vSphere 7 RDM to Shared VMDK Migration

VMware Storage

Table of contents

vSphere 7 RDM to Shared VMDK Migration

Introduction

With vSphere 7, one of the new core storage features is shared VMDKs on VMFS. With many clustering applications, such as Microsoft's Windows Server Failover Cluster (WSFC), SCSI-3 persistent reservations (SCSI3-PR) are required. SCSI3-PR allows multiple servers to share disks between them with the application managing IO priority between nodes. This requirement is one of the primary use cases for RDMs. In vSphere 6.7, VMware announced SCSI-3 PR support for vVols and validated support for WSFC. See more detail here. With vSphere 7, VMware has added SCSI-3 PR support on VMFS, allowing for shared VMDKs to be used with WSFC on VMFS, initially using FC connectivity. This move is yet another to reduce the requirement for RDMs in the virtual environment. To read more about shared VMDKs, and the requirements, please refer to the article here.

Preface

The process outlined here is storage vendor agnostic. Some of our storage partners may have other, specific methods to migrate off RDMs. As always, please make sure to have backups of all your data and systems. This process uses Storage vMotion to migrate the disks from RDM to shared VMDK.

A few notes about this demo

Non-shared disks do not have to be EZT, but they should be on a separate SCSI controller from the shared disks. In the demo, the primary controller is NVMe. Subsequently, the first SCSI controller, for the shared disks, is 0, not 1.

Preparing

To prepare for the migration, you need to capture all the shared disk details. What SCSI controller the disks are attached and on what channel of that controller. These details are critical and must be captured so they may be duplicated when reattaching the disks to the secondary nodes. In this example, you see disk 2 is on SCSI controller 0:1 and disk 3 is SCSI 0:2. These settings should be the same across all WSFC nodes.

Before you can migrate a shared disk to VMFS, you must prepare the destination datastore by enabling "Clustered VMDK" functionality. At this time, this feature is only available on datastores connected via FC, and the datastore must be VMFS6. Make sure your destination datastore has enough space for all the disks/VMs being migrated.

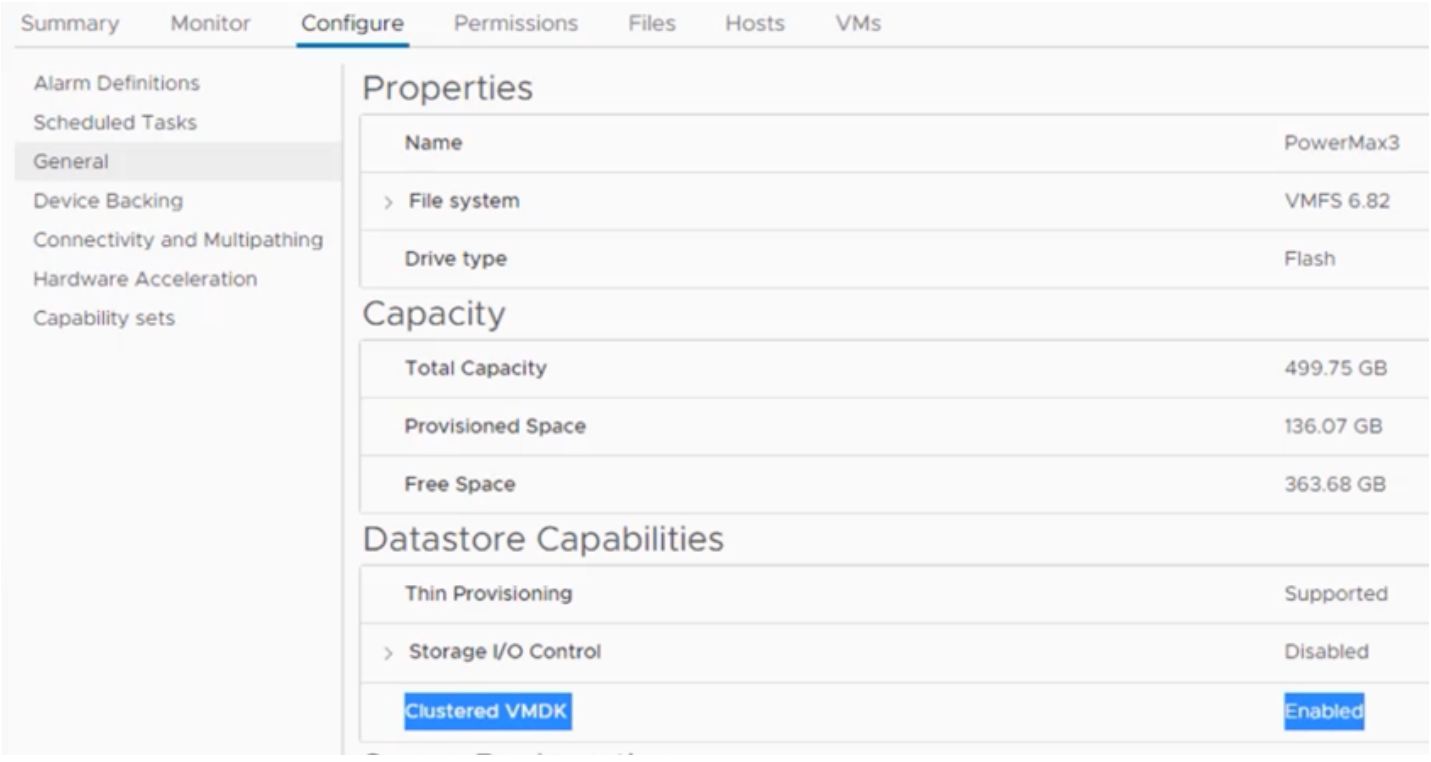

Figure 1

With “Clustered VMDK” enabled, you may proceed with the migration.

WSFC VMware Resources

Docs.VMware.com

- About Setup for Windows Server Failover Clustering on VMware vSphere
- vSphere WSFC Setup Checklist

VMware KB WSFC Articles

- Microsoft Windows Server Failover Clustering (WSFC) with shared disks on VMware vSphere 7.x: Guidelines for supported configurations (79616)
- Microsoft Windows Server Failover Clustering (WSFC) with shared disks on VMware vSphere 6.x: Guidelines for supported configurations (2147661)
- Configuring a shared disk resource for Windows Server Failover Cluster (WSFC) and migrating SQL Server Failover Cluster Instance (FCI) from SAN (RDMs) to vSAN (74786)

Blogs

- Hosting Windows Server Failover Cluster (WSFC) with shared disks on VMware vSphere: Doing it right!
- vSphere 7 – RDM to Shared VMDK Migration
- Using the Perennially Reserved Flag for WSFC RDMs.
- Migrating WSFC RDMs to vVols
- Deploying a Windows Server Failover Cluster (WSFC) in the Supported Configuration on VMware Cloud on AWS and vSAN
- Oracle on VMware Collateral – One Stop Shop

Migration

The WSFC service, and all VMs hosting nodes of the WSFC cluster, must be shut down. You cannot use storage vMotion on shared disks actively in use. Migration time is entirely dependant on the size of the disks and the network.

Next, you need to remove, NOT DELETE, all shared disks from all secondary nodes. Secondary nodes are nodes in the WSFC cluster, sharing disks from the primary node. DO NOT remove the shared disks from the primary node. You must leave the shared disks (RDMs) attached to the primary node for the migration to succeed.

When removing the shared disks from the secondary nodes, make sure you DO NOT check the “Delete files from datastore,” you only want to remove the disk from the VM, not delete it.

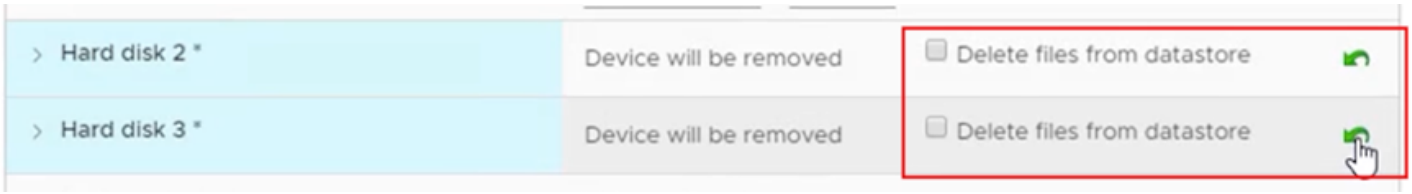

Figure 2

Once the shared disks have been removed from all secondary nodes, you may then initiate a storage vMotion of the primary node.

Initiate a migration choosing “Change storage only.”

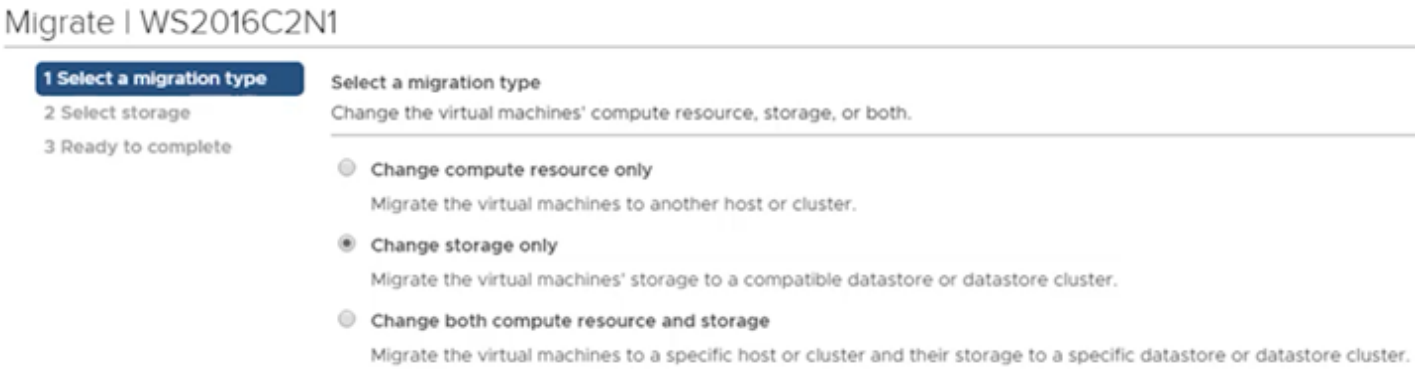

Figure 3

On the next screen, you will need to enable “Configure per disk.”

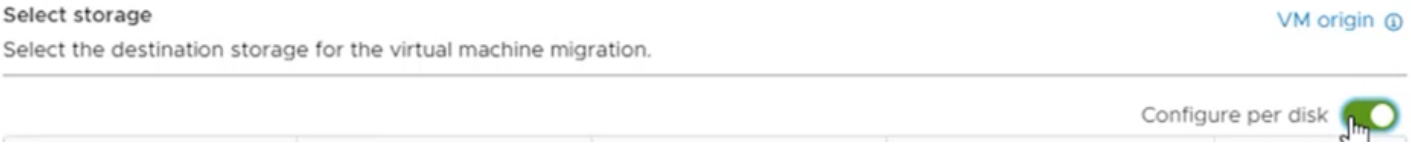

Figure 4

Then for each disk, you will select the destination datastore and configure each shared disk to use “Thick Provision Eager Zeroed.” EZT is required for shared VMDKs on VMFS. If you do not select EZT, the WSFC will fail to start. Your non-shared disks do not have to be EZT, they can be thin or LZT provisioned.

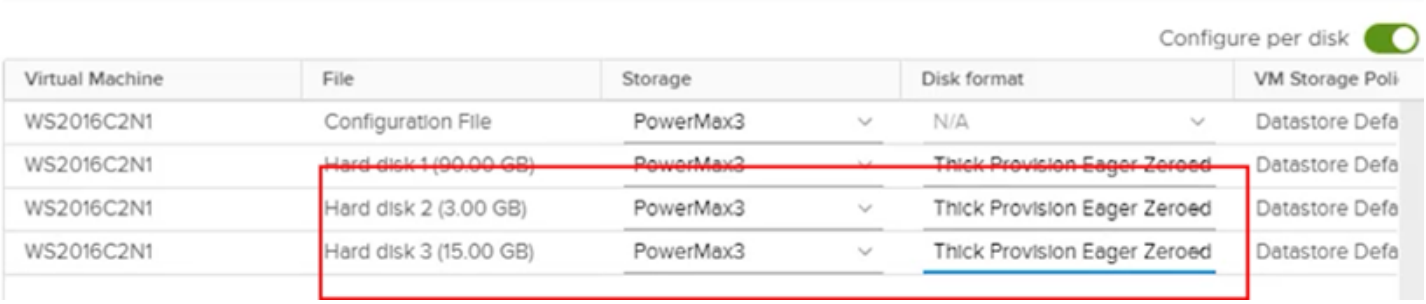

Figure 5

Once the primary node migration has completed, review the VM’s hardware to verify the shared disks, previously using RDMs, are now a standard VMDK located on the new destination datastore.

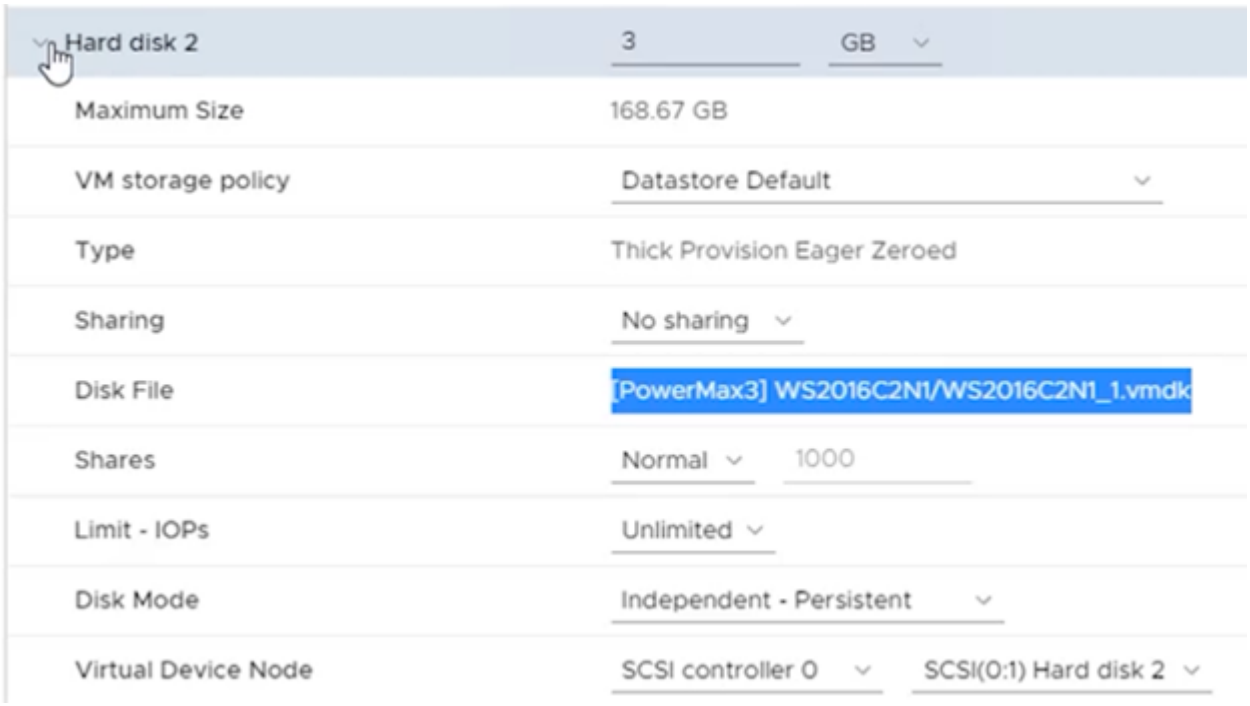

Figure 6

Now migrate all remaining secondary nodes in the WSFC cluster, making sure to follow the same process of "Configure per disk" and selecting the same destination datastore if the storage location of non-shared disks should be changed as well

With all the nodes migrated, you now must reattach the shared disks to all secondary nodes. Remember, all disks must be reattached to the same SCSI controller and channel previously used.

Go into the VM's hardware and under "Add New Device" select "Existing Hard Disk."

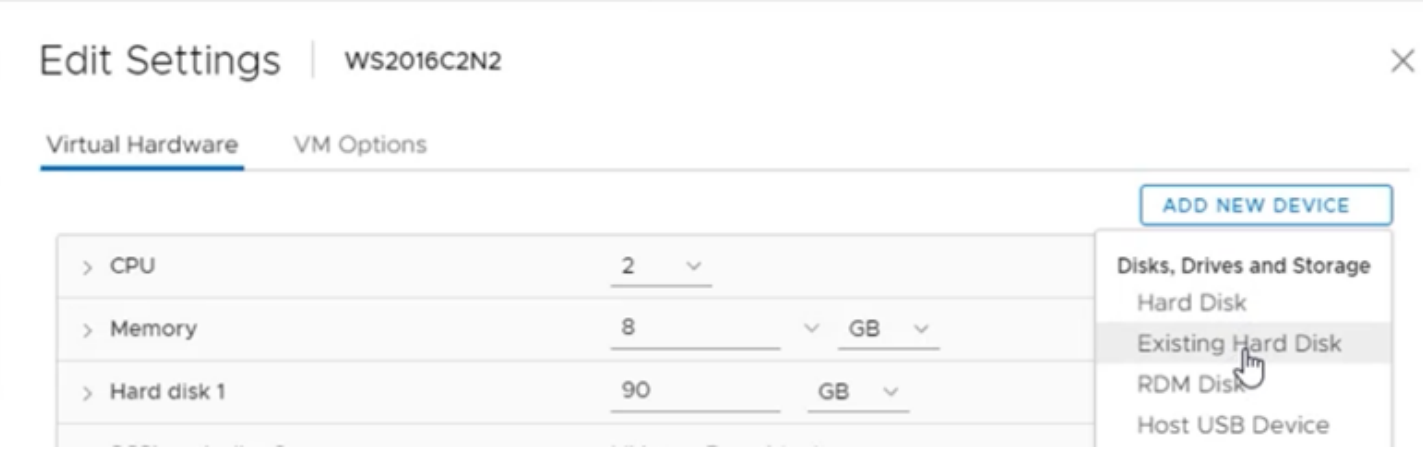

Figure 7

You will then navigate to the new datastore, find the primary node, and attach the disks in the exact same configuration previously used.

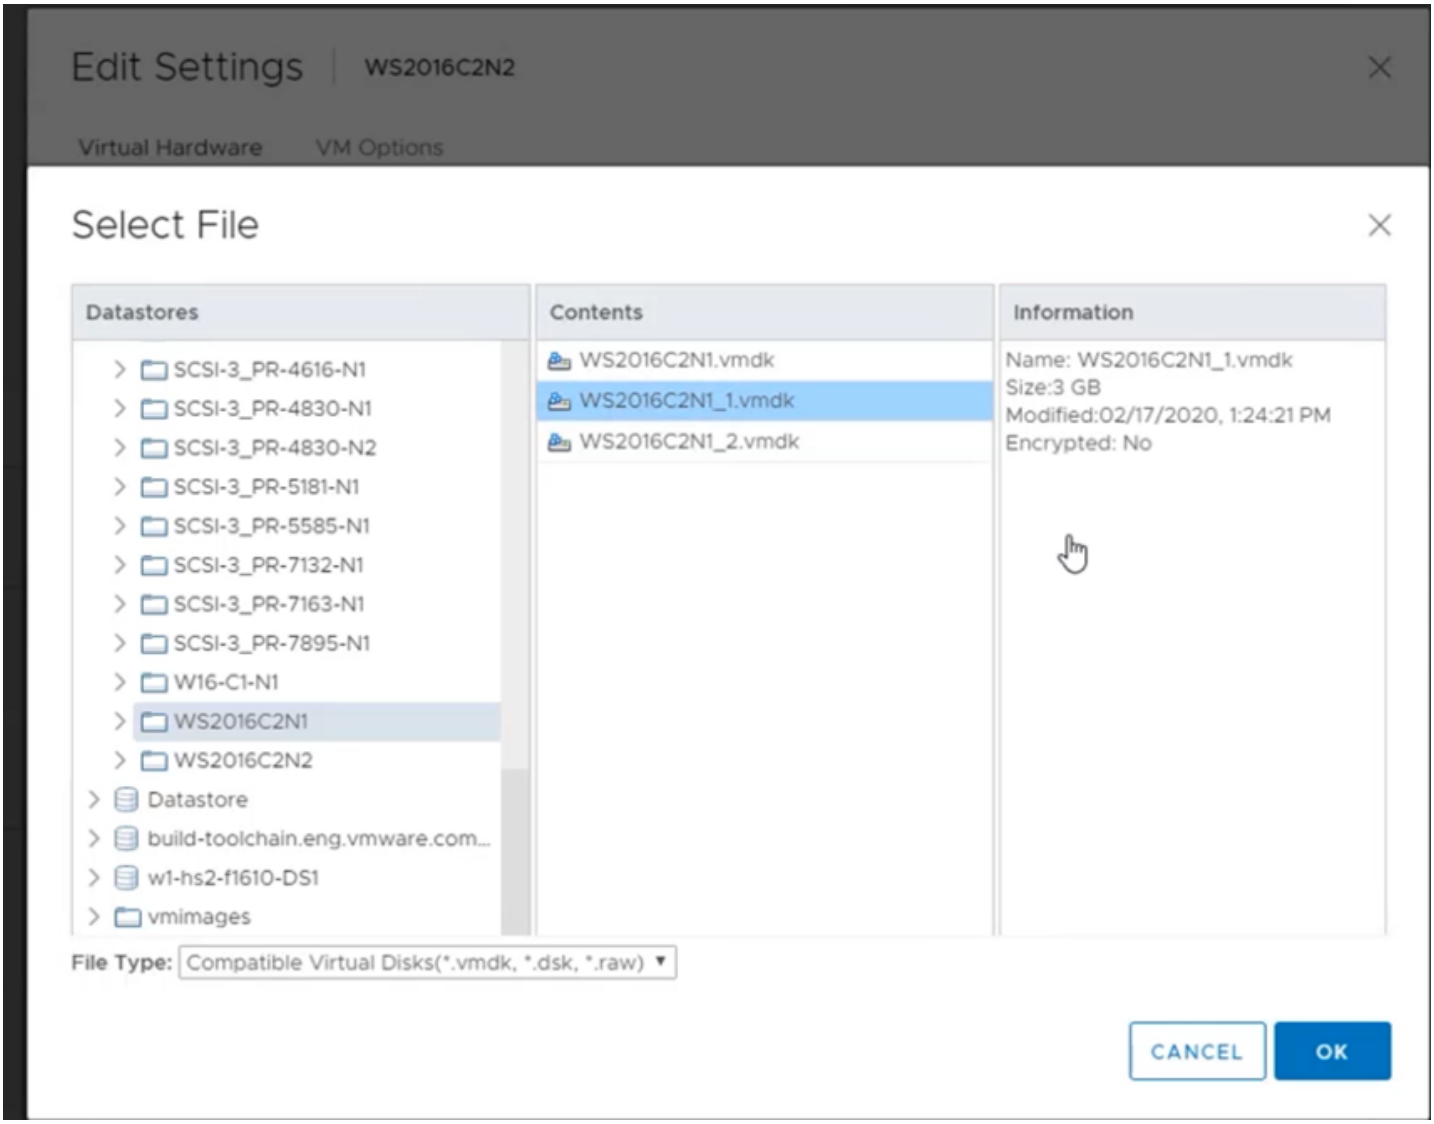

Figure 8

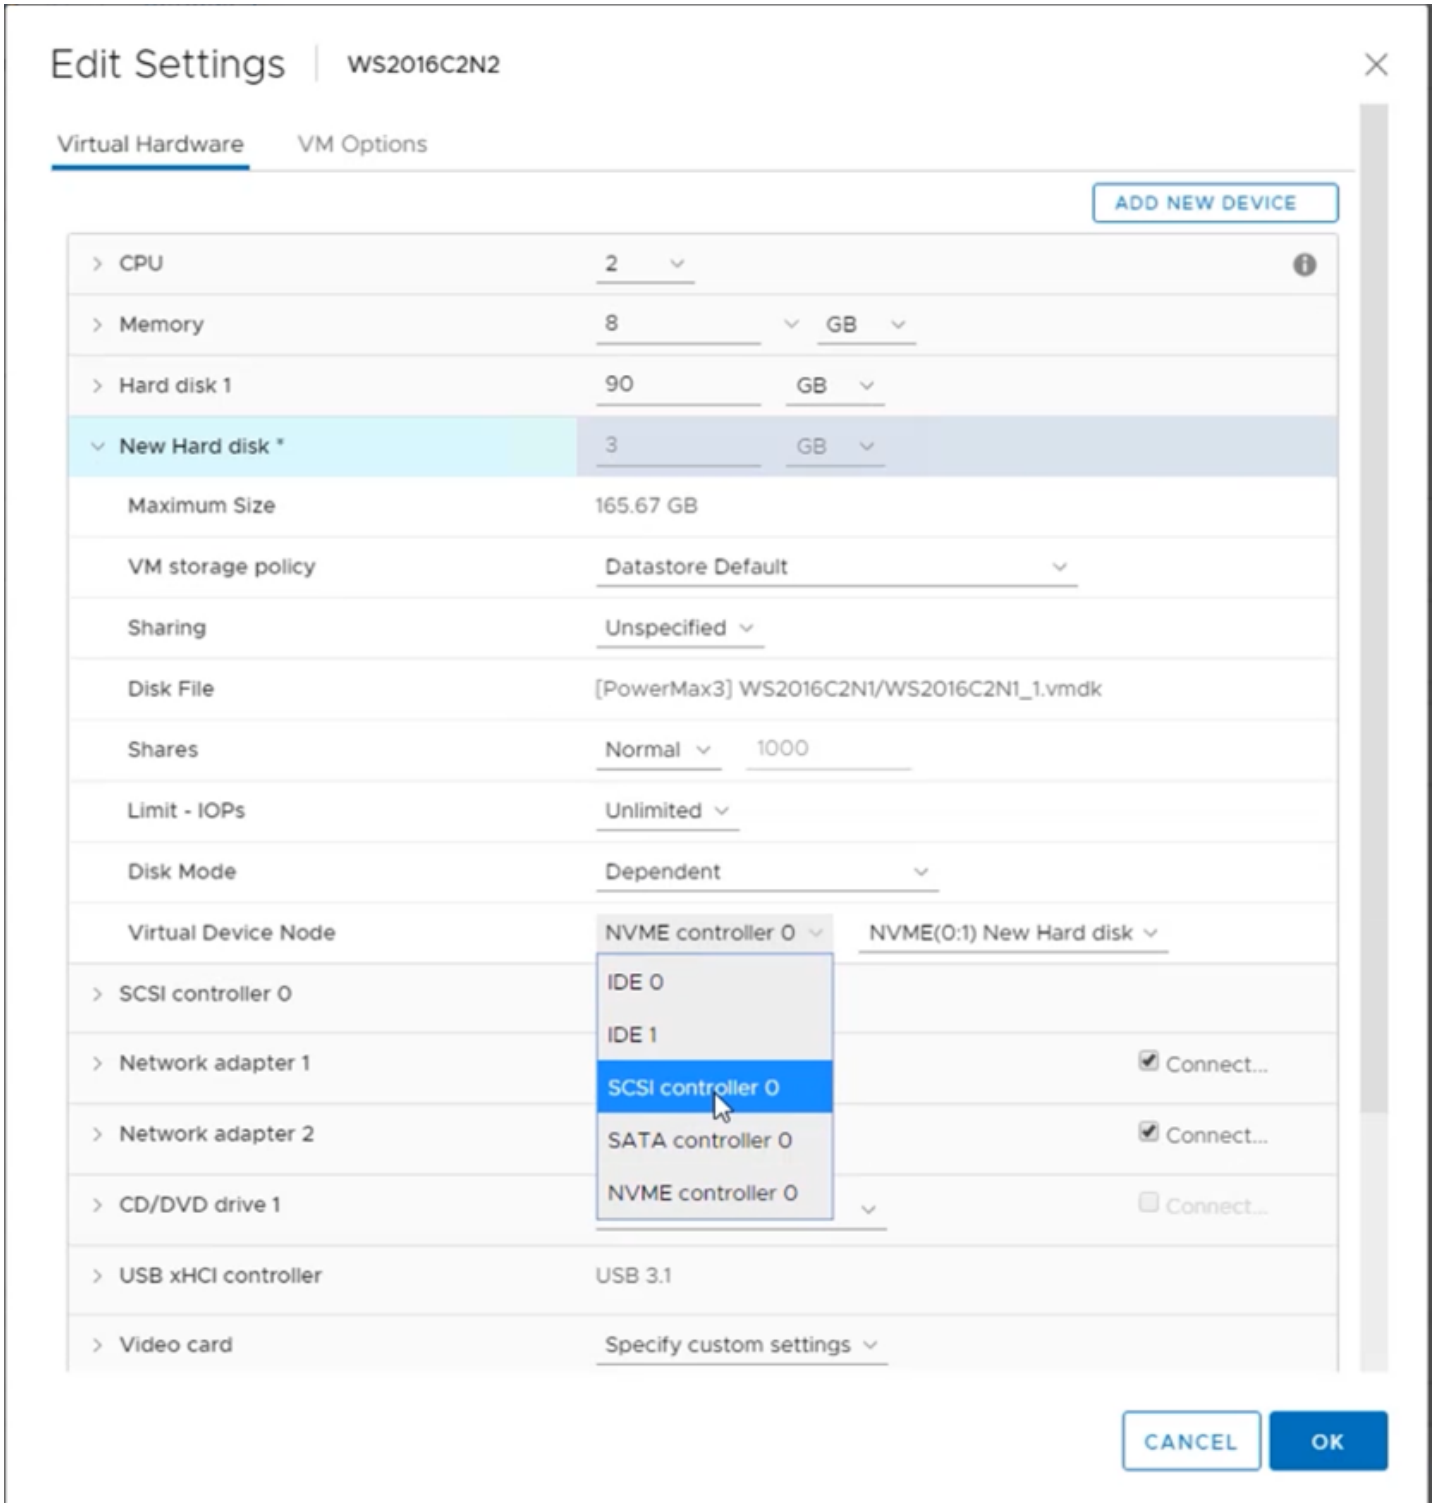

Figure 9

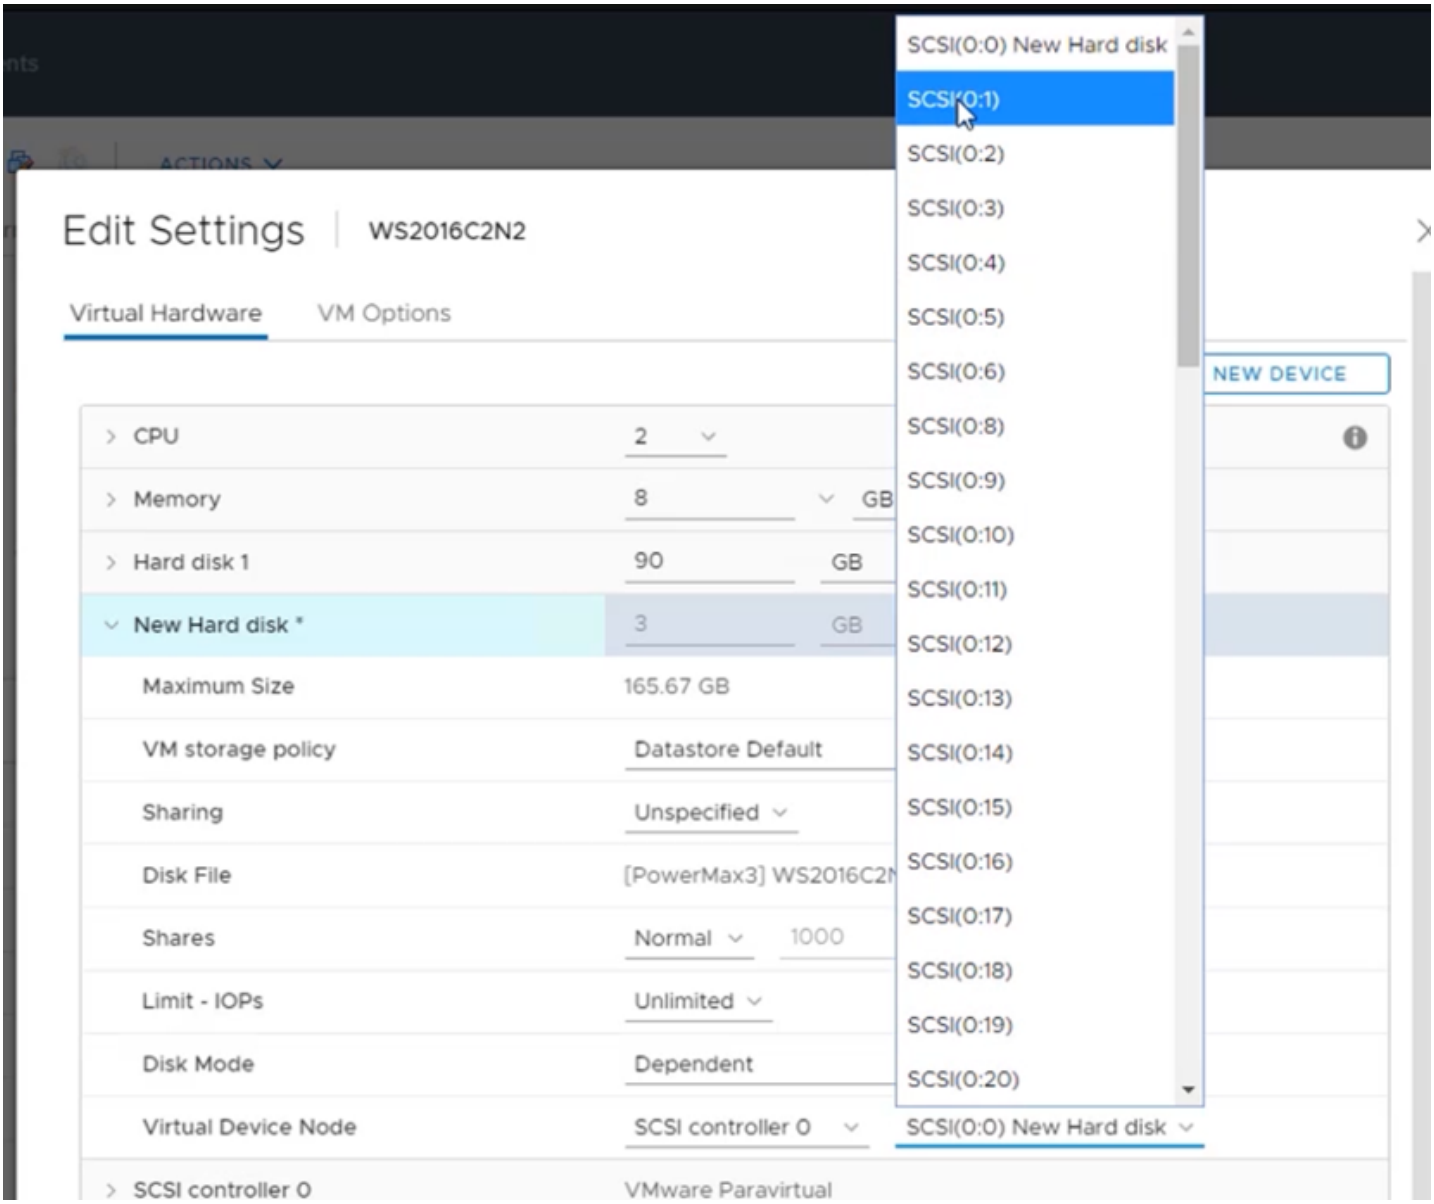

Figure 10

Here, you can see the shared disk is the same path and VMDK as the primary node disk.

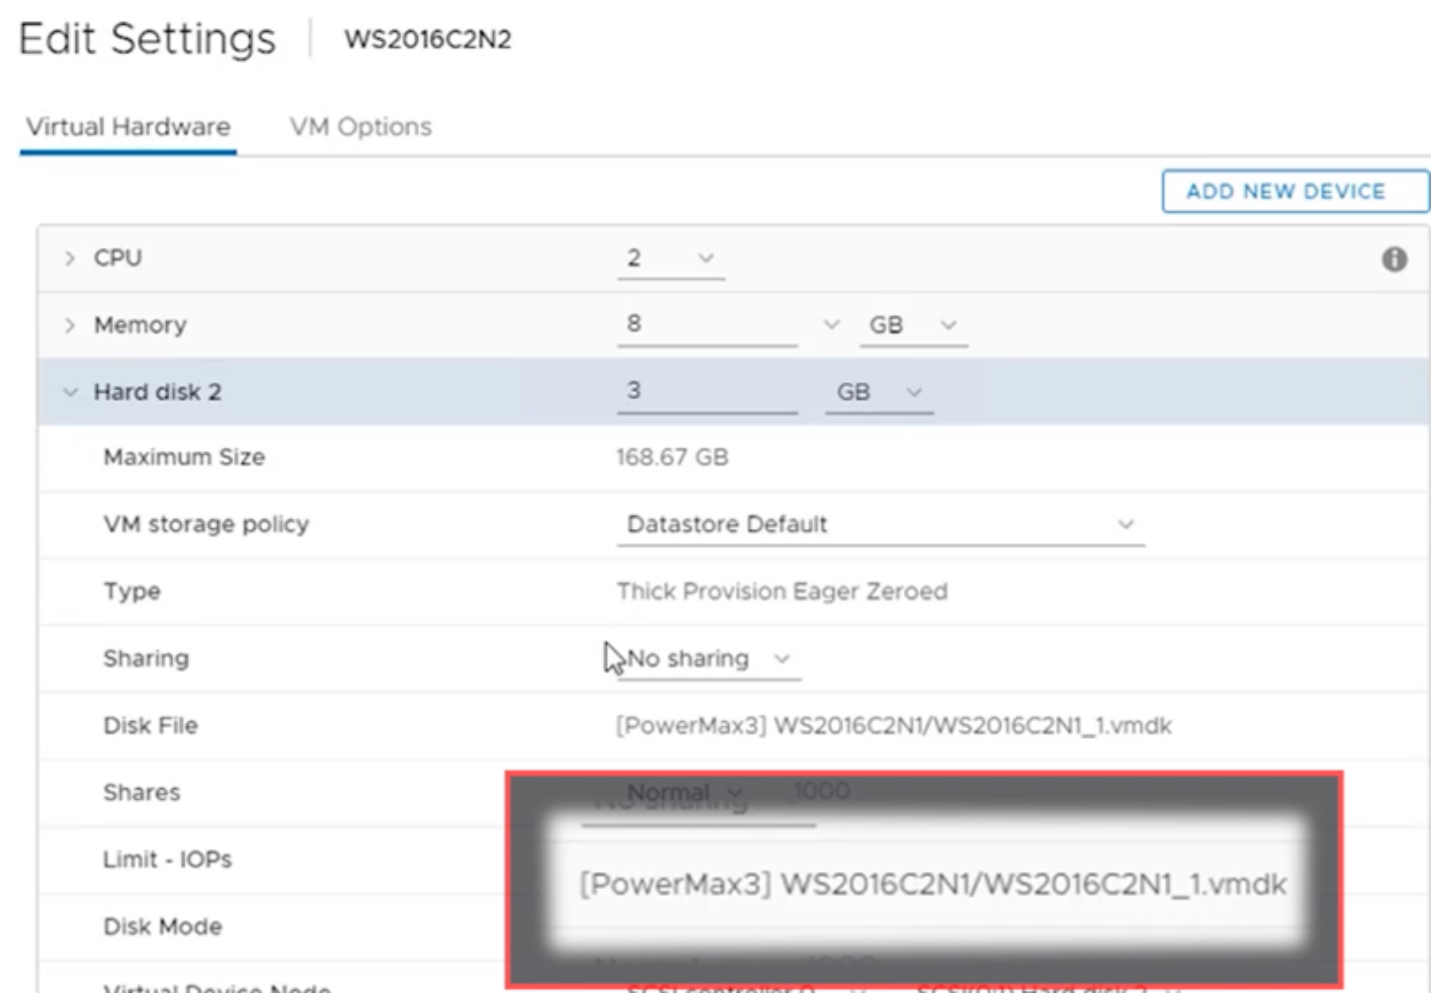

Figure 11

With all the shared disks reattached to all secondary nodes, you may now power on the WSFC cluster, starting with the primary node. With all nodes powered on, validate the WSFC cluster is back online and functioning correctly.

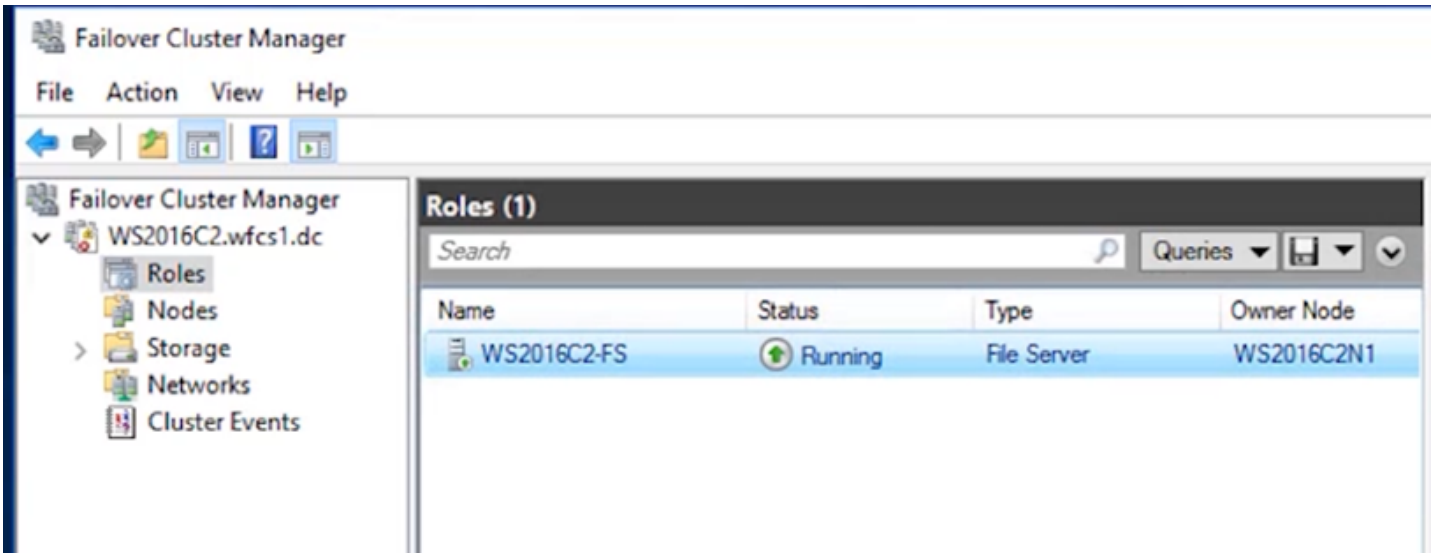

Figure 12

Remember, your RDMs still exist, they are not attached to any VM but have not been deleted. If the migration fails, you can reattach the RDMs in the original configuration as a fallback option.

Video of Migration