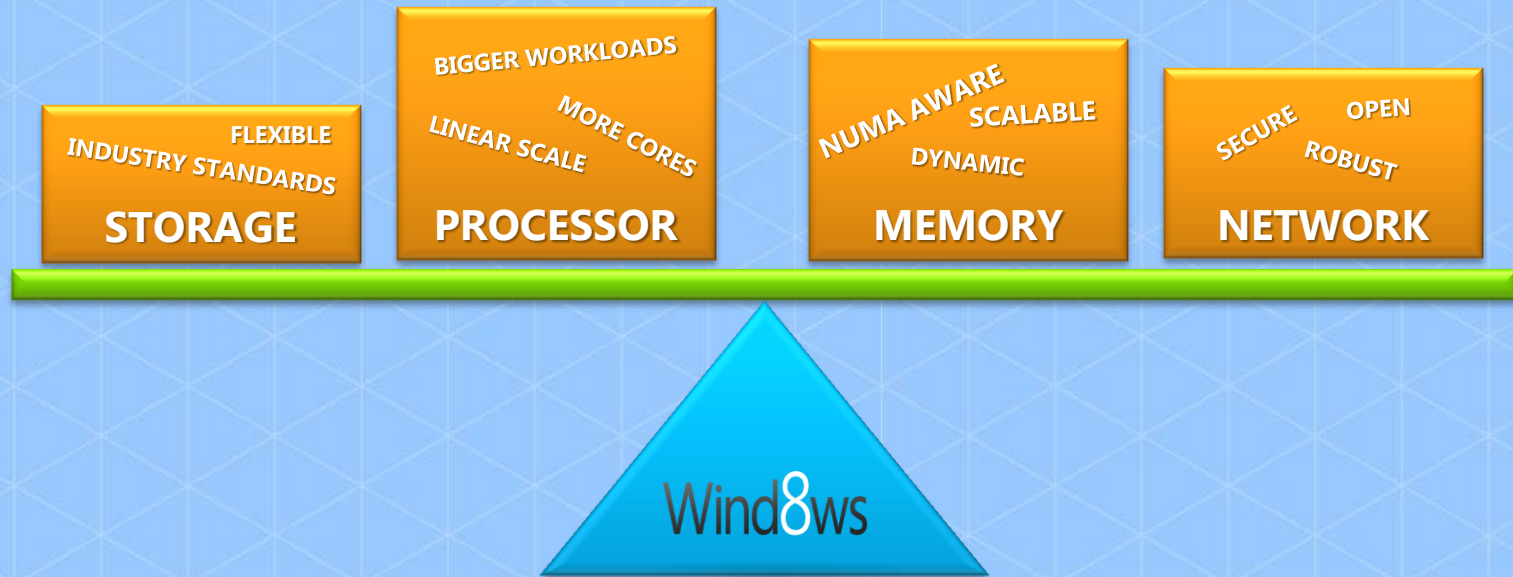
MISSION CRITICAL WORKLOADS

- Scale isn't just more virtual processors, it's **balance**
 - Compute
 - Memory
 - Network I/O
 - Storage I/O
- Server **resiliency, uptime** and **security** is paramount
 - More resources and capacity than ever
- **Maximize** hardware usage
- Virtualization benefits such as Live Migration should just work
 - **No tradeoffs**
- Virtual Machine performance must increase with cores
- Virtualize workloads considered “non-virtualizable”

Balanced Scale Up



Balanced Scale Up



Balanced Scale Up

CPU AND MEMORY

CPU Scale Comparison

	Windows Server 2008	Windows Server 2008 R2	Windows Server 8 Developer Preview
HW Logical Processor Support	16 LPs	64 LPs	160 LPs
Virtual Machine Processor Support	Up to 4 VPs	Up to 4 VPs	Up to 32 VPs
VP:LP Ratio	8:1	8:1 for Server 12:1 for Client (VDI)	Hyper-V scales with hardware capabilities
Max VMs/VPs per Host	128	384/512	1024/2048
Latest CPUs Instructions	Yes – at the time	Yes – at the time	Yes – now includes AMD: XSAVE, AES INTEL: AES/NI, PCLMULQDQ, XSAVEOPT

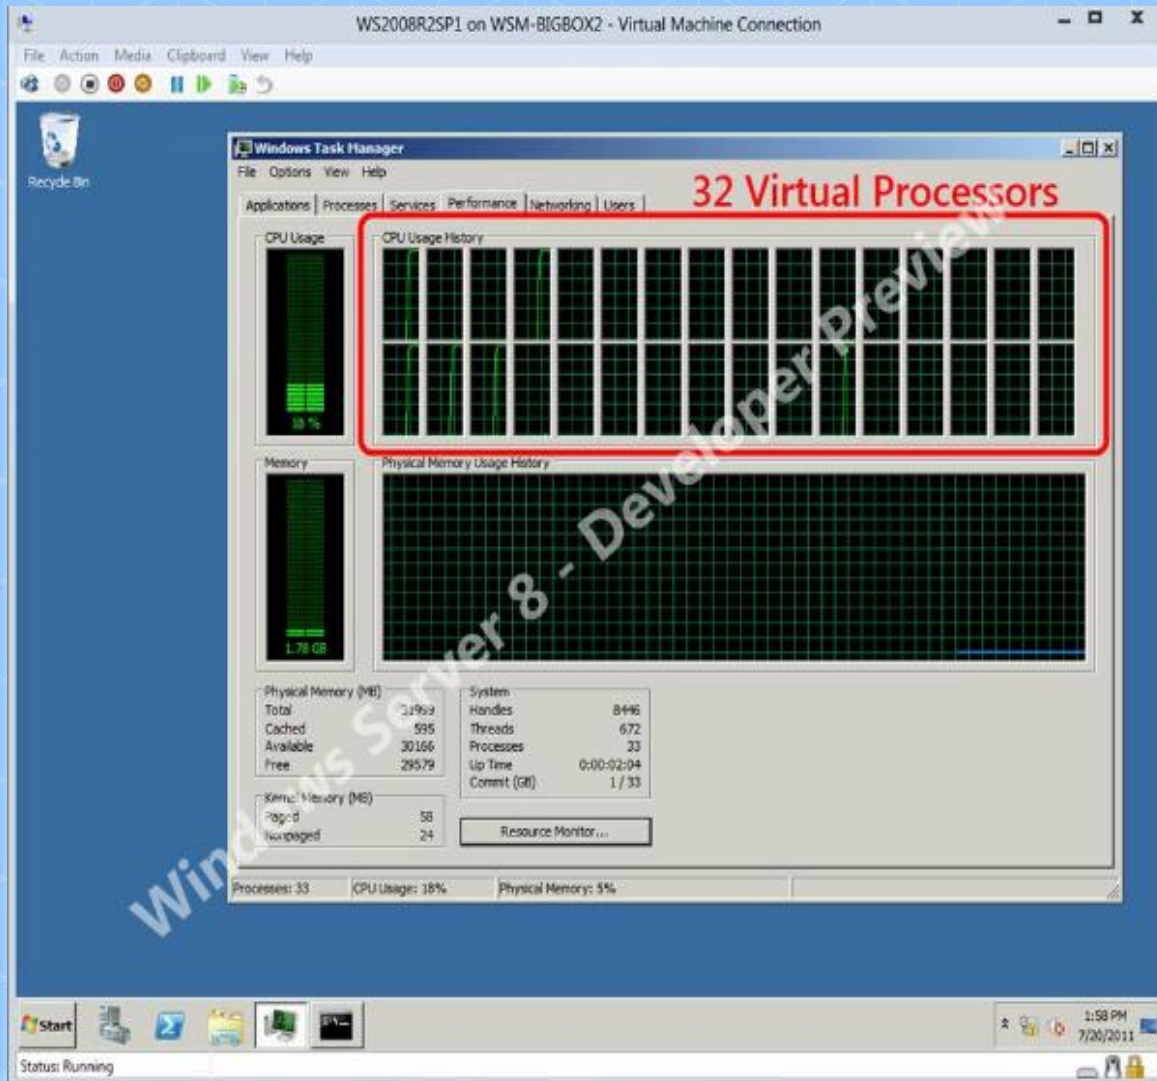
Memory Scale Comparison

	Windows Server 2008	Windows Server 2008 R2	Windows Server 8 Developer Preview
Physical Memory Support	1 TB	1 TB	2 TB
Virtual Machine Memory	Up to 64 GB	Up to 64 GB	Up to 512 GB
Guest NUMA Support	No	No	Yes
Runtime Changes	No	No	Yes

Hyper-V Scale Comparison

Cluster and Live Migration

	Windows Server 2008	Windows Server 2008 R2	Windows Server 8 Beta
Cluster Scale	16 Nodes up to 1000 VMs	16 Nodes up to 1000 VMs	63 Nodes up to 4000 VMs
Live Migration	Yes, one at a time	Yes, one at a time	Yes, with no limits. As many as hardware will allow.
Live Storage Migration	No. Quick Storage Migration via SCVMM	No. Quick Storage Migration via SCVMM	Yes, with no limits. As many as hardware will allow.
Servers in a Cluster	16	16	64



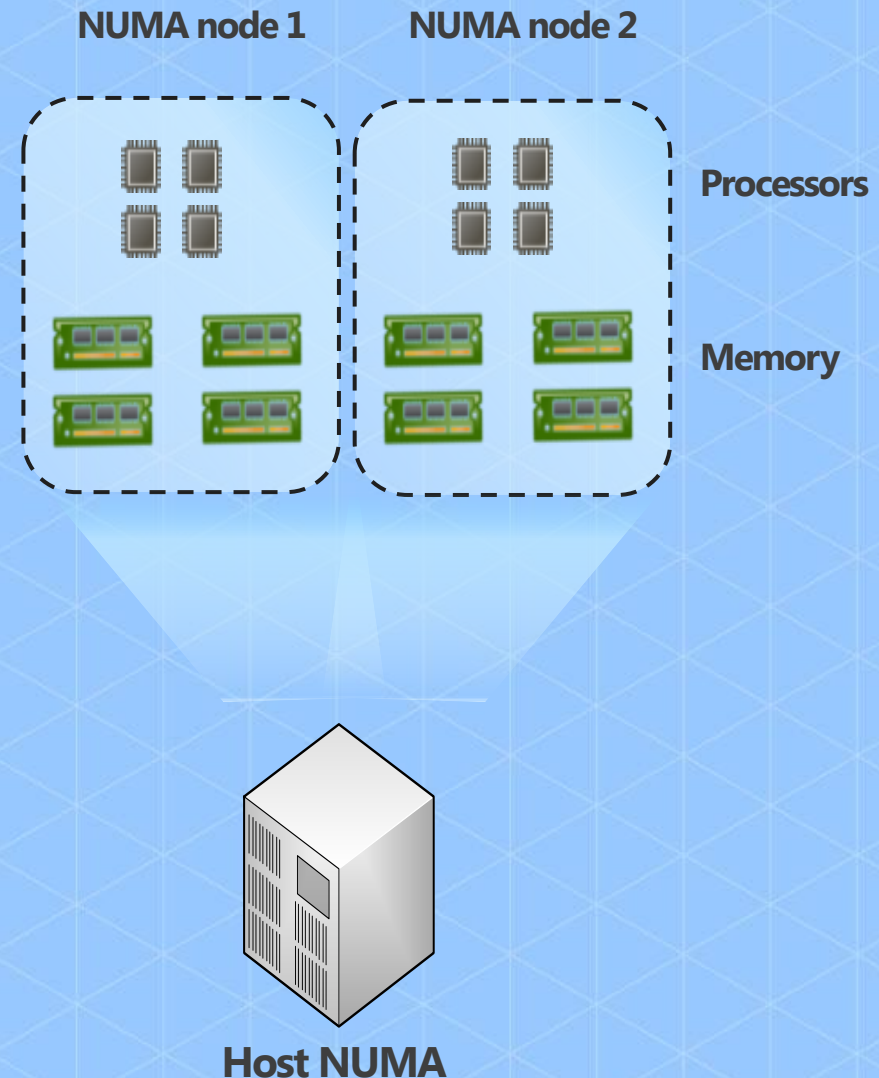
Demo

CPU, Memory, & VHDX



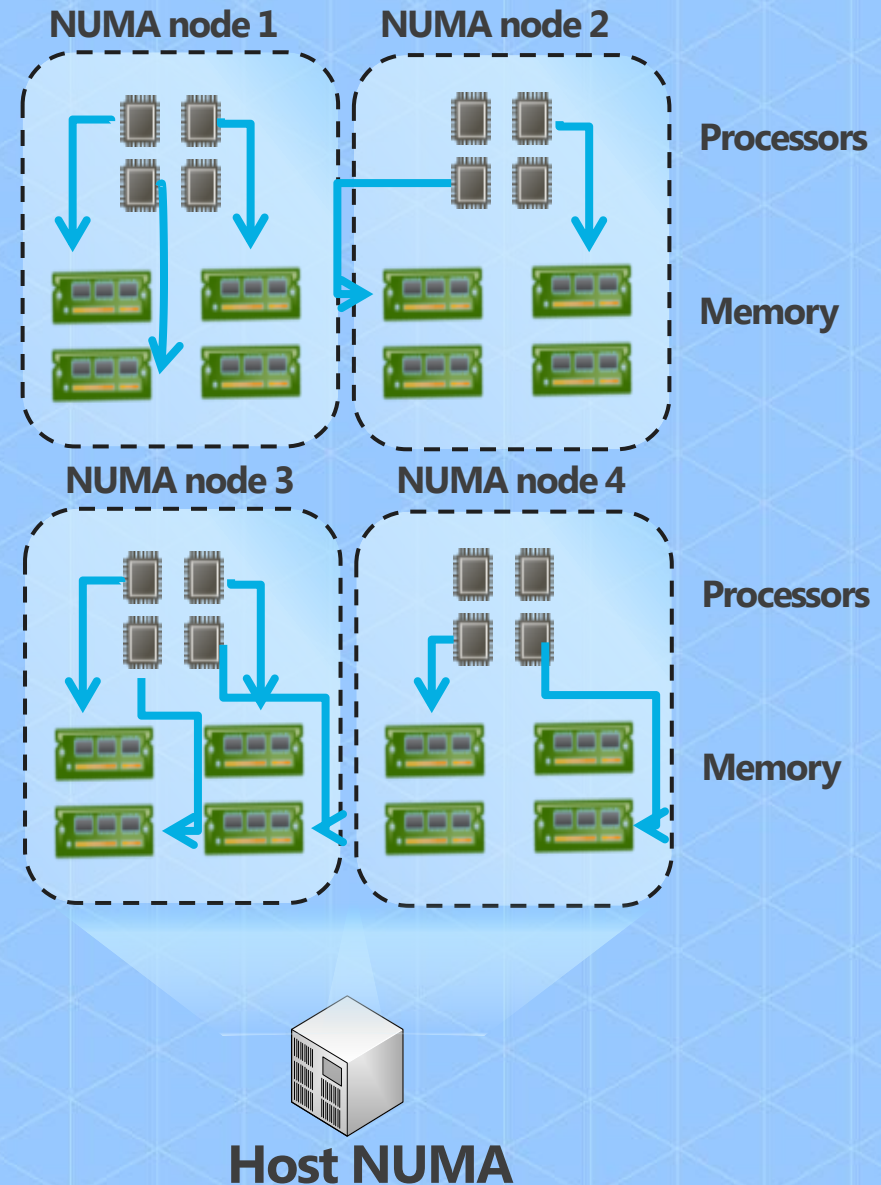
Scaling up: Physical NUMA

- NUMA (Non-uniform memory access)
 - Helps hosts scale up the number of cores and memory access
 - Partitions cores and memory into “nodes”
 - Allocation and latency depends on the memory location relative to a processor
- High performance applications detect NUMA and minimize cross-node memory access



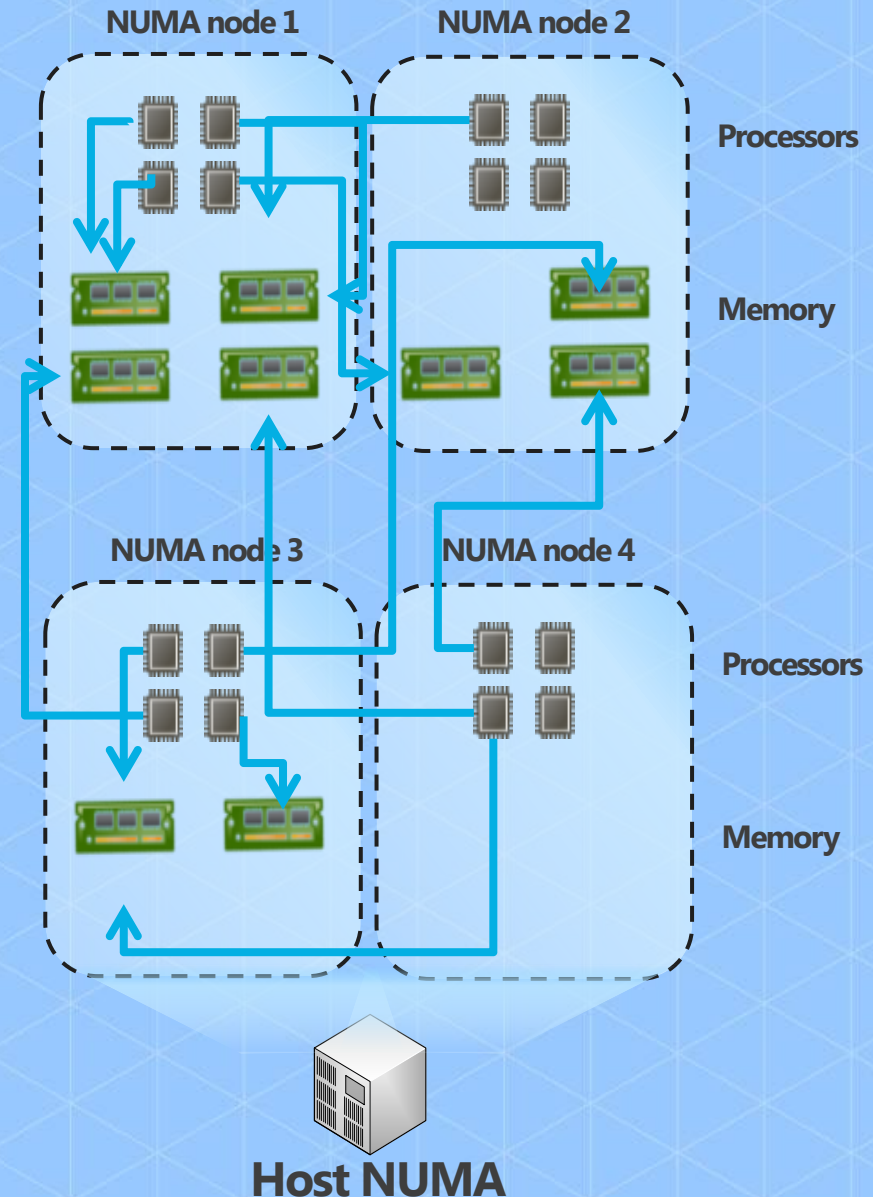
Scaling up: Physical NUMA

- This is optimal...
 - System is balanced
 - Memory allocation and thread allocations within the same NUMA node
 - Memory populated in each NUMA node




Scaling up: Physical NUMA

- This isn't optimal...
 - System is imbalanced
 - Memory allocation and thread allocations across different NUMA nodes
 - Multiple node hops
 - NUMA Node 2 has an odd number of DIMMS
 - NUMA Node 3 doesn't have enough
 - NUMA Node 4 has no local memory (worst case)



Scaling Up: Guest NUMA

- Guest NUMA
 - Presenting NUMA topology within VM
 - Guest operating systems & apps can make intelligent NUMA decisions about thread and memory allocation
- Guest NUMA nodes are aligned with host resources
- Policy driven per host – best effort, or force alignment



Demo


NUMA Settings



NETWORKING


Extensibility

Customers want specialized functionality with lots of choice ...

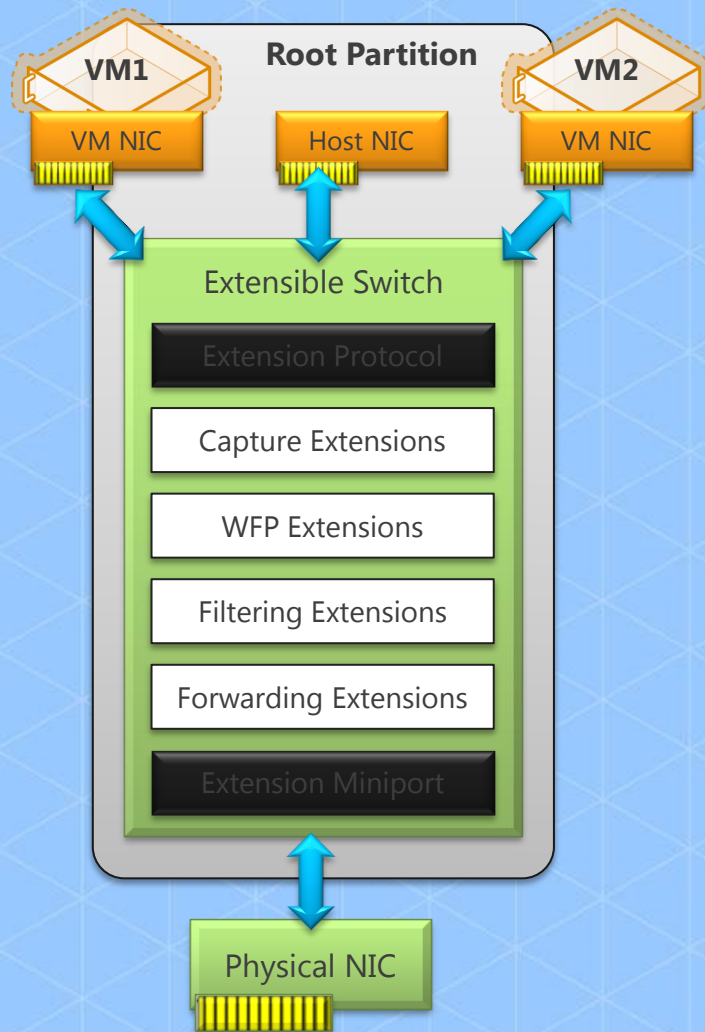


Extensibility

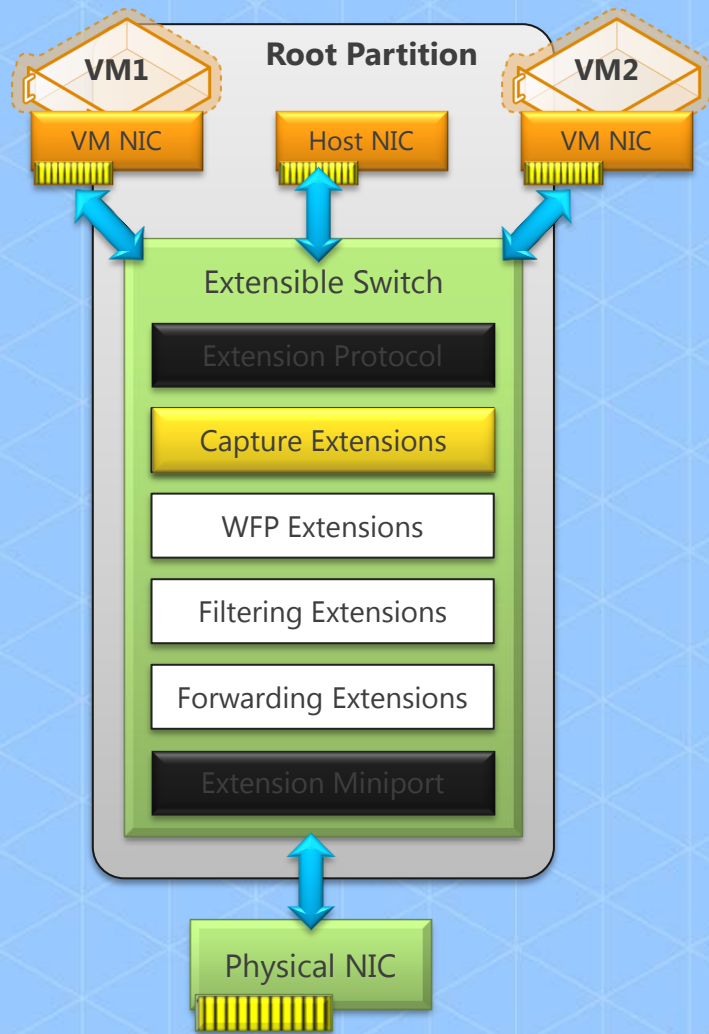
Customers want specialized functionality with lots of choice ...
... for firewalls, monitoring and physical fabric integration



Hyper-V Extensible Switch



Hyper-V Extensible Switch

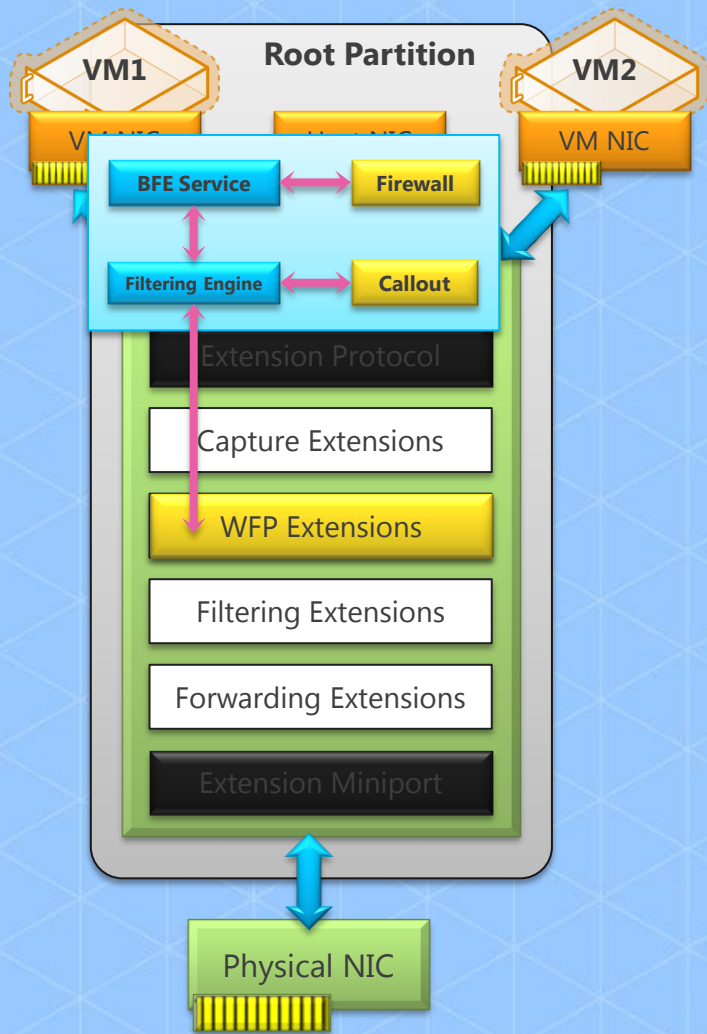


- Capture extensions can inspect traffic and generate new traffic for report purposes
- Capture extensions do not modify existing Extensible Switch traffic

Example: sflow by inMon



Hyper-V Extensible Switch

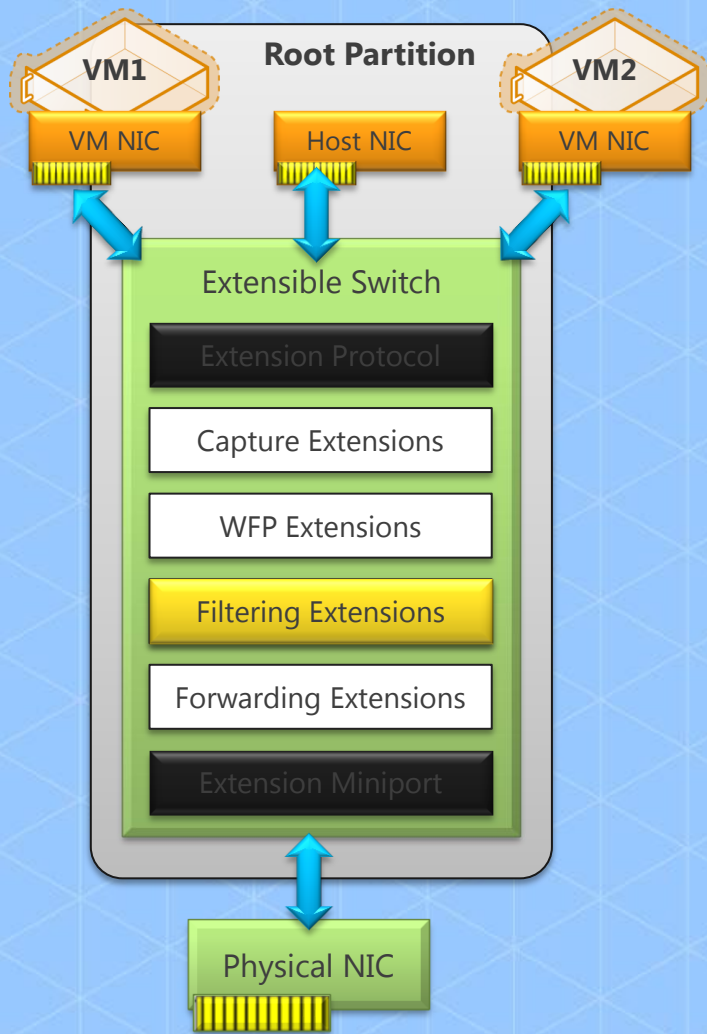


- Windows Filter Platform (WFP) Extensions can inspect, drop, modify, and insert packets using WFP APIs
- Windows Antivirus and Firewall software uses WFP for traffic filtering

Example: Virtual Firewall by 5NINE Software



Hyper-V Extensible Switch

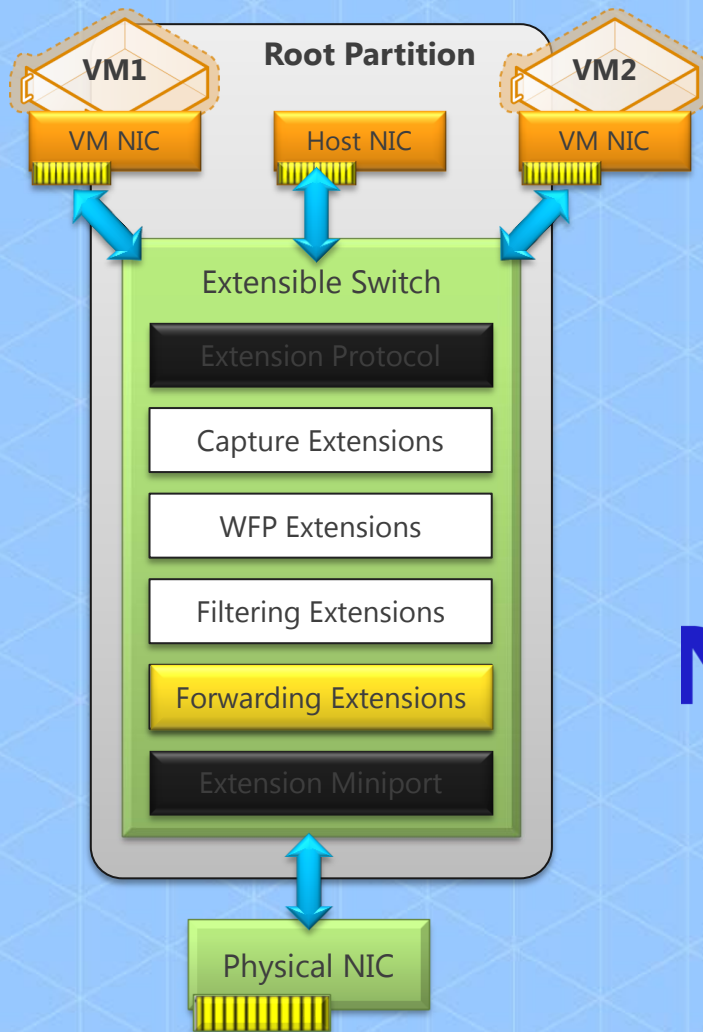


- Filtering extensions can also be implemented using NDIS filtering APIs

Example: VM DoS Prevention by Broadcom



Hyper-V Extensible Switch



- Forwarding extensions direct traffic, defining the destination(s) of each packet
- Forwarding extensions can capture and filter traffic

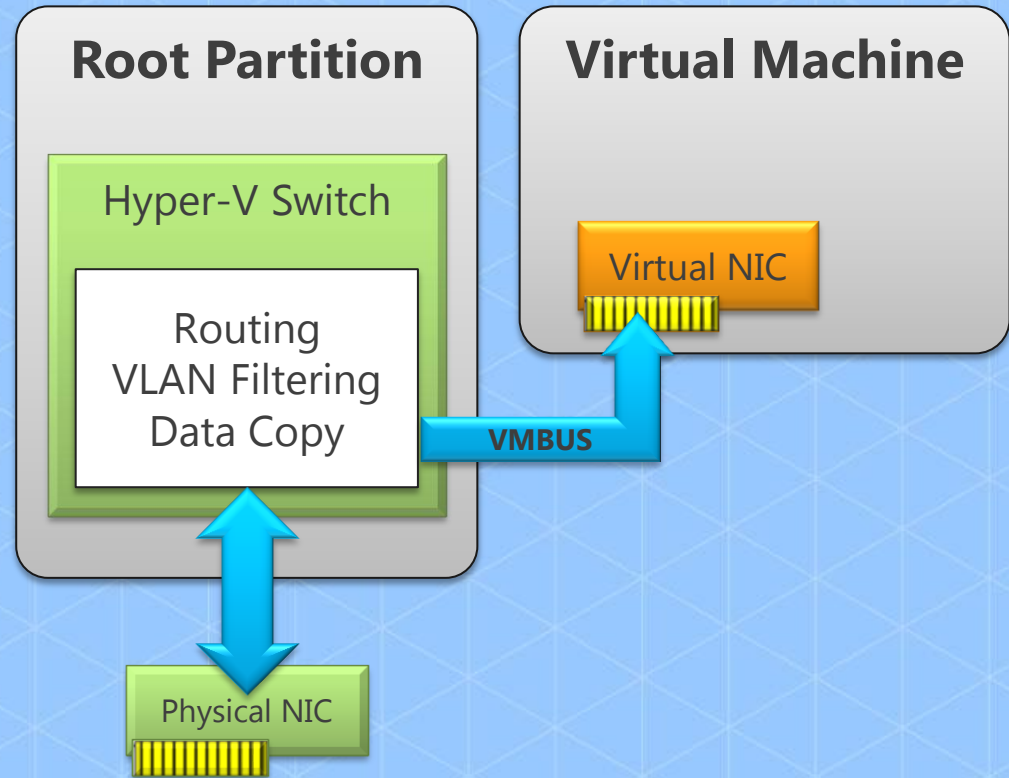
Examples: Cisco Nexus 1000V and UCS
NEC OpenFlow

NEC



Single-Root I/O Virtualization (SR-IOV)

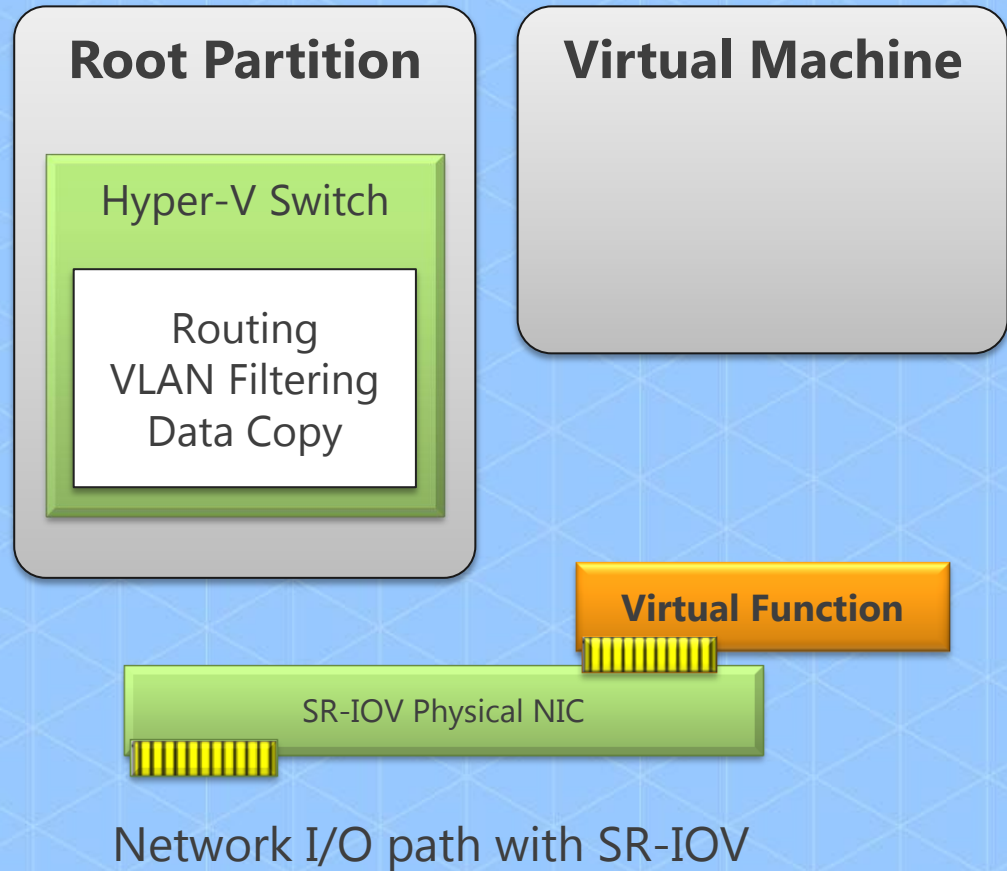
DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY



Network I/O path without SR-IOV

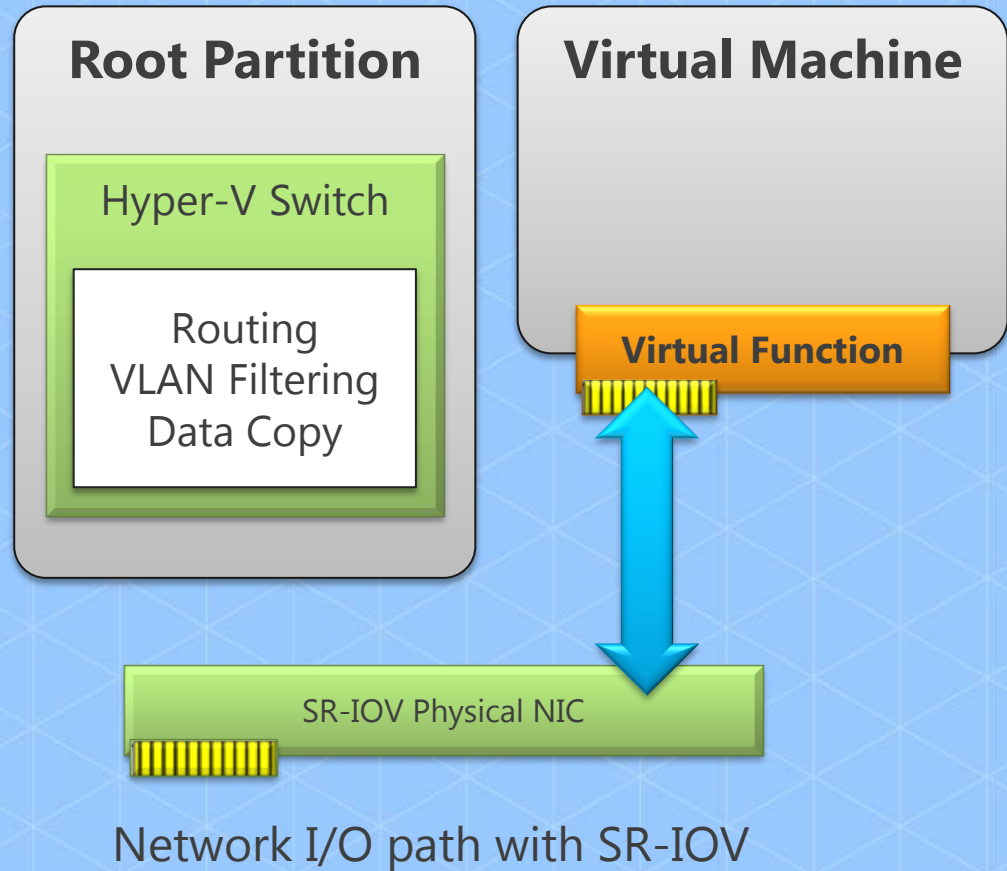
Single-Root I/O Virtualization (SR-IOV)

DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY



Single-Root I/O Virtualization (SR-IOV)

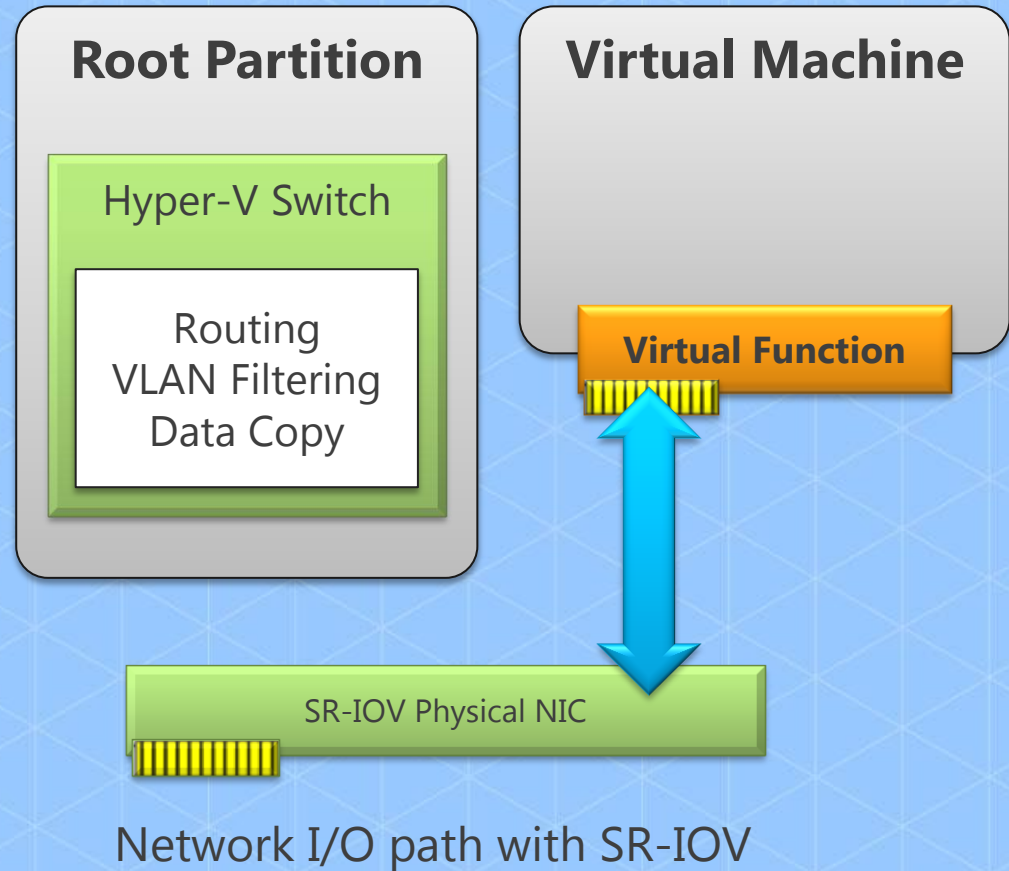
DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY



Single-Root I/O Virtualization (SR-IOV)

DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY


- Reduces CPU utilization for processing network traffic



Single-Root I/O Virtualization (SR-IOV)

DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY


- Reduces CPU utilization for processing network traffic
- Reduces latency of network path



Single-Root I/O Virtualization (SR-IOV)

DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY

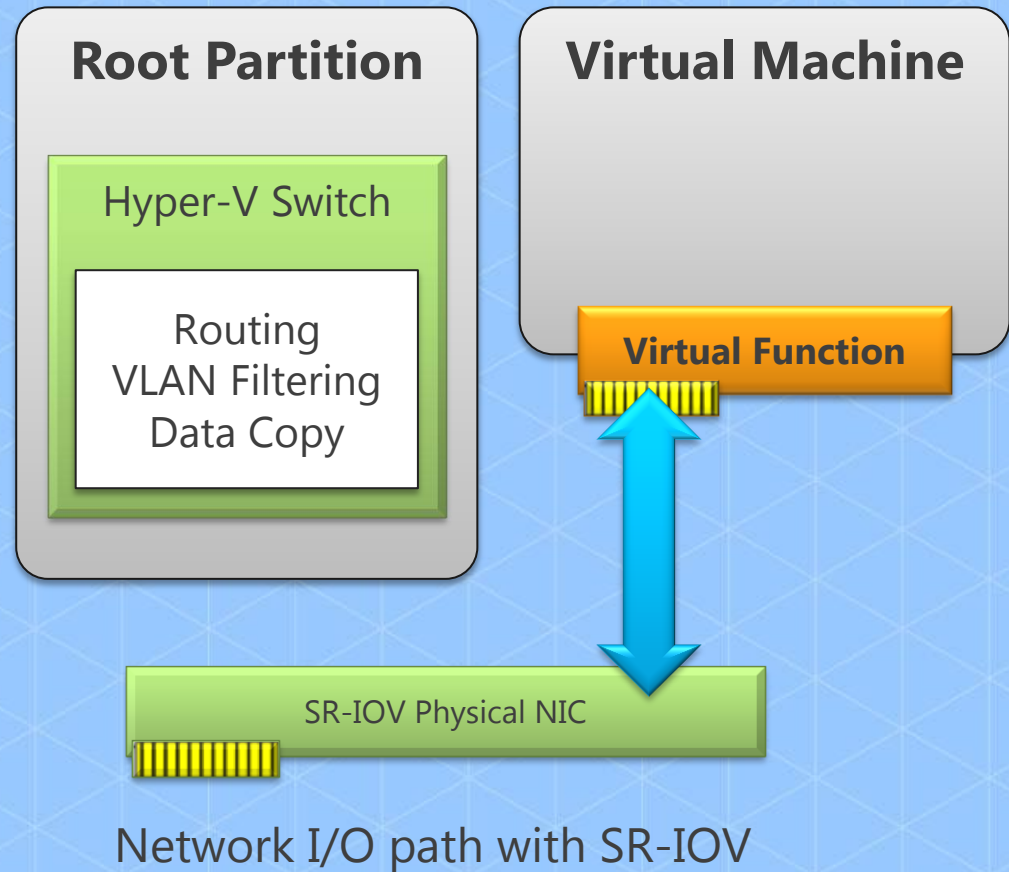
- Reduces CPU utilization for processing network traffic
- Reduces latency of network path
- Increases throughput



Single-Root I/O Virtualization (SR-IOV)

DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY


- Reduces CPU utilization for processing network traffic
- Reduces latency of network path
- Increases throughput
- Supports Live Migration



Single-Root I/O Virtualization (SR-IOV)


DIRECT DEVICE ASSIGNMENT TO VIRTUAL MACHINES WITHOUT COMPROMISING FLEXIBILITY

- Reduces CPU utilization for processing network traffic
- Reduces latency of network path
- Increases throughput
- Supports Live Migration
- Requires:
 - Chipset: Interrupt and DMA remapping: VT-d2 or IOMMU
 - Access Control Services (ACS) on PCIe root ports
 - Alternative Routing-ID Interpretation (ARI)
 - CPU: Hardware virtualization, EPT or NPT



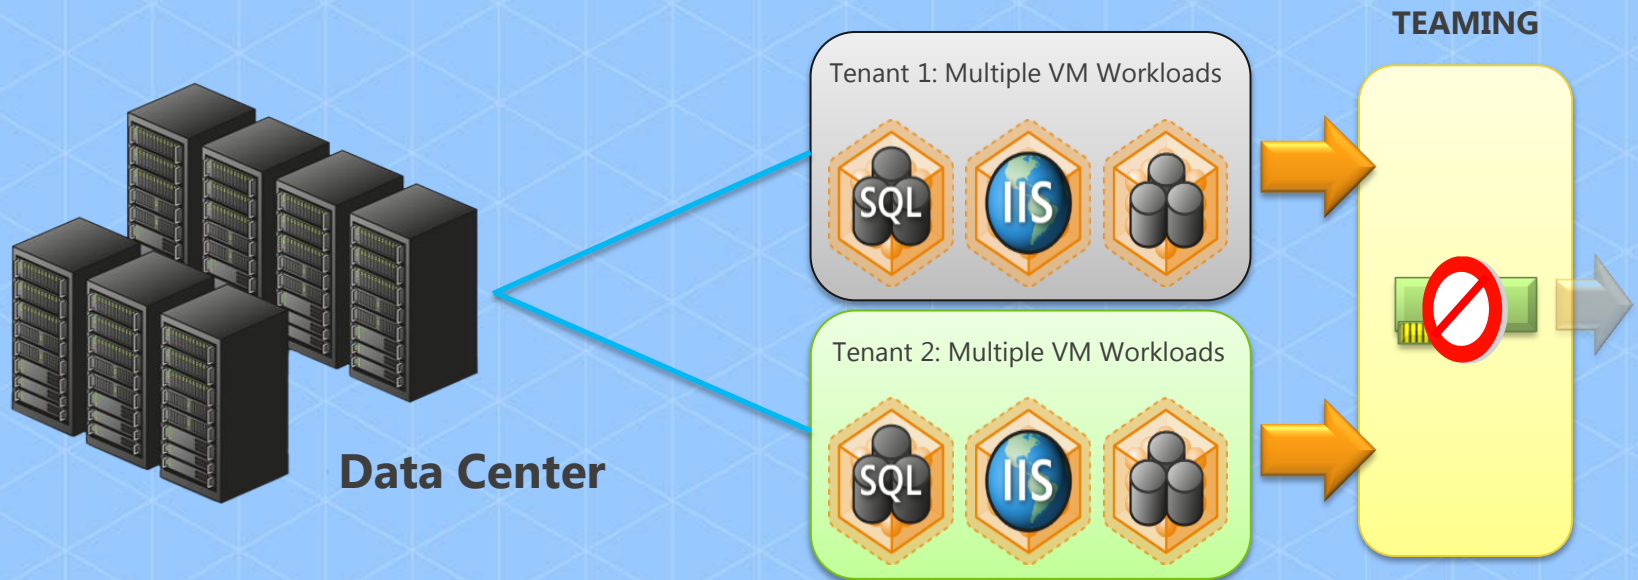
Reliability

Even when hardware fails ...



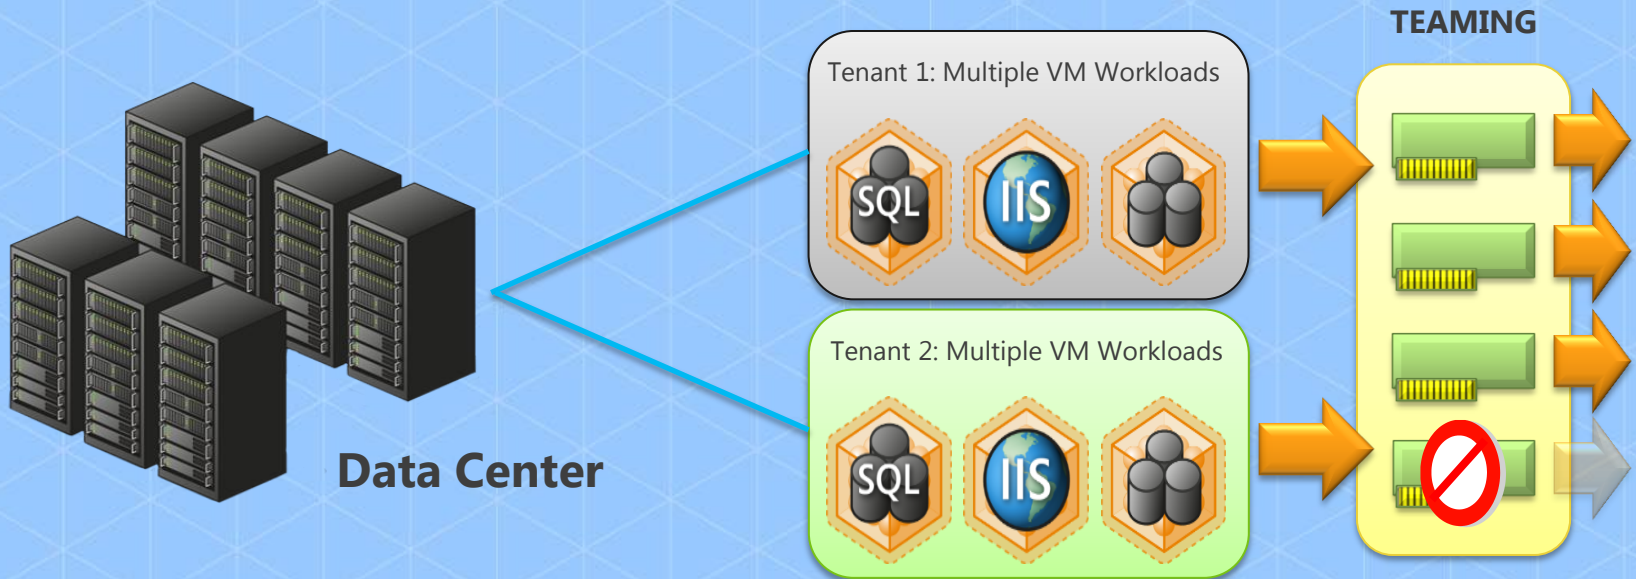
Reliability

Even when hardware fails ...

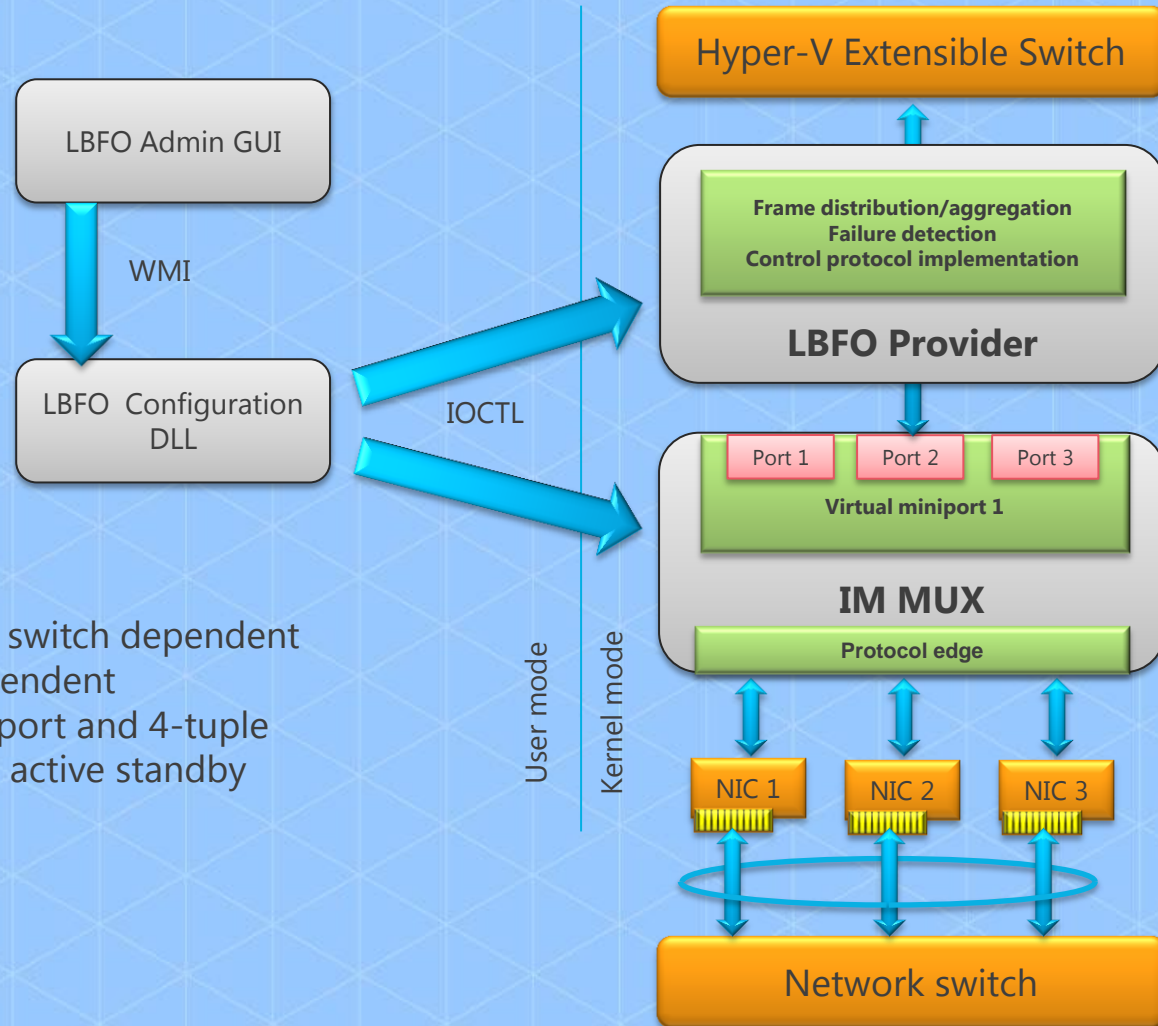


Reliability

Even when hardware fails ...
... customers want continuous availability



NIC Teaming



- Multiple modes: switch dependent and switch independent
- Hashing modes: port and 4-tuple
- Active active and active standby

Demo

Network Bandwidth and NIC Teaming



STORAGE

Dynamic, High Performance Storage

HYPER-V STORAGE: SCALING TO NEW HEIGHTS

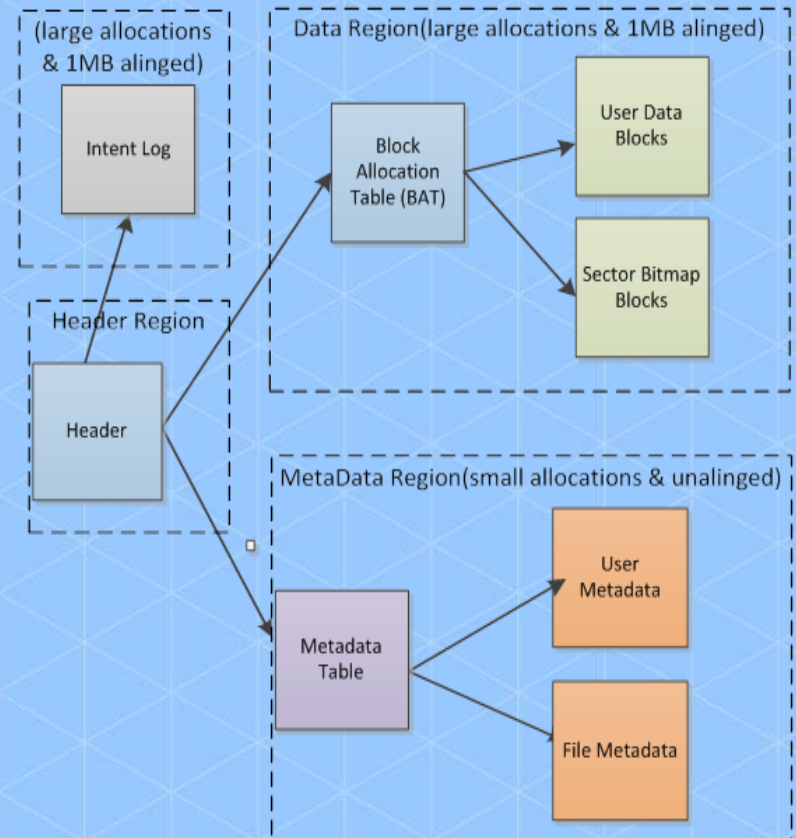
Dynamic, High Performance Storage

HYPER-V STORAGE: SCALING TO NEW HEIGHTS

- Live Storage Migration
- Virtual Fiber Channel
- Support for File Based Storage on SMB 2.2
- Online MetaOperations
 - Live VHD Merge
- Native 4K Disk Support
- New VHDX
- Offloaded Data Transfer (ODX) & Trim

VHDX: Highly Scalable

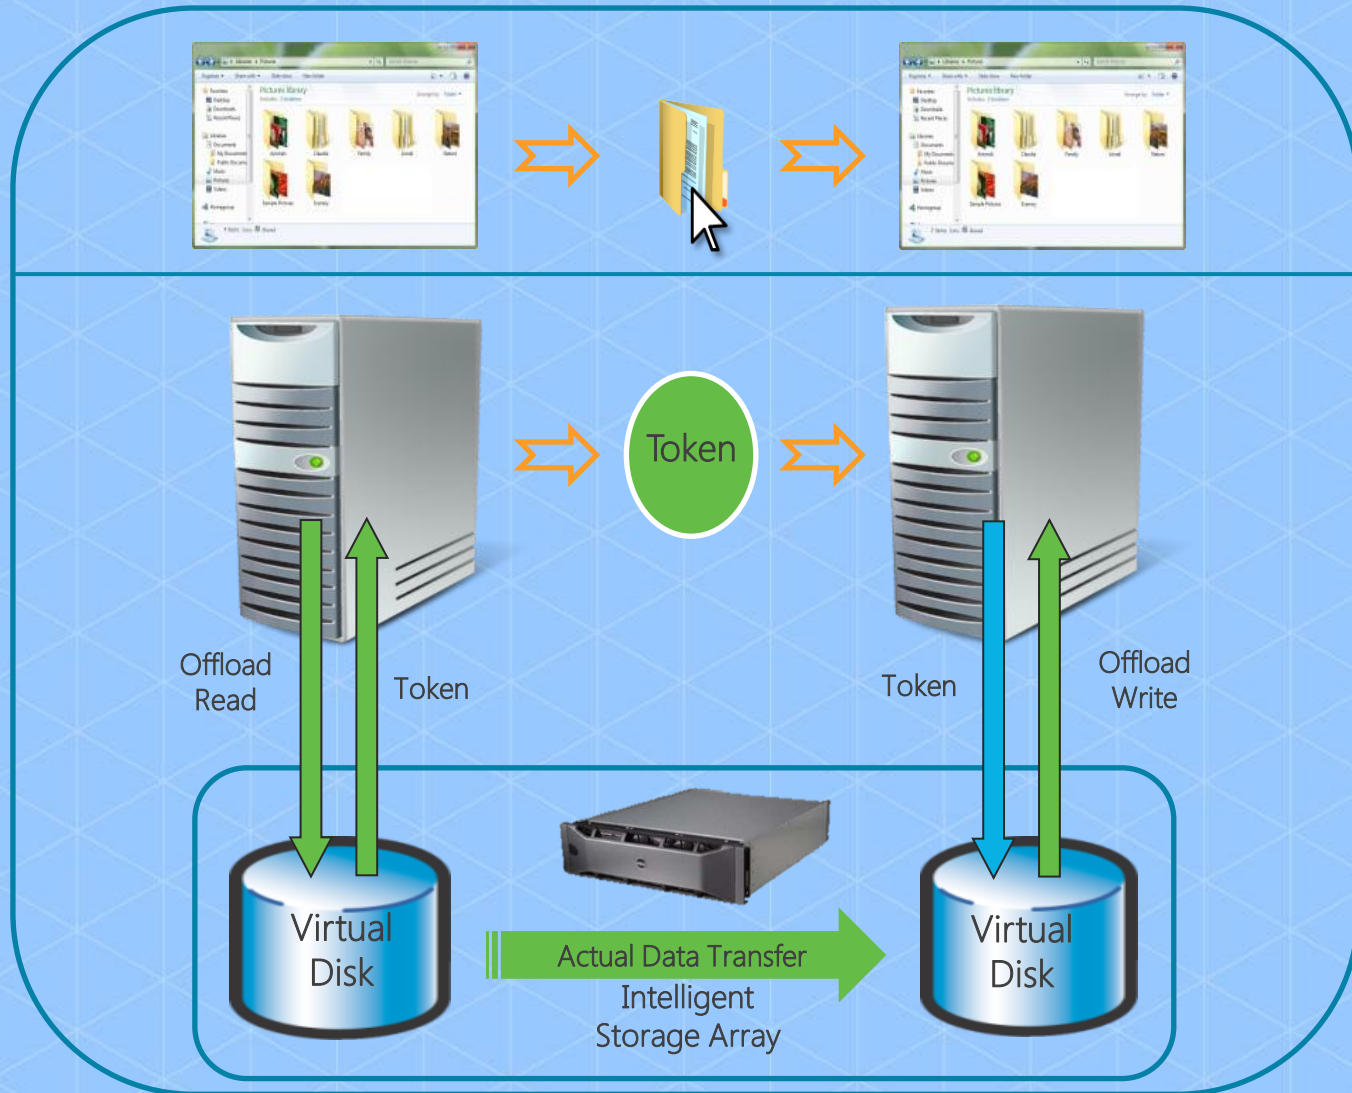
- >2TB disks
- Better performance (eliminate alignment issues)
- Resilient to corruption
- Embed user defined metadata
- Larger block sizes to adapt to workload requirements



Hyper-V Storage Comparison

	Windows Server 2008	Windows Server 2008 R2	Windows Server 8 Developer Preview
Live Storage Migration	No. Quick Storage Migration via SCVMM	No. Quick Storage Migration via SCVMM	Yes, with no limits. As many as hardware will allow.
VMs on File Storage	No	No	Yes, SMB
Guest Fiber Channel	No	No	Yes
Virtual Disk Format	VHD up to 2 TB	VHD up to 2 TB	VHD up to 2 TB VHDX up to 16 TB
VM Guest Clustering	Yes, via iSCSI	Yes, via iSCSI	Yes, via iSCSI or FC
Native 4k Disk Support	No	No	Yes
Live VHD Merge	No, offline.	No, offline.	Yes
Secure Copy Offload (ODX)	No	No	Yes

Offloaded Data Transfer (ODX)




Hyper-V ODX Support

- Secure Offload data transfer
 - Fixed VHD/VHDX Creation
 - Dynamic VHD/VHDX Expansion
 - VHD/VHDX Merge
 - Live Storage Migration
- Just one example...

Hyper-V ODX Support

- Secure Offload data transfer
 - Fixed VHD/VHDX Creation
 - Dynamic VHD/VHDX Expansion
 - VHD/VHDX Merge
 - Live Storage Migration
- Just one example...

Creation of a 10 GB Fixed Disk



POWERSHELL

Managing with PowerShell



- Use PowerShell to manage and monitor Windows Server 8 Clouds
 - Leverage the PowerShell community and skills and build your own tools
 - Simple and consistent APIs to manage and monitor
 - Write WMIv2 providers, get PowerShell for free!
- File Servers & Shares
 - Hyper-V configurations
 - Virtual Machines
 - Failover Clusters
 - Network Configuration
 - Hyper-V Extensible Switch
 - Hyper-V Replica

Demo

PowerShell

Speaker, Title
Group





CLUSTERING

Increased Scalability

- Increased scale up and scale out
 - 4x scale over Windows Server 2008 R2
 - Scale up to 64-nodes
 - Scale out to 4,000 VMs per cluster

- Windows Server 8 supports an industry best 4,000 VMs per cluster and can now massively scale out to 64 nodes in a cluster



SUMMARY

Leveraging Modern Hardware

CPU & Memory

- PHY: 160 Logical Processors
- PHY: Up to 2 TB of RAM
- VM: 32 VPs
- VM: 512 GB Memory
- Guest NUMA
- Removal of VP: LP Ratios

Networking

- Hyper-V Extensible Switch
- Single Root I/O Virtualization (SR-IOV)
- NIC Teaming

Storage

- New VHDX format
- Storage Offload (ODX)
- SMB2 Direct (RDMA)
- Guest Fiber Channel (MPIO)
- 4k Native Disk Support

Clustering

- Improved scale to 63 nodes and 4000 VMs per cluster
- Improved CSV with support for storage offloads, better backup, RDMA, LBFO and security

Microsoft[®]

© 2012 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries. The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.