

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

VMtoday

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

« [Storage Basics – Part I: An Introduction](#)
[Vote Now for Top Virtualization Bloggers](#) »

Storage Basics – Part II: IOPS

December 29th, 2009 | Author: [Joshua Townsend](#)

In [Part I of this series](#), I discussed the important of storage performance in a virtual environment (really any environment, virtual or not, where you want acceptable performance), and introduced some of the basic measures of a storage environment. In Part II, we will look more closely at what may be the most important storage design consideration in a VMware server-consolidation environments, many SQL environments, and VDI environments to name a few: IOPS.

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

If we stick with a single-disk-centric approach as we did in Part I, IOPS is quite simply a measure of how many read and write commands a disk can complete in a second. IOPS is an important measure of performance in a shared storage environment (such as VMware) and in high-transaction-rate workloads like SQL. Because hard drives are forced to abide by the laws of physics, the IOPS capabilities of a disk are consistent and predictable given a specific configuration. The formula for calculating IOPS for a given disk is pretty straight forward (please show your work):

$$\text{IOPS} = 1000 / (\text{Seek Latency} + \text{Rotational Latency})$$

Exact latencies vary by disk type, quality, number of platters, etc. You can look up the tech specs for most drives on the market. As an example, I have randomly chosen the [technical specifications of the Seagate Cheatah 15k.7](#) SAS drive. This particular drive has the following performance characteristics:

- Average (rotational) latency: 2.0msec
- Average read seek (latency): 3.4msec
- Average write seek (latency): 3.9msec

Using the read latency number, the math works out like this:

$$\frac{1000}{2.0 + 3.4} = 185 \text{ maximum read IOPS}$$

The maximum write IOPS will be a bit less (~169IOPS) because of the higher write seek latency. Writing is more 'expensive' than reading and therefore slower.

Fortunately, there are some widely accepted 'working' numbers, so you do not have to use this formula for each and every disk you might consider using. Because rotational latency is based on the rotational speed, we can use the published Rotations Per Minute (RPM) rating of the drive to guess-timate the IOPS capabilities. Typical spindle speeds (measured in RPM) and their equivalent IOPS are in the table below.

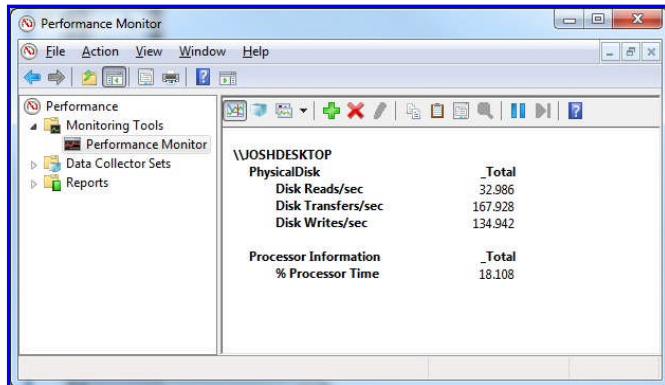
RPM.....IOPS

7,200 80

| | |
|--------|-------------|
| 10,000 | 130 |
| 15,000 | 180 |
| SSD | 2500 – 6000 |

While not a traditional spinning disk, I have also included Solid State Disks (SSD's) for reference as SSD's are starting to see increased market adoption. I have seen a wide range of sizing IOPS for SSD depending on the technology, type (SLC, MLC, etc.) Check out http://en.wikipedia.org/wiki/Solid-state_drive for an introduction, and ask your vendors for more in-depth technical information.

If you are brand-new to this (and you are still reading, congrats!), you can see how many IOPS your Windows computer is asking for by opening Performance Monitor and looking at the 'Disk Transfers/sec' counter under Physical Disk. This is a sum of the 'Disk Reads/sec' and 'Disk Writes/sec' counters as you can see in the screenshot below:



If you are after some stats for your VMware ESX environment, check out esxtop and looking for CMDs/s in the output. I published a couple articles on using esxtop [here](#) and [here](#). The numbers from PerfMon and esxtop get you pretty close but can be skewed by a few things we'll discuss in later posts.

Now that was fun and all, but let's get real: Single-disk configurations are uncommon in servers. As such, we'll part ways with our Simple Jack single disk approach to storage and begin to look at more real-world multi-disk enterprise-class storage configurations. A discussion of IOPS in a multi-disk array is a great way to start. From a very elementary perspective, you can combine multiple hard drives together to aggregate their performance capabilities. For example, two 15k RPM disks working together to server a workload could provide a theoretical 360 IOPS (180 + 180). This also scales out so ten 15k RPM disks could provide 1800 IOPS, and 100 15k RPM disks could provide 18,000 IOPS.

Designing your environment so that your storage can deliver sufficient IOPS to the requesting workload is of utmost importance. If you are working on a storage design, arm yourself with data from perfmon, top, iostat, esxtop, and vsesiStats. I typically gather at least 24 hours of performance data from systems under normal conditions (a few days to a week may be good if you have varying business cycles) and take the 95th percentile as a starting point. So from a very simple approach, if your data and calculations show a 1800 IOPS demand at the 95th percentile, you ought to have at least ten 15k RPM disks (or twenty-three 7.2k RPM SATA disks) to achieve performance goals. It's amazing how some simple data and a pretty little Excel spreadsheet can help you understand and justify the right hardware for the job.

Now before you go and start filling out that PO form for a nice new storage system based on these numbers there are a few more things we ought to discuss. RAID, cache, and advanced storage technologies will skew these numbers and need to be understood. Stay tuned to future articles in this series for more on those topics and more.

Finally, there has been a bunch of activity in the VMware ecosystem of vendors, bloggers, and twittering-type-folks around storage performance. As this here post sat in my drafts folder, Duncan Epping posted this gem of an article that pretty much included all of the content of this article, as well as future ones in my series: <http://www.yellow-bricks.com/2009/12/23/iops/>.

Do yourself a favor and read his post and the comments from his readers – both are filled with a ton of great information, including some vendor-specific implementations.

I was led to Duncan's article by a post by Chad Sakac on his blog: http://virtualgeek.typepad.com/virtual_geek/2009/12/whats-what-in-vmware-view-and-vdi-land.html. This is also a great read that covers some of the same information with a focus on VMware View/VDI and is also worth a few minutes of your time. Also check out <http://vpivot.com/2009/09/18/storage-is-the-problem/> for a rubber-meets-the-road post from Scott Drummonds on the importance of storage

performance vis-a-vis IOPS in a VMware-virtualized SQL environment.

Related posts:

1. [Storage Basics – Part VI: Storage Workload Characterization](#)
2. [Storage Basics – Part I: An Introduction](#)
3. [Storage Basics – Part V: Controllers, Cache and Coalescing](#)
4. [Storage Basics – Part III: RAID](#)
5. [Storage Basics – Part VII: Storage Alignment](#)

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

9 Responses to “Storage Basics – Part II: IOPS”

- [Tweets that mention Storage Basics – Part II: IOPS / VMtoday -- Topsy.com:](#)
[December 29, 2009 at 8:09 am](#)

[...] This post was mentioned on Twitter by VMware Planet V12n, joshuatownsend. joshuatownsend said: New VMtoday.com post: Storage Basics – Part II: IOPS <http://cli.gs/HgJ4A> #vmware [...]

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



- [Craig:](#)
[December 29, 2009 at 12:00 pm](#)

You miss out something regarding the block size and behavior which will significant impact the IOPS from time to time. Just my 2 cents

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



- [Joshua Townsend:](#)
[December 29, 2009 at 12:03 pm](#)

Craig – good observation. Block size/workload characterization comes in the next installment as I move from the basics into real-world scenarios. Thanks for reading! -Josh

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

- [Storage Basics – Part I: An Introduction / VMtoday:](#)
[December 29, 2009 at 12:07 pm](#)

[...] Storage Basics – Part II: IOPS Share and Enjoy: [...]

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



- [Andy:](#)
[January 4, 2010 at 2:03 pm](#)

Josh – As a SMB VMware customer, this series really helped fill in some gaps in my understanding of storage.

The Yellow Bricks article was helpful too. Please keep up the great work!

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

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

[...] nitty-gritty details and practical applications of these concepts in a VMware environment. In the second post of this series I introduced the concept of IOPS and explained how the physical characteristics of [...]

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

- [Storage Basics – Part IV: Interface / VMtoday:](#)
[March 23, 2010 at 1:23 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](#)

- [Storage Basics – Part VI: Storage Workload Characterization / VMtoday:](#)
[April 8, 2010 at 11:35 am](#)

[...] covered in Storage Basics Parts 1 through 5 was at a very elementary level. The math I used to do IOPS calculations, for example, is only true under very certain conditions. RAID controllers implement caching and [...]

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

- [Why Does Cloning A VM From Template Take A Long Time? / VM/ETC:](#)
[April 9, 2010 at 6:36 am](#)

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

[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 esxtop 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](#)

- [@Virtustream360](#) Thanks for the follow - nice to see you guys on Twitter!
- [@DellServerGeek](#) YUM!
- RT [@ericsiebert](#): @TrainSignal's new Pro Series video - my experience - <http://bit.ly/brlE35> <- Nice - I look forward to watching yours!
- [@stu](#) thanks for including me in your VMworld 2010 list! - <http://tweepml.org/VMworld-2010-San-Francisco/>
- RT [@jasonboche](#): North Korea is an idiot. <- Ha!
- RT [@cshanklin](#): I hope that if, one day, I accidentally post something to twitter that it's totally awesome.
- [@tscalzott](#) Quality.
- [@tscalzott](#) That was some impressive rain today!
- Please, sir. May I have another?
- RT [@veeam](#): Veeam FastSCP 3.0.3 with vSphere 4.1 support is available for download (and still FREE): <http://bit.ly/aaV3gc> <- time 2 thinapp!

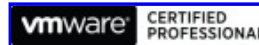
[Recent Bookmarks](#)

- [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 Education](#)
- [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

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


Virtualization Resources

 Windows Server 2008 H...
Rand Morimoto, Jeff Guillet...
\$32.84
★★★★☆


 Win Svr Std with o Hyp...
Microsoft Software




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




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




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



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




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



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

1 2 3 4 >

[Get Widget](#) [Privacy](#)

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