



GABE'S VIRTUAL WORLD

Your P.I. On Virtualization



[HOME](#) [HOT CASES](#) [COLD CASES](#) [TOP CASES](#) [COPYRIGHTS](#) [YOUR P.I.](#) [TWITTER](#) [RSS](#)

Using vscsiStats (the full how-to)

8 December, 2008

After reading verbeieren's [blog about vscsiStats](#), I was very interested to see what this tool could bring me. So I started to play with it and found that I needed some extra hints to get me started. Enough for another blogpost I thought.

With vscsiStats you can monitor the IO through the virtual SCSI controller of a VM (or multiple VMs). vscsiStats works from the console and outputs to the console. The data can be exported as a CSV file so you can import it into Excel or any other tool. The way vscsiStats works is that you first enable statistics collection on one or more VMs and then take samples at several moments. Never forget to stop the collection when you're done, because it will cause some performance degradation.

The vscsiStats command can be found in /usr/lib/vmware/bin and is not part of the path environment, which means that you either first change to this directory or refer to the full path everytime you use the command: /usr/lib/vmware/bin/vscsiStats.

Like said in the first part, you first have to start the collection of data. If you're troubleshooting a single VM, it's best to collect stats for just one VM. Before you can do that, you will have to find the Worldgroup ID of the VM. Strange thing is that when using the 'normal' method for this (vm-support -x) you will find slightly different WID's then what you need for vscsiStats. Better is to first run ./vscsiStats -l. This will give you a list of all available virtual machines and their disks. In my example I want to check my home exchange server called w2k3-ex01-64bit:

```
./vscsiStats -l
```

```
Virtual Machine worldGroupID: 1076,  
Virtual Machine Display Name: w2k3-ex01-64bit {  
Virtual SCSI Disk handleID: 8195  
Virtual SCSI Disk handleID: 8196  
Virtual SCSI Disk handleID: 8197  
}
```

1076 is the WID I need. Now to start collecting data for just this VM run:

```
./vscsiStats -s -w 1076
```

The result will be something like:

```
vscsiStats: Starting Vscsi stats collection for worldGro  
Success.  
vscsiStats: Starting Vscsi stats collection for worldGro  
Success.  
vscsiStats: Starting Vscsi stats collection for worldGro  
Success.
```

The option -s is to start collection, -w determines the WID of the VM you want to monitor. After stats are running for lets say 5 minutes, you can get a sample with:

```
./vscsiStats -w 1076 -p all -c
```


Search my blog

Search

Hot Cases

Memory management and compression in vSphere 4.1
Setting logfile location, swap file, SNMP and vmcore partition in ESX
VMware vSphere 4.1 released – What's new?
First vSphere Plugin for WordPress! by Nicholas Weaver
[RUMOR] VMware to buy Novell?
Must have apps for the iPad
Interview with Matt McSpirt – Partner technology advisor at Microsoft UK

Follow Virtualization

 Follow us on Twitter

Voted #8 blog at
[vSphere-land.com](#)

My Home Lab

CPU: 30,928 Mhz
RAM: 20,409 MB
Storage: 8,815 GB

ESX Hosts: 3
Resource Pools: 5
Virtual Machines: 34

VMotions: 73
Power States:
⏏(11) ⏏(23) ⏏(0)

last update
0 days, 15 hrs and 35 min.

plugin by nickapedia

Again -w is used for the WMD. Then with -p you can select the stats you would like to use. You can choose between: all, ioLength, seekDistance, outstandingIOs, latency, interarrival. When using the stats to put them in a nice histogram or spreadsheet, you can use -c to print the output as csv and ofcourse you could use ' > /tmp/outputfile.csv to redirect the output to file.

Now after you have collected these values for a few times and you're done with collecting, don't forget to stop the collection.

```
./vscsiStats -x
```

```
You will now see the following output:  
vscsiStats: Stopping all Vscsi stats collection for world  
Success.  
vscsiStats: Stopping all Vscsi stats collection for world  
Success.  
vscsiStats: Stopping all Vscsi stats collection for world
```



The tricky part is how to make some sensible info from it. I try to use excel to convert the data into a nice histogram and would love to have a macro for this. But until I have written one, I have to do it by hand. Keep in mind that these are histograms when looking at the following piece of data:

```
44, 4095  
1405, 4096  
10, 8191  
837, 8192
```

These numbers tell you that for 44 times, there was IO command with a blocklength of 4095, 1405 times of 4096, etc, etc. Be sure to interpret them correctly. Having too many blocks of 4095 and 8191 could point to a misaligned partition or disk in this case.

Hope this guide helped you working with the data. Should anyone have a nice excel macro to convert the data into histograms, please mail me or post it in the comments.

by Gabriele van Zanten



Add New Comment

Post as ...

Showing 3 comments

Sort by Oldest first

[Subscribe by email](#)

[Subscribe by RSS](#)

Tom Howarth 1 year ago

very interesting post. I particularly liked the reference to block sizes and misaligned disks in this case.

It is something I will add to the armory :D

Like Reply

Sven Huisman 1 year ago

Great post! I'm wondering, is there a tool to align a misaligned partition or disk without destroying the data on the disk?

Like Reply

Bouke Groenescheij 1 year ago

Very nice post!

@Sven: for VMFS - not a change, but using the VI client will align it correctly (and also creating a VMFS during installation). NTFS / EXT3 in a virtual disk: partition magic or something, you need to move the whole partition with its data to align it.

Like Reply

Reactions

gabvirtualworld 6 months ago

From [Twitter](#) via [UberVU](#)

@m3rc0 esxtop en eventueel kun je nog met vScsiStats aan de slag. <http://bit.ly/74OTol>

Trackback URL <http://disqus.com/forum>

[blog comments powered by DISQUS](#)

« [Prevent your LUNs running out of space](#) [GabesVirtualWorld now on iPhone and PDA](#) »