

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

VMtoday

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

« [Windows Server 2008 R2 & Windows 7 Freeze When Using SVGA Drivers](#)

[Storage Basics – Part II: IOPS](#) »

Storage Basics – Part I: An Introduction

December 28th, 2009 | Author: [Joshua Townsend](#)

I am increasingly finding that both my SMB and Enterprise customers are uneducated on the fundamentals of storage sizing and performance. As a result, storage is often overlooked as a performance bottleneck despite it being a vital component to consider in a virtualization implementation. Storage will only increase in importance as hosts are getting bigger, data volumes increase, and more workloads are virtualized. For some reason, most people can grasp the importance of CPU and memory performance constraints but storage performance is often overlooked and can be hard to explain to business users or executives.

[9](#)
[tweets](#)
[retwee](#)

Case in point – I have recently been called into some environments that were not performing well – these environments happened to be running Microsoft SQL, but could just have well been running any application or collection of virtual machines. Fingers were being pointed in all directions: at applications, at the virtualization layer, at a lack of memory, and DBA's were insisting that there were too few CPU's. The situation was getting political and emotional when I walked into it. A few minutes with Windows Perfmon was all I needed to identify storage performance as the root cause of the firestorm that had been ignited. Using a bit of data, I was able to turn the discussion from an emotional fight to a simple problem of physics and mathematics (and a bit of simple math could have avoided the problem in the first place).

I have seen this play out a few too many times and so decided to write-up this multi-part series on the basics of storage with a focus on storage performance. That said, a little math and physics is where we will start as we look at the basic building block of a storage environment: a hard disk drive. [Wikipedia defines](#) a hard disk drive as “a non-volatile storage device that stores digitally encoded data on rapidly rotating platters with magnetic surfaces.” Your computer, server, or VMware cluster uses hard drives to read and write data. Wikipedia also covers the history and atomic structure of a hard drive pretty well. For our purposes, the take away is that hard drives are physical objects, and as such, follow the laws of physics (duh) in the following measurable ways:

1.) Capacity, which is measured in bits or bytes and exponents thereof (MB, GB, TB, PB). This is how much data will fit on your disk, from simple text files to virtual disks, and everything in between. For example, if you have a 500GB SQL database, you darn well better have a hard drive that has a capacity of at least 500GB. This is a pretty simple concept, so I'll leave it there for now.

2.) Performance, which is measured in a couple ways:

- at the disk itself in Input-Output Per Second (IOPS) – a measure of how many read and write commands a disk can complete in a second

- interface throughput, measured in MBps or Gbps – a measure of the peak rate that a volume of data can be read from or written to disk

- latency – the amount of time between when you ask a disk (or storage system if you want to read ahead) to do something and when it can actually do it, very closely related to IOPS as you'll read in a forthcoming article in this series.

Each disk, array, and storage system has its own fixed set of measurements given a specific configuration. Knowing the physical capabilities of your storage system as measured in the above ways, and your systems storage requirements will go a long way towards a successful design and implementation of your storage environment. The remaining parts of this series will take a look at these performance characteristics a bit more in-depth and explain what happens as you introduce factors like RAID, cache, data reduction techniques such as snapshots and deduplication, and varying workloads.

Please keep in mind that while I have designed and implemented a variety of DAS, NAS, and SAN technologies from a host of vendors including Dell, EMC, IBM, and NetApp, I am by no means a storage expert. The information I will provide is generalized, over-simplified, and does not consider varying approaches from different storage vendors. Nonetheless, I hope you find this useful information if you are designing a solution, troubleshooting a performance issue or preparing to make a storage purchase.

Keep Reading:

[Storage Basics – Part II: IOPS](#)

Related posts:

1. [Storage Basics – Part II: IOPS](#)
2. [Storage Basics – Part V: Controllers, Cache and Coalescing](#)
3. [Storage Basics – Part VI: Storage Workload Characterization](#)
4. [Storage Basics – Part IV: Interface](#)
5. [Storage Basics – Part III: RAID](#)

Posted in [General IT](#), [Storage](#), [VMware](#) | Tags: [performance](#), [Storage](#), [VMware](#)

9 Responses to “Storage Basics – Part I: An Introduction”

- [Tweets that mention Storage Basics – Part I: An Introduction / VMtoday -- Topsy.com:](#)
[December 28, 2009 at 9:04 pm](#)

[...] This post was mentioned on Twitter by tscalzott, joshuatownsend. joshuatownsend said: New VMtoday.com post: Storage Basics – Part I: An Introduction <http://cli.gs/E568B> #vmware [...]

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

- [Twitted by PlanetV12n:](#)
[December 28, 2009 at 9:28 pm](#)

[...] This post was Twitted by PlanetV12n [...]

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

- [Storage Basics – Part II: IOPS / VMtoday:](#)
[December 29, 2009 at 7:51 am](#)

[...] Part I of this series, I discussed the important of storage performance in a virtual environment (really any environment, [...])

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

- [Storage Basics – Part III: RAID / VMtoday:](#)
[January 6, 2010 at 6:24 pm](#)

[...] on storage basics. I’ve had some good feedback from folks in the SMB space saying that the first couple posts in this series have been beneficial, so we’ll be sticking with some basic [...]

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

*John:*[February 11, 2010 at 7:34 pm](#)

Sadly I know of several outsource/consultants around here that would argue that you don't need fast disks or large vol/aggr (depending on vendor) to get good performance. I argued with my now ex-boss about why 20 7200 speed disks were worse then 20 15k or even 10k disks.. and he still said ahh they wont see any issues.. some people are just dumb..

Now had the SAN solution been for file storage I would have said no problem with using 7200K speed disks.. But this was for VMs, which included Exchange, Sql and oracle! Still my boss insisted it wouldn't matter..

[Reply](#)*Joshua Townsend:*[February 11, 2010 at 9:46 pm](#)

John – Amazing what a little data will do! I've had the same arguments, even when vendors come in and use their own tools to gather performance stats. Managers argue that the vendors skew their results to sell more disk. That's why I have worked hard on understanding who, what, where, when and how to capture and analyze performance statistics on my own (and write about it here). It's hard to argue with the raw data....

15k disks are not always the answer, but don't go buying a solution without asking questions and preparing with some good data!

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

- [IOPS « Virtual Persuasion:](#)

[February 12, 2010 at 4:53 pm](#)

[...] <http://vmtoday.com/2009/12/storage-basics-part-i-intro/> Categories: Uncategorized Comments (0) Trackbacks (0)
 Leave a comment Trackback [...]

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

- [Storage Basics – Part IV: Interface | VMtoday:](#)

[March 23, 2010 at 1:25 pm](#)

[...] parts I, II, and III of the Storage Basics series we looked at the basic building blocks of modern storage [...]

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

- [cmdln.org \(a sysadmin blog\) » Blog Archive » Analyzing I/O performance in Linux:](#)

[May 17, 2010 at 10:41 am](#)

[...] <http://vmtoday.com/2009/12/storage-basics-part-i-intro/> [...]

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

Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Additional comments powered by [BackType](#)

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-Al](#)
- [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
[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

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


keywords location

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


Virtualization Resources




[Windows Server 2008 H...](#)
Robert Larson, Janique Ca...
\$37.79
★★★★★




[Windows Server 2008 H...](#)
John Kelbley, Mike Sterlin...
\$31.49
★★★★★




[Windows Server 2008 H...](#)
Rand Morimoto, Jeff Guillet...
\$32.84
★★★★☆




[Win Svr Std with o Hyp...](#)
Microsoft Software
\$1,169.99




[POWERCHUTE NETWORK...](#)
APC
\$86.99



[MCTS: Windows Server ...](#)
William Panek (Paperback ...
\$40.49
★★★★★



[Microsoft Virtualization...](#)
Jason Kappel, Anthony Vel...
\$42.74



[virtualization.info](#)
Alessandro Perilli
\$0.99

1 2 3 4 >

 [Privacy](#)

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