

## What does 100% utilised actually mean?

When you look at processor utilisation data you will see that it is usually expressed as a percentage. Intuitively, 0% utilised means that no work is being done. Most systems report 100% as indicating that it is working flat out, occupying all its time with running software. There are a couple of notable exceptions: some proprietary monitoring packages take 100% to mean 100% of a single processor – therefore a fully utilised 6

processor server would register as 600% used. Compaq/Digital's VMS *monitor* utility also works this way. Fortunately all the tools that we'll use work as expected.

When we come to measure how busy a server is, we monitor it for a period of time and take the average utilisation. This smoothes out short bursts of activity and gives us a figure that people find easy to deal with, something like "69% utilised from 10:00 to 11:00".

There is a problem here. The figure given above of 69% could easily come about by having the server 90% busy for the first half-hour and 48% busy for the second. If you were planning to increase the server load by 25% this would change the utilisation in the first half hour from 90% to 112%,

leading to an overloaded server and potentially unhappy users.

There are 2 ways to prevent this: we could either monitor the server at more frequent intervals, or we could set the maximum usage threshold low enough to make this unlikely to occur. I say "unlikely" rather than "certain" as there is always a possibility that a short but intensive

```


OpenVMS Monitor Utility
TIME IN PROCESSOR MODES
on node TST104
11-NOV-2000 06:13:05.17
Combined for 3 CPUs      CUR   AVE   MIN   MAX
Interrupt State         7.12  4.87  2.21  9.37
MP Synchronization      2.03  1.38  0.50  2.94
Kernel Mode             22.46 14.39  6.86 25.42
Executive Mode          23.95 17.72  8.02 36.09
Supervisor Mode         0.01  0.02  0.01  0.03
User Mode               17.12 10.67  5.49 17.36
Compatibility Mode      0.00  0.00  0.00  0.00
Idle Time               228.83 150.17 95.20 228.83

```

**VMS CPU Monitoring output**

process could be run that is masked by the smoothing effect of taking averages.

Here is a graph of 10 minutes of CPU activity taken from a live system running a mixture of batch processes (with no user think time) and interactive applications.

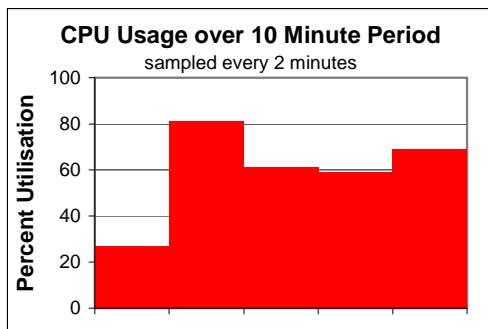


There are quite clearly times when the utilisation reaches 100%. Fortunately these are only for a few seconds at a time, otherwise the users whose applications are delayed may well start to complain and reach for their SLA documents

to see what level of responsiveness they are entitled to.

At this level of detail we would definitely advise caution before adding to the workload this server is running.

Now look at the same system monitored at the same time, with the same tool (vmstat) but with a 2 minute sampling period.



One could argue that there was 20% spare capacity on this server. The short spikes of 100% load during the last 5 minutes have gone completely – being averaged out by the troughs when there was less demand on the machine.

In fact, over the entire 10 minutes of this example, the average utilisation came out at, uncoincidentally, 69%.

Depending on the workload that a server is running, a short delay when it is at 100% utilisation for a few seconds may not matter. On the other hand it could be quite annoying. Imagine you had phoned a call-centre and during your dialog the call agent's PC froze for 5 seconds. Count it out for yourself and 5 seconds seems like a long time. From the call-agent's point of view, if this happened on every call he or she would soon be complaining about poorly running

applications (though maybe not in quite those terms!).

If we are to accept longer sampling periods, the threshold above which more work should not be added becomes lower. In this example, we could say that the server is already running at capacity at some times and therefore, when measured at 10 minute intervals the upper utilisation limit should be around 70%. Obviously this depends on the pattern of work on a given server and the tolerance of it's users.

Our experience has led us to a general conclusion that for workloads which are interactive in nature, a threshold of 85% CPU utilisation measured at 5 minute intervals gives an acceptable level of response. We do qualify this by also measuring the short-term CPU usage and determining if there is a pattern of 100% load that is masked by the longer term data.

One lesson that comes out of this is the conflict between the requirements of batch processes to consume as much resource as they can get and interactive users who require resources to be free so they can achieve acceptable response times for their work. While there are software solutions for this problem, the most reliable way to resolve this is to keep the two types of work separate – either in time, or by running them in dedicated environments.

Dalleon Systems provides services and consultancy for monitoring the performance of your servers. For more information visit our website at [www.dalleon.co.uk](http://www.dalleon.co.uk)