vmworld 2019

BCA2861BU

Supporting SQL Server Failover Cluster Instances

Without Those Pesky RDMs

Rob Girard - Principal TME @ Tintri by DDN
Shawn Meyers – Principal Architect @ House of Brick Technologies

#vmworldBCA2861BU

vmware

About Rob

Rob Girard

@robgirard www.linkedin.com/in/robgirard

- Principal Technical Marketing Engineer @Tintri
- Working in IT since 1997 with 12+ years of VMware experience
- vExpert, VCAP4/5-DCA, VCAP4-DCD, VCP2/4/5, MCSE, CCNA AND TCSE

About Shawn

Shawn Meyers

@1dizzygoose linkedin.com/in/shawnmeyers42

- SQL Server Principal Architect, practice lead
- Experience in VMware, Microsoft, SQL Server, storage infrastructure, performance tuning
- Working in IT since 1992, SQL Server since 1996, VMware since 2009
- vExpert, VCP 4/6.5, MCSE

About House of Brick

Click to edit optional subtitle

- Founded in 1998
- Oracle, Microsoft, VMware, and Cloud
- Focus on Business Critical Applications
- House of Brick Services
  - Oracle and MS license review and optimization
  - Application re-platforming and data migration
  - Virtualization of business critical applications
  - HA, DR, and private cloud architectures
  - Managed services for operations and license management

Apology

Abstract change

- Topics for VMworld with abstracts are submitted in April
- Talking about features in the next version can always be tricky
- We had to make some changes to present to you what is actually available

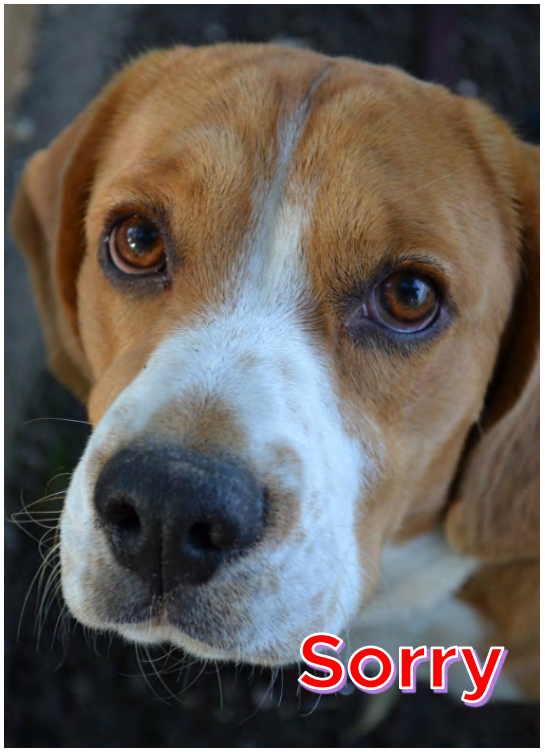

Abstract [ORIGINAL]

Finally! Shared VMDK Supporting SQL Server Without Those Pesky RDMs

VMware now allows VMDKs to support SCSI3 reservations. This cloud-friendly enhancement now allows you to do traditional SQL Server Failover Clusters (FCI) without RDMs or any other fancy technologies or understanding of the underlying physical infrastructure. Simply share the same VMDK you have come to love between 2 VMs on hosts that have access to the same datastore, and you're good to go!. In this deep dive session we will crawl through performance testing to validate performance with shared VMDK vs physical vs stand alone VMDK. This session will not only walk you through how to configure SQL FCI with a shared VDMK, but will also provide best practices to optimize, protect and monitor performance of VMs using shared VMDKs. A must see for Virtual SQL DBAs and virtual admins supporting SQL DBAs.

Abstract [Revised]

Supporting SQL Server Failover Cluster Instances Without Those Pesky RDMs

RDMs, in-guest mounts, cluster-in-a-box… these are the technologies that enable virtualized SQL Server Failover Cluster Instances (FCI) today. What if VMware supported SCSI3 reservations for VMDKs? This enhancement already allows shared VMDKs across VMs on different hosts in VMC and should be available for your on-prem environment in the near future. This session will delve into SQL Server FCI in a VMware environment using existing & future technologies. This session will not only walk you through how to configure SQL FCI using the future shared VDMK option, but will also provide best practices to optimize, protect and monitor performance of VMs using shared VMDKs. A must see for Virtual SQL DBAs and virtual admins supporting SQL DBAs.

Definitions

- AlwaysOn vs Always On – what is the difference
- WSFC – Windows Server Failover Cluster
- MCSC – Windows Server 2003 version of clustering
- FCI – Failover Cluster Instances
- AG – Availability Group
- RDM-V RDM-P
- Cluster-in-a-box
- VVols / vVols
- Storage terms, we will use a lot, assuming this audience knows all of these

Configure Shared Disk on Virtual Machine

History of SQL Server Clustering Part 1 of 2

A walk down memory lane

SQL Server 6.5 and Windows NT 4.0 is where it all started, called MSCS (Microsoft Clustering Server)

- Basically unusable
- SQL Server 2000 on Windows Server 2000 when it became something you could use
- Greatly improved on SQL Server 2005 and Windows Server 2008
- Improvements have been small since then and basically the same
- When physical, any workload you cared about needed to be a FCI

Mixed Virtual / Physical (called Hybrid)

- Way to start tinkering with virtual technologies
- Primary node was physical and backed by a virtual machine
- Complex storage zoning rules to allow the RDM-P and physical machine to share the same LUN

History of SQL Server Clustering Part 2 of 2

A walk down memory lane

Virtual Servers

- In the beginning RDMs were faster than VMDK
- FCI were easy to do as most storage in VMware was RDM
- This changed in vSphere 4.x
- SAN admins hated RDMs
- VMware admins hated RDMs
- Goal was to use VMware tools to provide HA support for SQL Servers and remove FCIs

Availability Groups

- Availability Groups still require WSFC but no shared storage
- Came around with SQL Server 2012 Enterprise Edition
- Many workloads work great with AGs
- Fastest workloads have problems with latency from AGs, requiring FCIs

Clearing Confusion

RDM support since the start of VMware, back in the GSX days.

- RDMs used to be faster than VMDK
- 2007 is VMDKs caught up, we should no longer ever hear that a RDM can be faster
- VMDKs have a small edge in performance now, but RDMs can still provide solid performance

Oracle RAC support came about with the advent of the multi writer flag in ESX and ESXi 4.0

- Allows for multiple VMs to write to the same VMDK
- This does not help SQL Servers as Microsoft is using a SCSI-3 reservation

SCSI-3 reservations

- Open SCSI standard, set in 1997
- This standard was when the command protocol was separated from the device/connectors
- Allows for reservations for locking
- Microsoft followed this open protocol

Quick High Availability Comparison

Classification	AG	FCI	VMware HA
Failover time	30 seconds	30 seconds	2-5 minutes
Rolling patching	Yes	Yes	No
Licensing impact	SQL Server Enterprise	Can use Standard Edition	Can use Standard Edition
Staff skills	Highest skill	Medium skill	Standard
Reliability/Simplicity	Most complex	Complex	Simplest
Storage space requirement	Double (May be reduced depending on storage)	Single	Single
Resource overhead	Double compute, memory and disk writes Read-Only secondaries can offload reporting	Double compute and memory, single disk writes	Single

vSphere Support of SQL Server FCI Clusters - Part 1 of 2

By version

- 5.5 is not supported by VMware if you are on it, time to upgrade
- 6.0+ Cluster in a box has always worked
- 6.0+ – In guest options
  - In guest I-SCSI
  - SMB 2.0, need to be 3.0 or higher to be functional
  - CSV – Cluster Shared Volume, relies on SMB

  6.0 - Added support for vMotion of VMs with RDM-P
  - Need to be on hardware version 11
  - Only WSFC is supported not MCSC, meaning Windows Server 2008+

vSphere Support of SQL Server FCI clusters – Part 2 of 2

By version

- 6.7 – vVols work with SCSI-3 reservations
- 6.7 - vSAN FCI on iSCSI service
- 6.7 U3 – vSAN FCI on native vSAN datastores
- vNEXT not released yet, news at the end of this session from VMware

In Guest Drive Setup

In-Guest Options

iSCSI Mounts

Pros

- No vSphere limitations
- Tried and true
- Enable FCI between bare-metal servers and VMs
- May leverage storage-level capabilities – varies by vendor

Cons

- Invisible from vSphere
- IQNs can be complex to setup and manage
- May require additional network considerations

In-Guest Options

SMB3 Shares

Pros

- No vSphere limitations
- Storage arrays are presenting the SMB share
- May leverage storage-level capabilities – varies by vendor

Cons

- Many DBAs are not familiar with this and are risk adverse
- Lack of troubleshooting knowledge on SMB share latency and performance
- May require additional network considerations

Cluster-in-a-box

Some say why, we say why not

Pros

- Application availability for in-guest patching and maintenance
- Easy to setup
- Can use a shared VMDK

Cons

- Portability issues (storage vMotion)
- Complicates host patching and maintenance
- Backups that rely on VMW Snapshots (independent:persistent) will not work
- vMotion will not work, need cold migrations and both guests need to move to same host
- Relies on VMware HA for host failures (short outage, but not outage-free)

VVols

Policy based storage

Pros

- Same performance as a VMDK
- All virtualization tools work, these VMs are no different than any other VM with vVols
- Policy driven performance
- All benefits of vVols come into play

Cons

- Support varies by storage vendor
- Many organizations will need to update storage software to utilize this once on vSphere 6.7

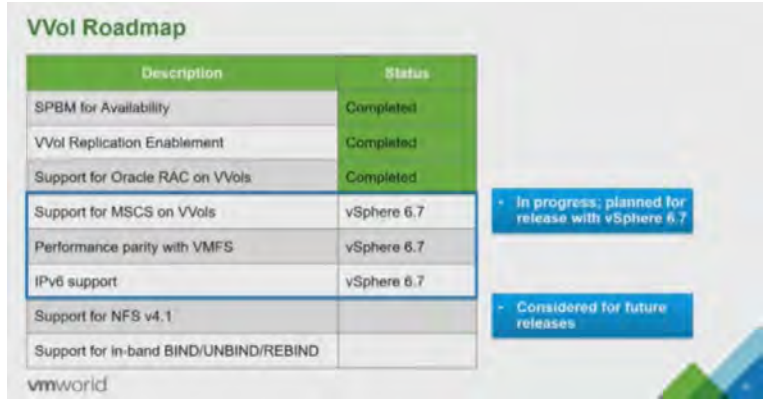

Setting up WSFC – Part 1 of 2

Vvols Specific Guidelines

- Be sure to use HW Version 13 or greater for your FCI VMs
- ESXi 6.7 supports VVol storage with up to 5 node WSFC
- Storage array must support SCSI persistent operations at the subsidiary LUN level
- ESXi 6.7 supports VVol storage for Windows Server 2008 SP2 and above releases
- All hosts must be running ESXi 6.7 or above
- WSFC on VVols can work with any type of disk, “Thin” as well as “Thick” provisioned disks
- The underlying transport protocol can be FC, ISCSI, or FCOE
- Only cluster-across-box (CAB) is supported

Setting up WSFC – Part 2 of 2

Vvols Specific Guidelines

- Cluster-in-a-box (CIB) and a mixture of CAB and CIB is not supported
- N+1 cluster configuration, in which oneESXihost has virtual machines which are secondary nodes and one primary node is a physical box is not supported
- This feature enables customers to move away from using Pass-through RDM (physical compatibility mode)
- WSFC on Vvols supports HA, DRS and vMotion

Source: https://blogs.vmware.com/virtualblocks/2018/05/31/scsi-3-vvols/#respond

vSAN v6.7 (initial release) - In-Guest iSCSI

Pros & Cons: Refer to previous in-guest iSCSI options

Reference:

- https://storagehub.vmware.com/t/vmware-vsan/sql-server-fci-and-file-server-on-vmware-vsan-6-7-using-iscsi-service/

  https://storagehub.vmware.com/t/vmware-vsan/sql-server-fci-and-file-server-on-vmware-vsan-6-7-using-iscsi-service/sql-server-and-file-server-virtual-machine-configuration/

vSAN v6.7U3 Native Shared VMDK

Released on August 20, 2019

Pros

- Vastly improved vSAN performance from previous versions of vSAN
- Runs native, no need to setup iSCSI initiators in vSAN

Cons

- New and so far unproven, but time will tell
- *Everything else omitted from this section*

Validation and WSFC Build

Official Supportability

https://kb.vmware.com/s/article/2147661

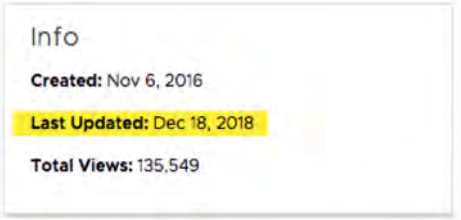

Windows Version	Minimum vSphere Version	WSFC	FCI Mode	Always On Availability Group
Windows 2016	vSphere 6.5 Patch 1 or vSphere 6.7	Yes	SQL 2017, SQL 2016	SQL 2017, SQL 2016
Windows 2012	vSphere 6.5 or vSphere 6.7	Yes	SQL 2017, SQL 2016	SQL 2017, SQL 2016
Windows 2016	vSphere 6.0 GA	Yes	SQL 2016	SQL 2016
Windows 2012	vSphere 6.0 GA	Yes	SQL 2016	SQL 2016

Official Supportability

https://kb.vmware.com/s/article/2147661

This table outlines VMware vSphere support for Microsoft clustering solutions:

	Microsoft Clustering on VMware	vSphere support	VMware HA support	vMotion DRS support	Storage vMotion support	WSFC Node Limits	Storage Protocols support						Shared Disk	
							FC	In-Guest OS iSCSI	Native iSCSI	In-Guest OS SMB	FCoE	NFS	RDM	VMFS
Shared Disk	WSFC with Shared Disk	Yes	Yes[1]	Yes[5]	No	5	Yes	Yes	Yes	Yes[4]	Yes	No	Yes[2]	Yes[3]
	Exchange Single Copy Cluster	Yes	Yes[1]	Yes[5]	No	5	Yes	Yes	Yes	Yes[4]	Yes	No	Yes[2]	Yes[3]
	SQL Clustering	Yes	Yes 1	Yes[5]	No	5	Yes	Yes	Yes	Yes[4]	Yes[8],9	No	Yes[2]	Yes[3]
Non shared Disk	Network Load Balance	Yes	Yes 1	Yes	Yes	Same as OS/app	Yes	Yes	Yes	N/A	N/A	No	N/A	N/A
	Exchange CCR	Yes	Yes[1]	Yes	Yes	Same as OS/app	Yes	Yes	Yes	N/A	N/A	No	N/A	N/A
	Exchange DAG	Yes	Yes[1]	Yes	Yes	Same as OS/app	Yes	Yes	Yes	N/A	N/A	No	N/A	N/A
	SQL AlwaysOn Availability Group	Yes	Yes[1]	Yes	Yes	Same as OS/app	Yes	Yes	Yes	N/A	Yes[7]	Yes[6]	N/A	N/A

Info

Created: Nov 6, 2016

Last Updated: Dec 18, 2018

Total Views: 135,549

vNext

What is to come

Stay tuned for more after the show…

More cool stuff is coming just not sure when

vmworld 2019

Thank you!

vmware