

Objective 3.2 - Demonstrate Advanced Performance Analysis Techniques

Written by Matthijs van den Berg

Monday, 23 March 2009 15:50

KNOWLEDGE

• Demonstrate the use of various performance tools

The VMware virtual infrastructure provides a couple of performance measurement tools. From the commandline this is esxtop. The GUI provides performance data via the Virtual Center Client.

○ CLI – esxtop

esxtop is a commandline tool to display performance data on a single ESX host. The type of performance data that can be collected are CPU, Memory, Disk and network. When you type

```
esxtop
```

into a command prompt a standard performance view is shown.

```
1:07:37pm up 29 min, 57 worlds; CPU load average: 0.01, 0.01, 0.01
PCPU(%): 2.63, 0.12, 1.54, 0.49 ; used total: 1.20
CCPU(%): 2 us, 0 sy, 98 id, 0 wa ; cs/sec: 288
```

ID	GID NAME	NWLD	%USED	%RUN	%SYS	%WAIT	%RDY	%IDLE	%OVRLP
1	1 idle	4	395.17	395.21	0.00	0.00	4.79	0.00	0.00
2	2 system	6	0.00	0.00	0.00	600.00	0.00	0.00	0.00
6	6 helper	23	0.02	0.02	0.00	2299.94	0.02	0.00	0.00
7	7 drivers	14	0.00	0.00	0.00	1399.95	0.00	0.00	0.00
8	8 vmotion	1	0.00	0.00	0.00	100.00	0.00	0.00	0.00
9	9 console	1	2.45	2.53	0.01	97.46	0.01	97.45	0.07
15	15 vmware-vmkauthd	1	0.00	0.00	0.00	100.00	0.00	0.00	0.00
16	16 VCS001	7	2.03	2.03	0.01	697.96	0.01	198.08	0.15

- This show for the most important processes and the VM's on the system the CPU usage. You can also change to Memory view of Network view by pressing:

- m – Memory view
- c – CPU view
- u – disk devices
- v – virtual disks
- n – network view

- Per view you can sort the data (press o or O), add and remove columns (press f or F). To save the new view press W. A filename needs to be provided, default the standard boot file name is proposed. Using this esxtop will default start with the newly saved settings.

There are many more command options to use, please see man esxtop for all possibilities. See some examples under "Skills and Abilities".

○ Virtual Center GUI

In the GUI you can view performance counters on several levels. For example on VM level, ESX server level, Resource pool level and cluster level. You can see real time and historical data on CPU, Disk, memory, network and system. The historical data depends on the collection level set. Default only limited data is visible. Read more further down.

• Understand configuration options for performance data collection (<http://inet-gw.mapgis.homeip.net>)

○ line graphs vs. stacked graphs

In the GUI you can choose to use line or stacked graphs via "Change chart option" option in the screen.

- Line

Menu - VCDX

- Introduction
- More Info
- - Enterprise Admin
- - Design Exam

VCDX Cloud

virtual server
vmware network
host file service
esxcfg click
configure machine
traffic vswitch
machines storage
command console
hosts cluster disk
client physical
configuration files port

Recent Blog Postings

UPDATE 20 Oktober 2009 - I have received word from VMware: This issue will be addressed in VC 4 Update 1. We expect Virtual Center 4 Update 1 to be released later this year, please be aware this is may be subject to change. The only work around at present is to Disconnect the ESX host from the

[Read more](#)

Add Banners

- Each instance shown separately

- Stacked

- Graphs are stacked on top of each other
- Only applies to certain kinds of charts, e.g.:
 - Breakdown of Host CPU MHz by Virtual Machine
 - Breakdown of Virtual Machine CPU by VCPU

- **Real-time vs. historical metrics**

VMware shows real-time statistics and save a preconfigured portion of that for historical metrics (see next bullet). Default following granularities are stored:

Time Interval	Data Fequency	Number of Samples
Past hour	20 seconds	180
Past day	5 minutes	288
Past week	30 minutes	336
Past month	2 hours	360
Past Year	1 day	365

Data is converted from a shorter interval to a longer interval.

You can change these settings via the Virtual Center Client:

- Administration
- Virtual Center Management Server Configuration
- Statistics
- Place or remove a checkmark in the boxes of the Interval Duration you would like to change.

- **statistics collection levels**

As shown above historical data is stored. However the detail level of the data varies per time interval. Default the followin level are used:

- Level 1 Basic Metrics like average CPU, mem, disk and network
- Level 2 All metrics for CPU, mem, disk and network, no devices
- Level 3 All metrics for all counter groups, no rollup types
- Level 4 All metrics

You can change these settings via the Virtual Center Client

- Administration
- Virtual Center Management Server Configuration
- Statistics
- Click Edit to change the statistics levels.

- When you need to change the collection level of a long interval you need to change the shorter intervals as well to al least that level of stats collection.

Use performance information to troubleshoot and resolve:

- **CPU Utilization issues**

You can have a over committed CPU. This means that the VM demands more CPU than the host can give it during a peri of time. Warnings for this are High Ready Time. More info: <http://inet-gw.maphis.homeip.net/>

- **Memory utilization issues**

When you VM Is running short on memory it will start swapping. VMware uses the ballooning driver to minimize swappin but cannot prevent it at all times. Monitor for swapping an ballooning to detect memory issues.

- **Disk utilization issues**

Disk performance is dependent on many factors:

- Filesystem performance
- Disk subsystem configuration (SAN, NAS, iSCSI, local disk)
- Disk caching
- Disk formats (thick, sparse, thin)

- Look at MB write and read and the latency for performance bottlenecks. Using disk cache or a custum que depth can increase performance.

- **Network utilization issues**

Network statistics show throughput and packets per seconds. Esxtop also shows dropped packets and can show bandwidth per pysical and virtual NIC.

SKILLS AND ABILITIES

- **Use esxtop to monitor the health of the ESX Server**

I have written a [blog posting about esxtop here](#). Be sure to know how to use this command, especially how to modify the data in the view (switch between CPU, memory, Disk and Network performance views), change the columns per view and how to save it to the default.

- **Use vm-support to capture performance snapshots of the ESX Server**

You can use vm-support -s to collect performance snapshots. These snapshots can also collect performance data based on timer. To do so you can use: vm-support -s -i <sleep interval> -d <duration>

- **Use guest OS performance analysis tools to determine performance characteristics within the virtual machine**

I find this a little strange question due to the rather large number of guest OSs supported. Nevertheless we can use the default troubleshooting tools within the guest OS like perfmon, top etc. Also 3rd party tools can be used to monitor and collect performance data.

- **Generate reports and collate data from VirtualCenter**

- **Alarms**

Alarms are used in Virtual Center Server to generate an event like sending an e-mail, send a SNMP trap or run a script. An alarm can apply to an ESX host or a virtual machine and when triggered it also shows up in the alarms tab in the lower left corner of the Virtual Center Client. The items you can monitor by default are selectable in the second tab of the Alarm Settings screen and depend on the host or Virtual Machine choice made in the first screen:

- Host CPU Usage
- Host Memory Usage
- Host Network Usage
- Host Disk Usage
- Host state
- Host Hardware Health
- Virtual Machine CPU Usage
- Virtual Machine Memory Usage
- Virtual Machine Network Usage
- Virtual Machine Disk Usage
- Virtual Machine State
- Virtual Machine Heartbeat

- Alarms can be configured on Datacenter, cluster, host or virtual machine level. To do so click one of those items, go to the tab Alarms and click Definitions.

- **Resource utilization**

To see how much resources are used you can use Virtual Center. When you click a cluster or a host and go to the tab Virtual Machines you see the resources used per virtual machine. Another option is the tab Resources Allocation that is only visible on the cluster level. This shows how much of a reservation is used.

- **Performance**

You can view performance counters in Virtual Center via the tab "Performance". This tab can be viewed on several levels: cluster, host and Virtual Machine.


- **Topology Maps**

A topology map shows how hosts, virtual machines, networks and datastores are connected to each other. You can view dependencies of these items in the Maps view by pressing the Maps button.

- **Diagnose resource utilization issues**

- **CPU ready time/wait time**

When a virtual machine has large periods of time where there is significantly more used CPU time than there is Ready CPU time.



This might indicate a overcommitted CPU. 3 Possible reasons for high ready time:

- CPU overcommitment
Possible solution: add more CPUs or VMotion the VM
- Workload variability
A bunch of VMs wake up all at once
Note: system may be mostly idle: not always overcommitted
- Reservation set on VM
4x2GHz host, 2 vcpu VM, limit set to 1GHz (VM can consume 1GHz)
Without limit, max is 2GHz. With limit, max is 1GHz (50% of 2GHz)
CPU all busy: %USED: 50%; %MLMTD & %RDY = 150% [total is 200%, or 2 CPUs]

o Memory ballooned/swapped

A little explanation what ballooning is: Ballooning is a technique to let VMware manage Virtual Machine memory. It is a driver / kernel extension installed with the VMware tools in the guest OS. When memory is low on the host it triggers the ballooning driver in the guest OS to inflate. Doing this the memory pressure in the guest OS increases. The guest OS notices this, and will start managing the memory in guest OS level, for example start paging to its virtual disk. The host is aware of the memory allocated by the ballooning driver and will use this for other purposes. Read a lot more on ballooning in this very good whitepaper: <http://www.stanford.edu/>

o Disk queue depth/locking

The queue depth is the number of outstanding requests between the HBA and storage controller. The default is 32 but this can be changed per ESX server. Because VMware handles a lot of I/Os for the guest operating systems sometimes increasing this value can increase storage performance. I have written a blog post on how to do this.

o Network dropped packets/

Dropped packets can indicate, amongst other things, a bad or too long cabling problem or a overloaded network. This can only be viewed via the CLI using esxtop. When switching to the network view while pushing n in esxtop. The percentage of packets dropped during transmit is shown in the "%DRPTX" column, the percentage of packages dropped during receive shown in the "%DRPRX" column.

TOOLS

- CLI
 - Esxtop
 - vm-support
- VI Client
 - Performance graphs
 - VirtualCenter management server configuration

< Prev

