

VMware Storage



Table of contents

grate a WSFC Cluster on RDMs to vVols	3
roduction	3
Preparation	3
Migration	8
Closing	20
Video	20



Introduction

Microsoft Windows Server Failover Clustering (WSFC) on VMware vSphere 6.x requires the use of SCSI-3 Persistent Reservations to be able to share disks between nodes in the cluster. One of the most common, not necessarily popular, ways to achieve this is using RDMs. RDMs allows SCSI commands to be directly passed to a LUN using SCSI-PRs. This allowed WSFC to directly interact with the array for the shared disks. With the release of vSphere 6.7, Virtual Volumes (vVols) added support for SCSI-3 PR and running WSFC. The question then arises, how can you migrate from using RDMs and to vVols? In this article, I will go over the steps and show you how easy it is.

Preparation

Before migrating, first and foremost, make sure you have a full backup of your WSFC. Next, it's critically important to capture all the details regarding your shared disks in the cluster.

In Figures 1 and 2, you can see these disks are physical LUNs. Disk 2 is on SCSI controller 1, channel 0 (1:0), and disk 3, is on the same controller but channel 1 (1:1). These details are crucial to ensure the cluster will resume without any issues after migrations.

Edit Settings ws2016-c1N2		×
Virtual Hardware VM Options		
		ADD NEW DEVICE
> CPU		0
> Memory	8 <u> </u>	
> Hard disk 1	91 <u>GB ~</u>	
v Hard disk 2		
VM storage policy	Datastore Default ~	
Sharing	No sharing 🖌	
Physical LUN	vml.0200fc0000624a937080076c6c13e4455b0	0011824466c61736841
Compatibility Mode		
Shares	Normal ~	
Limit - IOPs	Unlimited ~	
Virtual Device Node	SCSI controller 1 SCSI(1:0) Hard disk 2	,

Figure 1



Edit Settings ws2016-c1N2	
Virtual Hardware VM Options	
	ADD NEW DEVICE
> CPU	<u>2 ×</u> ()
> Memory	8 <u>GB ~</u>
> Hard disk 1	91 <u>GB ~</u>
> Hard disk 2	5 GB ~
V Hard disk 3	<u>100</u> <u>GB v</u>
VM storage policy	Datastore Default ~
Sharing	
Physical LUN	vml.0200f90000624a937080076c6c13e4455b00012156466c6173684 1
Compatibility Mode	
Shares	Normal v 1000
Limit - IOPs	
Virtual Device Node	SCSI controller 1 V SCSI(1:1) Hard disk 3 V

Figure 2

Reviewing the second node in the WSFC, we can see the RDM is the same for each disk and as are the SCSI controller configurations.



Edit Settings ws2016-c1N2	
Virtual Hardware VM Options	
	ADD NEW DEVICE
> CPU	<u>2 v</u> ()
> Memory	8 <u>GB ~</u>
> Hard disk 1	91 <u>GB ~</u>
✓ Hard disk 2	
VM storage policy	Datastore Default
Sharing	
Physical LUN	vml.0200fc0000624a937080076c6c13e4455b00011824466c61736841
Compatibility Mode	Dhustest xx yml 0200fc0000624a937080076c6c13e4455b00011824466c61736841
Shares	
Limit - IOPs	
Virtual Device Node	SCSI controller 1 👻 SCSI(1:0) Hard disk 2 🛩



Edit Settings ws2016-c1N2	
Virtual Hardware VM Options	
	ADD NEW DEVICE
> CPU	2 ~ 0
> Memory	8 <u>GB ~</u>
> Hard disk 1	91 <u>GB ~</u>
> Hard disk 2	5 GB ~
✓ Hard disk 3	100 <u>GB v</u>
VM storage policy	Datastore Default ~
Sharing	No sharing 🖌
Physical LUN	vml.0200f90000624a937080076c6c13e4455b00012156466c6173684
	vml.0200f90000624a937080076c6c13e4455b00012156466c6173684
Compatibility Mode	
Shares	Normal 🗸
Limit - IOPs	Unlimited ~
Virtual Device Node	SCSI controller 1 V SCSI(1:1) Hard disk 3 V

Figure 4

In the following figures, you can see the status of the cluster before migration.

🗟 Failover Cluster Mana	ger				
File Action View H	lelp				
🗢 🔿 🙍 🖬 🛛					
📲 Failover Cluster Mana	Roles (1)				
✓ WS2016-C1.vsanp	Search			2	Queries 🔻 🗔
Roles					
Nodes	Name	Status	Туре	Owner Node	Priority
🗸 📇 St 🚽 ge	NS2016-C1-FS	(f) Running	File Server	WS2016-C1N2	Medium
Disks	-	<u> </u>			
Pools					
Enclosures					
Networks					
Cluster Events					





		P.
Status	Assigned Vote	Current Vote
💿 Up	1	1
🛞 Up	1	1
	Status Up Up Up	Status Assigned Vote Up 1 Up 1 Up 1

Figure 6

Failover Cluster Mana WS2016-C1.vsanp	Disks (2) Search			PQuerie	s 🔻
Nodes	Name	Status	Assigned To	Owner Node	1
🗸 📇 Storage	📇 Cluster Disk 1	() Online	Disk Witness in Quorum	WS2016-C1N1	
Pooss	📇 Cluster Disk 2	Online	WS2016-C1-FS	WS2016-C1N2	
Enclosures					
Networks					
B Cluster Events	1				
Figure 7					

Now that you have gathered the vital information regarding the cluster and each node's disk configurations, you can shut down the WSFC.

Kailover Cluster N	lanager				
File Action View	Help				
🗢 🄿 🙍 🖬	?				
Hailover Cluster N	lana Cluster WS2016-C1.vsa	npe.vmware.	com		
✓ WS2016-C1.ve	200	_			
📷 Roles	Configure Role	ster WS	2016-C1		
Nodes	Validate Cluster	lustered ro	les and 2 n	odes.	
V 📇 Storage	View Validation Report	mware.com		Networks: Cluster Netwo	ork 1
	Add Node	016-C1N1		Subnets: 1 IPv4 and 0 I	Pv6
Enc		one in the las	t hour	Storage Spaces Direct	(S2D): Disabled
Netwo	Close Connection				
国 Cluster	Reset Recent Events				
	More Actions	> Config	gure Cluster	Quorum Settings	, or copy roles from a cluster
	View	> Copy	Cluster Roles	5	
	Refresh	Shut D	own Cluster	r	the web
	Properties	Destro	y Cluster		
	Help	Move	Core Cluster	Resources	•
	1 Cluster-Aware Updatin	Cluste	r-Aware Upo	dating	

Migration

Now that you have gathered the important information regarding the cluster and each node's disk configurations, you can shutdown the WSFC.

唱Fi	ailover Cl	uster Ma	nager					
File	Action	View	Help					
(=) 🖄	?						
醖 Fa	ilover Clu	uster Mar	Cluster WS2016-C1.vsan	pe.v	nware.com			
~ 8	WS2016	-C1 vean	Configure Pole	٦.				
	Role	ls	Configure Role	iste	er WS2016-C1			
	Mod Not	les	Validate Cluster	lus	tered roles and 2 no	odes.		
~	Stor	age	View Validation Report	mwa	are.com	Networks: Cluster Netw	vork	:1
		Der	Add Node	016	C1N1	Subnets: 1 IPv4 and 0	IPv	6
	E C		Additionen	me in the last hour Storage Spaces Direct		± (t (S2D): Disabled	
	Net	wo	Close Connection					
	Clus	ter	Reset Recent Events	_				
			More Actions		Configure Cluster	Quorum Settings		, or copy roles from a cluster
			View 2		Copy Cluster Roles	***		the Web
			Refresh		Shut Down Cluster			the web
			Properties		Destroy Cluster			
			Help		Move Core Cluster	Resources	>	
			Cluster-Aware Updating	L	Cluster-Aware Upd	lating		

Figure 8

Once the cluster has been shut down, you will then power off all the nodes in the WSFC. Because of the shared disks, there is no live option to migrate. After powering off all the nodes, you will then go to each secondary node, nodes accessing the primary node's disks, and remove the shared disks from their configuration. It is very important you <u>do not select</u> "Delete files from datastore", you are removing the shared disks from the secondary nodes only. Leave the shared disks (RDMs) attached to the primary node.

Edit Settings ws2016-c1N2		>
Virtual Hardware VM Options		
		ADD NEW DEVICE
> CPU	<u>2 ×</u>	0
> Memory	8 <u> </u>	
> Hard disk 1	<u>91</u>	
> Hard disk 2 *	Device will be removed Delete files from	n datastore 🛛 🖍
> Hard disk 3 *	Device will be removed Delete files from	n datastore
> SCSI controller 0	LSIL opic SAS	

Figure 9

Once you have removed all the shared disks from all secondary nodes, you may then initiate the storage migration of the primary node. Because the pRDMs are still attached to the primary node, when you do a storage migration, the data from the disk is copied to the new destination datastore.





Figure 10

When do the storage migration, you are only changing the storage.

Migrate WS2016-C1	N1
1 Select a migration type 2 Select storage	Select a migration type Change the virtual machines' compute resource, storage, or both.
3 Ready to complete	Change compute resource only Migrate the virtual machines to another host or cluster.
	Change storage only Migrate the virtual machines' storage to a compatible datastore or datastore cluster.
	Change both compute resource and storage Migrate the virtual machines to a specific host or cluster and their storage to a specific datastore or datastore cluster.

Figure 11

For the destination, select your vVols datastore with an array that supports SCSI-3 PR with vVols. With vVols you do not need to use the "Configure per disk" option.



Migrate WS2016-C1	IN1						
 1 Select a migration type 2 Select storage 	Select storage VN Select the destination storage for the virtual machine migration.						
3 Ready to complete						Configure per disk	
	Select virtual disk format:		Thin Provision				
	VM Storage Policy:			Keep existing VM storage policies ×			
	Name	Capacity	Provisioned	Free	Туре	Cluster	
	PS-VMFS-x20	1,000.75 GB	403.47 GB	792.47 GB	VMFS 6		
	PS-VMFS-x50	999.75 GB	242.54 GB	954.69 GB	VMFS 6		
	PS-vVols-x20	8 PB	689.94 GB	8 PB	vVol		
	PS-vVols-x50	8 PB	878.3 GB	8 PB	vVol		
	Unity-NFS	100 GB	181.78 GB	55.16 GB	NFS v3		
	Unity-VMFS	99.75 GB	1.41 GB	98.34 GB	VMFS 6		
	Unity-vVols	99 GB	76.36 GB	71.75 GB	vVol		
	bugs	370 TB	348.14 TB	21.86 TB	NFS v3		

Figure 12

vVols are automatically Thin Provisioned for all disks eliminating the need for per disk configuration.

Migrate WS2016-C1N	1				
 1 Select a migration type 2 Select storage 	Ready to complete Verify that the information is correct and click Finish to start the migration.				
3 Ready to complete					
	Migration Type	Change storage. Leave VM on the original compute resource			
	Virtual Machine	WS2016-C1N1			
	Storage	PS-vVols-x20			
	Disk Format	Thin Provision			
				P.L.OK	- Eliment
			CANCEL	BACK	FINISH

Figure 13

Click Finish to initiate and complete the migration of the primary node and its disks. Once the migration has completed, you can go into the VM's properties and verify the shared disks have been migrated. In Figures 14 and 15 you can see the shared disks are now on the vVols datastore and are no longer RDMs. Also notice the shared disks, now on the vVols datastore, are on the same SCSI controller channels where the RDM previously were.



E	dit Settings ws2016-c1N1		
、			
-		٢	
	> CPU	2 ~	
	> Memory	8 × GB ×	
	> Hard disk 1	91 GB ~	
	✓ Hard disk 2	5 GB ~	
	Maximum Size	62 TB	
	VM storage policy	VVol No Requirements Policy ~	
	Туре	Thin Provision	
	Sharing	No sharing V	
	Disk File	[PS-vVols-x20] rfc4122.829b0f73-552e-4deb-9583 ab8a7f20b0be/WS2016-C1N1_1.vmdk	
	Shares	Normal ~ _ 1000	
	Limit - IOPs	Unlimited ~	
	Disk Mode	Independent - Persistent	
	Virtual Device Node	SCSI controller 1	



Edit Settings ws2016-c1N1		>
Virtual Hardware VM Options		
	ADD NEW DEVICE	
> CPU	2 ~	0
> Memory	8 × <u>GB ×</u>	
> Hard disk 1	91 GB ~	
> Hard disk 2	5 <u>GB ~</u>	
✓ Hard disk 3		
Maximum Size	62 TB	
VM storage policy	VVol No Requirements Policy ~	
Туре	Thin Provision	
Sharing	No sharing ~	
Disk File	[PS-vVols-x20] rfc4122.829b0f73-552e-4deb-9583- ab8a7f20b0be/WS2016-C1N1_2.vmdk	
Shares	Normal ~ 1000	
Limit - IOPs	Unlimited ~	
Disk Mode	Independent - Persistent V	
Virtual Device Node	SCSI controller 1 × SCSI(1:1) Hard disk 3 ×	

Figure 15

Now that you have migrated the primary node with the shared disks, you can proceed with migrating all remaining secondary nodes in the WSFC. Follow the same process making sure to select the same destination datastore as the primary node.





Figure 16

Μ	igrate WS2016-C1N	12						
	1 Select a migration type 2 Select storage	Select storage V Select the destination storage for the virtual machine migration.						
	3 Ready to complete	Calast vistual dials formats			This Description		Configure per disk	
		Select virtual disk format: VM Storage Policy:		Keep existing VM storage policies				
		Name	Capacity	Provisioned	Free	Туре	Cluster	
		🛢 76-DS1	104.25 GB	5.3 GB	98.95 GB	VMFS 6		
		PS-VMFS-x20	1,000.75 GB	304.56 GB	826.57 GB	VMFS 6		
		PS-VMFS-x50	999.75 GB	242.54 GB	954.69 GB	VMFS 6		
		PS-vVols-x20	8 PB	894.18 GB	8 PB	vVol		
		PS-vVols-x50	8 PB	878.3 GB	8 PB	vVol		

Figure 17

Once you have completed all the secondary node migrations, you will then go into the secondary VM's configuration and re-attach all shared disks. To do this you "ADD NEW DEVICE" and select "Existing Hard Disk". When going through this process, make sure you add the disks in the same order and with the same configuration prior to the migration. This is a critical step and failing to use the exact same configuration may result in the secondary nodes not coming back online.



Edit Settings ws2016-c1N2		×
Virtual Hardware VM Options		
		ADD NEW DEVICE
> CPU	2 ~	Disks, Drives and Storage
> Memory	<u>8</u> <u>GB ~</u>	Existing Hard Disk
Hard disk 1	91 <u>GB Y</u>	RDM Disk
> SCSI controller 0	LSI Logic SAS	NVDIMM
SCSI controller 1	ISI Logic SAS	CD/DVD Drive



When you add an existing disk, the datastore explorer will open allowing you to navigate to the datastore and primary node folder. There you will select the appropriate shared disk to be added back to the secondary node.

Edit Settings ws2016-c1N2			×
Virtual Hardware VM Options			
Select File			×
Datastores	Contents	Information	
 PS-vVols-x50 Unity-vVols Unity-vVMFS Unity-NFS PS-vVols-x20 Catalog Catalog fcd jm-vrops8 jm-vrops8-old Ubuntu1 Win10-4 Wis2016-C1N1 WS2016-C1N2 WS2016-WSFC 76-DS1 	Image: WS2016-C1N1_1.vmdk Image: WS2016-C1N1_1.vmdk Image: WS2016-C1N1_2.vmdk	Name: WS2016-C1N1_1.vmdk Size:0 B Modified:05/06/2020, 4:12:36 PM Encrypted: No	
File Type: Compatible Virtual Disks(*.vmdk, *	.dsk, *.raw) ▼		
		CANCEL	

An absolutely key step is to make sure and select the correct SCSI controller and channel for that shared disk. This is the information you captured prior to the migration.

		ADD NEW DEVICE
> CPU	<u>2 ~</u>	0
Memory	8 <u> </u>	
> Hard disk 1	91 GB ~	
New Hard disk *		
Maximum Size	62 TB	
VM storage policy	VVol No Requirements Policy	
Sharing	Unspecified ~	
Disk File	[PS-vVols-x20] rfc4122.829b0f73-5 ab8a7f20b0be/WS2016-C1N1_1.vm	52e-4deb-9583- dk
Shares	Normal ~ 1000	
Limit - IOPs	Unlimited ~	
Disk Mode	Dependent ~	
Virtual Device Node	NVME controller 0 V NVME(0:	:0) New Hard disk $ imes $
> SCSI controller 0	IDE 0	

Figure 20



Edit Settings ws2016-c1N2			×
Virtual Hardware VM Options			
			ADD NEW DEVICE
> CPU	2 ~		0
> Memory	8	∽ <u>GB ∽</u>	
> Hard disk 1	91	<u>GB ~</u>	
✓ New Hard disk *			
Maximum Size	62 TB		
VM storage policy	VVol No Require	ements Policy 🗸 🗸	
Sharing	Unspecified ~		
Disk File	[PS-vVols-x20] rf ab8a7f20b0be/V	fc4122.829b0f73-552e-4deb-958 VS2016-C1N1_1.vmdk	33-
Shares	Normal ~ 10	000	
Limit - IOPs	Unlimited ~		
Disk Mode	Dependent		
Virtual Device Node	SCSI controller 1	 SCSI(1:0) New Hard dis 	k ¥
> SCSI controller 0	LSI Logic SAS	SCSI(1:0) New Hard dis SCSI(1:1)	sk
> SCSI controller 1	LSI Logic SAS	SCSI(1:2)	
		SCSI(1:3) SCSI(1:4) SCSI(1:5)	
		SCSI(1:6)	

Figure 21

Once you have added all the shared disks back to all secondary nodes, you may then power on the primary node in the WSFC. After the primary node has been powered on you may also power on your secondary nodes in the cluster. With the cluster powered on, go into the primary node and verify the WSFC and all disks are back online.



📲 Failover Cluster Manager					
File Action View Help					
🗢 🔿 🙍 🖬					
Railover Cluster Manager	Roles (1)				
WS2016-C1.vsanpe.vmware	Search			🔎 Queries 🔻 🗖	
Nodes	Name	Status	Туре	Owner Node	Priority
> 🤮 Storage	B WS2016-C1-FS	() Running	File Server	WS2016-C1N1	Medium
Networks					
in Cluster Events					
Figure 22					





You can also verify the disks in the disk management.



Tile Action	ment /iew Help						-		×
⇔ →									
Volume	Layout	Туре	File System	Status	Capacity	Free Spa	% Free		
	Simple	Basic		Healthy (R 450 MB	450 MB	100 %		
	Simple	Basic		Healthy (E 125 MB	125 MB	100 %		
- (C:)	Simple	Basic	NTFS	Healthy (B 90.42 GB	73.83 GB	82 %		
- C1-Vol0 (F:)	Simple	Basic	NTFS	Healthy (P 99.87 GB	96.82 GB	97 %		
- Qurrom (E:)	Simple	Basic	NTFS	Healthy (P 4.97 GB	4.91 GB	99 %		
				System re		jerne, crash bu	, r 11101	y r arcición	
Disk 1 Basic 4.97 GB Reserved	Qurrom (E:) 4.97 GB NTFS Healthy (Primary Partition)								
Disk 2 Basic 99.88 GB	C1-Vol0 (F:) 99.87 GB NTFS								

Figure 24

You can also do a test failover of a function or node to verify operation.



🖏 Failover Cluster Manager						
File Action View Help						
🗢 🔿 🙍 📰 🚺						
🗟 Failover Cluster Manager	Roles (1)					
VS2016-C1.vsanpe.vmware	Search			F	Queries 🔻	
Roles	News	Ortor	T		Ourse Niede	Director
in Nodes	Name	Status	Type		Owner Node	Phonty
Disks	WS2016	Start Role) Se	rver	WS2016-C1N1	Medium
Pools	0	Stop Role				
Enclosures Networks	22	Add File Share				
Cluster Events		Move	•	Best Possible	Node	
_	3	Change Startup Priority	, 😰	Select Norte		
		Information Details		- 0		
	13	Show Critical Events				
	< 3	Add Storage				>
	🗸 🛄 v 🖸	Add Resource	•		Preferred Owners:	Any node
		More Actions	•			
	Status: 🗙	Remove				
	Priority: Owner No	Properties				
	Client Access IP Addresses	Name: WS2016-C1-FS 10.156.175.24				

Figure 25

Navigating to the secondary node, you can see the active role for the file share has been moved and is running on the secondary node.



📅 Disk Management							-		×
File Action Vi	iew Help								
Volume	Layout	Туре	File System	Status	Capacity	Free Spa	% Free		
-	Simple	Basic		Healthy (R.	. 450 MB	450 MB	100 %		
-	Simple	Basic		Healthy (E.	. 125 MB	125 MB	100 %		
- (C:)	Simple	Basic	NTFS	Healthy (B.	. 90.42 GB	73.08 GB	81 %		
- C1-Vol0 (F:)	Simple	Basic	NTFS	Healthy (P.	. 99.87 GB	96.82 GB	97 %		
- Disk 0									
Basic				(C:)				
Online	450 MB Healthy (Recover	125 MB 90.42 GB NTFS Healthy (FEI System P) Healthy (Boot Page File Crac				mn Priman	Partition		
	reality (necovery random) reality (cri system r reality (boot, rage rile, clash bump, rimary						Fartition	·	
	1		1			- F			_
ODisk 1									
Basic									
Reserved	4.97 GB								
	1								
- Disk 2									
Basic	C1-Vol0 (F:)		6				//////	/////	
99.88 GB Reserved	99.87 GB NTFS	Dartition	///////////////////////////////////////		////////		//////	/////	
neserved	rieditiy (Primary	-artitionj					//////	/////	
									22
CD-ROM 0									~
Unallocated Primary partition									

Figure 26

Once you have verified everything and the WSFC is back online, remember, your RDMs still exist in your environment., they just are not being used. Once you are comfortable with the migration, you can reclaim these LUNs and space.

Closing

Migrating a WSFC from RDMs to vVols is a fairly straight forward and simple process. The biggest benefit is you no longer have to manage LUNs and RDMs! You can vMotion nodes within your vSphere cluster without worrying if the RDM is accessible on all vSphere hosts. Putting hosts in Maintenance Mode is also supported as long as the vVols datastore has been mounted to all hosts in the cluster, and of course, have more hosts than WSFC nodes.

For more detail on the specific supported features with Microsoft WSFC, please see KB 2147661 Microsoft Windows Server Failover Clustering (WSFC) with shared disks on VMware vSphere 6.x: Guidelines for supported configurations.

Video

Below is a video of the process showing all the steps to migrate from RDMs to vVols.



