

- [Home](#)
- [About](#)
- [Resume](#)
- [Virtualization Jobs](#)
- [Virtualization Resources](#)

## VMtoday

VMware News, Views, & How-To's from Josh Townsend

« [Free SAN Monitor for DS3300, MD3000i and others](#)

[EMC Virtual Storage Integrator Update](#) »

## Storage Basics – Part VII: Storage Alignment

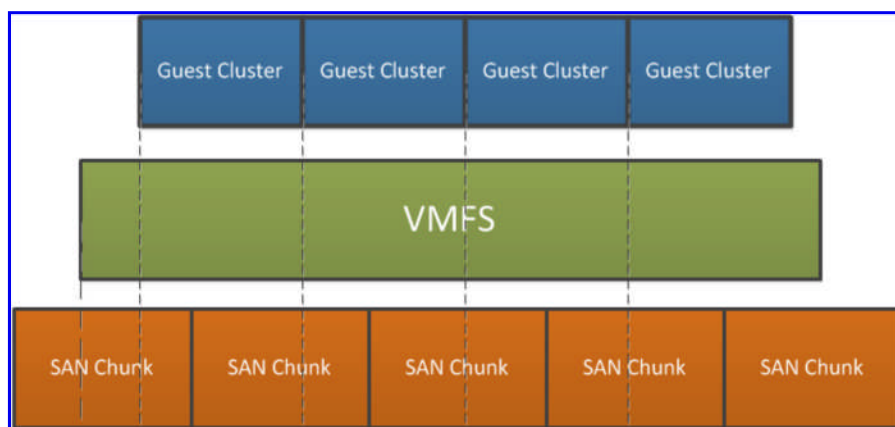
June 23rd, 2010 | Author: [Joshua Townsend](#)

At the risk of beating a dead horse, it's time to resurrect my Storage Basics series. I've recently had some great feedback on the series and figured I should round out a few of the concepts before I wrap it up. I want to cover a topic often discussed amongst virtualization professionals, but one I often find general practitioners and server admins not understanding: storage alignment. Storage alignment, or the lack of alignment, is not a new issue and is not unique to VMware or virtualization in general. However, the effects of misaligned storage can be more greatly felt in terms of reduced performance and strain on a storage system in shared, oversubscribed or high I/O environments. Many others in the virtualization and storage communities have already covered partition alignment (see [Duncan Epping](#), [Vaughn Stewart](#), and most recently [Chad Sakac](#)), but I feel it is an important enough topic for me to re-hash as part of this series.

[6](#)  
[tweets](#)  
[retwee](#)

### What is Storage Alignment?


Let's start with a quick overview of what storage alignment means. Quite simply, storage alignment refers to the positioning (starting offset) of the various pieces of a systems storage components – the physical disk sectors or array's chunks, the VMware File System (VMFS) in a VMware environment, and the guest file system's clusters within a partition – in relation to the layer directly under the element in question. A quick graphic often makes quick work of explaining this (I often whiteboard this concept for colleagues and clients):



As you can see, the starting offset of the VMFS partition does not correspond to the physical segmentation of the underlying disks (in this case, the chunks on a SAN – but could be conceptually replaced with the sectors of a single disk). Furthermore, the clusters (or blocks) of the guest VM are not aligned to the VMFS partition nor to the underlying storage. For traditional (physical) systems or VMware RDM's, the VMFS layer could be abstracted, but the result would be the same – the clusters of a partition would be misaligned to the underlying disk.

### What Does it Mean?

Quite simply, misaligned storage (both VMFS partitions and Guest File Systems) can lead to poor performance under certain conditions. How badly performance is impacted depends on the degree of I/O strain your server and storage are under, the caching mechanisms in your environment, and the architecture of your SAN. Again, a visual can help explain how misaligned storage can hurt you. For simplicity let's leave out the VMFS layer as we consider the following diagram (pardon my hasty Visio visualization):



What we see is that the target data in a tiny 16kb read request spans two 64kb chunks on our storage array. Any reads of that piece of data will result in twice the amount of data as would be minimally necessary being transferred to the host's storage stack. The net effect is an increase in the work the storage array must do – gobbling up **IOPS** that would otherwise be available for the real work of reading data, reducing throughput on the **interface**, and messing with **cache** algorithms and dedupe mechanisms on some arrays. In short, misaligned storage is an efficiency killer. Now add in the VMFS layer back in and you'll see how things get complicated.

If (and we're talking a big IF here) every bit of data you wanted to read spanned a chunk or sector boundary, you could experience half the expected performance due to misalignment. In reality, depending on your **workload** and storage technology your performance increase from properly aligning your storage will probably be somewhere between 10-30%.

## Want to dig deeper?

There have been some great resources published on this issue over the past few years on storage alignment. Major vendors have all begun pushing information on the problem – here are some of the best that I have found:

**Microsoft** has a Knowledge Base article (<http://support.microsoft.com/kb/929491>) that describes the problem and symptoms of misaligned partitions, how to determine if your partition is aligned, and the use of diskpart to create aligned partitions.

**Microsoft** also has an in-depth article on MSDN, including some performance numbers at <http://msdn.microsoft.com/en-us/library/dd758814.aspx>. Also check out Jimmy May's series Partition (Sector) alignment for SQL Server here: <http://blogs.msdn.com/b/jimmymay/archive/2008/10/14/disk-partition-alignment-for-sql-server-slide-deck.aspx>. One of the best descriptions of the complexities of the problem can be found in Jimmy's blog series.

**VMware** has an article here: [http://www.vmware.com/pdf/esx3\\_partition\\_align.pdf](http://www.vmware.com/pdf/esx3_partition_align.pdf). Be aware that this article is for Virtual Infrastructure 3, not vSphere 4.0. Some of the information is now a bit dated.

**Netapp** has a few documents to check out: <http://media.netapp.com/documents/tr-3428.pdf> (VI3), and <http://media.netapp.com/documents/tr-3749.pdf> (vSphere)

**EMC** covers alignment in their TechBooks for [Clariion](#), [Celerra](#), and [Symmetrix](#).

## Tools to Align Partitions:

Ok – so you've bought into this whole partition alignment thing as being a real issue. How to you fix it? Here are some tools:

1. MSInfo32.exe, wmic, and dmideg will show you misaligned partitions on Windows machines (check the Microsoft links above for usage info).
2. [Diskpart.exe](#) (or diskpart.exe on versions of Windows previous to 2003) creates aligned partitions on Windows systems. Diskpart cannot be used to realign a previously created partition, only to create new correctly aligned partitions.
3. [MBRScan/MBRAAlign](#) from NetApp can report on and realign existing virtual disks on a VMware ESX server. Also a nifty PowerShell script from NetApp to find if your partitions are aligned: <http://communities.netapp.com/docs/DOC-6175>
4. [vOptimizer](#) from Vizioncore can report on and realign existing virtual disks.
5. [GParted](#) can be used to create aligned partitions on both Windows and Linux machines, and to realign some existing

partitions.

6. VMware vCenter – VMFS datastores created using vCenter are aligned automatically. Note – Guest VMDK's are not aligned automatically by vCenter – you must manually create aligned partitions on your VMDK's or use a Guest OS that creates properly aligned partitions (Windows 2008 and later).

## Best Practices:

Before I wrap this installment up, here are some best practices for storage alignment in your environment:

- Create aligned partitions in your VMware templates. Do it once, do it right – every machine you deploy from the template will be aligned.
- Use caution with tools like Symantec Ghost. Ghost can take images of aligned partitions and misalign them when laying down on a new system.
- Use caution when performing P2V's using VMware vCenter Converter – it does not align guest disks on import. You might consider using Converter to perform a P2V of the system disk only, then create new VMDK's on the converted guest. Use Diskpart, gparted, or another tool to create aligned partitions on the new VMDK's and finally copy the data over to the newly virtualized server using a tool like Robocopy, [RichCopy](#), or rsync.
- SSD's are particularly sensitive to misalignment, leading to poor performance and excessive wear.
- Local VMFS volumes created by the ESX installer are not aligned. If you are using an installer-created local VMFS for anything where performance matters, you might consider re-creating it through vCenter.
- Watch out when attaching a data disk from an older VM to a new VM. For example, you are upgrading your SQL servers to Windows 2008 R2 from 2003. You decide to do a side-by-side upgrade, using the detach/attach method. You install (or better yet, deploy from template) a new Windows 2008 R2 VM, detach your databases from the old server, move your SQL data and log virtual disks from your 2003 VM to the new VM and attach the SQL DB's on the new server. Those old VMDK's may be misaligned! Consider using Robocopy, RichCopy or rsync to ensure an aligned disk.
- Check your storage vendors best practices for your particular environment (OS, workload, SAN, etc.).
- There is some debate on whether or not it is advised to align your OS partitions. There is no clear-cut answer on this as it depends so much on your environment and particular needs. For help in deciding if you should align your Guest OS drives, see the comments in the blogs by [Duncan Epping](#), [Vaughn Stewart](#), and [Chad Sakac](#).
- While working the [VMware User Group](#) booth at the Washington, DC Virtualization Forum 2010 I had a user ask me if rules and procedures for alignment on [4k sector disks](#) are different. I forgot to research it until just now, so I honestly don't know (please comment if you do know!). Check with your storage vendor if this is an issue for you.
- Finally, you can't realign partitions using tools like mbralign or vOptimizer in ESXi -Aaron Delp explains the problem here: <http://blog.aarondelp.com/2010/06/my-1-issue-with-vmware-esxi-today.html>.

I hope this is helpful for you in understanding the problem of storage alignment and how it can impact your environment. Comments or questions are welcomed!

### Related posts:

1. [Storage Basics – Part VI: Storage Workload Characterization](#)
2. [Storage Basics – Part V: Controllers, Cache and Coalescing](#)
3. [Storage Basics – Part III: RAID](#)
4. [Virtualization Bookmarks for August 28th](#)
5. [Storage Basics – Part II: IOPS](#)

Posted in [Storage Basics](#), [VMware](#) | Tags: [alignment](#), [disk](#), [IOPS](#), [ntfs](#), [performance](#), [SAN](#), [sql](#), [Storage](#), [throughput](#), [VMDK](#), [VMware](#)

 [6 Tweets](#)

## 7 Responses to “Storage Basics – Part VII: Storage Alignment”



- [joshuatownsend:](#)  
[June 23, 2010 at 8:32 am](#)

New VMtoday.com post: Storage Basics – Part VII: Storage Alignment <http://bit.ly/daRDxY> #vmware

*This comment was originally posted on [Twitter](#)*

[Reply](#) [More from author](#)

- [Tweets that mention Storage Basics – Part VII: Storage Alignment / VMtoday -- Topsy.com:](#)  
[June 23, 2010 at 8:48 am](#)

[...] This post was mentioned on Twitter by Todd Scalzott, Josh Townsend. Josh Townsend said: New VMtoday.com post: Storage Basics – Part VII: Storage Alignment <http://bit.ly/daRDxY> #vmware [...]

[Reply](#) [More from author](#)



- [PlanetV12n:](#)  
[June 23, 2010 at 8:57 am](#)

Storage Basics – Part VII: Storage Alignment (VMtoday.com) <http://bit.ly/ci3gvA>

*This comment was originally posted on [Twitter](#)*

[Reply](#) [More from author](#)



- [Harley Stagner:](#)  
[June 23, 2010 at 9:56 am](#)

From the esx3\_partition\_alignment pdf on page 3:

“Note: Aligning the boot disk in the virtual machine is neither recommended nor required. Align only the data disks in the virtual machine.”

If they go through the trouble of stating it in the document, it should probably not be done for the OS drives.

[Reply](#) [More from author](#)



- [Joshua Townsend:](#)  
[June 23, 2010 at 10:03 am](#)

Harley – I agree for most use cases. The general consensus seems to be that it is not worth the time and effort to correct existing VM boot disks, but any new VM's should be properly aligned.

I would recommend aligning Windows XP VMware View VM's boot disks (done automatically if using Vista or Windows 7), and any server VM's boot disk if (for whatever reason) the boot disk is used for heavy I/O.

Thanks for the comment!

Josh

[Reply](#) [More from author](#)



- [vmware\\_in\\_th:](#)  
[June 29, 2010 at 2:17 am](#)

Storage Basics – Part VII: Storage Alignment <http://is.gd/d8z63>

*This comment was originally posted on [Twitter](#)*

[Reply](#) [More from author](#)

- [PlateSpin Migrate 8 review « UP2V:](#)  
[June 30, 2010 at 10:08 am](#)

[...] machine demand a lot of disk resources. For more information on the importance of alignment <http://vmtoday.com/2010/06/storage-basics-part-vii-storage-alignment/> VMware Converter will not align your virtual disk files. PlateSpin Migrate will align the virtual [...]

[Reply](#) [More from author](#)

## Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Submit Comment

Additional comments powered by [BackType](#)

Search

### Recent Posts

- [VMware vExpert](#)
- [EMC Virtual Storage Integrator Update](#)
- [Storage Basics – Part VII: Storage Alignment](#)
- [Free SAN Monitor for DS3300, MD3000i and others](#)
- [Washington DC VMware User Group – April 29th](#)
- [Storage Basics – Part VI: Storage Workload Characterization](#)
- [Update: SVGA Drivers on Windows 2008 R2 and Windows 7](#)

### Tags

[3.5](#) [2009](#) [best practices](#) [client](#) [compatibility](#) [disk](#) [DL380](#) [ecosystem](#) [EMC](#) [ESX](#) [esxi](#) [esxstop](#) [fusion](#) [G3](#) [HP](#) [Hyper-V](#) [I/O](#) [IOPS](#) [memory](#) [Microsoft](#) [NetApp](#)  
[network](#) [Patch](#) [performance](#) [powershell](#) [reporting](#) [SAN](#) [security](#) [sizing](#) [sql](#) [statistics](#) [Storage](#) [Update 3](#) [vcenter](#) [VI3](#) [virtual](#) [virtual center](#)  
[virtualization](#) [Virtual Machine](#) [VM](#) [VMUG](#) [VMware](#) [vmworld](#) [vsphere](#) [windows](#)

### Categories

- [Documentation](#) (2)
- [EMC](#) (2)
- [General IT](#) (10)
- [Issues & Troubleshooting](#) (18)
- [Microsoft](#) (8)
- [NetApp](#) (3)
- [News](#) (6)
- [Personal](#) (3)

- [Product Releases](#) (2)
- [Storage](#) (8)
- [Storage Basics](#) (3)
- [Uncategorized](#) (3)
- [VMUG](#) (2)
- [VMware](#) (46)
- [VMware How To](#) (18)
- [VMworld](#) (3)

#### [Recent Tweets](#)

- No public Twitter messages.

#### [Recent Bookmarks](#)

- [Windows Guest Defragmentation, Take Two « Pivot Point](#)
- [Dan Weiss Blogsite » NetApp performance testing engineer tells it like it is](#)
- [Understanding how storage design has a big impact on your VDI! - Ruben Spruijt - BrianMadden.com](#)
- [Performance Troubleshooting VMware vSphere – Storage | virtual insanity](#)
- [PowerCLI – Where do I start ?! « Virtu-AI](#)
- [CPU/MEM Reservation Behavior » Yellow Bricks](#)
- [Refreshing Test Dev SQL Environments using NetApp Technologies · Rajeev Karamchedu](#)
- [Finding thin-provisioned virtual disks with PowerShell | VCritical](#)
- [ESXi4 Console - Secret Commands for Emergency Maintenance](#)
- [VMware Server on a Domain Controller \(DC\) « Derek's Musings](#)
- [How to Enable SSH in ESXi – VMware « Code Ghar](#)
- [VMware ESX, EMC CLARiiON Arrays, and Multiple Protocols - blog.scottlowe.org - The weblog of an IT pro specializing in virtualization, storage, and servers](#)
- [PVSCSI and Low IO Workloads « Pivot Point](#)
- [VDI.com - Why should I provision VDI VM's with less RAM and other resources than my physical machines?](#)
- [The old pain in the DAS « StorageTexan's Blog](#)

#### About Me



Hello, and thank you for visiting VMtoday. My name is Josh Townsend. I am a technology professional with a strong background in VMware Virtualization, Storage, and Microsoft technologies. I am a Sr. Systems Administrator at [Tiber Creek Consulting](#) in Fairfax, VA, and hold several technical certifications, including VMware Certified Professional. I am also a 2010 VMware vExpert.



I am also leader of the [Washington DC Metro Area VMware User Group \(VMUG\)](#).



The opinions expressed on this site are my own and may not reflect the views of my employer, VMware, or any other party unless otherwise stated.

Please feel free to follow me on Twitter  
[@joshuatownsend](#)

## Archives

- [July 2010](#) (2)
- [June 2010](#) (1)
- [May 2010](#) (1)
- [April 2010](#) (2)
- [March 2010](#) (3)
- [February 2010](#) (1)
- [January 2010](#) (4)
- [December 2009](#) (3)
- [November 2009](#) (3)
- [September 2009](#) (7)
- [August 2009](#) (2)
- [July 2009](#) (1)
- [June 2009](#) (6)
- [March 2009](#) (4)
- [January 2009](#) (4)
- [December 2008](#) (2)
- [November 2008](#) (2)

## VMware Related Sites

- [boche.net – VMware Virtualization Evangelist](#)
- [Chris Wolf](#)
- [Gabes Virtual World](#)
- [Gerben's Blog On Virtualization](#)
- [IPMer](#)
- [Ken's Virtual Realty](#)
- [Mike Laverick's RTFM Eduction](#)
- [Planet V12n](#)
- [Rene Jorissen's Booches.nl](#)
- [Run-Virtual](#)
- [Scott Lowe's Blog](#)
- [The VMguy](#)
- [vinternals](#)
- [Virtu-AI](#)
- [Virtualization.info](#)
- [VM /ETC](#)
- [VM Guru](#)
- [VMtoday Virtualization Resources](#)
- [VMware](#)
- [VMware Land](#)
- [VMware vLaunchpad](#)
- [vSphere Land \(Eric Siebert\)](#)

## Virtualization Jobs

### **Virtualization Jobs**

[Server Administrator](#)  
[New Albany, IN](#) - Floyd Memorial  
Hospital & Health Services

[DAS - SDCOE CLOUD](#)  
[Infrastructure Architect](#)  
[Louisville, KY](#) - IBM

[DRA - Advisory \(SCON\) IT](#)  
[Architect](#)  
[Louisville, KY](#) - IBM

[DRA - Sr. \(SCON\) IT Architect](#)  
[Louisville, KY](#) - IBM



[DRA - Executive \(SCON\) IT Architect](#)

[Louisville, KY - IBM](#)

[System x FTSS](#)

[Louisville, KY - IBM](#)

[Custom Migration Windows Architect](#)

[Architect](#)

[Louisville, KY - IBM](#)

[Field Technical Sales Specialist](#)

[\(Must Reside In the Boston / Hartford Area\)](#)

[Louisville, KY - IBM](#)

[Midrange Storage Virtualization Solutions Technical Consultant](#)

[Louisville, KY - IBM](#)

[Server Systems Administrator](#)

[Jeffersonville, IN - Heartland Payment Systems](#)

[Payment Systems](#)

[See All Jobs](#)

[Post a job for \\$25](#)

<input type="text" value="keywords"/>	<input type="text" value="location"/>
---------------------------------------	---------------------------------------

[Go](#)

Jobs by [SimplyHired](#) [+ Get Widget](#)

## Virtualization Resources



[Mastering Microsoft Vir...](#)

Tim Cerling, Jeffrey Buller ...

**\$37.79**



[Pro Hyper-V \(Ex...](#)

Harley Stagner (Paperbac...

**\$16.75**



[Microsoft Windows Ser...](#)

Microsoft Software

**\$3,476.99**



[Mastering Virtual Machi...](#)

Michael Michael, Hector Li...

**\$37.79**



[Virtualization For Dum...](#)

Bernard Golden (Paperbac...

**\$19.79**



[Windows Server 2008 H...](#)

John Kelbley, Mike Sterlin...

**\$31.49**



[MCTS: Windows Server ...](#)

William Panek (Paperback ...

**\$40.49**




[Virtualization: From th...](#)

Chris Wolf, Erick M. Halter...

**\$64.90**







Powered by [WordPress](#) and [WordPress Theme](#) created with Artisteer by [Josh Townsend](#).