

# **VMware vSphere Optimize and Scale [v5.5]**

Custom Lab 12 – Diagnosing CPU performance issues

© Copyright 2014 – Global Knowledge Network Training Ltd.

## Lab 12

### Diagnosing CPU performance issues

**Objective:** Use a performance troubleshooting checklist to detect possible CPU performance problems


In this lab, you will perform the following tasks:

1. Configure advanced performance charts
2. Prepare the lab environment
3. Check for performance issues in a Resource Pool
4. Check for CPU saturation of an ESXi host
5. Check for CPU saturation within a guest Virtual Machine
6. Determine if high guest CPU load causes performance issues

#### Task 1: Configure advanced performance charts

You configure an advanced CPU performance chart for the VMware® ESXi™ host and another CPU performance chart for the test virtual machine.



Student A and student B should do this task independently.

1. In the browser window, click the **vSphere Web Client** tab.
2. In the Hosts and Clusters inventory tree, select the ESXi host.
3. In the middle pane, click the **Monitor** tab and click the **Performance** button.
4. Select the **Advanced** link and collapse the links panel by clicking the  icon.
5. Click the chart **View** menu, select **CPU** and click the **Chart Options** link.
6. Configure and save the ESXi host CPU performance chart.
  - a. For Timespan, leave **Real-time** selected.
  - b. In the **Select object for this chart** panel, click the **None** button and select the ESXi host check box.
  - c. For the Chart Type, leave **Line Graph** selected.
  - d. In the **Select counters for this chart** panel, click the **None** button and select the **Usage** check box.
  - e. At the top of the Chart Options dialog box, click the **Save Options As** button, type **Lab-12a** in the **Chart options name** text box, and click OK.

#### **CAUTION**

Save the chart options as instructed. If the chart resets to a default set of options, you can restore the chart by loading the saved options using the Chart Options dialog box.

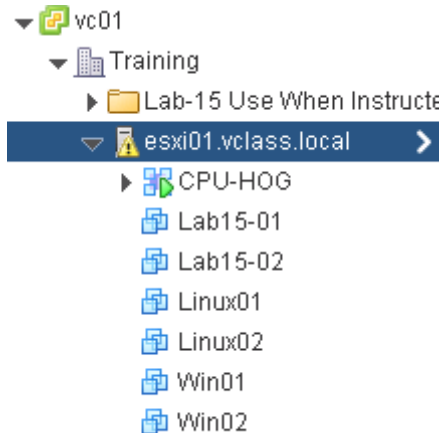
- f. Click **OK** to close the Chart Options dialog box.

7. Verify that the chart **View** selection is **Lab-12a**.
8. In the Hosts and Clusters inventory tree, expand the CPU-HOG vApp and select any one of the **Workload##** virtual machines.
9. In the middle pane, click the **Monitor** tab and click the **Performance** button.
10. Select the **Advanced** link and collapse the links panel by clicking the  icon.
11. Click the chart **View** button, select **CPU** and click the **Chart Options** link.
12. Configure and save the virtual machine CPU performance chart.
  - a. For Timespan, leave **Real-time** selected.
  - b. In the **Select object for this chart** panel, click the **None** button and select the **Workload##** check box.
  - c. For the Chart Type, leave **Line Graph** selected.
  - d. In **Select counters for this chart panel**, click the **None** button and select the **Max Limited, Ready** and **Used** check boxes.
  - e. At the top of the Chart Options dialog box, click the **Save Options As** button, type **Lab-12b** in the **Chart options name** text box, and click **OK**.
  - f. Click **OK** to close the Chart Options dialog box.
13. Verify that the chart View selection is **Lab-12b**.
14. In the Hosts and Clusters inventory tree, select the **Linux01** virtual machine.
15. In the middle pane, click the **Monitor** tab and click the **Performance** button.
16. Select the **Advanced** link and collapse the links panel by clicking the  icon.
17. Click the chart **View** button, select **CPU** and click the **Chart Options** link.
18. Configure and save the virtual machine CPU performance chart.
  - a. For Timespan, leave **Real-time** selected.
  - b. In the **Select object for this chart** panel, click the **None** button and select the **Linux01** check box.
  - c. For the Chart Type, leave **Line Graph** selected.
  - d. In **Select counters for this chart panel**, click the **None** button and select the **Max Limited, Ready** and **Used** check boxes.
  - e. At the top of the Chart Options dialog box, click the **Save Options As** button, type **Lab-12c** in the **Chart options name** text box, and click **OK**.
  - f. Click **OK** to close the Chart Options dialog box.
19. Verify that the chart View selection is **Lab-12c**.

## Task 2: Prepare the lab environment

In this task, you will stop the Linux01 virtual machine and adjust CPU-HOG vApp configuration. Student A and student B should do this task independently.

1. Close all tabs to your virtual machine consoles.
2. Shutdown the Linux01 and any windows virtual machines that might be running, but leave the CPU-HOG vApp running.
3. Make sure your environment state resembles the following screenshot



4. Select your ESXi host and in the middle pane, select the **Summary** tab.
5. In the **Hardware** section, make a note of the maximum Mhz your ESXi host can provide, rounding the number down to the nearest thousand. For example, if your host has 2 x CPUs running at 2.50 Ghz write down 5000. Max Mhz: \_\_\_\_\_
6. Right click the **CPU-HOG** vApp and select **Edit Settings**.
7. Expand **CPU resources** and set the CPU **Limit** to **HALF** of the Mhz value that you recorded in step 5. For instance if you recorded 5000 Mhz, set the **Limit** to 2500. Click **OK**.
8. Leave **resxtop** running in the PuTTY session to the vSphere Management Assistant.
9. Change the **resxtop** interval refresh rate to 20 seconds:
  - a. Switch to your PuTTY session that is running resxtop
  - b. Type **s**.
  - c. Type **20**.
  - d. Press Enter.

### Task 3: Check for performance issues within a resource pool

You examine advanced performance charts and `resxtop` CPU statistics to determine if a Resource Pool (in this case the CPU-HOG vApp) has CPU performance issues.

Student A and student B should do this task independently.

1. Record the advanced performance chart statistics for the host

- a. In the Hosts and Clusters inventory tree, select the ESXi host.
- b. Verify that the chart **View** selection is **Lab-12a** and correct the selection if necessary.

Perform this verification step each time when switching between the various performance charts.

- c. Scroll down to uncover the Performance Chart Legend counter value columns.
- d. Record the column values for Usage.

- Latest (Usage) \_\_\_\_\_
- Maximum (Usage) \_\_\_\_\_
- Average (Usage) \_\_\_\_\_

If the Average column is not displayed, resize one or more of the visible columns. The Average column is always the last column in the Legend and appears when a column resize occurs.

2. Record the `resxtop` statistics.

- a. Switch to the PuTTY (vSphere Management Assistant session) window.
- b. In the section above the per-object statistics table, find `PCPU USED (%)`
- c. Record the PCPU USED (%) AVG value. \_\_\_\_\_

3. Record the advanced performance chart statistics for virtual machines within the resource pool

- a. In the Hosts and Clusters inventory tree, expand the CPU-HOG vApp and select the same **workload###** virtual machines you used earlier.
- b. Verify that the chart **View** selection is **Lab-12b** and correct the selection if necessary.
- c. Scroll down to uncover the Performance Chart Legend counter value columns.
- d. Record the **Latest** column values for each metric.

- Max Limited \_\_\_\_\_
- Ready \_\_\_\_\_
- Used \_\_\_\_\_

The Max Limited metric represents the amount of time that a virtual machine was ready to execute, but has not been scheduled for CPU time because the VMkernel deliberately constrained it.

The values recorded above are the sum total of milliseconds observed over the chart refresh interval period (every 20 seconds). In order to be of more use, they can be converted to a percentage value. For each value recorded above, use the following formula to convert the summation into a percentage value and record the results below.

$$(\text{summation\_value} / (\text{chart\_refresh\_interval} * 1000)) * 100$$

*Note: The chart\_refresh\_interval depends on which type of chart you are currently viewing. For a list of all possible chart intervals see <http://kb.vmware.com/kb/2002181>*

- Max Limited (%) \_\_\_\_\_
- Ready (%) \_\_\_\_\_
- Used (%) \_\_\_\_\_

4. Record the `resxtop` statistics

- a. Switch to the PuTTY (vSphere Management Assistant session) window.
- b. Record the latest Used, Ready and Max-Limited values for the Workload VM you used in the previous step.

- %USED \_\_\_\_\_
- %RDY \_\_\_\_\_
- %MLMTD \_\_\_\_\_

5. Use the gathered statistics to answer the questions.

**Q1. Does the Latest CPU usage metric of the host show a value > 75%?**

*No, the Latest CPU usage metric will be at approximately 50%. At this time the host should not be CPU saturated because the CPU-HOG vApp was restricted to approximately half of the clusters available CPU resources. Maximum and Average cannot be relied upon as they both will be skewed by older stats. These two metrics would only become useful if the CPU issues had persisted for longer than 1 hour.*

**Q2. Do any of the workload virtual machines inside the resource pool (CPU-HOG vApp) have a Max Limited value higher than 0.00 or 0%?**

*Yes. All of the workload virtual machines will show quite a lot of "Max Limited" restrictions.*

**Q3. The Ready metric includes the Max Limited value in its calculation. Comparing Max Limited % to the Ready % values, how much of the Ready time is due to limitations being imposed on the VM by the VMKernel?**

*Nearly ALL of the Ready time shown will be due to the VMkernel imposed Max Limit restrictions.*

**Q4. Considering the workload virtual machines are running a CPU stress testing tool, are the USED percentage values lower than you would expect?**

*Yes. You would expect to see high used percentages if there were adequate resources available.*

The USED percentage of each virtual machine tells you how much CPU resource that VM was given. To determine how much the VM actually wanted to use, add the Used and Ready percentages together. You can use the percentage values either from the Performance Charts or the `resxtop` values. Record the result here: \_\_\_\_\_

**Q5. Based on the result recorded above, are the VMs in the resource pool being given the physical CPU resources they require?**

*No. For the majority of the snapshot period the vCPUs will be stuck in a Ready state.*

**Q6. Based on the information gathered so far, is there a CPU performance problem within the resource pool?**

Yes.

**Q7. Is the CPU performance issue due to host CPU saturation?**

No. The host CPU usage is averaging at around 50%.

**Q6. What if anything could be done to improve the problem?**

Increase the CPU limit on the resource pool or move some of the VMs elsewhere.

## Task 4: Check for CPU saturation of an ESXi host

You examine advanced performance charts and `resxtop` CPU statistics to determine if the ESXi host has CPU performance issues.

1. Return to the **vSphere Web Client** browser tab and power on the Linux01 virtual machine.
2. Right click the Linux01 virtual machine and select **Open Console**.
3. Whilst the Linux01 VM boots, return to the PuTTY window that is connected to the vSphere Management Assistant and observe the values for PCPU USED (%) Average and the %MLMTD for the Linux01 virtual machine. Although the host now has quite high CPU usage, there are no VMKernel imposed restrictions on this virtual machine.
4. Once the Linux01 virtual machine has booted, login as `root` using the standard lab password and run the `starttest2` script. This script is configured to mimic database activity using two separate threads. Allow the script to run for a few moments before continuing.
5. Record the advanced performance chart statistics for the ESXi host.
  - In the Hosts and Clusters inventory tree, select the ESXi host.
  - Verify that the chart **View** selection is **Lab-12a** and correct the selection if necessary.
  - Scroll down to uncover the Performance Chart Legend counter value columns.
  - Record the column values for Usage.
    - Latest (Usage) \_\_\_\_\_
    - Maximum (Usage) \_\_\_\_\_

*Note: Do not record the Average usage, as it shows the average usage over the chart interval (1 hour), which will remain low.*

6. Record the `resxtop` statistics for the host.
  - Switch to the PuTTY (vSphere Management Assistant session) window.
  - In the section above the per-object statistics table, find `PCPU USED (%)`
  - Record the PCPU USED (%) AVG value. \_\_\_\_\_
7. Use the gathered statistics to answer the questions.

**Q1. Is the average CPU usage greater than 75 percent?**

Yes.

**Q2. Is the maximum (peak) greater than 90 percent?**

*Maybe*

If the answer to either question 1 or 2 is yes, then the host *might* be suffering from CPU saturation.

8. In the PuTTY window connected to the vSphere Management Assistant, find the Linux01 virtual machine and observe the values for %USED, %RDY, %CSTP and %MLMTD. Record the values below.

- %USED \_\_\_\_\_
- %RDY \_\_\_\_\_
- %CSTP \_\_\_\_\_
- %MLMTD \_\_\_\_\_

9. Used the gathered statistics so far in this task to answer the following questions.

**Q1. Does the Linux01 virtual machine show more than 10% Ready?**

*Yes, the virtual machine should be showing a lot of %RDY*

It is generally accepted that a Ready value of anything up to 10% is acceptable. Anything above this point could indicate performance problems.

**Q2. Does the Linux01 virtual machine have any limitations imposed by the VMKernel?**

*No. The %MLMTD should read 0.00 since the Linux VM is not in a resource pool and there are no other limits imposed on it.*

**Q3. Add %USED and %RDY together. Record the value here \_\_\_\_\_. Is the result what you expected?**

*Depends on experience. The collapsed view shows the aggregate of all the virtual machines worlds. Since the Linux VM has two CPUs, the maximum achievable value for %USED will be close to 200% (the remaining worlds are for the most part idle and will not contribute much to the overall USED value). When adding %USED and %RDY together the value shown should be > than 100% but not quite 200% either.*

Return to PuTTY and expand the Linux01 group.

1. Return to PuTTY and type `e`.
2. Enter the GID value shown against the Linux01 virtual machine. Press Enter.

**Q4. Review the %CSTP metric. Does it show a value greater than 0.00 for either of the vCPUs?**

*It should do.*

The %CSTP (CO-STOP) metric represents the percentage of time that the virtual machine is ready to execute commands but was waiting for the availability of multiple CPUs. This metric is only relevant to SMP virtual machines (VMs with more than one vCPU) and will show 0.00 for VMs with a single vCPU. If this metric is showing values above 0.00 for either vCPU in the Linux01 VM this can indicate one of two things;

1. If a %CSTP value for either vCPU is persistently higher than the other, this can indicate that there might be too many vCPU's allocated to this VM, the vCPU with the higher %CSTP value is being used much more than the others.
2. If the CSTP values are present but do not show a bias to either vCPU, this could indicate that the ESXi host does not have enough CPU resources available to co-schedule the processors.

The combination of a low USED value with high RDY, a certain amount of CSTP and *no* MLMTD on the Linux01 VM (this particular VM is not in a resource pool) could indicate that there is host CPU saturation *and* that it is causing a performance problem (If RDY was low, we wouldn't necessarily need to worry about the other values). If the current PCPU USED average is also above 75%, there is a high chance of CPU saturation.



## Task 5: Check for CPU saturation within a guest virtual machine

You examine advanced performance charts and `resxtop` CPU statistics to determine if there is possibly CPU saturation present in a virtual machine.

1. Stop the `starttest2` script and run a command that will cause the guests CPU's to spike

- Return to the **Linux01** browser tab, activate the console and press CTRL+C to cancel the running script.
- Type the following command to cause a high amount of guest CPU activity.

```
for i in 1 2; do while : ; do : ; done & done
```

The script starts two background processes. Return to the **vSphere Web Client** tab.

2. Record the advanced performance chart statistics for the virtual machine.

3. In the vSphere Web Client tab, select the Linux01 virtual machine. The advanced performance charts should be displayed with your custom Lab-12c chart view.

4. Scroll down to uncover the Performance Chart Legend counter value columns.

5. Record the **Latest** column values for Used and Ready.

- Used (Summation) \_\_\_\_\_
- Ready (Summation) \_\_\_\_\_

Using the formula from Task 3 - Step 3, convert the summation values to percentages and record them below.

- Used (%) \_\_\_\_\_
- Ready (%) \_\_\_\_\_

6. Using the statistics gathered, answer the following questions.

**Q1. Is the Ready (%) value greater than 10%?**

*Yes. The value should be quite high (>50%)*

**Q2. Is the Used (%) value lower than what you might expect?**

*Yes. The Used value will be affected by RDY.*

**Q3. Does USED (%) + Ready (%) = (Number of vCPUs \* 100%)?**

*Maybe, it depends if CSTP is present*

If the answer is no, return to the PuTTY session and observe the %CSTP value for the Linux01 VM (collapse the world's first!)

**Q3a. Does %USED + %RDY + %CSTP = 200%?**

*Yes, more or less the totals of these will roughly equal 200%.*

%USED + %RDY + %CSTP is the amount of CPU requested by the guest VM. Without looking at any further metrics you can make the assumption that the guest VM is either close to or *is* CPU saturated. Since there is a low value for USED, coupled with a large amount of Ready and possibly CSTP present, the assumption can be made that the VM is CPU saturated but it is also not getting enough CPU entitlements from the host.

## Task 6: Determine if high guest CPU load causes performance issues

In this task you stop the CPU-HOG vApp and determine if saturated guest CPUs are always a cause for concern.

1. Return to the vSphere Web Client tab and select the CPU-HOG vApp.
2. From the **Actions** menu, select **Power Off**. Click **Yes** on the confirmation dialog.
3. Return to the PuTTY session connected to the vSphere Management Assistant and wait for all the Workload## VMs to disappear.
4. Wait a little while for resxtop to refresh and observe what happens to the values.

You should see that %USED increases to almost the maximum amount (approximately 200%). %RDY should decrease to acceptable levels (at or lower than 10%). %CSTP should also have reduced to very minimal levels. PCPU USED (%) for the host should be high at approximately 99%.

### Summary

Although the host CPU is peaking at almost 100% activity, the only active virtual machine for the most part is getting the resources it requires (at least within acceptable tolerances). Although the depicted scenario of a single VM on an ESXi host is a little unlikely, the point here is that high CPU activity observed on a guest VM is not always a cause for concern. Whilst it is pretty clear that the Linux01 VM is CPU saturated, the guest is not suffering from any VMware related performance issues. To investigate if the guest itself was having performance problems as a result of the CPU saturation, you would need to login to the guest OS to investigate further.

## Task 7: Clean-up for the next lab

1. Return to the console of the Linux01 VM and run the following command to terminate the CPU spike script:

```
killall bash
```

This command will terminate the running scripts and display an output similar to the following when complete:

```
[1]- Terminated      while ;; do
      ;;
done
[2]- Terminated      while ;; do
      ;;
done
```

### NOTE

It may be necessary to type the `killall bash` command twice, before the output above is displayed.

2. When the prompt returns, start the `starttest2` script in preparation for the next lab. Type `./starttest2`.
3. Return to the **vSphere Web Client** tab.
4. In the Hosts and Clusters inventory tree, select the **CPU-HOG** vApp.
5. Select **All vCenter Actions** -> **Remove from Inventory** from the **Actions** drop-down menu.
6. On the Confirm Remove dialog box, click **Yes**.