

VSE CPU Monitor Tool

This document describes details and setup information for the VSE CPU Monitor Tool (VSE CPUMON).

Note: The VSE CPUMON Tool is provided **as-is**, no support, no warranty.

The VSE CPUMON tool is intended to help customers to measure the CPU utilization of their VSE system over a period of time. Especially when you plan for a processor upgrade it is very important to know the CPU utilization of your VSE system over a day or a week. This helps you to estimate the size of the new processor.

VSE CPUMON is **not** intended to replace any existing monitoring product provided by vendors. Instead it provides only very basic monitoring capabilities on an overall VSE system level. No details about CPU usage of certain applications are provided.

VSE CPUMON periodically issues a TDSERV FUNC=TDINFO macro to get performance relevant data. The data provided by the macro is the same as command QUERY TD shows. The data from each measurement interval is printed to SYSLST in a comma separated format. Later on this data can be imported into a spreadsheet (see below).

VSE CPUMON provides the following data:

- Timestamp
- Number of CPUs defined
- Number of active CPUs
- Number of quiesced CPUs
- Number of standby CPUs
- Accumulated CPU time in milliseconds since last reset
- Accumulated nonparallel CPU time in milliseconds since last reset
- Accumulated spin time in milliseconds since last reset
- accumulated dispatcher cycles since last reset
- Accumulated number of SVCs since last reset
- CPU time in last measurement interval in milliseconds
- Non-parallel CPU time in last measurement interval in milliseconds
- CPU utilization in last measurement interval in percent
- Non-parallel CPU utilization in last measurement interval in percent
- NPS in last measurement interval (NP/TOT, Non-Parallel-Share)
- Maximum number of CPUs exploited (= 0.9 / NPS)
- Rolling average for CPU time in 8 measurement intervals

For technical questions please send an e-mail to zvse@de.ibm.com.

Installing VSE CPUMON

To install the VSE CPUMON tool you must upload the file CPUMON.BJB either to VM or to your VSE reader in BINARY and LRECL=80. The file contains a binary job that catalogs CPUMON.PHASE into a library of your choice (default is PRD1.BASE). The job contains a PAUSE statement that allows you to specify the destination library.

Running VSE CPUMON

CPUMON runs in a VSE partition (dynamic or static). The partition size must be at least 5 MB. CPUMON is started using

```
// EXEC CPUMON,PARM='nn [RESET] [DD:output] [SYSINFO] [XML] '  
/*
```

where 'nn' specifies the measurement interval in seconds. Default is 60 seconds. The tool can be stopped by entering the following command:

```
MSG xx,DATA=EXIT
```

where xx specifies the partition where VSE CPUMON is running.

Note: It is recommended to reset the counters before starting VSE CPUMON. You can do that using the following command:

```
SYSDEF TD,RESETCNT
```

This command can be included into a job stream that starts VSE CPUMON:

```
// EXEC DTRIATTN,PARM='SYSDEF TD,RESETCNT '  
/*  
// EXEC CPUMON,PARM='nn '  
/*
```

Alternatively, you can run CPUMON in *RESET mode*. In this case CPUMON will reset the counters at every interval. This can be useful when a very long timeframe is to be measured, and the internal VSE counters would overflow. To get more consistent data in this case, you can run the CPUMON Tool in *RESET mode*:

```
// EXEC CPUMON,PARM='nn RESET '  
/*
```

You can manually reset the counters and start/stop CPUs while VSE CPUMON is running, however the measured data might be inconsistent. VSE CPUMON will detect this and issues a corresponding message to the console. It also writes a special line to SYSLST to indicate a counter reset.

VSE CPU Monitor Tool

VSE CPUMON does not reset the counters (unless it is running in *RESET mode*) nor modify any VSE system parameters by itself.

Note: It is recommended that CPUMON runs in a partition with a priority just below VSE/POWER. Otherwise it can happen that the tool gets not enough CPU to collect the data.

Per default VSE CPUMON prints the output to SYSLST. You can specify a different output destination as follows:

```
// EXEC CPUMON,PARM='nn [RESET] DD:output '  
/*
```

where `DD:output` can be for example `DD:SYS009`. In this case `SYS009` can be assigned to a printer other than SYSLST.

To avoid a POWER queue file full condition, you should separate the listing after a specific number of pages. You can do that using the `RBS=nn` parameter in the POWER list card.

The output can also be directed to a VSAM ESDS cluster instead of SYSLST. To do that you must specify the DLBL name of the VSAM ESDS cluster as second parameter (e.g. `DD:MYCLUST`). The cluster must have a maximum record length of at least 250 bytes. To retrieve the data, you can FTP the contents of the VSAM cluster to your PC in text mode.

VSE CPUMON has been tested on VSE/ESA 2.3 and later. On VSE/ESA 2.3 it only works if Turbo Dispatcher is active, it does not run with Standard Dispatcher. From VSE/ESA 2.4 and later only Turbo Dispatcher is supported.

VSE CPUMON can print information about the processor it is running on. Use the `SYSINFO` parameter to do so:

```
// EXEC CPUMON,PARM='nn SYSINFO '  
/*
```

Note: This feature is only available on newer z Series, System z or z Systems servers, i.e. on IBM zSeries 890, 990 or newer servers, but not on IBM zSeries 800 or 900 if running z/VSE in an LPAR (i.e. not under z/VM).

For use with the zCP3000 capacity planning tool, CPUMON can produce output in XML format. Use the `XML` parameter to do so:

```
// EXEC CPUMON,PARM='nn XML '  
/*
```

Note: You can combine various parameters, like `XML`, `SYSINFO`, `RESET` and `DD:output`. For use with zCP3000, it is recommended to specify the `XML` and `SYSINFO` parameters.

Analyzing the measured data

To analyze the data measured by VSE CPUMON you simply transfer the listing created by VSE CPUMON to a directory on your PC. You can do that using IND\$FILE file transfer, TCP/IP FTP or any other file transfer method of your choice. The transfer must be done in text format.

Once the listing is in a directory on your PC, you open the file with a text editor (e.g. Notepad) and remove the surrounding JCL. The data starts with a line that looks as follows:


```
----- START OF DATA -----
Timestamp;#CPUs;#ActCPUs;#QuiescedCPUs;#StandbyCPUs;CPUTime;NPTime;....
```

The data ends just before the line

```
----- END OF DATA -----
```

Next step is to import this data into a spreadsheet. In the following the steps are explained for use with Microsoft Excel. If you are using another spreadsheet program, the steps are similar.

- 1.) Start Excel and open a new spreadsheet (Menu: File->New).
- 2.) Import the data Menu: Data->Import external data->Import data.
- 3.) Select the file you just downloaded and click "Open".
- 4.) In the Text Import Wizard choose "Delimited" and click "Next".
- 5.) Select "Semicolon" as delimiter and click "Next".
- 6.) Click on the "Advanced" Button and make sure you use '.' as "Decimal separator" (e.g. in Germany the default is ',').
- 7.) Finish the import wizard and let Excel import the data into the existing worksheet.
- 8.) Select the all rows in columns L, M and P (%CPU, %NP and %CPURollAvg). Right click the selection and choose "Format cells" in the context menu.
- 9.) Choose "Percentage" with 2 decimal places and click "OK".
- 10.) Select the all rows in columns L, M and P (%CPU, %NP and %CPURollAvg). To create a chart use Menu: Insert->Chart. Choose a Line chart-type and press "Next".
- 11.) Select the "Series" Tab and give the 3 series a name of your choice (e.g. %CPU, %NP and %CPURollAvg).
- 12.) To define the "Category (X) axis label" click on the range selector on the right. Select all timestamps in column A. Click "Finish" to create the chart.



Converting CPUMON XML data to CSV

In case you are using the XML option to produce data for use with zCP3000 capacity planning tool, but you want to also analyse the data in a spreadsheet program on your own, you can convert the XML data into CSV data with the XML2CSV tool.

The XML2CSV tool is contained in the XML2CSV subdirectory. It runs on the workstation and requires a Java Runtime Environment (JRE) version 1.5 or later. Please see here <http://www-03.ibm.com/systems/z/os/zvse/downloads/#install> for more information about checking if Java is installed and how to download Java if it is not installed.

To run XML2CSV, execute run.bat/run.cmd (Windows) or run.sh (Linux, Unix). If no arguments are passed a graphical dialog will prompt for the input file(s) containing the XML data to be converted. The next dialog will prompt for the output CSV file to be created.

Alternatively, you can run the XML2CSV tool in batch mode. Here you pass the name(s) of the input and output files as parameters:

```
run.bat <infile-1> [<infile-2> .... <infile-n>] [-o <outfile>]
```

If no output file is specified, then the output file will be created as output.csv in the same directory as the first input file is located.

Note: Multiple input files are combined/appended in sequence. Specifying multiple input files is useful when you have multiple XML files from the **same system**, but from different days or weeks. Do **NOT** specify multiple input files from **different systems**.