hp OpenView

reporter

concepts guide

Reporter Documentation Set

📖 Chapter 1: Introducing HP OpenView Reporter

HP OpenView Reporter creates Web-based reports from data of targeted systems it "discovers."   Discovery of a system can occur if the system is running OpenView agent and subagent software. Agent software includes HP OpenView Performance Agent for Windows systems (MeasureWare Agent) or for UNIX systems. Reporter can also generate reports on systems managed by HP OpenView Operations [OVO]. After Reporter has run through its discovery, it gathers data based on pre-defined and user-specified lists of metrics, then formats the collected data into Web page reports.

Reporter reports can relate to ALL systems, system groups, and individual systems. Data is graphed and displayed in a variety of formats that allow you to quickly assess performance and use of the agent-installed systems.
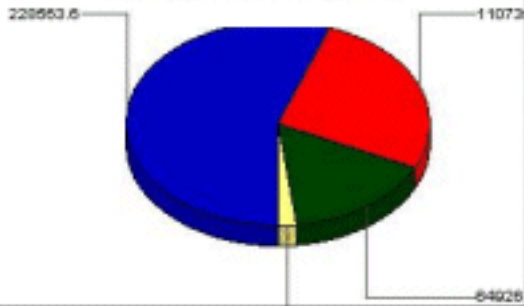
Top Applications by CPU for All Systems

This report was prepared: 9/6/01, 11:18:18 AM
Data is from 8/30/01 to 9/5/01

Top Applications by CPU

228563.6 — 11073

64925

Most Active Transactions
HP Perf Tools

Transaction Name	Completed	Aborted	Violations	SLO Percent
Scope_Get_Process_Metrics	55,693	0	2	0.00%
Scope_Get_Global_Metrics	11,147	0	13	0.12%
Scope_Log_Headers	54	0	0	0.00%
			15	0.02%

Top File System Utilizing Systems

Top Filespace Utilization Percent
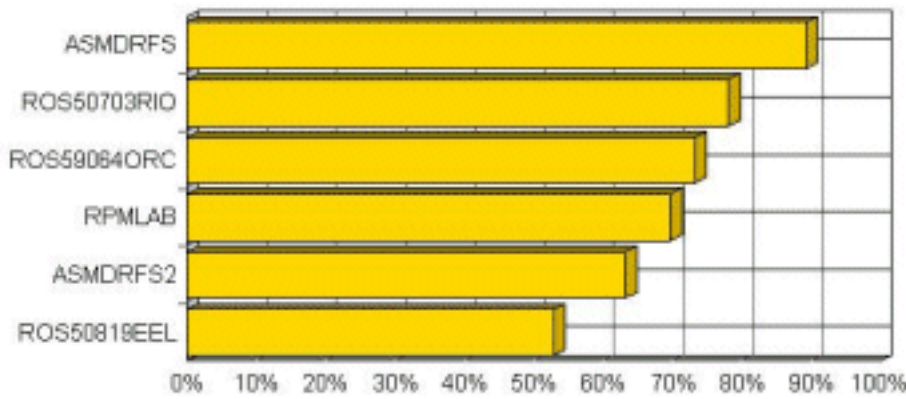
System Name: ASMDRFS

Even though Reporter is installed and runs on a Windows system, it can collect data from a variety of operating systems, such as HP-UX, Solaris, AIX, Linux, and Windows.

Customizations to Consider

After you have begin using Reporter, you may want to customize Reporter by organizing your systems into different system groups. You can assign reports to the new system groups and generate sets of reports that are immediately organized in a way that is relevant to your organization functions. Lastly, to create custom reports, you will need to purchase Crystal Reports to define new reports with data that you select (by creating metric lists).

In addition, even though Reporter has its own default database, which requires no configuration, it can be configured to store data in a more robust database such as Oracle. To customize data storage in this way, please see the *Installation and Special Configurations Guide*, where instructions are available for setting up MS SQL or Oracle as the Reporter database, as well as other custom configurations that you may find useful.

Where to Go from Here

If you haven't done so already, see Section 1 of the *Installation and Special Configurations Guide* for a description of what occurs when you install Reporter. After reading the material there, you will have an idea of how Reporter completes its daily routine. More on that topic is available in this Concepts Guide as well as how to use and customize the product. Along with the *Concepts Guide* and the *Installation and Special Configurations Guide*, within the Reporter interface, you will see Help buttons that can assist you in understanding Reporter dialogs and their required entries.

Thank you for purchasing Reporter; we are sure you will find it surprisingly easy to use and understand, with timely, relevant information that will keep you updated on a daily basis.

📖 Chapter 2: The HP OpenView Reporter Window

Reporter is made up of a number of programs that work together. These programs complete actions as follows: (1) discover systems that are running OpenView Performance agents for Windows systems or OpenView Operations (OVO) subagents on your network; (2) schedule periods of time during which data is collected on those systems or from the OVO database; and (3) generate reports that contain the collected data. When installed and configured, Reporter allows you to access reports through an automatically generated Web page that allows you to view information on the systems targeted for reporting.

The illustration below shows the three panes of the Reporter main window. Each component is described in the sections that follow.
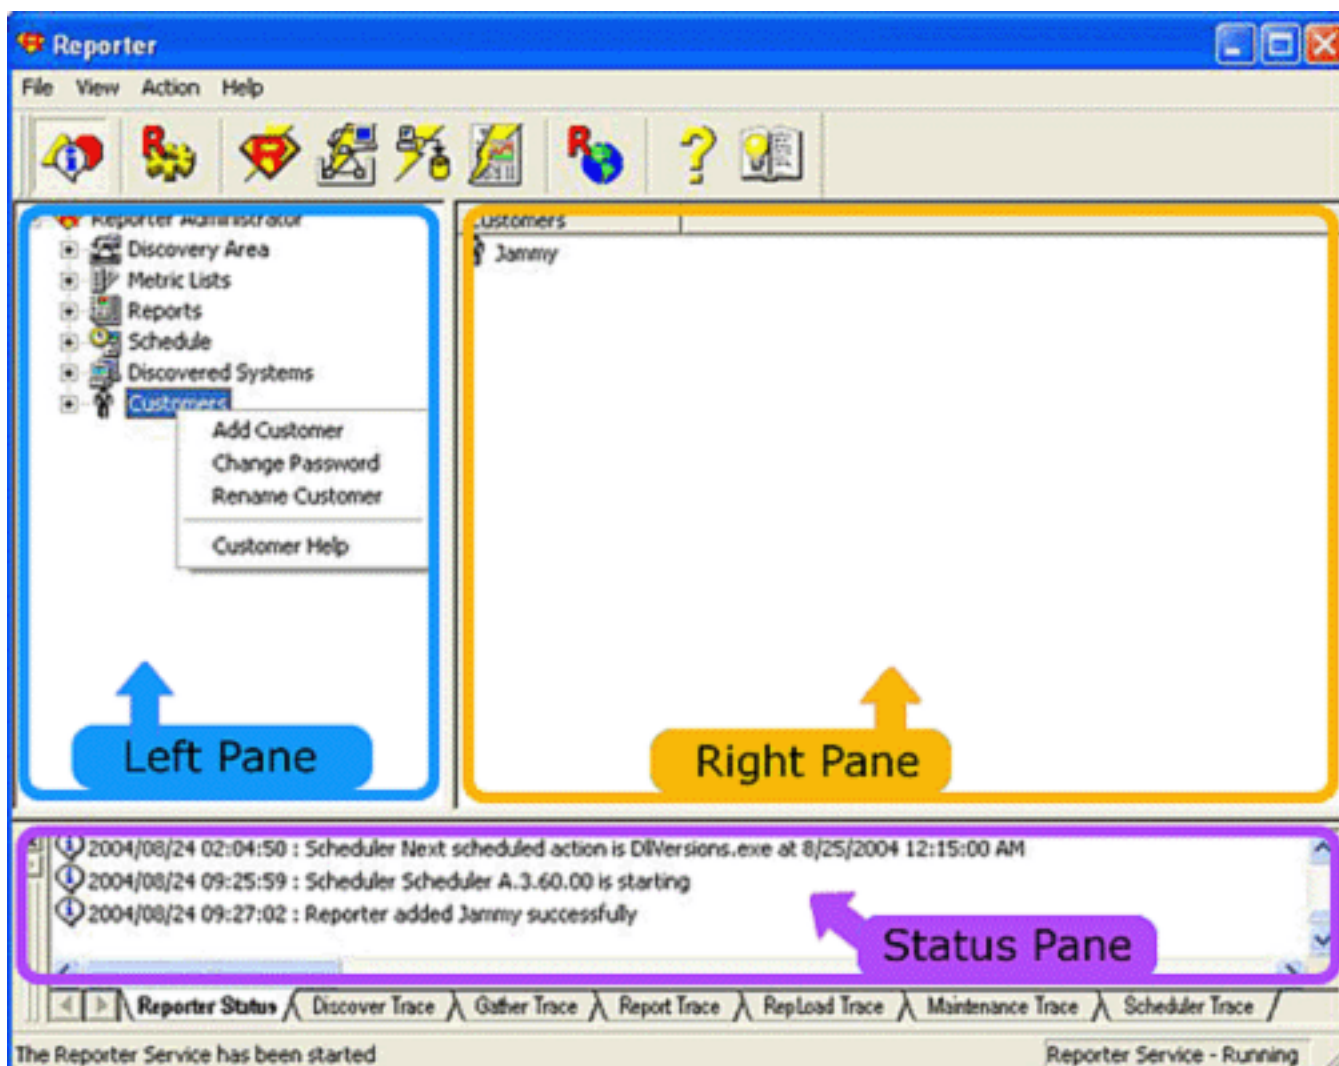
The Left Pane: The left pane contains the list of available items in a tree structure. Press the "plus" sign to expand part of the tree. You can select an item in the left pane to see more information about it in the right pane. You can also drag and drop items from the left pane to add them to selected items in the right pane. Use the left pane whenever you want to <u>display</u>, <u>add</u>, or <u>rename</u> items. You can perform these operations from the left pane:

- Drag and drop networks, domains, and systems to add to the Discovery List that the Reporter searches for systems with OpenView agents.
- Display available metric lists
- Display available reports
- Display the scheduled actions.
- Display the systems that have been discovered.
- Add new system groups, report definitions, metric lists and schedule actions
- Display configured customers and add new customers

The Right Pane: The right pane shows specifics about the item selected in the left pane. Depending on the item selected, you can <u>edit</u> and <u>delete</u> items from the right pane. You can perform these operations from the right pane:

- Delete systems, domains, and networks from the list that the Reporter searches for OpenView agents.
- Edit some settings and definitions
- View the metric lists, reports, system information, and system group for a specific system
- Assign systems, reports, and metric lists to system groups
- View the metrics contained in a metric list
- Assign or remove systems from customers

The Status Pane: The Status Pane is a read-only window that gives real-time information on the status of all the programs within the Reporter. To help navigate through the information displayed, each line entry is tagged as follows:

    Information Only

    Error

    Warning

    Trace information
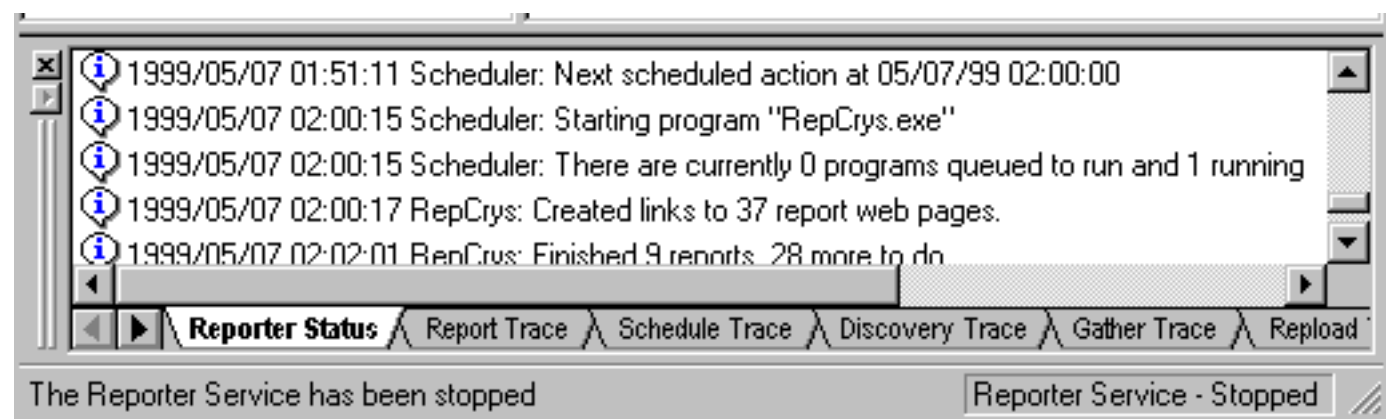
The status pane (window) offers flexibility by allowing you to do the following:

- **Disconnect or hide the status pane:** You can disconnect the Status pane from the Reporter window and move it anywhere on the Desktop. You can also hide the Status pane through the View menu or the toolbar.
- **Display trace file information:** You can show additional information on Reporter activities by selecting specific trace files from which to display information. The illustration above shows tracing file tabs starting with Reporter Trace (shows information related to report generation). To use this feature, see the online Help topic "Configure the Status Window."

The Reporter Toolbar

Reporter's toolbar allows you to quickly initiate the following actions:

Hides/shows the status pane.

Starts/Stops the Reporter Service (must be activated before any of the programs can start; the program tool buttons appear disabled (colorless) if the Reporter Service is inactive).

Starts all Reporter programs (actions) immediately.

Starts the Discovery programs immediately.

Starts the Gatherer programs immediately.

Generates reports immediately.

Shows the reports list on the Reporter Web page.

Opens online Help.

Opens the Web page that links to all Reporter documents (Reporter Documentation set) which includes: the Release Notes, *Concepts Guide*, and *Installation and Special Configurations Guide*.

The Reporter Hierarchy

The left pane of the Reporter main window contains a hierarchy that allows you to view systems and change configurations. The sections that follow describe in general terms each specific area of the hierarchy and how you can use it. For specific information for the tasks you can complete using this hierarchy, refer to the Reporter online Help.
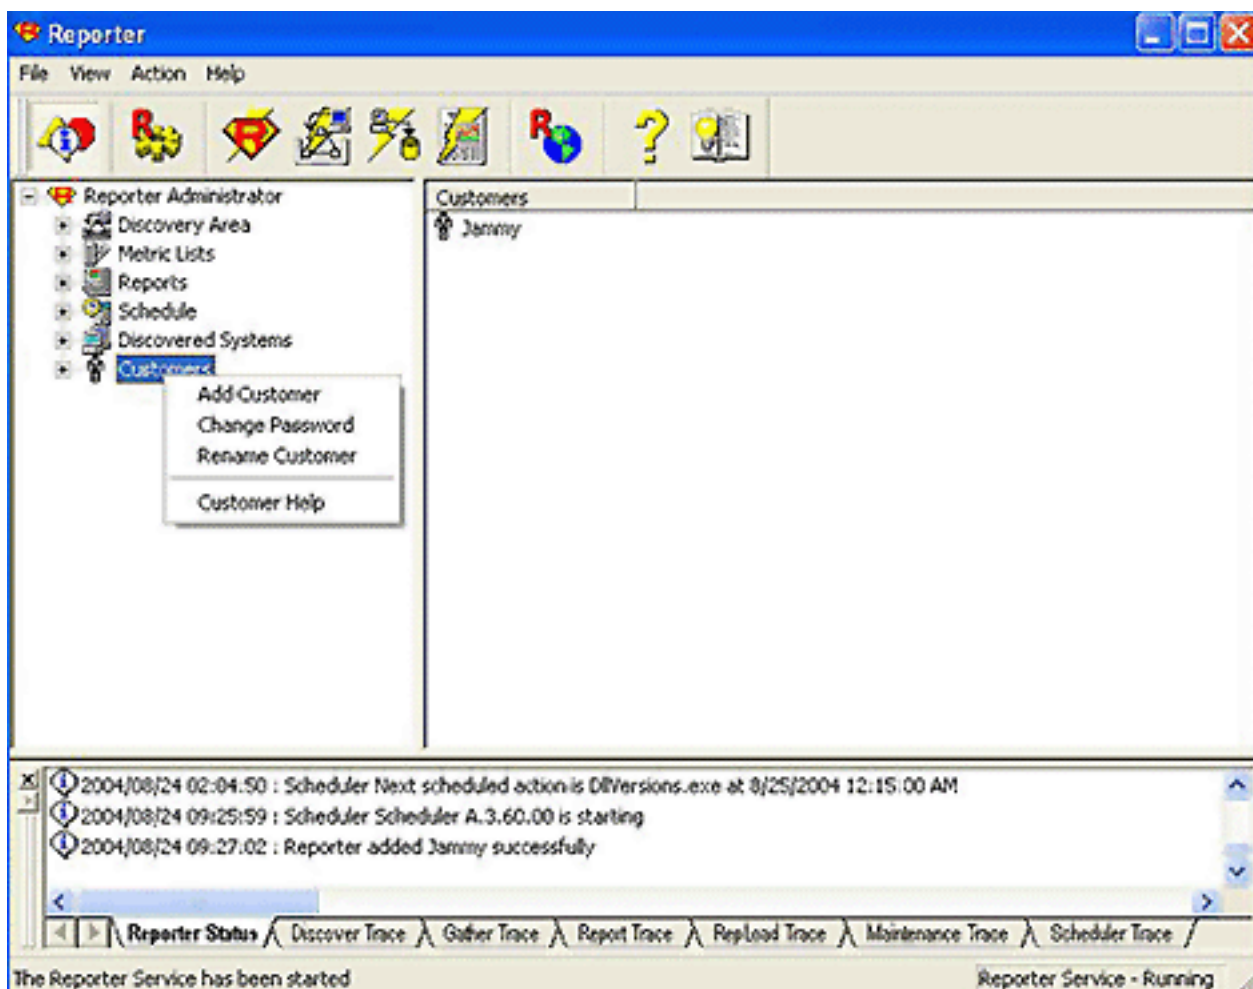
Discovery Area The Discovery Area is used to limit the search for OpenView Performance Agents. The networks, domains and systems that appear beneath the Discovery Area icon are those that are known to Windows networking. Discovery can be extended outside the networks known to Windows in several ways described in later sections.

The Discovery Area is not used for discovery of OpenViewOperations (OVO) systems. The OVO software keeps track of its systems in a central database. The Reporter copies this information from the OVO database.

Metric Lists The metric lists that appear beneath the Metric Lists icon are used to select the information that is gathered from system agents. This data is summarized and written into the Reporter database for easy reporting. Metric lists for OpenView Performance Agents are modifiable and new metric lists can be created for additional OpenView Performance Agent data sources (if you create new metrics, DO NOT include

spaces in them as they will not be processed correctly). The metric lists for OVO agents are not modifiable.

You can assign metric lists to systems or groups of systems to control which data is gathered from them. By default, the standard metric lists are assigned to be gathered from all systems.

Reports The Report Families that appear beneath the Reports icon can be expanded to see individual report definitions. You can add a report to a system, group of systems or to the special group ALL to specify which systems to include in the report.

The colored squares in the corners of the report icons indicate the kinds of groupings this report supports (all systems [red], group [blue], or single system [green]). If you click a report you should see its definition in the right pane.

Reports are formatted in hypertext markup language (HTML) and are viewable through your Web browser.

If you choose to use a supplementary reporting application (Crystal Reports™), you can create your own reports or modify the reports furnished with the Reporter.

Schedule The entries beneath the Schedule icon are program actions configured to start at specified times. They are named for the actions they complete; for example, the Discover schedule starts the discovery process, while the Gather schedule starts the data collection process. You can change schedule times for all scheduled actions, or you can change the schedule time for one scheduled action. See the "Schedule" section in "How Reporter Works" for details on how to make such changes.

Discovered Systems The systems that have been discovered appear in groups beneath the Discovered Systems icon. Systems may be grouped to make them easier to locate, and to provide collections for reporting purposes. A system can belong to more than one group and all systems belong to the group named ALL.

Systems are added to groups automatically as they are discovered. Systems can show up in different groups based on any of these two methods: (1) You can create new groups and then place systems into those groups manually within the Reporter window. You place systems into the groups by double-clicking the system group name so that it appears in the right pane. Then you can drag 'n drop individual systems from the left pane to the right pane where the group appears. (2) You can also set up automatic grouping of systems by defining parameters for grouping. Then when Reporter discovers a system, it can place this system into the appropriate group, based on the parameters you have defined. Instructions for this method are in Chapter 7, Advanced

Topics, By default, an automatic grouping is defined to add each system into a group based on its operating system.

You can select to place all newly discovered systems into a group named NEW to assist in manual grouping. If you are using OVO agents, you can copy the OVO nodegroups into the Reporter groups.

Customers The entries beneath the Customers icon show all configured customers. You can assign a system(s), report(s) to one or more customers for reporting purposes. When you expand a specific customer folder, you can view each system(s) and report(s) that has been assigned to that customer.

📖 Chapter 3: How HP OpenView Reporter Works

Reporter's main window shows you a hierarchy that contains Reporter's configurable components. These components are: The network area to search for systems that are running OpenView Performance agents, reports, scheduled actions and discovered systems. This section offers you more details on the configurable components listed below that are contained in Reporter's main window and also discusses how Reporter's actions (programs) are run and what happens during that time.

- [Discovery Area](#)
- [Metric Lists](#)
- [Reports](#)
- [Schedule](#)
- [Discovered Systems](#)
- [Customers (customer-segmented reports)](#)
- [When Reporter's Programs Work](#)
- [Re-scheduling Data Gathering for Specific Systems](#)

Discovery Area Before Reporter can access a system, it must be "discovered." Reporter is designed to discover Windows or UNIX systems through a couple of different methods. After a system is discovered, it is added to the Discovered Systems area of the Reporter hierarchy. Each time an existing system is re-discovered, its information is updated in the Reporter database.

To discover systems that are running OpenView agents, you have three methods to choose from: network browsing, single system, and bulk discovery. You use the Discovery Area (in the Reporter main window) for the first two methods. Bulk discovery is discussed fully in "Customizing Reporter," Chapter 4, of this document.

Discovering systems through **network browsing** relies on the Windows operating system, which lists all systems known to the Windows network. It is possible for Windows network to know many thousands of systems so you will probably want to restrict the systems to be searched. Click the **Discovery Area** icon in the left pane to see the list of

current networks and domains for network browsing. Under the Discovery Area in the left pane you can click the plus sign next to the icons to expand the hierarchy to see the networks, domains and systems known to the Windows operating system. This refers to Windows nodes only. UNIX nodes will not be discovered automatically. Drag to the right pane the network, domain, or system you would like searched.

NOTE: Systems that have aliases (multiple names) will have multiple appearances in the Discovery area. However, this situation does not create a problem as Reporter recognizes only one of those systems when discovery occurs. The system is reported on with the name used for OpenView agent data collection. If you set your trace level to "1," you can check the trace.Discovery file to determine if multiple system names were discovered for a single system. (To change trace level, select **Configure** from the File menu in the main window.)

Discovering **single systems** is not reliant on the Windows operating system. You can add individual systems to the Discovery List by right-clicking the Discovery Area icon and selecting **Add Single System** from the drop-down menu shortcut that appears. After you enter the system name, it is included in the list of systems available for Discovery in the right pane.

To remove an item that appears in the Discovery List (the right pane of the Reporter window), right-click the system name in the right pane and select **Delete** from the drop-down menu shortcut that appears.

Metric Lists Metric lists control what information is gathered from a system into the Reporter's database. A metric list groups metrics from a single metric class supplied by the OpenView agent for UNIX or Windows. The metric list can also select the degree of summarization (points every 5 minutes, hour, day, etc.) and how much data to gather and retain in the database. The shorter the interval, the more records collected. The default summarization level is one hour.

Reporter includes default metric lists that target data collection for standard OpenView agent metric classes and OpenView Operations messages. If you choose to create your own metric lists, you need to create new report definitions and templates to generate reports based on these customized lists. One metric naming requirement is that the name for the new metric NOT include embedded spaces. *NOTE: To create report templates, you must use Crystal Reports, a separate product that you can purchase from Business Objects.*

You can assign metric lists to systems or groups of systems by selecting the systems or groups from the "Discovered Systems" icon in the scope pane, then dragging the metric

list onto the right pane.

Reports Available reports are listed in families under the Reports icon. Expand the Reports hierarchy in the Reporter window and click on a report family to display in the right pane the reports it contains. You can also delete the report when it is in the right pane. Click on an individual report to display its definition details. You can edit the report definition details in the right pane including the Date Range and Shift Name.

Each report can include data from:

	All systems in the database (**All Report**—color code red)
	Selected system groups (**Group Report**—color code blue)
	Single systems (**System Report**—color code green)
	Combinations of the above

Each report definition must name a report template for one, two, or all three report types. Right-click the Reports icon and select **Add Report** from the drop-down item to create a new report definition.

In order to modify these report templates or add your own, you must purchase Crystal Reports™, which allows you to create additional templates. Chapter 4, "Generated Reports," describes how to use Crystal Reports to create templates to use with Reporter.

To view the last generated reports, click the Reports button on the tool bar or right-click the name Reports in the left pane and select **Show Reports**. You can also run a single report by right-clicking the report in the left pane and selecting **Run Report**.

You can assign reports to systems, to system groups, or to all systems. Click the system or group under the Discovered Systems icon in the left pane to display the system or group in the right pane and then drag the report onto the right pane. You can easily see if the report is meant for a single system (by the green square attached to it), a system group (blue square), or all systems (red square) and assign it accordingly.

Customer Enabled Reports: If the selection formula have any one of the following columns in the SQL query, the report will be customer enabled.

- CUSTOMER_NAME
- SYSTEMNAME
- SYSTEMID

Only customer enabled reports and reports with Group and All templates can be assigned to customers. Reports having System templates cannot be assigned to the customers.

Schedule The Schedule controls when actions occur within Reporter. Select the Schedule icon to display a list of scheduled actions (Reporter programs) in the right pane. You can restrict this list to actions of a particular type by clicking the plus sign next to the Schedule icon to expand the hierarchy and select one of the items under it. Normally all scheduling happens automatically, but you can modify scheduled actions when necessary. To modify a single scheduled action, double-click the entry for that action in the right pane. To add a new action, right-click the Schedule icon and select the **Add Schedule** from the drop-down menu shortcut that appears.

You can right-click the Schedule icon and select **Global Settings** to change the date for all scheduled actions. This is most useful when you want Reporter to delay all its scheduled actions until a future date. (Just choose Global Settings and change the Job Date to the desired date.) You can change the date and times for all executions of a particular type of action by right-clicking the action type (for example, Discover, Gather, Maintenance) under the Schedule icon and selecting **Global Settings** from the drop-down menu shortcut.

Discovered Systems This section shows systems that have been discovered by the discovery programs. Discovered systems are placed in groups for easier access and reporting. Systems can be assigned to groups as they are discovered as specified in the [Automatic Grouping facility](). By default, systems are automatically assigned to groups based on the name of their operating system. You can also manually assign systems to groups. A system can belong to any number of different groups. You can create your own groups and assign systems to them manually using drag and drop.

Special Groups:

- **ALL** - This group contains all discovered systems. Deleting a system from the ALL group removes all information about that system from the Reporter database.

- **NEW** - This group contains systems that are newly discovered. You can enable this group by selecting Options from the File>Configure menu. This group is useful if you are manually assigning systems to groups.
- **Nodegroups** – This group is part of the OVO configuration. OVO assigns systems to specific nodegroups. You can have these nodegroups copied to the Reporter groups by selecting Options from the File>Configure menu.

To add metric lists and reports to new system groups or to individual systems, click on the group or system under the Discovered Systems icon, then drag and drop the metric list or report from the left pane to the right.

To see more about a group or a system, select its name in the left pane. You can click the plus signs in the right pane to expand the information. Single systems contain useful information under the **System Information** icon.

You can delete a system from a group by selecting the group in the left pane and expanding the **Systems** icon in the right pane. Right-click the system and select **Delete** from the drop-down shortcut menu.

You can delete a group (except for the group ALL) by selecting Discovered Systems in the left pane, then selecting the group(s) to delete in the right pane and pressing the delete key (or right clicking and selecting the delete menu item).

Customers (customer-segmented reports)

This section shows the specific customers and discovered systems associated with the customer. Specific customers created under the Customers icon are custom-created by default, there are no defined customers. This feature comes into play when reports are generated. When you choose to use this feature by adding customers and assigning system(s), report(s) to them, a separate set of reports is generated for that customer and stored in a separate directory tree.

To view the system(s), report(s) assigned to a customer, expand the **Customers** node and then expand the node for the individual customer. A list of the system(s), report(s) owned by that customer will appear below the customer.

For more information on how to create customers and assign system(s), report(s) to customers, please refer to Chapter 5.

Reporter Features

Reporter has many features that operate automatically from the time you install it. Reporter has been preconfigured to automatically run specific programs each night. These programs allow Reporter to discover new systems, add them to the database, gather and store those systems' data within the data base, and generate reports based on the newly gathered data. Reporter can operate successfully with no user interaction since it contains a schedule that automates the running of these programs. However, you can change the time at which any program is run using the Edit Schedule Entry dialog:

Program Scheduling

Reporter's scheduling program (the Master Scheduler) is driven from a schedule table inside the central reporting database. Each entry in the schedule table describes when to run a program, what parameters to use, and how often to repeat the action. When no programs are scheduled to run, the scheduler is inactive.

The **Master Scheduler** has some special features that attempt to avoid overburdening the Reporter system with too many programs running at once. These features space out the running programs so that the system is still usable for other work.
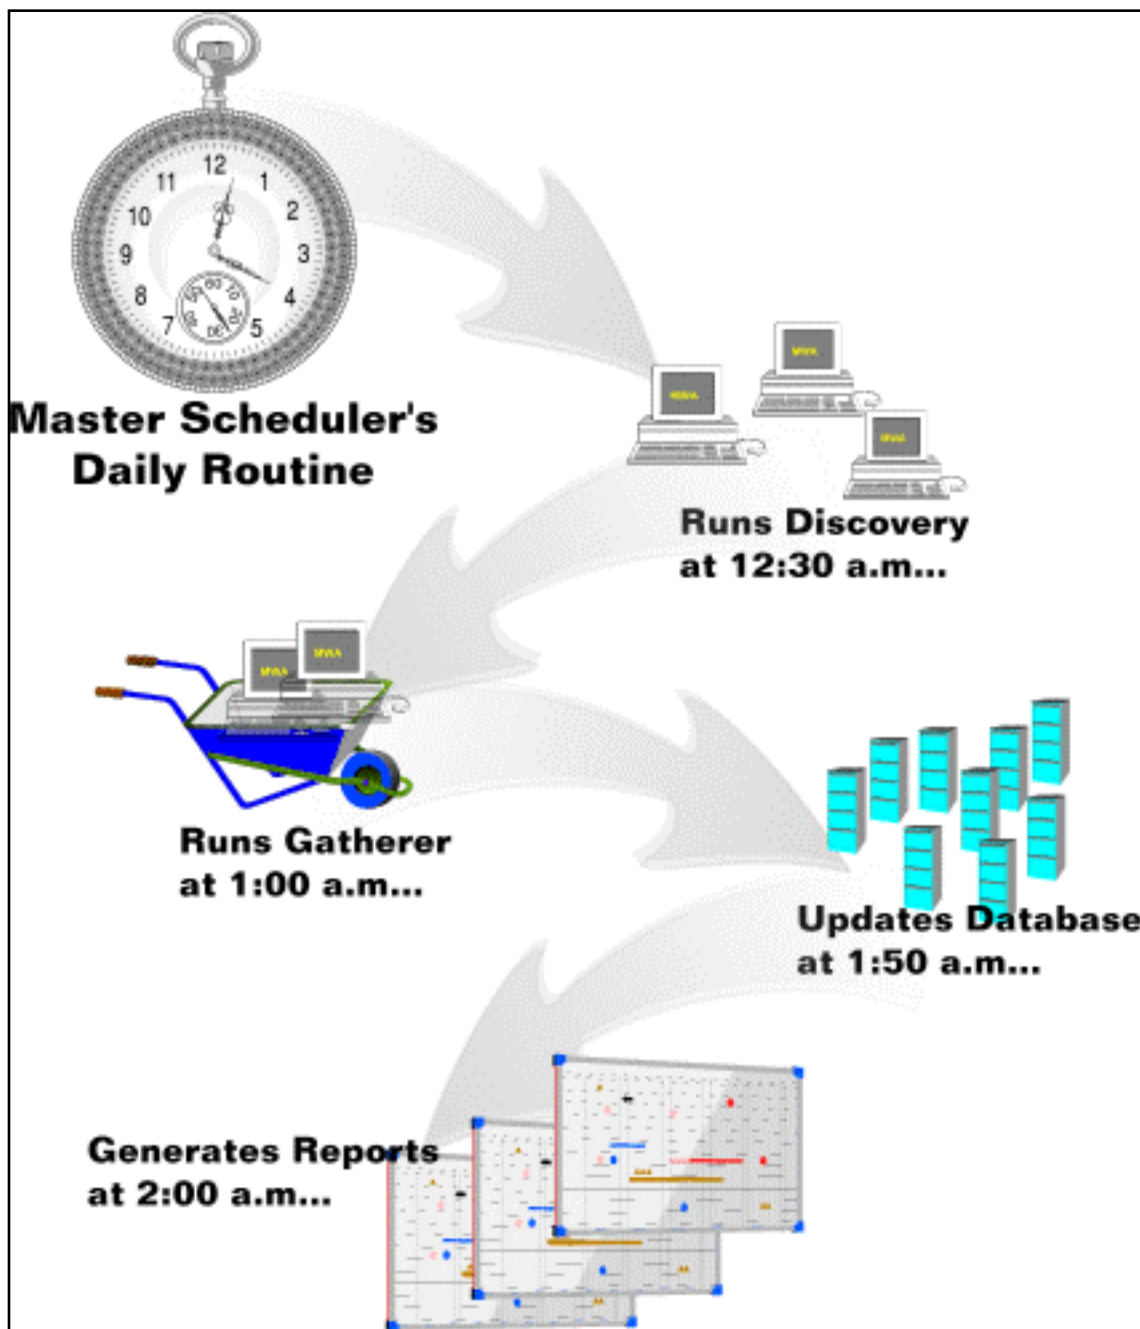
Discovery Programs

Each night at 12:30 a.m.. (default time) the Discovery programs run.

Discover_ITO Program

Discover_ITO finds system nodes in the OVO database. It copies this information and optionally the OVO nodegroup information into the Reporter database.

Discovery Program

Discovery examines the network and looks for targeted systems that are running OpenView agent. When Reporter finds a new system, it retrieves information about the system and its data sources. It adds this information to the Reporter database. It also schedules the execution of the Gather program to collect data for this system each night.

Gather Programs

Each night, starting at 1:00 a.m.(default time), the Gather programs run.

Gather_ITO Program (program name ITO is actually referring to OVO)

This programs reads information from the OVO database and writes summaries to the Reporter database. Because summaries take less space than the original information, they can be held for a longer period of time.

Gather Program

This program uses information in the Reporter database to determine what metrics to gather from the OpenView Performance agent data available in each system.This information is in the form of *metric lists* which define the metrics to be gathered from every data source. Reporter uses this information to collect data from each discovered system and adds this data to the Reporter database.

The Gather program executes every night, once for each discovered OpenView Performance agent system. It only gathers information that is not already in the Reporter database.

Database Maintenance Program

Each night at 1:50 a.m. (or following the completion of the last Gather program) the database maintenance program (RepMaint) starts.

RepMaint program

This program performs routine database maintenance such as deleting old data.

Report & Web Page Link Generation

From the data that is stored in the database and the report definitions, which determine the specific reports to generate for specific systems, Reporter is able to generate reports. In the last step, the report program builds a web page to link to all the HTML reports that Reporter generates. Most reports are generated during the night from summarized data saved in the Reporter database. You also have the option of generating reports immediately from this same data by selecting **Reports on Demand**, listed among the report links on Reporter's cover Web page.

Other reports, however, are generated directly from other databases, such as the OVO database. These reports show data which is not easily summarizable. You can identify **direct reports** by a non-blank Database Name field in the Report Definitions dialog box.

Modify the Gather Schedule to Multiple Times Daily

OpenView Reporter performs batch reporting of data every night. However, it is possible to generate reports more frequently. For example, you may have a few critical server systems on which you want to see data updated every hour. To generate reports more often, Reporter can read the collected data more frequently than once a day. You can make this happen by using the Edit Schedule Entry dialog to re-schedule *gather* to show the data up to the last hour and run the gather program every hour. The two parameters you can use are **-h** (for showing all the data up to the last hour) and *<system>* (for naming the system on which you want the data shown).

To re-schedule data gathering for a specific system to an hourly basis:

1. Open the Reporter main window and in the left pane double-click **Schedule** (or right-click the plus next to Schedule) to expand the folder.
2. Still in the left pane select **Gather**.
   (In the right pane each system's scheduled gather jobs appear. The system name is in the Parameters column.)
3. In the right pane, right-click the system whose gather job you want to change and select **Edit**.
4. In the Edit Schedule Entry dialog box, type -h in the Parameters box to precede the system name. In the example, " -h" precedes "haggin" to gather all collected data up to the last hour for system "haggin." Example: gather -h "haggin"
   In the Job Reset box, click the down-arrow and select **Hour** to restart the job every hour

📖 Chapter 4: Generated and Custom Reports

This section covers the following:

Reporter comes with a wide variety of report templates that you can see under the Reports icon. These templates are used to create report definitions based on the data gathered from discovered systems.

💡 NOTE: If you are using another OpenView product, the product may include a **Report Package** that is either automatically installed or can be added using the Reporter main window (choose **File**>**Configure**>**Report Packages**). Please see the online Help for details.

💡 NOTE: If you create your own reports by following the guidelines given in the **Create your own reports** section, and encounter a problem with those reports, HP support will not fix any design issue with crystal templates. HP support is responsible only for reports templates provided by HP.

Reports are grouped into **Report Families** for easy locations. You can add and delete report families and move reports between them using drag and drop. All reports belong to the ALL report family. Deleting a report from the ALL report family removes it entirely from the Reporter database. Deleting a report from any other report family affects only that family.

Report Type Color Coding: Reports are color coded with small squares to indicate if you can apply the report to **ALL** systems (**red** in the upper left corner), to a system **GROUP** (**blue** in the upper right corner), or to a single **SYSTEM** (**green** in the lower left corner).You can identify the report contents by the name displayed; for example the report Discovered Systems provides you with a list of all the systems that Reporter has discovered and on which reports are currently available.

Instructions are available at the end of this chapter for how to create custom reports using the Business Objects generating software Crystal Reports™ .

Generating Reports

Each night starting at 2:00 AM (following the update of the Reporter database by the RepMaint program), the report generation program (RepCrys*)* runs. This program uses a series of Crystal Reports™ templates to produce reports from the data in the Reporter database. These reports are stored in HTML format as .htm files and are suitable for viewing with web browsers. This program also creates a series of web pages that contain links that, when clicked, display the named report.

To view the latest generated reports, click the Reports button on the tool bar or right-click the name Reports in the left pane and select **Show Reports**. You can also run a single report by right-clicking the report in the left pane and selecting **Run Report**.

Generating Reports on Demand

On Reporter's cover Web page, which lists all available reports, the **Reports on Demand** selection allows you to launch a utility within a Web page where you can choose the time, date, and system on which to create an immediate report. With this feature you can show data that will assist you in investigating performance and availability issues for which you might need a quick answer.

Reports on Demand also lets you consolidate data in ways different from the standard reports you generate with Reporter every night. Just click the link and explore the very customized way in which you can generate a report.

Using On Demand Reports according to Customer (using Customer Segmented Views)

On Demand Reports can be used with [Customer Segmented Reports](#) to generate data for a specific customer. When you configure customers and assign systems to them in the Reporter main window, this information is saved in the database. After the links on the Reporter Web page are generated, each customer has his/her own link to run On Demand Reports and see reports on only that specific list of systems. This Reporter feature builds reports that show data only from the Reporter database.

On Demand Reports Status File

When On Demand Reports runs, warning and errors are logged to a text file called "status.PM." This file is located in *<installation_directory>*\data\. Each line in the status file is marked with the date and time that the warning/error was generated.

Enabling Tracing

If you need more detail than is contained in the status file, you can enable tracing for the On Demand Reports program (Analyzer.exe). When tracing is enabled, the output will to into a file named "trace.Analyzer".

To set the tracing level, open an MS-DOS console window and execute:

```
<installation_directory>\HPOV_IOPS\cgi-bin\Analyzer.exe -trace <trace_level>
```

where *<trace_level>* is a number in the range of 0 (no tracing) to 9 (maximum tracing detail).
You do not need to reboot or restart any programs. Tracing takes effect immediately once it is set.

Troubleshooting Java Graph Problems

To troubleshoot the Java graph problems, complete the following steps.

1. Open the Control Panel.
2. Select **Java Plugin** (the version you are using with On Demand Reports).
3. In the Java Plugin Control Panel, if the **Basic** tab is not already selected, select it.
4. Check **Show Java Console**.
5. Click **Apply**, then close.
6. Close and reopen the Web browser.
7. Generate a Java graph to reproduce the problem.
8. In the Java console window copy the text from this window and send it to your support person.

Location of Default and User Graph Template Files

Graph template files are named VPI_Graphs*<name>*.txt where "name" is the name you see in the standard graph form. Default graph template files are located in the *<installation_directory>*\newconfig. User-defined graph template files are located in *<installation_directory>*\data\.

Included Reports

The list of reports below are those that are included with Reporter, generated for discovered systems. When you add other OpenView products, those products generally include a report package that mean more generated reports with additional

information.

- **Discovered Systems**: This report displays a table of information about each system that is assigned to your report group.
- **Most Active Transactions**: This report shows the transaction rates for various applications running on the systems assigned to your reporting group. Graphs are provided to illustrate the information provided within the report table. The table columns show the following for each active application:

> **Completed** = the number of transactions completed successfully.
>
> **Aborted** = the number of transactions that have been aborted instead of completed.
>
> **Violations** = the number of transactions completed successfully that exceeded their Service Level Objective (SLO).
>
> **SLO Percent** = the number of Service Level Objective violations as a percent of those completed successfully.
>
> **SLO Limi**t = the Service Level Objective time limit in seconds.
>
> **Response Seconds** = the average time to complete a successful transaction.
>
> **Abort Seconds** = the average time taken by an aborted transaction.

- **System Up Time**: This report displays a series of tables showing Up-Time and Down-Time statistics for each system in your group. The data shown for All-Day Up Time percent is based upon 24-hour days for the range of dates in the database. The data is also broken out for each scheduled work shift.
- **Systems Performing Each Transaction**: The type of data being monitored is customized for your group. This report organizes information by Application, Transaction, then System. This data is useful when comparing application performance on different systems. The report may illustrate that hardware, software, or configuration differences are affecting system performance. The table columns show the following for each type of data being monitored:

> **Completed** = the number of transactions completed successfully.
>
> **Aborted** = the number of transactions that have been aborted instead of completed.
>
> **Response** = the average time (seconds) to complete a successful transaction.
>
> **SLO Percent** = the number of Service Level Objective violations as a percent of those completed successfully.

- **Top Applications by CPU**: This report shows the percent of CPU being used on each system assigned to your group, by application. The totals for all systems are shown in a pie chart. The data for each individual system is shown in a table.
- **Top Busy CPU Systems**: This report shows which systems (assigned to your group) consumed the most CPU cycles during the reporting interval. A graph shows the average CPU busy percentage for each system. Individual system graphs show more detailed information about each system's consumed CPU cycles over time.
- **Top Busy Disk Systems**: This report shows which systems (assigned to your group) performed the most disk transfers during the reporting interval. A graph shows the total number of disk transfers for each system. Individual system graphs show more detailed information about each system's disk transfer rates over time.
- **Top Busy Memory Systems** : This report shows which systems (assigned to your group) had the highest rate of disk-swapping activity during the reporting interval. High disk-swapping activity indicates that more RAM is desirable on the system. A graph shows the average number of page-outs per second for each system. Individual system graphs show more detailed information about each system's memory page-out rates over time.
- **Top Busy Network Systems**: This report shows which systems (assigned to your group) had the highest network packet transfer rate during the reporting interval. A graph shows the average number of network packet transfers per second for each system. Individual system graphs show more detailed information about each system's network packet transfer rate over time.
- **Top File System Utilizing Systems**: This report shows which systems (assigned to your group) consumed the highest percent of their file space during the reporting interval. A graph shows the average file space utilization for each system. Individual system graphs show more detailed information about each system's average file space utilization rate over time.
- **Transactions by System**: The type of data being monitored is customized for your group. This report organizes information by System, Application, then Transaction. This data is useful when comparing the performance of different applications on one system. The report may indicate that certain software needs to be reconfigured for optimal performance. The table columns show the following for each type of data being monitored:

> **Completed** = the number of transactions completed successfully.
>
> **Aborted** = the number of transactions that have been aborted instead of completed.
>
> **Response** = the average time (seconds) to complete a successful transaction.
>
> **SLO Percent** = the number of Service Level Objective violations as a percent of those completed successfully.

Adding Report Packages for other OpenView Products

Some OpenView products include report packages that, once installed, automatically load a set of report templates that Reporter uses to generate reports for the newly installed product. However, in some cases you will have to manually select

the product's report package within the Reporter main window.

To select the Report Package:

1. Open the Reporter main window and select **File** > **Configure** > **Report Packages**.
2. In the dialog that appears, select the package for the product you installed.
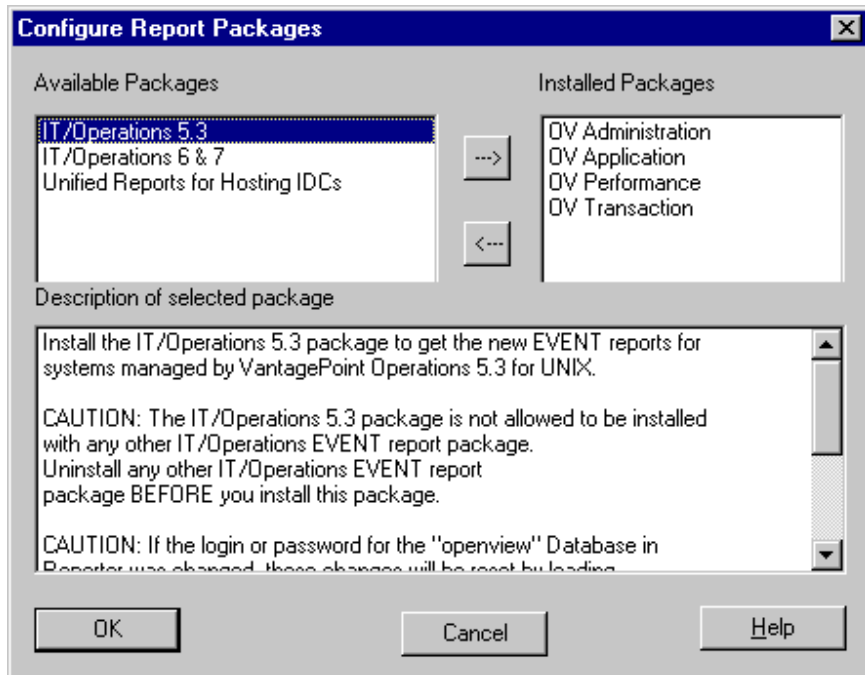
Please refer to the specifc OpenView product documentation for details on whether or not the package is automatically installed, or if you need to manually add it to Reporter (using the dialog shown above).

Displaying Reports as Web Pages

Reports can be viewed using any web browser such as Microsoft Internet Explorer^{TM} or Netscape Navigator^{TM}.

- To view the reports from the system which has Reporter installed, click the Show Reports button on the toolbar or right-click the name Reports in the hierarchy and select Show Reports.
- To view reports using a web server, create a link to the Reporter system URL; for example:
  http://<*reporter_server*>/HPOV_Reports/reports.htm
  NOTE: You must supply the fully qualified domain name of the Reporter system.
- For more information on this topic see the Advanced Topics chapter of this document.

Reporter Data

Accessing Reporter Data Reporter data is centralized on the system where Reporter is installed. You can configure more than one centralized area to gather data from the same OpenView agent.

Reporter Data Retention Reporter data is retained for seven days by default and most reports, by default, reflect this setting. You can change this setting by editing each metric list.You can generate reports, though, that have fewer days' data in them by editing the Report Definition and selecting the desired **Date Range**. In this way you can produce reports with differing amounts of data from the same database.

Where Reporter Data Resides All data accessed by Reporter is stored in a relational database that supports Open Database Connectivity (ODBC). Reporter includes a default database. The database does not have to reside on the Central Reporting System, as programs on that system can access it through ODBC calls. Check the Release Notes for a list of currently supported databases.

Creating Your Own Reports

As explained in the previous sections, Reporter includes a variety of reports that are automatically generated. You can make changes to a report definition, such as the date range you would like the data to cover, the shift on which you want data reported, etc. You may find that the default reports along with the changes you can make to their definitions are sufficient to meet your organization's needs. If this is the case, you can skip this section. But if you need to create your own custom reports, please read on.

New report templates are necessary when you want to report on data from new or modified metric lists. Custom reports allow you to display the data collected for these new lists. The procedures here cover how to change the report templates included with Reporter. If you change any default template, we recommend you retain the original report and template name and **save your changes under a new report and template name**.

Crystal Reports™ from Business Objects

To create or modify report templates, you need to first purchase Crystal Reports. We recommend the professional version, developed by Business Objects. For more information, go to the Web site at www.businessobjects.com.

Please refer to the Crystal Reports documentation for creating report templates. The steps below will supplement this information and describe how to integrate the templates into Reporter.

Step 1: Specify the database

When you create a new report template in Crystal Reports, you must designate the data source for the report's data. Even though more than one way may exist for accessing the Reporter database, we recommend you use the ODBC method. This method saves you from having to change your report templates if the underlying database for Reporter is moved.

- Open Crystal Reports.
- Select **New** from the File menu.
- In the Crystal report gallery, select **Standard**.
- Select the **SQL/ODBC** button.
- From the list of ODBC Server Types, select **ODBC-Reporter**.

Step 2: Choose a database table

After you have selected the database, you need to select the table that contains the data you want to use.

- Select the table(s) for your report.
  If you select more than one table, you are asked to indicate how they should be linked.

The following links are recommended:

(most).SYSTEMNAME links to	GROUPS.SYSTEMID
(most).SYSTEMNAME links to	SYSTEMS.SYSTEMID
GROUPS.SYSTEMID links to	SYSTEMS.SYSTEMID

The following tables are named according to the default metric list within Reporter. These tables store the data gathered for each metric list. If you add metric lists, additional tables are created from which you can generate other reports. Each database table is named according to the metric list name.

Database Table	Contains	Useful for
APPLICATION	One record (row) for each active user defined application on each system. Data points once each hour.	Reports showing which applications are most active.
GLOBAL	One record (row) for each hour showing overall activity, resources, etc.	General status of a system, resource usage, etc.
TRANSACTIONS	One record (row) for each Application Response Measurement (ARM) transaction once each hour.	Showing service level conformance and business transaction volumes.
DOWNTIME	One record (row) for each shift a system was available each day.	System Up and Down Time reports.

Additional tables control Reporter actions and maintain its status. These tables contain information generally not vital to your understanding of Reporter, but you might find some of the information interesting.

Database Table	Contains	Useful for
COMPLETED	One record (row) for each program which was completed in the operations of Reporter.	Tracking execution time of the various Reporter programs.
DATASOURCES	A list of the data sources and classes which were discovered on each OpenView agent.	Viewing the DataClasses and DataSources reports.
GROUPS	A list of systems belonging to each system group.	Viewing the names of the systems in a group.
SYSTEMS	A list of discovered systems.	See the DiscoveredSystems report. In addition, this table may be linked to other tables to provide additional information about a system.

Step 3: Choose fields (metrics) to use in the report

If you need more information on metrics before making your selections, refer to OpenView agent documentation for metric descriptions.

- Select the fields to include in the report.

Each table generated by a metric list contains the following fields:

Field	Description
ID	Sequentially assigned number for each row in the table. This is most likely not interesting for reports.
SYSTEMNAME	The name of the system that supplied the data. It matches the SystemID field in the Systems table.
DATETIME	The Date and Time at the beginning of the data sample period (for data points each hour, 8:00 represents the data from 8:00 to 8:59) This time is LOCAL TIME for the system which collected the data.

GMT	The Date and Time, specified in Greenwich Mean Time (also known as Universal Coordinated Time).
SHIFTNAME	A string with the name of the shift to which this record belongs, based on the first time covered by the record. This field may also contain "HOLIDAY" or "OFFSHIFT" to indicate holidays and times which are not covered in the user defined shifts.
"Instance Identifier"	For classes where more than one instance is logged in the same interval, the table includes the unique identifier for each instance, even if the identifier is not explicitly entered into the metric list. Examples of instance identifiers are APP_NAME for applications and TT_NAME for transactions. For classes where only one instance is logged each interval (GLOBAL, for example), no instance identifier is present.

Step 4: Design the template using Crystal Reports

Using Crystal Reports, you can explore the available formatting options. You will probably want to refer to the Crystal Reports documentation to take full advantage of Crystal Reports capabilities to design professional reports of presentation quality. Remember to preview reports with data from the Reporter database. After you are satisfied with the appearance of the report, move to the next section where you save the report template.

Creating Report templates that support Date Range

A report definition may specify a date range in order to select less data than is present in the database, based on dates. In order to support this feature, you must add a special selection in your report template.

- In Crystal Reports, choose Select Expert button.
- Highlight the field for DATETIME or whatever field contains the date.
- Adjust the Select Expert so that it chooses rows where the DATETIME "is" "in the period" "AllDatesToToday".

  When you use this report template, "AllDatesToToday" is replaced with the date range the user selects in the report definition.

Creating Report Templates that support Shifts

No special features are required in the report template to support user selected shifts. The only requirement is that the database table referenced by the report contains the field SHIFTNAME.

Step 5: Save the template

To save the template, preserve disk space, and ensure that the report is accessible to the Reporter system that generates the report:

- In Crystal Reports, on the File menu disable the Save Data with Report option.
- Save the template to a disk NOT mounted over the network.
  (Reporter generates reports using a background service that does not require you to remain logged onto the system. Network drives are usually mounted only when you log on.)

For easy access, save report templates in the same directory as the standard Reporter templates and use a similar naming scheme. The standard Reporter templates are saved to the directory <install_directory>\data\reports. The directory differs if you changed the default directory during installation.

Standard template names start with:

a_ for reports that apply to ALL systems
g_ for reports that apply to a GROUP of systems
s_ for reports that apply to SINGLE SYSTEMS

The Three Report Template Types

Reporter currently recognizes three different types of report templates: All (preceded by "a_"), Group (preceded by "g_"), and Single System (preceded by "s_"). The simplest template to create is ALL, which reports on all data in the database. If you do not want to complete the steps below, classify your report template as ALL, using an a_ prefix in the report template name. This classification provides you a simple means in the future for creating other templates. Later, you can modify the ALL template by adding selection criteria that allow it to be used as a Group or a Single System template as well.

Single System Reports: A Single System report template selects data from only one system from the database table. Exactly which system it selects is determined within the Reporter hierarchy when you select the system and drag the report onto that system. If you drag a report onto a single system, then that system is selected when the report is created. You can use the same template to create reports for different systems. Here are the steps required to construct a Single System report template.

- In Crystal Reports, choose Select Expert button.
- Highlight the field for SYSTEMNAME or SYSTEMID. (If these fields are unavailable then the report cannot be made into a single system report).
- Adjust the Select Expert so that it chooses rows where the SYSTEMNAME (or SYSTEMID) "is" "equal to" and then pick any available system name.

When you use this report template, the system name you have chosen is replaced with the one to which you drag the report. To follow a standard naming convention for your reports, save each Single System report template with an s_ prefix.

Group Reports: A Group report template selects only systems that belong to a selected group of systems. Creating a Group report template is the same as creating a Single System report template except that the System Name is present in all data tables but the group name is not. In order to allow group selection by group name, you must first link the "GROUPS" table to your data table. This can be done even after a report template is already working as an ALL template.

- Add the GROUPS table and link it into your reporting table (see *C*rystal Reports, 10.0, online Help: Database/Add Database to Report).
- Pick the table **GROUPS**.
- Adjust the links by dragging from the database table System Name field into the groups table SystemID field.
- If you see a link between the ID field in the two tables, remove it by selecting it and pressing the **Delete** button.

When you are finished, your linking diagram should look like this:
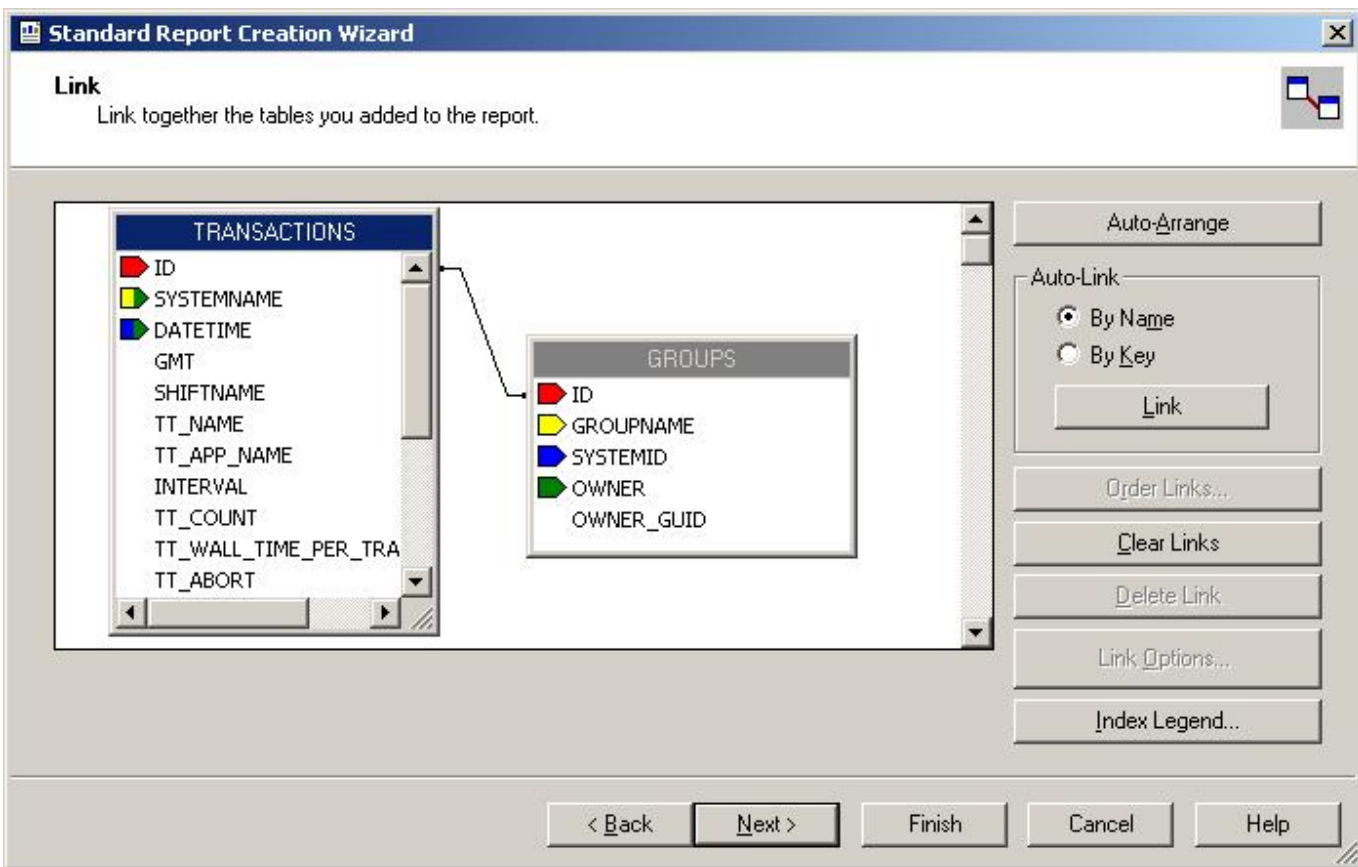
To finish:

- Go into the **Report/Select Expert** and choose the field **GROUPS.GROUPNAME** and "**is**" "**equal to**"
- Choose any valid group.
- Save the group report template with a g_ prefix.

Later in Reporter, when you drag this report onto a group, that group's name replaces the group name you chose above when you created the template.

Step 6: Add the Report Definition to Reporter

After you have created report templates, make them available to Reporter by right-clicking the Reports icon and selecting **Add Report**. In the dialog box that appears, fill in the fields (as explained below) and select **Add**, which displays the Add Report Definition dialog box.

The above illustration shows an example of a report definition that uses All, Group, and Single System report templates.

Report Name	The name that appears in Reporter's main window within the hierarchy under the Reports icon. Keep the name short to avoid scrolling to read it all. Make it descriptive of the report contents. You can use the same definition to produce reports for all systems (ALL), system groups (Group), individual systems (Single System), or any combination. Choose a report name that represents the entire family of templates.
Copy From	If you have already entered a report definition that is similar to the one you are now entering, you can save some typing by selecting that definition in the Copy From list.
Description	A label for the report that identifies it on the cover web page that lists reports. Enter a description that would help someone viewing the reports in a web browser understand the report contents.
Category	Allows sorting and grouping reports into categories. This field controls the layout of the web page links, by sorting all reports when the web page links are generated. A Category title is added each time a report in a new category is added. A special feature allows you to specify category sort sequences that are not strictly alphabetical. To use this feature, start a category with numbers and an optional space. These numbers are then used in determining the sort order but are not printed in the linked web page. You can also use numbers to specify the order of reports within each category. Generally categories are sorted as they are entered, but any leading numbers and spaces are removed prior to adding to the link web page.
All Template	The name of the Crystal Reports template for ALL type reports. (If you have not entered the special select criteria required for Group and System templates, then enter your report template name into this field. You can omit the leading *<install_directory>*\data if desired. Use the Browse... button to locate your template using the file Explorer-type window.

Group Template	The name of the Crystal Reports template for Group reports. (This report template must have a select criteria for **Group name is equal to …**). You can omit the leading *<install_directory>*\data if desired.
System Template	The name of the Crystal Reports template for Single System reports. (This report template must have a select criteria for **System Name is equal to … or SystemID is equal to …**). You can omit the leading *<install_directory>*\data if desired. (A report definition can use one, two, or all three types of templates.)
HTML Directory	The directory where the report web page is stored. If this directory does not exist when the reports are created, it is created at that time. By convention, the HTML directory should be a directory under *<install_directory>*\data\webpages directory. You can omit the leading *<install_directory>*\data if desired.
Date Range	To change a report to include data for a specific range of dates, select a date range from the the drop-down list.
Shift Name	To change a report to include data from specific time periods during a 24-hour day, select a shift from the drop-down list. See Customizing Shifts for more information on shifts.
Database Name	Most reports have their data summarized and placed in the Reporter database. In this case, this field should be blank. For reports which pull data from a database other than the Reporter database, select that database here. See Reporting From Other Databases for more information.
Max Run Time	After attempting to produce a report for 30 minutes, the report-generating program RepCrys moves on to the next report. If you have reports which take longer than the configured maximum minutes to generate, you should increase the number of minutes. You can also extend the time out for the RepCrys program itself. See Customizing the Reporter, timeouts for details.

Enhanced Report Generation for Group and System Reports

Reporter can now generate GROUP and SYSTEM reports automatically from **ALL** templates.

Assign the **ALL** template of a particular report to the desired report's GROUP and SYSTEM template as follows:

1. Open 'Reporter' main console.
2. Select the desired report by clicking the left pane of the Reporter's main console.
3. Click 'F5' to view the selected report's definition on the right pane of the main console.
4. Copy the template assigned to ALL reports to GROUP and SYSTEM reports.
5. Click 'Save Changes'.
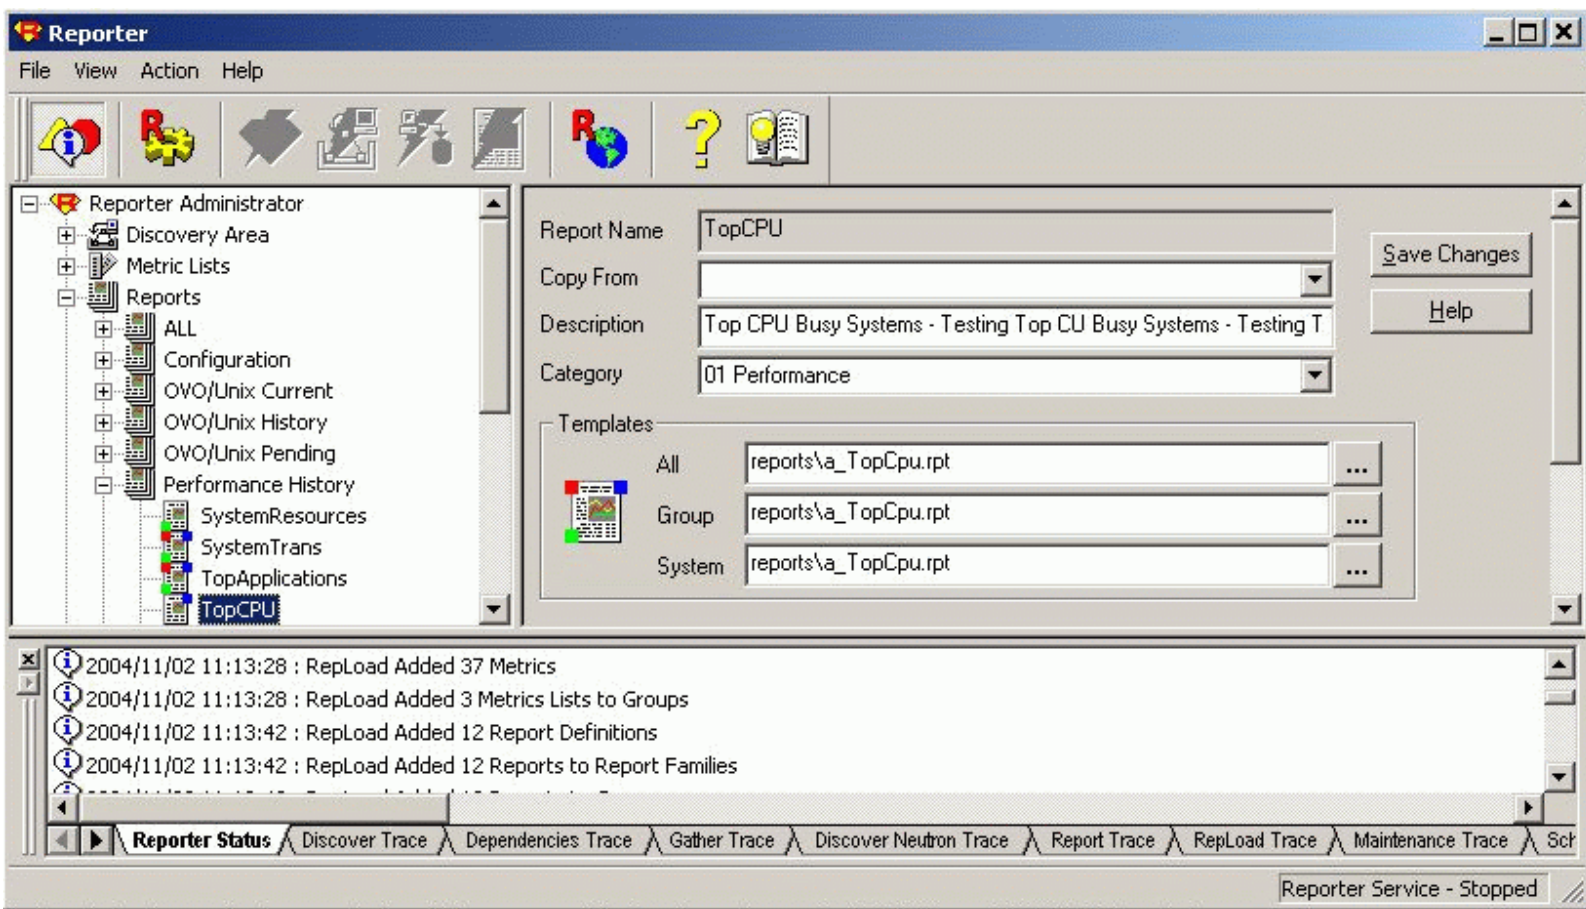
Run RepCrys to generate report. GROUP and SYSTEM reports is generated from the assigned ALL template.

📖 Chapter 5: Customizing Reporter

Reporter comes with a carefully chosen set of default actions and reports, based on customer input and market testing. In many cases these defaults will be exactly what you need and you can let the Reporter run with no changes at all. Reporter offers a high degree of flexibility in changing the default settings. This chapter describes areas you can change in the order you are most likely to use them.

- [Customizing Discovery](#)
- [Changing Times of Actions](#)
- [Creating New System Groups](#)
- [Adding Customer Report Groupings](#) (Customer Segmented Reporting)
- [Customizing What Data is Gathered from Systems](#) (Creating new metrics)
- [Timeouts](#)
- [Customizing Shifts](#)
- [Reporting From Other Databases](#)

Customizing Discovery

Discovery of Windows or UNIX systems running OpenView Performance Agent (also known as MeasureWare Agent) occurs as follows:

- (a) Reporter forms a list of system names.
- (b) Reporter probes each system to determine if it has an active OpenView agent.

You can modify the method of discovery in a number of ways, as described below.

NOTE: Customization of the discovery process for systems running only OpenView Operations (OVO) subagents is not available.

Initially, system discovery is limited to those agent-installed systems in the Reporter's local network domain (which is likely your local domain), which becomes the discovery list. You might want to broaden or restrict this list in order to report on more or fewer systems. To display the list, select **Discovery Area** in the left pane. The Discovery list, appears in the right pane, showing networks, domains, and single systems. When a network or domain is listed among those in the Discovery list (in the right pane), all

systems contained in that network or domain are searched.

When you expand each level of the Discovery Area, you see the potential **networks**, **domains**, and **systems**.
To add systems: In the left pane, drag and drop one or more systems from left to right pane to add to the Discovery Area list.
To delete systems: In the right pane, right-click the system and select Delete from the drop-down menu.

The networks, domains and systems found in the Discovery Area are derived from the Windows operating system and may not represent all the systems that are available to you. Reporter allows you to configure it to discover systems through four mechanisms:

1. Networks, domains and systems known to Windows networking are added to the discovery list using drag and drop.
2. Systems discovered in the OVO database are automatically searched for the presence of OpenView Performance Agent.
3. Single systems added to the discovery list, including those not known to Windows networking.
4. Lists created manually for discovering systems in an external source; such lists supply these system names/location in "bulk" for the discovery process.

In each case, discovery must be able to resolve the system name it is searching into a network address. You can determine if this is possible by running the Command Prompt on Windows and issuing a "ping" command for this system (ping system_name). If the ping command succeeds then the system_name is valid for discovery's use.

Targeting Single Systems Outside the Discovery Area (Add Single System)

Any system that you add can be outside the Windows network where the Reporter system resides if Reporter can still contact the system ('ping' the system). To add a system to the Discovery Area:

1. Right-click the Discovery Area and select **Add Single System**.
2. In the dialog that appears enter the system name.

When systems names are added using **Add Single System**, Reporter will issue a warning if no OpenView Performance Agent is found running on that system.

Targeting Multiple Systems Outside the Discovery Area (Bulk Discovery)

To discover multiple systems outside the Windows operating system, a method called **bulk discovery** is more efficient. This method works well for UNIX systems not identified within the Windows networking system, and so are unavailable to the Discovery Area. You could use the method described above, adding one UNIX system at a time, but this kind of repetitious exercise could be time consuming. With bulk discovery you create a list of potential system names in a text file. **See Reporter online Help** *"Bulk Discovery"* **for details.** Discovery looks for the existence of an OpenView Performance Agent on each system in the list.

You could create the list of systems by exporting all known system names from some other network topology scheme such as /etc/hosts on UNIX or the Network Node Manager software. Reporter substitutes this list of systems for the first phase of discovery.

NOTE: Unlike the Add Single System method, during bulk discovery Reporter generates no warning messages for systems on the list that go undiscovered because OpenView Performance Agent software is not found.

Changing Times of Actions

Initially, actions occur in the early morning hours (as described in How the Reporter Works). Actions are scheduled by default at these times to use otherwise idle time on your systems. In some cases, you might want to reschedule actions to occur at other times. For example, if many of the systems in your network are powered down each evening, these systems will not be available for discovery and data gathering if the Reporter runs at night. You might want to reschedule the discovery and gather programs to execute during the lunch hour.

You can change the time of any program that is already scheduled by using the Schedule icon in the left pane of the Reporter window. Select Schedule or one of the items under it to display scheduled programs in the right pane. Right click on the desired program and select Edit (or just double click the program you want). You can then edit the scheduling parameters for this one program execution.

To change the scheduling settings for an all program executions, click the plus sign next to the Schedule icon and right-click one of the programs below it. Select **Global Settings** to display a list of program settings. All programs (for example, all gather programs) are changed. The changes you make in this way apply to any programs for newly discovered

systems. For example, by changing the Job Time for Gather to 12:15 p.m., all the currently scheduled gather programs run during lunch. In addition, any new Gather programs (for newly discovered systems) are automatically scheduled for the new run time of 12:15 p.m.

You can change the following scheduling parameters using Global Settings:

Setting	Description
Job Date	Date the next program action is scheduled.
Job Time	Time the next program action is scheduled.
Job Reset	Frequency of which program repetition.
Run Minutes	Length of time program is allowed to run before it is terminated.
Max Concurrent	Maximum number of programs allowed to run at the same time.

NOTE: You may not see some settings for some programs because these settings either do not apply to the program or cannot be changed.

Creating New System Groups

Systems, when discovered, are automatically placed under the Discovered Systems icon. Systems can be arranged in groups to make them easier to locate and to facilitate reporting on groups of systems. A single system can reside in more than one group but all systems belong to the group called **ALL**.

You can assign systems to groups **manually**. Right-click the Discovered Systems icon and select Add Group. Select the group, then drag and drop systems from other groups into the right pane to add them to the new group. Organizing systems with similar characteristics into groups can be useful when performing actions such as gathering data and reporting on entire groups rather than on individual systems.

You can also set up systems to be grouped **automatically** when they are discovered. By selecting AutoGroups from the Configure submenu on the File menu, you see a dialog box you use to enter the parameters for automatic grouping. For a detailed description of this feature see Chapter 7: "Advanced Topics," Automatic System Grouping. By default Reporter automatically places each discovered system in a group based on its operating

system (HP-UX, Windows systems, etc.) If OpenView Operations (OVO) systems are included, all systems discovered in the OVO database are placed in the group **ITO**.

You can enable the use of **special groups**. Select Options from the Configure submenu on the File menu of the Reporter main window.

- As previously mentioned, all discovered systems will belong to the special group **ALL**. Deleting a system from any other group merely removes it from that group. Deleting a system from the ALL group will remove the system and all data from that system from the Reporter database.
- In this dialog box you can direct Reporter to automatically create the group **NEW** and add each newly discovered system to it. This feature is helpful if you are manually grouping systems. After assigning the systems in the NEW group to the groups you desire, delete the NEW group. When the NEW group reappears, you know you have new systems to assign.
- If you are also discovering OVO systems, you can enable the automatic creation of **OVO NodeGroups**. These group names are copied from the OVO database. The OVO node groups (and the systems within them) automatically appear under the Reporter Discovered Systems hierarchy.

Adding Customer Report Groupings

The Customers feature allows you to group systems according to Customer. Then, when reports are generated, each customer has a set of reports generated, using data from only those systems owned by that customer. Customer report sets are stored in *<install_dir>*\data\webpages\customers*<customer_name>*.
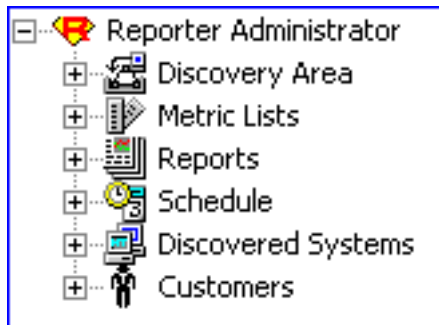
The customer reports feature is available within the left pane of the Reporter main window. There you see the top-level Reporter Administration folder, under which all other folders appear. Among them you see a Customer folder, which by right-clicking, allows you to create or delete new or existing customers.

After you have created a customer, a directory (folder) is created that will contain the reports for that customer.

NOTE: If a customer name includes a blank, an underscore character will be substituted for the blank in the directory name. For example, if a customer is entered as Wells Fargo, the directory is created as \wells_fargo\. Each set of customer reports is accessible from the main Reporter Web page.

To create a customer

1. In the Reporter window left pane, right-click **Customers** and select **Add Customer**.
2. Enter the customer name and password in the dialog box.

3. Select **Save**.
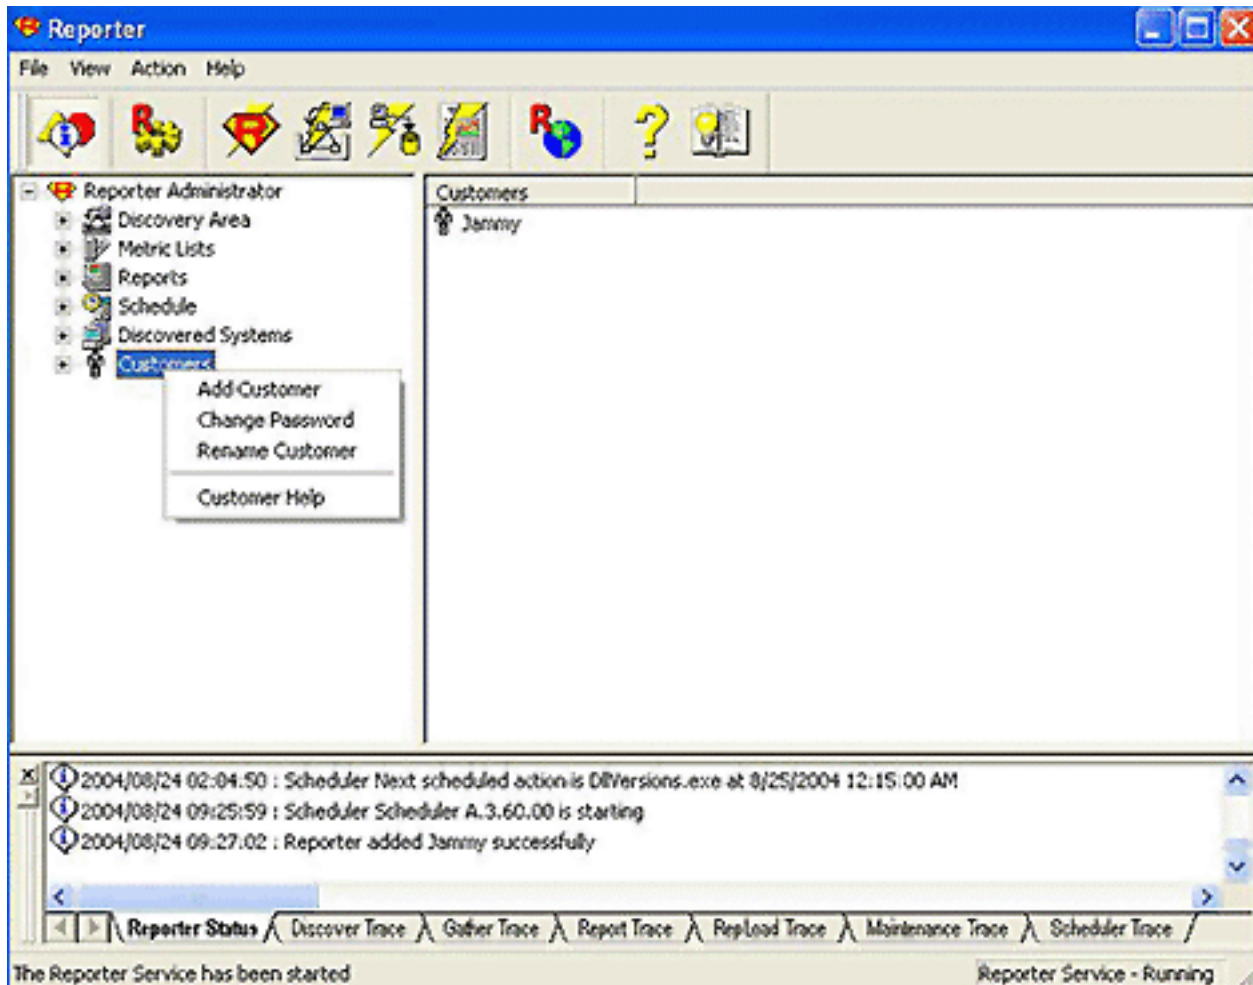
To rename a customer

1. Expand the **Customers** node to display all customers and select the customer.
2. Right-click the customer and select **Rename Customer**.
3. Enter a new customer name. All settings and systems associated with the customer remain the same when you rename the customer.

To delete a customer

1. Expand the **Customers** node to display all customers and select the customer.
   (All the customers are displayed in the right pane.)
2. Select the customer and press the **Delete** key. All settings and systems associated with the customer is removed.

To assign reports to a customer:

1. Double-click Customers to expand the tree in left pane and see the list of customer names. Select a customer name.
2. Double-click **Reports**. In the left pane, double-click the group containing the report you want to assign. Reporter expands the list to show reports in that group. Select the report or reports you

want to assign and drag them to the right pane to add them for the selected customer.

To delete a report from a customer:

1. Expand the **reports** node underneath the customer in the right pane.
2. Right-click the report and select **Delete**.
   Or, having selected the system, press the **Delete** key.

Reports are generated according to customer and group associations. In other words, group reports are generated for each customer but show only those systems owned by that customer.
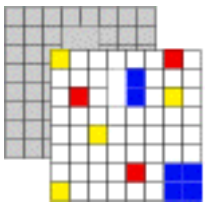
Customizing What Data Is Gathered From Systems

Reporter includes pre-defined metric lists that determine what data to gather from each system. Reporter also averages the data according to the summarization level (for example, hourly averages [the default], averages of 5, 10, 15, or 30 minutes, 3 hours, 6 hours, etc.).

Add metric lists: You can add your own metric lists by using the Metric List icon in the left pane. Reporter allows you to create metric lists for data from OpenView Performance Agents and associate them with systems or groups of systems from which data will be gathered.

NOTE: Metric names cannot contain embedded spaces.

Metric lists for the OVO systems are fixed and cannot be altered (although you can alter the number of days to retain data for the OVO metric lists, you should not add or delete metrics to these lists).

One obvious reason for creating a new metric list is to direct Reporter as to what data should be gathered from a new data source. If you want to access data from OpenView Performance Agent log files that originated from Data Source Integration (DSI) technology, then you can create a metric list to gather this DSI data and add it to the Reporter database.

Edit metric lists: You can edit existing metric lists (to gather additional metrics for example). You can create more than one metric list against the same OpenView Performance Agent data source if desired. For example, you can create one list to collect GLOBAL data at one point every hour and retain 7 days of data. Then you could create a second metric list to collect GLOBAL data at one data point every day and retain 365 days of data. Be careful in deleting metrics from existing metric lists because deleting metrics can cause some of the default reports to fail.

NOTE: Gathering additional data into the Reporter database does not cause this data to show up in one of the reports. To generate reports based on custom metric lists, you must purchase Crystal Designer. With Crystal Designer you can design report templates which display the data you desire. After creating the report template, you can add a report definition (in Reporter) that uses the new data from the custom metric list and uses the template to format the data into a web-based report.

You can assign metric lists to different systems or groups of systems by dragging the list from the left pane over into the right pane which is displaying the desired group or system information.

Timeouts

As Reporter completes its actions through the Master Scheduler that runs each program, it makes sure that any program that fails to execute properly does not prevent other programs from executing. Reporter accomplishes this by starting a timer when any program (such as gather) runs. If the timer runs out before the program completes its actions, Reporter assumes that the program has failed to properly execute and terminates it. Any time this situation occurs, you can verify its occurrence by looking at the status column on the schedule and on several of the default reports.

In some cases, the work that a program has to perform takes longer than expected. Gathering large volumes of data over a very slow network, for example, might cause the gather program to terminate before it has finished collecting data from a system. In such a case, you should increase the value of the Run Time parameter to avoid terminating the program before it can finish its work.

You can change timeouts on individual programs by selecting the Schedule icon in the left pane and double clicking the desired program in the right pane. To change the timeouts for all programs, select the program under the Schedule icon in the left pane, right-click it and select **Global Settings**. Check the box next to Run Minutes and enter the desired value. This change affects all the programs run for every system. The Run

Minutes value setting also applies to all new programs added for newly discovered systems.

Customizing Shifts

Most reports produced by the Reporter show data from the entire day for every day that has been gathered into the database. You can specify reports to show data for less than a 24-hour day by specifying a **shift** in the report definition. You can modify an existing report or copy a report definition to create a new one and select the shift you desire in the Shift Name box (in the Report Definition). Making such a change means that the report contains data only from the selected shift.

You can change shift times or define new shifts by selecting **Shifts** from the Configure submenu on the File menu. Default shift names are already present within the Configure Shifts dialog box and include **Prime**, **Swing**, and **Graveyard**. You can modify these shift times and also add your own shifts. You can also configure Holidays. Within the Report definition, however, if the data you are gathering for a particular report contains a date and time combination that is not covered by any defined shift or specified holiday, you can just specify **OFFSHIFT**.

You can define one set of shifts that applies to all systems in the database. Shifts should not overlap with each other since each data record must be uniquely identified with a single shift name. If shifts and holidays match the same day, the entire day will be marked Holiday.
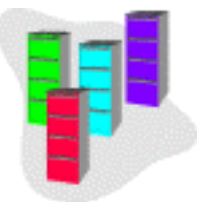
Reporting From Other Databases

You can configure Reporter to generate reports directory from another database, such the one used by HP OpenView Operations (also known as IT/Operations). Those reports generated from the OpenView Operations Oracle database show information not easily summarized by Reporter and that require templates different from the default set that Reporter generally uses.

To customize Reporter so that it can generate reports directly from another database

requires:

- The database be accessible to the Reporter system. Reporter programs run on the Windows platform and must be able to access the database from there. This may require loading additional database client software such as SQL Server Client or [SQL*Net](#).
- Report templates be created that reference the database directly (requires the use of the separately purchased Business Objects Crystal Reporter Writer).
- The database security (user ID and password) be entered into Reporter. (In the Reporter main window,
  select **File** > **Configure** > **Databases**.)
- The report definitions in Reporter include the database name.

📖 Chapter 6: Scalability and System Requirements

ℹ️ NOTE: The examples provided for Scalability in this chapter is legacy and not tested with Reporter 3.6.

This section addresses how to optimally configure a Reporter installation, based on particular environments and reporting needs. Information presented here can help in estimating the number of OpenView agents that a single Reporter system can successfully support. This information also assists in planning on how much disk space to allow for Reporter data, given a specific number of days and systems.

ℹ️ NOTE: When gathering data from an OpenView Operations (OVO) database, the number of systems does not impact Reporter's schedule as much as when gathering data from individual systems because OVO has already gathered the data.

This section covers system requirements, scheduling recommendations, and disk space use for both default and configured Oracle databases. Recommendations offered here are based on tests that used Reporter's default metric lists:

GLOBAL	10 metrics summarized into one data point every hour
APPLICATION	7 metrics summarized into one data point for each active application every hour
TRANSACTIONS	9 metrics summarized into one data point for each active transaction every hour

The following minimum hardware configuration is recommended:

CPU	500 MHz Pentium III
Main Memory	128 MB as a minimum
Disk Space	130 MB for product files and up to 1 GB or more for data

The above recommendations are affected by many different factors, which can directly impact the performance of the Reporter installation. These factors are listed below with descriptions following:

- [Number of agents](#) Reporter must query for data
- [Time Allotment for Reporter tasks](#)
- [CPU speed and main memory size of the Reporter system](#)
- [Network speed](#)
- [Network size of the networks used by the Discovery process](#)
- [Reporter Database location](#)
- [Gather concurrency settings](#)
- [Number of defined groups, customers, and reports](#)
- [Database disk space used (for default--MSDE Service Pack 3a)](#)
- [Database disk space used (Oracle)](#)
- [Database disk space used (Access)](#)
- [Number of other other HP OpenView products utilizing Reporter](#)
- [Database disk space used (number of days' worth of stored data)](#)
- [System Requirements to view Online help, End User documents and training videos](#)

Each of these factors is discussed in detail in this section. In addition to these factors, a short section on [Troubleshooting](#) is included.

Number of Agents

Reporter must contact each OpenView agent, collect the appropriate data, store it in the database, and later access the data to create reports. These processes take significantly longer for ten thousand agents than for just 100 agents.

Example: Using a Reporter system with the minimum recommended hardware described above and assuming a 100MB network with a local database and default configuration*, Reporter should be able to gather and produce reports as follows:

- for 6,200 systems in approximately 6 ½ hours
- for 11,500 systems in approximately 12 hours.**

*A default configuration is one where there are no customers defined, a minimal number of groups (3 or fewer) and only the standard, already defined reports.

**In all these tests, the reporting phase produced 52 reports. Data for more than 1000 systems was extrapolated from collected data. All tests were conducted on Windows 2000, Service Pack 2.

Time Allotment for Reporter to Complete Report-Generating Tasks

You may need to determine whether you can use Reporter as is or if you will need to reconfigure its schedule to fit in with your organization's available resources and daily work

routine. For the sake of this discussion, we are characterizing the running of Reporter into two levels of basic usage patterns. Report likely will dominate system use during these times.

Level 1, using MSDE Service Pack 3a: Using MSDE Service Pack 3a, Reporter's default database, requires no change to Reporter's late-night processing schedule and is directed at systems used for other tasks across the organization. In the default schedule Reporter starts its routine at 12:30 AM each night. This schedule allows at least 6 1/2 hours for Reporter to finish so that the processing system is idle by 7:00 AM each morning.

| Hardware	Systems supported / Database space required
• 500 MHz Pentium II, single processor • 128 MB memory	• 6,580 OpenView Performance Agents (or MeasureWare Agents)§ • 930 MB for database (approximate)
• 733 MHz Pentium III, dual processors • 692 MB memory	• 8100 OpenView Performance Agent (or MeasureWare Agents) • 1140 MB for database (approximate)
• 1 GHz MHz Pentium III, single processor • 512 MB memory • Gather Concurrency: 20	• 12,272 OpenView Performance Agents (MeasureWare Agents) • 1727 MB for database (approximate)

Level 2, using MSDE Service Pack 3a: Allow more time for Reporter to gather its data and generate reports. If, for example, you don't require updated reports until 12:30 PM each day, and Reporter is run on a dedicated system, then you can allow 12 hours for the reporter processing cycle. In this case the number of systems you can monitor will increase.

Hardware	Systems supported/Database space required
• 500 MHz Pentium II, single processor • 128 MB memory	• 12,400 OpenView Performance Agents (or MeasureWare Agents) • 1740 MB for database (approximate)
• 733 MHz Pentium-III, dual processors • 692 MB memory	• 15,400 OpenView Performance Agents (or MeasureWare Agents) • 2170 MB for database (approximate)
• 1 GHz Pentium III, single processor§ • 512 MB memory§ • Gather Concurrency: 20	• 23,010 (or OpenView Performance Agents (or MeasureWare Agents) • 3,400 MB for database (approximate)

The graphs below illustrate the results of tests that determined the typical number of hours required to complete Reporter tasks for a given number of systems. The Reporter phases that consume the most resources are discovering systems, gathering data, and producing reports.

You can dramatically affect these times by changing the report collection parameters such as the number of days in the database (7 days is assumed for these graphs) or the collection concurrency, as illustrated in the 1 GHz tests where it is set to 20.

500 MHz Pentium II (Local MSDE DB)

Dual 733 MHz Pentium III (Local MSDE DB)

Legend: Report, Discovery, Gather
Y-axis: Hours
X-axis: Number of Systems (Thousands)
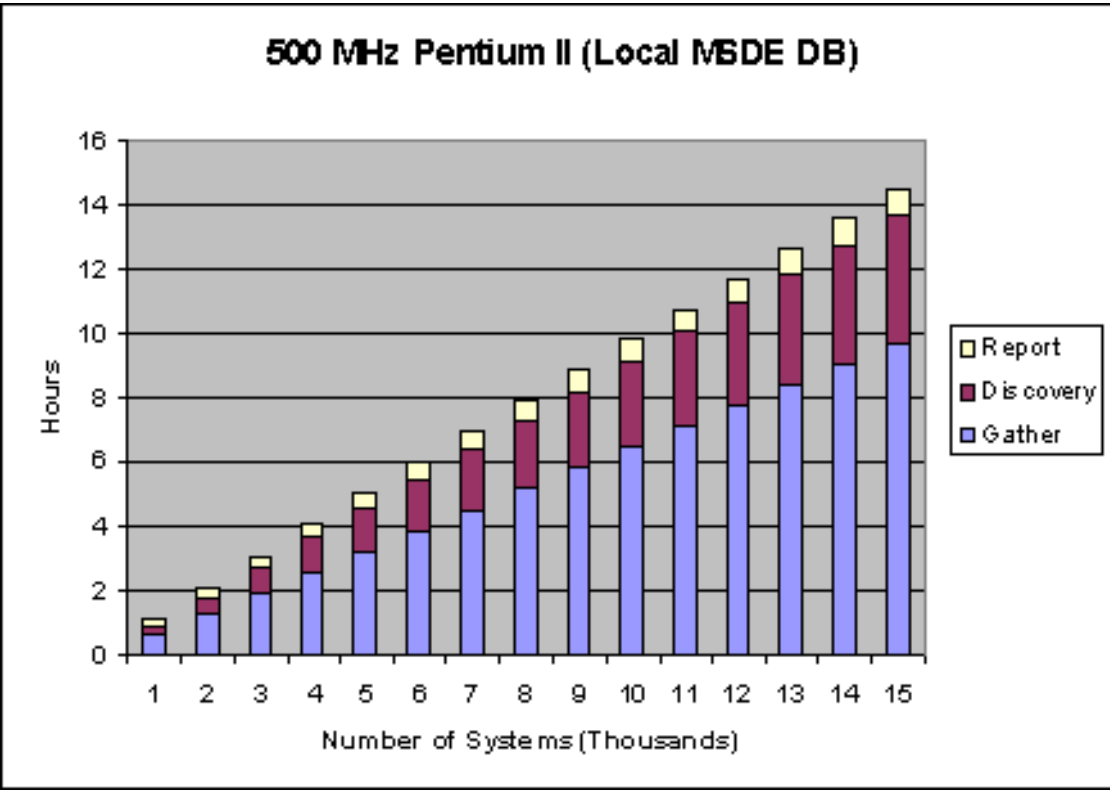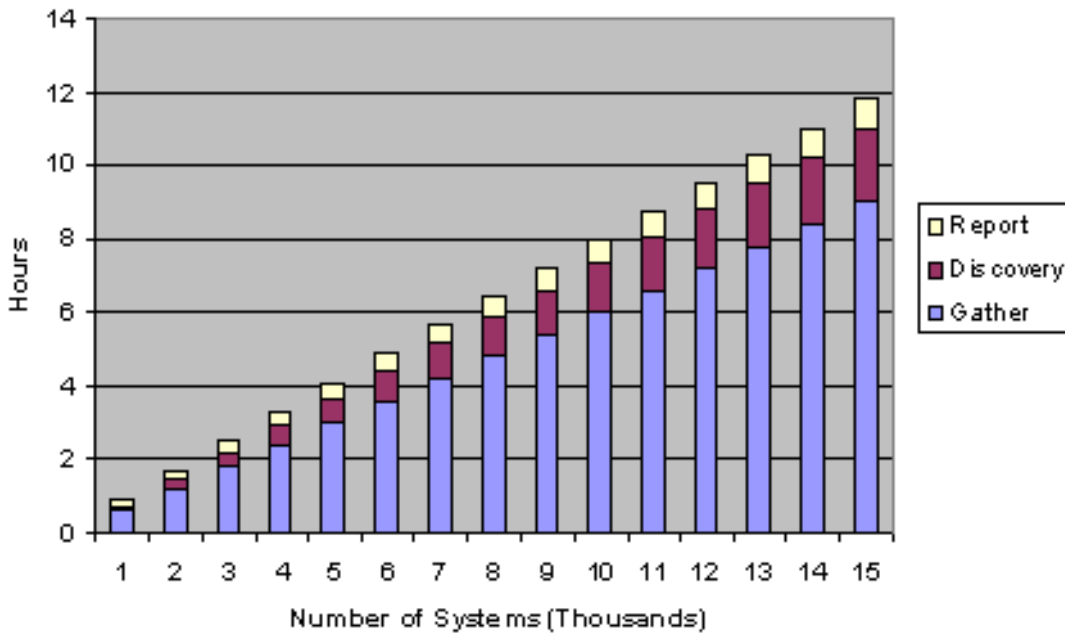
1 GHz Pentium III (Local MSDE DB)
Gather Concurrency 20

Legend: Report, Discovery, Gather
Y-axis: Hours
X-axis: Number of Systems (Thousands)
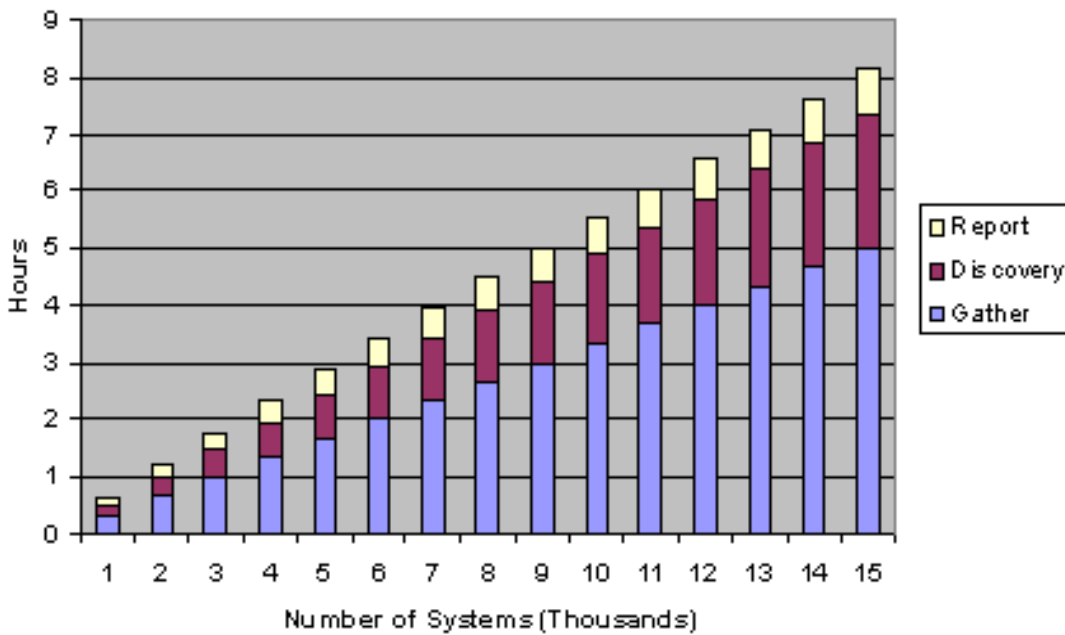
Level 1, Using Oracle: Requires no change to Reporter's default processing schedule, aimed at non-dedicated systems used for other tasks. For systems like these, Reporter's late-night processing does not interfere with the daily routine. Reporter uses the default nightly schedule, which starts at 12:30 AM. This level recognizes that Reporter must finish its work in 6 1/2 hours so that the system is idle by 7:00 AM.

Server Hardware	Client Hardware	Systems supported / Database space required
- Model UltraSparc 2 E3500 - 4 Way 400 MHz CPU - 2 GB Memory - 100 Mbps LAN	- 500 MHz Pentium II, single processor - 128 MB memory - 100 Mbps LAN	- 4150 OpenView Performance Agents (or MeasureWare Agents) - 1010 MB Minimum for database (approximate)
- Model UltraSparc 2 E3500 - 4 Way 400 MHz CPU - 2 GB Memory - 100 Mbps LAN	- 1 GHz Pentium-III, single processor - 512 MB memory - 100 Mbps LAN - Gather Concurrency: 20	- 7500 OpenView Performance Agents(or MeasureWare Agents) - 1840 MB Minimum for database (approximate)

500 MHz Pentium II (Oracle DB)
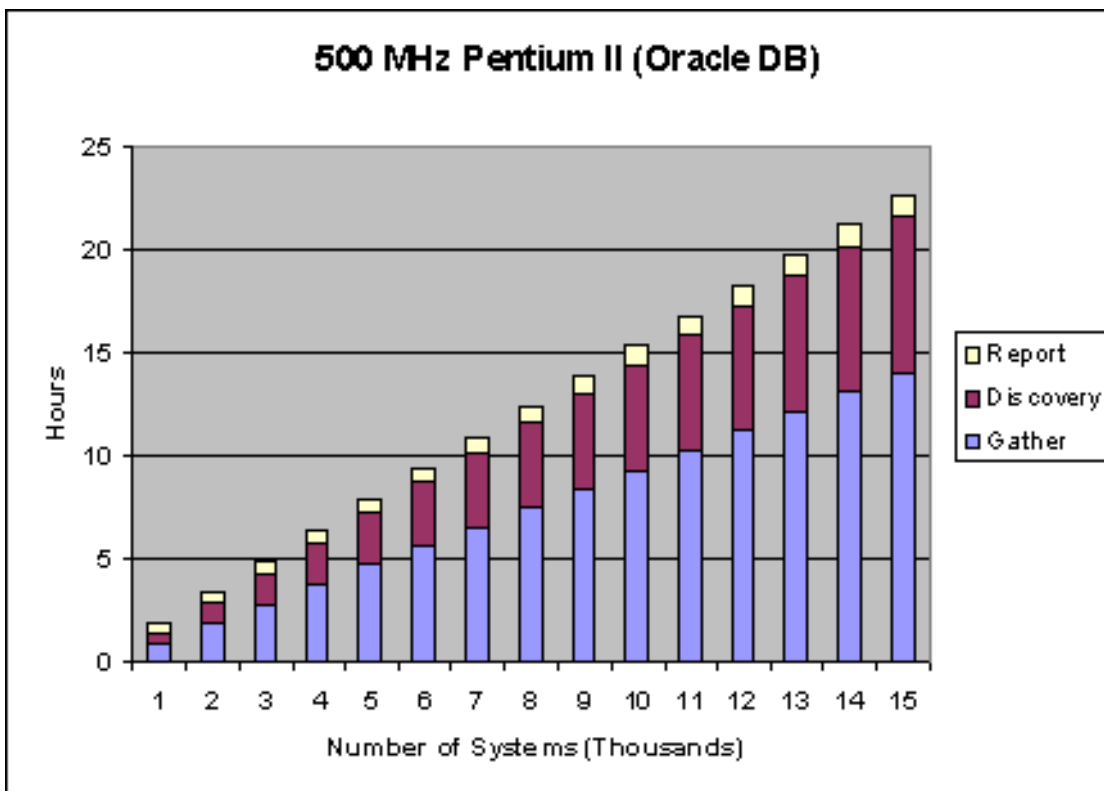
1 GHz Pentium III (Oracle DB)
Gather Concurrency 20
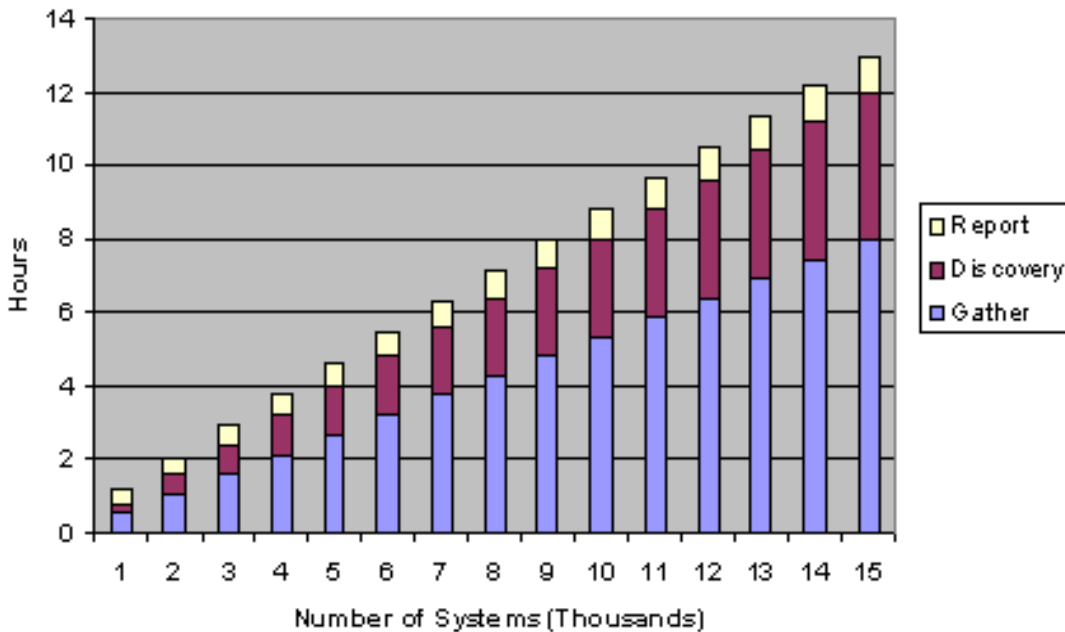
Level 2, using Oracle as the Reporter database: Allow more time for Reporter to gather its data and generate reports. If, for example, you don't require updated reports until 12:30 PM each day, and the Reporter is run on a dedicated system, then you can allow 12 hours for the reporter processing cycle. In this case the number of systems you can monitor will increase.

Server Hardware	Client Hardware	Systems supported / Database space required
• Model UltraSparc 2 E3500 • 4 Way 400 MHz CPU • 2 GB Memory • 100 Mbps LAN	• 500 MHz Pentium II, single processor • 128 MB memory • 100 MB Lan	• 7800 MeasureWare Agents (or VP Performance Agents) • 1900 MB Minimum for database (approximate)
• Model UltraSparc 2 E3500 • 4 Way 400 MHz CPU • 2 GB Memory 100 MB LAN	• 1 GHz Pentium-III, single processor • 512 MB memory • 100 Mbps LAN • Gather Concurrency 20	• 14375 OpenView Performance Agents (or MeasureWare Agents) • 3530 MB Minimum for database (approximate)

Level 1, using Access: Requires no change to Reporter's default processing schedule and is directed at those systems that are used for other tasks across the organization. For systems like these, Reporter's late-night processing cycle means that it does not interfere with the daily routine. In this situation, Reporter uses the default schedule, which starts up its routine at 12:30 AM each night. This level recognizes that Reporter must finish its work in 6 1/2 hours so that the system is idle by 7:00 AM each morning.

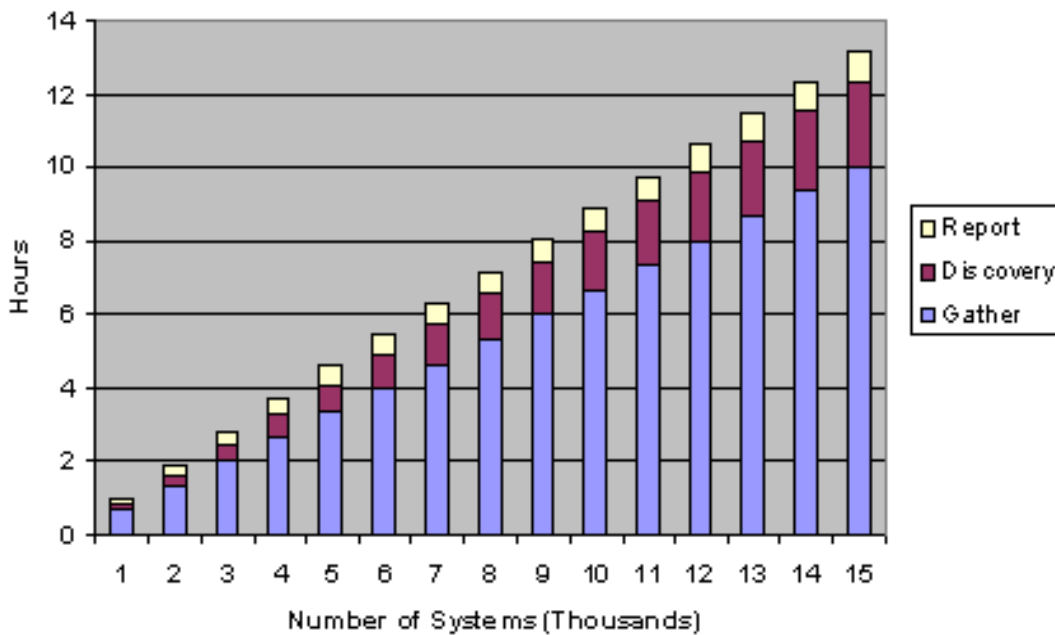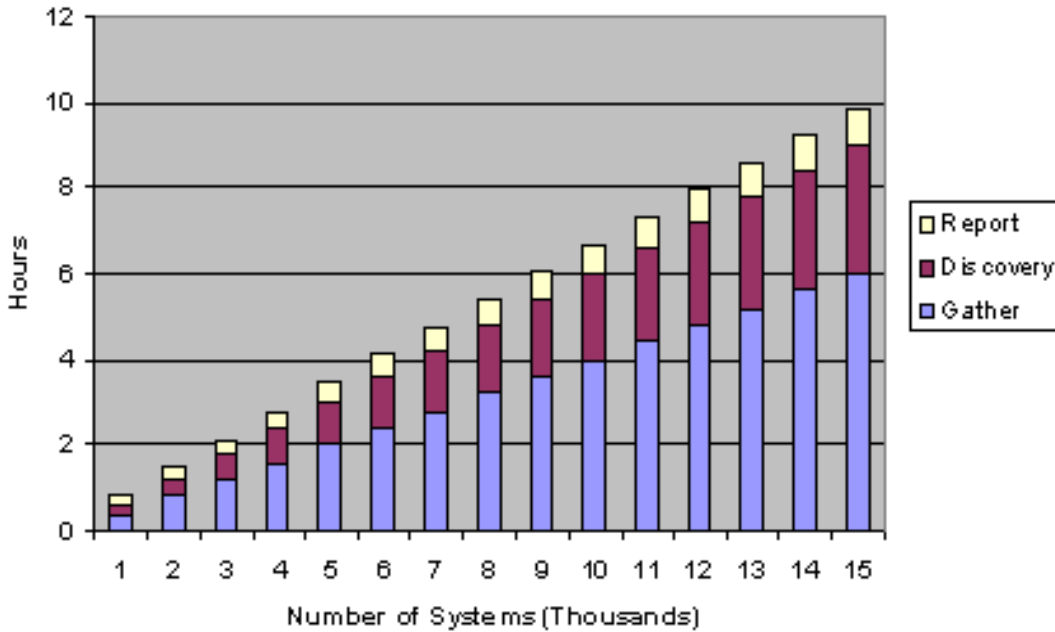| Hardware	Systems supported / Database space required
• 500 MHz Pentium-II processor • 128 MB memory • 100 Mbps LAN	• 5570 VP Performance Agents (or MeasureWare Agents) • 440 MB minimum for database (approximate)
• 733 MHz Pentium-III, dual processor • 692 MB memory • 100 Mbps LAN	• 7300 OpenView Performance Agent s (or MeasureWare Agents) • 572 MB for database (approximate)
• 1 GHz Pentium III single processor • 512 MB Memory • 100 Mbps LAN • Gather Concurrency: 20	• 10000 OpenView Performance Agents (or MeasureWare Agents) • 777 MB for database (approximate)
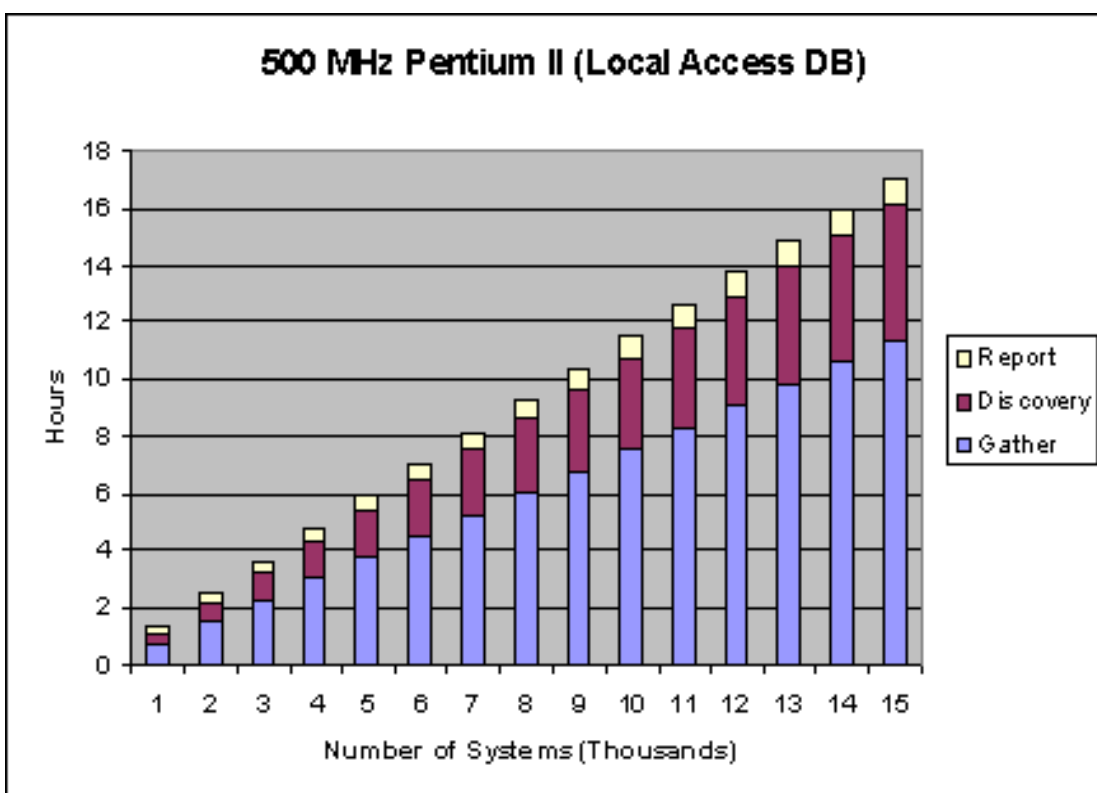
Level 2, using Access: Allow more time for Reporter to gather its data and generate reports. If, for example, you don't require updated reports until 12:30 PM each day, and the Reporter is run on a dedicated system, then you can allow 12 hours for the reporter processing cycle. In this case the number of systems you can monitor will increase.

Hardware	Systems supported / Database space required
• 500 MHz Pentium-II processor • 128 MB memory • 100 Mbps LAN	• 10500 OpenView Performance Agents (or MeasureWare Agents) • 830 MB minimum for database (approximate)

• 733 MHz Pentium-III, dual processor • 692 MB memory • 100 MB LAN	• 13750 MeasureWare Agents (or VP Performance Agents) • 1080 MB for database (approximate)
• 1 GHz Pentium III single processor • 512 MB Memory • 100 Mbps LAN • Gather Concurrency: 20	• 19166 OpenView Performance Agents (or MeasureWare Agents) • 1490 MB for database (approximate)

Dual 733 MHz Pentium III (Local Access DB)

1 GHz Pentium III (Local Access DB)
Gather Concurrency 20

Chart showing Hours vs Number of Systems (Thousands), with stacked bars for Report, Discovery, and Gather.

CPU Speed, Memory Size, and Gather Settings

The CPU speed and memory size of the Reporter system affects processing time. Although the minimum recommendations offer acceptable processing speed, you can gain significant improvements in performance by increasing memory and CPU speed on the Reporter system.

The following are necessary in your environment:

1. Increased CPU and memory to allow a higher number of concurrent gathers. If you are gathering from many thousands of agents, this can be a significant savings in time. Gather concurrency is discussed in detail below.
2. Increased CPU speeds up the reporting cycle. The report engine is CPU intensive so the more the better here. For example, the figures below show how memory and CPU affect the required processing time for the "TopNetwork" report for all systems:

500 MHz with 128 MB memory	32 seconds
1 GHz with 512 MB memory	12 seconds

If your Reporter installation regularly produces hundreds of reports, increasing CPU can dramatically improve the reporting phase performance.

Network Speed

Network speed factors into processing time. All Reporter tests were completed over a 100 Mbps

LAN. Because Reporter data must travel over a network to the Reporter system, a slower network lengthens the time required for discovery and gather. In addition, a remote database impacts the speed of the reporting phase. Remote databases are discussed below.

Network Size

The number of systems/agents on the networks which Reporter is discovering affects processing time. When Reporter discovers systems through "network browsing" every system on that network is contacted by Reporter and queried for a performance agent. The more systems, the longer the discovery cycle will run. Also, the A.03.50 version of Reporter is now attempting to discover OpenView Operations subagents as well as OpenView Performance Agents. Therefore, the discovery cycle of Reporter A.03 will run approximately twice as long as the A.02 version of Reporter against networks of the same size.

Reporter Database Location

If the Reporter database exists on the same system, less time is required for data retrieval. A non-local database requires more time for all the phases of Reporter. For example, the following table compares a Reporter installation with an MSDE Service Pack 3a local database against the same installation with a non-local Oracle database:

Reporter Phase	Remote Oracle Database	Local MSDE Database
Discovery	12	8
Gather	40	28
Reporting	20	10
Total	72	46

In this case, a non-local database increased the entire Reporter cycle by about 60% overall but with a 100% increase in the reporting phase. For this test, Reporter was installed on a 1GHz PIII with 512 MB of memory and configured with two groups (no defined customer reports), reporting against 750 systems. When the Gather concurrency was configured to 8 (default), each reporting cycle produced 49 reports.

Gather Concurrency

The number of concurrent Gathers affects how soon Reporter can move on on its subsequent tasks of storing the data and generating reports. The default concurrent number of gathers is 8, based on the minimum configuration above. However, you can get substantial performance improvements on faster systems by increasing the number of concurrent gathers.

For example, on a 1GHz/ 512MB system gathering from 750 agents, the following settings show

how improvement occurred when the number for concurrent gathers was increased:

- 8 Concurrent Gathers --- took 28 minutes
- 10 Concurrent Gathers --- took 22 minutes
- 20 Concurrent Gathers --- took 14 minutes

Number of Defined Groups and Customers

The number of system groups and defined customers also affects how long the report processing phase must run. For each group defined and linked with reports that have group templates, all the group reports will be generated. Also, for each customer defined, many of the reports, as well as any appropriate group reports (if the customer systems are also grouped) will be generated. By defining many groups and customers, you may be increasing the number of reports significantly, thus the reporting phase can run much longer. This table shows an example of how increasing the customers and groups defined can increase the number of reports produced.

Customers	Groups	Reports
0	0	21
1	0	33
2	0	45
0	3	52
0	4	64
0	5	76
2	1	78
10	1	262

Because report generation consumes much CPU and memory, reports are generated one at a time. On average, 200 reports take twice as long to generate as 100. Also report generation varies according to how much data is accessed per report. A report that accesses global data for every defined system would take longer to generate than a report on global data for only a selected number of systems, based on customer or group membership.

Database Disk Space Use (default)

The following disk space values assume the default metrics are collected and held for 7days. If you collect from more systems or hold more days' worth of data, the database will require more disk space and processing will take more time. Initial testing indicates that the default database, MSDE Service Pack 3, requires at least approximately 20.1 KB disk space per system per day. To store data for 1,000 systems for the default database for 7 days would require 141 MB (1,000 systems multiplied by 7 days multiplied by 20.1 KB).
The following chart shows the database disk space requirements according to the number of systems and the number of days' worth of data stored in the database.

MSDE Disk Space
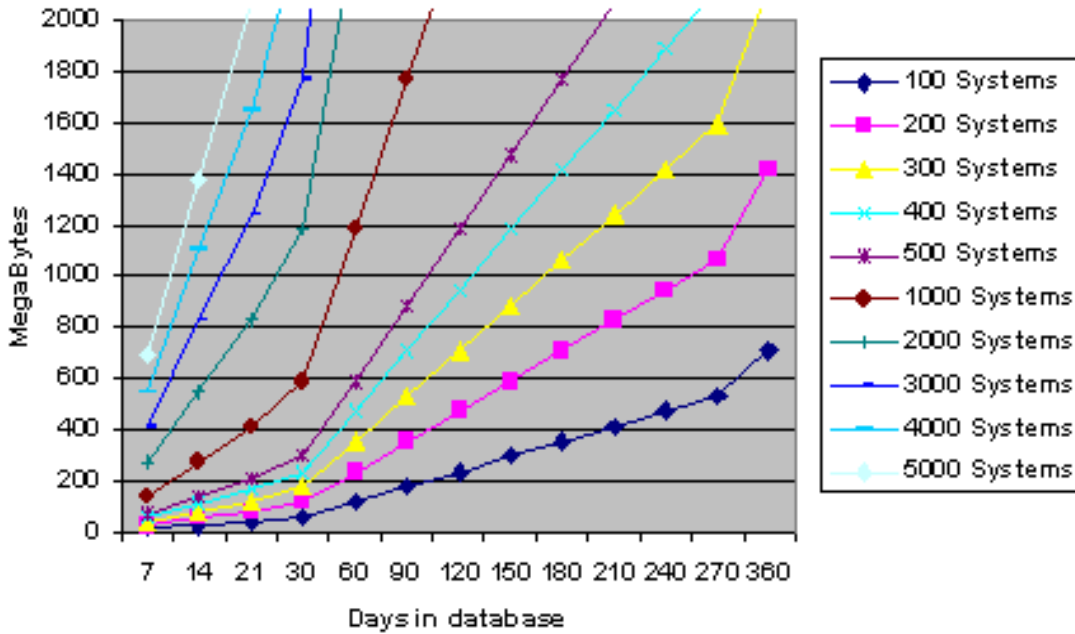
Database Disk Space Used (Oracle)

Initial testing indicates that the Oracle 8.1.7 database requires approximately 35KB per system per day. To store data for 3,000 systems for 7 days (default setting) would require 693 MB (3,000 multiplied by 7 days multiplied by 35KB).

The following chart shows the database disk space requirements according to the number of systems and the number of days' worth of data stored in the database.

Oracle Disk Space

Legend:
- 100 Systems
- 200 Systems
- 300 Systems
- 400 Systems
- 500 Systems
- 1000 Systems
- 2000 Systems
- 3000 Systems
- 4000 Systems
- 5000 Systems

Y-axis: GigaBytes (0, 10, 20, 30, 40, 50)
X-axis: Days in database (7, 14, 21, 30, 60, 90, 120, 150, 180, 210, 240, 270, 360)
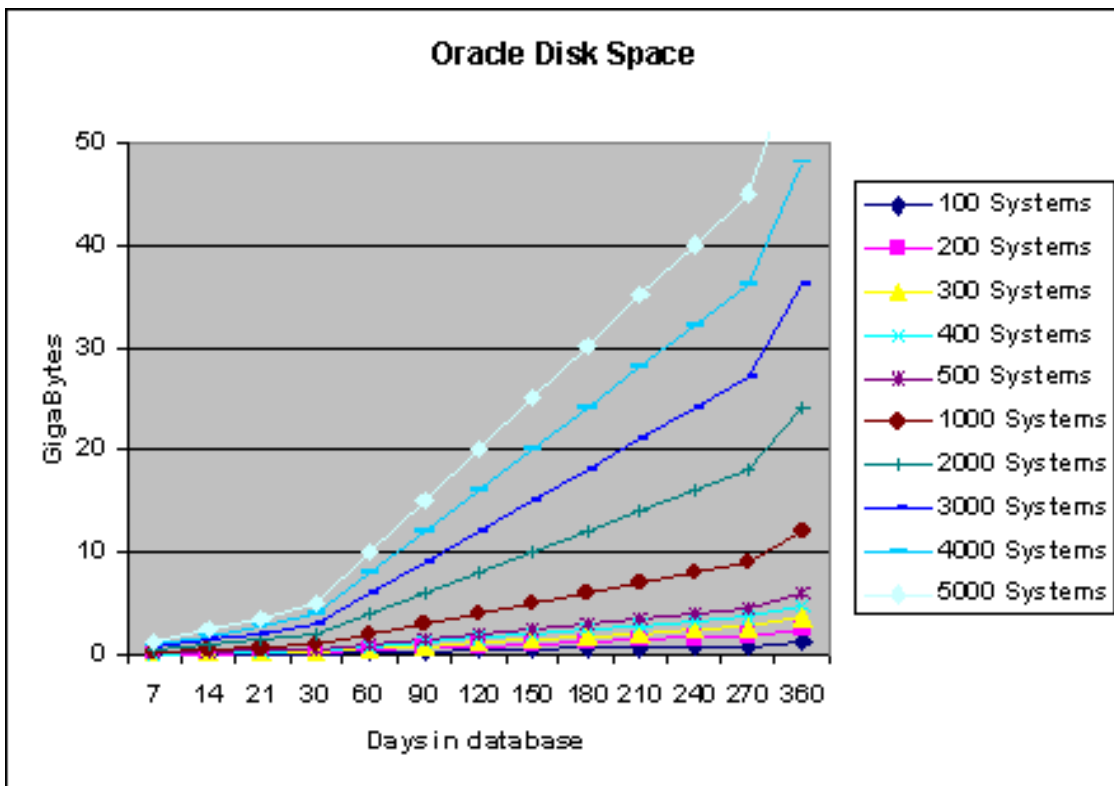
Database Disk Space Used (Access)

Initial testing indicates that the Access database requires approximately 11.2 Kilobyte per system per day. To store data for 2,000 systems for 7 days (default setting) would require 157MB (2,000 multiplied by 7 days multiplied by 11.2KB).

The following chart shows the database disk space requirements according to the number of systems and the number of days' worth of data stored in the database.

Access Disk Space
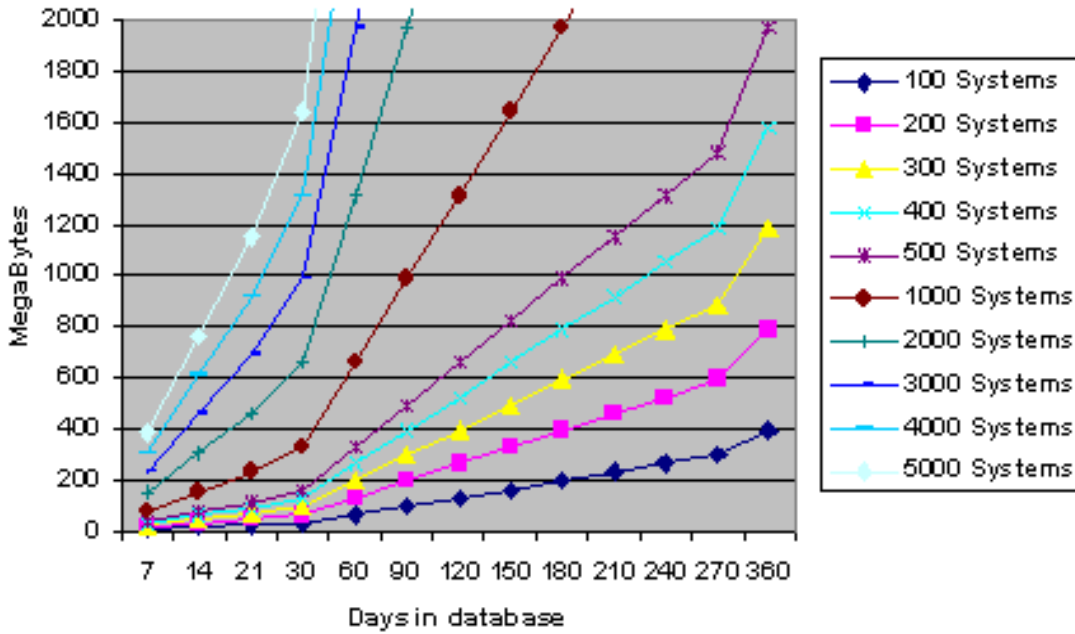
System Requirements to view Online help, End User documents and training videos

To view online help, documents and videos, you need to have the following software(s) and hardware(s) installed in the Reporter machine:

Software		
Component	**Minimum**	**Recommended**
Internet Browser	Microsoft Internet Explorer 5.5	Microsoft Internet Explorer 5.5 or later
Windows Media Player	6.4	9.0

Hardware		
Component	**Minimum**	**Recommended**
Sound Card	16-bit sound card.	Compatible 5.1 multichannel audio sound card.
Sound Output	Speakers or headphones.	5.1 multichannel speakers. For more information about device compatibility, see the documentation provided with your sound card.
Video adapter and monitor	Super VGA (800 x 600) or higher resolution.	Super VGA (800 x 600) or higher resolution.

Other HP OpenView Products

All testing in HP labs was performed on systems where Reporter was installed in a stand-alone configuration with no other HP OpenView products installed. As a result, some products, such as Internet Services, which uses Reporter's scheduler and report generator, will likely increase the the report generation time used for that phase during Reporter's daily activity schedule.

Troubleshooting Task Completion Times

If Reporter continues to take too much time to run through all its phases or tasks, determine which phase is taking the most time. By examining the **status.reporter** file, shown in the Reporter main window (bottom pane), note the starting and ending times for discovery, gather, and repcrys.

- If the **discovery** is taking too long, consider removing networks that are taking too long to browse. Add the systems from those networks manually, using the "bulk discovery" method instead (see Reporter's online Help for instructions). Once again, network speed could also be a factor here, especially when going over a WAN.
- If the **gather** is taking too long, consider increasing the gather concurrency and/or the speed of the CPU. Another factor to examine is network speed.
- If the **repcrys** is taking too long, consider upgrading the CPU speed. Other options include removing any unnecessary reports from Reporter, decreasing the number of customers and/or groups configured, or decreasing the number of days of data kept in the database. If the database is remote, you might consider moving it to the local system.

If you need more assistance with performance issues, please contact your Hewlett-Packard support representative.

📖 Chapter 7: Advanced Topics

This section includes topics for further customizing Reporter to your environment and for your reporting needs. You should know how to use components of Reporter's main window to implement the features discussed in this section.

The following topics are covered:

- [Making Reporter Web pages available on the network](#)
- [Securing Web pages](#)
- [Configuring Reporter to gather data through a firewall](#)
- [Backing up Reporter data](#)
- [Changing the scheduled times of Reporter actions](#)
- [Scheduling a batch file](#)
- [Excluding systems from reporting](#)
- [Automatic system grouping](#)
- [How Reporter handles system names & data sources](#)
- [Recovering a corrupted database](#)
- [Starting over](#)

Making Reporter Web Pages Available

(on the network)

After you have installed Reporter and you have the Web page reports, you may want to make these reports accessible. To do this, you can supply the Reporter system URL as a link from any published Web page. In this way users of a company's home Web page can examine reports generated by Reporter. This method of providing a link to Reporter allows those who routinely access the company Web page to use their Web browser to view reports through a variety of sources compatible with the Web such as Windows-based PCs, UNIX systems, Macintosh systems, and even Web-TV. Of course, the company home Web page security will probably be set up so that only authorized users can access it. If not, consider including a password to the Reporter Web page link.

NOTE: Configuration of links to the Reporter reports from other systems has been facilitated due to the Reporter 3.0 Web server software requirement.

To link Reporter's cover page from your company Web page:

1. Open the company Web page file with your web-authoring application.
2. Enter a hyperlink to the Reporter cover page as: http://<*reporter_server*>/HPOV_Reports/reports.htm

   ℹ️ NOTE: You must supply the fully qualified domain name of the Reporter system.

To link to a specific customer's cover reports page from your company Web page:

1. Open the company Web page file with your web-authoring application.
2. Enter a hyperlink to the configured customer's reports.htm Web page; for example:
   http://<*reporter_server*>/HPOV_Reports/Customers/<*customer_name*>/reports.htm

   ℹ️ NOTE: You must supply the fully qualified domain name of the Reporter system.

Securing Web Pages

If you want to restrict access to reports, use normal Web page security techniques to implement restrictions. As previously stated, you might assign a password to the Reporter link that restricts access to the main reporter Web page. You can also limit access to some reports by aping security to individual report Web pages.

To secure your Reporter Web pages:

- Create a report cover Web page to replace the reports.htm file.
- Save the file with a different name since the reports.htm file will be replaced each time reports are regenerated. Simply saving the current reports.htm file under another name, then editing it should produce the desired results.

ℹ️ NOTE: The addition of new reports to Reporter will be reflected in the reports.htm file and not in the copy you are using. You should edit your new Web page to include the new reports and whatever security you need for them. You can apply security at a lower level by creating your own pages to link into the group and single system Web pages. The Web page hierarchy used by Reporter is as follows (all files are relative to the directory \Program Files\HP OpenView\Data\webpages if you installed in the default directories):

1. The reports.htm file contains links to Web pages for all (All Systems), group (System Group), and single system (Single System) reports.
2. Each ALL report resides in its own directory with a Web page called default.htm. For example, the TopCPU report for all systems will be in file topcpu\default.htm.
3. The group_<groupname>.htm file contains links to all group reports for a single group of systems. Multiple files may be present, one for each group that has reports assigned to it. The <groupname> will be the name of the group of systems.

   Each group report resides in the directory named after the report, but in its own sub-directory named after the group itself. For example, the group report for TopCPU for systems in the NT group will be in file topcpu\NT\default.htm.

4. The single_<reportname>.htm file contains links to each single system which had a report of this type. Multiple of these files may be present, one for each single system report which was generated.

   Each single system report resides in the directory named after the report, but in its own sub-directory named after the system itself. For example, the Top Transactions report for system MyWebServer would be in the file toptrans\MyWebServer\default.htm.

The best approach to finding a particular report is to examine the reports.htm file and follow its links to the desired report. You can then link into that report directly from your own Web pages.

Configuring Reporter to Gather Data through a Firewall

Reporter Discovery and Gather programs can be configured to work with systems outside a firewall as long as those systems are running OpenView Operations Agent for UNIX, the OpenView Performance Agent for Windows systems, or the OpenView Operations subagent. Also, Reporter Discovery and Gather programs can be configured to work with OpenView Operations agents outside a firewall as long as the those systems run in DCE mode. No support is provided for systems using NCS.

For the latest technical information on firewall configurations for OpenView agents, please consult the HPOpenView Support website at http://support.openview.hp.com.

Backing Up Data

To prevent data loss, you should consider implementing two safeguards. First, ensure that Reporter processing does not occur during normal backup times. Second, set up a process that automatically backs up Reporter data.

Schedule Reporter Processing Outside of Backup Time

By default, Reporter performs all its processing between the hours of 12:30 AM to 3:00 AM (or later if you have many systems). This processing should not interfere with normal backup processing except for the amount of CPU and network traffic it consumes. Reporter automatically limits its processing so that it does not consume more than 50% of the CPU for long periods. Even so, in some cases this additional activity could slow the process of backing up system data if backup actions and regular Reporter actions are scheduled at the same time. If your backup occurs during the hours of 12:30 and 3:00 AM, you might want to change the schedule that Reporter follows in performing its actions (see below).

Backing up Reporter Data

By default, Reporter retains the last 7 days of data that it gathers from MeasurewWare Agent (or OpenView Performance Agents for UNIX)   OpenView Performance Agents usually retain 30 to 60 days of data locally. If for any reason the Reporter database is lost, the data can be re-gathered from OpenView Performance Agents as a part of normal Reporter processing.

If you are retaining more data on the Reporter system than you would want to re-capture from OpenView Performance Agents, you should probably periodically make a backup copy of the Reporter database. To back up data, you need to know what database software you are using, whether it be MSDE SP3 (Reporter A.03.5)or MSDE (Reporter A.03 default) or Access (Reporter A.02 default) or MSDE SP3a (Reporter A.03.6) .

To determine your database type, open the Control Panel (you will see SQL Server if using MSDE; you will see Microsoft Access Driver *.mdb if using Access):

1. Select Control Panel from the Windows Start/Settings menu.
2. Double-click ODBC in the Control Panel window .
3. Select the System DSN tabbed page.

Backing up the MSDE database

If you use the default MSDE database, backup procedures are available and described on the Microsoft Web site. The below-mentioned options use Microsoft utilities. Please refer to the Microsoft documentation for suportabilty issues or errors that may occur when using these procedures.

Option #1, for MSDE using SQL Client Tools :
If SQL 2000 Client Tools are installed, use SQL Enterprise Manager to back up your MSDE database.

Option #2, for MSDE using neither Access 2000 nor SQL Enterprise Manager:

If you have only MSDE installed, you can use the TSQL BACKUP DATABASE command and execute with Osql.exe (command line Query tool).
MSDN as well as the SQL online books provide detail on how to use the stored procedures outlined below. Create a backup/detach/restore, etc. procedures, by entering syntax, as described below:

NOTE: The steps below provide an example of how to use the various stored procedures with MSDE to perform a backup or restore. You may want to customize the steps for your particular environment. Some apitional things you might want to do is to create a daily backup job, or produce a daily backup report. Refer to Microsoft (MSDN) documentation on osql utility, BACKUP DATABASE and RESTORE DATABASE for other options and features. Be sure and verify that your backup and restore work correctly.

IMPORTANT: Stop Reporter services and make sure that no client tools are accessing the Reporter MSDE database before you attempt backing up or restoring the database.

Certain defaults have been chosen in the examples that follow. You may be required to change the directory name, user name, and password.

To stop the Reporter service

1. Select Control Panel > Administrative Tools > Services.
2. Locate the Reporter Service. Click Stop.
3. Close the Reporter Graphical User Interface.

Backup the MSDE database, syntax is as follows:

```
osql -S.\OVOPS -Usa -P
USE reporter
BACKUP LOG reporter WITH TRUNCATE_ONLY
EXEC sp_apumpdevice 'disk', 'Reporter_BKUP',
'C:\Program Files\HP OpenView\Data\Databases\backup\Reporter_1.bak'
BACKUP DATABASE reporter TO Reporter_BKUP WITH INIT, STATS
EXEC sp_dropdevice 'Reporter_BKUP'
go
```

Restore the MSDE database, syntax is as follows:

```
osql -S.\OVOPS -Usa -P
USE master
RESTORE DATABASE reporter FROM
DISK='C:\Program Files\HP OpenView\Data\Databases\backup\Reporter_1.bak' WITH RECOVERY,
REPLACE, STATS
go
```

You may need to use the following two examples if you want to move the database from one system to another:

Detach the Reporter MSDE database, syntax is as follows:

```
osql -S.\OVOPS -Usa -P
```

```
sp_detach_db reporter
go
```

Attach the Reporter MSDE database, syntax is as follows:

```
osql -S.\OVOPS -Usa -P
sp_attach_db reporter
'C:\Program Files\HP OpenView\Data\Databases\reporter_1.mdf',
'C:\Program Files\HP OpenView\Data\Databases\reporter_1.ldf'
go
```

Backing up the Access database

If you use an Access database, the gathered data and operational parameters for Reporter are kept in a single database. By default this database is a single file \Program Files\HP OpenView\Data\datafiles\Reporter.mdb. To back up Reporter data stored in Access, make a copy of the file.

Remember that when you are backing up the Reporter database, it is unavailable for other Reporter activities. If you back up the database during data gathering or discovery, messages may appear such as "WARNING: Unable to update the database after 20 tries." You can recover most Reporter activities so that you can recapture lost data during the next processing cycle. Even so, try to avoid repeated loss of data by performing database backup at times when Reporter is not busy.

NOTE: If you do full system backups, the database data will not be backed up if its datafiles are in use. You can verify successful backups by checking for copies of Reporter data files reporter_1.mdf and reporter.1_ldf in your backup log files.

Archiving Web Page Reports

In its normal routine Reporter replaces all report Web pages each time new reports are generated. Because of this you do not need to manage old report pages or worry about filling up your disk space with them. However, if you want to save past report Web pages, you can do so as follows.

1. All report Web pages are in the \<*inlstall_directory*>\data\webpages directory and its sub-directories. The simplest way to retain old report pages is to rename this directory before the new reports are generated.
2. Arrange a method for naming the old Web pages directory to retain the desired number of copies and to manage disk space consumption.
3. A sub-directory named images resides in the same directory as the webpages directory. The files in this directory contain graphics/pictures that are displayed in some Web pages. Because these files are not updated but are used mainly for labeling the report (such as in the report banner displayed at the top of the rpoert), you do not need to copy the files to the archive location.
4. Arrange for the renaming process to happen each time that you want to retain a set of Web reports. One automated way to do this is to construct a batch file that does the renaming of directories for you. Next, add an entry to Reporter's schedule to execute this batch file prior to generating a new set of

reports.

Changing Scheduled Times of Actions

Several reasons could exist for changing the default processing times for Reporter actions. You might want to avoid daily backup times as described above. You might also have an environment where the systems are unavailable at the default times. Many sites may power down their workstations each night. If you want to discover and gather data from these systems, you may have to reschedule these actions to occur during the day (at lunch hour for example).

The procedure below lets you change all currently scheduled actions to occur at the desired times, while still maintaining the sequence of events that are important to Reporter's processing cycle. Using this procedure also ensures that new actions are scheduled at the new times rather than the previous default times.

1. Stop the Reporter Service (to ensure no processing starts up due to a typing error).
2. Set the default times for new actions by selecting an item below the Schedule icon, right-clicking it, and selecting Global Settings from the short-cut menu.
3. Select the starting time for each standard action. Sequence times in this order (1) Discovery, (2) Gather, (3) Maintenance, and (4) Report (earliest times at the top of the list, later times toward the bottom). Do not be concerned about allowing enough time for each action to finish before starting the next action. As long as the times are not identical, the Reporter service ensures that actions occur in the proper sequence.
4. Restart the Reporter service.

Scheduling a Batch File to Run

You can also schedule a batch file to run through the OpenView Reporter Scheduler by completing the following steps:

1. Create the batch file and save it.
2. Open the Reporter main window.
3. In the left pane right-click Schedule and select Add Schedule.
4. In the Ap Schedule Entry window that appears in the Program Name box type:
   cmd.exe
5. In the Parameters text box type:
   /c "cmd.exe /c \temp\<*myfile*>.bat"

   NOTE: Be sure to enter the syntax in the Parameters text box exactly as it appears on the above line including the double quotation marks and substituting the directory where you have located your batch file and the batch file name.

Excluding Systems from Reporting

There are two ways to exclude systems from the reports produced by the Reporter.

1. Set the Discovery Area so that undesired systems are never discovered.
2. Set the System Information parameters for a system that has already been discovered to exclude that system from all reporting.

Use the first technique if you want to selectively exclude numbers of systems that are separated into different networks and domains. Use the second technique to exclude selected systems from within a discovery area that you otherwise want discovered normally. You might exclude your test systems when they reside inside your normal network domain using the second technique.

Procedures for setting the discovery area are documented in the main part of this document and in Reporter's online Help.

To exclude single systems from reporting:

1. Allow the system to be discovered along with the other systems in the discovery area.
2. Select the system in the Discovered Systems area of the left Reporter window pane (the right pane should display the details for this system).
3. In the right pane, double-click System Information.
4. Double-click Exclude in the list of items under System Information. Each time you double-click this item, it toggles between *Yes* and *No*. Accept the Exclude (Yes) status to exclude this system from further reporting.

What happens when a system is marked as excluded?

- A maintenance program removes information on this system from the Reporter database. The system should not appear in any reports that are generated after this information is removed.
- The system, however, still remains in the Discovered Systems area under the ALL group so that you can later select the system and change the Exclude setting if necessary.
- While excluded, the system still is available to the discovery process, but no information is copied from the system into the Reporter database. Excluded systems should not be present in any of the Reporter's standard reports.
- If you create any of your own reports using the SYSTEMS table, you should include a select statement which selects only systems whose exclude value is zero. (See the a_systems.rpt template file for an example.)

If you exclude a system by mistake, no permanent damage occurs. Just go back into the System Information (as explained above) and toggle the Exclude setting back to *No*.

- The maintenance program is scheduled to run at the normal maintenance time (usually 1:55 AM the next morning). If this program hasn't run yet, change the System Information Exclude value to (No) and no data will be removed.
- If the maintenance program has already run and all data has been removed, you can still remedy the situation. Remember, too, that you still have the historical data maintained by OpenView Performance Agent. To correct the situation, reset the System Information: Exclude value to (No) and Reporter will re-discover the system and gather data from it into the Reporter database. You need to ap this system back into the system groups you previously included it in and ap any single system reports or metric lists that were previously assigned to it.

Automatic System Grouping

You can automatically group systems using the automatic grouping feature. With this feature you can set Reporter to automatically group newly discovered systems based on a variety of system characteristics. Whenever a system is discovered, Reporter examines its characteristics and aps it to any group(s) it matches. Re-discovery of a system can ap the system to apitional groups if its characteristics or the group specifications change.

You can find this feature by selecting AutoGroups from the Configure submenu on the File menu in the Reporter main window. The following dialog box appears where you can add or alter specifications for automatic grouping.

Automatic Grouping

| Group Description | By Operating System | New | Delete |

Group Name: \1

Match Parameters
○ All   ● Any

Parameter: OS Name   Browse

Pattern: (.+)   Set   Clear

Numeric:

Currently Set Parameters:

OS Name: (.+)

OK   Save   Test   Help   Cancel

The Automatic Grouping dialog box lets you set automatic grouping specifications for newly discovered systems. You can use any number of specifications, each associated with the **Group Description**. Each specification links to the name of the group assigned to systems that match the specification. The **Group Name** box can contain a literal string and/or string substitution codes consisting of a back slash (\ ) and a number. Substitution codes are covered later. The remaining components of the dialog box follow.

Parameter box: allows you to select one system identification criteria from a drop-down list. Depending on the type of data you select, the next boxes are enabled or disabled. Selecting a character-based criteria allow pattern matching using the **Pattern** box. Selecting numeric-based criteria causes the boxes in the line labeled **Numeric** to be enabled. A summary of all criteria currently in effect for this specification is shown in the large box at the bottom of this section.

- **New** and Delete buttons: create and destroy autogrouping specifications.
- **Save** button: Saves the specification to the database. If you make changes to a specification but do not save them, you are prompted to save before selecting a new specification or exiting the screen.
- **Match Parameters** section lets you indicate whether a match will be made for systems that match any or all of the criteria shown:
  - All requires a system to match <u>every</u> defined criterion in order to belong to the group.
  - Any requires a system match any <u>single</u> defined criterion to belong to the group.
- **Browse** button: lists the current values for the parameter you selected. With this information, you can determine the exact format or range for the group you are defining. You can select a value from the browse list to copy it into the pattern or numeric boxes.
- **Set** button: After you enter the pattern or numeric match for the parameter you selected, select the **Set** button to ap it to the Currently Set Parameters list.
- **Clear** button: you can select **Clear** to remove a parameter from the list.
- Test button: lists the systems that match the parameters and the group name to which they would be assigned in the bottom part of the screen. Note that no changes are actually made to the groups in the database at this time. The Test button shows only what systems would match the specifications. You can see from the text that appears above the white box how many systems did and did not match the parameters.

Pattern Matches

Parameters that are character strings let you define patterns in the **Pattern** box for matches with systems that you want included in the group. The contents that you include in the Pattern box are recognized as *regular expressions* to allow a wide flexibility in pattern matching. At its simplest, a character string can successfully match a system if the system data contains that character string in any position.

See Regular Expressions for rules for using regular expressions.
See Pattern Examples for examples of patterns that use regular expressions.

One advantage of regular expressions is that you can specify sub-strings that match any or all of the data in the metric. You can use these sub-strings as variables in the **Group Name** using a back slash followed by a field number. The first matching sub-string replaces the "\1" field, the second the "\2" field, etc. For example, if you select the "Agent" as the Parameter and enter a pattern of "(.+)", it will will match any non-null string in its entirety. A system with

an agent metric containing "SCOPE A.02.00" would replace the Group Name "Agent(\1)" with the agent metric name so that the Group Name displayed below Discovered Systems in Reporter's hierarchy would be Agent(SCOPE A.02.00).

Numeric Matches

Parameters that are numbers let you match a number or range of numbers using the **Numeric** boxes.

Numeric	> ▼	128	AND ▼	< ▼	300

These boxes define two sets of numeric comparisons, joined by a conjunction. Each numeric comparison specifies an operator (=, >, <, >=, <=) for (Equal To, Greater Than, Less Than, Greater Than or Equal To, Less Than or Equal To) and a numeric value. The conjunction may be "AND" (both comparisons are true"), "OR" (either comparison is true") or a blank (the second comparison is not used).

You can specify a match of all values EXCEPT for one specified by specifying Numeric > value OR < value.

After entering or editing a pattern or numeric match , don't forget to press the **Set** button to ap it to the Group Specifications list.

Regular Expressions

The **Pattern** box in the Automatic Grouping dialog box is used to enter pattern-matching criteria in the form of regular expressions. Regular expressions have been used in many operating systems over the years and so may be familiar to you. Still, there have been some variations in what patterns are allowed, so a review of what is allowed might be useful. If you are not familiar with regular expressions, the material in this section may make you want to review the rules governing regular expressions in more detail in another source. The material

here shows simple to more complex concepts with examples sprinkled throughout. The examples are formulated specifically to match the use of regular expressions by Reporter; consequently other uses are not discussed.

Regular expressions: A regular expression (often abbreviated RE) consists of **regular characters** and characters with special meanings.

Regular characters: Any non-special character will match an exact copy of itself in the same case (upper/lower case). A string of these non-special characters will match the exact same string (with no extra characters of other "stuff" in the miple). Unless told otherwise, the pattern string may occur ANYWHERE in the data field. Note than **blanks** are regular characters and so are significant in the pattern.	Example*: Pattern "rose" will match data "roseville", "Red Red rose", and "arose"* but not "Rose" (upper case "R") and not "ro se" (blank in the miple).*
Line Anchors: You may specify that a pattern must start at the beginning of the field by starting with the special character **caret** "^". (carets may be used as special characters elsewhere as well, but as the first character in a pattern, it means the rest of the pattern must start at the beginning of the data field. Similarly, you may specify that a pattern must end at the end of the field by making the last character in the pattern a **dollar sign** "$".	Example: *Example: Pattern "^rose" will match data "roseville", but not "Red Red rose", and "arose".* *Pattern "rose$" will match "arose" and "Red Red rose" but not "roseville".* *Pattern "^rose$"will match "rose" but not "arose" "Red Red rose" or "roseville".*
Match any character: The Period will match any single character.	Example: Pattern will match "rose", "rise", "rose" "r#se" but not "roose" or "ro se"
Match a character sequence: Repeat the period to match a sequence of any characters.	Example: *Pattern "..." matches any three characters. "abc", "123", "a c", but not "ab" (only TWO characters)*
Repeats: Several special characters are used to specify that part of a pattern may be repeated. For example, following a character by the **asterisk** "*" will match zero or more occurrences of that character.	Example: *Pattern "a*" will match "a", "aa", "aaa", even "" (zero occurrences of "a") but not "a a" or "aba".*
Wildcard as plus: plus "+" works the same as the asterisk except that it matches ONE or more occurrences of the character.	Example: *Pattern "a+" will match "a", "aa", "aaa" but NOT "" (zero occurrences).*
Wildcard as question mark: question mark "?" will match exactly ONE occurrence of the character, or nothing at all but not multiple occurrences of the character.	Example: *Pattern "a+" will match "a", ""a", "aaa" but NOT "" (zero occurrences).*

Repeating Patterns: The repeat characters work with single character patterns, but they may also be used to match repeating patterns which are enclosed in **parenthesis** "**()**" and ranges which are enclosed in **square brackets** "**[]**". These constructs are explained more fully later in this section.	Example*: Pattern "(abc)*" will match "abc", "abcabc" "abcabcabc" and ""* *Pattern "[0-9]+" will match any contiguous sequence of one or more numbers*
Ranges: A range of characters may be specified inside of **square brackets** "**[]**". A range of characters may be specified by simply entering them between the brackets. "[0123456789]" would match any single character than was a number. "[ab]" would match either an "a" or a "b". A short hand may be used to specify all characters falling between two others is to specify the first character, a **dash** "**-**" then the last character. Thus "[0-9]" is the same as "[0123456789]". Any characters may be used as long as the follow the ASCII collating sequence. "[a-zA-Z]" matches any alpha character. It may be easier to specify a range of characters that you do NOT want to match. TO do this, place a **caret** "**^**" as the first character then specify the characters to exclude. "[^xyz]" would match any character EXCEPT "x", "y", or "z". It is important to remember that a range matches a SINGLE CHARACTER in the data string. If you want to match more than one character then you should follow the range specification by one of the repeat characters	Example*: The Pattern "[0123456789]" would match "0" or "8" but not "45" (two characters, it would match the first one "4" but not both of them).* Example*: The pattern "[0-9]+" would match any string of digits (but at least one digit). It would match "1","1234", "668888" but not "" (must have one digit and not "555-1212" (the "-" isn't a digit In this case the pattern would match just the "555"))* Example*: The pattern "a[0-9]*" would match "a" followed by zero or more digits. It would match "a", "a0", "a77655"* Example*: The Pattern "a[^b].*" would match "a" followed by a character which is NOT "b" followed by anything (zero or more occurrences of any character). It would match the string "ace", "ac", "a9", and even "actuarial" but it would not match "abcradabra" or "a" (range matches ONE character).*
OR: You may specify more than one regular expression that may be used to match a given field by separating them by a **vertical bar** "**|**". Remember that each regular expression is evaluated against the entire field contents so if you need anchor characters you must repeat them.	Example*: The Pattern "abc|def" would match either "abc" or "def"*

Special Characters as regular characters: The following are special characters: ".", "*", "?", "+", "(", ")", "{ ", "} ", "[", "]", "^", "$", "|", "\". If you want a special character to match itself, precede it with a back slash "\" and it becomes non-special.	Example: *The pattern "[0-9]+\.[0-9]+\.[0-9]+\.[0-9]+" matches the four parts of an IP address, separated by periods.* Example: *The pattern "…\\…." matches any three characters, a back slash then four characters.*

Group Names & Sub-strings

Sub-Strings are useful in automatically generating group names based on the variable contained in the sub-string. Any sub-string that matches a pattern can be used as part of the Group Name that appears in the Reporter hierarchy under Discovered Systems. The Group Name should have a **back slash** "\" followed by a **number** to indicate where the value for the sub-string should be substituted. The number indicates which sub-string value to substitute (in case there is more than one sub-string in your pattern).

Substring Matches: Regular expressions can be built up of single character patterns as described above. In apition they can be broken into sub-strings using **parenthesis** "()". The characters inside the parenthesis are treated as a separate regular expression that is matched against the data string. As illustrated in the preceding section, this can be combined with the repeat characters to match repeating patterns of characters in the data.

Example 1: *The Pattern "^(ab.)" will match any data string starting with "ab" and then having at least one more character. Thus it would match "abc", "ab9", "abcradabra" (with the sub-string matching only the "abc" part of "abcradabra") but not "ab" (nothing to match the period) and not "dabc" or " abc" since the caret indicates the pattern must start the field.*

Now, if you specify a Group Name of "Special - \1" then the actual group names generated would be:

Parameter Value	Group Name
abc	Special - abc
ab9	Special - ap9
abcradabra	Special - abc

Example 2: *Suppose you wanted to create a group for any non-null Operating System Name and you wanted that group to be exactly the same as the Operating System Name itself. Specify a Pattern for the Operating System Name like "(.+)" which says to match a sub-string (parenthesis) of any character (the period) repeated one or more times (the plus). Then*

specify the Group Name as "\1" which will substitute the first matched sub-string as the Group Name.

Example 3: *Suppose your System Ids consist of three characters for the Area Code, any one character then four more characters for the region. You want to group names consisting of these two fields in the reverse order.*

Match specification: *The pattern "^(…).(….)$". Note that by anchoring the string at the beginning and end of the field (caret and dollar sign) we will not match any data field that deviates from this pattern. "aaa-rrrr" would match but "aaa-rrrrz" and "aaa - rrrr" would not.*

Group name result: *Next the Group Name "\2\1" specifies a group named with the two matched sub-strings, but with the second one first.*
Parameter Value "aaa-rrrr" => Group name "rrrraaa"

Nonprinting Characters: Although their use in the Reporter Automatic Grouping function is not likely, the pattern field will accept the following escape codes for non-printing characters:

Escape Sequence	Character Code	Meaning
\a	0x07	Bell
\f	0x08	Form Feed
\n	0x0C	Newline
\r	0x0D	Carriage Return
\t	0x09	Tab
\v	0x0B	Vertical Tab

Pattern Examples

The following character strings illustrate some typical constructions for the regular expressions which are allowed in the pattern field:

Pattern	Use
. (period)	Matches any character
[0-9]	Matches character in the range of digits
[a-zA-Z]	Matches an alphabetic character
[a-zA-Z0-9]	Matches an alphanumeric character
non special character	Matches the same character anywhere in the field

^	Requires match to occur at the beginning of the field
$	Requires match to occur at the end of the field
*	Matches any repeat of previous pattern 0 or more times
+	Matches any repeat of previous pattern 1 or more times
?	Matches previous pattern or nothing at all
( )	Specifies to treat the enclosed expression as a sub-string. Sub-strings as variables can be substituted in the Group Name using \1, \2, \3, etc., for first, second, third sub-string.
|	Allows separate matching; functions the same as OR (matches either pattern)
\	Escape character, treats special character as non-special.

This table contains some patterns that might be useful in Reporter

Parameter	Pattern	Matches
any	.+	Non-empty field
any	(.+)	Non-empty field, allow sub-string to substitute in the group name in place of "\1"
any	(.*)	Anything, including an empty field. Allows sub-string substitution in the group name
System ID	^(...)	Allows grouping of systems based on the first three characters of their system ID. Group Name might be "Application - \1"
IP address	^([0-9]+)\.([0-9]+)\.([0-9]+)\.[0-9]+$	IP address and allows sub-strings for the first three parts of the address to be used in the Group Name as \1, \2, \3.
Agent	^SCOPE	Any system that has OpenView Performance Agent installed.
OS Name OS Release OS Version	^NT$ ^4\.0 ^SvcPk 3	Any system running NT 4.0 with Service Pack 3. NOTE: The Match All radio button should be selected to avoid matching any system that matches ANY of these three specifications.
Agent	^SCOPE.*([A-Z]\.[0-9][0-9]\.[0-9][0-9])	OpenView Performance Agent-installed software, allows a group name such as: "OpenView \1" to group systems by the OpenView agent release installed.
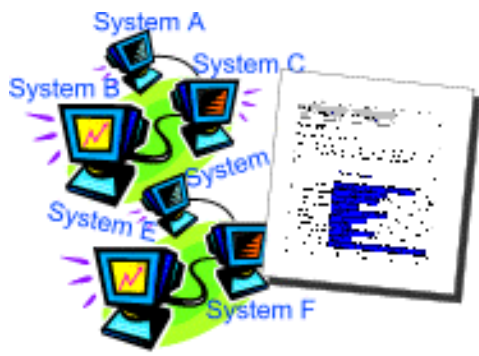
How Reporter Handles System Names and Data Sources

(and what you can do in problem situations)

The reports generated by Reporter list systems by their network names. However, what happens if a system's network name is not unique or if one system has more than one network name? Further, what happens if data from one system is transferred to another system? Can Reporter sort out these situations correctly? For the most part, yes. Reporter automatically resolves these potential problems because OpenView Performance Agent log files identify the original source of data by including a unique system ID.

However, other situations can occur that you may need to be aware of so that you can avert problems. These situations involve duplicate network names and duplicate system IDs within log files. This section covers what you can do, and what you might look for in a complex network environment where it is easy to duplicate names or values inadvertently.

Network Names Identify Systems in Reports

Reporter's pre-configured reports use each system's network name to identify the system. However, you can, to some degree, control the system names in the reports by controlling what network names are used to discover each system. Note that system names must be usable by the name server to link to a valid network IP address. You cannot use any arbitrary name, but you can choose from equivalent names, such as the short name, the fully qualified name, or the IP address itself.

NOTE: After Reporter discovers a system, that network name is fixed. The network name cannot be replaced, even if the system uses a second or third network name. If you want to change network names, you must delete the system entirely or start over with an empty database and discover again, using the desired network names.
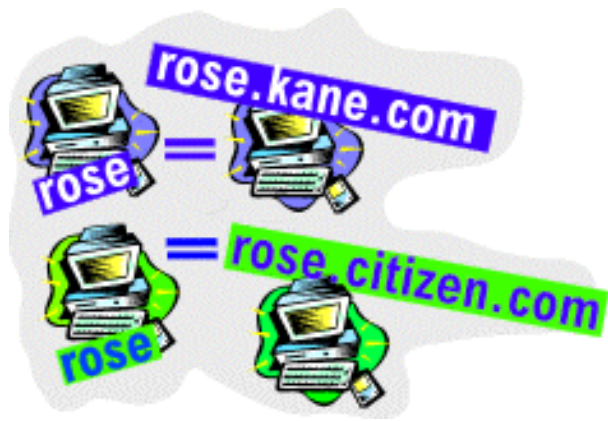
Duplicate Network Names for Different Systems
(what you can do to resolve)

When you open Reporter's main window and use the Discovery Area to discover systems in the Microsoft Windows systems network domains, the names that appear in your reports match those that you see in the expanded Discovery Area tree. The names are usually short names like *rosebud*. If the short form of any name occurs more than once in the network, Reporter will discover the first one and ignore the second. For this reason, you need to resolve this situation before the Discovery process begins.

What you can do to resolve duplicate network names before running Discovery:

In cases where network names duplicate each other in the short form, you can resolve this problem by using Add Single System or Bulk Discovery and enter the names in their longer, fully qualified form. Both methods require that you manually enter names. The name that you enter must be a valid name on the name server so that it can be resolved with its network IP address. For example, enter rosebud.citizen.kane.com and rosebud.flowers.are.us instead of allowing Reporter to automatically discover the first short network name of rosebud and ignore the second duplicate name (for a different system).

NOTE: You can test the validity of a name by entering the command ping <system name> in the Command Prompt window. If ping finds the correct system, you can use that name in Add Single System or Bulk Discovery.

Multiple Network Names for the Same System

(no resolution necessary)

If multiple network names refer to the same system, Reporter resolves the situation. This situation can occur when Reporter discovers a single system that uses network identifiers, such as rosebud, rosebud.citizen.kane.com and 123.13.45.221. Even though Reporter might discover this single system multiple times because of its different network names, Reporter associates these multiple network names with one OpenView Performance Agent system ID. Reporter assigns data to the first network name and ignores the others.

NOTE: To view the network name and the OpenView Performance Agent system ID for any system, select the system in the Discovered Systems area of Reporter's main window. Expand the System Information icon in the right pane. SystemID contains the network name for this system and SystemName contains the OpenView Performance Agent system ID value.

OpenView Performance Agent System IDs Link Performance Data and Systems

When Reporter has discovered a OpenView Performance Agent -installed system, it looks within the log file set to identify the source of the collected data through the metric: GBL_SYSTEM_ID. This is the system ID, which is stored in the database (even though you do not see it in any of the default reports). With a unique system ID value in each log file set, no problems occur.
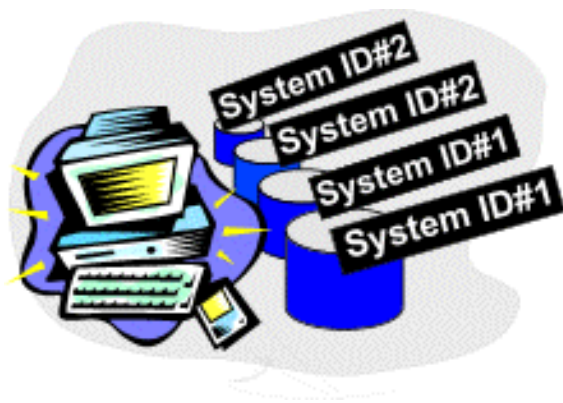
Multiple Log files on the Same System

Reporter can determine if log files (such as a raw log file and an extracted log file) are from the same system. By checking the system ID (GBL_SYSTEM_ID metric), Reporter recognizes that these log files originate from the same system and ignores the duplicate.

If two log file sets have the same system ID (duplicate values for the GBL_SYSTEM_ID metric), Reporter recognizes the first and that becomes the only log file set (and system) for which a report is generated. If you are missing a report on a system, possibly the system ID metric value for that system is duplicated by another system.

How to tell if you have this problem: From the Reporter window select Configure from the File menu and set the trace level to a value of 1 or higher. Next, allow the Discovery process to execute one complete cycle, then examine the file trace.discovery in the

<install_directory>\data directory. Look for messages with the string "Found Duplicate System IDs". Determine if the data sources listed are actually log files from the same system. If not, you may need to change the system ID as described in the following paragraph.

What you can do to resolve duplicate system IDs: Open the OpenView Performance Agent parm file for the missing system. *You need to change the name for the duplicate(s) by editing the parm file on that (those) system(s). (See the OpenView Performance Agent documentation for details).*

NOTE: If the default value is not present (no system ID is present in the parm file), check the system ID according to the operating system.
For example:
UNIX systems use the nodename, obtained by entering the command: uname -n.
Windows NT systems use a name obtained from an API call: GetComputerName.
Windows 3.x and Windows 95 systems use the PC label assigned to the system during MeasureWare Desktop installation.
MPE/iX systems use the variable listed as HPSYSNAME.
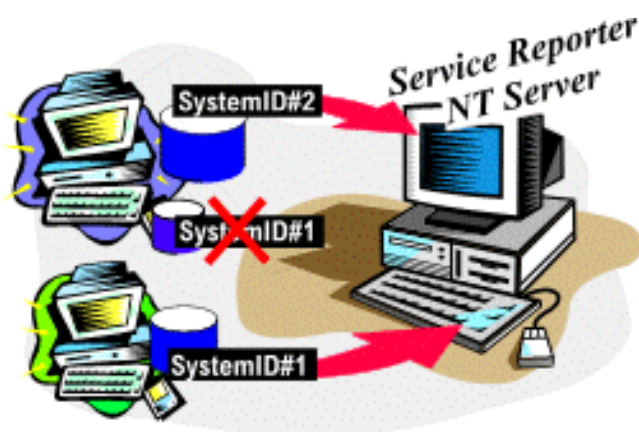
Automatic Update of Proxied Systems to Direct Sources of Data

If during the discovery process, Reporter discovers a duplicate system ID for a log file originating directly from a system and the system already discovered is being proxied, Reporter replaces the entry for the indirect data source with the new, non-proxied system. Existing data, groups, reports, etc., are updated from the old proxied identifier to the new direct one.

As illustrated above, Reporter could originally access only System #2, which stored its own log files and those from System #1. Later when Reporter gained access to System #1, it automatically ignored the second log file on System #2, and retrieved that data directly from System #1.

Using the OpenView Performance Agent System ID in Reports

If you create your own reports or edit the standard reports using Crystal Designer (purchased

separately), you can use the OpenView Performance Agent system ID as follows:

Each Reporter database table created from a metric list has the same name as that metric list. At minimum, the table contains these four standard items:

- ID: a sequential number used to guarantee uniqueness on each record.
- SYSTEMNAME: the network name for the system (as used by discovery).
- DATETIME: the local time on the OpenView Performance Agent system for the record
- GMT: Greenwich Mean Time for the record.

The **SYSTEMS** table contains discovered information about each system. Each discovered system has one record. Within that record, fields of   interest are:

- SYSTEMID: the network name for the system (as used by discovery).
- SYSTEMNAME: the OpenView Performance Agent GBL_SYSTEM_ID value.
- PROXYNAME: the network name of the system which is hosting the proxied log files for this system.

To use the OpenView Performance Agent system ID name (the value for GBL_SYSTEM_ID) in a report, you can link from a metric list table using the **SYSTEMNAME** item into the Systems table using the **SYSTEMID** item.

Create Report Expert

Data | Links | Fields | Sort | Total | TopN | Graph | Select | Style

GLOBAL
- ID
- SYSTEMNAME
- DATETIME
- GMT
- SHIFTNAME
- INTERVAL
- GBL_CPU_TOTA
- GBL_DISK_PHYS
- GBL_SWAP_SPA
- GBL_MEM_PAGE

SYSTEMS
- NETWORKID
- DOMAINID
- SYSTEMID
- SYSTEMNAME
- PROXYNAME
- DISCOVERYDATETIME
- OSNAME
- OSRELEASE
- OSVERSION
- MACHINETYPE

Options... | Delete | Tables... | Arrange | Smart Linking

<< Back | Next >> | Cancel | Preview Report | Preview Sample

Recovering a Corrupted Database

The current default database, MSDE Service Pack 3a, and the previous default database, Access, have different database recovery procedures. If you have configured a database other than the default, please see the database documentation accompanying the database product.

- For a Reporter system that uses MSDE as its database, you can proceed according to the software available on your system.
  - If you have Access 2000, you can start Access Project and from the Database Utilities menu, select Compact and Repair Database.
  - If you have SQL Server Client Tools, you can use SQL Enterprise Manager (see the product

documentation).

- For a Reporter system that uses Access, you can recover the damaged database by completing the following steps:
    1. Select Stop Reporter Service from the Action menu.
    2. Repair the database: Open the Control Panel, choose ODBC, select the System DSN tab and highlight the Reporter database driver (*.mdb). Select the Configure... button and in the next window the Repair... button.

If this procedure fails or if the database is still unusable, please see the section that follows to create a new database.

Starting Over

The current Reporter release uses MSDE Service Pack 3a as its default database, while previous Reporter releases used Access. Removing all data from either   Reporter database differs. The differences are noted in the procedure below.

Recreating MSDE Database

The following steps cover deleting an existing MSDE Reporter database and recreating a new MSDE Reporter database. The script that you run must be executed from the *<install_directory>*\bin directory where newdb.exe resides. Also, the Reporter DSN must name a MSDE database.

1. To stop the Reporter service, open the Reporter main window, select Action>Stop Reporter Service. (or select the Start/Stop Reporter Service toolbar button ![icon]).
2. If the Reporter main window is open on your Desktop, close it.
3. Open an MS-DOS command window.
4. Use change directory (cd) command to change to the *<install_directory>*\bin directory.
5. Type the following command at the command prompt: cscript RecreateMSDEDB.vbe
6. Check the *<install_directory>*\Data\status.Reporter file for the status of newdb.

Recreating the Access Database

1. To stop the Reporter service, open the Reporter main window, select Action>Stop Reporter Service. (or select the Start/Stop Reporter Service toolbar button ![icon]).
   (You can also do this from the Control Panel Services area by highlighting Reporter Service and pressing the Stop button.)
2. If the Reporter main window is open on your Desktop, close it.
3. Reporter Access database users: If you want to save existing data, rename the database; if you do not

need to save existing data, delete the file \<*install_directory*\data\datafiles\Reporter.mdb.

4. For Windows 2000 systems: Select Start>Settings>Control Panel>Administrative Tools>Data Sources (ODBC)
   or
   For Window NT: Select Start>Settings>Control Panel>Data Sources (ODBC).
5. Select the System DSN tab and within the System Data Sources: select Reporter and click the Configure… button.
6. In the window that appears that click the Create … button.
7. In the New Database window browse to the \<*install_directory*>\data\datafiles\ directory and in the Database Name text box type Reporter and clickOK.
8. Now create the Reporter tables within the database by running the <*install_directory*>\bin\NewDB.exe program.
9. Change the default configuration in the Discovery area as necessary in the Reporter main window (dragging and dropping additional systems to the Discovery Area).
   (To create new system groups, wait until after the Discover process has completed so that you can add the newly discovered systems to those groups.)
10. To restart the Reporter service, in the Reporter main window select Action>Start Reporter Service.

Reporter immediately begins the processing cycle as it did when it was first installed.

Recreating the Oracle Database

1. To stop the Reporter service, open the Reporter main window, select Action>Stop Reporter Service. (or
   select the Start/Stop Reporter Service toolbar button).
2. If the Reporter main window is open on your Desktop, close it.
3. On the Windows system where Reporter is installed, copy the \<*install_directory*>\newconfig\oracle\DropNewdb.sql file to the UNIX system in directory: $ORACLE_HOME/dbs/
4. On the UNIX system where the Oracle database is installed, verify that for the current Oracle session, the Oracle SID=REPORTER.
5. Log on as oracle and at the oracle> prompt, enter svrmgrl to start the Oracle Server Manager program.
6. At the SVRMGR> prompt, enter connect internal.
7. Enter the following to remove data from the database: @ORACLE_HOME/dbs/dropnewdb.sql
8. <u>Return to the previous procedure and complete remaining steps</u>.

INDEX

hp OpenView   reporter

Documentation Set

- **Concepts Guide**

  Contains information, both basic and advanced, on Reporter features, interface, use, and daily operation.

- **Installation and Special Configurations Guide**

  Contains an overview of the installation process and detailed instruction sets for configuring alternative database software to store and retrieve Reporter data. Also included are the setting requirements for these special configurations

- **Release Notes**

  Contains new feature updates added to the current version as well as requirements for installation and use.

Legal Notices

Warranty.

Hewlett-Packard makes no warranty of any kind with regard to the material in this web site, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental, or consequential damages in connection with the furnishing performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend .

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices.

Trademark Notices.

Java^{TM} is a trademark of Sun Microsystems, Inc.