

# Module 4: Hyper-V for VDI

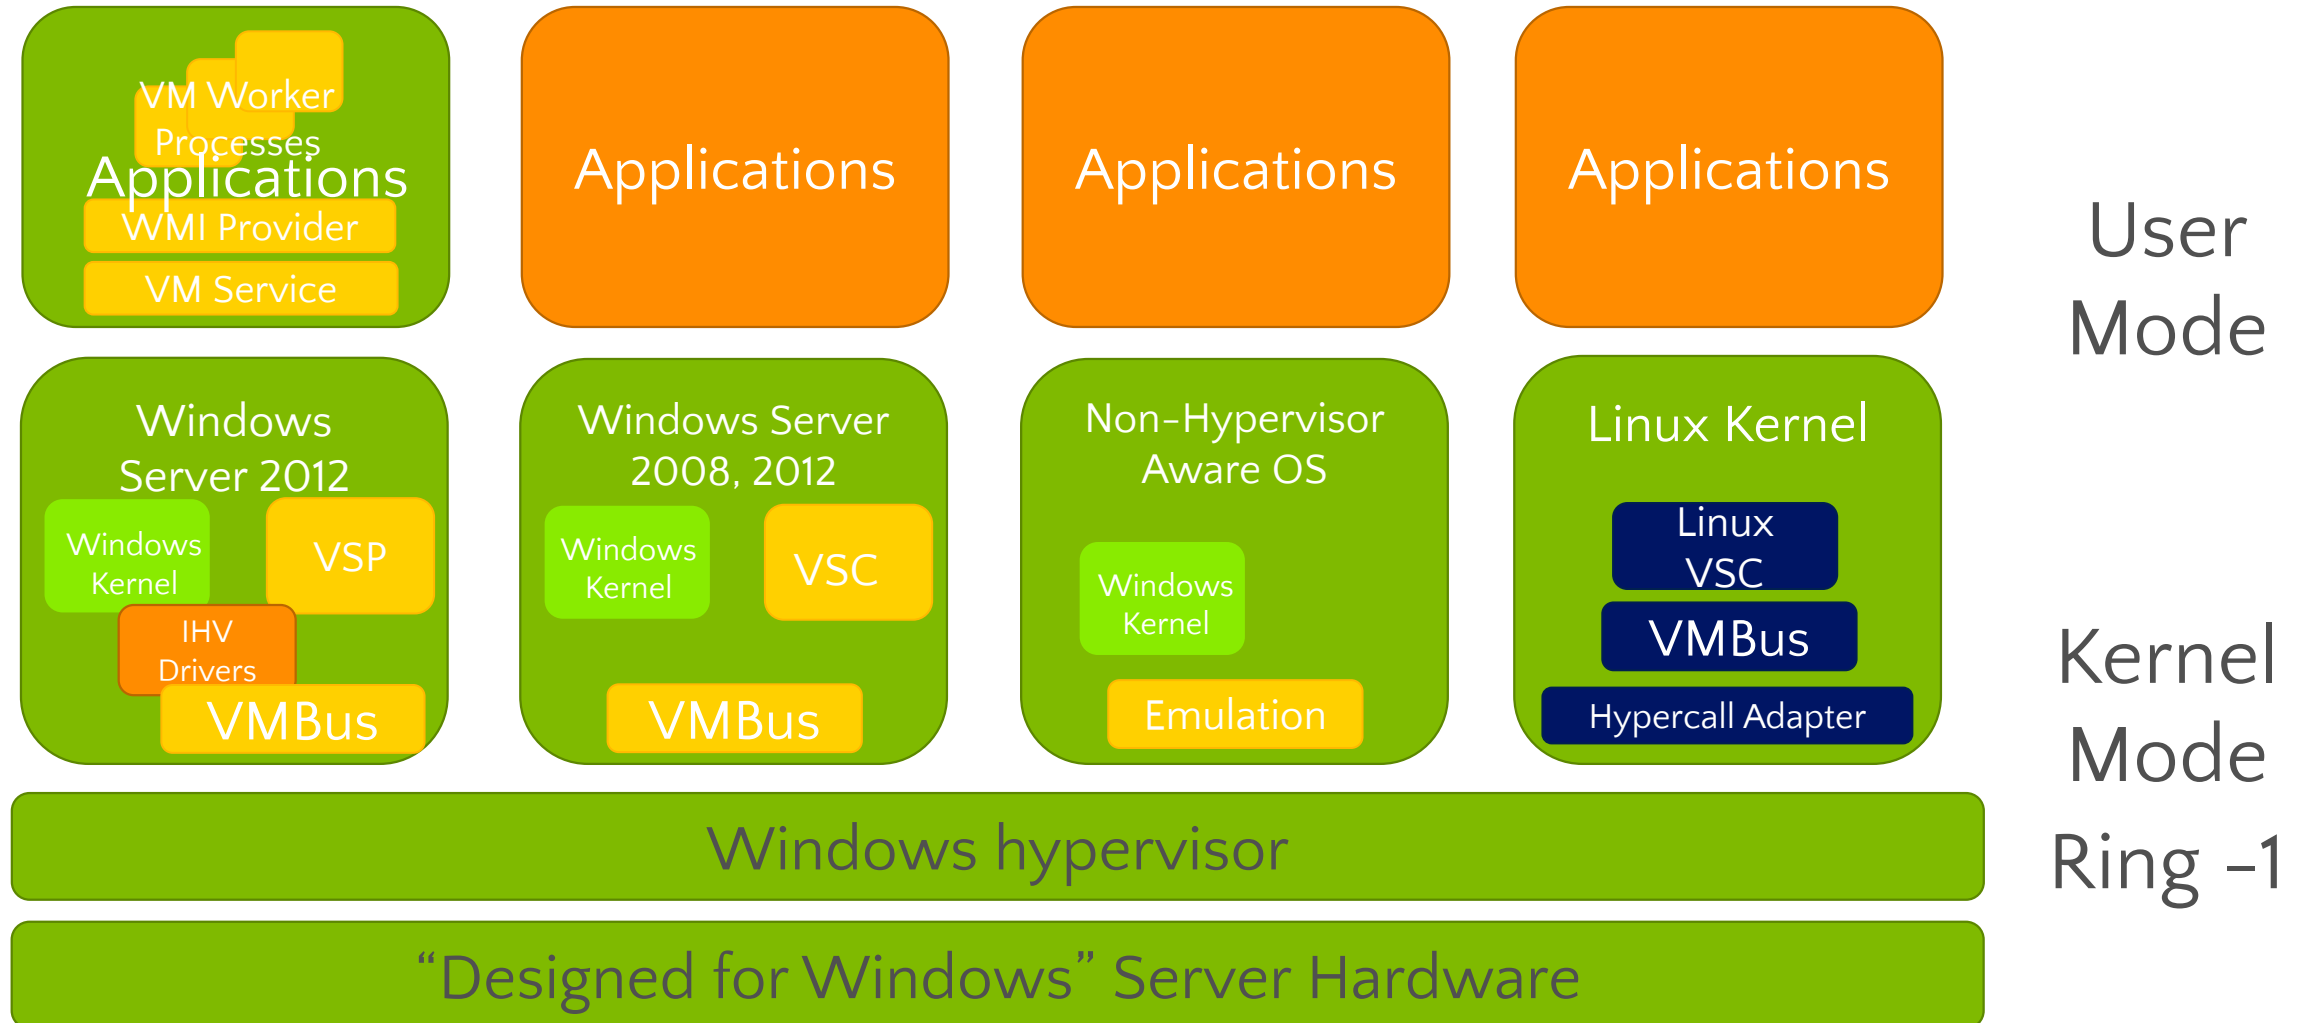
# Agenda

- Introduction to Hyper-V
  - Technical requirements and architecture
- Hyper-V for VDI components
  - VDI over SMB, SAN, DAS
  - Storage tiers for pooled desktops
  - Clustering and storage spaces

# Hyper-V Architecture

Parent Partition


Child Partitions



# Hypervisor: Monolithic vs. Microkernelized


## Monolithic hypervisor

- Simpler than a modern kernel, but still complex
- Contains its own drivers model
- Used by competitors



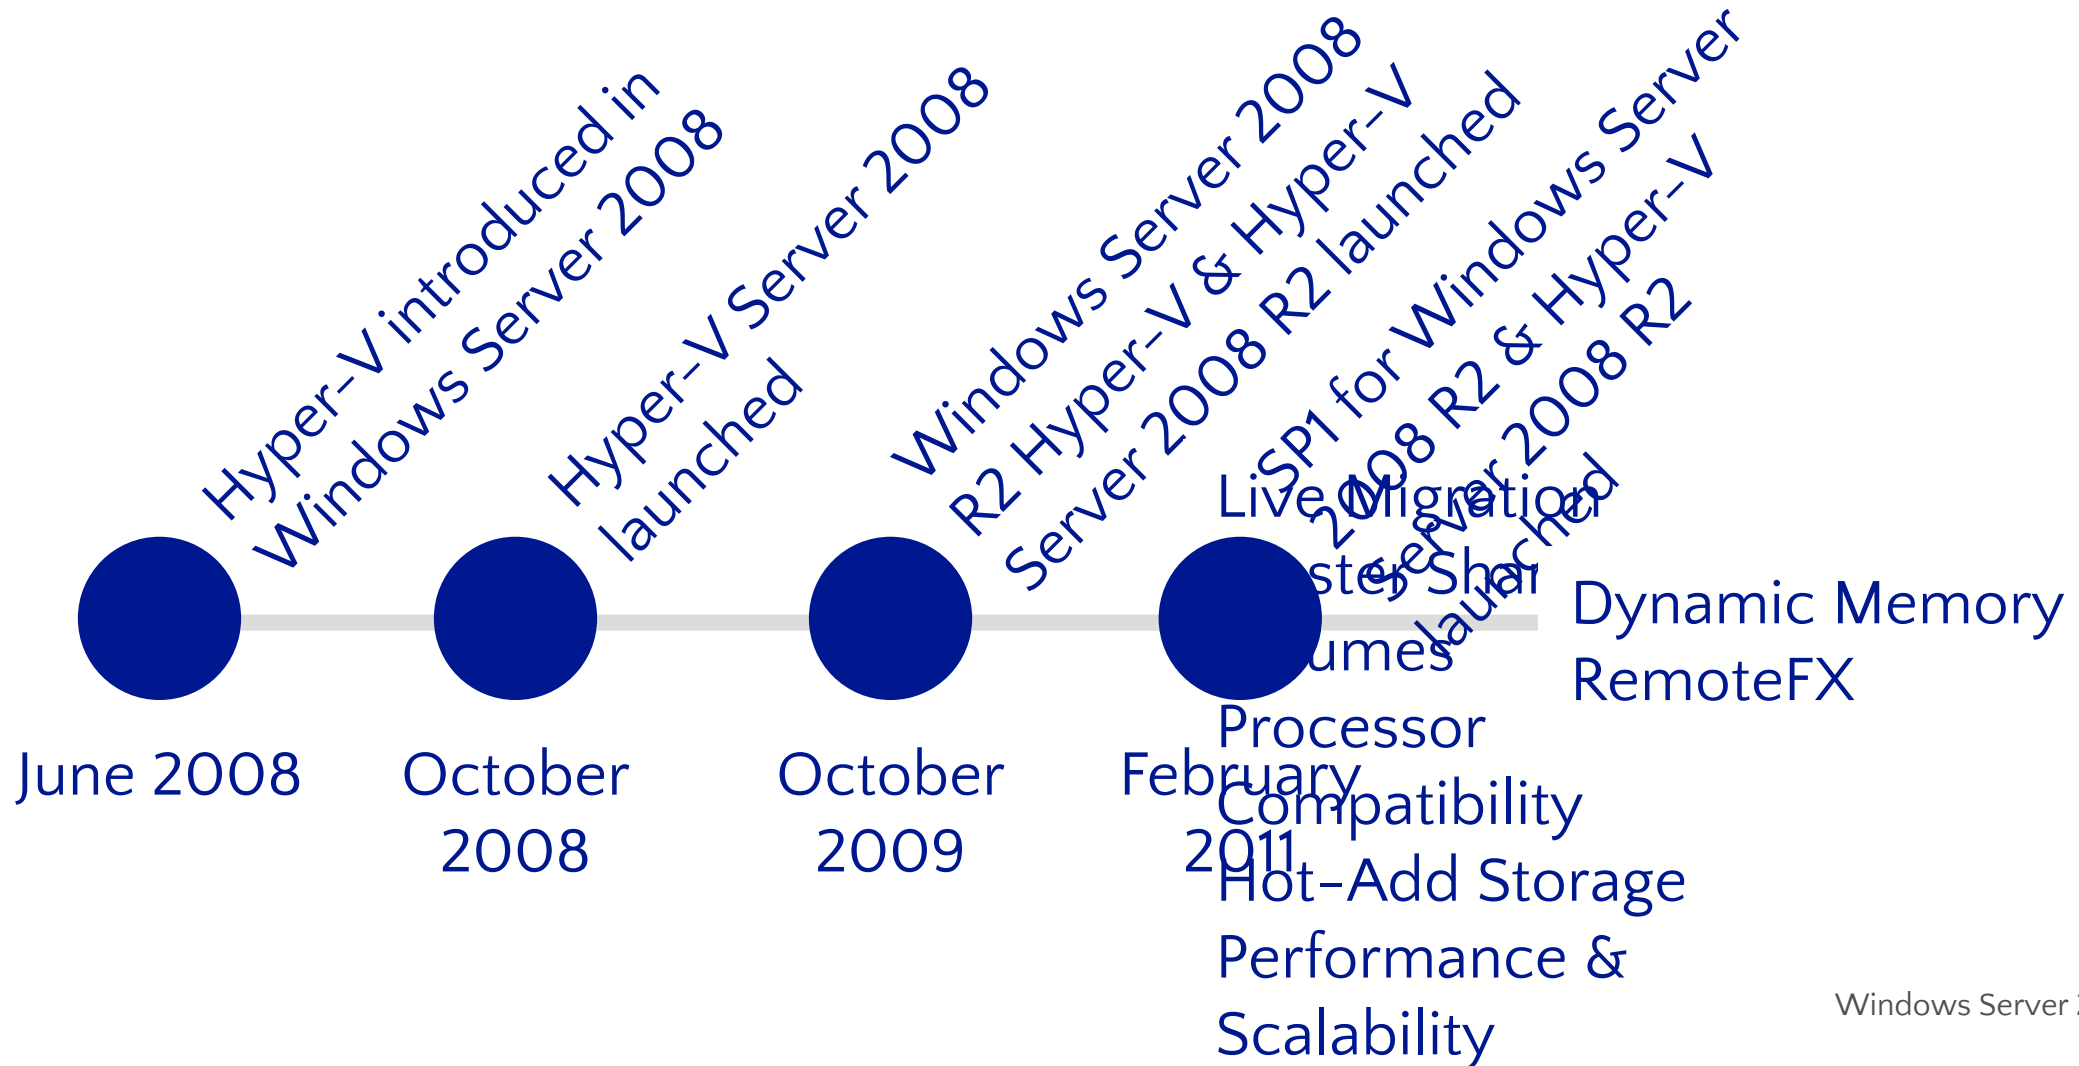
## Microkernelized hypervisor

- Simple partitioning functionality
- Increase reliability and minimize TCB
- No third-party code in HV
- Drivers run within guests
- Hyper-V





# Before Windows Server 2012 Hyper-V



# Hyper-V Scalability Improvements

System	Resource	Windows Server 2008 R2 Hyper-V	Windows Server 2012 Hyper-V	Improvement Factor
Host	Logical Processors	64	320	5×
	Physical Memory	1TB	4TB	4×
	Virtual CPUs per Host	512	2,048	4×
VM	Virtual CPUs per VM	4	64	16×
	Memory per VM	64GB	1TB	16×
	Active VMs per Host	384	1,024	2.7×
	Guest NUMA	No	Yes	-
Cluster	Maximum Nodes	16	64	4×
	Maximum VMs	1,000	8,000	8×



# Hyper-V for VDI components

# Windows Server 2012 Storage

BEST  
VALUE  
FOR VDI



Use direct-attached storage, network-attached storage, and clustered or SAN storage

Provide configuration options to optimize for tiered storage

Reduce storage cost while maximizing I/O operations per second



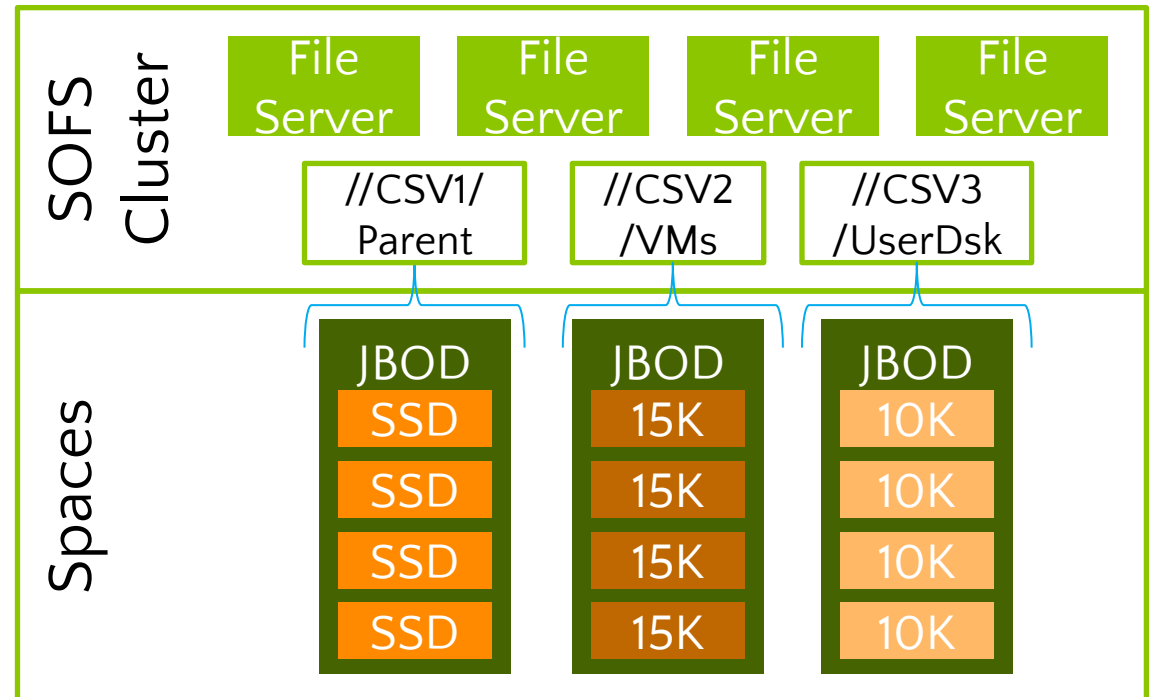

# VDI Storage in Hyper-V

## New in Windows Server 2012

- Storage Spaces
- SMB 3.0, Hyper-V over SMB support
- Cluster Shared Volume
- Scale-Out File Server

## VDI Storage Configuration

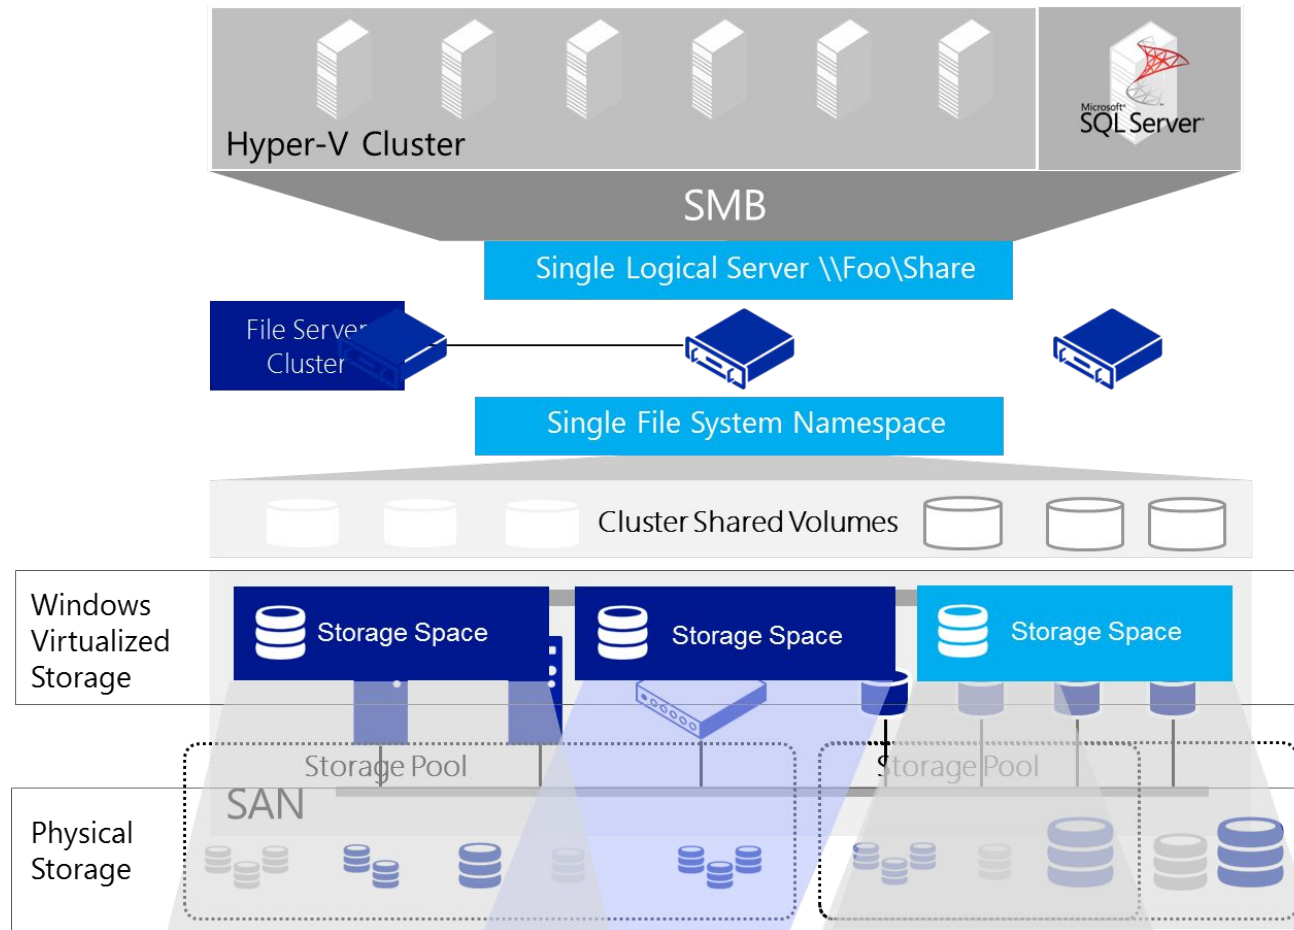
- Direct Attached, Central SMB, or Central CSV / SAN
- Storage location of user profile disks, parent VHD, & virtual machines can be configured separately for each collection to enable tiered storage for optimal performance and savings
- Optimum use of high IOPS vs. high volume storage



# Windows Server 2012 File Server

- Not your traditional file server
- Enables many new architectures and technologies
  - SMB 3.0
  - RDMA
- Introduces file server for “server data”
  - Server data is an extended file lock

# Server Message Block 3.0



- New flexible storage solution for virtual or cloud infrastructure
- Virtually no application downtime for planned maintenance and unplanned failures
- Data store for Microsoft SQL Server databases and Hyper-V workloads
- Highly available scale-out file server
- Active/active access
- Built-in encryption support
- Single namespace for file system

# New in SMB 3.0

SMB Transparent Failover

SMB Scale Out

SMB Multichannel

SMB Direct

SMB Directory Leasing

SMB Encryption

SMB-Specific PowerShell  
cmdlets

Performance Counters for  
server applications

Improved performance optimizations

- Server Message Block overview:
  - <http://technet.microsoft.com/en-us/library/hh831795.aspx>

# Hyper-V over SMB

- Can store virtual machine files (configuration, VHD, snapshots) in files shares over the SMB 3.0 protocol
- Is supported for both stand-alone and clustered servers that use Hyper-V with shared file storage for the cluster
- Can support scale-out file shares and clusters
- Can leverage SMB Multichannel



# Why Should I Care?

## **Ease of storage provisioning and management**

- You can manage file shares instead of storage fabric and logical unit numbers (LUNs)
- Previously all Hyper-V clusters required shared storage using Fibre Channel, Fibre Channel over Ethernet, iSCSI, or Serial-Attached SCSI

## **Increased flexibility**

- You can dynamically migrate virtual machines or databases in the data center

## **Ability to take advantage of existing investment in a converged network**

- You can use your existing converged network with no specialized storage networking hardware

## **Reduced capital expenditures**

- Capital expenses (acquisition costs) are reduced

## **Reduced operating expenditures**

- You can reduce operating costs because there is no need for specialized storage expertise


# Requirements

- One or more computers running Windows Server 2012 with the File Services role installed
- One or more computers running Windows Server 2012 with the Hyper-V role installed (separate from the file server)
- A common Active Directory infrastructure
  - The servers running Active Directory Domain Services (AD DS) do not need to run Windows Server 2012
  - Constrained delegation
- Supported configurations
  - Standalone Hyper-V servers (not a high-availability solution)
  - Hyper-V servers configured in a failover cluster

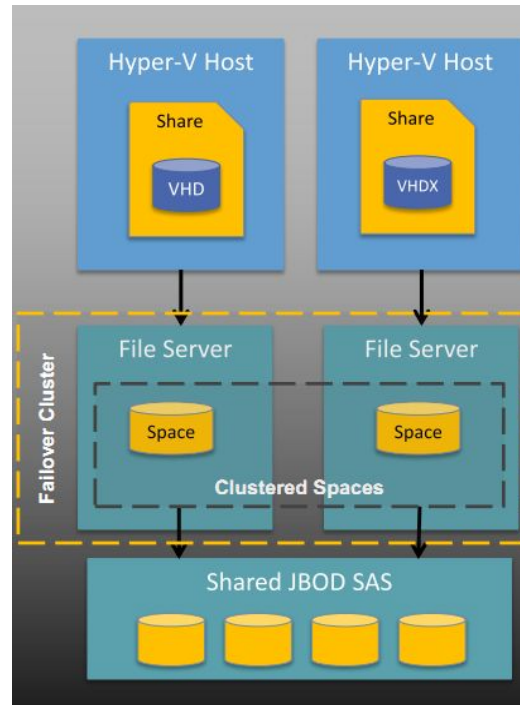
Although not required, failover clustering is supported on the Hyper-V side, the File Services side, or both. They must be separate clusters.

# SMB Hyper-V Cluster

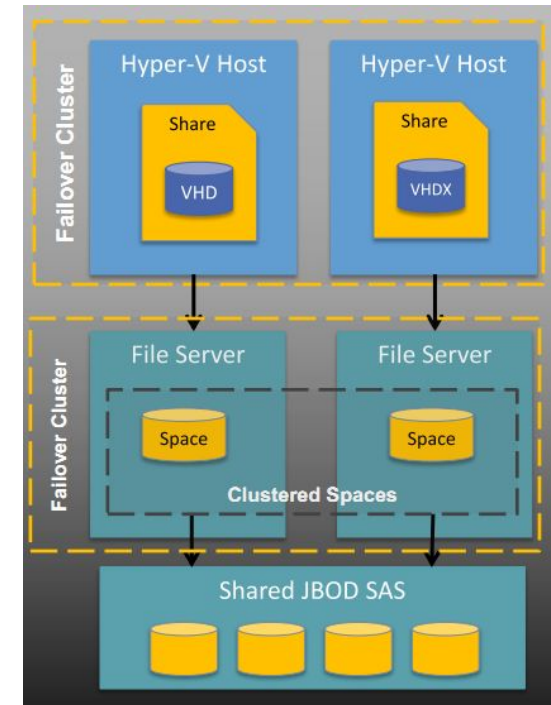
Standalone File Server/  
Clustered Hyper-V



Clustered File Server/  
Standalone Hyper-V



All Clustered



# Clustering Hyper-V with SMB

- Cluster is a “storageless” cluster.
- Configure a file server witness quorum model.

# CSV Cache

- Block level read cache for CSV
- Takes up to 20% of file server memory
- Drastic reductions in IO from reads
- Configure **SharedVolumeBlockCacheSize** cluster property
- Set **CSVEnableBlockCache** to 1 for each CSV Volume



***Microsoft***<sup>®</sup>

Be what's next.<sup>™</sup>