

RMF's PM of OS/390 Java Edition Hand's on Lab

SHARE Technical Conference
Minneapolis
Session 2553

Jürgen Holtz and Oliver Benke
IBM Deutschland Entwicklung GmbH
Schönaicher Str 220
D-71032 Böblingen, Germany
E-mail: benke@de.ibm.com

Trademarks

The following terms trademarks of the IBM Corporation in the United States or other countries:

MVS

RMF

OS/390

z/OS

zSeries

z/Architecture

Tivoli is a registered trademark of Tivoli Systems Inc.

UNIX is a registered trademark of The Open Group.

Linux is a registered trademark of Linus Torvalds.

The following terms are trademarks of other companies:

Java

Windows

Windows 9x/NT

Netscape

Microsoft

Introduction

Performance Monitoring of OS/390, or **PM of OS/390** for short, is a workstation-based function that allows you to monitor the various resources in an OS/390* system and even in a sysplex in a very flexible way. PM of OS/390 provides filtered data in a graphical format within DataViews, allows you to manage multiple DataViews as one entity by providing the concept of a PerfDesk, and offers the means to make definitions such as DataViews and PerfDesks persistent for reuse at any time.

The new PM of OS/390 Java Edition is now provided on a product level and runs in contrast to its predecessor on any Windows** platform, such as Windows 9x/NT**, or on Linux. The dependency on OS/2* disappeared while at the same time appearance and functionality of PM of OS/390 have been greatly improved.

This tutorial introduces the Java Edition of PM of OS/390 and helps you to get started using PM of OS/390 in order to monitor the performance of your OS/390 system, your parallel sysplex, or many sysplexes in your enterprise. The following seven lessons guide you step by step through the terminology and usage of PM of OS/390. The examples shown in this document are based on the sysplex *My Sysplex* that has been created during the installation of PM of OS/390. You are encouraged to carry them out on your installation!

Lesson 1 - Overall Concept

All you need to know about the InstallShield-based installation of RMF PM is that you have to start the *pmwinv2.exe* file downloaded from our web page, execute it and type in your TCP/IP host name or IP number after the install program asked you for it.

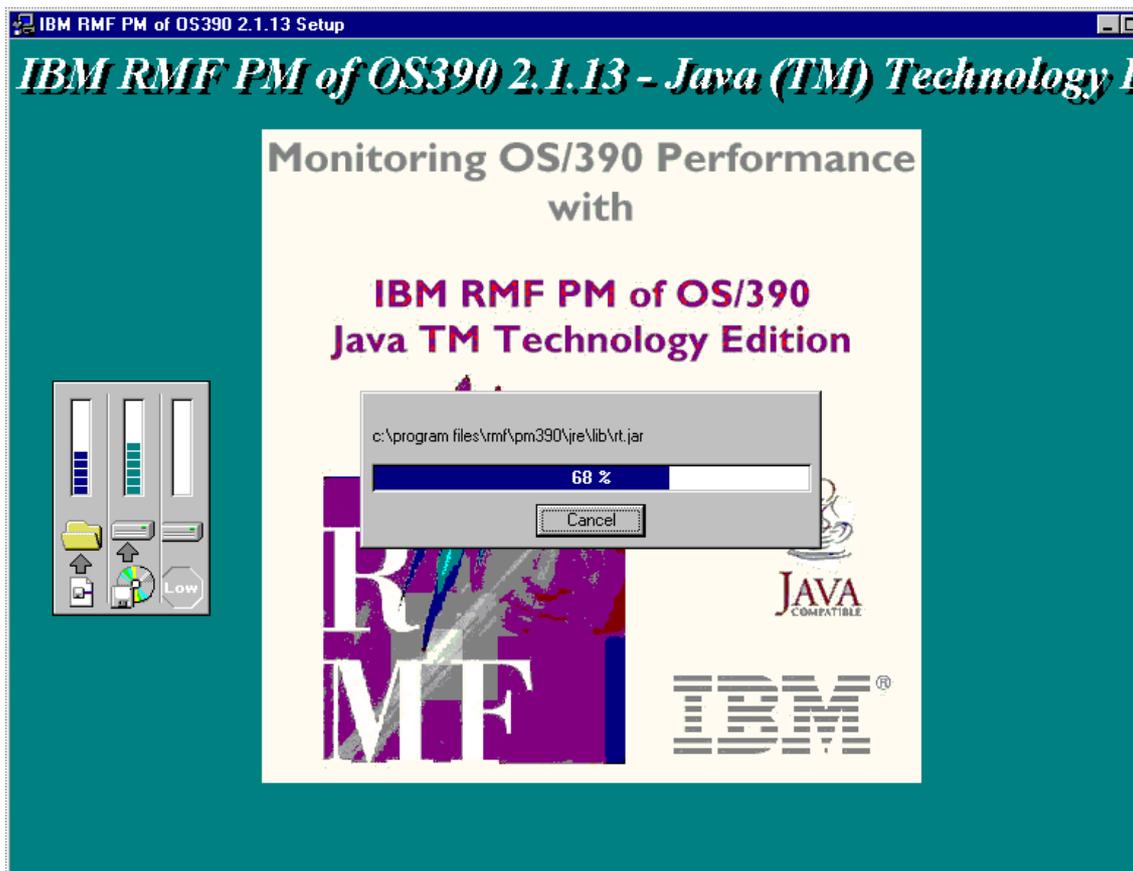


Figure 1

It is also assumed that the distributed data server, that you have uploaded to the host before, is started and active on one system within each sysplex that you want to monitor.

If these prerequisites are not fulfilled, please check the appendix for a quick reference of what is necessary to do or follow the detailed instruction guidelines shipped with this function.

During the installation you were prompted to specify a program folder for PM of OS/390. The default is "RMF Workstation Products". Navigate to it by clicking on the **Start**-button, then on **Programs**, and last click on **PM of OS/390** within the program folder created. When you start

SHARE Minneapolis

PM of OS/390 the first time a screen similar than the one below is shown:

You are prompted to logon to the sysplex *My Sysplex*. Logon with a valid User ID and Password, then click on **Ok** to get to the screen shown on the next page. Make sure that the User ID specified exists on the host whose IP address has been specified during installation. For details how to setup a sysplex, refer to the next chapter.

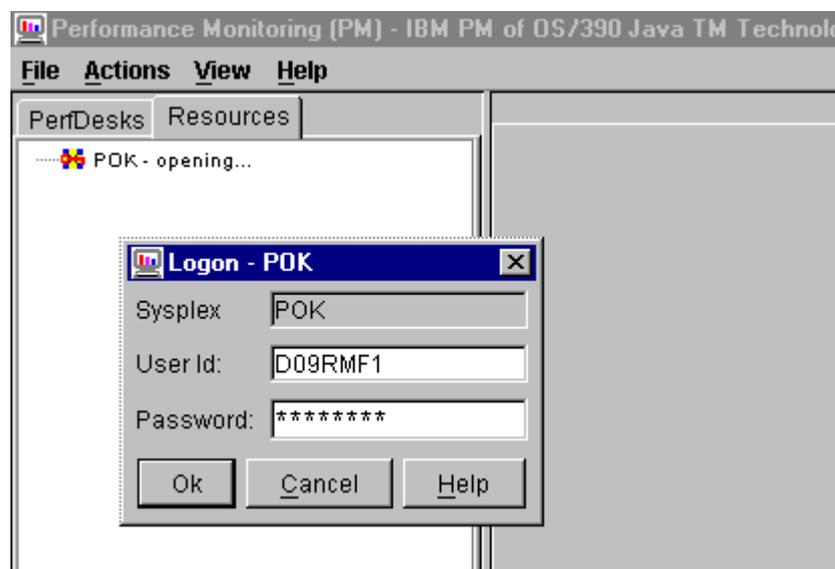


Figure 2

The window entitled PM of OS/390 is divided into two parts. On the left side a notebook with two tabs is shown - one for the PerfDesks, that means basically the various graphical views and their folders that you have defined, and the other one for the resources PM of OS/390 has access to. The first time you start PM of OS/390, the PerfDesk section only shows one item - the *Samples* PerfDesk folder. Later on you will see how to create your own PerfDesk folders to organize all your views based on functional or organizational requirements.

On the right side a notebook is provided with a tab for each PerfDesk that is currently open. For example, the figure on the next page shows you the *Sysplex - Overview* PerfDesk which consists of four DataViews. A DataView displays a bar chart of performance metrics -so called **series**- and controls the retrieval of the data from the host. This allows you to have different refresh intervals in the DataViews as well as to start or stop data sampling in one DataView without impacting other DataViews.

SHARE Minneapolis

Below the bar chart in a DataView, the legend classifies the series by color and name. If you add multiple metrics, each one is presented in a different color. The control panel below or left of each DataView allows you to start or stop data sampling for this DataView just by clicking on the colored status button  on the lower left corner. A green color indicates active sampling, while a red color indicates no data sampling.

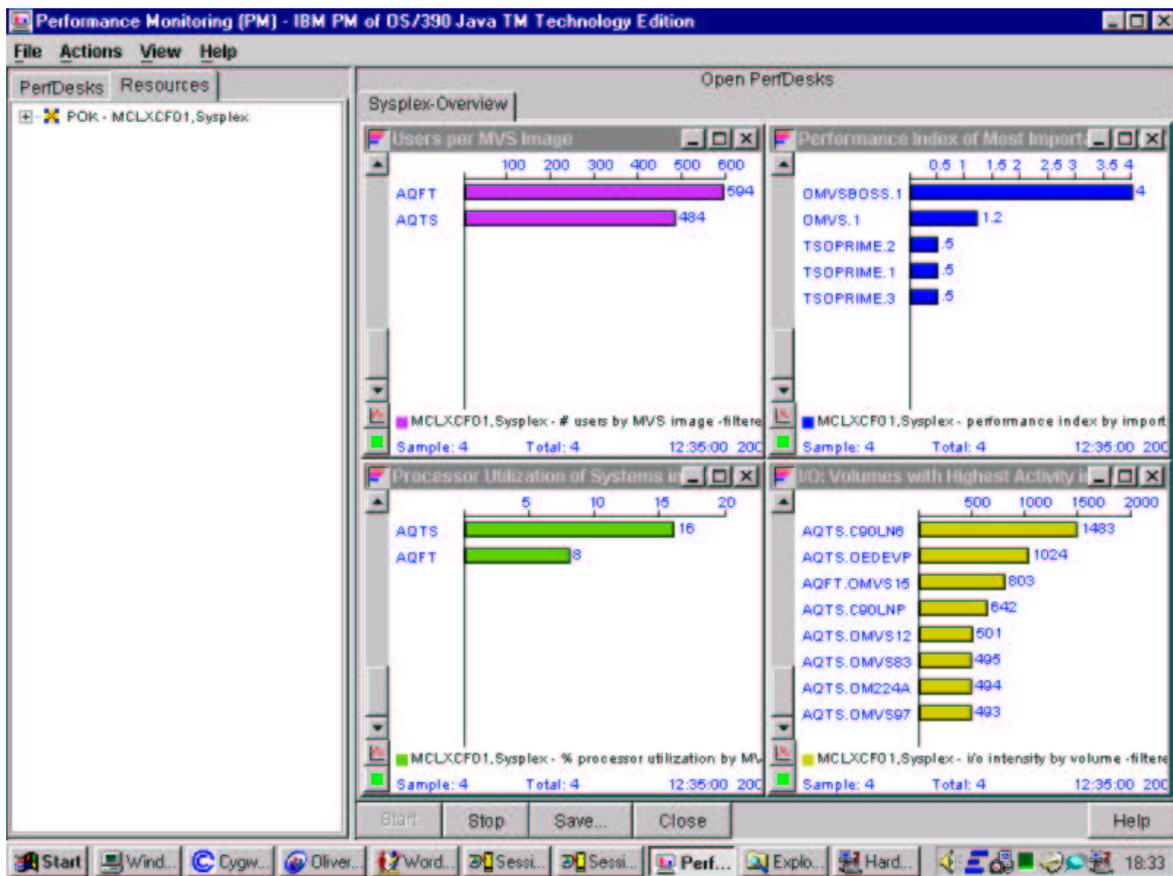


Figure 3

If you move the slider you can navigate back and forth on the time axis. The current sample as well as the related time stamp indicate the time you are currently looking at

SHARE Minneapolis

In the figure above, the DataView with the title *Processor utilization of systems in Sysplex* shows you the systems in that sysplex with the highest processor utilization. The DataView consists of one series that is called `% processor utilization by mvs image`. The sample has been collected at 12:35:00 and you see that system AQTS reports the highest CPU utilization (16%). A series like this is called *value-list* (or *list-valued counter*) since for one metric (here, processor utilization) a list of values is reported (here, by `mvs image`). As a general rule, value-list series can always be recognized by the preposition **by**.

Similarly, the DataView with the title *I/O: Volumes with highest activity in Sysplex* shows you the value-list `i/o intensity by volume1` for the seven volumes with the highest I/O intensity in the sysplex. As you look at the labels underneath the bars you see the system name that is placed in front of each volume label. This tells you from which system this I/O intensity has been measured. In the example above, the highest I/O intensity volume is `C90LN6` on AQTS with a value of 1483. The way the labels are created in general is the subject of lesson 3 where the resource naming concept is described.

Besides value-list series, PM of OS/390 also provides *single-value series* (or *single counter*). On a chart, a single-value series is shown for several intervals at once, that means that instead of certain instances on the x-axis the sample time is shown instead. So, single-value series can be seen as trends for a certain metric. Like with value-list series, it is possible to add multiple single-value series to one DataView, however, it is not possible to mix the type of series within a DataView.

¹ I/O intensity is defined as the product of I/O activity rate and I/O response time.

Lesson 3 - Sysplex Connectivity

PM of OS/390 is a client/server application. The client on the workstation is periodically asking for data provided by one or more servers on remote hosts. This lesson shows you how to define such remote hosts to PM of OS/390. Before that, however, I would like to show you a feature that allows you to change the look and feel of this application. When you select the **View** menu item on the action bar, a pop-up menu is opened. By default, the item **System Look and Feel** is selected which means that all the graphical controls are rendered in a way common to the system you are currently working at.

If you prefer to have the controls rendered in the typical Java look, select the menu item **Java Look and Feel**. Since the system look and feel dialog controls such as sliders or tabs have a consistent look and feel across all applications on your workstation, they are usually easier to recognize and use. With system settings, let us now start to define a remote host.

What role is a remote host playing? On the remote host the server part of PM of OS/390 has to be installed and running. This host can run in any WLM mode. If it happens to be a member in a sysplex, it also represents all the other systems in that sysplex at the same time. So, in other words, one definition is sufficient for a sysplex. Isolated systems need their own definition. In this sense they have to be understood as a special case of sysplex, consisting just of one system.

SHARE Minneapolis

To define a remote host², select the menu item **File** on the action bar and then click on **New**, then on **Sysplex...** . The dialog shown in Figure 4 is displayed in response to that

New Sysplex specifies a description of the system or the sysplex that appears as the first level entry with the resource tree shown on the **Resources** notebook page.

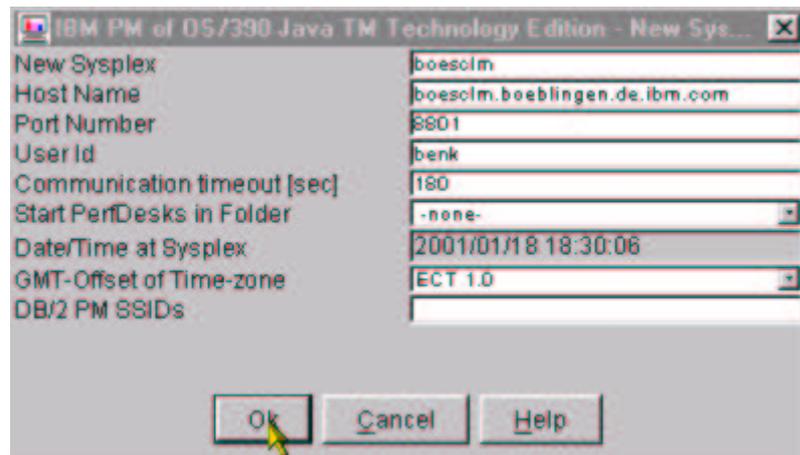


Figure 4

1. **Host Name** specifies the TCP/IP host name of a system where the distributed data server (DDS) runs. You can enter a symbolic name that can be resolved by a name-server or your hosts file, or an IP address like 9.164.182.251. Remember that this host represents the whole sysplex if it is a member in asysplex configuration.
2. **Port Number** specifies the TCP/IP port the distributed data server is listening to. Make sure that this port number matches the port passed to the data server during startup. The default port is set to 8801.
3. **User ID** specifies the TSO user id on behalf of which PM of OS/390 opens asysplex by default. This user id is related to the system represented by the host name specified above. Of course, any other valid user id can be entered on the logon panel when opening a sysplex.
4. You can choose whether or not a PerfDesk should be started immediately after logon to the sysplex by selecting a PerfDesk Folder in the combo box **Start PerfDesks in Folder**. You have seen the *Sysplex Overview* PerfDesk as an example which is started when *My Sysplex* is opened. This PerfDesk is an item in the *Samples* PerfDesk Folder.
5. In order to display the correct time values on the data views, you finally have to adjust the host GMT offset for your time zone. The offset should be preset with the correct value that

² PM of OS/390 has already created a remote host for you during installation. It is *My Sysplex* that you saw in Lesson 1. For this definition, the default Port Number 8801 is assumed. You can change the settings for *My Sysplex* at any time by selecting in a row first the menu item **File**, then **Change Settings**, and last **Sysplex...** .

SHARE Minneapolis

corresponds to your locale. For remote domains, however, you need to check whether the offset is correct and change it if not.

You finish the definition by clicking on button **Ok**. Note, that the **Ok** button can be selected only when for each of the items 1 to 4 some input has been specified. If you take a look at the **Resources** notebook page you will notice the new sysplex.

By the way, in the rare case where you would like to get rid of a sysplex definition, click the right mouse button on the respective item on the **Resources** notebook page and select menu item **Delete**.

Lesson 3 - Resources

In this lesson you will learn to navigate through the resources tree that is shown when you click on the **Resources** notebook tab. PM of OS/390 has a certain concept of what resources are and how they are organized in the sysplex.

A resource is any facility of a computing system or an operating system required by a job or task. The resources handled by PM of OS/390 are based on a data model that defines the hierarchy of these resources. The top resource is the Sysplex. It is the root of the resource hierarchy tree. Below the Sysplex in the hierarchy are the resource types Coupling Facility and Image. Children of Coupling facility are resources of type CFStructure while the children of Image are resources like I/O-Subsystem, Processor, or Storage.

In PM of OS/390, resources are named according to the following convention: each resource is identified by its

1. upper level qualifier
2. resource name
3. resource type

The resource name is written in the following style:

<upper level qualifier>,<resource name>,<resource type>

Examples of resource descriptors are:

SYS2,DATA01,Volume

This descriptor describes a volume named DATA01 that is mounted in MVS image SYS2.

SYS1,Image

This descriptor represents the resource type Image and its SMF name SYS1. The upper level qualifier isn't specified here since each MVS image is unique in a sysplex.

For some resources, only one instance exists within its containing resource. The convention for these resources is that the resource name is simply omitted. For example, there is only one SQA or only one processor³ in an MVS image so this resource would be identified by:

SYS2,SQA or by SYS2,Processor

Before you can look at the resources in your sysplex, you have to open a connection to that sysplex. *My Sysplex* is open already. You can recognize this by the [+] sign in front of the sysplex description. Figure 7 shows the resource configuration of the *My Sysplex*. To open any other sysplex, go to the **Resources** notebook page, click on the sysplex description with the right mouse button, and select **Open**.

The resources are ordered hierarchically in a configuration tree. The nodes of this tree can be expanded interactively, starting at the root node. The root node represents the top level resource

³ RMF does not differentiate between single processors but only looks at the total processing capacity in an SMP environment.

SHARE Minneapolis

Sysplex. Even if the remote host associated with a description is an isolated system, it is still modeled as a sysplex.

The links between the nodes of the resource hierarchy represent relationships. However, the semantics of these relationships are not uniform for the whole tree. In most cases, the relationship is 'contained', which means that the resource represented by the higher level node is an aggregate of the resources represented by its lower level nodes.

The [+] sign before the sysplex description or before any resource indicates that you can expand it into its contained resources. Just click on the [+] beside the sysplex description and the contained resources should be listed indented below. Vice versa, if a [-] sign is shown beside a resource you can click on it to collapse its contained resources. A resource without a [+] before its name has no contained resources and thus cannot be expanded any further.

For example, go to *My Sysplex* and click on [+] beside "My Sysplex - MCLXCF01 , Sysplex". The configuration tree for the resource *My Sysplex* is expanded and you see the MVS images as well as the coupling facilities configured in that sysplex. When you click on [+] beside "AQTs , Image" in turn, the resource tree of system AQTs is expanded and all resources contained in that system are shown.

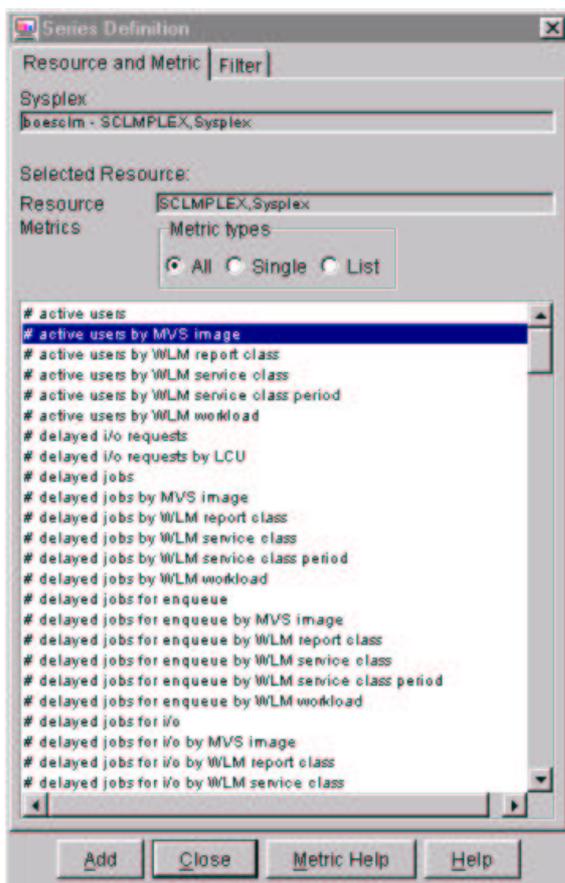


Figure 5

Lesson 4 - Creating DataViews

Now you should be familiar with the most important concepts that belong to PM of OS/390. This is the point in time where you can start to create your own DataViews. The following case study assumes that you have a sysplex and that one goal of operations is to monitor the activity on the sysplex and its contained systems continuously.

How do we start with the definition of such a DataView? The first consideration is to identify the metrics that you need in order to tell whether a system has high or low activity. There are several possibilities how to define what 'activity' really means. In the following example, 'activity' is interpreted as

- Number of active jobs in a system
- Number of delayed jobs in a system

You might have your own understanding of 'activity'. That is fine. Once we have gone through this example, you can try to create your special 'Activity DataView'.

Depending on the number and type of the metrics, you also have to choose more than one DataView for that purpose. But this is not really a problem since PM of OS/390 helps you to manage a set of DataViews together with the help of a PerfDesk. As far as PerfDesks are concerned, you will learn more in the next lesson.

The second consideration is whether you would like these metrics grouped together in a value-list, in other words *by Image*, or whether you just want to create one or more standalone DataViews for a certain system. Since our case study assumes that we want to monitor the activity in a sysplex, the first option is the better choice.

You select the sysplex resource on the Resources notebook page and click the right mouse button for the context menu.

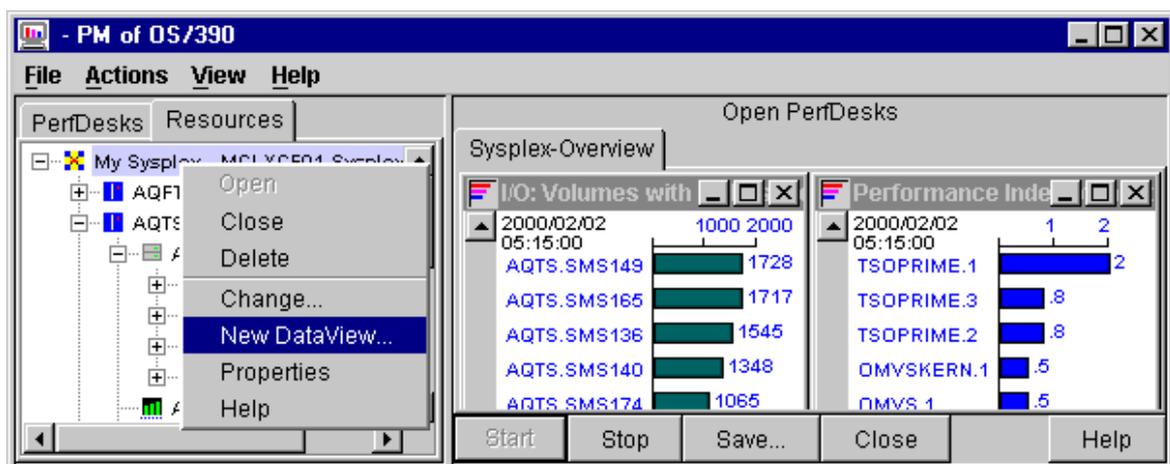
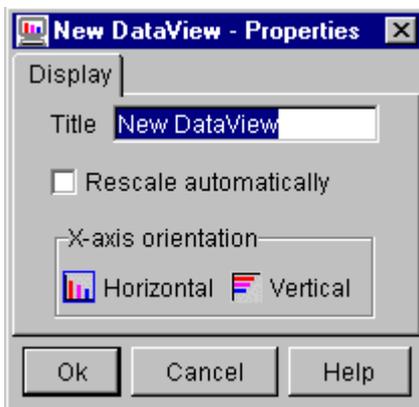


Figure 6

Note: you can also display properties defined for some resources that are not available as metrics in DataViews. The resource type Sysplex, for example, has the active service definition and service policy together with their installation or activation date and time as attributes. Properties are normally static information like CPU type while metrics are measured information like PIs.

The *New DataView...* menu item gets you to the dialog shown on figure 6. Here you can enter the title of the DataView. The default title “New DataView” is selected and can be changed just by typing over it. Also, when the DataView you intend to create is a value-list it might make sense to arrange the bars in vertical rather than horizontal orientation. This is the case, for example, when looking at resources with long names, such as “SYS2,CICSD1,Volume.” In our case study, the resource names are just the system names themselves, so the default (horizontal orientation) is appropriate.



The check box *Rescale automatically* -when checked- directs the DataView to adjust the scale on the y-axis according to the highest value in that DataView. While it allows you to get a higher resolution when several values are quite close together it alerts you on the other side since even low values appear over dimensioned on the DataView.

Figure 7

Now, click on *Ok* to get to the actual Series Definition dialog. You can choose from a rich set of metrics that are related to the resource selected.

If you have just created a new DataView, the dialog shows you all possible metrics for the resource, value-list metrics as well as single value metrics. Earlier I have mentioned that you cannot combine both kind of metrics in one DataView. To reduce the number of metrics shown, select the *List*-button in the Metric Types group in order to only show value-list metrics. From this list in turn select:

1. # active users by mvs image. The *Add*-button is activated once you have selected a metric.
2. Click on *Add*. Notice that the number of metrics in the list box is now reduced showing only those items that belong to the same list type, here by Image.
3. # delayed jobs by mvs image.
4. Click again on *Add*, then *Close*.

The DataView that has been created is not yet started. To start data sampling, click on the red button in the lower left corner of the DataView.

SHARE Minneapolis

You might think that the number of active and delayed users by MVS system is not sufficient for your purpose. It would be good, if each system's processor utilization is provided as well. There are two choices:

The first alternative is to create a second DataView in the same way as the previous one. If the proportions of the metrics "processor utilization" and "# users/jobs" in your installation is about the same, the second alternative is to add the new series to the existing DataView. However, I recommend to separate different kind of metrics in different DataViews.

Nevertheless, for the purpose of demonstration, here is the second alternative: move the mouse cursor somewhere to the background of the bar chart and click the right mouse button. The DataView context menu shown in a popup window allows you to add any number of series.

To add the processor utilization, click on **Add Series...** , specify

```
% processor utilization by mvs image
```

in the *Series Definition* dialog, click on **Add**, and then on **Close**.

The new series is immediately added to the DataView and by clicking on the red button, new values are sampled.

The context menu can also be used to change the properties of a DataView (title, orientation, and automatic rescaling) at any time. Similarly, you can clear the whole DataView just by selecting the **Remove Values...** menu item. If you want to remove a specific series from the DataView, place the mouse cursor on one of the bars representing that series, click the right mouse button and then **Remove Series**.

Now, please try to add a DataView by yourself.

SHARE Minneapolis

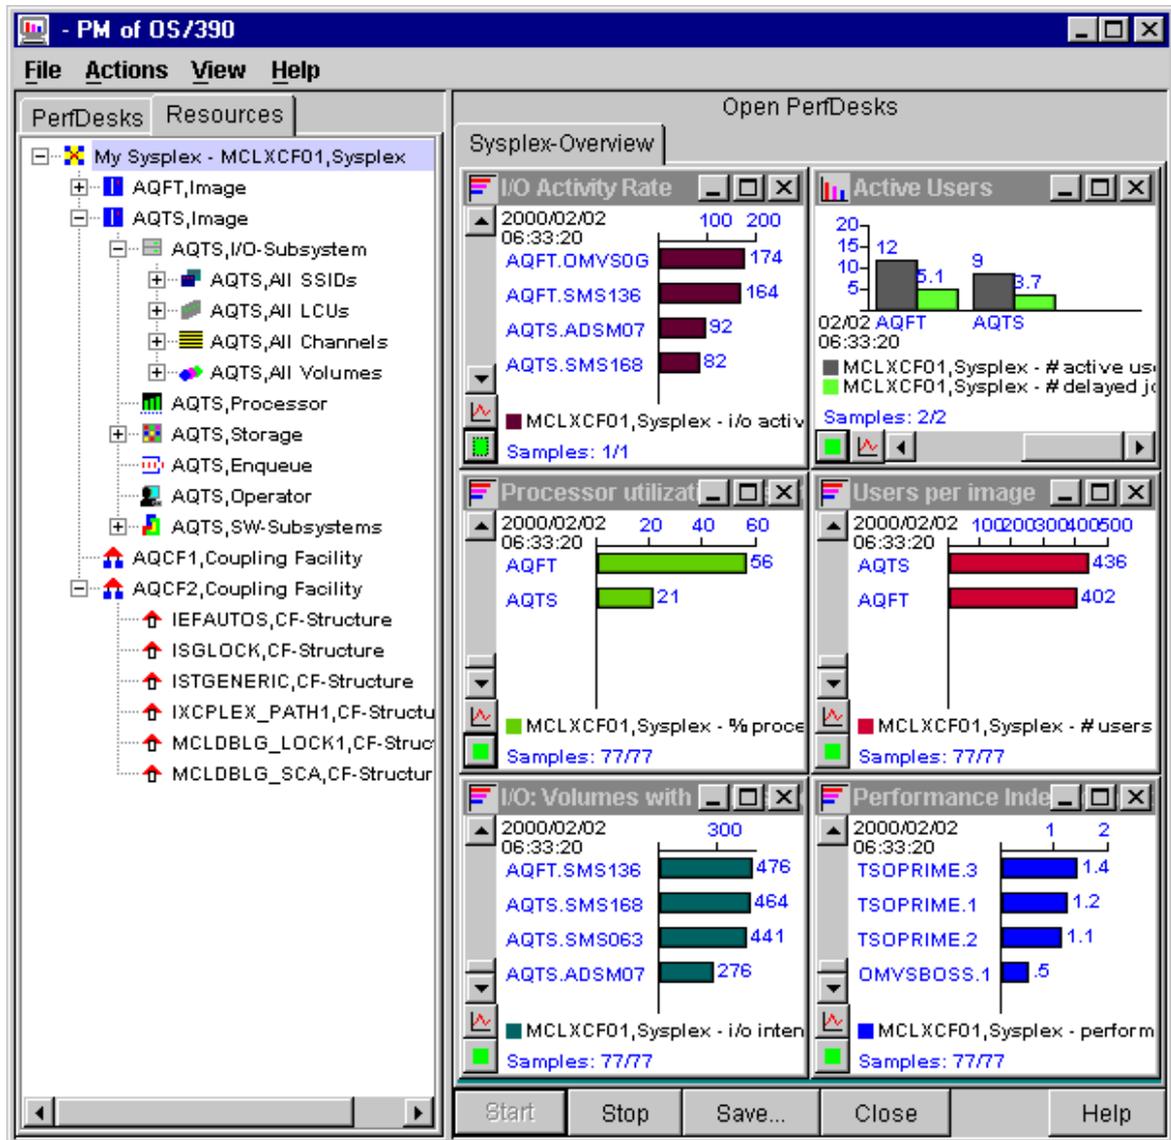


Figure 8

The new DataView should monitor the five volumes with the highest I/O activity rate in the sysplex.

Lesson 5 - Creating PerfDesks

We have seen a PerfDesk as a collection of DataViews so far. But PerfDesks are really more than that. With PerfDesks it is possible to start or stop data sampling in all the DataViews at once. A PerfDesk also functions as a container for its contained DataViews. Thus, only the DataViews of the selected PerfDesk are shown while other PerfDesks and their views remain in the background. To manage PerfDesks, PM of OS/390 provides a separate notebook page called *PerfDesks*.

When you click on the *PerfDesks* notebook tab, a tree-view is shown that displays all existing PerfDesk Folders. So far, only one exists, namely the *Samples* folder. If you click on the [+] sign, the PerfDesks within that folder are shown. Since you only added new DataViews to an existing PerfDesk, only one PerfDesk is listed - *Sysplex - Overview*. Again, by clicking on the [+] sign, the PerfDesk is expanded and its DataViews are listed.

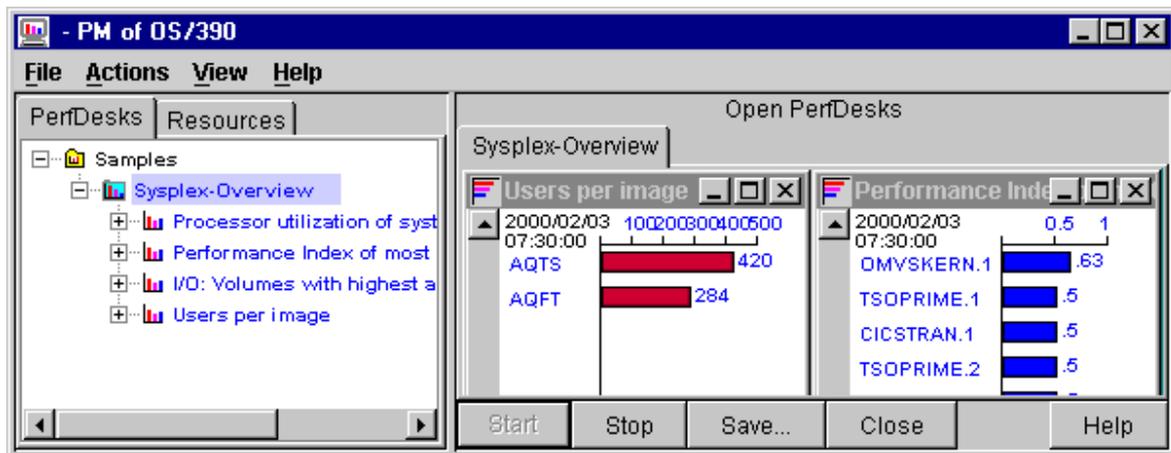


Figure 9

No te, the new DataViews that we have created in the last lesson are not yet listed underneath the current PerfDesk. Only when the PerfDesk is saved, DataViews are persistently added to it. To save all the DataViews that you have created so far in a new PerfDesk select the *Save...*-button on the bottom of the Open PerfDesks notebook. A dialog is shown that asks you for the location where this PerfDesk should be saved. Locations are **PerfDesk Folders** represented by a folder / sub directory in your system. If folders already exist they are shown in the list. The first time you do this, only *Samples* will be listed.

If you want to override the current PerfDesk, just click *Ok* now. If you want to create a new PerfDesk in the same folder, give it a name in the *Save PerfDesk* field.

For organizational purposes it also makes sense to create a new PerfDesk Folder. You might keep folders for example to separate PerfDesks for different sysplexes or systems or in order to manage different kind of views To create a new folder, click on the *Create folder...*-button and specify a name for this folder.

SHARE Minneapolis

If you would like to change the PerfDesk name later on you can do so just by getting the PerfDesk's context menu when you click somewhere in the blank area of a PerfDesk and then select ***Change Description...***. Alternatively, select the PerfDesk you would like to rename in the PerfDesk notebook, then in a row click on ***File***, on ***Change Settings...***, and last on ***PerfDesks...***. In both cases you are prompted for the new name and finish with ***Ok***.

On the ***PerfDesks*** notebook page introduced above, you can now see that the tree-view has changed according to the specifications you have made. The context menu for the entries within this tree-view allows you to open PerfDesks or to delete them. PerfDesk folders can be removed in a similar way. Beware because in this case all contained PerfDesks are removed as well.

When you open a PerfDesk, the PerfDesk is added to the Open PerfDesks notebook but the DataViews do not automatically start sampling. To start sampling on all DataViews at once, click the ***Start***-button on the bottom of the notebook. Notice, that the status button of each DataView switches from red to green in this case. Vice versa, data sampling is stopped for all DataViews when you click the ***Stop***-button. Last but not least, to close a PerfDesk use the ***Close***-button.

Finally, you can instruct PM of OS/390 to start one or a set of PerfDesks automatically whenever a sysplex is opened. In lesson two, which dealt with the definition of the remote hosts, I briefly mentioned this mechanism. Recall, that when asysplex is defined, you can specify whether or not PerfDesks in a certain folder are started automatically. So, if you would like to start a number of PerfDesks for any sysplex or standalone system, go to the ***File*** action bar item and in a row click first on ***Change Settings...*** then select ***Sysplex...***. The ***Change Sysplex*** dialog now displayed looks familiar to the one you have seen in lesson two. Select the folder containing all the "auto-" PerfDesks in the group box ***Start PerfDesks in Folder***.

The automatic startup of PerfDesks is effective only when you start PM of OS/390. In this case also the data sampling of the DataViews is started. However, PerfDesks are not started automatically when you open asysplex from the ***Resources*** notebook page. This implies when you have specified a startup PerfDesk Folder in a sysplex definition as described above, this change is effective only after the next invocation of PM of OS/390.

Lesson 6 - Performance Analysis

What can you do if, for example, your activityPerfDesk shows that unacceptable delays occur on one system or the activity rate of a volume is so high that this might be an indication for a severe problem in the near future? In such a case you need to acquire more information. You need to know what the main delays are that cause the overall system delay Or you are interested in the jobs using the device with the high activity rate. PM of OS/390 helps you to find the answers for to these questions.

The most interesting function in PM of OS/390 is its capability to navigate through the large number of resources or users that are using the resources or that are delayed by these resources in an intelligent manner. With this **Analysis** function you can **drill down** a problem. If your sysplex activity PerfDesk showed a very high activity rate, you can ask PM of OS/390 to show the jobs that are using the associated volume most, i.e. the jobs that are responsible for the activity rate. On the other side, PM of OS/390 would show you the jobs that are impacted most because they need to wait for that volume.

The analysis also provides you **context** information about the resource you are looking at. The context of a volume, for example, is among other things the breakdown of response time into IOS queue time, connect time, disconnect time, and pending time.

The third type of analysis is a **roll-up** to broaden the current view. Rather than drill down the problem, such an analysis step allows you to also look at the environment the problem is located in. All three types together are the means for you to locate the performance bottleneck very fast and to act accordingly before the problem gets severe.

To start the analysis, place the mouse cursor somewhere on the bar showing an interesting metric, for example a bar indicating a high device activity rate and click the right mouse button. The series' context menu is shown and the first item you see is called **Analysis...** Click on it and the *PM of OS/390 Analysis in <sysplex>* dialog is displayed. To make it simple you can also double-click on the interesting bar to get to the dialog.

The top part of this dialog tells you the resource you are currently looking at. In the example above, we have a DataView *Sysplex I/O Activity* providing the volumes with the highest activity rate in sysplex MCLXCF01, so the current resource type is Sysplex and the name of the sysplex is MCLXCF01.

The metric we have selected is shown below - also what the current value is and to which resource it belongs. As you can see, it is VolumeSMS168 on system AQTS.

In the middle and bottom part of the dialog you can select

1. the sample time, i.e. either the time for which you selected the analysis or the time for which the most recent sample data is available
2. the analysis type, i.e. what you want to look at next.
3. whether or not PerfDesks opened in a previous analysis step should be closed. This is especially useful if you repeat the analysis several times and to avoid having too many opened PerfDesks at a time. Therefore the default is to close previous PerfDesks.

SHARE Minneapolis

For example, to see what the average response time is, how long the IOSQ time is, or to see other characteristic metrics for that volume, take the default selection

AQTS,SMS168,Volume - Context

and click on ***Ok***. The analysis type

AQTS,SMS168,Volume - Context by Job

provides you the jobs on system AQTS that use volume SMS168 most in that sample plus the jobs on system AQTS that are delayed most since someone else is using the volume.

The analysis type differs from resource to resource. PM of OS/390 only offers analysis types that are related to the resource you have selected.

SHARE Minneapolis

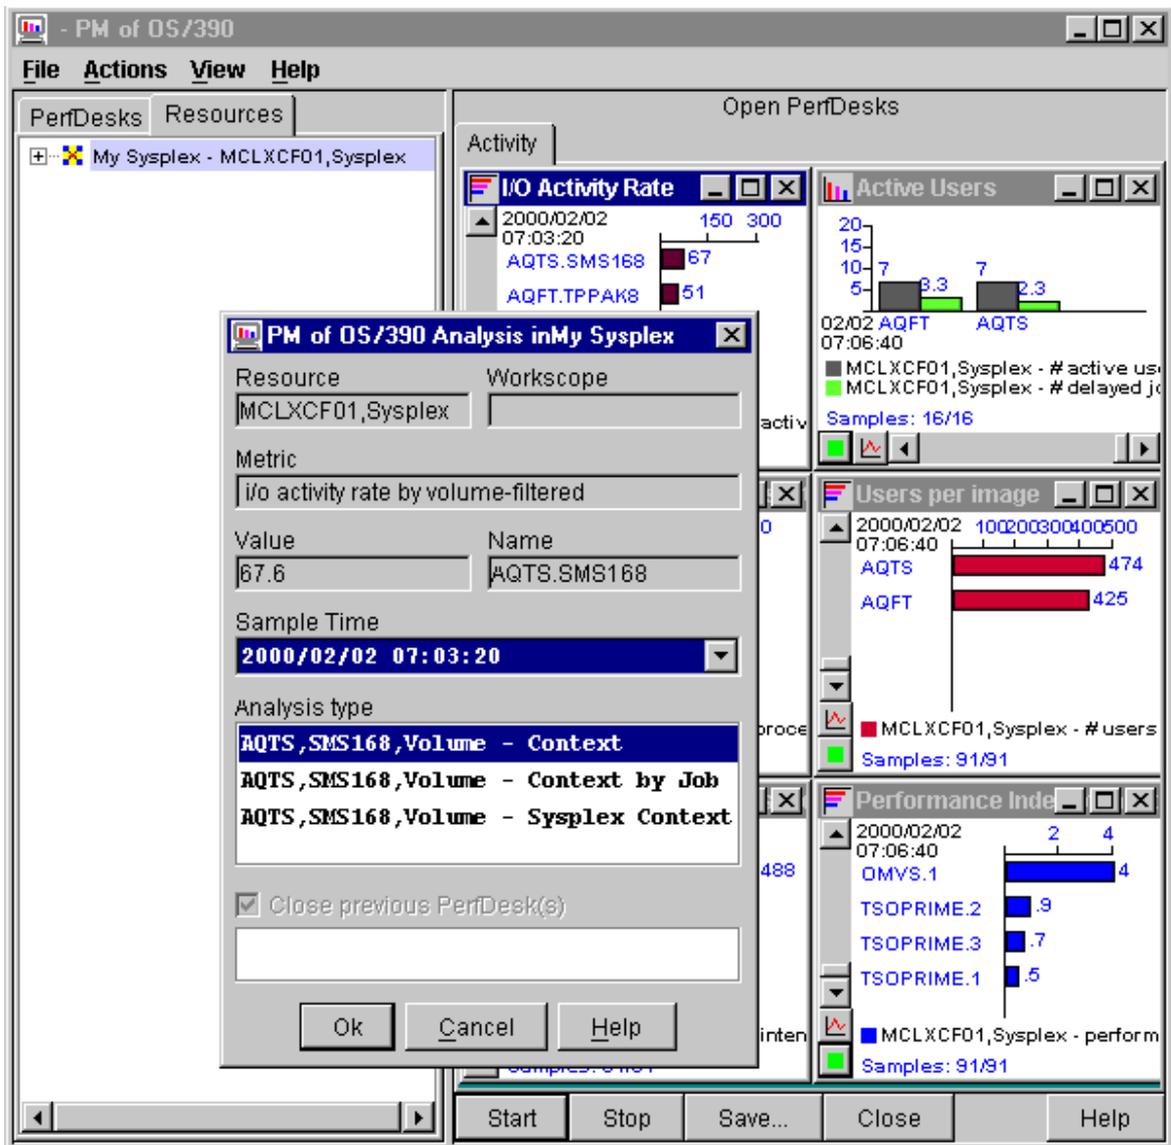


Figure 10

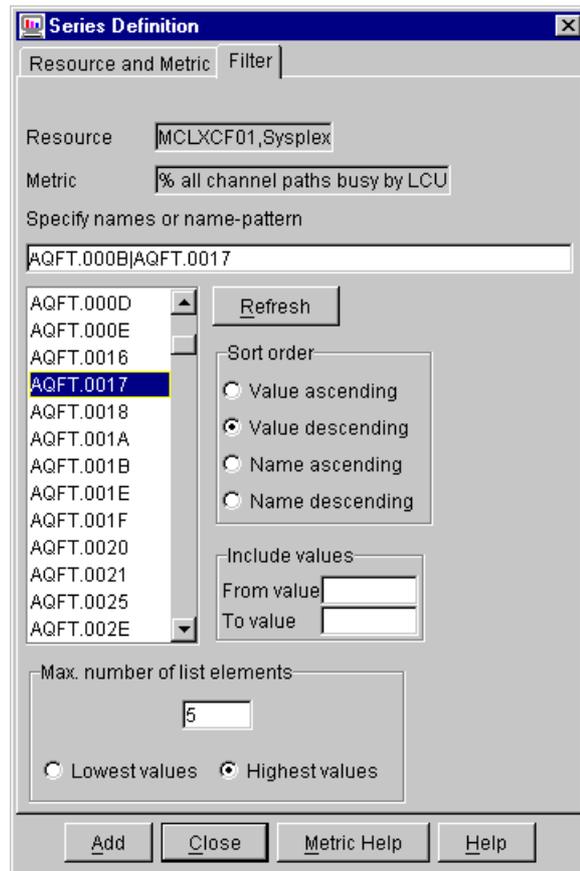


Figure 11

Lesson 7 - Features

In the last lesson of this tutorial, I would like to introduce some interesting goodies that help you to make your job more efficient.

Sorting and Filtering

The first thing to mention is the powerful sorting and filtering capabilities for list-value DataViews. So far we only made use of the existing default handling. The defaults take care that list-values are sorted in descending order, that means the most interesting ones are shown first. The default number of items shown is five. Nevertheless you might need to change the default, for example to concentrate only on a list of jobs or volumes that obey a specific naming convention. Also the number of five items might be too restrictive for you and you would like to see 10 items in a value-listDataView.

You can modify the default filter only when you create a DataView or when you add a new series to an existing DataView. Before you click on button **Add** to add the series to the DataView, click on the **Filter** notebook tab to change the sorting and the filters according to your needs. A dialog such as that shown in figure 12 comes up:

SHARE Minneapolis

As this example shows, you can reduce the number of logical control units(LCUs) monitored by specifying a name-pattern.

Clicking on **Refresh** causes PM of OS/390 to get the current list of contained resources from the host, in this case the list of LCUs in sysplex MCLXCF01. It may take a while when you do this but since this list is unlikely to change very often, you probably have to do it only once.

You can determine the sort order, whether the items in the DataView should be sorted by the value or maybe by the name, whether the list should be presented in ascending order or in descending order.

You can also specify boundaries to only provide items in the DataView that fall within the bounds of the specified thresholds. This keeps uninteresting LCUs out of the DataView and lets you concentrate on the important ones only.

Once the thresholds are set you can finally choose how many elements from the filtered list you would like to see and whether the low-value range or the high-value range of list elements is shown on the DataView.

Historical Data Collection

You can get historical data as long as it is available for RMF Monitor III data gatherer.

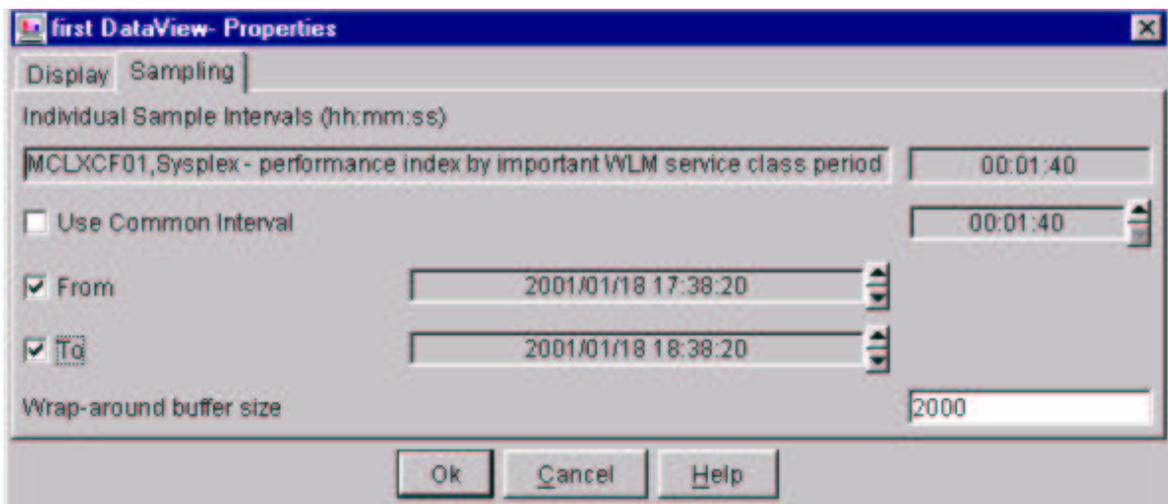


Figure 12

Please select the DataView for which you would like to get historical data, press the left mouse button, select "Properties" and select "Sampling".

SHARE Minneapolis

Activate the “From” control field and choose the starting time for the historical data on the right. You can optionally specify an end-time by activating the “To” control field. After you’ve restarted the DataView, RMF PM starts to get the historical data from the Distributed Data Server DDS. Please note that the DDS generates the data based on the “raw” performance data into the Monitor III gatherer which are “cooked” using the Monitor III reporter. This can be a relatively expensive operation on the whole sysplex (depending on which data you’d like).

Plotting and Saving Series

A very useful function is the capability to save the values and names of a series within a *.WK1-spreadsheet file. This allows you to export the data into your favorite spreadsheet application for performance analysis or capacity planning over a longer period in time

Go to the DataView that contains the series that you would like to save and click on the  button to invoke the *Save Plot of Selected Series* dialog. Figure 13 is an example of what you get once you have invoked this dialog.

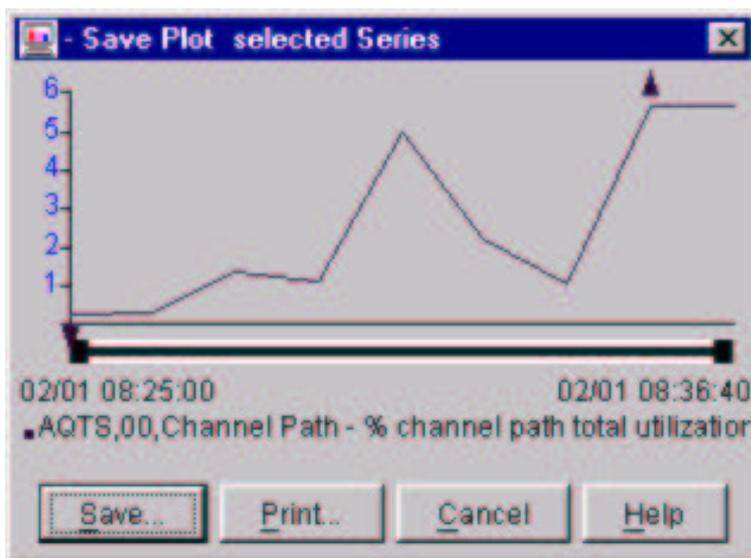


Figure 13

The chart on top of figure 13 represents the course of one of the series shown in the DataView over the time.

In the case of a value-list, the resource with the highest value during the monitoring period is picked and shown. If there are multiple series in such a DataView, the first one is picked automatically. However, you are free to select any other series in that DataView and/or the

⁴ The *.WK1 format specification allows a maximum of 8000 values in one spreadsheet file.

SHARE Minneapolis

resources that are shown on the chart. For single value DataViews, all series are picked at once and their course is shown on the graph.

Once you selected the series and resources that you would like to save, you can also limit the time window. For example if the series are available from 8.00 a.m. to 4.00 p.m. and you are only interested in saving the data collected around noon, change the **From** and **To** time values below the chart accordingly. To expand such a smaller window for better readability, press the button labeled **Zoom**. To condense, click on the same button - now labeled **Total**.

Finally, when all selections are made, press the **Save...**-button and a *File Save* dialog appears where you can enter the name of the *.WK1 file that PM of OS/390 should create. The number of rows in the resulting spreadsheet corresponds with the contents shown in field# **Samples**. The **Print...**-button, finally, allows to print the chart shown on the window on any printer installed on your workstation.

Tile DataViews

When you have created a couple of DataViews, the PerfDesk might become badly arranged. To arrange all DataViews in the PerfDesk properly, click on the **View** action bar menu item and then select **Tile Data Views**. This option distributes the space within the PerfDesk equally to all DataViews defined on that PerfDesk.

Finding Highest and Lowest Values

For each individual series in aDataView it is possible to search for the highest or lowest value for such a series since PM of OS/390 has been started. Once you selected that function with the context menu of a series (right mouse button-click on a bar) the DataView moves forth or back on the x-axis and shows you the data at that time.

Getting Help

Since the Java 2 platform has no concept how to handle **Help**-requests consistently and transparently for an application, PM of OS/390 uses the assistance of a web browser to display a help file once you press the **Help** button. If you want to change the browser used by PM of OS/390, select **View** from the action bar, then **Options**, and finally **Help Browser**. On the dialog, either select Netscape** Navigator or Microsoft** Internet Explorer.

Changing the Local Time Zone

If the times shown in the DataView does not match the local time of your workstation, you can change the local time zone by clicking in a row first on the action bar item **View**, then on **Options**, and finally on **Set Timezone**. Set a time zone to match the contents of field **GMT-offset and Timezones** with your local time. Click **Ok** when you are done.

Linux for S/390

If you install the RMF data gatherer for Linux, you can connect to Linux images just the same way you are connecting to OS/390 or z/OS Sysplexes. We currently don't have the concept of a cluster for Linux so you have to connect to each single Linux image.

The Linux data gatherer is available for x86Linux and for Linux for S/390, Kernel release 2.2. Please visit the RMF homepage and try out the tools.

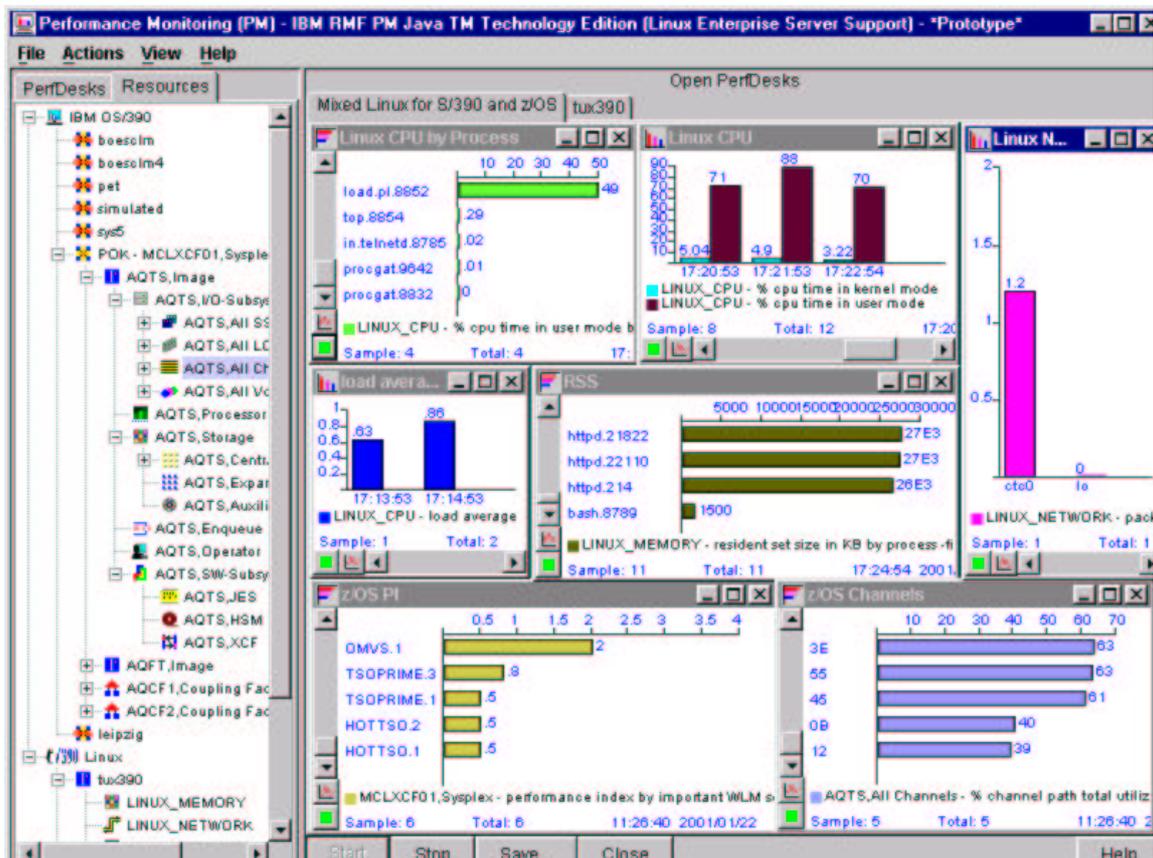


Figure 14

Summary

This document explained the most important concepts of PM of OS/390. This should help you to get started with PM of OS/390 and to use the application efficiently at your site. Refer to the extended publications described in “Bibliography” below for further details and assistance.

Appendix

Installation

Detailed installation instructions can be found in file *rmf.htm* located within the PM of OS/390 install directory. For your convenience, I have put a summary inside here:

1. Download PM of OS/390 from the RMF website
<http://www.ibm.com/servers/eserver/zseries/rmf/rmfhtmls/pmweb/pmweb.htm>
2. You will find the self-extracting file *gpmwinv2.exe* in the directory that you have specified in the download dialog of your browser. A detailed installation instruction can be found on the web page, as well as RMF PM versions for Linux and for IBM OS/2 Warp 4.
3. For customers with OS/390 V2R6 or higher, the host code will be distributed through the normal RMF maintenance channels. Customers who are on OS/390 V1R3 or OS/390 V2R4 will find the host code on the RMF website for download. Instructions, how to install the code on the host can be found in the file *readme.htm* after the download.
4. Make sure the following checklist for the host is complete
 - Your operating system level is OS/390 V1R3 or higher
 - Required PTFs are installed. A complete list can be found in *readme.txt*.
 - TCP/IP is configured and active
 - UNIX*** System Services is configured and the user that submits the host server has an OMVS segment and a home directory in the HFS
 - RMF and RMF Monitor III gatherer (RMFGAT) are active
5. Start PM of OS/390 from the program folder you have specified during installation.

Bibliography

RMF.HTM	located in the install path of PM of OS/390
Resource Measurement Facility	
User's Guide	SC28-1949