Performance Management

This chapter describes the CiscoWorks features that help you manage the performance level of your network, including managing individual devices, lines, and interfaces. This chapter also describes how to monitor CiscoWorks performance by:

- Creating Polling Tables Using Device Polling
- Using Polling Summary
- Using the Real-Time Graphs
- Using the Path Tool
- Using Show Commands to View Router Data
- Loading Private MIBs on HP OpenView or NetView for AIX
- Writing Reports with the Sybase Data Workbench (DWB)

Performance Management Applications

Several CiscoWorks applications help manage the performance of the SNMP devices in your network. A brief description of each application discussed in this chapter follows:

- Device Polling—Probes and extracts data about the condition of your network devices.
- Polling Summary—Views polling data, and stops and starts polling.
- Real-Time Graphs—Monitors the behavior of device interfaces or other network elements suspected of operating in a degraded mode and displays them in a graph.
- Path Tool—Graphically displays a route of the path from a source device to a destination device.
- Show Commands—Provides data similar to the router EXEC **show** commands output.
- Sybase DWB—Allows you to access the Sybase Data Workbench application to write reports.

Use these applications to collect network data as a baseline before your network develops problems. These applications are discussed in detail in the following sections.

Creating Polling Tables Using Device Polling

This section describes how to set up a customized polling table (or table group) for devices and interfaces on the internetwork using Device Polling.

Device Polling Overview

CiscoWorks allows you to probe and extract information about the condition of your networks by using the polling feature. Information acquired from these polls is stored in the database for further analysis. The construction and use of polling configurations allows you to compare relative performance and status of devices and interfaces on the network. You can poll devices individually, in groups, or poll all devices at once based on your disk space availability. Use multiple polls to collect a wide range of data at varying polling interval rates.

For consistency, this publication uses the term *object* as a replacement for such terms as Management Information Base (MIB) variables, MIB object instances, and so on. These terms are used interchangeably in this guide.

You must have security privileges to read and write to the Device Polling window. If you do not have write privileges, you might not be able to apply or delete table changes. You can use the **Change Domain** command to change to a different domain in order to access other devices for polling. You can use the **Change User** command to change to a username with different privileges. Use the **Privilege** command on the Security menu to check your privileges for Device Polling.

Device Polling Process

The polling process consists of the following general procedures. Each task is described in greater detail later in this chapter.

1 Construct a polling table (or polling group) in Device Polling that contains the polling configuration components. (When you complete construction of the table, select **Activate Changes** to notify the polling daemon of these changes.)

Note Activate Changes occurs automatically when you exit Device Polling using the **File>Exit** command.

2 Access the polling table in the Polling Summary application and begin or end polling the devices defined in the configuration (if you have not yet done so).

The MIB objects selected in the table are read by the polling table and written to the database where they are retained for future use. Subsequent data polls do not overwrite previous polled data. Data is stacked with all collected information available for analysis.

To browse poll data in a text or graph format, select the data segment and use the **Browse Data** command.

This command sends the data segment to the Results Browser application.

3 While browsing the data segment, you can view text or convert the data to a graph using the SNM Results Browser. Manual intervention is required to move the polling data to the HP OpenView or NetView for AIX grapher.

Note   You can only use the Browse Data command, the Results Browser application, and the SNM Results Browser on SunNet Manager platforms.

---

4   Use the Polling Summary Tools menu to export polling data or create reports.

Figure 4-1 illustrates the device polling concept. The polling daemon CiscoWorks uses, *nmpolld*, queries and stores information in the Sybase database. Both the Device Polling and Polling Summary applications use the polling daemon.
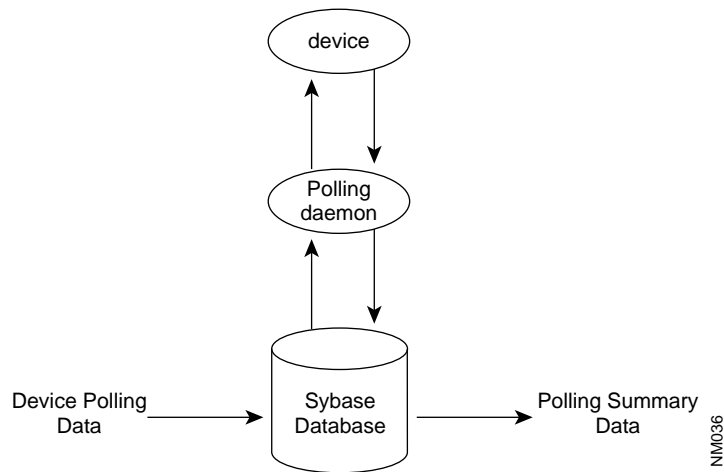

Figure 4-1   Device Polling and Sybase Database Concept

Device Polling Tables

A Device Polling table consists of the following elements: a polling frequency, a set of MIB objects, and a list of devices. The information specified in the table is collected during the polling process. A collection of information from the polling process is called a *data segment*. A segment is defined as the data selected from the start of polling to the stop of polling, regardless of the frequency used to sample. For instance, if you start polling at 1:00 p.m. and stop at 2:00 p.m. on the same day, the data segment will contain one hour of polling data. If you had set a polling frequency of 1 minute, you would have obtained 60 samples in the segment.

Figure 4-2 illustrates the polling table construction concept. You can select MIB objects to collect information on a device and choose a polling interval for the poll. The polling interval has a significant impact on the amount of data collected in a given time frame.
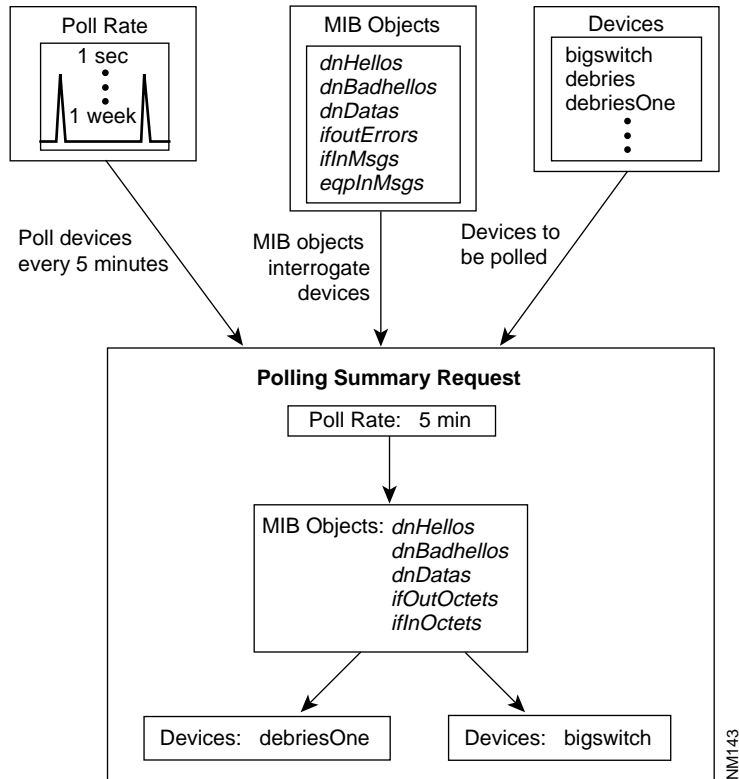
```
  ┌─────────────────┐      ┌─────────────────┐      ┌─────────────────┐
  │   Poll Rate     │      │   MIB Objects   │      │    Devices      │
  │   ┌─────────┐   │      │  ┌───────────┐  │      │ ┌─────────────┐ │
  │   │ 1 sec   │   │      │  │ dnHellos  │  │      │ │ bigswitch   │ │
  │   │    •    │   │      │  │ dnBadhellos│ │      │ │ debries     │ │
  │   │    •    │   │      │  │ dnDatas   │  │      │ │ debriesOne  │ │
  │   │ 1 week  │   │      │  │ ifoutErrors│ │      │ │      •      │ │
  │   └─────────┘   │      │  │ ifInMsgs  │  │      │ │      •      │ │
  └─────────────────┘      │  │ eqpInMsgs │  │      │ └─────────────┘ │
                           │  └───────────┘  │      └─────────────────┘
                           └─────────────────┘
```

Poll devices every 5 minutes   MIB objects interrogate devices   Devices to be polled

Polling Summary Request

Poll Rate: 5 min

MIB Objects: *dnHellos*
              dnBadhellos
              dnDatas
              ifOutOctets
              ifInOctets

Devices: debriesOne   Devices: bigswitch

NM143

Figure 4-2 Polling Table Construction Concept

Note   The amount of data collected varies depending upon the type and number of queries made, so be sure to monitor the disk space available for storage of polled information. For additional information about database and transaction log space, refer to Chapter 8, "Database Administration."

Running Multiple Pollers

The Polling daemon allows you to run multiple daemons. You can run any number of processes on a single system, or you can run processes on several systems.

The poller ID assigned to a polling group defines which poller process will perform the polling for that poll group. Refer to the manual page for the specific rules.

The Polling daemon on separate machines can be started or stopped by running the Process Manager on the various machines. For example, you can run the Process Manager on one workstation to start and stop a single Polling daemon and at the same time run the Process Manager on another workstation to start and stop a single Polling daemon on that machine. The Process Manager does not assist in starting or stopping more than one Polling daemon on a single machine.

Following are two caveats of starting and stopping multiple Polling daemons:

- Starting—The Process Manager will start only one Polling daemon. To start more than one Polling daemon, you must do it manually, or you can add commands to the *$NMSROOT/etc/nmstartup* file so that the Polling daemons can be started up every time the workstation comes up. CiscoWorks has already put one line in *nmstartup* to start and stop the Polling daemon. You can copy and paste that line into the *$NMSROOT/etc/nmstartup* file in order to start a second Polling daemon.

- Stopping—You can stop the Polling daemons from the command line by issuing a **kill -TERM** command. You also can use the Process Manager, but we recommend that you use the command line method.

Note Even though you instruct the Process Manager to stop the Polling daemon, the Process Manager may still indicate that the Polling daemon is running. This does not indicate that either the Process Manager or the Polling daemon is working improperly. It may be due to multiple daemons.

Device Polling Window

Before you can begin polling devices for data, you must set up your polling table in the Device Polling window.

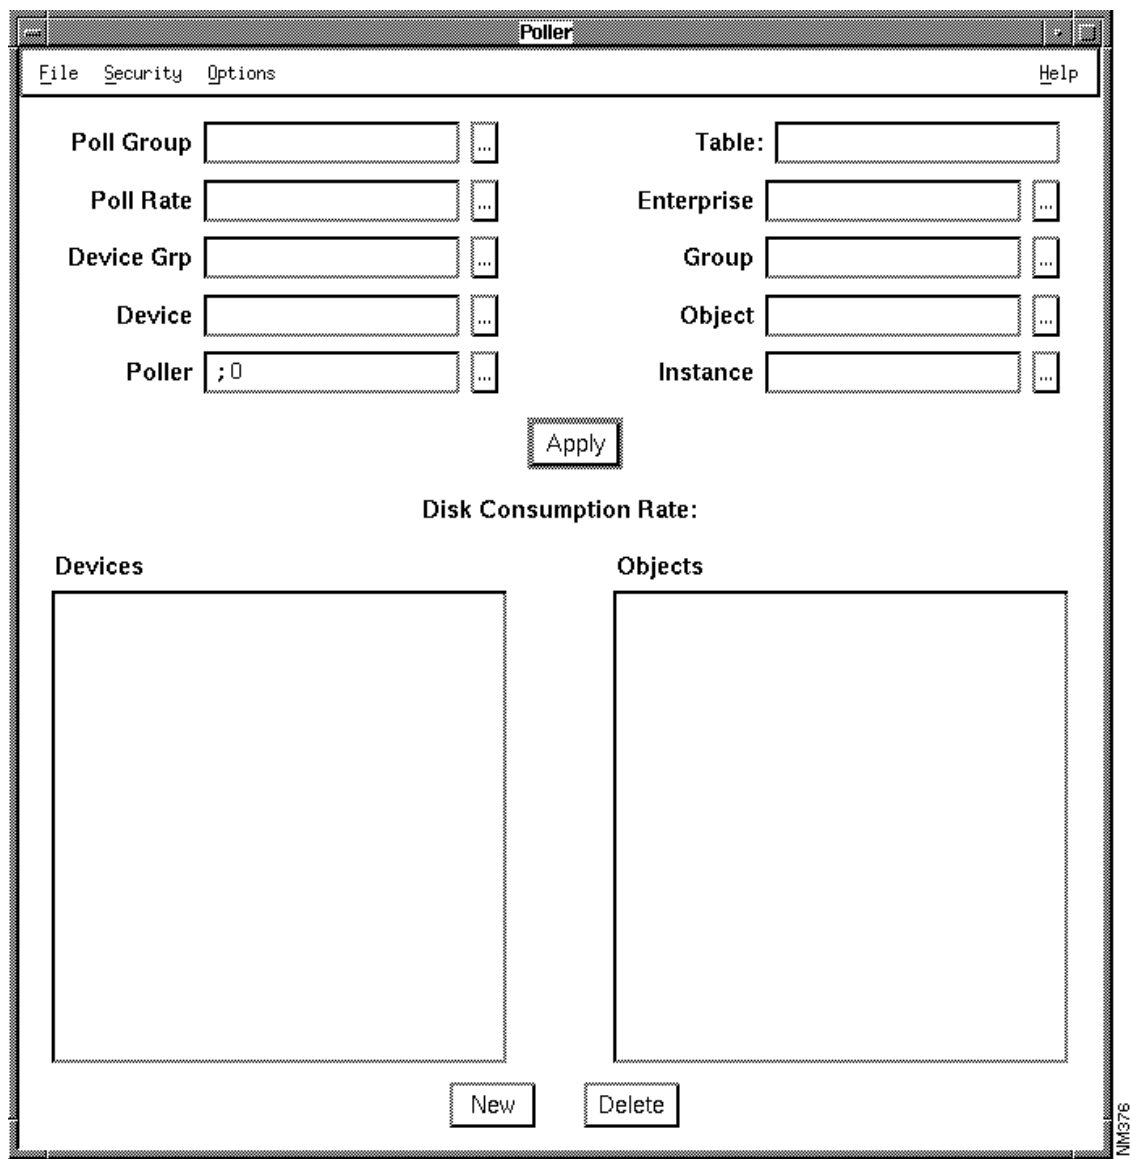Figure 4-3 illustrates the Device Polling window.

Figure 4-3 Device Polling Window

Table 4-1 describes the components of the Device Polling window.

Table 4-1   Device Polling Window Components

Component	Subcomponent	Description
File	Print	Prints a snapshot of the window.
	Exit	Exits the current window.
Security	Change Domain	Changes the domain.
	Change User	Changes the user ID to another user.
	Privilege	Views the current user ID privileges.
Options	Activate Changes	Sends a signal to the Polling daemon to reread the database so that current changes are included in the poll.
	Destroy Table	Deletes the selected polling table.
Help	On Version	Displays the CiscoWorks application version information.
	On Device Polling	Displays information about the current window.
Poll Group	sample sample_load sample_traffic sample_mix sample_error	Describes the table poll group name. For example, CiscoWorks supplies several example tables named *sample**.
Poll Rate	Default = No polling	Sets the polling interval for the last table name selected. Polled data segments are date and timestamped with start poll times.
Device Grp	All Devices	Lists the devices you are choosing from.
	Device A—Device ?	
	Device ?—Device Z	
Device		Lists all devices currently in the device group.
Poller	;0	Selects which systems will actually perform the current poll. You can have multiple systems acting as polling engines.
	others	
	Default = ;0	The first poller started on every machine.
Table	Default is the same as the poll group name	Describes the Sybase database table name of the current poll. This name can be the same as the poll group name.
Enterprise	snmp-mibII cisco hp	Lists the enterprise MIB objects.
Group	(second level)	Lists of related MIB groups from the specified enterprise group above.
Object	(third level)	Lists of MIB object names (also referred to by Cisco as *variables*) from the specified MIB group above. Objects applied here define objects available in the group.
Instance	Pick menu for if_tables instances	Describes the list of keys supplied by the MIB object tables to poll. Only lists interfaces if you use the if_tables. Otherwise, the pick menu will be empty, and you must supply your own key for the MIB object selected. If no interface is chosen, all interfaces (keys) will be polled. Only objects in a table need or can have an instance.
Apply		Applies current changes for this polling table to the database. This does not notify the polling daemon. Use the **Activate Changes** command to complete the polling setup.

Component	Subcomponent	Description
Disk Consumption Rate		Describes the disk space growth rate (in kilobytes per day) to hold the data for this type of poll.
New		Erases all information currently in the window. The form remains, ready to accept new information. This command does not affect data in the database.
Delete		Erases the devices selected from the database.

MIB Object Requirements

The Object field lists the MIB objects you can select to poll. MIB objects define the type of information about devices and interfaces that will be collected with the poll.

When selecting objects, remember the following requirements:

- For protocol-specific MIB objects, for example, *xnsInput*, you must have the corresponding protocol enabled in the device. Otherwise, no data will be stored for the MIB objects.

- For SunNet Manager platforms only:
  If you have modified SNM schemas, you should check that the variables are still in the schemas.

Disk Space Consumption Rate

As each object is selected, the Disk Consumption Rate field is adjusted to reflect an estimate of how fast the disk space will be consumed in order to store the polled data. The algorithm uses the poll rate, number of devices, and number of objects currently selected in the table to estimate the space required to hold one day (24 hours) of polling. The result of the calculation is reflected as whole kilobytes required per day.

Multiply the disk space consumption rate by the length of time you plan to poll. This equals the number you compare with the available disk space. If you do not want to change the polling table, you might have to increase your database space allocation. Before deciding to adjust disk space, review the polling interval and duration and the number of objects to determine if any can be adjusted. To increase the database space, refer to Chapter 8, "Database Administration," for more information.

Creating a New Polling Table

In order to set up polling, access the Device Polling window and complete the necessary information. If you are not familiar with MIB objects, using Device Polling will be more difficult. Contact your network administrator for more information on MIB objects.

To access the Device Polling window to set up a new polling table, perform the following steps:

Step 1 Select **Device Polling**.

On SNM, select **Tools>Device Polling**.

On HP OpenView or NetView for AIX, select **Monitor>Device Polling**.

The Device Polling window appears. (See Figure 4-3.)

For a description of the device polling fields, refer to Table 4-1.

Step 2   Create a new poll group by clicking on **New**.

The **New** button clears the poll group data. You must press the **New** button before you enter new information.

If you want to use one of the *sample* tables as a template, open the Poll Group pick menu and choose a *sample* table.

Step 3   Create the new poll group name by entering the name in the Poll Group field.

Note   CiscoWorks is based on the Sybase SQL-based relational database. Poll group or table names can include only the alphanumeric characters A through Z, 0 through 9, and underscore (_). A dot (.) is not allowed. For more information on naming tables, refer to your Sybase documentation.

Only one poll group is used with each data configuration. The Poll Group field defines what tables are in the poll groups.

CiscoWorks includes several example poll group tables in the pick menu called *sample**. You can use these tables as templates. Modify them to your specific polling needs, but rename the poll group to describe the type of polling you are performing.

Step 4   To use a sample template to create a new poll table, select the sample table and rename it by entering a new group and table name and click on **Apply**.

If you overwrite or delete one of the *sample** tables, you can recreate it using the **makesample** program. The **makesample** script adds the *sample** files back to your directory structure. Refer to Chapter 8, "Database Administration," for information on the **makesample** script.

Step 5   Create the new table name or use the poll group name as the table name by default. To create a new table name, enter the name in the Table field.

Your poll group name and table name will default to the same name. You may want to change the table name.

Caution   You can only enter a table name prior to pressing the **Apply** button. After you press **Apply**, the table is created, and the table name cannot change.

Step 6   Select the poll rate for this poll. There are two methods of selecting a poll rate:

- Select a default value from the Poll Rate pick menu.

- Enter a value directly in the Poll Rate field of the Device Polling window and press Return. The poll rate recalculates after you press **Apply**.

Use seconds, minutes, days, or weeks as the unit of measure. If you enter 21 days in the field, the data will be converted to 3 weeks, the largest unit of measure for the integer interval entered.

Caution   Be careful when selecting a poll interval because you need to watch your disk space use.

Step 7   Select the Poller field and pull down the pick menu to one of the poller options.

The default is ;0. If you use the ;0 default, the first poller will start working on every machine. For more information on running more than one polling daemon, refer to the section "Running Multiple Pollers" earlier in this chapter.

Step 8   Select the Enterprise field and select one of the enterprise MIB objects from the pick menu.

For example, if you are using the *sample* table, pull down to *cisco*.

Step 9   Select the Group field and pull down the pick menu to one of the MIB groups. Select a MIB object.

For example, if you are using the *sample* table, pull down to *ifTraffic*.

Step 10   Select the Object field and pull down the pick menu to select one of the MIB objects you want to include in this poll group.

For example, if you are using the *sample* table, pull down to *ifInOctets*.

The Object field lists all the MIB objects that are available. To select multiple objects for the object poll group, click on **Apply** and then select another object.

The objects selected apply to the last table name entered. If you want an object to apply to a different table name, select the table first.

Step 11   If you have fewer than 60 devices in your device group, continue to the next step.

If you have more than 60 devices in your device group, select the Device Group field and open the pick menu to select one of the listed device group options you want to include in this poll.

By selecting a specific device group, your Device field list becomes more defined. For example, if you choose the group of devices from X to Z, your device list will include only those devices beginning with the letters X, Y, and Z.

Step 12   Select the Device field and use the pick menu to select one of the devices you want to include in this poll.

This list contains all devices in alphabetical order that are present in the CiscoWorks database. To select multiple devices for this poll group, click on **Apply** and select another device.

Step 13   After you select the devices from the Device field, select the Instance field and open the pick menu to select an instance associated with the device.

If you are using tables, for example *if_tables*, select the instance from the pick menu.

For Sun SPARCstations only:
If you are not using *if_tables*, you must know the format of Sun MIB object tables found in the schema file to enter the instance values. If you do not know the format, skip the Instance field. This will cause the poll group to poll all instances. Only objects in a table need or can have an instance. You can skip this field if your object does not belong to an object table.

Step 14   When you click on **Apply**, all data is sent to the database. The scroller windows display what appears in the database.

Step 15   Select **Options>Activate Changes** to send the changes to the Polling daemon.

If you have chosen a poll rate, the **Activate Changes** command informs the Polling daemon (*nmpolld*) to start polling this poll table.

Caution **Activate Changes** occurs automatically when you exit Device Polling using the **File>Exit** command.

This completes the process of creating a new polling table.

Adding Instance Data to an Existing Device and Object Pair

An instance defines which row to go to in a particular MIB table. Use the default instance selection in Device Polling to poll all objects on all devices. If you want to select specific devices and MIB objects to limit your poll group, add an instance to a device and object pair.

If you want to add an instance to a device and object pair, add the instance in the Device Polling window.

To add an instance to a device and object pair, perform the following steps:

Step 1 Specify the device and object pair. To select a device and object pair, select the device name and then the object in the scroll windows.

Step 2 For *if_tables*, select an instance using the pick menu.

For all other keys, enter the key manually. If you do not know your key, contact your network administrator.

Step 3 Click on **Apply**.

The new table with the instance added is now available for polling. The device will now be polled for this specific object and specific instance. An *I* marks the device and object pair you selected for polling. An asterisk (*) marks the polled instance in the instance list.

Step 4 Select **Options>Activate Changes** to send the changes to the polling daemon.

Step 5 Repeat steps 1 through 4 for each instance you want to add to a device and object pair.

Note When troubleshooting the Device Polling daemon, *nmpolld*, and the dataserver is up, access the Log Manager centralized log. Check the log messages for a message that *nmpolld* was reinitialized after Sybase restarted.

Modifying an Existing Poll Group Table

Device polling has the following limitations:

- You cannot delete an object that already exists.
- You cannot change an existing poll group or table name.

To delete an object, delete the poll group and start again.

To modify an existing poll group table effectively, perform the following steps:

Step 1 From the Poll Group pick menu, select the poll group name.

When an existing table name is selected, the configuration displays in the window. If you selected the table to change a data element, such as the polling interval, perform the necessary changes.

Step 2   Save your changes to the database by clicking on **Apply**.

Step 3   Select **Options>Activate Changes** to notify the Polling daemon (*nmpolld*) of changes.

   If you do not perform this step, any changes you have made are unknown to *nmpolld* and cannot be polled for.

Step 4   To switch from one table to another in the Device Polling window, pull down the Poll Group selection and select a new table.

Deleting a Poll Group Table

After you finish with a poll group and want to remove it from the database, delete it from the Device Polling window.

To delete a poll group, perform the following steps:

Step 1   To select the poll group you want to delete, select the poll group name from the pick menu.

Step 2   Select **Edit>Destroy Table**.

   A confirmation window appears asking you to confirm the deletion. Answer accordingly.

Step 3   Select **Options>Activate Changes** to notify the Polling daemon (*nmpolld*) of changes.

The Poll Group is now deleted from the Polling Summary table and the database.

Deleting a Device from a Poll Group

To delete a particular device from a poll group, delete the device from the Device Polling window.

To delete a device from a poll group, perform the following steps:

Step 1   To select the poll group you want to change, select the poll group name from the pick menu.

Step 2   From the Devices scroller, select a device.

Step 3   Click on **Delete**.

   The new table with the device removed is now available for polling.

Step 4   Select **Options>Activate Changes** to notify the Polling daemon (*nmpolld*) of changes.

The device is now deleted from the poll group.

Deleting Instance Data from a Poll Group Table

If you want to delete a particular instance from a poll group, delete the instance from the Device Polling window.

To delete an instance from a poll group, perform the following steps:

Step 1   To select a device and object pair, from the scroll windows, select the device name and then the object.

Step 2   Select the instance you want to delete from the pick menu in the Instance field.

Step 3   Click on the **Delete**.

   The new table with the instance removed is now available for polling.

Step 4   Repeat steps 1 through 3 for each instance you want to delete from the device polling group.

Using Polling Summary
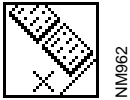
The Polling Summary application enables you to perform the following tasks:

- Start and stop polling

- Examine polling group data

- Delete polling data or tables

- Export poll data to flat files

- Run a variety of reports from the window or from the command line

These tasks are described in the following sections.

Polling Summary Window

The Polling Summary window is the interface used to stop, start, browse and export poll data, and run reports.

Figure 4-4 illustrates the Polling Summary window. Table 4-2 describes its components.
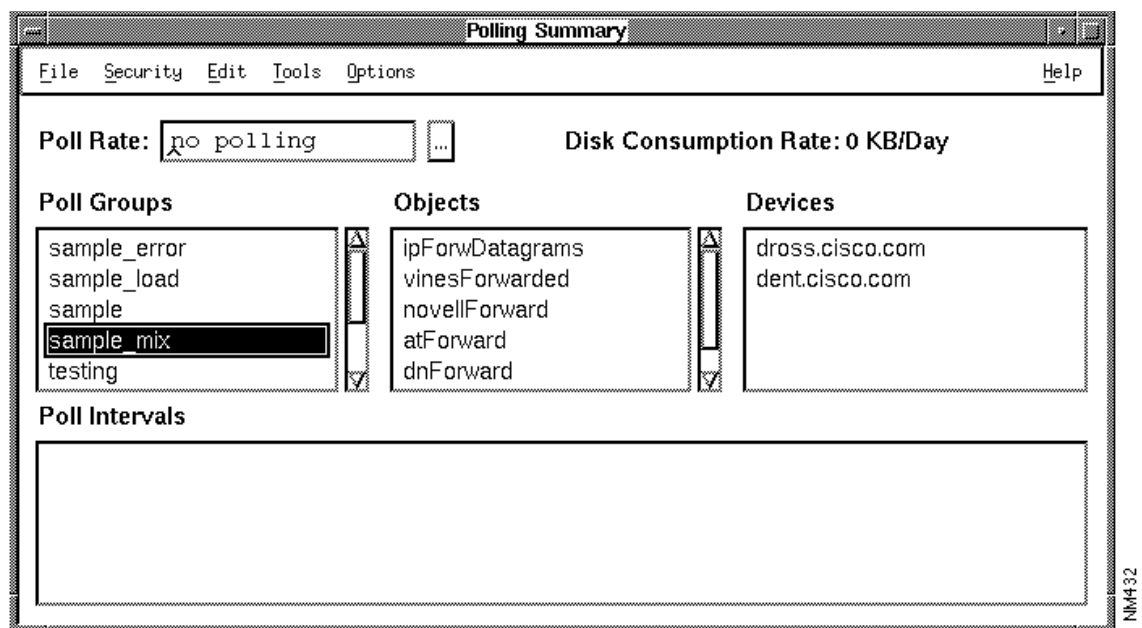

Figure 4-4   Polling Summary Window

Table 4-2   Polling Summary Window Components

Component	Subcomponent	Description
File	Print	Prints a snapshot of the window.
	Exit	Exits the current window.
Security	Change Domain	Changes your domain.
	Change User	Changes your user ID to another user.
	Privileges	Views the current user ID privileges.
Edit	Delete Poll Interval	Deletes the selected poll interval.
	Delete All Data	Deletes all data in the selected poll group.
Tools		Allows you to run reports and export poll data to a file. You can also add custom commands. Menu items include:
		Export Data
		Configuration Summary Report
		Device Inventory Report
		Errors by Device Report
		Device Load Report
		Top/Bottom 10 IF Traffic Report (Gross)
		Top/Bottom 10 IF Traffic Report (Normalized)
		Traffic by Interface Report
		Traffic Mix Report
		Help for Reports
		Schedule Polling
Options	Activate Changes	Signals the device polling daemon to reread the database to get polling changes.
	Browse Data (for SNM only)	Uses SNM Browser to view selected data reports or to graph data using Grapher.
	Destroy Table	Deletes the selected polling group.
Help	On Version	Displays the CiscoWorks version information for this application.
	On Polling Summary	Displays a manual page on the current window.
Poll Rate	no polling 1, 2, 5, 10, 30 seconds 1, 2, 5, 10, 30 minutes 1, 2, 5, 12 hours 1, 2, 14 days 1 week	Starts and stops the polling process. Sets the polling time interval for the last table name selected. Any change is reflected in the database.
Disk Consumption Rate		Describes the disk growth rate required (in kilobytes per day) to hold the data for this type of poll.
Poll Groups scroller		Lists the poll groups available with the Device Polling window. These are also known as *tables*.
Objects scroller		Lists MIB objects defined in Device Polling.
Devices scroller		Lists all devices in the poll group defined in Device Polling. You may need to use Sync w/Sybase to locate new devices.
Poll Intervals scroller		Lists data reports that contain poll report data. Includes the start time and length of poll interval for the data report.

Starting or Stopping Your Poll Tables

After you set your polling criteria in the Device Polling window, you can access the Polling Summary application to turn polling off, start additional polling tables, view poll data, and so on.

To start or stop polling your poll group table, perform the following steps:

Step 1   Select **Polling Summary**.

On SNM, select **Tools>Polling Summary**.

On HP OpenView or NetView for AIX, select **Monitor>Polling Summary**.

The Polling Summary window appears. (See Figure 4-4.)

Step 2   From the Poll Groups scroll window, select a group you want to poll.

Step 3   If the table is not being polled currently, select a polling frequency. Then select **Options>Activate Changes**.

While the Polling daemon (*nmpolld*) starts polling your devices in the background based on your entries in the Device Polling window, you can continue to work with the Device Polling window.

If your table is currently being polled, continue to the next step.

Step 4   After you collect the necessary data, stop the polling activity on this group from the poll rate pick menu by selecting **No Polling**.

Step 5   To signal the polling daemon to stop polling, select **Options>Activate Changes** to notify the Polling daemon (*nmpolld*) of changes or exit the Polling Summary window.

> ⚠ **Caution** If you do not stop a polling group process, it continues to run indefinitely and decreases your available disk space. If your disk fills up, you may loose useful information. Remember to use the **No Polling** command from the Poll Rate pick menu and the **Options>Activate Changes** command.

Examining Polling Data

After you complete your polling activity, you can examine your data in text, graph, or report form. Data is retrieved from the polled information segments obtained in polling operations. To view polling data, you use the SNM Results Browser and Grapher on SunNet Manager or the graphers on HP OpenView or NetView for AIX. To view your polling data on the HP or NetView for AIX graphers, you will need to export the data to a spreadsheet and then import it into the grapher.

You can also run reports using the Polling Summary Tools menu. For more information on the automated SQL reports that are shipped with Polling Summary, refer to "Viewing Polling Data Using the Polling Summary Tools" later in this section.

Viewing Polling Data Using the Browser or Grapher (SunNet Manager only)

With the SunNet Manager (SNM) Results Browser and Grapher features, you can review your data report information. You can browse text information using the Results Browser. From the Results Browser, you can send data to the SNM Grapher.

For example, you can compare input broadcast packets against the total number of input packets to identify broadcast storms, as well as review and analyze interface performance or traffic patterns by observing in and out packets over specific periods.

Refer to your SNM documentation for more information on using the Results Browser and Grapher features.

To view your polling data using the browser or grapher, perform the following steps:

Step 1 Select a poll group.

Selecting a poll group updates the devices and interfaces in the Devices scroll window that will be available and, in turn, data segments that can be viewed or graphed. A graph can consist of multiple segment columns and multiple devices.

For example, CiscoWorks ships several *sample* tables with the Device Polling application. Click on the *sample* table.

Note Group selection triggers access to data contained in Devices and Variables browsers. Prior to group selection, any attempts to obtain information from these browsers are ignored. If a group name you want is not present, it has not been created, has been created but was never active, or has not completed a polling segment. Use the Device Polling window in the Device Polling application to verify the existence of the group.

Step 2 To modify the poll rate interval, select your desired poll rate from the pick menu, or enter the poll rate in the field and press Return.

Step 3 To view the data collected in your poll, from the Polling Summary window select the data report in the Polling Intervals scroller.

Step 4 Select **Options>Browse Data**.

The SNM Results Browser window appears with the data report (also known as the *data stream* in SNM) listed in the scroller. For more information on the Results Browser, refer to the *SunNet Manager 2.0 Reference Guide*.

Step 5 To view the text of the data report, click on the report name.

The data report text appears in the lower portion of the results browser.

Step 6 To view the data report in a graph format, click on the report name and from the popup menu select **Graph**. Continue to pull the **Graph** menu to the right and pull down to the specific graphing options you desire.

For this exercise, select **All Attributes**. An SNM Grapher window appears with the data report name in the scroller. Seconds later, a graph appears.

Note Although multiple graphs can be viewed simultaneously, initially only one poll group can be used with one Grapher window. For more detailed window information, such graphing table data or altering graph properties, refer to the *SunNet Manager 2.0 Reference Guide*.

Step 7 To change the properties on your graph, return to the Results Grapher window, select the report name, and click on **Properties**. (See Figure 4-5.)

For more detailed information on customizing the SNM Graph Properties window, refer to your *SunNet Manager 2.0 Reference Guide*.
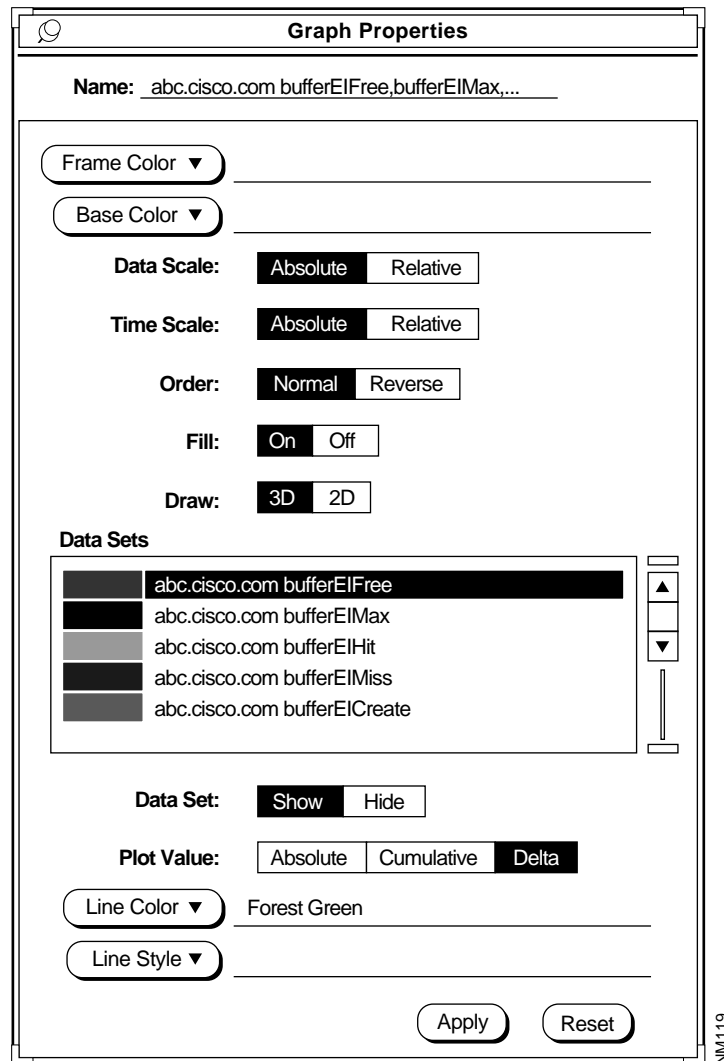

Figure 4-5 SNM Graph Properties Window

Viewing Polling Data Using the HP OpenView or NetView for AIX Graphers

The Device Polling application allows you to probe and extract data about the condition of devices on your network. In general, device polling stores data in polling tables. You can then graph the data with your platform grapher.

To view polling data using your platform grapher, complete the following steps:

Step 1 Export your polling data using Polling Summary.

For information on exporting polling data, refer to the section "Exporting Polling Data to Flat Files," later in this section.

Step 2 Move the ASCII exported data into a third-party spreadsheet, then import the information into the */usr/OV/databases/snmpCollect* directory.

Refer to your HP OpenView or NetView for AIX documentation on how to move ASCII data into this directory.

Step 3 With your network map open, select **Monitor>MIB Values** from the menu.

Step 4 Select **Graph Collected Data**>**SNMP**>**All**.

Deleting Polling Data or Tables

During polling, CiscoWorks stores any poll data in segments we refer to as *polling intervals.* If you use polling extensively, you may need to delete polling data as a way to maintain disk space on your workstation. Instructions on how to delete selected polling intervals or entire sets of polling data follows.

The following caution applies when using CiscoWorks on any platform.

Caution When you delete a polling table, be sure to follow the instructions given in this guide. Do not enter a DROP TABLE statement from SQL or ISQL, because using the statement has a different effect than using the Device Polling and Polling Summary applications.

Deleting Selected Polling Data

In order to delete polling data (or intervals), perform the following steps:

Step 1 From the Polling Summary window, select a poll group.

Selecting a poll group updates the objects, devices, and polling intervals that are available with the selected poll group.

Step 2 Select a poll interval to delete by clicking on the data segment in the Polling Intervals scroller.

If you do not select a poll interval, you will receive an error message when you attempt to delete the poll interval.

Step 3 Select **Edit>Delete Poll Interval**.

A confirmation box displays asking you to confirm the deletion.

Step 4 To confirm the deletion of the selected poll interval, click on **OK**. To cancel the deletion, click on **Cancel.**

If you deleted the poll interval, the Polling Summary window no longer displays the data segment and removes the information from the database.

Deleting All Polling Data

In order to delete all polling data (or intervals) associated with a poll group, perform the following steps:

Step 1 From the Polling Summary window, select a poll group.

Selecting a poll group updates the objects, devices, and polling intervals that are available with the selected poll group.

Step 2 Select **Edit>Delete All Data**.

A confirmation box displays asking you to confirm the deletion.

Step 3   To confirm the deletion of the selected poll interval, click on **OK**. To cancel the deletion, click on **Cancel.**

If you deleted all polling data, the Polling Summary window no longer displays the data segment and removes the information from the database.

Note   You can also use the command line interface by running the *nmpolldelete* utility. The *nmpolldelete* utility helps with deletion of old polling data without destroying the poll table. This utility is located in the *$NMSROOT/etc* directory and includes a manual page.

Viewing Polling Data Using the Polling Summary Tools

The Tools menu in the Polling Summary application contains the several utilities including an export tool and automated SQL reports. The *reports.nmstool* file in the *$NMSROOT/lib* directory contains the details of each report.

If you have an overabundance of polling data, use the *$NMSROOT/etc/nmpollsummarize* utility. The *nmpollsummarize* utility creates a summary table which summarizes data in a polling interval into averages. A poll table using a 10-second poll rate can be summarized into hourly records for a more compact file that saves on storage space. For more information on *nmpollsummarize*, refer to the CiscoWorks online manual page.

⚠ **Caution**   Do not edit the *reports.nmstool* file. If you need to alter this file, copy it first and edit your custom file.

These automated reports help you display a variety of information including:

- Configuration Summary Report—Provides the configuration version each device is running.

- Device Inventory Report—Provides general device data by domain.

- Errors by Device Report—Provides the total number of errors and discards for each system for each device and interface.

- Device Load Report—Provides information on CPU load, discards, queues, buffers, and so on.

- Top/Bottom 10 IF Traffic Report (Gross)—Provides information on the interfaces with the most and least traffic levels.

- Top/Bottom 10 IF Traffic Report (Normalized)—Provides information on the interfaces with the most and least traffic levels based on the time the interfaces were up.

- Traffic by Interface Report—Provides information on the number and percentage of bytes sent and received per interface.

- Traffic Mix Report—Provides the number and percentages of packets forwarded by a device per protocol.

- Help for Reports—Provides detailed information on each report including a description, the SNMP variables used in the report, output format, report sort order, command syntax, and column descriptions.

- Schedule Polling—Allows you to start or stop polling from the command line or cron file. Creates a batch program to run with the Scheduler utility. Remember to create two cron files in the Schedule window; one file to start the poll and one to stop the poll.

Note The automated reports/utilities can be run from the Polling Summary window or from the command line using *$NMSROOT/bin/nmreport*. For command line syntax, refer to CiscoWorks software online help.

Exporting Polling Data to Flat Files

If you want to display polling data in some other format, such as a spreadsheet, the Polling Summary application offers an option that allows you to export polling data into an ASCII flat file. You can then import the ASCII file into other applications, such as a spreadsheet application, that displays your data in a more tabular format.

Note You can also use the command line interface by running the *nmpollexport* utility. The *nmpollexport* utility dumps polling data to a file that can be imported to a spreadsheet. This utility is located in the *$NMSROOT/etc* directory and includes a manual page.

To export polling data to an ASCII file, perform the following steps:

Step 1 From the Polling Summary window, select a poll group.

Selecting a poll group updates the objects, devices, and polling intervals that are available with the selected poll group.

Step 2 Select a poll interval to export by clicking on the data segment in the Polling Intervals scroller.

If you do not select a poll interval, you will receive an error message when you attempt to export the poll interval data.

Step 3 Select **Tools>Export Data**.

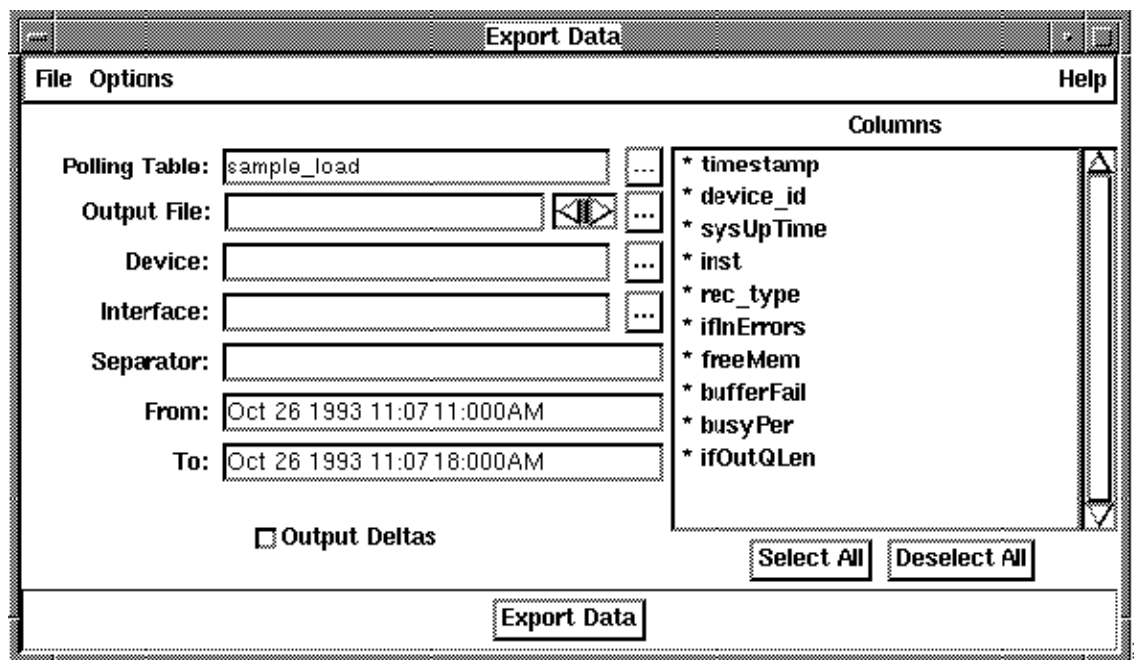The Export Data window displays. (See Figure 4-6.)

Figure 4-6 Export Data Window

Step 4 To toggle between exporting polling or summary tables, select **Options>Export Poll Table** or **Options>Export Summary Table**.

The Export Data window displays again with the appropriate field change.

Step 5 To toggle between exporting data to a UNIX file or a Sybase table, select **Options>To UNIX File** or **Options>To Sybase Table**.

Step 6 Enter data into the fields required.

For information on the data field, refer to Table 4-3.

Step 7 Select the type of output you want to export, absolute or deltas, by clicking on the **Output Deltas** toggle.

Step 8 To deselect columns that you do not want to be included in the output file, select an entry in the Columns scroll window.

To deselect all columns, click on **Deselect All.**

Step 9 To return selected columns to the output file, select the column names or to return all columns, click on **Select All**.

Step 10 Click on **Export Data** to export the data to a file.

The data from the Sybase database tables is exported to the file you designated. You can use this file to import information into other applications, for example, a spreadsheet application.

Table 4-3 describes the components of the Export Data window.

Table 4-3   Export Data Window Components

Component	Subcomponent	Description
File	Print	Prints a snapshot of the window.
	Exit	Exits the current window.
Options	Toggle Buttons:	Allows you to select which export method to use:
	Export Poll Table	• Export a poll table to a UNIX file.
	Export Summary Table	• Export a poll table to another Sybase table.
	To UNIX File	• Export a summary table to a UNIX file.
	To Sybase File	• Export a summary table to a Sybase table.
		The summary tables are created using the **nmpollsummarize** command.
Table Pick Menu	Polling or Summary	Contains the polling or summary table name which contains the data. Allows you to change to a poll or summary table from the pick menu.
Output File	Pick Menu	Allows you to name the file which will contain the exported data.
Device	Pick Menu	Restricts output to a specific device in the poll group. Enter the name of the device without the domain name or click on the pick menu.
Interface	Pick Menu	Restricts output to a specific interface of a device. To use this field you must also name a device in the Device field. Enter the interface index (*ifIndex*) value of the interface or click on the pick menu.
Separator		Allows you to enter any character to replace the default Tab character that separates fields.
From		Restricts output to be later than a specified date/time. Enter the date and time in the following format: Jan 1 1994 1:00AM
To		Restricts output to be earlier than a specified date/time. Enter the date and time in the following format: Jan 1 1994 1:00AM
Output Deltas toggle button		Causes counter and timetick values to be exported as delta values rather than absolute values.
Columns list		Displays the columns in the poll table named in the Table field. An asterisk (*) to the left of the column name indicates that this column will be included in the output file. Clicking on an entry in the Columns listbox toggles asterisk on and off.
Select All		Returns all asterisks in the Columns list to on.
Deselect All		Turns all asterisks in the Columns list off. Turning the asterisks off prevents the column from being included in the output file.
Export Data		Turns all asterisks in the Columns list back on. Turning the asterisks off prevents the column from being included in the output file.

Using the Real-Time Graphs

When you use the Real-Time Graphs application to manage the performance of your network, you can observe real-time information via a two- or three-dimensional graph (depending on your platform grapher). CiscoWorks enables you to graph data about router health, interface, and traffic information.

Note   You can use the Real-Time Graphs application from the menus on your NMS platform. However, the grapher windows will look different on the different platforms, as illustrated later in this chapter.

Real-Time Graphs Window

Figure 4-7 illustrates the Real-Time Graphs window.
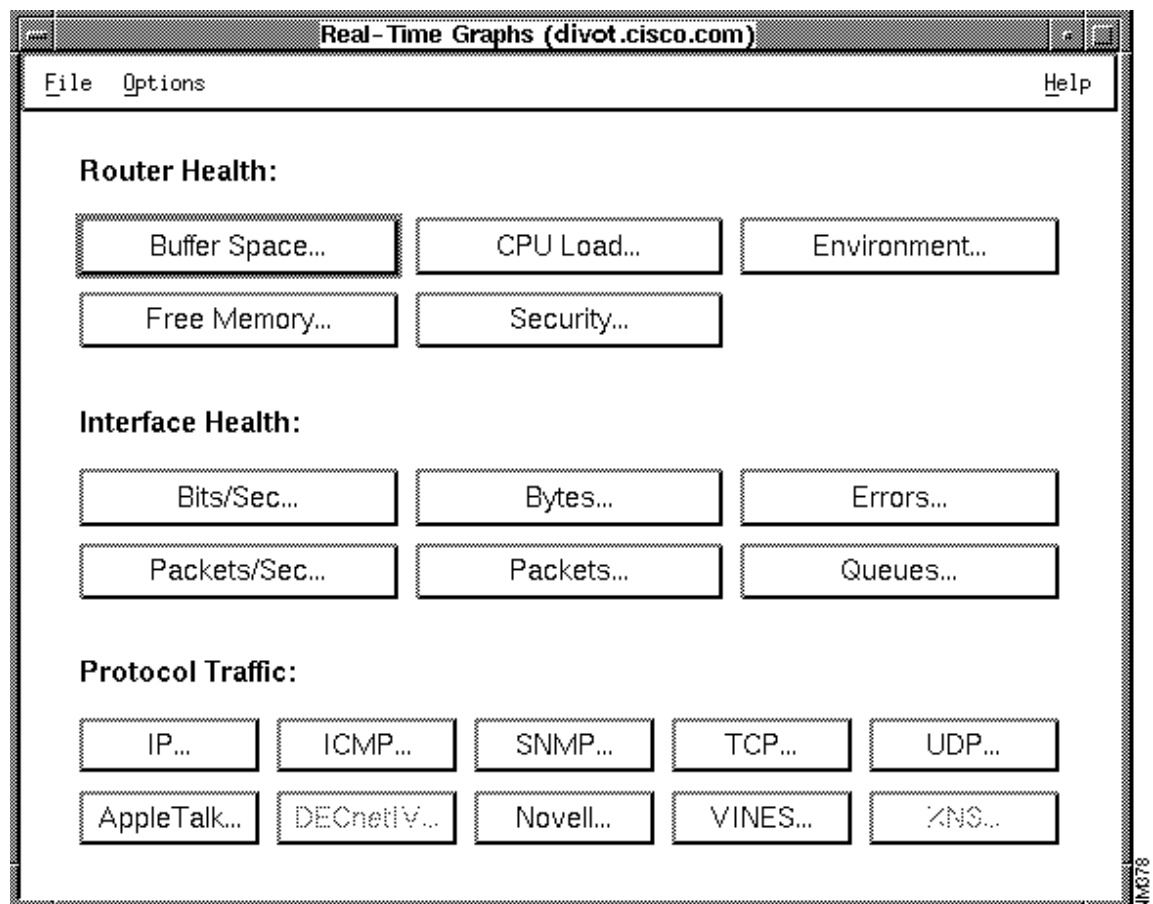

Figure 4-7   Real-Time Graphs Window

Table 4-4 describes the components of the Real-Time Graphs window.

Table 4-4   Real-Time Graphs Window Components

Component	Subcomponent	Description
File	Print	Prints a snapshot of the window.
	Exit	Exits the current window.
Option	Set Polling Frequency	Changes polling rate. Frequency of analysis performance is represented in seconds. Can be set using slider or entering in the polling interval field. Default = 2 seconds.
Help	On Version	Displays the CiscoWorks application version information.
	On Real-Time Graphs	Displays a manual page on the current window.
Router Health		Refer to Table 4-5 for detailed descriptions of the router health buttons.
Interface Health		Refer to Table 4-6 for detailed descriptions of the interface health buttons.
Protocol Traffic[1]		Refer to Table 4-7 for detailed descriptions of the protocol traffic buttons.

1. If a button is grayed out, the selected device does not have this capability. For example, currently only Cisco AGS+ and Cisco 7000 with Microcode Version 2.0 have the Environment router health button capability.

Graphing Your Real-Time Device Data

The Real-Time Graphs application observes the behavior of devices suspected of being in degraded mode or introducing erratic behavior in traffic patterns, error status indications, or statistics.

Note   On SNM, the Real-Time Graphs application monitors and graphs a single device. You can monitor multiple devices simultaneously by opening more than one Grapher. In addition, you can merge graphs to present the data in one graph. The Real-Time Graphs application uses the SNM graphing utility, so for information on customizing your graph, see the *SunNet Manager 2.0 User's Guide*.

To create a graph with real-time device data, perform the following steps:

Step 1   Click on the device.

Step 2   Select **Real-Time Graphs**.

On SunNet Manager, select **Tools>Real-Time Graphs**.

On HP OpenView or NetView for AIX, select **Monitor>Real-Time Graphs**.

The Real-Time Graphs window appears. (See Figure 4-7.) As you can see from the figure, the device named *divot* does not have the DECnet IV or XNS protocols activated.

If you receive an SNMP error message, check the device's SNMP configuration under the Options menu. HP OpenView 3.3 uses public as the default for the set community string. If you have a device or devices that use other community strings, you will need to configure them appropriately using **Options>SNMP Configuration**.

On SNM, you can customize your current graph. For instructions, refer to your *SunNet Manager 2.0 User's Guide*. The Grapher is an SNM feature and is not covered in this publication. On HP OpenView or NetView for AIX you can customize your graph using their platform grapher. Refer to your NMS documentation for more information.

Step 3 To gather data on interface health, click on the appropriate button. Table 4-6 describes the buttons and MIB object descriptions that are polled when you press the button.

Step 4 To gather data on router health, click on one of the button choices for data.

Table 4-5 describes the buttons and MIB object descriptions that are polled when you press the button. Refer to the *Cisco MIB User Quick Reference* for a description of MIB objects.

Step 5 To gather data on protocols, click on the appropriate button. Table 4-7 describes the buttons and MIB object descriptions that are polled when you press the button.

The real-time graphs you create will use the polling frequency that is set in the Polling Frequency window. You can enter a new polling frequency selecting **Options>Set Polling Frequency**, entering the new frequency in the window, and clicking on **OK**.

Step 6 After you click on a button, a grapher window appears.

Figure 4-8 shows the real-time graph for buffer space on SNM. Graphs on SNM look similar to the one shown in Figure 4-8.
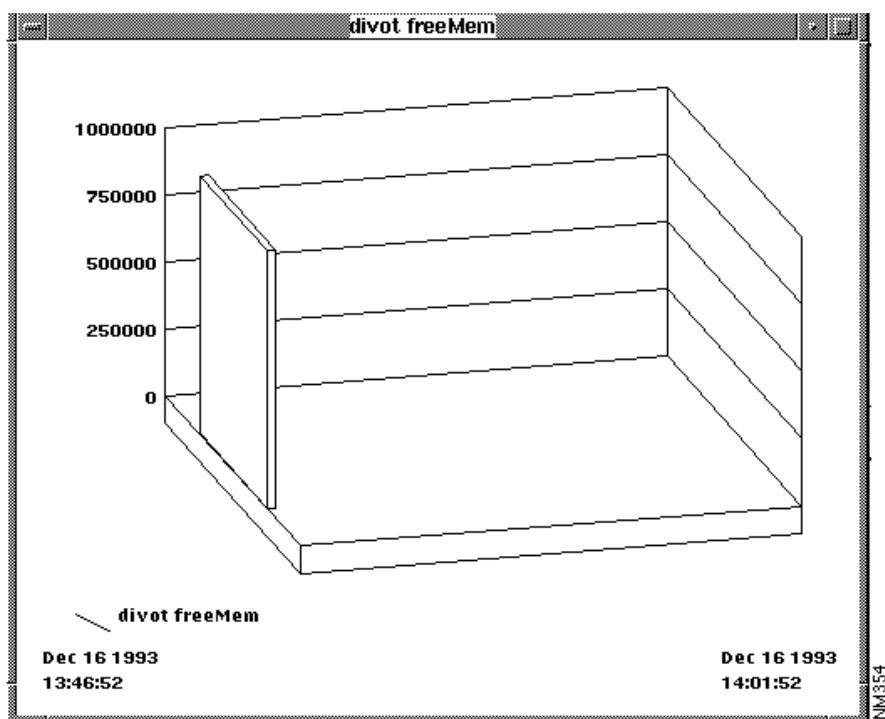

Figure 4-8 Real-Time Grapher Window for Buffer Space Statistics on SunNet Manager

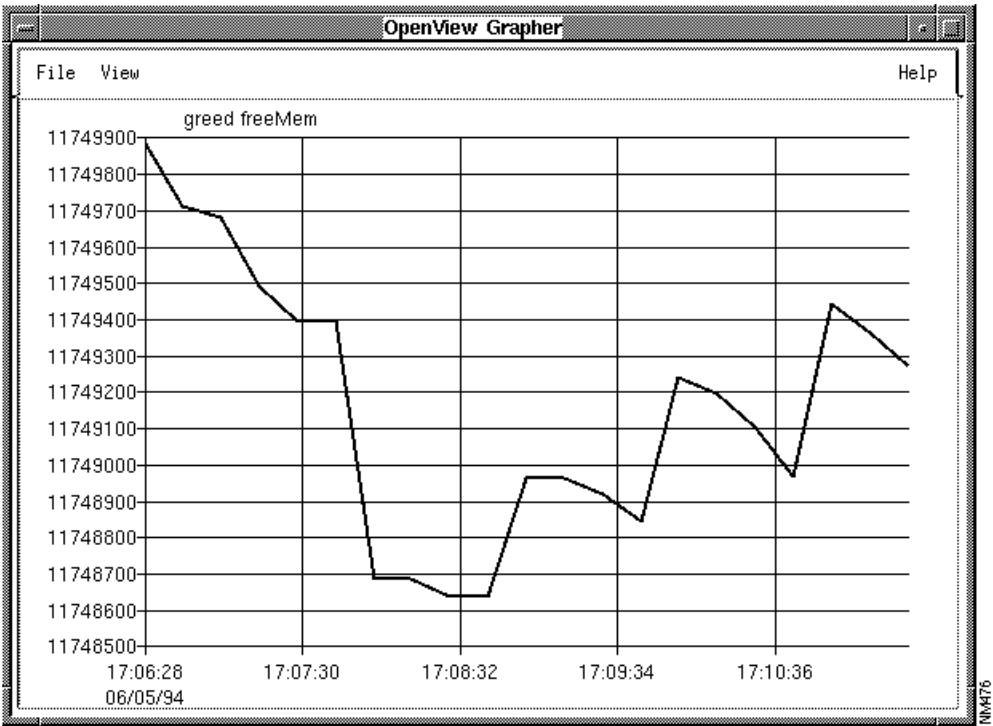Real-time graphs on HP OpenView look like the one shown in Figure 4-9.

Figure 4-9 Real-Time Grapher Window for Free Memory on HP OpenView

If you are using SNM, a Results Grapher window also appears (Figure 4-10), but is hidden behind the grapher window. The devices and objects you selected appear in the scroller window.
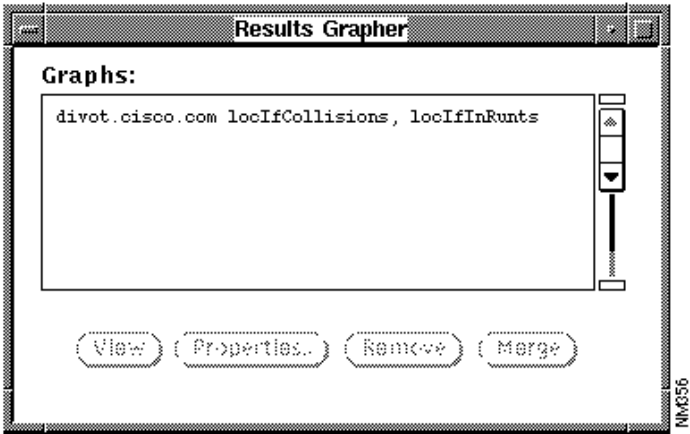

Figure 4-10 Results Grapher Window on SunNet Manager

Viewing or Changing Graph Properties (SunNet Manager only)

If you want to change the appearance of your real-time graph or delete the graph, use the SNM Results Grapher window. The following tasks can be performed using the SNM Results Grapher:

- View different graphs listed in the Results Grapher browser.

- Change graph properties, such as color, plot values, and scaling parameters.

- Remove or halt real-time graphs data collection.

- Merge two or more graphs.

For more information on the SNM grapher, refer to the *SunNet Manager 2.0 User's Guide*.

Caution   You must remove graphing requests explicitly. Just exiting the Real-Time Graphs application will allow graphs to continue to polling for information. Since the graphing process requires CPU resources, you may want to delete graphs from the Results Grapher window to save these resources for other uses.

Table 4-5, Table 4-6, and Table 4-7 describe the command buttons in the Real-Time Graphs window. (See Figure 4-7.)

Table 4-5   Router Health Buttons

| Buttons	Description	MIB Object Names
Buffer Space	Displays the following buffer elements: number of free buffers, maximum number, hits, misses, creates, buffer allocation failures, and buffer create failures due to no free memory.	*bufferElFree* *bufferElMax* *bufferElHit* *bufferElMiss* *bufferElCreate* *bufferFail* *bufferNoMem*
CPU Load	Displays CPU busy percentage for one- and five-minute averages and for the last five-second period.	*avgBusy1* *avgBusy5* *busyPer*
Environment	Displays AGS+ router internal temperature and airflow statistics.	*envInternalTemperature (envTestPt1Measure)* *envAirflowTemperature (envTestPt2Measure)*
Free Memory	Displays the amount of free memory in bytes.	*freeMem*
Security	Displays the total number of packets dropped due to access control failures.	*ipNoaccess* For DECnet, *dnNoaccess* For AppleTalk, *atNoaccess*

Table 4-6   Interface Health Buttons

Buttons	Description	MIB Object Names
Bits/Sec	Displays the 5-minute average of input and output bits per second for Cisco-specific devices.	*locIfInBitsSec* *locIfOutBitsSec*
Bytes	Displays the 5-minute average of input and output bits per second. For routers with Software Release 9.1, displays all protocols on an interface.	*ifInOctets* *ifOutOctets*
Errors	For Cisco-specific devices, displays number of input packets with various characteristics.	For Ethernet, 802.3 CSMA/CD, and starLAN: *locIfCollisions* *locIfInRunts* *locIfInGiants* *locIfInCRC* *locIfResets* *locIfRestarts* For FDDI and Token Ring: *locIfInRunts* *locIfInGiants* *locIfInCRC* *locIfResets* *locIfRestarts* For serial (Cisco only): *locIfInFrame* *locIfInOverrun* *locIfInIgnored* *locIfInAbort* *locIfResets* *locIfRestarts* *locIfCarTrans*
	For any non-Cisco devices, displays number of input and output errors with various characteristics on an interface.	*ifInErrors* *ifOutErrors*
Packets/Sec	Displays the 5-minute average of input and output packets per second on an interface.	*locIfInPktsSec* *locIfOutPktsSec*
Packets	For any device, displays the input and output packets on an interface. For routers with Software Release 9.1, displays all protocols on an interface.	*ifInNUcastPkts* *ifOutNUcastPkts* *ifInUcastPkts* *ifOutUcastPkts*
Queue	For any Cisco-specific device, displays the number of packets dropped because the input or output queue was full.	*locIfInputQueueDrops* *locIfOutputQueueDrops* *ifOutQLen* (non-Cisco devices)

Table 4-7   Protocol Traffic Buttons

Button	Description	MIB Object Names
AppleTalk	Displays cumulative AppleTalk packet activity.	*atInput*
		atOutput
		atForward
		atLocal
		atBcastin
		atBcastout
		atNoencap
		atInmult
IP	Displays cumulative IP packet activity.	*ipForwDatagrams*
		ipInDelivers
		ipInReceives
		ipOutRequests
		ipInDiscards
		ipOutDiscards
		ipOutNoRoutes
ICMP	Displays cumulative ICMP activity.	*icmpInMsgs*
		icmpOutMsgs
		icmpInErrors
		icmpOutErrors
		icmpInEchos
		icmpOutEchos
		icmpInEchoReps
		icmpOutEchoReps
SNMP	Displays cumulative SNMP packet activity.	*snmpInTotalReqVars*
		snmpInTotalSetVars
		snmpInTraps
		snmpOutTraps
		snmpInPkts
		snmpOutPkts
		snmpOutGetResponses
		snmpOutSetRequests
TCP	Displays cumulative TCP packet activity.	*tcpActiveOpens*
		tcpPassiveOpens
		tcpAttemptFails
		tcpEstabResets
		tcpInSegs
		tcpOutSegs
		tcpRetransSegs

Button	Description	MIB Object Names
UDP	Displays cumulative UDP packet activity.	*udpInDatagrams*
		udpOutDatagrams
		udpInErrors
		udpNoPorts
DECnet IV	Displays cumulative DECnet IV packet activity.	*dnForward*
		dnReceived
		dnHellos
		dnHellosent
		dnDatas
		dnNoroute
		dnToomanyhops
Novell	Displays cumulative Novell packet activity.	*novellInput*
		novellOutput
		novellForward
		novellLocal
		novellBcastin
		novellBcastout
		novellNoroute
		novellNoencap
VINES	Displays cumulative VINES packet activity.	*vinesInput*
		vinesOutput
		vinesForwarded
		vinesLocaldest
		vinesBcastin
		vinesBcastout
		vinesBcastfwd
		vinesNoroute
		vinesEncapsfailed
XNS	Displays cumulative XNS packet activity.	*xnsInput*
		xnsForward
		xnsLocal
		xnsBcastin
		xnsBcastout
		xnsErrin
		xnsErrout

Using the Path Tool
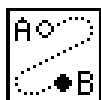
The Path Tool application enables you to monitor the performance our severity thresholds over a period of time and monitor the changes accordingly so you can collect baseline data. This application is described in the section "Locating Device Routing Paths" in Chapter 3, "Fault Management."

NM972

Using Show Commands to View Router Data

You can use the Show Commands application to monitor system status, IP information, and traffic information. This information helps you to determine how to change and improve the efficiency of your network environment.

For more detailed descriptions of all **show** commands, refer to the *Router Products Configuration and Reference* publication refer to Volume 1 for all but protocol-specific **show traffic** commands and **show** commands.

For detailed information on the Show Command window components, refer to "Using Show Commands to View Router Data," in Chapter 3, "Fault Management."

To use the Show Commands application, perform the following steps:

Step 1   Click on a network device.

Step 2   Select **Show Commands**.

On SNM, select **Tools>Show Commands**.

On HP OpenView or NetView for AIX, select **Diagnose>Show Commands**.

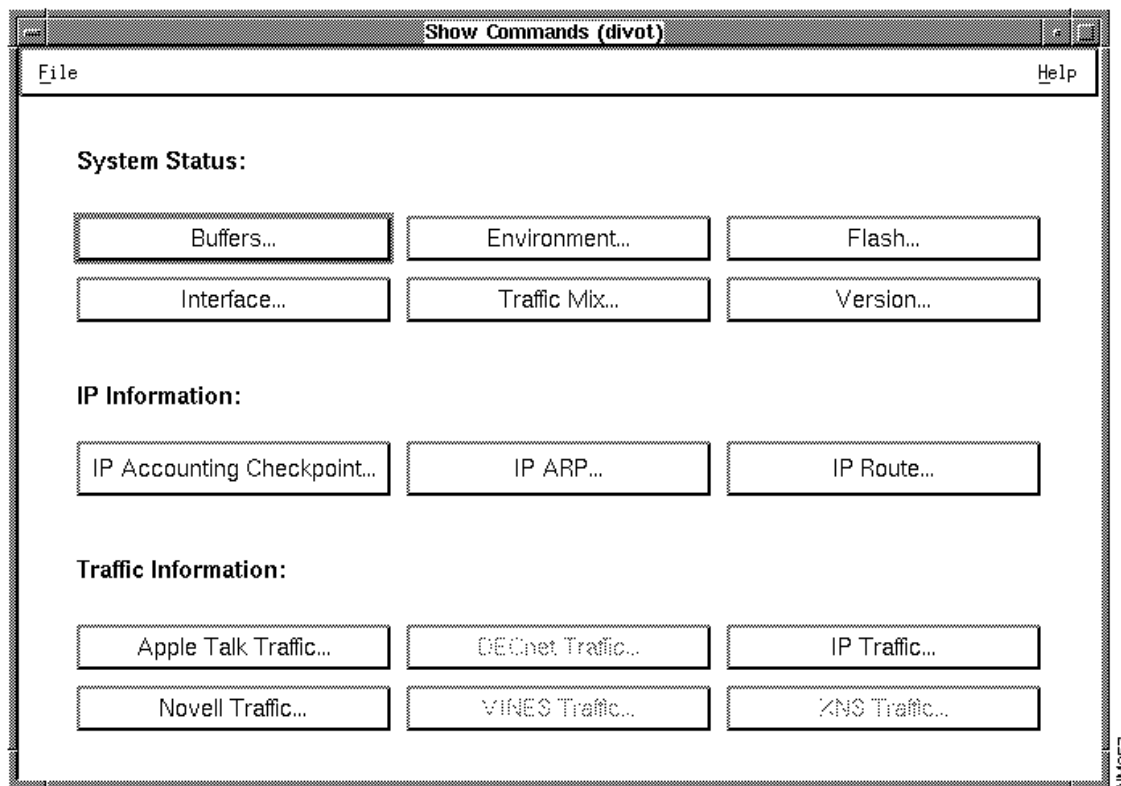Figure 4-11 illustrates the Show Commands window.

Figure 4-11 Show Commands Window

Step 3   To request specific system status, IP information, or traffic information, click on the desired Show Command button.

Each show window is described in detail in the section "Using Show Commands to View Router Data," in Chapter 3.

Step 4   To exit this window, select **File>Exit**.

Note   The **Show** command information provided using the CiscoWorks software differs slightly from the information provided by **show** command performed directly at the router console. A different appearance may result if information is not available through the SNMP protocol.

Loading Private MIBs on HP OpenView or NetView for AIX

With CiscoWorks on HP OpenView or NetView for AIX, you can load a private MIB into the menu bar. To accomplish this, you need to place the MIB file onto the network management platform, then load the MIB file into CiscoWorks.

By default, the Cisco MIB is loaded after you install CiscoWorks. To load a private MIB into HP OpenView or NetView for AIX and CiscoWorks, perform the follow steps:

Step 1   Select **Options>Load/Unload MIBs:SNMP**.

You see a list of currently loaded MIBs.

Step 2   Click on **Load** to see a list of MIB files available in the directory */usr/OV/snmp_mibs*.

Step 3   Select a MIB file to load.

Step 4   Click on **OK** to load the specified file.

This action places the MIB file into the platform. To view the newly added MIB, select the **Browse MIB:SNMP** command under the Monitor menu.

Step 5   To load a private MIB from HP OpenView or NetView for AIX into CiscoWorks, run the *makemib* script as described in Appendix A of the *CiscoWorks User Guide.*

Writing Reports with the Sybase Data Workbench (DWB)

The CiscoWorks **Sybase DWB** application allows you to access the Sybase Data Workbench (DWB) utilities. While Sybase DWB provides several applications, this section only describes how to use the report-writing application in DWB. For more information on Sybase DWB, refer to the Sybase documentation.

There are several ways to access Sybase DWB: through the SNM Tools menu, through the Misc menu, or through the command line. These methods for accessing Sybase DWB are described in the following sections.

Note   The CiscoWorks system administrator should immediately change the default to a password appropriate to your operation using the instructions in the *Sybase System Administration Guide*. For more information on the Data Workbench application, refer to the Sybase documentation. The *Sybase Data Workbench User's Guide* explains the DWB application. The *Sybase Report Workbench User's Guide* explains the report writer accesses through DWB.

For complete information on the Sybase DWB applications, refer to the Sybase manuals listed in the section "About This Guide." For an online manual page, enter **syman dwb**. Sybase contains a documentation directory that has several helpful files called *$SYBASE/doc*.

Note   Sybase Data Workbench windows in this section display OpenLook windowing adornments.

Invoking Sybase DWB from the Menu

To invoke the Sybase DWB application, perform the following steps:

Step 1   Select **Sybase DWB**.

On SNM, select **Tools>Sybase DWB**.

On HP OpenView or NetView for AIX, select **Misc>Sybase DWB.**

The Sybase Data Workbench Login window displays. (See Figure 4-12.)
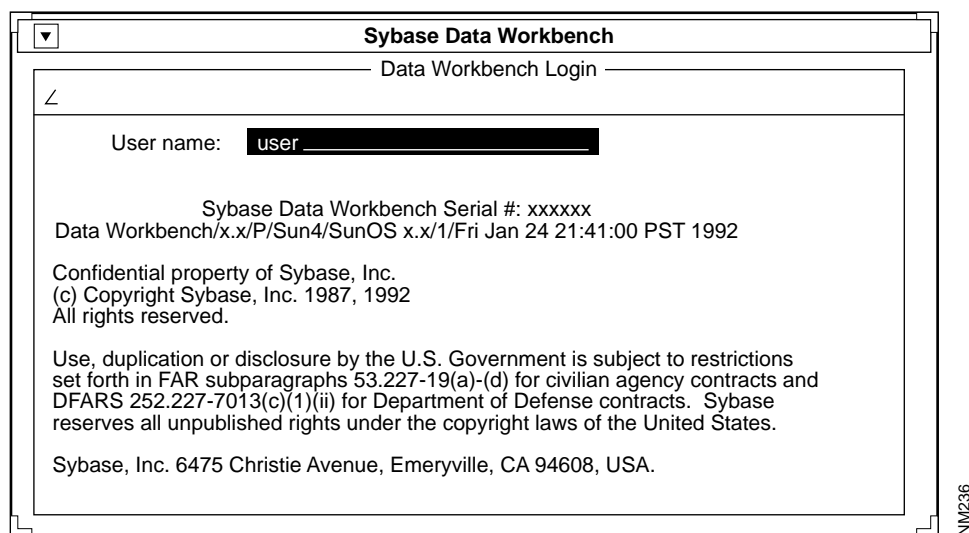
Figure 4-12 Sybase Data Workbench Login Window

Step 2   Use the username *user* to log in.

Step 3   Press Return to enter the default password.

If your password entry is incorrect, the system will display a confirmation window. To close the window, press Return. Retry your password entry.

Step 4   If you decide not to reenter the Data Workbench and want to iconify the window, click on the DWB window menu icon in the upper left corner of the window.

If your password entry is correct, you are immediately transferred to the DWB main window. (See Figure 4-13.)
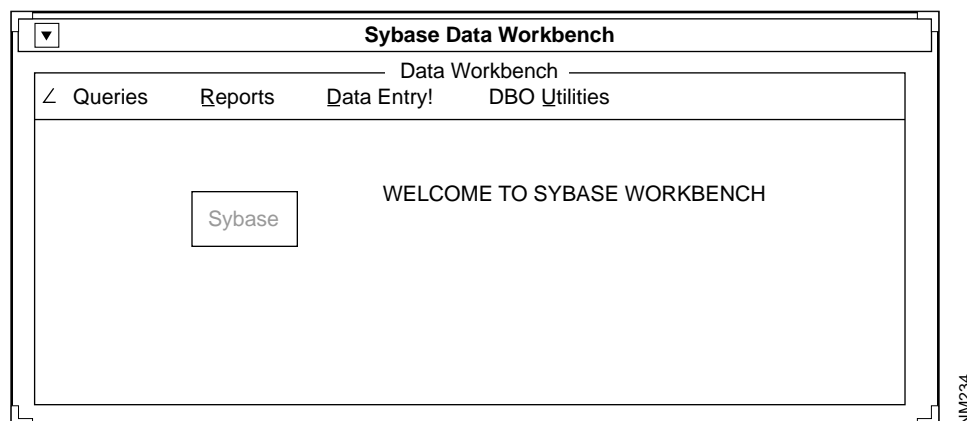
Figure 4-13 Sybase Data Workbench

Note If you want to exit any window in Data Workbench, click on the DWB window menu icon. Select the **Close** command, or click outside of the menu.

Invoking Sybase DWB from the Command Line

Before running this application from the command line, alter your terminal environment variable. We recommend that you access the application through the menu because it alters the variable automatically. If you need to use the command line, follow these steps:

Step 1 To display the current terminal environment, enter the following at the command line:

```
% echo $TERM
```

Make a note of your terminal default. Enter this variable when you finish using Sybase DWB to return your terminal to use other CiscoWorks applications.

Step 2 To set your terminal variable to run using your NMS workstation, enter the following at the command line:

- For SNM or HP OpenView on a Sun, `set term=xterm_c.sun`

- For HP OpenView on HP-UX, `set term=xterm_c.hp`

- For NetView for AIX, `set term=xterm`

Step 3 At the command line, enter the following:

```
% cd $NMSROOT/sybase/bin
```

**Step 4** Then enter the Sybase DWB command:

```
% dwb
```

The DWB Login window appears. (See Figure 4-12.)

Step 5 Use the username *user* to log in.

Step 6 Press Return to enter the default password.

Step 7 For detailed instructions on how to use this application, continue with the section "Using Report Writing."

Step 8 When you have finished using DWB from the command line, reset your terminal default to its original state by entering the following and pressing Return:

```
% set term=original term setting
```

Using Report Writing

The Reports menu on the Sybase Data Workbench window enables you to run and print a report on any table created with the **Device Polling** application.

Following are two examples of how to use the report-writing feature:

- Visual query language (VQL)—VQL is a visually oriented front-end to SQL. VQL relies on menus and a point-and-select method to create the equivalent SQL.

- Structured query language (SQL)—An industry-standard database query language, SQL is commonly used with most well-known databases.

Running a Sample Report with VQL

This section describes how to run sample reports that are available in the CiscoWorks database.

There are three sample reports in the CiscoWorks database: *devices, totals,* and *utilization*. These reports correlate to the Sybase polling tables named *sample**.

A brief description of the sample reports follow:

- *devices*—Reports which network devices are in the database.

- *totals*—Calculates totals for numbers of packets, errors, broadcasts, maximum traffic, and uptime for each device.

- *utilization*—Calculates link utilization for each device interface.

To run the sample reports, perform the following steps:

Step 1 Access **Sybase DWB** as described in the previous section, "Invoking Sybase DWB from the Menu."

Step 2 In the main window of DWB, select **Reports>Existing**. (See Figure 4-14.)

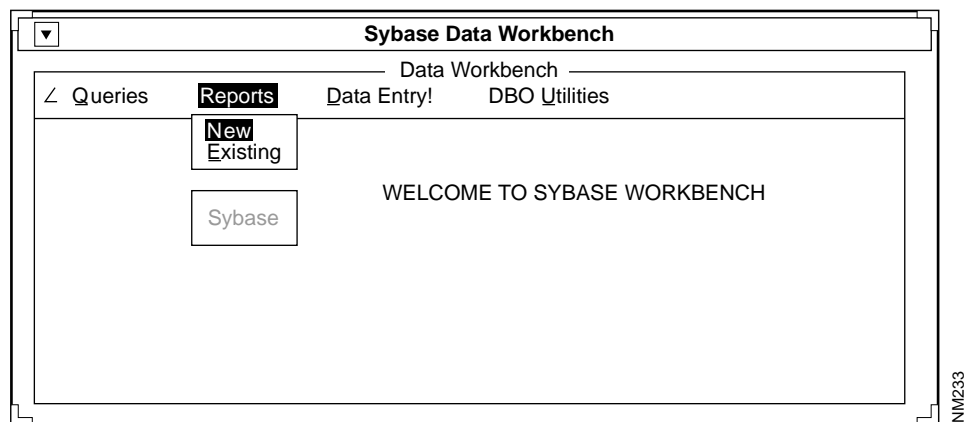**Figure 4-14** Sybase Data Workbench—Reports Menu

Step 3 Select one of the reports. Three reports appear: *devices*, *totals*, and *utilization*.

Step 4 Select **Actions>Run**.

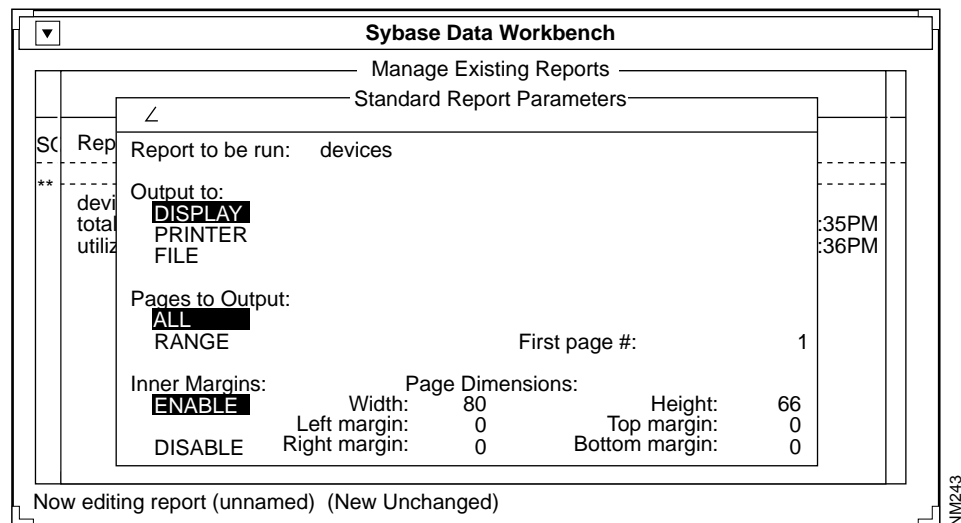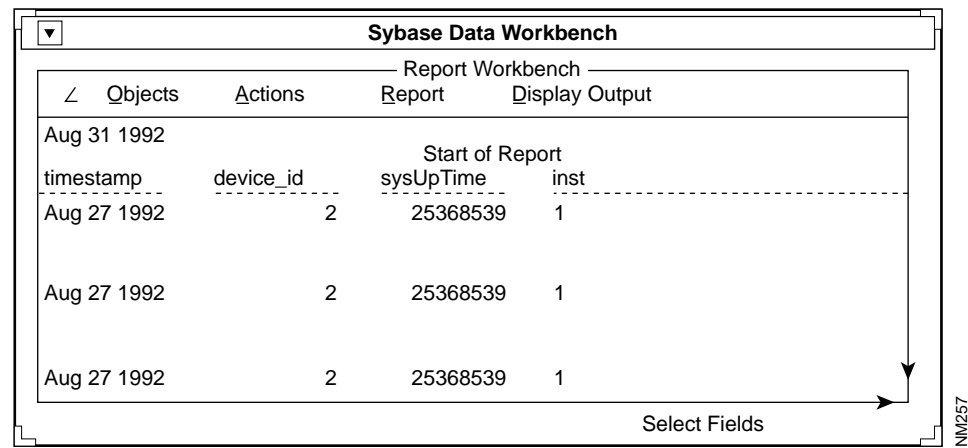The Standard Report Parameters window appears. (See Figure 4-15.)

Figure 4-15 Sybase Data Workbench—Standard Report Parameters Window

Step 5 Accept the default settings to run the report to the display.

Step 6 Click on the DWB window menu icon and select **Close>Apply**.

A report similar to the one in Figure 4-16 displays on your screen.

Figure 4-16 Sybase Data Workbench—Standard Report

Creating a New Report with VQL

This section describes how to create a new report. To create a new report, define the report you want to run and then format it using the **Device Polling** and **Polling Summary** applications. In a custom polling table, the columns containing the polled object values are labeled var[1-n], where n is the number of polled MIB objects. The *genmibview* script in the *$NMSROOT/etc* directory enables you to create a view of the polling table that substitutes MIB object names for current column names such as var1 and var2. Use the *genmibview* script on polling tables prior to CiscoWorks 2.0. Tables created with CiscoWorks 2.0 contain MIB object names for current columns automatically.

The following examples use the default VQL layout. For details on how to customize the format of your reports, refer to the Sybase *Data Workbench User's Guide*

To create a new report, complete the following steps:

Step 1 Define the report in the Polling Summary application.

For information on setting your polling table, refer to the section "Creating Polling Tables Using Device Polling." For information on performing polling, refer to the section "Using Polling Summary" in this chapter.

Step 2 To check that your SYBASE environment variable is set correctly, enter the following at the command line:

```
hostname% printenv SYBASE
```

The system should respond with the location of SYBASE on your system, typically */usr/sybase*.

Step 3 To run the *genmibview* script, enter the following at the command line:

```
% cd $NMSROOT/etc
% genmibview -Uuser -Ppassword tablename
```

The *genmibview* script creates separate tables that contain the MIB variable names instead of variable names such as var1, var2, and so on. The new table name will append *_view* to the original table name. For example, if the table name is *traffic*, the new table is called *traffic_view.*

Step 4 Select **Reports>New**.

The SQL Batch Editor window appears. (See Figure 4-17.)

```
┌─────────────────────────────────────────────────────────────────┐
│ ▼              Sybase Data Workbench                             │
│        ┌──────────── Report Workbench ────────────┐             │
│  ∠     │          ──── SQL Batch Editor ────       │             │
│        │  ∠ VQL!      File      Edit                │             │
│ SQL    │                                            │  ─ ─ ─ ─   │
│        │              Name:   test_                 │             │
│ *** N  │   Maximum output width:  9999              │             │
│        │ Character field maximum width:   40        │             │
│        │           Break style:  Style 2   Style 1  │             │
│        │             SQL text:                      │             │
│        │ SELECT timestamp, device,_id, sysUpTime, inst, rec_type,│
│        │ ifOutOctets, ifInOctets, ifOperStatus, ifAdminStatus    │
│        │ FROM test 1_view                           │             │
│        │ COMPUTE AVG(ifOutOctets), AVG(ifInOctets), AVG(ifOperStatus),│
│        │ AVG(ifAdminStatus)                         │             │
│        └────────────────────────────────────────────┘             │
│  Now editing report:  (unnamed)  (New Unchanged)                 │
└─────────────────────────────────────────────────────────────────┘
```

Figure 4-17 **SQL Batch Editor Window**

Step 5 Select the **VQL!** menu.

Step 6 Select **Use>Table**.

The Use Table window appears over the center of the VQL window.

Step 7 At the *Name* prompt, enter the filename of the table that contains the variables for which you would like to run a report.

Step 8 Click on the DWB window menu icon in the upper left corner, select the **Close>Apply**.

The contents of the table appear in the window.

Step 9 Select all of the data columns in the table.

Step 10 Select **Add>Result Columns**. (See Figure 4-18.)
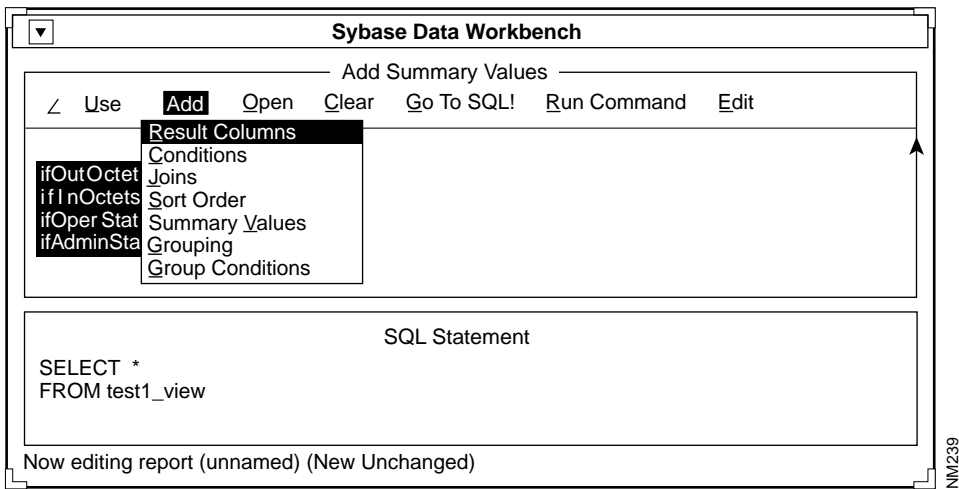
Figure 4-18 Sybase Data Workbench—Results Columns Command

Step 11 Select **Add>Summary Values**.

The Add Summary Values window appears.

Step 12 Select the items to be averaged, then select **Function>Average**. (See Figure 4-19.)
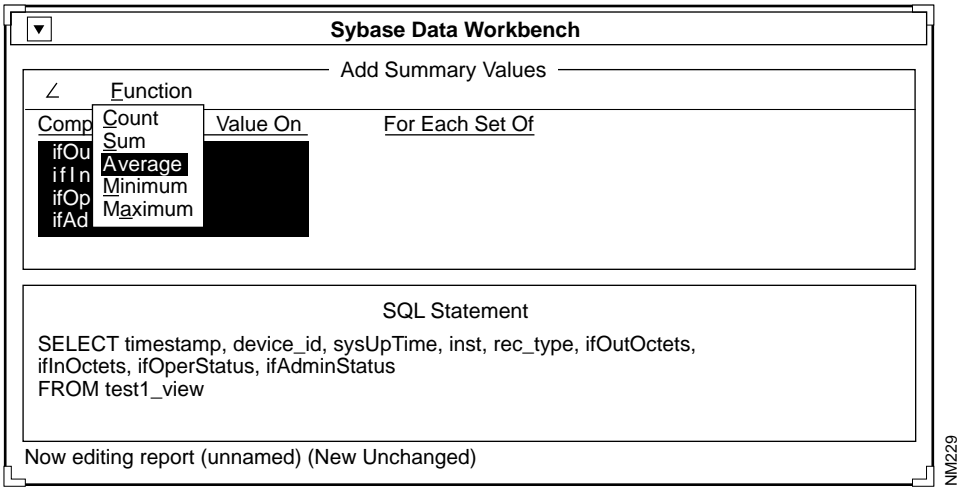
Figure 4-19 Sybase Data Workbench—Add Summary Values Window

Step 13 Click on the DWB window menu icon in the upper left corner and select the **Close** option to exit.

The SQL Batch Editor displays.

Step 14 At the *Name* prompt, enter a name to identify your SQL batch statement.

Step 15 Click on the DWB window menu icon in the upper left corner, and select **Close>Apply**. (See Figure 4-17.)

The Report Workbench window appears.

Step 16 Select **Report>Processing**.

Ensure that the filename is entered at the Process prompt.

Step 17 Select the **Compile** menu. (See Figure 4-20.)

Figure 4-20 Sybase Data Workbench—Compile Menu

Step 18 Click on the DWB window menu icon in the upper left corner and select **Close>Apply**.

Step 19 Select **Report>Run**.

The Standard Report Parameters window appears.

Step 20 Accept the default settings to run the report to the display.

Step 21 Click on the DWB window menu icon in the upper left corner, and select **Close>Apply**.

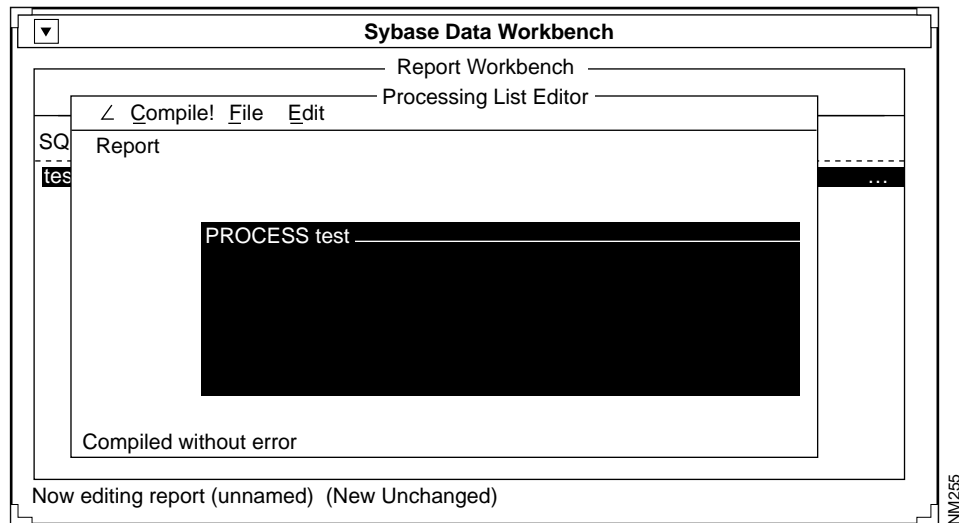A report similar to the one in Figure 4-15 appears on your screen.

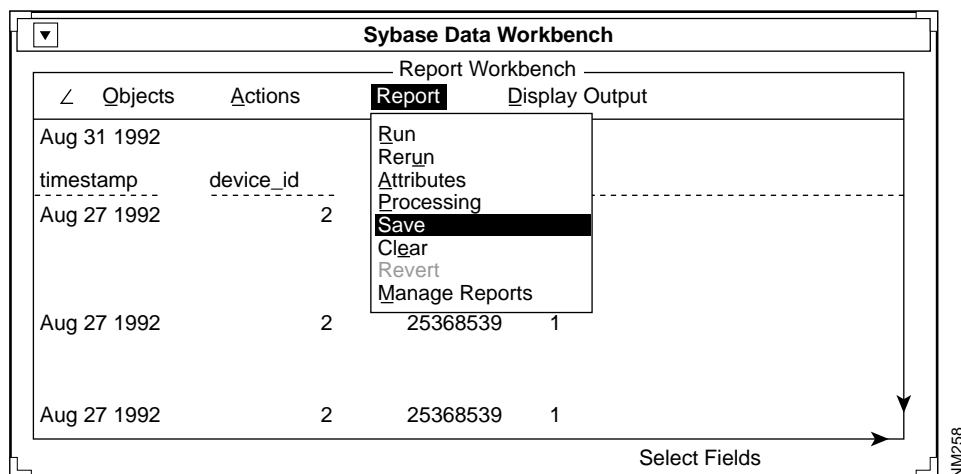**Step 22** To save the report, select **Report>Save**. (See Figure 4-21.)

```
┌──────────────────────────────────────────────────────────┐
│ ▼                  Sybase Data Workbench                  │
│              ┌──── Report Workbench ────┐                 │
│   ∠  Objects     Actions    │Report│      Display Output  │
│                           ┌─────────────┐                 │
│ Aug 31 1992               │Run          │                 │
│                           │Rerun        │                 │
│ timestamp     device_id   │Attributes   │                 │
│                           │Processing   │                 │
│ Aug 27 1992           2   │Save         │                 │
│                           │Clear        │                 │
│                           │Revert       │                 │
│                           │Manage Reports│                │
│ Aug 27 1992           2        25368539    1              │
│                                                           │
│                                                           │
│ Aug 27 1992           2        25368539    1              │
│                                          Select Fields    │
└──────────────────────────────────────────────────────────┘
```

Figure 4-21 Sybase Data Workbench—Saving the Report

Step 23 At the **Save edited report as** prompt, enter the name of the report.

Step 24 Click on the DWB window menu icon in the upper left corner and select **Close>Apply**.

Step 25 Click on **Exit Report Workbench** and press Return.

Running a Sample Report with SQL

Several SQL report-writing examples are presented in this section. Using these examples, you can construct your custom reports, print them on a printer, or view them on the screen.

A data field contains a single item of information such as device name, device type, or postal code. Related columns of data fields are collected into tables. These collections of related information are arranged in records that relate to entities such as addresses, names, device names, and so on. These horizontal collections of records are also called *rows*. In the database, all fields are in columns and all records in rows.

For instance, the following columns are in the locations table:

- address_id
- location
- street
- street_two
- city
- state
- country
- zip_code

Static Table Queries

The sections that follow describe how to use columns from static tables.

Creating Reports

You can create reports using columns of data in the static database. For example, to create a quick report on devices of a certain type including type, name, and serial number, enter the following query in the SQL parameter window:

```
select device_id, device_type, device_name, serial_number from
devices where device_type = 0
```

Running the Poll

Run the poll by selecting **Sybase DWB**. From the Queries menu, select the SQL option. When the SQL window appears, you can enter the poll directly into the window, or read it in from a file.

Reading the File

To read a file into the window, select **Recall Command**>**From Host File**. You will be prompted for the name of the file in which you have stored the poll.

Processing the Poll

Regardless of how you enter the poll, process it by selecting **Run Command**. You can display the results on screen, send them to a printer, or store them in a file. If you want to experiment, a screen display is usually most convenient.

After the results display, you can use the Format option to format each column of your report. For detailed information, refer to the *Sybase Data Workbench User's Guide*. While in the display, you can scroll back and forth or up and down among the results data.

If you want to produce hardcopy reports with pagination, headers, and footers, rather than just columns of text, use Report Workbench to format the report. The Report Workbench will take your running poll and allow you to lay out the results for output to a printer using a WYSIWYG editor.

After you successfully run a poll and specify the set of data you are interested in, exit the Display Window (click the X icon) and then exit the SQL editor.

Select **Reports>New**. This opens the SQL Batch Editor of the Report Workbench. Name the file in which you stored the poll and bring up the default report.

For a description of how to format your results for output to the printer, refer to the *Sybase Data Workbench User's Guide*.

These examples show a simple version of the results in an online poll in DWB:

```
device_id        device_type      device_name      serial_number
5                0                dross            NULL
```

Often, data you want is distributed across multiple tables. If you want to know more about the device type than just its number, you can include a description in your report by combining data from the devices and device_types tables. By naming both tables in your FROM clause, you expand the WHERE clause to find only the relevant type description by linking (the technical term is "joining") the two tables on their shared columns.

The poll request to do this follows:

```
select device_id, devices.device_type, device_name, device_types.
type_desc from devices, device_types where device_id = 5 and
devices.device_type = device_types.device_type
```

Notice that a full name in tablename.columnname format is required whenever a column name is not unique in a poll. The results would look something like this, depending on how you formatted them:

```
device_id          device_type     device_name     type_descr
5                       0           dross           default
```

You can have as many tables as you want. A poll (and results) drawing data from three tables could look like this:

```
select device_name, device_types.type_desc, vendor_name
from devices, device_types, vendors
where device_type =0
and devices.device_type = device_types.device_type
and devices.vendor_id = vendors.vendor_id
device_name        type_descr      vendor_name
dross                default        default
```

Running Custom Table Queries

You can also use columns from static tables in conjunction with columns from custom tables you create. These custom tables include real-time data for the MIB variables you select. Because no two networks are alike, the type of diagnostic or accounting information will be specific to your environment.

These examples are included only as a starting point for your own experimentation. If you need information about device utilization levels and uptime percent, you could produce this report using Device Polling, Polling Summary, and Sybase DWB:

```
For the period From: Mar 27 1992  5:44PM   To: Mar 28 1992  11:59AM
Device           Interface       Pct Util
--------------- --------------- --------------------
 abc             Ethernet0                   83.607993
 abc             Ethernet1                   75.815578
 abc             Ethernet10                  36.269430
 abc             Ethernet11                  83.352942
 abc             Ethernet12                  36.269430
 abc             Ethernet13                  36.269430
 abc             Ethernet14                  72.564532
 abc             Ethernet15                  82.645540
 abc             Ethernet2                   74.855613
 abc             Ethernet3                   97.716789
```

To obtain these results, use polling tables to collect information from the network and store it in the database. Then use SQL queries to display data from the database.

In the following example, the device is 5; the interfaces are 5, 11, and 12; and the MIB variables are *ifInOctets*, *ifOutOctets*, and *ifSpeed*; the polling interval is 10 minutes; and the name of the new table is *device_usage*.

The text of the poll that creates this report is as follows:

```
/* Report %Utilization for each interface in 15 minute intervals */
/*
 * Get the 15 minute totals for each interface that is operational
 */
select device_id, convert(int, inst) inst, sum(ifInOctets) ifInOctets, sum(ifOutOctets)
ifOutOctets, max(ifSpeed) ifSpeed, min(timestamp) interval, datepart(dd, timestamp),
datepart(hh, timestamp) hourint, convert(int, datepart(mi, timestamp)/15) minint
into #t1
from sample_view
where rec_type = 0
group by device_id, convert(int, inst), datepart(dd, timestamp), datepart(hh, timestamp),
convert(int, datepart(mi, timestamp)/15)
having max(ifOperStatus) = 1 and max(ifAdminStatus) = 1
go

/*
 * Compute utilization levels for each interface. The computation is split
 * into serial and non-serial interfaces (full duplex and half duplex) by
 * using interface type.
 */
select substring(device_name, 1, 14) Device, interval, substring(interface_name, 1, 14)
Interface,100 * convert(float, ifOutOctets)/ (15.0 * 60.0) /convert(float, (ifSpeed/8))
"% Utilization"
from #t1, devices, interfaces
where #t1.device_id = devices.device_id
and #t1.inst = interfaces.interface_id
and #t1.device_id = interfaces.device_id
and interfaces.interface_type in (2,3,4,5,16,17,18,19,20,21,22,23)
union
select substring(device_name, 1, 14) Device, interval, substring(interface_name, 1, 14)
Interface,100 * convert(float, (ifInOctets + ifOutOctets))/ (15.0 * 60.0) /convert(float,
(ifSpeed/8)) "% Utilization"
from #t1, devices, interfaces
where #t1.device_id = devices.device_id
and #t1.inst = interfaces.interface_id
and #t1.device_id = interfaces.device_id
and not (interfaces.interface_type in (2,3,4,5,16,17,18,19,20,21,22,23))
order by substring(device_name, 1, 14), interval, substring(interface_name, 1, 14)
go

/*
 * Clean up temporary tables
 */
drop table #t1
go
```

Specifying MIB Object Values

When you create a polling table in Device Polling, you specify the MIB object values that you want to collect. These values are collected by the poller and added to the database as the values are received. When polled values arrive at the poller in different packets, the values are written to the database in different records. Therefore, you might find that the data for a given poll is spread over two or more records, with NULL appearing where data was not available.

The following example shows how polling data is received over time and added to the database:

```
Polling:
                time     sysUpTime       var1      var2      inst
                x            100          10                    1
                x+1          101                     15        1
                x+2          102          10                    2
                x+10         110          20                    1
                x+11         111          25                    2
Database storage:
                rec_type sysUpTime       var1      var2      inst
                1            100          10                    1
                1            101                     15        1
                1            102          10                    2
                0            8            10        1
                0            1            15                    2
```

You cannot perform row operations without first grouping and aggregating your data. You might want to use the timestamp column to perform groupings and then use the AVG or MAX (for absolute values) or SUM (for delta values) to aggregate. You can use the following GROUP BY clause to group by device, inst, and 15-minute intervals:

```
select ...
group by device_id, convert(int, inst), datepart(dd, timestamp),
datepart(hh, timestamp), convert(int, datepart(mi, timestamp)/15)
```

Writing SQL Reports for Custom Polling Tables

To write SQL-based reports for analyzing data collected in polling tables, you must be familiar with CiscoWorks polling applications, Simple Network Management Protocol (SNMP), and Structured Query Language (SQL).

Table 4-8 shows the columns that are always present in a polling table.

Table 4-8   Columns in a Polling Table

Field Name	Field Size	Field Type
timestamp	datetime	Timestamp for each row of data polled
device_id	int	Join to devices table
sysUpTime	int	System up time in 100ths of a second
inst	char (255)	Index into MIB object table
rec_type	tinyint	0 = normal record, 1 = start record, 2 = not used, 3 = restart record, 4= irregular records

The remaining columns are determined by your definition of the polling table and contain the values of the polled MIB objects.

If the polled MIB objects are indexed by *ifIndex*, the *inst* column contains the appropriate value (for example, *ifIndex* or *locIfxxx)*. You can join the *inst* column with the *interface_id* column in the Interfaces table. To join the two columns, use the following SQL syntax:

```
select ...
where ...
and convert(int, inst) = interfaces.interface_id
```

When you start polling a device, CiscoWorks creates a record with rec_type = 1. All the values in this record are absolute. All other records have rec_type = 0 and contain a mix of delta and absolute values. When you stop polling a device, you do not get any additional rows. A new rec_type value of 4 has been added to account for any anomalies or irregularities, which includes absolute values instead of delta values.

You can determine whether a MIB object will contain an absolute or delta value by getting its data type from the *$SNMHOME/agents/cisco.schema* file and identifying the value for each type. For example, *sysUpTime* uses timeticks as its record type, so the data value will be displayed as a delta value (the difference in value between the current poll and the last poll) if rec_type = 0.

For information on identifying the value for a data type, refer to Table 4-9.

Table 4-9   Values for Types of Records

Record Type	Value
timeticks	delta
counter	delta
gauge	absolute
int	absolute

Computing sysUpTime Over a Polling Interval

To calculate time-dependent statistics for a device, calculate the total device uptime during the polling period. Create a temporary table containing the uptime for each polled device by using the following SQL commands:

```
select device_id, sum(sysUpTime)/100 uptime into #t1 from <table>
where rec_type = 0 group by device_id
```

The SQL commands create a temporary table called *#t1* that contains device IDs and the uptime values for each device. The value is divided by 100 so that uptime will be in seconds (*sysUpTime* is in 100ths of a second).

Using ifOperStatus and ifAdminStatus

If you are polling interface values, include *ifOperStatus* and *ifAdminStatus* in your polling values to determine whether a given interface is operational. If you are using a GROUP BY clause in your SQL, add a HAVING condition to the group to check whether the interface was operational. Add the following commands to the group:

```
select ...
group by ...
having max(ifOperStatus) = 1 and max(ifAdminStatus) = 1
```

The possible values for *ifAdminStatus* are listed in Table 4-10.

Table 4-10   Values for ifAdminStatus

Value	Meaning
1	up
2	down
3	testing

These values are also described in RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II.*

Calculating Traffic Values on Serial Lines

Because serial lines are full duplex, there is no single way to perform a traffic calculation. The following calculation provides the best view:

```
traffic = max(ifInOctets, ifOutOctets)
utilization = traffic/line bandwidth
```

In general, however, it is often less complex and more direct to look at traffic as the amount of data transmitted by a given device onto a line. This is consistent with the router calculation that obtains a load value. In this case, use the following calculation (where the utilization is a function of which end of the serial line that you are measuring):

```
traffic = ifOutOctets
utilization = traffic/line bandwidth
```

If you use *ifInOctets+ifOutOctets* as a traffic measure for serial lines, your utilization levels appear to be greater than 100 percent.

To determine whether you have a serial line, check the interface_type column in the Interfaces table. Table 4-11 lists the serial interface types and the protocols associated with each interface. For detailed information on interfaces and protocols, refer to RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II..*

Table 4-11 Serial Interface and Protocol Type

Serial Interface Type	Protocol Type
2	1822
3	1822
4	ddnx25
5	RFC 877-x25
16	lapb
17	sdlc
18	ds1
19	e1
20	ISDN
22	serial
23	PPP

To select serial interfaces, use the following with the WHERE clause:

```
select ...
where ...
and interfaces.interface_type in (2,3,4,5,16,17,18,19,20,21,22,23)
```

Improving Report Format

To improve the appearance of the output from your queries you can make simple with the Format option after you run a poll from the Data Workbench. For more complex changes in the appearance of the output use the Report Writer. You also can save your queries and reports and run them later from UNIX without using Data Workbench.

When you run the poll and display the results on the Data Workbench, the Format menu item lets you change the appearance of the results. Following are some of the things you can change:

- Column width
- Column headings
- Data alignment
- Column display order

The Format option is explained in more detail in the *Sybase Data Workbench User's Guide.* For hardcopy reports, make any additional changes with Report Writer.

Using Report Writer to Present Your Data

After you write and save your poll, you can focus on the presentation of data on the printed page. From the Data Workbench, select **Reports>Define**>**Setup**. Then select **Queries**, and from Queries, select **SQL**. The poll you run appears in the window.

In the window, you can recall queries from files and save them to files. When you use the **Run Report** command, a layout of all the elements of the report displays. You can edit the page layout and the data display, add trim, and change control break processing of the data. **Run Report** allows you to display the default layout on screen.

The poll will collect the set of data to be reported, and the report writer will format the current set of data in the way you describe. When you exit from the Reports window, you will be prompted to save your report in the database. Enter the report name and confirm your exit from the report menu.

For complete information about the **Report** command, refer to the *Sybase Data Workbench User's Guide.*

Running Reports from UNIX

You can run reports created and saved by the Data Workbench from the UNIX command line with the **runrw** command. The output of the report is automatically sent to your default printer. It is located in the *$SYBASE/bin* directory.

The following is the command syntax for the sample report called *storm_report*:

runrw *storm_report* **-U***yourname* **-P***password*

To run a report owned by another user, enter the following at the command prompt:

```
% runrw user name.devicerpt -Uyourname -Ppassword
```

For more information about the **runrw** command, enter the following command:

syman runrw