



EPV Graph for z/OS Plus Installation and Customization in Unix/Linux



Supporting
EPV Graph for z/OS Plus V11

March 2013



All the mentioned trademarks belong to their respective companies.

EPV Technologies contact details:

EPV Technologies
Viale Angelico, 54
00195 Roma
Tel. 06 86210880
Fax. 06 86387461
E-mail: epvtech@epvtech.com
WEB: <http://www.epvtech.com>



Contents

- 1 Overview - 5 -
- 2 Architecture..... - 7 -
- 3 Installation..... - 8 -
 - 3.1 Preparing products and password folders in Unix/Linux..... - 8 -
 - 3.2 Verifying the HTML publishing folders..... - 8 -
- 4 Customization (mandatory)..... - 10 -
- 5 Scheduling..... - 11 -
 - 5.1 Collecting data once a day - 11 -
 - 5.2 Collecting data in continuous mode using EPV agents - 11 -
- 6 Manual Customization (optional) - 12 -
 - 6.1 Customizing the CONFIG.PL parameters - 12 -
 - 6.2 Setting SHIFTS - 12 -
 - 6.3 Assigning a user name to each CEC serial number - 12 -
 - 6.4 Excluding CECs from rolling 4-hour average analysis - 12 -
 - 6.5 Changing the series colors in EPVGRAPH HTML pages..... - 12 -
 - 6.6 Changing the number of days to report daily..... - 12 -
 - 6.7 Setting your country holidays - 13 -
 - 6.8 Modifying batch throughput - 13 -
 - 6.9 Including or excluding systems..... - 13 -
 - 6.10 Publishing on the IBM HTTP Server on zSeries - 13 -
 - 6.11 Exporting views in Microsoft Excel - 14 -
 - 6.12 Customizing the user interface..... - 14 -
 - 6.13 Customizing the Recovery procedure - 14 -
- 7 Using the EPV Customization GUI - 15 -
 - 7.1 ADVANCED customization..... - 15 -
 - 7.2 Customizing the product’s parameters..... - 17 -
 - 7.3 Additional customization steps - 17 -
- 8 User defined graphs - 18 -
 - 8.1 Preparing sequential text files for the UGE - 18 -
 - 8.2 Preparing the menu structure for the UGE..... - 19 -
 - 8.3 Setting graph parameters..... - 20 -
- 9 Customer support - 22 -
- Attachment A – Config parameters - 23 -
- Related documentation..... - 26 -



About this manual

This manual is intended to help anyone wanting to install and customize EPV Graph for z/OS Plus V11 in a Unix/Linux environment.

Changes

Technical changes or additions to the text are indicated by a vertical line to the left of the change.

Terminology

A “view” is an EPV report presented in an HTML page.



1 Overview

Enterprise Performance Vision (EPV) **Graph for z/OS Plus**, EPVGRAPH in the following, is a product designed to provide the system manager an immediate and clear vision of company's systems and workload.

EPVGRAPH is integrated with the suite of EPV products, adding graphical functionality which permits intuitive and simple analysis of:

- **Workload activity:** EPVGRAPH provides an immediate vision of the daily and historical CPU utilization; by means of drill down functionality it is possible to navigate from enterprise level to CEC, partition, subsystem down to the single address space level; EPVGRAPH provides information on the CPU utilization in MIPS for GCP, zAAP and zIIP; GCP graphs allow to instantaneously switch from MIPS to MSU by simply choosing the metric to analyse;
- **Memory Utilization:** EPVGRAPH provides an immediate vision of the daily and historical memory utilization; by means of drill down functionality it is possible to understand which service or report classes are using the memory.
- **Throughput activity;** EPVGRAPH provides the daily and historical productivity in terms of transactions for IMS, CICS, DDF, TSO, MQ, WEBSphere and Batch, at enterprise and detail level for each workload type and system;
- **WLM Importance:** EPVGRAPH provides a daily vision of the CPU utilization classified by WLM importance. By means of drill down functionality it is possible to navigate from enterprise to CEC and partition level;
- **Disk Space:** EPVGRAPH provides an immediate vision of the daily and historical disk space utilization; by means of drill down functionality it is possible to understand which storage groups are using the disk space and how much disk space is still available.
- **Tape Space:** EPVGRAPH provides an immediate vision of the daily and historical tape space utilization; by means of drill down functionality it is possible to understand which pools are using the tape space and how much tape space is still available.
- **WLC:** EPVGRAPH reports the used MSU, calculated as a “4 hours rolling average”, accordingly with the IBM WLC policy. EPVGRAPH allows to analyse the charged MSU by day and partition. Additional graphs are provided to analyse Group Capacity effectiveness and the possibility to reduce WLC costs by exploiting zAAP and zIIP and WLM controls;
- **IPLS:** EPVGRAPH provides an immediate historical vision of all the IPLS performed; by drill down functionality it is possible to understand in which day of the week and in which time the IPL has been performed.
- **MF Indexes:** EPVGRAPH allows to analyse processor cache effectiveness; the following indexes are automatically calculated and tracked in daily and monthly trends:
 - %L1M, percentage of miss in the Level 1 cache;
 - RNI, Relative Nest Intensity;
 - CPI, cycles per instruction;
 - %PRB, percentage of problem state instructions over total instructions;
 - %PNL, percentage of penalty cycles over total cycles.



All the graphs are produced automatically without the need of any additional component. The product creates all graphs automatically in a standard format and maintains a simple, intuitive, and efficient navigation.

The graphics produced can be used immediately in word processors, spreadsheets, studies, reports, documentation, etc.



2 Architecture

EPVGRAPH is an application designed to add value to the EPV for z/OS product by providing a rich set of standard graphs that are correlated between themselves.

The architecture is composed of three components:

1. The **Data Collection Interface (DCI)** – it reads the EPV databases and prepares a series of text files used as input for the GPE component.
2. The **Graphics Production Engine (GPE)** – it reads the sequential files created by the DCI component and creates graphs.
3. The **User Graphics Engine (UGE)** – it reads user files and definitions and creates personalized user graphs.

The DCI, designed to acquire the data needed for EPVGRAPH, is a series of PERL programs that must be executed after EPV for z/OS has completed loading the data. A series of text files are created and then used as input to the GPE.

Both the GPE and the UGE are written in Perl.

The logical daily flow of the product is the following:

1. Extraction of the necessary data from the EPV for z/OS databases through the DCI;
2. Transfer of the text files produced from the previous step to the system hosting the GPE;
3. Production of the graphs and HTML pages through the GPE;
4. Transfer of the text files produced by the user to the system hosting the UGE (optional).
5. Production of the user graphs and HTML pages through the UGE (optional).



3 Installation

The EPV zParser product is a prerequisite of EPVGRAPH so from here on:

- “*\$EPVPATH*” should be substituted with the path to the EPV home folder created at EPV zParser installation;
- “*\$ProfileName*” should be substituted with the name of the user profile created at EPV zParser installation.

Warning: if you performed the EXPRESS customization as described in the “EPV Plus Installation and EXPRESS Customization” the EPVGRAPH product is already installed. No other action is required.

To install EPVGRAPH you need to perform the following steps:

1. Preparing products and password folders;
2. Verifying the HTML publishing folders.

3.1 Preparing products and password folders in Unix/Linux

Products and password folders should have already been prepared when installing the EPV zParser product which is a prerequisite to EPVGRAPH. Please refer to the EPV zParser documentation.

If for any reason you need to update the EPVGRAPH product, you have to copy the supplied /PRODUCTS/EPVZGRAPH_VXX folder (where XX stays for the version number of the product) from the EPV Installation CD to the \$EPVPATH/PRODUCTS folder.

If you need to update the license key please copy the LICENSE_EPVZGRAPH.EPV file in the \$EPVPATH/PASSWORD folder.

3.2 Verifying the HTML publishing folders

On the system that will host the HTML pages you have to verify that the following folders exist under \$EPVPATH/USERPROFILE/\$ProfileName/COMMON/HTM:

- **JAVA** folder containing JAVA code, scripts and sheets for the user interface;
- **IMG** folder containing images used by EPV (EPV logo, Microsoft Excel icon, etc.);
- **ZOSGRAPH** folder containing all the HTML pages produced (empty at installation);
- **ZOSTEXT** folder containing all the txt files produced by ZGRAPHLOAD.

The following HTML pages have also to exist under \$EPVPATH/USERPROFILE/\$ProfileName/COMMON/HTM:

START.HTML, HOME.HTML, ABOUT.HTML and CONTRIBUTORS.HTML.



If any of these folder and files doesn't not exist you must copy what is missing from the ./PRODUCTS/EPVUI folder of the EPV Installation CD to \$EPVPATH.



4 Customization (mandatory)

There are no mandatory steps to perform.



5 Scheduling

Scheduling the EPVGRAPH provided procedures allows you to process data and produce the HTML reports.

5.1 Collecting data once a day

When your environment is set to collect data once a day the whole EPV process is run by the ALLPHASES.sh procedure, created during EPV zParser installation, which is located in the \$EPVPATH/USERPROFILE/\$Profilename/COMMON/PROCS folder.

Inside that procedure you should find (or add if missing) a call to NIGHTBATCH.sh. NIGHTBATCH.sh calls the NIGHTBATCH_ZGRAPH.sh procedure which is located under the \$EPVPATH/USERPROFILE/\$Profilename/EPVZGRAPH/PROCS folder.

You can schedule this procedure to run daily at a specific hour using the standard task schedulers provided in Unix/Linux or you can start it manually during the day.

5.2 Collecting data in continuous mode using EPV agents

When your environment is set to collect data in a continuous mode using the EPV agents, the daily consolidation process is automatically run by the POSTZPARSER.sh procedure, created during the EPV zParser installation.

It is located in the \$EPVPATH/USERPROFILE/\$Profilename/COMMON/PROCS folder.

Inside that procedure you should find (or add if missing) a call to the NIGHTBATCH_ZGRAPH.sh procedure

This procedure is also located in the \$EPVPATH/USERPROFILE/\$Profilename/EPVZGRAPH/PROCS folder.



6 Manual Customization (optional)

In the following some optional steps, to fit specific user needs, are described.

6.1 Customizing the CONFIG.PL parameters

You can go into the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/CONFIG.PL` file and customize all the parameters and thresholds as desired.

You'll find a detailed description of each parameter and threshold in attachment A.

6.2 Setting SHIFTS

All Trend views provide HTML tables for different shifts. The member named **SHIFT.PL** located in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` library contains our default shift values which should be modified with your standard company values. When migrating to a new version, it is very important to control your shift definitions in the **SHIFT** member before loading new data.

6.3 Assigning a user name to each CEC serial number

You can assign a user name to each CEC serial number by customizing the **UEXGCEC.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder. The settings take place during the EPV HTML phase.

6.4 Excluding CECs from rolling 4-hour average analysis

You can exclude from WLC graphs the CEC you're not interested in (e.g. old machines). In order to do that, you've to customize the **UEXNO4AV.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder.

6.5 Changing the series colors in EPVGRAPH HTML pages

You can change the order of the series colors in the EPVGRAPH HTML pages by customizing the **COLORS.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder.

6.6 Changing the number of days to report daily

By default EPVGRAPH creates daily reports for the last two days.

If you want to change this behavior you have to customize the **RPTDATES.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder.



6.7 Setting your country holidays

To set your country holidays (e.g: Eeaster, Christmas) you have to customize the **UEXHOLY.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` library. In the array `@HOLIDAYS` you have to put all the country holidays, while in the array `@EVENT` you have to put the days you want to consider as holidays.

6.8 Modifying batch throughput

By default EPVGRPAH selects only the job batch assigned to the JOB workload when creating the throughput graphs; to modify the workload you can customize the **UEXJOBWK.PL** file provided in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder.

6.9 Including or excluding systems

To include or exclude some systems from EPVGRAPH you have to customize the **UEXSYS.PL** file in the `$EPVPATH/USERPROFILE/$Profilename/EPVZGRAPH/USEREXIT` folder.

6.10 Publishing on the IBM HTTP Server on zSeries

To publish the EPV HTML pages on the IBM HTTP Server on zSeries, the following steps are needed:

- a) Verify that all files and directories are in uppercase except for the `.class` suffix.
- b) Transfer all the files in the IMG directory in binary mode.
- c) Transfer the `.pdf` files in the ZOSDOC directory in binary mode.
- d) Transfer the following files in the JAVA directory in binary mode:

- *.JS
- *.class
- *.CSS

Transfer all the other files in ascii mode using the following quote command¹:

```
quote site sbdataconn=(IBM-1047,ISO8859-1)
```

IMPORTANT NOTE:

Beyond these general rules the correct FTP mode (binary or ASCII) depends on how customers set their web server. Depending on the **addtype** statement in **httpd.conf** the different file types have to be transferred in binary or ASCII mode.

Normally `.JS` files are used without converting to ebcdic (the conversion needs to be done again when transferring the `.JS` to the client; it takes some time and is not good for performance).

However if the `.JS` files are defined as ebcdic they need to be transferred in ascii.

¹ IBM-1047 has to be eventually substituted with the DefaultFsCp value set in `httpd.conf` (if different from IBM-1047).



The .CSS should in any case be defined using an **8bit addtype** and transferred in binary mode to work properly.

6.11 Exporting views in Microsoft Excel

It is possible to load an EPVGRAPH table to a Microsoft Excel spreadsheet. This functionality is based on ActiveX. In order to use it you need to set the security of your browser at a low level, or better insert the EPVGRAPH website in the list of the trusted sites.

6.12 Customizing the user interface

EPV pages are produced following HTML standards. It is possible to customize the appearance of pages (font, colours, etc.); for a detailed description of the styles used see the “EPV V10 User Interface” manual.

6.13 Customizing the Recovery procedure

By default EPV Graph for z/OS creates the daily reports for the last two days.

In cases you have activated the Recovery procedure in the EPV for z/OS product (see the EPV for z/OS Installation and Customization manual) and you want to create the daily Graph HTML pages older than two days without the need to modify manually the CONFIG.PL parameters (\$RPTEPVS, \$RPTEPVE), you

have to check the **\$RECOVER** and the **\$MAXRPTDAYS** parameters in the RECOVER.PL member located under the *\$EPVPATH/USERPROFILE/\$ProfileName/COMMON* folder.

The **\$RECOVER.PL** parameter which activates the Recovery procedure has to be setted to YES (DEFAULT: NO).

The **\$MAXRPTDAYS** parameter which indicates the maximum number of HTML reporting days (DEFAULT:2) has to cover the period that you want to maintain in the HTML reports.

The recovery procedure dynamically changes the above CONFIG.PL parameters internally in the program to permit the creation of the html pages without modifying the original CONFIG.PL parameters.

WARNING: The Recovery procedure doesn't change any of the trend configuration parameters. In order to maintain the recovered data in the HTML trend reports you have to verify that the trend parameters in the CONFIG.PL member cover the period you want to maintain.



7 Using the EPV Customization GUI

The EPV Customization GUI has been designed to run on a Windows system, however you can create a profile which can be used to run the EPV products on Unix or Linux.

The created profile can be modified manually after moving it to the Unix/Linux environment.

To start working with the EPV Customization GUI, you must enter the SETUP folder and run the EpvInst.exe program.

7.1 ADVANCED customization

You will get the windows in Figure 1 where you have the possibility to choose the customization mode.

To proceed with the ADVANCED customization you have to select ADVANCED and press CONTINUE (for the EXPRESS customization please refer to the “EPV Plus Installation and EXPRESS Customization in Windows” manual).

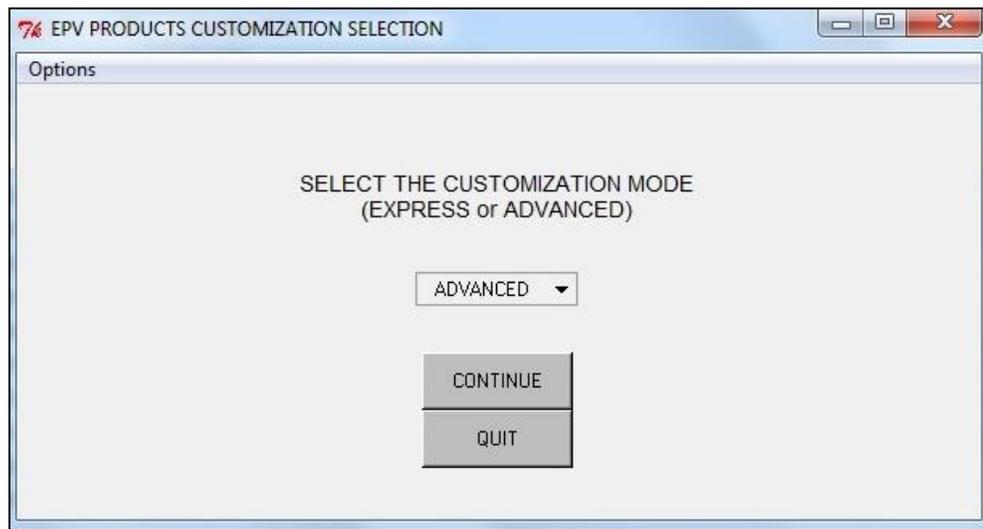


Figure 1

In the next window you are asked to select a profile.

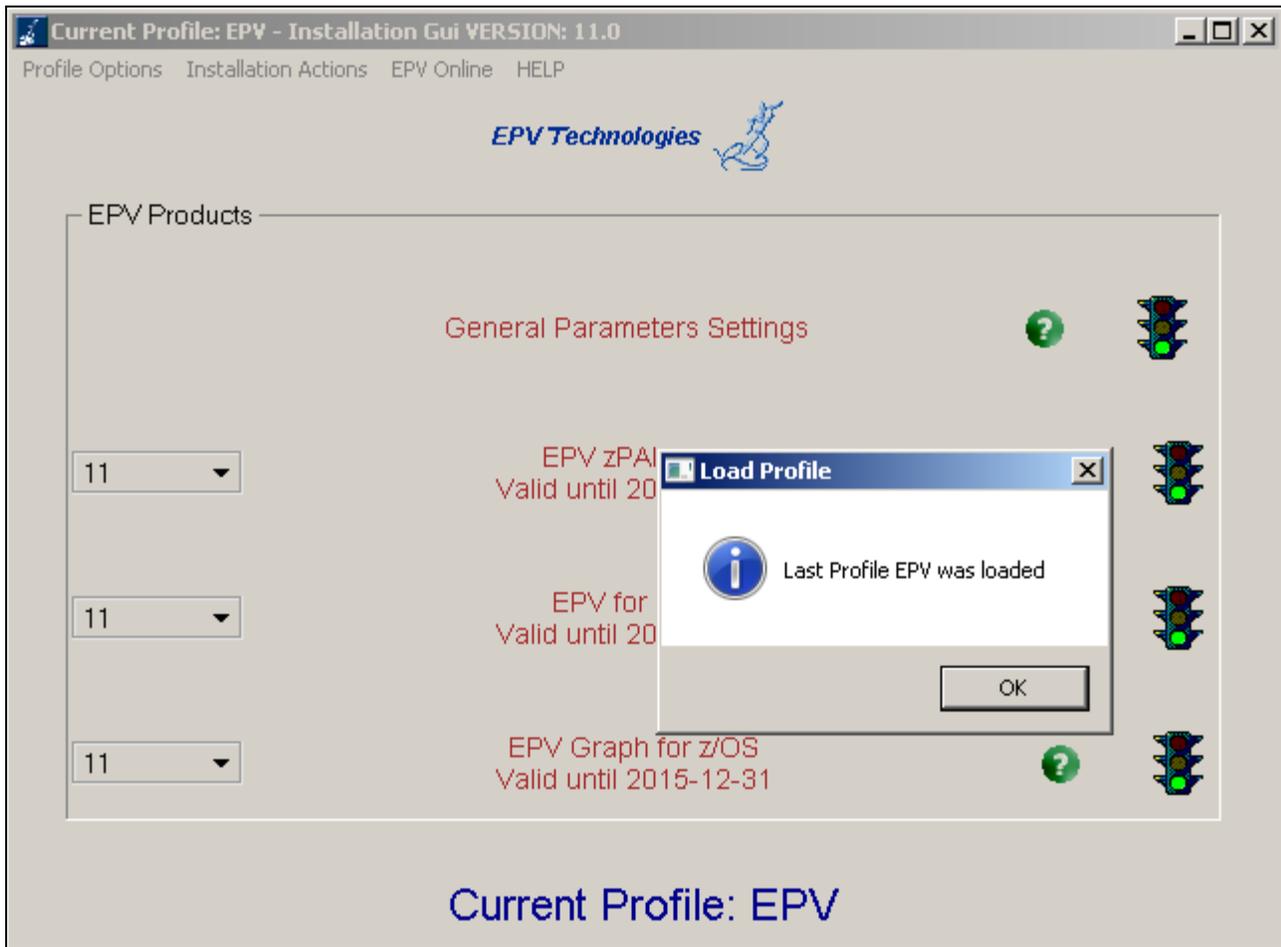


Figure 2

Click OK and choose an existing profile (normally created during the EPV zParser installation).



7.2 Customizing the product's parameters

If you have a valid license for the EPVGRAPH product you first need to select the correct version from the combo box located at the left side of the main panel and then choose the EPVGRAPH for z/OS entry.

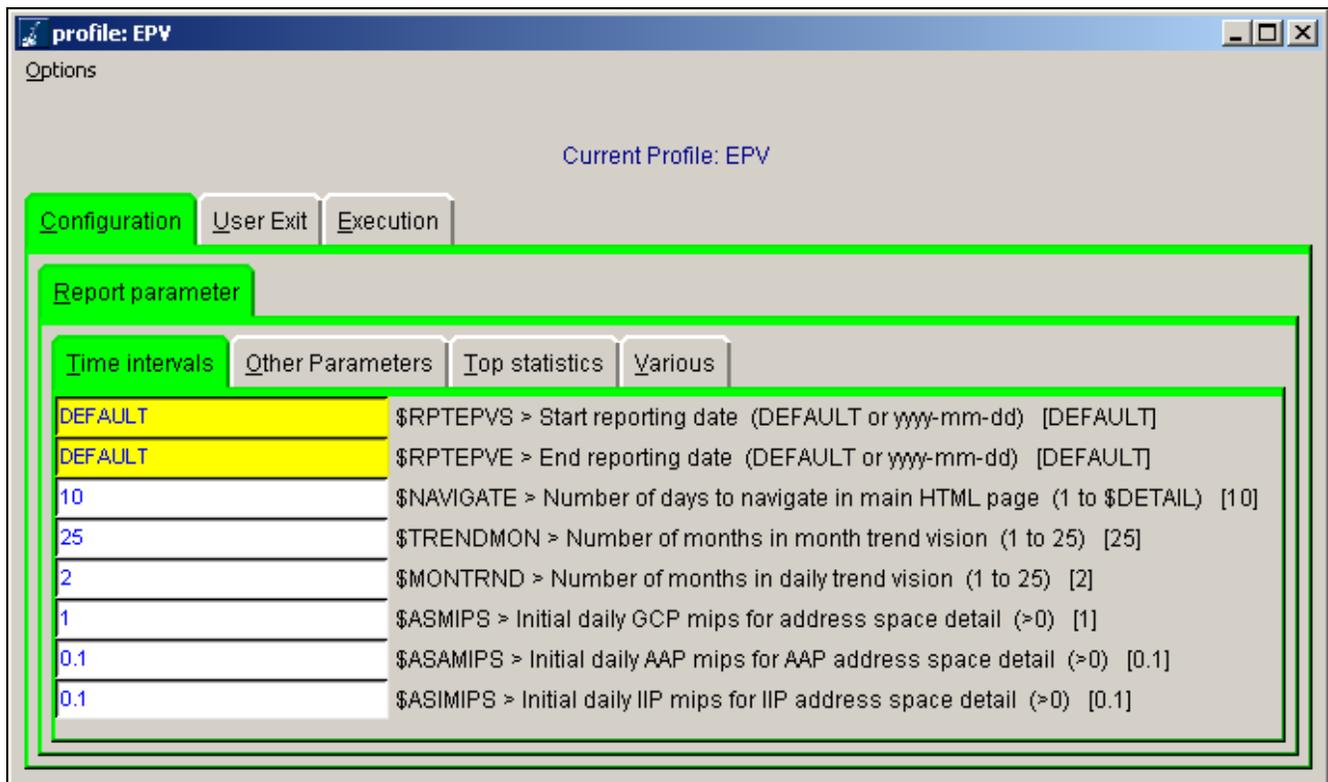


Figure 3

Through this panel you can customize all the parameters needed by EPVGRAPH, change the default settings for the user exits and so on. For a detailed description of all the parameters managed by this utility please see the Attachment A at the end of this manual.

Finally save the customizations you have done by clicking on the 'Options' menu and choosing the 'Save and Return' option; it will close the panel and bring you again to the main window.

7.3 Additional customization steps

After you moved your profile to a Unix or Linux system you have to perform the following additional customization steps:

- change the permission of all folders and files, inside it, to at least 755;
- remove all the CR (Carriage Return) inside all the .sh files in the profile; you can do that by using the standard DOS2UNIX utility; we suggest to run the following command on all the profile folder ($\$EPVPATH/USERPROFILE$): `find . -type f -name '*.sh' -exec dos2unix '{}' \.`



8 User defined graphs

In order to automatically create user graphs you have to perform the following steps:

- Preparing sequential text files for the UGE
- Preparing the menu structure for the UGE
- Setting graph parameters

8.1 Preparing sequential text files for the UGE

Sequential text files, which will be the input of the **USERGRAPH.PL** program, must be created by the user and must contain the following fields:

- **Interval:** The value plotted on the X axis.
- **Value:** The value plotted on the Y axis
- **Series:** The series value

Warning: The column separator must be the tab character.

Example:

| <i># Interval</i> | <i>Value</i> | <i>Series</i> |
|-------------------|--------------|---------------|
| 0 | 15.5 | LPAR1 |
| 1 | 17.0 | LPAR1 |
| | | |
| 23 | 10.0 | LPAR1 |
| 0 | 12.01 | LPAR2 |
| 1 | 28.2 | LPAR2 |
| ... | | |
| 23 | 15.81 | LPAR2 |

Figure 4

Warning: The user is responsible that all the interval periods are the same for all the series of data. As an example for a monthly year trend file, all the systems must have 12 entries. One for each month. For non-existing periods the user can provide null values.



8.2 Preparing the menu structure for the UGE

Once created, sequential files have to be transferred to the following folder:
\$EPVPATH/USERPROFILE/\$ProfileName/COMMON/HTM/ZOSTEXT/USERLIB.

Each file in this folder will originate an entry in the first level of the EPVGRAPH User defined menu having the same name as the file.

The user can also create additional subfolders under **USERLIB**. These subfolders will automatically generate additional levels in the User defined menu.

Each level of the EPVGRAPH menu will show an entry for each file name included in the folder.

In the following example two files have been created in the /CPU_UTILIZATION/2013-04-11 subfolders:

- *\$EPVPATH/USERPROFILE/\$ProfileName/COMMON/HTM/ZOSTEXT/USERLIB/CPU_UTILIZATION/2013-04-11/CPUD.TXT*
- *\$EPVPATH/USERPROFILE/\$ProfileName/COMMON/HTM/ZOSTEXT/USERLIB/CPU_UTILIZATION/2013-04-11/CPUM.TXT*

Figure 5 shows an example of the menu structure resulting from that.



Figure 5



8.3 Setting graph parameters

The UGE produces the user graphs based on the definitions you set in the **Userdef.PL** file located in the `USERPROFILE\${ProfileName}\EPVZGRAPH\USEREXIT` directory.

The **Userdef.PL** file should contain all the information needed for the user graphical engine to create the user graphs. The **gbuild** macro located in the **Userdef.PL** file creates one graph for each input data file provided. To produce more graphs you have to replicate and customize it.

The following parameters need to be customized for each **gbuild** macro

- **Path:** The path where the input data file is located;
- **InputTXT:** Name of the input data file;
- **AsseX:** The column position in the input data file where data for the X axis starts;
- **AsseY:** The column position in the input data file where data for the Y axis starts;
- **Group:** The column position in the input data file where the series variable starts;
- **FirstRow:** The number of the first line in the input data file where data is valid;
- **LableX:** X AXIS description;
- **LableY:** Y AXIS description;
- **Title1:** Table title;
- **Title2:** Image title;
- **Title3:** Sub Title;
- **TypeX:** Type variable x (CHAR/ INT);
- **TypeY:** Type variable y (CHAR/ INT);
- **TypeG:** Type variable group;
- **Y_Max:** Maximum value for the Y axis;
- **Y_Min:** Minimum value for the Y axis;
- **xl_b_skip:** Labels to skip on X axis;
- **yl_b_skip:** Labels to skip on Y axis;
- **Grid:** Flag for grid Y/N;
- **X_labels_vertical:** X label vertical or horizontal (0/1);
- **Width:** Graph width;
- **Height:** Graph height;

An example of the **gbuild** macro is provided in the following figure.



```
&gbuild(Path => '$SEVPATH/USERPROFILE/$ProfileName/COMMON/HTM/ZOSTEXT/USERLIB',  
InputTXT => 'file.txt',  
FirstRow => '1',  
LabelX => 'DATE',  
LabelY => 'CPU',  
Title1 => 'Graph 1',  
Title2 => 'Image 1',  
Title3 => 'Graph SubTitle',  
AsseX => '1',  
AsseY => '2',  
Group => '0',  
TypeX => 'CHAR',  
TypeY => 'INT',  
TypeG => 'CHAR',  
Y_Max => 640000,  
Y_Min => 0,  
xlb_skip => 0,  
y lb_skip => 0,  
Grid => 'N',  
Width => '750',  
x_labels_vertical => 0,  
eight => '550'  
);
```

Figure 6



9 Customer support

For any technical problem with or question about EPV for z/OS write an email to:

epv.support@epvtech.com

For any other issue about EPV for z/OS please write an email to:

epv.info@epvtech.com



Attachment A – Config parameters

| PARAM | LABEL | RANGE | DEFAULT |
|------------|--|--------------------------------------|------------------------|
| \$RPTEPVS | Start reporting date | DEFAULT (yesterday) or yyyy-mm-dd | DEFAULT (yesterday) |
| \$RPTEPVE | End reporting date | DEFAULT (today) or yyyy-mm-dd | DEFAULT (today) |
| \$NAVIGATE | Number of days to navigate in main HTML page | 1 to \$DETAIL | 10 |
| \$TRENDMON | Number of months in month trend vision | 1 to 25 | 25 |
| \$MONTRND | Number of months in daily trend vision | 1 to 25 | 2 |
| \$ASMIPS | Initial daily GCP mips for address space detail | >0 | 1 |
| \$ASAMIPS | Initial daily AAP mips for AAP address space detail | >0 | 0,1 |
| \$ASIMIPS | Initial daily IIP mips for IIP address space detail | >0 | 0,1 |
| \$ASCPUDET | set to Y if you want the Address Space CPU Detail | Y,N | Y |
| | | | |



| | | | |
|------------|---|---------|------------------|
| \$ASAAPDET | set to Y if you want the Address Space CPU AAP Detail | Y,N | Y |
| \$ASIIPDET | set to Y if you want the Address Space CPU IIP Detail | Y,N | Y |
| \$ALLIPLS | Set to 1 if you want to generate all monthly IPLS | 0 to 1 | 0 |
| \$WIMPDET | Minimum level of WLM importance to plot | 1 to 6 | 6 |
| \$GMTOP | Top mips job | 1 to 30 | 5 |
| \$GMTOPIN | Range of hours for top jobs | 0 to 23 | 9,10,11,12,15,16 |
| \$GMAPPTOP | Top App for trend | 1 to 30 | 10 |
| \$DSKTOPSG | Storage Group Top disk usage | 1 to 50 | 10 |
| \$TAPTOPOL | Pools Top page usage | 1 to 50 | 10 |
| \$DBDEC | Decimal point to store in EPV database | 1 to 3 | 3 |
| \$DEBUG | Debug detail level | 0 to 3 | 2 |
| \$FMT | number format (E/U/S) Europe/Usa/Swiss | E,U,S | E |



| | | | |
|-------------------|--|-----------|----|
| | | | |
| \$TopThrTrnf | Max occurrence on legend (x row) | 1 to 9 | 9 |
| \$GMT | Normalize user data to GMT | YES, NO | NO |
| \$GMTOFF | GMT offset to normalize the user hour | -12 to 12 | 0 |
| \$ZGRAPH_MAX_TASK | Number of tasks to process in parallel mode | 1 to 6 | 1 |



Related documentation

The following manuals complement the information provided in this manual:

- *EPV V10 User Interface*