

**EMC ControlCenter Navisphere
Analyzer
Version 6.X**

ADMINISTRATOR'S GUIDE

**P/N 069001158
REV A01**

Copyright © EMC Corporation 2002. All rights reserved.

Printed July, 2002

EMC believes the information in this publication is accurate as of its publication date. However, the information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication require an applicable software license.

Trademark Information

EMC², EMC, MOSAIC:2000, Symmetrix, CLARiiON, and Navisphere are registered trademarks and EMC Enterprise Storage, The Enterprise Storage Company, The EMC Effect, Connectrix, ControlCenter, EDM, SDMS, SRDF, Timefinder, PowerPath, InfoMover, FarPoint, EMC Enterprise Storage Network, EMC Enterprise Storage Specialist, EMC Storage Logic, Universal Data Tone, E-Infostructure, Celerra, Access Logix, SnapView, and MirrorView are trademarks of EMC Corporation.

All other trademarks mentioned herein are the property of their respective owners.

Preface	ix
Chapter 1 About the EMC Navisphere Analyzer	
About Analyzer.....	1-2
Navisphere 6.X Management Environments.....	1-3
Analyzer Dependencies	1-6
Chapter 2 Analyzing Performance	
Starting a Navisphere Manager or Analyzer Session.....	2-2
Menu Bar Options for Analyzer	2-2
Real-Time and Archive File Display	2-4
Analyzer Charts	2-6
Performance Survey Chart	2-8
Performance Summary Chart	2-11
Performance Detail Chart	2-17
IO Size Distribution Charts (LUN Only).....	2-23
Performance Characteristics.....	2-28
Zooming and Resizing Charts	2-34
Exporting and Printing Charts.....	2-37
Customizing Chart Display.....	2-38
Capturing Performance Data in Archive Files	2-45
Discovering LUN Device Names	2-46
Chapter 3 Understanding Performance Data	
Understanding Workload	3-2
Identifying a Performance Problem	3-5
Eliminating Bottlenecks	3-7

Working with the Storage-System Cache 3-9
Addressing Server Problems..... 3-13
Index.....i-1

2-1	Menu Bar Options for Analyzer	2-2
2-2	Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks	2-3
2-3	Performance Survey Menu Options	2-10
2-4	Performance Summary Chart Tree Root Node Right-Click Menu Options.....	2-13
2-5	Performance Summary Chart Tree Panel Right-Click Menu Options .	2-13
2-6	Performance Summary Chart Panel Right-Click Menu Options	2-13
2-7	Performance Detail Chart Tree Root Node Right-Click Menu Options	2-20
2-8	Performance Detail Chart Tree Right-Click Menu Options	2-20
2-9	Performance Detail Chart Right-Click Menu Options	2-21
2-10	IO Size Distribution Summary Chart Left Panel Right-Click Menu Option	2-25
2-11	IO Size Distribution Summary Chart Right-Click Menu Options	2-25
2-12	IO Size Detail Chart Left Panel Right-Click Menu Options	2-27
2-13	IO Size Distribution Detail Chart Right-Click Menu Options	2-27
2-14	Basic Performance Characteristics	2-28
2-15	Advanced Performance Characteristics	2-30

1-1	Storage Domains on the Internet	1-4
1-2	Navisphere Management Software Layers	1-5
2-1	Analyzer Tree Controls	2-7
2-2	Sample Performance Survey Chart	2-9
2-3	Sample Performance Summary Chart	2-12
2-4	Performance Summary Chart before Sort	2-15
2-5	Performance Summary Chart after Sort by Maximum	2-16
2-6	Sample Performance Detail Chart Showing Utilization of SP and LUN	2-18
2-7	Performance Detail Display before Highlight	2-19
2-8	Performance Detail Display after Highlight	2-19
2-9	Sample IO Size Distribution Summary Chart	2-24
2-10	Sample IO Size Distribution Detail Chart	2-26
2-11	Chart Before Zoom	2-35
2-12	Chart With Zoom Area Selected (Left Mouse Button)	2-36
2-13	Chart After Zoom (Release of Left Mouse Button)	2-36
2-14	Customize General Dialog Box	2-39
2-15	Customize Performance Survey Dialog Box	2-40
2-16	Customize Summary Chart Dialog Box	2-42
2-17	Customize Detail Chart Dialog Box	2-43
2-18	Customize Archives Box	2-44
2-19	Retrieve Archive Dialog Box	2-45
3-1	Performance Analysis Flowchart	3-6

This manual describes using the EMC ControlCenter™ Navisphere® Analyzer disk-array performance monitor, called Analyzer.

Audience and Objectives

This manual is for people who want to monitor and tune disk-array storage systems. After reading the manual, you will be able to use Analyzer to monitor storage-system performance. The manual assumes

- ◆ You are familiar with the operating system running on the servers and with the concepts and operation of the storage systems.
- ◆ You understand Navisphere Manager Agent selection and tree structures and how to use them. They are explained in the *EMC ControlCenter Navisphere Manager Administrator's Guide*.
- ◆ The disk-array storage systems have been connected to servers, and that each server is using storage-system disks.
- ◆ Analyzer is installed. Installing Analyzer is explained in the *EMC ControlCenter Navisphere Analyzer Installation Guide*.

Organization of the Manual

- | | |
|-----------|--|
| Chapter 1 | Introduces Navisphere Analyzer. |
| Chapter 2 | Explains starting and using Analyzer. |
| Chapter 3 | Explains interpreting Analyzer performance data. |

Related Manuals	<i>EMC ControlCenter Navisphere Analyzer 6.X Installation Guide,</i> P/N 069001182
	<i>EMC FC4700-2 Storage System Configuration Planning Guide,</i> P/N 014003087
	<i>EMC CX-Series Storage System Configuration Planning Guide,</i> P/N 014003113

Conventions Used in This Guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.



WARNING

A warning contains information essential to avoid a hazard that can cause severe personal injury, death, or substantial property damage if you ignore the warning.



DANGER

A danger notice contains information essential to avoid a hazard that will cause severe personal injury, death, or substantial property damage if you ignore the message.

Typographical Conventions

EMC uses the following type style conventions in this guide:

Format Conventions We use the following format conventions in this manual:

Convention	Meaning
This type	Indicates text (including punctuation) that you type verbatim, all commands, pathnames, and filenames, and directory names.
<i>This type</i>	Represents a variable; an item for which you substitute a valid value.
<code>This type</code> or <code>this type</code>	Represents a system response (such as a message, prompt, or program listing).
[]	Encloses optional entries. Do not type the brackets.
x → y	Represents a menu path. For example, Change Parameters → Change Storage System Parameters tells you to select the Change Parameters item from the menu that appears, and then select the Change Storage System Parameters item from the next menu that appears.
↵	Represents the Enter key. (On some keyboards this key is called Return or New Line.)

Where to Get Help For questions about technical support and service, contact your service provider.

If you have a valid EMC service contract, contact EMC Customer Service at:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

Follow the voice menu prompts to open a service call, then select CLARiON Product Support.

Sales and Customer Service Contacts For the list of EMC sales locations, please access the EMC home page at:

<http://www.emc.com/contact/>

For additional information on the EMC products and services available to customers and partners, refer to the EMC Powerlink Web site at:

<http://powerlink.emc.com>

Your Comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to **techpub_comments@emc.com** with your opinions of this guide.

About the EMC Navisphere Analyzer

The EMC Navisphere® Analyzer performance monitor application measures the performance of disk-array storage systems.

This chapter describes the following topics:

- ◆ About Analyzer 1-2
- ◆ Navisphere 6.X Management Environments 1-3
- ◆ Analyzer Dependencies 1-6

About Analyzer

Navisphere Analyzer gathers storage-system performance statistics and presents them in various types of chart. These charts can help you find and anticipate bottlenecks in the disk storage component of a computer system.

Analyzer 6.X has two parts: a provider program and a UI (graphical user interface).

The Analyzer provider must be installed on any storage system you want to analyze. It controls, accumulates, and manages access to performance information on SPs, LUNs, and disks. The new information displaces the old. The information stored covers approximately the previous 24 hours.

The Analyzer UI interprets and displays the information. The UI must be installed on a storage system along with Navisphere Manager, but it need not be installed on every storage system. If the UI is installed on just one system in a storage domain, you can run it from that system and interpret that performance data accumulated on any storage system in the domain.

Analyzer can display the performance data directly (in *real-time mode*) or as a file containing past performance data (in *archive mode*.) It can display real-time and archive file data simultaneously. You can compare charts from different storage systems (using real-time data) or compare different times on one storage system (using real-time and archive data).

Navisphere 6.X Management Environments

Navisphere Analyzer 6.X works within storage domains administered by the web-based Navisphere Manager.

Each storage system that runs Navisphere 6.X storage management server software can be assigned to a storage domain — a group of storage systems on an internet or intranet you define using Manager.

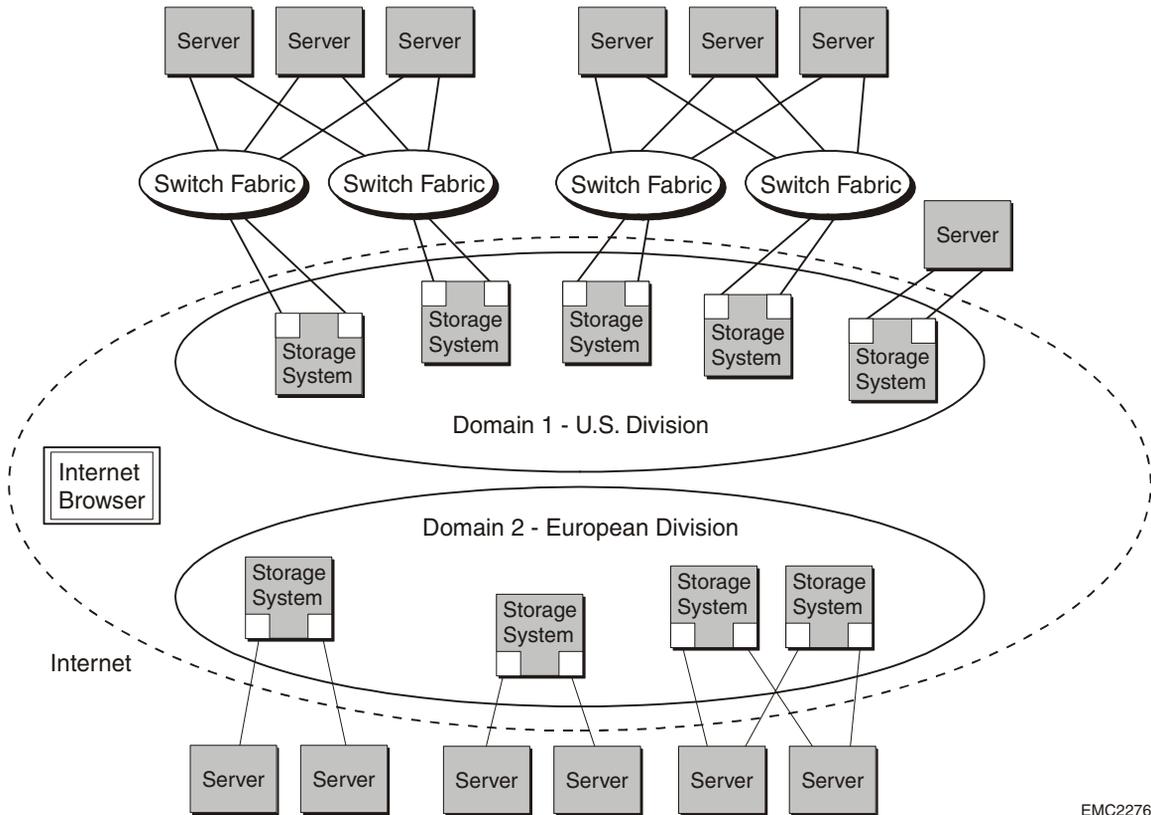
You can create one or more domains for any installation, provided that a storage management server can be a member of only one storage domain. Each domain must have at least one node (SP) that has Manager installed.

Each storage system in the domain is accessible from any other in the domain. Using an internet browser, you point at a storage system that has Manager installed. The security software then prompts you to log in. After logging in, depending on the privileges of your account, you can monitor, manage, analyze performance data for and define user accounts for any storage system in the domain. Storage systems outside the domain are not viewable from the domain.

You can run the internet browser on any supported PC or laptop with a network controller and browser software. At least one storage system — ideally at least two for higher availability — in a domain must have Manager (the Manager UI, user interface) installed.

The privilege of monitoring, managing, and creating accounts on systems in a domain depends on the type of account used when you log in.

The following figure shows an intranet that connects nine storage systems. It shows two domains, a US Division domain with five storage systems (four systems on SANs) and a European Division with four storage systems. The 13 servers that use the storage may be connected to the same or a different network, but the intranet shown is the one used to manage the storage systems.



EMC2276

Figure 1-1 Storage Domains on the Internet

Storage Management Software

A storage system with Navisphere software that lets it be managed by the web-based Navisphere Manager can belong to a domain. If such a storage system also contains Manager, you can log into the system that has Manager and then you can run Analyzer on any storage system in the domain that has the Analyzer provider installed.

You access Manager using an internet browser running on a computer called a client on a network connected to the storage application server.

The following figure shows two storage systems in a domain. If you log into the system that has Manager installed, you can run Analyzer on either storage system.

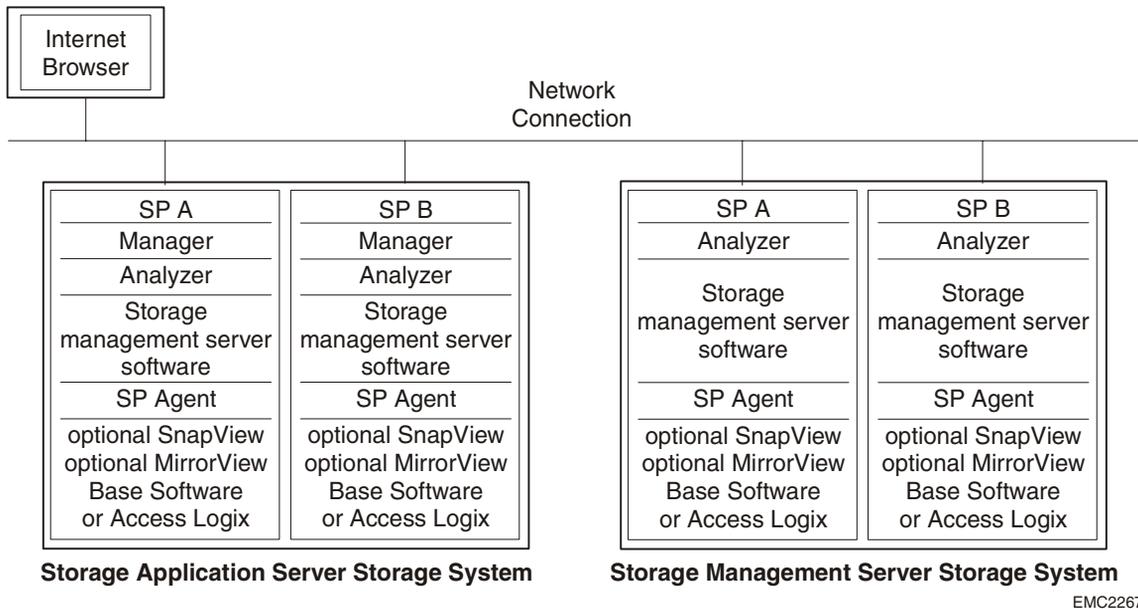


Figure 1-2 Navisphere Management Software Layers

Analyzer Dependencies

Analyzer requires the following on the management station:

- ◆ Navisphere Manager UI and Storage Management Server software of the same revision are on at least one of the storage systems in the domain of the storage system you want to analyze
- ◆ Analyzer UI is on the same storage system as the Manager UI
- ◆ Analyzer Provider is on the storage system you want to analyze
- ◆ JRE (Java Runtime Environment) of the correct revision (see Release Notes) on the client from which you run the browser.

For the latest requirements, see the Release Notes supplied with the software.

What Next? To start Analyzer, continue to Chapter 2.

This chapter explains using Analyzer to examine aspects of performance. It assumes you understand using the Navisphere Manager trees. Topics are

- ◆ Starting a Navisphere Manager or Analyzer Session.....2-2
- ◆ Menu Bar Options for Analyzer2-2
- ◆ Analyzer Charts2-6
- ◆ Performance Survey Chart2-8
- ◆ Performance Summary Chart..... 2-11
- ◆ Performance Detail Chart2-17
- ◆ IO Size Distribution Charts (LUN Only)2-23
- ◆ Performance Characteristics2-28
- ◆ Zooming and Resizing Charts.....2-34
- ◆ Exporting and Printing Charts.....2-37
- ◆ Customizing Chart Display2-38
- ◆ Capturing Performance Data in Archive Files.....2-45
- ◆ Discovering LUN Device Names.....2-46

Starting a Navisphere Manager or Analyzer Session

This section explains how to start a real-time Analyzer session.

1. Use a browser (Internet Explorer or Netscape) to point to the IP address of a storage system that has the Manager provider and GUI installed.
2. After a few moments, the software displays a log in dialog box.
3. Log into the storage system.

After a few moments, Manager starts and displays the Enterprise Storage tree.

The Storage tab display shows all storage systems in the domain.

Menu Bar Options for Analyzer

The Enterprise menu bar has a Tools option that has an Analyzer choice. When you select Tools -> Analyzer, the software presents the following choices.

Table 2-1 Menu Bar Options for Analyzer

Option	Use to
Open archive	Open an archive file that was previously saved on the system from which you ran the browser. (See Retrieve Archive below). Each archive file you open remains open until you close it or exit from Analyzer.
View archive	Display the data, in the chart type you specify, from a previously opened archive file. If multiple files are open, it displays a list box from which you can choose one or more files. This option appears dimmed if no archive file is open. View Archive offers a more flexible alternative to using the drop down menus on existing charts. It is also the only way to view data from an open archive file after all previously displayed charts from the archive file have been closed.
Close archive	Close all charts displaying data from an archive file, and releases the memory for any computations needed for display. If multiple files are open, it displays a list box from which you can choose one or more files. This option appears dimmed if no archive file is open. Use the close option when you have finished viewing an archive file's contents.

Table 2-1 Menu Bar Options for Analyzer (continued)

Retrieve archive	Capture the performance information recorded by the selected storage system and opens dialog boxes that let you specify a time period and a file to contain the performance data. (The SPs in each storage system automatically maintain an archive of performance data over the past 24 hours. Recent information overwrites older information.) The default extension for archive files is .nar .
Create Device to LUN File	Create a file that matches the LUN names as displayed by Navisphere to the specific LUN device names for your operating system. You can use these names to write custom management scripts.
Customize	Change display items such as colors, legends, and the initial screen Analyzer displays.

Analyzer can display performance information on LUNs, SPs, disks, and SnapView™ snapshot sessions, if any. The options available when you right-click on a LUN, SP, or disk icon are as follows.

Table 2-2 Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks

Icon Description	Menu Option	Displays
SP A or SP B	Performance Summary	Performance Summary window (performance range).
	Performance Detail	Performance Detail window (line chart).
LUN <i>lunID</i> [<i>RAIDtype</i>]	Performance Survey	Performance Survey window (graphs the five major performance characteristics of one or more LUNs).
	Performance Summary	Performance Summary window (performance range).
	Performance Detail	Performance Detail window (line chart).
	IO Size Distribution Summary	Bar chart showing numbers of I/Os of various sizes.
	IO Size Distribution Detail	Line chart showing numbers of I/Os of various sizes over time.
Disk <i>diskID</i>	Performance Summary	Performance Summary window (performance range).
	Performance Detail	Performance Detail window (line chart).

If other Navisphere applications are installed on the management station, you will see additional menu options. For information on these options, see the on line help Index or the manual for the application.

Real-Time and Archive File Display

Analyzer 6.X has two parts: a provider program that accumulates performance information and a UI (graphical user interface) that interprets and displays the information.

Analyzer can display data in *real-time mode*. Or it can display a file containing past performance data in *archive mode*. You can capture performance data in an archive file at any time and store it where you want on the server from which you connected with Analyzer.

Analyzer can display real-time and archive file data simultaneously. You can compare charts from different storage systems (using real-time data) or compare different times on one storage system (using real-time and archive data).

To obtain performance data (to display in real-time mode or save as an archive file) from any storage system, you must work through (point your browser at and log into) a storage system that has Manager and Analyzer installed.

To Display Real-time Data

1. Log into a storage system in a domain as explained earlier.
2. From the **Storage** tab, Expand the **Storage Systems** tree and select a storage system.
3. Select one or more storage system LUNs, SPs, or disks in the Enterprise **Storage** tree view. Generally, LUNs and SPs provide broader workload metrics than disks.

For performance data to be gathered on a LUN, disk, or SP, the Analyzer provider must be installed on the storage system that owns (or is) the selected item.

4. Right-click the selected item.
5. On the pull-down menu, select **Performance Survey** (LUN only); **Performance Summary**; **Performance Detail**; **IO Size Distribution Summary** (LUN only); or **IO Size Distribution Detail** (LUN only).
6. After pausing for a few moments to gather data, Analyzer will start to display a chart of the selected type for the selected components. A graph, if you selected one, does not appear immediately.

Analyzer displays values as data is generated at pre-determined intervals. Closing a chart discards all accumulated data.

You can display other real-time charts or open archive files (below) as needed.

To Display Archive Data

For archive file display, Analyzer requires at least one Analyzer archive file accessible locally.

1. On the Analyzer menu bar, select **Open archive**
2. From the standard file selection dialog box, select an archive file and click **Open**. Archive files are stored on the server from which the browser was run, in the directory specified, when the archive data was retrieved and stored. If you log in via a browser from a different host, you cannot access the archive files.

You can retrieve an archive file of the latest 24-hour data on a storage system as explained in *Capturing Performance Data in Archive Files* on page 2-45.

The screenshot shows a dialog box titled "Log File Display Time". It contains four text input fields with the following values:

- Earliest Sample Time: Jan 18, 2000 10:45:48
- Latest Sample Time: Jan 19, 2000 21:10:17
- Set Display Start Time: Jan, 18, 2000 10:45:48
- Set Display End Time: Jan, 19, 2000 21:10:17

At the bottom of the dialog are two buttons: "OK" and "Cancel".

3. Specify a time period to display or select the entire archive period by clicking **OK**.

The range of start times is limited to the period included in the archive file. Analyzer will create and display a graph with data points based on the polling interval used when the data was recorded.

By default, Analyzer displays a Performance Survey chart using the archive file. (You can change the initial display with the **Customize** option.) You can display real-time charts or open other archive files simultaneously.

Analyzer can display up to 255 data samples for each component, beginning with the first sample at the start time entered. If there are more than 255 points between the start and end times, Analyzer scales the data to show the entire range. For example, if the range includes 300 points, Analyzer displays every other point.

Analyzer Charts

This section explains Analyzer tree controls and chart types. Note that some of the screen captures shown here appear slightly different depending on your operating system and browser.

Using the Analyzer Tree Controls

Analyzer has its own tree, displayed in the upper left panel of most screens, that lets you control the display. The Analyzer tree offers several types of selection option:

- ◆ Expand/collapse — A plus sign (+) at the start of a line, when clicked, will expand to show all subsidiary items. A minus sign (-) means that the line has been expanded and clicking it will collapse the display and hide the subsidiary items.
- ◆ Select/Highlight — When you select an item in the tree, the performance characteristic display in the lower left panel changes to show only those characteristics that pertain to that object.
- ◆ Checkbox — Selecting a checkbox displays the performance metric for the item in the chart. The Select All right-mouse button choice selects all checkboxes; Clear All clears them all.

The following figure shows the three control types.

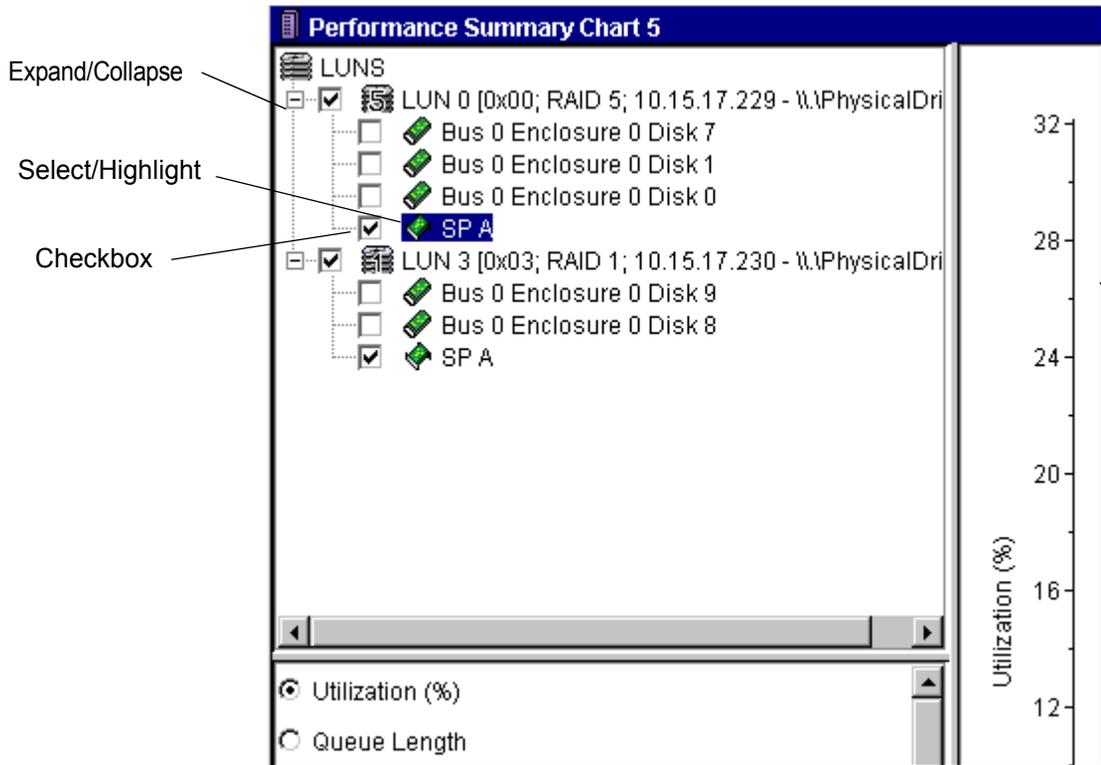


Figure 2-1 Analyzer Tree Controls

Chart Types

Analyzer has five types of chart:

- ◆ Performance Survey (LUN only), page 2-8
- ◆ Performance Summary, page 2-11
- ◆ Performance Detail, page 2-17
- ◆ IO Size Distribution Summary (for LUNs only), page 2-23 and
- ◆ IO Size Distribution Detail (LUNs only), 2-26

To choose a chart type, you select the objects you want, and then right-click to display a menu.

Each Analyzer chart is a unique window in the main Navisphere application. The Performance Summary and Performance Detail charts include their own tree, whose organization differs from the Enterprise Storage trees. Each Analyzer tree has a checkbox that you

can select or clear. Each component also has an icon that shows the device type and the device name.

For any chart that has a tree, Analyzer displays information for only those components whose checkboxes are selected. The display depends on real-time or archive-file data.

Analyzer takes the data directly from the storage system or from an archive file and displays the data graphically in the type of chart you specify.

Performance Survey Chart

The Performance Survey chart provides a simple, at-a-glance overview of performance for selected LUNs. It is an overview rather than an analysis tool.

The display has five small thumbnail graphs that show the essential performance characteristics of each LUN. For each characteristic, you can set a threshold value. The charts change color to indicate how long the threshold has been violated.

To display a Performance Survey chart, right-click one or more LUNs, and select **Performance Survey** from the menu. To display this chart from a Summary or Detail chart, right-click on the root node in the Tree window and select **Performance Survey**.

The LUN characteristics displayed are

- ◆ Utilization (%)
- ◆ Total Throughput (I/O/Sec)
- ◆ Total Bandwidth (MB/Sec)
- ◆ Response Time (ms)
- ◆ Queue Length

You can change the display to sort by any of these characteristics, descending or ascending, by left-clicking in the column header. Click again to restore the prior order.

For details on a graph, double-click it. This opens a Detail chart containing the same LUNs with the pertinent LUN and performance characteristic preselected.

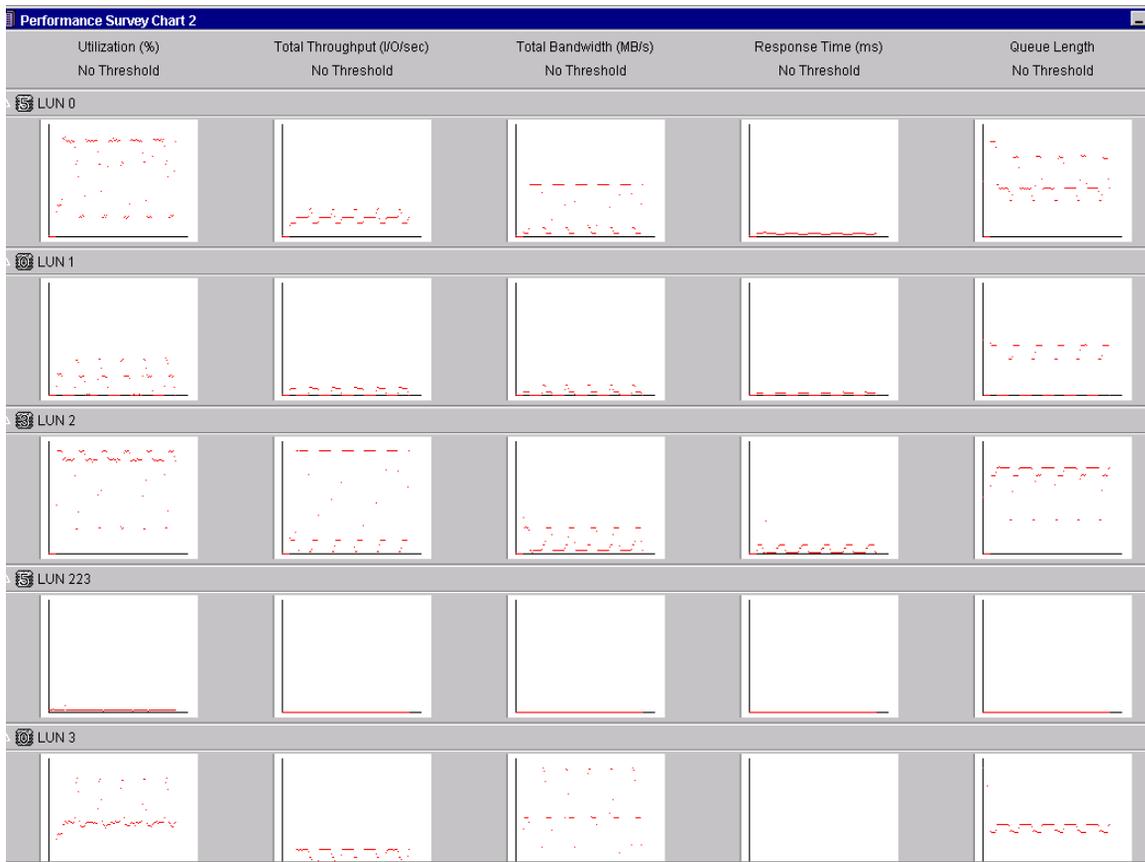


Figure 2-2 Sample Performance Survey Chart

LUN Display Graph

Each thumbnail graph displays the most recent 128 points. For each performance characteristic, you can specify one threshold value and an interval.

Each graph's vertical scale depends on the maximum value displayed for all LUNs. In other words, all Utilization (or Throughput, or Bandwidth, or Response Time, or Average Queue Length) graphs on a single Performance Survey Chart have the same vertical scaling.

This helps you compare the characteristic across all LUNs being displayed.

If you specify a threshold value and enable thresholds, Analyzer displays a dashed line (default color blue) for the threshold value. The graph border then becomes green, yellow, or red, depending on the threshold and interval settings:

- ◆ Green means that none of the plotted values on the graph exceed the threshold value.
- ◆ Yellow means some of the plotted values exceed the threshold value but not consecutively for the interval you defined.
- ◆ Red means that consecutive values exceed the threshold value for the interval you defined.

To set thresholds, see the section “Setting Performance Survey Chart Threshold Values” on page 2-40.

Performance Survey Menu

To display the menu for the Performance Survey chart, right-click the chart for the desired LUN. The Performance Survey menu options follow.

Table 2-3 Performance Survey Menu Options

Menu Option	Displays
Performance Summary Chart	Performance Summary chart.
Performance Detail Chart	Performance Detail window (line chart).
IO Size Distribution Summary Chart	Line chart showing numbers of I/Os of various sizes that occurred with the LUN. (LUN only)
IO Size Distribution Detail Chart	Line chart showing numbers of I/Os of various sizes that occurred with the LUN over time. (LUN only)
Print	Display a standard Print dialog box, so you can send an image of the chart to the printer you selected and configured in the dialog box.
Help	An explanation of the chart.

Performance Summary Chart

The Performance Summary chart shows the activity, based on the performance characteristic selected, for each device whose checkbox is selected. Each device display is cumulative and shows the range of all values over the last 255 data samples (real time) or during the period captured in the archive file.

The Performance Summary chart shows a value range over an entire period, while the Detail chart (described next) tracks the progression of values at intervals during the period.

The chart has three parts: Analyzer tree, to select devices; performance characteristics list, to select one performance characteristic (such as Utilization or Response Time); and the chart display space, which graphs the performance characteristics of the devices.

The chart display uses the following symbols by default:

- ◆ A red line for the range of values, minimum to maximum
- ◆ A green diamond for the latest value
- ◆ A pink circle for the average value

To display a Performance Summary Chart, use the steps explained in section *Real-Time and Archive File Display* on page 2-4.

A sample Performance Summary chart for a LUN, an SP, and the LUN's constituent disks looks like the following.

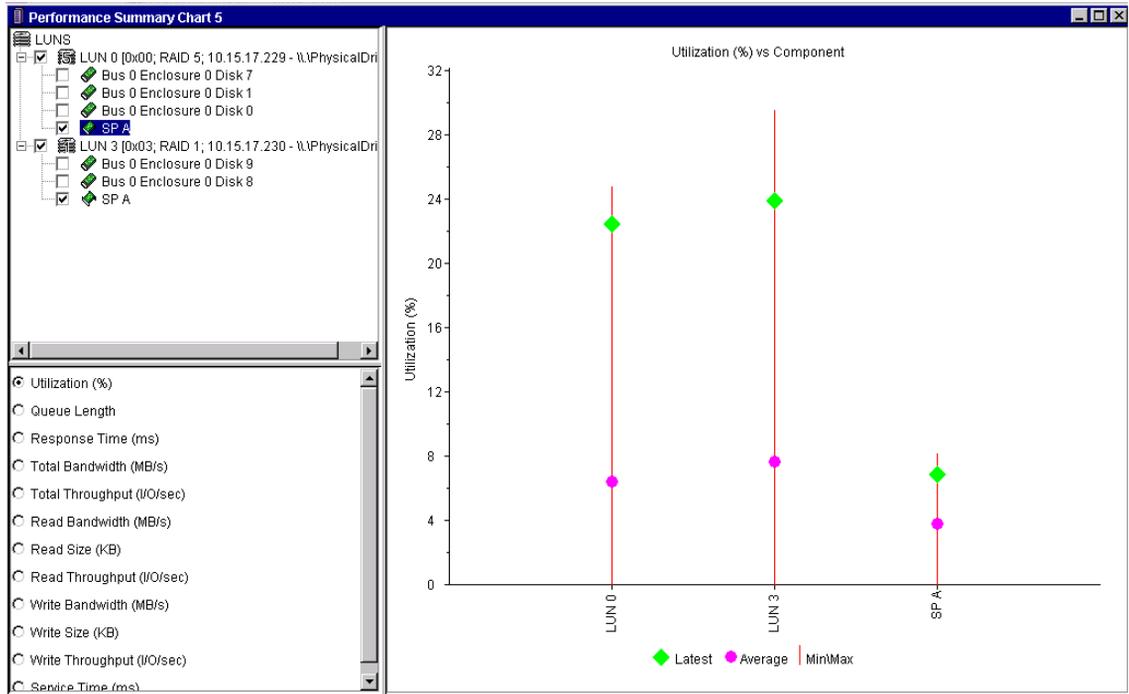


Figure 2-3 Sample Performance Summary Chart

You can select a device by clicking it. To remove a device's values from the chart, clear the device checkbox.

The performance characteristics list lets you select any single listed characteristic; the one shown is Utilization. The Performance Detail chart, described next, lets you select multiple characteristics. All the performance characteristics are explained starting on page 2-28.

Performance Summary Chart Menus

The Summary chart provides three types of right-click menu. Right-click in the tree panel root node to choose a different kind of chart; right-click in the tree panel outside of root node to select and deselect objects; or right-click in the chart area to export or print a chart.

The three right-click menu options are as follows.

Table 2-4 Performance Summary Chart Tree Root Node Right-Click Menu Options

Menu Option	Displays
Performance Survey	Performance Survey chart
Performance Summary	Performance Summary chart.
Performance Detail	Performance Detail chart (line chart).
IO Size Distribution Summary (LUN only)	Line chart showing the number of I/Os of various sizes.
IO Size Distribution Detail (LUN only)	Line chart showing the number of I/Os of various sizes over time, displayed at each interval.
Remove Object	Confirmation box that lets you delete the object (and any components) from the tree and chart.
Properties	Information on the device properties, such as model and capacity.
Help	An explanation of the chart.

Table 2-5 Performance Summary Chart Tree Panel Right-Click Menu Options

Menu Option	Use to
Select All	Select all objects.
Deselect All	Deselect all objects.
Help	Display an explanation of the chart.

Table 2-6 Performance Summary Chart Panel Right-Click Menu Options

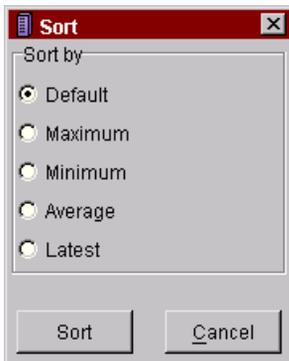
Menu Item	Use to
Export	Export Chart Data to File, Chart Data to Clipboard, or Graph to File. For details on exporting chart data, see <i>Exporting and Printing Charts</i> on page 2-37.
Print	Display a standard Print dialog box, so you can send an image of the chart to the printer you selected and configured in the dialog box.
Help	Display an explanation of the chart.

The Performance Summary chart also lets you change the vertical performance metrics themselves. When you right-click on a metric point, Analyzer displays a pop-up menu with the following choices:

Menu Option	Use to
Expand (LUN only)	Select all objects.
Collapse (LUN only)	Deselect all objects.
Sort	Lets you sort the displayed metrics by maximum, latest, or other values.
Hide	Remove an object from the chart.

Sorting Performance Summary Metrics

You can sort the values displayed on a Performance Summary chart by right-clicking on a chart data point. Analyzer displays a pop-up (explained above) from which you can choose Sort. Analyzer then displays the Sort dialog box (following) from which you select the order.



To restore the original (alphabetical) order, select the Default choice.

The Sort feature can be very helpful as you view a display with many values. You can use it with any Performance Summary chart.

The following figures show a Performance Summary chart before and after a sort by maximum values. Note that the performance metrics are grouped by LUNs, SPs and finally Disks, instead of being interspersed.

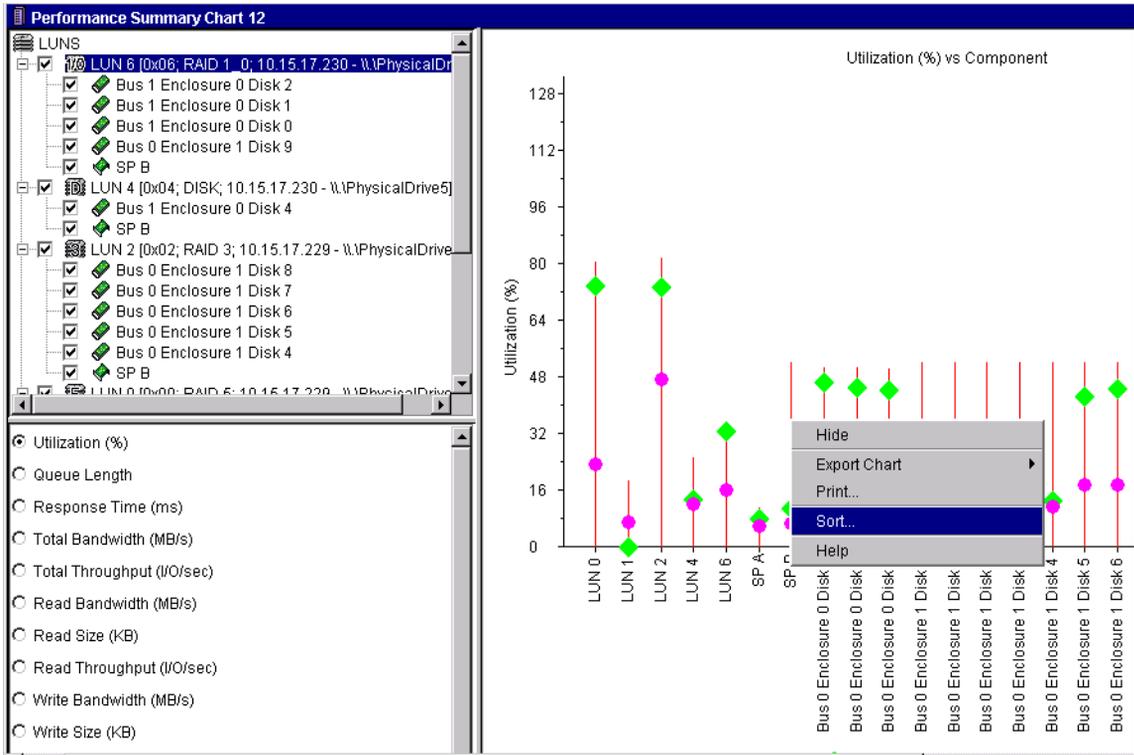


Figure 2-4 Performance Summary Chart before Sort

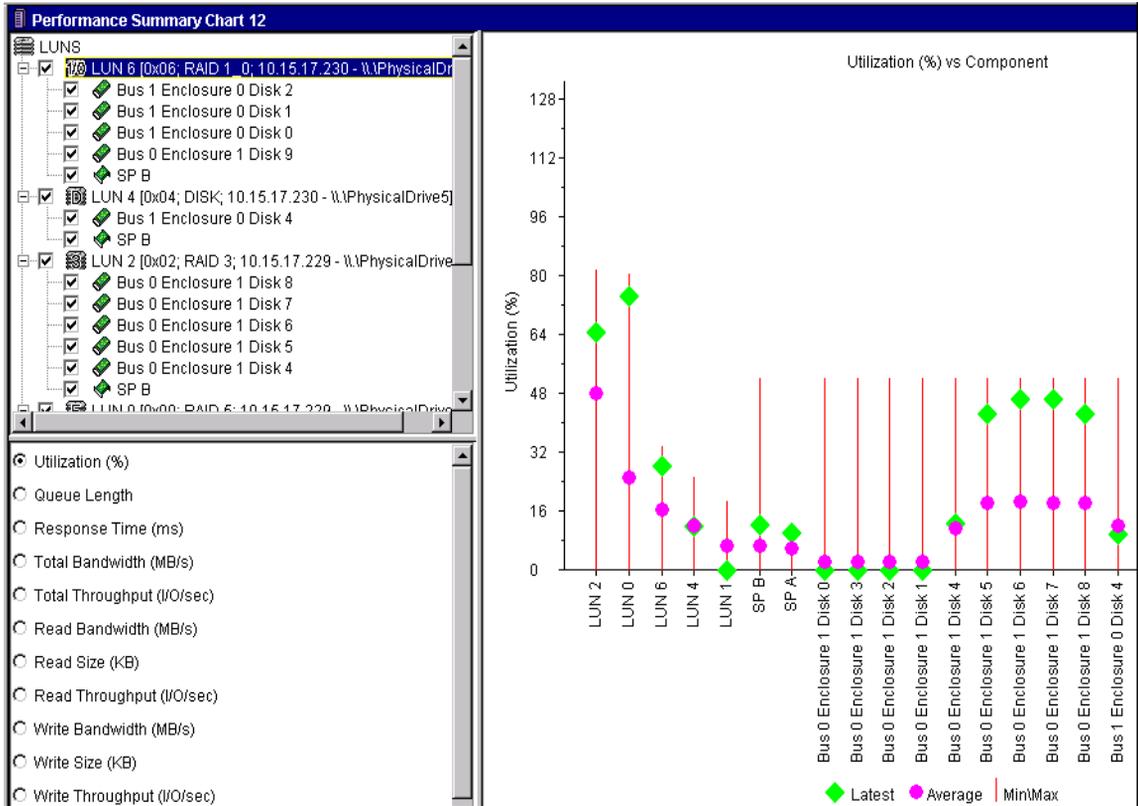


Figure 2-5 Performance Summary Chart after Sort by Maximum

Performance Detail Chart

The Performance Detail chart is a line chart that shows the performance characteristic (or characteristics) of the components selected when you selected the chart. Unlike the Summary chart, you can select more than one performance characteristic at a time. A sample line chart showing Utilization of the SP, the LUN, and the disks in the LUN follows.

The Detail chart shows the progression values at intervals during a period, unlike the Summary chart, which shows the range of all values seen during a period.

Time values are displayed on the X-axis until 20 samples are on display. After the 20th sample, the total number displayed remains 20. However, the time values are scaled so that each represents

(total-time-period-specified) / 20

All sample points are still displayed on the chart at the times at which they were taken.

Analyzer can plot a maximum of 255 samples at a time. Each sample after the 255th replaces the oldest sample.

With a Detail chart, you can select multiple performance characteristics; for example, you can select both Utilization and Total IOs/second. You can also select a time span to expand (zoom) with the left mouse button as explained on page 2-34.

To display a Performance Detail Chart, use the steps explained in section *Real-Time and Archive File Display* on page 2-4.

A sample performance detail chart follows.

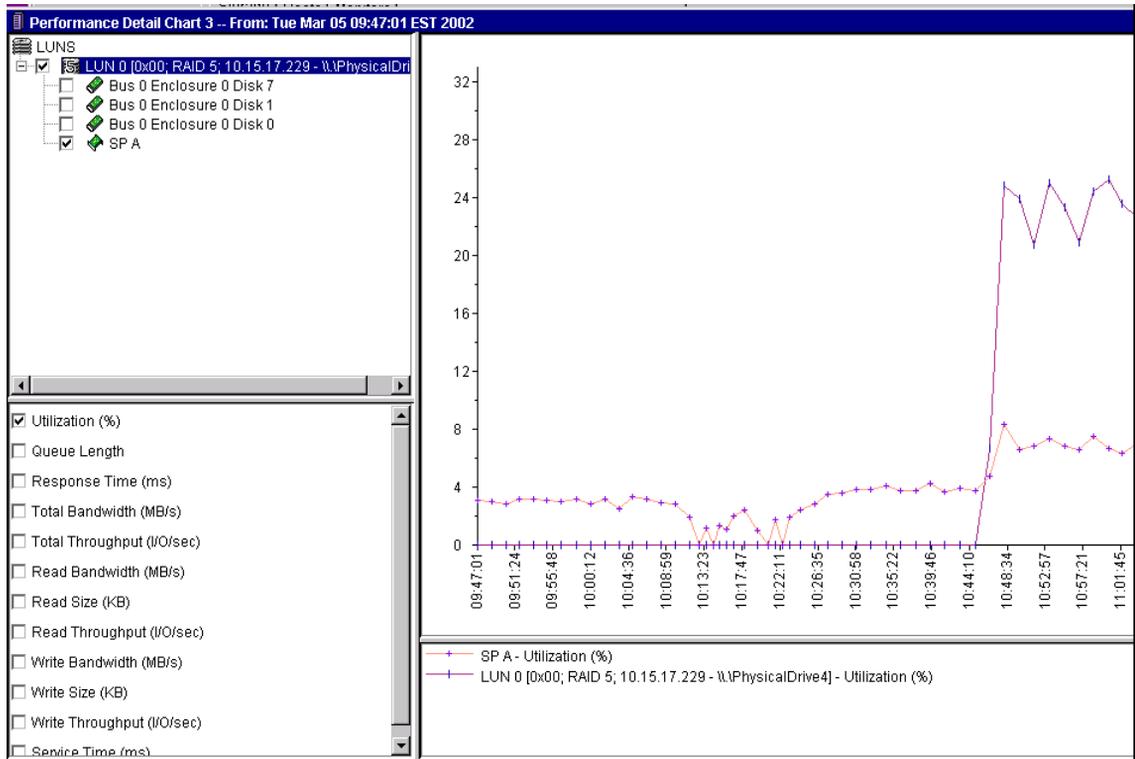


Figure 2-6 Sample Performance Detail Chart Showing Utilization of SP and LUN

You can direct Analyzer to chart multiple objects or a single object in the tree. To chart multiple objects, click each object checkbox.

Highlighting a Performance Detail Metric

You can highlight (make bold) a Performance Detail metric by clicking the legend shown below the chart. Highlighting is very helpful when you want to call attention to a specific metric. You can highlight a different metric by clicking its legend.

You can select only one line at a time. To remove highlighting, click in the Analyzer tree or performance characteristic window.

The following figures show a Performance Detail display before and after the LUN 6 Total Bandwidth metric has been highlighted.

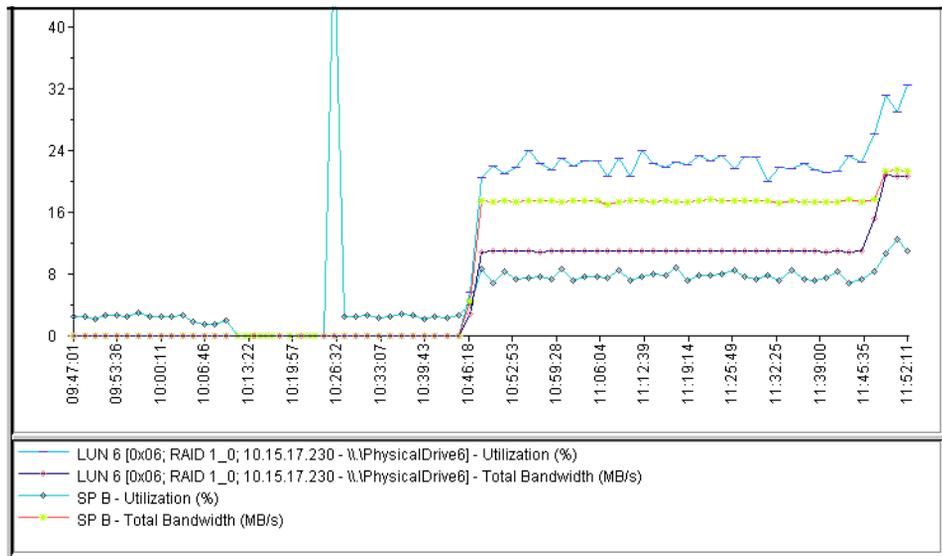


Figure 2-7 Performance Detail Display before Highlight

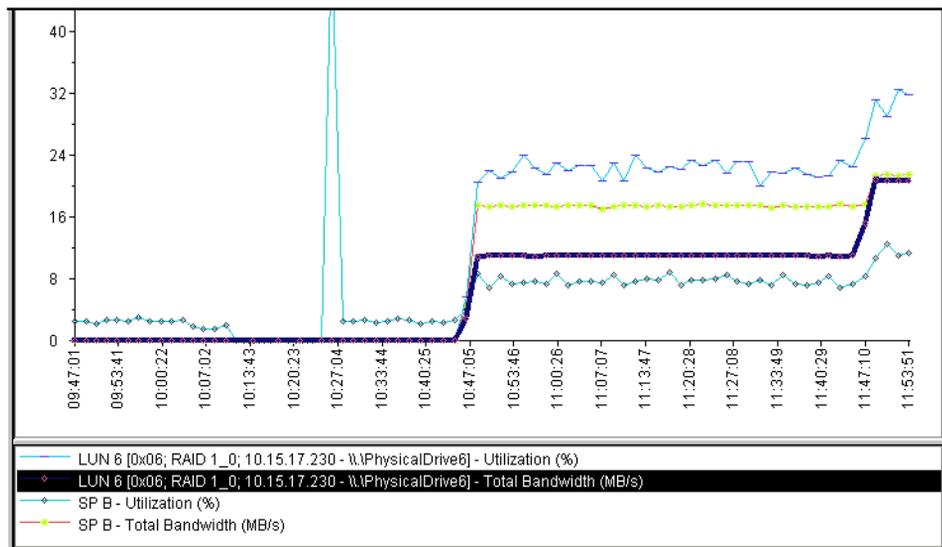


Figure 2-8 Performance Detail Display after Highlight

Performance Detail Menus

The Detail chart provides three types of right-click menu. Right-click in the tree panel root node to choose a different kind of chart; right-click in the tree panel outside of root node to select and deselect objects; or right-click in the chart area to export or print a chart.

The three right-click menu options are as follows.

Table 2-7 Performance Detail Chart Tree Root Node Right-Click Menu Options

Menu Option	Displays
Performance Survey	Performance Survey chart
Performance Summary	Performance Summary chart.
Performance Detail	Performance Detail chart (line chart).
IO Size Distribution Summary (LUN only)	Line chart showing the number of I/Os of various sizes.
IO Size Distribution Detail (LUN only)	Line chart showing the number of I/Os of various sizes over time, displayed at each interval.
Remove Object	Confirmation box that lets you delete the object (and any components) from the tree and chart.
Properties	Information on the device properties, such as model and capacity.
Help	An explanation of the chart.

Table 2-8 Performance Detail Chart Tree Right-Click Menu Options

Menu Option	Use to
Select All	Select all objects
Deselect All	Deselect all objects.
Help	Display an explanation of the chart.

Table 2-9 Performance Detail Chart Right-Click Menu Options

Menu Item	Use to
Reset Zoom	Restore the display of a chart zoomed on axis X or Y to its original size. Appears only if a chart has been zoomed.
Hide/Show Legend	Hide or display legends (map of line and marker colors and shapes).
Configure Detail Chart	Display the Configure Detail dialog box, which lets you customize the color and line style of each line that represents an item-property combination. See the following section for details.
Export	Expands to a menu with Export Chart Data to File, Chart Data to Clipboard, or Graph to File. For details on exporting chart data, see <i>Exporting and Printing Charts</i> on page 2-37.
Print	Display a standard Print dialog box, so you can send an image of the chart to the printer you selected and configured in the dialog box.
Help	Display an explanation of the chart.

Performance Detail Chart Data Tool Tips

You can obtain data on any device whose behavior is displayed in the Detail chart by moving the mouse cursor to the data marker and hovering for a few moments.

Analyzer will display device data until you move the mouse cursor. This feature lets you view the specific numeric value for the point along with device information and time stamp.

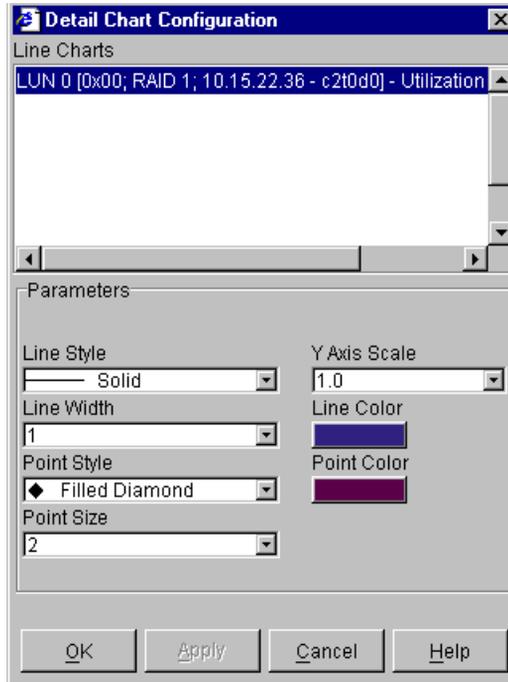
Configuring Detail Chart Displays

For Performance and IO Size Detail charts, you can change Detail Chart display qualities such as color and line style as follows:

To Change Details on a Performance Detail Chart

1. Right-click anywhere in the Performance Detail chart.
2. Click the option **Configure Detail Chart**.

The **Configure Detail Chart** dialog box appears as follows:



If you want to change multiple line parameters, you must click **Apply** after setting each line's parameters.

To Change the Chart Configuration

1. Select one or more charts in the display area using multiple select. When you select more than one chart, only the Y-Axis scale, point size, and line width options remain available; the others are greyed out.
2. Select a line style.

Line Style lets you specify dashed or solid lines. You can select solid (default) or choose from the list box.
3. Select a line width.

Line sets the relative line width: 1 (default), 2, 3, 4, or 5.
4. Select a point style.

Point Style lets you specify a graphic to mark each data point. You can select none (default), or choose a graphic from the list box.

5. Select a point size.

Point Size sets the size of the point graphic (**Dot, Box, Triangle**, and so on, as selected above). You can choose values from 0 (invisible point display graphic) through 11 points.

6. Select a Y-axis scale.

The **Y-axis** scale sets the order of magnitude for the display data along the Y-axis (vertical axis). The default scale for the Y-axis is 1. When you change the Y-axis scale value, Analyzer converts the display data to the new scale

For example, if the Y-axis numbers range from 0-800 and you change the Y-axis scale to 10, the display data will now range from 0 - 8000.

7. Select a line color.

Line Color shows a palette from which you can choose line color. To display the palette, click the button.

8. Select a point color.

Point Color shows a palette from which you can choose point color. To display the palette, click the button.

9. To see the chart display in the display area, click **Apply** or, if you are satisfied with the display, click **OK**.

After you click **OK**, the chart settings you specify become standard for Detail chart display after you save them.

IO Size Distribution Charts (LUN Only)

The IO Size Distribution charts show how many I/Os within given size ranges occurred for a LUN. There are two charts: a summary and a detail chart.

The IO Size Distribution Summary chart shows the total number of reads and writes of specific size ranges. This kind of chart is also known as a histogram.

The IO Size Distribution Detail is a line chart which shows read and write totals. The Detail chart can include as many 255 data sampling intervals.

To display an IO Size Distribution chart, use the steps explained in section *Real-Time and Archive File Display* on page 2-4.

The following figure shows an IO Size Distribution Summary chart. A detail chart appears on page 2-26.

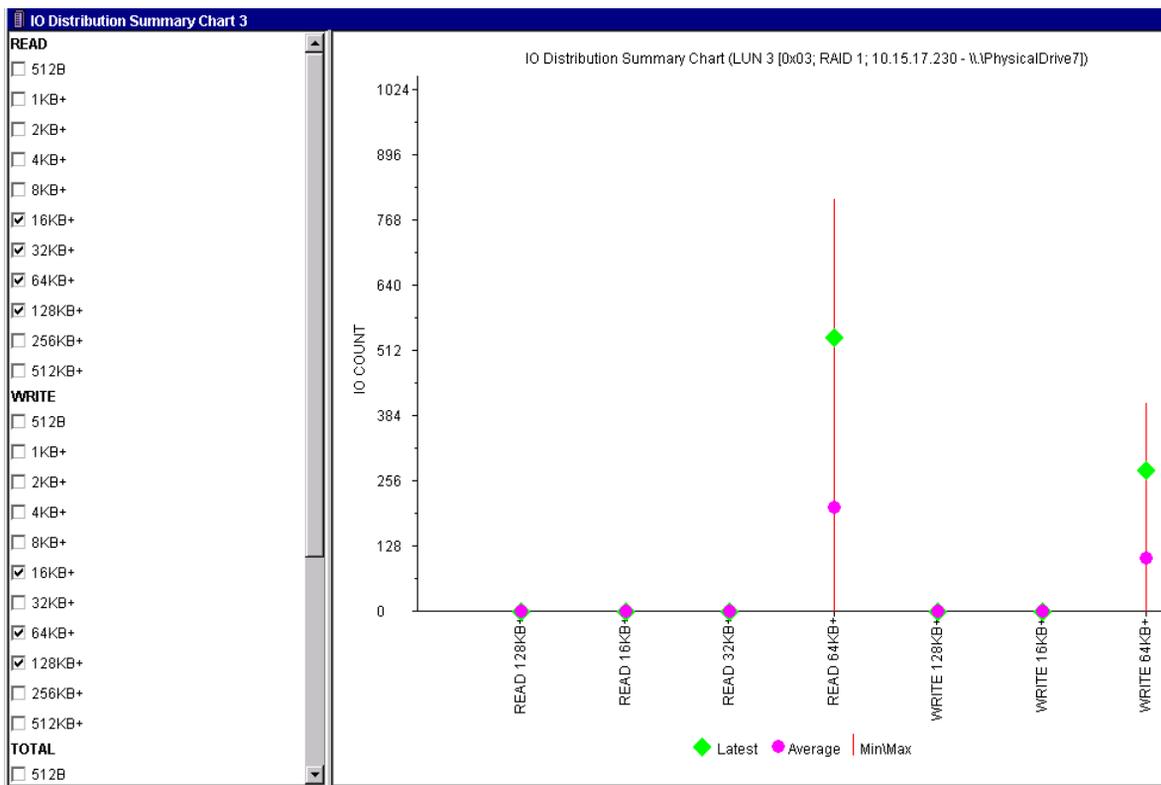


Figure 2-9 Sample IO Size Distribution Summary Chart

This display shows various read and write I/O size distributions at 8K, 16K, 32K, 64K, and 128K.

IO Size Distribution Summary Menus

To display the menu for an IO Size Distribution Summary chart, right-click in the IO size characteristics list. The menu options follow.

To display the menu, click anywhere in the Summary chart left panel. As with other charts, clicking in the chart area lets you export or print.

Table 2-10 IO Size Distribution Summary Chart Left Panel Right-Click Menu Option

Menu Option	Use to
Performance Summary Chart	Display a Performance Summary chart.
Performance Detail Chart	Display a Performance Detail chart
IO Size Distribution Detail	Line chart showing the number of I/Os of various sizes over time, displayed at each interval.
Select All	Select all IO sizes
Deselect All	Deselect all IO sizes.
Help	Display an explanation of the chart.

To display the menu for the IO Size Distribution chart, right-click anywhere in the chart area. The menu shows the following choices.

Table 2-11 IO Size Distribution Summary Chart Right-Click Menu Options

Menu Item	Use to
Export	Expands to a menu that lets you export the Chart Graph to File, Chart Data to File, or Data to Clipboard. For details on exporting and printing charts, see page 2-37.
Print	Print the chart on the default printer.
Help	Display an explanation of the chart.

IO Size Distribution Detail Chart

A sample IO Size Distribution Detail chart follows.

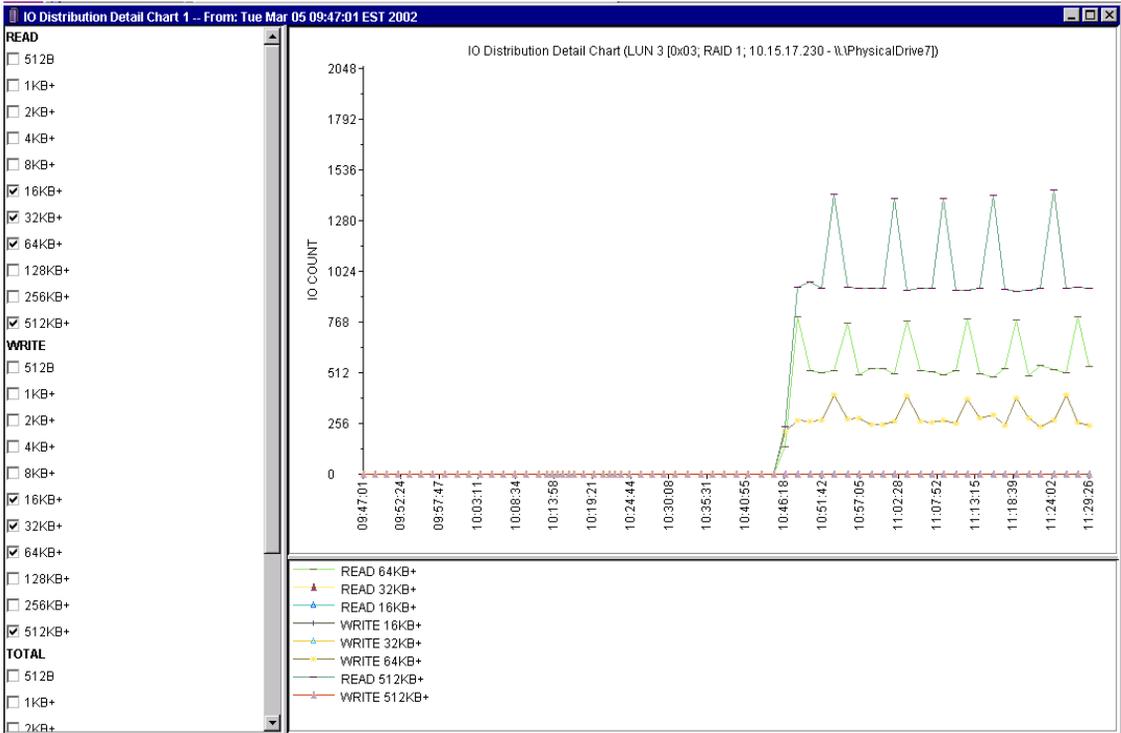


Figure 2-10 Sample IO Size Distribution Detail Chart

The previous display shows little activity for 32-, 64-, and 512-Kbyte reads and writes until about 10:40.

With a Detail chart, you can use the zoom feature as explained on page 2-34.

IO Size Distribution Detail Chart Menus

To display the tree menu, click anywhere in the Detail chart left panel. As with other charts, clicking in the left panel gives you a choice.

Table 2-12 IO Size Detail Chart Left Panel Right-Click Menu Options

Menu Option	Use to
Performance Summary Chart	Display Performance Summary chart.
Performance Detail Chart	Display Performance Detail chart
IO Size Distribution Summary	Line chart showing the number of I/Os of various sizes.
Select All	Select all IO sizes
Deselect All	Deselect all IO sizes.
Help	Display an explanation of the chart.

To display the menu for the IO Size Distribution chart, right-click anywhere in the chart area. The menu shows the following choices.

Table 2-13 IO Size Distribution Detail Chart Right-Click Menu Options

Menu Item	Use to
Reset Zoom	Restore a zoomed display to its original size. Appears only if chart has been zoomed.
Hide/Show Legend	Hide or show the legend key display that identifies IO sizes.
Configure Detail Chart	Display the Configure Detail dialog box, which lets you customize the color and line style of each line that represents an item-property combination. See page 2-22 for details.
Export	Expands to a menu that lets you choose to export Chart Data to File, Chart Data to Clipboard, or Graph to File. For details on exporting, see <i>Exporting and Printing Charts</i> on page 2-37.
Print	Display a standard Print dialog box so you can send an image of the chart to the printer you select.
Help	Display an explanation of the chart.

You can sort the values displayed on an IO Size Distribution Detail chart by right-clicking on a chart data point. Sort details appear on page 2-14.

Performance Characteristics

The Analyzer Performance characteristics apply to disks, SPs, LUNs, and/or snapshot sessions. Analyzer lets you select only those characteristics that pertain to a specific device; for example, it displays Average Seek Distance for Disks only, not for SPs or LUNs.

There are two categories of performance characteristics. The categories available to customers are Basic (default) and Advanced. The Advanced category includes both Basic and Advanced characteristics.

To change between the Basic and Advanced characteristic display, use the sequence Tools → Analyzer → Customize, and check or clear the **Advanced** box.

The Basic performance characteristics are shown in Table 2-14, the Advanced Characteristics in Table 2-15. The tables show characteristic names for all devices, followed by the letters D (applies to disks), S (applies to SPs), L (applies to LUNs), and SS (applies to SnapView snapshots). For example, Utilization (D/S/L) means that the Utilization characteristic applies to disks, SPs, and LUNs.

Table 2-14 Basic Performance Characteristics

Characteristic	Description
Utilization (D/S/L)	For Disk, SP, LUN (D/S/L). Describes the fraction of a certain observation period that the system component is busy serving incoming requests. An SP or disk that shows 100% (or close to 100%) utilization is a system bottleneck since an increase in the overall workload will not affect the component throughput; the component has reached its saturation point. Since a LUN is considered busy if any of its disks is busy, LUN utilization usually presents a pessimistic view. That is, a high LUN utilization value does not necessarily indicate that the LUN is approaching its maximum capacity.
Queue Length (D/S/L)	The average number of requests within a certain time interval waiting to be served by the component, including the one in service. An (average) queue length of zero indicates an idle system. If three requests arrive at an empty service center at the same time, only one of them can be served immediately; the other two must wait in the queue, resulting in a queue length of three.

Table 2-14 Basic Performance Characteristics (*continued*)

Characteristic	Description
Response Time (ms) (D/S/L)	The average time, in milliseconds, required for one request to pass through a system component, including its waiting time. The higher the queue length for a component, the more requests are waiting in its queue, thus increasing the average response time of a single request. For a given workload, queue length and response time are directly proportional.
Total Bandwidth (MB/s) (D/S/L)	The average amount of data in Mbytes that is passed through a system component per second. Larger requests result usually in a higher total bandwidth than smaller requests. Total bandwidth includes both read and write requests.
Total Throughput (IO/s) (D/S/L)	The average number of requests that pass through a system component per second. Since smaller requests need a shorter time for this, they usually result in a higher total throughput than larger requests. Total throughput includes both read and write requests.
Read Bandwidth (MB/s) (D/S/L)	The average number of Mbytes read that were passed through a component per second. Larger requests usually result in a higher bandwidth than smaller ones.
Read Size (KB) (D/S/L)	The average read request size in KBytes. This number indicates whether the read workload is oriented more toward throughput (I/Os per second) or bandwidth (MB/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.
Read Throughput (IO/s) (D/S/L)	The average number of read requests passed through a component per second. Since smaller requests need less processing time, they usually result in a higher read throughput than larger requests.
Write Bandwidth (MB/s) (D/S/L)	The average number of Mbytes written that were passed through a component per second. Larger requests usually result in a higher bandwidth than smaller ones.
Write Size (KB) (D/S/L)	The average write request size in Kbytes. This number indicates whether the read workload is oriented more toward throughput (I/Os per second) or bandwidth (MB/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.
Write Throughput (IO/s) (D/S/L)	The average number of write requests passed through a component per second. Since smaller requests need less processing time, they usually result in a higher write throughput than larger requests.
Average Seek Distance (GB) (D)	Average seek distance in gigabytes. Longer seek distances result in longer seek times and therefore higher response times. Defragmentation might help to reduce seek distances.

Table 2-14 Basic Performance Characteristics (*continued*)

Characteristic	Description
Service Time (ms) (D/S/L)	Time, in milliseconds, a request spent being serviced by a component. It does not include time waiting in a queue. Service time is mainly a characteristic of the system component. However, larger I/Os take longer and therefore usually result in lower throughput (IO/s) but better bandwidth (MB/s).
Reads From Snapshot LUN (SS)	With the optional snapView software only: The number of read requests on snapshots during this snapshot session.
Reads From Snapshot Source LUN (SS)	With the optional SnapView™ software only: the number of reads during this Snapshot session from the source LUN. It is calculated by the difference between Total Reads in Session and Reads From Cache.
Writes To Snapshot Source LUN (SS)	With the optional SnapView software only: The number of writes during this Snapshot session to the source LUN (on the pertinent SP).

Table 2-15 Advanced Performance Characteristics

Characteristic	Description
Utilization (D/S/L)	For Disk, SP, LUN (D/S/L). Describes the fraction of a certain observation period that the system component is busy serving incoming requests. An SP or disk that shows 100% (or close to 100%) utilization is a system bottleneck since an increase in the overall workload will not affect the component throughput; the component has reached its saturation point. Since a LUN is considered busy if any of its disks is busy, LUN utilization usually presents a pessimistic view. That is, a high LUN utilization value does not necessarily indicate that the LUN is approaching its maximum capacity.
Queue Length (D/S/L)	The average number of requests within a certain time interval waiting to be served by the component, including the one in service. An (average) queue length of zero indicates an idle system. If three requests arrive at an empty service center at the same time, only one of them can be served immediately; the other two must wait in the queue, resulting in a queue length of three.
Response Time (ms) (D/S/L)	The average time, in milliseconds, required for one request to pass through a system component, including its waiting time. The higher the queue length for a component, the more requests are waiting in its queue, thus increasing the average response time of a single request. For a given workload, queue length and response time are directly proportional.

Table 2-15 Advanced Performance Characteristics

Characteristic	Description
Total Bandwidth (MB/s) (D/S/L)	The average amount of data in Mbytes that is passed through a system component per second. Larger requests result usually in a higher total bandwidth than smaller requests. Total bandwidth includes both read and write requests.
Total Throughput (IO/s) (D/S/L)	The average number of requests that pass through a system component per second. Since smaller requests need a shorter time for this, they usually result in a higher total throughput than larger requests. Total throughput includes both read and write requests.
Read Bandwidth (MB/s) (D/S/L)	The average number of Mbytes read that were passed through a component per second. Larger requests usually result in a higher bandwidth than smaller ones.
Read Size (KB) (D/S/L)	The average read request size in KBytes. This number indicates whether the read workload is oriented more toward throughput (I/Os per second) or bandwidth (MB/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.
Read Throughput (IO/s) (D/S/L)	The average number of read requests passed through a component per second. Since smaller requests need less processing time, they usually result in a higher read throughput than larger requests.
Write Bandwidth (MB/s) (D/S/L)	The average number of Mbytes written that were passed through a component per second. Larger requests usually result in a higher bandwidth than smaller ones.
Write Size (KB) (D/S/L)	The average write request size in Kbytes. This number indicates whether the read workload is oriented more toward throughput (I/Os per second) or bandwidth (MB/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.
Write Throughput (IO/s) (D/S/L)	The average number of write requests passed through a component per second. Since smaller requests need less processing time, they usually result in a higher write throughput than larger requests.
Used Prefetches (%) (L)	The indication of prefetching efficiency. To improve read bandwidth, two consecutive requests trigger prefetching, thereby filling the read cache with data before it is requested. Thus, sequential requests will receive the data from the read cache instead of from the disks, which results in a lower response time and higher throughput. As the percentage of sequential requests rises, so does the percentage of used prefetches.
Read cache Hit Ratio (L)	The fraction of read requests served from both read and write caches vs. the number of read requests to this LUN. The higher the ratio, the better the read performance.

Table 2-15 Advanced Performance Characteristics

Characteristic	Description
Read Cache Hits/s (L)	The number of read requests that were satisfied by either write or read cache within a second. A read cache hit occurs when recently accessed data is referenced while it is still in the cache.
Read Cache Misses/s (L)	The rate of read requests that could not be satisfied by the SP cache and therefore required a disk access.
Dirty Pages (%) (S)	Percentage of cache pages owned by this SP that were modified since they were last read from or written to disk. In an optimal environment, the dirty-pages percentage will not exceed the high watermark for a long period.
Flush Ratio (S)	The fraction of the number of flush operations performed compared to the number of write requests. A flush operation is a write of a portion of the cache to make room for incoming write data. Since the ratio is a measure for the back-end activity compared to front-end activity, a lower number indicates better performance.
Write cache Hit Ratio (L)	The fraction of write requests served from the write cache compared to the total number of write requests to this LUN. The higher the ratio, the better the write performance.
Write Cache Hits/s (L)	The number of write requests per second that were satisfied by the write cache since they have been referenced before and not yet flushed to the disks. Write cache hits occur when recently accessed data is referenced again while it is still in the write cache.
Write Cache Misses/s (L)	The number of write requests per second that could not be satisfied by the write cache since their data was not currently in the cache from a previously disk access.
Mbs Flushed/ s (MB/s) (S)	The number of megabytes per second written from the write cache to the disks. The value is a measure of back-end activity.
Forced Flushes/ s (L)	Number of times per second the cache had to flush pages to disk to free space for incoming write requests. Forced flushes indicate that the incoming workload is higher than the back end workload. A relatively high number over a long period of time suggests that you spread the load over more disks.
High Water Flush On (S)	Number of times, since the last sample, that the number of modified pages in the write cache reached the high watermark. The higher the number, the greater the write workload coming from the host.
Idle Flush On (S)	Number of times, since the last sample, that the write cache started flushing dirty pages to disk due to a given idle period. Idle flushes indicate a low workload.

Table 2-15 Advanced Performance Characteristics

Characteristic	Description
Low Water Flush Off (S)	Number of times, since the last sample, that the number of modified pages in the write cache reached the low watermark, at which point the SP stops flushing the cache. The higher the number, the greater the write workload coming from the host. This number should be close to the High Watermark Flush On number.
Write Cache Flushes/s (S)	Number of times per second that the write cache performed a flush operation. A flush operation is a write of a portion of a cache for any reason; it includes forced flushes, flushes resulting from high watermark, and flushes from an idle state. This value indicates back-end workload.
Disk Crossing (%) (L)	Percentage of requests that require I/O to at least two disks compared to the total number of server requests. A disk crossing may involve more than two disks; that is, more than two stripe element crossings. Disk crossings relate to the LUN stripe element size. Generally, a low value is needed for good performance.
Average Busy Queue Length (D/S/L)	Average number of requests waiting at a busy system component to be serviced, including the request that is currently in service. Since the queue length is counted only when the component is busy, the value indicates the frequency variation (burst frequency) of incoming requests. The higher the value, the bigger the burst, and the longer the average response time at this component.
Average Seek Distance (GB) (D)	Average seek distance in gigabytes. Longer seek distances result in longer seek times and therefore higher response times. Defragmentation might help to reduce seek distances.
Service Time (ms) (D/S/L)	Time, in milliseconds, a request spent being serviced by a component. It does not include time waiting in a queue. Service time is mainly a characteristic of the system component. However, larger I/Os take longer, and therefore usually result in lower throughput (IO/s) but better bandwidth (MB/s).
Reads From Snapshot Cache (SS)	With the optional SnapView software only: The number of reads during this session that have resulted in a read from the snapshot cache (instead of a read from the source LUN).
Reads From Snapshot Source LUN (SS)	With the optional SnapView software only: The number of reads during this snapshot session from the source LUN. It is the difference between Total Reads in Session and Reads From Cache.
Reads From Snapshot Copy (SS)	With the optional SnapView software only: The number of read requests on snapshot copies during this snapshot session.

Table 2-15 Advanced Performance Characteristics

Characteristic	Description
Writes To Snapshot Source LUN (SS)	With the optional SnapView software only: The number of writes during this snapshot session to the source LUN (on the pertinent SP).
Writes To Snapshot Cache (SS)	With the optional SnapView software only: The number of writes to the source LUN this session that triggered a copy-on-write operation (the first write to each snapshot cache chunk region).
Writes Larger Than Cache Chunk Size (SS)	With the optional SnapView software only: The number of writes to the source LUN during this session that were larger than the chunk size (they have resulted in multiple writes to the cache).
Cache Chunks Used in Snapshot Session (SS)	The number of chunks that this session has used.

Zooming and Resizing Charts

Analyzer lets you zoom into (expand) any type of chart along its time sequence (X Axis), numeric value (Y Axis), or both axes.

To zoom into a chart:

1. Click in the desired start point, move the cursor to the end point, and drag to display the zoom box. You can drag the zoom box sideways to expand the X axis, vertically to expand the Y axis, or drag diagonally to expand the area in both axes.
2. Release the mouse button.

The chart expands to display the range you selected. You can use the vertical and horizontal scroll bars as needed to display different parts of the chart.

To reset after zooming (that is, restore the chart display to its original state), right-click in the chart area and select the **Reset Zoom** option.

The following three figures show a graph before zooming, the same graph with a zoom area selected, and the graph after zooming.

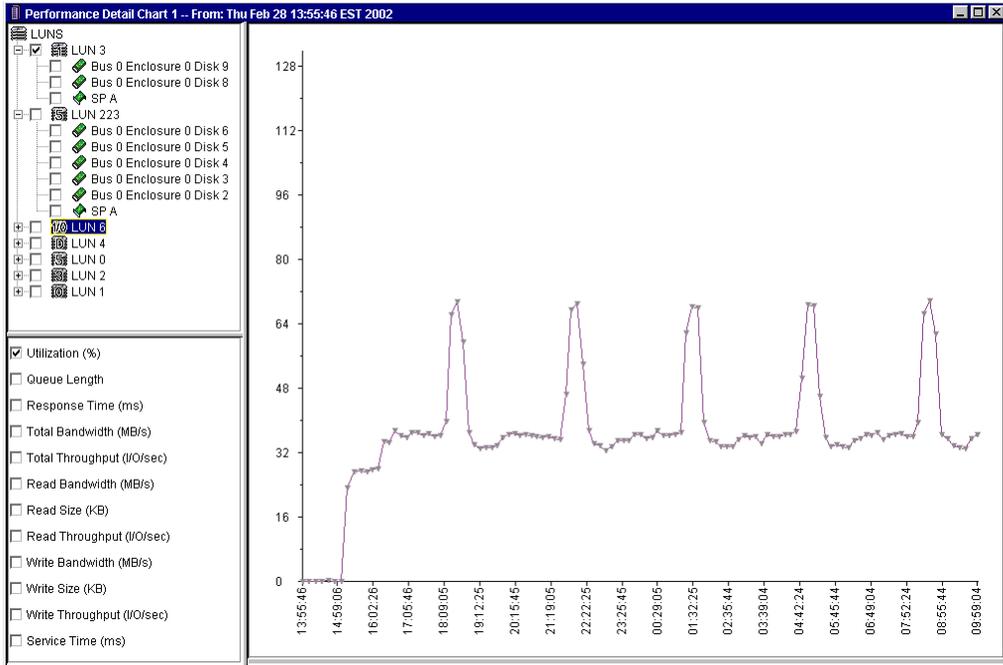


Figure 2-11 Chart Before Zoom

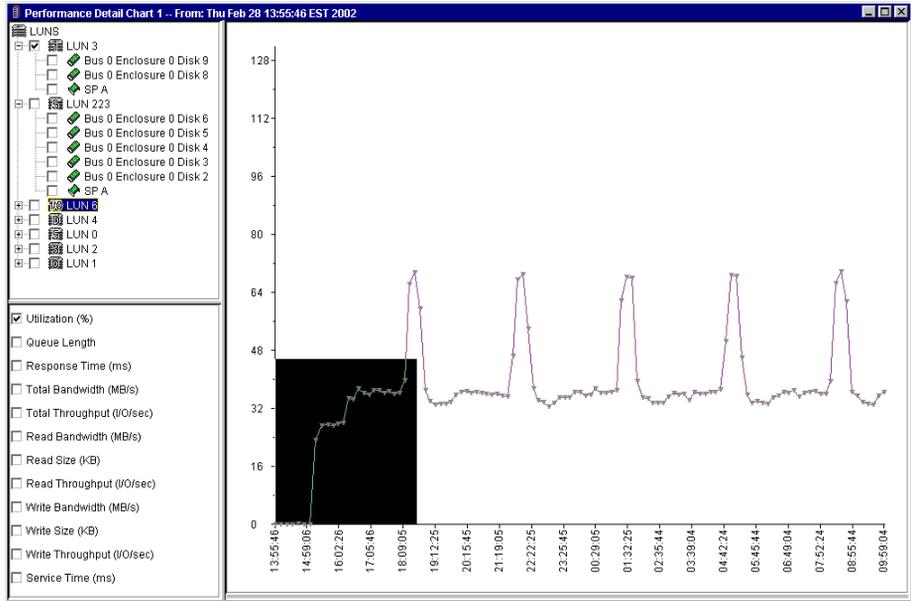


Figure 2-12 Chart With Zoom Area Selected (Left Mouse Button)

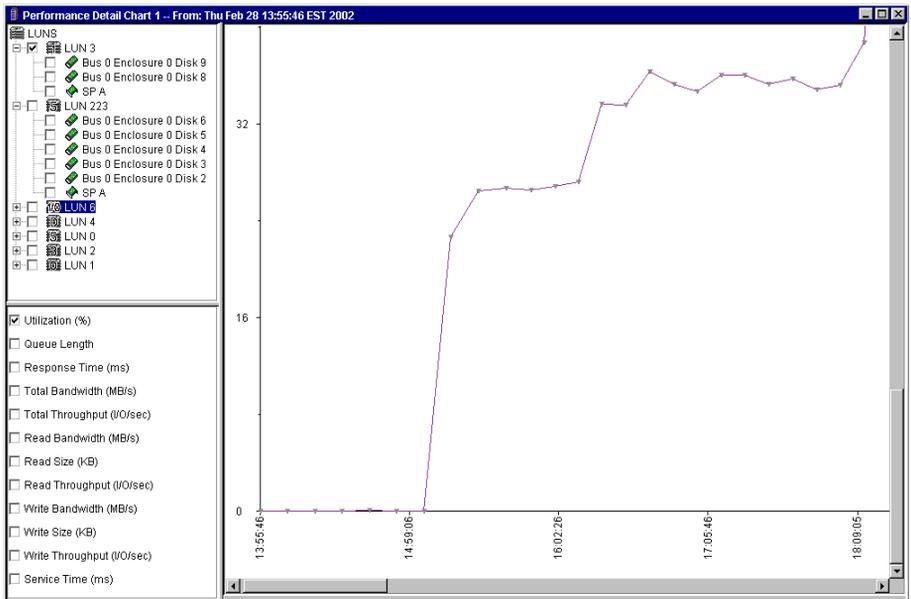


Figure 2-13 Chart After Zoom (Release of Left Mouse Button)

Exporting and Printing Charts

You can export chart data to a file, export chart data to the clipboard, or print a chart as follows (for all but Performance Summary charts).

To Export Chart Data to a File

1. Move the mouse cursor to the chart area, and right-click to display the pull-down menu.
2. Select **Export Chart** → **Export Chart Data to File**.

The **Save As** dialog box appears.

3. Specify the filename and select the location on the local host.

Analyzer saves the file as a comma-delimited file which has an extension of **.csv**. The default filename is **ChartData.csv**.

Below is sample output from a Summary Chart display.

```
LUN 0x01 [RAID5;toolsdev1 - \\.\PhysicalDrive1],SP A,Enclosure 0
Disk 6,Enclosure 1 Disk 0,Enclosure 1 Disk 1,Enclosure 1 Disk 2,
Current Values:,99.9823,31.9149,44.3425,48.1707,41.4634,43.5976,
Average Values,99.981,28.5642,45.4579,44.3035,44.5652,43.5922,
Max Values:,99.9859,32.2188,49.3902,50.9146,48.6322,47.2561,
Min Values:,99.9734,20.4545,41.2844,38.9058,38.7195,40.2439
```

A typical spreadsheet program displays data in the following format.:

	LUN 0x00 [RAID1/0; toolsdev1]	SP B	Enclosure 0 Disk 0	Enclosure 0 Disk 1
Current Value:	0.952381	11.538462	0	0.952381
Average Value:	0.469667	9.01991	0.152207	0.31746
Max Values	0.952381	11.538462	0.456621	0.952381
Min Value:	0	7.33945	0	0

The Analyzer uses the same format for **Export Data to Clipboard**.

To Print a Graph

1. Move the mouse cursor to the chart area, and right-click to display the pop-up menu.
2. Select **Print**.

Analyzer Display a standard Print dialog box so you can send an image of the chart to the printer you select.

Customizing Chart Display

You can tailor Analyzer chart display items, such as Performance Survey thresholds, colors, and legends, using the **Customize** choice from the Analyzer menu bar (sequence **Tools** → **Analyzer**).

Your custom settings persist until you change them again using the **Customize** option.

The items you can customize are as follows.

For all charts (**General** tab):

- ◆ Advanced or basic performance characteristics
- ◆ Color of background
- ◆ Color of axis lines
- ◆ Color of text

For Survey charts (**Survey Charts** tab)

- ◆ Thresholds
- ◆ Color of threshold line and point indicator
- ◆ Sort order based on latest values or maximum values.

For Summary Charts (**Summary Charts** tab):

- ◆ Color of average point indicators (the circles that show average values)
- ◆ Color of latest point indicator (the diamonds that show latest values)
- ◆ Color of min/max lines (the vertical lines that show value ranges)

For Detail Charts (**Detail Charts** tab):

- ◆ Data markers, display or omit (the small circles at the junction of chart lines)
- ◆ Data legends, display or omit (the captions at the base of figures that define conventions, such as point shape and color.

For archive files (**Archive** tab):

- ◆ Initial chart display, Performance Survey, Summary, or Detail
- ◆ Initial selection of all tree objects at first display

The initial **Customize** dialog box, displayed when you select **Tools** → **Analyzer** → **Customize**, looks like this.

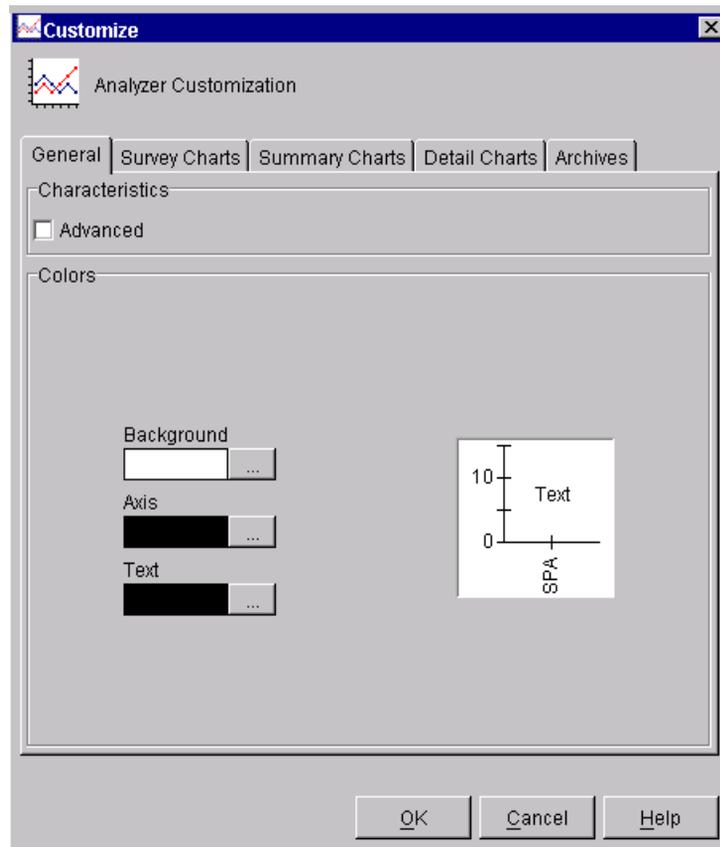


Figure 2-14 Customize General Dialog Box

Characteristics section — lets you specify the category of performance characteristics displayed by default: **Advanced** or Basic (default).

Colors section — The **Background** box lets you specify the color of all chart backgrounds; the **Axis** box lets you specify the color of Axis line; the **Text** box lets you specify the color of text. For any box, click the button beside the color sample to display a color selection dialog box; then select a color.

Click **OK**, and, to confirm the custom change, click **OK** again.

To customize display of Summary, Detail, or Archive charts, click the appropriate tab.

Setting Performance Survey Chart Threshold Values

You set threshold values to show when performance characteristics meet or exceed significant levels.

Each Survey Chart characteristic has a separate set of controls for configuring its thresholds. You display the dialog box for setting thresholds using the **Survey Charts** tab. The dialog box appears as follows.

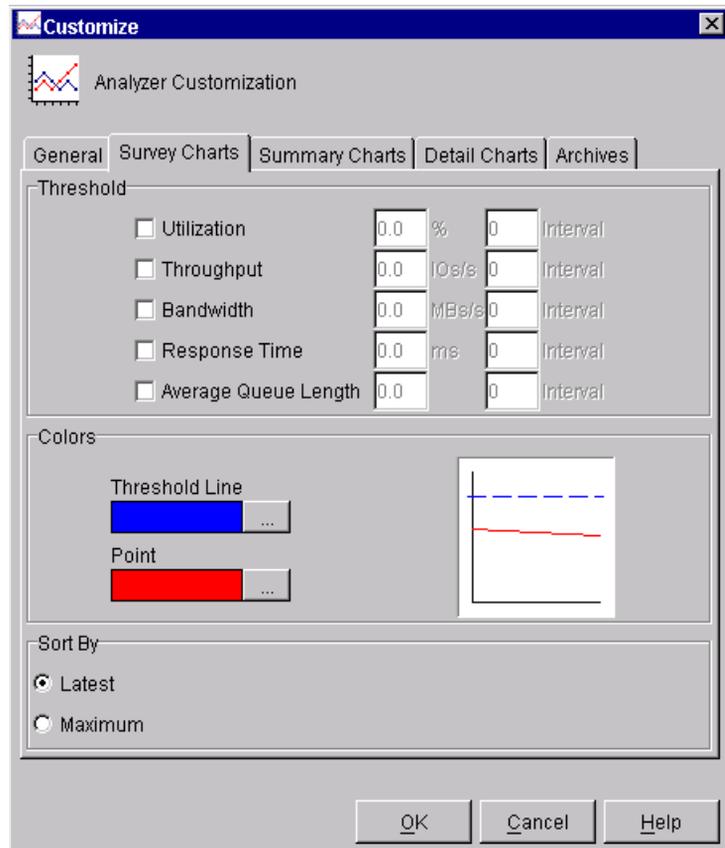


Figure 2-15 Customize Performance Survey Dialog Box

Threshold section — for each characteristic, the first box sets the value at which Analyzer will change the display color. The **Interval** value indicates the number of consecutive samples at which it will change the display color.

Colors section — **Threshold Line** lets you select the color of the line that represents the threshold value you set. **Point** lets you set the color of the data points generated by the software.

Sort By section — Sets the sort order of **Latest** (default) or **Maximum**. These settings determine the sort order for the chart columns in the **Survey** window. Sorting a performance characteristic by **Latest** compares the LUNs using the most recent data value plotted. Sorting a performance characteristic by **Maximum** compares the LUNs using the highest point plotted.

Click **OK**, and, to confirm the custom change, click **OK** again.

How to Set Useful Thresholds for Performance Surveys

Useful threshold values vary with storage systems and applications. A useful threshold produces an obvious indicator that lets you quickly identify potential problem LUNs. We suggest the following procedure.

1. Select a time when a representative load is being placed on a storage system, but when performance is acceptable. Open a Survey or Detail chart for the LUNs and find the maximum values currently being displayed for each of the five threshold parameters.
2. Set the threshold values 10%-20% higher than the current maximum values. Set the number of samples to 2 or 3.
3. Analyze the storage system again, using the **Survey** chart, when the load is somewhat higher — perhaps during a peak load or a time when performance has degraded noticeably. Note how many graphs changed to red or yellow. If they all changed to red or yellow, then the threshold values are set too low (or you have achieved the miracle of a perfectly balanced storage system). If no graphs changed, then the thresholds may be set too high.
4. If you find the Survey charts are too sensitive to transient spikes in the load — all the graphs tend to go from green to red — then increasing the number of samples will make the Survey charts more useful.

5. In general, adjust the values and the sample counts until you can quickly identify the busiest LUNs during periods of heavy load.

Setting Custom Summary Chart Colors

You can change the default colors for display of average, latest, and min/max display using the **Summary Charts** dialog box as follows.

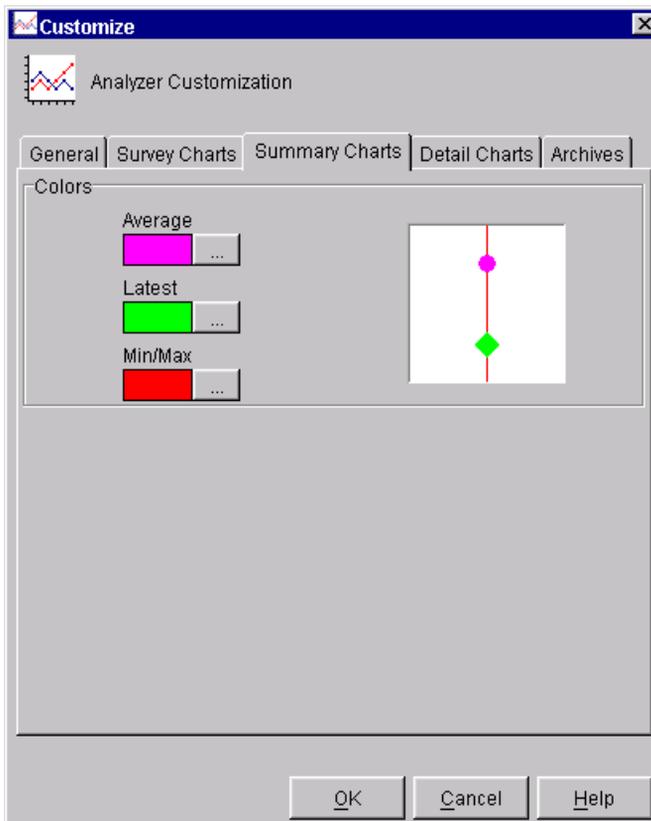


Figure 2-16 Customize Summary Chart Dialog Box

Colors section — The **Average** box lets you specify the color of the average point indicator; the **Latest** box lets you specify the color of the Latest indicator; and the **Min/Max** box lets you specify the color of the min/max line. For any box, click the button beside the color sample to display a color selection dialog box; then select a color.

Click **OK**, and, to confirm the custom change, click **OK** again.

Setting Custom Detail Chart Display Specifications

You can change the default display -- turning off display of data markers and/or the initial legend -- for detail charts using the **Detail Charts** dialog box as follows.

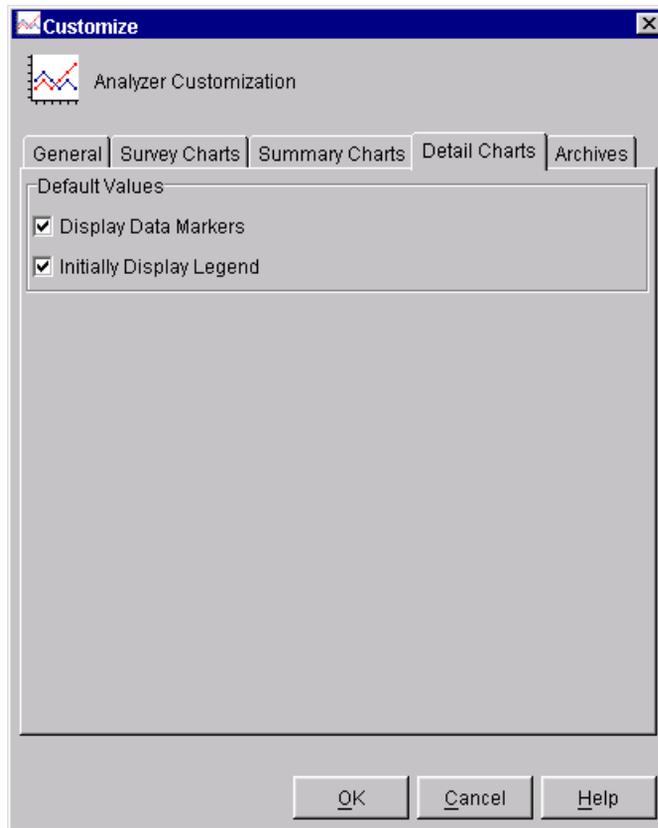


Figure 2-17 Customize Detail Chart Dialog Box

Default Values section — **Display Data Markers** lets you select display (default) or not. The data markers are the symbols that represent the individual data points. **Initially Display Legend** lets you select display of the data legend (default) or not. The data legend is a section below the chart that identifies the plot lines, their respective line color, and their data marker (if applicable).

Click **OK**, and, to confirm the custom change, click **OK** again.

Setting Custom Archive Chart Display Specifications

You can change the default display for Archive charts using the **Archives** dialog box as follows.

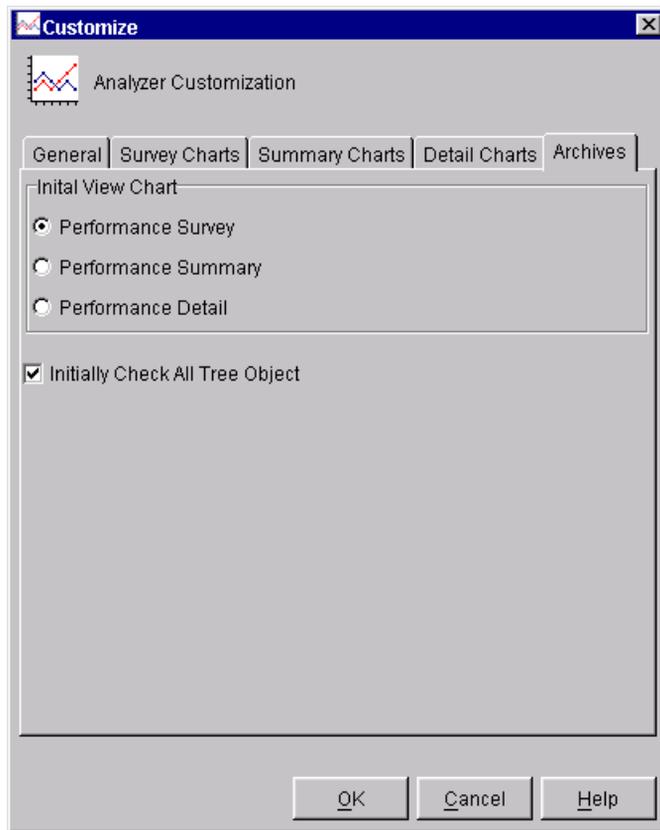


Figure 2-18 Customize Archives Box

Initial View Chart — Lets you specify the chart that Analyzer displays first: **Performance Survey** (default), **Performance Summary**, or **Performance Detail**.

Initially Check All Tree Objects — Specifies whether the software automatically selects all objects in an Archive when it opens the chart.

Click **OK**, and, to confirm the custom change, click **OK** again.

Capturing Performance Data in Archive Files

Generally, real-time charts are most useful as for monitoring. For detailed study, you may want to use archives.

Each storage system SP that uses Navisphere 6.X maintains a record -- called an archive -- of performance data during approximately the previous 24 hours. The archive is stored in the SPs, with each SP maintaining its own information. As the storage system runs, new information displaces the old. When you capture data from the archive, Analyzer captures the entire stored period of about 24 hours. Later, when you open and view the archive, you can specify part or all of the recorded period.

Select the storage system SP whose archive data you want to capture, and then from the Analyzer menu bar select **Retrieve Archive**. Analyzer then displays a dialog box that lets you specify the location and name for the file.

If you omit an extension from the filename you specify, Analyzer appends the extension.nar.

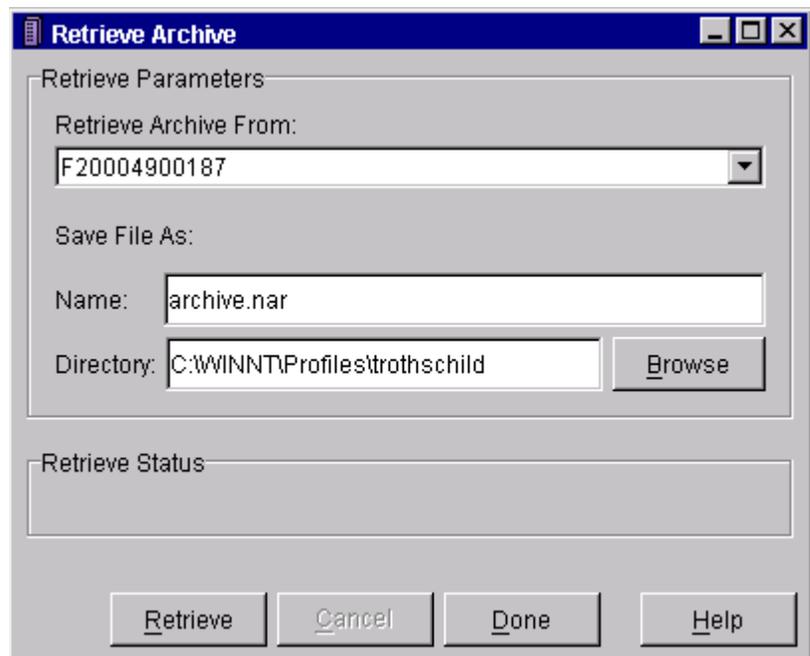


Figure 2-19 Retrieve Archive Dialog Box

To help you find archive files later on, you may want to create a standard directory for them on the system from which you normally run the browser; for example, **C:\performance\archive_files.** or **/usr/performance/archive_files.** You can use the SP hostname or other conventional name for the archive files themselves.

To specify a directory (which must already exist), use the **Browse** button to find and specify the directory, and then click **OK.**

Displaying archive files is explained starting on page 2-5.

Discovering LUN Device Names

Analyzer can create a file that provides useful information about the storage-system LUNs seen by Navisphere Manager, including the LUN operating system device name. Knowing the operating system device name can help you write management scripts (using Navisphere CLI commands) that perform maintenance and other tasks.

The file Analyzer creates with the LUN device names is called a device-to-LUN file. Each line of the file consists of the operating-system-dependent device name followed by the Navisphere LUN name.

To create a device-to-LUN file, select **Create Device To LUN File** on the Analyzer menu. Analyzer displays a dialog box that prompts for the required information and creates the required file.

The **Select Management Host** list box displays the management hostnames. Select the host from whose managed LUNs you want the file created.

The **Save File To** box lets you specify (via typing or using the **Browse** button) the file's path and filename.

The **Browse** button displays a file selection browse dialog. When you click **OK**, the path/file name entered will be returned to the Path/Filename field.

The **Save** button creates the file and closes the dialog box.

The **Cancel** button closes the dialog box without creating a file.

The **Help** button displays context-sensitive help.

Device-to-LUN File Format

The format of the file is one line of ASCII text for each mapping entry. Each line appears in the following format:

device-name LUN-specifier

where

device-name is the operating system device name. This varies with the operating system. For example,

For Solaris: **c2t1d1** (/dev/(r)dsk)

For Windows, formatted: **E:**

For Windows, unformatted: **\\.\PhysicalDrive1**

The separator after *device-name* is one space.

LUN-specifier identifies the LUN. It has the form

cabinetname_lunname [raid-type]

where

cabinetname is, by default, the cabinet serial number. However, you can rename cabinets in Navisphere, so this may be a user-defined name.

lunname is **LUN n** where n is the LUN number, based on order of creation in the cabinet.

raid-type is RAID0, RAID1/0, RAID1, RAID3, RAID5, or Disk (individual disk unit).

A sample line from a device-to-LUN file is

```
c2t1d1 F20010500278_LUN6[RAID5]
```

What Next?

To use Analyzer information to diagnose the cause of performance problems, continue to the next chapter.

Understanding Performance Data

This chapter describes how to solve performance problems using Navisphere Analyzer. Major topics are

- ◆ Understanding Workload3-2
- ◆ Identifying a Performance Problem3-5
- ◆ Performance Analysis Flowchart.....3-6
- ◆ Working with the Storage-System Cache3-9
- ◆ Addressing Server Problems3-13

Understanding Workload

Before trying to solve any performance problem, you should understand the workload that your application is applying to the storage system. This will help you correlate your knowledge of your application to Analyzer statistics.

Application workload has two important factors: workload features and intensity. The main features are

- ◆ I/O size
- ◆ Read/Write ratio
- ◆ Locality of reference (seek area, hot spots, and random/sequential access patterns)

These workload features are not constant; they vary with time depending on the work the application performs, and they vary from LUN to LUN.

To use Analyzer to examine the workload features, select the LUNs in the **Storage** tree and create Performance Summary charts to show their characteristics. Read and Write, throughput and bandwidth will tell you about the read/write traffic both in terms of data and I/O requests. Clicking an individual LUN will let you look at a time series of the quantity of interest.

Use this information along with your knowledge of your application to make sure the storage system is performing as expected. For example, suppose you are running a file server that is mostly used for data retrieval. If you find that the LUN workload is write-intensive, this indicates that something is wrong with the storage system and further investigation is needed.

I/O size data is also very important, particularly the quantity of traffic at each I/O size. For example, if you set the buffer size of a file system to 16 Kbytes, then you would expect a large amount of 16-Kbyte traffic. If the amount of data traffic is small, then a volume manager or I/O driver might be combining or splitting requests. Average Read and Average Write size data are useful but can be misleading. An average I/O size of 20 Kbytes may be mostly 16-Kbyte I/O with some 256-Kbyte I/Os; this is very different from half 8 Kbytes and half 32 Kbytes that would give the same result. To distinguish sizes, Analyzer lets you look at the I/O size distribution (histogram), which records the number of I/O of each size.

You cannot directly derive application program locality data using Analyzer. Such detail requires more complex analysis that cannot be achieved in real time. You can draw some locality conclusions from Analyzer cache data and disk seek data.

The quantity of work that is applied to a LUN is also important. A single-threaded application presents requests one at a time to the LUN. If the thread supplying requests waits between I/O operations (“think time”), then the LUN will have periods of inactivity. During these lightly loaded conditions, LUN utilization is a good measure of load. As load increases, the utilization rises until it approaches 100%. At this point, a single thread cannot request any more work from the LUN and the throughput is at its maximum. If the application is multithreaded, the load and throughput can continue to increase; at this point, Queue length is a better measure of applied load.

LUN response times measure the way that the storage system responds to an applied load. As the load increases, response times do not change until one of the components that implements the LUN reaches saturation (100% utilization). When this happens, response time will rise linearly with applied load. Unfortunately, response time depends on both load level and load features because writes take longer than reads, and large I/Os take longer than small ones. LUN response time is an excellent measure of how the storage system is affected by changing loads, but you must identify the factor of the workload that is changing.

If the system is to perform at its *theoretical* maximum, the workload intensity must be evenly distributed across all of the LUNs. In practice, this is difficult to achieve. Workload intensity varies with time as does the distribution of work between LUNs. The variations can occur over widely different time scales. Business cycles such as end of month, end of quarter, and even end-of-year-processing may be very different from normal day-to-day operation. Operational schedules such as backups and data-mining operations can cause problems. The user load associated with logging in and at break period can also vary intensity and change distributions.

Analyzer lets you look at the LUN workloads and determine

- ◆ the load at critical times of day
- ◆ which LUNs have most traffic
- ◆ which LUNs are idle

This might help you decide how to move data from one LUN to another to better distribute the load. With some thought, you might

be able to eliminate some bottlenecks by shifting the schedules for some system operations. After implementing any easy or obvious schedule improvement, you can use Analyzer to understand the impact of your changes. But remember, distributing data to solve one problem may cause another at a different time.

Storage Systems Under Light Loads

Under light load situations, indicated by low disk and SP utilizations, the response times experienced by the requests are mostly due to the processing times at the various system components. The queuing that can occur at higher loads does not influence the response times. For requests serviced from the cache, the response times are typically a few milliseconds. For requests serviced by the disks, the response times can be as high as 25 milliseconds. Under the following situations with light loads, you may be able to reduce response time by reconfiguring the read or write caches.

- ◆ Single-threaded sequential reads: If the cache is configured properly, prefetching will allow almost all the requests to be serviced from the cache. This will be indicated by a read cache hit ratio of 100%.
- ◆ Locality of reference by reads: If the size of the most frequently used data set is smaller than the read cache size, most of the requests will be serviced from the cache.

For write operations, data is written to the cache. Under light loads, a small write cache might suffice to absorb the data. Beyond that, tuning the cache configuration properties is unlikely to influence performance.

Configuring the cache involves setting the read and write cache sizes, and the cache page sizes. Some of the general guidelines for this are provided later in this chapter.

Storage Systems with Bottlenecks

There are several areas of computer system operations that pertain to bottlenecks:

- ◆ Operating cycles (for resource-hungry procedures such as backups),

- ◆ Business cycles (for periodic demanding applications such as data reduction for reports), and
- ◆ Time-of-day cycles (such as a 9:00 am surge when users log in).

To use Analyzer effectively with bottlenecks, first you need to determine when the bottleneck(s) occur, so that you can record a meaningful Analyzer log file that includes them.

Then note the components that have high utilizations (SPs and disks), click them to expand and show the time course. In general, if the utilization of a component is higher than 70%, the device is operating near its capacity.

Do not use LUN utilization values to detect bottleneck conditions. Since a LUN is considered busy if at least one of its disks is busy, LUN utilization usually presents a pessimistic view. Thus, a high LUN utilization value does not necessarily indicate that the LUN is operating at its maximum capacity.

If the utilization of components such as the disks and SP is low, the problem is likely at the server. See the section *Addressing Server Problems* on page 3-13.

Identifying a Performance Problem

After appraising your workload and deciding on obvious solutions, you can try a procedural approach to identify performance problems. The following flowchart guides you toward this goal.

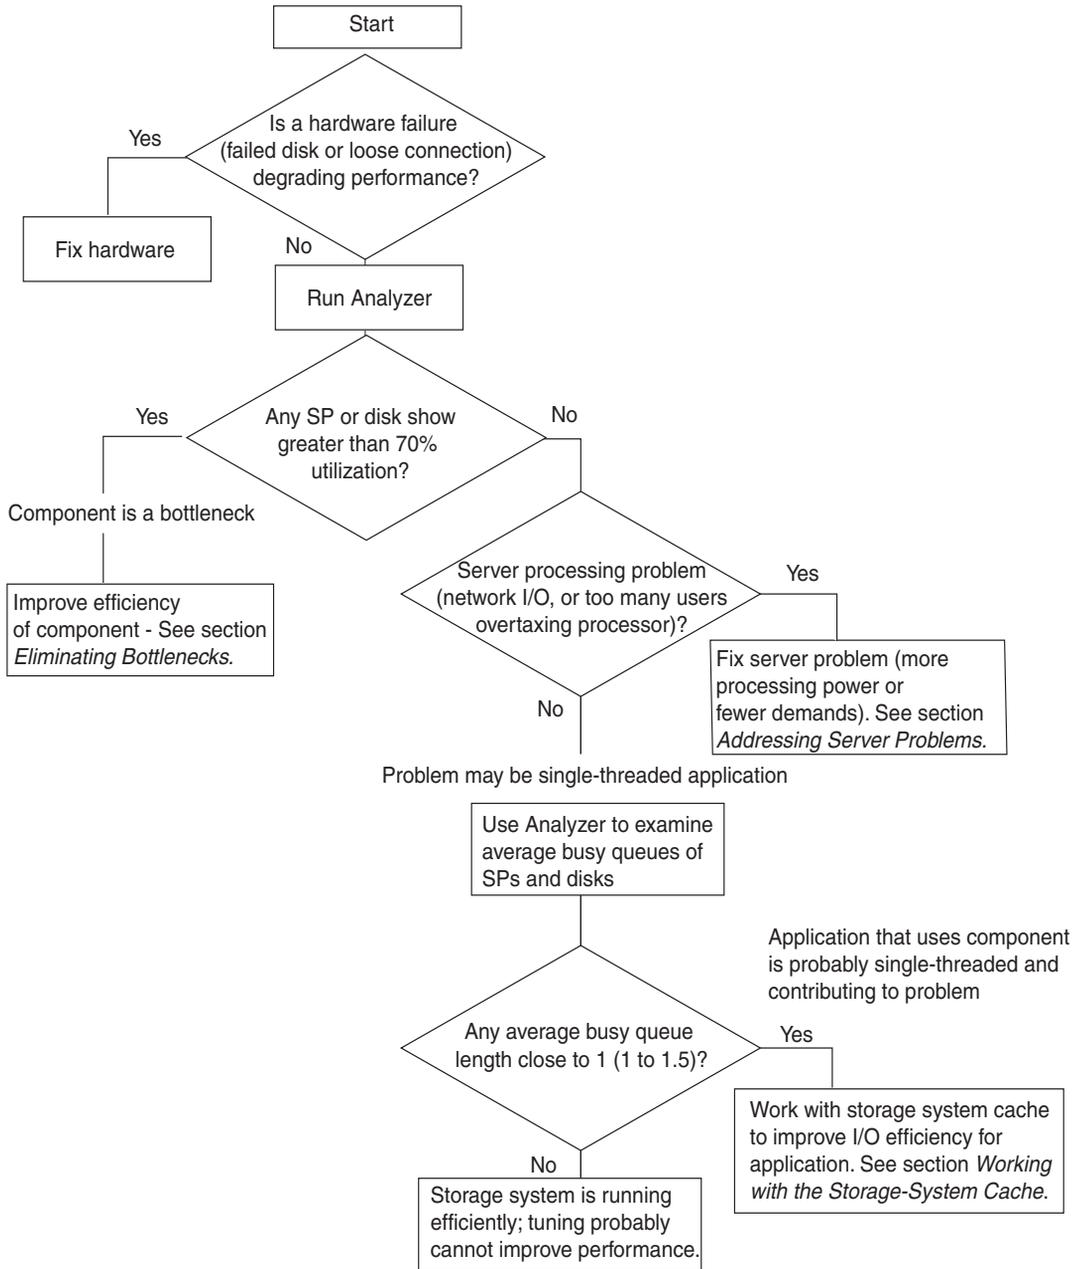


Figure 3-1 Performance Analysis Flowchart

Eliminating Bottlenecks

Generally, if the disks have higher utilization than the SP, then the disks are the bottleneck. You can work with the cache or change the RAID type to push the SP harder, as described later in this section.



CAUTION

Changing the RAID type of a LUN requires you to unbind the LUN, which eliminates all data stored on it. Before unbinding, back up all data that you want on the LUN; then after rebinding with the new RAID type, load the backed-up data.

If the SP utilization is higher than the disks, then the SP is the bottleneck. You should consider moving some load off the SP (perhaps by moving some LUNs to another SP), or changing the schedules of some operations.

When you identify the component with the highest utilization, you identify the performance bottleneck and limiting factor. If this component is running at maximum capacity, system performance is limited by the component and any improvement in system performance requires eliminating the bottleneck. This can be achieved in several ways:

- ◆ Reduce the demand on the component. This might be as simple as scheduling some operations to another time of day. You can reduce disk workload by using caching so that data can be accessed in the cache rather than from disk. You may also be able to reduce workload by choosing a more appropriate RAID type (but see the caution above).
- ◆ Move some workload to another component. You can do this by moving files and databases (or index tables of databases) to different LUNs. If the disk module utilizations within a striped LUN are not equal (within 15% of each other), then the I/O distribution of the striping needs to be improved. Reducing the stripe element size increases the chances of localizing I/O to one disk, but it also can have the detrimental effect of increasing the number of stripe-crossing I/Os.
- ◆ Add components. If Analyzer shows disks as the bottleneck, then you can improve performance by adding more disks or LUNs. But if Analyzer shows an SP as a bottleneck, adding more disks

will not help performance. For example, adding disks will not help if disk utilization is 40% and the SP utilization is 65%, since the SP is the limiting factor.

Disk-Based Factors - RAID Type and Stripe Element Size

You may be able to reduce the effect of a disk bottleneck by using a different RAID type for one or more LUNs:

- ◆ For general-purpose high throughput and response time with high availability, use RAID 5 with five disks and write caching. The RAID 5 stripe-element size should be larger than most I/O sizes, but smaller than the disk buffer size. Ideally, each read request and small write request should be less than the stripe-element size. Small, partial stripe writes, where the I/O sizes are smaller than half the stripe size, are expensive. EMC recommends a stripe element size of 64 Kbytes (128 sectors), the default.
- ◆ For applications that use small, randomly distributed I/Os where write intensity is very high for long enough to saturate the cache, use a mirrored RAID type: RAID 1 or RAID 1/0. RAID 1/0 — and to a lesser extent RAID 1 — offer slightly higher availability than RAID 5 or RAID 3. RAID 1 is suitable for a system disk.
- ◆ For a single-threaded application that performs large, sequential I/Os, consider RAID 3.
- ◆ For throughput and response time without high availability, use RAID 0. RAID 0 yields better write performance than RAID 5 by avoiding the overhead of computing and writing parity data. RAID 0 read performance is identical to RAID 5. The stripe-element size issues are the same as for RAID 5. The number of disks in a single LUN does not influence performance significantly.

Working with the Storage-System Cache

For best performance with most applications, each SP should have its maximum amount of cache memory and you should use the default settings for the cache properties. Analyzer shows how the cache affects the storage system, and lets you tune the cache properties to best suit your application.

A storage-system cache has two parts: a read cache and a write cache. The read cache uses a read-ahead mechanism that lets the storage system prefetch data from the disk. Therefore the data will be ready in the cache when the application needs it. The write cache buffers and optimizes writes by absorbing peak loads, combining small writes, and eliminating rewrites.

You can change read cache size, write cache size, and cache page size to achieve optimal performance. The best sizes of the read and write caches depend on the read/write ratio. A general norm for the ratio of reads to writes is two reads per write; that is, reads represent 66% of all I/Os.

Since the contents of write cache are available for read operations as well, you should allocate most of the available SP memory to the write cache. However, since the write cache is flushed after a certain timeout period, a read cache is also required to hold active data for longer periods of time.

Read cache size: The read cache holds data that is expected to be accessed in the near future. If a request for data that is in the cache arrives, the request can be serviced from the cache faster than from the disks. Each request satisfied from cache eliminates the need for a disk access, reducing disk load. If the workload exhibits a “locality of reference” behavior, where a relatively small set of data is accessed frequently and repeatedly, the read cache can improve performance. In read-intensive environments, where more than 70% of all requests are reads, the read cache should be large enough to accommodate the data set that is most frequently accessed. For sequential reads from a LUN, data that is expected to be accessed by subsequent read requests is read (prefetched) into the cache prior to being requested. Therefore, for optimal performance, the read cache should be large enough to accommodate prefetched data for sequential reads from each LUN.

Write cache size: Write cache serves as a temporary buffer where data is stored temporarily before it is written to the disks. Cache writes are

far faster than disk writes. Also, write-cached data is consolidated into larger I/Os when possible, and written to the disks more efficiently. (This reduces the expensive small writes in case of RAID 5 LUNs.) Also, in cases where data is modified frequently, the data is overwritten in the cache and written to the disks only once for several updates in the cache. This reduces disk load. Thus, the write cache absorbs write data during heavy load periods and writes them to the disks, in an optimal fashion, during light load periods. However, if the amount of write data during an I/O burst exceeds the write cache size, the cache fills. Subsequent requests must wait for cached data to be flushed and for cache pages to become available for writing new data.

The write cache provides sustained write speed by combining sequential RAID 5 write operations and writing them in RAID 3 mode. This eliminates the need to read old data and parity before writing the new data. To take advantage of this feature, the cache must have enough space for one entire stripe of sequential data (typically 64 KBytes \times [*number-of-disks* - 1], or, for a five-disk group, 256 KBytes) before starting to flush. Note that the sequential stream can be contained in other streams of sequential or random data.

Cache page size: This can be 2, 4, 8, or 16 Kbytes. As a general guideline, we suggest 8 Kbytes. The ideal cache page size depends on the operating system and application. Analyzer can help you decide which size performs best.

Using Analyzer to Monitor Cache Performance

Analyzer provides counters to monitor cache performance, as follows.

Write Cache Counters

The write cache counters show forced flush rates, dirty cache pages, cache hit ratio, and high and low watermark flush rates. They are as follows.

Forced flush rate: Forced flushing occurs when a write request arrives and there are no empty pages in the cache. Forced flushes are undesirable since they incur the overhead of cache management with little chance of optimizing writes. A high forced flush rate indicates that the write cache is not large enough to accommodate the modified data. If the workload imposes heavy loads for long periods of time, the cache will eventually run out of empty pages and cause forced

flushes. However, if the workload consists of short bursts, the cache size should be large enough to accommodate the write data during a burst. The cache content will be flushed during the reduced load periods. You might reduce the forced flush rate either by increasing the cache size or by reducing the cache page size to reduce memory fragmentation. As a general rule, if the number of forced flushes is more than half the number of writes, performance will be better without the write cache.

% Dirty pages: Dirty pages are pages that have been modified in the write cache but have not yet been flushed to disk. The dirty pages percentage count should be lower than the high watermark under most situations. A consistently high percentage of dirty pages might indicate a balance between the incoming load and the rate at which cache pages are flushed. A larger cache might be required to absorb any I/O burst. Alternatively, you might resolve the problem by decreasing the cache page sizes to reduce memory fragmentation in cache, if you suspect fragmentation. A high percentage of dirty pages does not necessarily indicate a problem. However, when the forced flush rate is high, the dirty pages percentage is likely to be high as well.

Write cache hit ratio: A high cache hit ratio indicates good performance. A low value might be caused by large I/O sizes, since large writes bypass the cache. In this case, the forced flush rate will be low and cache tuning will not help. A low hit ratio value might also be caused if the cache is too small for the incoming workload. In this case, the forced flush rate will be high. You might resolve this problem by increasing the write cache size, if possible, or by reducing the cache page size to reduce cache memory fragmentation.

High and low watermark flush rates: The high watermark specifies the maximum value for dirty pages beyond which the cache page flushing starts. This is done so that empty cache pages will be available for subsequent write requests. The low watermark specifies the minimum value of dirty pages below which the cache page flushing stops, triggered by the high watermark. (Note that there are other conditions that also trigger cache page flushing, but the low watermark does not stop cache flushing under other conditions). A high value for a high watermark flush rate indicates a heavy incoming workload. However, a high value does not necessarily indicate a performance problem.

The default values for the watermarks yield reasonably good performance under most situations. However, the high and low

watermarks can be useful tuning parameters. A high value for the dirty pages percentage as the high watermark results in less frequent cache page flushing, but leaves fewer empty pages in the cache to accommodate any subsequent I/O burst. A lower value can accommodate I/O bursts better, but causes more frequent cache page flushing. Thus, the loads imposed on the disks are different under the two situations. Likewise, a lower value for the low watermark is better suited to accommodate I/O bursts, but imposes a steady load on the disks. This may influence performance differently, depending on the load imposed on the disks by read requests. A higher value for the low watermark allows the more frequently modified pages to stay in the cache, allowing more data rehits, which in turn reduce the load on the disks. An optimal value for the high and low watermark depends on the nature of the incoming workload.

Read Cache Counters

The read cache counters show prefetch effectiveness, cache hit ratio, and hit rate. These counters are as follows:

% Used prefetches: A high value indicates good performance for sequential reads. In general, it is difficult to detect a sequential stream of references to a LUN if there are other active reference streams to that LUN. However, if there is a single sequential stream addressing a LUN and the Used prefetches percentage is low, this indicates a performance problem. You might improve performance either by increasing the read cache size or decreasing the cache page size.

Read cache hit ratio: The hit ratio is a measure of the efficiency of your tuning and should exceed 20%. The read cache hit ratio should be high, particularly for sequential reads from a LUN. The prefetch occurs at the speed of the disks. If the arrival rate of requests is high and the I/O sizes are large, the prefetch process might not be able to keep up with the load, causing a low cache hit ratio.

For random reads from a LUN, the cache hit ratio will be high only if

- ◆ the requests exhibit good locality of reference (accessing a limited set of data during a phase), and
- ◆ the read cache is large enough to hold the active data set (working set).

For random reads, or if the working set is too large, the read cache will not provide any advantage. Therefore, a low read cache hit ratio indicates a performance problem only if:

- ◆ the access stream is sequential and the load is low, or
- ◆ the access stream is random but has a good locality of reference behavior.

Read cache hit rate: The hit rate shows the impact of caching on your application. It should be high if the read cache hit ratio is high and the workload is relatively heavy. A low cache hit rate with a high cache hit ratio indicates that the workload is not driving the cache to its maximum capacity. If the cache hit ratio is low, the cache hit rate will also be low.

Addressing Server Problems

Workload analysis is often the key to solving problems that appear to be caused by the storage system but that you can solve by changing the way the server accesses the storage system.

Server Throttles

Often, software designers face problems of resource allocation. There are not enough resources to maintain as many parallel activities (paths, threads, tasks, queued requests, and so on) as desired and the designer resorts to an execution throttle. The ultimate effect of these throttles is to limit the number of concurrent requests submitted to the storage system. A storage system can process multiple requests in parallel. By reducing the number of concurrent requests, the storage system reduces the throughput. On average, 16 concurrent threads are required to reach 90% maximum throughput of a five-disk RAID 5 LUN. Reducing concurrency below this level will considerably reduce the system throughput, resulting in longer response times seen by the application.

Analysis of LUN queue lengths can provide the clues needed to diagnose server throttles. You might find that throttles are imposed on each LUN, some LUNs, SPs, buses and even the entire storage system.

Data Alignment

The alignment of large I/Os (larger than 256 Kbytes) and small I/Os (smaller than eight Kbytes) does not normally impact performance. I/O sizes between these can unduly affect performance. With random

I/O, a stripe crossing can require additional seek time. Even though this additional seek time occurs in parallel and does not directly affect response time, it causes extra busy time for a second disk. The second disk is not available for other work; this reduces the maximum throughput of the LUN, which can consequently lead to increased response times.

Misalignment can be caused by file systems, table spaces, and other storage objects not being aligned to stripe boundaries. When this happens, the I/O is misaligned, and in sequential operations, every *n*th I/O will cause a stripe crossing. To resolve this kind of problem, you must find some kind of padding mechanism to realign the objects in question, or change a block or buffering factor to improve alignment.

Sometimes the application has no concept of alignment, and I/Os are generated with random starting addresses. This is called *unaligned I/O* and unless you can create some kind of alignment, there is nothing to be done.

Comparing LUN I/O sizes and disk I/O sizes can give you clues about alignment, as can the **DiskCrossing** counter in the characteristics list. Precise analysis requires a Fibre Channel analyzer or trace facility and consultation with a qualified storage-system engineer.

A

- advanced performance characteristics 2-28
- Agent, about 1-6
- alignment, data 3-13
- Analyzer
 - charts, introduced 2-6
 - menu option 2-2
 - right-click menu options 2-2
 - starting 2-2
- archive file
 - charts 2-4
 - extension 2-45
 - viewing 2-5
- audience for manual 1-ix

B

- bottlenecks 3-4
- bottlenecks, eliminating 3-6

C

- cache
 - page size 3-10
 - performance 3-9, 3-10
- characteristics, performance 2-28
- chart
 - Analyzer, introduced 2-6
 - archive file and real time 2-4
 - data, exporting 2-37
 - display, customizing 2-38
 - IO Detail 2-23
 - menu
 - IO Size Distribution 2-25

- Performance Detail 2-17
 - resizing, zooming 2-34
- comma-separated values (csv) 2-37
- component objects, selecting 2-12
- csv (comma-separated values) 2-37
- customizing chart display 2-38

D

- data, alignment on stripe and performance 3-13
- device menu
 - Performance Detail 2-20
 - Performance Summary 2-12
- dirty pages
 - counter 2-32
 - meaning 3-11
- disk
 - factors in performance 3-8
 - menu 2-3
- domain, introduced 1-2

E

- environments, Manager 1-2
- exporting chart data 2-37

F

- file, archive, extension 2-45
- forced flush rate 3-10

H

- high watermark flush rate 3-11

hit ratio

- read cache 3-12
- write cache 3-11

host performance issues 3-13

I

IO Size Distribution chart 2-23

J

jpeg (.jpg) images 2-25, 2-37

L

light loads 3-4

load, light 3-4

locality of reference 3-4

low watermark flush rate 3-11

LUN (logical unit)

- bottleneck, identifying 2-41
- device names 2-46
- menu 2-3
- operating system device names 2-46
- performance (Performance Survey chart) 2-8
- performance survey chart 2-8
- Survey chart 2-41

M

Manager

- environments 1-2
- starting 2-2

menu

- IO Size Distribution chart 2-25
- option, Analyzer 2-2
- Performance Summary 2-12
- Performance Survey 2-10

menus

- for storage-system components
 - disk 2-3
 - LUN 2-3
 - SPs 2-3

metrics, sorting 2-14

N

nar (.nar) extension to archive filename 2-45

Navisphere Manager, see Manager

O

operating system device names of storage-system LUNs 2-46

P

page

- dirty 3-11
- size, cache 3-10

performance

- characteristics
 - advanced 2-28
 - characteristics, basic 2-28
 - data, understanding, see also Chapter 3
 - problem, identifying 3-5
 - survey chart 2-8

Performance Detail chart 2-17

Performance Summary chart 2-11

Performance Survey chart 2-8

- menu 2-10

pictures from charts (exporting) 2-37

printing graphic displays 2-37

Q

queue length characteristic

- counter 2-30
- using 3-3

R

RAID factors in performance 3-8

read cache

- counters 2-31
- performance 3-9

real-time charts 2-4

resizing charts 2-34

right-click menu options 2-2

S

server

- performance issues 3-13
- throttle 3-13

sorting based on latest or maximum values 2-38

sorting by performance characteristics 2-8

sorting metrics 2-14

SP (storage processor) menu 2-3

starting Analyzer 2-2
storage system
 LUN device names 2-46
stripe size factors and performance 3-8

T

thresholds
 LUN Survey chart 2-11
 useful values 2-41
throttle, server 3-13
tool tips 2-21

U

utilization characteristic

counter 2-28
using 3-3, 3-5, 3-7

W

watermark flush rates 3-11
workload, light 3-4
workload, understanding 3-2
write cache
 counters, using 3-10
 performance 3-9
 size 3-9

Z

zooming charts 2-34

