



VMware ESX Server 2

Using esxtop to Troubleshoot Performance Problems

The VMware `esxtop` tool provides a real-time view (updated every five seconds, by default) of ESX Server worlds sorted by CPU usage. The term *world* refers to processes running on the VMkernel. There are three types of worlds:

System: The worlds that are needed to perform various system services. These include one idle world per physical CPU that runs when there is nothing else to run on that physical CPU, helper worlds for performing asynchronous tasks and driver worlds.

Service Console: The world for the service console. It always runs on physical CPU0.

Virtual Machine: The world for each virtual CPU. This is the world you look at when troubleshooting.

Also, `esxtop` displays information about the state of the physical server running an ESX Server. It lists CPU utilization for each physical processor, memory utilization, and disk and network bandwidth for each network and disk device available to the ESX Server machine.

Furthermore, `esxtop` lists CPU and memory utilization for each individual VMkernel world. Memory utilization is characterized by the type of memory (for example, shared, private, or swapped) that is being consumed. These CPU and memory statistics let you monitor the resource utilization for each of your virtual machines.

The following sections describe the process for troubleshooting a suspected performance problem at the virtualization layer (i.e., the ESX Server layer):

- [Using esxtop on page 2](#)
- [Other Tools on page 6](#)



Using esxtop

This section describes the necessary steps to troubleshoot your ESX Server machine performance problems using `esxtop`. The steps are:

- [Getting Started](#)
- [Starting esxtop](#)
- [Examining CPU Usage](#)
- [Assessing Memory Usage](#)
- [Assessing Disk and Network Usage](#)
- [Exiting esxtop](#)

Getting Started

Do the following before you start to troubleshoot a problem using `esxtop`:

1. Log on to the VMware Management Interface for the ESX Server machine in question. Refer to the online document, [Logging Into the VMware Management Interface](#), for details. In the status monitor, under **Virtual Machines**, note the virtual machine IDs (or VMIDs) for all virtual machines running on the server.

Virtual Machines (6)			HB	Display Name
				ts-eservices eroom w2k server (DON'T DELETE) Powered on PID 7797 VMID 143
				Linux 100MB Powered off
				w2k3entopen.dsk Powered off
				supp073-VirtualCenter 1.0.0 Powered off
				Debian31 with Apache Suspended
				supp072-VirtualCenter 1.1.1 Powered on PID 2345 VMID 142

Partial screen shot of the VMware Management Interface showing virtual machine IDs

2. Make certain you have an secure shell (SSH) client. Windows users can get a free SSH client from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>.
3. If you have ESX Server version 2.0.x, refer to the VMware Knowledge Base Answer ID 1078 for instructions on downloading and installing the VMware performance monitoring tools, `esxtop` and `vmkusage`. ESX Server version 2.1 and higher include `esxtop` and `vmkusage`. See [Using vmkusage to Isolate Performance Problems on page 6](#) for a description of `vmkusage`.

Starting esxtop

Perform the following steps to start and set up `esxtop`.



1. Using a secure shell (SSH), log on to the ESX Server machine as root.
2. Enter `esxtop` in the SSH command line. The `esxtop` display appears.

```

10:09am up 22:09, 16 worlds, load average: 0.03, 0.01, 0.00, 0.00
PCPU:  3.49%,  1.95% :  2.72% used total
LCPU:  3.07%,  0.42%,  1.91%,  0.04%
MEM: 850944 managed(KB), 270336 free(KB) : 68.23% used total
SWAP: 1047552 av(KB), 0 used(KB), 1037080 free(KB) : 0.00 MBr/s, 0.00 MBw/s
DISK vmhba0:6:0: 0.00 r/s, 0.00 w/s, 0.00 MBr/s, 0.00 MBw/s
DISK vmhba0:0:0: 0.00 r/s, 7.57 w/s, 0.00 MBr/s, 0.02 MBw/s
NIC vmnic1: 0.00 pTx/s, 14.55 pRx/s, 0.00 MbTx/s, 0.01 MbRx/s
NIC vmnic0: 0.00 pTx/s, 14.55 pRx/s, 0.00 MbTx/s, 0.01 MbRx/s

```

VCPUID	WID	WTYPE	%USED	%READY	%EUSED	%MEM
129	129	idle	59.86	0.00	59.86	0.00
128	128	idle	50.83	0.00	50.83	0.00
131	131	idle	45.77	0.00	45.77	0.00
130	130	idle	38.14	0.00	38.14	0.00
127	127	console	2.31	0.02	2.31	0.00
142	142	vmm	2.29	0.36	2.29	35.00
143	143	vmm	0.76	0.22	0.76	15.00
132	132	helper	0.02	0.22	0.02	0.00
140	140	driver	0.00	0.00	0.00	0.00
139	139	reset	0.00	0.00	0.00	0.00
138	138	reset	0.00	0.00	0.00	0.00
137	137	helper	0.00	0.00	0.00	0.00
136	136	helper	0.00	0.00	0.00	0.00
135	135	helper	0.00	0.00	0.00	0.00
134	134	helper	0.00	0.00	0.00	0.00
133	133	helper	0.00	0.00	0.00	0.00

Note: The `esxtop` tool includes several interactive commands. To view a list of the interactive commands, enter `h`.

3. Enter the `f` command. The Field Select page appears.

```

Current Field Order: ABCDEfghijklmnoPqrstuvwX
Toggle fields with a-x, any other key to return:

```

- * A: VCPUID = VCPU Id
- * B: WID = World Id
- * C: WTYPE = WORLD Type
- * D: %USED = CPU Usage
- * E: %READY = CPU Ready
- F: %SYS = CPU System
- G: %WAIT = CPU Wait
- H: CPU = CPU Last Used
- I: AFFINITY = CPU Affinity
- J: HTSHARING = HT Sharing
- K: MIN = CPU Shares Min
- L: MAX = CPU Shares Max
- M: SHARES = CPU Shares Allocated
- N: EMIN = CPU Shared Effective Min
- * O: %EUSED = CPU Effective Usage
- * P: %MEM = MEM Usage
- Q: UNTCHD = MEM Untouched (MB)
- R: SWPD = MEM Swapped (MB)
- S: SWAPIN = MEM SwapIn (MB/s)
- T: SWAPOUT = MEM SwapOut (MB/s)
- U: MCTL = MEM Ctl (MB)
- V: SHRD = MEM Shared (MB)
- W: PRVT = MEM Private (MB)
- X: OVRHD = MEM Overhead (MB)

4. Enter `r` to toggle on the SWPD field.



- Press any key other than a through x to see the `esxtop` display again.

```

1:32pm up 1 day, 1:32, 16 worlds, load average: 0.04, 0.03, 0.03, 0.01
PCPU:  3.36%,  4.18% :  3.77% used total
LCPU:  3.12%,  0.24%,  1.87%,  2.31%
MEM: 850944 managed(KB), 271360 free(KB) : 68.11% used total
SWAP: 1047552 av(KB), 0 used(KB), 1037080 free(KB) : 0.00 MBr/s, 0.00 MBw/s
DISK vmhba0:6:0: 0.00 r/s, 0.00 w/s, 0.00 MBr/s, 0.00 MBw/s
DISK vmhba0:0:0: 0.00 r/s, 6.77 w/s, 0.00 MBr/s, 0.02 MBw/s
NIC vmnic1: 0.00 pTx/s, 23.70 pRx/s, 0.00 MbTx/s, 0.02 MbRx/s
NIC vmnic0: 0.00 pTx/s, 23.70 pRx/s, 0.00 MbTx/s, 0.02 MbRx/s

```

VCPUID	WID	WTYPE	%USED	%READY	%EUSED	%MEM	SUPD
130	130	idle	50.65	0.00	50.65	0.00	0.00
128	128	idle	48.63	0.00	48.63	0.00	0.00
131	131	idle	48.08	0.00	48.08	0.00	0.00
129	129	idle	45.05	0.00	45.05	0.00	0.00
127	127	console	3.09	0.02	3.09	0.00	0.00
142	142	vmx	2.31	0.28	2.31	38.00	0.00
143	143	vmx	1.97	0.00	1.97	11.00	0.00
140	140	driver	0.00	0.00	0.00	0.00	0.00
139	139	reset	0.00	0.04	0.00	0.00	0.00
138	138	reset	0.00	0.00	0.00	0.00	0.00
137	137	helper	0.00	0.00	0.00	0.00	0.00
136	136	helper	0.00	0.00	0.00	0.00	0.00
135	135	helper	0.00	0.00	0.00	0.00	0.00
134	134	helper	0.00	0.00	0.00	0.00	0.00
133	133	helper	0.00	0.00	0.00	0.00	0.00
132	132	helper	0.00	0.00	0.00	0.00	0.00

Note: You can also run `esxtop` in batch mode. For example, use the command:

```
[root]# esxtop -b -n iterations > logfile.
```

For detailed command reference information, enter `man esxtop` on the SSH command line.

Examining CPU Usage

This section describes how to assess system CPU loading, percentage of individual CPU use and individual virtual machine CPU use.

Load Average Line

Examine the load average on the first line to determine the amount of use for all physical CPUs on the ESX Server machine. The load averages are displayed for five-second, and one-, five- and fifteen-minute intervals.

A load average of 1.00 means that the ESX Server machine's physical CPUs are fully utilized, and a load average of 0.5 means they are half utilized. On the other hand, a load average of 2.00 means that you either need to increase the number of CPUs or decrease the number of virtual machines running on the ESX Server machine because the system as a whole is overloaded.

PCPU Line

Examine the PCPU line for the percentage of individual physical CPU use for CPU0 and CPU1 respectively (for a dual-processor machine). The last value is the average percentage for all of the physical CPUs.

As a rule of thumb, 80.00% is a desirable usage percentage, but bear in mind that different organizations have varying standards with respect to how close to capacity they run their servers. 90% should be considered a warning that the CPUs are approaching an overloaded condition.



You can enter the interactive `c` command to toggle the display of the `PCPU` line. If hyper-threading is enabled, the `LCPU` line appears whenever the `PCPU` line is displayed. The `LCPU` line shows the logical CPU use.

Virtual Machine CPU Usage

A virtual machine world is listed as `vmm` in the `WTYPE` column. The world ID (`WID`) corresponds to the `VMID` in the VMware Management Interface Status Monitor (see [Getting Started on page 2](#)). For virtual machines with one virtual CPU (VCPUs), the `VCPUID` and `WID` is the same. For virtual machines with two VCPUs, there are two `VCPUIDs` associated with one `WID`. For example:

```
VCPUID  WID  WTYPE ...
135     135  vmm ...
136     135  vmm ...
```

Use the `WID` and `VMID` values as cross references to identify a specific virtual machine's display name.

Use the following steps to assess virtual machine CPU usage.

1. Examine the `%READY` field for the percentage of time that the virtual machine was ready but could not get scheduled to run on a physical CPU. Under normal operating conditions this value should remain under 5%.
2. Examine the `%USED` field for the percentage of physical CPU resources used by a VCPU. If the physical CPUs are running at full capacity, you can use `%USED` to identify a virtual machine that is using a large amount of physical CPU resources.
3. Examine the `%EUSED` field for the percentage of the maximum physical CPU resource usage a virtual machine is currently using. The `%EUSED` value is calculated as follows:

$$\%EUSED = \%USED * ((\# \text{ of VCPUs} * 100) / \text{max})$$

In this formula:

of VCPUs is the number of VCPUs configured in a virtual machine.

max is the maximum percentage of CPU resources allocated to a virtual machine. The default is 100.

The following table shows how `%EUSED` increases as the value of *max* decreases for a virtual machine with a constant physical CPU usage of 40% (i.e., `%USED` = 40).

<code>%USED</code>	Number of VCPUs	Maximum % of Physical CPU Allocated	<code>%EUSED</code> (in Bold)
40	1	100	$40 * ((1 * 100) / 100) = \mathbf{40}$
40	1	80	$40 * ((1 * 100) / 80) = \mathbf{50}$
40	1	50	$40 * ((1 * 100) / 50) = \mathbf{80}$
40	1	40	$40 * ((1 * 100) / 40) = \mathbf{100}$ The virtual machine is using 100% of physical CPU resources.

`%EUSED` is an useful indicator of how close a virtual machine is to saturating the physical CPUs. If a virtual machine consistently saturates the physical CPUs (i.e., it uses 100% of the physical CPU resources), you can fix it by either:

- Decreasing the number of virtual machines running on an ESX Server machine.



- Moving the virtual machine to a different ESX Server machine that has more physical CPU resources, and increasing the value of `max`.

Assessing Memory Usage

Look at the percentage of maximum memory actively used by a virtual machine listed in the `%MEM` column. Also, note the amount of swapped memory used by a virtual machine listed the `SWPD` column. Using some swap space is not necessarily bad, because the pages swapped out could be inactive; however, the use of swap space is a clue that you could be losing performance due to active swapping. If the swap percentage increases, there could be a performance problem. If the maximum system memory or swap space is exceeded, additional virtual machines will not power up and you may need to reconfigure memory and swap space.

Assessing Disk and Network Usage

Examine the `DISK` and `NIC` lines for disk activity. If the throughput is either not meeting expectations or approaching the maximum capacity of the hardware device, there is a potential for a performance bottleneck.

Exiting `esxtop`

Enter `q` to exit `esxtop`.

Other Tools

Using `vmkusage` to Isolate Performance Problems

The `vmkusage` tool displays historical graphs that show physical server and virtual machine system statistics. These graphs show the most recent data, as well as daily and weekly views. The tool generates the graphs as Web pages that you can view by going to `http://<ESXservername>.<your_company>.com/vmkusage`.

The data is collected every minute, and the graphs are refreshed every five minutes. The Web pages generated by the `vmkusage` tool show CPU utilization and memory utilization, as well as network and disk bandwidth. You can use this information to isolate performance problems.

Using the VMware Management Interface to Detect CPU and Memory Performance Problems

The VMware Management Interface is a Web-based management tool that lets you monitor the status of virtual machines running on an ESX Server machine. Refer to the *ESX Server 2 Administration Guide* for a complete description of using the VMware Management Interface.

To use the VMware Management Interface:

- You must know the ESX Server machine's name or IP address.
- You must have a valid user name and password on that server.

Status Monitor

The status monitor provides system and virtual machine status information that is averaged over five minutes and is refreshed every 90 seconds. Status indicators that are consistently red indicate a system or virtual machine configuration or capacity problem.



Memory

The Memory page displays system and virtual machine memory and swap space information. As with the status monitor, information in red indicates a problem with the system or virtual machine configuration or capacity.