LUND
PERFORMANCE SOLUTIONS

WHAT YOU NEED TO KNOW. **WHEN** YOU NEED TO KNOW IT.

Performance Gallery Gold

User Guide

Legal Notices

Restricted Rights Legend

LUND PERFORMANCE SOLUTIONS

240 2nd Avenue SW

Albany, OR 97321

USA

Copyright Notices

Printed in the United States of America.

Trademark Notices

Performance Gallery Gold version F.04

Whitney Olsen, Laura Bryngelson, Rodica Lupea 6/22/05

TABLE OF CONTENTS

1

INTRODUCTION

Welcome to Performance Gallery Gold

Performance Gallery Gold™ is a PC-based software tool that uses the data gathered by one of the following host-based collector programs to report many critical aspects of system performance.

- Meta-View Performance Manager™ for MPE/iX, the HP e3000 system performance data collection agent by Lund Performance Solutions.
- Meta-View Performance Manager™ for HP-UX, the HP 9000 system performance data collection agent by Lund Performance Solutions.
- Meta-View Performance Manager™ for Solaris, the system performance data collection agent by Lund Performance Solutions.
- Performance SeNTry®, the Windows NT system performance data collection agent by Demand Technology Software.

NOTE Utilizing Performance Gallery Gold's Real-Time monitoring for MPE/iX and Unix systems requires that the Meta-View Data Service (mvdatad for Unix and mvdataj for MPE/iX) be included with Meta-View Performance Manager agent installation on the host.

Monitoring and tracking system performance is an integral part of system administration. Performance Gallery Gold transforms performance data records into a variety of full-color graphs and tables that show system administrators precisely how their system handles specific aspects of its workload, such as processor utilization, memory utilization, and disk performance.

Product Requirements

The Windows-based Performance Gallery Gold program requires the following system hardware and software.

Hardware

- IBM-compatible personal computer with a Pentium 90 or later
- 32 MB of RAM (64 MB is recommended)
- 12 MB of free hard disk space (25 MB required for install)
- Double-speed CD-ROM drive
- 256-color VGA or better video
- Two-button mouse (a three-button mouse is recommended)
- Connectivity to the host system

NOTE While .-smf files are indexing, Performance Gallery Gold requires twice as much RAM as the size of the file. For instance, if you load an 8 MB .smf file, 16 MB of RAM are required to index it.

Software

- Windows NT 4.0 with Service Pack 6 or Windows 98/ME/2000/XP for PWS connections.
- Microsoft Internet Explorer 4.x (or higher) Web browser.

  Internet Explorer should be installed on your system to enable specific system files to be updated.
- One of the following host data collectors:
  - Meta-View Host for MPE/iX by Lund Performance Solutions
  - Meta-View Host for HP-UX by Lund Performance Solutions
  - Meta-View Host for Solaris by Lund Performance Solutions
  - Performance SeNTry by Demand Technology Software
- FTP (file transfer protocol) or terminal emulation software to transfer data log files from the HP e3000 or HP 9000 system to the PC (not necessary for Windows NT systems)

Product Support

Lund Performance Solutions Product Support

The Performance Gallery Gold and Meta-View Performance Manager products by Lund enable administrators to collect, monitor, analyze, and report HP 3000, HP 9000, and Windows NT system performance trends faster and more effectively than ever before.

Lund Performance Solutions Main Offices

Postal Address

Lund Performance Solutions

240 2nd Avenue SW

Albany OR 97321 USA

Internet URL

Visit the Lund Performance Solutions Web site at **www.lund.com**.

Telephone Number

For customer and technical support, call **(541) 812-7600**, Monday through Friday during the hours of 7:00 A.M. to 4:00 P.M. Pacific time, excluding holidays.

Fax Number

Transmit fax messages to **(541) 812-7611**.

E-mail Addresses

Send e-mail messages to:

- Sales Team **info@lund.com**
- Technical Support Team **support@lund.com**
- Consulting Team **consulting@lund.com**
- Documentation Team **documentation@lund.com**

Lund Performance Solutions Sales Team

Lund Performance Solution's professional sales team is ready to answer your sales and customer support questions Monday through Friday during the hours 7:00 A.M. to 4:00 P.M. Pacific time, excluding major holidays.

Please contact your sales representative for information about the latest Lund Performance Solutions products, the Lund Software Subscription Plan, upgrade options and prices, and more.

Lund Performance Solutions Technical Support Team

At Lund Performance Solutions, we are working hard to provide you with intuitive software products. Additionally, we try to provide superior online and printed documentation. However, should you find yourself with a technical question that you cannot answer with the tools provided, please contact our technical support team.

NOTE You must be a registered user to access Lund Performance Solutions' support services. Lund Performance Solutions' support services are subject to Lund Performance Solutions' prices, terms, and conditions in place at the time the services are used.

E-mail Tech Support

Ask questions and receive detailed answers from the Technical Support Team by sending an e-mail message to **support@lund.com**. Please include the product version number with your question. You will receive a reply by e-mail.

Telephone Tech Support

You can reach the Technical Support Team by phone at **(541) 812-7600**, Monday through Friday during the hours 7:00 A.M. to 4:00 P.M. Pacific time, excluding major holidays.

When you call, please be at your computer, have your documentation in hand, and be prepared to provide the following information:

- Product name and version number.
- Type of computer hardware you are using.
- Software version number of Windows.
- Exact wording of any messages that appear on your screen.
- What you were doing when the problem occurred.
- How you tried to solve the problem.

Lund Performance Solutions Documentation Team

Lund Performance Solutions makes every effort to produce the highest quality documentation for our products, and we welcome your feedback. If you have comments or suggestions about our online Help or printed guides, send an e-mail message to **documentation@lund.com** or contact your account manager.

Lund Consulting Services IT Consultants

Lund Consulting Services, a division of Lund Performance Solutions, offers strategic IT solutions and expert support to a wide range of businesses. Our team of experienced IT professionals provides onsite consulting, training, and project management services to help businesses optimize their computer resources and achieve long-lasting success.

For information about Lund Consulting Services, please review our Web site, send an e-mail message to **consulting@lund.com**, or contact your account manager.

Demand Technology Software Product Support

Lund Performance Solutions has formed a strategic partnership with Demand Technology Software of Naples, Florida, to introduce an exceptional Windows NT performance monitoring solution for Windows NT networks, servers, and workstations. This joint product integrates the collector technology of Demand Technology's product, Performance SeNTry, with the data management and graphical reporting capabilities of Performance Gallery Gold. For more information about Demand Technology Software, visit their Web site at **www.demandtech.com**.

Demand Technology Software will provide both technical support and certified training for the Performance Gallery with Performance SeNTry bundled software. Lund Performance Solutions will provide sales and consulting services for this product (refer to "Lund Performance Solutions Sales Team" on page 3 and "Lund Consulting Services IT Consultants" on page 4).

Demand Technology Software Main Offices

Postal Address

Demand Technology Software

1020 8th Avenue South, Suite 6

Naples FL 34102 USA

Internet URL

Visit **www.demandtech.com**.

Telephone Number

For customer and technical support, call **(941) 261-8945**, Monday through Friday during the hours of 7:00 A.M. to 4:00 P.M. Eastern time, excluding major holidays.

Fax Number

Transmit a fax message to **(941) 261-5456**.

E-mail Address

Send an e-mail message to **support@demandtech.com**.

Demand Technology Software Technical Support Team

The development and support teams at both Demand Technology Software and Lund Performance Solutions work hard to provide you with intuitive software products and superior online and printed documentation. However, should you find yourself with a technical question that you cannot answer with the tools provided, please contact the Technical Support Team at Demand Technology Software.

E-mail Tech Support

Ask questions and receive detailed answers from the Technical Support Team by sending an e-mail message to **support@demandtech.com**. Please include the product serial number with your question. You will receive a reply by e-mail.

Telephone Tech Support

You can also contact the Technical Support Team by phone at **(941) 261-8945**, Monday through Friday during the hours 7:00 A.M. to 4:00 P.M. Eastern time, excluding major holidays. When you call, please be at your computer, have your documentation in hand, and be prepared to provide the following information:

- Product name and version number.
- Type of computer hardware you are using.
- Software version number of Windows.
- Exact wording of any messages that appear on your screen.
- What you were doing when the problem occurred.
- How you tried to solve the problem.

Demand Technology Software Certified Trainers

Demand Technology Software presents NT system performance training courses at various locations throughout the year. The Certified Trainer Program is designed for trainers from all educational areas, including academia, consulting, and business.

For information about Demand Technology Software's certified training or to receive an application, please review our Web site at **www.demandtech.com** or send an e-mail message to **support@demandtech.com**.

Product Documentation

User's Guide

This user's guide accompanies the Performance Gallery Gold software as a guide for the new user and as a quick reference for experienced users.

This guide assumes that you have a working knowledge of the Windows operating system.

Document Conventions

The user interface in Performance Gallery Gold is designed to accommodate a variety of work styles. Some users prefer using standard menu commands; others are more comfortable with shortcut menus, keyboard commands or toolbar buttons.

Generally, standard menu commands appear in the procedural sections of this document. Alternative methods of accessing menus, controls, and commands are discussed in “Keyboard Commands” on page 275.

Online Help System

In the online Help system, you will find explanations of the many features of Performance Gallery Gold, as well as tips to guide you through the program’s basic functionality.

Reference Materials

- ***Performance SeNTry User’s Guide.*** Copyright © 1999 by Demand Technology Software, Inc., Naples, Florida, USA.
- ***Meta-View Performance Manager - User’s Guide to Meta-View Agent and Host for MPE/iX.*** Copyright © 2002-2005 by Lund Performance Solutions, Albany, Oregon, USA.
- ***Meta-View Performance Manager - User’s Guide to Meta-View Agent and Host for HP-UX.*** Copyright © 2002-2005 by Lund Performance Solutions, Albany, Oregon, USA.
- ***Meta-View Performance Manager - User’s Guide to Meta-View Agent and Host for Solaris.*** Copyright © 2002-2005 by Lund Performance Solutions, Albany, Oregon, USA.
- ***Windows NT Performance Counters Online Help System.*** Copyright © 1996 by Microsoft Corporation, Redmond, Washington, USA.

2

PRODUCT OVERVIEW

Product Description

Performance Gallery Gold is an exceptional performance monitoring and reporting solution for HP e3000 systems (MPE/iX), HP 9000 systems (HP-UX), Sun® systems (Solaris 2.6/Sun OS 5.6), and Windows NT/2000 networks, servers, and workstations. Performance Gallery Gold enables system administrators to track how their system handles specific aspects of its workload, including:

- Processor utilization
- Memory utilization
- Demand paging
- Disk performance
- Cache performance
- Network traffic

Performance Gallery Gold analyzes, then reports system performance data collection sets in full-color graphs and tables, which make it easy to spot system bottlenecks and monitor overall system performance and usage.

Performance Gallery Gold is designed for use by system administrators. However, the usefulness of this product will benefit an entire organization, from online users to top management. It is an ideal tool for businesses that want to make informed decisions about their MPE/iX, HP-UX, Sun Solaris 2.6/Sun OS 5.6, or Windows NT systems.

New Features (F.04)

The following new features and enhancements are included in this version of Performance Gallery Gold.

New Continuous Update Features

- Ability to communicate in Continuous Update mode with Meta-View Agents.

- Improved responsiveness of Performance Gallery Gold when running in Continuous Update mode for more than one Meta-View Agent by using distinct data retrieval threads.
- Ability to run in Continuous Update mode against Meta-View Agent for Windows.

New Chart Features

- Connect the label (hover stats) and its corresponding data point with a line in 3D charts.
- Table columns can now be resized manually by dragging them with the mouse.
- A new workbook for table charts was developed.
- Hitting the ENTER key while in the Template Properties or Global Chart Options dialog generates "apply changes".
- Ability to move the mouse over a point on a graph chart and see what the numbers are behind the data point.

Enhanced File Handling

- The window size and position on the Desktop are saved on exit and restored when restarted.
- Ability to open historical data directly from running Meta-View Agents.
- When running on Windows XP, program will use Visual Styles.

New Licensing Feature

- The need for licensing and demo codes has been removed in keeping with our other client software packages. Licensing will remain in Meta-View Agent.

Components

Data Files

A ***data file*** is a collection of information about a computer system that has been gathered at time intervals specified by the user. This file usually includes information about the processor, memory, and disk statistics, and it may include process information as well.

Performance Gallery Gold for Windows NT systems utilizes the following data collection files:

- *.pfg (Performance Gallery Gold) data files collected by Meta-View Performance Manager (HP-UX, Sun Solaris 2.6/Sun OS 5.6, and MPE/iX).
- *.smf (System Management Facility) data files collected by the Performance SeNTry collection agent (Windows NT/2000).

Charts

A ***chart*** is a visual means of representing data. With Performance Gallery Gold, a chart can be constructed in the form of a table, a two-dimensional graph, or a three-dimensional graph. Tables can be either "time-indexed" (to show values over time) or a "snap shot" (to show values at a specific time). Data displayed in two-dimensional (2-D) graphs can be line, area, bar or pie graphs. The data items in 2-D graphs, with the exception of pie graphs, can be stacked. Three-dimensional graphs can be either bar or surface graphs.

Please refer to “Graph Properties” on page 116 and “Table Properties” on page 132 for more information about charts.

Desktops

A ***desktop*** is the current state of the Performance Gallery Gold display area, which includes the size and position of the main program window and all of the chart windows currently open in the application, saved to disk as a Performance Gallery Gold desktop file (*.dsk). The chart types included on the desktop are stored in such a way that specific charts are retained, (such as CPU Utilization) as well as any added chart titles and global chart options. By saving and loading a desktop, a previous session can be quickly and easily restored with minimal work duplication. One desktop may also be designated as the ***default desktop***, to be used whenever a data file is initially loaded and no charts are currently open.

Please refer to “Desktop Menu” on page 100 for more information about desktops.

Templates

A Performance Gallery Gold chart ***template*** specifies the design and contents to be applied to a specific graph or table. It includes all non-data information about a chart, such as the chart title, the chart type, the y axis scale, any limits that should be applied, the data elements to be plotted, etc.

Please refer to “Template Menu and Commands” on page 109 for more information about templates.

Host Data Collectors

Meta-View Performance Manager for MPE/iX Systems

MPE/iX performance data is logged in batches by the Meta-View Performance Manager software from Lund Performance Solutions. This application collects, extracts, and exports performance data to Performance Gallery Gold as a *.pfg (Performance Gallery Gold) data file.

Please refer to the *Meta-View Performance Manager - User’s Guide to Meta-View Agent and Host for MPE/iX* for information.

Meta-View Performance Manager for HP-UX Systems

HP-UX performance data is collected either continuously or in batches by Meta-View Performance Manager from Lund Performance Solutions. This application uses two programs, MVLOGD and MVLOGX, to log and export performance data records (*.pfg data files) to Performance Gallery Gold for numerical and graphical analysis.

Please refer to the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for HP-UX* for information.

Meta-View Performance Manager for Sun OS Systems

Sun Solaris performance data is collected either continuously or in batches by Meta-View Performance Manager from Lund Performance Solutions. This application uses two programs, MVLOGD and MVLOGX, to log and export performance data records (*.pfg data files) to Performance Gallery Gold for numerical and graphical analysis.

Please refer to the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for Solaris* for information.

Performance SeNTry for Windows NT Systems

Data files from Windows NT/2000 systems, *.smf files, are collected using Performance SeNTry, a Windows NT performance data collector from Demand Technology Software (see "Demand Technology Software Product Support" on page 5). Performance SeNTry collects a variety of Windows NT performance objects and counters using standard Win32 API calls. The program calculates object counter values based on the counter type and writes them at regularly scheduled intervals to an ASCII format, comma delimited data file. The *.smf data files can be analyzed and reported with Performance Gallery Gold.

Please refer to the *Performance SeNTry User's Guide* for information.

Procedural Summaries

MPE/iX Systems

Collecting, Extracting, and Exporting MPE/iX Data

1. Create and configure the workload definition file (see "Exporting MPE/iX Performance Data" on page 20). This step is optional—it is not required to run the collector job. This step can be ignored or performed at another time.
2. Stream the data collector job (see "Collecting MPE/iX Performance Data" on page 17).
3. Extract the performance data via a job stream or an interactive (online) terminal session (see "Extracting MPE/iX Performance Data" on page 18).

4. Export the data records from the MPE/iX to the PC (see “Extracting MPE/iX Performance Data” on page 18).

Reporting MPE/iX Data

1. Load an MPE/iX data collection file (.pfg) (see “Open Data File” on page 76).
2. Open a pre configured MPE/iX chart (see “Open Chart” on page 149).
3. Interpret the chart data with the help provided in “MPE/iX System Performance” on page 173.
4. Modify the chart to pinpoint specific activity or to improve the chart’s readability (see “Menus and Commands” on page 75 and “Template Menu and Commands” on page 109).
5. Compare the chart to related performance chart(s) using preset links (see “Links Properties” on page 144).
6. Save the open chart(s) as a desktop file (see “Save Desktop” on page 102).
7. Print the chart(s) (see “Print All” on page 83).
8. Export the chart(s) (see “Export” on page 86).

Unix Systems

Collecting, Extracting, and Exporting Unix Data

1. Create and configure the workload definition file (see “Exporting Unix Performance Data” on page 29). This step is optional—it is not required to run the collector job. This step can be ignored or performed at another time.
2. Start the batch logging process (see “Unix Performance Data” on page 23).
3. Extract the performance data (see “The Historical Performance Data Extraction Utility” on page 27).
4. Export the data records from the Unix to the PC (see “The Historical Performance Data Extraction Utility” on page 27).

Reporting Unix Data

1. Load a *.pfg data collection file (see “Open Data File” on page 76).
2. Open a pre configured Unix chart (see “Open Chart” on page 149).
3. Interpret the chart data with the help provided in “Unix System Performance” on page 203.
4. Modify the chart to pinpoint specific activity or to improve the chart’s readability (see “Menus and Commands” on page 75 and “Template Menu and Commands” on page 109).
5. Compare the chart to related performance chart(s) using preset links (see “Links Properties” on page 144).
6. Save the open chart(s) as a desktop file (see “Save Desktop” on page 102).
7. Print the chart(s) (see “Print All” on page 83).

8 Export the chart(s) (see “Export” on page 86).

Windows NT/2000 Systems

Collecting Windows NT/2000 Performance Data

Refer to the *Performance SeNTry User's Guide* for more information regarding the following procedures.

1 Define the data collection set and its parameters.

2 Assign the collection set to the Windows NT/2000 systems you want to monitor.

3 Activate the collection set.

Reporting Windows NT/2000 Performance Data

1 Load an *.smf data collection file (see “Open Data File” on page 76).

2 Open a pre-configured Windows performance chart (see “Open Chart” on page 149).

3 Interpret the chart data with help provided by Mark Friedman of Demand Technology Software (see “Windows NT/2000 System Performance” on page 225).

4 Modify the chart to pinpoint specific activity or to improve the chart's readability (see “Menus and Commands” on page 75 and “Template Menu and Commands” on page 109).

5 Compare the chart to related performance chart(s) using preset links (see “Links Properties” on page 144).

6 Save the open chart(s) as a desktop file (see “Save Desktop” on page 102).

7 Print the chart(s) (see “Print All” on page 83).

8 Export the chart(s) (see “Export” on page 86).

3

PRODUCT STARTUP

Starting Performance Gallery Gold

From the Windows Start button, select **Programs**, **Performance Gallery Gold**, **Performance Gallery Gold** to launch the Performance Gallery Gold program.

Accessing the Online Help Facility

Help Topics

To access the online indexed Help, select **Help Topics** from the **Help** menu. Select a book or page from the Help **Contents** tab, or use the **Index** or **Search** tabs. You can also add frequently-accessed topics to the **Favorites** tab. To add a favorite topic:

1. Open the topic in the browser window by using the **Contents** or **Index** tabs.
2. Click the **Favorites** tab.
3. Click **Add**.

The next time you open **Help Topics**, that topic will be saved in the **Favorites** tab.

Context-sensitive Help

To access context-sensitive Help on a specific dialog box, click the **Help** button in the dialog box. Help about the dialog box will display in a separate window.

To get context-sensitive Help elsewhere in Performance Gallery Gold, do either of the following:

- Position your mouse over the item in question and press the F1 function key. This is especially helpful for open dialog boxes that do not contain a **Help** button, as the Help Toolbar button will not open dialog boxes and cannot be selected if a dialog box is already open.

  OR

- Select the Help toolbar button  for use with toolbar buttons or menu items. This will change the cursor to a question mark. Move the question mark pointer over an item and click to get Help for that item.

NOTE The online Help was designed with and for Internet Explorer v4.0. Using an older version of IE4 or a different browser may result in the pages displaying improperly.

Exiting Performance Gallery Gold

To exit the Performance Gallery Gold program, select **Exit** from the **File** menu.

4

DATA COLLECTION, EXTRACTION, AND EXPORT

MPE/iX Performance Data

Collecting MPE/iX Performance Data

Data collection on MPE/iX can be accomplished with or without the Meta-View Performance Manager for MPE/iX software from Lund Performance Solutions.

Streaming the Data Collector Job with Meta-View

24-hour Collection

To establish 24-hour continuous data collection, do the following:

From the MPE/iX terminal, type **:Stream mvmonj.pub.lps**

8-hour Collection

To establish 8-hour continuous data collection, do the following:

From the MPE/iX terminal, type **:Stream mvjob.pub.lps**

Streaming the Data Collector Job without Meta-View

24-hour Collection

To establish 24-hour continuous monitoring, do the following:

1. From the MPE/iX terminal, view the **perfcolj.pub.lps** file.
2. Verify the **Duration of job in minutes** value is **0** (zero). If it isn't set to **0**, use a text editor to change the value.
3. From the MPE/iX prompt, type **:Stream perfcolj.pub.lps**.

These changes will cause the **perfcolj** job to automatically restream itself every 24 hours at midnight. If your backup jobs kill this job, you will need to restream it as part of a daily routine.

8-hour Collection

To establish 8-hour continuous data collection, do the following:

1. From the MPE/iX terminal, view the **perfcolj.pub.lps** and verify that the **Duration of job in minutes** value is **480**. If it isn't set to **480**, use a text editor to change the value.
2. Comment out the **Stream perfcolj** line near the bottom of the file. Do not remove the line.
3. From the MPE/iX prompt, type **:Stream perfcolj.pub.lps**.

Changing the Sampling Interval Length

By default, the performance data sampling interval for all data collection job streams is 300 seconds (five minutes). Every five minutes, an average of all performance indicators for the current cycle is recorded in the log file.

Extracting MPE/iX Performance Data

The Performance Gallery Gold program performs MPE/iX data extraction in two steps:

1. The program looks at the collected data and selects only the records that meet the criteria set by the user.
2. It creates a file to be downloaded to the PC.

Data extraction can be accomplished via a job stream or an online (interactive) session with Meta-View Performance Manager software by Lund Performance Solutions.

Extracting MPE/iX Data Using a Job Stream

Set the following parameters in the **mvlogxj.pub.lps** file (or the **perfredj.pub.lps** file), as prompted:

- Enter the name of the MPE file designated to hold the extracted data.
- Enter the date (mm/dd/yy) to start the data extraction.
- Enter the time (hh:mm) to start the data extraction.
- Enter the date (mm/dd/yy) to end the data extraction.
- Enter the time (hh:mm) to end the data extraction.

Extracting MPE/iX Data Using an Online (Interactive) Session

1. From the MPE/iX terminal, type **Run mvlogx.pub.lps**.
2. Set the following parameters, as prompted:
    - Enter the start date (mm/dd/yy) of the desired data sample. Press the **Enter** key. By default, the program will retrieve the start date of the earliest sample record.
    - Enter the start time (hh:mm) of the data sample. Press the **Enter** key. By default the program will retrieve the start time of the earliest sample record.

NOTE Perform steps 3-6 only if there is a need to modify the current parameters. Other wise, skip to step 8.

3 From the MVLOGX log screen in Meta-View Performance Manager, select **OPTIONS** (F1).

4 If you want to average your data files, select **5) Data break configuration menu**.

Edit **1) Duration in Minutes of each Sample**.

Changing this value will average your data before you export it to Performance Gallery Gold. For instance, if you set **percolj** to sample every 5 minutes, and you type 10 into **Duration in Minutes of each Sample**, then samples will still be taken every 5 minutes, but those sample will then be averaged into 10 minute files for downloading to Performance Gallery Gold.

5 From the MVLOGX Main Options menu, select **11) Performance Gallery configuration (SUBMENU)**. Ensure that **option 1 - Performance Gallery Version i**s set to **2 - Gold.**

6 From the Performance Gallery configuration submenu, adjust parameters as needed:

- Enter the number of the parameter you wish to change.
- Type **N** (No) to disable the parameter.
- Type **Y** (Yes) to enable the parameter.
- Type a new extract file limit.

7 When the options are configured as needed, select **EXIT OPTIONS** (F8).

8 Select **UTILITY KEYS** (F5).

9 Select **PERFORM GALLERY** (F7).

10 Set the following parameters, as prompted:

- Enter the file name of the Performance Gallery Gold data file to be downloaded to the PC.
- Enter the date (mm/dd/yy) to start the data collection, or accept the date shown in parentheses.
- Enter the time (hh:mm) to start the data collection, or accept the default start time shown in parentheses.
- Enter the date (mm/dd/yy) to end the data collection, or accept the end date shown in parentheses.
- Enter the time (hh:mm) to end the data collection, or accept the default end time shown in parentheses.

11 The program will display "Working_" as it scans the appropriate log files and writes the records to the file you created.

The file will be placed in the same group in which the user is logged.

The extraction may take a few moments to several minutes to complete, depending on the size of the file and the speed of the processor.

12 Select **MAIN KEYS** (F8).

13 Select **EXIT MVLOGX** (F8). This step completes the data extraction process.

Exporting MPE/iX Performance Data

The transfer of data from the MPE/iX system to the PC requires:

- A PC connected to the host MPE/iX system.
- FTP (file transfer protocol) or terminal emulation software such as Reflections, Business Session for Windows, MiniSoft32, Or AdvanceLink.

Refer to the documentation provided with your terminal emulation software for instructions to download from the MPE/iX system to the PC ***using binary format***. The name of the file on the PC should include the extension ".pfg." The extension can be added during the download process, or the file can be renamed after it is transferred to the PC.

Establishing MPE/iX Workloads

Performance Gallery Gold offers several chart types that show performance measurements by workload. Workloads, or applications, are groups of related programs and/or online terminal sessions. You can use the three pre-defined MPE/iX workloads or create your own workloads definition file (see "Creating an MPE/iX Workload Definition File" on page 20).

Predefined MPE/iX Workloads

Performance Gallery Gold uses three predefined (default) workloads:

- JOB (includes all batch job streams)
- SESSION (includes all interactive (online) terminal sessions)
- SYSPROCS (includes all system processes)

If you prefer to use more specific workloads, you may create a workload definition file (see the next section, "User-defined MPE/iX Workloads.").

User-defined MPE/iX Workloads

Defining specific workloads is a good way to track specific processes, such as processes run by a specific department in your business. However, user-defined workloads are optional—they are not required to run Performance Gallery Gold.

Creating an MPE/iX Workload Definition File

User-defined workloads are created in the **mvwkdef.pub.lps** file. Use your editor to create an **mvwkdef** file (**quad.util.lps** is included on the distribution tape).

The basic format of the **mvwkdef** file requires three items for each workload:

- The name of the workload (up to 10 characters).
- The type of the processes included in the workload (JOB, SESSION, SYSPROCS, or BOTH).
- The user and/or program specifications (a list of one or both of the following):
    - USER (the account by which MPE knows you (your user or login ID))
    - PROG (the MPE fully-qualified program file name)

Workload Definition File Configuration Rules

1. A workload name of up to 10 characters is required.
2. A workload type specification is necessary to indicate which types of processes to include or exclude from the workload definition. This makes it possible to create two workloads for processes that run in both interactive and batch modes. For example:
    - JOB includes only batch processes.
    - SESSION includes only interactive (online) terminal sessions.
    - SYSPROCS includes only system processes.
    - BOTH includes both job and session processes, but not system processes.
3. Workloads must be separated by one or more blank lines within a definition file.
4. Comments can be included on any line if preceded by an exclamation point (!).
5. Either a user or a program specification is required. The specification must be entered one per line and must be one of the following three types:
    - A program specification (PROG=PROGRAM.GROUP.ACCOUNT).
    - A user specification (USER=SESSION, USER.ACCOUNT,GROUP). The session name and logon group are optional.
    - The MPE logical device number (LDEV=nnn) or range of device numbers (LDEV=nnn-nnn).

   The symbol "@" can be used as a wildcard for any of the criteria, just as it can with normal MPE/iX rules (partial or full).

NOTE The LDEV specification means that you can capture activity on a terminal-by-terminal basis, or even within a range of terminals. Use this option carefully!

6. There is virtually no limit to the number of user, program, and LDEV specifications allowed for each workload group.
7. The name and type of specification line is required. All other lines are optional. In order to be considered part of a workload group, a process must satisfy the program, user, and LDEV specifications, if all three are present.

- If more than one program specification lines are included, a program needs to satisfy only one of the program specifications to be included in the group.
- If no program specifications are entered, all process programs are considered to be in the group, unless the process is somehow disqualified by the user or LDEV specification.
- The user and LDEV specifications are resolved in the same way. For instance, the lines in the following example should be entered into the **mvwkdef** file to define the workload called "WORKTEST."

```
WORKTEST                    !Workload name
SESSION                     !Only terminals
PROG=@.PUB.MFG
PROG=MONEND.PUB.QTR
USER=JANE, MGR.MFG
USER=@,@.QTR
LDEV=50-60
```

Figure 4.1 *Workload definition file (example)*

For a process to be included in the WORKTEST workload, it must satisfy only one program specification, one user specification, and one LDEV specification. Each is considered to be an "and" condition. For instance, a program, INVEN01.PUB.MFG, run by MGR.QTR at LDEV 56 would be counted in the WORKTEST workload.

8 When selecting any of the following system-type processes, the program group and the account must be specified as "@."

- Specify command interpreter processes by the program file name "ci" (PROG=ci).
- Specify spooler processes by the program file name "sp" (PROG=sp).
- All other system processes can be identified by name.

To strip out command interpreters from the catch-all sessions workload, create a separate workload with the program name "ci" to track the response times for sessions. It will reflect what the users actually experience.

9 A process can belong to only one workgroup. If a process meets the criteria for two or more groups, it will be assigned to the first workload in the file for which it qualifies.

Unix Performance Data

Collecting Unix Performance Data

Performance Gallery Gold uses data records gathered in batches by the Meta-View Performance Manager for HP-UX, and the Meta-View Performance Manager for Solaris, both by Lund Performance Solutions.

Meta-View Performance Manager for both HP-UX and Solaris comprises the following components:

1 **Meta-View Agent**

Meta-View Agent is the succesor to Lund's SOS product. It retains the powerful collectors and host-based capabilities of SOS on HP-UX, Linux and Solaris computers. Meta-View Agent uses the following processes to collect, log and serve performance data from the host systems to the Meta-View Host, Web and Mobile clients.

- Meta-View Data Collector - MVMID
- Meta-View Log Daemon - MVLOGD
- Meta-View Log Extractor - MVLOGX
- Meta-View Data Daemon - MVDATAD
- Meta-View Host Client - MVHOST

2 **Meta-View Host**

The Meta-View Host client provides system administrators with comprehensive, real-time and historical performance data in detailed ASCII terminal displays.

Meta-View automatically collects and presents performance data in real time for monitoring, or it can show an "instant replay" of historical performance data to aid in trobleshooting and trending analysis.

For more information about the Meta-View Performance Manager for HP-UX or the Meta-View Performance Manager for Solaris software, please refer to the product user manual or contact your account manager at Lund Performance Solutions.

The Historical Performance Data Logging Utility

The Meta-View Performance Manager application suite includes a data logging utility called MVLOGD. MVLOGD enables the user to collect historical system performance data for analysis of performance problems and trends. The data is collected and stored in SL (system log) files for later use by MVHOST or MVLOGX.

Data can be logged three ways:

- A single, one-time-only, session. The length of the data collection period is determined by the user.

- Repeatedly, by retreaming the collector. The logging job must be stopped manually with the **lpskill** command.
- Scheduled, using the **cron** facility. For example, Monday through Friday, 6:30 AM to 6:30 PM.

SL Files

MVLOGD creates one logical file record for every batch interval. The default interval is 10 minutes (600 seconds). The log file is saved in the Meta-View Performance Manager log directory and named using the format *SLyyjjjs*.

- *SL* represents the Meta-View log file.
- *yy* represents the current year.
- *jjj* represents the Julian day of the year.
- *s* represents the sequence of the log (up to 26 characters, from a through z).

Starting the Logging Process

To begin the collection process, type **mvlogd** at the shell prompt of your home directory. It is assumed that your path statements are set up properly.

Configuring MVLOGD Run Time

By default, MVLOGD will run, collecting host data, until midnight (23:59). The run time can be configured tu run repeatedly or at specific intervals by using either the **-c** command line switch or the **cron** facility.

-c Command Line Switch

Enter **mvlogd -c** at the shell prompt of your home directory to modify the MVLOGD run time configuration. The effects of the command line switches vary depending on whether you are using the default or advanced configuration parameters.

cron

The -c switch should not be used in conjunction with cron. cron should be used if the user wants MVLOGD to run only during specific hours. For example, to run MVLOGD from 8AM to 6PM, configure it to run 10 hours, then configure a cron job to start it every day at 8AM. The functions of each command line switch is summarized in the following table.

Table 4.1 *MVLOGD command line switch functions*

Switch	Default Configuration	Advanced Configuration
-c	Log continuously (24 hours) by restarting at 00:00 hours.	Log until the Run Time value has expired, and then automatically restream.

Switch	Default Configuration	Advanced Configuration
-h	Display all available command line switches.	
-o	Display the default configuration.	Display the configuration parameters in the .mvlogdrc file.

Viewing Default Configuration Parameters

To view the default configuration parameters for your system, type **mvlogd -o** from your home directory. The default parameters are described in Table 4.2.

Table 4.2 *MVLOGD default configuration parameters*

Parameter	Description
Enter duration of job in minutes (0)	The collector will run until midnight. When used with the -c command switch, the collector will automatically restream itself immediately.
Interval time in seconds (600)	The program will take a sample and write a log record every 10 minutes.
Company name ( )	The company name is blank, by default.
Display advice messages (Y)	MVLOGD will display advice messages.

Setting Advanced Configuration Parameters

To create a custom configuration parameter file:

1. Create a custom file (.mvlogdrc) with your editing program, listing the parameters as described:
   - RunTime

     The amount of time (minutes) Meta-View Performance Manager will monitor your system's processes. A RunTime of "0" means "until midnight."
   - CycleTime

     The amount of time (seconds) between samples.
   - CompanyName

     Your company name (added to the title of each log report). (This can be the name of your system or another subheading, if desired.)
2. Place the .mvlogdrc file in your home directory to enable batch logging parameters.

3 To change the configuration, edit the parameters in the .mvlogdrc file. For example:

- To collect data in eight-hour batches, change the batch run time to 480 minutes by typing **RunTime=480**.
- To shorten the interval time to five minutes (300 seconds), type **CycleTime=300**.
- To add the name of your company (or another subheading) to the title of each log report, type **CompanyName=<your company's name>**.

Configuration Variables

The .mvlogdrc configuration variables are outlined in the following table.

Table 4.3 *MVLOGD default configuration variables*

Variable	Type	Min	Max	Default	Description
RunTime	integer	0	1440	0	Duration of job in minutes
CycleTime	integer	10	3600	600	Interval time in seconds
CompanyName	string	N/A	N/A	<blank>	Company name
DisplayAdvice	Y/N	N/A	N/A	Y	Display advice messages
ProcLog	Y/N	N/A	N/A	Y	Log processes
ProcCPUThreshold	integer	0	100	0	CPU percentage required for process display
LogOnlyActProc	Y/N	N/A	N/A	Y	Log only active processes
LogInteractProc	Y/N	N/A	N/A	Y	Display attached processes
LogNonInteractProc	Y/N	N/A	N/A	Y	Log non-interactive processes
LogDeadProc	Y/N	N/A	N/A	Y	Log processes that died
ProcLogonFilter	reg exp	N/A	N/A	.*	Process logon filter

Variable	Type	Min	Max	Default	Description
ProcSortOpt	integer	1	8	4	Process sort options: 1 - sort by PID# 2 - sort by Logon Terminal 3 - sort by Workload group 4 - sort by CPU time 5 - sort by Disc I/O 6 - sort by Term reads 7 - sort by Priority 8 - sort by Wait State
ProcSortAscend	Y/N	N/A	N/A	N	Log processes sorted in ascending order
ProcLogLimit	integer	1	127	10	Maximum number of processes to be logged per interval

The Historical Performance Data Extraction Utility

MVLOGX is the historical data counterpart to MVHOST. It provides the means for reviewing performance data stored in the log files that MVLOGD has collected. The user interface is similar in many ways to MVHOST. The main difference is that the MVLOGX screens do not display current samples of online performance data. Instead, they display historical data collected by MVLOGD.

The primary functions of MVLOGX are:

- To browse through the data recorded in your log files using a variety of screen reports. This is usually done to identify periods of system activity that may require further analysis.
- To prepare logged performance data from the log files for Performance Gallery Gold.

Getting Started

To run MVLOGX:

1 From your home directory, enter **mvlogx** (lowercase) at the prompt. The initial MVLOGX screen will display (see Figure 4.2)

NOTE To view the screen without soft function keys, add the -k command key switch (type mvlogx -k at the prompt).

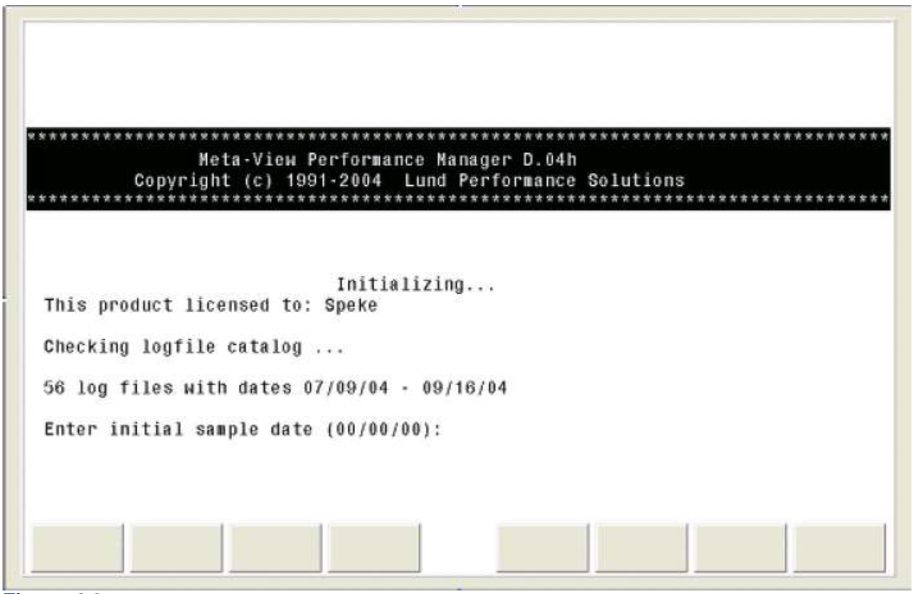

Figure 4.2 *MVLOGX Initial Screen*

2 When MVLOGX is run for the first time on a system, it creates a catalog of the SL files located in the working directory. The catalog is saved as a file (.sllogcat) in the same location as the SL files. At subsequent startups, MVLOGX will check to see if the catalog exists. It will create a new catalog if the previous catalog cannot be found.

3 Enter the following information as prompted:

   a The start date of the initial sample using the format **mm/dd/yy**. If you do not provide a date, the program will retrieve the earliest sample date recorded by default. Press **Enter**.

   b The start time using the format **hh:mm**. If you do not provide a start date, the program will retrieve the start time of the earliest sample data recorded, by default.

The first MVLOGX data report, the **CPU Summary** screen, will display. Information about each MVLOGX report is provided in the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for HP-UX.*

Exporting Unix Performance Data

The transfer of data from the Unix system to the PC requires the following:

- A PC connected to the host system.
- FTP (file transfer protocol) or terminal emulation software such as Reflection, Business Session for Windows, MiniSoft32, or AdvanceLink.

Refer to the documentation provided with your terminal emulation software for instructions to download from the Unix to the PC using ***binary data format***. The name of the file on the PC should include the extension ".pfg." The extension should already be automatically added to the file, or the file can be renamed after it is downloaded to the PC.

Exporting Data to Performance Gallery Gold

Setting the Performance Gallery Configuration

If needed, you can change the configuration of the export file with the options provided in MVLOGX.

1. From any MVLOGX report display, type **o** (lowercase) to access the **MVLOGX MAIN OPTION MENU**.
2. Select the **Performance Gallery configuration (SUBMENU)** option.
3. Select the **Export Data configuration (SUBMENU)** option.
4. Check the configuration of the export data. Make modifications as needed. Press the Enter key to exit the submenu.
5. From the Performance Gallery configuration submenu, select the **Export Thresholds configuration (SUBMENU)** option.
6. Set the export thresholds. (The options in the **Export Thresholds configuration** submenu are described in the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for HP-UX*.) Press the Enter key to exit the submenu.
7. Press the Enter key to exit the **Performance Gallery configuration** submenu.
8. Press the Enter key again to exit the **MVLOGX MAIN OPTION MENU**.

Creating the PGG Export File

From any MVLOGX report display, type **P** (uppercase) to start the process. Respond to the following prompts:

- Enter Performance Gallery export file (the .pfg file extension will be automatically appended)
- Enter start date for Performance Gallery (mm/dd/yy)

- Enter start time for Performance Gallery (hh:mm)
- Enter end date for Performance Gallery (mm/dd/yy)
- Enter end time for Performance Gallery (hh:mm)

MVLOGX will read the log files that meet the date and time criteria entered, then write the eligible data to the specified export file. If a directory path is not given for the export file, the file will be written to the current working directory.

1. The program will display "Working_" as it scans the appropriate log files and writes the records to the export file you created.
   - The file will be placed in the current working directory.
   - The extraction may take a few moments to several minutes to complete, depending on the size of the file and the speed of the processor.
2. Select **MAIN KEYS** (F8).
3. Select **EXIT MVLOGX** (F8). This step completes the extraction process.

Establishing Unix Workloads

A workload is a group of similar, identifiable transactions on the host system performed by individual users and programs. Workloads can be grouped by:

- Applications
- User login
- Departmental processes

A workload may be as simple as one user running one program, or as complex as entire departments running many programs.

Identifying and Characterizing Workloads

Make sure workloads are homogeneous. A homogeneous workload consists of processes of a similar type, function, and priority.

Averaging is meaningless for workloads made up of dissimilar transactions. For example, if an average accounts receivable transaction takes 200 milliseconds of the CPU's time, while general ledger transactions average 500 milliseconds, taking an average of the two does not provide a meaningful average for either transaction.

Identifying Workloads

Input from management and system users is essential in identifying and defining workloads. Interview managers and users to determine how the system is used and to identify distinct functions, such as order entry, telemarketing, or accounting. Break down the various departmental functions into essential components, based on your desired result. Identify groupings that will provide you with the needed information. These grouped components make up your workloads.

Characterizing Workloads

Once you have identified your workloads, use the following guidelines to further refine your workload definitions:

1 Limit the components of any workload to users or transactions with service demands of comparable magnitude and similar balance across the system. Do not mix heavy-CPU/low-I/O transactions with light-CPU/heavy-I/O transactions.

2 Do not mix interactive processes and batch processes in the same workload. System resources, priorities, and think times are different for interactive and batch processes.

3 Use separate workloads for specific divisions, branches, or departments as needed.

4 Identify workloads by user logon, if possible.

Creating a Workload Definition File

Once you have identified and refined you workloads, enter the data in a workload definition file.

Workload Definition File

User-defined workloads are created in /opt/lps/lib/workdefs.

Workload Groups

Four workload groups are defined by default. These four workloads should always exist.

Table 4.4 *The default workload groups*

Workload	Description
INTERACT	The INTERACT workload contains any processes attached to a terminal (interactive processes). The INTERACT workload group should be configured by the user.
DAEMON	The DAEMON workload contains any daemon processes. By default, this workload group is configured to include any process not attached to a terminal and owned by the root user. The DAEMON workload group should be configured by the user to reflect the system.
BATCH	The BATCH workload contains any batch job processes. By default, this is configured to include any process outside of the DAEMON workload group that is not attached to a terminal. The BATCH workload group should be configured by the user to reflect the system.

Workload	Description
DEFAULT	The DEFAULT workload contains any process that does not match any other workload definition. Note that initially, this will be an empty workgroup (no processes will match), because at least one of the other defaults will include any possible process. However, since those workload groups are configurable, this workload group must exist. The DEFAULT workload cannot be modified. It guarantees a process will fall into at least one workgroup by matching any process that does not fall into any other workgroup definition.

Response Time Calculations

Response time calculations are performed for processes, workload groups, and on a system-wide basis. They include:

- Transactions
- Response time
- Average response time per transaction

The definitions of *transaction and response time* are dependent upon the type of workload group to which the process belongs. The response time calculations are described below.

INTERACT Response Time Calculations

- Transactions = Terminal reads + Process deaths
- Response time = Process live time - Think time

Process live time is the amount of time a process is alive during the interval. If the process was created before the interval and doesn't die until after the interval, process live time will equal the interval time.

Think time is defined as the amount of time a process is waiting for user input.

DAEMON Response Time Calculations

- Transactions = Voluntary context switches + Process deaths
- Response time = Process live time - Process resource wait time

Process resource wait time is the amount of time a process is waiting on resources (such as CPU) verses waiting on an event (such as a terminal input).

BATCH Response Time Calculations

- Transactions = Process deaths
- Response time = Process live time

Mix Response Time Calculations

The mix response times vary, based on whether the process is attached to a terminal or not.

- If the process is attached to a terminal, the INTERACT response time calculations are used.
- If the process is not attached to a terminal and it has a high Nice value, the BATCH response time calculations are used.
- If the process is not attached to a terminal and it has a default or low Nice value, the DAEMON response time calculation is used.

Workload Definition Requirements

The workdefs file requires the following information for each workload:

1. The name of the workload, up to ten characters.
2. The type of process or processes included in the workload, such as INTERACT, DAEMON or BATCH.
3. The user or program specification, including one or more of the following:
    - USER (your user ID or logon ID)
    - PROG (the name of the executable program file)
    - TTY (the device name of your terminal)
    - GROUP (the user group identification)

Workload Definition File Configuration Guidelines

Use the following guidelines to create or edit workload definition files:

1. Separate workloads by one or more blank lines.
2. Include comments on any line, if desired, preceded by an exclamation character (!).
3. A workload-type specification is needed to indicate the types of processes to include or exclude from the workload definition. This makes it possible to create two workloads for processes that run in both interactive and batch modes. (Refer to Table 4.4 on page 31for more information.)
4. Program and user specifications are specified by:
    - PROG=program name.
    - USER=user name/group name. System group names are valid specifications. Check the /etc/group file for a list of existing group names.

    For more information about group names, refer to the Unix manpage, *regexp* (Regular Expressions).
5. Device file specifications, such as TTY=tty0p2, are also valid. You can capture activity on a terminal-by-terminal basis, or for multiple terminals.
6. There is no limit to the number of user, program, and tty specifications allowed for each workload.

7 Name and type specification lines are required. All other lines are optional.

8 To be included in a workload group, a process must satisfy the program, user, and tty specifications, if all three are present.

- If one or more program specification lines are included, a program needs to satisfy only one of these to be included in the group.
- If no program specifications are entered, all process programs are included in the group, unless the process is somehow disqualified by the user or tty specifications.

9 A process can belong to only one workload group. If it fits the criteria for two or more groups, it is assigned to the first workload in the file for which it qualifies.

10 Four workloads appear by default: INTERACT, DAEMON, BATCH, and DEFAULT. Processes that do not fit into user-defined workloads will be included in one of these pre-defined workloads.

Windows NT/2000 Performance Data

Performance SeNTry, from Demand Technology Software, can be used to collect performance data from Windows NT/2000 servers and workstations. Please follow the directions provided in the Performance SeNTry User's Guide to collect and store 9x/NT/2000 performance data.

Collecting Windows NT/2000 Performance Data

Performance SeNTry data collection files are generated and assigned names by the host collection service. The data file name (systemname.yyyymmddhhmm.smf) denotes the following information about the data file:

- The name of the system (systemname)
- The year the data collection cycle started (yyyy)
- The month the data collection cycle started (mm)
- The day the data collection cycle started (dd)
- The hour and minute the data collection cycle started (hhmm)

The Performance SeNTry data files are saved with an SMF filename extension in the default destination folder, **C:\Ntsmf\data**.

5

PERFORMANCE GALLERY GOLD QUICK TOUR

Overview

The Quick Tour in this chapter introduces first-time users to the features and functions of Performance Gallery Gold for Windows 2000/NT, MPE/iX, and Unix. For users upgrading from earlier versions of the software, these tours provide a quick and easy way to get acquainted with the changes provided in the current version.

Before beginning the tour, install Performance Gallery Gold along with the appropriate sample files on your personal computer as described in “Product Startup” on page 15.

The Charts

We will open four charts during this tutorial which are designed to reveal different aspects of your system’s performance. A single Performance Gallery Gold chart cannot tell you everything you need to know about your system performance. For instance, a high Processor Queue Length, at first glance, may indicate a CPU bottleneck and hence, a need for a faster processor. But perhaps the CPU Utilization never surpasses 50%. Obviously, the CPU is not being overworked. In this case, it is likely that several processes or services were started all at the same time. Purchasing a faster processor would not help in this situation; it would be a waste of your company’s money.

CPU Utilization

This chart displays how much of your Central Processor Unit (CPU) is being used. If this metric reaches 85%, the system will begin to slow down noticeably. If the CPU Utilization remains at 85% for extended periods of time, or spikes at certain times every working day, it can be an indication of a few things:

- Too many users for the CPU. You may need to upgrade to more or faster processors.
- You need to schedule batch and low priority processes or services to run at times (at night, for instance) when your users will not be using the system, thereby spreading the CPU usage around the clock.
- You may need to add more memory.

For a more detailed look at CPU Utilization, see “CPU Utilization by Process” on page 228.

Logical Disk Response Time

This chart displays how long it takes from the time a request for data (read or write) is made to the time the data is read from or written to the disk. Most requests can be satisfied within 11 ms (0.011 sec). If your response times are consistently above 0.025 seconds, it may indicate the need to:

- Increase memory. If the data is found in memory, then there is no need to access the disk, and the response time will improve.
- Upgrade to a faster disk. Faster disks transfer data faster.
- Defragment your hard disk. If the data on your disk is spread throughout the disk, it will take longer to find the data you are interested in.
- Invest in other hardware solutions. See "Logical Disk Response Time" on page 245 for more information on these solutions.

File Cache Activity by Type

This chart displays how many times per second a request for data was satisfied from the memory or cache, (also known as read/hits). If the read/hit ratio is high, that means data was successfully cached and accessed, thereby reducing the number of times the disk needed to be accessed to find data (disk I/O). Adding memory, defragmenting your system, or spreading out processes and services over a longer period of time may improve the File Cache read/hit ratio. See "File Cache Activity by Type" on page 239 for more information.

Processor Queue Length

This chart shows how many processes (threads) are waiting for service in the NT dispatcher ready queue. Each processor should have a queue length of no more than 2 threads. A long queue length could indicate that:

- Several processes are waking up and attempting to run all at the same time. These could be spread out to run at different times so as not to bunch up.
- The CPU is overloaded and is causing a bottleneck of data.

MPE/iX and Unix

Charts for Unix and MPE/iX systems that resemble those previously described are:

- CPU Utilization
- Disk Service Time by Drive
- Memory Manager/Rd Hit%
- Disk I/O Queue Length

Getting Started

This section demonstrates how to start Performance Gallery Gold, how to open a data file, and how to open a chart. After these steps are completed, we can study and analyze the charts in order to assess the performance data.

Starting Performance Gallery Gold

1 On your **Start** menu, click **Programs**.

2 Click on **Performance Gallery Gold**. Your program file may be nested in another folder, as shown in Figure 5.1.

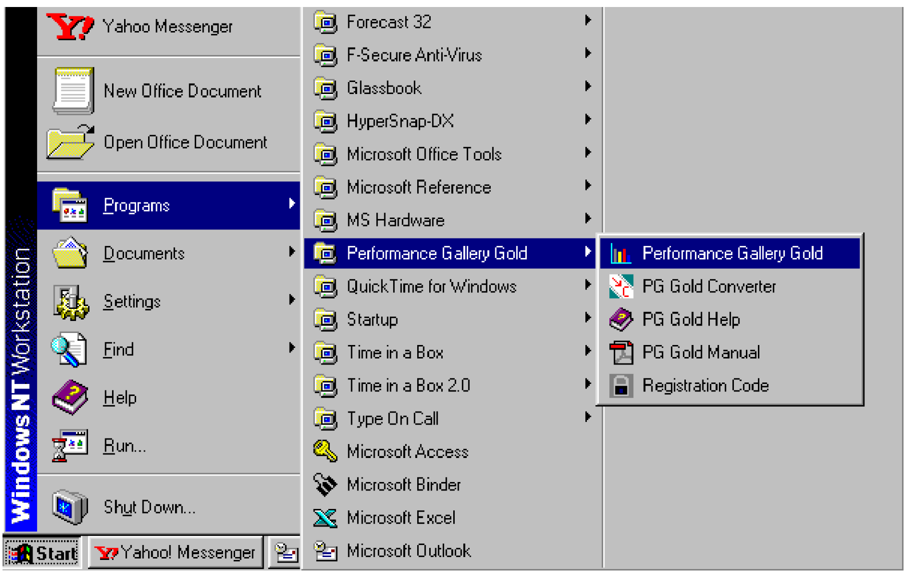

Figure 5.1 *Starting up Performance Gallery Gold*

Opening a New Data File

Performance Gallery Gold creates charts of your system performance based on the SeNTry (*.smf) or *.pfg files transferred from the host system.

Your software includes a sample data file for use with this tour.

To load the data file:

1 On the **File** menu, click **Open Data File**. The **Open** dialog box displays (see Figure 5.2).

2 In the **Files of type** list, click **SeNTry (*.smf)**.

3 Click once on the data file, **NTSystem.199902201434.smf** (the actual name may vary), to select the file. If you have your own Meta-View Performance Manager generated .pfg file, open it instead.

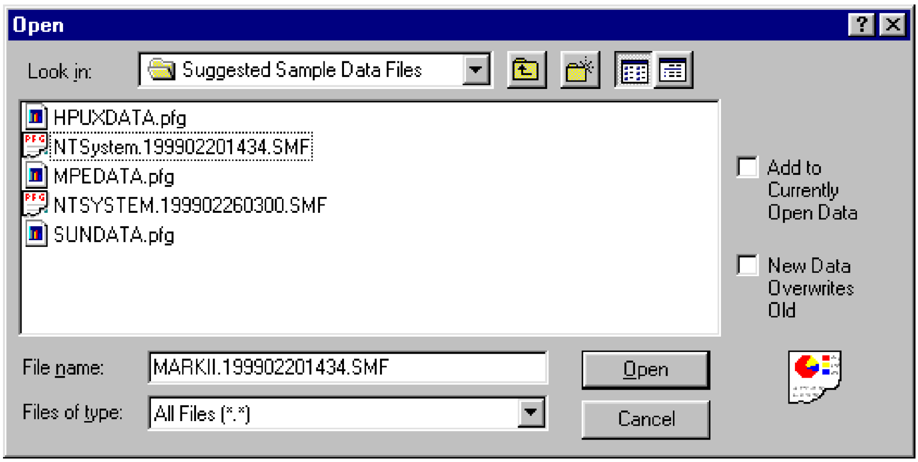

Figure 5.2 *Open Data File dialog box*

4 Click **Open** to open the data file. The name of the data file will be visible at the top of the Performance Gallery Gold screen.

Assigning an Alias

Once you have opened a data file, you may be given the opportunity to rename the data file. This capability is most useful when data files from multiple systems have been loaded (see "System Alias Manager" on page 80 for more information regarding Multiple Host Capability). For our purposes here, you may simply click **Cancel** in the **System Alias Manager** dialog box.

Opening a New Chart

Now that the data file has been opened, Performance Gallery Gold gives you options as to what information the chart displays. For this exercise, you will open several key performance charts. Make sure the Performance Gallery Gold window is maximized.

To open the charts:

1 On the **Chart** menu, click **Open Chart**.

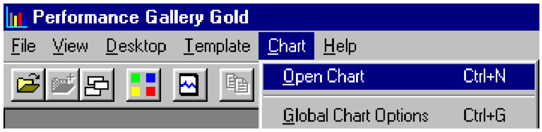

Figure 5.3 *Chart menu with Open Chart selected*

2 In the **General** tab of the **Open Chart** dialog box, select the chart titled **CPU Utilization**.

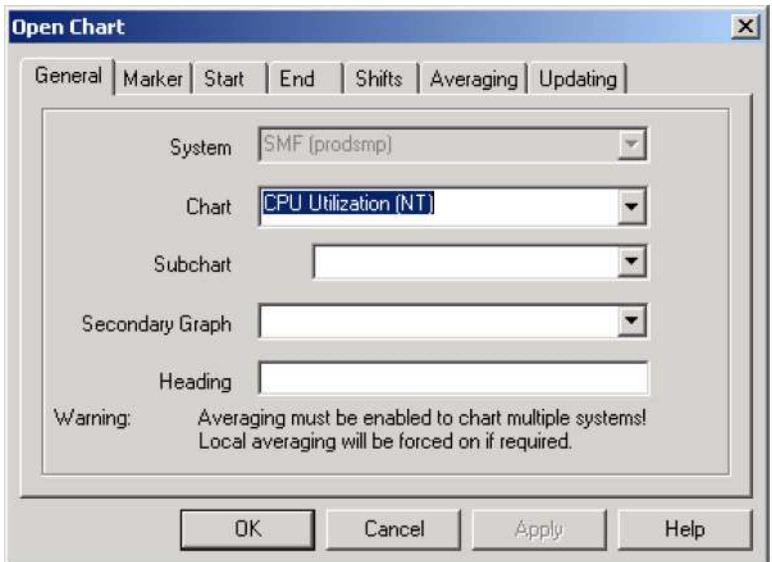

Figure 5.4 *Open Chart dialog box with CPU Utilization selected*

3 Click **OK**. The chart opens and the **Edit** and **Window** menus are added to the menu bar.

4 Repeat steps 1-3, opening the following charts:

- Logical Disk Response Time
- File Cache Activity by Type
- Processor Queue Length

5 In the **Window** menu, click **Tile Horizontal**. The four open charts will tile so that each one has it's own quadrant (see Figure 5.5 on page 40).

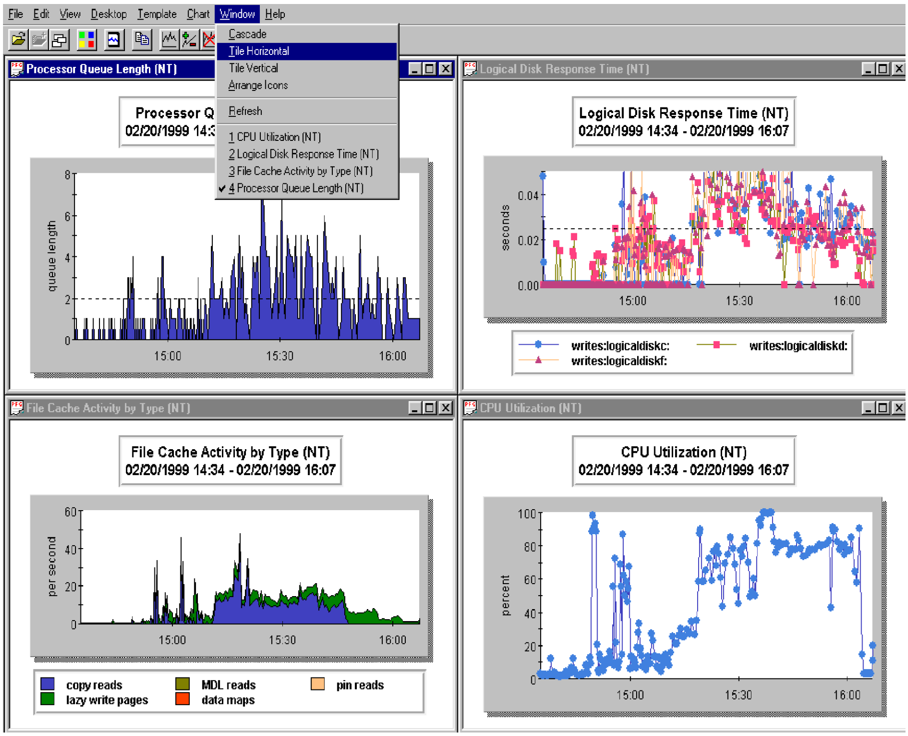

Figure 5.5 *All charts opened and tiled*

Analyzing the Charts

In this section, we will look at the attributes of the four open charts. Once we see their default settings, we will modify parts of each graph to better acquaint you with all the options Performance Gallery Gold offers and where these options are found. Changes can be made to charts to focus in on certain time periods, to narrow the charts down to certain data elements, or simply to make them more readable and aesthetically pleasing.

CPU Utilization Chart

This chart displays how much of your Central Processor Unit (CPU) is being used. If this metric reaches 85%, the system will begin to slow down noticeably.

Identify Chart Properties

Graph Attributes

Take a closer look at the chart. Maximize the **CPU Utilization** chart.

- The left side of the graph, (the y axis) shows how much (what percent) of the CPU was being used during the sample intervals.
- The x axis displays the length of the reporting period. In this case, the reporting period is from 14:34 to 16:07.

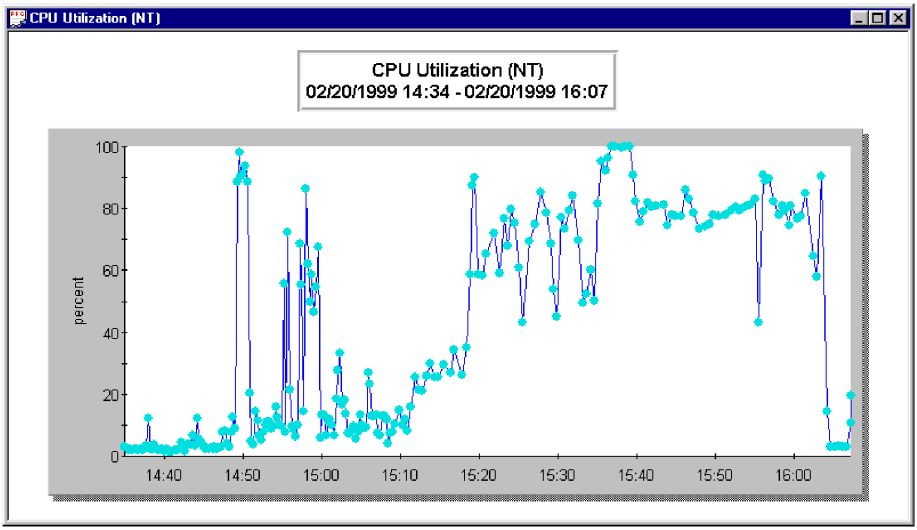

Figure 5.6 *Example CPU Utilization graph from sample data*

System Performance Conclusions

According to this chart alone, this system shows heavy usage (higher than 85%) during the reporting period. If this pattern was consistent throughout several and longer reporting periods, it could indicate the need for performance tuning or even an upgrade. Because of the limited reporting time, it would be difficult to accurately gauge the need for an upgrade. We recommend monitoring your system 24 hours a day for a week in order to gather a wider range of data upon which to make decisions.

Template Properties

1. On the **Template** menu, click **Properties**.
2. On the **General** tab, you'll see that the selected **Chart Type** is **Graph** and the **Timeline Type** is **Continuous**.
3. Click the **Graph** tab.

- **Y Axis Labeling** is defined as **percent**.
- **Graph Type** is **Line**.
- **Scale** is **Percent**.
- **Stacking** is not enabled, as indicated by the lack of a check mark in the check box.
- A **Marker** has not been defined.

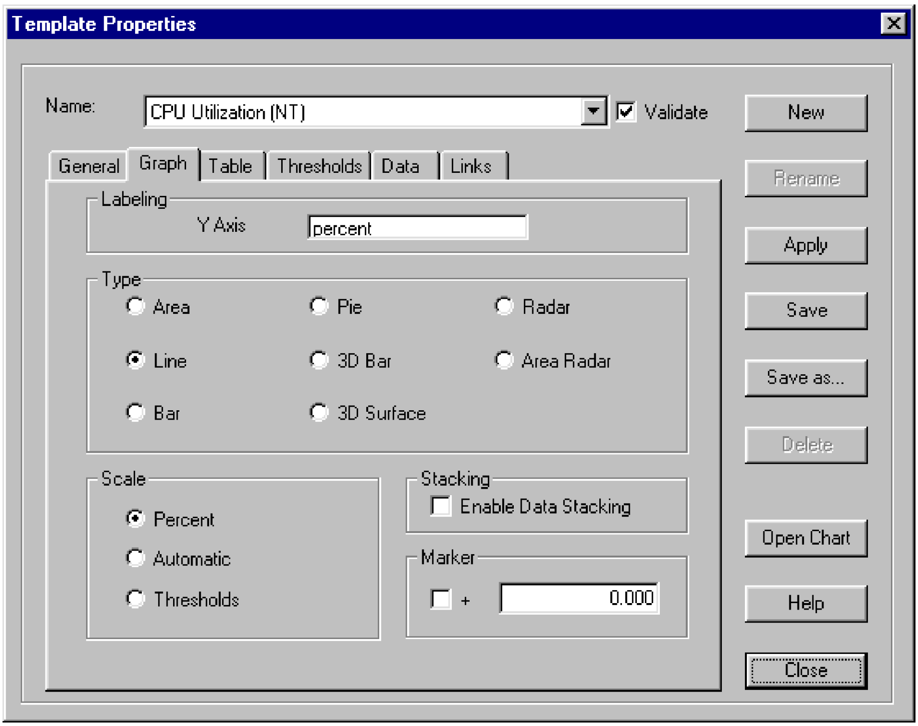

Figure 5.7 *Template Properties, Graph tab*

4 Since this is a graph, skip the **Table** tab and click on the **Thresholds** tab. Neither a maximum nor a minimum threshold has been defined, as indicated by the lack of any check marks in the check boxes. Information regarding the **Table** tab can be found in "Table Properties" on page 132.

5 Click the **Data** tab (see Figure 5.8).

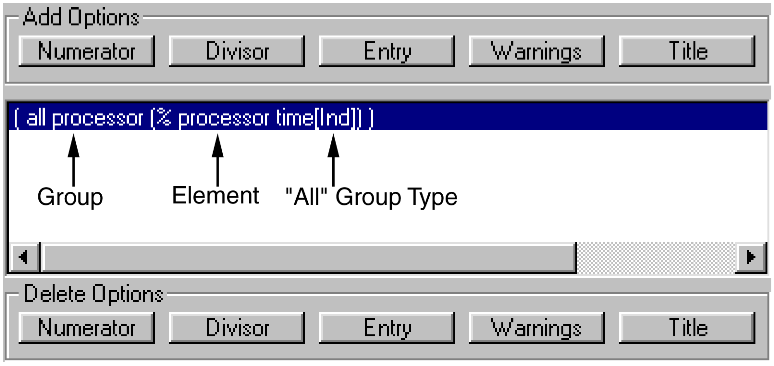

Figure 5.8 *Computation window in the Data tab*

- **Group** has been defined as **all processor**.
- **Element** is **% processor time**.
- **Individual** has been selected from the **"All" group type**.
- **Sign** is defined as positive.
- The data set does not have its **Warnings** enabled (expressed by the lack of **<cc>** in the computation window).

6 Click **Warnings** in the **Add Options** window.

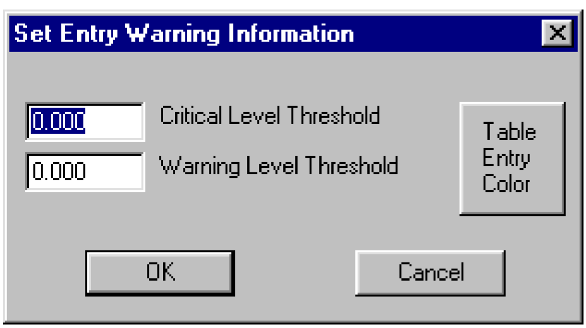

Figure 5.9 *Set Entry Warning Information dialog (example)*

- The **Critical** and **Warning Level Thresholds** are set to **0.000**.

7 Click **OK** to close the **Set Entry Warning Information** dialog box.

8 Click the **Links** tab. The CPU Utilization chart is linked to the Processor Utilization by Category (NT) and Processor Queue Length (NT) charts. This means that by using the shortcut menu (right mouse click), you can access these charts.

9 Click **Close** at the bottom right of the **Template Properties** dialog box.

Modify Chart

1 On the **Chart** menu, click **Modify Chart**. Make sure that CPU Utilization is selected in the **Chart** drop-down menu.

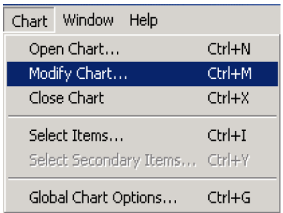

Figure 5.10 *Chart menu with Modify Chart selected*

2 Click through the tabs. None of the local settings have been enabled.

3 Click **OK** to close the **Modify Chart** window.

Global Chart Options

1 On the **Chart** menu, click **Global Chart Options**.

2 On the **Timeline Options** tab:

Figure 5.11 *Global Chart Options: Timeline Options tab*

- **All** is selected for the **X Axis Range**.

- **X Axis Points** are set to **Sample**.
- The **Blanking Limit** is **1441** minutes (24 hours).
- No days are excluded.

3 Click the **Global Date/Time Select** tab.

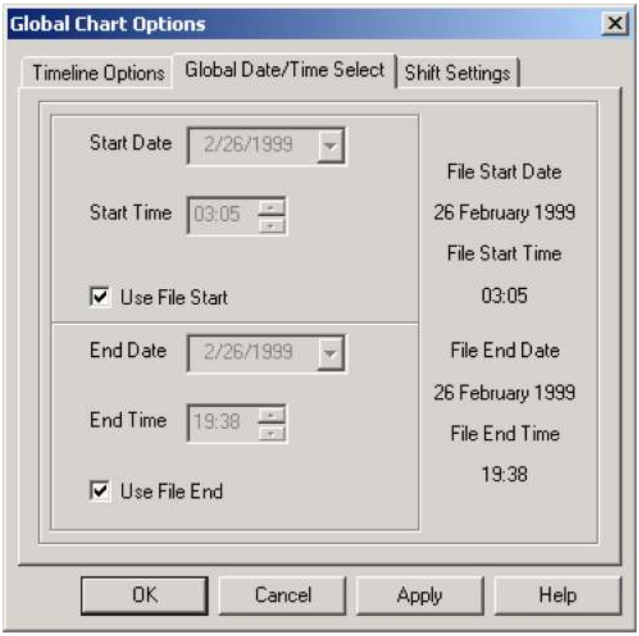

Figure 5.12 *Global Chart Options: Global Date/Time Select tab*

Note that:

- **Use File Start** is enabled.
- **Use File End** is enabled.

Because these are both enabled, we are currently seeing all the data that is available for this data file. From this point, we can only shorten the length of the reporting period displayed in our graphs, not expand it.

4 Click the **Shift Settings** tab (see Figure 5.13).

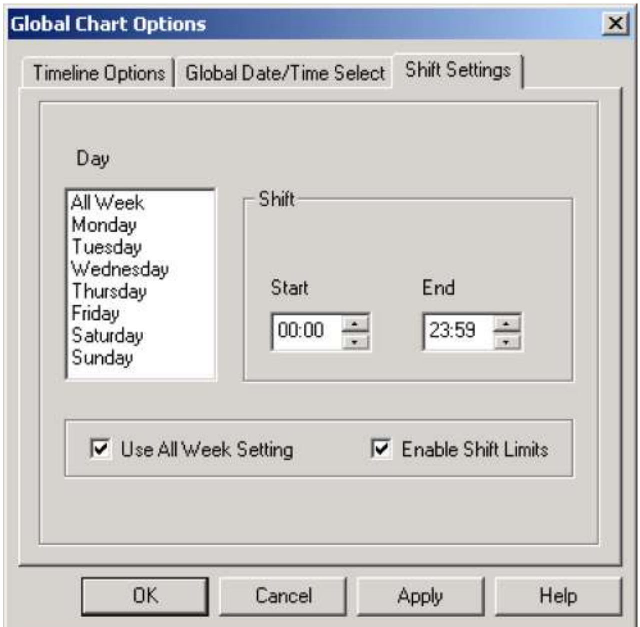

Figure 5.13 *Global Chart Options: Shift Settings tab*

Note that:

- **Use All Week Setting** is enabled.
- **Enable Shift Limits** is selected. The shift has been defined as beginning at 00:00 and ending at 23:59.

5 Close the **Global Chart Options** dialog box.

Exception Window

1 On the **View** menu, click **Exception Window**.

The exception window shows that at 22 different points during the reporting period, the CPU utilization surpassed the Warning Level threshold ; and that at 2 points within that same time period, the CPU utilization surpassed the Critical Level threshold .

2 Close the **Chart Exception** window when you have finished (see Figure 5.14).

<Chart Exceptions>

Host	Date/Time	Chart	Type	Data Set	Value
SMF	02/26/1999 15:05.00	CPU Utilization (NT).	critical	processor1	42.830000
SMF	02/26/1999 12:55.00	CPU Utilization (NT).	warning	processor1	24.670000
SMF	02/26/1999 12:50.00	CPU Utilization (NT).	warning	processor1	20.450000
SMF	02/26/1999 10:20.00	CPU Utilization (NT).	warning	processor1	24.720000
SMF	02/26/1999 15:45.00	CPU Utilization (NT).	warning	processor0	24.270000
SMF	02/26/1999 15:35.00	CPU Utilization (NT).	warning	processor0	20.210000
SMF	02/26/1999 15:25.00	CPU Utilization (NT).	warning	processor0	20.840000
SMF	02/26/1999 15:10.00	CPU Utilization (NT).	warning	processor0	20.270000
SMF	02/26/1999 15:05.00	CPU Utilization (NT).	critical	processor0	43.990000
SMF	02/26/1999 15:00.00	CPU Utilization (NT).	warning	processor0	22.020000
SMF	02/26/1999 14:55.00	CPU Utilization (NT).	warning	processor0	22.180000
SMF	02/26/1999 14:50.00	CPU Utilization (NT).	warning	processor0	23.660000
SMF	02/26/1999 14:45.00	CPU Utilization (NT).	warning	processor0	20.760000
SMF	02/26/1999 14:40.00	CPU Utilization (NT).	warning	processor0	20.790000
SMF	02/26/1999 14:30.00	CPU Utilization (NT).	warning	processor0	20.450000
SMF	02/26/1999 14:25.00	CPU Utilization (NT).	warning	processor0	22.370000
SMF	02/26/1999 14:15.00	CPU Utilization (NT).	warning	processor0	20.940000
SMF	02/26/1999 14:00.00	CPU Utilization (NT).	warning	processor0	22.720000
SMF	02/26/1999 13:30.00	CPU Utilization (NT).	warning	processor0	20.370000
SMF	02/26/1999 13:25.00	CPU Utilization (NT).	warning	processor0	21.150000
SMF	02/26/1999 13:15.00	CPU Utilization (NT).	warning	processor0	20.850000
SMF	02/26/1999 12:55.00	CPU Utilization (NT).	warning	processor0	28.090000
SMF	02/26/1999 12:50.00	CPU Utilization (NT).	warning	processor0	25.030000
SMF	02/26/1999 12:45.00	CPU Utilization (NT).	warning	processor0	22.100000

Figure 5.14 *CPU Utilization Chart Exception window (example)*

Modify the CPU Utilization Chart

To change how the graph looks and how it displays information, we are going to change the **Template Properties**, **Modify Chart** settings, and **Global Chart Options**. Remember that settings defined in the **Global Chart Options** affect all open charts but do not override any properties defined in the **Modify Chart** options. **Modify Chart** settings affect only the graph you select in the **Chart** pulldown menu on the **General** tab; and enabling any local settings in this menu overrides the template settings defined in the **Template Properties** dialog boxes.

Template Properties

1 Click the **Data** tab. For this exercise, you're going to change the **Warning** and **Critical Level Thresholds**. Highlight the data item:

   (all processor (% processor time[Ind]))

2 Click **Warnings** in the **Add Options** window.

   a Change the **Critical Level Threshold** to 90.000.

   a Change the **Warning Level Threshold** to 80.000.

   b Click **OK** to close the **Set Entry Warning Information** dialog box.

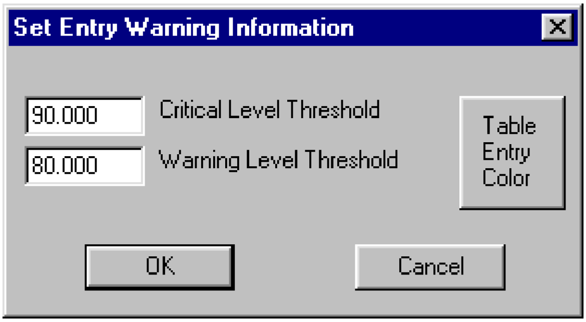

Figure 5.15 *Set Entry Warning Information dialog box*

3 Click **Apply** to save your changes for the current session to the current template. Click **Save As** to save the altered template under a different name for use in future sessions.

4 Click **Close** to exit the **Template Properties** dialog box.

Modify Chart

Since this graph is fairly dense, you can change the averaging to clear it up. Note, however, that by choosing to average the chart based on longer periods of time, it may become harder to pinpoint the exact moments or under what circumstances your system performance begins to suffer.

1 In the **Chart** menu, click **Modify Chart**.

2 Click the **Averaging** tab.

3 Click in the **Local Averaging On** check box.

4 Type **1** in the **Minutes** text box.

5 Click the **Marker** tab.

6 In the **Marker** section, click in the check box to enable a local horizontal marker.

7 Type **85** into the **Marker** text box. At this level (85%), you are using the CPU to it's full potential. If the CPU usage is higher than 85%, it can lead to bottlenecks and slow data processing.

8 Click **OK**. You should see a considerable difference in the graph (see Figure 5.16). All the data points that had been collected at 15-30 second intervals have now been averaged into one minute interval data points. Also notice that now only three data points are above the 85% threshold line.

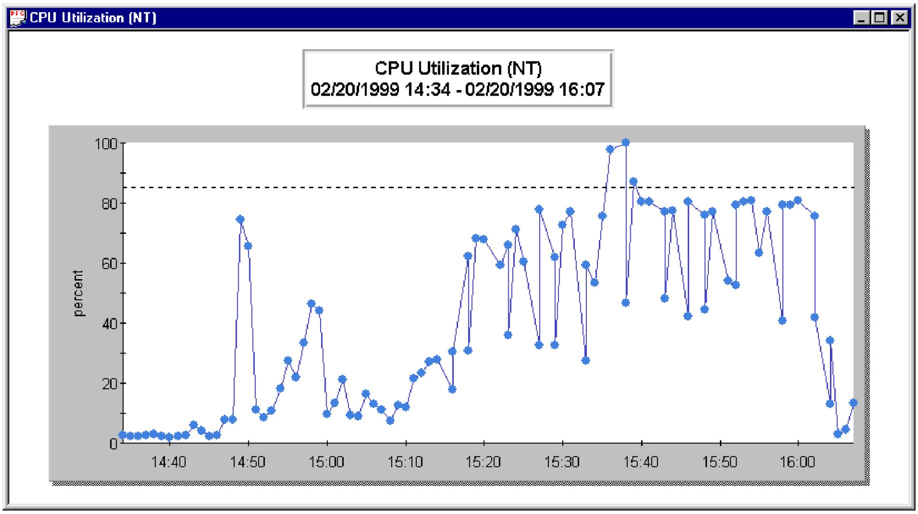

Figure 5.16 *CPU Utilization graph with Averaging set to 1 minute and a marker at 85%.*

Global Chart Options

Because there isn't much change in the data points between 14:34 and 14:45, you can elect not to enable **Use File Start**, and instead begin the chart data at a point where the data begins shifting.

1. On the **Chart** menu, click **Global Chart Options**.
2. Click the **Global Date/Time Settings** tab.
3. Click in the **Use File Start** check box to disable it. The absence of a check mark indicates that function has been disabled. The **Start Date** and **Start Time** labels also become active (no longer grayed out).
4. Using the **Start Time** standard Windows control, change the start time to 14:46 (see Figure 5.17).

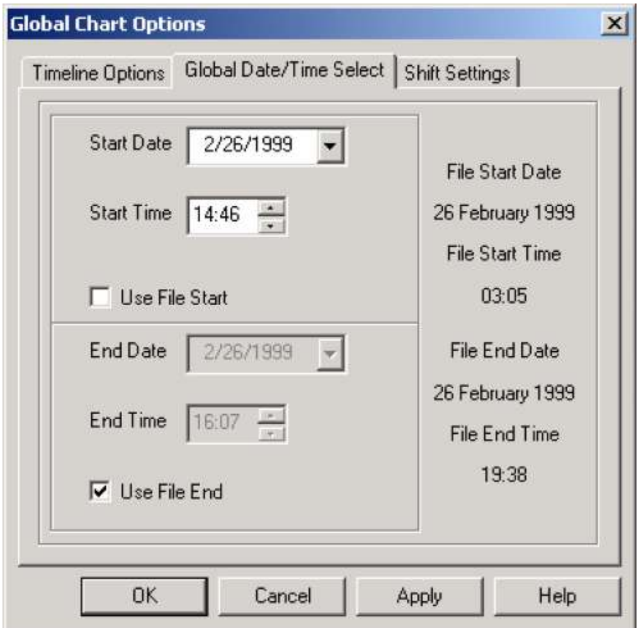

Figure 5.17 *Global Chart Options, Global Date/Time Select*

5 Click **Apply**.

6 Click **OK**.

The graph is reformatted to show data points from 14:46 to 16:07.

Exception Window

On the **View** menu, click **Exception Window**.

Because the warning and critical thresholds were added, this window now shows that at nearly 30 different points during the reporting period, the CPU utilization surpassed the Warning or Critical threshold limit. Close the **Chart Exception** window. Minimize the CPU Utilization chart when you have finished.

Logical Disk Response Time Chart

This chart displays how long it takes from the time a request for data (read or write) is made to the time the data is read from or written to the disk. Most requests can be satisfied within 11 ms (0.011 sec).

Identify Chart Properties

Because we made a global change to the CPU Utilization chart, all of the three remaining charts now show data from 14:46 to 16:07. To return the remaining charts to their original reporting periods:

1. On the **Chart** menu, click **Global Chart Options**.
2. Click the **Global Date/Time Select** tab.
3. Enable **Use File Start** by clicking in the checkbox. This causes all four of the graphs to return to the default start settings.

We also have a section of the Performance Gallery Gold window that is not being utilized. To fix this, do one of the following:

- On the **Window** menu, click **Tile Horizontal**.
- On the **Window** menu, click **Tile Vertical**.
- On the **Window** menu, click **Cascade**.

You will see a difference in how the three charts are tiled (the CPU Utilization chart remains minimized). Choose whichever option you prefer.

Graph Attributes

Take a closer look at the chart. Maximize the Logical Disk Response Time chart.

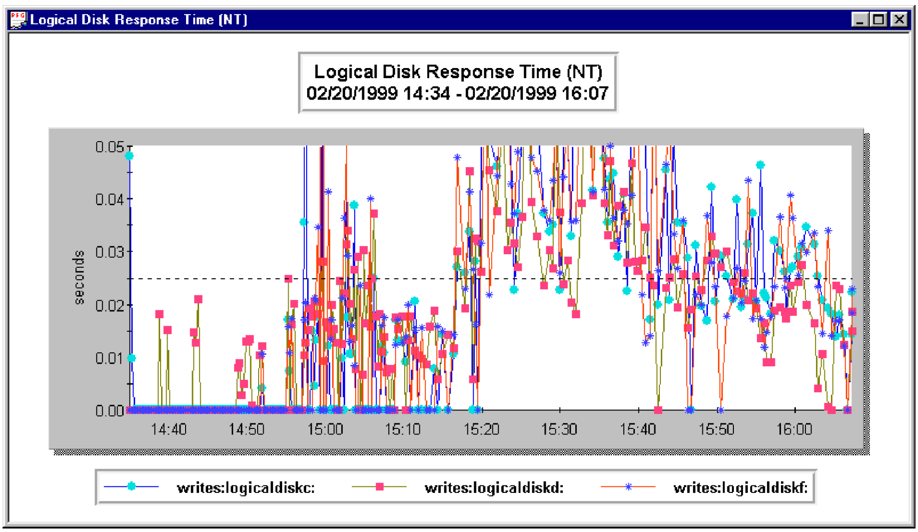

Figure 5.18 *Logical Disk Response Time graph*

- The left side of the graph, (the y axis) shows the average elapsed time, in seconds, of I/O requests during the sample interval for each physical disk (C, D, and F).

- The x axis displays the length of the reporting period. In this case, the reporting period is from 14:34 to 16:07.

System Performance Conclusions

This system shows very long response times (longer than 0.025 seconds) for all three disks at several times during the reporting period. If this pattern was consistent throughout several and longer reporting periods, it could indicate the need for performance tuning or even an upgrade. Because of the limited reporting time and the small sampling intervals, it would be difficult to accurately gauge the need for an upgrade.

Template Properties

The following table displays the attributes for this chart found in the **Template Properties** dialog box.

Table 5.1 *Template Properties attributes of the chart*

General tab	
Chart Type	Graph
Timeline Type	Continuous
Template Warnings	(none)
Graph Tab	
Labeling (Y Axis)	Seconds
Type	Line
Scale	Thresholds
Stacking	Not enabled
Marker	Enabled at 0.025 seconds
Table Tab	
N/A	N/A
Thresholds Tab	
Maximum	Enabled Drop sets always above enabled
Minimum	Not enabled
Data Tab	
Group - reads	all logicaldisk

Group - writes	all logicaldisk
Element - reads	avg. disk sec/read
Element - writes	avg. disk sec/write
"All" group type	Individual
Sign	Positive
Warnings	Not enabled
Links Tab	
Link 1	Logical Disk Utilization
Link 2	Logical Disk Detail
Link 3	Logical Disk Avg Q Len Statistics

NOTE To see how the avg. disk sec/read and avg. disk sec/write computations relate to the Logical Disk Response Time chart, refer to “Logical Disk Response Time” on page 245.

Close the **Template Properties** dialog box.

Exception Window

1 On the **View** menu, click **Exception Window**.

Because warnings were not enabled for this chart, the **Chart Exceptions** window is empty.

2 Close the **Chart Exceptions** window.

Modify the Logical Disk Response Time Chart

As with the CPU Utilization chart, we can alter this chart to see different data in different formats. As you probably have noticed, seeing individual data elements in this graph is difficult. To view individual elements, do the following:

1 Right-click in the graph area. This opens the shortcut menu.

2 Click on **Select Items**.

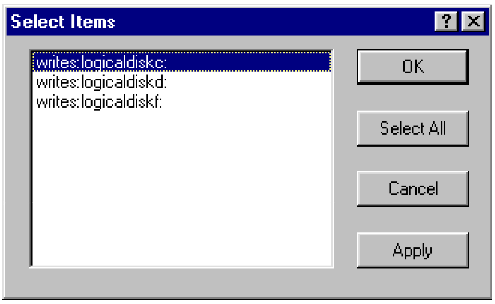

Figure 5.19 *Select Items dialog box*

a In the **Select Items** dialog box, click **writes:logicaldiskc**. Clicking on one selection automatically deselects all the other options.

b Click **Apply**.

c Click **OK**.

The graph has now been reduced to displaying only disk C data. This can be repeated for any of the three disks. To re-select all disks:

1. Right-click in the graph area.
2. Click on **Select Items**.
3. In the **Select Items** dialog box, click **Select All**.
4. Click **Apply**.
5. Click **OK**.

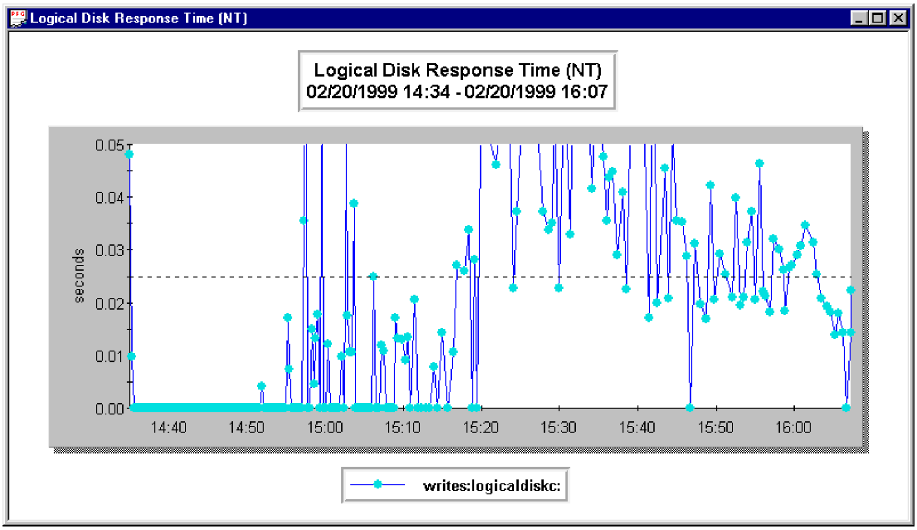

Figure 5.20 *Logical Disk Response Time graph with only one item selected*

Template Properties

Sometimes, data is better viewed in a table format, since the precise data point values are displayed. By changing this graph to a table, we will be better able to see the exact values of each data point for all disks. To change the graph to a table:

1. Click **Properties** on the **Template** menu.
2. On the **General** tab under **Chart Type**, click **Table**.
3. Click **Apply**.
4. Click the **Data** tab. We will enable warnings and color coding for this table.
5. Highlight the first equation, **reads: (all logicaldisk (avg. disk sec/read[Ind]))**, in the computation window and click **Warnings** under **Add Options**.
6. In the **Set Entry Warning Information** window, change the **Critical Level Threshold** to 0.03, and the **Warning Level Threshold** to 0.025. Click **OK**.
7. Click the second equation, **writes: (all logicaldisk (avg. disk sec/write[Ind]))**, in the computation window and repeat step 6.

While setting these thresholds will cause exceptions to appear in the exception window, you need to enable the color coding option before the values in the table will appear color coded.

1. Click the **Table** tab.
   - Enable color coding by clicking in the check box next to **Color Coded** in the **Options** section.

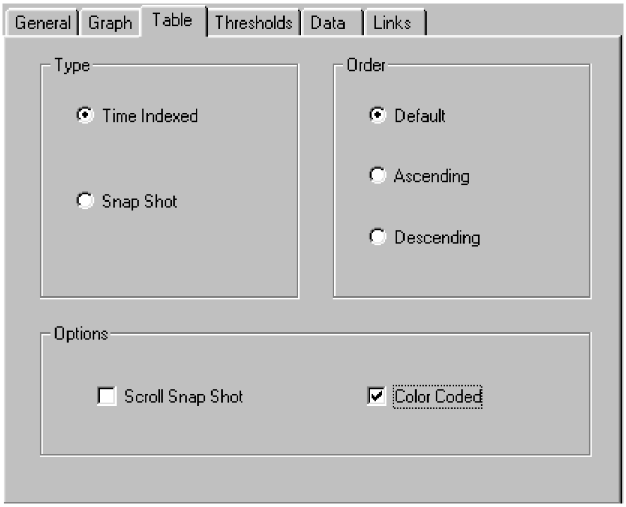

Figure 5.21 *Table tab with Color Coded enabled*

2 Click **Apply**.

3 Click **Close**.

Using the horizontal scroll bar, scan through the table data. You will be able to see the exact values for all data points. You'll also be able to see at a glance which values exceed the warning threshold (in yellow) and which exceed the critical threshold (in red).

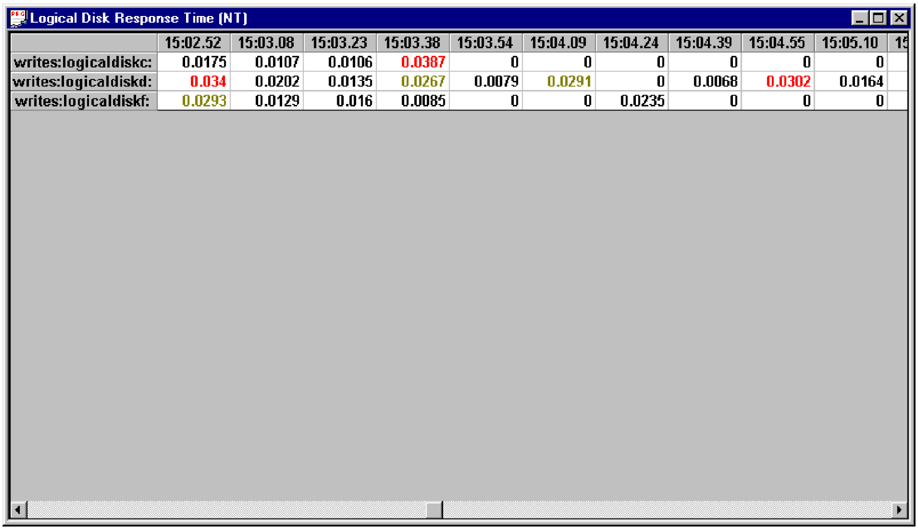
Logical Disk Response Time (NT)

	15:02.52	15:03.08	15:03.23	15:03.38	15:03.54	15:04.09	15:04.24	15:04.39	15:04.55	15:05.10	15
writes:logicaldiskc:	0.0175	0.0107	0.0106	0.0387	0	0	0	0	0	0	
writes:logicaldiskd:	0.034	0.0202	0.0135	0.0267	0.0079	0.0291	0	0.0068	0.0302	0.0164	
writes:logicaldiskf:	0.0293	0.0129	0.016	0.0085	0	0	0.0235	0	0	0	

Figure 5.22 *Logical Disk Response Time table with Color Coding*

Because we now have critical and warning thresholds defined, open the exception window (see “Exception Window” on page 88) to see that this data has generated an exorbitant number of exceptions, showing that the response time is an issue with this system. Close the exception window after you have finished.

Minimize the Logical Disk Response Time chart.

File Cache by Activity Type Chart

This chart displays how many times per second a request for data was satisfied from the memory or cache, also known as read/hits. If the read/hit ratio is high, that means data was successfully cached and accessed, thereby reducing the number of times the logical disk needed to be accessed to find data (disk I/O).

Identify Chart Properties

Since we minimized the other two charts, we now have a section of the Performance Gallery Gold window that is not being utilized. To fix this, do one of the following:

- On the **Window** menu, click **Tile Horizontal**.
- On the **Window** menu, click **Tile Vertical**.
- On the **Window** menu, click **Cascade**.

You will see a difference in how the two remaining graphs are tiled. Choose whichever option you prefer.

Graph Attributes

Take a closer look at the chart. Maximize the File Cache Activity by Type chart.

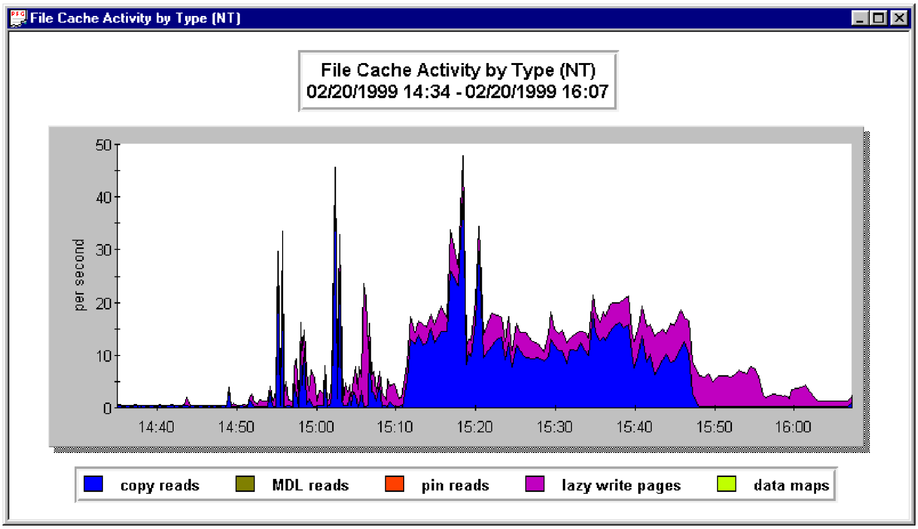

Figure 5.23 *File Cache Activity by Type graph*

NOTE For more detailed information on the File Cache by Activity Type chart, see "File Cache Activity by Type" on page 239.

- The left side of the graph, (the y axis) shows how many times per second that individual data element (copy reads, MDL reads, etc.) requests were satisfied from the cache without having to access the disk. This reduces disk I/O operations.
- The x axis displays the length of the reporting period. In this case, the reporting period is from 14:34 to 16:07.

System Performance Conclusions

Any conclusions based on this graph alone would be incomplete or incorrect without having detailed knowledge of the system in question. You can see from this graph that at approximately 15:18, the total read/hits per second was around 46. Because this graph doesn't show how many attempts were made, we really don't know if 46 is good or bad. If 100 attempts were made, then your read/hit percentage is only 46%; whereas, if only 46 attempts were made, then all were successful. Other graphs (such as CPU Utilization and File Cache Copy Read Hit %) must be looked at in conjunction with this one to be able to make important upgrade decisions. In the following sections, we will modify this chart to make it more useful.

Template Properties

The following table displays the attributes for this chart found in the **Template Properties** dialog box.

Table 5.2 *Template Properties attributes for the chart*

General Tab	
Chart Type	Graph
Timeline Type	Continuous
Template Warnings	No template links defined. This is not important to this graph; it only means that if you want to open a related graph, you will have to go to the Open Chart dialog box. You can choose to add template links to the Links tab.
Graph Tab	
Y Axis Labeling	Seconds
Graph Type	Area
Scale	Automatic
Stacking	Enabled
Marker	Not enabled
Table Tab	
N/A	(Tables are referenced in "Table Properties" on page 132.)
Thresholds Tab	
Maximum Threshold	Not enabled
Minimum Threshold	Not enabled
Data Tab	
Group - all data elements	nt cache
Element - reads	<data element>/sec
"All" group type	N/A

Sign	Positive
Warnings	Not enabled
Links Tab	
Chart is linked to:	None

Close the **Template Properties** dialog box.

Exception Window

1. On the **View** menu, click **Exception Window**.

   Because color coding was not enabled for this chart, the **Chart Exceptions** window is empty.

2. Close the **Chart Exceptions** window.

Modify the File Cache Activity by Type Chart

There are several things that can be done to this graph to make it more useful.

Template Properties

1. On the **Template** menu, click **Properties**.
2. Click the **Thresholds** tab.
   a. In the **Minimum** section, click in the **Enable Threshold** check box.
   b. Click in the **Drop sets always below** check box.
   c. Change the **Value** to **1.000**.

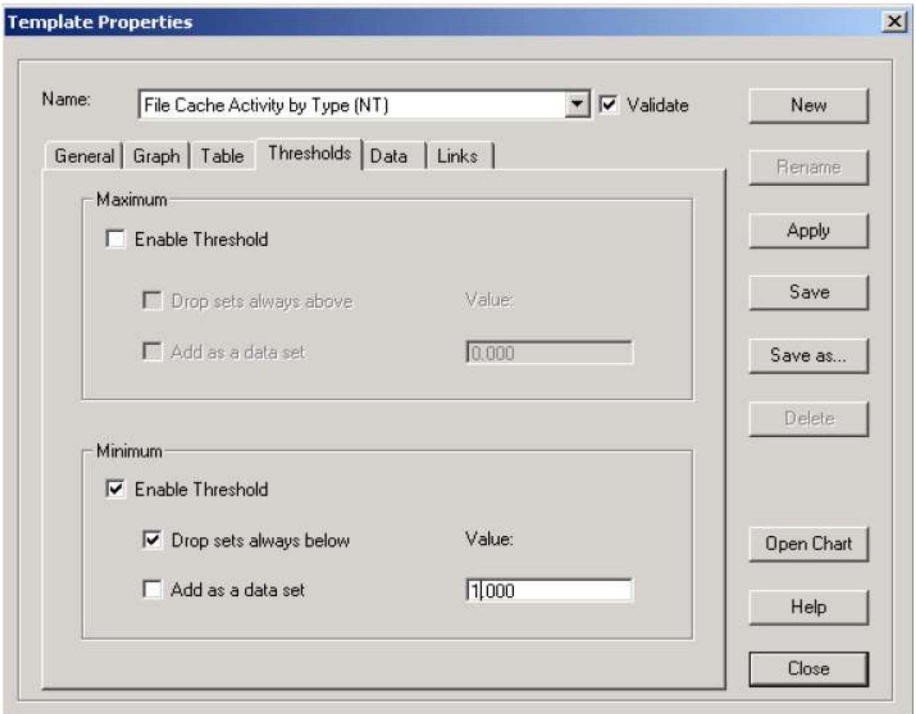

Figure 5.24 *Thresholds tab with Minimum Threshold enabled*

3 Click **Apply**. **MDL reads**, **pin reads**, and **data maps** have now been dropped from the graph.

4 Click the **Data** Tab.

In the computation window, highlight **lazy pages** and click **Entry** in the **Delete Options** window.

5 Click **Apply**.

6 Click **Close**.

The graph now only contains data for copy reads, which we will compare to other graph data to glean meaning from the data seen here.

Modify Chart

We are going to add a second chart that relates directly to the File Cache Activity by Type chart. This second chart will show us the percentage of successful copy read hits, rather than just a number without a reference point.

1 In the **Chart** menu, click **Modify Chart**. Ensure that **File Cache Activity by Type** is selected in the **Chart** drop-down menu.
2 In the **Secondary Graph** scroll box on the **General** tab, select **File Cache Copy Read Hit %**.
3 Click **OK**.

A second data set and a second Y axis is added to the right side of the graph. The new data set shows the percentage of copy read requests that were satisfied from the file cache without accessing the disk. This new data set is hidden at points by the copy reads data, so changing the type of graph for the copy reads graph will be necessary.

In the **Template** menu, click **Properties**.

1 Make sure that **File Cache Activity by Type** is selected in the **Name** window.
2 Click the **Graph** tab.
3 Change the **Graph Type** from **Area** to **Line**.
4 Click **Apply**. Do not close the **Template Properties** window.

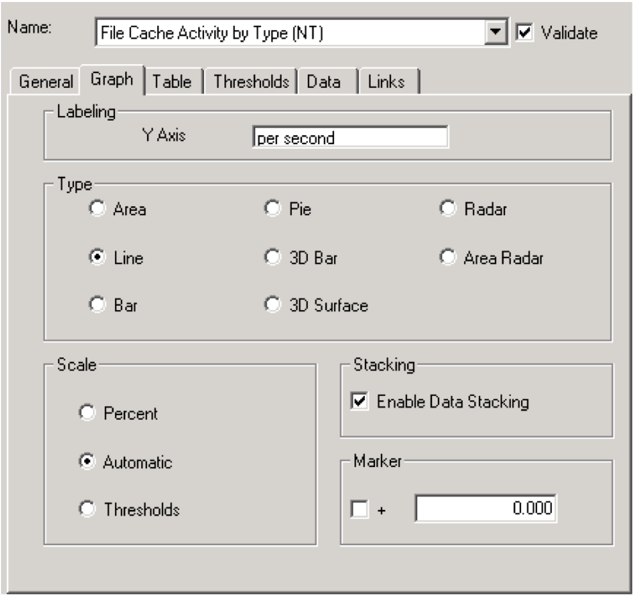

Figure 5.25 *Graph tab with Type changed to Line graph*

We cannot add a marker to the secondary graph, but it can be given warning thresholds to display in the exception window.

1 Click the **General** tab. Change the **Name** to **File Cache Copy Read Hit %**.
2 Click the **Data** tab.

3 Click the **Warnings** button in the **Add Options** window.

   a Change the **Critical Level Threshold** to **90.000**.

   b Change the **Warning Level Threshold** to **80.000**.

   c Click **OK**.

4 Click **Apply**.

5 Click **Close**.

Now we can get a better idea of how the system is performing. For most systems, read hits should be successful at least 80% of the time. By combining these two graphs, we can see that at around 15:20, there were about 42 successful copy read hits per second and that those 42 hits translate to 85% of the total attempts. So during this sample interval, around 50 requests were made of the file cache, and 42 of those requests were satisfied without having to perform a disk I/O operation.

View Options

If it is difficult to see the copy reads line against the backdrop of the copy reads hit %:

1 In the **View** menu, click **Options**.

2 In the **Appearance** tab, increase the line thickness in the **Line Thickness** scroll box.

3 Disable **Mark Data Points** by clicking in the check box (see Figure 5.26).

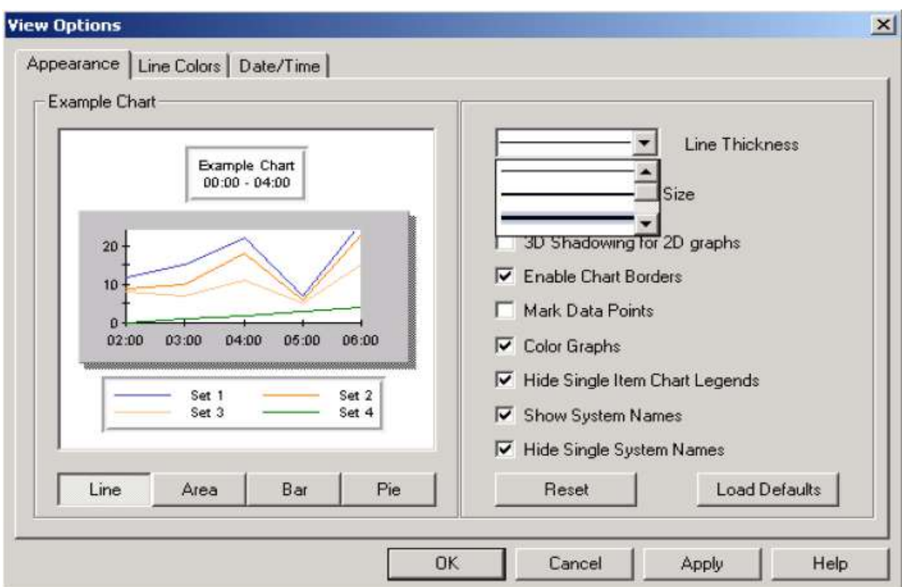

Figure 5.26 *View Options: Appearance tab with Data Points disabled and Line Thickness menu extended*

If the color of the primary chart too closely resembles that of the secondary chart, you can change the color of either one in two ways:

Option One:

1 Click the **Line Colors** tab.

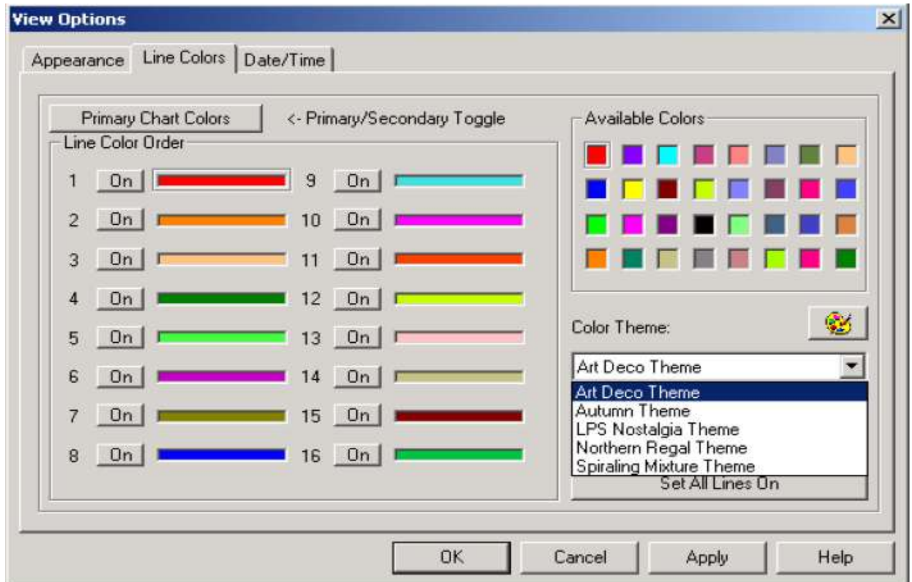

Figure 5.27 *Line Colors tab with the first color selected and the Color Theme menu extended*

2 Click on the first color box in the **Line Color Order** window.

3 Find a color in the **Available Colors** window that contrasts both with the color of the other data element and with the background of the graph window. For instance, if the copy read hits % area graph is black, you wouldn't want to choose a really light yellow or gray, because it would be too difficult to see against the white background of the graph window. A bright green or neon yellow or pink is usually a good contrast for charts printed in color. For charts printed in black and white, a medium gray (50%) for the area chart and a thick black line for the line chart works well.

4 Click **Apply**.

5 Click **OK**.

Option Two:

1 Click on the **Line Colors** tab.

2 Click on the **On/Off** toggle switch for each **Line Color Order** option (turning them to **Off**) until the first color that is turned **On** is a good contrast to both the other data element and the background of the graph window.

3 Click **Apply**.

4 Click **OK**.

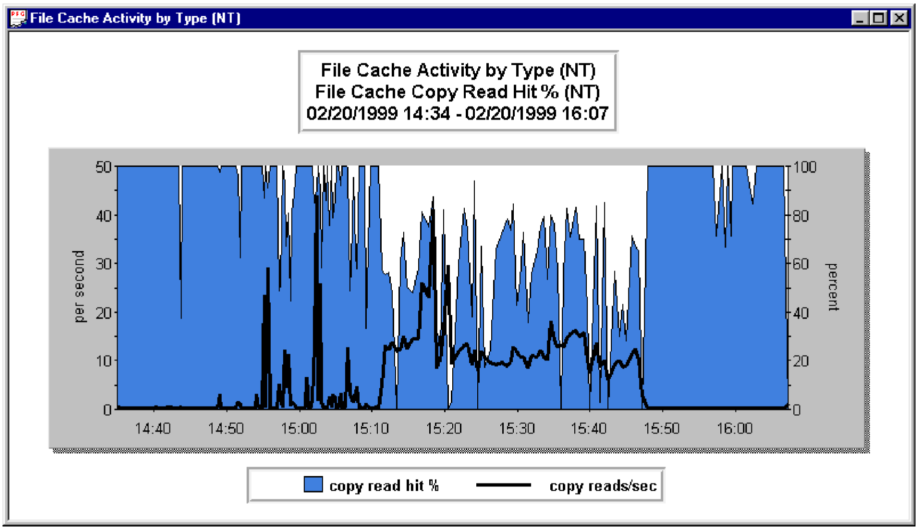

Figure 5.28 *Dual Y axis graph*

After you have examined this graph, minimize it.

Processor Queue Length Chart

This chart shows the number of processes (threads) that are waiting for service in the NT dispatcher ready queue. Each processor should have a queue length of no more than 2 threads.

Identify Chart Properties

Graph Attributes

Take a closer look at the chart. Maximize the Processor Queue Length chart.

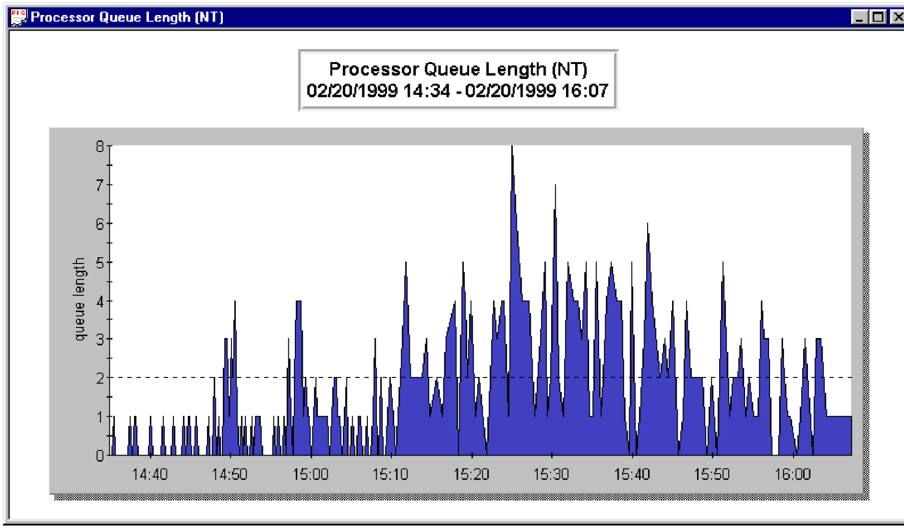

Figure 5.29 *Processor Queue Length area graph*

NOTE For more detailed information on the Processor Queue Length, see "Processor Queue Length" on page 227.

- The left side of the graph, (the y axis) displays the number of threads that are currently waiting for service in the NT dispatcher ready queue. At one point during the sample interval, the queue length reached as high as 8.
- The x axis displays the length of the reporting period, which was from 14:34 to 16:07.

System Performance Conclusions

At several times during this reporting period, and especially between 15:10 and 15:50, the queue length far exceeded the ideal threshold of 2. Combine this information with the information we gleaned from the previous graphs: during the same time:

- The CPU Utilization chart showed increased CPU usage;
- The Logical Disk Response Time graph showed increased disk response times;
- The number of copy read hits increased, while the percentage of those being successful dropped;

It's clear that the high usage during this time has caused bottlenecking.

Template Properties

The following table displays the attributes for this chart found in the **Template Properties** dialog box.

Table 5.3 *Template Properties attributes for the chart*

General Tab	
Chart Type	Graph
Timeline Type	Continuous
Template Warnings	No template links defined. You can choose to add template links to the Links tab.
Graph Tab	
Y Axis Labeling	Queue length
Graph Type	Area
Scale	Automatic
Stacking	Not enabled
Marker	Enabled at 2.000
Table Tab	
N/A	(Tables are referenced in “Table Properties” on page 132.)
Thresholds Tab	
Maximum Threshold	Not enabled
Minimum Threshold	Not enabled
Data Tab	
Group	nt system
Element - reads	processor queue length
"All" group type	N/A
Sign	Positive
Warnings	Not enabled
Links Tab	

Chart is linked to:	None

Close the **Template Properties** dialog box.

Exception Window

1 On the **View** menu, click **Exception Window**.

Because warnings were not enabled for this chart, the **Chart Exceptions** window is empty.

2 Close the **Chart Exceptions** window.

Modify the Processor Queue Length Chart

Template Properties

1 In the **Template** menu, click **Properties**.

2 Click the **Data** tab.

3 Select the computation in the window and click **Warnings** in the **Add Options** window.

   a Change the **Critical Level Threshold** value to **3.000**.

   b Change the **Warning Level Threshold** value to **2.000**.

   c Click **OK**.

4 Click the **Links** tab. In the **Link 1** window, select **Logical Disk Activity (totals)**.

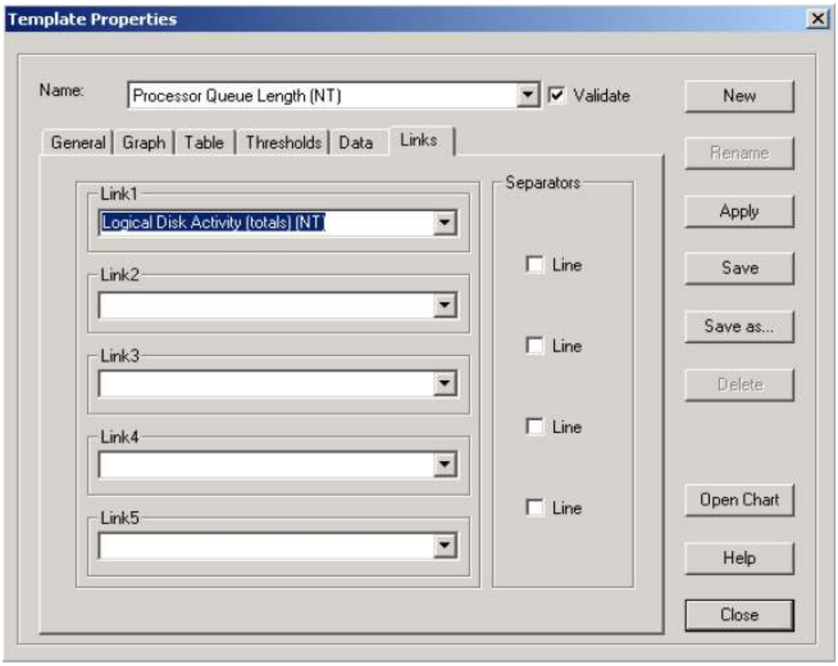

Figure 5.30 *Processor Queue Length Links tab*

5 Click **Apply**.

6 Click **Close**.

Exception Window

In the **View** menu, click **Exception Window**. There will be exceptions from three of our four charts. Initially, these are organized based on the time of the exception.

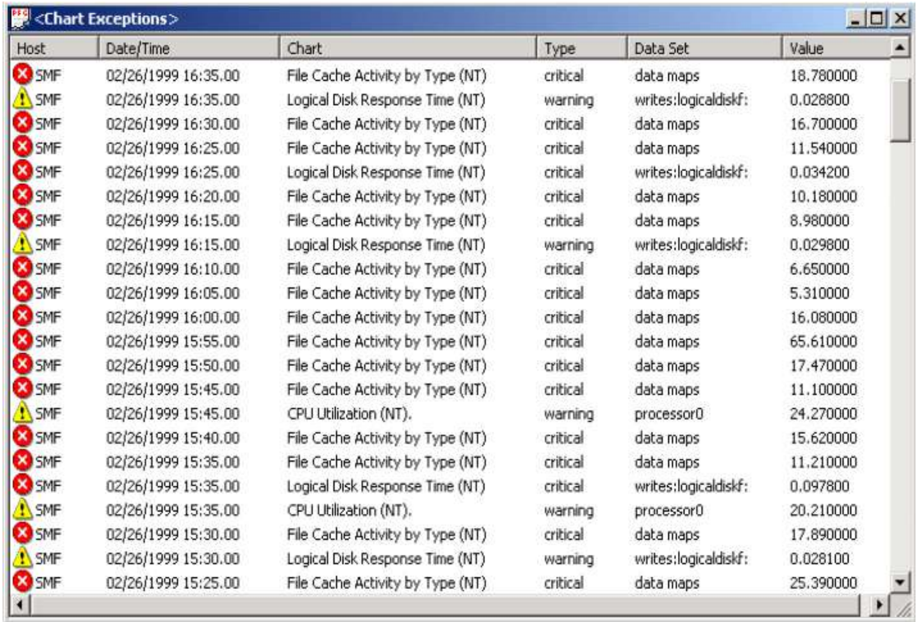

Host	Date/Time	Chart	Type	Data Set	Value
SMF	02/26/1999 16:35.00	File Cache Activity by Type (NT)	critical	data maps	18.780000
SMF	02/26/1999 16:35.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.028800
SMF	02/26/1999 16:30.00	File Cache Activity by Type (NT)	critical	data maps	16.700000
SMF	02/26/1999 16:25.00	File Cache Activity by Type (NT)	critical	data maps	11.540000
SMF	02/26/1999 16:25.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.034200
SMF	02/26/1999 16:20.00	File Cache Activity by Type (NT)	critical	data maps	10.180000
SMF	02/26/1999 16:15.00	File Cache Activity by Type (NT)	critical	data maps	8.980000
SMF	02/26/1999 16:15.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.029800
SMF	02/26/1999 16:10.00	File Cache Activity by Type (NT)	critical	data maps	6.650000
SMF	02/26/1999 16:05.00	File Cache Activity by Type (NT)	critical	data maps	5.310000
SMF	02/26/1999 16:00.00	File Cache Activity by Type (NT)	critical	data maps	16.080000
SMF	02/26/1999 15:55.00	File Cache Activity by Type (NT)	critical	data maps	65.610000
SMF	02/26/1999 15:50.00	File Cache Activity by Type (NT)	critical	data maps	17.470000
SMF	02/26/1999 15:45.00	File Cache Activity by Type (NT)	critical	data maps	11.100000
SMF	02/26/1999 15:45.00	CPU Utilization (NT).	warning	processor0	24.270000
SMF	02/26/1999 15:40.00	File Cache Activity by Type (NT)	critical	data maps	15.620000
SMF	02/26/1999 15:35.00	File Cache Activity by Type (NT)	critical	data maps	11.210000
SMF	02/26/1999 15:35.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.097800
SMF	02/26/1999 15:35.00	CPU Utilization (NT).	warning	processor0	20.210000
SMF	02/26/1999 15:30.00	File Cache Activity by Type (NT)	critical	data maps	17.890000
SMF	02/26/1999 15:30.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.028100
SMF	02/26/1999 15:25.00	File Cache Activity by Type (NT)	critical	data maps	25.390000

Figure 5.31 *Exception window showing exceptions from three different charts, organized by the time of the exception (default)*

We want to look at processor queue length exceptions, but trying to weed them out from all the others can be difficult. To change the way the exception window is organized, click the gray **Chart** button at the head of the list. All the entries are now arranged according to chart type. Close the exception window when you have finished.

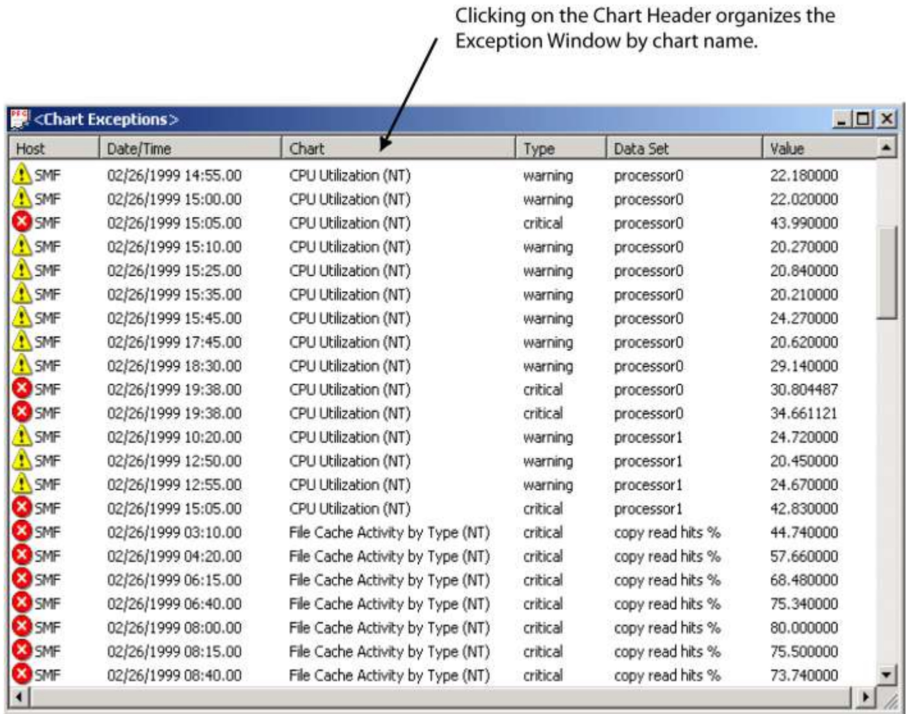

Host	Date/Time	Chart	Type	Data Set	Value
SMF	02/26/1999 14:55.00	CPU Utilization (NT)	warning	processor0	22.180000
SMF	02/26/1999 15:00.00	CPU Utilization (NT)	warning	processor0	22.020000
SMF	02/26/1999 15:05.00	CPU Utilization (NT)	critical	processor0	43.990000
SMF	02/26/1999 15:10.00	CPU Utilization (NT)	warning	processor0	20.270000
SMF	02/26/1999 15:25.00	CPU Utilization (NT)	warning	processor0	20.840000
SMF	02/26/1999 15:35.00	CPU Utilization (NT)	warning	processor0	20.210000
SMF	02/26/1999 15:45.00	CPU Utilization (NT)	warning	processor0	24.270000
SMF	02/26/1999 17:45.00	CPU Utilization (NT)	warning	processor0	20.620000
SMF	02/26/1999 18:30.00	CPU Utilization (NT)	warning	processor0	29.140000
SMF	02/26/1999 19:38.00	CPU Utilization (NT)	critical	processor0	30.804487
SMF	02/26/1999 19:38.00	CPU Utilization (NT)	critical	processor0	34.661121
SMF	02/26/1999 10:20.00	CPU Utilization (NT)	warning	processor1	24.720000
SMF	02/26/1999 12:50.00	CPU Utilization (NT)	warning	processor1	20.450000
SMF	02/26/1999 12:55.00	CPU Utilization (NT)	warning	processor1	24.670000
SMF	02/26/1999 15:05.00	CPU Utilization (NT)	critical	processor1	42.830000
SMF	02/26/1999 03:10.00	File Cache Activity by Type (NT)	critical	copy read hits %	44.740000
SMF	02/26/1999 04:20.00	File Cache Activity by Type (NT)	critical	copy read hits %	57.660000
SMF	02/26/1999 06:15.00	File Cache Activity by Type (NT)	critical	copy read hits %	68.480000
SMF	02/26/1999 06:40.00	File Cache Activity by Type (NT)	critical	copy read hits %	75.340000
SMF	02/26/1999 08:00.00	File Cache Activity by Type (NT)	critical	copy read hits %	80.000000
SMF	02/26/1999 08:15.00	File Cache Activity by Type (NT)	critical	copy read hits %	75.500000
SMF	02/26/1999 08:40.00	File Cache Activity by Type (NT)	critical	copy read hits %	73.740000

Figure 5.32 *Exception Window organized by chart name*

Shortcut Menu

In the **Template Properties** window, we linked this chart to the Logical Disk Activity (totals) graph. Because of this link, we can quickly open the related chart.

1 Right-click in the area of the Processor Queue Length chart.

2 At the bottom of the menu, click **Logical Disk Activity (totals)**.

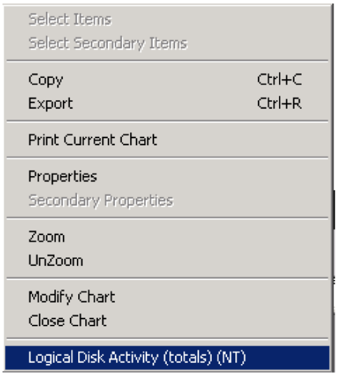

Figure 5.33 *Shortcut menu with the linked chart selected*

A new graph opens for perusal.

Saving the Desktop

If you have found these graphs useful, it is likely that you will want to use them again later without having to repeat everything you've done in this chapter. By saving this desktop, you can use these same charts for other data files.

1 In the **Desktop** menu, click **Save Desktop**.

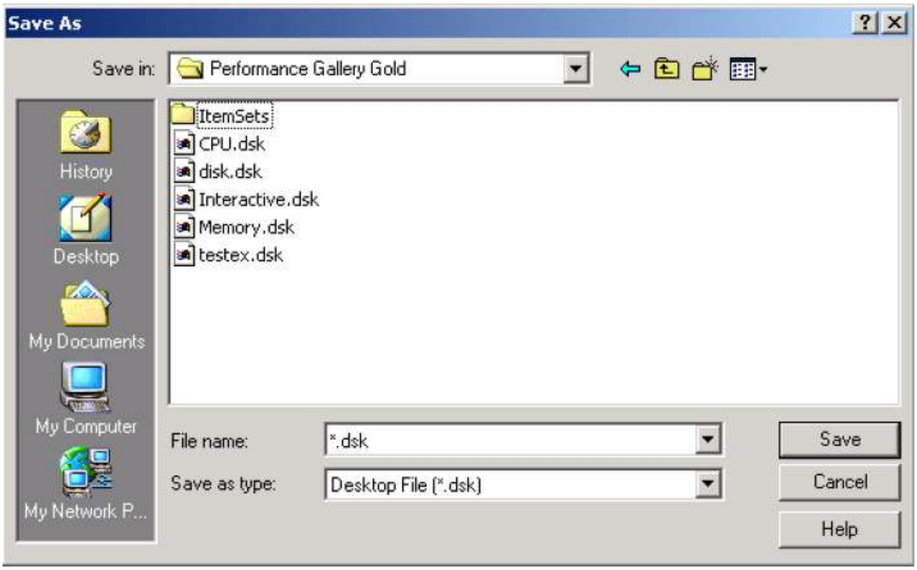

Figure 5.34 *Save Desktop dialog box*

2 Enter a name in the **File name** window.

3 Click **Save**.

Next time you would like to use these charts for a new data file:

1 In the **File** menu, select **Open...**. Locate the file you would like to open and click **Open**.

2 In the **Desktop** menu, click **Load Desktop...**. Select the desktop you would like to use and click **Open**. The last five desktops you used should also be listed in the **Recently Used Desktops** list in the **Desktop** menu (see “Recently Used Desktops” on page 103).

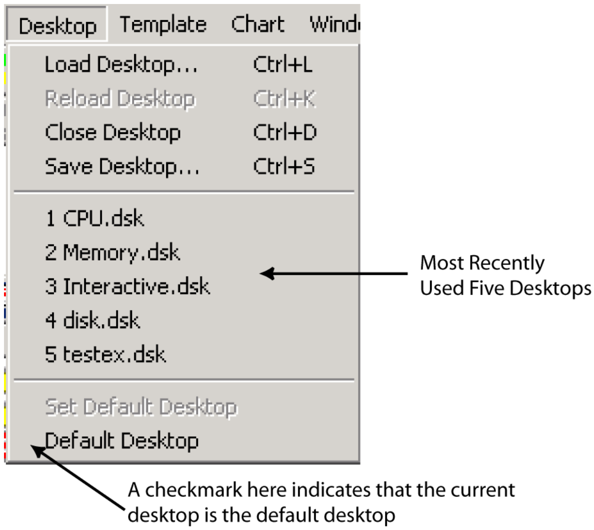

Figure 5.35 *Desktop menu*

If you would like to use this same desktop every time you use Performance Gallery Gold (without loading a desktop from the **Open** dialog box), you can set this desktop to be your default desktop.

1. In the **Desktop** menu, click **Set Default Desktop**.
2. A dialog box will appear asking if you want to set the current desktop as the default. Click **OK**.

The next time you open Performance Gallery Gold and open a data file, the charts in the default desktop will automatically open.

Exiting Performance Gallery Gold

When you are done with your Performance Gallery Gold session, click **Exit** in the **File** menu. The program will close. It may prompt you to save any modified templates. Click **No**.

6

Menus and Commands

Overview

There are eight major menus within Performance Gallery Gold. There are six main menus available at program startup. When a chart is opened, the **Edit** menu and **Window** menu are displayed, for a total of eight menus. The drop-down menus are located in the menu bar in the program window.

Figure 6.1 *Performance Gallery Gold menu bar*

File Menu

The **File** menu (see Figure 6.2) contains options to open and close a data file, or close all the data files opened, to set the system alias manager, to view the most recently-used data files, to view the current data file statistics, to print the graph(s), and to exit the Performance Gallery Gold session.

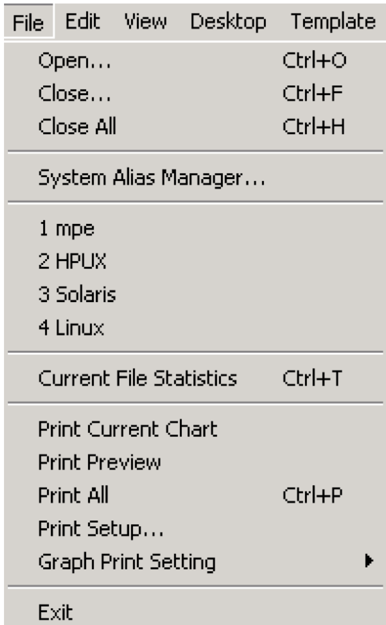

Figure 6.2 *File menu*

Open Data File

Use the **Open...** menu option to open (load) a Performance Gallery Gold (*.pfg) or Performance SeNTry (*.smf) data file. System performance data files are downloaded from the host data collector to the PC.

To open a data file, do the following:

1. On the **File** menu, click **Open...**.

   As an alternative, click the File Open toolbar button , or press the Ctrl+O shortcut keys.

2. Select the appropriate file group from the **Files of type:** drop-down menu:
   - Perfgal Files (*.pfg)
   - Sentry Files (*.smf)
   - All Files (*.*)
3. Select a data file from the list box, or type the name of the file in the **File name** text box.
4. Enable the following options as needed:

   a. Select the **Add to Currently Open Data** check box. (See "Add to Currently Open Data" on page 77) b Select the **New Data Overwrites Old** check box. (See “New Data Overwrites Old” on page 78)

5 Click **Open**.

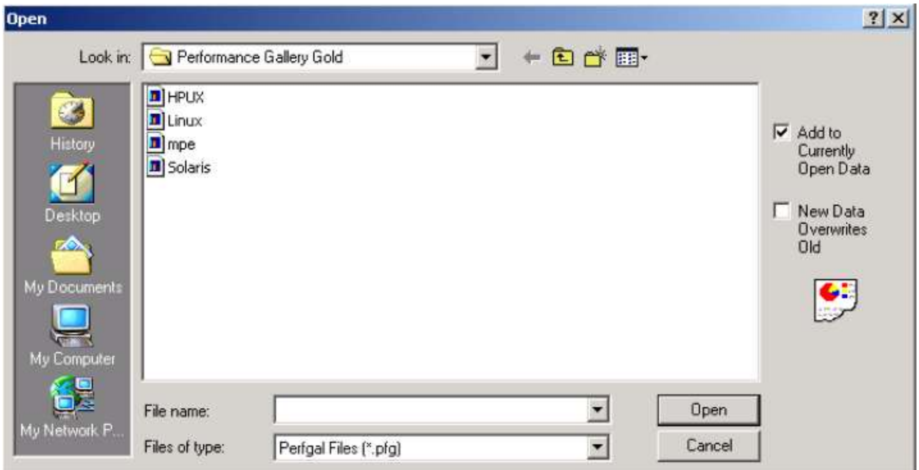

Figure 6.3 *Open dialog box*

NOTE In the "Open" dialog box, when "Add to Currently Open Data" option is not checked, the "New Data Overwrites Old" checkbox is grayed-out.

Add to Currently Open Data

This option behaves in two different ways, depending on the files that are being opened:

Multiple files from the same host

Multiple data files from the same host system can be loaded and utilized as a data table by the Performance Gallery Gold program. The program combines the data files into one composite data table which is viewed in a top-down format.

The **Add to Currently Open Data** option determines how a data file is loaded.

- When the **Add to Currently Open Data** check box in the **Open** (data file) dialog box is selected, the data file is added to the data table of currently open data file(s), provided those files were collected on the same host system as the new file.

  A data file that has been loaded as part of a composite data table cannot be removed from the data table. The data file can, however, be loaded again independently.

- When the **Add to Currently Open Data** check box is cleared (not selected), the composite data table (if one exists) is purged from memory, and the new data file is loaded and utilized independently of any other data file or data table.

To load additional data files on top of the currently open data file(s) from the same system, select the **Add to Currently Open Data** check box in the **Open** dialog box.

To view the top-down order of the composite data table, view the **Current File Statistics** dialog box (see "Current File Statistics" on page 82).

Multiple files from different hosts

In order for multiple data files from different host systems to be loaded and visible, the **Add to Currently Open Data** option must be selected, otherwise attempting to open a second data file will result in the second data file simply taking the place of the first. This will be evident in the **Open Chart** dialog box, when the **System** text box is grayed out and only the last data file loaded is present.

New Data Overwrites Old

The **New Data Overwrites Old** option is used in conjunction with the **Add to Currently Open Data** provided the new data file (to be loaded) is from the same host system as the currently open data file(s).

- When both options are selected from the **Open** (data file) dialog box, the new data is added to the open data table. In instances when the collection periods (time ranges) overlap, the most recent data will overwrite the earlier data.
- When the **New Data Overwrites Old** check box is cleared and the **Add to Currently Open Data** check box is selected, only the new data from new collection periods (time ranges not already included in the currently open data table) will be loaded and added to the data table.

To view the top-down order of the composite data table, view the **Current File Statistics** dialog box (see "Current File Statistics" on page 82).

NOTE The user has the ability to cancel a file load using the Cancel button in the dialog box (see Figure 6.4). This dialog box is visible if the file is big enough and it takes a while to load.

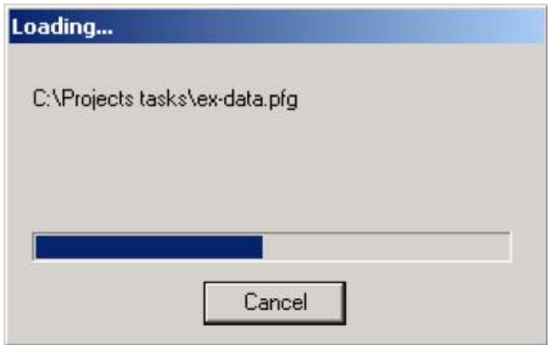

Figure 6.4 *Loading a file dialog box*

Close Data File

To close the active data file or composite data table, select **Close...** in the **File** menu, or press the Ctrl+F shortcut keys. If data from multiple systems is loaded, a second dialog box (see Figure 6.5) will appear which allows you to choose which system's data file you would like to close.

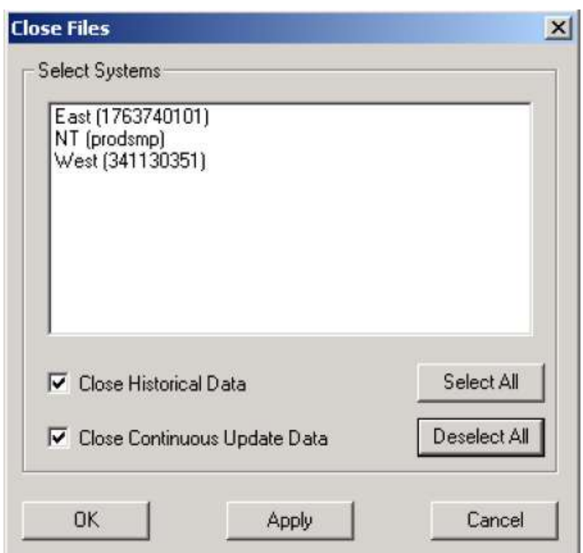

Figure 6.5 *Close Files dialog box*

Close All Data

To close all data files currently open with Performance Gallery Gold, select **Close All** in the **File** menu, or press the Ctrl+H shortcut keys.

System Alias Manager

The System Alias Manager dialog box has been added to assist in managing data files from multiple systems.

Because Performance Gallery Gold uses data from Meta-View Performance Manager, (which does not and cannot pass along information like system names) the only information available to differentiate one system from another is by the HP SUSAN number. This is not very user-friendly. As a result, the System Alias Manager has been developed in order to allow the user to assign user-friendlier names to the SUSAN numbers when a new data file is loaded. Figure 6.6 on page 81 displays a sample System Alias Manager dialog box with data files from two different systems loaded.

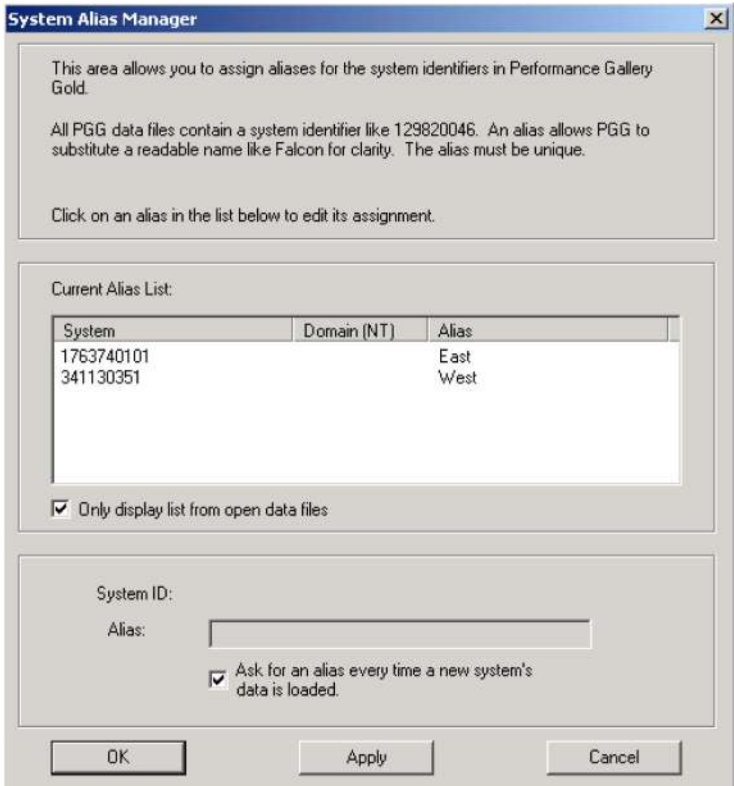

Figure 6.6 *System Alias Manager dialog box*

Two files from MPE/iX machines have been loaded. They have pretty similar SUSAN numbers, which would have made identifying and differentiating the systems within Performance Gallery Gold very difficult. Instead, both were given new aliases based on their machine names. There are two options that affect the display options of the Alias Manager:

- Only display list from open data files

  If this is checked, then only SUSAN numbers and assigned aliases from currently open data files will be displayed in the **Current Alias List**. To see all previously opened and aliased data files, uncheck this box.

- Ask for an alias every time a new system's data is loaded

  If this is checked, then as soon as a new data file is opened within Performance Gallery Gold, this dialog box will appear. If this is unchecked, then the data file will open with only the SUSAN number in the header of any opened graphs.

You may assign an alias to a data file at any time:

1 Click on **System Alias Manager...** in the **File** menu.

2 Select the file in the **Current Alias List**.

3 Type an alias in the **System ID: Alias:** text window. This must be a unique alias.

4 Click **Apply** if you want to continue assigning aliases, or **OK** to apply changes and close the **System Alias Manager** dialog box.

Most Recently Used Data Files

The **File** menu lists up to four of the most recently used data files. These files are numbered chronologically, **1** being the most recently used. Simply select a recently used data file in the **File** menu to open it.

Current File Statistics

To view a list of all data files that make up the composite data table in the top-down view order, click **Current File Statistics** in the **File** menu. The **Current Data File Statistics** dialog box provides the following information for each open data file:

- The name of the data file.
- The host system from which the data file was collected.
- The start time of the data collection/extraction period.
- The end time of the data collection/extraction period.

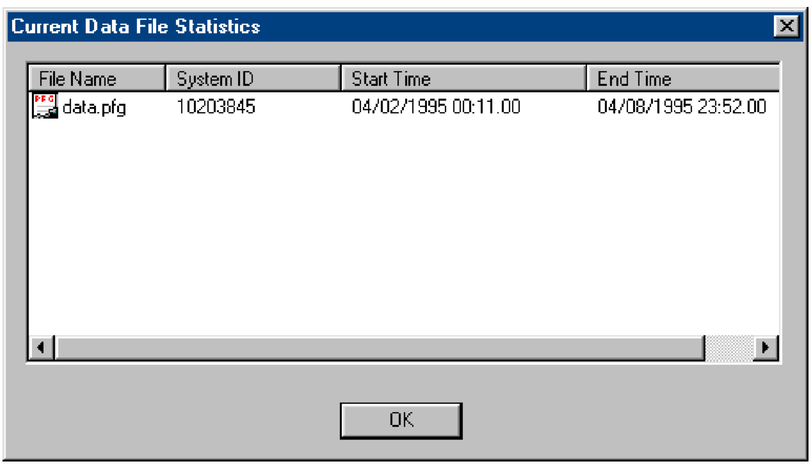

Figure 6.7 *Current Data File Statistics dialog box*

Print Current Chart

To send the currently-selected graph or table (or selected cells from a table) to a printer, click **Print Current Chart** in the **File** menu.

Before the chart is actually sent to the printer, a **Print** dialog box will prompt you to check the print settings. Adjust as necessary, then click **OK** to send the print request.

NOTE A table with a range so extensive that it requires scrolling to view some cells will be printed on multiple pages.

Print Preview

To preview the print job prior to printing, send the selected material to the **Print Preview** screen.

1. In the **File** menu, click **Print Preview**, or click the **Print Preview** toolbar button . One or more charts must be selected.
   - Click **Zoom In** to examine the page in closer detail.
   - Click **Zoom Out** to view the entire page.
   - Click **Next Page** to view the next page if necessary.
   - Click **Previous Page** to view any previous pages, if necessary.
   - Click **Two Page** to view two consecutive pages, side by side.
2. After previewing the print job, either click **Close** to return to the program window, or select the **Print...** command to send the chart(s) to the printer.

Print All

To print all charts currently open, click **Print All** in the **File** menu, click the **Print** toolbar button or use the shortcut keys Ctrl+P.

NOTE The Print All command is not applicable to tables.

Multiple graphs will be printed one per page unless you specify otherwise (see “Graph Print Setting” on page 84).

Print Setup

Specify the parameters for your print request, such as the paper size and orientation, in the **Print Setup** dialog box. To display the **Print Setup** dialog box, click **Print Setup...** in the **File** menu.

Graph Print Setting

Choose **Graph Print Settings** in the **File** menu to choose how many graphs will print on a page.

1 Graph per Page

Print 1 Graph per Page is the default print setting associated with the **Print All** command. To print each of the selected graphs, one per page, click **1 Graph per Page** in the **Graph Print Setting** menu.

2 Graphs per Page

Select **2 Graphs per Page** in the **Graph Print Settings** menu to print two charts on one page. One chart will be printed on the upper half, and the second chart will printed on the lower half of the page.

4 Graphs per Page

Click **4 Graphs per Page** in the **Graph Print Setting** menu to print four graphs per page; one graph in each quadrant of the page.

Exit

To exit the Performance Gallery Gold program, click **Exit** in the **File** menu.

Edit Menu

The **Edit** menu options enable the user to copy a chart to the Clipboard, copy a table to the Clipboard, clear any Performance Gallery Gold images from the Clipboard, and export the Performance Gallery Gold charts to other applications.

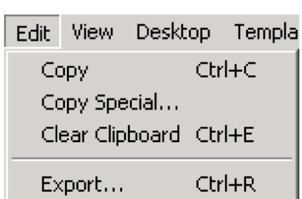

Figure 6.8 *Edit menu*

Copy

To move a copy of the active chart to the Clipboard, click **Copy** in the **Edit** menu, or use the toolbar button  or shortcut keys (Ctrl+C).

Use another graphics tool or word processor (such as Microsoft Excel, Paintbrush, or Word for Windows) to retrieve the copied chart from the Clipboard. This function can also be used to copy the information contained in the Exception Window onto the clipboard to be pasted into another program (Word, Excel, etc.)

Copy Special

The user can copy a table (with or without headers) to the clipboard, by selecting the cell(s) of the table and then clicking **Copy Special...** in the **Edit** menu. The Copy Special dialog box displays (see Figure 6.9).

Figure 6.9 *Copy Special dialog box*

Choose to include or not the headers of the table and also the desired delimiter, and click **OK**. Use another word processor (such as Microsoft Excel, or Word for Windows) to retrieve the copied table from the Clipboard.

Clear Clipboard

Unlike most Windows-based applications, the **Clear Clipboard** (or Ctrl+E shortcut keys) option in the **Edit** menu deletes all Performance Gallery Gold images from the Clipboard. This command does not affect any open chart(s) or the contents of the Clipboard associated with other programs.

Export

Performance Gallery Gold charts can be exported to other applications.

Graphs can be exported as:

- Windows or OS/2 Bitmap (*.bmp) files
- JPEG (*.jpg) files
- Portable Network Interface (*.png) files

Tables and the exception window can be exported as:

- Microsoft Excel Workbook (*.xls) files
- Tabbed Text (*.txt) files
- HTML (*.htm) or HTML Data Only (*.htm) files

To export a graph or table, do the following:

1. Select an open graph or table (select either specific cells or the entire table).
2. In the **Edit** menu, click **Export...** or use the Ctrl+R shortcut keys.
3. In the **Select Image Dimensions** dialog box, set the dimensions of the chart by selecting the appropriate button, or accept the default dimensions shown in the **Width** and **Height** boxes, which display the original dimensions of the chart.

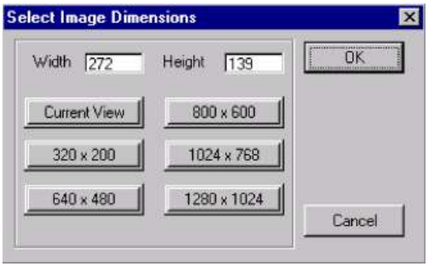

Figure 6.10 *Select Image Dimensions (for export) dialog box*

NOTE The dimensions must be within the range of 80 to 8000 pixels. If you attempt to set the width or height dimensions using values outside this range, the program will automatically default to the nearest value within the acceptable range (80 or 8000). For example, if you try to set the width to 60 pixels, the program will default to 80 pixels.

---

4 Click **OK**. A **Save As** dialog box will prompt you to save the graph or table to disk in the file format of your choice.

To export the exception window, do the following:

1 Select the exception window.
2 In the **Edit** menu, click **Export** or use the Ctrl+R shortcut keys.
3 Type a file name in the **File Name:** text box.
4 Select a file type from the **Save as Type:** text box.
5 Click **OK**.

View Menu

The **View** menu options affect the program window display.

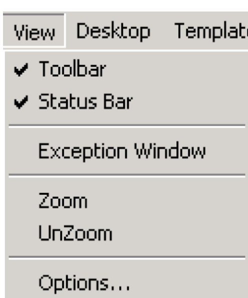

Figure 6.11 *View menu*

Toolbar

The **Toolbar** option displays a set of buttons used to carry out common menu commands. Toolbar buttons can change, depending on which window or view is currently selected.

Figure 6.12 *Performance Gallery Gold toolbar*

To hide the toolbar, clear the check mark next to the word **Toolbar** in the **View** menu. The absence of a check mark next to the **Toolbar** menu item indicates the option in disabled.

To display the toolbar, click the **Toolbar** option in the **View** menu again. The check mark next to the menu item indicates the option is enabled.

Status Bar

The status bar at the bottom of the main window displays information about the currently selected command, the active dialog box, the standard keys on the keyboard, or the current state of the program and the keyboard.

To hide the status bar, clear the check mark from beside the **Status Bar** option under the **View** menu. The absence of a check mark next to the menu item indicates that the option is disabled.

To display the status bar, click the **Status Bar** option in the **View** menu.

Exception Window

Performance Gallery Gold provides a quick and easy way to view system performance faults. If warning and critical thresholds have been set in the **Template** menu (see "Warnings" on page 141), the warnings will be displayed in the exception window. The exception window displays the following information:

- The name of the host.
- The time and date the exception occurred.
- The name of the chart containing the exception data.
- Whether the exception is a "warning"  or "critical"  exception.
- Which data set, if applicable, contained the exception.
- The exact value of the data point that caused the exception.

Double-clicking on any exception will bring the chart containing the exception to the front of the Performance Gallery Gold window.

The information in the **Chart Exceptions** window can be organized at the touch of a button by clicking on any of the column headers. For example, if you want to view exceptions based on when they occurred, earliest time and date at the top, click on the Time/Date column header. Clicking on the header again will cause the data to organize based on the time and date of each exception, but with the latest exception listed first.

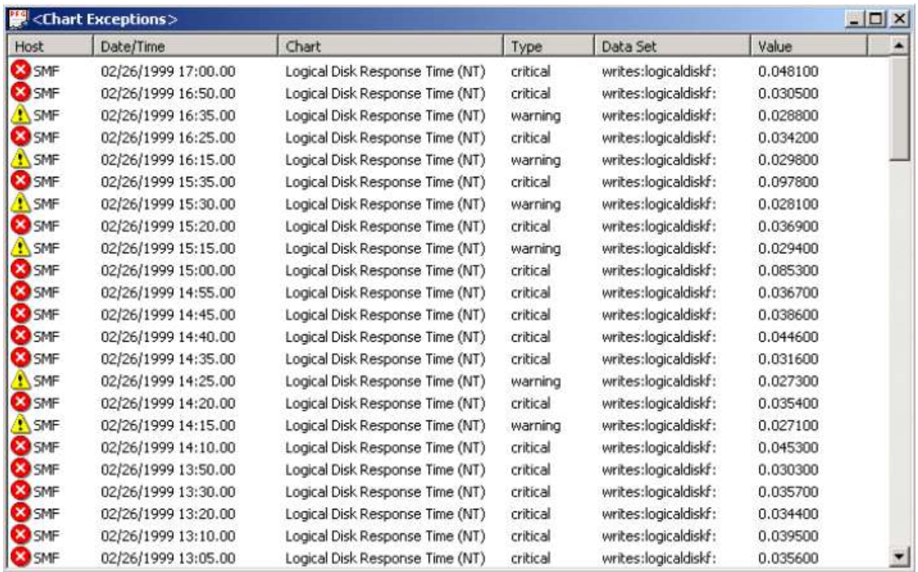

Host	Date/Time	Chart	Type	Data Set	Value
SMF	02/26/1999 17:00.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.048100
SMF	02/26/1999 16:50.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.030500
SMF	02/26/1999 16:35.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.028800
SMF	02/26/1999 16:25.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.034200
SMF	02/26/1999 16:15.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.029800
SMF	02/26/1999 15:35.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.097800
SMF	02/26/1999 15:30.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.028100
SMF	02/26/1999 15:20.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.036900
SMF	02/26/1999 15:15.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.029400
SMF	02/26/1999 15:00.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.085300
SMF	02/26/1999 14:55.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.036700
SMF	02/26/1999 14:45.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.038600
SMF	02/26/1999 14:40.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.044600
SMF	02/26/1999 14:35.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.031600
SMF	02/26/1999 14:25.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.027300
SMF	02/26/1999 14:20.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.035400
SMF	02/26/1999 14:15.00	Logical Disk Response Time (NT)	warning	writes:logicaldiskf:	0.027100
SMF	02/26/1999 14:10.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.045300
SMF	02/26/1999 13:50.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.030300
SMF	02/26/1999 13:30.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.035700
SMF	02/26/1999 13:20.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.034400
SMF	02/26/1999 13:10.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.039500
SMF	02/26/1999 13:05.00	Logical Disk Response Time (NT)	critical	writes:logicaldiskf:	0.035600

Figure 6.13 *Exception Window*

To access the exception window, click **Exception Window** in the **View** menu. The presence of a check mark next to the menu item indicates that the exception window is currently open.

To close the exception window, clear the check mark from beside the **Exception Window** menu item in the **View** menu, or click the close button in the top right corner of the window.

To export the **Chart Exceptions** window, ensure that the exception window is currently selected and click on **Export...** in the **Edit** menu. Select the preferred format and enter a title in the **File Name** field. Click **Save**.

Zoom

Selecting **Zoom** from the **View** menu allows you to view a portion of a graph more closely. To do this:

1. Select **Zoom** from the **View** menu or click the **Zoom** toolbar button .
2. Select the area to be enlarged by clicking the left mouse button and holding to select a rectangular area. Release the left mouse button to enlarge the chosen area. This can be repeated as often as necessary.

To return a graph to its original view, click **Unzoom** or press **R** on the keyboard.

To turn off the Zoom feature, click the **Zoom** toolbar button again, or deselect it in the **View** menu.

Unzoom

Clicking **Unzoom** from the **View** menu (or clicking the toolbar button ) returns a graph to its original view. If does not, however, turn off the Zoom feature.

Options

Selecting **Options...** in the **View** menu displays the **View Options** window. In this window, changes can be made to the appearance of the charts through the **Appearance** tab, while changes to the color scheme can be made through the **Line Colors** tab. The way the dates are displayed can also be modified in the **Localized Time Display** tab. The changes made to **View Options** are stored in the system's Registry, so the next time you log on and open Performance Gallery Gold, these settings will be saved, regardless of who else logged on and made changes.

The **Appearance** tab displays first when the **View Options** window opens.

Appearance

The **Appearance** tab controls the general appearance of all open charts, except for the color of the data lines (see "Line Colors" on page 93).

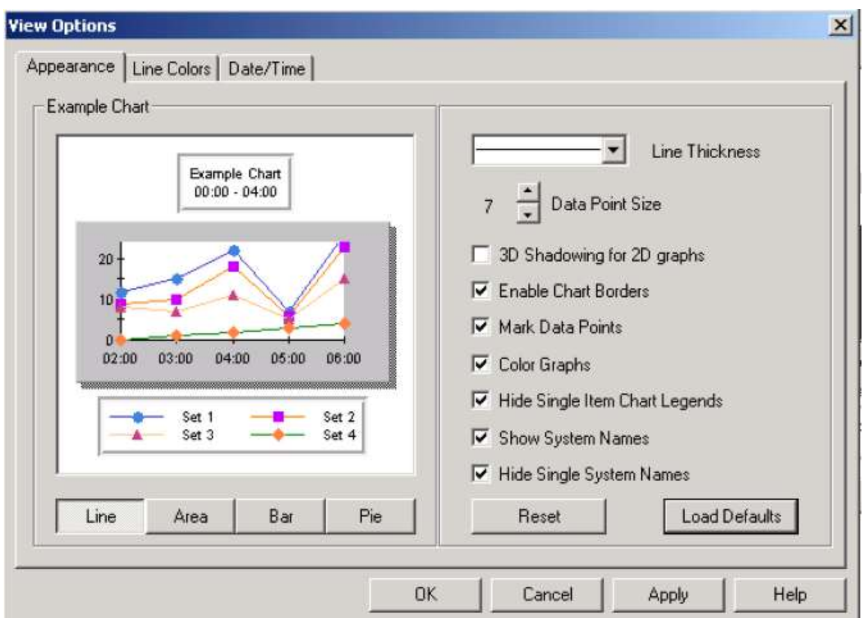

Figure 6.14 *View Options dialog box: Appearance tab*

Example Chart

The **Appearance** tab contains an Example Chart window that displays how the selected changes will look. How your graphs are affected by these options depends on what type of graph(s) are currently open.

The following nine changes can be made to the open graphs:

- Line Thickness
- Data Point Size
- 3D Shadowing for 2D graphs
- Enable Chart Borders
- Mark Data Points
- Color Graphs
- Hide Single Item Chart Legends
- Show System Names
- Hide Single System Names

There are four types of charts that can be modified in this window:

- Line
- Area
- Bar
- Pie

Note, however, that all four graphs cannot be modified using all nine appearance options. If a change is selected that does not apply to the example graph displayed, the window will default to a style of graph that does support the modification. For instance, if the bar chart button is currently selected, and you elect to reduce the size of the data points, the example window will display a line graph instead, since a bar chart does not utilize data points.

Line Thickness

The **Line Thickness** pull down menu displays five line choices, from 1 point to 5 points in thickness. This option can apply only to line charts, because none of the other three charts contain data lines. Click **Apply** to see the changes to all open line graphs.

Data Point Size

As with line thickness, this feature applies only to line charts. This function controls the size of the data points within a range of 1 to 99. The default size of the data points is 7. Use the up and down arrows to select the correct data point size for your graph(s) and click **Apply**. Changes will affect all open line charts.

3D Shadowing for 2D Graphs

This feature creates a three-dimensional shadow for two-dimensional bar graphs and pie charts. It does not affect line or area graphs. The presence of a check mark in the check box indicates this feature has been activated. Click **Apply** to modify all open bar or pie charts.

Enable Chart Borders

Enabling chart borders creates a shadow effect around the entire chart area. This is a feature that affects all four chart types and is the default choice when opening the **View Options** window. To remove chart borders from all open graphs, clear the check mark from beside the **Enable Chart Borders** option and click **Apply**.

Mark Data Points

This option displays or hides data points (the circles, triangles, etc.) used to identify different data elements in a line graph. Disabling this option causes line graphs to have data lines uninterrupted by data points, so the graph may look cleaner, but sampling intervals are not displayed. The presence of a check mark next to this option indicates that it has been enabled. Click **Apply** to modify all open line graphs.

NOTE This option must be enabled if Color Graphs is disabled in order to be able to differentiate between multiple data lines. If the Color Graphs option is disabled, the Mark Data Points option will automatically re-enable.

Color Graphs

All open graphs are affected by the **Color Graphs** option. Disabling this option causes all graphs to utilize different patterns, shades, and combinations of black and white. Line graphs containing multiple data lines can be adversely affected if displayed in black and white, because it is difficult to differentiate one black line from another as they cross or overlap. By default, **Color Graphs** is enabled, unless changes have been made during a previous session. To change all open graphs from color to black and white, clear the check mark from beside the Color Graph option and click **Apply**.

Hide Single Item Chart Legends

By default, the legend for single item graphs is hidden. However, by clearing the check mark from beside this option, you can elect to display the legend.

Show System Names

By default, Show System Names is enabled. To disable this option, clear this check box.

Hide Single System Names

By default, if only one system's data is loaded, the system name is hidden. However, by clearing the check mark from beside this option, you can elect to display the system name.

Reset

Clicking **Reset** will undo any changes made since the last time **Apply** or **OK** was selected.

Load Defaults

Clicking **Load Defaults** will return the **Appearance** tab to the factory-selected defaults, regardless of how many changes have been made. The following options are the factory-selected defaults:

- Previewed Graph = Line
- Line Thickness = 1st (thinnest line)
- Data Point Size = 7
- 3D Shadowing for 2D graphs = Disabled
- Chart Borders = Enabled
- Mark Data Points = Enabled
- Color Graphs = Enabled
- Hide Single Item Chart Legends = Enabled
- Show System Names = Enabled
- Hide Single System Names = Enabled

Line Colors

The Performance Gallery Gold program assigns a specific color to each possible data line in a line graph, or areas in area, bar, and pie graphs, using up to sixteen colors in a set sequential order. The color sequence is called the global "Line Color Order," and it can be viewed or modified in the **Line Colors** tab in the **View Options** window.

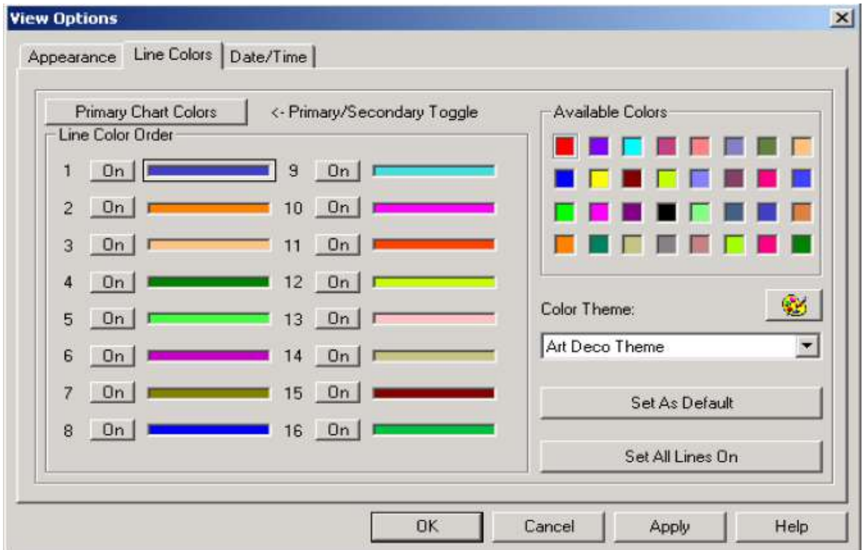

Figure 6.15 *View Options dialog box: Line Colors tab*

To open the **Line Colors** dialog box, click **View Options** in the **View** menu. Click on the **Line Colors** tab. The **Line Colors** dialog box can also be opened by clicking the Line Colors toolbar button  or by pressing the Ctrl+A shortcut keys.

The program will recycle the line/area colors in instances when a graph requires more data items than are currently available in the color sequence. For example, if a chart contains 24 data items and 14 of the colors in the **Line Color Order** group are **On**, the first 10 colors in the sequence will be used twice.

The user can select one color sequence for primary graphs, and select an entirely different color sequence for secondary graphs (see "Primary/Secondary Chart Colors" on page 95).

The user can also exclude one or more colors from the color sequence, or replace colors with colors from the **Available Colors** group box. The set of 32 colors in the **Available Colors** group box is from one of four preset color themes. Each of the Performance Gallery Gold color themes is made up of 32 unique RGB (combinations of Red, Green, and Blue) colors. The user can select one color from any of the four themes to replace a color in the line color order, or replace all sixteen colors in the line color order with the first sixteen colors in any color theme.

All **Line Color Order** properties with the exception of the color themes are global, meaning they will affect all charts opened in the current Performance Gallery Gold session. The currently-selected **Color Theme** property will apply only to the currently-selected chart, unless that color theme is set as the global default.

Primary/Secondary Chart Colors

Primary chart colors are the line and area colors selected for primary charts. Secondary chart colors apply to secondary charts, and to the color of the data markers for the primary data lines.

To display the primary and secondary chart colors, do the following:

1 In the **Select Line Colors** dialog box, locate the **Primary Chart Colors** command button.

   This button acts as a toggle switch between the **Primary Chart Colors** and **Secondary Chart Colors** properties. When the **Primary Chart Colors** button is displayed, the line/area colors displayed in the **Select Line Colors** dialog box are primary chart colors.

2 To display the color properties for the secondary chart, click the **Primary Chart Colors** button. It will toggle to "**Secondary Chart Colors**."

3 To return to the primary line/area colors, click the **Secondary Chart Colors** button.

Line Color Order

The global line color order includes the currently selected colors in the **Line Color Order** group box.

Figure 6.16 *Line Color Order group box.*

The program applies the currently selected line/area (including bar and pipe charts) colors in ascending order to each data item in a graph (denoted by a line or filled area).

To exclude a specific color from the **Line Color Order**, do the following:

1 Click the **ON** button that corresponds to the color to be switched **OFF**. (When selected, the **ON** button will switch to read **OFF**.) The first and third colors in the Line Color Order shown in Figure 6.16 are **OFF**.

2 Click **Apply**.

3 Repeat steps 1 and 2 to exclude other colors from the sequence, as needed.

4 Click **OK**.

To return a color to the sequence, do the following:

1 Click its **OFF** button once to switch it **ON**.

2 Click **OK**.

To return all 16 colors to the sequence, do the following:

1 Click **Set All Lines On**.

2 Click **OK**.

Available Colors

The **Available Colors** group box in the **Select Line Colors** dialog box contains a set of 32 colors from the currently selected color theme (see Figure 6.18 on page 96). Any color in the **Line Color Order** group box can be replaced with any color shown in the **Available Colors** group box. This type of change is global, meaning it will apply to all graphs opened in during the current Performance Gallery Gold session.

To replace one or more line/area colors, do the following:

1 In the **Line Color Order** group box in the **Select Line Colors** dialog box, select the color sample of the color you wish to replace. The third color in the sequence is selected in the following example (see Figure 6.17).

Figure 6.17 *Selected line/area color sample*

2 Select a new color from the **Available Colors** group box (Figure 6.18).

Figure 6.18 *Available colors group box*

3 Click **Apply**.

4 Repeat steps 1-3, as needed to replace other colors.

5 Click **OK**.

NOTE It is possible to use one color option more than once in the line color sequence, either intentionally or by mistake. Make sure to check the entire sequence carefully.

To return to the default line color sequence, click **Set As Default**, then click **OK**.

Custom Colors

Performance Gallery Gold allows you to create and use your own colors within graphs. To open the customer colors palette, click on the **Pallet** icon . This opens the **Line Colors** dialog box.

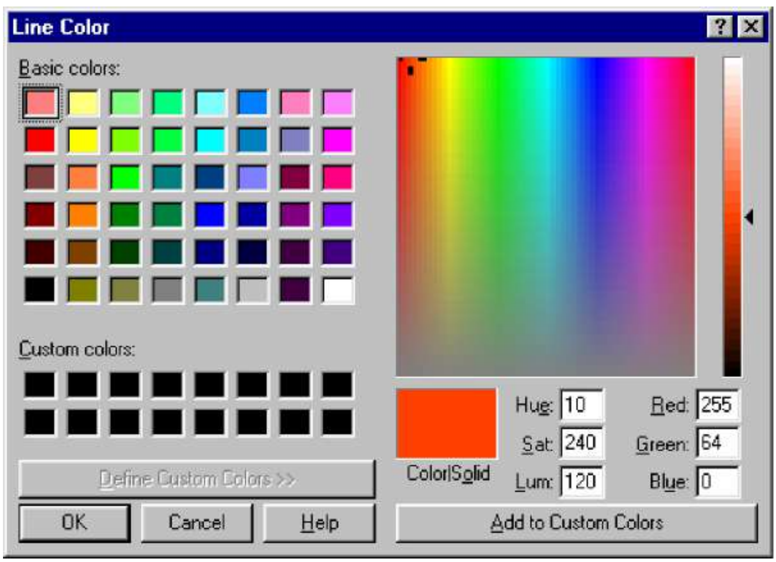

Figure 6.19 *Custom colors Line Color dialog box*

To create custom colors:

1 Select a color from the **Basic colors** selections, or click anywhere in the rainbow-colored palette box.

2 Fine-tune colors by moving the small black arrow at the far right of the dialog box up or down. The new color will appear in the box labeled **Color|Solid**.

3 Click **Add to Custom Colors**. The new color will be added to the **Custom colors** area and will be available for use any time Performance Gallery Gold is opened.

4 Click **OK** to exit the dialog box and to apply changes to the selected line in the Line Color Order area of the Line Colors tab.

Color Themes

Performance Gallery Gold offers the following five color themes:

- Art Deco (default)
- Autumn
- LPS Nostalgia
- Northern Regal
- Spiraling Mixture

Each of the five color themes is made up of 32 RGB colors. The first 16 colors of a theme (the 4x4 block of color tiles in the **Available Colors** group box) can be applied globally as the new default colors, or you can select any of the 32 colors in any theme to replace any color in the line color order.

To replace the current color theme with a new color theme, do the following:

1. In the **Select Line Colors** dialog box, select a theme from the **Color Theme** drop-down menu.
2. Click **Set as Default**. This will save the new theme as the global default.
3. Click **OK**.

To replace a single line/area color with a color from another theme, do the following:

1. In the **Select Line Colors** dialog box, select the color to be replaced in the **Line Color Order** group.
2. Select a theme from the **Color Theme** drop-down menu. All 32 colors of that theme will display in the **Available Colors** group box (see Figure 6.18).
3. Select a color from the **Available Colors** group box to replace the currently selected color in the **Line Color Order** group box.
4. Click **OK**.

Date/Time

Performance Gallery Gold allows you to change the date and time formats displayed in the charts in order to better serve our customers around the globe.

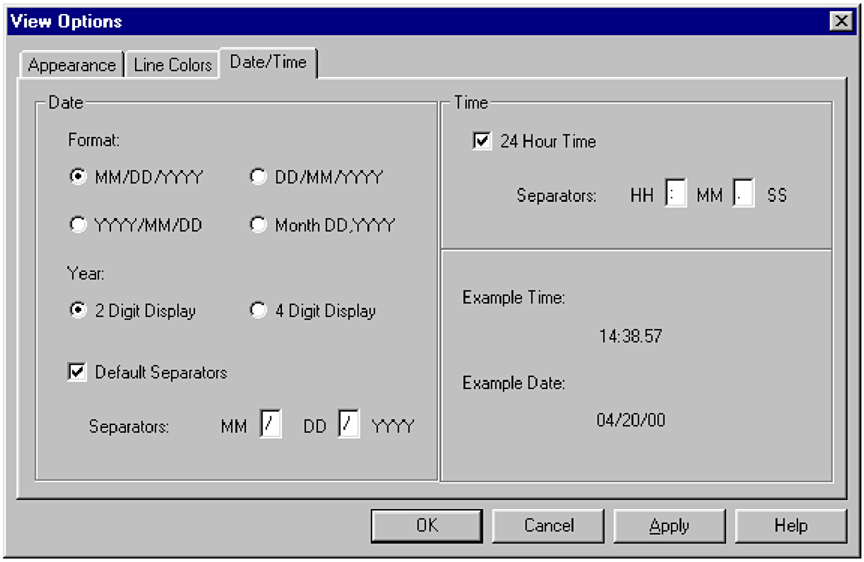

Figure 6.20 *View Options dialog box: Date/Time tab*

Date

Within the **Date** area of the **Date/Time** tab, there are options to change the format of the date:

- MM/DD/YY
- YYYY/MM/DD
- DD/MM/YYYY
- Month DD, YYYY

These options can be further modified by selecting to display a two or four digit year. The default separators for the date are slashes (/), or in the case of the Month DD, YYYY display option, a comma (,) is used, but this can be modified as well.

1. In the **View Options** dialog box, click the **Date/Time** tab.
2. Click in the **Default Separators** check box to remove the check mark.
3. Click in and type a new separator (i.e., a period (.), comma (,), or a hyphen (-)) in each of the windows between **MM**, **DD**, and **YY**.

Your changes will be displayed in the **Example Date** area. If you are happy with the changes, click **Apply** to apply the changes to your charts. Click **OK** to close the **View Options** dialog box.

Time

The times in your graphs can be displayed in two ways, either as a 24-hour clock or using A.M., and P.M., designations. By default, Performance Gallery Gold uses a 24-hour clock. To change your charts to utilize a 12-hour, A.M./P.M., clock, do the following:

1. In the **View Options** dialog box, click the **Date/Time** tab.
2. Click in the check box next to **24 Hour Time**.
3. If you want to change the separators, (from a colon (:) to a period (.), for instance) click in the separator windows and type a new separator.

The absence of a check mark indicates that the 24-hour display has been disabled. Your changes will also appear in the **Example Time** area in the **Date/Time** tab. If you want these changes to be applied to your charts, click **Apply**. Click **OK** to close the **View Options** dialog box.

Desktop Menu

The **Desktop** menu contains options to load, reload, save, and close a desktop file (see Figure 6.21). The menu also provides five links to the most recently used desktop files and options to set and load a default desktop file. For information about desktop files, refer to “Desktops” on page 11.

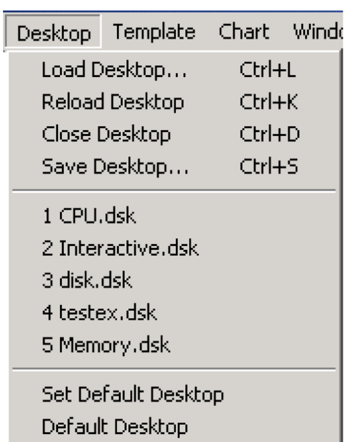

Figure 6.21 *Desktop menu*

Load Desktop

To load a previously saved Performance Gallery Gold desktop file, do the following:

1 In the **Desktop** menu, click **Load Desktop...**. As an alternative:

- Click the Load Desktop toolbar button 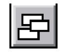
- Type the Ctrl+L shortcut keys

2 In the **Open** dialog box, select a desktop file (*.dsk).

3 Click **Open**

NOTE If multiple host data files are currently open, a dialog box may appear asking which data file the desktop should be applied to (see Figure 6.22). Simply choose from the list of open files and click OK.

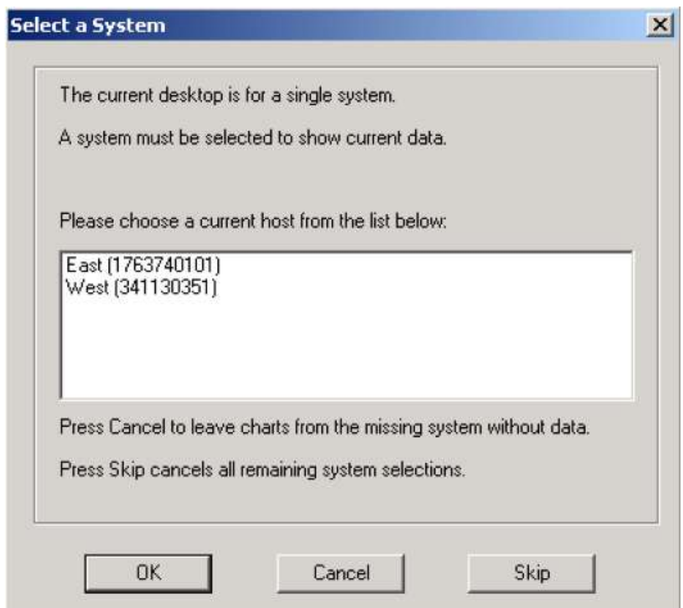

Figure 6.22 *Select a System dialog box*

See Figure 6.23 for an example of active Performance Gallery Gold desktop. The current desktop file name is displayed in the program's caption bar.

Reload Desktop

Clicking **Reload Desktop** will reload the currently loaded desktop in its saved form. Any unsaved changes made to the desktop will be lost when this is selected.

Close Desktop

To close an open desktop file, click **Close Desktop** (or use Ctrl+D shortcut keys) in the **Desktop** menu. The program will prompt you to save the active desktop.

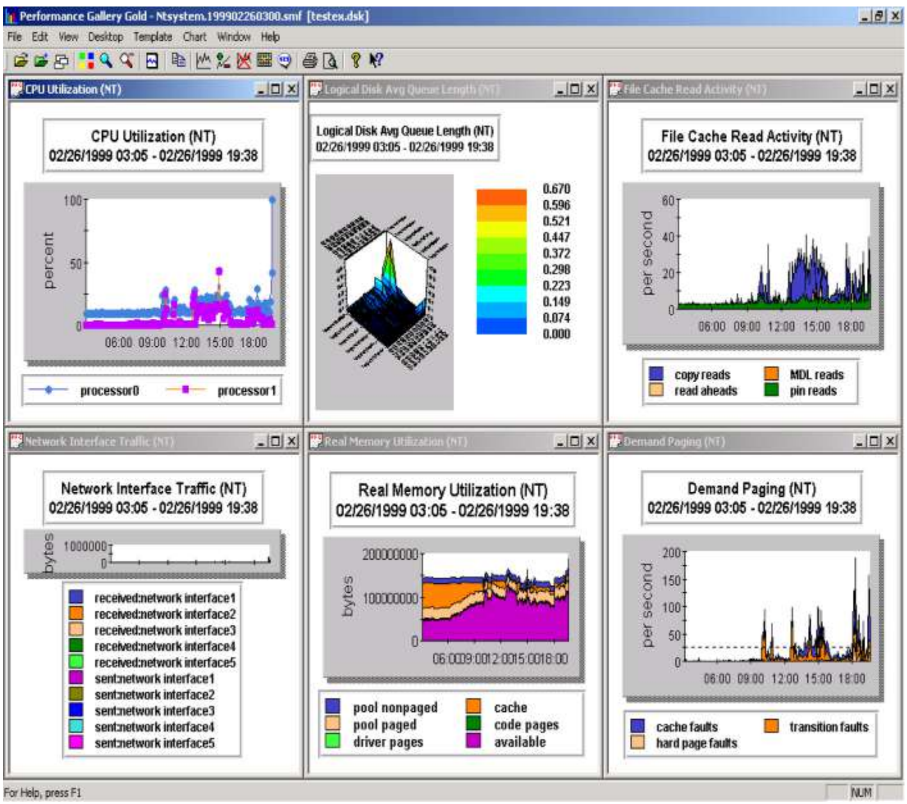

Figure 6.23 *Example of active Performance Gallery Gold desktop*

Save Desktop

The **Save Desktop...** menu option is a unique feature that enables you to create and save a group of performance graphs and tables as a "desktop."

To save an active desktop, do the following:

1 In the **Desktop** menu, click **Save Desktop...** or use the Ctrl+S shortcut keys.

2 In the **Save As** dialog box, type a file name for the desktop you wish to save in the **File name** text box. Type only the file name; the extension ".dsk" will be appended automatically.

   All desktop files are saved in the default destination file **C:\...\Performance Gallery Gold**, unless specified otherwise.

3 Click **Save**.

Recently Used Desktops

The Performance Gallery Gold program remembers the last five desktop files opened and displays their titles in the **Desktop** menu. These desktop files are numbered chronologically, **1** being the most recently used. Simply select the needed recently-used file in the **Desktop** menu to open it.

Set Default Desktop

To designate a specific desktop file as the default desktop, do the following:

1 In the **Desktop** menu, click **Set Default Desktop**. A **Desktop Query** dialog box will display the question, "Set the current desktop to be the default desktop?"

2 To save the open chart(s) or desktop file as the default desktop, click **OK**.

NOTE A desktop must be saved before it can be set as the default desktop.

Default Desktop

To load the default desktop file, select **Default Desktop** in the **Desktop** menu. The check mark next to the menu item indicates the Default Desktop option is enabled. The default desktop will load automatically the next time a data file is opened.

NOTE If a desktop other than the default desktop is currently in use, the program will continue to display that desktop instead of loading the default desktop.

To disable the default desktop, clear the check mark from beside the **Default Desktop** option in the **Desktop** menu. The absence of a check mark next to the menu item indicates that the option has been disabled.

Template Menu

In addition to providing a variety of pre-configured graphs and tables, the Performance Gallery Gold graphics program enables you to customize charts. The **Template** menu provides the tools to create new charts and modify pre-configured charts. Utilization of these tools is discussed in "Template Menu and Commands" on page 109.

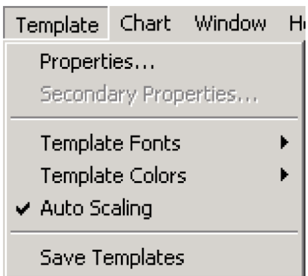

Figure 6.24 *Template menu*

Chart Menu

The **Chart** menu contains options to open, modify, and close a chart, among several other chart display options. For information regarding the Performance Gallery Gold **Chart** menu, see "Chart Menu and Commands" on page 149.

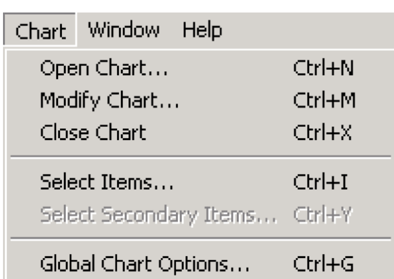

Figure 6.25 *Chart menu*

Window Menu

The **Window** menu options are standard Windows-based options that modify the viewing arrangement of multiple charts. These arrangements can be saved as default settings in a desktop file.

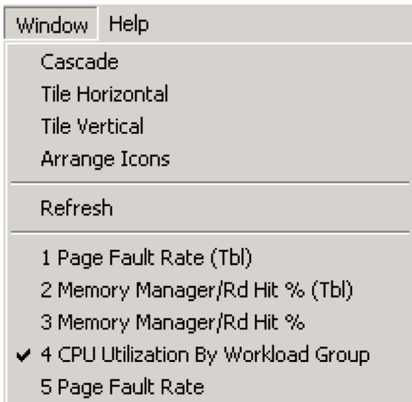

Figure 6.26 *Window menu*

Cascade

To arrange all open charts in the order in which they were opened, cascading from the upper left to the lower right portion of the program window, click **Cascade** in the **Window** menu.

Tile Horizontal

To arrange all open graphs horizontally in "tiles" of proportionate size, using the full extent of the program window, click **Tile Horizontal** in the **Window** menu.

Tile Vertical

To arrange all open graphs vertically in "tiles" of proportionate size, using the full extent of the program window, click **Tile Vertical** in the **Window** menu.

Arrange Icons

To display the icons of any minimized charts, left to right across the bottom of the program window in chronological order, click **Arrange Icons** in the **Window** menu.

Refresh

To refresh the Performance Gallery Gold program window and all open charts, click **Refresh** in the **Window** menu.

Active Charts

In the lower portion of the **Window** menu is a list of all the charts currently open in the program window. To bring a specific chart to the front, select it in the **Window** menu.

Help Menu

Figure 6.27 *Help menu*

Help Topics

The Performance Gallery Gold Help system is a quick-reference of helpful topics. These topics are indexed and can be recalled by performing a keyword search. To access the indexed Help system, click **Help Topics** in the **Help** menu.

Context-sensitive Help

To get context-sensitive Help on a specific dialog box, click the **Help** button in the dialog box. Help about the dialog box will display in a separate window.

To get context-sensitive Help elsewhere in Performance Gallery Gold, do either of the following:

- Position your mouse over the item in question and press the F1 function key. This is especially helpful for open dialog boxes that do not contain a **Help** button, as the Help Toolbar button will not open dialog boxes and cannot be selected if a dialog box is already open.

  OR

- Select the Help toolbar button  for use with toolbar buttons or menu items. This will change the cursor to a question mark. Move the question mark pointer over an item and click to get Help for that item.

NOTE The online Help was designed with and for Internet Explorer v4.0 or higher. Using an older version of IE4 or a different browser may result in the pages displaying improperly.

About Performance Gallery Gold

To display the current version and copyright information for your Performance Gallery Gold software, select **About Performance Gallery Gold** in the **Help** menu or click the About toolbar button .

7

Template Menu and Commands

Overview

In addition to providing a variety of pre-configured graphs and tables, the Performance Gallery Gold graphics program enables you to customize charts. The **Template** menu provides the tools to create specialized charts, either by modifying an existing chart template, or by building an entirely new chart template.

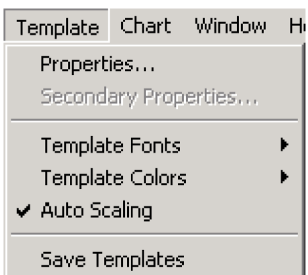

Figure 7.1 *Template menu*

Properties

To create custom chart templates, click **Properties** in the **Template** menu, or click the Properties toolbar button .

The **Template Properties** dialog box displays six tabs: **General**, **Graph**, **Table**, **Thresholds**, **Data**, and **Links**; and eight buttons: **New**, **Rename**, **Apply**, **Save**, **Save As...**, **Delete**, **Open Chart**, **Help**, and **Close** (see Figure 7.2).

Clicking in the **Validate** check box narrows the options displayed in the **Template Name** drop down menu to only those templates that can be used with the type of data file currently open. For instance, if the current data file is from a Unix system, only Unix-compatible templates will be displayed. Leaving this box unchecked, however, allows you to modify templates for any system without having to open a data file first. This setting is saved between Performance Gallery Gold sessions.

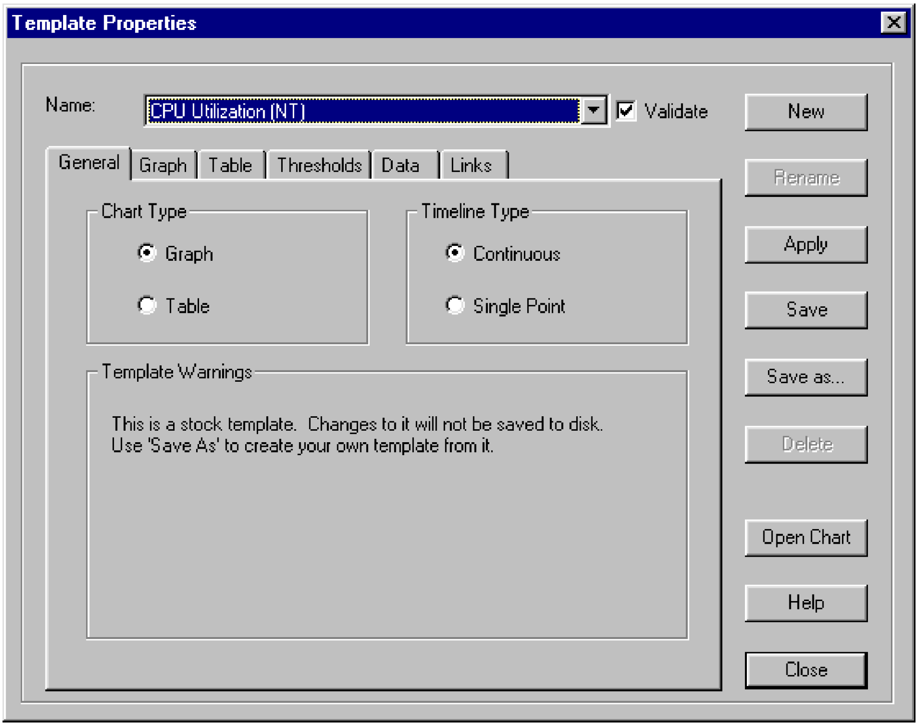

Figure 7.2 *Template Properties dialog box*

Buttons

New

Clicking this button clears the **Name** drop down window so that a different template may be chosen. If you want to create your own template with a new name, click **New**, then **Save As...** to give the new template a name. In order to create a new template, any existing template data should be cleared. If modifying an existing template, select the template to modify in the **Template Name** drop-down list box.

NOTE If you select New, then begin altering the template without clicking Save As..., then you may lose all of your changes if you accidently choose another template from the drop down menu before saving. We recommend giving your new template a name before making any changes to it.

Rename

This button opens the **Set Template Name** dialog box. From this dialog box, user-defined templates can be renamed. Performance Gallery Gold templates cannot be renamed, so the **Rename** button is only active if the chart selected in the Template Name window is a user-defined template.

Apply

Click **Apply** to apply your changes to any template without closing the **Template Properties** dialog box. These changes are lost once Performance Gallery Gold is closed unless you opt to save your template changes upon closing Performance Gallery Gold.

Save

Close temporarily saves changes to the template and closes the **Template Properties** dialog box. If changes were made to a user-defined template, you will be prompted to save these changes again before closing Performance Gallery Gold.

Save As...

If you want to make changes to an existing Performance Gallery Gold template and save it to be used during another session, use the **Save As...** button to give the template a new name.

NOTE When saving a new or modified template, the customized template isn't saved in place of an existing template. The template to be saved should be renamed. Modified templates should be saved using the Save As... command.

Delete

Use the **Delete** button to remove user-defined templates from the drop down menu. Again, because Performance Gallery Gold templates cannot be modified, this button is only active if a user-defined template is selected in the **Name** window. Use this button with caution, as it will not ask you to verify your delete decision.

Open Chart

This button opens the Open Chart dialog box (see "Open Chart" on page 149), and allows you to open new charts without closing out of the Template Properties dialog box

Help

Clicking on the **Help** button from the **Template Properties** box opens the online context-sensitive help that describes the **Template Properties** dialog box.

Tabs

General Tab

Under the **General** tab are three sections: **Chart Type**, **Timeline Type**, and **Template Warnings**.

- To set the type of chart to be graphical or tabular, select the appropriate chart type in the **Chart Type** area (see "General Properties" on page 114).
- To condense the chart data into a single data point - an average of the eligible data, select **Single Point** in the **Timeline Type** area. **Continuous** is selected by default. For more information on Single Point and Continuous timelines (see "Timeline Type" on page 114).
- Any warnings regarding template properties are displayed in this section (see "Template Warnings" on page 116).

Graph Tab

Under the **Graph** tab are five options: **Labeling**, **Type**, **Scale**, **Stacking**, and **Marker**.

- The **Labeling** text box displays the Y axis label of the currently-selected graph (percent, ms, kb, pages, etc.).
- Select one of the eight graph types found in the **Type** area (see "Graph Properties" on page 116).
- Select one of the three y axis scale options in the **Scale** area (see "Scale" on page 128).
- Check the **Enable Data Stacking** check box if chart data should be displayed without overlap (applicable to line, area, bar, radar, and area radar charts only, as described in "Stacking" on page 129).
- To add a horizontal marker as a reference line, check the corresponding check box, and enter a value in the **Marker** field (see "Marker" on page 131).

Table Tab

Under the **Table** tab are three options: **Type**, **Order**, and **Options**.

- Select one of the two table formats available in the **Table** area: **Time Indexed** and **Snap Shot** (see "Type" on page 132).

- Select the order in which to display the chart data in the **Order** area (see “Order” on page 134).
- The **Scroll Snap Shot** and **Color Coded** options control how the data is displayed in the table (see “Options” on page 134).

Thresholds Tab

Under the **Thresholds** tab there are two options: **Maximum**, and **Minimum**.

- In the **Maximum** options section select the **Enable Threshold** option and then select or set the other options as needed (see “Maximum and Minimum Thresholds” on page 135).
- In the **Minimum** options section select the **Enable Threshold** option and then select or set the other options as needed (see “Maximum and Minimum Thresholds” on page 135).

Data Tab

Under the **Data** tab are seven options: **System**, **Group**, **Element**, **"All" group type**, **Sign**, **Add Options**, and **Delete Options**. The **Data** tab also contains a large computation text box.

- Select a system from the **System** drop-down list box. This represents the currently open host data file(s) (see “System” on page 138).
- Select a group from the **Group** drop-down list box. This represents a group for which a particular element exists (see “Group” on page 138).
- Select an element of the group from the **Element** drop-down list box. This represents a particular item in the group (see “Element” on page 138).
- To select an entire set of groups (for example, all drives, all processors), select all and the group of to be monitored (such as, "All discs" or "All workloads") in the **Group** drop-down list box. To set how the data is calculated and displayed, click the appropriate option in the **All group type** section. (Refer to “All group type” on page 139 for more information.)
- To set the value of the divisor and numerator (see “Add Options” on page 140 and “Delete Options” on page 142) as being either negative or positive, select the appropriate sign in the **Sign** area (see “Sign” on page 140).
- To add an entry to the currently selected template, click **Entry**, located under the **Add Options** area (see“Add Options” on page 140).
- To delete an entry from the currently selected template, click **Entry**, located under the **Delete Options** area (see“Delete Options” on page 142).

Links Tab

Under the **Links** tab there are two option sets: **Link (1-5)**, and **Separators**.

- To make other relevant charts readily available, select the appropriate charts in the **Link 1** through **Link 5** drop-down list boxes (see “Links 1-5” on page 144).
- To place a separator between links in the right-click shortcut menu, check the corresponding **Separator** check box (see “Separators” on page 145).

General Properties

The **General** tab includes three sections: **Chart Type**, **Timeline Type**, and **Template Warnings**. The **Chart Type** and **Timeline Type** options control the essential format of the data.

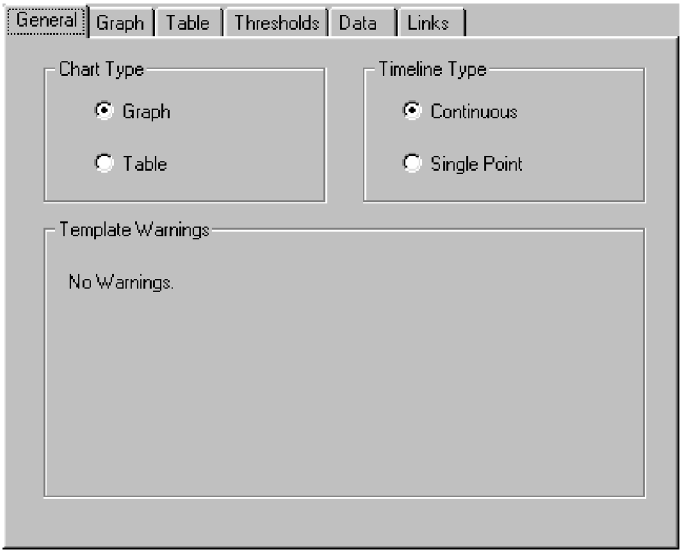

Figure 7.3 *Template Properties dialog box: General tab*

Chart Type

The **Chart Type** menu area offers two options for presentation: **Graph**, and **Table**.

Graph

Graph allows for a graphical presentation of data in several different formats. These formats include **Area**, **Line**, **Bar**, **Pie**, **3D Bar**, **3D Surface**, **Radar**, and **Area Radar**. These graph types are described in "Graph Properties" on page 116.

Table

Table allows for a tabular presentation format. (For more information on tables, refer to "Table Properties" on page 132.)

Timeline Type

The **Timeline Type** area allows data to be displayed as a continuous graph or a single, combined data point.

Single Point

A single point graph is displayed as an average. All of the data points that are logged are then averaged into a single point.

A graph using the **Single Point** option may look like the example in Figure 7.4.

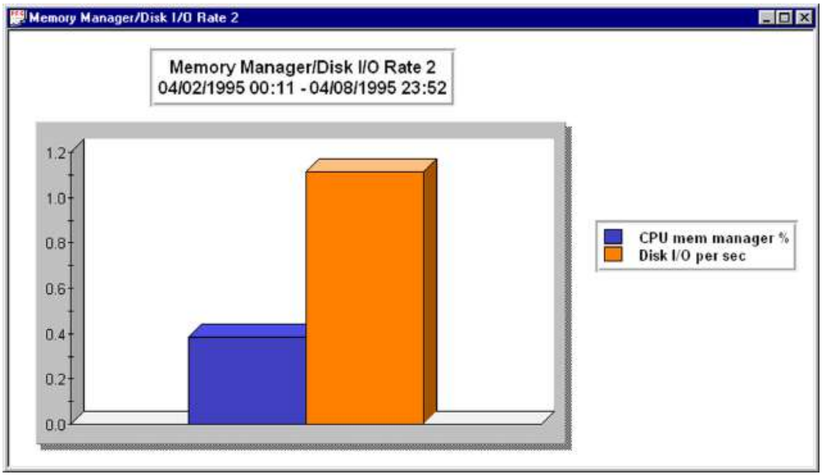

Figure 7.4 *A single point bar graph*

Continuous

A continuous graph displays the data points from each collection interval. This kind of graph shows trends and variations over the collection period, rather than the average of all data points of each data set for the entire collection period.

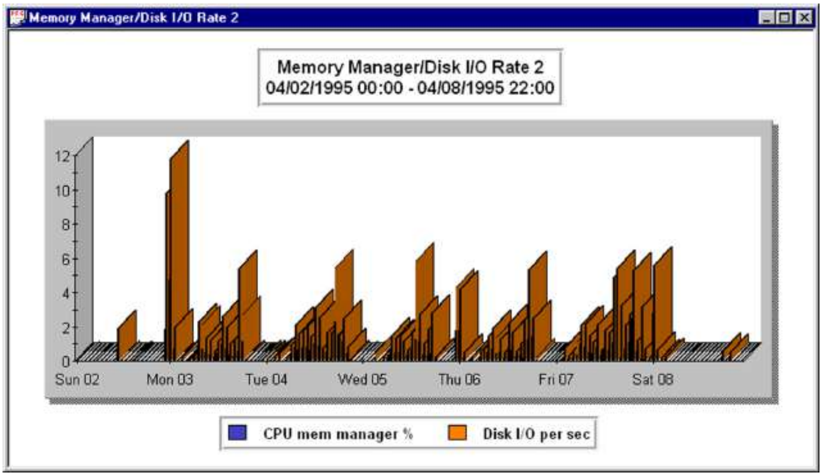

Figure 7.5 *A continuous bar graph*

Template Warnings

The third section of the **General** tab displays template warnings. If you have a Performance Gallery Gold template open, you will be alerted that you will not be able to save changes to the template (see Figure 7.2). If you click **New** to create a new template, all of the following warnings will be displayed until you give your new template a name, links, and data:

- No template name specified
- No template links defined
- This template contains no data

If a chosen template lacks defined links (see "Links Properties" on page 144) or data (see "Data Properties" on page 137), then the appropriate message will display in the **Template Warnings** dialog box.

Graph Properties

The **Graph** tab offers several different options for **Labeling**, **Graph Type**, **Scale**, **Stacking**, and **Marker**.

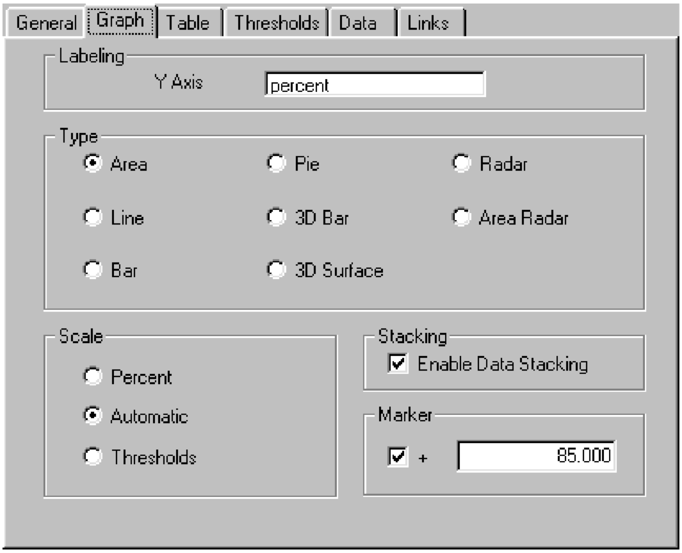

Figure 7.6 *Template Properties dialog box: Graph tab*

Labeling

The **Y Axis** text box displays the default Y axis label for the currently selected graph.

To change the Y axis label:

1. In the **Graph** tab of the **Template Properties** dialog box, click in the **Labeling Y Axis** text box.
2. Type the new label (percent, kb, queue length, pages, etc.).
3. Click **Apply**. After all modifications have been made, click **Close**.

Graph Type

Among the possible graph types that can be displayed are the following: **Area**, **Line**, **Bar**, **Pie**, **3D Bar**, **3D Surface**, **Radar**, and **Area Radar**.

Area Graphs

Area graphs are primarily used to plot changes in quantity or position and are useful for illustrating trends. A few examples of how this type of graph my be useful follow:

- Checking CPU workloads.
- Comparing between multiple data sets.
- Comparing storage device performance.

Area graphs can be displayed in two separate fashions: Overlapped or Stacked. (For more information on stacked graphs, see “Stacking” on page 129.)

An example of a stacked area graph may look like this:

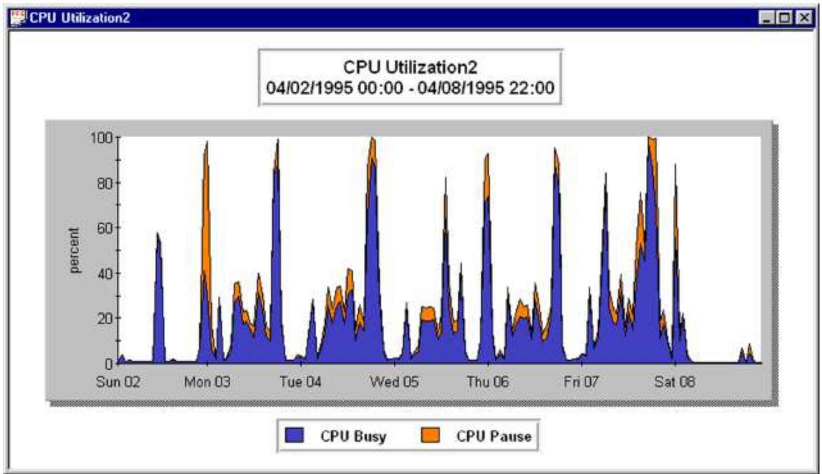

Figure 7.7 *A stacked area graph*

The CPU Pause data is stacked on top of the CPU Busy data, and all data points are visible.

An overlapped area graph is usually not used to view graph data due to possible obscuring of data. Line graphs are generally used instead. This is an example of an overlapped area graph:

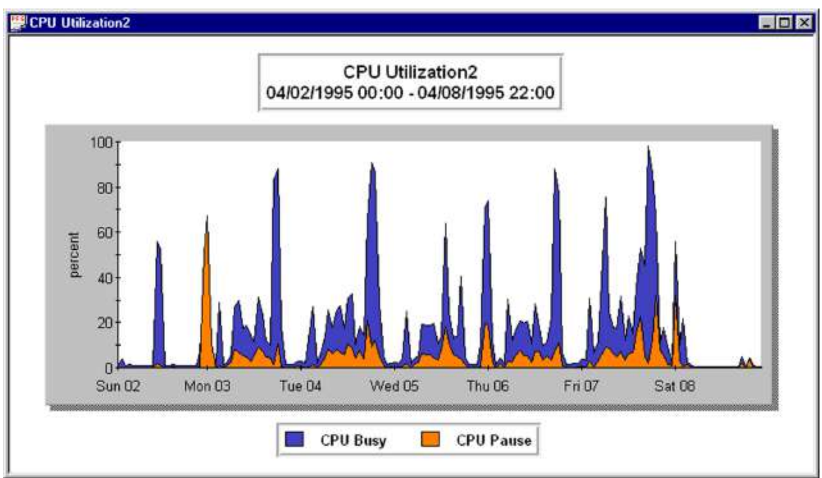

Figure 7.8 *An overlapped area graph*

Line Graphs

Line graphs present much the same data as area graphs, however they are not filled. Line graphs may be more practical than other types of graphs when:

- The data points from different data sets overlap to the extent that individual data points are hidden from view.
- Viewing only a few data sets.

An example of a line graph that is not stacked is shown in Figure 7.9.

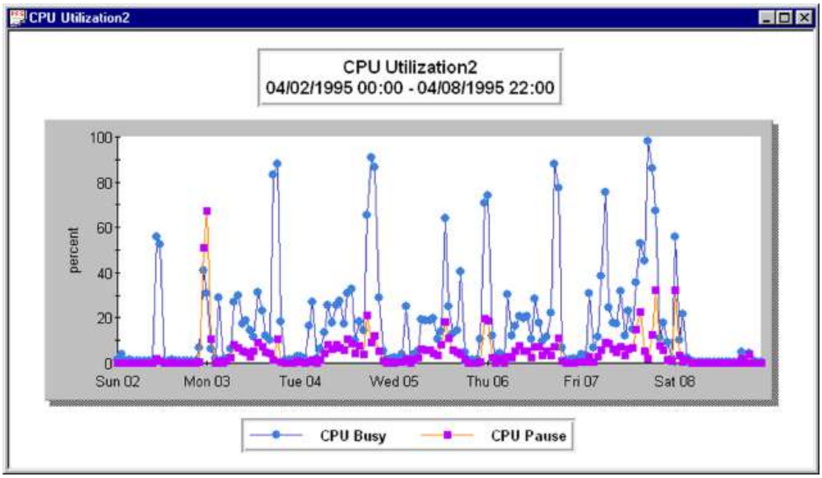

Figure 7.9 *An unstacked line graph*

Bar Graphs

Bar graphs are primarily used to:

- Compare items at one particular time or over a period of time.
- View changes in one item over a period of time.
- Compare portions of a single item.

The horizontal and vertical axes represent the two elements being illustrated, such as time and quantity, and can be viewed either stacked or unstacked (see "Stacking" on page 129). A bar graph that is stacked may look like the example shown in Figure 7.10.

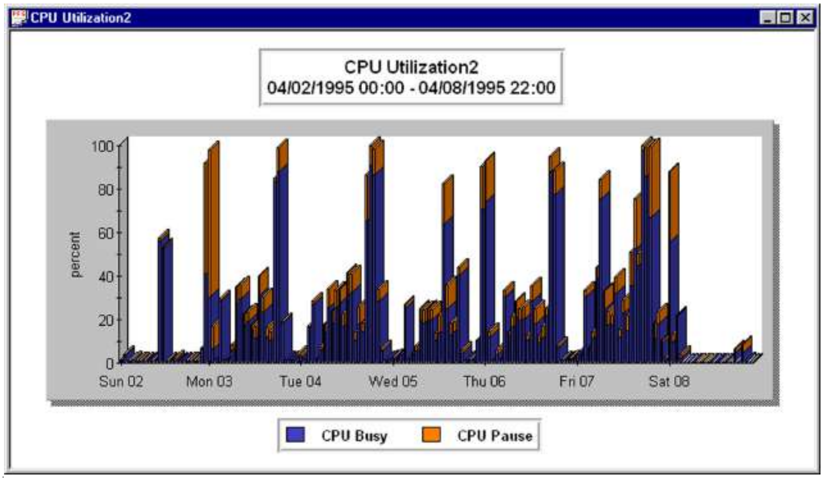

Figure 7.10 *A stacked continuous bar graph*

Occasionally, there may be a need to view two data sets side-by-side as opposed to stacked. When the **Stacking** option is disabled, the graph may be viewed in this manner (Figure 7.11).

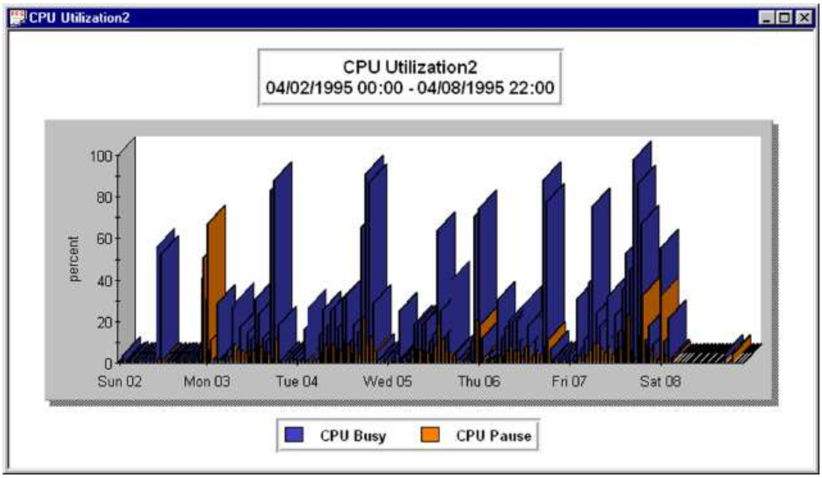

Figure 7.11 *An overlapped (unstacked) bar graph*

Pie Graphs

Pie graphs use a circle to represent a whole unit. A "whole unit" would normally represent percentages or a quantitative analysis. Examples of this could be:

- CPU Utilization by process or workload
- Comparing Memory Usage to Cache Usage
- Comparing Disk Pause and Read Hit Percentages

Pie charts can only be displayed when using the **Single Point Timeline** type (see "Timeline Type" on page 114). An example of a pie graph is shown in Figure 7.12.

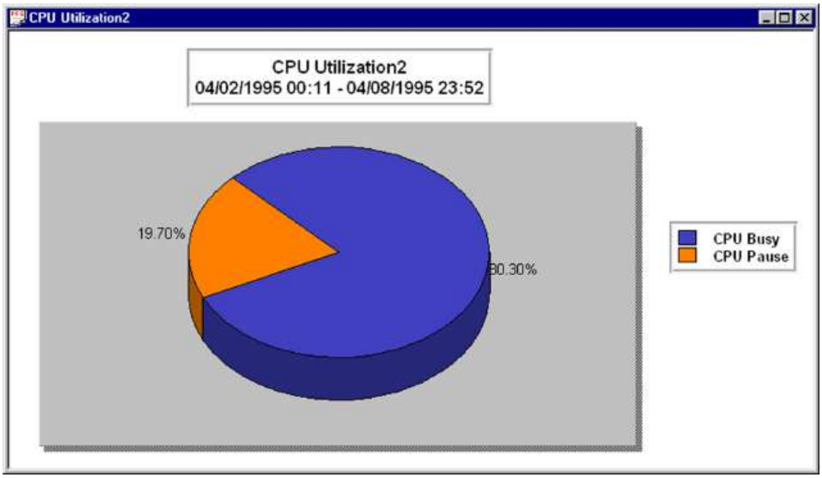

Figure 7.12 *A pie chart*

Radar Graphs

Radar graphs are line graphs wrapped around a center point. They are most useful when looking at either 12-hour or 24-hour data, as it presents data in such a fashion as to be in a "clock" format. This makes seeing intervals of spiking easier. An example of a radar graph with data stacking enabled is shown in Figure 7.13.

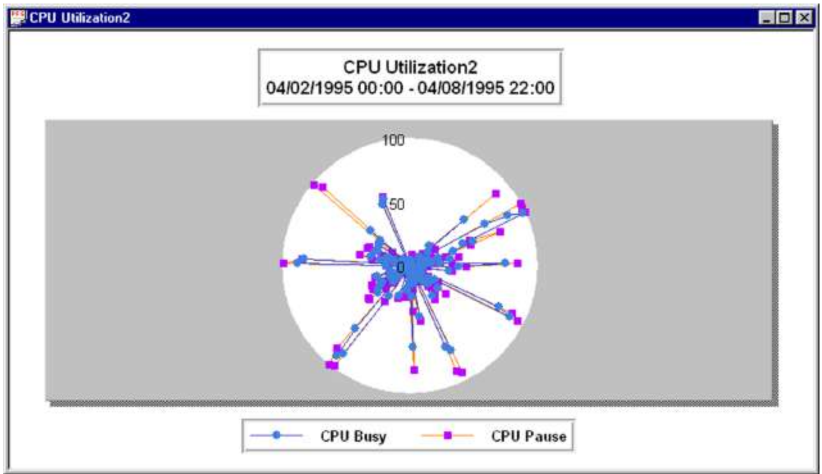

Figure 7.13 *A stacked radar graph*

The figure on the left represents a conceptual model of a radar graph. Each degree of the circle represents a percentage of time (in this case, hours), the outermost circle represents the maximum (in this case, 100), and the center circle represents a halfway point, presented here as 50.0 for clarity.

A radar graph is essentially a line graph (see “Line Graphs” on page 120) "wrapped" around a center point.

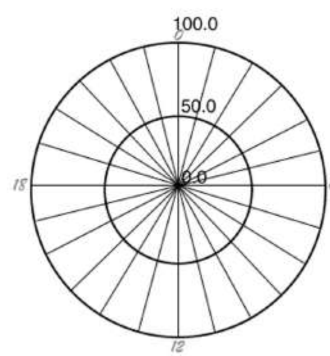

Figure 7.14 *Example of a radar graph with stacking disabled.*

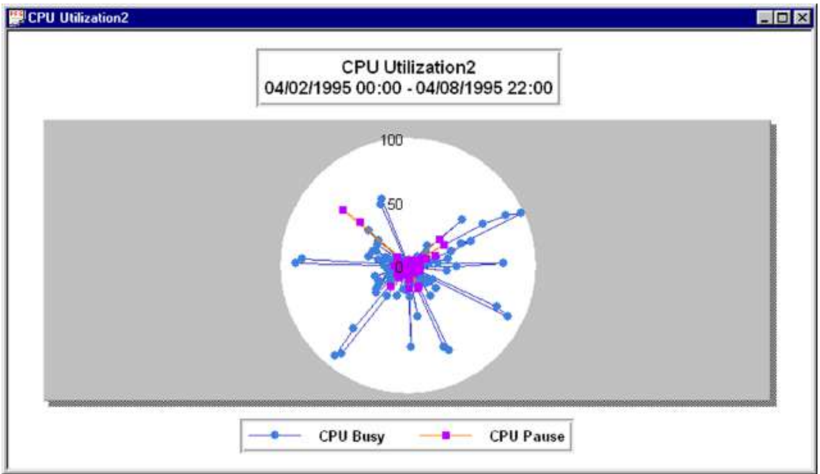

Figure 7.15 *An unstacked radar graph*

Area Radar Graphs

Area radar graphs are used in the same fashion as radar graphs. In some circumstances area radar graphs may be easier to read due to the fact that they are filled. Figure 7.16 shows an example of an area radar graph with stacking enabled.

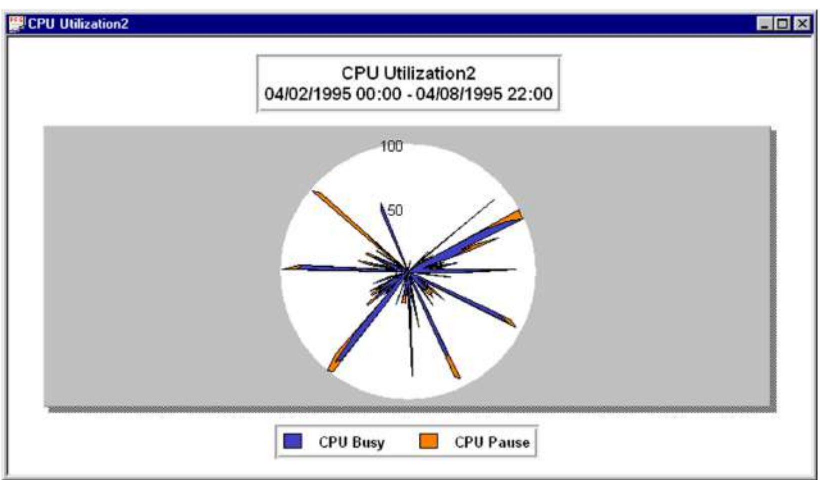

Figure 7.16 *An area radar graph*

Figure 7.17 shows an example of an area radar graph with data stacking disabled.

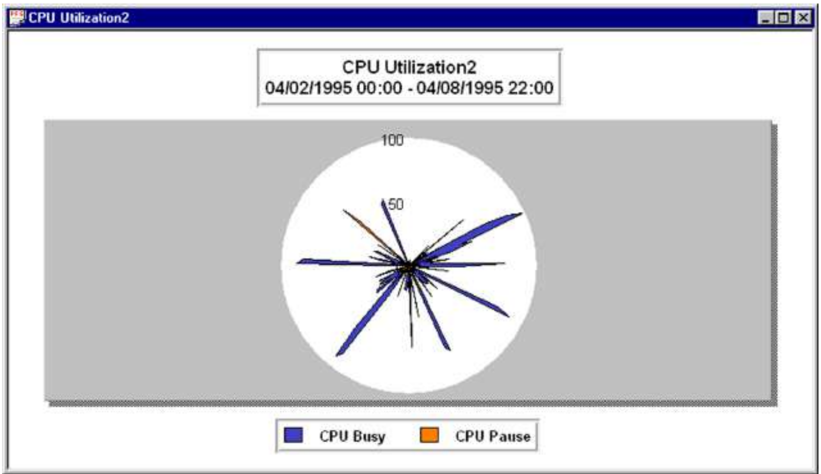

Figure 7.17 *An unstacked radar graph*

3D Graphs

Performance Gallery Gold offers three-dimensional (3D) bar and 3D surface graphing capabilities. Three-dimensional graphs have three axes: x, y, and z:

- The x axis represents the range of the graph. In most cases the range is the time or length of the reporting period divided into logical and equidistant segments (minutes, hours, days).
- The y axis represents the scale (the quantity of the data over the given reporting period).
- The z axis represents the categories or groups of data (processors, drives, applications).

Figure 7.18 identifies each of the three axis.

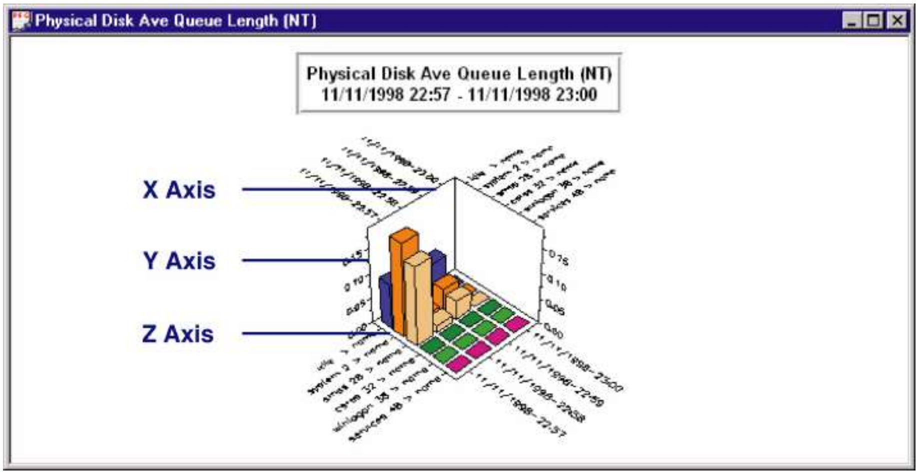

Figure 7.18 *X, Y, and Z axes of a three-dimensional graph*

Pre-configured, three-dimensional graphs can be opened from the **Open Chart** dialog box (see "Open Chart" on page 149).

3D Bar Graphs

3D bar graphs are much the same as a normal bar graph, with one significant exception: it utilizes a third axis. This allows for data to be displayed using three labeled chart axes. Here are some ways that 3D bar graphs may be beneficial:

- Viewing multiple processor loads over multiple days
- Viewing Processor Utilization with multiple jobs
- Monitoring multiple drives over a span of time

A 3D bar graph can be configured to display data as a continuous timeline display (see Figure 7.19) or in a single point (see Figure 7.20).

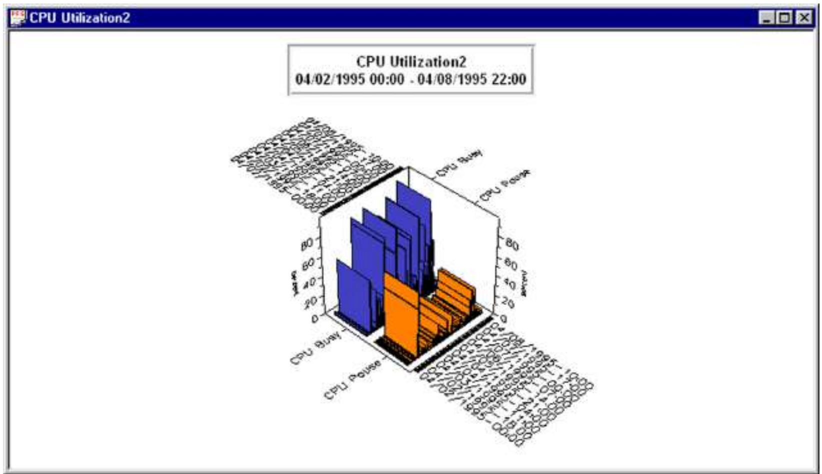

Figure 7.19 *Continuous data displayed in a 3D bar graph*

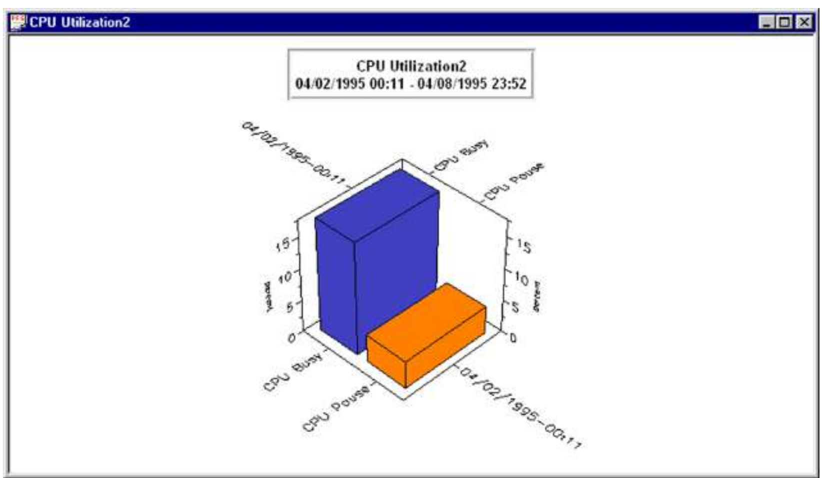

Figure 7.20 *Single point data displayed in a 3D bar graph*

3D Surface Graphs

3D surface graphs with built-in color coding are used to compare varying levels of activity, for example:

- CPU Utilization
- Disk I/O
- Disk Queue Length by Drive

Figure 7.21 shows an example of a 3D surface graph of CPU Utilization.

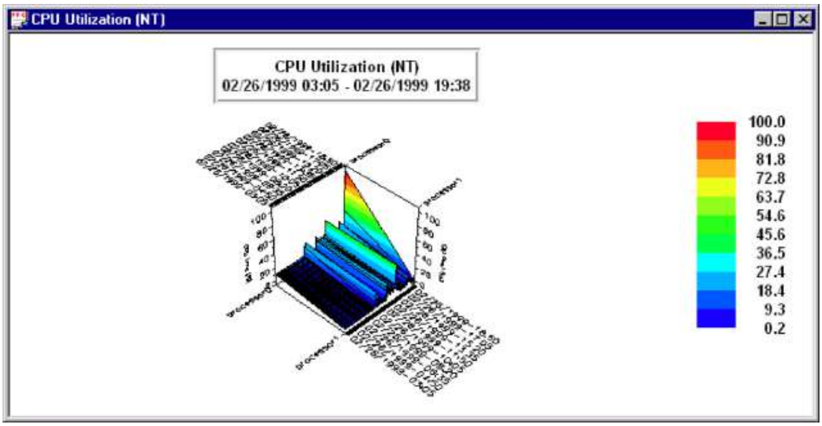

Figure 7.21 *A 3D surface graph*

Scale

A graph's scale (along the y axis) can be tailored to better "fit" its data. There are three scaling options:

- **Percent** scaling allows for a graph to be displayed from 0 to 100 of the original scaling. For example: If the graph being viewed has fairly tight increments (for instance, if one point cannot be discerned from the next) scaling could assist in viewing the data.
- **Automatic** scaling allows for a graph to be displayed in a scale that fits all data in the graph appropriately (for example, the highest point on the graph is the highest point viewable, while the lowest point on the graph is the lowest point viewable).
- **Thresholds** scaling allows for a graph to be displayed in a scale defined by the threshold values located under the **Thresholds** tab (see “Thresholds Properties” on page 135).

Stacking

Data elements in an area, line, or bar graph can be "stacked" vertically, which means that the Performance Gallery Gold program will display all of each data element without overlap. Rather than every data set starting over at zero, data sets are added together, one on top of the other, to equal the total. For instance, In the next figure, Disk Utilization by Drive is displayed unstacked. Each of the five drives is represented as its own data set starting at zero and reaching to its individual disk utilization value, in this case, ranging from 1 to 2.5%. The unstacked graph focuses on the individual data sets (drives) and their values.

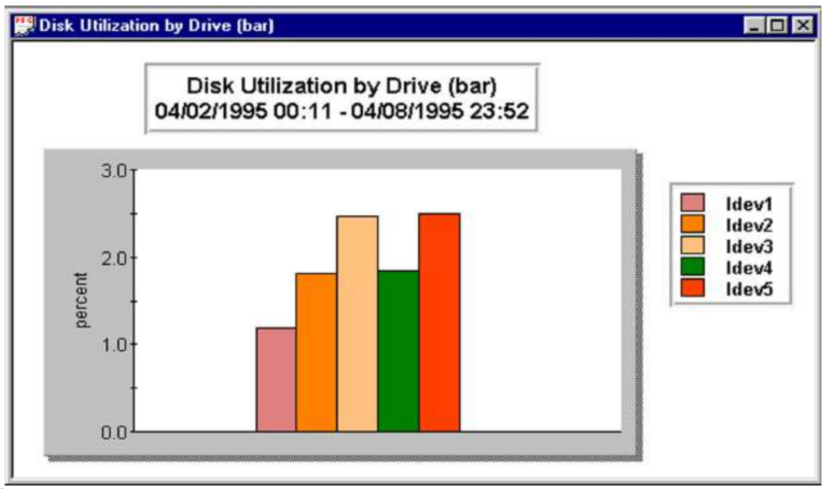

Figure 7.22 *A bar graph displaying unstacked data sets*

Enabling the stacking option using this same data shifts the focus of the graph from each data set value to how they combine to form the "big picture." In the next example (Figure 7.23), each data set is still separated (by color) and an accurate representation of how each drive is performing is still presented. However, this graph shows that these five data sets combined to utilize almost 10% of available disk space.

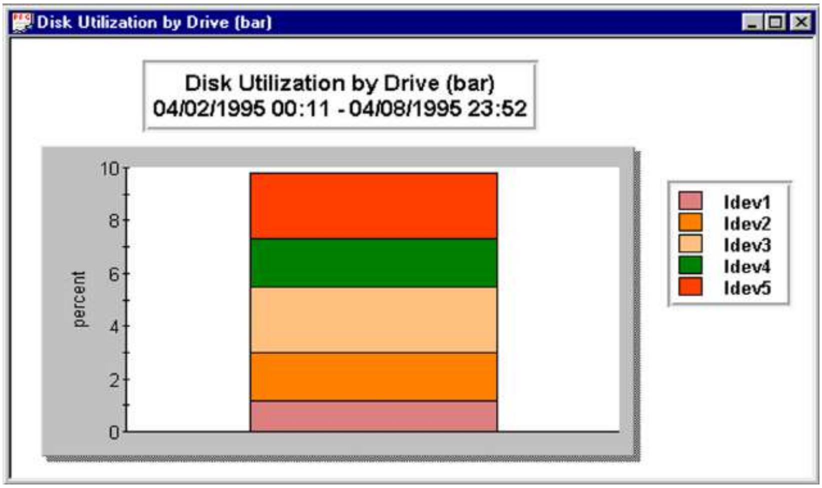

Figure 7.23 *A bar graph displaying stacked data sets*

To activate the stacking option:

1 In the **Template Properties** dialog box, click the **Graph** tab.

2 Near the bottom right corner of the **Graph** tab, select the check box next to **Enable Data Stacking**.

To deactivate this feature, clear the check mark from the check box beside **Enable Data Stacking**.

Marker

Markers are used to identify levels. If a marker is placed at the 10.0 mark, a dotted line is displayed across the graph at the 10.0 mark, helping to identify peaks, valleys, and performance levels accurately. Figure 7.24 shows an example of a marker.

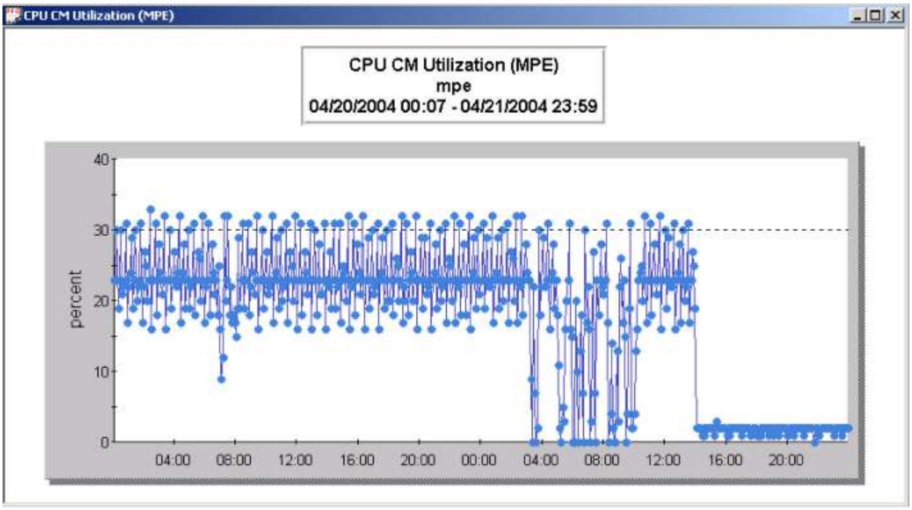

Figure 7.24 *A graph with a marker at 30.0 percent*

To add a horizontal marker to a graph:

1 In the **Template Properties** dialog box, click on the **Graph** tab.

2 Enter a numerical value (+/- n.nnn) in the **Marker at:** text box.

3 Click **Apply**.

To eliminate the horizontal marker, simply clear the **Marker at:** check box.

NOTE A horizontal marker can only be added to a primary graph, not to a secondary graph.

Table Properties

Tables are used to display data in an organized, numerical layout, making some types of comparisons easier. Data in tables can be exported in several different formats, including:

- HTML (*.htm)
- HTML Data Only (*.htm)
- Tabbed Text (*.txt)
- Excel (*.xls)

The properties of a data table template are selected in the **Table** tab in the **Template Properties** dialog box.

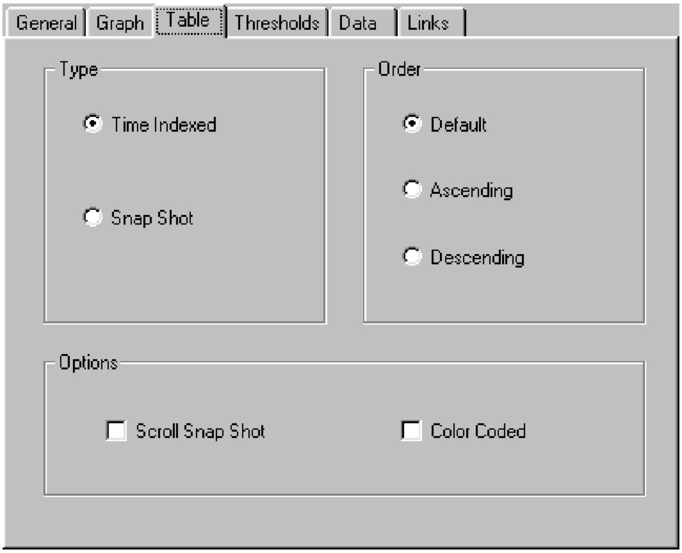

Figure 7.25 *Template Properties dialog box: Table tab*

Type

There are two table types available in Performance Gallery Gold: **Time Indexed**, and **Snap Shot**.

Time Indexed

By default, the range of a Performance Gallery Gold table is indexed over a period of time. The data is placed in columns according to sample dates or times (the x axis on most graphs) and is determined by the overall length of the data collection interval. Each individual sample date or time is provided in the column headings. The **Time Indexed** option allows for data that spans over a period of time to be viewed (see Figure 7.26).

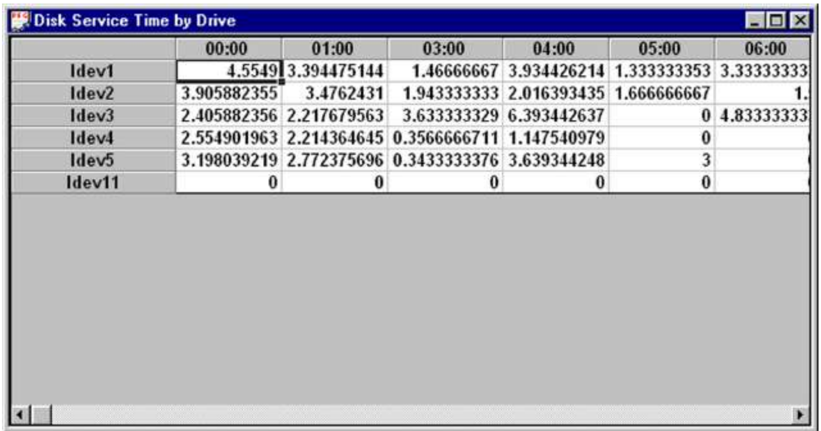

	00:00	01:00	03:00	04:00	05:00	06:00
ldev1	4.5549	3.394475144	1.46666667	3.934426214	1.333333353	3.33333333
ldev2	3.905882355	3.4762431	1.943333333	2.016393435	1.666666667	1.
ldev3	2.405882356	2.217679563	3.633333329	6.393442637	0	4.83333333
ldev4	2.554901963	2.214364645	0.3566666711	1.147540979	0	
ldev5	3.198039219	2.772375696	0.3433333376	3.639344248	3	
ldev11	0	0	0	0	0	

Figure 7.26 *A time-indexed table*

To present chart data in a time-indexed table:

1. In the **Table** tab in the **Template Properties** dialog box, select **Time Indexed** in **Type** options.
2. Click **Apply**. If all modifications are made, click **Close**.

Snap Shot

A table using the **Snap Shot** option allows for data to be viewed as an average of all of the time logged in the current data file:

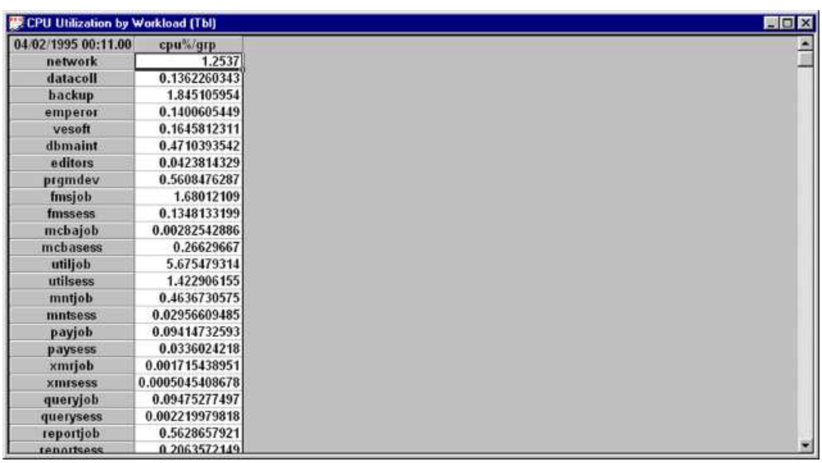

04/02/1995 00:11.00	cpu%/grp
network	1.2537
datacoll	0.1362260343
backup	1.845105954
emperor	0.1400605449
vesoft	0.1645812311
dbmaint	0.4710393542
editors	0.0423814329
prgmdev	0.5608476287
fmsjob	1.68012109
fmssess	0.1348133199
mcbajob	0.00282542886
mcbasess	0.26629667
utiljob	5.675479314
utilsess	1.422906155
mntjob	0.4636730575
mntsess	0.02956609485
payjob	0.09414732593
paysess	0.0336024218
xmrjob	0.001715438951
xmrsess	0.0005045408678
queryjob	0.09475277497
querysess	0.002219979818
reportjob	0.5628657921
reportsess	0.2063572149

Figure 7.27 *A snap shot table*

To show a "snap shot" of a data sample as the average of the data sample at a specific point in time:

1 In the **Table** tab of the **Template Properties** dialog box, select **Snap Shot** in **Type** options.
2 Click **Scroll Snap Shot** (see "Scroll Snap Shot" on page 134) in the **Options** section to add a horizontal scroll bar to the table window.
3 Click **Apply**. If all modifications are complete, click **Close**.

Order

By default, the table data order is determined by the Performance Gallery Gold program. To display the data in either ascending or descending order, enable the **ascending** or **descending** option button accordingly.

Ascending

To view the table data in standard alphabetical order by data line title:

1 In the **Template Properties** dialog box, click the **Table** tab.
2 In the **Order** section of the **Table** tab, select **Ascending**.
3 Click **Apply**, then **Close**.

Descending

To view the table data in reverse alphabetical order by data line title:

1 In the **Template Properties** dialog box, click the **Table** tab.
2 In the **Order** section of the **Table** tab, select **Descending**.
3 Click **Apply**, then **Close**.

Options

Scroll Snap Shot

The **Scroll Snap Shot** option is used to view data one time index at a time. This makes viewing individual times easier, as only data from a single instance in time is viewed at a time.

Color Coded

Normal, warning and critical level thresholds of a table can be designated and assigned color codes from the **Data** tab (see "Warnings" on page 141). To enable this color coding feature for tables, the **Color Coded** option must also be enabled in the **Table** tab.

Thresholds Properties

Upper and lower limits, or *thresholds*, can be applied to a chart template for the following purposes:

- To disqualify any data points and/or entire data sets which fall outside the threshold range, as defined by the user.
- To apply a solid, horizontal reference line to mark the minimum and/or maximum acceptable level or range of a graph (markers cannot be added to a table).

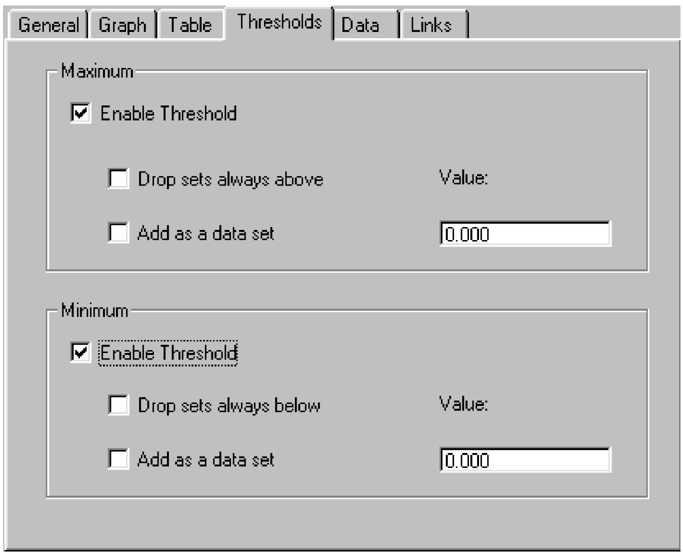

Figure 7.28 *Template Properties dialog box: Thresholds tab*

Maximum and Minimum Thresholds

The **Thresholds** tab is divided into **Maximum** and **Minimum** threshold properties, with **Maximum** options occupying the top half of the **Thresholds** tab. Both thresholds are defined using the following options:

- Enable Threshold
- Drop sets always above/below
- Add as a data set
- Value

Enable threshold

The status of the particular threshold. The default value is disabled. This must be enabled in order for the values entered in the **Maximum** or **Minimum Value:** text boxes to take effect on the tables or graphs.

Drop sets always above/below

To exclude data elements that remain above the maximum threshold value throughout the entire measurement interval:

1. In the **Graph** tab in the **Template Properties** dialog box, set the **Y Axis Scale** to **Automatic** (see "Scale" on page 128).
2. In the **Thresholds** tab, select **Enable Threshold**.
3. Select the **Drop sets always above** check box.
4. Set the maximum threshold **Value**.
5. Click **Apply**.
6. If you are done modifying the open graphs, click **Close**.

To exclude data elements that remain below the minimum threshold value throughout the entire measurement interval:

1. In the **Graph** tab in the **Template Properties** dialog box, set the **Y Axis Scale** to **Automatic** (see "Scale" on page 128).
2. In the **Thresholds** tab, select **Enable Threshold**.
3. Select the **Drop sets always below** check box.
4. Set the minimum threshold **Value**.
5. Click **Apply**.
6. If you are done modifying the open graphs, click **Close**.

Add as a data set (for graphs)

Because only one "official" marker can be added to a chart, using the **Add as a data set** option allows for a second or third reference line to be added, working around the one marker rule. This is only useful for unstacked line graphs.

To add a solid reference line to show the maximum threshold in a graph:

1. In the **Thresholds** tab in the **Template Properties** dialog box, select **Enable Threshold**.
2. Select the **Add as a data set** check box.
3. Set the maximum threshold **Value**.
4. Click **Apply**. If there are no other modifications, click **Close**.

To add a solid reference line to show the minimum threshold in a graph:

1. In the **Thresholds** tab in the **Template Properties** dialog box, select **Enable Threshold**.
2. Select the **Add as a data set** check box.
3. Set the minimum threshold **Value**.
4. Click **Apply**. If there are no other modifications, click **Close**.

Value

To restrict the scale of the active chart template to maximum and/or minimum threshold values:

1. In the **Template Properties** dialog box, click the **Thresholds** tab.
2. In the **Maximum** and/or the **Minimum** options select **Enable Threshold**.
3. Find the **Maximum Value:** and **Minimum Value:** text boxes, and do one or both of the following:
   - Select the **Maximum Value:** text box and replace the default maximum threshold of 0.000 with a new numeric value (+/- n.nnn).
   - Select the **Minimum Value:** text box and replace the default minimum threshold of 0.000 with a new numeric value (+/- n.nnn).

Data Properties

The **Data** properties tab, located on the **Properties** dialog box in the **Template** menu, contains options to do the following:

- Define new data elements for a chart.
- Remove one or more data elements from a chart.
- Modify data elements within a chart.

The computation of each data entry is displayed in the large text box in the lower half of the **Data** tab (see Figure 7.29). Use the UP ARROW and DOWN ARROW keys or click on an entry to select the data entry you wish to modify.

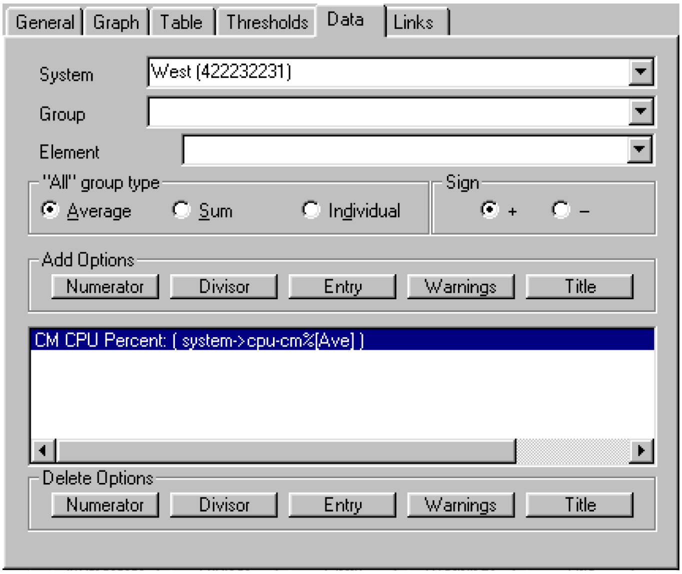

Figure 7.29 *Template Properties dialog box: Data tab*

System

All currently open host data files will be available from this drop-down menu. The host of the currently selected graph or table will appear here by default. If you want to make changes to a different system's template, select a new system from the drop-down list.

Group

A *data group* is an instance or set of instances of which a data element is a member. For example, in the calculation **(nt server (bytes transmitted/sec))**, "*nt server*" is the data group.

To enter a data group for a specific data entry:

1. In the **Data** tab in the **Template Properties** dialog box, either type in a group name in the **Group** text box or select one in the **Group** drop-down menu.
2. Click **Apply** and continue to define chart data, or click **OK** to save the chart template.

Element

A *data element* is an actual data item that is present in one or more of the data groups. For example, **(all processor ((%processor time [Ind]))**, where "*%processor time*" is the data element.

Select an element in the **Element** drop-down menu after the group has been defined, or type a new element within the **Element** text box.

Strings

Strings are also acceptable data elements. For example, **(process->($) stringname)**, where "*($)_*" is the prefix and "*stringname*" represents the name of the string.

After entering a data group, either choose a string from the **Element** drop-down menu or type the characters of the string in the **Element** text box. When entering a string manually, be sure to precede the string name with the dollar character in parentheses ($), followed immediately by a single space. The string prefix is not displayed in the element entry in the computation box at the bottom of the **Data** tab.

NOTE Because a string is a non-numeric value, you will not be able to manipulate it mathematically. For instance, you will not be able to add a numerator or divisor to a string, nor will you be allowed to change it to a positive or negative value.

Placeholders

When an All Group category (see “All group type” on page 139) is entered in the Group box, you can type a number sign (#) in the data **Element** text box as a placeholder for the actual number of elements in the group. For instance, **(workloads->#[Ind])**, where "*#*" will represent the actual number of workloads.

All group type

"All" group refers to a set of group data, such as *all workloads* or *all cache*. This data can be calculated as an average or a sum of a group combined, or each component can be represented individually. For example, in the calculation **(all logicaldisk (avg. disk sec/read [Ind]))**, "*all logicaldisk*" is the group set of all logical disks and "*[Ind]*" shows that the avg. disk sec/read counter for each individual logical disk is included in the chart data.

To select a group set:

1 In the **Data** tab in the **Template Properties** dialog box, either type a group set name in the **Group** text box or select one from the **Group** drop-down menu. Group sets start with "All..." and are listed alphabetically.

2 In the **All Group Type** box, select one of the following:

- **Average** to use the average of all data within the group set in the calculation
- **Sum** to use the sum of all data within the group set in the calculation
- **Individual** to include each individual data component in the calculation (for example, each processor in the group set of all processors)

3 Click **Apply** or **Save As** and continue to define chart data or click **Close** to close the chart template.

Sign

The **Sign** option allows you to change a data element to a positive or negative numeric value. This will not apply if the data element is a string.

The **Sign** option is located in the **Data** tab. Once the data element is defined, select either **Positive** or **Negative** option.

Add Options

The **Add Options** dialog box enables you to create and add a data item to the chart data calculation.

To add a new entry to the currently-selected chart template:

1. In the **Data** tab in the **Template Properties** dialog box, select the appropriate template name from the **Name** drop-down menu, or type in a new name.
2. Type a group name in the **Group** text box, or select one in the **Group** drop-down menu.
3. Select a data element in the **Element** drop-down menu.
4. Click **Entry** in the **Add Options** dialog box. The new data entry is placed at the end of the list of computations in the large text box in the **Data** tab (see Figure 7.29).
5. Make sure the new computation is selected. In the **Add Options** dialog box, click on **Title** and type a data entry title in the **Enter Data Line Title** dialog box.
6. Click **OK**. The new title is added to the beginning of the computation and separated from the computation with a colon (:). This change will also occur in the chart legend.

Numerator

Click the **Numerator** button in the **Add Options** dialog box (**Data** tab, see "Data Properties" on page 137,) to add the positive or negative data group and element to the numerator of the currently-selected data entry. By default, the numerator value is +1.0.

Divisor

In the **Data** tab under **Table Properties**, click on the **Add Options Divisor** button to add the positive or negative data group and element to the divisor of the currently selected data entry. By default, the divisor value is +1.0.

Entry

Clicking on the **Add Options Entry** button adds the new data entry to the end of the list of computations in the text box.

Warnings

Tables

Performance Gallery Gold is able to color code the data entries in table templates in order to help you identify data that exceeds warning and critical level thresholds.

- Any data that exceeds the warning level threshold is color-coded yellow.
- And data exceeding the critical level threshold is color-coded red.
- Any data that is "normal" (not in the warning or critical ranges) will be displayed in black or another user-defined color.

To set the color code thresholds for the currently-selected table template:

1. In the computation text box in the **Data** tab, select the data entry to color code.
2. Click the **Warnings** button in the **Add Options** dialog box.
3. In the **Set Entry Warning Information** dialog box, type the numerical value (+/- nn.nnn) of the warning level threshold (for example, 85.000) and the value of the critical level threshold (for example, 95.000).
   - If the critical and warning level threshold values are equal to one another (critical level threshold = warning level threshold), Performance Gallery Gold will not employ the threshold color codes.
   - If the warning level threshold value is greater than the critical level threshold value, Performance Gallery Gold will assume the values of the data are in descending order and will color code the data appropriately.
   - If the critical level threshold value is greater than the warning level, Performance Gallery Gold will assume the data is in ascending order and will color code the data appropriately.

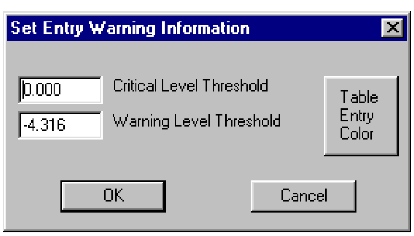

Figure 7.30 *Set Entry Warning Information dialog box*

4. To select a color to represent the data entries within the normal range (which are normally black), click the **Table Entry Color** button in the **Set Entry Warning Information** dialog box. The **Entry Color** dialog box will display.
5. Click any of the **Basic colors** sample boxes.
6. Click **OK**.

Any data that is color-coded in the table will also show up in the **Chart Exceptions** window (see “Exception Window” on page 88).

Graphs

Data entries that exceed the Warning and Critical level settings are placed into the **Chart Exceptions** window for at-a-glance viewing (see “Exception Window” on page 88).

- Any data that exceeds the warning level threshold is denoted in the **Chart Exceptions** window with a standard Windows "warning" icon .
- And data exceeding the critical level threshold is denoted in the **Chart Exceptions** window with a standard Windows "critical" icon .
- Any data that is "normal" (not in the warning or critical ranges) will not be displayed in the **Chart Exceptions** window.

To set warnings for the currently-selected graph template:

1. In the computation text box in the **Data** tab, select the data entry to color code.
2. Click the **Warnings** button in the **Add Options** dialog box.
3. In the **Set Entry Warning Information** dialog box, type the numerical value (+/- nn.nnn) of the warning level threshold and the value of the critical level threshold.
4. Click **OK**.

Title

To give a data entry computation a title:

1. Select the data entry in the computation dialog box.
2. Click on the **Add Options: Title** button.
3. Type the data entry name in the **Title** text box.
4. Click **OK**.

Delete Options

The **Delete Options** dialog box enables you to delete parts of a computation, color code settings, the title of the computation, or the data entry itself. This feature, combined with the **Add Options** feature, allows users to build and modify data entry calculations piece by piece. Select the computation that requires editing, and click on any of the buttons in the **Delete Options** dialog box.

Numerator

Click the **Delete Options Numerator** button in the **Data** tab to clear the active numerator from the chart data computation. When the **Numerator** button is selected, the program actually replaces the numerator with +1.0. Selecting this option immediately removes the numerator from the equation.

Divisor

Click the **Delete Options Divisor** button in the **Data** tab to clear the active divisor from the chart data computation. When the **Divisor** button is selected, the program actually replaces the divisor with +1.0. Selecting this option immediately removes the divisor from the equation.

Entry

To remove the entire equation from the computation dialog box, select the entry in question and click **Delete Options: Entry**. Performance Gallery Gold will prompt you for verification before removing the data entry.

Warnings

Click the **Delete Options Warnings** button to disable the warning and critical level thresholds for the currently-selected data entry. Before the program clears the **Set Entry Warning Information** dialog box, you will be prompted for verification.

Title

To delete the title of any data entry computation, select the computation, and click the **Delete Options Title** button. Click **OK** when the **Delete Title** dialog box appears. The title will be deleted from the computation and from the legend titles.

Links Properties

Performance Gallery Gold is able to store up to five links from one chart to other available charts that use the same data records. This feature enables the user to view a chart, then select relevant charts from a shortcut (right-click) menu.

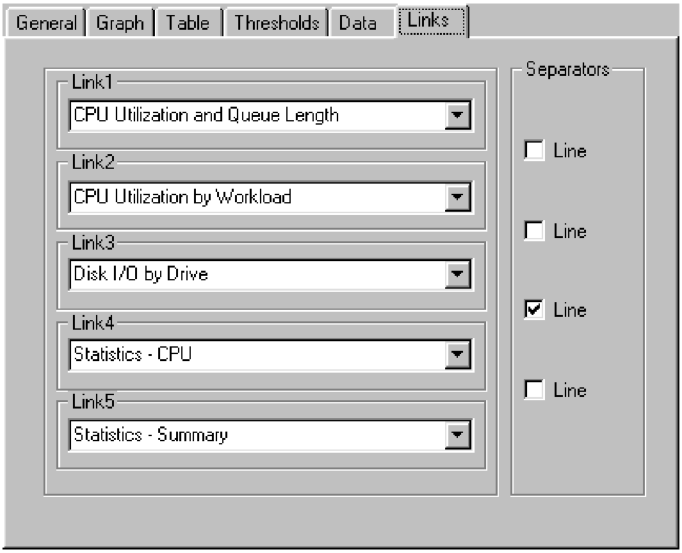

Figure 7.31 *Template Properties dialog box: Links tab*

Links 1-5

To establish links from the currently-selected chart to related charts:

1. In the **Links** tab in the **Properties** dialog box, select a chart in the **Link 1** drop-down menu. The menu lists all charts that use the same data files as the currently-selected chart.
2. Apply step 1 to set up **Link 2**, **Link 3**, **Link 4**, and **Link 5**, as needed.

Links for the currently-selected chart are listed in the shortcut (right-click) menu. To activate a link:

1. Position the cursor over the open chart displayed in the program window.
2. Right-click to open the shortcut menu. The links are always the last menu items in the shortcut menu.
3. Choose a chart link from the shortcut menu. The chart will be loaded and placed in the program window, overlapping the first chart.

To remove a link from the chart template, simply clear the **Link (1-5)** check box in the **Links** tab.

Separators

Separating lines can be established between two consecutive links to aid in their logical grouping within the shortcut menu. Set a separator between two links by selecting the corresponding **Line** check box in the right-hand side of the **Links** tab. Figure 7.31 shows a separator between **Link 3** and **Link 4**.

Secondary Properties

The procedures to modify a secondary graph template are the same procedures used to modify a primary graph template.

1 In the **Template** menu, click **Secondary Properties** to open the **Template Properties** dialog box for the currently-selected secondary graph.

2 Apply the instructions (as described in “Properties” on page 109).

Template Fonts

To change the font, the font style and size, and the special effects of a chart’s title, legend, and index:

1 In the **Template** menu, click **Template Fonts**, then click one of the following:

- Title Font
- Index Font
- Legend Font

2 In the **Font** dialog box, select your choice of font type, style, size, and special effects.

3 Click **OK**.

NOTE If you want the chart to be easily portable, use common font types such as Arial or Times New Roman.

Template Colors

To change the color of a chart's backdrop, graph key, and data field, utilize the **Template Colors** options.

1 In the **Template** menu, click **Template Colors**, then click one of the following:

- Backdrop color
- Graph Key Color
- Data Field Color

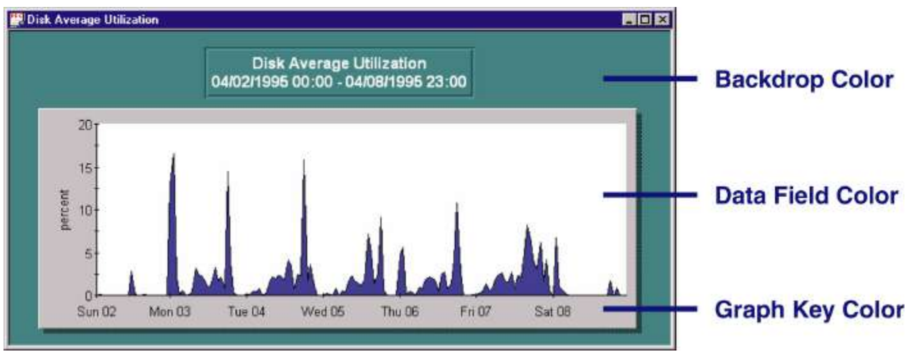

Figure 7.32 *Template backdrop, data field, and graph key colors*

2 In the **Color** dialog box, select one of the 48 color samples.

3 Click **OK**.

NOTE If you choose a dark template color, you may want to adjust the color of the template font(s) to optimize contrast and readability, as in the example shown in Figure 7.32.

Auto Scaling

The **Auto Scaling** option, when selected in the **Template** menu, will automatically increase or decrease the font point sizes to accommodate the size of the chart's window.

To disable the automatic scaling option, clear the check mark from beside the **Auto Scaling** menu item in the **Template** menu. The absence of a check mark indicates the option is disabled.

Save Templates

Select **Save Templates** to save all customized templates to the **tmplates.usr** file in the program file. If you modify a pre-configured chart, you will need to click **Save As** (see “Save As...” on page 111) and give the template a new name, which will be available the next time you open a chart.

NOTE Changes made to pre-configured chart templates cannot be saved by the user. This precautionary measure is in place to preserve the pre-configured charts.

8

Chart Menu and Commands

Overview

The **Chart** menu contains options to open, modify, and close a chart, and several other chart display options. A Performance Gallery Gold "chart" can be either a graph or a table.

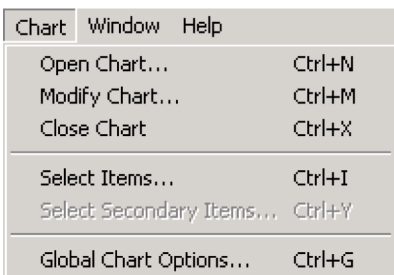

Figure 8.1 *Chart menu*

Open Chart

The **Open Chart...** option in the **Chart** menu allows you to open charts, subcharts, and secondary charts, and to modify the time periods to which these charts are limited.

To open a Performance Gallery Gold chart, click **Open Chart...** in the **Chart** menu or click the Open Chart toolbar button.

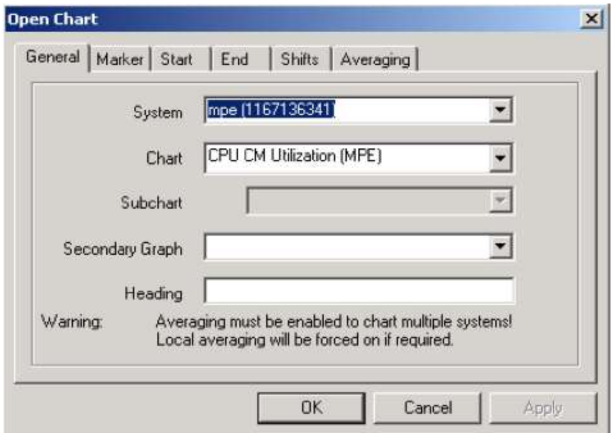

Figure 8.2 *Open Chart dialog box*

The **Open Chart** dialog box displays seven tabs: **General**, **Marker**, **Start**, **End**, **Shifts**, **Averaging**, and **Updating**. Selecting **All Systems** requires data averaging to be on. If this option is selected without enabling averaging (see "Averaging Tab" on page 151), the chart will automatically enable data averaging. The warning at the bottom of this dialog box serves as a reminder.

NOTE The Open Chart menu is identical to the Modify Chart menu, except that the Modify Chart menu affects only charts that are open and currently selected.

General Tab

The **General** tab contains five options: **System**, **Chart**, **Subchart**, **Secondary Graph**, and **Heading**, followed by **Warning**.

- All open host data files will be contained in this list. Select for which host data file you'd like to open a graph.
- To open a chart, select one from the **Chart** drop-down menu.
- If the selected chart supports subcharts, the **Subchart** drop-down list will become active (switch from gray to white), and a subchart can be selected.
- To open two data sets on the same graph with two Y axes, select a secondary graph from the **Secondary Graph** drop-down list.
- To include a heading (other than the names of the charts), type a heading in the **Heading** text box.

Below the **Heading** text box there is a **Warning** containing a note related to the system(s) and the chart that user opened.

Marker Tab

The **Marker** tab contains three options: **Enable Marker**, **Marker Value**, and **Marker Color**.

- To place a marker in your chart, click in the **Enable Marker** check box.
- To set the value of the marker, type a value in the **Marker Value** text box.
- Click the **Marker Color** command button to select a color for the chart marker.

Start Tab

The **Start** tab contains options to modify a chart's start time: **Use Local Start**, **Start Date**, and **Start Time**.

- Clicking in the **Use Local Start** check box enables the user to change the start time and date displayed in the chart.
- Use the **Start Date** standard Windows control to select a start date for the currently-selected chart.
- Use the **Start Time** standard Windows control to select a start time for the currently-selected chart.

End Tab

The **End** tab contains options to modify a chart's end time: **Use Local End**, **End Date**, and **End Time**.

- Clicking in the **Use Local End** check box enables the user to change the end time and date displayed in the chart.
- Use the **End Date** standard Windows control to select an end date for the currently-selected chart.
- Use the **End Time** standard Windows control to select an end time for the currently-selected chart.

Shifts Tab

The **Shifts** tab contains options to exclude data outside of a defined shift: **Enable Shifts**, **Shift Start**, and **Shift End**.

- Clicking in the **Enable Shifts** check box enables the user to define a shift.
- Use the **Shift Start** standard Windows control to select a start time for the shift.
- Use the **Shift End** standard Windows control to select an end time for the shift.

Averaging Tab

The **Averaging** tab contains five options to modify how the data points are displayed: **Local Averaging On**, **Sample**, **Hour**, **Day**, and **Minutes**.

- To enable local averaging, click in the **Local Averaging On** check box.
- Clicking the **Sample** option removes all averaging from the chart.
- Click **Hour** to average all the data points in a sixty-minute timespan to a single data point.
- Click **Day** to average all the data points within a 1440 minute time span into one data point.
- To set the averaging to a different amount of time, type the amount of time (in minutes) in the **Minutes** text box.

General

The **General** tab controls which primary or secondary charts will be opened, which, if any, subcharts to include, and how to label the charts.

Chart

To open a Performance Gallery Gold chart:

1. In the **Chart** menu, click **Open Chart**, click the Open Chart toolbar button, or use the Ctrl+N shortcut keys.
2. In the **System** drop-down menu, select a system or all systems opened.
3. In the **Chart** drop-down menu, select a graph or table.
4. If you do not want to modify this chart at this point, Click **OK**. A new chart opens.

Subchart

Charts that show system performance for members of an All group (see "All group type" on page 139), such as *all discs* or *all workloads*, can be further defined to show a subset of their normal scope, such as a specific disc or a particular workload, respectively,

NOTE Pie graphs and 3D graphs cannot be subcharted or used as subcharts.

To create a subchart:

1. In the **Open Chart** dialog box, select a system **System** drop-down menu and then select a graph or table from the **Chart** drop-down menu. If the **Subchart** window doesn't activate, (change from gray to white,) then the selected chart doesn't support subcharts.
2. In the **Subchart** drop-down menu, select a subchart data item. For example, a subchart of the CPU Utilization graph is the CPU Utilization graph for a specific processor.
3. Click **OK**.

The resulting subchart shows the data pertaining to the subchart data item only. For example, the subchart in Figure 8.3 shows the CPU Utilization for the specified processor, *processor 1*.

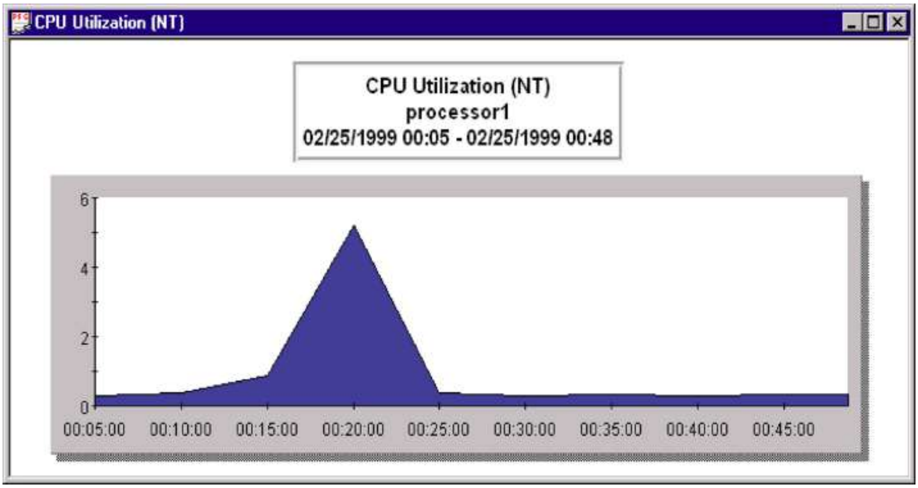

Figure 8.3 *Subchart example*

To undo a subchart:

1 Clear the **Subchart** text box in the **Modify Chart** dialog box.

2 Click **OK**.

Secondary Graph

A secondary graph can be combined with a primary graph to make one graph with dual y axes. In Figure 8.4, the CPU Utilization graph is primary (along the left y axis) and the Copy Read Hits%/ Page Faults graph is secondary (along the right y axis).

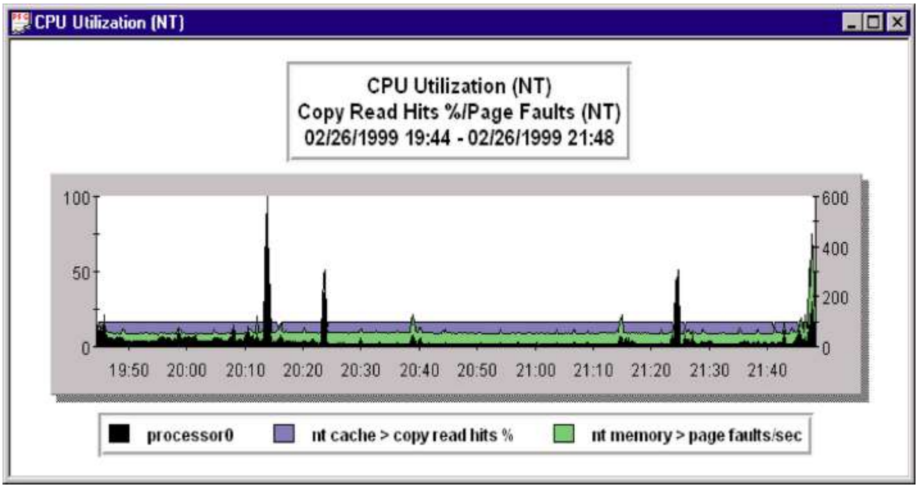

Figure 8.4 *Primary and secondary graphs (dual y axes)*

NOTE Secondary graphs are not applicable to tables, 3D graphs, or pie charts.

To open a graph with dual y axes:

1. In the **Open Chart** or **Modify Chart** dialog box, select a system from the **System** drop-down menu and then select a primary graph that supports secondary charts from the **Chart** drop-down menu.
2. In the **Secondary Graph** drop-down menu, select a secondary graph.
3. A heading to be displayed below the title of the secondary graph is optional. Type the heading in the **Heading** text box, located just below the **Secondary Graph** drop-down menu. This heading would replace the name of the secondary graph template.
4. Click **OK**.

To remove a secondary graph from a primary graph:

1 Clear the **Secondary Graph** box in the **Modify Chart** dialog box.

2 Click **OK**.

Heading

A heading line can be placed below the graph title in the graph title box (see Figure 8.5). In the following example, the heading line is "Your Heading."

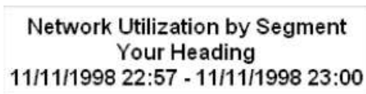

Figure 8.5 *Graph heading example*

To add a single line heading to the title box of the currently-selected graph:

1 In the **Open Chart** dialog box, type the heading (up to 35 characters) in the **Heading** text box.

2 Click **OK**.

To remove the heading from a graph title:

1 Clear the **Heading** text box in the **Modify Chart** dialog box.

2 Click **OK**.

Marker

A **marker** is a horizontal line that marks a specific point on the y axis of a graph. A horizontal marker that is set in the **Open Chart** dialog box will override any other markers included in the graph's template (see "Marker" on page 131). When the marker setting is disabled, or when the graph is closed, the program will revert to the settings in the graph template.

The example in Figure 8.6 shows a horizontal marker at y= 85.

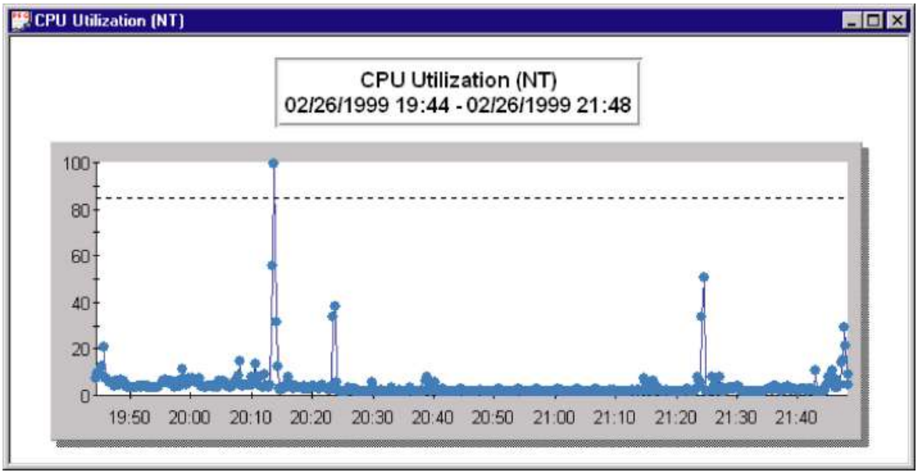

Figure 8.6 *Horizontal Marker (y=85)*

To set a marker in the currently-selected graph:

1. In the **Open Chart** dialog box, click the **Marker** tab. Select the **Enable Local Marker** check box. Enabling a local marker overwrites any existing template marker (see "Marker" on page 131).
2. Enter a numerical value in the **Marker Value** text box. The default marker setting is y = 0.
3. Click **OK**.

By default, the marker is black. To change the marker color:

1. In the **Marker** tab in the **Open Chart** or **Modify Chart** dialog box, verify that the **Enable Local Marker** check box is selected and a numerical value is entered.
2. Click **Marker Color**.
3. In the **Marker Line Color** dialog box, select a color sample.
4. Click **OK** in the **Marker Line Color** dialog box.
5. Click **OK** in the **Modify/Open Chart** dialog box.

To remove a horizontal marker:

1 In the **Modify Chart** dialog box, select the graph from the **Chart** drop-down menu (the currently selected chart should automatically be selected in the drop-down menu).
2 Click the **Marker** tab.
3 Clear the **Enable Local Marker** check box.
4 Click **OK**.

Reporting Periods

With Performance Gallery Gold, you can adjust the length of the reporting period or impose data collection "shifts" to find system performance trends. For example:

- To determine what activities lead to a CPU bottleneck late last Thursday afternoon, you could examine the data collected in the 24-hour period from 5:00 P.M. on Wednesday through 5:00 P.M. on Thursday by modifying the start date and time and the end date and time of the reporting period.
- To monitor system activity that occurs during regular business hours only, you can enable a shift setting from 7:00 A.M. through 6:00 P.M. for example, and exclude data collected during off hours.

Start

The **Start Date** and **Start Time** properties of a chart specify when the reporting period begins (when x=0). This function will not work for radar, 3D, or pie graphs.

- By default, the **Start Date** is always the earliest (oldest) performance record in the active data file.
- By default, the **Start Time** is always the time of the earliest (oldest) performance record in the active data file.

To change the **Start Date** and/or **Start Time** properties:

1 In the **Open Chart** or **Modify Chart** dialog box, click the **Start** tab.
2 Click the **Use Local Start** check box. A check mark should appear in the check box.
3 Adjust the **Start Date/Start Time** properties by using the **Start Date/Start Time** standard Windows controls to adjust the date/time.
4 Click **OK**.

To return to the default **Start Date** and **Start Time**, clear the **Use Local Start** check box and click **OK**.

End

The **End Date** and **End Time** properties of a chart specify when the reporting period begins (when x=0). This function will not work for radar, 3D, or pie graphs.

- By default, the **End Date** is always the latest (most recent) performance record in the active data file.
- By default, the **End Time** is always the time of the latest (most recent) performance record in the active data file.

To change the **End Date** and/or **End Time** properties:

1. In the **Open Chart** or **Modify Chart** dialog box, click the **End** tab.
2. Click the **Use Local End** check box. A check mark should appear in the check box.
3. Adjust the **End Date/End Time** properties by using the **End Date/Start Time** standard Windows controls to adjust the date/time.
4. Click **OK**.

To return to the default **End Date** and **End Time**, clear the **Use Local End** check box and click **OK**.

Shifts

The **Shift Start** and **Shift End** settings specify the beginning and end of an imposed shift. This program feature restricts the chart data to a specific range of time (a shift) within the data sample, which can, as a result, exclude unwanted data from the chart, such as data collected during the late-night hours.

NOTE The Performance Gallery Gold program executes shift filters according to an order of operations. The program will execute a shift established in Global Chart Options first (see "Global Chart Options" on page 161), and then execute a shift defined in the Open Chart or Modify Chart dialog box.

To define a shift:

1. In the **Open Chart** or **Modify Chart** dialog box, click the **Shift** tab.
2. Select the **Enable Shifts** box.
3. Adjust the **Shift Start** and **Shift End** properties by using the **Shift Start/Shift End** standard Windows controls to adjust the properties.
4. Click **OK**.

To disable the shift settings, clear the **Enable Shifts** check box and click **OK**.

Averaging

Data collected by the host data collector (see "Welcome to Performance Gallery Gold" on page 1) is averaged based on user-defined configurations. The **Averaging** tab allows users to change the averaging time period, or to turn off the averaging feature altogether.

To modify a chart's averaging settings:

1 In the **Chart** menu, click **Open Chart** for a new chart, or **Modify Chart** to change an open chart.

2 Click the **Averaging** tab.

3 Click in the **Local Averaging On** check box.

4 Select one of the following time settings:

- **Sample**

  Clicking **Sample** turns off all averaging for the currently-selected chart. The **Minutes** text box will read "0."

- **Hour**

  Clicking **Hour** causes the data to be averaged in hourly increments. The **Minutes** text box will read "60."

- **Day**

  Clicking Hour causes the data to be averaged on a daily basis. The **Minutes** text box will read "1440."

- **Minutes**

  If you want to see the data averaged during different time settings, enter a numerical value (in minutes) in the **Minutes** text box. For example, if you want all data averaged every two hours, enter 120 minutes in the **Minutes** text box. If you just want to see the average over two days of data collecting, type 2880 minutes in the **Minutes** text box.

Modify Chart

Options available to modify a chart are the same as the options in the **Open Chart** dialog box (see "Open Chart" on page 149). Use the **Modify Chart** menu if you want to make changes to a chart that is already open.

Close Chart

To close the active chart, select **Close Chart** in the **Chart** menu.

- As an alternative, click the Close Chart toolbar button .
- Or use the Ctrl+X shortcut keys.

Select Items

The **Select Items** feature of Performance Gallery Gold specifies which data items to include in a particular chart for as long as that chart remains active. This option is applicable to 2-D graphs and tables only.

To select specific data items for an active 2-D graph or table:

1 In the **Chart** menu, select **Select Items**.

The **Select Items** dialog box displays a list of all data items available for the graph (see Figure 8.7).

- The selected (shaded) data items in the list box are included.
- The non-selected data items are currently excluded.
- Clicking on any one data item deselects all of the other data items.

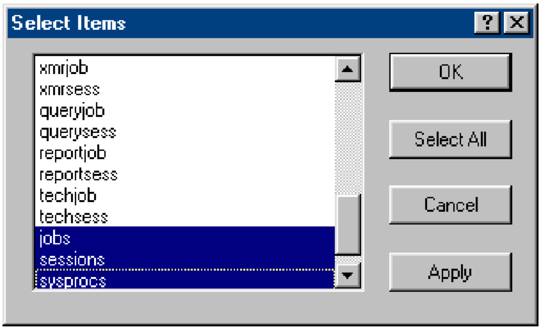

Figure 8.7 *Select Items dialog box*

2 To include a non-selected data item in the active graph, select that data item and click **Apply**. To include more than one data item, press and hold the **Ctrl** key and single-click on each entry you want to be included in the chart. Click **Apply** when all the necessary data items are selected.

Select Secondary Items

The **Select Secondary Items** feature enables the user to specify which data items to include in a particular secondary graph for as long as that graph remains open. This feature applies to two-dimensional graphs only.

To select specific data items for the currently-active secondary graph:

1 In the **Chart** menu, select **Select Secondary Items**.

The **Select Items** dialog box displays a list of all data items available for the secondary graph.

- The selected (shaded) data items in the list box are included in the chart.
- The non-selected data items are currently excluded.
- Clicking on any one data item deselects all of the other data items.

2 To include a non-selected data item in the active graph, select that data item and click **Apply**. To include more than one data item, press and hold the Ctrl key and single-click on each entry you want to be included in the chart. Click **Apply** when all the necessary data items are selected.

Global Chart Options

The global chart options apply to all charts for the duration of the Performance Gallery Gold session, or until the global chart properties are manually reset, even if the data file is closed or replaced. The settings selected in the Global Chart Options are saved in the system's Registry, so the settings in Global Chart Options will be saved until the next time the user logs in. However, if a desktop file is opened after a global property is set, the program will use the default chart properties saved in the desktop file, even after the desktop file is closed.

To open the **Global Chart Options** dialog box, click **Global Chart Options** in the **Chart** menu.

- As an alternative, click the Global Chart Options toolbar button .
- Or use the Ctrl+G shortcut keys.

The **Global Chart Options** dialog box displays three tabs: **Timeline Options**, **Global Date/Time Select** and **Shift Settings** (Figure 8.8).

The global options available in each of these tabs are described and illustrated in the following sections of this manual.

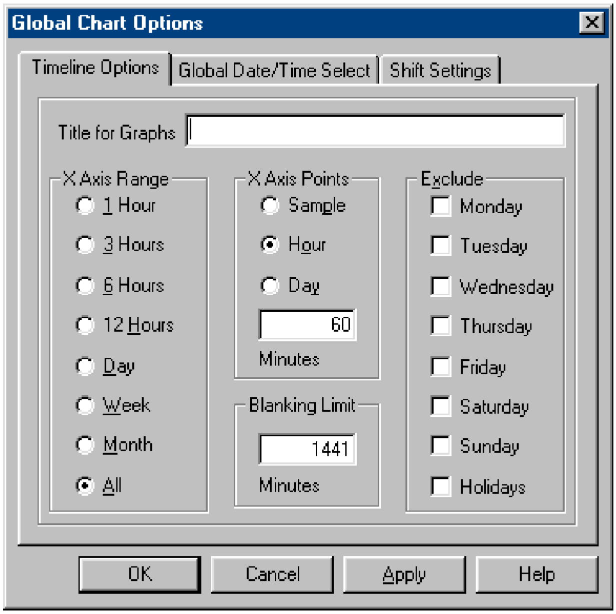

Figure 8.8 *Global Chart Options dialog box*

Timeline Options

The global timeline options affect the following graph and table properties:

- The graph title (applies to graphs only).
- The range of the x axis (applies to graphs only).
- The number of points on the x axis (applies to both graphs and tables).
- The maximum blanking limit (applies to graphs only).
- The days of the reporting period to be included or excluded (applies to both graphs and tables).

Set global timeline options in the **Timeline Options** tab in the **Global Chart Options** dialog box.

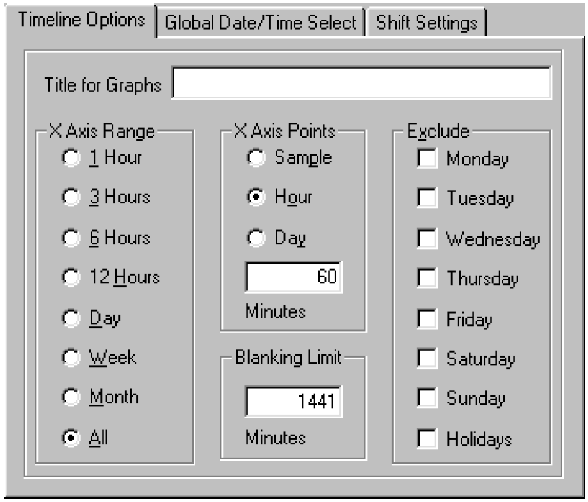

Figure 8.9 *Global Chart Options dialog box: Timeline Options tab*

Title for Graphs

The graph title is placed in the first line of the graph title box. The following example shows the graph title, "Lund Performance Solutions."

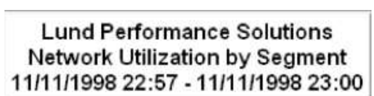

Figure 8.10 *Global graph title*

To add a global title to the top line of the graph title box:

1 In the **Timeline Options** tab in the **Global Chart Options** dialog box, type a title (up to 50 characters) in the **Title for Graphs** text box.

2 Click **Apply**.

3 Click **OK**.

To remove a graph title, clear the **Title for Graphs** text box and click **OK**.

X Axis Range

By default, the global x axis range of a graph or table property is **All**, meaning the chart will show *all* of the data sample in the x axis range.

To select a specific global x axis range:

1 In the **Timeline Options** tab in the **Global Chart Options** dialog box, select an **X Axis Range** property.

For example:

- Select **1 Hour** to display one hour of data.
- Select **Day** to display twenty-four consecutive hours of data.
- Select **Week** to display seven consecutive days of data.
- Select **All** to display all data collected in the measurement interval.

2 Click **OK**.

NOTE If you collect data once per hour, only one data point will be included in each graph. For area graphs or line graphs with data points turned off, this will result in a blank graph, as a graph needs two data points to make a line or an area.

X Axis Points

The **X Axis Points** property specifies what each data point on the x axis represents, for instance:

- A sample
- An hour
- A day

NOTE The X Axis Points property must coincide with the range specified in X Axis Range (see “X Axis Range” on page 163). For example, if the X Axis Range is only 12 Hours, the X Axis Points should not be set to Day, because, at that setting, the data point would be an average of the data collected in one day (24 hours). In other words, the length of the reporting period (the x axis range) must be greater than the X Axis Points setting, because each graph must contain at least two data points.

To define the points of the x axis range:

1 In the **Timeline Options** tab in the **Global Chart Options** dialog box, click on either **Sample**, **Hour**, **Day,** or enter a value in the **Minutes** text box.

Table 8.1 *X axis points options*

Option	Description
Sample	Select **Sample** to display each data point on the x axis as an average of the data collected in one collection cycle (determined in the parameters of the host data collector).
Hour	Select **Hour** to display each data point as the average of the data collected in one hour.
Day	Select **Day** to display each data point as the average of the data collected in one day (24 consecutive hours),
Minutes	Select **Minutes** to display each data point as the average of the data collected in a time period other than Sample, Hour, or Day. For instance, 120 minutes if you want the data to be averaged over a two hour time period.

2 Click **OK**.

NOTE The X Axis Points property is global. However, when the X Axis Points property is set for a specific data file and that file is closed while a second data file is opened, the program will adjust these properties, if needed, to fit the new data sample.

Blanking Limit

The *blanking limit* is the maximum number of minutes between data input allowed by the program. If no data is recorded within the set period of consecutive minutes (the blanking limit), the program will show a break in the graph where the absence of data occurs.

To set the blanking limit value:

1 In the **Timeline Options** tab in the **Global Chart Options** dialog box, type a blanking limit value (1-9,999,999) in the **Blanking Limit** text box.

2 Click **OK**.

Exclude

System performance data can sometimes be more useful when the number of days in the reporting period is limited, for example, to only working days (Monday through Friday). The

Exclude feature will disqualify data pertaining to selected weekdays and/or holidays from the data sample.

Holidays are pre-defined in the **C:\...\Performance Gallery\Holidays.dat** file. Pre-defined holidays include:

- New Year's Day
- Independence Day
- Thanksgiving Day
- Memorial Day
- Labor Day
- Christmas Day

NOTE The Holidays.dat file can be modified. For instructions to modify the file, open the MS Word document in C:\...\Performance GalleryGold\Holidays.doc.

To exclude specific days from graphs and tables globally:

1 From the **Timeline Options** tab in the **Global Chart Options** dialog box, select one or more days from the **Exclude** group.

2 Click **OK**.

The graph shown in Figure 8.11, for example, excludes the data collected on Saturdays, Sundays and holidays.

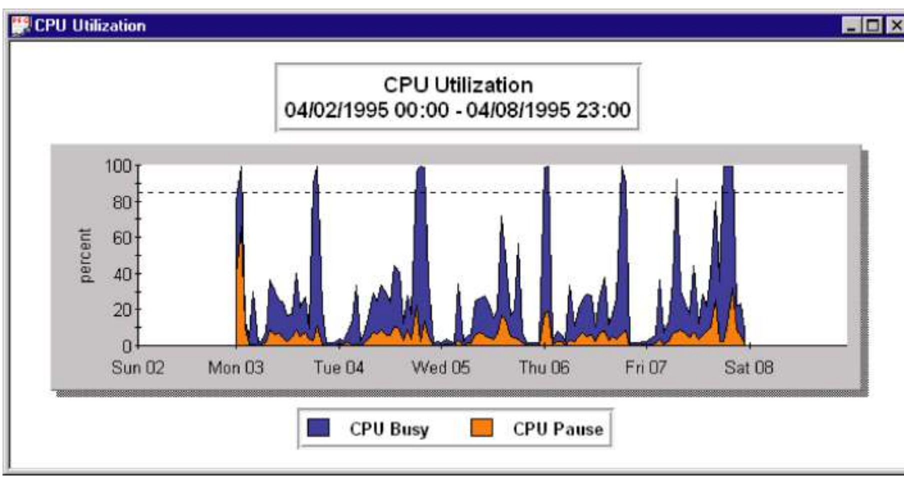

Figure 8.11 *Graph demonstrating excluded data (Saturdays, Sundays, and holidays excluded)*

Global Date/Time Select

By default, Performance Gallery Gold uses the entire range of performance data collected in the open data file.

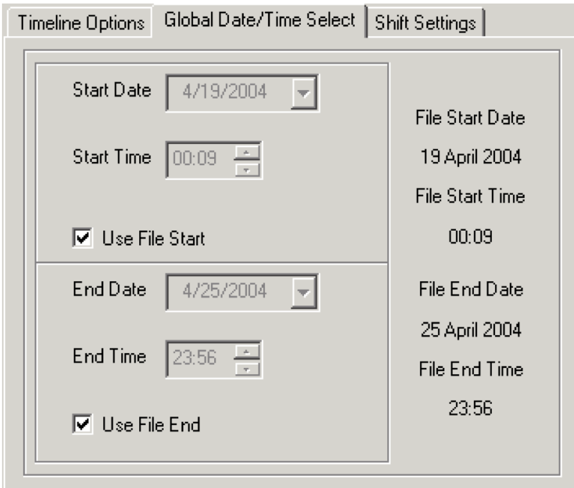

Figure 8.12 *Global Chart Options dialog box: Global Date/Time Select tab*

Start Date/Time

The **Global Date/Time Select** properties enable the user to further define the range of the data file by specifying an exact **Start Date** and **Start Time** of the reporting period.

To define the **Start Date** and **Start Time** of the active data file:

1. Click in the **Use File Start** check box to remove the check mark and disable the **Use File Start** option.
2. In the **Global Date/Time Select** tab in the **Global Chart Options** dialog box, adjust the **File Start Date**:
   - Use the **Start Date** standard Windows control.
   - Use the **Start Date** spin box.
3. Adjust the **File Start Time** in the same manner.
4. Click **OK**.

Use File Start

Use File Start must be disabled before the start date or start time can be adjusted. To use a start date/time other than the default file start date/time, clear the check mark from beside **Use File Start**.

End Date/Time

The **Global Date/Time Select** properties enable the user to further define the range of the data file by specifying an exact **End Date** and **End Time** of the reporting period.

To define the **End Date** and **End Time** of the active data file:

1. Click in the **Use File Start** check box to remove the check mark and disable the **Use File Start** option.
2. In the **Global Date/Time Select** tab in the **Global Chart Options** dialog box, adjust the **File End Date**:
   - Use the **End Date** standard Windows control.
   - Use the **End Date** spin box.
3. Adjust the **File End Time** in the same manner.
4. Click **OK**.

Use File End

Use File End must be disabled before the end date or end time can be adjusted. To use a end date/time other than the default file end date/time, clear the check mark from beside **Use File End**.

Shift Settings

Restricting the hours of a data sample (the active data file) to specific shifts can sometimes reveal short-term spikes or dips in a chart that may otherwise be undetected. Applying this global chart option may also simplify cluttered graphs and make them more readable.

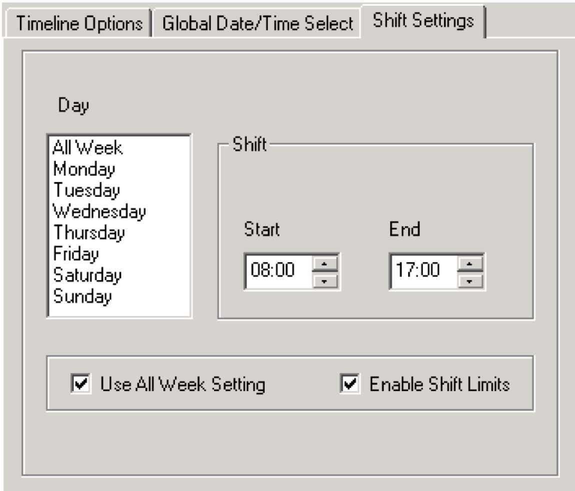

Figure 8.13 *Global Chart Options dialog box: Shift Settings tab*

Shift Day

The **Shift Day** list box contains all the days of the week and an option called **All Week**. Click on one of these choices to define a shift for that day, (or for the week if **All Week** is selected).

Shift Start

To define the start time of a shift, use the **Shift Start** standard Windows control.

Shift End

To define the end time of a shift, use the **Shift End** standard Windows control.

NOTE While the Shift Start and Shift End can be adjusted, the changes will not take effect on the charts unless the Enable Shift Limits check box is enabled.

Use All Week Setting

To set the **Use All Week** shift as the global default shift:

1. In the **Shift Settings** tab, select the **Enable Shift Limits** check box, if it is not already selected.
2. Select **All Week** in the **Shift Day** group box.
3. Set the **Shift Start** and **Shift End** times.
4. Select the **Use All Week Setting** check box.
5. Click **OK**.

Enable Shift Limits

To specify the hours of an imposed shift setting:

1. In the **Shift Settings** tab in the **Global Chart Options** dialog box (see Figure 8.13), select the **Enable Shift Limits** check box.
2. Select one of the options listed in the **Shift Day** group box, either a specific day of the week (**Monday**,..., **Sunday**) or **All Week**.
3. Use the **Shift Start** standard Windows control to define the beginning of the global shift (hh:mm). By default, the shift will start at 00:00 (midnight).
4. Use the **Shift End** standard Windows control to define the end of the shift (hh:mm). By default, the shift will end at 23:59.
5. Click **Apply**.
6. Apply steps 3 through 6 to set additional shifts, as needed. If you want to enable different shifts for each day, make sure the **Use All Week** option in disabled.
7. Click **OK**.

To disable shift limits, clear the **Enable Shift Limits** check box in the **Shift Settings** tab and click **OK**.

Hover Stats

The user has the ability to move the mouse over a point on a graph chart and see what the numbers are behind the data point.

1. Click **Open Chart...** in the **Chart** menu and choose a graph chart.
2. In the toolbar, click on the Show/Hide hover stats button  to enable the hover stats.

   Hovers stats use a lighter yellow color as background.
3. Move the mouse on the chart and find the numbers that are behind the data points.

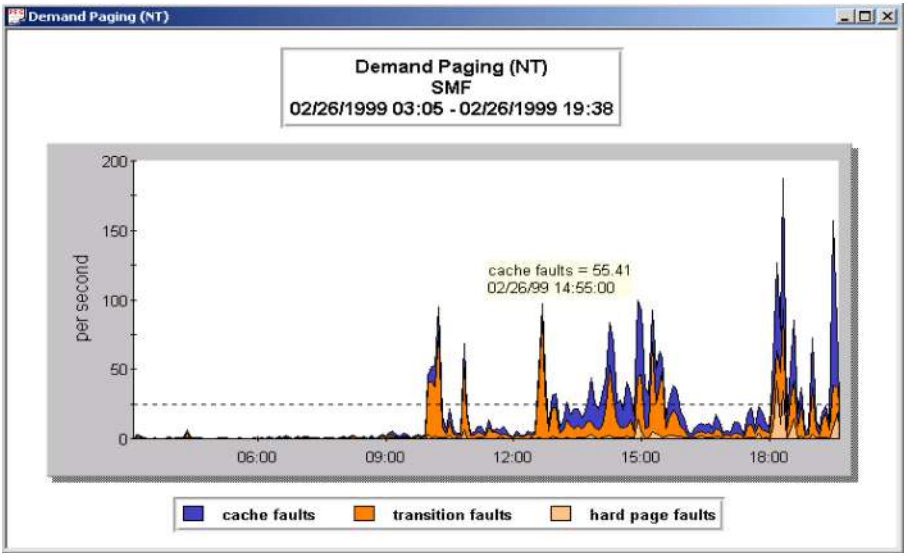

Figure 8.14 *Example of a chart displaying hover stats*

9

MPE/iX System Performance

Processor Performance

A Queue and B Queue Percentages

Table 9.1 *A Queue and B Queue Percentages Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	**Problematic**	**Unacceptable**
A and B Queue %	less than 5	5 to 8	greater than 8

The *A queue* is usually reserved for high-priority system processes. The *B queue* is reserved for lower-priority system processes and very high-priority user processes.

This information is reported in the Performance Gallery Gold **CPU Utilization by Subqueue (MPE)** stacked area graph. In the example shown in Figure 9.1, the data items have been restricted to show just the A and B queue data.

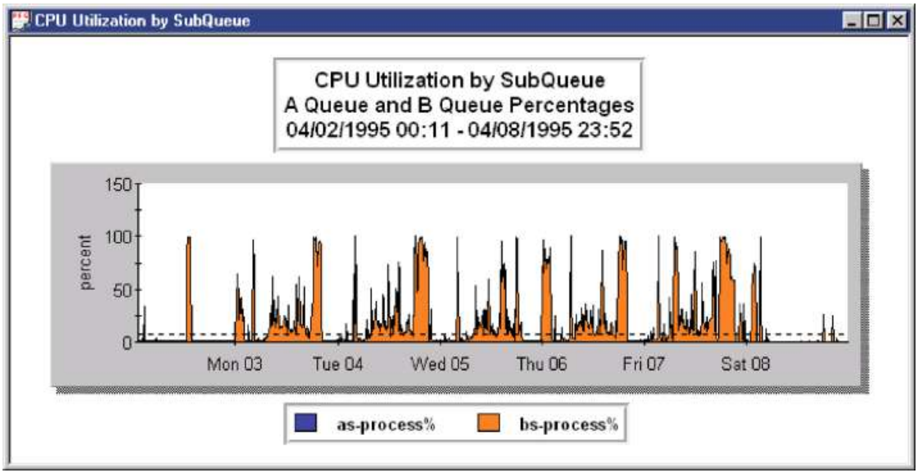

Figure 9.1 *Example CPU Utilization by Subqueue graph: A and B Queue percentages*

CPU Busy Percentage

Table 9.2 *CPU Busy Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Busy %	less than 50	50 to 85	greater than 85

The *CPU Busy %* pulse point shows the percentage of time the CPU spent executing user and system processes, interrupts, overhead, and context switches, instead of being in a pause or idle state.

This information is reported in the Performance Gallery Gold **CPU Utilization (MPE)** stacked area graph (see the data element titled "**CPU Busy**").

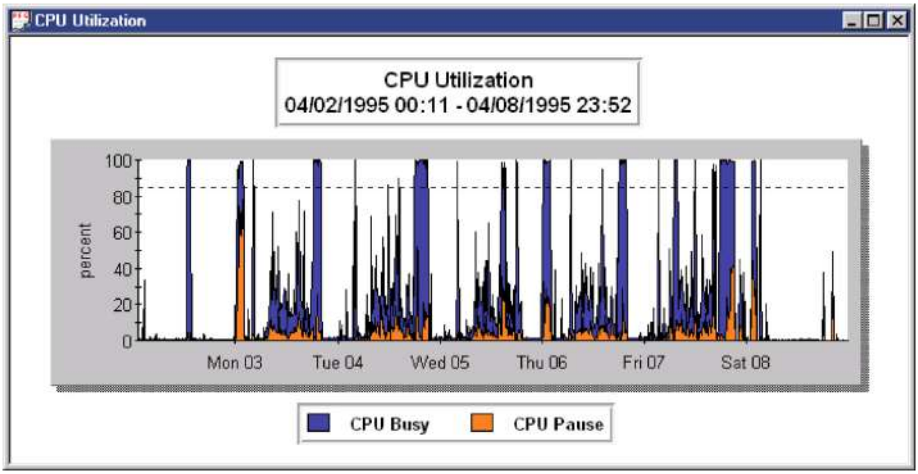

Figure 9.2 *Example CPU Utilization (MPE) stacked area graph: CPU busy*

CPU Compatibility Mode Percentage

Table 9.3 *CPU Compatibility Mode Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Compatibility Mode %	less than 10	10 to 50	greater than 50

The *CPU Compatibility Mode %* pulse point shows the percentage of time the CPU spent servicing compatibility mode code.

CPU Compatibility is one point of resource drain due to mode switching activity. In order to accommodate non-native mode program execution, MPE/iX must translate code into its "native tongue." This costs the CPU time. There are three metrics you will want to monitor:

- The percentage of all CPU time spent in compatibility mode. If this number is high, you should reduce CPU utilization by compiling your busiest programs into native mode.
- CM to NM switch rate (the number of compatibility mode to native mode switches per second). This rate can be relatively high, but it still costs the CPU time.
- NM to CM switch rate (the number of native mode to compatibility mode switches per second). At the global level, keep the number of NM to CM switches to a minimum; they are more costly.

This information is reported in the Performance Gallery Gold **CPU CM Utilization (MPE)** line graph.

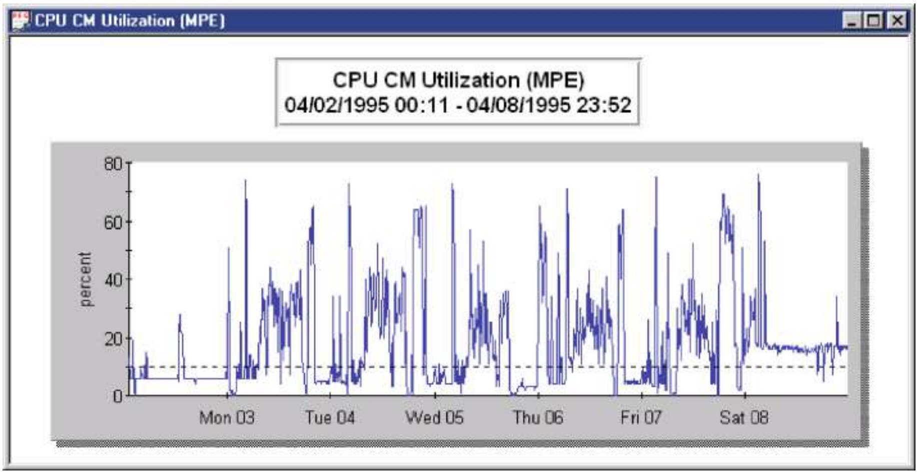

Figure 9.3 *Example CPU Compatibility Mode Utilization (MPE) line graph*

CPU Queue Length

Table 9.4 *CPU Queue Length Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Queue Length	less than 5	5 to 15	greater than 15

The *CPU Queue Length* pulse point shows the average number of processes in the CPU request queue.

The concept of queue length can be compared, for example, to customers standing in a bank teller line.

- If there are no customers waiting when you arrive, then the queue length is zero and you will be served right away.
- If every time you arrive, there are customers in line ahead of you, you will have to wait before you are helped.

It is the same with disk I/O requests, they are served one at a time. The “cost” of disk I/O requests can be either “cheap” or “expensive.” Disk I/O’s that require more work of the disk device are worse for performance.

This information is reported in the Performance Gallery Gold **CPU Utilization and Queue Length (MPE)** line graph (see the data element titled **"Queue Length"**).

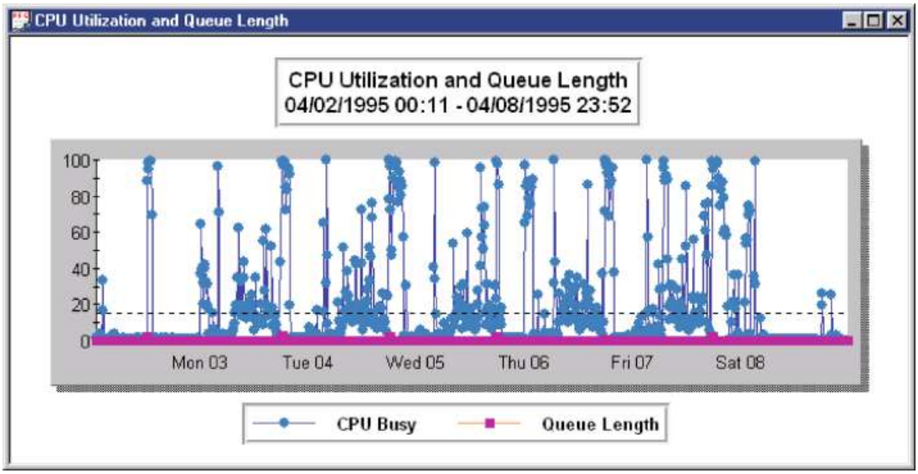

Figure 9.4 *Example CPU Utilization and Queue Length (MPE) line graph: Queue Length*

A queue length value exceeding 1.0 or greater is not permissible. (Please keep in mind that specific busy situations and unique rush hours can cause queue length to temporarily spike above 1.0.) Watch carefully for any sign of a trend.

If you have consistently high queue lengths, you should answer the following questions:

- Is the disk arm movement excessive due to heavily hit files?

  The disk drive's arm mechanism moves across the disk. If the mechanism moves widely over the device's platter to jump from one data set to another, the consequence for those particular I/O's is longer queues.

- Do files that depend on one another reside on the same disk drive?
- Are your files balanced?
- Could there be an application design problem?
- Is the current disk drive model too slow for your particular application?

High Priority Percentage

Table 9.5 *High Priority Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
High Priority %	less than 50	50 to 85	greater than 85

The *High Priority %* pulse point shows the percentage of time the CPU spent executing high-priority system and user processes, interrupts, and overhead.

This information is reported in the Performance Gallery Gold **CPU Utilization by Subqueue (MPE)** stacked area graph, modified to show only the high-priority system and user processes, interrupts, and overhead data elements.

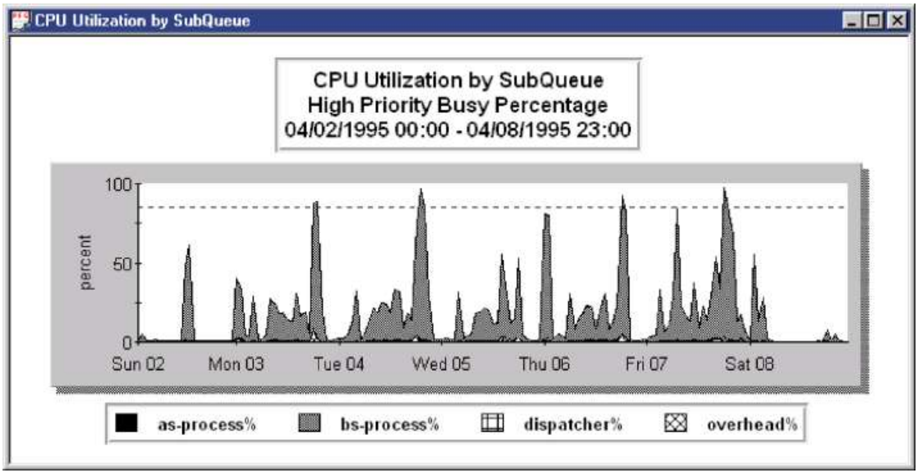

Figure 9.5 *CPU Utilization by Subqueue (MPE) stacked area graph: Hi-Pri Busy %*

ICS/OH+Dispatcher Busy Percentage

Table 9.6 *ICS/OH+Dispatcher Busy Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
ICS/OH+Dispatcher %	less than 10	10 to 15	greater than 15

The *ICS/OH+Dispatcher %* pulse point shows the percentage of time the CPU spent on overhead and interrupts.

This information is reported in the Performance Gallery Gold **CPU Utilization by Subqueue (MPE)** stacked area graph, modified to show only the overhead and dispatcher data elements.

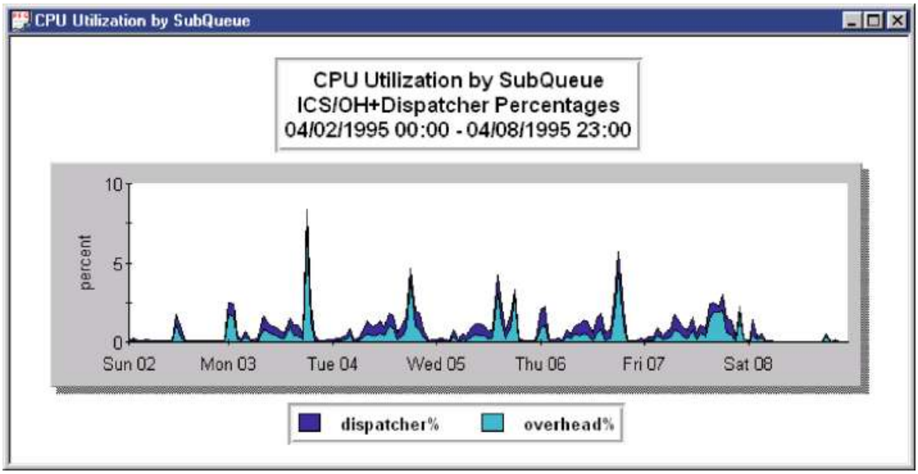

Figure 9.6 *Example CPU Utilization by Subqueue graph: ICS/OH+Dispatcher Busy%*

Memory Performance

CPU Memory Manager Percentage

Table 9.7 *CPU Memory Manager Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Memory Manager %	less than 4	4 to 8	greater than 8

The *CPU Memory Manager %* pulse point shows the percentage of time the CPU spent cycling through main memory, probing for space to satisfy requests. A large number of cycles will increase the percentage of CPU used to service memory needs.

This information is reported in the Performance Gallery Gold **Memory Manager** graphs (see the data element titled **"CPU mem manager %"**).

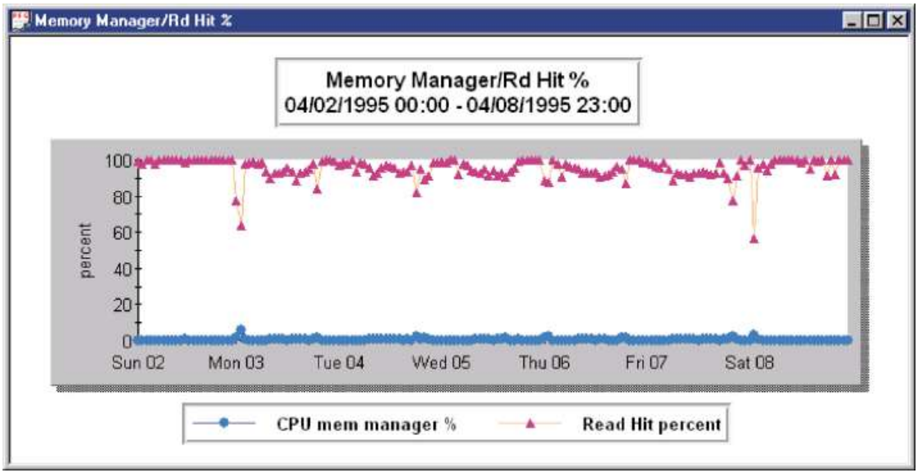

Figure 9.7 *Example Memory Manager/Rd Hit % line graph: CPU Memory Manager %*

Page Fault Rates

Table 9.8 *Page Fault Rates Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Page Faults/second			
Small, single processor HP 3000 series models 920, 922, 925, 932, 935	less than 4	4 to 8	greater than 8
Medium, max. 2-way HP 3000 series models 917, 927, 937, 947, 918, 928, 929, 939, 949	less than 8	8 to 12	greater than 12
Moderate, max. 2-way HP 3000 series models 950, 955, 957, 967, 977, 987, 960, 968	less than 13	13 to 19	greater than 19
Large, max. 2-way HP 3000 series models 959, 978, 980, 988, 990	less than 20	20 to 40	greater than 40
Larger, max. 4-way HP 3000 series models 959, 969, 979, 989, 992, 995, 996, 997	less than 40	40 to 60	greater than 60
Even larger max. 6-way HP 3000 series models 969, 979, 989, 992, 995, 996, 997	less than 100	100 to 150	greater than 150
Very large max. 8-way HP 3000 series models 969, 979, 989, 992, 995, 996, 997	less than 150	150 to 200	greater than 200
Note: Performance ranges for HP 3000 series models 996/900-996/1200 may vary depending upon the application.			

The *Page Fault Rate* pulse point shows the number of instances per second that page faulting occurred. A page fault is counted when a process needs a memory object (code or data) that is absent from main memory.

This information is reported in the Performance Gallery Gold **Memory Read Hit % / Page Fault Rate (MPE)** line graph (see the data element titled **"Page Faults"** in Figure 9.8).

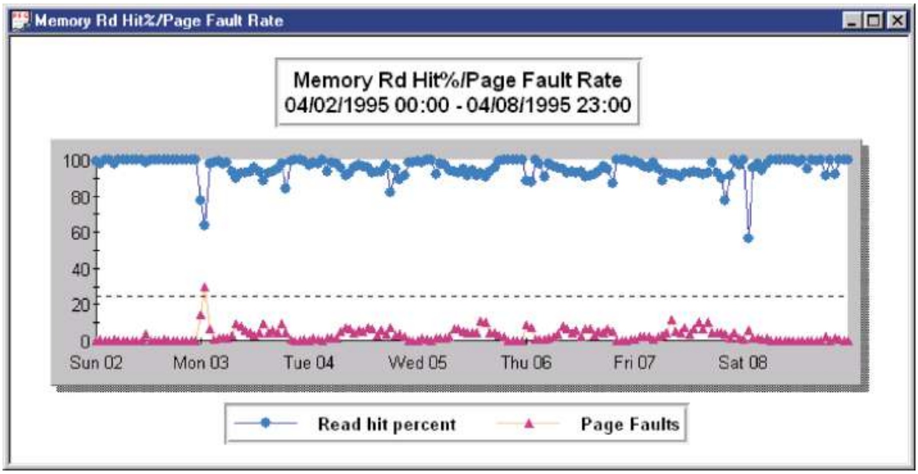

Figure 9.8 *Example Memory Read Hit % / Page Fault Rate (MPE) line graph: Page Faults*

Swaps per Launch Ratio

Table 9.9 *Swaps per Launch Ratio Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Swaps per Launch	less than 0.4	0.4 to 0.8	greater than 0.8

The *Swaps per Launch* pulse point shows the ratio of swaps to launches. A high ratio can indicate a shortage of memory. To improve performance, either reduce the demands for CPU or increase the available memory.

This information is reported in the Performance Gallery Gold **Memory Manager / Swaps per Launch (MPE)** line graph (see the data element titled **"Swaps/Launch"**).

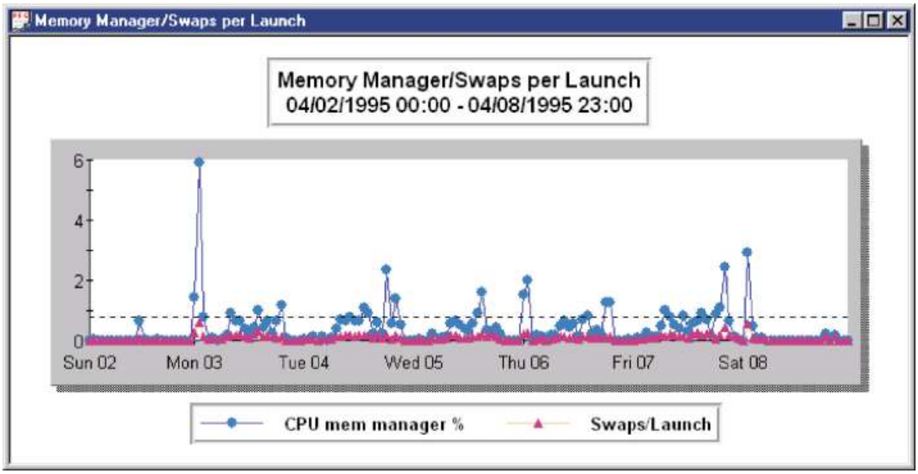

Figure 9.9 *Example Memory Manager / Swaps per Launch line graph: Swaps/Launch*

Disk Performance

CPU Pause Percentage

Table 9.10 *CPU Pause Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Pause %	less than 5	5 to 15	greater than 15

The *CPU Pause %* pulse point shows the percentage of time the CPU spent waiting for disk I/O's to complete.

This information is reported in the Performance Gallery Gold **CPU Utilization (MPE)** stacked area graph (see the data element titled "**CPU Pause**").

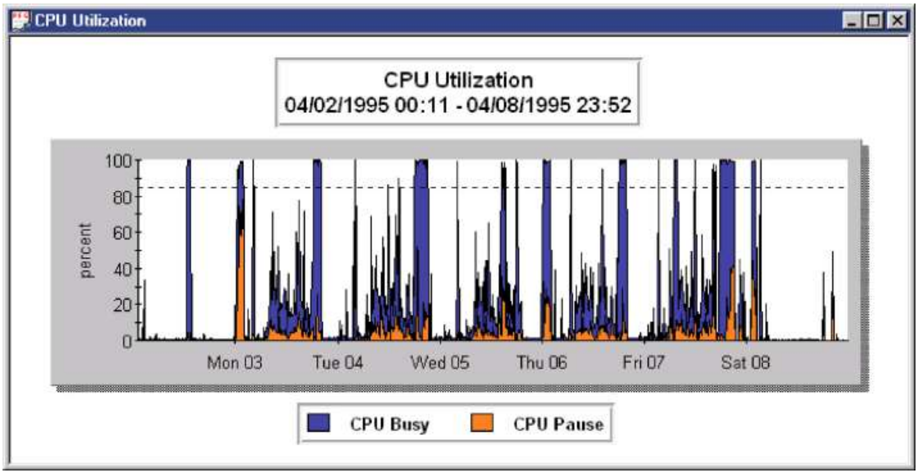

Figure 9.10 *Example CPU Utilization (MPE) stacked area graph: CPU Pause*

Disk I/O Rate

Table 9.11 *Disk I/O Rate Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Disk I/O's/second	less than 10	10 to 25	greater than 25

The *Disk I/O Rate* pulse point shows the total number of disk I/O's (writes and reads) per second.

This information is reported in the Performance Gallery Gold **Disk I/O Rates (MPE)** stacked area graph. The total disk I/O's is the sum of the reads and writes.

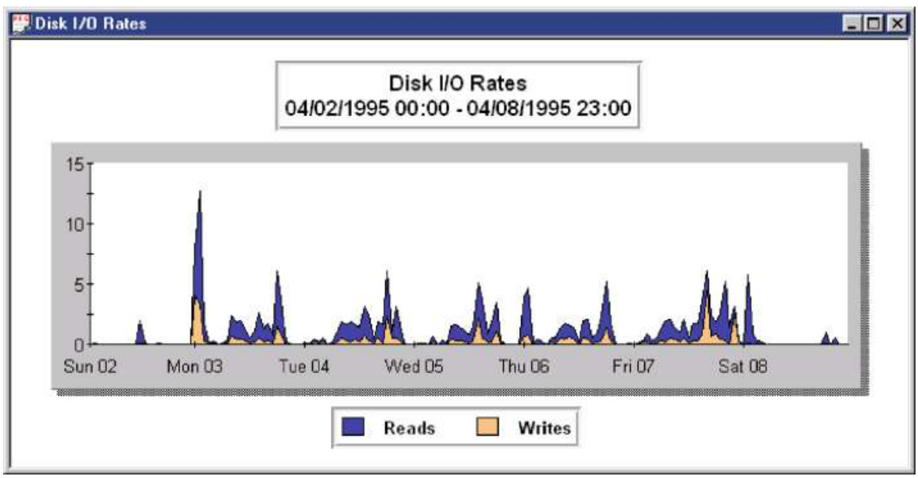

Figure 9.11 *Example Disk I/O Rates (MPE) stacked area graph*

A consistently high rate may indicate a bottleneck. To verify, check the *CPU Pause* indicator.

Disk Queue Length

Table 9.12 *Disk Queue Length Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Disk Queue Length	less than 0.5	0.5 to 1.0	greater than 1.0

The *Disk Queue Length* pulse point shows the average number of processes in the request queue for a particular disk drive.

This information is reported in the Performance Gallery Gold **Disk Queue Length by Drive (MPE)** 3-D surface graph.

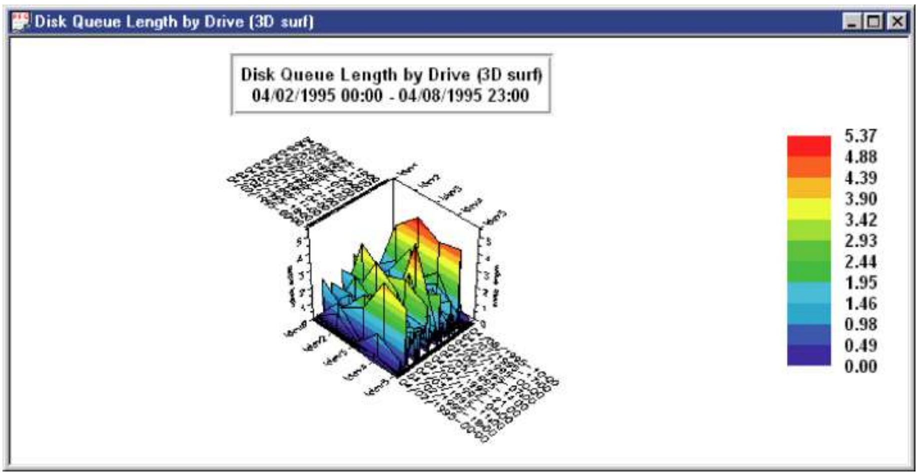

Figure 9.12 *Example Disk Queue Length by Drive (MPE) 3-D surface graph*

If the average queue length is high (greater than 0.5), investigate and remedy the cause (excessive disk arm movement, poor file management, poor application design, or inadequate disk drive speed).

Read Hit Percentage

Table 9.13 *Read Hit Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Read Hit %	greater than 95	85 to 95	less than 85

The *Read Hit %* pulse point shows the percentage of time that requests for information were satisfied in main memory.

This information is reported in the Performance Gallery Gold **Memory Manager / Read Hit % (MPE)** line graph (see the data element titled "Read Hit %").

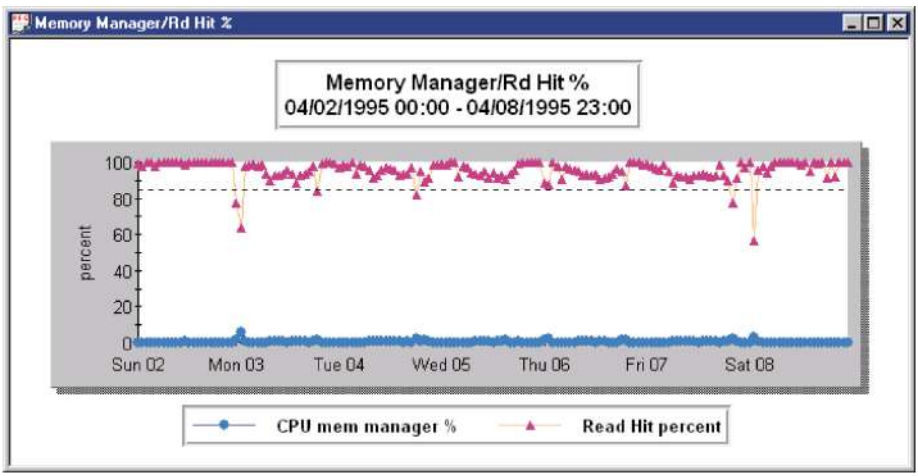

Figure 9.13 *Example Memory Manager / Read Hit % (MPE) line graph: Read HIt %*

A low Read Hit percentage may indicate a data locality problem or a shortage of memory

File Space

Performance Gallery Gold reports File Space information in 6 different templates: File Space Free (Stacked Area and Pie graphs), File Space Used, File Space Used by Device, File Space Utilization, and File Space Utilization by Device. These charts can help identify fragmented devices. The amount of free space required varies from system to system.

File Space Free

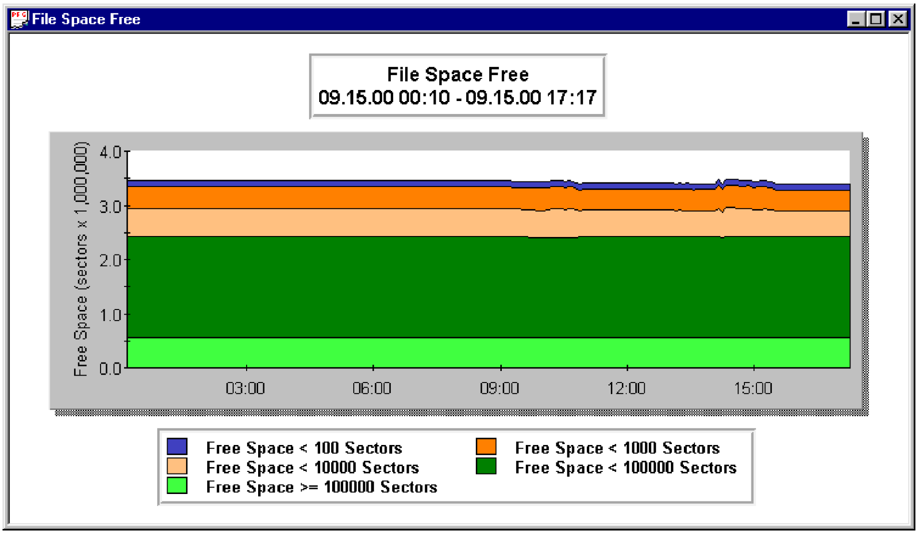

Figure 9.14 *Example File Space Free Stacked Area Graph*

The File Space Free graph (also available in a pie graph) displays the total amount of contiguous free space on your system. It is defined in sectors (256b) and categorized in five ways:

- <100 - The number of sectors of free space in chunks that are smaller than 100 sectors in size.
- <1000 - The number of sectors of free space in chunks that are 100 -999 sectors in size.
- <10,000 - The number of sectors of free space in chunks that are 1000 -9999 sectors in size.
- <100,000 - The number of sectors of free space in chunks that are 10,000 -99,999 sectors in size.
- >=100,000 - The number of sectors of free space in chunks that are 100,000 or more sectors in size.

In the example in Figure 9.14, out of the available 3.5 million sectors, approximately 2.5 million sectors are in contiguous chunks of 100,000 sectors or more. This system is not heavily fragmented. If your system displays that most of your free space is contained in chunks of 1000 or fewer sectors, your system is probably heavily fragmented.

File Space Used

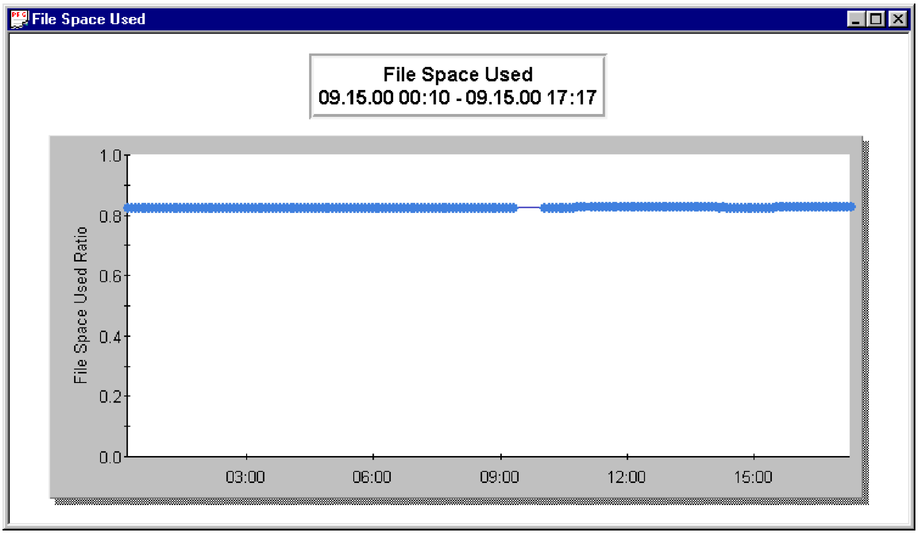

Figure 9.15 *Example File Space Used graph*

The File Space Used graph represents how much file space is being used and how much is free space. The area below the line is the percentage (0 -100%) disk space currently in use, while the area above the line is free space.

The default threshold for this graph is set to 80%, leaving 20% free space. This threshold changes based on the system. If your system runs many batch jobs at night, it may need 50% free space to contain temp files, while a system that doesn't run jobs may be fine with only 20% free space. The system in the example in Figure 9.15 has approximately 18% free space available.

File Space Used by Device

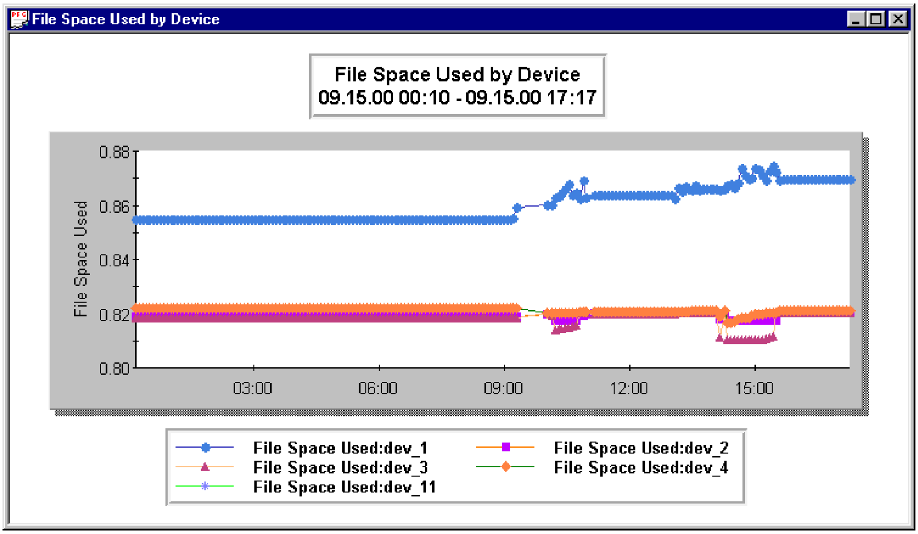

Figure 9.16 *Example File Space Used by Device graph*

The File Space Used by Device chart is an unstacked line graph that displays what percentage (0-100%) of space each device is using and how much is free space. In the example in Figure 9.16, Ldev 1 consistently uses around 86% of its file space, while Ldevs 2, 3, and 4 all use around 82%. If your system doesn't run batch jobs at night, requiring room for many temporary files, this level of free space may be acceptable.

File Space Utilization

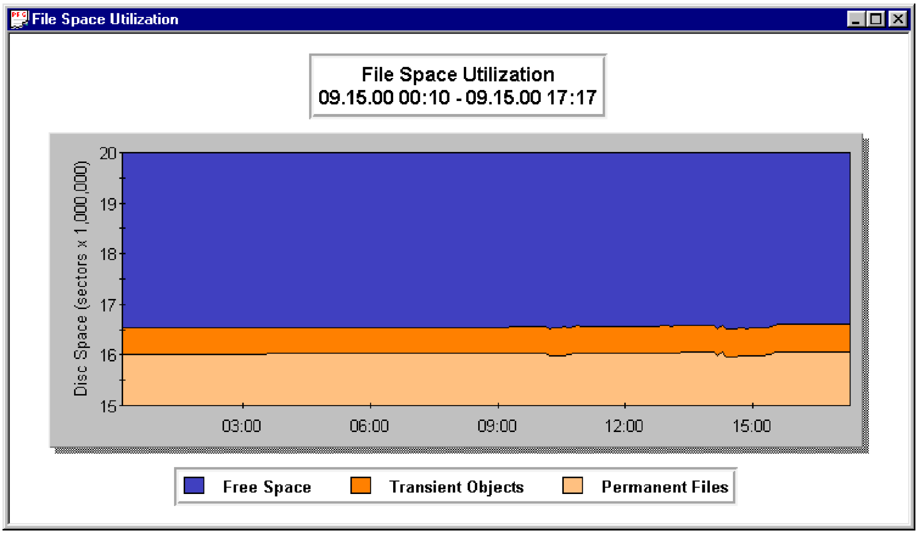

Figure 9.17 *Example File Space Utilization graph*

The File Space Utilization chart is a stacked area graph that displays how file space is being used. The graph shows how much of the available space, in sectors, is free space, transient (temporary) objects, or permanent files. From the example in Figure 9.17, this system has 3.5 million sectors of free space, 0.5 million sectors used for transient objects, and 16 million sectors used for permanent files. Again, the thresholds for this data changes depending on the system and what it is used for.

File Space Utilization by Device

File Space Utilization by Device (Tbl)

09.15.00 00:10.00	Trans Used	Perm Used	Trans Avail	Perm Avail	Max Trans %	Max Perm %	Largest Free Spa
dev_1	511168.0	3001216.0	597328.0	597328.0	91.0	91.0	57360.0
dev_2	0.0	4335072.0	959072.0	959072.0	100.0	100.0	113920.0
dev_3	0.0	4332080.0	962064.0	962064.0	100.0	100.0	115328.0
dev_4	0.0	4352704.0	941440.0	941440.0	100.0	100.0	113728.0
dev_11	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Figure 9.18 *Example File Space Utilization by Device table*

The File Space Utilization by device table displays all of the disk space information for each of your system's disc devices. This table includes the following for each device:

- Sectors of space used by transient and permanent files.
- Sectors of space available for use by transient and permanent files.
- The maximum percentage of the total space available for transient and permanent files.
- The largest chuck of free space.
- The number of sectors of free space in chunks that are smaller than 100 sectors in size.
- The number of sectors of free space in chunks that are 100 -999 sectors in size.
- The number of sectors of free space in chunks that are 1000 -9999 sectors in size.
- The number of sectors of free space in chunks that are 10,000 -99,999 sectors in size.
- The number of sectors of free space in chunks that are 100,000 or more sectors in size.
- Total amount, in sectors, of free space.

Process Information

The Process Information table provides detailed information regarding what processes were running at the time the data was sampled. This table is also conveniently linked from many Performance Gallery Gold templates. For example, right-click on an open CPU Utilization graph, and scroll to the bottom of the shortcut menu to **Process Information (Tbl)**.

For more in depth information regarding process information, see Chapter 6 of the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for MPE/iX.*

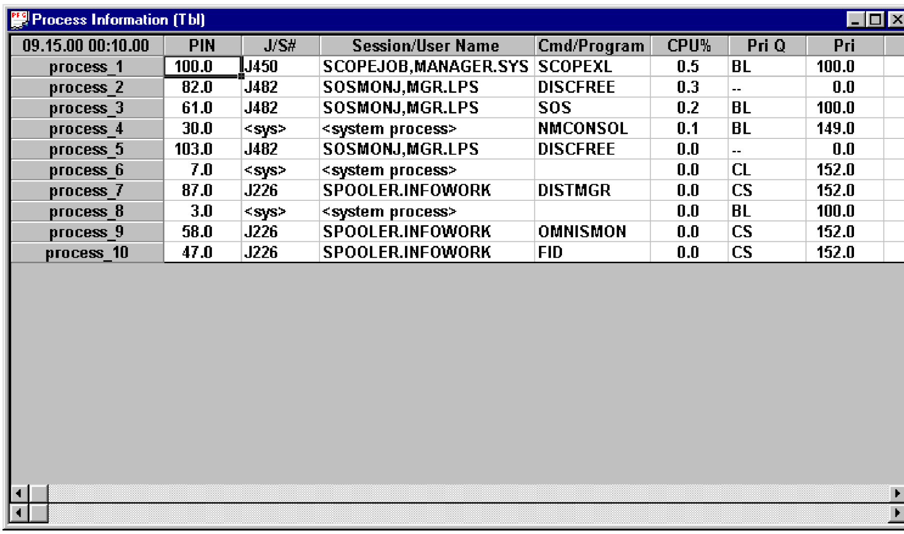
Process Information (Tbl)

09.15.00 00:10.00	PIN	J/S#	Session/User Name	Cmd/Program	CPU%	Pri Q	Pri
process_1	100.0	J450	SCOPEJOB,MANAGER.SYS	SCOPEXL	0.5	BL	100.0
process_2	82.0	J482	SOSMONJ,MGR.LPS	DISCFREE	0.3	--	0.0
process_3	61.0	J482	SOSMONJ,MGR.LPS	SOS	0.2	BL	100.0
process_4	30.0	<sys>	<system process>	NMCONSOL	0.1	BL	149.0
process_5	103.0	J482	SOSMONJ,MGR.LPS	DISCFREE	0.0	--	0.0
process_6	7.0	<sys>	<system process>		0.0	CL	152.0
process_7	87.0	J226	SPOOLER.INFOWORK	DISTMGR	0.0	CS	152.0
process_8	3.0	<sys>	<system process>		0.0	BL	100.0
process_9	58.0	J226	SPOOLER.INFOWORK	OMNISMON	0.0	CS	152.0
process_10	47.0	J226	SPOOLER.INFOWORK	FID	0.0	CS	152.0

Figure 9.19 *Example Process Information table*

PIN

This acronym stands for the **Process Identification Number** (**PIN**). A process is defined as the unique execution of a program. These processes can be executed by MPE/iX, a user or a batch job. Each process on the system is uniquely identified by its own **PIN**. You may locate processes best by knowing this number. A single job or session can have many processes associated with it.

J/S#

This is the job or session number associated with the particular process. If the process is a system type (not originating from a user job or session), **<sys>** will appear in this column.

Session/User Name

This is the logon sequence as initiated by the user or job minus the logon group. Once again, if the process was spawned by MPE/iX and not with a session or job then **<system process>** will be shown here.

Cmd/Program

This is the program or last MPE/iX command executed by the user. Some system type program names will be uniquely identified, such as "Spooler." If the process is a Command Interpreter Process, a **CI:** will appear in this column. Following **CI:** will be the last MPE/iX command the user or job issued. Notice that the actual command will only appear for root level **CI** process and not for subsequent **CI** processes in that process tree.

CPU%

This percentage reflects the amount of CPU capacity consumed by this process during the current interval.

Pri Q

This is the particular MPE/iX Dispatcher subqueue in which the process is executing. The possible letters for the dispatch subqueue are in the form of Nx. "N" is the subqueue and "x" is either "L" for linear or "S" for circular. Linear or circular denotes whether the process will remain at a fixed priority or if it will be changed by the operating system. Linear is fixed and circular can be changed.

AL

A very high priority linear queue. This queue is generally reserved for highest priority MPE/iX system processes that need assurance of adequate CPU time.

BL

A high priority queue. This queue is used by some lower priority MPE/iX system processes and by some very high priority user processes. For example, logging on a system with a "PRI=BS" parameter will allow your terminal to receive more CPU attention than those in lower queues. You should be cautious when running processes in this queue. If a looping condition takes place, often the only remedy is to restart the system! This is because processes in the A and B queues generally will not give up control of the CPU until they are through with it. This queue is linear.

CS

This subqueue is the one in which normal interactive sessions run. When you log on at a terminal, your Command Interpreter Process (the process that allows you to dialogue with an MPE/iX prompt) is assigned a priority of 152 in the C queue unless the default queue settings have been altered. As your process uses more CPU time than the average last 100 transactions (the SAQ value described in the Global and Disk Module), your priority is decremented (increased numerically - logically lower in priority). The net effect is that HOG interactive transactions are penalized. They have less chance of getting CPU time. Short transactions are rewarded by maintaining a higher priority. It is by this method that MPE/iX tries to fairly allocate resources among competing processes.

DS

This subqueue is generally used for high priority batch jobs. The rules for this and the E queue are described below and are similar to that of the C queue. Processes fall in priority as they exceed the filter values. In the CS queue this is the dynamically calculated SAQ value. For the D and E queues these values are the MINQUANTUM and MAXQUANTUM.

ES

This subqueue is typically used for lower priority batch jobs. Processes running at low priority will only get table scraps of CPU time. Processes running at higher priorities leave leftovers for these lower priority processes.

Pri

This is the absolute priority number that the MPE/iX dispatcher uses to determine who will get the CPU's attention next. This number is ultimately used to determine the CPU's next process customer.

#Rd

This is the absolute number of logical disk reads (usually not the same as physical) performed by this process during the current interval.

#Wr

This is the absolute number of logical writes (usually not the same as physical) performed by this process during the current interval.

#I/O

This is the sum of #Rd and #Wr.

LDV

This is the logical device number of the device at which the process was created. Batch jobs will show streams dev#, usually 10. System processes will show 0.

#Tr

This is the current number of terminal transactions (possibly equivalent to terminal reads) performed by the process to a particular terminal device. Under certain conditions this number will represent the actual number of user transactions (posting a payment, inquiring on an account, etc.). An inaccurate situation will be displayed if multiple carriage returns per screen are used for data entry. VPLUS status checks are not counted by the measurement interface (which Meta-View accesses). Therefore, transaction counts for VPLUS applications will be accurate. This number represents a consistent transaction count for VPLUS applications and is a fairly accurate count for character mode transactions. The best way to tell if terminal reads and transactions are equivalent is to test them. A user can enter a certain number of transactions defined from the user's standpoint and track that activity via Meta-View to see if there is any discrepancy. Refer to Chapter 6, "Process Summary by Application Workload" section, in the *Meta-View Performance Manager - User's Guide to Meta-View Agent and Host for MPE/iX* for a detailed discussion on how to capture statistics (response, transactions, etc.) on an entire group of programs and users.

PRes

This number represents the terminal read response time for interactive users. This is the response time for the user from the time C/R or ENTER is pressed to the first prompt at which the user enters a new transaction. Thus, only those processes performing terminal reads will have a positive number in this category, otherwise a **0** will appear.

Wait

This item represents the state of the process at the instant that Meta-View Performance Manager took the picture of the system at the end of the last sampling interval. When processes are "stuck," this number can help you find out why. Keep in mind that this wait state indicator is only a first line of defense. If you suspect an impedance problem, you should take a look at the process' wait state breakdown (i.e., the next item in this table) or go to the Process Detail screen to further analyze that process. Keep in mind that the wait state of a process (over even a few seconds) can change radically. These states are defined in Table 9.14.

Table 9.14 *Wait States*

Wait State	Definition
Bio	Waiting for non-disk I/O to complete
CPU	Currently active in the CPU resource
Dea	This process has terminated and will not show next interval. Its last will and testament (statistics) are shown
Disk	Waiting for Disk I/O to complete
FS	Waiting for activation by its father or son process
Imp	Waiting due to some resource being unavailable. An example is database locks, lack of system table entries, etc.
Mem	Waiting for a segment(s) to be brought into memory
Msg	Waiting for message file I/O, port sendmail or port receivemail
Pre	This process has been preempted by a higher priority process

Wait State	Definition
RIN	Waiting for a RIN to become available
Tim	Waiting for a timer
Trd	Waiting for a terminal read to complete
Twr	Waiting for a terminal write to complete
Oth	Waiting for a miscellaneous condition to complete

CP

This wait state is the percentage of the process' response time due to servicing by the CPU. That is, it takes time away from the CPU to perform the commands of processes.

ME

This wait state is the percentage of the process' response time that is due to time spent waiting for missing memory segments to return to main memory. When a process tries to continue to run but cannot because of missing necessary memory segments, that process is blocked. Memory fault stop time is counted in this category.

DI

This wait state is the percentage of the process' response time due to waiting for missing data to be brought into main memory from disk. An I/O brick wall occurs when a process wants to continue running, but cannot because necessary user-requested data is missing from disk. Since a process is literally stopped and the CPU is taken away when a physical disk access is performed, it is absolutely necessary to minimize this percentage.

IM

This wait state is the percentage of the process' response time due to being impeded by various lock and latch control mechanisms. This category includes many stop reasons. An impede occurs when a process tries to gain access to a software table or control structure and cannot because other processes arrived first. TurboIMAGE access is one of the most common sources of impedes. When a process wants to gain entry to a particular dataset and another process has that set locked via the DBLOCK intrinsic, then the waiting process is counted as having been impeded. It must wait until the prior process is finished with its current operation before it can continue.

Also, any file may have only one disk request outstanding. That is, in order for a process to access even a simple MPE/iX flat file, it must first gain control of that file's control block. This access is not by the FLOCK intrinsic as is the case in the RIN wait state bucket. Rather, only one user (regardless of programmatic locking) can gain access at a time. Other sources of impedes include unavailable system table entries, terminal buffers, etc.

PR

This wait state is the percentage of the process' response time due to preemption by other processes. A preemption occurs when a process is forced to give up use of the CPU because a higher priority process is ready to execute.

RI

This wait state is the percentage of time the process is waiting for a RIN.

TW

This wait state is the percentage of time the process is waiting for terminal writes to complete. Since terminal output is usually buffered, this will only accumulate time if the system runs out of terminal buffers or if the program is blocking on terminal output.

IO

This wait state is the percentage of time the process is waiting for non-disk I/O to complete (e.g., tape drive activity). Datacomm overhead is accumulated in this bucket as well.

TI

This wait state is the percentage of time the process is waiting for a programmatic timer (such as the PAUSE intrinsic) to complete.

FS

This wait state is the percentage of time the process is waiting on a father and/or son wait.

MS

This wait state is the percentage of time the process is waiting on a message file, port, or sendmail/receivemail wait.

OT

This wait state is the percentage of time the process is waiting on other events not covered by the above definitions.

CPU:ms

This number represents the amount of CPU milliseconds consumed by the process for the current interval. These milliseconds are the time that the process spends at the CPU receiving service. Current means the interval specified by the I:nn:nn at the top banner line. A cumulative total can be found on the Process Detail screen.

CPU/Tr

This value is the number of CPU milliseconds used by the process per each terminal transaction. This will always be blank for batch jobs because batch jobs do not perform terminal transactions. This number is calculated by dividing the total number of terminal transactions into the total amount of CPU used for the current interval.

D/Tr

The number in this column represents the number of physical Disk I/O's that were performed per user terminal read (Tr). If you have redefined a terminal read to mean a user transaction, then this number will reflect the average number of disk I/O's performed by the process per user transaction. If you have not, then this number will reflect the number of disk I/O's per terminal read.

C/N

The number in this column represents the number of compatibility to native mode switches incurred by the process during the current interval.

N/C

The number in this column represents the number of Native Mode to Compatibility Mode switches incurred by the process during the current interval.

%CM

This statistic represents the average amount of time the CPU spends in the Compatibility Mode program code for this process.

Launches

This statistic represents the total number of times the process was launched (given the CPU).

Summary of MPE/iX Pulse Points

For your convenience, all of the MPE/iX pulse points have been summarized in the following table.

Table 9.15 *MPE/iX Pulse Points Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	**Problematic**	**Unacceptable**
Processor Performance			
A and B Queue %'s	less than 5	5 to 8	greater than 8
CPU Busy %	less than 50	50 to 85	greater than 85
CPU CM Utilization %	less than 10	10 to 50	greater than 50
CPU Queue Length	less than 5	5 to 15	greater than 15
High Priority %	less than 50	50 to 85	greater than 85
ICS/OH+Dispatcher %	less than 10	10 to 15	greater than 15
Memory Performance			
CPU Memory Mgr %	less than 4	4 to 8	greater than 8
Page Faults/second			
Small, single processor HP 3000 series models 920, 922, 925, 932, 935	less than 4	4 to 8	greater than 8
Medium, max. 2-way HP 3000 series models 917, 927, 937, 947, 918, 928, 929, 939, 949	less than 8	8 to 12	greater than 12
Moderate, max. 2-way HP 3000 series models 950, 955, 957, 967, 977, 987, 960, 968	less than 13	13 to 19	greater than 19
Large, max. 2-way HP 3000 series models 959, 978, 980, 988, 990	less than 20	20 to 40	greater than 40

Performance Indicator	Performance Ranges		
	Normal	**Problematic**	**Unacceptable**
Larger, max. 4-way HP 3000 series models 959, 969, 979, 989, 992, 995, 996, 997	less than 40	40 to 60	greater than 60
Even larger max. 6-way HP 3000 series models 969, 979, 989, 992, 995, 996, 997	less than 100	100 to 150	greater than 150
Very large max. 8-way HP 3000 series models 969, 979, 989, 992, 995, 996, 997	less than 150	150 to 200	greater than 200
Note: Performance ranges for HP 3000 series models 996/900-996/1200 may vary depending upon the application.			
Swaps per Launch	less than 0.4	0.4 to 0.8	greater than 0.8
Disk Performance			
CPU Pause %	less than 5	5 to 15	greater than 15
Disk I/O's/second	less than 10	10 to 25	greater than 25
Disk Queue Length	less than 0.5	0.5 to 1.0	greater than 1.0
Read Hit %	greater than 95	85 to 95	less than 85

10

UNIX SYSTEM PERFORMANCE

The following Unix pulse points are provided by John Herberg of Lund Performance Solutions.

Processor Performance

CPU Busy Percentage

Table 10.1 *CPU Busy Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU Busy %	less than 60	60 to 85	greater than 85

The *CPU Busy %* pulse point shows the percentage of time the CPU spent executing the following activities instead of being in a pause or idle state:

- Processing user and system process code.
- Managing main memory.
- Scheduling and dispatching processes (interrupts).
- Processing context switches and overhead (external device activity).

This information is reported in the Performance Gallery Gold **CPU Detail (UX)** line graph (see the data element titled **"cpu>cpu-busy%"**).

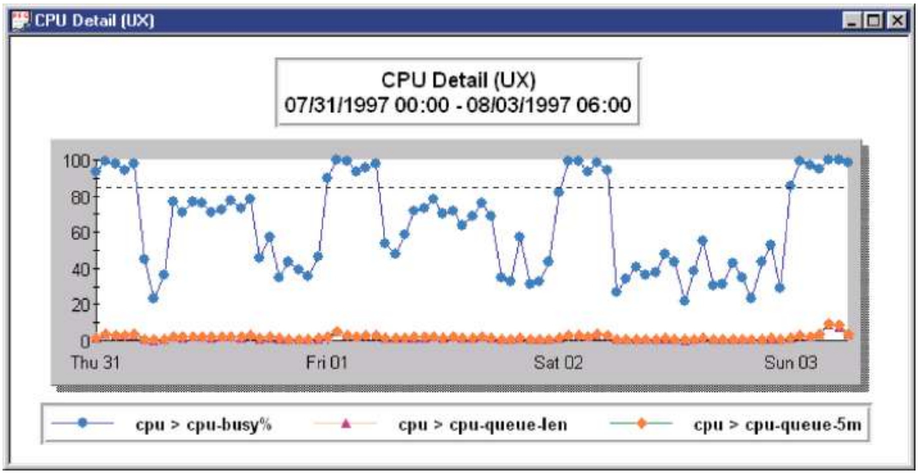

Figure 10.1 *Example CPU Detail (UX) line graph: cpu> cpu-busy%*

CPU High Priority Busy Percentage

Table 10.2 *CPU High Priority Busy Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
CPU High Priority Busy %	less than 60	60 to 85	greater than 85

The *CPU High Priority Busy %* pulse point shows the percentage of time the CPU spent executing high-priority system and user processes, interrupts, and overhead.

This information is reported in the Performance Gallery Gold **CPU Utilization (UX)** line graph (see the data element titled **"CPU high pri"**).

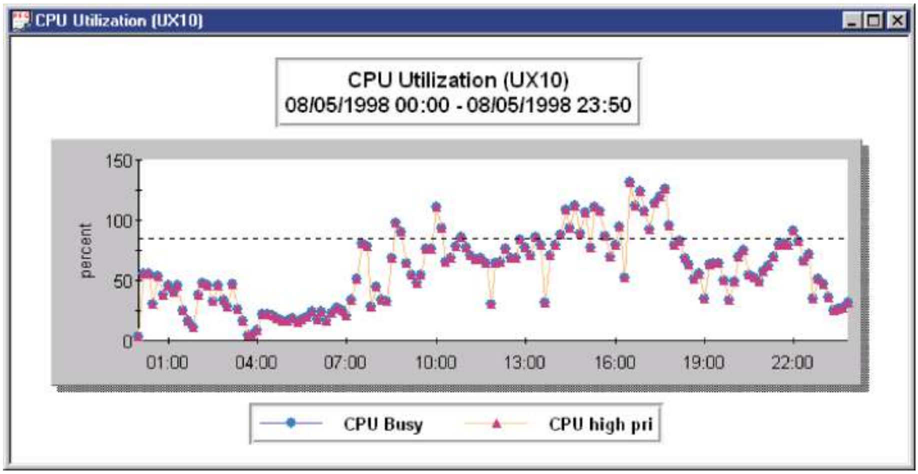

Figure 10.2 *Example CPU Utilization (UX) line graph: CPU high pri*

Real Time Processing Percentage

Table 10.3 *Real Time Processing Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Real Time Processing %	less than 5	5 to 10	greater than 10

The *Real Time Processing %* pulse point shows the percentage of time the CPU spent serving online, interactive sessions ("real time" user processes). These processes run at a fixed high-priority status.

This information is reported in the Performance Gallery Gold **CPU Utilization by Category (UX)** stacked area graph (see the data element titled **"Real"**).

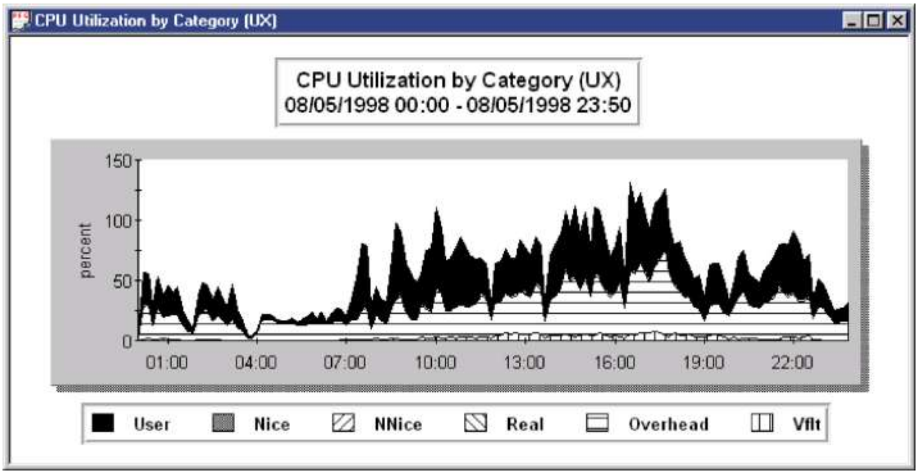

Figure 10.3 *Example CPU Utilization by Category (UX) stacked area graph: Real*

Run Queue Average

Table 10.4 *Run Queue Average Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Run Queue Average	less than 5	5 to 10	greater than 10

The *Run Queue Average* pulse point shows the average number of executable processes that waited for the CPU during a collection interval.

This information is reported in the Performance Gallery Gold **CPU Queue Detail (UX)** line graph (see the data elements titled **"Queue 5 min average"**).

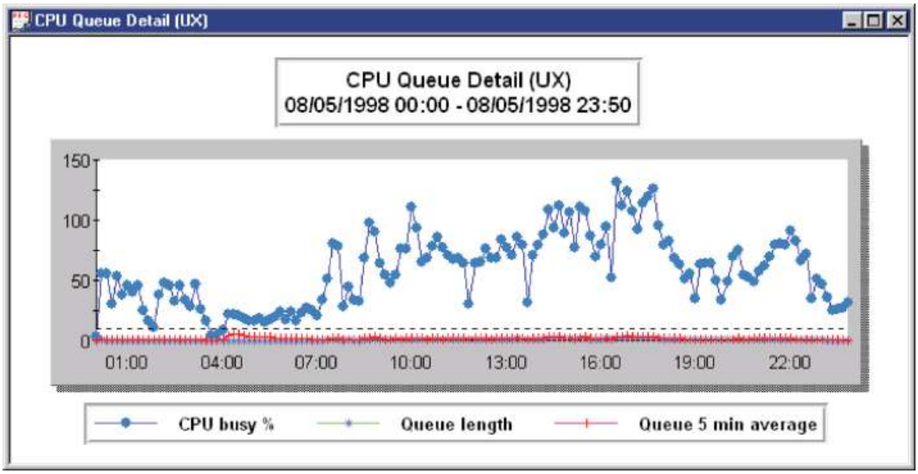

Figure 10.4 *Example CPU Queue Detail (UX) line graph: Queue 5 min average*

System Processing Percentage

Table 10.5 *System Processing Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
System Processing %	less than 10	10 to 20	greater than 20

The *System Processing %* pulse point shows the percentage of time the CPU spent executing system calls or operating in kernel mode.

This information is reported in the Performance Gallery Gold **CPU Utilization by Workload (UX)** subchart of system processes.

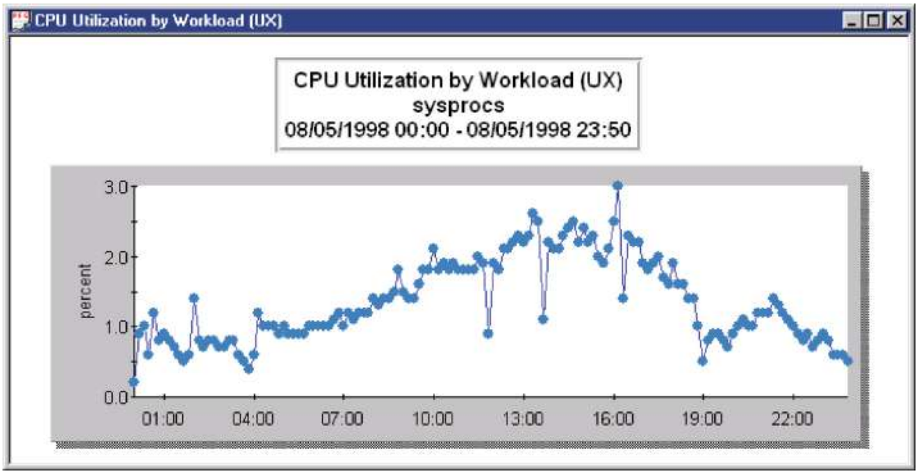

Figure 10.5 *Example CPU Utilization by Workload (UX) subchart (sysprocs)*

Memory Performance

Memory Used Percentage

Table 10.6 *Memory Used Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Memory Used %	less than 80	80 to 90	greater than 90

The *Memory Used %* pulse point shows the average percentage of main memory used during the collection interval.

This information is reported in the Performance Gallery Gold **Memory Used % / Deactivations Rate (UX)** line graph (see the data element titled **"Memory used percent"**).

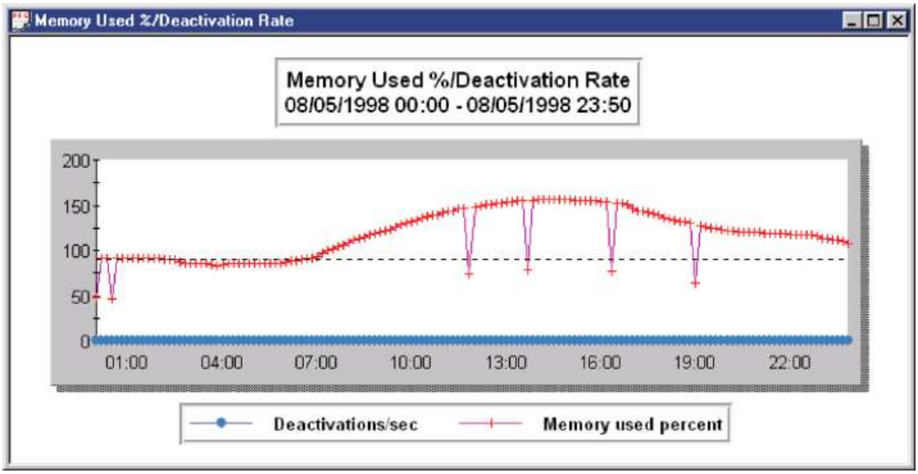

Figure 10.6 *Example Memory Used % / Deactivations Rate line graph: Memory used %*

Page Outs per Second

Table 10.7 *Page Outs per Second Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Page Outs/second	less than 5	5 to 10	greater than 10

The *Page Outs/second* pulse point shows the number of instances per second that a page out occurred during the collection interval. A page out is performed to move the least-needed pages from memory by writing them to swap space or to the file system. A page out occurs when physical memory becomes scarce.

This information is reported in the Performance Gallery Gold **Memory Rd Hit % / Page Fault Rate (UX)** line graph (see the data element titled **"Page outs/sec"**).

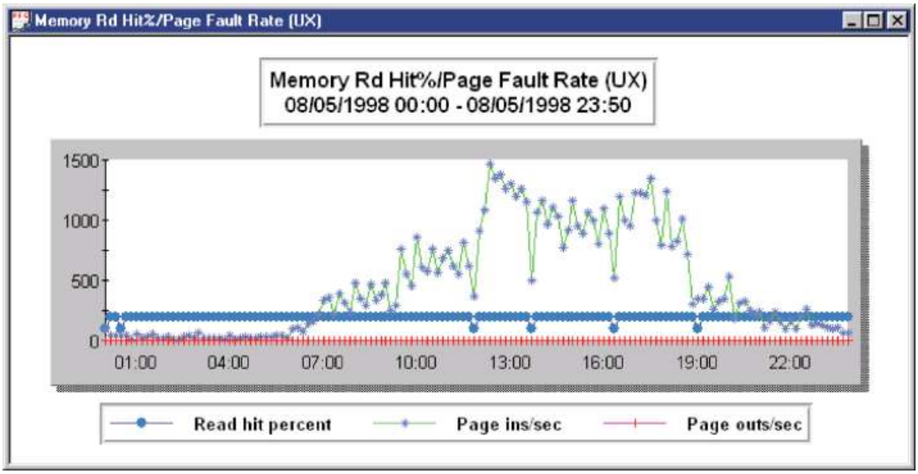

Figure 10.7 *Example Memory Rd Hit % / Page Fault Rate (UX) line graph: Page outs/sec*

Deactivations per Second

Table 10.8 *Deactivations per Second Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	**Problematic**	**Unacceptable**
Deactivations/second	less than 2	2 to 5	greater than 5

The *Deactivations/second* (swap outs/second) pulse point shows the number of processes swapped out of memory to disk in order to satisfy extreme memory shortages.

This information is reported in the Performance Gallery Gold **Memory Used % / Deactivations Rate (UX)** line graph (see the data element titled "**Deactivations/sec**").

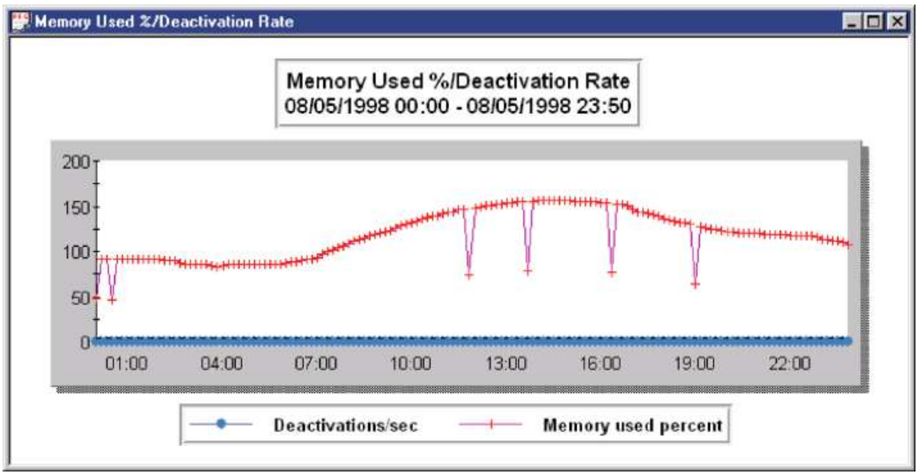

Figure 10.8 *Example Memory Used % / Deactivations Rate line graph: Deactivations/sec*

Disk Performance

Disk Queue Length

Table 10.9 *Disk Queue Length Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Disk Queue Length	less than 1	1 to 3	greater than 3

The *Disk Queue Length* pulse point shows the average number of processes in the request queue for a particular disk drive.

This information is reported in the Performance Gallery Gold **Disk I/O Queue Length (UX)** line graph.

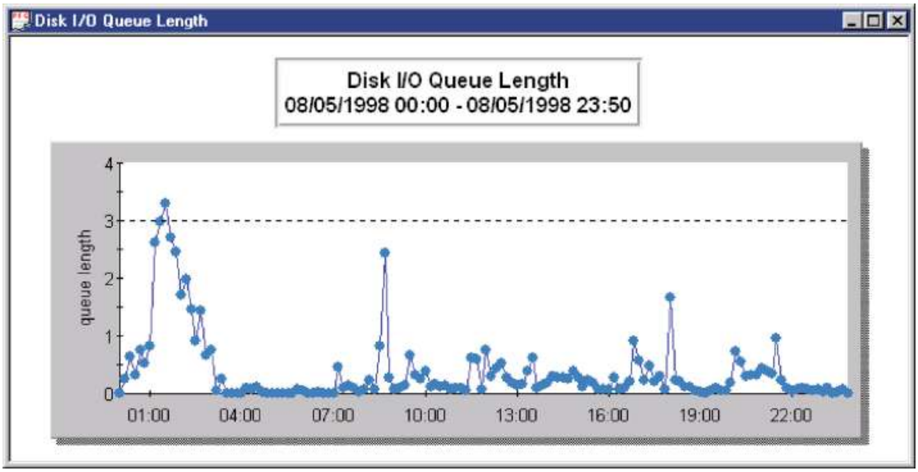

Figure 10.9 *Example Disk I/O Queue Length (UX) line graph*

Read Hit Percentage

Table 10.10 *Read Hit Percentage Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Read Hit %	greater than 90	90-80	less than 80

The *Read Hit %* pulse point shows the percentage of time that requests for information were satisfied in main memory. The Read Hit % value should be above 90. A percentage less than 90 could indicate a data locality problem or a shortage of memory.

This information is recorded in the Performance Gallery Gold **Memory Rd Hit % / Page Fault Rate (UX)** line graph (see the data element titled **"Read hit percent"**).

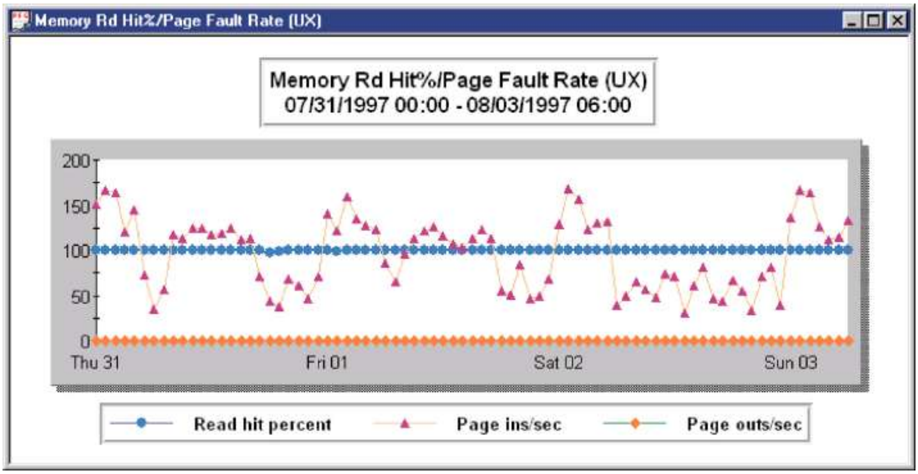

Figure 10.10 *Example Memory Rd Hit % / Page Fault Rate (UX) line graph: Read hit percent*

Disk I/O Rate

Table 10.11 *Disk I/O Rate Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	Problematic	Unacceptable
Disk I/O's per second	less than 40	40 to 60	greater than 60

The *Disk I/O Rate* pulse point shows the number of disk I/O (reads and writes to disk) per second.

This information is reported in the Performance Gallery Gold **Disk I/O by Drive (UX)** graph. This stacked area graph plots the disk I/O rate recorded for each disk drive.

To view the disk I/O rate for a particular disk drive, do the following:

1. Open the **Disk I/O by Drive (UX)** chart.
2. From the **Chart** menu, select **Modify Chart**.
3. From the **Modify Chart** dialog box, select the disk drive from the **Subchart** drop-down menu.
4. Click **OK**.

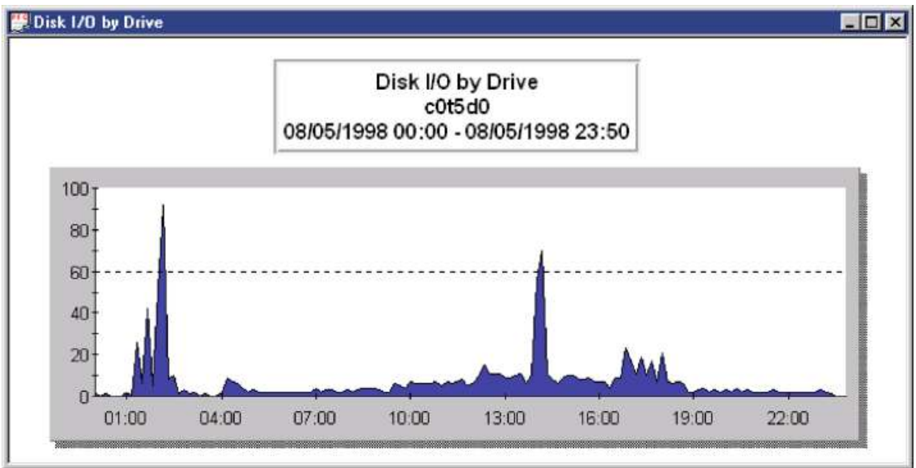

Figure 10.11 *Example Disk I/O Rate (UX) subchart, c0t5d0 disk*

Summary of Unix Pulse Points

For your convenience, all of the Unix pulse points have been summarized in the following table.

Table 10.12 *Unix Pulse Points Performance Ranges*

Performance Indicator	Performance Ranges		
	Normal	**Problematic**	**Unacceptable**
Processor Performance			
CPU Busy %	less than 60	60 to 85	greater than 85
CPU High Pri Busy %	less than 60	60 to 85	greater than 85
Real Time Processing %	less than 5	5 to 10	greater than 10
Run Queue Average	less than 5	5 to 10	greater than 10
System Processing %	less than 10	10 to 20	greater than 20
Memory Performance			
Memory Used %	less than 80	80 to 90	greater than 90
Page Outs/second	less than 5	5 to 10	greater than 10
Deactivations/second	less than 2	2 to 5	greater than 5
Disk Performance			
Disk Queue Length	less than 1	1 to 3	greater than 3
Read Hit %	greater than 90	90-80	less than 80
Disk I/O Rate	less than 40	40 to 60	greater than 60

Examples of HP-UX Charts

Following are some examples of HP-UX charts and detailed information related to them.

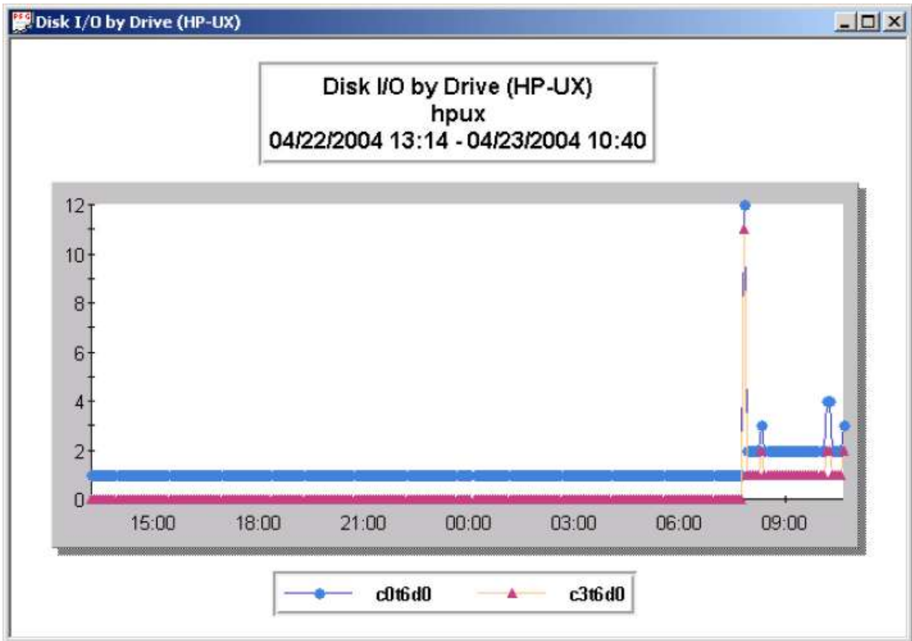

Figure 10.12 *Example of Disk I/O by Drive HP-UX Chart*

Figure 10.12 displays information related to Disk I/O by drive. There are no significant spikes during the period in study. There is a spike of 12 at 7:50 A.M., on April 23, 2004, on the c0t6d0 drive and also a spike of 11 at the same time on the c3t6d0 drive, but these spikes are not of concern. As we can see in Table 10.12, spikes less than 40 are normal. The other spikes that occur during the period in study are less than 4 on both drives and they are not of concern. The chart is study is a line chart.

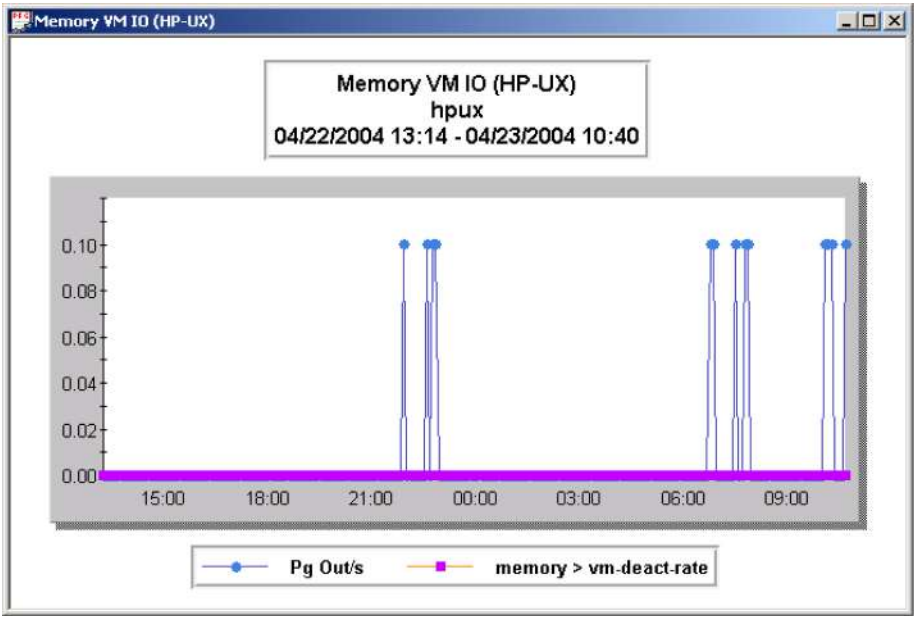

Figure 10.13 *Example of Memory VM IO HP-UXChart*

Figure 10.13 displays information related to Page Outs per second and Deactivation Rate. There are a few spikes of 0.1 per second for **Page Out/s**, during the period in study but there are not of concern. Spikes less than 5 per seconds are normal for Page Outs. There are no spikes for **deactivations** during the period in study. The chart above is a line chart.

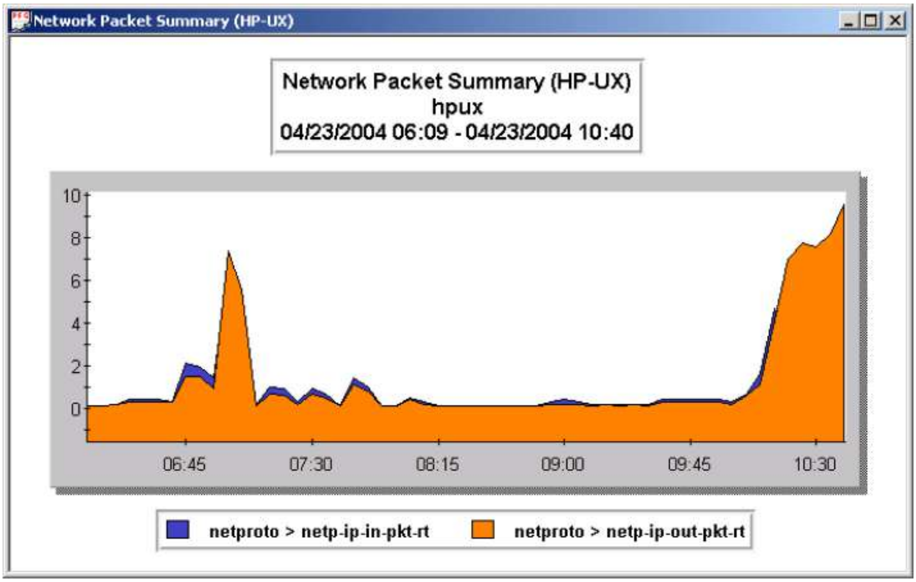

Figure 10.14 *Example of Network Packet Summary HP-UX Chart*

Figure 10.14 displays information about Network Packet Summary. There are a few spikes of both **netp-ip-in-pkt-rt** and **netp-ip-out-pkt-rt** elements of the Network Packets but the significant spikes occur at 7:00 A.M. and at 10:40 A.M. respectively. The chart in study is an area chart.

Examples of Linux Charts

Following are some examples of Linux charts and detailed information related to them.

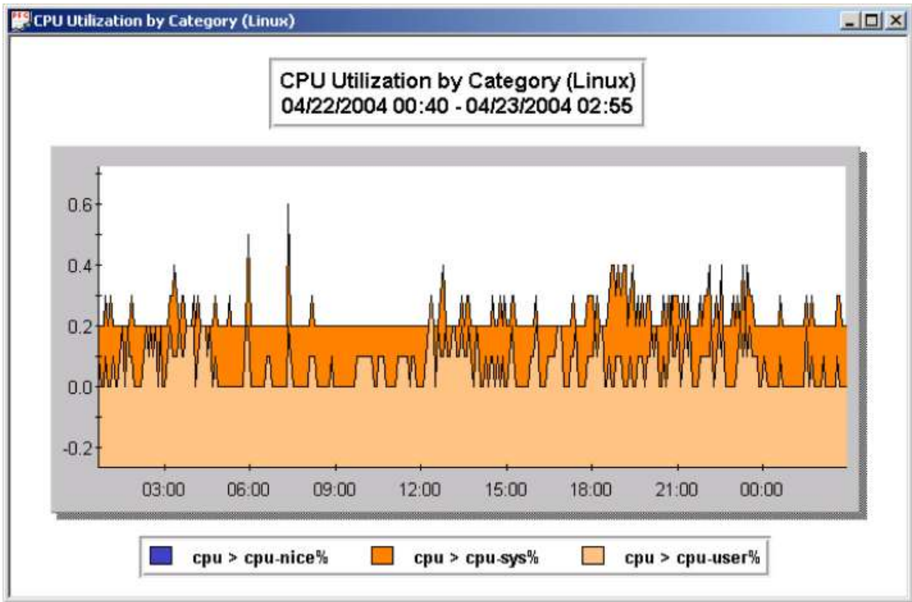

Figure 10.15 *Example of CPU Utilization by Category Linux Chart*

Figure 10.15 displays information related to Real Time Processing %. There are no problematic or critical spikes during the period in study. No spikes greater than 0.5 occurs and these spikes are normal. There is no spike of **cpu-nice%** element. The chart in study is an area chart.

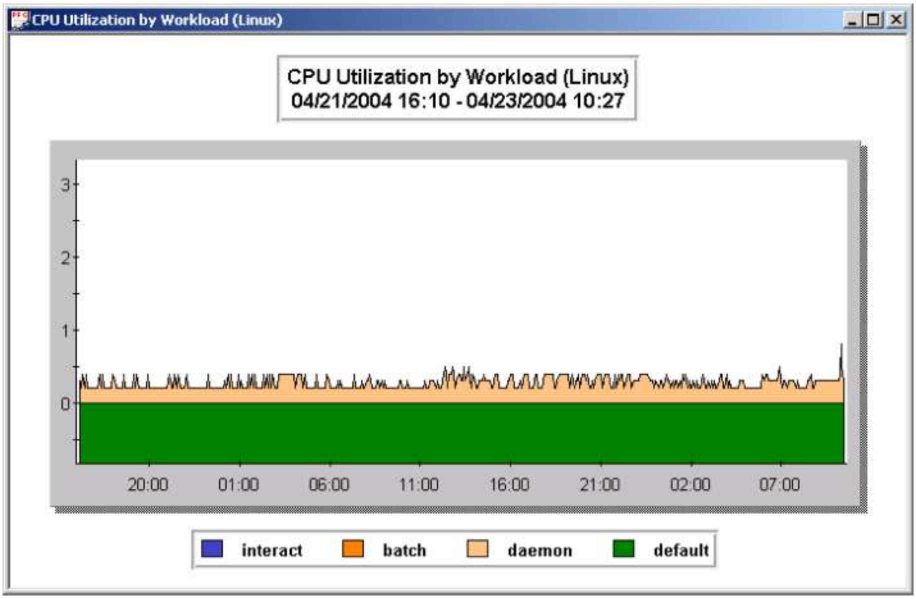

Figure 10.16 *Example of CPU Utlization by Workload Linux Chart*

Figure 10.16 displays information related to System Processing %. There are spikes only of the **daemon** workload and these spikes are normal, less than 1. There are no problematic or critical spikes during the period in study. The chart above is an area chart.

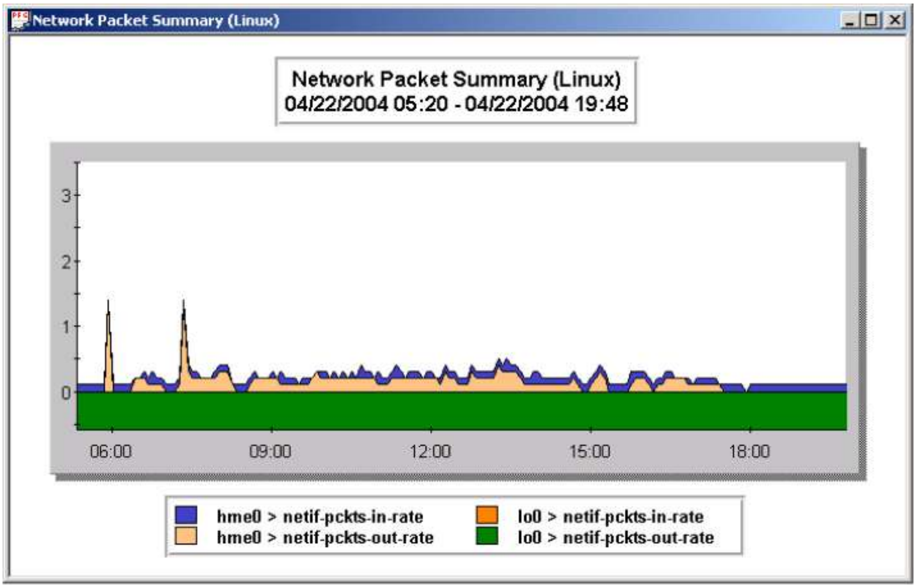

Figure 10.17 *Example of Network Packet Summary Linux Chart*

Figure 10.17 displays information related to Network Packet Summary. There are only a few spikes during the period in study. Two significant spikes occur around 6:00 A.M. and at 7:21 A.M. on **hme0 > netif-pckts-out-rate** element. These spikes are not of concern. The chart in study is an area chart.

Examples of Solaris Charts

Following are some examples of Solaris charts and detailed information related to them.

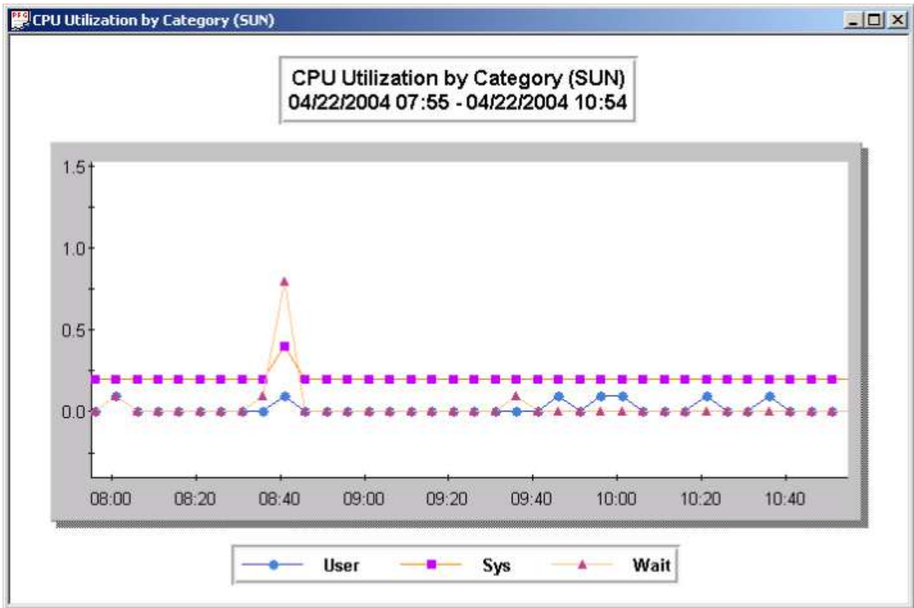

Figure 10.18 *Example of CPU Utilization by Category Solaris Chart*

Figure 10.18 displays information related to Real Time Processing %. There are no problematic or critical spikes during the period in study. At 8:40 A.M. there is a spike on all the three elements of CPU, **User**, **Sys**, and **Wait** but this spike is not of concern. The chart above is a line chart.

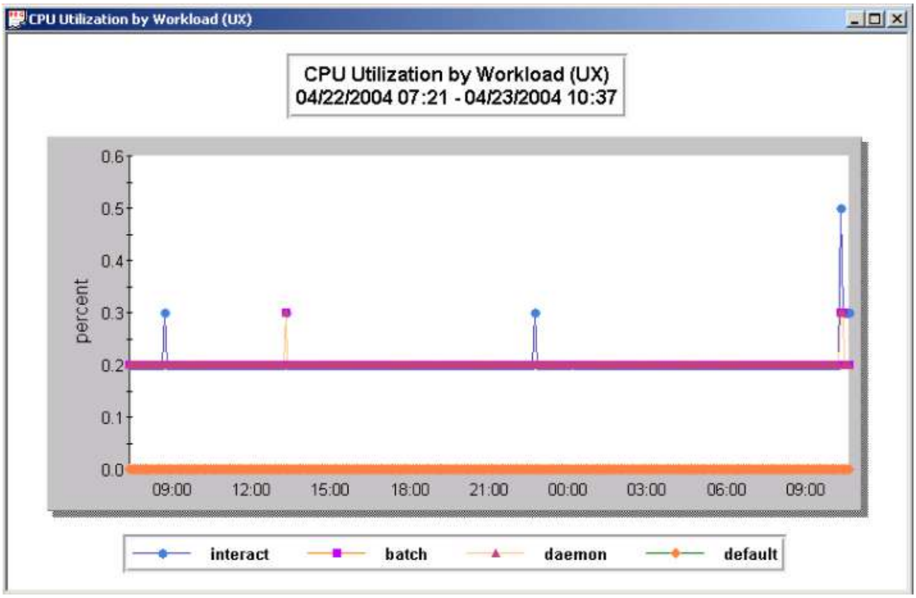

Figure 10.19 *Example of CPU Utilization by Workload Solaris Chart*

Figure 10.19 displays information related to System Processing %. There are a few spikes of the **interact** and **batch** workloads and these spikes are normal, less than 1. There are no problematic or critical spikes during the period in study. The chart above is a line chart.

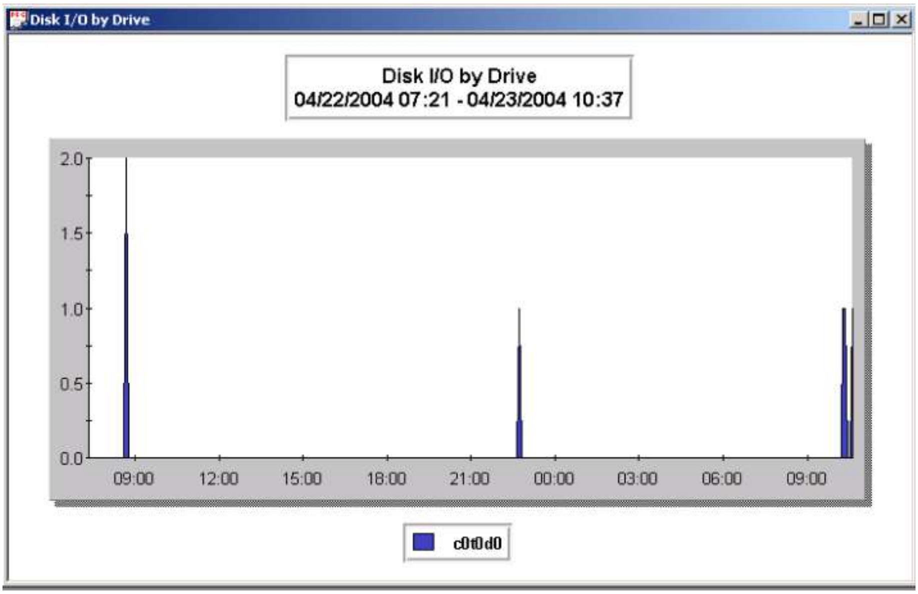

Figure 10.20 *Example of Disk I/O by Drive Solaris Chart*

Figure 10.20 displays information related to Disk I/O by drive. There are no significant spikes during the period in study. There is a spike of 2 at 8:41 A.M., on April 22, 2004, spikes of 1 at 10:42 P.M. on April 22, 2004, at 10:20 A.M. and at 10:37 A.M. on April 22, 2004, but these spikes are not of concern. Spikes less than 40 are normal. The chart is study is an area chart.

11

Windows NT/2000 System Performance

The Windows NT/2000 system performance information in the Performance Gallery Gold User's Guide is provided by Mark Friedman of Demand Technology Software.

Overview

Performance Gallery Gold works with Demand Technology Software's Performance SeNTry collection agent to report many crucial aspects of Windows NT/2000 performance. Performance Gallery Gold accepts one or more SMF (System Management Facility) collection files from a single computer as input and produces a variety of charts and graphs that can be used to pinpoint performance bottlenecks. It contains a variety of predefined chart templates that are designed to get you up and running quickly.

NOTE Refer to "Windows NT/2000 Objects and Counters" on page 255 for definitions of the Windows NT counters.

Processor Performance

The Windows NT processor charts focus on the System performance object, the Processor performance object, and the processor utilization of specific processes. The thread is actually the dispatch able unit on NT, but Performance Gallery Gold cannot report thread activity. Thread CPU activity is summarized to the individual process. NT supports both single-processor and multiprocessor configurations.

NT supports what is known as *symmetric multiprocessing*, which means, by default, any thread can run on any processor. The tuning options for a large NT Server multiprocessing configuration include defining hard processor affinity, which restricts the processors that threads are eligible to execute on. Using these tuning options demands a good understanding of the CPU resource usage pattern of various processes.

Processor Utilization

The Performance Gallery Gold **Processor Utilization Breakdown (NT)** graph separates overall processor utilization into four categories:

% Privileged Time counter

Privileged CPU time accumulates when NT operating system services run. These services include any Win32 services, including the Windows Graphical Device Interface (GDI) functions.

% User Time counter

User CPU time accumulates when application programs run and operating system services are not used.

% Interrupt Time counter *

Interrupt CPU time accumulates when interrupt service routines associated with device drivers run at high priority.

% DPC Time counter *

"Well-behaved" NT device drivers rely on deferred procedure calls (DPC's) for the bulk of their processing. Unlike interrupt service routines that mask off lower-level interrupts, DPC's run with interrupts enabled.

* Both interrupt services routines and deferred procedure calls run at higher priority than user and kernel threads.

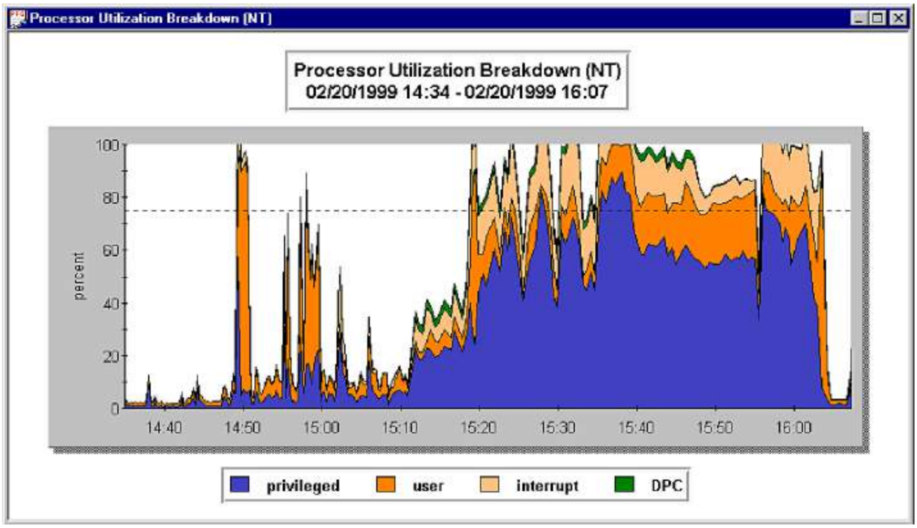

Figure 11.1 *Example Processor Utilization Breakdown (NT) graph*

The **Processor Utilization Breakdown (NT)** graph has a reference line at 75%. Systems with processors that run consistently greater than 75% busy can be reaching capacity limits. Verify by reviewing the Performance Gallery Gold **Processor Queue Length (NT)** graph to see how many threads are delayed in the NT dispatcher ready queue. If the processor queue length value is also large, a CPU bottleneck is probably impacting application performance.

Processor Queue Length

The Performance Gallery Gold **Processor Queue Length (NT)** graph shows the number of threads that are currently waiting for service in the NT dispatcher ready queue. The queue length is an instantaneous count of the number of ready threads at the end of the measurement interval (an integer value is reported), not an average. Of the threads in the ready queue, one thread per processor is in the running state. For instance, a thread from the Performance SeNTry collector process that is taking the measurements will show up in the running state.

NT has a single dispatcher queue, where ready threads wait, that services all processors. The reference line in the **Processor Queue Length (NT)** graph is drawn at two ready threads. A good rule of thumb is to have no more than two ready and waiting threads *per processor*. On a single processor, queue lengths of 6 or 8 ready threads are cause for concern, but 6 to 8 ready thread are okay on a four-way multiprocessor.

High processor utilization (greater than 75% CPU busy) coupled with a long ready queue indicates a CPU bottleneck. A long ready queue without high overall processor utilization may mean that a large number of timer-activated threads all happen to wake up at the same time that the Performance SeNTry collector does. In that case, the long processor queue length measured is a by-product of that behavior, not a sign of trouble.

Processor Utilization by Processor

The Performance Gallery Gold **Processor Utilization by Processor (NT)** graph shows processor utilization by both usage category and individual processor.

Windows NT/2000 is a symmetric multiprocessing system, therefore, processor utilization statistics are normally very similar across all processors. However, if you use processor hard affinity performance options, it may be necessary to examine these statistics on a processor-by-processor basis.

System Configuration

The Performance Gallery Gold **System Configuration (NT)** table shows hardware and software configuration information for the Windows NT system being monitored.The fields are described in the following table.

Table 11.1 *Windows NT hardware and software configuration data*

Parameter	Description
Domain	The security administration domain.
Computer	The computer name.
OS type	The type of operating system. The OS type for Windows NT Workstation is "winnt." The OS type for Windows NT Server is "lanmannt."
Version	The level of the Windows NT software.
CPU type	The AT-compatible processor, normally.
Installed memory	The total amount of memory installed, in bytes.
#CPU's	The number of processors installed.
CPU speed	The clock rate of the processor (not available for Alpha systems).

CPU Utilization by Process

The Performance Gallery Gold **CPU Utilization by Process Table (NT)** shows which applications are the heaviest consumers of processor resources. Processes include both applications that are started manually and applications that are started automatically by the Service Control Manager. (Services are equivalent to daemons in Unix or started tasks in MVS, and they run in the background.)

Some common processes are described in the following table.

Table 11.2 *Windows NT processes*

Process	Description
system	The operating system.
csrss.exe	The Client-server subsystem, which is the component responsible for desktop Windows management.
services.exe	Includes a number of system services packaged in a single executable program, including the network Computer Browser service, the network Redirector service, the network File Server service, the Alerter service, the Messenger service, and the Event log service.
lsass	The security remote validation service.
osa	A service associated with the Microsoft Office Suite.
rpcss	The service that supports NT's Remote Procedure Call interface.
smss	The Session Manager subsystem, which is involved in all security and authorization.
spoolss	The printer Spooler service.
winlogon	The application that controls access to the desktop. For example, the application that responds when you type Ctrl+Alt+Delete.

Memory Performance

The amount of random access memory (RAM) installed is one of the critical elements of Windows NT/2000 performance. This component is often referred to as *real memory*, which is contrasted with *virtual memory*, the logical view of memory that an application is granted. Each individual process views memory as a 4-gigabyte (GB) virtual address space. The upper half of each virtual address space is reserved for the NT operating system and various authorized components. The lower half of each virtual address space is available for use by each individual process. The system's virtual addresses are common to each process. while the per process areas are unique.

The virtual address spaces of active processes (including the System process) define a logical set of memory addresses that is normally considerably larger than the amount of RAM you actually have installed. The Virtual Memory Manager component of the NT operating system attempts to manage real memory, so the active pages of processes reside there, while inactive pages are stored on disk-resident paging files.

Virtual memory locations are mapped to real memory locations using a series of page tables. It is a function of the operating system to build and maintain a page table for each process that executes. Windows NT also tracks which real memory locations are in use and by which process.

Virtual memory systems are very dynamic. Only the active pages of a process need to reside in real memory, which is something that changes over time. When there is not enough real memory to go around, NT trims inactive pages from active processes and attempts to redistribute them based on current activity. NT will swap the pages of an inactive desktop application to the paging file on disk, for example, to free up real memory for some active process. *Paging* is disk I/O activity that occurs whenever an application references a valid virtual memory location that is not currently present in real memory.

Real memory is also used extensively for caching in Windows NT. *Caching* refers to storing frequently-used objects from the disk inside memory to speed up access. Caching is important, because memory access is thousands of times faster than disk access. The NT cache is an important part of NT's ability to serve as a file server for network clients. Other applications, like Internet Information Server, MS SQL Server, and MS Exchange, also rely on caching frequently-used objects. The frequently-accessed objects cached by these applications are Web pages, database data, and messages. There is also a sense that real memory serves as a cache for the frequently-accessed virtual pages of executing processes.

Real memory often has a bigger impact on NT performance than any other single component. This is because a number of applications rely on caching frequently-accessed objects in memory. Virtual memory systems can degrade quite suddenly; one minute they can be running fine and the next minute they can be hopelessly slow. Often, if does not take a major change to transition from a satisfactory state to an unsatisfactory one.

The two most important indicators of real memory performance are 1) the Available Bytes counter, which reports how much real memory is currently available for use, and 2) measures of demand paging activity to disk.

Real Memory Utilization

The Performance Gallery Gold **Real Memory Utilization (NT)** chart is a stacked area graph that shows the number of bytes currently allocated in the six system real memory pools. Each data point in the graph represents a specific counter; an instantaneous count of the number of bytes in the pool at the end of the measurement interval. The measurements are reported as bytes, because NT supports different page sizes on different hardware (4K pages for systems with an Intel processor, 8K pages for systems with an Alpha processor).

Real Memory Counters

The following six memory counters show how NT is using real memory to complete various system functions.

Available Bytes

The number of available bytes is a very important indicator to track. The Available Bytes counter tracks how much "free" real memory exists on the system during the specified measurement interval. More specifically, it represents the sum of three lists:

- The Standby List of stolen pages in transition.
- The Free List of aged, trimmed pages.
- The Zero List of free pages that are zero-filled.

New allocations are always made from the Zero List. A kernel thread called the "Balance Set Manager" is dispatched once per second to "trim" excess pages from each process address space. Initially, trimmed pages are put on the Standby List. As they age, trimmed pages are moved to the Free List. (If the page in real memory is modified, it must be written to the paging file before it can be moved to the Free List.) The low-priority system zero page thread is responsible for zeroing out pages on the Free List. The page trimming process is illustrated in Figure 11.2.

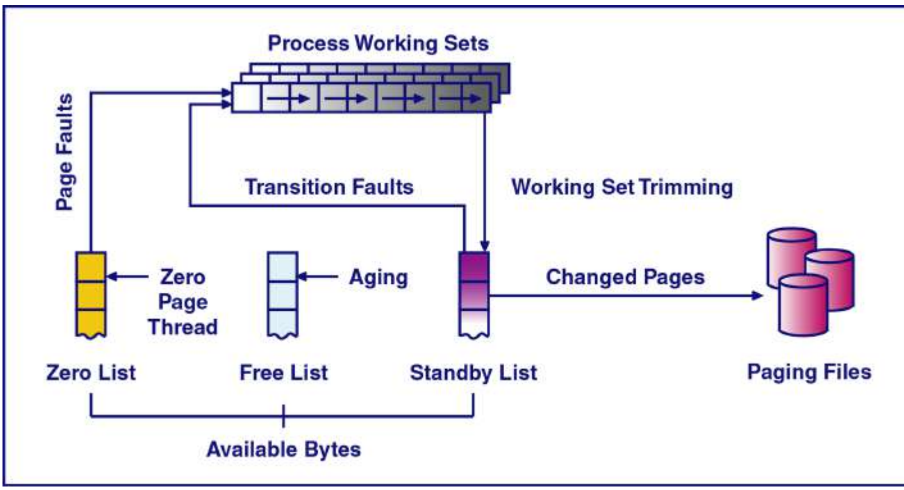

Figure 11.2 *Working set trimming process diagram*

As long as the amount of available memory is greater than 4-megabytes (MB), process working sets (resident pages of an active process) are allowed to grow. If the available memory falls below 4 MB, the system can become constrained by a storage of real memory. if available memory falls below 1 MB, the system is likely to be constrained and, consequently, process working sets are trimmed much more aggressively.

To confirm a shortage of memory, refer to the Performance Gallery Gold **Hard Page Fault Rate (NT)** graph.

Pool Non-paged Bytes

The Pool Non-paged Bytes counter tracks the memory that is not pageable. Many kernel threads allocate memory from this pool.

Pool Paged (resident) Bytes

The Pool Paged Bytes counter tracks the system memory that is pageable but happens to reside in real memory at the moment, presumably because the pages are active.

System Code Resident Bytes

The System Code Resident Bytes counter tracks the system memory where OS code is loaded. Many OS kernel routines must reside in real memory.

System Driver Resident Bytes

The System Driver Resident Bytes counter tracks the system memory allocated by device drivers that is currently residing in real memory. During I/O interrupt processing, device driver interrupt service routines must use virtual memory addresses that reside in real memory. I/O memory buffers also must be locked in real memory while they are being accessed by the devices.

System Cache Resident Bytes

The System Cache Resident Bytes counter tracks the number of bytes of pageable operating system code in the file system cache. The NT memory cache is mapped to a 512 MB segment of the system's 2 GB virtual memory address space. It is then allocated to active files in 256-kilobyte (KB) chunks. This is where the portions of the file cache that are residing in memory are counted. The cache bytes also includes some additional memory usage by system threads. Internally, the system process's working set is known as the cache.

Real Memory Calculation

To determine the real memory used by application processes at a given time, do the following:

1. Calculate the sum of real memory bytes used by all six memory counters.
2. Subtract this number from the total real memory installed on the system.
3. Refer to the Performance Gallery Gold **Memory Usage by Active Process (NT)** graph to see how different applications are using the available real memory.

Memory Usage by Active Processes

The Performance Gallery Gold **Memory Usage by Active Process (NT)** graph shows the usage of real memory by active processes.

NOTE By default, the Performance SeNTry collector filters out intervals where a process was not active, so not all running processes are accounted for in this graph.

The working set of a process consists of all its virtual memory pages that are currently resident in real memory. The process Working Set counter is an instantaneous value for the process at the end of the measurement interval.

Pages from a shared DLL that are currently resident in memory are counted as part of every process address space where the DLL is loaded. As a result, resident pages from shared DLL's can be counted two, three, four, or more times, so the sum of all the process working sets and the amount of system memory used is usually much greater than the amount of real memory installed on the system.

Virtual Memory Usage (commit%)

The Performance Gallery Gold **Virtual Memory Usage (NT)** graph shows the percentage of committed bytes in use.

NT uses virtual memory. Once virtual memory is committed by application processes, NT allocates space for the virtual memory on one of its paging files. Frequently-accessed, pageable virtual memory pages tend to reside in real memory. However, when more virtual memory is committed than exists inside the computer (its real memory), paging activity occurs. The mapping performed by the operating system of virtual addresses to real memory locations is transparent except when a program references a virtual address that is not resident in real memory. That causes a page fault to occur, which NT must handle.

Virtual memory pages are committed when NT reserves space for them on one of its paging files. A range of virtual addresses can be reserved without being committed. For example, NT reserves about 512 MB of the system's 2 GB virtual address range for the file cache. These pages are not committed until the cache manager allocates them to satisfy file requests.

Virtual memory allocations are limited by the amount of real memory and the amount of space on disk allocated for paging files. (NT supports up to sixteen paging files, one per logical volume, each of which can be as large as 1 GB.) This value is known as the *virtual memory commit limit*. When the number of committed virtual memory bytes reaches 90% of the commit limit, NT issues an "Out of Virtual Memory" message. Respond by creating another paging file or by increasing the size of the existing paging files. NT will automatically extend the size of the paging file, if possible, when the percentage of committed bytes remains above 90% for an extended period of time.

The **Virtual Memory Usage (NT)** graph includes a reference line at 70% (y=70). Systems that run consistently above 70% are more likely to experience excessive paging rates. The paging rate can be referenced from the Performance Gallery Gold **Demand Paging (NT)** and **Hard Page Fault Rate (NT)** graphs.

Demand Paging

The Performance Gallery Gold **Demand Paging (NT)** graph breaks down the page fault rate for an NT system according to three page fault categories:

- Hard faults that lead to paging file I/O.
- Soft, or transition, faults that are resolved without resorting to I/O operations.
- Cache faults which lead to I/O application files.

Each of these categories is described in detail in the following sections.

Hard Page Fault Rate

A hard page fault occurs when a process attempts to access a virtual memory location that is not resident in real memory. A hard page fault requires operations to the physical disk to be performed by the NT memory manager while the program that incurred the fault is forced to wait. A hard page fault is detected and resolved in the following manner:

1. Access to an invalid virtual memory locations causes an exception.
2. The NT memory manager examines the exception and determines that the reference is valid and the page does not reside in real memory.
3. The memory manager then "grabs" a real memory page from the Free List and calls the I/O manager to copy the page from the paging file to the designated memory location.
4. Once the page is resident in real memory, the program that incurred the page fault is restarted.

There is no NT memory counter that tracks the number of hard page faults per second; the Page Reads/sec counter is the closest equivalent. To determine the number of hard page faults, calculate:

Hard page faults = Page faults/sec - (Cache faults + Transition faults)

NOTE This computation, while logically correct, sometimes leads to a negative number.

NT performs bulk paging operations that are more efficient than executing a single disk I/O at a time. When NT encounters a page fault, it reads beyond the address specifically requested and takes several contiguous pages off the disk in a single disk operation.

- The Page Reads/sec counter tracks the number of times NT initiates a paging operation to read one or more pages from disk. Generally, multiple pages are input for each read request.

- The Pages Input/sec counter tracks the number of individual pages read from the disk. This counter can be referenced from the Performance Gallery Gold **Paging Activity (total) (NT)** graph.

When the disk (where the page fault resides) is very busy, page fault requests can accumulate. To check the efficiency of the NT paging operations, refer to the Performance Gallery Gold **Clustered Paging I/O Operations (NT)** graph.

The **Demand Paging (NT)** graph has a reference line at 25 hard page faults per second (y=25). The reference line can be adjusted (see “Marker” on page 131). The number of hard page faults you system can tolerate is determined by:

- The speed and number of paging file disks, each with its own capacity limits.
- The percentage of disk I/O bandwidth that can be devoted to paging.

NOTE Any I/O bandwidth involved in system paging activity cannot also be available to applications.

Soft (transition) Fault Rate

When a page fault occurs, the NT Memory Manager checks the contents of the Standby List to see if the page was among the pages trimmed recently by the Balance Set Manager. If the page requested is located on the Standby List, the Memory Manager is able to add it to the process working set immediately; no time-consuming I/O is involved. This type of page fault is referred to as a "soft" or "transition" fault. Soft/transition faults are just a by-product of the working set trimming practice that NT uses; they are not normally a performance issue.

Cache Fault Rate

By default, application I/O requests are diverted to the NT file cache, which represents a specific region of virtual memory. File segments are mapped to this area of virtual memory in 256 KB chunks. The range of the system’s virtual memory associated with the file cache contends for real memory like any other application process. A reference to a file location mapped to virtual memory and not in real memory causes an NT file cache fault.

When a cache fault occurs, the memory manager calls the I/O manager to generate a read request to the file. While a high rate of cache faults may indicate a shortage of real memory on a file server, many cache faults are simply the result of new files being accessed by different local processes and network clients. To determine the effectiveness of the file cache, refer to the Performance Gallery Gold **File Cache Hit% by Type (NT)** graph. There is little that can be done to tune the file cache in Windows NT beyond purchasing more RAM.

Clustered Paging I/O Operations

The Performance Gallery Gold **Clustered Paging I/O Operations (NT)** graph can help determine the efficiency of NT's paging operations (to disk).

Windows NT is designed to perform bulk paging operations, because the paging file disks are used most efficiently when multiple requests are processed in a single operation. Although bulk paging operations are efficient, they take longer than individual paging requests.

Page reads are the result of page faults. When NT discovers a page fault in a data page, it automatically reads several neighboring pages. It does the same for a process code page, but is even more aggressive about prefetching the pages from the disk. Trimmed pages that are subsequently modified must be updated on disk before the page in memory can be used by another process address space. The NT kernel modified page writer thread will wait until a number of changed pages accumulate before writing the modified pages to the paging file in bulk.

- Pages written per write

  The pages written per write data in the **Clustered Paging I/O Operations (NT)** graph is computed by dividing the number of pages output per second by the number of page writes per second.

- Pages input per read

  The pages input per read data is computed by dividing the number of pages input per second by the number of pages read per second.

In situations when the paging file disk is very busy and the processor is not, consider the following adjustments to improve system performance:

- Increase the I/O bandwidth for paging operation by adding paging files.
- Reduce the number of hard page faults by adding memory.

Paging Activity (total)

The Performance Gallery Gold **Total Paging Activity (NT)** graph shows the number of page inputs (Pages Input/sec counter) and outputs (Pages Output/sec counter) per second that occurred during the specified measurement interval.

The number of hard paging operations that a system can tolerate is determined by:

- The speed and number of paging file disks, each with its own capacity limits.
- The percentage of disk I/O bandwidth that can be devoted to paging.

The graph has a reference line at 100 pages per second (y=100). This line can be adjusted as appropriate for your system (see "Marker" on page 131). Consider that the I/O bandwidth that is absorbed in system paging activity is not available to application programs.

Paging Operations

The Performance Gallery Gold **Paging Operations (NT)** graph shows the total number of page read and write requests that occurred in the specified interval.

- The writes data element in the graph represents the Pages Output/sec counter.
- The reads data element represents the Pages Input/sec counter.

The number of pages read for input is greater than the number of page read operations, because NT performs bulk paging operations. The number of pages output is greater than the number of page write operations for the same reason.

The total number of pages input and output can be referenced in the Performance Gallery Gold **Paging Activity (total) (NT)** graph. The efficiency of NT's bulk paging operations can be referenced in the **Clustered Paging I/O Operations (NT)** graph.

Memory Utilization Index

The Performance Gallery Gold **Memory Utilization Index (NT)** graph produces a memory contention index that can be useful in predicting when an NT system might experience a memory bottleneck. This index is designed to help identify potential conflicts caused by a memory bottleneck and correct them before they become serious.

The memory utilization index is the virtual memory bytes allocated in the system pageable pool, divided by the memory bytes that are resident in the same pool. As system activity increases, the number of bytes in the pageable pool tends to increase. If the amount of real memory available to back this pool is limited, the ratio of virtual bytes allocated to real memory bytes consumed will increase. The value computed serves as a memory contention index, because the virtual memory pages in the pageable pool all contend for real memory.

An increase in the memory contention index is often accompanied by an increase in hard page faults. Combine the Performance Gallery Gold **Memory Utilization Index (NT)** graph and the **Hard Page Faults Rate (NT)** graph (see "Secondary Graph" on page 154) to see the relationship between the two performance indicators. The combined, dual-y-axes graph can be changed to a table (see "Chart Type" on page 114), then exported to MS Excel (see "Export" on page 86) and displayed as a scatter diagram (see Figure 11.3). The example scatter diagram includes a linear regression trend line to show the relationship between the two sets of measurements.

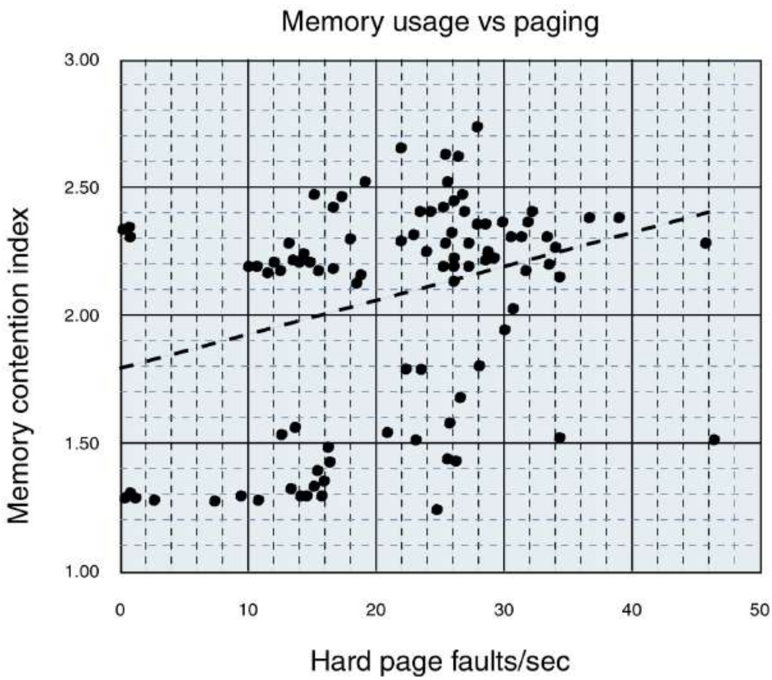

Figure 11.3 *Example Memory Usage Index vs. Hard Page Faults Rate scatter diagram*

File Cache Performance

On an NT Server machine configured to run as a file server, one of the biggest consumers of real memory is often the file cache. If the file cache is too small, performance will be poor, because there are too many accesses from disk and not enough accesses from memory. The most important way to tune a file server is to ensure it has an adequate amount of real memory to use for caching.

NT reports a variety of file cache statistics. The most important are the various cache hit ratios—the percentage of accesses resolved from memory compared to all file accesses. Keeping these hit ratios high is critical to system performance. Outside of one tuning parameter called "LargeSystemCache," which is a drastic change from the default working set management policy, there is very little that can be done to tune the file cache other than adding more RAM to the system. Because the file cache can never be larger than 512 MB, there is even an upper limit on how much RAM to add.

File Cache Activity by Type

The Performance Gallery Gold **File Cache Activity by Type (NT)** graph provides an overview of all NT file cache activities, including several different types of cacheable read requests and the amount of changed pages written to disk by the lazy writer thread.

The performance of the NT system file cache is critical when an NT system is used as a file server. File segments that are cached in real memory are accessed without having to access the physical disk. The best overall measure of the effectiveness of the file cache is to monitor the cache hit percentages. Cache hits occur when a request for a file was satisfied from the cache.

The **File Cache Activity by Type (NT)** graph shows the hit ratios for the following counters

- Copy Reads/sec.
- Data Maps/sec counter.
- Pin Reads/sec counter
- MDL Reads/sec counter
- Lazy Write Pages/sec counter

Each of these counters are described in the following sections.

Copy Reads/sec counter

When application files are accessed, they are read into the NT file cache first, then data buffers are copied from the NT file cache into the process virtual address space. Thus, normal application file requests become copy read requests. The copy reads data element shows how many copy read requests per second were satisfied from the cache without having to access the file. When application files are accessed sequentially, the file NT cache performs read-ahead operations to keep the copy read hit ratio high.

Data Maps/sec counter

The data mapping function of the cache is used primarily by the NTFS file system to cache master file table (MTF) entries. The NTFS master file table stores information about files and directories contained in the file system. This is how many times per second that data map requests are satisfied from the file cache without accessing the disk.

Pin Reads/sec counter

MFT entries that are mapped in the cache are pinned when they are modified. NTFS (NT file server) pins the changed MFT entries so the file system can exercise direct control over when the changed pages are flushed to disk. Following a change, but before the change is written to disk, a pinned entry is highly likely to be re-accessed by the file system. This data element is how many times per second those requests were satisfied without having to perform disk I/O operations.

MDL Reads/sec counter

Memory descriptor lists (MDL's) are buffers used by devices that support "scatter-gather" direct memory access (DMA) I/O operations. Scatter-gather devices reference multiple data areas in real memory within a single logical I/O request. The file server component of Windows NT also uses MDL's to improve the efficiency of large file requests. This data element shows how many times per second MDL reads were satisfied without having to access the disk.

Lazy Write Pages/sec counter

See the description for the Lazy Write Pages/sec counter on page 240.

File Cache Lazy Writer

When writes are issued to a cached file, NT does not immediately write back the changed data to the corresponding area where the file is stored permanently on physical disk. Instead, the action is deferred. This form of deferring write-back caching is called lazy writing.

There are several potential performance benefits associated with a lazy write policy. Because current data is available in the cache, an application that needs to re-read that information subsequently can get to the data without having to access the disk again. Applications like a word processing program that scroll through a large data file can benefit from this. In addition, if the application modifies the data again, the original write operation can lead to more efficient physical disk access. Suppose the application program subsequently changes a block near the original modification. Now, when changed pages of the file are flushed to disk, the physical disk can be accessed in a very efficient manner.

Keep in mind that because deferred write-back caching is used, application write requests do not translate into physical disk write operations immediately. When enough changed pages in the file cache accumulate, a cache manager lazy writer thread is dispatched to write a bunch of changed pages to disk in bulk.

The Performance Gallery Gold **File Cache Lazy Writer (NT)** graph includes data from the Lazy Write Flushes/sec counter and the Lazy Write Pages/sec counter.

Lazy Write Flushes/sec counter

The rate of the lazy write flushes per second.

Lazy Write Pages/sec counter

The number of lazy write pages per second that the lazy write thread writes to disk.

When the lazy write cache function is working efficiently, multiple changed pages are written to disk during each flush. On the other hand, if write activity is very low, there will be very few pages to flush, no matter how long NT defers the request. This graph does not present a stacked view, because the number of lazy write pages is usually greater than the number of lazy write flushes.

File Cache Read Activity

The Performance Gallery Gold **File Cache Read Activity (NT)** graph shows four different types of cached read requests.

Copy Reads/sec counter

When application files are accessed, they are read into the NT file cache first, then data buffers are copied from the NT file cache into the process virtual address space. This data element represents the total number of copy read requests.

MDL Reads/sec counter

Memory descriptor lists (MDL's) are buffers used by the devices that support "scatter-gather" direct memory access (DMA) I/O operations. Scatter-gather devices reference multiple data areas in real memory within a single logical I/O request. The file server component of Windows NT also uses MDL's to improve the efficiency of large file requests. This is the total number of MDL read requests.

Read Aheads/sec counter

When application files are accessed sequentially, the file NT cache performs read ahead operations to keep the copy read hit ratio high. This data element represents the total number of read ahead requests initiated by the file cache.

Pin Reads/sec counter

Master file table (MFT) entries mapped in the cache are pinned when they are modified. NTFS (NT file server) pins changes MFT entries so the file system can exercise direct control over when changed pages are flushed to disk. Following a change, but before the change is written back to disk, a pinned entry is likely to be accessed by the file system. This data element represents the total number of pinned read requests.

Cached File System Mapping Requests

The data map function of the NT file cache was designed mainly for use by installed file systems. The NTFS file system uses the data map function to cache master file table (MTF) entries. The NTFS master file table is where NT stores information about what directories and files are contained in the file system. This is also referred to as file system *metadata*.

The Performance Gallery Gold **Cached FS Mapping Requests (NT)** graph includes data from the Data Maps/sec counter and the Data Map Pins/sec counter.

Data Maps/sec counter

The data maps element reflects the rate of data mapping requests by the installed NT file system(s). MFT entries mapped in the cache are pinned when they are modified. NTFS pins change MFT entries so the file system can exercise direct control over when these changed pages are flushed to disk.

Data Map Pins/sec counter

The data map pins element is the rate at which data-mapped segments in the cache are pinned.

Normally, file system metadata changes quite frequently as new files are allocated, old files are deleted, and current (active) files grow and change.

File Server Performance

NT Server is often deployed as a network file and print server. Support for file server services is built into all version of Windows NT. This support is associated with the Server service, which resides in the **services.exe** program that runs as a service. Network file sharing uses a session-oriented wire protocol known as server management blocks (SMB's). SMB defines a protocol for accessing a shared disk resource, querying the contents of disk directories, and accessing individual files for reading or updating. The Server applet in the Control Panel enables you to view the status of the active file server sessions.

File Server Activity

File server activity is reported through various counters reported under the Server performance object, with more detailed performance data under the Server Work Queues object.

The Performance Gallery Gold **File Server Activity (NT)** graph tracks NT file server activity and the number of concurrent server sessions. The main indicators of file server activity are the number of bytes transmitted (the Bytes Transmitted/sec counter) and received (the Bytes Received/sec counter).

File Server Work Queues

Open the Performance Gallery Gold **File Server Work Queues (NT)** graph to view detailed file server performance statistics.

File Server Request Rate

File server requests are sent by network client using the SMB (server management block) wire protocol. The Performance Gallery **File Server Request Rate (NT)** graph tracks the rate of SMB requests serviced by the server component on the specified machine. The server request rate can be used in conjunction with the statistics in the Server Work Queues object to estimate the service time for requests handled by the machine.

Open the **File Server Request Rate (NT)** graph to view the rate at which SMB (server management block) requests were sent to the server from the clients across the network.

Logical Disk Performance

Disk performance is a critical area, because mechanical disks are relatively slow compared to other electronic components of a computer system. Disk performance can also be complicated, because there are so many options available to improve disk performance, including:

- Replacing slow disks with faster disks.
- Implementing software to defragment disks.
- Using the hardware caching option to speed disk access.
- Employing different disk configuration options, such as striping data across multiple physical disks.

Disk performance monitoring in Windows NT is performed at both the physical and logical disk level. It is performed by a component called "diskperf" (see “How diskperf works” on page 243). The NT disk performance statistics do a good job of characterizing the I/O workload, and they can be used to spot performance problems, although they are probably not detailed enough to tell you why a performance problem is occurring. To diagnose disk performance problems in NT can require very specialized and very specific knowledge of your disk hardware performance characteristics.

How diskperf works

diskperf is a special I/O filter driver program called *diskperf.sys*. Envision the I/O manager of Windows NT as a series of layers or a stack beginning with the file system and working down to the disk device driver layer, and, finally, to the SCSI miniport driver layer, which is provided by the maker of your SCSI interface board. I/O request packets (IRP) represent an I/O request. they are passed down through various layers of the I/O manager until they reach the SCSI miniport driver, where they are turned into SCSI disk commands. When the physical disk completes the command requested, status is returned back through the layers of the I/O manager stack all the way back to the original requesting application.

The diskperf filter driver gathers statistics on I/O requests. The module counts the number of I/O requests, the number of bytes transferred, and whether the request was to read or write data. In addition, diskperf times how long the request takes. The time diskperf measures is the amount of time the IRP spent during processing at the device itself, but also includes any time spent in the software layers below diskperf. Think of it as the round trip time to the disk and back, or the response time of the disk.

By default, the diskperf statistics gatherer is not enabled. Without it, no disk performance data is collected and all the performance counters associated with the logical and physical disk objects are zero. You must enable diskperf and reboot the system in order to collect performance statistics.

diskperf can be positioned in one of two places in the I/O manager stack that affects how it gathers statistics. The Windows NT disk administrator has basic options to format your hard drive. It also contains more advanced volume manager capabilities. These advanced features include:

- Volume sets that can be extended dynamically.
- Disk striping for better performance.
- Disk mirroring.
- Disk striping with parity (RAID 5) for fault tolerance.

An optional component called *ftdisk* (fault tolerant disk driver) implements these volume manager features. If any advanced volume manager functions are configured using the disk administrator, *ftdisk.sys* is loaded automatically during system initialization. diskperf.sys can be loaded either above or below ftdisk.sys in the I/O manager stack as illustrated in Figure 11.4.

- If you specify diskperf -y when you turn on disk performance monitoring, then diskperf.sys is loaded above ftdisk.sys in the I/O manager stack.
- If you specify diskperf -ye, then diskperf.sys is loaded below ftdisk.sys in the I/O manager stack.

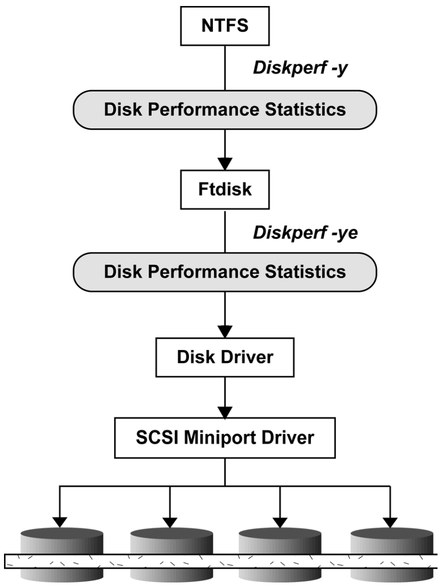

Figure 11.4 *I/O Manager stack diagram*

Since what ftdisk.sys does is redirect a logical disk request to its appropriate physical disk representation, the place where diskperf.sys is loaded influences the measurement data that is collected. For example, suppose you create a logical disk, D:, which is a mirrored disk constructed from two physical partitions—one on physical disk 1 and the other on physical disk 2. ftdisk.sys is the component responsible for: 1) taking a single logical disk request to write records to the D: disk, and 2) sending the SCSI miniport driver two identical write commands—one to physical disk 1 and the other to physical disk 2.

- If diskperf.sys is loaded above ftdisk.sys, diskperf will record the logical disk statistics correctly, but it will not understand that two physical disk requests were issued for each write operation requested.
- If diskperf.sys is loaded below ftdisk.sys, diskperf will measure the two physical disk requests that were issued correctly, but it will not understand that only one logical disk write operation was requested.

The logical disk and physical disk statistics that diskperf collects are identical. If ftdisk.sys is loaded, diskperf.sys can collect logical disk statistics, or correct physical disk statistics, but not both. You are likely to see some very strange performance statistics if you access the wrong object.

Logical Disk Response Time

The Performance Gallery Gold **Logical Disk Response Time (NT)** graph shows the average elapsed time of I/O requests, from the viewpoint of *diskperf* (see “How diskperf works” on page 243). The fields reported here correspond to the Avg. Disk sec/Read counter and Avg. Disk sec/Write counter under the Logical Disk performance object. This is the total response time of I/O requests, so it includes both disk service time and queue time spent waiting for NT software components and the disk hardware.

A reference line is included in the **Logical Disk Response Time (NT)** graph at 25 milliseconds (Y=0.025 seconds). This is an arbitrary boundary, but a good place to begin worrying about disk performance problems. If the response time is consistently above the 25 ms threshold, it would be worthwhile to investigate the cause. Access the link to the **Logical Disk Utilization (NT)** graph to see how poorly the disks are performing. Unless the high response times are associated with high rates of activity, the problem is probably not severe enough to warrant serious attention.

Disk Performance Expectations

A reasonable service time expectation for today’s popular disks is that most requests can be serviced in 10-15 ms or less. Disk service time is generally made up of the following components:

Seek Time

Seek time is the time spent positioning the read write arm of the disk over the proper track. In sequential file access, the arm is already positioned in the right location (unless another application I/O request "steals" the arm), and there is no seek time component. At the opposite extreme, there are long seeks from the beginning of the disk to the end, which might take 20 ms or more. On average, a seek of one-third the distance across the platter normally takes 8-10 ms, depending on the make and model of the disk.

Since *sequential* is a popular mode of access, figure that your workload needs to perform average seeks only 50% of the time.This reduces the expected average seek time per I/O request to about 5 ms. Optimizations like defragmenting your hard drive regularly will increase the number of sequential accesses performed and will improve I/O performance.

Rotational Delay or Latency

Once the arm is positioned to the right disk track, it is necessary to wait until the sector requested rotates under the read/write head before the operation can continue. Popular disks rotate at 5,400-10,000 RPM’s, which is about 6-12 ms per rotation. Sometimes the disk does not have to rotate very far before the right sector is reached. Other times you have to wait much closer to a full disk revolution. On average, you can expect a half-rotation delay per I/O, which is about 3-6 ms.

Protocol Delay

Protocol delay refers to the SCSI or ATA commands that must be sent back and forth across the interface bus to establish communication with the disk and tell it what to do. The SCSI protocol requires about 0.5 ms

Data Transfer

Finally, the disk is ready to read or write the data as requested. How fast the device can transfer data is a function of 1) how fast the disk spins and 2) how much data is recorded per track. Disks today are capable of moving data at a rate of 10-25 MB per second.

File Size

Another factor in data transfer time is the size of the request. In NTFS, valid sector sizes range from 512 bytes to 4 KB. Today's disks can transfer even the larger 4 KB blocks in less than 1 ms. Keep in mind that many areas of Windows NT are optimized to perform bulk requests and prefetch data in anticipation of its use. Both demand paging and file cache prefetching requests tend to be bulk requests. In general, this results larger blocks transferred in a single I/O request.

If you add up the various components of I/O service time, you get:

Seek time + Latency delay + Protocol delay + Data transfer = I/O service time

A reasonable service time expectation is about 11 ms for an average disk on an average day. Of course, there can be many workload factors that are above or below average. The service time expectation calculated above ignores all other delay factors, like what happens when there are multiple requests active at the same time for a single disk, or when the disk shares access to the SCSI bus with other devices. On a busy system, all of these factors can easily contribute to queuing delays that are as long as the actual service time. So, you might not want to investigate the various I/O tuning strategies until the response time measures are consistently above 20-25 ms per request, on average.

Disk Hardware Performance Options

Besides simply buying faster disks—ones that spin faster also transfer data faster—there are many other hardware solutions that improve disk performance.

Actuator-level Buffers

A feature common to many high-performance disks is a built-in, dedicated, *actuator-level buffer*. During a read request, buffered devices transfer data from the track requested directly into this built-in buffer memory. If a subsequent request asks for data that is adjacent to the original request, the data can probably be found in the actuator buffer. A buffer hit means only protocol time and data transfer time are needed to satisfy the request. There is no seek time or latency delays associated with a buffer hit. Also, most devices can transfer data from the buffer at full interface speed, which might be 40 MB per second, while the device transfer speed is usually somewhat slower.

Caching Controllers

Similar to actuator-level buffers, *caching controllers* boost performance when the data requested is found in the cache. This allows the request to be serviced without having to access the physical disk. Again, there is no seek time or disk latency delays involved, just protocol and data transfer time.

Disk Arrays

Disk arrays are two or more disks grouped together and accessed in parallel to speed up performance. This grouping can be performed two ways:

- Via software using the NT disk administrator. Array operations are implemented by the **ftdisk** fault tolerant disk driver module, **ftdisk.sys**.
- Via hardware using disk striping. Disk striping spreads a single logical request across multiple blocks of data. However, in most Windows NT environment, disk striping adds little performance value, because of the small sector size used by NTFS.

RAID (Redundant Array of Independent Disks)

RAID refers to disk arrays in which information is also replicated to make the configuration fault tolerant. In a RAID configuration, a single drive in the array can fail without losing data permanently. The simplest form of RAID organization is disk *mirroring*, also called RAID level 1. Disk striping is sometimes combined with mirroring to achieve high performance with the benefit of fault tolerance. This is called RAID level 0/1.

RAID 5 is the most popular fault tolerant disk organization. RAID 5 is equivalent to the striping with parity option in disk administrator. Instead of making a full copy of the original data on a separate disk as in mirroring, RAID 5 creates a single parity error correction code sector for each corresponding set of data sectors. The parity data is sufficient to recreate the original data following a single disk failure. RAID 5 is a high-availability feature, first and foremost. Disk performance actually suffers under RAID 5, because of the extra steps that must be taken to maintain the parity information during writes. Watch out for this RAID 5 write performance penalty. Battery-packed controller cache is used in more expensive RAID 5 subsystems to buffer writes and to mask the RAID 5 write penalty.

Logical Disk Detail

The Performance Gallery Gold **Logical Disk Detail (NT)** table is a report that summarizes the most important disk performance measurements. Use this table to look for problem disks that have a combination of both a high activity rate and a high response time.

The data fields in the **Logical Disk Detail (NT)** report are described below.

Read RT (response time)

The read response time data is based on the Avg. Disk sec/Read counter. This is the same value reported in the **Logical Disk Response Time (NT)** graph. It includes disk service time and all queuing delays.

Reads/sec

The reads per second data element is the rate of disk read requests, based on the Disk Reads/sec counter. It is the activity rate for read requests during the measurement interval.

Avg Read Bytes

The average read bytes data element is the size, in bytes, of the average read request. It is calculated by dividing the Disk Read Bytes/sec counter by the Disk Reads/sec counter. The average block size is significantly larger than the file system sector size, which reflects the use of bulk requests in NT, which are more efficient than individual disk sector accesses.

Write RT (response time)

The write response time data is based on the Avg. Disk sec/Write counter. This is the same value reported in the **Logical Disk Response Time (NT)** graph. It includes disk service time and all queuing delays.

Writes/sec

The writes per second data element is the rate of disk write requests, based on the Disk Writes/sec counter. It is the activity rate for write requests during the interval.

Avg Write Bytes

The average write bytes data element is the size, in bytes, of the average write request. It is calculated by dividing the Disk Write Bytes/sec counter by the Disk Writes/sec counter. The average block size is significantly larger than the file system sector size, which reflects the use of bulk requests in NT.

in System

This data element corresponds to the Avg. Disk Queue Length counter, which is the product of the overall I/O request rate (Disk Transfers/sec counter) and the average response time value in the Avg. Disk sec/Transfer counter. The calculation uses Little's Law, a basic tenant of queuing theory. A more detailed explanation of this calculation is provided in the description for the Q length data (below).

Q Length

Queue length is an instantaneous value based on the Current Disk Queue Length counter. It is the number of requests that are currently active, including any I/O request that is in service at the disk at the time the disk performance data was collected. For example, a current disk queue length value of 2 indicates that one request is active and one is currently waiting.

% Free Space

This data element is the percentage of the free disk space at the time of the data collection, based on the % Free Space counter.

MB's Free

This data element is the free space on the disk expressed in terms of absolute space, based on the Free Megabytes counter. The Throughput data element is the sum of the Disk Read Bytes/sec counter and the Disk Write Bytes/sec counter.

Logical Disk Utilization

Windows NT reports three disk performance counters, which are described as measures of disk busy:

- % Disk Time counter
- % Disk Read Time counter
- % Disk Write Time counter

The official Microsoft explanation for % Disk Time is that it is "the percentage of elapsed time that the selected disk drive is busy *servicing* [emphasis added] read or write requests." This is a misleading explanation. Strictly speaking, these counters do not report disk utilization, which is why they sometimes behave in odd ways. The Performance Gallery Gold **Logical Disk Utilization (NT)** graph reports the three counters associated with % Disk Time and shows their relationship to the Avg. Disk Queue Length counter.

Disk Utilization

Disk utilization can normally be found by measuring the activity rate and the average disk service time. % Disk Time could then be calculated as the product of the Disk Transfers/sec counter and disk service time. This is known as the Utilization Law in queuing theory.

Disk utilization = Requests/sec x Service time

However, NT's diskperf program does not measure disk service time. Instead, it measures disk response time, which is service time plus queue time.

Disk response time = Service time + Disk queue time

Disk Average Queue Length

The product of the Disk Transfers/sec counter and the Avg. Disk sec/Transfer counter (disk response time), according to the well-known Little's Law formula, is the average number of outstanding disk requests—also known as average queue length.

 # in system = Requests/sec x Response time (Little's Law)

Masquerading as utilization, the % Disk Time counters are artificially capped at 100%. This is because it is impossible in queuing theory (and in reality) for a server to be greater than 100% utilized. This can sometimes lead to absurd results. For example, say the Logical Disk D: is 2.49% busy from reads, 100% busy from writes, and 100% total busy, according to the % Disk Time counters. The product of Disk Transfers/sec and Avg.Disk sec/Transfer can be greater than 100%, because the Avg. Disk sec/Transfer counter measures response time.

Three corresponding Avg. Disk Queue Length counters were introduced in NT 4.0 to try to clear up this confusion. These queue length counters are calculated using Little's Law, without the capping. The Avg. Disk Queue Length counters are logically consistent while the purported disk busy counters are not.

How should the Avg. Disk Queue Length counters be interpreted? They literally represent the average number of outstanding requests to the disk, including any requests that are currently in service. A value of 2.79 for Avg. Write Disk Queue Length means, on average, that one request is in service at the disk while almost two requests are always waiting. For reality testing, the calculated values of the Avg. Disk Queue Length counter should be compared to the measured values of the Current Disk Queue Length as shown in the **Logical Disk Detail (NT)** table. But also keep in mind that the measured values of Current Disk Queue Length are likely to be systematically under-sampled (especially on uniprocessors). This is due to the fact that the disk driver Interrupt Service Routine and Deferred Procedure Calls that manage the device queue run at a higher dispatching priority than the Performance SeNTry collection service. By the time the Performance SeNTry collection service is eligible to run, the disk driver software may have already dispatched the next I/O request.

Logical Disk Average Queue Length

The Performance Gallery Gold **Logical Disk Avg Q Len Statistics (NT)** table provides the raw values for the (false) Disk% Busy counters and the (correct) Avg. Queue Length counters. Both are computed as the product of the activity rate and the response time. The (false) disk % busy counters are capped at 100%, while the (correct) Avg. Queue Length counters are not.

Additional measurement anomalies are introduced by the *ftdisk* fault tolerant disk driver (see "Logical Disk Performance" on page 242). These anomalies can be explored in the **Logical Disk Avg Q Len (NT)** table.

The Performance Gallery **Logical Disk Avg Queue Length (NT)** graph is a three-dimensional temperature chart (3D surface chart) that enables you to easily spot disk performance problems across a large server disk farm. The color-coded temperature chart legend shows the range of measurement data broken down into ten equal ranges. The red, orange, and yellow peaks identify devices with lengthy queuing delays.

Physical Disk Performance

A set of Performance Gallery Gold charts nearly identical to the logical disk charts is available for reporting attached physical disk performance. Physical disk performance statistics are gathered by the NT diskperf driver program (see "Logical Disk Performance" on page 242). The same exact measurement data is provided with the same quirks. The only difference is the logical disk measurements include statistics on free disk space—these statistics are not available for the physical disk.

Redirector Performance

The Redirector component of Windows NT is used to transform (or redirect) networked file requests to use the network instead of a local disk. It is the client-side of file server requests. In fact, the sum of all Redirector requests for network clients is precisely equal to the sum of all file server requests, assuming all network clients are running Windows NT.

Network Activity (redirector)

The Performance Gallery **Network Activity - Redirector (NT)** graph shows the total amount of network traffic associated with redirected file requests. The value shown in the graph corresponds to the Bytes Total/sec counter in the Redirector performance object. From this graph, you can link to the following charts:

- **Redirector Errors by Type (NT)** - shows Redirector error statistics.
- **Redirector Bytes Received/Sent (NT) -** breaks down bytes transmitted into reads and writes.
- **Redirector File Operations (NT) -** shows Redirector file operations.

Network Traffic Performance

When the Windows NT Network Monitor Agent is installed on a system, a variety of network utilization statistics can be collected. In general, the networking statistics are gathered at the network interface card level, with the network monitor collecting basic information on the number and types of network requests, including the number of bytes associated with each request. Statistics are available for each of the different networking protocols that Windows NT supports, including TCP/IP, NetBEUI, IPX, and AppleTalk. Additional statistics associated with TCP/IP are available only if the SNMP (Simple Network Management Protocol) service is installed.

Network traffic per segment data is collected by running the network interface card (NIC) in promiscuous mode, where the NIC notifies the host whenever a packet is received. Using a standard Ethernet hub, the network segment is logically a continuous loop of wire, where all stations see all packets. It is only necessary to gather statistics at one station to be able to summarize activity on the entire segment. On a switched network, each station sees only traffic specifically intended for it.

Due to their performance characteristics, Ethernet networks can degrade very rapidly when multiple stations contend for access to the shared wire. If two stations attempt to access the wire at the same time, a collision occurs and both stations need to retransmit. Characteristically, utilization tends to increase very quickly once collisions start to occur. Unfortunately, statistics for the number of collisions are only available by running the full Network Monitor. However, Windows NT makes it very easy to gather statistics utilization.

Network Utilization

The Performance Gallery Gold **Network Utilization by Segment (NT)** graph reports the specific network segment(s) with which the station is associated. Utilization is calculated by capturing the number of bytes received and bytes transmitted per segment and dividing by the rated capacity of the network interface card.

The scope of the network segment is defined by the networking hardware. On a standard Ethernet hub, all attached stations are part of the segment, and each station sees all packets. Using switched Ethernet, a station sees only packets specifically addressed to it, so each segment consists of a single station.

The **Network Utilization by Segment (NT)** graph has a reference line at 50% utilization (y=50). With Ethernet segments, be careful about collisions caused by multiple stations that need access to the wire concurrently, because Ethernet has no arbitration phase. Stations that need to transmit data simply wait for the wire to be free, then they place their packet on the wire. When two stations can attempt to do this at the same time, a collision occurs and both stations must retransmit. This causes a characteristic "bulge" in utilization, once the rate of traffic reaches 40% to 50% of capacity. Consequently, you would normally not want to see Ethernet network segment utilization running at a sustained level in excess of 40-50%.

This rule does not apply when the bulk of network traffic consists of a single session between one station and another, as in bulk file transfers or backup operations. For example, since the two stations performing a bulk file transfer are involved in a sort of "conversation" and must wait until the receiver acknowledges the receipt of messages, there is no contention for the wire. It is possible to driver Ethernet utilization to nearly 100% in these circumstances without incurring collisions.

From the **Network Utilization by Segment (NT)** graph you can link to more detailed network interface, error statistics, and TCP/IP statistics.

Network Interface Traffic

The Performance Gallery Gold **Network Interface Traffic (NT)** graph shows network traffic, broken down into bytes sent and bytes received at the network interface level. The data elements in the graph correspond to the Bytes Sent/sec counter and Bytes Received/sec counter in the Network Interface object.

System Activity

The Performance Gallery Gold **System Activity (NT)** graph reports the Context Switches/sec counter and the Total Interrupts/sec counter. These are two separate, but related, performance indicators. Both measures are *relative* indicators of system activity.

In contrast to most other Performance Gallery Gold pre-configured charts, this area graph does not show a stacked data view. Since interrupts cause context switches to occur, the two measurements are directly related. However, the number of context switches will always be greater than the number of interrupts, because some context switches are unrelated to the servicing of interrupts.

Context switching occurs whenever a running thread voluntarily relinquishes the processor, or when a running thread is preempted by a higher-priority ready thread following an interrupt. When the new thread executes in a different address space, the internal Intel processor Task State Segment (TSS) register must be reloaded and some internal processor caches may be invalidated. Context switches also occur when a user-mode thread calls an NT executive service or a Win32 subsystem service. In the latter case, a privileged kernel mode thread is assigned to service the request. A context switch occurs and is counted, but since the TSS does not have to be reloaded, the performance impact of this event is trivial.

NT device drivers service interrupts whenever peripherals need to notify the processor that some external event has occurred—usually that an I/O request has completed. Look for sudden unexplained changes in the interrupt rate that may be caused by a malfunctioning interface board or device.

When CPU utilization is running consistently at 100% with very few interrupts, check to see if a program is caught in an endless loop.

12

Windows NT/2000 Objects and Counters

Cache Object

The *Cache* performance object consists of counters that monitor the file system cache, an area of physical memory that stores recently used data as long as possible to permit access to the data without having to read from the disk. Because applications typically use the cache, the cache is monitored as an indicator of application I/O operations. When memory is plentiful, the cache can grow, but when memory is scarce, the cache can become too small to be effective.

Cache Object Counters

Copy Reads/sec

Copy Reads/sec is the frequency of reads from pages of the file system cache that involve a memory copy of the data from the cache to the application's buffer. The LAN Redirector uses this method for retrieving information from the file system cache, as does the LAN Server for small transfers. This is a method used by the disk file system as well.

Data Map Pins/sec

Data Map Pins/sec is the frequency of data maps in the file system cache that resulted in pinning a page in main memory, an action usually preparatory to writing to the file on disk. While pinned, a page's physical address in main memory and virtual address in the file system cache will not be altered.

Data Maps/sec

Data Maps/sec is the frequency that a file system such as NTFS, maps a page of a file into the file system cache to read the page.

Lazy Write Flushes/sec

Lazy Write Flushes/sec is the rate at which the Lazy Writer thread has written to disk. Lazy Writing is the process of updating the disk after the page has been changed in memory, so that the application that changed the file does not have to wait for the disk write to be complete before proceeding. More than one page can be transferred by each write operation.

Lazy Write Pages/sec

Lazy Write Pages/sec is the rate at which the Lazy Writer thread has written to disk. Lazy Writing is the process of updating the disk after the page has been changed in memory, so that the application that changed the file does not have to wait for the disk write to be complete before proceeding. More than one page can be transferred on a single disk write operation.

MDL Reads/sec

MDL Reads/sec is the frequency of reads from the file system cache that use a Memory Descriptor List (MDL) to access the data. The MDL contains the physical address of each page involved in the transfer, and thus can employ a hardware Direct Memory Access (DMA) device to effect the copy. The LAN Server uses this method for large transfers out of the server.

Pin Reads/sec

Pin Reads/sec is the frequency of reading data into the file system cache preparatory to writing the data back to disk. Pages read in this fashion are pinned in memory at the completion of the read. While pinned, a page's physical address in the file system cache will not be altered.

Read Aheads/sec

Read Aheads/sec is the frequency of reads from the file system cache in which the cache detects sequential access to a file. The read aheads permit the data to be transferred in larger blocks than those being requested by the application, reducing the overhead per access.

Logical Disk Performance Object

The *Logical Disk* performance object consists of counters that monitor logical partitions of hard or fixed disk drives. Performance Monitor identifies logical disks by their drive letter, such as C.

Logical Disk Object Counters

% Disk Read Time

% Disk Read Time is the percentage of elapsed time that the selected disk drive is busy servicing read requests.

% Disk Time

% Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.

% Disk Write Time

% Disk Write Time is the percentage of elapsed time that the selected disk drive is busy servicing write requests.

% Free Space

% Free Space is the ratio of the free space available on the logical disk unit to the total usable space provided by the selected logical disk drive.

Avg. Disk Queue Length

Avg. Disk Queue Length is the average number of both read and write requests that were queued for the selected disk during the sample interval.

Avg. Disk sec/Read

Avg. Disk sec/Read is the average time in seconds of a read of data from the disk.

Avg. Disk sec/Transfer

Avg. Disk sec/Transfer is the time in seconds of the average disk transfer.

Avg. Disk sec/Write

Avg. Disk sec/Write is the average time in seconds of a write of data to the disk.

Current Disk Queue Length

Current Disk Queue Length is the number of requests outstanding on the disk at the time the performance data is collected. It includes requests in service at the time of the snapshot. This is an instantaneous length, not an average over the time interval.

Multi-spindle disk devices can have multiple requests active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests are experiencing delays proportional to the length of this queue minus the number of spindles on the disks. This difference should average less than 2 for good performance.

Disk Read Bytes/sec

Disk Read Bytes/sec is the rate bytes are transferred from the disk during read operations.

Disk Reads/sec

Disk Reads/sec is the rate of read operations on the disk.

Disk Transfers/sec

Disk Transfers/sec is the rate of read and write operations on the disk.

Disk Write Bytes/sec

Disk Write Bytes/sec is the rate bytes are transferred to the disk during write operations.

Disk Writes/sec

Disk Writes/sec is the rate of write operations on the disk.

Free Megabytes

Free Megabytes displays the unallocated space on the disk drive in megabytes. One megabyte = 1,048,576 bytes.

Memory Object

The *Memory* performance object consists of counters that describe the behavior of physical and virtual memory on the computer. Physical memory is the amount of random access memory on the computer. Virtual memory consists of space in physical memory and on disk. Many of the memory counters monitor paging, which is the movement of pages of code and data between disk and physical memory. Excessive paging, a symptom of a memory shortage, can cause delays which interfere with all system processes.

Memory Object Counters

% Committed Bytes in Use

% Committed Bytes in Use is the ratio of Memory: Committed Bytes to the Memory: Commit Limit. (Committed memory is physical memory in use for which space has been reserved in the paging file, should it be written to disk. The commit limit is determined by the size of the paging file. If the paging file is enlarged, the commit limit increases, and the ratio is reduced.)

Available Bytes

Available Bytes is the amount of physical memory available to processes running on the computer, in bytes. It is calculated by summing space on the Zeroed, Free, and Standby memory Lists. Free memory is ready for use; Zeroed memory is pages of memory filled with zeros to prevent later processes from seeing data used by a previous process. Standby memory is memory removed from a process's working sets (its physical memory) on route to disk, but is still available to be recalled.

Pool Non-paged Bytes

Pool Non-paged Bytes is the number of bytes in the non-paged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated. Memory: Pool Non-paged Bytes is calculated differently than Process: Pool Non-paged Bytes, so it might not equal Process: Pool Non-paged Bytes:_Total.

Pool Paged Bytes

Pool Paged Bytes is the number of bytes in the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used. Memory: Pool Paged Bytes is calculated differently than Process: Pool Paged Bytes, so it might not equal Process: Pool Paged Bytes:_Total.

System Cache Resident Bytes

System Cache Resident Bytes is the number of bytes of pageable operating system code in the file system cache. This value is a component of Memory: System Code Resident Bytes which represents all pageable operating system code that is currently in physical memory.

System Code Resident Bytes

System Code Resident Bytes is the number of bytes of operating system code currently in physical memory that can be written to disk when not in use. This value is a component of System Code Total Bytes, which also includes operating system code on disk. System Code Resident Bytes (and System Code Total Bytes) does not include code that must remain in physical memory and cannot be written to disk.

System Driver Resident Bytes

System Driver Resident Bytes is the number of bytes of pageable physical memory being used by device drivers. It is the working set (physical memory area) of the drivers. This value is a component of Memory: System Driver Total Bytes, which also includes driver memory that has been written to disk. Neither System Driver Resident Bytes nor System Driver Total Bytes includes memory that cannot be written to disk.

Transition Faults/sec

Transition Faults/sec is the number of page faults resolved by recovering pages that were being used by another process sharing the page, or were on the modified page list or the standby list, or were being written to disk at the time of the page fault. The pages were recovered without additional disk activity. Transition faults are counted in numbers of faults, but because only one page is faulted in each operation, it is also equal to the number of pages faulted.

Network Interface Object

The *Network Interface* performance object consists of counters that measure the rates at which bytes and packets are sent and received over a TCP/IP connection. It includes counters that monitor connection errors.

Network Interface Object Counters

Bytes Received/sec

Bytes Received/sec is the rate at which bytes are received on the interface, including framing characters.

Bytes Sent/sec

Bytes Sent/sec is the rate at which bytes are sent on the interface, including framing characters.

Processor Object

The *Processor* performance object consists of counters that measure aspects of process activity. The processor is the part of the computer that performs arithmetic and logical computations, initiates operations on peripherals, and runs the threads of processes. A computer can have multiple processors. The processor object represents each processor as an instance of the object.

Processor Object Counters

% DPC Time

% DPC Time is the percentage of time that the processor spent receiving and servicing deferred procedure calls (DPC's) during the sampling interval. (DPC's are interrupts that run at a lower priority than standard interrupts.) % DPC Time is a component of % Privileged Time because DPC's are executed in privileged mode. They are counted separately and are not a component of the interrupt counters.

% Interrupt Time

% Interrupt Time is the percentage of time the processor spent receiving and servicing hardware interrupts during the sampling interval. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards, and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended during interrupts. Most system clocks interrupt the processor every 10 milliseconds, creating a background of interrupt activity.

% Privileged Time

% Privileged Time is the percentage of non-idle processor time spent in privileged mode. (Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and all memory. The alternative, user mode, is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The operating system switches application threads to

privileged mode to obtain operating system services.) % Privileged Time includes time servicing interrupts and DPC's. A high rate of privileged time might be attributable to a large number of interrupts generated by a failing device.

% User Time

% User Time is the percentage of non-idle processor time spent in user mode.

User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems.

The alternative, privileged mode, is designed for operating system components and allows direct access to hardware and all memory. The operating system switches application threads to privileged mode to obtain operating system services.

Working Set

Working Set is the current number of bytes in the working set of this process. The working set is the set of memory pages touched recently by the threads in the process. If memory in the counter is above a threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from working sets. If they are needed, they will then be soft-faulted back into the working set before they leave main memory.

Redirector Object

The *Redirector* performance object consists of counters that monitor network connections originating at the local computer.

Redirector Object Counter

Bytes Total/sec

Bytes Total/sec is the rate the Redirector is processing data bytes. This includes all application and file data in addition to protocol information, such as packet headers.

Server Object

The *Server* performance object consists of counters that measure communication between the local computer and the network.

Server Object Counters

Bytes Received/sec

The number of bytes the server has received from the network. Indicates how busy the server is.

Bytes Transmitted/sec

The number of bytes the server has sent on the network. Indicates how busy the server is.

System Object

The *System* performance object consists of those counters that apply to more than one component of the computer.

System Performance Object Counters

Context Switches/sec

Context Switches/sec is the combined rate at which all processors on the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service. It is the sum of Thread: Context Switches/sec for all threads running on all processors in the computer and is measured in numbers of switches. There are context switch counters on the System and Thread objects.

A

PROGRAM MESSAGES

Table A.1 *Program Messeges*

Message	Category	Description
Cannot add entries to a string!	Template Notification Message	Chart data character strings (character data) cannot have other items added to them.
Cannot add string entries in a calculation!	Template Notification Message	Chart data character strings cannot be used in mathematical calculations.
Cannot add string entries to the divisor!	Template Notification Message	Chart data character strings cannot be used in mathematical calculations.
Chart entry full! Only 5 items allowed.	Template Notification Message	Performance Gallery Gold allows up to five entries in either the numerator or the divisor fields of a chart data entry.
Critical error in the current entry's structure!	Template Memory Error	An invalid memory error was encountered that had no other error message assigned to it.
Error - file save should not be executed!	File Save Error	This message should not be encountered as data files in Performance Gallery Gold can only be loaded.

Message	Category	Description
Error defining groups! Groups defined out of order.	Data File Load Error	The data description (header) area of the data file contains blocks of information separated by section headers and delimited by end statements. This message occurs when a statement is either incorrect or encountered out of order. The header section is either corrupt or it did not convert correctly when the file converter was used.
Error defining instances!	Data File Load Error	The data description (header) area of the data file contains blocks of information separated by section headers and delimited by end statements. This message occurs when a statement is either incorrect or encountered out of order. The header section is either corrupt or it did not convert correctly when the file converter was used.
Error in conversion - attempt failed!	File Conversion Error	An unexpected error occurred that could not be categorized. The conversion will abort and the file will not be loaded.
Error loading itemlist - file not present or corrupt!	File Conversion Error	The itemlist file for the selected data file could not be found in the Performance Gallery Gold directory. This file is required for the conversion to succeed.
Error loading template definition file!	File I/O Error	An unexpected error occurred that could not be categorized. The template load will abort.

Message	Category	Description
Error loading the default desktop!	File Load Error	An error occurred while loading the desktop file that cannot be attributed to the operating system and has no other error-specific message. This message is usually encountered with corrupted desktop files or ones that have been transferred between machines using the ftp protocol set to ASCII and not binary code.
Error loading the selected desktop!	File Not Found	The desktop file that was specified does not exist.
Error writing export file to disk!	File I/O Error	An unexpected error occurred that could not be categorized. The export will abort and the file will not be saved.
Error: File Save attempted on a read only file!	File I/O Error	The converter was called on a file save rather than a file load basis. This error should not occur as the converter is called via the data file open command.
File is not a Performance Gallery data file!	Data File Load Error	The identification information in the file that was specified does not match the Performance Gallery Gold series format.
File rename after conversion failed!	Post Conversion Error	The converted file is initially stored with a .pf2 extension and when the conversion is complete, the original file is renamed with a .pfb extension and the converted file is renamed with a .pfg extension. If a .pfb file with the same name already exists in the target directory, or if rename rights are not available on the original file (Windows NT/2000 only), the rename attempt will fail.

Message	Category	Description
File rename after template save failed!	Post Save File Error	The template file is initially stored with a temporary extension. When the operation is complete, the original file is moved out of the way, the new file is put in its place, and the original file is deleted. If rename rights are not available on the original file (Windows NT/2000 only), the rename attempt will fail.
Invalid extension! Unable to export file.	File I/O Error	In the file type pull-down menu for the export graph, there are several available file formats. If the format that is returned by the file Save dialog is not one of the formats supported by Performance Gallery Gold, this error will occur. This error should not be encountered.
Invalid file type. Unable to load desktop file!	Desktop Load Error	The file that was specified is either corrupt or not a Performance Gallery Gold desktop file.
Invalid height or width entry!	Graph Export Error	This error message is displayed when an invalid graph size is entered from the Export dialog box. This error should not occur, because the Export dialog accepts only valid entries regarding graph export size.
Invalid new template name!	Template Save Error	This message is usually encountered when the template name has been left blank or when invalid characters (such as unprintable characters, the "\" character, or other command characters) are included in the template name.

Message	Category	Description
Invalid section definition in the data file!	Data File Load Error	The data description (header) area of the data file contains blocks of information separated by section headers and delimited by end statements. If a statement is encountered out of order or is incorrect, these messages occur. The header section is either corrupt or did not convert correctly when the file converter was used.
Master data table memory allocation failed! Out of memory. Try extracting less data in MVLOGX or increasing your page file size.	File Load Aborted	There was not enough memory to create the master data table that contains all of the file information. Either more memory (RAM or page file) must be added, or less data must be loaded.
Master data table memory allocation failed! Out of memory. Try increasing your page file size or loading fewer data files.	File Load Aborted	There was not enough memory to create the master data table that contains all of the file information. Either more memory (RAM or page file) must be added or less data must be loaded.
No Type 5 records present!	File Load Failed	The Type 5 records are required by Performance Gallery Gold so the data can be properly indexed and interpreted.

Message	Category	Description
Not all charts could be found!	Desktop Load Error	The desktop contains a list of charts that are based on a template. User-defined templates are stored in a file named tmplates.usr. This file is unique to the machine that it is created on and it is not stored in the desktop. If the desktop is moved between machines or installations of Performance Gallery Gold, the user templates may not be present and Performance Gallery Gold will not be able to load all of the charts.
Only files from the same system may be concurrently loaded!	File Load Aborted	When the multiple file load feature is used, only files from the same host machine can be loaded together. If the Add to Current Data check box in the file Open dialog box is checked and files from different machines are selected to be loaded together, this message will occur.
Only one individual'All' type allowed per entry!	Template Notification Message	Performance Gallery Gold restricts the number of "All" groups displayed individually to one per entry.
Process information present but not loaded! Process ID field not present.	File Load Warning	In order for process information to be loaded, it must have a unique identifier. The identifier that Performance Gallery Gold uses is a combination of the process name and process ID. Both fields must be present for the process data to be loaded.
Source is not a perfgal D.01 data file!	File Conversion Error	The identification information in the file that was specified does not match the Performance Gallery D.01 series format.

Message	Category	Description
Template file version higher than expected!	File I/O Error	The template file that was loaded contains a version number that is higher than the current version is allowed to load. This generally means that a version of Performance Gallery Gold that is newer than the one reporting this error message created the template file.
The assigned alias already exists! Unable to update the alias entry.	Alias Manager error	The alias entered into the System Alias Manager dialog box has been used before. Assign a new name to the system.
The demo code has expired on this program!	Security Check Failed	Demonstration license codes for Performance Gallery Gold operate for a period up to 99 days. After this period has passed, a new code is required.
The file contains data that is too recent to convert! The license does not allow converting that is this recent.	File Conversion Aborted	This message occurs only on systems that have the licensed version of Performance Gallery Gold. This is due to the fact that the converter will only convert data files containing data that is older than 90 days after the license issue date.
The license code is not present or valid!	Security Check Failed	The license code that was entered was incorrect. The most common cause of this message is that the company name, which is encoded in the license, was entered differently from the one displayed on the paper or message containing the code. Both the company name and the license code must match exactly with those received.
Unable to add new chart element definition!	Template Memory Error	Not enough memory was available to store the information to the template.

Message	Category	Description
Unable to add new chart variable!	Template Memory Error	Not enough memory was available to store the information to the template.
Unable to add new template object!	Memory Error	Not enough memory was available to add a new template in the template list.
Unable to add sample data!	Memory Error	Not enough memory was available to load the data sample into memory.
Unable to allocate new group!	Data File Memory Error	The memory that Performance Gallery Gold requested is not available or the PC is out of memory.
Unable to allocate new instance!	Data File Memory Error	The memory that Performance Gallery Gold requested is not available or the PC is out of memory.
Unable to allocate static information!	Data File Memory Error	The memory that Performance Gallery Gold requested is not available or the PC is out of memory.
Unable to allocate version information!	Data File Memory Error	Not enough memory was available to store the information to the data table. The memory that Performance Gallery Gold requested is not available or the PC is out of memory.
Unable to assign alias!	Data File Memory Error	The memory that Performance Gallery Gold requested is not available or the PC is out of memory.
Unable to assign sub group!	Data File Memory Error	The memory that Performance Gallery Gold requested is not available or the PC is out of memory.

Message	Category	Description
Unable to change chart name - duplicate found!	Template Modify Error	Template names must be unique.
Unable to change chart name - name is preset!	Template Modify Error	The chart templates that are included with Performance Gallery Gold cannot be modified and then saved with the same name. Modified templates must be saved under a different name. A chart already exists with the same name.
Unable to find group matching with element data!	Master Data Error	A memory error occurred when searching through the list of data descriptors. The master index to the data has been corrupted or encountered a memory error.
Unable to initialize the template dialog object! Internet Explorer 3.x or higher may not be present. The application must exit.	Property Sheet Initialization	Microsoft Internet Explorer 3.x or higher includes an updated comctl32.dll file that is required by many programs, including Performance Gallery Gold. If the updated file is not present, tabbed dialog boxes will not be available and will cause the Performance Gallery Gold program to exit.
Unable to load template information!	Template Object Error	The template was not found in memory or was corrupted and could not be retrieved.
Unable to locate discovery records!	File Load Failed	Discovery records were not enabled in the data file. These are required to describe the data so Performance Gallery Gold can correctly index the information.

Message	Category	Description
Unable to open new graph frame window! The application must exit.	Property Sheet Initialization	There was not enough memory available to create the primary display for Performance Gallery Gold. This message is rarely encountered, and it typically occurs on systems with an upgrade edition of Windows 95. The two common causes are that Windows has not been configured correctly or the system did not have a memory optimization before the installation of Windows.
Unable to read data line.	File Load Error	An error occurred while attempting to read a line of data. The line is either too long for Performance Gallery Gold to read, or the data contains unexpected characters due to a newer file format or file corruption.
Unable to read sample data/ time stamp!	Data Error/Data File Load Error/ File Conversion Error	At least one block of data contains a time value that is not valid.
Unable to save desktop file!	Desktop Save Error	An error occurred while saving the desktop file that cannot be attributed to the operating system and has no other error-specific message.
Unable to update template items!	Template Data Error	An invalid memory error occurred when attempted to manipulate the memory assigned to a template.

Message	Category	Description
Unexpected END statement in the data file!	Data File Load Error	The data description (header) area of the data file contains blocks of information separated by section headers and delimited by end statements. If a statement is encountered out of order or is incorrect, these messages occur. The header section is either corrupt or did not convert correctly if the file converter was used.
Unexpected section in data file! End statement expected.	Data File Load Error	The data description (header) area of the data file contains blocks of information separated by section headers and delimited by end statements. If a statement is encountered out of order or is incorrect, these messages occur. The header section is either corrupt or did not convert correctly if the file converter was used.
Could not connect to pwscomm.dll	Continuous Update	The pwscomm.dll file provides PGG with the ability to talk to the remote systems. If it is not present, no connection attempts will succeed.
Could not parse remote system time data	Continuous Update	The data provided by the remote system must be broken down into its individual items. This is done by the Microsoft XML Parser. This message indicates an error in the parsing that may be due to data corruption or a problem with the XML parser.

Message	Category	Description
Unable to connect to remote system.	Continuous Update	The remote system was present but did not accept the attempt to retrieve data. The remote system may not have the data providing service (PWS) installed or configured correctly. Additionally, there may be a networking problem between the machine running PGG and the remote system that is preventing this connection.
Unable to retrieve latest data sample time	Continuous Update	All data provided by the remote system should contain information detailing the time that it was collected under. If the time cannot be found, the new data cannot be added to the data currently loaded in PGG.

B

KEYBOARD COMMANDS

File Menu Commands

Table B.1 *File Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Open Data File		Alt + F + O	Ctrl + O
Close Data File		Alt + F + C	Ctrl + F
Close All Data		Alt + F + A	Ctrl + H
System Alias Manager		Alt + F + S	
Current File Statistics		Alt + F + T	Ctrl + T
Print Current Chart		Alt + F + H	
Print Preview		Alt + F + V	
Print All		Alt + F + P	Ctrl + P
Print Setup		Alt + F + R	
1 Graph per Page		Alt + F + G + 1	
2 Graphs per Page		Alt + F + G + 2	
4 Graphs per Page		Alt + F + G + 4	
Exit		Alt + F + X	

Edit Menu Commands

Table B.2 *Edit Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Copy		Alt + E + C	Ctrl + C
Copy Special		Alt + E + S	
Clear Clipboard		Alt + E + E	Ctrl + E
Export		Alt + E + R	Ctrl + R

View Menu Commands

Table B.3 *View Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Toolbar		Alt + V + T	
Status Bar		Alt + V + S	
Exception Window		Alt + V + E	
Zoom		Alt + V + Z	
Unzoom		Alt + V + U	
View Options		Alt + V + O	

Desktop Menu Commands

Table B.4 *Desktop Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Load Desktop	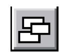	Alt + D + L	Ctrl + L
Reload Desktop		Alt + D + R	Ctrl + K
Close Desktop		Alt + D + D	Ctrl + D
Save Desktop		Alt + D + S	Ctrl + S
View the last used desktop files		Alt + D + the underlined file number to the left of the desktop file name	
Set Default Desktop		Alt + D + F	
Default Desktop		Alt + D + K	

Template Menu Commands

Table B.5 *Template Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Template Properties		Alt + T	
Template Fonts		Alt + T + F	
Title Font		Alt + T + F + T	
Legend Font		Alt + T + F + L	
Index Font		Alt + T + F + I	
Template Colors		Alt + T + C	
Backdrop Color		Alt + T + C + B	
Graph Key Color		Alt + T + C + K	

Command	Toolbar Button	Keyboard Commands	
Data Field Color		Alt + T + C + D	
Auto Scaling		Alt + T + A	
Save Templates		Alt + T + S	

Chart Menu Commands

Table B.6 *Chart Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Open Chart		Alt + C + O	Ctrl + N
Modify Chart		Alt + C + M	Ctrl + M
Close Chart		Alt + C + E	Ctrl + X
Select Items		Alt + C + S	Ctrl + I
Select Secondary Items		Alt + C + I	Ctrl + Y
Global Chart Options	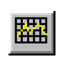	Alt + C + G	Ctrl + G

Window Menu Commands

Table B.7 *Window Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Cascade		Alt + W + C	
Tile Horizontally		Alt + W + T	
Arrange Icons		Alt + W + A	

Command	Toolbar Button	Keyboard Commands	
Refresh		Alt + W + R	
View Open Charts		Alt + W + the underlined file number to the left of the file name	

Help Menu Commands

Table B.8 *Help Menu Commands*

Command	Toolbar Button	Keyboard Commands	
Help Topics		Alt + H + H	
About Performance Gallery Gold		Alt + H + A	

C

Commands to Maneuver Charts

RECOMMENDATION If you have a two-button mouse, or if your three-button mouse is not configured to use the middle button, press both the left and right buttons simultaneously to simulate a middle button.

Maneuvering 3D Graphs

Table C.1 *Maneuvering 3D Graphs Commands*

Action	Result
Press and hold the X key while dragging the 3-D graph using the middle mouse button.	Rotates the graph on its x axis.
Press and hold the Y key while dragging the 3-D graph using the middle mouse button.	Rotates the graph on its y axis.
Press and hold the Z key while dragging the 3-D graph using the middle mouse button.	Rotates the graph on its z axis.
Drag the 3-D graph using the middle mouse button.	Rotates the graph on all axes.
Press and release the T key.	Returns the graph to its original twist (rotation).
Press and hold the SHIFT key while dragging the graph using the middle mouse button.	Repositions the graph within the program window.

Action	Result
Press and hold the Ctrl key while using the middle mouse button to drag the image outward, toward the edge of the program window.	Zooms in.
Press and hold the Ctrl key while using the middle mouse button to drag the image inward, toward the center of the program window.	Zooms out.
Press and hold the Ctrl key and use the left mouse button to select the area to enlarge, then release the left mouse button.	Enlarges the specified portion of the graph.

Maneuvering 2D Charts

Table C.2 *Maneuvering 2D Graphs Commands*

Action	Result
Press and hold the SHIFT key while dragging the graph using the middle mouse button.	Repositions the graph within the program window.
Press and hold the Ctrl key while using the middle mouse button to drag the image outward, toward the edge of the program window.	Zooms in.
Press and hold the Ctrl key while using the middle mouse button to drag the image inward, toward the center of the program window.	Zooms out.
Press and hold the SHIFT key and use the left mouse button to select the rectangular area to enlarge. Release the left mouse button.	Enlarges the specified portion of the graph.
Press and release the R key.	Returns the graph to its original size and position.

D

Automating Performance Gallery Gold

Performance Gallery Gold (PGG) contains a flexible, command-line driven automation facility. These capabilities are most commonly accessed in two ways. The first is by customizing a program shortcut to run PGG with command-line parameters on an as-needed basis. The second is through a batch file that may be scheduled to run at given time(s) using the built-in scheduling facilities provided with Windows 95 (with the Plus Pack) or Windows 98/ME/NT/2000.

Customizing a Shortcut

The following steps may be used to customize a Performance Gallery Gold shortcut:

1. Right-click on the Performance Gallery Gold shortcut and select **Properties**. A dialog window will pop up showing the contents of the shortcut.
2. There will be two tabs at the top of the dialog, **General** and **Shortcut**. Make sure that the **Shortcut** tab is the one that is currently being displayed. If it is not, left-click on the **Shortcut** tab.
3. The **Target** field should be highlighted. Left-click in the field and move to the end of the text. Insert a space (all commands should be separated by spaces) and type in the desired automation commands (see "Automation Commands.", below).
4. When all of the commands are entered, click **OK**.

The next time the shortcut is used, the automation commands will be executed.

Batch Files

Batch files are text files that end with the .bat extension and contain commands that may be issued without user interaction. They may be scheduled (see "Scheduling for the Batch Files.") to execute at given times and are very flexible. For simplicity, this explanation will only cover the basics of using batch files to execute Performance Gallery Gold (PGG) in automated mode, as well as adding an FTP command to the batch file to transfer files automatically to another machine (usually a PC or server containing a web page that references the text and images exported by PGG).

The following steps may be used to create a batch file for Performance Gallery Gold automation:

1. Start **Notepad** (under **Start/Programs/Accessories/Notepad**) or any text editor and open a new text file (Notepad will start with a new text file automatically).
2. Enter the path and executable name of the currently installed copy of PGG. This parameter must be enclosed in quotation marks if it contains any spaces or special characters other than the colon and reverse slash. This file will only have one line of text. Do not insert a hard return. Refer to “Example Automation Settings” on page 290.
3. Insert a space (all commands should be separated by spaces) and type in the desired automation commands (see "Automation Commands.", below). Commands may be inserted in any order. Refer to “Example Automation Settings” on page 290.
4. Save the file with the .bat extension.

FTP in Batch Files

File Transfer Protocol (FTP) is commonly used to transfer files between computers on the Internet. A command line version of FTP is included with Windows 95/98/ME/NT/2000. As a command line executable, FTP may be included in a batch file to be run either by itself or in combination with Performance Gallery Gold.

This capability can be used to transfer either the data file to the Windows machine running Performance Gallery Gold (PGG) or to transfer the resulting images and tables produced by PGG to a desired web server.

NOTE The following instructions assume that the machines that will be conducting the transfer are properly configured for TCP/IP networking and that the user has a FTP-capable login on the source machine to the desired file(s) and that the source machine has an operational FTP server service running.

The following steps may be used to transfer a data file from a machine generating a data file to the machine that will process it using Performance Gallery Gold's automation facility:

1. Start **Notepad** (under **Start/Programs/Accessories/Notepad**) or any text editor and open a new text file (Notepad will start with a new text file automatically).
2. Enter the **cd** (change directory) command followed by a space and then the name of the directory that you want the file to placed in, as in the following example:

```
cd "C:\My Documents"
```

3. Press **Enter** to go to the next line.
4. Enter in the FTP command followed by the **-i** switch and the **-s:<filename>** switches. The **-i** switch disables prompting for the transfer of multiple files, otherwise user interaction would be required to perform multiple file transfers. The **-s** switch specifies a file that contains instructions for the FTP command - essentially a batch file specific to FTP. This will be the next file that we create. The following example displays the correct format:

```
ftp -i -s:"C:\My Documents\PGGftp.txt"
```

5. Save the file with the .bat extension.
6. Close the file and open a new, blank one. This file will contain the instructions to give the FTP command in order to transfer the file.
7. Enter in the open command followed by the name of the machine to get the data file from. Normally, it is good form to use the full name of the machine, including the domain name in this parameter.

   ```
   open ftp.example.com
   ```

8. Press **Enter** to go to the next line.
9. Next type the name of the user that will be logging in to the remote machine.
10. Press **Enter** to go to the next line.
11. Type the password of the user mentioned above.
12. Press **Enter** to go to the next line.
13. Type **bin** to set the transfer mode to binary (this not required if the default setting for the FTP client is binary or if the data file is an .smf text file).
14. Press **Enter** to go to the next line.
15. Type the **get** command with two parameters in the following format:

   ```
   get datafile.pfg mydatafile.pfg
   ```

   The first part is the name of the file to get from the source machine. The second is the name that will be given to the file when it is saved on the local machine. Wildcard characters are allowed, but the first parameter cannot match more than one file on the source machine or an error will be generated.
16. Press **Enter** to go to the next line.
17. Type **bye**, which tells FTP to disconnect from the remote machine and exit FTP.
18. Finally, save the FTP instruction file with the name that you gave it in the original batch file (PGGftp.txt in the above example). Be sure that the original batch file is either located in the same directory as this file, or have the original batch file list the complete path to the FTP instruction file in its call to the FTP program.

The instructions above illustrate how to execute a FTP file "pull" running on the client machine using a set of batch files. It is also possible to execute a file "push" where the server runs the batch files and sends the data to the client machine. The client machine must have an FTP server for this to work, however.

Scheduling for the Batch Files

This section will cover scheduling of the automation batch file(s) under Windows NT/2000. Scheduling is also possible using the **Task Scheduler** present in Windows 95 Plus! Pack and Windows 98, but it will not be covered in this document.

Windows NT Scheduling

Scheduling is handled by the Scheduler service in Windows NT. For any scheduled task to execute, this service must be running.

Verifying Services

To verify that the service is running and configured correctly, take the following steps:

1. Open the **Control Panel** under **Start/Settings**.
2. Double-click on the **Services** icon to open its dialog window.
3. Scroll down to the **Schedule** service in the list view window of the dialog.
4. If the service is stopped, click on the service to highlight it and then click on the **Start** button.
5. If the service is not set to start automatically (it may be set at manual or disabled), click on the service to highlight it and click on the **Startup...** button. A dialog will pop up containing the startup settings of the dialog.
   a. Select a startup type of **Automatic**.
   b. Click on the **OK** button to close the service properties dialog.
6. Close the **Services** dialog.

NOTE The user must have administrator rights on the machine to perform changes to the services settings in Windows NT.

Scheduling Tasks

The following steps may be used to schedule a task in Windows NT:

1. Open a command line window under **Start/Programs/Command** Prompt.
2. Type in the desired task in the format below:

```
at <time> /every:<day of the week> <command>
```

- The time should be in an hh:mm format.
- The day of the week should be as follows:

  M - Monday

  Tu - Tuesday

  W - Wednesday

  Th - Thursday

  F - Friday

Sa - Saturday

Su - Sunday

- Multiple days of the week can be specified by using commas (i.e. M,Tu,W).
- The command field should contain the full name and path of the file to be scheduled.

For a Windows dialog-based way to schedule a task in Windows NT, the WinAT program contained in the Windows NT Resource Kit may be used. WinAT also allows the editing and removal of individual commands from the scheduler, where as the AT command does not. For more information on the WinAT program, see the Windows NT Resource Kit documentation.

To delete all currently scheduled tasks, type **at /delete** at the command prompt.

Windows 2000 Scheduling

Scheduling is handled by the Task Scheduler service in Windows 2000. For any scheduled task to execute, this service must be running.

Verifying Services

To verify that the service is running and configured correctly, take the following steps:

1. Open the **Control Panel** under **Start/Settings**.
2. Open the **Administrative Tools** window by double- clicking on its icon.
3. Open the **Computer Management** dialog by double-clicking on its icon.
4. Double-click on the **Services and Applications** entry in the left pane of the dialog window to expand the selection and show the entries under it.
5. Click on the **Services** entry in the left pane of the dialog window to load the services list into the right pane of the dialog window.
6. Scroll down to look at the **Task Scheduler** service listed in the right pane of the dialog window.
7. If the service is stopped, single-click on the service to highlight it and then click on the **Start** item in the **Action** pull-down menu.
8. If the service is not set to start automatically (it may be set at manual or disabled), single-click on the service to highlight it and click on **Properties** in the **Action** pull-down menu.
   a. A dialog will pop up containing the settings for the service.
   b. The dialog should start off in the **General** tab. If it does not, click on **General** to view its options.
   c. In the **General** tab, pull down the **Startup** type list box and select **Automatic**.
   d. Click on **OK** to close the dialog.
9. Close the **Computer Management** window when the service has been properly configured.

NOTE The user must have administrator rights on the machine to perform changes to the services settings in Windows 2000.

Now the machine is ready to execute scheduled tasks.

Scheduling Tasks

The following steps may be used to schedule a task in Windows 2000:

1. Open the **Control Panel** under **Start/Settings**.
2. Open the **Scheduled Tasks** window by double clicking on its icon.
3. Double-click on **Add Scheduled Task** to start the wizard for adding new tasks to the scheduler. The wizard dialog will pop up.
4. In the wizard, click on **Next** after you have read the first page containing an introduction. The second page of the wizard will pop up and ask for the task that is to be scheduled.
5. To schedule a batch file (like those used for Performance Gallery Gold automation), click on **Browse** to pop up the **Select Program To Schedule** dialog.
   a. Browse the file system to select the desired file in the **Select Program to Schedule** dialog.
   b. Once the file has been selected, click on **Open** to finish selecting the file.
6. Now select how often the file should be executed using the wizard. Do not be concerned that this page only offers day, week, and month options for its settings. The next dialog will ask for specific information like which days of the week, at what times, and so on.
7. Click on **Next** to go to the page of the wizard.
8. Select the specific days and times for the task to be run.
9. Click on **Next** to go to the login page of the wizard.

   The login page of the wizard will ask for a login and password for the scheduled command to use. This is for security purposes and to allow automated tasks to run as a different user that may be monitored and audited without being confused with a user or other process. Typically, however, the current user's login and password are used.
10. Click on **Next** to go to the confirmation page.Now the final dialog will pop up. It confirms the basic information of the task that is about to be scheduled.
11. Click on **Finish** if the information is correct or use the **Back** button to correct it or the **Cancel** button to exit without scheduling the task.

Automation Commands

Under automation, Performance Gallery Gold will first evaluate the entire command line, then begin loading the given data file (if any). Next, it will set the various non-file related options, load the specified desktop or default desktop (if any), and attempt to output all chart windows to disk using the chart's template name for the file name. Graphs will be stored as .jpg, and tables will be stored as .htm files. If the batch command has been issued, the program will then exit.

The automation commands for Performance Gallery Gold (PGG) are issued via the command line upon program execution. The following commands are currently available:

<data file>

List all the data files you want loaded. Up to 20 data files may be specified. If multiple files are specified, there must be a space between the filenames. The data file parameter must be enclosed in quotation marks if it contains any spaces or special characters other than the colon and reverse slash.

NOTE Any of the following parameters following the program executable name may contain a path and file name. **If a path and file name is specified, then it must be enclosed by quotation marks if it contains any spaces or special characters other than the reverse slash!**

Table D.1 *Automation Commands*

Parameter	Description
-batch	This option tells the program to exit after the automation command completes.
-d:<desktop file>	This parameter specifies the desktop file for PGG to use upon start up. If no desktop file is specified and data file(s) have been listed in the automation command, PGG will load the default desktop. If no desktop file is specified and no default desktop has been selected, no charts will be exported by the automation command.
-e:<export directory>	This item specifies the directory to output the exported charts to. This directory must exist and the current user must have write privileges to it.
-h:<heading>	This item specifies a setting for the global heading that is displayed at the top of all graphs. The heading parameter must be enclosed in quotation marks if the heading contains any spaces or other non-alphanumeric characters.

Parameter	Description
-notoverlap	This item directs PGG to load multiple data files with the non-overlapping option on (see “New Data Overwrites Old” on page 78).
-overlap	This item directs PGG to load multiple data files with the non-overlapping option off (see “New Data Overwrites Old” on page 78).
-r:<export resolution>	This item sets the graph export resolution in x and y pixel values. For example: **-r:640x480** The separator between the two pixel values should be a lower-case x.
-x:<X axis resolution>	This item sets the Global Chart Options x axis resolution item using these possible values: -x:sample -x:hour -x:day

Example Automation Settings

Example 1:

```
"C:\Program Files\Lund Performance Solutions\Performance Gallery
Gold\perf.exe" data1.pfg data2.pfg -notoverlap -d:"c:\temp\auto.dsk" -batch
-r:640x480 -e:"C:\Temporary Files"
```

The above command runs the copy of Performance Gallery Gold (**perf.exe**) that is present in the **C:\Program Files\Lund Performance Solutions\Performance Gallery Gold** folder.

The command directs PGG to load two data files, **data1.pfg** and **data2.pfg**, that are contained in the same directory as the one that the command is run in.The data file names are not enclosed in quotation marks, as they do not contain any spaces or special characters other than the colon and reverse slash.

Then the **-notoverlap** command tells PGG to not allow any overlapping times between files; so any data in the data2.pfg file that contains overlapping information will have its data used for the overlapping time period (files are listed in the order that they will be loaded into memory).

Then the **-d** command tells PGG to load a desktop file called auto.dsk located in the C:\temp directory. The desktop file name is enclosed in quotes.

Next, the **-batch** command tells PGG to exit after completing the automation commands.

After that, the **-r** command specifies that the chart images exported to disk will be 640 by 480 pixels in dimension.

Finally, the **-e** command specifies that the image files that are generated by the automation process will be placed in the C:\temporary files directory. The directory path is enclosed in quotation marks since the path contains a space.

Example 2:

```
perf.exe sample1.pfg -d:test.dsk -batch -r:800x600 -e:".\output"
```

The above command runs a copy of Performance Gallery Gold (**perf.exe**) that is either in the same directory as the one that the command line is run in, or is in the system path (see the Windows OS documentation for details).

The command then directs PGG to load the **sample1.pfg** data file contained in the same directory as the one that the command is run in.

Then the **-d** command directs PGG to load the desktop file called test.dsk. The desktop file is located in the same directory as the one the command line is run in. The desktop file name in not enclosed in quotation marks, as it does not contain any spaces or special characters other than the reverse slash. There is no path listed as it is located within the same directory as the one the command line is run in.

After that, the **-r** command specifies that the exported chart images will be 800 by 600 pixels in dimension.

Finally, the **-e** command specifies that the image and text files that are generated by the automation process will be placed in to an example folder titled "output." The directory path is enclosed in quotation marks because the path contains a special character other than the reverse slash.

INDEX

Symbols

A

B

C

D

E

F

G

H

J

L

M

O

P

R

X

Y

Z