

HP OpenView Performance Manager Administrator's Guide

For HP-UX, Solaris and Windows

Version 6.01

Contents

HP OpenView Performance Manager Administrator's Guide	1
For HP-UX, Solaris and Windows.....	1
Version 6.01	1
Contents	2
Conventions used in this document	5
OVPM Administrator Tasks	6
Initial Setup of OVPM	6
About OVPM Port Settings and Firewalls.....	7
Running OVPM in a clustered environment.....	7
Selecting the Language that OVPM will Use	8
Defining Holidays.....	9
OVPM Command	10
Installing OVPM License Key.....	11
Using the Administrator User Interface.....	12
Accessing your OVPM Server Remotely	13
Configuring Datasources	14
Data from Systems (OVPA, OVO-EPC, SPIs).....	15
Secure communications with OVO 8	16
Batch Entry of System Lists	17
Proxied Log Files.....	19
HP OpenView Reporter Integrations.....	20
Configuring for Integration of OVPM with HP OpenView Reporter and OVIS	20
Integration Issues between OVPM and HP OpenView Reporter or OVIS	22
Using OVIS Restricted Views	22
Excluding Groups of Systems from HP OpenView Reporter	23
Configuring for Other Databases.....	23
Database Configuration Items.....	25
Flat File Data Source	27
Customers, Users and Logons.....	30
Customers	30
Defining OVPM Users.....	32
Logons.....	33
Logons, simplified use cases.....	33
Security	34
Enabling Client Authentication Security	35
Changing the Web Server Port Numbers.....	36
Using Secure Sockets Layer (SSL or https).....	37
Alternate Ways to Access OVPM.....	38
Linking Directly to OVPM	39
Running Graphs in Batch.....	40
Exporting to Text File.....	41

OVPM Maintenance	43
Managing Output files and Disk Space	43
Daily Maintenance	44
Advanced Work With Graph Template Files	45
Editing Graph Template Files Directly	45
Making User-Designed Graphs Permanent	48
Customizing the Look of the OVPM Interface.....	49
Using a Style File to Set Defaults	50
Using Skins to Affect Presentation Styles	54
Using Client Settings to Set Defaults.....	56
TroubleShooting	57
Troubleshooting the OVPM client.....	57
Graphs are not Updated	57
Problems with the Java Plug-in.....	58
Java Screens Are Cut Off.....	59
No Data is Shown in a graph	59
Enable Java Logging.....	61
OVPM does not Respond.....	61
No Names are in the System Name Box (Web Forms) or Data Sources Window (Java interface)	61
No Graphs Display in Java Interface	62
Data Overlaps in Area graphs	62
Agent not Found on a System.....	62
No Configuration Data Found in Selected Range.....	63
Some clients can't connect to OVPM	63
WARNING: Duplicate name in Manifest: Depends-On	65
sun.plugin.cache.DownloadException.....	65
Troubleshooting the OVPM Server	65
Step 1: Check if your server is configured to run OVPM properly?	66
Step 2: Check the Application Server.....	67
Step 3: Test What DOES Work	67
Step 4: Check the Status File	68
Step 5: Enable Tracing.....	69
Unable to Find the First Web Page.....	69
Unable to Find Graph.....	70
Keyword Unrecognized	70
Database Open Error.....	70
Troubleshooting Database Connections	70
Unresolved Host Name	71
Troubleshooting OVO for Windows Graphs.....	71
Detail Specifications	72
Metrics	72
Date Ranges	72
Number of Points	75
Filters	76
TableFilters	78

Labels	78
Parameters	79
Date Format Strings	80
Time Format Strings	80
Keywords	81

©Copyright 2003-2007 Hewlett-Packard Development Company, L.P.

Conventions used in this document

When describing the location of files on the HP OpenView Performance Manager (OVPM) server system, the following abbreviations will be used. (Note that slashes are "/" on UNIX and "\" on Windows, the / will be used in examples throughout this guide).

<code><install_dir></code>	<p>The directory where OVPM was installed.</p> <p>On HP-UX and Solaris systems this directory is <code>/opt/OV</code>.</p> <p>On Windows, the directory may vary but by default it is <code>C:\Program Files\HP OpenView</code>.</p>
<code><data_dir></code>	<p>The directory chosen for OVPM to place its working files.</p> <p>On HP-UX and Solaris, this directory is <code>/var/opt/OV</code>.</p> <p>On Windows the directory may vary but by default it is <code>C:\Program Files\HP OpenView\data</code>.</p>
<code><java_dir></code>	<p>The directory where the Java 1.4 SDK is installed.</p> <p>On HP-UX and Solaris this is <code>/opt/OV/nonOV/jdk/a</code>.</p> <p>On Windows this can vary but is usually <code>C:\Program Files\HP OpenView\nonOv\jdk\a</code></p>
<code><tomcat_dir></code>	<p>The directory where the Tomcat application server is installed.</p> <p>On HP-UX and Solaris, this directory is <code>/opt/OV/nonOV/tomcat/a</code></p> <p>On Windows the directory may vary but by default it is <code>C:\Program Files\HP OpenView\nonOV\tomcat\a</code></p>
<code><system></code>	<p>The name of the system where the OVPM server is running. This is normally used from a web browser in a URL and should be a name that the network name servers can resolve to the OVPM system.</p>
<code><stdout.log></code>	<p>Where the tomcat log files are installed.</p> <p>on Windows the directory may vary but by default it is: <code>C:\Program Files\HP OpenView\nonOV\tomcat\a\logs</code></p>
<code><conf></code>	<p>Where the configuration files are installed.</p> <p>on windows the directory may vary but by default it is: <code>C:\Program Files\HP OpenView\nonOV\tomcat\a\conf</code></p>

OVPM Administrator Tasks

OVPMP presents a web-based user interface to its end users. This interface can be used to display data in a variety of formats such as graphs, tables and spreadsheets. The end user has a large degree of control using the web-based interfaces, but there are still a few items that require setup on the OVPMP server system.

These items include setting the security policy, defining logins, setting up connections to the data sources that OVPMP uses to get its data, managing OVPMP trials, permanent license keys, and performing advanced customizations whenever they are required. This document covers such topics and should be used by someone that is in charge of administering the OVPMP server, henceforth called the "OVPMP Administrator".

Initial Setup of OVPMP

OVPMP has considerable flexibility in configuring its operation. You don't need to configure each feature in order to start using OVPMP. Most features come pre-configured to work in a default manner. You can change the OVPMP configuration at any time in order to override the default operations. Details about each of these configuration options are presented later in this manual.

OVPMP presents data from other data sources. In order to work properly, you must configure OVPMP to know the location of those data sources.

By Default: OVPMP for Windows will detect HP OpenView Internet Services (OVIS) and HP OpenView Reporter if they are installed on the same system when OVPMP is installed. Otherwise each data source must be configured before it can be accessed by OVPMP.

OVPMP will allow the data from some data sources to be segregated into different customers. For some data sources like OVIS this segregation is a requirement. Other data sources like HP OpenView Reporter allow customer segmentation of the data as an option. For other data sources such as HP OpenView Performance Agent (OVPA) customer segregation is not a feature but can be added using OVPMP when these data sources are accessed.

By Default: If OVIS or HP OpenView Reporter is configured as a data source, then any customers they define will be known to OVPMP. Otherwise no data segregation by customer is enabled until customers are configured into OVPMP. For more details see [HP OpenView Reporter Integration](#), [Configuring for Integration of OVPMP with HP OpenView Reporter and OVIS](#) and OVIS Integration using OVIS Restricted Views.

OVPMP uses a web based user interface and so does not require a logon to the server system in order to be used. When used in this manner, modified graphs

and user preferences are shared by all users. Optionally, OVPM can be configured to define users that each have their own private areas for saving graphs and user preferences.

By Default: No users are enabled in OVPM. All OVPM clients will share the same area for their graphs and user preferences. For more details see [Customer, Logons and Users](#).

As an added security measure, OVPM allows some control over the URL used by web users to locate the OVPM server. The port number can be overridden using the `ovpm config` command. A list of authorized clients for OVPM can be configured using either the client network names or their IP addresses (See [Security](#) for more details.)

By Default: Port 8080 will be used to find the OVPM web pages and port 8443 will be used for secure connections. Any client system is authorized to connect to the OVPM server.

About OVPM Port Settings and Firewalls

When OVPM runs, it provides a web interface. In order to access the OVPM web pages from systems other than the one where OVPM is installed, you must have your environment configured to allow this type of access. Access to the standard OVPM user features is accomplished through a single TCP port (8080 by default). If you want to run the OVPM Administrator Interface then you should also enable the TCP port for HTTPS (8443 by default).

The exact mechanisms for enabling this type of interface will vary depending on the type of system, type of firewall, etc.

For example, on Microsoft Windows XP, Service Pack 2:

1. Select Start > Control Panel > Network Connections.
2. Right click on "Local Area Connection" and select Properties.
3. Under the Advanced tab, select the Windows Firewall "Settings..." button.
4. Select the Exceptions tab.
5. If you do not already have the "tomcat5.exe" and "tomcat5w.exe" programs in the Programs and Services list:
 - a. Press the "Add Program..." button.
 - b. Use the "Browse..." button to locate the Tomcat executables. Usually located at `<tomcat_dir>\bin\tomcat5.exe` and `<tomcat_dir>\bin\tomcat5w.exe`
6. Check the Tomcat5.exe and Tomcat5w.exe programs and exit.

Running OVPM in a clustered environment

OVPM can be used in a clustered environment where multiple processors share access to common disk space and one or more systems will act as a "fail-over" processor for the other when there is a hardware failure. For example, if you had

a two-node cluster, and node A shuts down, the cluster management system detects that node A is not available, and a cluster fail-over occurs. The cluster management system would start OVPM running on node B. However, if node A was running, and OVPM stopped, OVPM would not automatically failover to node B.

To run OVPM on a cluster, OVPM must be installed on the primary node (also called processor) and on each node that will be used in a fail-over. In order to maintain OVPM configuration information and saved user graphs in the event of a fail-over, these files must be placed on a disk system that is shared by all the processors.

For more details on installing in a cluster refer to the OVPM Installation, Upgrade and Migration Guide.

Selecting the Language that OVPM will Use

The OVPM product is Internationalized so that it can adapt some features such as date, time and number formats to match the language on the OVPM server system. Generally the user interface and documentation will remain in English but date, time, and number formatting can be adjusted.

If your version of OVPM has been Localized to a specific language, then it can interact in that language as well as English. Localization typically includes translated web pages, online help, user documentation, and item selections for the user interface drop-down menus.

You select the language that OVPM should use with the language settings on the OVPM server system before you start OVPM.

- On Microsoft Windows, set the system language using the control panel "Regional Options" feature. Set the default locale for the system to the language you want to use. You may have to reboot to have this change take effect.
- On Unix, export the environment variables LANG and potentially LC_ALL to a locale that is installed on your server system. You can get a list of installed locales with the command `locale -a`.

Once the language is set, restart the OVPM service.

If you are running a localized version of OVPM, then you should set the language to match the localization. Otherwise OVPM will interact in English.

The clients of OVPM should be running in the same language as the OVPM server. If they are not, then they can see a mixture of date, time and number formats depending on whether the number is formatted on the server or on the client. The Java User Interface executes on the client system and will attempt to format data according to the language of the client system. No translation is done to match data coming from the server with that of the client so certain features

may not function properly when the client and server systems are running in different languages.

Regarding Time Zones:

When the OVPM server is in one time zone and the agent is in another, the OVPM server's time zone takes precedence. When the OVPM server, agent and client are in 3 different time zones, the OVPM server's time zone takes precedence.

NOTE: OVPM does not support working across multiple locales. If the client is in a different locale than the OVPM server, users may experience unexpected behavior.

Defining Holidays

When drawing graphs, a user can choose to exclude data from holidays. The definition of which days are holidays is a responsibility of the OVPM Administrator. The list of holidays can be set globally or it can be set individually for each customer. Holidays are set by creating a text file called `VPI_Holidays.txt`. The syntax of this file is described below. If a `VPI_Holidays.txt` file is found under the customer configuration directory (`<data_dir>/conf/perf/VPI_CUST_<customername>/VPI_Holidays.txt`) then this file is used for all logons to this customer. If no customer holidays file is present, then the global holidays file at `<data_dir>/conf/perf/VPI_Holidays.txt` is used. If neither the customer nor the global Holidays file is found, then the skip holidays feature has no effect.

The `VPI_Holidays.txt` files is a text file consisting primarily of dates, one per line. The file can contain blank lines and comment lines (starting with a #). The first line in the file can specify the order of the date fields in the rest of the file. By default, Dates are expected in the order of Year, Month, Day. Any non-numeric separator except blank is allowed to separate the three fields. Years can be four digits or less.

2004,12,25 is the same as 2004/12/25 is the same as 4.12.25

If you want to enter dates in an order other than Year,Month,Day, enter a special first line containing the characters "Y", "M" and "D" in the order you want to enter them. Repeated characters are allowed as are separator characters. Valid model lines are:

MDY
MM/DD/YY
dmy

All dates in the rest of the file should consist of three numeric fields in the order of this model.

Anything following the third numeric field is ignored and can be used for comments.

Here is an example Holidays file:

MM/DD/YY
1/1/04 New Year's Day
1/19/04 Martin Luther King Day
2/16/04 President's Day
5/31/04 Memorial Day
9/6/04 Labor Day
11/25/04 Thanksgiving Day
11/26/04 Friday following Thanksgiving
12/27/04 Monday after Christmas Day
12/28/04 Company Floater
1/3/05 New Year's Day
1/17/05 Martin Luther King Day
2/21/05 President's Day
5/30/05 Memorial Day
7/4/05 Independence Day
9/5/05 Labor Day
11/24/05 Thanksgiving
12/23/05 Company Floater
12/26/05 Christmas Day

OVPM Command

The System Administrator can use the `ovpm` command with various options to manage OVPM. On Windows, the `ovpm` command must be entered through a MS-DOS command window or through a Windows script file. The OVPM command is:

```
<install_dir>/bin/ovpm options
```

Options include the following:

`status` - to check the OVPM execution status

`start` - to start the OVPM application service

`stop` - to stop the OVPM application service

`restart` - to stop and then start the OVPM application service

`license` - to check or update the product license, displays the AutoPass GUI.

`admin` - (OVPM admin) to start the OVPM Administrator's Interface

`trace [log] [off]` - to start or stop generating detailed trace files.

`config [-port portnumber]` - to configure OVPM, by default port 8080 is used. You must restart the OVPM application for changes to take affect.

Installing OVPM License Key

A license key password is required to use OVPM. When the software is installed and used for the first time, a 60-day trial license will be granted. Within this 60-day period, you must obtain a permanent license key password or a trial evaluation extension to continue using the product.

To obtain a permanent license, run the following command:

```
<install_dir>/bin/ovpm license
```

Your license status will be displayed and then you will receive a command line prompt: `License Command>`. If OVPM is running with a trial license, it will display a message indicating `Instant On: Yes`. Instant On is the trial license feature. If the permanent product license is already installed, it will display a message indicating "HP OpenView Performance Manager has a permanent license".

If you have an internet connection and wish to obtain a permanent password, enter the command `gui` on the OVPM server. This will bring up the HP OpenView AutoPass GUI. On HP-UX and Solaris, you may have to export the `DISPLAY` variable to direct the AutoPass gui to the proper display window before running the `ovpm license` command.

When the AutoPass GUI first runs, click **OK** to continue past the Temporary Evaluation screen (this dialog will always display, regardless of your license status). Then perform the following steps:

1. In the Internet Connection dialog, select either Direct Connection or enter the Proxy information required.
2. Select Install Permanent Password in the dialog and enter the HP OpenView purchase order information from your Entitlement Certificate.
3. Select the LTU (License To Use) number that corresponds to the number of licenses purchased. Note: Only one license is needed per installation on a server.
4. Fill out the customer information form.
5. When finished, the password keys will be installed automatically.

If you used the www.webware.hp.com website to request a license password and have received it via e-mail, you will need to install it manually. Follow these steps:

1. Extract the password from the e-mail attachment and save the license keys in a plain ASCII text file on the system where OVPM is installed.
2. Enter this command: `<install_dir>/bin/ovpm license`
3. After displaying your current license status, you will receive a `License Command>` prompt. Enter `install`.
4. You will then be prompted for the name of the file created in step 1. Enter the full path to the file. If this process is successful, you will see a message "Password successfully added"; otherwise, an error message will be displayed.

If you need to request a Trial (Evaluation) License Extension, access the webware website to get contact information for evaluation software passwords. Contact the License Center in your region to request a Trial (Evaluation) License Extension. The Product Number and Name will be "TRIAL-OVPM5". This is a one time 60-day extension. You will receive your password key by email and it can be installed manually using the process described above.

For more complete details about licensing, see the `AutoPass_guide.pdf` document in the `<install_dir>/paperdocs/PM/C` directory.

Using the Administrator User Interface

The OVPM Administrator User Interface is used to configure the following local to OVPM:

- Customers
- Users
- Groups
- Systems

If OVPM is connected to the HP OpenView Reporter product then Customers, Groups and Systems are obtained from HP OpenView Reporter.

If OVPM is connected to the OVIS product, then customers are obtained from OVIS.

In order to edit these non-local items, you should use the HP OpenView Reporter and OVIS interfaces directly. The OVPM Administrator Interface can be used to augment the information obtained from HP OpenView Reporter and OVIS and to define Users.

The OVPM Administrator's Interface shows only Customers, Users, and Systems that are defined locally by OVPM. When running the OVPM Java and Web Forms interfaces, the list of Customers, Users and Systems from HP OpenView Reporter and OVIS are merged with the local definitions and so may present a different list from those found in the OVPM Administrator's Interface.

The OVPM Administrator's Interface is accessible through two means: command line (a Java application) and by Web Browser (a Java applet).

The OVPM Administrator's Interface may be run locally on the system where OVPM is installed by invoking the command:

```
<install_dir>/bin/ovpm admin
```

On Windows, you can use the shortcut: Start / Programs / HP OpenView / Performance Manager / Performance Manager Administration.

You can run the Administrator Interface remotely using a web browser by referencing the secure web URL:

```
https://<system>:8443/OVPM/PMAdmin.htm
```

where <system> is the name of the OVPM system and 8443 is the Secure Sockets Layer (SSL) port number configured to the Tomcat server. In either case you must provide the password for the OVPM administrator before you will be allowed to view or edit any items.

For specific instructions on using the OVPM Administrator Interface, press the Help button once the interface is running.

Accessing your OVPM Server Remotely

You can access the "HP OpenView Administrator Client" on another OVPM server through the OVPM Administrative Client (Java Application) or from a Web browser **IF you have a SSL port configured on the remote server.**

To access another OVPM server using the HP OpenView Administrator Client:

1. Launch the Administrator's Interface. From the command line type: ovpm admin or select: Start - Programs - HP OpenView - performance manager - Performance Manager Administration.
2. The HP OpenView Administrator Client program displays.
3. Select File - Connect...
4. The Connect to OVPM Server displays.
5. Change the host name to the fully qualified name of the server you want to access.
6. Change Port number to the SSL port number for the server you want to access.
7. Select OK.
8. The HP OpenView Administrator Client should change to the HP OpenView Administrator Client of the new server.
9. Launch a Web browser on the local OVPM server system.
10. In the browser's address box type the name of the remote server. For example:

```
https://moon:8443/OVPM/PMAdmin.htm
```

then press Enter or OK.
11. The Connect to OVPM Server displays.

Note: When you connect to a remote server through a web browser, the Menu bar of the HP OpenView Administrator Client shows the menu items: Refresh, Save, Help instead of the menu items: File, Help. The Refresh and Save in the menu indicates you are on the remote system. To return to your local system you must change the system name in the Web Address box.

Configuring Datasources

OVPM includes no data collectors of its own but it can display data from a variety of sources. Before using OVPM you should configure it to locate its data sources. This section of the manual covers how to configure the data sources.

OVPM can present data from the following OpenView product's data sources and databases:

- **HP OpenView Reporter** contains a database that consolidates data from a variety of other performance products including OVPA and OVO agents. It discovers systems that have one of these agents configured and groups them according to rules set by the system administrator. You can use HP OpenView Reporter to configure customers and the systems that they control. When used with Performance Manager, this list of systems, groups and customers will be shared by both products.
- **OVIS** probes network services such as web pages (HTTP, HTTPS), email, network name servers, etc. It saves information about the availability and response time from these probes into a database. Performance Manager can display this information alone, or in conjunction with data from other sources.
- **HP OpenView Operations (OVO)** deploys its HP OpenView Operations agent (OVO agent) to systems it is monitoring. Starting with OVO agent version 7, this agent includes an Embedded Performance Component (OVO-EPC) that can log information about the general health and performance of a system.
- **OVPA** provides for a robust set of performance information about a system. It efficiently logs system health information plus the information necessary to drill down into details showing exactly what was happening on that system during times of interest. Previous releases of OVPA have been known as "MeasureWare (MWA), VantagePoint Performance Agent (VPPA), and Performance Collection Software (PCS).
- Oracle and SQL Server databases can be used as data sources providing that they meet the requirements of providing time-series data. OVPM presents the data as defined by the database tables and views and will not perform joins of multiple tables. Views or stored procedures may be used to accomplish table joins and other functions. Each table or view should provide a metric of a type of date/time that indicates the time of the record.
- OVPM is not a general purpose database reporting tool, but it can provide limited display of time-series data from generic databases. Instructions for configuring these generic database datasources is covered later in the [Configuring for Other Databases](#) section.

Performance Manager supports access to at most one system running HP OpenView Reporter and one system running Internet Services. It supports access to a large number of systems that are running OVPA or OVO agent. The names of these systems can be obtained from the HP OpenView Reporter product or they can be entered locally using the OVPM Administrator's Interface or loaded from a file.

- [Data from Systems](#)
- [Secure communications with OVO 8](#)
- [Batch Entry of Systems](#)

- [Proxied Log Files](#)
 - [HP OpenView Reporter and OVIS](#)
 - [OVPA](#)
 - [OVO Agent](#)
 - [Other Databases](#)
-

Data from Systems (OVPA, OVO-EPC, SPIs)

OpenView provides several ways to collect performance data from computer systems. OVO provides **OVO-EPC** to collect lightweight performance data used to evaluate the general health of a system. It logs 5 weeks of data that can be retrieved and displayed by tools like OVPM. As the data is logged, data is also evaluated so that conditions requiring operator attention are immediately recognized and brought to the attention of the OVO console.

OVPA collects in-depth performance data for a system and logs a configurable data interval for later analysis. OVPA also analyzes its performance data and can perform corrective actions or send notifications to other tools like Openview Operations for their attention.

OVO-EPC and OVPA both log their data on the system being monitored. This allows very large scalability without burdening the network or central servers. The data are all available by connecting to the systems as needed.

Application data can be captured by HP OpenView Smart Plug-Ins (**SPIs**). A wide variety of SPIs are available to monitor things like Databases, Mail Servers, major applications, etc. These SPIs utilize one of the performance agents (OVO-EPC or OVPA) to log and manage their data.

The HP OpenView Reporter product provides a data warehouse to centrally store summary information from the OVO-EPC and OVPA collectors. By design, HP OpenView Reporter discovers which systems have performance agents and then gathers data summaries each night and saves them in its central database. It then generates reports that group and rank the systems' data in a variety of ways. HP OpenView Reporter generates new reports each night that provide a management level overview of how systems in the enterprise are performing.

OVPM can access and display data from all these data sources but it does not generate or archive any data on its own. When drawing a graph showing data from OVPA, for example, OVPM will use datacomm to connect to the system where OVPA is running and retrieve data directly from the OVPA log files. This data is considered **near real-time** and can be only minutes old. Once the OVPM graph is discarded, the data are not preserved by OVPM. If the graph is re-drawn, the OVPM will access the system and retrieve the data again, thus retrieving any new data that has been logged.

When displaying data from the HP OpenView Reporter product, data is accessed from the HP OpenView Reporter central database. Data can be retrieved for a single system, or for multiple systems at once. The granularity of the data is limited by whatever summarization was done by HP OpenView Reporter as it

gathered its data. Since data is not obtained directly from the systems, only data that is gathered into the HP OpenView Reporter database can be shown from the HP OpenView Reporter data source. This is typically data up until the previous midnight. For this reason, the HP OpenView Reporter data are considered **non real-time**. Displaying data from the HP OpenView Reporter database can insure that the production systems are not being asked to retrieve performance data during their peak usage periods.

In order to communicate with OVO-EPC or OVPA, the OVPM product must know the name or IP address of the desired system and what type of datacomm it is using. OVPA versions 3 and earlier use datacomm based on DCE or NCS protocols. OVO-EPC and OVPA version 4 and later use datacomm based on HTTP protocol (called internally "BBC" or "CODA"). It is possible for a system to have either, both, or none of the performance datasources running.

Some data sources for OVPM require that you provide a system name. Users can type in a system, or the OVPM Administrator can enter system names using a batch utility or the OVPM Administrator GUI. But the easiest way for OVPM to get a list of system names and what data sources they provide is for OVPM to get this information from HP OpenView Reporter. HP OpenView Reporter discovers systems and their agents and keeps them organized in its database. If OVPM connects to HP OpenView Reporter, then it can access this system list. See [Configuring for HP OpenView Reporter](#) for details on how to configure OVPM to connect to HP OpenView Reporter.

If you do not have HP OpenView Reporter available, or if you want to display data from systems that are not known to HP OpenView Reporter, then you can configure systems locally in OVPM. Use the [OVPM Administration Interface](#) to configure systems directly in OVPM.

When systems are configured, they will be displayed in a "pick list" or "system selection tree" in the OVPM user interfaces. This is the easiest way to select systems and offers such features as grouping of systems for easy location and restricting access to systems to specific customer logins. If a system has not been pre-defined, then users can manually enter the system name (providing they have not been restricted based on their logon.) Manually entered systems are not saved and must be re-entered each time they are used. Any systems that will be used repeatedly should be permanently entered by the OVPM Administrator.

Secure communications with OVO 8

Currently only the OVO agent version 8.0 and later offers the option for doing secure communications with OVPM. In the future other agents should get this capability. Using secure communications prevents anyone from examining the data that you transfer across the network as well as restricting access to the agents to specifically enabled installations of OVPM.

Using secure communications will occur a bit slower than regular communications due to the effort to encrypt and decrypt the data as it is moved

so you should only use it when you need it. See also the section on [Security](#) for more information on securing your OVPM installation.

To use secure communications to the OVO 8 agent, follow these steps:

1. Follow the OVO instructions for installation and configuration of secure datacomm.
2. From the OVO console, deploy an OVO 8 agent to the system where OVPM is installed. This will mark that system as belonging to the group of systems that are owned by this OVO installation and so allow OVPM to communicate with other agents.

At this time, we are not supporting using the same OVPM system to connect to OVO agents that were deployed by different OVO consoles. While this may be possible, you should seek qualified consulting help to configure it.

3. On each system where you want secure communications to be used, change the configuration to set `SSL_SECURITY` to `ALL`. Follow the OVO procedures or logon to each system and execute the command:

```
/opt/OV/bin/ovconfchg -namespace coda -set  
SSL_SECURITY ALL
```

See the HTTPS Agent guide in the OVO documentation set for alternate ways to enable secure communications.

OVPM will now automatically use secure communications (https) to any agent that is configured to require it. OVPM will use regular communications (http) to any agent that is not configured for secure communications, or does not support it. Other than the slightly slower communications to secure agents, the operations on secured agents will be identical to regular agents.

Batch Entry of System Lists

A facility has been provided to enter lists of systems from an external source. You provide the list of systems in two separate files

systemsCODA.txt Contains a list of systems that are running the OVO agent (version 7 and later) and OVPA,(version 4 or later). These agents all use the HP BBC datacomm libraries.

systemsMWA.txt Contains a list of systems that are running the OVPA using DCE or NCS datacomm. OVPA 4 agents using the BBC datacomm should be listed in the systemsCODA.txt file. The MeasureWare Agent and the VantagePoint Performance Agent should also be listed in this file.

Sample files can be found in the `<install_dir>/newconfig/OVPM` directory. These files can be edited using a text editor or generated by a script or program from another source. The syntax of the files is as follows:

- Blank lines and lines beginning with a pound sign # are considered comments.
- A line that starts with `GROUP:` specifies a group name. Systems following this line in the file are considered to be a member of the group. A succeeding `GROUP:` line will start a new group. The group name should follow the `GROUP:` keyword.
Example: `GROUP: Web Servers` starts a group called Web Servers.
- If no group is entered, then the system name will be available, but it will not belong to a group. This is a perfectly acceptable thing to do.
- All other non-comment lines are considered to be system names. A system name should be valid for network addressing. Try the `ping` command if you are unsure about whether a system name is valid. IP addresses can be entered, but are often confusing to users when they want to select the system.
- To indicate a log file that is being proxied by another OVPA system, precede the original system name with the name of the hosting system in square brackets.
`[hostname]originalsystem`
indicates that *originalsystem* log files are mounted under OVPA on the *hostname* system.
- A system may belong to more than one group; to do this, list the system after each of the desired group specifications. If a system is entered more than once in the same group it will only be displayed once in that group.

A sample systemsMWA.txt file:

```
#####  
#####  
# systemsMWA.txt file  
#  
# This file contains the names of systems that are running  
the OpenView  
# Performance Agent (OVPA) using DCE or NCS datacomm. OVPA  
4 agents using  
# the BBC datacomm should be listed in the systemsCODA.txt  
file.  
# The MeasureWare Agent and the VantagePoint Performance  
Agent should also  
# be listed in this file.  
#  
# The systems*.txt files are used to populate the drop down  
box on the  
# OpenView Performance Manager forms.  
#  
# Blank lines and lines starting with a pound sign "#" will  
be treated as  
# comments.  
#  
# GROUP: lines identify grouping of systems
```

```

#
# If you do not have the OpenView Reporter installed, then
you should
# enter the names of the systems to be used in this file
#####
#####
GROUP: Agent MWA
WebServer1
WebServer2
AppServer.city.company.com

GROUP: Web Servers
WebServer1
WebServer2

GROUP: Proxied Logs
[MyHost]bigbird

```

If a system has more than one agent (both MWA and Coda), add the system to the same groups in the `systemsMWA.txt` and `systemsCODA.txt` files. If you add the system to a group that exists in only one file, then unpredictable results may occur when selecting a group and drawing a graph in the Web Forms interface.

1. Edit the files to include the systems you want listed
2. Place these files into the `<data_dir>/conf/perf` directory.
3. Restart the OVPM server.
4. Next, run the OVPM Administrator Interface. The systems should be shown in the list of systems. Perform any manual edits or re-grouping desired, then press the **Save** button.

NOTE: After you press **Save**, the system information from the `systemsMWA.txt` and `systemsCODA.txt` files are saved in the `OVPMSystems.xml` file. The `.txt` files are backed up (converted to `.bak` files) and not processed anymore. This avoids double processing the system names each time the OVPM server starts. The `.bak` files can be deleted if desired. If you want to add new systems by using batch entry method, repeat the steps 1 to 4.

If you upgrade MeasureWare Agent (MWA) to OVPA in your environment, OVPM can automatically configure the data source information if you run the following command:

```
ovpmbatch updatesystem
```

After you run the command, the `OVPMsystems.xml` file (present in `<data_dir>/conf/perf` directory) captures the configuration information. You need not restart OVPM for the changes to take effect. This utility should be run when OVPM usage is minimal as this operation is time and resource consuming.

Proxied Log Files

When an OVPA log file is transferred to another system and then mounted under that system's OVPA agent, we refer to it as being "proxied" by the second

system. A proxied log file is automatically detected and configured by HP OpenView Reporter when it discovers the hosting system.

You can add a proxied log file manually just like you do when you add a system. Make a temporary addition using the Java interface **Add** button or by entering the system identifier in the Web Forms interface type-in systems field. Make a permanent addition by using OVPM Administrator Interface.

The syntax for a proxied log file is [*hostsystem*]*actualsystem* where *hostsystem* is the name of the system that is hosting the log file and *actualsystem* is the name of the system that generated the original log file.

HP OpenView Reporter Integrations

When OVPM is used together with HP OpenView Reporter, several integrations are enabled. When OVPM is installed on the same Windows system as HP OpenView Reporter, it is automatically integrated. If HP OpenView Reporter is installed after OVPM or is installed on a different system, then you will have to modify the **OVPMconfig.ini** to locate the HP OpenView Reporter database. OVPM does not require the installation of HP OpenView Reporter, but the two products are designed to work well together. Major integration areas consist of the following:

- HP OpenView Reporter discovers systems with HP OpenView performance agents and with OVO_EPC. HP OpenView Reporter also groups these systems by operating system (although the HP OpenView Reporter administrator can alter this and add groupings of their own). The systems and groups defined in HP OpenView Reporter are automatically available in OVPM. upon integration.
- HP OpenView Reporter (3.0 and later) allows segregation of systems by customer. Once customers have been set up in HP OpenView Reporter, OVPM can use these definitions to produce customer specific graphs and tables. This capability can be used to segment a user's access to their own systems. (Refer to HP OpenView Reporter and OVIS documentation for more information.)

Configuring for Integration of OVPM with HP OpenView Reporter and OVIS

To integrate with HP OpenView Reporter or OVIS, OVPM must be able to connect to the HP OpenView Reporter database which is used by HP OpenView Reporter and OVIS. In order to access its database from OVPM, the HP OpenView Reporter database must be either Oracle or SQL Server (the default MSDE database is acceptable).

HP OpenView Reporter. To integrate with the HP OpenView Reporter product, edit the [REPORTER] section of the file `<data_dir>/conf/perf/OVPMconfig.ini`. Remove the comment marks ";" from the example lines and replace the remaining information to match your installation.

Example:

```
;;If the OpenView Reporter product is installed:
;
[REPORTER]
SYSTEM = Belleview
DSN = reporter
TYPE = sqlserver
USER = openview
PASSWORD = openview
PORT = 1433
MINCONN = 5
MAXCONN = 10
CONNEXPIRE = 60
SLEEPTIME = 2000
MAXRETRY = 30
```

OVIS. To integrate with the OVIS product, edit the [INTERNET_SERVICES] section of the file `<data_dir>/conf/perf/OVPMconfig.ini`.

```
;; If the OpenView Internet Services product is installed:
;
[INTERNET_SERVICES]
SYSTEM = SunStroke
DSN = ovis
TYPE = oracle
USER = openview
PASSWORD = openview
PORT = 1521
MINCONN = 5
MAXCONN = 10
CONNEXPIRE = 60
SLEEPTIME = 2000
MAXRETRY = 30
```

For an explanation of each of these items, see the [Database Configuration Items](#) tables.

NOTICE: Modifications to the `OVPmconfig.ini` file will not take effect until the `ovpm` service is restarted. Issue the `ovpm restart` command after modifying the OVPM configuration.

Integration Issues between OVPM and HP OpenView Reporter or OVIS

If using HP OpenView Reporter or OVIS 5.20 or OVIS 6.0 and you had OVPM 4.x installed then upgraded to OVPM 6.0, the configuration items such as the list of service groups and customers are not updated until after the configuration refreshes each night, or until OVPM is stopped and restarted. This means the configurations are not updated each time a URL request is made.

See the OVPM Release Notes for more details on integration issues with HP OpenView Reporter and OVIS.

Using OVIS Restricted Views

OVIS can be used with two security models. The default provides data segregation by customer but does not enforce password protection. In this mode one customer could examine another's data providing they knew the logon. By logging on as "All Customers", a user would have the ability to display data from all customers at once.

OVIS also has a security mode called "Restricted Views". In this mode customers should not be allowed to see the data from another customer without knowing that customer's password. Typically in this mode the "All Customers" logon would similarly be protected by a password.

When used with OVIS, OVPM can run in either security mode to match the OVIS settings. These considerations also apply to the use of the "Graphs" feature in OVIS since both products share a common graphing component. Some manual synchronization between OVIS and OVPM is required starting with OVPM release 5.0.

- When configuring OVIS to run in restricted mode you should:
 1. In OVIS, Add a password to the "All Customers" customer name.
Passwords for all other customers are recommended but not required.
 2. In OVIS, enable Restricted Views
 3. In OVPM, set the password for the blank customer to match the OVIS "All Customers" password.
- When configuring OVIS to run in unrestricted mode you should:
 1. In OVIS, disable Restricted Views
 2. In OVIS, Remove any password from the "All Customers" customer name.
 3. In OVPM, set the password for the blank customer to a blank.

When OVPM runs, it reads the OVIS configuration for customers, passwords and service groups. It then refreshes this information each morning at 3:00 AM. If you make any changes to the OVIS configuration in these areas (customers, passwords, service groups), then OVPM and the OVIS "Graphs" function will not

reflect those changes until the following day. If you want changes to take effect sooner, then you should restart the OVPM server. (command "ovpm restart").

Excluding Groups of Systems from HP OpenView Reporter

When the HP OpenView Reporter product is being used in very large configurations, the number of systems and groups it provides to OVPM may slow down the initialization of the Java User Interface. A configuration with 5,000 systems, each of which is defined in four different groups will result in 20,000 items in the OVPM Java interface System Selection tree. It may take a minute or more to load this many systems into the tree. Once loaded, the Java interface should perform normally.

To speed the loading of the OVPM Java interface, you can exclude some groups of systems from being provided to OVPM by HP OpenView Reporter. Create or edit a file called `OVPMExclude.txt` in the `<data_dir>/conf/perf` directory. Add each HP OpenView Reporter Group to be excluded from OVPM, one per line to this file. Blank lines and lines starting with the # character are treated as comments and ignored.

Sample `OVPMExclude.txt` file

```
#####  
#####  
# The following groups, while extremely useful when using  
the OpenView  
# Operations product, are not necessary or helpful when  
they are copied  
# through the OpenView Reporter product and into OpenView  
Performance  
# Manager. Therefore, exclude them from OVPM.  
#####  
#####  
[ ITO]AM_NT_PST  
[ ITO]AM_NT_QIP  
[ ITO]AM_NT_RF
```

The HP OpenView Reporter database is processed when the OVPM server starts, then once a day after that. To see the effects of changes to the `OVPMExclude.txt` file you should restart the OVPM server (using the command `ovpm restart`.)

Configuring for Other Databases

OVP can access the HP OpenView Reporter database and the OVIS database. In addition it has the ability to access other databases providing they fit the following criteria:

- Microsoft MSDE, Microsoft SQL Server, or Oracle 8 or 9.
- Tables provide time series data and include a Date/Time type field that provides the proper date and time for each record.
- The data can be retrieved using simple tables or user provided views. Table joins are not performed by OVP.
- Databases are not overly complex. Databases with more than a few hundred tables or that contain tables with more than a few hundred metrics will take an inordinate amount of time to process and should be avoided.
- Contain a reasonable amount of data. OVP will pull unsummarized data from the database. Tables containing very dense data that require hundreds of thousands of records to be processed in order to produce a graph may not be good performers using OVP. If this is the case, consider providing a view (stored procedure) to have the data compacted by the database before it is picked up by OVP.

Steps to configure a generic database as an OVP datasource:

1. Make sure the database is running and contains some data.
2. Insure that remote access is allowed for the database and that you have valid configuration data as shown below.

Edit the `OVPconfig.ini` file. You may add up to ten generic databases by adding sections [DB1] through [DB10]. Add the database connection information following the appropriate section field. For example:

```
[DB1]
NAME = MyDatabase
SYSTEM = where.is.it.com
DSN = datasource
TYPE = oracle
USER = george
PASSWORD = ofthejungle
SHOWTABLES = SHEP FELLA URSELA *CREATURES
HIDETABLES = REPTILE_CREATURES
PORT = 1521
MINCONN = 5
MAXCONN = 10
CONNEXPIRE = 60
SLEEPTIME = 2000
MAXRETRY = 30
```

All of these items are described in the [Database Configuration Items](#) table.

3. Stop and restart OVP. Access the OVP home page and then select either the Web Forms Interface or the Java interface. Check the OVP logs for indications of any problems opening the databases.

Your generic database should now be available for use.

In the Web Forms interface, the Design tab will show a data source like "MyDatabase (DB)". Select this data source to begin designing graphs against your database.

In the Java interface, the Design window's Sources tab will show a DB item at the root of the sources tree. Expand it and you should see the list of generic database names (MyDatabase in this case). Select the database you want, then select Add Metrics and start designing your graphs.

Database Configuration Items

Here are the meanings of each option in the [REPORTER], [INTERNET_SERVICES] and [DB1] - [DB10] sections:

SYSTEM	System name where the database lives
DSN	Datasource name (usually HP OpenView Reporter). NOTE: For Oracle Databases this item should be set to the SERVICE_NAME as found in the TNSNAMES.ORA file. See example below *.
TYPE	MUST be either "sqlserver" or "oracle"
USER	User name for db access
PASSWORD	User's password
PORT	Port number where the database server listens. On systems that host an oracle database, this is USUALLY 1521, but it could be different. If HP OpenView Reporter is on an SQL Server database, you must run <code>svrnetcn.exe</code> on the database system to find out the port number: <ol style="list-style-type: none"> 1. Run <code>svrnetcn.exe</code> (possibly located in <code>C:\Program Files\Microsoft SQL Server\80\Tools\Binn</code> on the system where the database is installed). 2. Click once on TCP/IP in the enabled protocols list to highlight it. 3. Click on Properties and it will show you the port number.
MINCONN	Minimum number of database connections to maintain. The default is 1.
MAXCONN	Maximum number of database connections to maintain. The default is 1.
CONNEXPIRE	Number of seconds before a connection expires or the maximum number of seconds a connection can be checked out before it is automatically closed and recycled. A value of 0 means a connection can be checked out indefinitely. The default is 60 seconds.
SLEEPTIME	Number of milliseconds to sleep between retries when a thread is unable to get a free connection. The default is 2000 milliseconds.
MAXRETRY	The maximum number of times a thread should try to get a free database connection before it gives up. The default is 30.

* DSN Example for Oracle:

Generated by Oracle configuration tools.

```
DATABASE.EXAMPLE.COM =  
  (DESCRIPTION =  
    (ADDRESS_LIST =  
      (ADDRESS = (PROTOCOL = TCP)(HOST = example.com)(PORT  
= 1522))  
    )  
    (CONNECT_DATA =  
      (SERVICE_NAME = database)  
    )  
  )
```

In this example, the DSN should be set to "database".

The following items are only used with Generic databases, sections [DB1] - [DB10]:

NAME	Database Name. Provides a recognizable name in the user interface for [DB1]-[DB10]. If NAME is not provided then the section name (such as DB1) is used.
SHOWTABLES	<p>Specifies which database tables and views will be visible for this data source. The value should be a list of table or view names separated by blanks. In place of a name, you can enter a partial name plus a wildcard, either an asterisk * preceding or following a string. IOPS* will match any table or view that starts with IOPS. *XYZ will match any table or view that ends with XYZ. *ABC* will match any view or table that contains ABC anywhere in its name. All comparisons are sensitive to upper or lower case.</p> <p>If neither SHOWTABLES or HIDE TABLES items are provided then all tables and views will be visible.</p>
HIDE TABLES	<p>Specifies which database tables and views will not be visible for this data source. The value will be a series of table/view names or patterns separated by blanks, as described in the SHOWTABLES item.</p> <p>If HIDE TABLES is specified but SHOWTABLES is not, then all tables and views except those that match the HIDE TABLES list will be visible.</p> <p>If a table matches the SHOWTABLES and the HIDE TABLES patterns then it will be hidden. Example: SHOWTABLES = IOPS* HIDE TABLES = IOPS_DETAIL_DATA would make all tables starting with IOPS visible EXCEPT for IOPS_DETAIL_DATA.</p> <p>If neither SHOWTABLES or HIDE TABLES items are provided then all tables and views will be visible.</p>

Web Server Configuration Items

Below are the meanings of the options in the WEB_SERVER section:

PORT	This is the port number for the OVPM web server. This port number is set by the "OVPM config -port" command. The default port is 8080.
TRACELEVEL	This level is set by the OVPM Trace command. When tracing is turned on, it is set to level 3. When tracing is turned off, it is set to level 0. More detailed tracing can be turned on by manually setting this level to 5. This should be done by Hewlett Packard support when troubleshooting a problem.
REFRESH_RATE	<p>This rate configures how much time would need to elapse before the OVPM server would contact the agent to refresh a cached datasource connection. In the OVPM5 version, the value was hard-coded to 2 minutes. This value can be increased as a potential performance improvement.</p> <p>This value is in milliseconds (e.g. for 5 minutes, multiple 5*60*1000). If this rate is not specified in the OVPMConfig.ini file, then the default will be set to 5 minutes.</p> <pre>;; Web Server Configuration ; [WEB_SERVER] PORT = 8080 TRACELEVEL = 1 AGENT_REFRESH = 300000</pre>

Flat File Data Source

The OVPM Flat File data source is a mechanism designed to allow quick and easy addition of external data into the OVPM product, and particularly the OVPM Planner function. The file format is essentially an ASCII text file with one record per line and each field in a line separated by a tab, or a comma. A special header line at the top of the file defines the contents of the file including the metric names and orders.

The flat file format supports multiple metrics in one file. Only one class will be seen in OVPM for each file data source. It will be called "FILEDATA". Multiple Instances can be supported by simply including a field that can be used to differentiate each one.

The files used by this facility must reside on the OVPM server disk. If a client system wants to access one of their own files, they must arrange to have it transported to the OVPM server system.

Example File:

MDY TIMEp
,TempOutside,WindChill,HiTemp,LowTemp,WindSpeed,WindHi,"Win
dDir",Rain

1/01/05 12:00a, 41.3, 38.6, 41.9, 40.8, 5.0,
16.0, ESE ,0.04

1/01/05 2:00a, 41.6, 41.6, 42.1, 38.9, 4.0,
10.0, ESE ,0.00

1/01/05 4:00a, 39.5, 39.5, 40.4, 38.2, 3.0,
8.0, E ,0.01

1/01/05 6:00a, 40.3, 40.3, 40.5, 40.0, 2.0,
7.0, E ,0.00

1/01/05 8:00a, 40.1, 40.1, 40.6, 39.4, 2.0,
5.0, E ,0.00

Flat File Format details:

- Each file contains ASCII data consisting of one or more numerical and textual fields.
- Each line in the file represents one data sample, except for the first line which is a header line.
- Blanks at the beginning or end of a field are ignored.
- Blank lines and lines starting with a pound sign "#" are ignored.

Header Line:

- The first line in the file is a header line. It identifies the fields in the file. Each field in the header line contains the date/time specifier or a metric name. These fields are separated by a delimiter character (tab or comma). A metric name enclosed in double quotes indicates that this field contains textual data. An unquoted metric name indicates the field is numeric.
- Metric names should not contain any delimiter characters (space, tab or comma).
- The first field in a flat file must be the date/time field. The date/time specifier in the header line consists of the characters "Y","M","D" in the order of the date field:
(YMD = Year, Month, Day)
(MDY = Month, Day, Year), etc.

The time specifier consists of the string "TIME" and optionally an AM/PM indicator. If time is in 24 hour format then "TIME" is all you need. If the time is in 12 hour format then enter the indicator that will be found in the time strings to indicate "PM" (For example: TIMEPM). Note that TIME is not a separate column but is considered a part of the DATE/TIME column.

Some sample date/time header fields:

Header	Data
MDY TIMEPM	12/31/05 11:59 PM
YMD TIME	2005-12-31 23:59
MDY	12.31.2005

Data lines:

- Each field is separated from the others by a delimiter character (tab or comma).
- Textual fields that contain delimiters must be enclosed in double quotes. If no special characters are in the field then the double quotes are optional.
- Numerical data can be integer or floating point but will be treated as floating point inside of OVPM.
- The first field in each line must be the date and optionally the time of the sample. Dates can be specified in any order of Year, Month and Day separated by non-numeric, non-delimiter characters (as specified in the header line).

Time format is Hour then Minute, separated by a non-numeric, non-delimiter character. If time is in a 12 hour format, then a trailing AM/PM indicator can be present.

- There is no limit to the number of date lines than can be accepted except that performance will suffer as the number of data lines is increased. 10,000 lines is a reasonable expectation for the maximum practical business units file.

File names and locations:

For ease in locating and processing the flat files, they should be named with no embedded spaces and terminate with the file extension ".bus" (for Business Units). Example: "MyOwnData.bus".

By default, OVPM will look for .bus files in the logon customer's home directory.

`<data_dir>/conf/perf` for the default logon
`<data_dir>/conf/perf/VPI_Cust_logonname` for the "*logonname*" customer logon.

Other file names and locations are possible, but will not have as much support in the OVPM user interfaces.

Customers, Users and Logons

OVPm allows you to configure customers or users as a way of segregating data to only be viewed or modified by the individual customer or user. This is controlled via **logons** with optional **password**. The configuration of customers and users can serve a variety of purposes:

- Data may be segregated so that one OVPm **Customer** is not allowed to see the data for a different customer. This is most often used in situations such as ISPs where one OVPm might be shared by different customers. The use of Customers is required for OVIS access and is optional for HP OpenView Reporter access. The use of customers to segregate data for individual system access is also optional for OVPm by itself. Customers can be protected by passwords to prevent unauthorized access.
- When someone using OVPm creates a new graph, they can save it on the server system. This graph is saved in a particular area (sandbox) based on their logon. If everyone is using the same logon, then they can all see, modify, and delete any user created graph definition. Each customer is a different logon so each customer has a separate sandbox for their graphs. You can create additional unique logons by defining **Users** and thus give each person their own sandbox. In this way, one person will not see or be able to modify another person's graph definitions. Users can be protected by passwords to prevent unauthorized access.
- You may have people using OVPm that you want to prevent from using advanced OVPm features such as saving, or even temporarily creating graph templates. You can define **Users** with restricted access to accomplish these goals.
- You may have a combination of the above situations. In this case, you can use both Customers and Users at the same time.

Customers

OVPm data can be segregated into different **customers**. Each customer's data can be protected by an optional logon password. A customer will not be allowed to view another customer's data.

The way data is segmented by customer differs based on the data source. The **OVIS** data source segregates all its service data by customer names. The **HP OpenView Reporter** product allows grouping performance data gathered from different systems into customers. This is an optional feature starting with HP OpenView Reporter version 3.0. When OVPm is used to view data from OVIS or HP OpenView Reporter, it can honor the customer segregation used by those products.

The OVPA and OVO agents do not currently provide a customer segregation facility. OVPM can add customer segregation to systems with these agents in two ways:

1. If OVPM is configured to use the HP OpenView Reporter product and HP OpenView Reporter has segmented systems into different customers, then OVPM will segregate data from the agents on those systems in the same way.
2. OVPM can itself be configured to define customers and the systems that should be associated to each one.

The Default Customer

OVPM always defines a default customer that can access all of the data. This customer is specified as `All Customers` or by a blank customer name. This customer may be configured to require a password in order to prevent unauthorized access.

When using OVPM with OVIS, the OVPM blank customer password should be set the same as the OVIS "All Customers" password. This is especially true when OVIS is running in restricted views and when OVPM is installed on the same system as OVIS.

Using OVIS with Customers

OVIS always segregates its data by customer. Once OVIS is added as a datasource to OVPM, the OVIS customer names can be used in the OVPM logon to view that customer's data. Use the OVIS Configuration Manager to configure Restricted Views in order to set passwords for each customer. In order to enforce these passwords during the OVPM logon, set the password for the blank customer to a non-blank value using the OVPM Administrator Interface. Using this password will allow access to all OVIS data. If you logon to OVPM using a valid OVIS customer and password then you will see only data that belongs to that customer.

To enable OVPM Customer Segmented Data with HP OpenView Reporter:

1. Configure OVPM to use the HP OpenView Reporter database.
2. Customer passwords can be set using the HP OpenView Reporter user interface (HP OpenView Reporter 3.5 and later) or by using the OVPM Administrator Interface. In order to enforce these passwords during the OVPM logon, set the password for the blank customer to a non-blank value using the OVPM Administrator Interface.
3. Use HP OpenView Reporter to add customers.
4. Use HP OpenView Reporter to assign systems to individual customers.

NOTE: When OVPM is connected to HP OpenView Reporter, then the HP OpenView Reporter definitions of customers will be honored against OVPA and OVO-EPC data sources as well. You can use the OVPM Administrator Interface to add systems to an existing HP OpenView Reporter customer by defining the customer exactly the same as it is in HP OpenView Reporter. You cannot remove a system from a customer that HP OpenView Reporter has defined. You must use the HP OpenView Reporter interface instead.

Defining OVPA and OVO Agent Customers Directly in OVPM:

You can configure customers for systems running the OVPA or the OVO agent by using HP OpenView Reporter as described above. You can also configure customers for this system directly on the system where OVPM is installed using the OVPM Administrator Interface (see [Admin GUI](#)). If customers are defined both in HP OpenView Reporter and directly, then OVPM will use a combination of both configurations.

When you access OVPM using the Web Forms interface, you should see two fields requesting the Logon (customer name) and password. Return to the **Home** web page to specify a different customer, user or password. When you run the Java interface, you should see a **Logon** button at the bottom left of the Display screen. Press this button to change the customer name, user and password. See the [Logons](#) section for more information on logon options.

Defining OVPM Users

Each unique logon to OVPM will have access to their own private area to save and edit graph definitions. When initially installed, OVPM has only one logon defined (blank customer) and all clients share the same area to save graphs (sandbox). If customers are defined, either by HP OpenView Reporter or using the OVPM Administrator Interface, each customer has their own sandbox. This means that a graph created by one customer cannot be seen or modified by a different customer. The pre-defined graphs included with OVPM can be read but not modified by any customer.

Define customers and users using the [OVPM Administrator Interface](#).

You can create additional sandboxes by defining **Users**. Each user will see the pre-defined graphs plus any graphs that they have saved. They will not be able to see or modify graphs created by any other user. User logons can be protected with passwords and can be configured for special capabilities inside OVPM.

A user can be attached to a customer, and thus will be restricted to seeing only data that belongs to that customer, or a user can be attached to the blank customer and be allowed to access all data.

Users can be defined with restricted access:

Normal	A user will be allowed to draw any pre-defined graph. They will also be allowed to edit and save new and modified graph definitions. All saved graph definitions will be placed into the user's private Sandbox.
Read Only	A user will be allowed to draw any pre-defined graph. They will also be allowed to edit and create graphs locally. They will not be allowed to save any graph definitions or modifications they have made.
Display Only	A user will be allowed to draw any pre-defined graph. They will not be allowed to create their own graph definitions or to modify any existing graph definition, not even locally. The user has access to the Display tab but not the Design

	tab.
Expert	Expert users have access to OVPM features of a more advanced nature. Advanced features include additional items in the Client Configuration feature and access to the Expert tab in the Java interface Design mode where you can configure more advanced graph design features.

Logons

When OVPM is configured with no customers and no users logons are not necessary. The OVPM home page and the Java interface should not present the logon option. Note that if Internet Services is configured for OVPM, it always uses customers so the logon fields will be present in this case.

If customers or users are configured for OVPM, either directly using the OVPM Administrator Interface or indirectly through HP OpenView Reporter or Internet Services, then OVPM should allow you to enter a logon and password. For the Web Forms interface, two fields are added to the OVPM home page in the Web Forms section for entry of Logon and Password. For the Java interface, a **Logon** button is shown at the lower left of the first (Display) window. Pressing the **Logon** button displays a dialog where the logon and password can be entered.

Each logon has its own area to save graphs. Clients using the same logon will share the same graph area.

Logons can include the following:

Logon	Password	Comments
A blank	The password for the blank customer	Can see all data
customer name	customer password	Can see only this customer's data
user name	user password	If user belongs to a customer, then can see only this customer's data. If user belongs to the blank customer (admin user) then can see all data.

Logons, simplified use cases

OVPM allows flexibility in the way that customers and users are defined. It may be simpler to choose one of the schemes below that most closely matches how you want to use OVPM.

Wide Open

If you aren't worried about security and only a few clients will be using OVPM; or, if the clients won't be creating and saving a lot of their own graph definitions, you can use the default settings for OVPM. In this case, there are no customers

defined, all clients get to see all the data that is available and no logons are required.

To implement this scheme, set the blank customer password to blank using the OVPM Administrator Interface.

Customers Only

If you want to segregate data into different groupings (customers) but you aren't worried that different clients at a customer will require their own graphs, you can use customers only. Clients of one customer will not be allowed to see graphs designed by a different customer.

To implement this scheme, define the customers and the data they can access using the OVIS or HP OpenView Reporter products, or using the OVPM Administrator Interface. Assign customer passwords or not, as you desire. For each logon, instruct your clients to use their customer name and customer password. You should take care to set the password for the blank customer in the OVPM Administrator Interface to a non blank. Otherwise, any client could logon to any customer using the blank password.

Users Only

If your data doesn't break out into different groupings, but you would like different users to have private areas to save their own graphs, then define users only. Use the OVPM Administrator Interface and enter the names of all your users under the blank customer. Instruct your clients to logon using their user name and user password. They will all be able to see all the data in OVPM, but they will have separate areas to save their own graphs.

Customers and Users

Then there is the case where you need to set both customers and users. You should define how the data is segregated by customer using OVIS, HP OpenView Reporter, or the OVPM Administrator Interface. Define users in the OVPM Administrator Interface and assign passwords as you desire. You can have a mix of admin users defined under the blank customer name that can access all the data plus users that are defined under a specific customer and are only allowed to see that customer's data. Take care to use unique names for all users and customers. Tell your clients to use their user name and user password (for users with private graph storage areas or special restrictions) or use the customer name and customer password for a graph storage area shared by other clients for this customer.

Security

OVPM security can be configured in a variety of ways:

1. **Physical access to the OVPM server.** Certain sensitive OVPM functions require physical or logon access to the system where OVPM is installed. Depending on the operating system, additional system logon security may also be required. For example, reconfiguring OVPM on a Unix server requires logon to the system with root or superuser credentials.
2. **Network access to the OVPM server.** Access to the OVPM server system as a client requires network access. This access can be restricted through the use of segregated networks or firewalls.
3. **Client authorization lists.** OVPM can be configured to use [Client Authorization Security](#) where each attempt to access the OVPM server is cross-checked against a list of authorized system names or IP addresses. Any client not in this list is denied access.
4. OVPM [Customer and User logons](#) plus passwords can restrict access and/or restrict data visibility.
5. Changing the OVPM server [default port number](#).
6. Using [Secure Sockets \(https\)](#) communication from the clients.

Enabling Client Authentication Security

OVPM allows optional authentication of client connections, for example by Web Browsers or by Service Information Portal (SIP). The authentication capability allows you to specify which clients are allowed to connect to OVPM.

- Unauthorized users will get the error message:
Unauthorized access from "[IP Addr] hostname".
- This information is also logged to the OVPM Status log file (`<data_dir>/log/ovpm.0.log`) on the OVPM server.

Authentication is enabled by the presence of a file called `authip_vpi` in the data directory (`<data_dir>/conf/perf`) on the OVPM server. The `authip_vpi` file lists hosts from which connections are to be permitted.

- In the absence of this file, no authentication is performed; that is, any client will be allowed to connect.
- If the `authip_vpi` file exists in the data directory, then its contents determine which hosts are permitted client connections. Clients running on the same host as OVPM are automatically authenticated, and so do not need an entry.
- An empty `authip_vpi` file dictates that only clients running on the OVPM server can connect.

The OVPM client authentication capability requires that your network be able to resolve the client entries in the `authip_vpi` file. Depending upon the nature of the entries, this may require name services such as those provided by DNS, or NIS. A good test is to ensure that you can successfully *ping* each `authip_vpi` entry from the OVPM host. Client authentication will work through a firewall, with the same requirement that the client entries in the `authip_vpi` file be pingable from the OVPM host.

The `authip_vpi` file must conform to the following format:

- One client host may be listed per line.
- Client entries can be in any one of the following formats:
 - Fully qualified domain name
 - Alias
 - IP address in IPv4 dotted quad format
 - Wildcarded domain name (with * matching anything in the same position)
 - Wildcarded IP address (with * matching anything in the same position)

- Client entries may start in any column.
- Client entries can have no embedded spaces.
- Client entries can have trailing spaces.
- Anything following a # in a line is taken as a comment and ignored.
- Blank or zero-length lines are ignored.

Thus, given a system with:

```
IP Address:      123.456.789.1
alias:          testy
network address: testbox.group1.thecompany.com
```

any one of the following entries in the `authip_vpi` file would enable clients from the "testbox" host to connect:

```
123.456.789.1
123.*.*.*
testy
testbox.group1.thecompany.com
*.*.thecompany.com
```

Changing the Web Server Port Numbers

If you want to change the http Port number for the OVPM web server, you should first stop the OVPM web server `ovpm stop` then use the `ovpm config` command as shown below:

```
ovpm stop
ovpm config -port <portno>
ovpm start
```

When changing the web server port number, take care to choose a port number that is not being used by another program on the system.

The default port for Secure Sockets Layer (SSL or https) is port 8443. To change this port number is only slightly more involved. Edit the file `<install_dir>/OVPM/conf/ovpmsserver.xml` on Unix or `<Tomcat_dir>/conf/server/xml` on Windows and replace all occurrences of 8443 in the "SSL Coyote HTTP/1.1 Connector" entry with the desired port number. When finished, issue the `ovpm config` command (no parameters are required) and restart `ovpm`.

Note:Note that with OVPM 5.0 and beyond, the OVPM web pages will be returned using the same server name that was used to initially contact OVPM. In this way, an OVPM server can be accessed by different names on each side of a name translating firewall without having to reconfigure OVPM. This mechanism also allows Secure Sockets Layer (SSL) communication to OVPM by specifying https: instead of http: and using the SSL port number (8443 by default.)

Using Secure Sockets Layer (SSL or https)

When OVPM is installed, it configures itself to allow both http and https access. This involves enabling port 8080 for http and port 8443 for https. See [Changing the Web Server Port Numbers](#) for information on how to change these port numbers.

The OVPM Administrator Interface requires secure communications in order to safely transport OVPM configuration information. If you don't want to leave https enabled, then you can enable it prior to running the OVPM Administrator Interface then disable it afterward. To disable secure communications;

On Unix Edit the file

`<install_dir>/OVPM/conf/ovpmsserver.xml` Comment out the section for the "SSL Coyote HTTP/1.1 Connector". Then issue the commands `ovpm config` and `ovpm restart`.

On Windows Edit the file `<Tomcat_dir>/conf/server.xml`

Comment out the section for the "SSL Coyote HTTP/1.1 Connector". Then issue the command `ovpm restart`.

To start OVPM using http:

```
http://<system>:8080/OVPM
```

OVPM should continue to run in normal (non-secure) mode.

To start OVPM using https:

```
https://<system>:8443/OVPM
```

OVPM should continue to run in secure mode.

Using SSL or https with the Tomcat application server requires obtaining a server certificate for Tomcat. This involves running the Java SDK program called

keytool. Instructions for this procedure are in the default Tomcat `server.xml` file. The certificate should be obtained when OVPM is installed.

You may decide to replace the certificate obtained at OVPM install for a variety of reasons:

- When connecting to OVPM using https, the certificate is compared against the system name to insure that they match. If the names are not exact, then most browsers issue a warning so that the user can decide whether to continue or abort the connection. If the name used by the OVPM install isn't the same as that used by most of your https users, then you might want to get a different certificate that does match.
- The certificate allows many fields for providing information about who issued the certificate and is warranting that what it enables is safe. The OVPM install certificate leaves most of these fields blank. The fields are provided to users to help them decide whether to continue with the https access or not. You may want to fill in these fields with your own information.
- The certificate password is the default as suggested by the Tomcat documentation. You might want to change it to avoid unauthorized changes.

To replace the tomcat SSL certificate:

1. Remove the existing certificate.
`<java_dir>/bin/keytool -delete -alias tomcat`
The password is changeit.
2. Create a new certificate.
`<java_dir>/bin/keytool -genkey -alias tomcat -keyalg RSA`
Using the password changeit for both the certificate and the keystore.
3. Restart the OVPM server.

Alternate Ways to Access OVPM

Normally OVPM is accessed using a web browser. Simply access the URL to the OVPM home page:

`http://<system>:8080/OVPM`

where `<system>` is the actual system name where OVPM is installed.

From there the user can access online help, and either of the two standard OVPM user interfaces, Web Forms interface and Java interface.

- Additional URLs exist that may be used to provide [shortcuts and additional features](#)
- Graphs and tables can be created directly from the OVPM server through direct [commands or script files](#).

- Data from any data source can be [exported to text files](#) for further processing by another program.

Linking Directly to OVPM

You can run OVPM from a web browser on the systems where OVPM is installed or on other systems. Simply access the URL to the OVPM home page:

`http://<system>:8080/OVPM`

where <system> is the actual system name where OVPM is installed.

To access OVPM using a secure connection, use https:

`https://<system>:8443/OVPM`

`http://<system>:8080/OVPM/PMJavaUI.htm`

Starts the Java interface directly, bypassing the OVPM home page.

`http://<system>:8080/OVPM/ReleaseNotes.html` Displays the OVPM Release Notes.

`http://<system>:8080/OVPM/Help/C/Administrator.htm` Displays this document, the Administrator's manual, in web page format.

`http://<system>:8080/OVPM/Analyzer?-GETNEWFORM:%20-`

`NUMBEROFMETRICLINES:%201%20 -`

`FORMTEMPLATE:%20VPI_StandardForm.txt%20-`

`GRAPHTEMPLATE:%20"VPI_GraphsAgents.txt"%20 -ONLYNONDSN:%20-`

`SKIN:%20%20-CUSTOMER:%20" "%20-PASSWORD:%20" "`

is quite a mouthful but it runs the Web Forms interface directly. By slightly modifying this URL, you can exert additional control over the web form.

<code>-GRAPHTEMPLATE:%20"VPI_GraphsAgents.txt"</code>	Specifies the initial graph template to display.
<code>-SKIN:%20"Shifts"</code>	Specifies using the "Shifts" style of web form.
<code>-SKIN:%20"StartDate"</code>	Specifies using the "StartDate" style of web form.
<code>-CUSTOMER:%20"customer"</code>	Specifies the logon name (this can be customer or user name)
<code>-PASSWORD:%20"password"</code>	Specifies the logon password
<code>-FORMTEMPLATE:%20VPI_DesignForm.txt</code>	Specifies starting with the Design Tab instead of the Display Tab. You should probably also specify <code>NUMBEROFMETRICLINES:%208</code> to specify 8 metric lines instead of only 1

Running Graphs in Batch

Normally, OVPM is driven through the web server by requests from web browsers or other applications. In some cases you might want to have OVPM pre-generate graphs so that they can be handled in other ways. You can periodically generate graphs and email them or move them into other applications. Graphs that take an exceptionally long time to create may be generated in off-peak times, then accessed directly. This section describes how to run OVPM from a command line interface or script.

Summary of technique:

Run the `ovpmbatch` script and supply parameters on the command line to select the graph and other necessary information. The minimum amount of information required is a graph template file and graph name. Depending on the datasources used in the graph, additional information such as system name may be required. For example:

```
<install_dir>/bin/ovpmbatch  
GRAPHTEMPLATE=Agents GRAPH="Global History"  
SYSTEMNAME=testsys1
```

where:

`GRAPHTEMPLATE=Agents` is the graph template file.
`GRAPH="Global History"` is the graph name. It is enclosed in double quotes since it contains a blank.
`SYSTEMNAME=testsys1` specifies the system where the graph gets its data. This parameter may be required or not, depending on the graph.

The output of the `ovpmbatch` script is written to a working area that is available through the OVPM web server. The URL to reference this graph is displayed when the `ovpmbatch` command completes successfully. Example:

```
Location:  
http://localhost:8080/OVPM/datafiles/VPI_All_Cu  
stomers/ Global_History/7_day/admin/graph.htm
```

To locate the output directly, replace the first part of the string, up to `datafiles/` with "`<data_dir>/datafiles/ovpm`" on Unix or "`<install_dir>www\webapps\OVPM\datafiles`" on Windows. In this example,

```
<data_dir>/datafiles/ovpm/VPI_All_Customers/  
Global_History/7_day/admin/graph.htm
```

is the name of the HTML for the web page. Additional files will be in the same directory.

You can direct the output to a directory of your own choosing by specifying a parameter `outputdir=<directory>`. The directory will be created if it does

not already exist. If `outputdir` is specified, then any output that would usually go to `stdout` will go to a file named `stdout.txt` in the specified directory.

See [keywords](#) for details on other parameters.

HTML Graphs

When `ovpmbatch` creates a regular graph, it will create an HTML file and any necessary png image files and then display the location of the HTML file.

With this command entered:

```
ovpmbatch parameters
```

ovpmbatch outputs:

```
http://<servername>:<port>/OVPM/.../graph.htm
```

This location is in a format that can be used to access the web page from Output Files and Disk Space for more details on how to use this information.

Should another user generate the very same graph, the graph files will be replaced. As a best practice, either specify an `outputdir` or move the files immediately after they are created.

HTML Tables

`ovpmbatch` writes table output in HTML text to the `graph.htm` file.

Java Graphs and Tables

You can use the `ovpmbatch` facility to create and display a Java graph or table. the display of the graph to the proper system. When displaying a graph, the `ovpmbatch` command will wait until the graph is closed before continuing.

To display a Java graph or table, use the regular parameters to generate a graph or table, then add one more parameter: `JAVAGRAPH=2`.

```
export DISPLAY=mysystem:0.0
<install_dir>/bin/ovpmbatch
GRAPHTEMPLATE=Agents GRAPH="Global History"
SYSTEMNAME=testsys1 JAVAGRAPH=2
```

Exporting to Text File

A special feature of OVPM has been provided specifically for batch operation. Data for a graph or table can be written in text file format by overriding the graph type. The following types of text files are provided.

XML When OVPM creates XML output it will print it to standard out. To use this stand-alone, you should redirect stdout to a disk file.

```
ovpmbatch parameters >myfile.xml
```

or

```
ovpmbatch outputdir=mydir parameters  
(file stdout.txt is created)
```

The first line left in the file should be `<?xml version="1.0" ?>`.

You can force output from any regular graph into XML by including the parameter `GRAPHTYPE=XML` in the `ovpmbatch` command line.

CSV Comma Separated Values (`GRAPHTYPE=CSV`). Each item in the file is separated from the previous item by a comma. This form is acceptable for importing to many other applications such as spreadsheets and databases. If your locale formats numbers using commas as the decimal separator, then for CSV output, the decimal separator will be replaced with a period to avoid confusion with the field separators.

TSV Tab Separated Values (`GRAPHTYPE=TSV`). Each item in the file is separated from the previous item by a tab character. This form can also serve as input to many applications and it is more visually appealing as the data tends to align into columns for viewing.

In these text file types, the end of a record is specified by a newline character. Numeric values are displayed as digits, possibly with a decimal separator. Text values are enclosed in double quotes. Date and time values are displayed according to the settings on the OVPM Server for local date and time.

To specify text file output, include a parameter `GRAPHTYPE=CSV` or `GRAPHTYPE=TSV`. Output is written to standard out so you should redirect it to a disk file. For example, enter the `ovpmbatch` command and then the parameters as shown below:

```
<install_dir>/bin/ovpmbatch  
graphtemplate=Agents graph="Global History"  
SYSTEMNAME=testsys1 GRAPHTYPE=CSV >myfile.csv
```

```
<install_dir>/bin/ovpmbatch  
graphtemplate=Agents graph="Global History"  
SYSTEMNAME=testsys1 GRAPHTYPE=TSV >myfile.tsv
```

The number of records that can be processed by a web browser is limited. Graphs cannot be drawn with more than 1024 points for HTML and 2048 points for Java. Most web browsers tend to balk at tables with more than a few thousand records. For this reason OVPM limits the number of records in a table to no more than 5000. Individual graphs and tables can specify a limit using the `NUMBEROFPOINTS` keyword. You can override this value by including `NUMBEROFPOINTS=value` in your parameters to `ovpmbatch`.

Since text file output is not constrained by web browser performance, the record limit for CSV and TSV output can be increased to as many as one million records. To increase the maximum number of records exported, specify the parameter `NUMBEROFPOINTS=` followed by the desired number from 1 to 1000000. You may need to specify the `POINTSEVERY=` parameter to get the summarization level you desire.

OVPM Maintenance

Administrative tasks include managing output files and disk space for your OVPM system. Some of the daily maintenance is automated.

Managing Output files and Disk Space

When HTML graphs are generated, the HTML and PNG files are written to a directory based on the parameters used to generate the graph. The location of the HTML file is written to standard out when running `ovpmbatch` unless the `outputdir` parameter is specified. The command:

```
<install_dir>/bin/ovpmbatch
graphtemplate=Agents graph="Global History"
systemname=mysys.home.com
```

displays a line like:

```
Location:
http://localhost:8080/OVPM/datafiles/VPI_All_Cu
stomers/
GlobalHistory/7_Days/mysyshomecom/graph.htm
```

You can invoke the web page by asking a browser to display the URL consisting of everything after Location:.

To locate these files on your system disk, replace the first part:

```
"Location:
http://localhost:8080/OVPM/datafiles/"
```

with: "`<data_dir>/datafiles/ovpm/`" for Unix systems or
with:

```
"<install_dir>\www\webapps\OVPM\datafiles" for
Windows.
```

In the example shown above, the files would be located in:

```
<data_dir>/datafiles/ovpm/VPI_All_Customers/  
GlobalHistory/7_Days/mysyshomecom/
```

- The HTML file is named graph.htm
- The first graph image is in file graph.png
- Additional graph images, if present, will be named graph1.png, graph2.png etc.

These files are deleted periodically to reclaim disk space. OVPM will perform daily maintenance to reclaim disk space for files and directories that are no longer needed and other cleanup such as closing any datasources that haven't been accessed in the last 24 hours.

This maintenance is scheduled to run each night at 3:00 AM local time. See [Daily Maintenance](#) for more information.

Daily Maintenance

OVPM will perform daily maintenance to reclaim disk space for files and directories that are no longer needed and other cleanup such as closing any datasources that haven't been accessed in the last 24 hours. This maintenance is scheduled to run each night at 3:00 AM local time.

The graph files described in the section [Managing Output Files and Disk Space](#) will be automatically cleaned up if they have not been accessed in 24 hours. If you wish to preserve any of these files, you should back them up to a different directory.

The disk space cleanup will also be done for customers and users that are no longer configured as valid customers or users. For each customer and user, there are directories that contain the customer or user graphs. These directories and the files that reside in them will be automatically removed if the customer or user is no longer configured.

If you wish to preserve these files after deleting a customer or user, they are located in the following directories.

- For each customer, there is a configuration directory
<data_dir>/conf/perf/VPI_CUST_<customername>
- For the blank customer, the user directories are in
<data_dir>/conf/perf/VPI_USER_user
- For users for a specific customer, the user directories are in
<data_dir>/conf/perf/VPI_CUST_<customername>/VPI_USER/user

The following messages will be logged to the OVPM status file (see the [Troubleshooting the OVPM Server](#) section for information on how to check the status file).

```
Perform daily datasource maintenance:
Perform Refresh Config
Found 145 systems unique to the HP OpenView Reporter
database in 0.54 seconds
End of datasource maintenance: cleaned up 8 files/directories,
17561 bytes.
```

Advanced Work With Graph Template Files

Most graph specifications can be entered using one of the OVPM user interfaces. There are still a few specifications that are not supported by these interfaces and require [manual editing of the graph template files](#).

Even in this case, the majority of the work in creating a graph template should be done

using the standard OVPM user interface with manual editing reserved for just those

items that require it. These features are typically advanced in nature and require a higher level of understanding of how OVPM works than is expected for the typical

OVPM end user.

When using the OVPM user interfaces, graphs are saved into an area that allows the user

that created them to modify or even delete them. If it is desirable to make a user generated graph available to all users, but in a read-only mode like the OVPM pre-defined

graphs, then the OVPM Administrator can [move a user defined graph into permanent graph area](#).

The OVPM server is driven using a collection of parameters. These parameters are what

make up the graph definition templates and are used in the URLs for generating new graphs, as

well as being used in the batch facility. A proper definition of the [OVPM keywords and their uses](#) is helpful in these areas.

Editing Graph Template Files Directly

Typically you would use the Java or Web Forms interface to create graphs. You can create completely new graphs, or you can also start from a predefined or a user-created graph, modify it as needed and save it as a new graph.

Note that you cannot replace or delete any of the predefined graphs. Users from the web can only modify graph templates that have been saved using the Web interface.

OVPm also allows you to create graphs using a text editor such as vi or notepad, in those rare cases when you need to use features not supported by Java, such as calculated metrics. You can use these features only by manually editing the graph definition. While you can create graphs entirely manually, it saves time to create the basic graph definition in the user interface, save it, then add the more complex features manually.

For users defined as expert (see [Defining Users](#)) many of these features are available from the Java interface Expert tab in Design mode.

The graph templates are all named `VPI_Graphs*.txt` and are located in different directories on the OVPm server system:

Pre-defined graphs	<code><install_dir>/newconfig/OVPm</code>
User graphs for the blank logon	<code><data_dir>/conf/perf</code>
Users graphs for the blank customer, non-blank user	<code><data_dir>/conf/perf/VPI_USER_user</code>
Customer graphs	<code><data_dir>/conf/perf/VPI_CUST_customer</code>
User graphs for users belonging to a customer	<code><data_dir>/conf/perf/VPI_CUST_customer/VPI_USER_user</code>

The web browser interface can read from all of these directories, but can only write files to the `<data_dir>` directory. As a convenience for the Web Forms interface, all files in the data directory will start with "VPI_GraphsUser" so that they are displayed in the Web Forms interface as "User...". This prevents the user from creating a graph template file with the exact same name as any predefined graph template file. The Java interface does not display the graph template filename but instead displays a string called "Family".

Caution:

If editing graph template files by hand, you must take care to preserve the following conditions:

- All graph names must begin with `VPI_Graphs` and have the `.txt` extension at the end. Any files not so named will not be recognized as graph template files. The files must reside in either the `newconfig` or the `data` directories as described above.
- No graph template file in the `newconfig` directory should have a name that begins with `VPI_GraphsUser` since this prefix is restricted for use by user-created files in the `data` directory.

- Graph template files in the `data` directory may be created without the `VPI_GraphsUser` prefix (using `VPI_Graphs` only), but then make sure that graph template files with the same name do not exist in the `newconfig` and the `data` directories.
- Avoid creating two graph template files that use exactly the same string in the `FAMILY: parameter` inside the file. While OVPM can handle this situation, it can be confusing in the Java interface when you have two identical graph families.
- **Warning:** Avoid updating the `VPI_Graphs` template files in the `newconfig` directory. If you wish to change the default graphs, make a copy using a different file name, then use the new file as your template. This is because these default files will be replaced with product updates.

The format of a graph template file is a plain text file. Lines should be less than 2000 characters in length and terminated by a newline character. No other formatting information should be saved in the file. Lines in the template file fall into one of the following categories:

Comments	Blank lines or lines starting with a pound sign # are treated as comments and ignored.
Single Line Keywords	Most lines contain a keyword (all upper case, ending with a colon :), one or more blanks, then optionally a value.
Multi Line Keywords	<p>Some keywords allow values covering multiple lines. One keyword starts the entry (such as <code>GRAPHHEADER:</code>) and all successive lines until the terminating keyword (such as <code>END_GRAPHHEADER:</code>) are considered part of the value. The keywords must be on their own line, e.g:</p> <pre> GRAPHHEADER: This is my graph header END_GRAPHHEADER: </pre> <p>The following will not work:</p> <pre> GRAPHHEADER: My header END_GRAPHHEADER: </pre>

The keywords that define a single graph or table fall between the `GRAPH:` and `END_GRAPH:` keywords. Some keywords are allowed before the start of the first graph or table definition. These keywords apply to the graph template file as a whole (`FAMILY:` for example).

Refer to the [Keywords](#) document for a description and usage of all keywords allowed in a graph definition file. You might also want to examine the predefined graph templates supplied with OVPM to gain more insight into how to create graphs.

Making User-Designed Graphs Permanent

Graphs that have been designed and saved using the OVPM user interface are written to the `data` directory (the [table](#) in the [Editing Graph Template Files Directly](#) section, describes how to locate files for various logons.

Any graph definitions in the `<data_dir>` directory can be modified or deleted by any user of OVPM providing they know the proper logon and password. An OVPM user may not read, modify or save graphs that belong to a different logon.

Predefined graphs are not modifiable using the OVPM user interface. They are located in the `newconfig` directory that can be read, but not modified by OVPM interface users.

You can "promote" a user-designed graph to a non-modifiable state by simply moving it from the `data` to the `newconfig` directory. However, whenever you do so, you must follow these rules:

- These actions can only be performed by a user that is logged in to the system where OVPM is installed. The actions are performed using system commands and a text editor (such as vi or notepad) and cannot be performed over the web. This insures that anyone modifying "non-modifiable" graph templates has passed the physical and logon security of the OVPM system.
- All the graphs in a graph template file have the same security. You cannot secure some graphs and not others unless the graphs are in different files. If you only want to secure one graph, then you should use a text editor to copy it into a different file.
- If you modify any of the predefined graphs, please keep a backup copy for troubleshooting purposes. Make a copy and change the file name so that it does not match the `VPI_Graph<name>.txt` naming scheme and the original `VPI_Graph<name>.txt` will "disappear" from the OVPM user interfaces. See [warning](#).
- No graph template file in the `newconfig` directory should have the same name as a file in the `data` directory. Since all graph template files created by the Web interfaces are forced to have a name of `VPI_GraphUser<name>.txt`, avoid creating any graph templates in the `newconfig` directory that have the User designation.

Do NOT just copy the user graph template file from the data directory to the newconfig directory without renaming it to remove the User part of the filename.

- It is generally not a good practice to use the same FAMILY: name in two different graph template files. It can be confusing in the Java interface.

To move graphs:

1. Once a graph has been saved, open the user template file using a text editor such as vi or notepad. The user template file is saved into the `<data_dir>` directory based on the user logon. (See [table](#))
2. Locate the desired graph and copy all text between and including the GRAPH: and the END_GRAPH: lines.
3. Create a new file, or add this text to an existing file in the predefined graphs area (newconfig). This file should be called `VPI_Graph<name>.txt` where `<name>` identifies the general areas of all the graphs in the file. See [warning](#).
4. Save the file into the `<install_dir>/newconfig/OVPM` directory, not the `<data_dir>/conf/perf` directory where the user template file is located.

In summary: Locate the user-generated graph in the data directory, make the modifications to filename and family as described above, then save it in the newconfig directory. The change should be immediately apparent to the web users, although currently active users may need to refresh their display to see it.

Customizing the Look of the OVPM Interface

While OVPM comes fully configured, it does allow for customization of certain aspects to adjust the look and feel to individual tastes. Many such changes can be set for individual logons. The OVPM Administrator can change:

- Default graph settings such as background color, subtitle, etc.
 - The appearance of the Java interface usually matches the client that is running it but you can force it to look like Windows, Motif, Metal.
 - The Web Forms interface can be altered to display your company logo, change the color scheme, or even to remove fields you don't use.
-

Using a Style File to Set Defaults

Style files override graph parameters or provide default values. They are used to reset parameters for all graphs without having to modify each graph individually. There may be several `VPI_Style.txt` files on the system but the one that is used depends on how you are logged on to OVPM. OVPM will search for the style file in the following order:

1. The logon user file in the user's "sandbox".
`<data_dir>/conf/perf/VPI_CUST_customer/VPI_USERuser/VPI_Style.txt`
2. The logon customer "sandbox".
`<data_dir>/conf/perf/VPI_CUST_customer/VPI_Style.txt`
3. The "localized skin" file in
`<install_dir>/newconfig/OVPM/language/VPI_SKINskin/VPI_Style.txt`
4. The localized default file in
`<install_dir>/newconfig/OVPM/language/VPI_Style.txt`
5. The default "skin" file in
`<install_dir>/newconfig/OVPM/VPI_SKINskin/VPI_Style.txt`
6. The default file in `<install_dir>/newconfig/OVPM/VPI_Style.txt`

Only the first Style file found will be used. All others will be ignored.

- If you want to set a style for everyone on the system, set it in the default file.
- If you want to set a style for all users in one customer, set it in the customer "sandbox".
- If you want to set a style for one user only, set it in the user "sandbox".
- If you want to set a style that applies only when a specific web-forms "Skin" is used, set it in that Skin directory.

The `VPI_Style.txt` file should contain [keywords](#) and their settings that will override all settings from the graph templates. **Caution:** Using keywords other than those provided in the following table could produce unpredictable results.

Examples of uses for a Style file include:

Setting a default graph and table background	GRAPHBACKGROUND: White
Disabling markers from all line graphs	NOMARKERS: or MARKERS: No
Disabling toolbar from Java graphs	NOTOOLBAR: or TOOLBAR: No
Setting the default datasource and graph template	DATASOURCE: IS GRAPHTEMPLATE: VPI_GraphsVPIS.txt
Setting the	JAVAGRAPHPOINTS: 2048

<p>maximum number of points in a Java Graph</p>	<p>The maximum number of data points in a Java Graph is 2048 points. This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Each graph can define its own NUMBEROFPOINTS value. The actual limit will be the smaller of the graph definition and this limit value.</p> <p>CAUTION: Setting this value to a large number can allow Java graphs to be generated with enough data points to cause memory problems on some clients. It can also delay the presentation of graphs when a large number of data points are processed.</p>
<p>Setting the maximum number of points in a Java Table</p>	<p>JAVATABLEPOINTS: 5000</p> <p>The maximum number of data points in a Java Table is 5000 points. This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Each graph can define its own NUMBEROFPOINTS value. The actual limit will be the smaller of the graph definition and this limit value.</p> <p>CAUTION: Setting this value to a large number can allow Java tables to be generated with enough data points to cause memory problems on some clients. It can also delay the presentation of tables when a large number of data points are processed. In some cases, tables with many rows and many columns could cause the application to consume large amounts of CPU and even to 'freeze'.</p>
<p>Setting the maximum number of points in an HTML Graph</p>	<p>HTMLGRAPHPOINTS: 1024</p> <p>The maximum number of data points in an HTML Graph is 1024 points. This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Each graph can define its own NUMBEROFPOINTS value. The actual limit will be the smaller of the graph definition and this limit value.</p> <p>CAUTION: Setting this value to a large number can allow HTML graphs to be generated with enough data points to cause memory problems on some clients. It can also delay the presentation of graphs when a large number of data points are processed.</p>
<p>Setting the maximum number of points in an HTML Table</p>	<p>HTMLTABLEPOINTS: 5000</p> <p>The maximum number of data points in an HTML Table is 5000 points. This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Each graph can define its own NUMBEROFPOINTS value. The actual limit will be the smaller of the graph definition and this limit value.</p> <p>CAUTION: Setting this value to a large number can allow HTML tables to be generated with enough data points to cause memory problems on some clients. It can also delay the presentation of tables when a large number of data points are processed. In some cases, tables with many rows and many columns could cause the application to consume large amounts of CPU and even to 'freeze'.</p>

Forcing the Java look and feel	JAVALOOK: system metal motif windows																				
Setting the default Graph SubTitle	GRAPHSUBTITLE: <i>string</i> Substitution variables can be placed anywhere in the graph subtitle string. They will be replaced by the actual values when the graph is drawn. <table border="1" data-bbox="524 569 1372 1094"> <thead> <tr> <th>Variable</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>@@[SYSTEM]</td> <td>The name of the system begin graphed</td> </tr> <tr> <td>@@DATERANGE</td> <td>Total time shown in the graph (7 Days)</td> </tr> <tr> <td>@@POINTSEVERY</td> <td>The amount of time summarized into each data point (10 Min)</td> </tr> <tr> <td>@@SHIFT</td> <td>The shift times (8:00-17:00 or "All Day")</td> </tr> <tr> <td>@@SKIPDAYS</td> <td>The list of days that are being skipped (SAT,SUN)</td> </tr> <tr> <td>@@SHOWDAYS</td> <td>The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)</td> </tr> <tr> <td>@@NOHOLIDAYS</td> <td>Indicates whether Holidays are being skipped or not (blank or "No Holidays")</td> </tr> <tr> <td>@@STARTTIME</td> <td>The beginning time in the graph (3/1/04 3:00 PM)</td> </tr> <tr> <td>@@STOPTIME</td> <td>The last time in the graph (3/8/04 3:00 PM)</td> </tr> </tbody> </table> To make the subtitle use smaller text and allow more characters to be displayed, add <small> at the beginning and </small> at the end. For example: <small>Points Every @@POINTSEVERY with Shift @@SHIFT @@SHOWDAYS @@NOHOLIDAYS</small>	Variable	Description	@@[SYSTEM]	The name of the system begin graphed	@@DATERANGE	Total time shown in the graph (7 Days)	@@POINTSEVERY	The amount of time summarized into each data point (10 Min)	@@SHIFT	The shift times (8:00-17:00 or "All Day")	@@SKIPDAYS	The list of days that are being skipped (SAT,SUN)	@@SHOWDAYS	The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)	@@NOHOLIDAYS	Indicates whether Holidays are being skipped or not (blank or "No Holidays")	@@STARTTIME	The beginning time in the graph (3/1/04 3:00 PM)	@@STOPTIME	The last time in the graph (3/8/04 3:00 PM)
Variable	Description																				
@@[SYSTEM]	The name of the system begin graphed																				
@@DATERANGE	Total time shown in the graph (7 Days)																				
@@POINTSEVERY	The amount of time summarized into each data point (10 Min)																				
@@SHIFT	The shift times (8:00-17:00 or "All Day")																				
@@SKIPDAYS	The list of days that are being skipped (SAT,SUN)																				
@@SHOWDAYS	The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)																				
@@NOHOLIDAYS	Indicates whether Holidays are being skipped or not (blank or "No Holidays")																				
@@STARTTIME	The beginning time in the graph (3/1/04 3:00 PM)																				
@@STOPTIME	The last time in the graph (3/8/04 3:00 PM)																				
Setting the default metric label format	DEFAULTLABEL: <i>string</i> The default metric label string is used when a graph is drawn that has not specified individual metric labels. The defaultlabel string sets the style for generating labels by including text, plus any or all of the following substitution variables: <table border="1" data-bbox="524 1528 1372 1801"> <thead> <tr> <th>Variable</th> <th>will be replaced by</th> </tr> </thead> <tbody> <tr> <td>@@[LABEL]</td> <td>The metric label specified by the data source. (CPU %) If the data source does not provide a label, then the metric name will be used.</td> </tr> <tr> <td>@@[METRIC]</td> <td>The metric name (GBL_CPU_TOTAL_UTIL)</td> </tr> <tr> <td>@@[CLASS]</td> <td>The metric class (GLOBAL)</td> </tr> <tr> <td>@@[SYSTEM]</td> <td>The system name supplying the metric</td> </tr> </tbody> </table>	Variable	will be replaced by	@@[LABEL]	The metric label specified by the data source. (CPU %) If the data source does not provide a label, then the metric name will be used.	@@[METRIC]	The metric name (GBL_CPU_TOTAL_UTIL)	@@[CLASS]	The metric class (GLOBAL)	@@[SYSTEM]	The system name supplying the metric										
Variable	will be replaced by																				
@@[LABEL]	The metric label specified by the data source. (CPU %) If the data source does not provide a label, then the metric name will be used.																				
@@[METRIC]	The metric name (GBL_CPU_TOTAL_UTIL)																				
@@[CLASS]	The metric class (GLOBAL)																				
@@[SYSTEM]	The system name supplying the metric																				

	<table border="1"> <tr> <td data-bbox="521 197 792 275">@[DATASOURCE]</td> <td data-bbox="792 197 1375 275">The data source for this metric (MWA, CODA, etc.)</td> </tr> <tr> <td data-bbox="521 275 792 527">@@metric</td> <td data-bbox="792 275 1375 527">The value of the <i>metric</i> metric from the same data source and class. For example, @@BYDSK_DEVNAME will label the metric with the value of the BYDSK_DEVNAME metric. If the BYDSK_DEVNAME metric value were "0", Disk @@BYDSK_DEVNAME would produce a label of Disk 0</td> </tr> </table> <p>EXAMPLE: @@[SYSTEM]:@[CLASS]:@[METRIC] might produce a label like "mysys.net.com:GLOBAL:GBL_CPU_TOTAL_UTIL" while the default @@[LABEL] might produce a label like "CPU %".</p> <p>Graphs that only contain one metric will automatically show the system</p>	@[DATASOURCE]	The data source for this metric (MWA, CODA, etc.)	@@metric	The value of the <i>metric</i> metric from the same data source and class. For example, @@BYDSK_DEVNAME will label the metric with the value of the BYDSK_DEVNAME metric. If the BYDSK_DEVNAME metric value were "0", Disk @@BYDSK_DEVNAME would produce a label of Disk 0																			
@[DATASOURCE]	The data source for this metric (MWA, CODA, etc.)																							
@@metric	The value of the <i>metric</i> metric from the same data source and class. For example, @@BYDSK_DEVNAME will label the metric with the value of the BYDSK_DEVNAME metric. If the BYDSK_DEVNAME metric value were "0", Disk @@BYDSK_DEVNAME would produce a label of Disk 0																							
Setting the Date Format	<p>DATEFORMAT: <DateString></p> <p>Normally dates are formatted according to the language on the system where the OVPM server is installed. You can override this format by entering a new one in this field. The format must contain the following three fields in order to be accepted. The order will determine the format for dates on graphs and tables plus the format of dates expected when inputting in a local format. Other characters will be repeated exactly as shown providing they are not one of the special formatting characters "GyMdhHmsSEDFwWakKz&apos;".</p> <table border="1"> <tr> <td data-bbox="521 1150 613 1266" rowspan="2">Year</td> <td data-bbox="613 1150 704 1228">y or yy</td> <td data-bbox="704 1150 1375 1228">Exactly two digits will be used. 1998 = "98" while 2005 = "05"</td> </tr> <tr> <td data-bbox="613 1228 704 1266">yyyy</td> <td data-bbox="704 1228 1375 1266">Four digits will be used.</td> </tr> <tr> <td data-bbox="521 1266 613 1413" rowspan="2">Month</td> <td data-bbox="613 1266 704 1344">M</td> <td data-bbox="704 1266 1375 1344">The Month will be shown without leading zeros. February = "2".</td> </tr> <tr> <td data-bbox="613 1344 704 1413">MM</td> <td data-bbox="704 1344 1375 1413">The Month will be shown with leading zeros. February = "02".</td> </tr> <tr> <td data-bbox="521 1413 613 1528" rowspan="2">Day</td> <td data-bbox="613 1413 704 1491">d</td> <td data-bbox="704 1413 1375 1491">The day of the month will be shown without leading zeros.</td> </tr> <tr> <td data-bbox="613 1491 704 1528">dd</td> <td data-bbox="704 1491 1375 1528">The day of the month will be shown with leading zeros.</td> </tr> </table> <p>Example: February 4, 2005 would be shown as:</p> <table> <tr> <td>Format</td> <td>Date</td> </tr> <tr> <td>M/d/yy</td> <td>2/4/05</td> </tr> <tr> <td>MM/dd/yy</td> <td>02/04/05</td> </tr> <tr> <td>M-dd-yyyy</td> <td>2-04-2005</td> </tr> </table>	Year	y or yy	Exactly two digits will be used. 1998 = "98" while 2005 = "05"	yyyy	Four digits will be used.	Month	M	The Month will be shown without leading zeros. February = "2".	MM	The Month will be shown with leading zeros. February = "02".	Day	d	The day of the month will be shown without leading zeros.	dd	The day of the month will be shown with leading zeros.	Format	Date	M/d/yy	2/4/05	MM/dd/yy	02/04/05	M-dd-yyyy	2-04-2005
Year	y or yy		Exactly two digits will be used. 1998 = "98" while 2005 = "05"																					
	yyyy	Four digits will be used.																						
Month	M	The Month will be shown without leading zeros. February = "2".																						
	MM	The Month will be shown with leading zeros. February = "02".																						
Day	d	The day of the month will be shown without leading zeros.																						
	dd	The day of the month will be shown with leading zeros.																						
Format	Date																							
M/d/yy	2/4/05																							
MM/dd/yy	02/04/05																							
M-dd-yyyy	2-04-2005																							
Setting the Time Format	<p>TIMEFORMAT: <TimeString></p>																							

Normally times are formatted according to the language on the system where the OVPM server is installed. You can override this format by entering a new one in this field. The format must contain the following fields in order to be accepted. The order will determine the format for times on graphs and tables plus the format of times expected when inputting in a local format. Other characters will be repeated exactly as shown providing they are not one of the special formatting characters "GyMdhHmsSEDFwWakKz'".

24 Hour	H	Hour of the day, based on a 24 hour clock (0-23) with no leading zeros.
	HH	Hour of the day, based on a 24 hour clock (0-23) with leading zeros.
12 Hour	h	Hour of the day, based on a 12 hour clock (0-11) with no leading zeros.
	hh	Hour of the day, based on a 12 hour clock (0-11) with leading zeros.
Minute	m	Minute of the hour with no leading zeros.
	mm	Minute of the hour with leading zeros.
AM/PM	a	AM/PM indicator (in the local language) This item is optional and should only be used when the hours are in 12 hour format.

Example: 8:05 PM (20:05 hours) would be shown as:

Format	Time
h:m a	8:05 PM
hh:mm a	08:05 PM
HH:mm	20:05

Many of these parameters may also be set by expert users using the [Client Configuration](#) feature. If a style file and a client configuration option are both set, the style file will take precedence. For this reason, if you want to allow client configuration options, you should not set corresponding values in the style files.

Using Skins to Affect Presentation Styles

As the OVPM Administrator, you can affect look and feel of the graphs and web forms seen when this product runs. The easiest way to do this is to create a new "skin". Multiple skins can exist on an OVPM installation to allow different users to see different appearances. Files for each skin must reside in a directory named `<install_dir>/newconfig/OVPM/VPI_Skin*` where * is the name of the skin.

Skins do not affect the look and feel of the Java interface, only the Web Forms Interface. Use the `VPI_Style.txt` file and the `JAVALOOK` parameter to control the Java interface look and feel.

The `skin` directory should contain only those files that should be different from the default files. The default files are all located in `<install_dir>/newconfig/OVPM`. The Style file and the Form Template files may be used in a skin directory.

- **Style** files can override settings on the graphs and tables. The default style is to obey all specifications in the graph definitions. Use a style file to set a common background color or disable markers on line graphs, etc.
- **Form** template files define the Web Forms that the users will see when they run OVPM.

Modifying the form template files in a skin can be complex. A form file is a mixture of HTML and special substitution variables. (Substitution variables all start with @@.) If you are experienced working with HTML then you might find it easy to replace the OVPM logo with your company logo. More complex modifications are possible but are beyond the scope of standard product use and thus cannot be supported by HP.

If you are adventurous and are willing to assume your own support, then here are some hints toward creating your own skin.

1. First find an existing skin that is close to what you want then copy its directory and all the files it contains to a new directory. Skin directory names all start with `VPI_Skin` and are found in the `<install_dir>/newconfig/OVPM` directory. Make sure to name your new directory similarly.
2. Locate the template file containing the web form that you want to change.

<code>VPI_StandardForm.txt</code>	The Display form is used to draw existing graphs.
<code>VPI_DesignForm.txt</code>	The Design form is used to design new graphs or to edit existing ones.
<code>VPI_PromptForm.txt</code>	The Prompt form is used when prompting for an instance name on a graph with <code>INSTANCE=*PROMPT*</code> .
<code>VPI_Graph.txt</code>	This file sets how HTML graphs are displayed.
<code>VPI_Table.txt</code>	This file sets how HTML tables are displayed.
<code>VPI_GraphJava.txt</code>	This file sets how Java graphs and tables are displayed.
<code>VPI_Style.txt</code>	The Style File sets defaults for many things. See the preceding section for details.
<code>VPI_Analyzer.txt</code>	This file exists only in the default directory <code><install_dir>/newconfig/OVPM</code> and is used to generate the OVPM home page. The home page is generated by the <code>ovpm config</code> command and also whenever returning to the home page from the Web Forms Display tab or Design tab.

3. Remove any template files that you don't want to customize and where the default file is acceptable.
4. Edit the files carefully.

Do not use an HTML editor to edit the template files. Their use of substitution variables can confuse most such editors. Instead use a regular text editor such as notepad or vi.

5. Test your results by selecting your skin in the OVPM home page. You may have to regenerate the home page first by entering one of the other skins then selecting Home. To troubleshoot the results, have the web browser display the web page source. This will show the HTML after the substitution variables have been replaced.
6. If you don't get the results you desire then please don't call HP for help. Seek the services of a consultant that is well versed in HTML and OVPM.

Using Client Settings to Set Defaults

The Client Settings feature allows expert users to set default values similar to the [Style File](#). While Style Files can affect all clients, the Client Configuration affects only accesses to OVPM from one client browser. These settings will take effect regardless of the OVPM logon used. Client Configuration information is preserved in a client "Cookie" and so is preserved from one OVPM session to the next. Cookies should be enabled in the client browser for the settings tab to work as expected. If cookies are disabled, the changes made in the settings tab will not be saved.

Some features in the Client Settings are only available to Expert users such as the ability to limit the number of points for graphs and tables. Once set, these values will be preserved for all accesses from this client, even if the logon is by a non-expert user. See [Defining Users](#) and [Administrator Interface](#) for more on how to configure users and to set or unset the "expert" flag for them. The default logon (blank) will not have the option to set the Client Settings parameters. Expert users should see the "Settings" button at the bottom of the Java interface or see a "Settings" tab in the web forms interface.

Once the client configuration is set, it will be preserved even if the OVPM logon is changed to one that does not allow access to expert mode features. For details about the parameters that can be set from the Client Settings, see the online help from the Client Settings screen.

If a parameter is set in both the client settings and in the active "Style" file then the value from the Style file will be used. In this way, a client setting cannot override a parameter once it is set in the Style file. Changing the OVPM logon or skin may select a different style file where the client settings value could then be accepted.

TroubleShooting

Whenever a problem occurs, the first step is to try to understand what any error message is trying to tell you. If you need more information, or if you are investigating a problem that someone else is having, there are several general steps you can take. If the general steps do not resolve the problem, look at the list of specific error situations described at the end of this section to see if you can recognize your symptoms then take the recommended actions.

Troubleshooting tips are broken down into two sections:

- Problems with the [client](#).
- Problems with the [OVPM server](#).

Troubleshooting the OVPM client

The following troubleshooting tips are focused on OVPM from the client perspective. Often problems can be caused by a misconfigured web browser or an improper plug-in on the client. Such problems are usually localized to a specific client.

OVPM Works From Some Clients, Not From Others

If you can run OVPM from a web browser on the same system that OVPM is installed on, but not from other systems then look for Network Connectivity or security issues.

- Check for personal Firewalls that could be preventing other systems from accessing the web ports on the OVPM server system.
- Check for external network security settings that prevent other systems from contacting the OVPM server. Can you "ping" the OVPM server from the client systems? Try doing a "telnet *system* 8080" to see if the OVPM port number is responding.
- If you are getting an *Unauthorized access from "[IP Addr] hostname"* message then see the section on [Enabling Client Authentication Security](#) for rules on using the `authip_vpi` security file.

Graphs are not Updated

The computer stores copies of frequently accessed pages in the Web browser's memory cache or disk cache. Users should set their Web browser to check for newer versions of stored pages by setting it to "every visit to the page". This setting ensures that the graph images will properly update on the system. Otherwise, the browser will cache the graph images and not update when you select minor changes and redraw your graph. Here's how to do it:

Internet Explorer

To specify how often Internet Explorer checks the network for page revisions:

1. Open IE browser. Select Tools -->Internet Options. Under Temporary Internet Files, click the **Settings** button.
2. Check the "Every visit to the page" radio button. Click **OK**.

Netscape and Mozilla

To specify how often Netscape and Mozilla check the network for page revisions (so that stale pages are not kept in the cache too long):

1. From the Edit menu, select Preferences.
2. Click the Advanced category, and then click **Cache**.
3. Select the "Every Time" radio button. The page you see will always be the latest one.

Problems with the Java Plug-in

Note that Java Plugins prior to 1.4.2 are no longer supported

- Invoking the Java graphs feature will require downloading the Java Plug-in to the system where the browser is executing. You may be prompted to download this plug-in the first time your browser invokes this feature. If not, you can manually download the plug-in from <http://java.sun.com/j2se/1.4.2/download.html>.
- To check the version of Java Plug-in you are using:
 - Enable the Java Console from your browser
 - Start (or attempt to start) the OVPM Java interface
 - The first few lines of the Java Console will list the actual Java plug-in version being used. This is often not the version that was expected when more than one plug-in is installed.
- On Unix systems, it is very important to follow the installation instructions provided during the download, and to make sure that the prerequisites are met in order to get the plug-in installed properly. Often a system reboot is required in order for the plug-in to begin working properly.
- On HPUX the plug-in can be downloaded from <http://www.hp.com/products1/unix/java/java2/jpi/downloads/index.html>.

The plug-in for HP-UX requires a special version of Netscape that can be downloaded from <http://www.hp.com/go/netscape>.

- The first time you draw a Java graph, you should see a message from your browser asking whether you want to accept the certificate of the Java applet.

This message is your assurance that the Java applet is a part of this product and has not been altered in any way. The message offers options to grant or reject (Deny) the applet permissions. Permissions can be granted for the current session only (Grant this session) or always (Grant always). It is best if you select "Grant always" then the message will not bother you again. If you deny permissions, the Java applet will still run, but some interactions with your system, such as copy and paste to the clipboard, will not be available.

Or this message may indicate that the security certificate has expired. Security certificates are only marked as valid for a short period of time. This message is only a warning that this time has passed. You should ignore this message and accept the certificate.

Java Screens Are Cut Off

This may happen if you are using an improper version of the Java plug-in. You may also see a wide gray band to the right of the display. The solution is to uninstall all versions of the Java Plug-in (use Add/Remove Programs from the Control Panel). Then, manually install the Java plug-in, version 1.4_2 or higher from the web.

No Data is Shown in a graph

It is possible to draw a graph that shows no data when you expect data to be displayed.

- If you get an error message indicating "Could not make connection to agent" or "The metrics in graph <graphname> were not found on this datasource", then the message indicates that either OVPM could not contact the system or datasource requested, or that it contacted it, but the list of classes and metrics available does not match those required in the graph definition.
- Make sure you have selected a system that provides the metrics for this graph and that the system is operational and accessible. If you are attempting to draw a graph provided by a Smart Plug-In, make sure the SPI is properly deployed and collecting data. You may have to wait as much as 10-15 minutes before data is available from a newly deployed agent or SPI.

- Use the **System Information** function to determine that the system you are attempting to use is responding and to determine what date ranges are available for each data class.
- Check the graph definition and compare it to the datasource you selected to insure they are compatible. Using the OVPM System Information function against the system you are attempting to draw data from can often provide clues. Check to see that the classes and metrics in the graph definition are being provided by the system and that the class contains data in the selected date range.

Check the graph definition using the Java interface:

1. Open the graph in the Design tab of the OVPM Java interface. Examine the metrics in the Metrics tab.
2. Select the datasource and system if appropriate from the Sources tab.
3. Back on the Metrics tab, select **Add** then compare the metrics in the graph definition to those found on the datasource.

Check the graph definition using the Web Forms interface:

4. Go to the Design tab and select **Open Graph** for the same template file and graph name.
 5. Select the Data Source you are using. Any metric not supported by that datasource will have its name blanked in the metrics column of the form.
- If you are presenting a graph versus a table or XML document, then only numeric items can be displayed. Non-numeric items are shown in the metric list with quoted names. Make sure that at least one numeric metric is selected for a graph.
 - If you are drawing a graph, but that graph is empty:
 - Verify that the Date Range and Ending selections are not pointing to a time when no data exists. Use System Information feature to verify. Change Date Range to **all** and re-draw the graph. If you get a graph this time, examine it to determine what dates and times have data available.
 - If it appears that the date and time range you selected have data now, but not when you draw for a smaller date range, or if you still get a blank graph, you may have the situation where data is present sporadically, or at intervals greater than your Points Every setting.

Change the graph type to **table**, and set Points Every to **unsummarized**. When you draw this graph, you should see each data point from the datasource. If this table is blank, then there truly does not appear to be any data in the date range you selected; try date range **all** again.

- When data from the datasource is present less frequently than your Points Every setting, the data may not be visible in the graph. This happens because you need 2 consecutive points to draw a line. You can:
 - Change your graph settings to select a marker for each data point.

- Change the Missing setting to **Previous** or **Zero**.
 - Change your **Points Every** setting to be equal to or greater than the distance between each data point.
-

Enable Java Logging

If the problem is in drawing a java graph and you suspect problems in the Java Components of the application, you should enable the Java Console. On Windows, you can set the Java Console to start each time a Java applet is executed. Start the Java Plug-In from the control panel and select "Show Java Console". Your dialog may vary depending on the version of the Java Plug-In you have installed.

Alternately, if the Java applet is running, you may see an icon in your windows Taskbar. You can enable the Java Console on demand, by right-clicking this icon and selecting "Show Console". On other browsers, the Java console is enabled using a menu item such as "tools".

You should see a new window appear called Java Console. It will display a trace of information from the java application and applets. Check this trace for error messages.

OVPm does not Respond

- Sometimes, OVPm will not let you proceed until you take some action in a dialog box that seems to have vanished from your screen. On Windows, press ALT-TAB to bring up the dialog box. ALT-TAB allows you to toggle between windows and will bring up any window that may seem to have disappeared, but is simply hidden behind other screens. Otherwise on Unix, try minimizing or moving other windows to expose the hidden dialog box.
 - When Analyzer is timing out, OVPm will display a blank gray box. Try to determine what is causing the time-out and fix it. Then restart OVPm.
-

No Names are in the System Name Box (Web Forms) or Data Sources Window (Java interface)

- In the Web Forms interface, the contents of the System Name box will vary based on the types of datasources that are in the selected report template. System Names will be available when the datasource is MWA, or Coda. No system names will be generated for the datasource HP OpenView Reporter or Internet Services. If more than one datasource is

in the template file, then the System Name box will contain a mixture of names for all those datasources.

- In the Java interface, all system names that are configured are displayed. If customers are defined, then only those systems that match the current logon customer will be shown. If an invalid customer or password is entered, then no systems will be shown. When logged in to the Java interface as a customer or one of its users, you will not be allowed to Add systems. This is normal since customers are restricted to displaying data only from those systems configured for them.
- No system names can be found. OVPM does not automatically discover system names. Instead, system names must be supplied by the [HP OpenView Reporter](#) product or by the [OVPM Administrator](#).

No Graphs Display in Java Interface

In the Java interface, OVPM displays a list of graphs only after you select a datasource. The graphs displayed also depend on the setting of the Sources field at the top of the Java interface Display tab:

Sources	Shows Graphs for
Individual Systems	OVPA, OVO-EPC, HP OpenView Reporter (Single system and as a whole)
System Groups	HP OpenView Reporter as a whole (not single systems)
Services	Internet Services

Complex graphs, using multiple data sources, will be displayed regardless of the Sources setting.

Data Overlaps in Area graphs

Area graphs should either be stacked or the Fill styles should be other than solid in order to see overlapping data.

Agent not Found on a System

When drawing a graph, you may get a message similar to *"MeasureWare Agent not found on "XYZ" (Socket connection was refused)."* This indicates that OVPM was not able to connect to the agent running on that system.

- First, check the agents on the system; check to insure that an agent of the selected type is actually running on that system (it may not be

- installed, or it may have been stopped). Selecting the OVPM System Information function against the system is a good way to see if you can communicate with the system and what data sources it is providing.
- Next, check to see that the type of datasource configured for this system (at the top of the System Information report) matches the actual datasources on the system. If not use the [Administrator Interface](#) to make corrections.
 - Try drawing a graph from Performance History (Agents template file).
 - Try other graphs.

No Configuration Data Found in Selected Range

A Configuration Details graph shows "no data found in the selected range ..."
Use these tips to draw this type of graph for different agents:

- Use the **System Information** function to determine that the system is responding and has data in the class for the time period selected.
- MWA (OVPA) datasource: In Duration, select Points Every=day and a large date range to capture the time when a record was logged. MWA logs a configuration record only when the agent is started. This will be different for each system.
- Coda (OVO-EPC) datasource: In Duration, select a very small date range, such as 1 hour, to capture the data.
- Sometimes graphs drawn in a Netscape browser do not refresh even when the browser is configured to not cache pages. To work around this problem, you can force page reloads by holding down the Shift key, and clicking the "Reload" button on the browser window. In extreme cases it is necessary to right-click on the image, select View Image, then right-click on the image and select reload.

Some clients can't connect to OVPM

In order for a client to connect to OVPM, it must be able to resolve the name of the OVPM server system. Here are some steps to check out client to host name resolution:

1. **Verify that DNS is operational.**

Use the `nslookup` and `ifconfig` commands to verify network connectivity. For example, the `nslookup` error "*Trying DNS...Non-authoritative answer..*" implies name resolution problems that may affect OVPM.

2. **Work around using `/etc/hosts`**

If the Domain Name Service is not working properly in your environment, you can edit the file `/etc/hosts` (on Unix) or `/system32/drivers/hosts` (on Win32) to override it. After configuration changes, flush the cache by running `ipconfig /flushdns` before testing.

3. Try OVPM on a different port

If your network isn't allowing access to OVPM using the default network port number, You can reconfigure OVPM to use a different port:

```
ovpm config -port 7180
```

Configures OVPM to use port 7180

Restart OVPM after changing the port number.

4. Check local access

Determine whether the network is causing OVPM access problems by running a web browser on the same system as the OVPM server. If you can access OVPM and execute the Web Forms interface then you can assume that OVPM itself is working properly and any problems from specific clients is likely due to network configuration problems.

Test OVPM locally using the `ovpmbatch` command on the OVPM server. The command `ovpmbatch info=1` should return a long list of XML information. Issue the command

```
ovpmbatch graphtemplate=Agents graph="Global History"  
systemname=system
```

where *system* is a system that has either OVPA or OVO agent installed. After a few seconds you should see a line like:

```
Location:  
http://server:8080/OVPM/datafiles//VPI_All_Customers/  
Global_History/...
```

indicating that the graph was generated.

5. If only the Java interface is broken

In order to run the Java interface or to display a Java Graph or Table, the proper Java plug-in must be installed on the client system. In some cases this is easy: Just start the Java interface or graph from a Web browser client and if a Java plug-in is not present, you will be prompted to allow its download from the web.

On other clients, the process may be more involved. Most Unix web browsers will not automatically download the Java plug-in when it is needed. Instead you will have to go to the Sun Java download web page at <http://java.sun.com/j2se/1.4.2/download.html> and select the download to match your system.

For HP-UX clients you must insure that you have the proper web browser (Netscape must be provided from HP, see web site <http://www.hp.com/go/netscape>) For Netscape and Mozilla download the Java plug-in from HP's web site <http://www.hp.com/products1/unix/java/java2/jpi/downloads/index.html>.

Insure that you have the proper Java Plug-in installed. See the ReleaseNotes for the latest information on which plug-ins are supported. You can tell which plug-in you are using by enabling the Java Console (this process differs from browser-to-browser). The plug-in version is printed to the console whenever it is first started.

OVPM Version	Supports Java Plug-ins
A.04.02 and A.04.04 for Windows	1.3.x
C.04.05 for HP-UX and for Solaris C.05.00 for all platforms	1.3.x 1.4.0, 1.4.2 NOT 1.4.1
C.06.00 for all platforms	1.4.2 and 1.5

WARNING: Duplicate name in Manifest: Depends-On

This particular error message can show up in the Tomcat status file (catalina.out) many times as the OVPM application is started. It can also show up inside the Java Console when the OVPM Java interface is started. It is a "feature" of Java 1.4.2 that reports a redundancy inside one of the Java files but otherwise causes no harm. Please ignore these messages if you see them.

sun.plugin.cache.DownloadException

A series of these error messages, accompanied by java stack marker traces can occur when the OVPM Administrator interface is run from a web browser. They can also occur when the OVPM Java interface is run using secure communication (https). These errors are merely reporting failed attempts to locate certain configuration files in several of their alternate locations. The files are eventually located and everything works properly so please ignore these messages if you see them.

Troubleshooting the OVPM Server

The following troubleshooting steps are focused on the OVPM server.

Step 1: Check if your server is configured to run OVPM properly?

OVPM runs as a Java Servlet under the Tomcat application server on the OVPM server system. As such the Java Virtual Machine runs the Tomcat application server plus the OVPM servlet code at the same time. As requests are processed, Java threads are created to allow for parallel processing and a more even response time. Each thread requires an allocation of virtual memory from the Java Virtual Machine when it runs. This memory is recycled once the thread terminates.

The system where the OVPM server is running must be configured to allow enough room for the Java Virtual Machine to do its job in managing the OVPM requests. The amount of each resource required will vary depending on the number and type of accesses to OVPM. Light and infrequent requests will likely finish before the next request is started and resources usage is low. Multiple simultaneous requests and requests that require very large amounts of data handling will require more resources.

Symptoms of insufficient resources include a general slow-down during busy times. In severe cases, the Java Virtual Machine itself can crash due to resource starvation. In this case the OVPM server stops and no further requests are processed. Look in the Tomcat status file for indications as to what resource was lacking. Most common problems come from a shortage of threads or virtual memory.

The Tomcat status file is found at `<tomcat_dir>/logs/catalina.out` on Unix systems and `<tomcat_dir>/logs/stdout.log` on Windows. Where `tomcat_dir` is `<install_dir>/nonOV/tomcat/a/`

Special Note for HP-UX systems only:

The default HP-UX configuration, specifically kernel parameters and system tunables, does not support certain types of Java applications well. Applications that typically need to be tuned by hand are large, server-side Java applications that are long running, have large numbers of threads, and have a large ratio of client to server connects, such as OVPM.

Java Out-of-Box is a stand-alone bundle that upon installation will install startup (RC) scripts, modify kernel parameters, rebuild the kernel, and reboot the system. During startup, the startup scripts will modify system tunables, thus providing better "out of the box" behavior for Java. For more information, visit this website:

<http://www.hp.com/products1/unix/java/java2/outofbox/index.html>.

HPjconfig is a Java configuration tool for tuning your HP-UX 11 PA-RISC & Itanium system kernel parameters to match the characteristics of your application. HPjconfig provides kernel parameter recommendations tailored to

your Java enterprise services (Web server, Application Server, etc.) and HP-UX hardware platform. It offers save and restore functions for easy distribution of tailored recommendations across your user base.

<http://www.hp.com/products1/unix/java/java2/hpjconfig/index.html>.

Step 2: Check the Application Server

Issue the command `<install_dir>/bin/ovpm status` to determine if the OVPM application server is running and to list the last few lines in the tomcat and OVPM status files. If the application server has failed, examine the tomcat status file `<install_dir>/OVPM/logs/catalina.out`.

If the application server is down, or appears to be hung, perform the following procedure to restart it:

1. While logged on as root or superuser on the OVPM server, issue the command to stop OVPM:

```
<install_dir>/bin/ovpm stop
```

This may generate errors indicating the application is not responding if it is not running or is otherwise hung.

2. Once the Java Virtual machine has terminated, restart OVPM.
3. Optional: watch the application server log file to insure it comes fully alive.

```
<install_dir>/bin/ovpm tail tomcat
```

Step 3: Test What DOES Work

1. Test the Tomcat application server where OVPM is running itself by using a browser to reference the URL: `http://<system>:8080`

where `<system>` is the actual system name where OVPM is installed.

- o If the application server is running properly, then you should see a web page showing "hp OpenView Software" and the HP logo.
 - o If this page displays properly, the application server is functional. If not, request assistance to solve the problem.
2. Test the OVPM home page by using a browser to reference the URL: `http://system:8080/OVPM`.

This is the OVPM home page. If it does not display, or displays incorrectly, then the OVPM configuration has not been completed properly. Re-run the OVPM configure command:

```
<install_dir>/bin/ovpm config [-port  
portno]
```

where `portno` is the port number (normally 8080). Once this is done, repeat this test.

On Windows, check the `<Tomcat_dir>/conf/server.xml` file to insure that it references the same port numbers that you are using.

3. Test the Web Forms interface by pressing the **Web Forms Interface** button on the OVPM home page. This assures that the OVPM application itself can be launched and will respond to commands.
4. Test the Java interface by returning to the OVPM home page and pressing the **Java Interface** button. If everything works up to this point, but the Java Interface does not come up, perform the following steps:
 - a. Use a browser to reference the URL:

```
http://<system>:8080/OVPM/Analyzer?-info2
```

You should be rewarded with a lot of XML data on your screen. If you are not using a current browser, then the XML may not actually be visible. Use the View Source option on the browser.

- b. Use a browser to reference the URL:

```
http://<system>:8080/OVPM/Analyzer?-info
```

You should be rewarded with even more XML data on your screen. If you are not using a current browser, then the XML may not actually be visible.

- c. If both these steps work, but the Java interface does not come up, then the problem most likely lies in the Java Plug-in installation on the system running the browser. Try another system or work on getting the correct plug-in installed.

Step 4: Check the Status File

As OVPM runs, it logs any warnings or errors to a text file called `ovpm.0.log`. This file is located in the `<data_dir>/log` directory. This file may contain many lines of text. To see the most recent entries, use the `tail` command:

```
<install_dir>/bin/ovpm status
OR
tail -50 <data_dir>/log/ovpm.0.log
```

Most warnings are considered minor and should not be of a concern unless they occur frequently. Errors may be simple such as the inability to connect to the system requested. More severe errors might indicate a problem where the product is not able to execute due to an expired trial period etc. Each line in the status file is marked with the date and time when the message was generated.

As the status file reaches one megabyte in size, it will be renamed to `ovpm.1.log` and a new `ovpm.0.log` file will be started. Only two log files are retained. If there is already an `ovpm.1.log` file present when the `ovpm.0.log` fills up, the existing `ovpm.1.log` file will be deleted to make room for new data.

Step 5: Enable Tracing

If you need more detail than is contained in the status file, you can enable tracing for the execution of the main OVPM program. To generate a trace log file for OVPM, execute the command:

```
<install_dir>/bin/ovpm trace log
```

There is no need to reboot or restart any programs. Tracing will take effect immediately once it is set. The trace data will be written to file `<data_dir>/log/Analyzer.trc`. This is a binary file that should be transmitted to your support personnel, taking care to specify binary mode file transfer.

Trace files can fill rapidly and consume considerable disk space so you should only log trace data for short periods of time.

To interactively "browse" the trace data without generating a log file, (not recommended unless told to do so by a support representative), issue the commands:

```
<install_dir>/bin/ovpm trace  
<install_dir>/bin/ovpm tail trace
```

You must press CONTROL-C to terminate when running in this mode.

Caution: The trace records contain much detail that is designed strictly to assist a trained support representative. No assistance will be given for understanding or supporting the information they contain.

To stop tracing, execute the command:

```
<install_dir>/bin/ovpm trace off
```

Unable to Find the First Web Page

When OVPM is installed or re-initialized, it creates a web page called `index.htm` in the web server `OVPM` directory. If you attempt to start OVPM and get an error message, it may be caused by the inability to display this web page.

- Check for the file named `<install_dir>/OVPM/webapps/OVPM/index.htm` on Unix or `<install_dir>\www\webapps\OVPM\index.htm` on Windows. If it is not present or if it has no text in it, try to create it by running the script: `<install_dir>/bin/ovpm config`
- If the file is present and not empty, try to open it using your web browser on the local system. If this succeeds, but accessing it through the web server fails, then investigate whether your web server is running properly. Try accessing a simple web page in the same directory.

Unable to Find Graph

If you are unable to find a graph when using the Web Forms interface it may be because the Graph Name includes more than 1 consecutive space. Do not use consecutive spaces in the graph name. (In the Web Forms interface this graph name "This is my graph" can be found, but not "This is my graph".)

Keyword Unrecognized

If you see the message "Keyword Unrecognized", the OVPM Analyzer version is out of date or someone has entered an invalid graph definition item.

Database Open Error

You can get this error when accessing the HP OpenView Reporter, Internet Services or Generic Database datasources. Insure the database information is properly entered into the `OVPMconfig.ini` file and that database security is set properly.

Following are some typical error messages with their causes:

DBDriver exception Io exception: The Network Adapter could not establish the connection

The host or port is misconfigured in the `OVPMconfig.ini` file.

DBDriver exception Io exception: Connection refused ...

The DSN is misconfigured in the `OVPMconfig.ini` file.

DBDriver exception ORA-01017: invalid username/password; logon denied

The username or password is incorrect in the `OVPMconfig.ini` file.

Troubleshooting Database Connections

1. Try to *ping* the system running the database server. If the *ping* command cannot locate the database server system, then neither can OVPM. You may replace the system name with its IP address if that is accessible.

```
ping <systemname>  
ping <ipaddress>
```

2. Test whether the port number you have configured has a listener process running. Execute the following command:

```
telnet <systemname> <port>
```

If you get a connection, press **return** and then type exit to terminate the telnet session. If you get a "connection refused" message, then there is no server listening on the port. Consult your database administrator to

insure that you are using the correct port number and that the database listener process is running.

Unresolved Host Name

If you can access the OVPM home page properly, but attempts to launch either the Java interface or Web Forms interface give an unresolved host name error, then there may be a problem with the way your server system is configured.

- If you can see the URL that the browser was unable to find, examine the first part. It should be a fully qualified version of your server name (such as *myserver.division.company.com*).
- If the URL has only the short version of the server name (*myserver*) or something equally bizarre such as *www.myserver.com*, then you should check your server system network identification information.

OVPM will attempt to determine the protocol (http or https), system name and port number from your initial connection and then continue to use them for all subsequent accesses.

Troubleshooting OVO for Windows Graphs

Problem: If full version of OVPM is installed on OVO management server (for Windows) and super user password of the newly installed OVPM is changed, OVO fails to generate graphs.

Possible Cause: OVO generates graphs by using graphing component of OVPM, which is shipped with OVO. After you install a full version of OVPM on OVO management server, the graphing component of newly installed OVPM starts generating graphs for OVO instead of the graphing component shipped with OVO.

If you change super user password of this new version of OVPM, the change does not get synchronized with OVO, and OVO graph generation fails.

Solution: Perform the following steps on OVO management server:

1. Right-click **My Computer**.
2. Click **Manage**. The Computer Management window appears.
3. Expand **Services and Applications**.
4. Expand **Internet Information Services**.
5. Expand **Web Sites**.
6. Right-click on **Default Web Site** and click **Properties**.
7. Select ISAPI Filters Tab.
ISAPI filter is enabled if you see the following entries:
 - o AnalyzerISAPI
 - o jakarta_ovpm
8. To disable ISAPI filters, click AnalyzerISAPI and jakarta_ovpm, and then remove them.

- Restart OVO server.

Detail Specifications

[Metrics](#)

[Date Ranges](#)

[Number of Points in the Graph](#)

[Filters](#)

[Labels](#)

[Parameters](#)

[Date Format Strings](#)

[Time Format Strings](#)

[Keywords](#)

Metrics

Metrics may be supplied by exactly matching one of the available class and metric values from the datasource. The metric comparison is case sensitive.

Calculated Metrics are also allowed. These metrics are formed by performing calculations on other metrics. A calculated metric name consists of other metric names, constants, plus the operators **+**, **-**, *****, **/**, **(**, **and** **)**.

The calculation on metrics is performed left to right, not by mathematical precedence rules. If the calculation contains a combination of multiplication/division and addition/subtraction, parentheses must be used. For example:

- $3+(2*10) = 23$
- $3+2*10 = 50$

If using parentheses, make sure they are well formed and balanced. Also, do not include spaces in your calculated metric. There is no limit to the length of the calculated metric string, other than the performance penalty paid for calculating overly complex equations.

An example of a calculated metric:

```
CLASS:    APPLICATION
METRIC:   (APP_CPU_TOTAL_TIME*100.0)/INTERVAL
```

Date Ranges

DATERANGE:, BEGINDATE:, ENDDATE: (or ENDDATEMONTH:, ENDDATEDAY:, ENDDATEYEAR:, ENDDATEHOUR, ENDDATEMINUTE:), STARTTIME:, STOPTIME:

If none of DATERANGE: ENDDATE: STARTDATE: STARTTIME: nor STOPTIME: is supplied, the default will be "7 Days" ending "now".

Data will be selected from the datasource based on the first date/time type of metric in the class. Usually this is the DATETIME metric and it represents local time.

The beginning and ending date and time of the included data can be set in one of three ways:

- Set the beginning and ending date and time explicitly using the STARTTIME: and STOPTIME: parameters. In this case, any values entered for DATERANGE:, BEGINDATE: and ENDDATE: will be ignored. STARTTIME: and STOPTIME: use a locale neutral format for specifying dates and times (YYYY,MM,DD HH:MM:SS).
- Set the beginning and ending date and time explicitly using the BEGINDATE: and ENDDATE: parameters. In this case, any values entered for DATERANGE: will be ignored. BEGINDATE: and ENDDATE: use a format for specifying date and time that is set by the locale of the system where Performance Manager is installed.
- Set the DATERANGE: to determine how large a time interval to select and select ENDDATE:, (or its alternate form consisting of any of ENDDATEMONTH:, ENDDATEDAY:, ENDDATEYEAR:, ENDDATEHOUR or ENDDATEMINUTE:) to determine the end of that interval. (The beginning of the interval will be calculated from these two values.) This is the method most commonly used.

DATERANGE: and ENDDATE:

DATERANGE selects a number and a units identifier. Acceptable units identifiers are *Minutes, Hours, Days, Weeks*

- "7 Days" is the default date range. A special date range, "all", selects all the data available from the datasource.
- ENDDATE sets the end of the time interval chosen with DATERANGE. The default end date, "now" sets the current date and time as the end of the interval. **Last** indicates the date and time when data was last added to the logfile.
- You can also specify a date and time. The format depends on the date format settings on the OVPM server system; but it will be Year, Month and Day in some order followed by Hour in 24-hour format, and Minute. Any non-numeric value, including spaces, can be used to separate each field. Fields may be omitted from right to left. Specifying a date but no time will select the end of the day (24:00). Note: The time stamp in HTML graphs does not always represent the time selected for ENDING; this is because the time is rounded off based on the Date Range and Points Every combination.

Instead of entering ENDDATE: you can provide the ending date using a series of parameters, one for each component of the date and time. If any of these parameters is specified, then the ENDDATE: value is ignored.

KEYWORD	DESCRIPTION	DEFAULT VALUE
ENDDATEYEAR:	Year 1970-2037	Current Year
ENDDATEMONTH:	Month 1-12	Current Month
ENDDATEDAY:	Day 1-31	Current Day of the Month
ENDDATEHOUR:	Hour 0-24	Midnight (24:00)
ENDDATEMINUTE:	Minute 0-59	0

STARTTIME: and STOPTIME:

The values in these fields will set the starting and stopping date and time for the data retrieval. These values will take precedence over the BEGINDATE: and ENDDATE: values.

The syntax for these dates is a series of numbers, separated by non-numeric characters and does not vary based on the settings of the OVPM server system. The numbers must be in the following order: Year, Month, Day, Hour, Minute, Second. Year may not be abbreviated to only two characters. Less than six fields may be provided, with the rightmost fields defaulting as follows:

```

Month= 2
Day= 1
Hour = 0
Minute = 0
Second = 0

```

For example, to specify July 4, 2000 at midnight:

```

2000,7,4,0,0,0
or
2000,7,4

```

BEGINDATE: and ENDDATE:

If STARTTIME: and STOPTIME are not provided, then the values in these fields will set the starting and stopping date and time for the data retrieval.

The syntax for these dates depends on the settings on the system where OVPM is installed. The syntax will always be Month, Day, and Year, in some order, followed by Hour and Minute. The values may be separated by any non-numeric character. The syntax does not have to match the locale specifications exactly. Hour is always 24 hour time, regardless of locale settings and no AM or PM indicator is allowed. Enclose the values in double quotes if they contain any blanks.

For example, to specify July 4, 2000 at midnight in the USA:

"7/4/00 24:00"
or
"7/4/2000"

Number of Points

POINTSEVERY and NUMBEROFPOINTS

For all graphs except "Pie" and "baseline", the number of data points along the X-axis may be specified in one of two ways.

- Specifying "**NUMBEROFPOINTS:**" followed by a number between 1 and *<maxpoints>* will place this number of points on the X-Axis. (This value may be rounded down slightly to insure each data point contains an integral number of 5-minute intervals). The default number of points is 50 or as specified in the graph definition.

<maxpoints> will vary depending on the type of graph being drawn:

Graph Type	Max Points
pie	1
Gauge	1
web form graph	1024
web form table	5000
Java graph	2048
Java Table	5000
ASCII (TSV, CSV)	1,000,000
XML	5000
Excel	65,000
Forecast	5000

The maximum number of points in a Java graph or table may be overridden using the JAVAGRAPHPOINTS and JAVATABLEPOINTS parameters in a style file (VPI_Style.txt).

CAUTION: some web browsers may have trouble displaying tables with more than 1000 rows. Java may hang and/or take excessive CPU when a large number of points is graphed or displayed in a table.

- Specifying "**POINTSEVERY:**" followed by one of the following codes will divide the date and time range (as supplied in the DATERANGE or STARTTIME/STOPTIME values) into as many points as it takes so that each point is the desired time range. Adjustments may be made to insure the number of points falls between 1 and Max Points.

auto	Automatically calculate the number of points for each data point (default)
5 min	One point for every five minutes
15 min	One point for every fifteen minutes
30 min	One point for every thirty minutes
hour 1 hour	One point for every hour
3 hours	One point for every three hours
6 hours	One point for every six hours
12 hours	One point for every twelve hours
day 1 day	One point for every twenty four hours
<number>	<number> is a user supplied number of minutes for each point
raw	Data will be presented unsummarized from the datasource

Filters

Filters allow you to select subsets of data. A filter is applied to each data record before summarization takes place. Data records that do not meet the filter specifications will not be used in the final calculation.

A filter consists of a metric name, a comparison symbol and a value to compare against. Metric names must come from the same class as the metric the filter is being applied to. The value to compare against should be a literal, or the special case described below. Comparison codes are:

=	Equal
>	Greater
<	Less
>=	Greater or Equal
<=	Less or Equal
!=	Not Equal
!>	Not Greater (Same as <=)
!>=	Not Greater or Equal (Same as <)
!<	Not Less (same as >=)
!<=	Not Less or Equal (Same as >)
~	Like (text with leading or trailing "*" wildcards)
!~	Not Like (text with leading or trailing "*" wildcards)

For example, the following filter could be used to select data for a single system from the HP OpenView Reporter datasource (SYSTEMNAME is used for HP OpenView Reporter only).

```
FILTER: SYSTEMNAME=MYOWNSYS
```

Multiple filters can be combined through the use of the && symbol.

```
FILTER: SYSTEMNAME=MYOWNSYS&&APP_NAME=Other
```

would only select data where the system name was "MYOWNSYS" and the application name was "Other".

For HP OpenView Reporter and other datasources you can use:

```
FILTER: APP_NAME=Other
```

This would only select data where the application name was "Other".

In the **special case** of the filter parameter, the value "@" ("*metricname=@*") will match all values from the datasource, but each distinct value will be displayed as a separate metric.

```
FILTER: APP_NAME=@
```

Will select all data, but each application name will form its own metric for graphing. The '@' sign may be used when multiple filters are combined together; for example,

```
FILTER: SYSTEMNAME=@&&APP_NAME=@
```

Will select all data, but form a unique metric for each system and application name combination.

Note: The special case '@' filter is most useful when the filter metric is textual.

"Like" comparisons (using the "~" and "!~" codes for "Like" and "Not Like") are textual comparisons that allow wild cards.

- If the value starts with an asterisk "*", then the data will be "like" if it contains the rest of the value at its end. For example "**metric~*abc**" would match data where metric values "Iknowmy**abc**", "**abc**" but not "**abcd**".
- If the value ends with an asterisk "*", then the data will be "like" if it contains the first part of the value at its beginning. For example: "**metric~abc***" would match metric values "**abcdefg**" and "**abc**" but not "Iknowmy**abc**".
- If the value begins and ends with an asterisk "*" then the data will be "like" if it contains the rest of the value anywhere. For example: "**metric~*abc***" would match metric values "Iknowmy**abc**", "**abc**", "**abcdefg**" and "Iknowmy**abc**swell".
- Asterisks not at the beginning or end of a value are treated as textual data, that is they are not treated as wild cards. For example: "**metric~abc*d**" will only match metric value "**abc*d**" it will not match "abcd" or "abdxd".
- "Not Like" comparisons "!~" pass data where the "like" comparison fails.

NOTE: You must prefix the filter string with *metricname=@&&* for any comparison symbol other than =.

Alternate Keyword Syntax

Filters can be specified using the single FILTER: keyword as illustrated above. This is the most general syntax and it allows multiple filter criteria. Alternately, filters in the form of "metric comparison value" can be specified using three keywords, one for each part of a simple filter:

SFMETRIC: Metric name

SFCOMPARE: Comparison code

SFVALUE: Comparison value

kp>The following commands are equivalent:

```
FILTER: APP_NAME=other SFMETRIC: APP_NAME SFCOMPARE: =  
SFVALUE: other
```

TableFilters

Table Filters are very similar to regular filters but with a few differences. They apply only to tables and are ignored for all graphs. Unlike filters, tablefilters are applied after the data is gathered and summarized. Regular filters are applied before the data is summarized. Using the Java Table, the table filters can be altered after a table is displayed in order to change which data is displayed without having to go back to the data sources to retrieve new data. This makes a convenient way to work quickly with tabular data and makes it easier to find what you are looking for.

Table filters can be combined using the "&&" (**AND**) connector like regular filters. They can also be combined with the "||" (**OR**) connector. Using &&, all filter specifications must be true in order for the filter to be true. Using ||, if any filter specification is true then the entire filter is true.

The alternate form for specifying filters (SFMETRIC: SFCOMPARE: SFVALUE:) is not available for table filters.

The metrics used in a Table Filter don't have to be in the same data class or even data source as they do with regular filters. The metrics only have to be in the table. Metrics are identified by metric name so if you have more than one metric with the same name in the table then the first occurrence will be used by the table filter.

Labels

You can assign a label to any metric. If you do not, then the default metric label specification is used on the graph. The default metric label will normally be any label supplied by the data source for this metric, or the metric name. The default metric label can be changed by the OVPM administrator.

When using the metric filter value @, you can also use in the label the string '@@' followed by the metric name to substitute the value of that metric into the label.

```
FILTER: APP_NAME=@  
LABEL: Application @@APP_NAME
```

will form labels like "Application " followed by the actual application name. For example "Application Other" and "Application Program Development".

You can also combine different metric names in a label:

```
FILTER: SYSTEMNAME=@&&APP_NAME=@  
LABEL: @@SYSTEMNAME: @@APP_NAME
```

will form labels like "MYOWNSYS:Other" and "TheOtherSystem:Program Development". This feature is available only in the Java interface, not in the Web Forms interface.

Parameters

While processing the template file, substitutions can be made using up to nine parameters named PARAMETER1 through PARAMETER9. This substitution is performed on each template line before that line is processed. Parameters may be used to substitute for any part of the template file syntax (including filter values, metric and class names etc.)

The template file should set a default value for a parameter using the "PARAMETER" statement. Note the syntax for this statement is "PARAMETER" followed by exactly one digit from 1 to 9, followed by a colon. "PARAMETER1:", "PARAMETER2:,... "PARAMETER9:".

Once a parameter default has been provided, it will be substituted into succeeding template lines wherever the syntax "@@PARAMETERn" is found. (n is a digit from 1 to 9). For example: the following statements are equivalent:

```
FILTER: SYSTEMNAME="MYSYSTEM"  
and  
PARAMETER1: MYSYSTEM  
FILTER: SYSTEMNAME=@@PARAMETER1
```

Using parameters allows easy generalization of templates so that a different graph template is not required for each parameter value.

Parameters can be overridden in the URL or command line interfaces. This allows different parameter values to be supplied each time the program is run.

Date Format Strings

Normally dates are formatted according to the language on the system where the OVPM server is installed. You can override this format by entering a new one using the DATEFORMAT keyword. Normally this is used with the SAVECONFIG command to save a client configuration or in a Style file.

The format must contain the following three fields in order to be accepted. The order will determine the format for dates on graphs and tables plus the format of dates expected when inputting in a local format. Other characters will be repeated exactly as shown providing they are not one of the special formatting characters "GyMdhHmsSEDFwWakKz'".

Year	y or yy	Exactly two digits will be used. 1998 = "98" while 2005 = "05"
	yyyy	Four digits will be used.
Month	M	The Month will be shown without leading zeros. February = "2".
	MM	The Month will be shown with leading zeros. February = "02".
Day	d	The day of the month will be shown without leading zeros.
	dd	The day of the month will be shown with leading zeros.

Example: February 4, 2005 would be shown as:

Format	Date
M/d/yy	2/4/05
MM/dd/yy	02/04/05
M-dd-yyyy	2-04-2005

Changes to the date format will be see on any new graphs that are drawn.

Time Format Strings

Normally times are formatted according to the language on the system where the OVPM server is installed. You can override this format by entering a new one using the TIMEFORMAT keyword. Normally this is used with the SAVECONFIG command to save a client configuration or in a Style file.

The format must contain the following fields in order to be accepted. The order will determine the format for times on graphs and tables plus the format of times expected when inputting in a local format. Other characters will be repeated

exactly as shown providing they are not one of the special formatting characters "GyMdhHmsSEDFwWakKz'".

24 Hour	H	Hour of the day, based on a 24 hour clock (0-23) with no leading zeros.
	HH	Hour of the day, based on a 24 hour clock (0-23) with leading zeros.
12 Hour	h	Hour of the day, based on a 12 hour clock (0-11) with no leading zeros.
	hh	Hour of the day, based on a 12 hour clock (0-11) with leading zeros.
Minute	m	Minute of the hour with no leading zeros.
	mm	Minute of the hour with leading zeros.
AM/PM	a	AM/PM indicator (in the local language) This item is optional and should only be used when the hours are in 12 hour format.

Example: 8:05 PM (20:05 hours) would be shown as:

Format Time

h:m a 8:05 PM

hh:mm a 08:05 PM

HH:mm 20:05

Keywords

These specifications are used to direct the Performance Manager run time engine to perform its various tasks.

When using ovpmbatch all keywords must take a parameter, even those that do not normally take a parameter. In those cases, set the keyword's value to "1" (such as "SUMFROMRAW=1").

These keywords may be specified in slightly different syntax, depending on their usage:

Web Browser	<p>-KEYWORD: <i>value</i></p> <p>Note that the dash preceding a keyword and the colon following it are required.</p> <p>If a keyword does not require a value, then it can be omitted. -SKIPDAY1 :</p> <p>Values that include blanks must be enclosed in quotes. -GRAPH: "Global History"</p> <p>Multiple keywords are separated by blanks. -GRAPHTEMPLATE: Agents -GRAPH: "Global History"</p>
URL Syntax 1	<p>-KEYWORD: %20<i>value</i>%20</p> <p>Same syntax as above, except that %20 (hex code for space) is used in place of a space. This is done automatically by most web browsers but</p>

	must be done manually if a URL is generated any other way. - GRAPHTEMPLATE:%20Agents%20-GRAPH:%20"Global%20History"
URL Syntax 2	KEYWORD=value& When using this syntax, a blank is replaced by a plus "+". Any special character in a keyword or value including equal sign "=" and ampersand "&" must be replaced by their hex codes. GRAPHTEMPLATE=Agents&GRAPH=Global+History&
ovpmbatch	KEYWORD=value The "=value" is required for ovpmbatch, even for keywords that do not require a value. For example: SKIPDAY1=1 GRAPHTEMPLATE=Agents GRAPH="Global History"

The following specifications apply to the entire graph. They cannot be specified independently for each metric.

General Specifications		
KEYWORD	VALUE	Descriptions
PARAMETER<n>: -p<n>	Parameter Value	Where <n> is a number from 1 to 9. Defines one default parameter value. Parameter values in a graph template file can be overridden by parameters supplied at run time. (See Parameters)
CUSTOMER:	Customer Name Or "All Customers"	The customer name restricts data presented to that which belongs to a particular customer. A different password may be required to access each customer's data. Customers are configured using the interfaces for other products such as Internet Services. If "All Customers" is entered and customer restricted views are enabled, then the administrator's password must be supplied.
PASSWORD:	Password	If customer restricted views are enabled for a data source, then a password may be required in order to access any of that customer's data. This can be a different password for each customer, or a special Administrator's password (blank user password). In order to specify all customers, the administrator's password must be supplied whenever restricted views are enabled.
GROUP:	Group name	Some data may be subdivided into groups. Specifying a group name will restrict data in the presentation to that which belongs to this group. Groups may have different meanings, depending on the data source. For example: Internet Services provide groups of services whereas HP OpenView Reporter

		provides groups of systems.																		
AUTOREFRESH:		Specifies that a graph should periodically refresh itself (every 2 minutes). Baseline graphs do not support this feature.																		
GRAPHTEMPLATE: (or first unnamed value)	Template file name	Identifies the name of the file containing the graph specifications.																		
GRAPH: (or second unnamed value)	Graph Name	Identify the name of the graph within the graph template file. Default is the first graph in the file.																		
GRAPHTITLE:	Title String	Specifies the title string for the graph. *																		
GRAPHSUBTITLE:	Second Title String	Specifies the second title string for the graph. *																		
	<p>* Keywords can be added to the Title or SubTitle fields to substitute the actual values for:</p> <table border="1"> <tr> <td>@@[SYSTEM]</td> <td>The name of the system being graphed</td> </tr> <tr> <td>@@DATERANGE</td> <td>Total time shown in the graph (7 Days)</td> </tr> <tr> <td>@@POINTSEVERY</td> <td>The amount of time summarized into each data point (10 Min)</td> </tr> <tr> <td>@@SHIFT</td> <td>The shift times (8:00-17:00 or "All Day")</td> </tr> <tr> <td>@@SKIPDAYS</td> <td>The list of days that are being skipped (SAT,SUN)</td> </tr> <tr> <td>@@SHOWDAYS</td> <td>The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)</td> </tr> <tr> <td>@@NOHOLIDAYS</td> <td>Indicates whether Holidays are being skipped or not (blank or "No Holidays")</td> </tr> <tr> <td>@@STARTTIME</td> <td>The beginning time in the graph (3/1/04 3:00 PM)</td> </tr> <tr> <td>@@STOPTIME</td> <td>The last time in the graph (3/8/04 3:00 PM)</td> </tr> </table> <p>You can also make the subtitle smaller to allow more text by adding <small> at the beginning of the subtitle and </small> to the end. Example: <small>@@DATERANGE with points every @@POINTSEVERY @@SHOWDAYS</small></p>		@@[SYSTEM]	The name of the system being graphed	@@DATERANGE	Total time shown in the graph (7 Days)	@@POINTSEVERY	The amount of time summarized into each data point (10 Min)	@@SHIFT	The shift times (8:00-17:00 or "All Day")	@@SKIPDAYS	The list of days that are being skipped (SAT,SUN)	@@SHOWDAYS	The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)	@@NOHOLIDAYS	Indicates whether Holidays are being skipped or not (blank or "No Holidays")	@@STARTTIME	The beginning time in the graph (3/1/04 3:00 PM)	@@STOPTIME	The last time in the graph (3/8/04 3:00 PM)
@@[SYSTEM]	The name of the system being graphed																			
@@DATERANGE	Total time shown in the graph (7 Days)																			
@@POINTSEVERY	The amount of time summarized into each data point (10 Min)																			
@@SHIFT	The shift times (8:00-17:00 or "All Day")																			
@@SKIPDAYS	The list of days that are being skipped (SAT,SUN)																			
@@SHOWDAYS	The list of days that are NOT being skipped (MON,TUE,WED,THU,FRI)																			
@@NOHOLIDAYS	Indicates whether Holidays are being skipped or not (blank or "No Holidays")																			
@@STARTTIME	The beginning time in the graph (3/1/04 3:00 PM)																			
@@STOPTIME	The last time in the graph (3/8/04 3:00 PM)																			
YAXISTITLE:	Title String	Specifies the label for the left Y-Axis.																		
YAXISTITLERIGHT:	Title String	Specifies the label for the right Y_Axis.																		
GRAPHTYPE:	line	For line or multi-type graphs.																		
	area	For filled area graph.																		
	bar	For bar graph.																		
	hbar	For horizontal bar graph.																		

	mixed	Metrics can be any combination of line, bar or area.
	pie	For pie graph.
	table	For table with time vertical and metrics horizontal.
	htable	For table with time horizontal and metrics vertical.
	gauge	For automotive type gauge, one per metric.
	xml	For XML data representation.
	csv	Text file with Comma Separated Values
	tsv	Text file with Tab Separated Values
	excel	Microsoft Excel Spreadsheet (requires Excel on the client system)
	baseline	For baseline report representation.
GRAPHBACKGROUND:	<color>	Color of graph background. The default is transparent for Java, i.e. same as application background.
GRAPHMETRICSPERGRAPH:	<number>	The maximum number of metrics on a single graph. The default is 8.
GRAPHMULTIPLEGRAPHS:	Yes or No	Defaults to NO. If Yes, then creates additional graphs if the total number of metrics to be graphed exceeds the GRAPHMETRICSPERGRAPH: value. Ignored for Java graphs. Text editors do not display the default values. To change the default settings, you need to specifically include the keywords.
DATERANGE:	<daterange>	Range of date/times to include in the graph (See Dateranges).
BEGINDATE: ENDDATE:	<datetime>	Earliest and latest date and time to include in the graph. This parameter works in conjunction with the DATERANGE: parameter to set the time covered in the graph. The format of the ENDDATE: field depends on the settings on the Performance Manager Server system. Year, Month and Day are entered first, although the specific order will match the server system's locale. Hour and Minute are optional, but if present, must follow the three date fields. Any non-numeric separator may be used between fields. In the USA, December 25, 2001 at one minute before midnight is entered "12/31/01 23:59".

STARTTIME:	<datetime>	Earliest date/time to include in the graph. Format YYYY,MM,DD,HH,MM,SS.
STOPTIME:	<datetime>	Latest date/time to include in the graph. Format YYYY,MM,DD,HH,MM,SS. If both STARTTIME and STOPTIME are provided, then this data range will be used and DATERANGE and ENDDATE will be ignored if they are also entered. (See Dateranges)
SHIFT:	HH:MM- HH:MM	The shift starting and stopping times using a 24 hour clock. All Day is specified as 00:00-24:00.
SHIFTSTART:	HH:MM	The shift starting time using a 24 hour clock (00:00 to 24:00).Use in conjunction with SHIFTSTOP
SHIFTSTOP:	HH:MM	The shift stopping time using a 24 hour clock (00:00 to 24:00).Use in conjunction with SHIFTSTART
SKIPDAYS:	1234567 SAT,SUN,... <i>Local day of week</i>	The days of the week to be excluded from the graph. Numbers 1 through 7 correspond to days Sunday through Saturday respectively. Localized days of the week can be used only when OVPM has been localized for the locale. See also SKIPDAY1 through SKIPDAY7.
SKIPDAY1: SKIPDAY2: SKIPDAY3: SKIPDAY4: SKIPDAY5: SKIPDAY6: SKIPDAY7:		The days of the week to be excluded from the graph. SKIPDAY1 will exclude Sunday, SKIPDAY2 will exclude Monday, etc. See also SKIPDAYS. You should avoid using SKIPDAYS and SHOWDAYS in the same request.
SHOWDAYS:	1234567 SAT,SUN,... <i>Local day of week</i>	The days of the week to be included from the graph. Numbers 1 through 7 correspond to days Sunday through Saturday respectively. Localized days of the week can be used only when OVPM has been localized for the locale. See also SHOWDAY1 through SHOWDAY7.
SHOWDAY1: SHOWDAY2: SHOWDAY3: SHOWDAY4: SHOWDAY5: SHOWDAY6: SHOWDAY7:		The days of the week to be included from the graph. SHOWDAY1 will include Sunday, SHOWDAY2 will include Monday, etc. See also SHOWDAYS. You should avoid using SKIPDAYS and SHOWDAYS in the same request.
NOHOLIDAYS:		Exclude holidays from the graph. Holidays are defined in a configuration file that is maintained by the administrator.

GRAPHSIZE:	small	Draws graph smaller than normal. Use PIXELSWIDE: and PIXELSHIGH: for more precise control. All graphsize values are ignored for Java graphs.
	medium	Draws graph at normal size. Use PIXELSWIDE: and PIXELSHIGH: for more precise control.
	large	Draws graph at larger than normal size. Use PIXELSWIDE: and PIXELSHIGH: for more precise control.
PIXELSWIDE:	<number>	The width of the graph image file in pixels (default=600). Use GRAPHSIZE: for simpler control. Ignored for Java graphs.
PIXELSHIGH:	<number>	The height of the graph image file in pixels (default=500) Use GRAPHSIZE: for simpler control. Ignored for Java graphs.
GRID:	none or horiz or vert or both	Specifies whether horizontal and/or vertical grid lines should be drawn.
GRIDHORIZ:		Horizontal grid lines should be drawn.
GRIDVERT:		Vertical grid lines should be drawn.
SORTBY:	value, reverse or label	Sorts metrics by their average value (highest to lowest), by their average value (lowest to highest) or by their label string. Default is to sort metrics by the sequence in the specifications.
SORTBYVALUE:		Sorts metrics by their average value.
SORTBYLABEL:		Sorts metrics by their label string.
STACKED:		Area and Bar graphs should "stack" the value of one metric on top of the previous metrics.
3D:	<blank> yes true 1	Presents graphs with a three dimensional effect. Any value other than blank, yes, true or 1 will disable three dimensional effect (default.)
NUMBEROFPOINTS:	<number>	The number of data points to present (default is 50 points). This value is used with the POINTSEVERY specification to determine how much to summarize the data for each data point. <ul style="list-style-type: none"> For line, and area, <number> should be from 1-255, and 1-1024 for java. HTML horizontal bar graphs are limited to 20 points. For table and xml graphs, <number> should be from 1-

		<p>5000; but take care, some browsers behave badly when tables exceed 1000 rows!</p> <ul style="list-style-type: none"> For pie and baseline graphs, this specification is ignored (pies always show one point for each metric, baselines always show one point for each hour). <p>(See PointsEvery)</p>
JAVAGRAPHPOINTS:	<number>	<p>The maximum number of data points in a Java Graph from 255 to 2048. (Default is 1024 points.) This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Large values can cause delays of several minutes when displaying Java Graphs with many points. This item is most often used in a Style File.</p>
JAVATABLEPOINTS:	<number>	<p>The maximum number of data points in a Java Table from 50 to 2048. (Default is 1024 points.) This value is used to limit the NUMBEROFPOINTS in order to balance detail versus speed of display. Large values can cause delays of several minutes when displaying Java Tables with many points. In some cases, tables with many rows and many columns could cause the application to consume large amounts of CPU and even to 'freeze'. This item is most often used in a Style File.</p>
POINTSEVERY:	<pointsevery>	<p>Selects how much data should be summarized into each data point. This value is used in conjunction with the NUMBEROFPOINTS value to determine the level of summarization to use. (See NumberOfPoints)</p>
FORCEPOINTSEVERY:	<pointsevery>	<p>Specifies that the pointsevery value should be forced to this value. When this is specified it will override any pointsevery value selected by the user when drawing a graph.</p>
SUMFROMRAW:	See *NOTE	<p>Specifies that data from the MeasureWare data source be retrieved unsummarized from the system, then summarized by Performance Manager. This feature should only be used in cases where MeasureWare summarization of the data is not desirable. Using this feature can cause large volumes of data to be transferred from the system and can slow down the</p>

		resulting graph. This feature only has an effect on MeasureWare data sources. *NOTE: This keyword has no parameters, but must be used in the form "SUMFROMRAW=value" to maintain ovpmatch required parameter pattern of "key=value".
MINY:	<real number>	The minimum value on the left (default) Y-Axis. For graph type <i>gauge</i> , selects the minimum value on the gauge scale.
MAXY:	<real number>	The maximum value on the left (default) Y-Axis. For graph type <i>gauge</i> , selects the maximum value on the gauge scale.
MINYR:	<real number>	The minimum value on the right (optional) Y-Axis.
MAXYR:	<real number>	The maximum value on the right (optional) Y-Axis.
XML:		Returns data for the selected graph or table in XML format. Additional information may be supplied depending on the graph type. Graph type <i>gauge</i> will return the gauge settings (min, max, green, yellow, red, units etc.) Graph type "baseline" will return statistical values for each hour of each day of the week.

The following specifications are used with forecast type graphs

Planner (Forecast) Graph Specifications		
FORECASTTYPE:	smooth	Draws the original metric plus a smoothed metric line. Smoothing is performed by doing a seven point running average (3 before, the point itself and 3 after) on each data point and repeating a maximum of 10 times. No forecasting is performed for smoothed data (and so the "Forecast Ending" date is not utilized.)
	linear	Perform a least squares linear regression calculation to fit a straight line to the original data. This forecast line is then extended past the end of the data until it reaches the "Forecast Ending" date
	exp	While the Linear method fits most data, there are some metrics that are known to behave in a non-linear manner. For example, as system load increased, response times typically degrade in an exponential fashion. For this type of metric, the exponential

		curve method might be more appropriate.
	s	Some data is known to be limited to a certain range of values. We know that CPU Percent should never go below 0% or over 100%. Linear and exponential forecasts are not constrained and so will predict values that are unbounded. The "S-Curve" forecast will constrain a forecast line to lie between the "S-Curve Min" and "S-Curve Max" values by fitting to an "S" shaped curve.
	bus	All other trending methods attempt to predict the behavior of a metric as a function of time. If the metric is changing as a function of time then these methods are appropriate. If you know that a metric is not smoothly varying over time but is instead tracking some other metric, then the Business Units method can be used. See the Business Unit keywords BUDATASOURCE, BUSYSTEMNAME, BUEXTENDED, BUCLASS and BUMETRIC.
FORECASTDATEDAY: FORECASTDATEMONTH: FORECASTDATEYEAR:	<1-31> <1-12> <1970-2037>	The end of the forecast period. This date should always be after the graph End Date. The forecast ending date is not used for Trending Method "Smoothing".
SHOWVALIDATE:		Shows any data following the End Date, and prior to the Forecast Ending Date as validation data. It will be excluded from the modeling method but shown in order to verify whether the forecast matches the actual data.
INTERVALTYPE:	none	Do not draw upper and lower confidence limit curves for a forecast
	confidence	Upper and lower confidence lines will be drawn showing the range in which the average of the predicted values (the trend line) will fall for the given percentage of the time.
	prediction	Upper and lower prediction lines will be drawn showing the range in which the predicted values (all points) will fall for the given percentage of the time. The range between the prediction limits will always be larger than the confidence limits since they apply to all points and not just their averages.

CONFIDENCE:	<number>	The percentage used in displaying the confidence or prediction interval. Typical values are 70, 75, 80, 90, 95, 96
SEASONHOD:		Automatically determine whether Hour-of-the-Day seasonality will improve the forecast and include it if it does.
SEASONDOW:		Automatically determine whether Day-of-the-Week seasonality will improve the forecast and include it if it does.
SEASONWOM:		Automatically determine whether Week-of-the-Month seasonality will improve the forecast and include it if it does.
SEASONMOY:		Automatically determine whether Month-of-the-Year seasonality will improve the forecast and include it if it does.
BUDATASOURCE:	reporter mwa coda is db file	Specifies the data source for the Business Units metric. This is only used when the FORECASTYPE is "bus". See the DATASOURCE keyword for more details.
BUSYSTEMNAME:	<system name>	Specifies the system name for the Business Units metric. This is only used when the FORECASTYPE is "bus". See the SYSTEMNAME keyword for more details.
BUEXTENDED:	<datasource:system name>	Specifies the data source and system name for the Business Units metric. This is only used when the FORECASTYPE is "bus". See the DSEXTENDED keyword for more details.
BUCLASS:	<class name>	Specifies the class for the Business Units metric. This is only used when the FORECASTYPE is "bus". See the CLASS keyword for more details.
BUMETRIC:	<metric name>	Specifies the metric name for the Business Units metric. This is only used when the FORECASTYPE is "bus". See the METRIC keyword for more details.

The following specifications can be made prior to any metric, in which case they become the default specification. They can also be overridden on a metric-by-metric basis if desired.

Default Metric Specifications

DSN:	<name>	String used to identify a Named Data Source.
DATASOURCE:	reporter	The HP OpenView Reporter database is the datasource. The system name is optional.
	mwa	A MeasureWare agent is the datasource. In this case, the name of a system running the agent must be provided in the SYSTEMNAME specification.
	CODA	A CODA Agent such as VPPA for Linux, is the datasource. In this case, the name of a system running the CODA agent must be provided in the SYSTEMNAME specification.
	is	The OVIS product is the datasource. In this case, no system name needs to be provided in the SYSTEMNAME specification. If one is entered, it will be ignored.
	db	A generic database, provided by the user, is the datasource. In this case, the system name is interpreted as being the name of the database. This name should match one of the "Other Databases" configured in the HP OpenView Reporter product. Specify the database name using the SYSTEMNAME specification.
	file	A "flat file" data source. Specify the file name using the SYSTEMNAME specification.
		If no Datasource is provided, but one or more system names are provided, then the datasource will be chosen based on the agents on each system. If multiple agents are available on a system, the one used will be chosen in the following order [MWA, CODA].
SYSTEMNAME:	<system name>	<p>Specifies the name of the system that is running the MeasureWare agent. (This specification is required when DATASOURCE is set to MWA or CODA.) If multiple system names are entered (before any metrics are defined) then the entire graph will be repeated for each system.</p> <p>If the datasource specified is db then this field should contain the name of the database, as entered in the HP OpenView Reporter Configure Databases function.</p>
DSEXTENDED:	<datasource:systemname>	This keyword allows specifying the

		DATASOURCE and SYSTEMNAME together. The value supplied consists of the DATASOURCE value, a colon, and the SYSTEMNAME value. If no SYSTEMNAME value is required, then the DATASOURCE value alone will suffice. Valid extended data source values include DSEXTENDED: MWA:mssystem, DSEXTENDED: is, DSEXTENDED: db:mydatabase.
CLASS:	<i><class name></i>	Specifies the name of a class of metrics that is provided by the named data source. If the datasource does not expose this data class, then this metric will be ignored. Note that this specification does not have to be repeated for each metric if that metric is in the same class.
MISSINGDATA:	hidden	Data points where this metric is not available are not graphed. Note that unless markers are specified, it can appear that single data points surrounded by missing data are themselves missing (you cannot draw a line using only one point). Use markers in this situation.
	zero	Data points where this metric is not available are graphed with a value of zero.
	previous	Data points where this metric is not available are graphed with the value equal to their last non-missing value.
DATAINTERVAL:	<i><number></i>	Signifies that data is logged "sparsely" with zero values not logged. When specified, the <i><number></i> indicates the number of minutes between expected records. This specification will affect the averaging of metric values by forcing the number of records in a summary to include those that were expected, but not logged.
FILTER:	<i><filterstring></i>	Specifies a data filter that is applied to the data. This usually results in restricting the values for this metric to only those that meet the filter specification (See filters).
SFMETRIC:	<i><filtermetric></i>	The metric name used in a simple filter syntax of : <i><metricname> <comparecode> <value></i> . SFMETRIC:, SFCOMPARE: and SFVALUE: are used together instead of the more generic FILTER: (See filters).
SFCOMPARE:	<i><comparecode></i>	The comparison code used in a simple

		<p>filter syntax of: <code><metricname> <comparecode> <value></code>. Comparison codes must be one of the following:</p> <p><code>= != > >= < <= ~ !~</code></p> <p>SFMETRIC:, SFCOMPARE: and SFVALUE: are used together instead of the more generic FILTER: (See filters).</p>
SFVALUE:	<code><filter value></code>	<p>The value to be compared in a simple filter syntax of: <code><metricname> <comparecode> <value></code>. SFMETRIC:, SFCOMPARE: and SFVALUE: are used together instead of the more generic FILTER: (See filters).</p>
INSTANCE	<code><value></code> <code>*PROMPT*</code>	<p>For keyed multi-instance data classes, select only the records where the key value is equal to the value specified. If <code>*PROMPT*</code> is specified, then the user will be prompted to select from a list of available values when the graph is drawn.</p>
RANGEYELLOW:	<code><real number></code>	<p>For graph type <i>gauge</i>, selects the point where the gauge changes from green to yellow. If RANGERED is a larger value than RANGEYELLOW then the gauge will be drawn with green at the left and red at the right. Otherwise it will be drawn with red at the left and green at the right.</p>
RANGERED:	<code><real number></code>	<p>For graph type <i>gauge</i>, selects the point where the gauge changes from yellow to red. If RANGERED is a larger value than RANGEYELLOW then the gauge will be drawn with green at the left and red at the right. Otherwise it will be drawn with red at the left and green at the right.</p>
YAXISTITLE:	<code><units string></code>	<p>For graph type <i>gauge</i>, selects the Units label shown at the bottom of the gauge.</p>

The following specifications can only be made in reference to a particular metric. If the next metric does not specify them, then they will revert to their default values.

Metric Specifications		
METRIC:	<code><metricname></code>	Specifies the name of a metric that is provided by the data source. If this metric is not exposed inside the CLASS, then the metric will be ignored (See Metrics).
LABEL:	<code><label string></code>	Specifies the label to use when identifying this metric on the graph. The label can contain literal values and references to special substitution variables.
	Variable	will be replaced by

		<p>@@[LABEL]</p> <p>The metric label specified by the data source. (Such as "CPU %") If the data source does not provide a label, then the metric name will be used.</p>
		<p>@@[METRIC]</p> <p>The metric name (Such as "GBL_CPU_TOTAL_UTIL")</p>
		<p>@@[CLASS]</p> <p>The metric class (Such as "GLOBAL")</p>
		<p>@@[SYSTEM]</p> <p>The system name supplying the metric (Such as "mysys.net.com")</p>
		<p>@@[DATASOURCE]</p> <p>The data source for this metric (Such as "MWA", "CODA", or "RepSys")</p>
		<p>@@<i>metric</i></p> <p>The value of the <i>metric</i> metric from the same data source and class. For example, "@@BYDSK_DEVNAME" will label the metric with the value of the BYDSK_DEVNAME metric. If the BYDSK_DEVNAME metric value were "0", "Disk @@BYDSK_DEVNAME" would produce a label of "Disk 0" (See Labels).</p>
		<p>EXAMPLE: @@[SYSTEM]:@@[CLASS]:@@[METRIC] might produce a label like "mysys.net.com:GLOBAL:GBL_CPU_TOTAL_UTIL" while the default @@[LABEL] might produce a label like "CPU %".</p> <p>If no label is specified, then the metric will be labeled according to the default label specification. By default this is the same as @@[LABEL] but can be changed by editing configuration files on the OVPM server system.</p>
COLOR:	<color>	The color to use when graphing this metric. The default is to choose unique colors for each metric.
LINEWIDTH:	<number>	The number of pixels wide to draw the line for this metric. This specification is only valid when the GRAPHTYPE is line and the LINSTYLE is solid(does not apply to Java). The default is 1.
LINESTYLE:	solid	Solid line (LINEWIDTH is valid-does not apply to Java). Only used when GRAPHTYPE is line.
	dashed	A dashed line used when GRAPHTYPE is line.
	dotted	A dotted line used when GRAPHTYPE is line.
	dash-dotted	A dash-dotted line used when GRAPHTYPE is line.
	dash-dot-dotted	A dash-dot-dotted line used when GRAPHTYPE is line.

	bar	A bar. Only used when GRAPHTYPE is line or bar. Can be filled in in Java.
	area	A filled area. Only used when GRAPHTYPE is line or area.
FILLSTYLE:	solid	Solid pattern for pie segments and bars (affects Java Graphs only).
	25%	A twenty five percent filled pattern for pie segments, areas and bars (affects Java Graphs only).
	50%	A fifty percent filled pattern for pie segments, areas and bars (affects Java Graphs only).
	75%	A seventy five percent filled pattern for pie segments, areas and bars (affects Java Graphs only).
	horizontal	A horizontal pattern for pie segments, areas and bars (affects Java Graphs only).
	vertical	A vertical pattern for pie segments, areas and bars (affects Java Graphs only).
	45degrees	A forty five degree pattern for pie segments , areas and bars (affects Java Graphs only).
	135degrees	A 135 degree pattern for pie segments, areas and bars (affects Java Graphs only).
	diagonal	A diagonal pattern for pie segments, areas and bars (affects Java Graphs only).
	cross	A cross pattern for pie segments, areas and bars (affects Java Graphs only).
MARKER:	none	No markers are drawn on the line.
	auto	Automatically selects line markers for Java graphs. Only used when GRAPHTYPE is line.
	square	Draws an open square at every valid point. Only used when GRAPHTYPE is line.
	circle	Draws a circle at every valid point. Only used when GRAPHTYPE is line.
	triangle	Draws a triangle at every valid point. Only used when GRAPHTYPE is line.
	dot	Draws a sphere at every valid point. Only used when GRAPHTYPE is line.
	cube	Draws a cube at every valid point. Only used when GRAPHTYPE is line.
	diamond	Draws a diamond at every valid point. Only used when GRAPHTYPE is line.
	star	Draws a star at every valid point. Only used when GRAPHTYPE is line.
	cross	Draws a cross marker at every valid point. Only used when GRAPHTYPE is line.
	vertical	Draws a vertical marker at every valid point. Only used when GRAPHTYPE is line.
	horizontal	Draws a horizontal marker at every valid point. Only used

		when GRAPHTYPE is line.
SUMTYPE:	auto	Use the summation type provided by the data source (MWA and CODA) otherwise use "Average"
	average	Combine numeric values by summing then dividing by the number of values (mean). Non-numeric data will use the last value.
	sum	Combine numeric values by adding the values.
	minimum	Combine numeric values by taking the smallest value.
	maximum	Combine numeric values by taking the largest value.
	last	Combine numeric or non-numeric values by taking the last value received.
	count	Count the number of records from the data source that were combined into each data point.
YAXISRIGHT:		This metric should be drawn to scale with the right hand Y-Axis. If no metrics select YAXISRIGHT then the Right hand Y-Axis will not be drawn.

The following specifications apply only to tables. They indicate special post-processing that is done to the table data before it is displayed. When using the Java Table these specifications can be dynamically altered on order to get different views of the data without having to go back to the data source. If these settings are used in a table template, they will set the initial settings for Java tables and as final settings for HTML tables.

Table Specifications							
TABLEFILTER:	<p>Eliminate some of the data from the table based on the values of its fields. This functions much like the "FILTER:" parameter except that the FILTER: is done before the data is summarized and TABLEFILTER: is done after. The TABLEFILTER: also allows filter conditions to be separated by the double vertical bar " " in order to indicate either the preceding condition OR the next one can be true. This is in addition to the double ampersand "&&" indicating both the preceding condition AND the next one must be true. The TABLEFILTER: specification applies to the entire table and should precede any METRIC: specifications. (See filters)</p>						
TABLEHIGHLIGHT:	<p>This specification can be used to change the color of an item in a table. It should follow the METRIC: specification to which it applies.</p> <p>Syntax: <i>Condition Value Color [&& ...]</i> Example: >80 Red</p> <p>Multiple specifications may be specified separated by "&&". The last specification that is true will set the color to be used.</p> <p>Example: >5 Yellow && >10 Magenta && >80 Red</p> <table border="1" style="margin-left: 20px;"> <tr> <td>normal</td> <td>for <=5</td> </tr> <tr> <td>Yellow</td> <td>for >5 and <=10</td> </tr> <tr> <td>Magenta</td> <td>for >10 and <= 80</td> </tr> </table>	normal	for <=5	Yellow	for >5 and <=10	Magenta	for >10 and <= 80
normal	for <=5						
Yellow	for >5 and <=10						
Magenta	for >10 and <= 80						

	Red for > 80
Example: >80 Red && >10 Magenta && >5 Yellow	
	normal for <=5
	Yellow for >5

The following specifications carry special meanings as described:

Special Purpose Specifications		
JAVAGRAPHs:		Use the Java Interactive graphing facility for all graph types except gauge, baseline, xml.
ENDOFMETRIC:		This specification can be used in a graph template file to mark the end of specifications for one metric. A metric specification is also terminated whenever a new metric is started (DATASOURCE:, CLASS: or METRIC: keyword) and by the end of the graph.
END_GRAPH:		This specification is used to end one graph definition in a graph template file. This allows multiple graph definitions to be placed in a single template file.
FORMTEMPLATE:	<i><file name></i>	Specifies the name of a form template file that is used to construct a web page form. The default form template file is "VPI_Form.txt". All form template files are found in the <install_dir>newconfig directory. Only used when GETNEWFORM is specified.
NUMBEROFMETRICLINES:	<i><number></i>	When creating a new web page form, this is the number of individual metric specifications to create. Only used when GETNEWFORM is specified.
GETNEWFORM:		<p> Note: Some of these parameters are mutually exclusive.</p> <p>Informs the run time engine that instead of supplying a graph or table, it should return the HTML specifications for a web page form. These forms can then be used to request graphs and tables. If the FORMTEMPLATE parameter is specified, it will determine which form is loaded. Used in conjunction with NUMBEROFMETRICLINES, which indicate the number of metrics to be displayed on this form.</p>
*DRAWGRAPH:		Informs the run time engine that it should generate a graph or table.

*SAVECONFIG:		Informs the run time engine that it should save the current client configuration settings into a cookie on the client's system.
*SAVEGRAPHCMD:		Informs the run time engine that it should save the current graph parameters into the graph template file specified in the SAVETEMPLATENAME: parameter ("User" by default). The SAVEGRAPHNAME parameter must be supplied, and non-empty, to specify the name of the graph in the template file. If a graph with this name already exists, it will be overwritten. If not, then a new graph will be added to the template file.
*LOADGRAPHCMD:		Informs the run time engine that it should read the graph parameters from the graph template file specified in the LOADTEMPLATENAME: parameter. This command is used to load a form to match a saved graph template. The LOADGRAPHNAME: parameter must be supplied, and non-empty, to specify the name of the graph to be read from the template file.
*DELETEGRAPHCMD:		Informs the run time engine that it should remove the graph from the graph template file specified in the SAVETEMPLATENAME: parameter ("User" by default). The SAVEGRAPHNAME: parameter must be supplied, and non empty, to specify the name of the graph to be removed from the template file.
*LOADTEMPLATENAME:		The template name to be used in the LOADGRAPHCMD operation. If this parameter is empty or not supplied, then the template specified in the SAVETEMPLATENAME or the GRAPHTEMPLATE: parameter will be used. This alternative parameter name is furnished in order to simplify the LOAD/SAVE operations on complex forms where the GRAPHTEMPLATE: parameter may have other uses.
*SAVETEMPLATENAME:		The template name to be used in the SAVEGRAPHCMD or DELETEGRAPHCMD: operations. If this parameter is empty or not supplied, then the template specified in the LOADTEMPLATENAME or the GRAPHTEMPLATE: parameter will be used. All saved and deleted graphs will be written to the data directory under the name "VPI_GraphsUser<name>.txt" where <name> is the name specified in the

		SAVETEMPLATENAME: parameter. This alternative parameter name is furnished in order to simplify the LOAD/SAVE operations on complex forms where the GRAPHTEMPLATE: parameter may have other uses.
*LOADGRAPHNAME:		The graph name to be used in the LOADGRAPHCMD operations. If this parameter is empty or not supplied, then the template specified in the SAVEGRAPHNAME: or GRAPH: parameter will be used. This alternative parameter name is furnished in order to simplify the LOAD/SAVE operations on complex forms where the GRAPH: parameter may have other uses.
*SAVEGRAPHNAME:		The template name to be used in the SAVEGRAPHCMD and DELETEGRAPHCMD operations. If this parameter is empty or not supplied, then the template specified in the LOADGRAPHNAME: or GRAPH: parameter will be used. This alternative parameter name is furnished in order to simplify the LOAD/SAVE operations on complex forms where the GRAPH: parameter may have other uses.
*RESETFORM:		Informs the run time engine to redraw the form with no user-selected information. Use this option to "start over" on a form.
*DISCOVERSYSTEM:		Generate a System Information Report by interrogating the system for data sources it provides.
SKIN:	<skin name>	Specifies the skin directory that contains the style settings, special forms, and graph web page specifications for the "look and feel". See Using "Skins" to affect presentation styles .
NOMARKERS:		Primarily used in a Style file to turn off line markers on all graphs. This will override any markers set in pre-defined graph templates or in custom graphs.
MARKERS:	on off	Primarily used in a Style file and the user configuration information to turn line markers on or off for all graphs. This will override any markers set in pre-defined graph templates or in custom graphs.
NOTOOLBAR:		Primarily used in a Style file to turn off the toolbar on java graphs.
TOOLBAR:	on off	Primarily used in a Style file and the user configuration information to turn the toolbar on or off for java graphs.

FAMILY:	<i><Family Name></i>	Used in a graph template file to supply a more descriptive name for the graphs it contains. Be sure to choose a family name that uniquely identifies this set of graphs. Only one FAMILY: parameter should be used in each graph template file.
CATEGORY:	<i><Category Name></i>	Used in a graph template file to indicate a general category for the following graphs. Each graph will be associated with the category until a different CATEGORY: parameter is encountered.
DESCRIPTION:	<i><string></i>	Used in a graph template file to supply a descriptive string for a graph. A different DESCRIPTION can be provided for each graph.
GRAPHHEADER:		This parameter may only be used in a graph template file. It begins a section of one or more text lines that will be placed above the HTML graph. Terminate this section with a END_GRAPHHEADER: parameter. Accessible in Java from the Property page,(Annotations).
GRAPHTRAILER:		This parameter may only be used in a graph template file. It begins a section of one or more text lines that will be placed below the HTML graph. Terminate this section with a END_GRAPHTRAILER: parameter. Accessible in Java from the Property page,(Annotations).
DATEFORMAT:	<i><dateString></i>	This parameter overrides the date format used. Normally the date format is taken from the settings on the OVPM server. This keyword allows the user to specify their own format. This keyword is normally used with the SAVECONFIG command to set the client configuration options. See the Date Strings for more information on how to set this string.
TIMEFORMAT:	<i><timeString></i>	This parameter overrides the time format used. Normally the time format is taken from the settings on the OVPM server. This keyword allows the user to specify their own format. This keyword is normally used with the SAVECONFIG command to set the client configuration options. See the Time Strings for more information on how to set this string.
-init		Creates the initial htm file, "OVPM.htm" and customizes it to match the installation environment.

