



Storage Block Alignment with VMware Virtual Infrastructure

Correctly Aligning LUNs and .vmdk Files to Network Appliance™ Storage

Tim Coste, Network Appliance, Inc.

July 2007 | TR-3593

Executive Summary

This technical report discusses how to correctly align LUNs and VMware .vmdk virtual disks in ESX to Network Appliance storage blocks.

Table of Contents

1. Introduction	3
1.1 LUN Alignment Basics.....	3
1.2 LUN/Virtual Disk Alignment with VMware Infrastructure 3 or ESX 2.5	4
1.3 Correct LUN Alignment	4
1.4 Counters that indicate Improper Alignment.....	5
2. LUN and .vmdk Alignment Steps	6
2.1 Aligning Virtual Disks (.vmdks) to a Network Appliance VMFS Datastore	6
2.2 Aligning Virtual Disks to a Network Appliance NFS Datastore	9
2.3 Aligning a Raw Device Mapping (RDM) to a Network Appliance Storage System.....	10

1. Introduction

In order to get the best performance from a storage system, the storage blocks on the LUN must be aligned to the file system on the operating system. This is typically done by setting the OS type when the LUN is created on the storage system. This document discusses LUN block alignment and the special considerations needed in a virtualized environment.

1.1 LUN Alignment Basics

There are many cases where, by default, a file system block will not be aligned to the storage array. In Figure 1 below, you can see that the blocks do not align. This would mean that for each random read or write, two blocks would need to be read or written. This can negatively impact the performance of the storage array. Sequential writes will also be affected, although to a lesser extent.



Figure 1) Blocks not aligned. The blocks of this storage system are not aligned to the operating system's file system.

In a nonvirtualized environment, the block alignment is done by selecting the appropriate LUN protocol type when the LUN is created. This will align the storage to the file system, as seen in Figure 2.

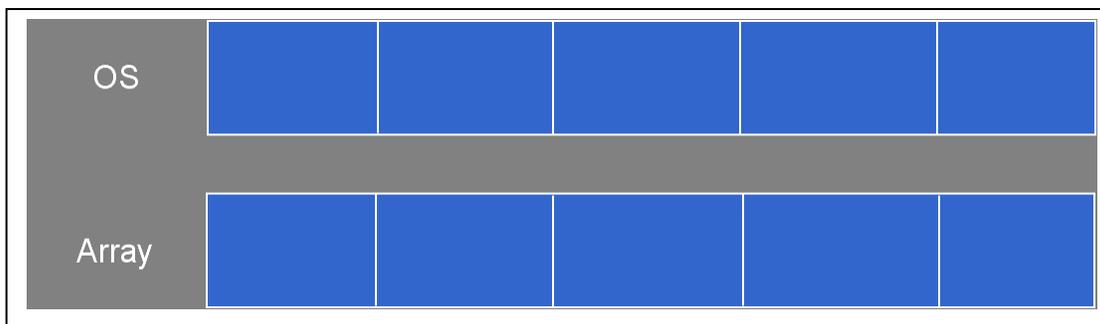


Figure 2) Blocks aligned. The blocks of this storage system are correctly aligned to the operating system's file system.

1.2 LUN/Virtual Disk Alignment with VMware Infrastructure 3 or ESX 2.5

Virtualization products such as VMware add another layer of complexity to alignment, as seen in Figure 3. Not only must the VMFS Datastore be correctly aligned to the storage blocks, but the guest OS file system must also be aligned.



Figure 3) Virtualized blocks not aligned. The blocks of this storage system are not correctly aligned to the virtualized file system or the guest operating system.

1.3 Correct LUN Alignment

Selecting the correct LUN protocol type will align the VMFS file system to the storage array blocks, but the guest OS still needs to be aligned (see Figure 4). This process is discussed in Chapters 2 and 3.

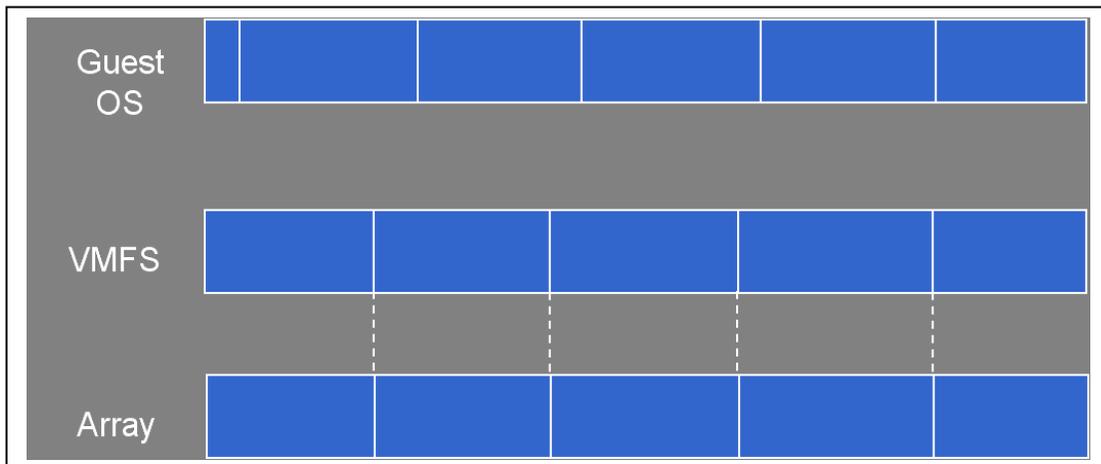


Figure 4) Virtualized host correctly aligned to storage blocks. The blocks of this storage system are correctly aligned to the virtualized file system, but the guest operating system is still unaligned.

1.4 Counters that indicate Improper Alignment

There are various ways of determining if you do not have proper alignment. Using perfstat counters, under the wafld_susp section, "wp.partial_writes", "pw.over_limit", and "pw.async_read" are indicators of improper alignment. The "wp.partial write" is the block counter of unaligned I/O. If more than a small number of partial writes happen, then WAFL® will launch a background read. These are counted in "pw.async_read"; "pw.over_limit" is the block counter of the writes waiting on disk reads.

Using Data ONTAP® 7.2.1 or newer, there are some per LUN counters to track I/O alignment:

- lun:read_align_histo: 8bin histogram for reads that tracks how many 512b sectors off the beginning of a WAFL block an I/O was. Reported as a % of reads.
- lun:write_align_histo: Same for writes.
- lun:read_partial_blocks: % reads that are not a multiple of 4k.
- lun:write_partial_blocks: same for writes.

If any of these except for read/write_align_histo[0] is nonzero, you had some misaligned I/O.

2. LUN and .vmdk Alignment Steps

2.1 Aligning Virtual Disks (.vmdks) to a Network Appliance VMFS Datastore

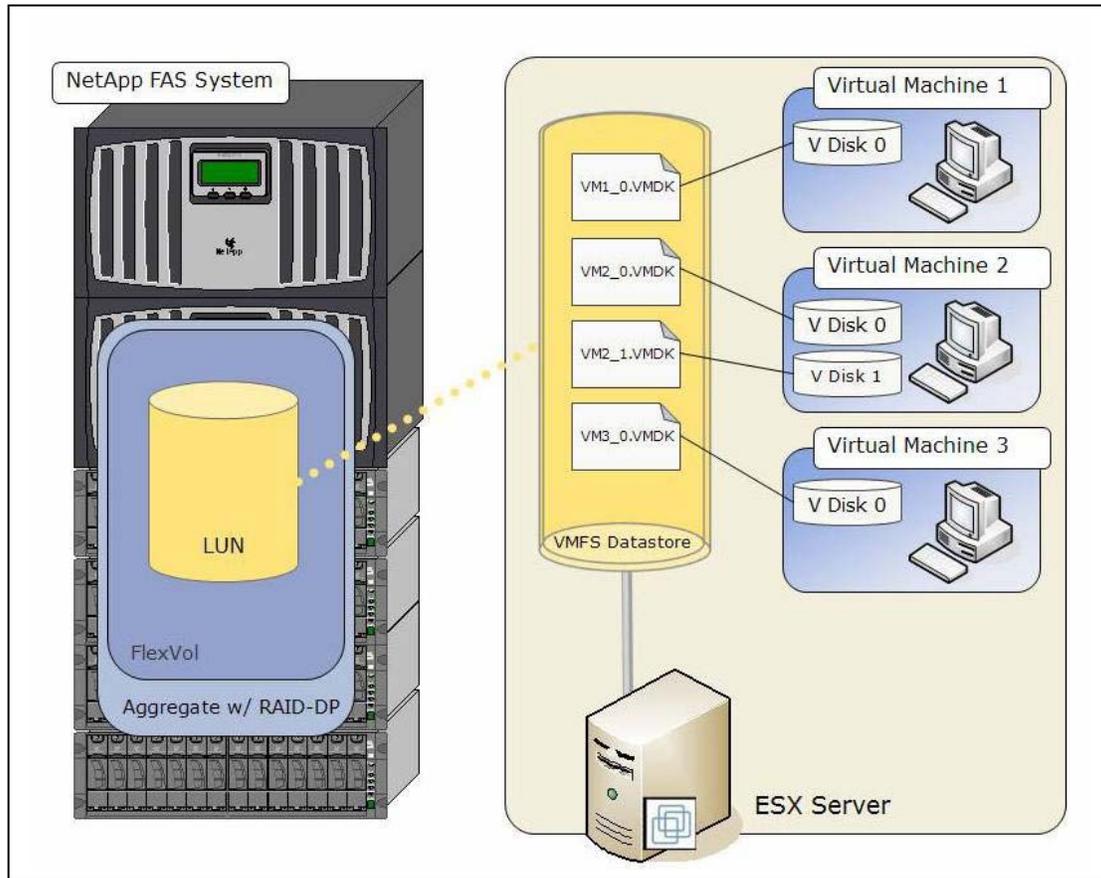


Figure 5) VMFS Datastore on NetApp storage.

In VMware Infrastructure 3 (VI3), the Network Appliance storage system can be aligned to the VMware VMFS file system by selecting VMware from the LUN Protocol Type when the LUN is created. This creates the block alignment as seen in Figure 4 above. To correctly align the Windows® guest operating system to the storage blocks, use the following steps.

1. Create a Virtual Machine using Virtual Infrastructure Client.
2. After the Virtual Machine is created, we need to add a second hard drive. Add a second hard drive using the following commands:
 - a. Right-click the Virtual Machine and select “Edit Settings...”
 - b. Click “Add...”
 - c. Select “Hard Drive” and click “Next>.”
 - d. Select “Create a new virtual disk” and click “Next>.”
 - e. Use the disk capacity you would like your secondary drive to be, then click “Next>.”
 - f. Click “Next>”, then click “Finish,” and then click “OK.”

This creates a second hard drive named “{Virtual_Machine_Name}_1-flat.vmdk” This vmdk may be removed from this Virtual Machine and presented to another virtual machine to perform the next section. Or Windows can be installed on the first drive of this Virtual Machine with the next section performed on the second drive just created.

From within Windows in the Virtual Machine, if you run “System Information” (msinfo32.exe in Start-> Run) and select Components → Storage → Disks, scroll to the bottom and you will see the Partition Starting Offset information. This number needs to be perfectly divisible by 4096. The default .vmdk, you will see the Partition Starting Offset set to 32,256 as seen in Figure 6. $32,256 / 4096 = 7.875$, and thus this file system is not correctly aligned.

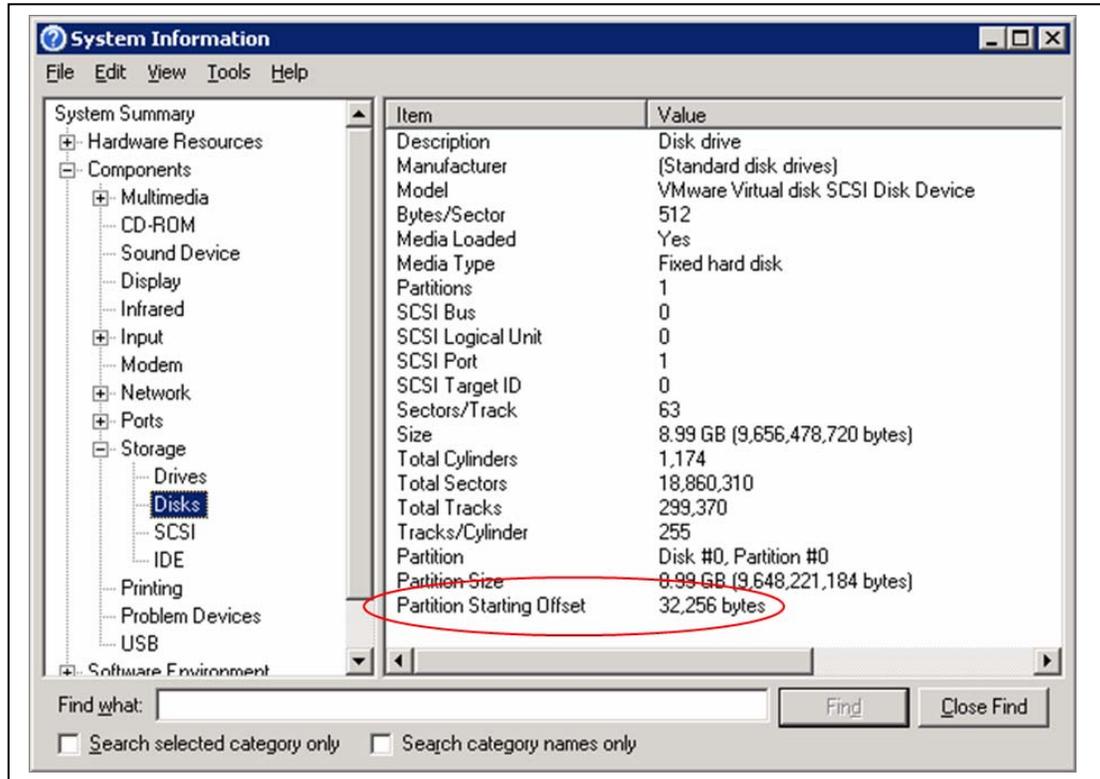


Figure 6) Partition Starting Offset of an unaligned NTFS file system on a .vmdk. The disk is not aligned to the storage blocks as this number needs to be perfectly divisible by 4096.

To correctly set the Partition Starting Offset, use the following steps:

1. Start a Command Prompt.
2. C:\>diskpart
3. DISKPART> list disk
Two disks should be listed.
4. DISKPART> select disk 1
This selects the second disk drive.
5. DISKPART> list partitions
This step should result in a message stating “There are no partitions on this disk to show.” This message confirms the disk is blank.
6. DISKPART> create partition primary align=64

Viewing the second disk using “System Information,” the Partition Starting Offset is now 65,536, as seen in Figure 7. $65,536 / 4096 = 16$, and thus this file system is properly aligned to the Network Appliance storage blocks.

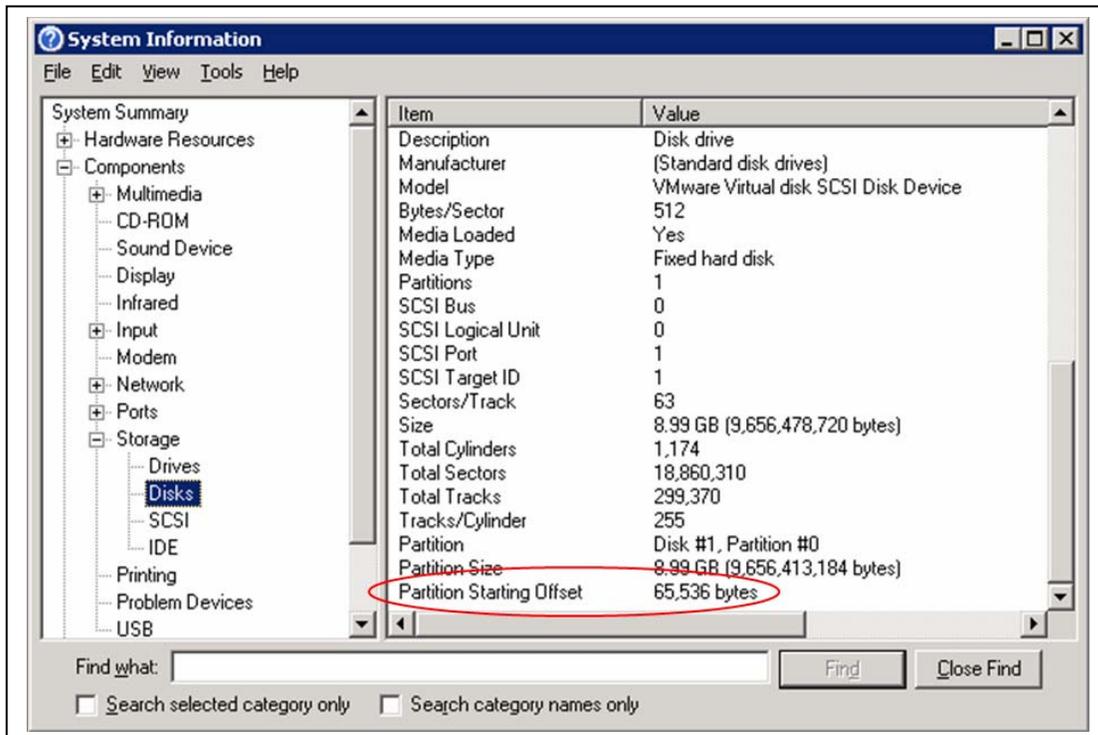


Figure 7) Partition Starting Offset of a correctly aligned NTFS file system on a .vmdk. The disk is correctly aligned to the storage blocks as this number is perfectly divisible by 4096.

This second drive can now be removed from this Virtual Machine and presented to another Virtual Machine as Drive 0 for Windows to be installed. It is suggested that a copy of this .vmdk be saved as a “Gold Copy.” This “Gold Copy” can then be manually copied to a virtual machine directory on VMFS for proper alignment. Another method for this would be to install Windows on this correctly aligned .vmdk, then select the “Clone to Template” or “Convert to Template” option in the Virtual Infrastructure Client. The Virtual Machines deployed from this correctly aligned .vmdk template will also be correctly aligned.

2.2 Aligning Virtual Disks to a Network Appliance NFS Datastore

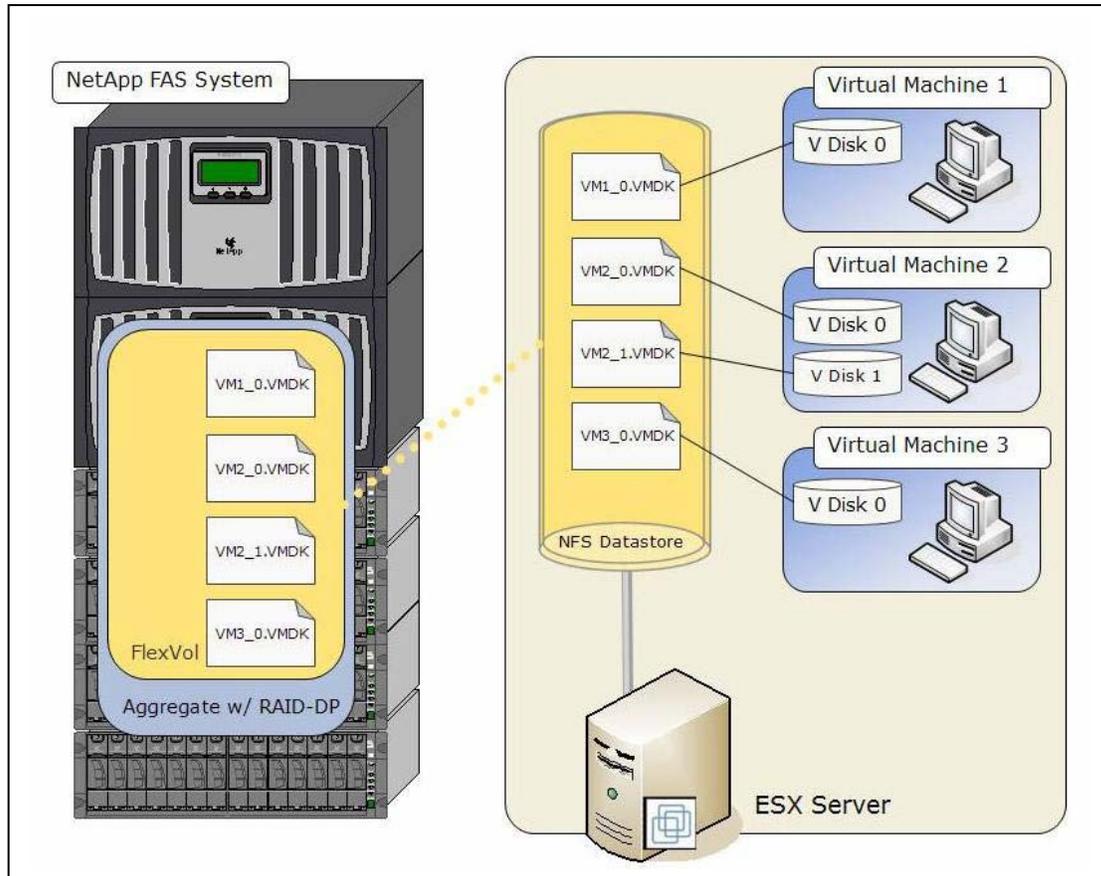


Figure 8) NFS Datastore on NetApp storage.

Since a .vmdk is a virtual block device, a .vmdk still needs to be correctly aligned to an NFS Datastore mounted from a Network Appliance NFS share. To correctly align a .vmdk file on an NFS Datastore, use the steps listed in Section 2.1, Aligning Virtual Disks (.vmdks) to a Network Appliance VMFS Datastore.”

2.3 Aligning a Raw Device Mapping (RDM) to a Network Appliance Storage System

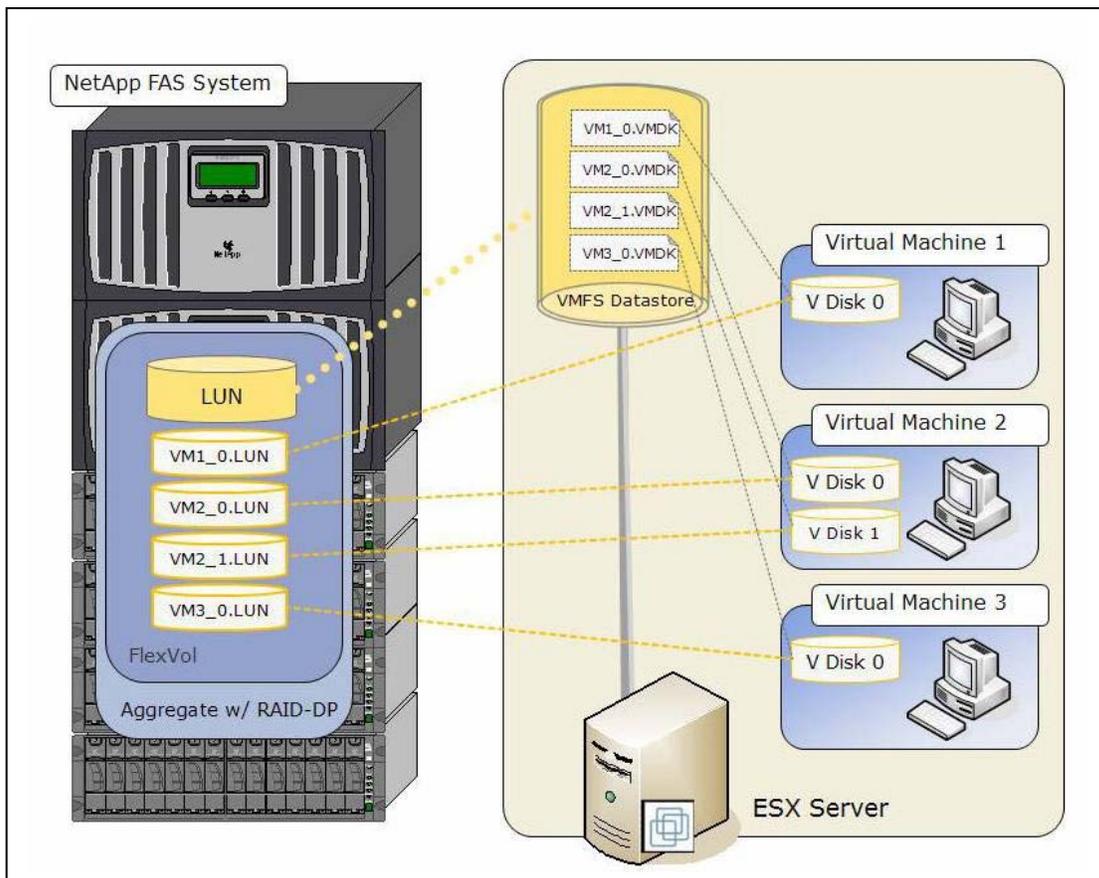


Figure 9) Raw Device Mappings (RDMs) on NetApp storage

When using Raw Device Mappings (RDMs), the file system will be correctly aligned if the protocol type of the NetApp LUN is set to the guest operating system type. If you are using Windows as a guest OS, set the LUN Protocol Type to Windows on the NetApp storage controller, and so on. Do not use the VMware LUN Protocol Type for RDMs. No other steps are necessary for proper alignment.