

**EMC Navisphere
Analyzer
Version 6.X**

ADMINISTRATOR'S GUIDE

**P/N 069001158
REV A04**

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.EMC.com

Copyright ©2002, 2003 EMC Corporation. All rights reserved.

Published December, 2003

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, CLARiiON, Navisphere, and PowerPath are registered trademarks and Access Logix, Application Transparent Failover, EMC ControlCenter, FLARE, MirrorView, SAN Copy, and SnapView are trademarks of EMC Corporation.

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

Preface	ix
 Chapter 1 About the EMC Navisphere Analyzer	
About Analyzer	1-2
Analyzer Dependencies	1-3
 Chapter 2 Analyzer Session	
Starting a Navisphere Manager or Analyzer Session.....	2-2
Displaying Data	2-3
Summary	2-6
 Chapter 3 Analyzing Performance	
Discovering LUN Device Names	3-2
Menu Bar Options for Analyzer	3-3
Real-Time and Archive File Display	3-7
Analyzer Charts	3-9
Performance Survey Chart	3-11
Performance Summary Chart	3-14
Performance Detail Chart	3-21
IO Size Distribution Charts (LUN Only).....	3-28
LUN IO Disk Detail Charts (LUN Only).....	3-34
Performance Characteristics.....	3-38
Zooming and Resizing Charts	3-45
Exporting and Printing Charts.....	3-48
Customizing Chart Display.....	3-50
Capturing Performance Data in Archive Files	3-55
Displaying Archive File Data	3-56

Chapter 4	Understanding Performance Data	
	Resetting Performance Logging	4-2
	Understanding Workload	4-3
	Identifying a Performance Problem	4-6
	Eliminating Bottlenecks	4-8
	Working with the Storage-System Cache	4-10
	Addressing Server Problems	4-14
Chapter 5	Command Line Interface (CLI) and Client Tools	
	navicli.jar	5-2
	analyzer -set	5-6
	analyzer -logging	5-7
	archivedump.jar	5-8
	archiveretrieve.jar	5-16
Appendix A	Customer Support	
	Overview of Detecting and Resolving Problems	A-2
	Troubleshooting the Problem	A-3
	Before Calling the Customer Support Center	A-4
	Documenting the Problem	A-5
	Reporting a New Problem	A-6
	Sending Problem Documentation	A-7

3-1	Menu Bar Options for Analyzer	3-3
3-2	Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks, Storage Groups, RAID Groups, Storage Systems	3-4
3-3	Performance Survey Menu Options	3-13
3-4	Performance Summary Chart Tree Root Node Right-Click Menu Options	3-16
3-5	Performance Summary Chart Tree Panel Right-Click Menu Options	3-17
3-6	Performance Summary Chart Panel Right-Click Menu Options	3-17
3-7	Metric Point Right-Click Menu Options	3-18
3-8	Performance Detail Chart Tree Root Node Right-Click Menu Options	3-25
3-9	Performance Detail Chart Tree Right-Click Menu Options	3-25
3-10	Performance Detail Chart Right-Click Menu Options	3-26
3-11	IO Size Distribution Summary Chart Left Panel Right-Click Menu Option	3-30
3-12	IO Size Distribution Summary Chart Right-Click Menu Options	3-30
3-13	IO Size Detail Chart Left Panel Right-Click Menu Options	3-32
3-14	IO Size Distribution Detail Chart Right-Click Menu Options	3-32
3-15	LUN IO Disk Detail Chart Tree Root Node Right-Click Menu Options	3-36
3-16	LUN IO Disk Detail Chart Tree Right-Click Menu Options	3-36
3-17	LUN IO Disk Detail Chart Right-Click Menu Options	3-36
3-18	Basic Performance Characteristics (Alphabetical Order)	3-38
3-19	Advanced Performance Characteristics (Alphabetical Order)	3-40

3-1	Analyzer Tree Controls	3-10
3-2	Sample Performance Survey Chart	3-12
3-3	Sample Performance Summary Chart	3-15
3-4	Performance Summary Chart Before Sort	3-19
3-5	Performance Summary Chart After Sort by Maximum	3-20
3-6	Sample Performance Detail Chart Showing Utilization of SP and LUN	3-22
3-7	Performance Detail Display Before Highlight	3-23
3-8	Performance Detail Display After Highlight	3-24
3-9	Sample IO Size Distribution Summary Chart	3-29
3-10	Sample IO Size Distribution Detail Chart	3-31
3-11	Sample LUN IO Disk Detail Chart	3-35
3-12	Chart Before Zoom	3-46
3-13	Chart With Zoom Area Selected (Left Mouse Button)	3-47
3-14	Chart After Zoom (Release of Left Mouse Button)	3-48
4-1	Performance Analysis Flowchart	4-7
A-1	Problem Detection and Resolution Process	A-2

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this manual may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this manual, please contact your EMC representative.

About This Manual

This manual describes the tasks for monitoring storage system performance using EMC Navisphere® Analyzer, a member of the EMC ControlCenter™ family of products. Each major section includes introductory information and a general procedure for completing a task. This manual is not intended for use during the actual setup of monitoring storage systems so the steps in the procedures purposely do not include screen captures of the dialog boxes.

The introductory information and *detailed* steps for each procedure appear in the Navisphere Analyzer online help so you have the complete information available when you actually set up, configure, and manage storage systems, should you require help.

Audience

This guide is part of the Navisphere documentation set, and is intended for use by those who want to monitor and tune disk-array storage systems.

Readers of this guide are expected to be familiar with the following topics:

- ◆ The operating system running on the servers and the concepts and operation of the storage systems.
- ◆ Navisphere Manager Agent selection and tree structure and how to use them. They are explained in the *EMC Navisphere Manager Administrator's Guide*.

The manual assumes the following:

- ◆ 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 Navisphere Analyzer Installation Guide*.

Organization

- | | |
|------------|---|
| Chapter 1 | Introduces Navisphere Analyzer. |
| Chapter 2 | Leads you through a session with Analyzer. |
| Chapter 3 | Details features of Analyzer. |
| Chapter 4 | Explains how to interpret Analyzer performance data. |
| Chapter 5 | Describes the Analyzer command line interface (CLI) and client tools. |
| Appendix A | Reviews the EMC process for detecting and resolving software problems, and provides essential questions that you should answer before contacting the EMC Customer Support Center. |

Related Documentation

Related documents include:

EMC Navisphere Analyzer 6.X Installation Guide, P/N 069001182

EMC Navisphere Manager Version 6.X Administrator's Guide, P/N 069001125

EMC FC4700-2 Storage System Configuration Planning Guide, P/N 014003087

EMC Storage Systems CX400-Series and CX600-Series Configuration Planning Guide, P/N 014003113

EMC CLARiiON CX300, CX500, and CX700 Storage Systems Configuration Planning Guide, P/N 300-001-273

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.

Typographical Conventions.

This manual uses the following format conventions:

This typeface	Indicates text (including punctuation) that you type verbatim, all commands, pathnames, filenames, and directory names. It indicates the name of a dialog box, field in a dialog box, menu, menu option, or button.
<i>This typeface</i>	Represents variables for which you supply the values; for example, the name of a directory or file, your username or password, and explicit arguments to commands.
<code>This typeface</code>	Represents a system response (such as a message or prompt), a file or program listing.
x > y	Represents a menu path. For example, Operations > Poll All Storage Systems tells you to select Poll All Storage Systems on the Operations menu.
[]	Encloses optional entries.
	Separates alternative parameter values; for example: <i>LUN-name LUN-number</i> means you can use either the LUN-name or the LUN-number.

Finding Current
Information

The most up-to-date information about the EMC Navisphere Analyzer software is posted on the EMC Powerlink website. We recommend that you download the latest information before you start the Analyzer software.

To access EMC Powerlink, use the following link:

<http://powerlink.emc.com>

After you log in, select **Support > Document Library** and find the following:

- ◆ *EMC Navisphere Analyzer 6.X Release Notes*
- ◆ The latest version of this manual that is applicable to your software revision
- ◆ *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems*, which provides a checklist of the tasks that you must complete to install your storage system in a storage area network (SAN) or direct attach configuration.

Where to Get Help

For questions about technical support, call your local sales office or 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 and select the applicable 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 Analyzer1-2
- ◆ Analyzer Dependencies1-3

About Analyzer

Navisphere Analyzer works within storage domains that the web-based Navisphere Manager administers. It gathers storage-system performance statistics and presents them in various types of charts. 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 user interface (UI).

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. This manual uses the term LUN to refer to LUNs and metaLUNs, except where explicitly noted.

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 the 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).

Analyzer Dependencies

Analyzer requires the following:

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

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

What Next?

For a sample session with Analyzer, continue to Chapter 2. For Analyzer details, skip to Chapter 3.

This chapter leads you through a sample session with Analyzer.
Topics are

- ◆ Starting a Navisphere Manager or Analyzer Session2-2
- ◆ Displaying Data.....2-3
- ◆ Summary2-6

Starting a Navisphere Manager or Analyzer Session

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

While you are running Analyzer, you can get context-sensitive help at any time by pressing the **Help** button in the dialog box. For general help, click the **Help** button on the Navisphere toolbar.

1. Use a supported browser to point to the IP address of a storage system or management station that has the Navisphere Management Server and the Management UI installed.
2. After a few moments, the software displays a login dialog box.
3. Log in to 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. Continue to the next section.

Displaying Data

Analyzer can display data in real time, or it can display an archive file containing past performance data. You can capture performance data in an archive file at any time and store it where you want on the host to which you connected with Analyzer.

To Display Real-time Data

1. Log in to 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 an SP.

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

Also, statistics logging must be enabled for both SPs. For information on enabling statistics logging, see the Navisphere Manager manual.

4. Right-click the selected item.
5. On the pull-down menu, select **Performance Detail**.

After pausing for a few moments to gather data, Analyzer starts to display a chart showing the activity on the LUNs attached to the SP. Analyzer displays new values as it accumulates data.

6. To close the chart, click the close (X) button at the upper right. Closing a chart discards all accumulated data.

Continue to the next section for information on archive data.

Capturing Performance Data in Archive Files

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

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 copy of the 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.

To capture performance data in an archive file:

1. From the Analyzer menu bar, choose **Tools > Analyzer > Retrieve Archive**.

Analyzer displays the **Retrieve Archive** dialog box that lets you specify the location and name for the file.

2. Select the storage system or legacy array management station from the **Retrieve Archive From** box.
3. To specify from which SP(s) you want to retrieve data, click the SP A and/or SP B checkboxes.
4. For now, take the default for the directory name in the **Save File As** field.
5. Click **Retrieve**.

Analyzer starts recording performance data from the specified SP to the file you specified.

6. Watch the **Retrieve Status** section until it shows **Done**.
7. Click **Done** to close the dialog box.

Customizing Archive Display

By default, when you open an archive, Analyzer displays a Performance Survey chart. You can change the default Archive display and the settings for other chart types using the **Customize Display** dialog box.

To change the default Archive display:

1. Select **Tools > Analyzer > Customize**.
Analyzer displays the **Customize** dialog box.
2. Click the **Archives** tab.
Analyzer displays the **Archives** dialog box.
3. Select **Performance Summary**, and click **OK**.

You have changed the default display from Performance Detail to Performance Summary.

Opening an Archive and Displaying Archive Data

For archive file display, Analyzer requires access to at least one Analyzer archive file. Archive files are stored on the host 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.

1. On the Analyzer menu bar, select **Tools > Analyzer > Open Archive**.
2. From the standard file selection dialog box, select an archive file and click **OK**. You will see a display selection dialog box.
3. Select the entire archive period by clicking **OK**.

Analyzer displays a Performance Summary chart using the archived performance data.

Displaying Other Chart Types

Analyzer has six types of chart:

- ◆ Performance Survey (LUNs only)
- ◆ Performance Summary
- ◆ Performance Detail
- ◆ IO Size Distribution Summary (for LUNs only)
- ◆ IO Size Distribution Detail (LUNs only)
- ◆ LUN IO Disk Detail

To choose a chart type

1. In the LUNs tree, right-click the first LUN.
Analyzer displays a menu with chart selections.
2. Click the chart type you want, for example, **Performance Summary**.

Analyzer displays the specified chart type.

3. Close the chart.

Continue to the next section to close the archive.

Closing an Archive

1. On the Analyzer menu bar, select **Tools > Analyzer > Close Archive**.
2. In the checkbox(es), select the file(s) you want to close.
3. Click **Close Archive**.

Closing an archive file also closes all displayed charts based on that file.

Summary

You have logged in, displayed a real-time chart, retrieved archive data in a file, customized a display, selected a different chart display, and closed an archive.

For details on these operations, and information on other Analyzer operations such as printing charts and exporting chart data, continue to Chapter 3.

This chapter explains the details of using Analyzer. It expands on topics covered in Chapter 2 and introduces some new ones. Topics are

◆ Discovering LUN Device Names.....	3-2
◆ Menu Bar Options for Analyzer	3-3
◆ Real-Time and Archive File Display.....	3-7
◆ Analyzer Charts	3-9
◆ Performance Survey Chart	3-11
◆ Performance Summary Chart.....	3-14
◆ Performance Detail Chart	3-21
◆ IO Size Distribution Charts (LUN Only)	3-28
◆ LUN IO Disk Detail Charts (LUN Only)	3-34
◆ Performance Characteristics.....	3-38
◆ Zooming and Resizing Charts.....	3-45
◆ Exporting and Printing Charts.....	3-48
◆ Customizing Chart Display	3-50
◆ Capturing Performance Data in Archive Files.....	3-55
◆ Displaying Archive File Data	3-56

Discovering LUN Device Names

Analyzer can create a file that provides useful information about the storage-system LUNs that Navisphere Manager sees, 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 standard file selection dialog box. 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 Windows, formatted: **E:**

- ◆ For Windows, unformatted: \\.\PhysicalDrive1
- ◆ For Solaris: **c2t1d1** (/dev/(r)dsk)

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 RAID 0, RAID 1/0, RAID 1, RAID 3, RAID 5, or Disk (individual disk unit).

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

```
c2t1d1 F20010500278_LUN6 [RAID5]
```

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 choices shown in Table 3-1.

Table 3-1 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 currently opened archive files. If multiple files are open, it displays a list 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.

Table 3-1 Menu Bar Options for Analyzer (*continued*)

Close archive	Close all charts displaying data from currently open archive files, and release 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.
Retrieve archive	Capture the performance information recorded by the selected storage system and open 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.
Data Logging	Set the polling interval for archive files.
Customize	Change display items such as colors, legends, and the initial screen Analyzer displays.

Analyzer can display performance information on LUNs, metaLUNs, SPs, disks, and SnapView™ snapshot sessions, if any. The options available when you right-click a LUN, metaLUN, SP, or disk icon, and select **Analyzer**, are shown in Table 3-2.

Table 3-2 Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks, Storage Groups, RAID Groups, Storage Systems

Icon Description	Menu Option	Displays
SP A or SP B	Performance Survey	Performance Survey window (performance overview).
	Performance Summary	Performance Summary window (performance range).
	Performance Detail	Performance Detail window (line chart).
	Logging > Reset (Administrator or Manager only)	Resets performance logging properties to their default values and restarts performance logging.

Table 3-2 Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks, Storage Groups, RAID Groups, Storage Systems (*continued*)

Icon Description	Menu Option	Displays
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.
	LUN IO Disk Detail	LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
metaLUN <i>lunID</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).
Storage Group SG <i>storage group number</i>	Performance Survey	Performance Survey window (performance overview).
	Performance Summary	Performance Summary window (performance range).
	Performance Detail	Performance Detail window (line chart).

Table 3-2 Enterprise Storage Main Window Right-Click Options for SPs, LUNs, Disks, Storage Groups, RAID Groups, Storage Systems (*continued*)

Icon Description	Menu Option	Displays
RAID Group <i>ID</i> [<i>RAID type</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).
	LUN IO Disk Detail	LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
Storage System	Performance Survey	Performance Survey window (performance overview).
	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 online help index or the documentation 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 (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 to 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 in to) a storage system that has Manager and Analyzer installed.

To Display Real-Time Data

1. Log in to 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, metaLUNs, SPs, disks, hosts, RAID Groups, or Storage Groups in the Enterprise **Storage** tree view. Generally, LUNs and SPs provide broader workload metrics than disks.

For performance data to be gathered on these storage components, 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 **Analyzer** and then one of the following choices: **Performance Survey** (LUN only); **Performance Summary**; **Performance Detail**; **IO Size Distribution Summary** (LUN only); **IO Size Distribution Detail** (LUN only); or **LUN IO Disk Detail** (LUN only; not metaLUN).

After pausing for a few moments to gather data, during which time status messages appear on the status bar at the bottom of the chart window, Analyzer will start to display a chart of the selected type for the selected components.

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 retrieve and open archive files (described later) as needed.

Changing the Polling Interval for Archive Files

If you have Manager or Administrator privileges, you can change the polling interval for archive files. The polling intervals automatically return to their default values after 36 hours, when the SP is rebooted or the Management Server is restarted.

Changing the interval affects the duration of data. The default of 10-minute intervals results in about 26 hours of data. Changing the interval to 5 minutes results in about 13 hours. Likewise, changing the interval to 20 minutes will double the duration.

1. On the Analyzer menu bar, select **Data Logging**.
2. In the **Target** box, select the target system from a list of systems in the domain that have Analyzer Provider installed.
3. In the **Values** section, type the polling value (number of seconds) you want in the **Archive Interval** field, and click **OK**.

Setting the value to something other than the default value may have significant performance impact.

If you are not a manager or administrator, when you click **Apply** or **OK**, a message states that you are not privileged to make the change, and no change occurs to the polling interval.

Analyzer Charts

This section explains Analyzer tree controls and chart types. Note that some of the screen captures shown here may 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:

- ◆ 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.

Figure 3-1 shows the three control types.




Figure 3-1 Analyzer Tree Controls

Chart Types

Analyzer has six types of chart:

- ◆ Performance Survey (LUN only), page 3-11
- ◆ Performance Summary, page 3-14
- ◆ Performance Detail, page 3-21
- ◆ IO Size Distribution Summary (for LUNs only), page 3-28
- ◆ IO Size Distribution Detail (LUNs only), page 3-31
- ◆ LUN IO Disk Detail (LUNs only), page 3-34

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. When you select a chart, a status bar appears at the bottom of the chart window. A status message appears on the left side of the bar, and a small icon appears on the right side of the bar that will move when an action occurs.

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 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 3-2 shows a sample Performance Survey chart.




Figure 3-2 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 Values” on page 3-51.

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 in Table 3-3.

Table 3-3 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.
IO Size Distribution Detail Chart	Line chart showing numbers of I/Os of various sizes that occurred with the LUN over time.
LUN IO Disk Detail Chart	LUN IO Disk Detail chart for the RAID Group that belongs to the LUN. (LUN only; not metaLUN)
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 3-7.

A sample Performance Summary chart for a LUN, an SP, and the LUN's constituent disks is shown in Figure 3-3.




Figure 3-3 Sample Performance Summary Chart

The tree has a **LUN** tab and a **RAID Group** tab, and each tab contains a checkbox tree. If you check an object on one of the tabs, the object in the other tab is also checked. The **LUN** tab contains individually selectable items, with **LUNs** as the tree root node and individual metaLUNs and LUNs as the top-level branches. Each metaLUN has individual component nodes that contain individual LUN nodes. Component nodes do not have a checkbox since they do not have performance information. MetaLUN nodes can also have an SP node and SnapView session node. You can select an item by clicking it. To remove an item's values from the chart, clear the checkbox.

The **RAID Group** tab contains items with and without checkboxes depending on the availability of performance information. The tree is organized with RAID Groups as the top-level branches, LUNs as

the second-level, and SPs, snap sessions, and disks as leaves. MetaLUNs do not appear in this tab.

The performance characteristics list lets you select any single listed characteristic; the one shown is Queue Length. The Performance Detail chart, described next, lets you select multiple characteristics. All the performance characteristics are explained starting on page 3-38.

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, print, or sort a chart.

The three right-click menu options are shown in Tables 3-4 through 3-6.

Table 3-4 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).
LUN IO Disk Detail	LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
Select All	You can select all Items - LUNs, Disks, SPs, or Snap Sessions - if at least one Snapshot session appears in the LUN tree.
Deselect All	You can deselect all Items - LUNs, Disks, SPs, or Snap Sessions - if at least one Snapshot session appears in the LUN tree.
Help	An explanation of the chart.

Table 3-5 Performance Summary Chart Tree Panel Right-Click Menu Options

Menu Option	Use to
Performance Survey	Display the Performance Survey chart.
Performance Summary	Display the Performance Summary chart.
Performance Detail	Display the Performance Detail chart (line chart).
IO Size Distribution Summary	Display a line chart showing numbers of IOs of various sizes that occurred with the LUN.
IO Size Distribution Detail	Display a line chart showing numbers of IOs of various sizes that occurred with the LUN over time, displayed at each interval.
LUN IO Disk Detail	Display the LUN IO Disk Detail chart for the RAID Group that belongs to the LUN. (Not for metaLUNs.)
Remove Object	Delete the object (and any components) from the tree and chart.
Properties	Display information on the device properties, such as model and capacity.
Help	Display an explanation of the chart.

Table 3-6 Performance Summary Chart Panel Right-Click Menu Options

Menu Option	Use to
Export Chart	Export Chart Data to File, Chart Data to Clipboard, Graph to File, or Graph to Clipboard. For details on exporting chart data, see <i>Exporting and Printing Charts</i> on page 3-48.
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 a metric point, Analyzer displays a pop-up menu with the following choices:

Table 3-7 Metric Point Right-Click Menu Options

Menu Option	Use to
Expand (LUN only)	Select all objects that belong to the LUN.
Collapse (LUN only)	Deselect all objects that belong to the LUN.
Sort	Sort the displayed metrics by minimum, maximum, average, latest, or default values. Sorts first by object, all LUNs, SP, disks together, and then by what you selected.
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 a chart data point. Analyzer displays a pop-up (explained above) from which you can choose Sort. Analyzer then displays the Sort dialog box 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.

Figures 3-4 and 3-5 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 3-4 Performance Summary Chart Before Sort




Figure 3-5 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. Figure 3-6 shows a sample line chart showing Utilization of the SP, the LUN, and the disks in the LUN.

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 3-45.

To display a Performance Detail Chart, use the steps explained in section *Real-Time and Archive File Display* on page 3-7.

A sample Performance Detail chart is shown in Figure 3-6.




Figure 3-6 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's 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.

Figures 3-7 and 3-8 show a Performance Detail display before and after the LUN 6 Total Bandwidth metric has been highlighted.




Figure 3-7 Performance Detail Display Before Highlight




Figure 3-8 Performance Detail Display After Highlight

Performance Detail Menus

The Detail chart tree has a **LUN** tab and a **RAID Group** tab, and each tab contains a checkbox tree. If you check an object on one of the tabs, the object in the other tab is also checked. The **LUN** tab contains individually selectable items, with LUNs as the tree root node and individual metaLUNs and LUNs as the top-level branches. Each metaLUN has individual component nodes that contain individual LUN nodes. Component nodes do not have a checkbox since they do not have performance information. MetaLUN nodes can also have an SP node and SnapView session node. You can select an item by clicking it. To remove an item's values from the chart, clear the checkbox.

The **RAID Group** tab contains items with and without checkboxes depending on the availability of performance information. The tree is organized with RAID Groups as the top-level branches, LUNs as the second-level, and SPs, SnapView sessions, and disks as leaves. MetaLUNs do not appear in this tab.

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 shown in Tables 3-8 through 3-10.

Table 3-8 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.
LUN IO Disk Detail	LUN IO Disk Detail chart for the RAID Group that belongs to the LUN. (Not for metaLUNs.)
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 3-9 Performance Detail Chart Tree Right-Click Menu Options

Menu Option	Use to
Select All	Check all checkboxes in the tree.
Deselect All	Clear all checkboxes in the tree.
Help	Display an explanation of the chart.

Table 3-10 Performance Detail Chart Right-Click Menu Options

Menu Option	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 Chart	Expand to a menu with Export Chart Graph to File, Chart Graph to Clipboard, Data to File, or Data to Clipboard. For details on exporting chart data, see <i>Exporting and Printing Charts</i> on page 3-48.
Freeze/Thaw	Freeze the chart in its current state or return (thaw) a chart window to an active state with the tree and performance characteristics list enabled.
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. 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 Width 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 (**Triangle**, **Square**, and so on, as selected above). You can choose values from 1 through 9 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 chart is a line chart which shows read, write, and total IO sizes. The Detail chart can include as many as 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 3-7.

Figure 3-9 shows an IO Size Distribution Summary chart. A detail chart appears on page 3-31.




Figure 3-9 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 3-11 shows the choices that appear when you right-click in the left panel of the chart.

Table 3-11 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	Show the number of I/Os of various sizes over time, displayed at each interval.
LUN IO Disk Detail (Disabled for metaLUNs)	Show a LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
Select All	Select all values (all objects), reads (all read IO sizes), writes (all write IO sizes), or totals (all total IO sizes) in the IO Size list.
Deselect All	Deselect all values (all objects), reads (all read IO sizes), writes (all write IO sizes), or totals (all total IO sizes) in the IO Size list.
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 choices listed in Table 3-12.

Table 3-12 IO Size Distribution Summary Chart Right-Click Menu Options

Menu Option	Use to
Export Chart	Expand to a menu that lets you export the Chart Graph to File, Chart Graph to Clipboard, Data to File, or Data to Clipboard. For details on exporting and printing charts, see page 3-48.
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 is shown in Figure 3-10.




Figure 3-10 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 3-45.

IO Size Distribution Detail Chart Menus

To display the tree menu, right-click anywhere in the Detail chart left panel. As with other charts, clicking in the left panel gives you a choice as shown in Table 3-13.

Table 3-13 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	Show the number of I/Os of various sizes.
LUN IO Disk Detail (Disabled for metaLUNs)	Show LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
Select All	Select all values (all objects), reads (all read IO sizes), writes (all write IO sizes), or totals (all total IO sizes) in the IO Size list.
Deselect All	Deselect all values (all objects), reads (all read IO sizes), writes (all write IO sizes), or totals (all total IO sizes) in the IO Size list.
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 choices are listed in Table 3-14.

Table 3-14 IO Size Distribution Detail Chart Right-Click Menu Options

Menu Option	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 3-27 for details.
Freeze/Thaw	Freeze the chart in its current state or return (thaw) a chart window to an active state with the tree and performance characteristics list enabled.

Table 3-14 IO Size Distribution Detail Chart Right-Click Menu Options (*continued*)

Menu Option	Use to
Export Chart	Expand to a menu that lets you choose to export Chart Graph to File, Chart Graph to Clipboard, Data to File, or Data to Clipboard. For details on exporting, see <i>Exporting and Printing Charts</i> on page 3-48.
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 3-18.

LUN IO Disk Detail Charts (LUN Only)

The LUN IO Disk Detail chart shows the portion of disk performance that is the result of activity on a specific LUN. This chart is most useful when multiple LUNs occupy a single RAID Group and you need to determine which LUNs are having the greatest effect on the disks' performance. The values shown are either those since the Analyzer started receiving data (for a real-time chart) or those at the time data was recorded in the archive file (in archive-file chart).

You can select multiple performance properties at a time.

The LUN IO Disk Detail chart displays the RAID Group that belongs to the LUN.

The chart has a RAID Group tree, Performance Characteristic list, and a chart display area.

To display a LUN IO Disk Detail chart, use the steps explained in section *Real-Time and Archive File Display* on page 3-7.

Figure 3-11 shows a sample LUN IO Disk Detail chart.




Figure 3-11 Sample LUN IO Disk Detail Chart

LUN IO Disk Detail Menus

The LUN IO Disk Detail chart provides three types of right-click menu. Right-click in the tree panel root node to choose a different kind of chart (as listed in Table 3-15); right-click in the tree panel outside of root node to select and deselect objects (as shown in Table 3-16); or right-click in the chart area to export or print a chart (as shown in the list of menu options in Table 3-17).

Table 3-15 LUN IO Disk Detail Chart Tree Root Node Right-Click Menu Options

Menu Option	Use to
Performance Survey	Display Performance Survey chart.
Performance Summary	Display Performance Summary chart.
Performance Detail	Display Performance Detail chart (line chart).
LUN IO Disk Detail (LUN only)	Show LUN IO Disk Detail chart for the RAID Group that belongs to the LUN.
Remove Object	Confirm that you have deleted the object (and any components) from the tree and chart.
Properties	Obtain information on the device properties, such as model and capacity.
Help	Display an explanation of the chart.

Table 3-16 LUN IO Disk Detail Chart Tree Right-Click Menu Options

Menu Option	Use to
Select All	Check all checkboxes in the tree. Appears on the root node as well as on individual RAID Groups and LUNs.
Deselect All	Clear all checkboxes in the tree. Appears on the root node as well as on individual RAID Groups and LUNs.
Help	Display an explanation of the chart.

Table 3-17 LUN IO Disk Detail Chart Right-Click Menu Options

Menu Option	Use to
Reset Zoom	Restore the display of a chart zoomed on an X or Y axis 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.
Freeze/Thaw	Freeze the chart in its current state or return (thaw) a chart window to an active state with the tree and performance characteristics list enabled.

Table 3-17 LUN IO Disk Detail Chart Right-Click Menu Options (*continued*)

Menu Option	Use to
Export Chart	Expand to a menu with Export Chart Graph to File, Chart Graph to Clipboard, Data to File, or Data to Clipboard. For details on exporting chart data, see <i>Exporting and Printing Charts</i> on page 3-48.
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 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.

The two categories of performance characteristics 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 3-18 and the Advanced Characteristics in Table 3-19, in *alphabetical order*.

The characteristics in these tables are followed by the letters D (applies to disks), S (applies to SPs), L (applies to LUNs), M (applies to metaLUNs), and/or SS (applies to SnapView snapshots). For example, Utilization (D/S/L/M) means that the Utilization characteristic applies to disks, SPs, LUNs, and metaLUNs.

Table 3-18 Basic Performance Characteristics (Alphabetical Order)

Characteristic	Description
Average Seek Distance (GB) (D)	Average seek distance in Gbytes. Longer seek distances result in longer seek times and therefore higher response times. Defragmentation might help to reduce seek distances.
Queue Length (D/S/L/M)	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.
Read Bandwidth (MB/s) (D/S/L/M)	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/M)	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 (Mbytes/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.

Table 3-18 Basic Performance Characteristics (Alphabetical Order) (*continued*)

Characteristic	Description
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.
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 from the source LUN during this snapshot session. It is calculated by the difference between Total Reads in Session and Reads From Cache.
Response Time (ms) (D/S/L/M)	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.
Service Time (ms) (D/S/L)	Time, in milliseconds, that a component spends servicing a request. 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 (Mbytes/s).
Total Bandwidth (MB/s) (D/S/L/M)	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/M)	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.
Utilization (D/S/L/M)	The fraction of a certain observation period when 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, a high LUN utilization value does not necessarily indicate that the LUN is approaching its maximum capacity.
Write Bandwidth (MB/s) (D/S/L/M)	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.

Table 3-18 Basic Performance Characteristics (Alphabetical Order) (continued)

Write Size (Kbytes) (D/S/L/M)	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 (Mbytes/second). For a finer distinction of I/O sizes, use an IO Size Distribution chart.
Write Throughput (IO/s) (D/S/L/M)	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.
Writes To Snapshot Source LUN (SS)	With the optional SnapView software only: The number of writes to the source LUN (on the pertinent SP) during this snapshot session.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)

Characteristic	Description
Average Busy Queue Length (D/S/L/M)	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 Gbytes. Longer seek distances result in longer seek times and therefore higher response times. Defragmentation might help to reduce seek distances.
Chunks Used in Snapshot Copy Session (SS)	The number of chunks that this session has used.
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.
Disk Crossing (%) (L/M)	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.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)
(continued)

Characteristic	Description
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.
Forced Flushes/ s (L/M)	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 should spread the load over more disks.
Full Stripe Writes/s (L/M)	Average number of write requests per second that spanned a whole stripe (all disks in a LUN). This performance data is applicable only to LUNs that are part of a RAID 5 or RAID 3 Group. For these RAID types, full stripe writes are most efficient since data and parity can be written out to the disks without having to pre-read any old data or parity first.
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.
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.
LUN Read Crossings/s (M)	The number of LUN crossings per second that a read request to a metaLUN caused. Since metaLUNs consist of multiple LUNs, a single read request can access disk drives that belong to two or more conventional LUNs.
LUN Write Crossings/s (M)	The number of LUN crossings per second that a write request to a metaLUN caused. Since metaLUNs consist of multiple LUNs, a single write request can access disk drives that belong to two or more conventional LUNs.
MBs Flushed/ s (MB/s) (S)	The number of Mbytes per second written from the write cache to the disks. The value is a measure of back-end activity.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)
(continued)

Characteristic	Description
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.
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 Cache Hit Ratio (L)	The fraction of read requests served from both read and write caches versus the number of read requests to this LUN. The higher the ratio, the better the read performance.
Read Cache Hits/s (L/M)	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/M)	The rate of read requests that the SP cache could not satisfy and therefore required a disk access.
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 (Mbytes/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.
Reads from Read Cache/s (L/M)	Average number of read requests per second that were satisfied by the read cache only. Reads from read cache occur when data that has been recently read or prefetched is re-read while it is still in the read cache. This is a subset of read cache hits which includes requests satisfied by either the write or the read cache.
Reads from Snapshot Cache (SS)	With the optional SnapView software only: During this session, the number of reads that have resulted in a read from the snapshot cache (instead of a read from the source LUN).
Reads from Snapshot Copy (SS)	With the optional SnapView software only: The number of read requests on snapshot copies during this snapshot session.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)
(continued)

Characteristic	Description
Reads from Snapshot Source LUN (SS)	With the optional SnapView software only: The number of reads from the source LUN during this snapshot session. It is the difference between Total Reads in Session and Reads from Cache.
Reads from Write Cache/s (L/M)	Average number of read requests per second that were satisfied by the write cache only. Reads from write cache occur when recently written data is read again while it is still in the write cache. This is a subset of read cache hits which includes requests satisfied by either the write or the read cache.
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.
Service Time (ms) (D/S/L/M)	Time, in milliseconds, that a component spends servicing a request. 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 (Mbytes/s).
Total Bandwidth (MB/s) (D/S/L/M)	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/M)	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.
Used Prefetches (%) (L/M)	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.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)
(continued)

Characteristic	Description
Utilization (%) (D/S/L)	The fraction of a certain observation period when 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, a high LUN utilization value does not necessarily indicate that the LUN is approaching its maximum capacity.
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 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.
Write Cache Hit Ratio (L/M)	The fraction of write requests that the write cache satisfied without requiring any disk access, 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/M)	The number of write requests per second that the write cache satisfied without requiring any disk access. Write cache hits are either requests that have been referenced before and are not yet flushed to the disks (rehits) or new write requests that did not trigger any forced flushes.
Write Cache Misses/s (L/M)	The number of write requests per second that the write cache alone could not satisfy and that, therefore also required additional disk access. Examples of write cache misses are write requests that bypass the write cache due to their size or write requests that trigger forced flushes.
Write Cache Rehit Ratio (L/M)	The fraction of write requests that the write cache satisfied since they have been referenced before and not yet flushed to the disks, compared to the total number of write requests to this LUN.
Write Cache Rehits (L/M)	The number of write requests per second that the write cache satisfied 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.

Table 3-19 Advanced Performance Characteristics (Alphabetical Order)
(continued)

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

Zooming and Resizing Charts

Analyzer lets you zoom in to (magnify) any type of chart along its time sequence (X axis), numeric value (Y axis), or both axes.

To zoom in to 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.

Figures 3-12 through 3-14 show a graph before zooming, the same graph with a zoom area selected, and the graph after zooming.




Figure 3-12 Chart Before Zoom




Figure 3-13 Chart With Zoom Area Selected (Left Mouse Button)




Figure 3-14 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 (for all but Performance Survey 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 Data to File**.
The **Save** dialog box appears.
3. Specify the filename and select the location on the local host.

Analyzer saves the file as a comma-delimited file with 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 Chart

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

Analyzer displays 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
 - Anti-aliasing
 - Color of background
 - Color of axis lines
 - Color of text
- ◆ For Survey charts (**Survey Charts** tab):
 - Display of the metaLUNs' component LUNs
 - Thresholds
 - Color of threshold line and point indicator
 - Sort order based on latest values or maximum values
- ◆ For Summary Charts (**Summary Charts** tab):
 - Color, style, and size of average point indicators
 - Color, style, and size of latest point indicator
 - Color, style, and width of min/max lines (the vertical lines that show value ranges)
- ◆ For Detail Charts (**Detail Charts** tab):
 - Color of the freeze border
 - 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)
 - Auto Freeze/Thaw

- ◆ For archive files (**Archives** tab):
 - Initial chart display, Performance Survey, Summary, or Detail, or LUN IO Disk Detail
 - Initial selection of all tree objects at first display

Setting Display Values for All Charts

When you initially select **Tools > Customize**, the **Customize** dialog box displays the **General** tab.

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

From the **Drawing** tab you can select or deselect the **Antialiasing** checkbox to enable or disable antialiasing for all chart windows. Antialiasing reduces the jagged distortions in curves and diagonal lines, so that the lines appear smoother. Antialiasing is turned off during scrolling to make scrolling more efficient.

From the **Colors** tab the **Background** box lets you specify the color of all chart backgrounds; the **Axis** box lets you specify the color of the 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 **Apply** to save the current settings and apply them to any currently opened windows.

Click **OK** to set the new values you specified and close the dialog box.

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

Setting Performance Survey Chart Values

To set Performance Survey chart values, select **Tools > Customize** and then select the **Survey Charts** tab.

The dialog box that appears includes the following sections:

- ◆ **MetaLUNs** section — check the **Include Component LUNs** checkbox if you want Performance Survey charts that contain metaLUNs to show the metaLUNs and their component LUNs. If you do not select the checkbox, only the metaLUNs appear in the chart window.

- ◆ **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 for 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 **Apply** to save the current settings, and click **OK** to close the dialog box.

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.

Within the **Colors** section, the **Average Point** box lets you specify the color, style, and size of the average point indicator; the **Latest Point** box lets you specify the color, style, and size of the Latest indicator; and the **Min/Max Line** box lets you specify the color, style, and width of the minimum/maximum line. For any box, click the button beside the color, style, or size sample to display a selection dialog box; then select a color, style, size, or width.

Click **Apply** to save the current settings, and click **OK** to close the dialog box.

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.

In the **Colors** section you can change the color of the freeze border for all detail charts. The default color is light blue.

The **Default Values** section includes the following checkboxes:

- ◆ **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).
- ◆ **Auto Freeze/Thaw** lets you automatically freeze a detail chart (LUN IO Disk Detail, Performance Detail, or IO Distribution Detail) when you zoom it. When you select **Reset Zoom**, (assuming that you have not already selected Thaw), the chart will return to its active state (thaw).

Click **Apply** to save the current settings, and click **OK** to close the dialog box.

Setting Custom Archive Chart Display Specifications

You can change the default display for Archive charts using the **Archives** dialog box. The **Initial View Chart** lets you specify the chart that Analyzer displays first: **Performance Survey** (default), **Performance Summary**, **Performance Detail**, or **LUN IO Disk Detail**.

The **Initially Check All Tree Objects** checkbox lets you specify whether the software automatically selects all objects in an Archive when it opens the chart.

Click **Apply** to save the current settings, and click **OK** to close the dialog box.

Capturing Performance Data in Archive Files

Generally, real-time charts are most useful for monitoring. For a 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.

In the **SP A** and **SP B** checkboxes, select the storage system SP(s) 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.

For single-SP systems or Windows servers that have the Storage Management Server software and user interface on them, the **SP A** and **SP B** checkboxes are disabled.

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

To help you find archive files later, 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 next.

Displaying Archive File Data

For archive file display, Analyzer requires at least one Analyzer archive file to be accessible. 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 follows:

1. On the Analyzer menu bar, select **Open archive**.
2. From the standard file selection dialog box, select an archive file and click **OK**.

Analyzer displays a dialog box that lets you select the time range for display.

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 more than 255 points exist 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.

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

- ◆ Resetting Performance Logging.....4-2
- ◆ Understanding Workload4-3
- ◆ Identifying a Performance Problem4-6
- ◆ Eliminating Bottlenecks4-8
- ◆ Working with the Storage-System Cache4-10
- ◆ Addressing Server Problems4-14

Resetting Performance Logging

You can reset logging only if you have Administrator or Manager privileges.

While the Analyzer Provider runs performance logging all the time, you may want to reset it.

1. From the Enterprise Storage tree, right-click the SP whose logging you want to reset.
2. Select **Logging > Reset**.
3. Click **Yes** to continue.

Resetting logging sets all the performance logging properties back to their default values and restarts performance logging without restarting the Management Server.

All data collected before you reset performance logging is lost.

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.

Workload Features

The main workload features are

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

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

Read/Write Ratio

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

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 amount of I/O for each size.

Locality Information

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.

Quantity of Work

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.

Intensity

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

The several areas of computer system operations that pertain to bottlenecks are

- ◆ 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), expand them 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 SPs is low, the problem is likely at the server. See the section *Addressing Server Problems* on page 4-14.

Identifying a Performance Problem

After evaluating your workload and deciding on obvious solutions, you can try a procedural approach to identify performance problems. Figure 4-1 guides you toward this goal.

Note that many of the performance characteristics (such as queue length) are explained in the section *Performance Characteristics* on page 3-38.




Figure 4-1 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 to keep 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 x *[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: Cache size page 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.

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 can better accommodate I/O bursts, but imposes a steady load on the disks. This may influence performance differently, depending on the load read requests impose on the disks. 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 other active reference streams to that LUN exist. However, if a single sequential stream addressing a LUN exists 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 the storage system appears to cause 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. Not enough resources exist 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 you need to diagnose server throttles. You might find that throttles occur 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 8 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.

When file systems, table spaces, and other storage objects are not aligned to stripe boundaries, 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 you can do.

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.

Command Line Interface (CLI) and Client Tools

This chapter describes the Analyzer Command Line Interface (CLI) commands and the Analyzer client tools. It also explains the EMC Navisphere CLI commands **navicli** and **navicli.jar**.

CLI commands for the EMC Navisphere™ software are explained in the CLI Reference manual.

The basic CLI command variations are also documented in the **mancli.txt** file.

Topics include:

◆ navicli.jar	5-2
◆ analyzer -set	5-6
◆ analyzer -logging	5-7
◆ archivedump.jar	5-8
◆ archiveretrieve.jar	5-16

navicli.jar

Sends status or configuration requests to a storage system via command line

Description The **java -jar navicli.jar** command sends storage-system management and configuration requests to a storage system via the Internet.

User Access Anyone that can log in to the host running Navisphere CLI 6.X or later.

Format The **navicli.jar** command is used as follows:

```
java -jar navicli.jar -help
or
java -jar navicli.jar
or for NetWare
java -jar [-ns | -nsac] navicli.jar -help
```

```
[-address address]
[-AddUserSecurity]
[-nopoll]
[-parse]
[-password password]
[-port port]
[-RemoveUserSecurity]
[-scope 0 | 1]
[-timeout seconds]
[-user username]
[-xml]
CMD [optional_command_switches]
```

where

where the **navicli** switches are

-help

Displays the help screen and does not start the **navicli** process. To start the **navicli** process, use one or more of the switches that follow instead.

-address *address*

Specifies the IP address or network name of the targeted SP on the desired storage system. The default, if you omit this switch, is **localhost**.

-AddUserSecurity

Directs the CLI to add user security information to the security file on this host. You must use the **-password** *password* and **-scope** switches.

-nopoll

Directs the feature provider not to issue a poll request. This switch significantly increases performance when dealing with large or multiple storage systems. The feature provider automatically polls unless this switch is specified.

When the **-nopoll** switch is set, **get** commands may return stale data and **set** commands may erase previously changed settings. Use caution when the **-nopoll** switch is set.

-ns | **-nsac**

For NetWare systems, creates a new screen for the Java application. One of these switches is needed with NetWare for CLI command output to appear on the screen. **-ns** creates a new screen for the Java application. **-nsac** also creates a new screen for the Java application, and automatically closes the screen when the application terminates.

-parse

Directs the CLI to validate the command. The CLI verifies the command syntax and displays a message stating whether the command was valid. The CLI takes no other action.

-password *password*

Specifies the password on the storage system you want to log in to. The password is case sensitive. You can omit this if you are using the security file.

-port *port*

Sets the port number (type) of the storage system. The default is 80.

-RemoveUserSecurity

Directs the CLI to remove security information about the current user from the security file on this host.

-scope 0 | 1

Specifies whether the user account on the storage system you want to log in to is local or global. A **0** (default) indicates global; a **1** indicates local.

A global account is effective throughout the domain. When the administrator creates a global account, the software copies the definition of this account to the domain directory, which makes it accessible on all storage systems in the domain.

A local account is effective on only the storage systems for which the administrator creates the account. The user can log in to only those storage systems on which he has a local account.

-timeout *seconds*

Sets the time-out value in seconds. The default is 600 seconds.

-user *username*

Specifies the username on the storage system you want to log in to. You can omit this if your username has been added to the security file.

-xml

Specifies command output in XML format. Use the **-o** (override switch) when specifying **-xml** on commands that require confirmation. Otherwise, the XML output will contain your confirmation string.

CMD

The Analyzer *CMDs* are **set**, **logging**, **archivedump**, and **archiveretrieve**.

The CLI will set the exit code before terminating. It sets the exit code to zero when the command executes successfully or to a value greater than zero when the command fails. Failures include command line parsing errors as well as errors encountered during the execution of the command.

CMD Switches The *CMD* switches are described on the pages that follow.

Spaces in Arguments

Normally, each argument to a CLI command consists of numbers, letters, and other valid printable characters for the operating system on which the CLI is running.

If a filename or other argument includes a space, you must enclose the entire string that includes the name in quotation marks ("xx xx").

analyzer -set

Modifies the performance logging properties

Description	The navicli.jar analyzer command with the -set switch lets you modify the performance logging properties.
User Access	You must have a user account on the storage system on which you want to execute the command.
Format	analyzer -set is used with java -jar navicli.jar (described on page 5-2) as follows. analyzer -set [-narinterval <i>seconds</i>] [-default] or analyzer -messner -set [-narinterval <i>seconds</i>] [-default] where -narinterval <i>seconds</i> Changes the polling interval for archive data. When you do not specify this switch, the interval is set to the default value of 600 seconds. Setting this value lower than the agent-polling rate will cause redundant data to go into the logger. Modifying this value affects the duration of the data in the file. For example, at 600 seconds an archive should have approximately 26 hours of data. If you lower the rate to 300 seconds, the archive file would have approximately 13 hours of data. -default Resets the values of all performance properties back to their default values.
Conventions and Recommendations	None.
Example	analyzer -set -narinterval 1200 This command sets the archive interval to 20 minutes. analyzer -set -default This command sets the archive interval back to 10 minutes.
Output	None if the command succeeds; status or error information if it fails.

analyzer -logging

Resets performance logging

Description The `navicli.jar analyzer` command with the `-logging` switch lets you reset performance logging.

User Access You must have a user account on the storage system on which you want to execute the command.

Format `analyzer -set` is used with `java -jar navicli.jar` (described on page 5-2) as follows.

`analyzer -logging [-reset]`

where

`-reset`

Stops performance logging. It deletes all the data that the Analyzer Provider has collected, resets all the performance logging properties, and deletes any existing archives. Then it starts performance logging.

Providing no options will return the current state of logging.

Conventions and Recommendations None.

Example `analyzer -logging -reset`

This command resets performance logging properties back to their default values and restarts performance logging.

Output None if the command succeeds; status or error information if it fails.

archivedump.jar

Dumps Analyzer archive files

Description The **archivedump** Java command line application dumps Analyzer 6.X archive files in a format that you specify.

User Access You must have a user account on the storage system on which you want to execute the command.

Format The **archivedump.jar** command is used as follows.

```
java -jar archivedump.jar -data filename [-header y | n]
[-object s | ml | l | d | ss] [-format codes] [-overwrite y | n] [-out filename]
[-stime start time] [-ftime end time] [-timezone codes] [-delim codes]
[-enquote codes] [-eol codes] help stats archive filename
[-timezone codes]
```

where

-data

Tells the dumper to dump out data from an archive file. You must provide the archive filename following the **-data** command.

-header

Specifies whether or not to output a column header.

y = Output a header

n = Do not output a header

-object

Specifies the objects to collect performance statistics for and the order in which to output them.

s = SP

ml = metaLUN

l = LUN

d = Disk

ss = Snap sessions

-format codes

Specifies which Performance Characteristics to output and the order in which they will appear. The **-format** switch requires one or more of the codes listed below. If you do not specify this switch, all performance characteristics will be output in the order listed below.

Performance Characteristic	Object Type	Code
Object Name	All objects	on
Owner Array Name	All objects	oan
Poll Time	All objects	pt
Utilization (%)	LUN, metaLUN, disk, SP	u
Queue Length	LUN, metaLUN, disk, SP	ql
Response Time (ms)	LUN, metaLUN, disk, SP	rt
Total Bandwidth (MB/s)	LUN, metaLUN, disk, SP	tb
Total Throughput (I/O/sec)	LUN, metaLUN, disk, SP	tt
Read Bandwidth (MB/s)	LUN, metaLUN, disk, SP	rb
Read Size (KB)	LUN, metaLUN, disk, SP	rs
Read Throughput (I/O/sec)	LUN, metaLUN, disk, SP	rio
Write Bandwidth (MB/s)	LUN, metaLUN, disk, SP	wb
Write Size (KB)	LUN, metaLUN, disk, SP	ws
Write Throughput (I/O/sec)	LUN, metaLUN, disk, SP	wio
Full Stripe Writes/s	LUN, metaLUN	fsw
Used Prefetches (%)	LUN, metaLUN	up
Read Cache Hits/s	LUN, metaLUN	rch
Read Cache Misses/s	LUN, metaLUN	rcm
Reads From Write Cache/s	LUN, metaLUN	rwc
Reads From Read Cache/s	LUN, metaLUN	rfr
Read Cache Hit Ratio	LUN, metaLUN	rchr

Write Cache Hits/s	LUN, metaLUN	wch
Write Cache Misses/s	LUN, metaLUN	wcm
Write Cache Rehits/s	LUN, metaLUN	wcr
Write Cache Hit Ratio	LUN, metaLUN	wchr
Write Cache Rehit Ratio	LUN, metaLUN	wcrr
Forced Flushes/s	LUN, metaLUN	ff
Disk Crossings (%)	LUN, metaLUN	dcp
Average Busy Queue Length	LUN, metaLUN, disk, SP	abql
Disk Crossings/s	LUN, metaLUN	dc
Service Time (ms)	LUN, metaLUN, disk, SP	st
LUN Read Crossings/s	metaLUN	lrc
LUN Write Crossings/s	metaLUN	lwc
Average Seek Distance (GB)	disk	asd
Dirty Pages (%)	SP	dp
Flush Ratio	SP	fr
MBs Flushed (MB/s)	SP	mbf
High Water Flush On	SP	hwfo
Idle Flush On	SP	ifo
Low Water Flush Off	SP	lwfo
Write Cache Flushes/s	SP	wcf
Reads From Snapshot Cache	Snap Session	rfsc
Reads From Snapshot Copy LUN	Snap Session	rfsnl
Reads From Snapshot Source LUN	Snap Session	rfssl
Writes To Snapshot Source LUN	Snap Session	wssl
Writes To Snapshot Cache	Snap Session	wsc
Writes Larger Than Snapshot Cache Chunk Size	Snap Session	wltscs
Chunks Used in Snapshot Copy Session	Snap Session	cuscs

Read 512B	LUN, metaLUN	r512b
Read 1KB+	LUN, metaLUN	r1kb
Read 2KB+	LUN, metaLUN	r2kb
Read 4KB+	LUN, metaLUN	r4kb
Read 8KB+	LUN, metaLUN	r8kb
Read 16KB+	LUN, metaLUN	r16kb
Read 32KB+	LUN, metaLUN	r32kb
Read 64KB+	LUN, metaLUN	r64kb
Read 128KB+	LUN, metaLUN	r128kb
Read 256KB+	LUN, metaLUN	r256kb
Read 512KB+	LUN, metaLUN	r512kb
Write 512B	LUN, metaLUN	w512b
Write 1KB+	LUN, metaLUN	w1kb
Write 2KB+	LUN, metaLUN	w2kb
Write 4KB+	LUN, metaLUN	w4kb
Write 8KB+	LUN, metaLUN	w8kb
Write 16KB+	LUN, metaLUN	w16kb
Write 32KB+	LUN, metaLUN	w32kb
Write 64KB+	LUN, metaLUN	w64kb
Write 128KB+	LUN, metaLUN	w128kb
Write 256KB+	LUN, metaLUN	w256kb
Write 512KB+	LUN, metaLUN	w512kb
Total 512B	LUN, metaLUN	t512b
Total 1KB+	LUN, metaLUN	t1kb
Total 2KB+	LUN, metaLUN	t2kb
Total 4KB+	LUN, metaLUN	t4kb
Total 8KB+	LUN, metaLUN	t8kb

Total 16KB+	LUN, metaLUN	t16kb
Total 32KB+	LUN, metaLUN	t32kb
Total 64KB+	LUN, metaLUN	t64kb
Total 128KB+	LUN, metaLUN	t128kb
Total 256KB+	LUN, metaLUN	t256kb
Total 512KB+	LUN, metaLUN	t512kb

-out filename

Specifies a file to which to write the data. Follow the switch with a single filename. Use quotes around the filename if it has spaces.

-overwrite

Works with the **-out** switch. If you use the **-overwrite** switch without the **-out** switch, it is ignored

y = Without a warning, overwrite the existing file specified by the **-out** switch.

n = Display an error message if the file that the **-out** switch specified already exists.

If you try to overwrite a read-only file, an error message appears.

-stime

Specifies an archive poll time from which to start iterating through the data. The **-stime** switch requires an argument that specifies a time in the following format: *mm/dd/yyyy hh:mm:ss*, where *mm* is the month of a year (1-12), *dd* is the day of a month (1-31), *yyyy* is the four-digit year, *hh* is the hour of the day (0-23), *mm* is the minute in an hour (0-59), and *ss* is the second in a minute (0-59). You must enclose the time in double quotes.

The **-stime** switch works in combination with the **-ftime** switch, described below. If the start time is greater than the end time, an error message appears. If you do not specify a start time, this switch defaults to the earliest poll time in the archive file.

If you use the **-stime** switch with the **-timezone** switch, you must apply the time zone you specify to the start time.

If the start time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the start time format is not valid, an error message appears.

-ftime

Specifies an archive poll time from which to stop iterating through the data. The **-ftime** switch requires an argument that specifies a time in the following format: *mm/dd/yyyy hh:mm:ss*, where *mm* is the month of a year (1-12), *dd* is the day of a month (1-31), *yyyy* is the four-digit year, *hh* is the hour of the day (0-23), *mm* is the minute in an hour (0-59), and *ss* is the second in a minute (0-59). You must enclose the time in double quotes.

The **-ftime** switch works in combination with the **-stime** switch, described above. If the start time is greater than the end time, an error message appears. If you do not specify an end time, this switch defaults to the latest poll time in the archive file.

If you use the **-ftime** switch with the **-timezone** switch, you must apply the time zone you specify to the end time.

If the end time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the end time format is not valid, an error message appears.

-timezone codes

Applies a particular time zone to the performance characteristic, Poll Time, included in the output data. By default, Analyzer displays this characteristic in the local time of the client system.

local = local time (default)

utc = universal time

-delim *codes*

Specifies a delimiter that will separate the output data.

cm = comma (default)

sc = semicolon

cl = colon

tb = tab

sl = slash

bs = backslash

pd = period

ws = white space

If you use a slash or a period as a delimiter, you must also use the **-enquote** switch.

-enquote

Encloses each performance characteristic within a set of special characters.

no = none (default)

qu = quotes

ap = apostrophe

ac = accent character

-eol

Specifies a particular end-of-line character. The output reaches an end of line when all performance characteristics for one specific object at a particular time have been output. The underlying client system determines the default end-of-line character.

cr = carriage return

lf = line feed

-help

Displays command usage.

-stats *archive filename* [**-timezone** *codes*]

Displays information about the archive file, including the start time, end time, total number of objects, total number of LUNs, total number of SPs, total number of disks, total number of metaLUNs, and total number of SnapView sessions.

-timezone *codes*

The **-timezone** switch applies a particular time zone to the start and end time the **-stats** switch returns. By default it displays both times in the local time of the client system.

local = local time (default)

utc = universal time

Conventions and Recommendations

None.

Examples **java -jar archivedump.jar -data archive.nar -header n**

This command tells the dumper to dump out data from the archive file, **archive.nar** without outputting a header.

java -jar archivedump.jar -data archive.nar -objects s,ml,l,d

This command tells the dumper to dump out data from the archive file, **archive.nar** for the SP, metaLUN, LUN, and disk.

Output Varies depending upon which switch you use.

archiveretrieve.jar

Retrieves Analyzer archive files

Description The **archiveretrieve** Java command line application retrieves archive files from storage systems.

User Access You must have a user account on the storage system on which you want to execute the command.

The **navicli.jar** is not required for the **archiveretrieve** command to run.

Format The **archiveretrieve.jar** command is used as follows.

```
java -jar archiveretrieve.jar [-User username] [-Password password]
[-Scope 0 | 1] -Address IP address [-Port port number] [-Timeout
seconds] [-Help] -File filename -Location directory path [-Overwrite 0 | 1]
[-Retry number] [-V]
```

where

-User *username*

Specifies the username that the security provider on the storage system you are logging into will attempt to validate.

-Password *password*

Specifies the password that the security provider on the storage system will attempt to validate.

-Scope 0 | 1

Specifies whether the account entry is global or local.

0 = global

1 = local

-Address *IP address*

Specifies the IP address or network name of the targeted SP on the storage system you indicate. The default is the local host.

-Port *port number*

Sets the port number to the target. The default is 80.

-Timeout *seconds*

Sets the timeout value in seconds. The default is 600 seconds.

-File *filename*

Specifies the name of the archive file once it is retrieved to the client system.

-Location *directory path*

Specifies the directory location of the archive file once it has been retrieved to the client system. If you do not specify a location, Analyzer uses the current working directory.

-Overwrite *0 | 1*

Specifies whether or not to overwrite an existing archive file on the client system.

0 = Do not overwrite existing file

1 = Overwrite existing file

-Retry *number*

Specifies the number of times to retry a retrieve before giving up completely. A retry will occur only when the actual moving of the archive file from the storage system to the client system fails.

-V

As the retrieve process progresses, the status of each stage appears.

-Help

Displays command usage.

Conventions and Recommendations

None.

Example `java -jar archiveretrieve.jar -User user -Password password
-Address 1.1.1.1 -File archive.nar -Location c:\temp -Overwrite 1
-Retry 2`

This command retrieves an archive from the specified address to the file **archive.nar** in **c:\temp** and overwrites the existing file.

Output None.

Customer Support

This appendix reviews the EMC process for detecting and resolving software problems, and provides essential questions that you should answer before contacting the EMC Customer Support Center.

This appendix covers the following topics:

- ◆ Overview of Detecting and Resolving ProblemsA-2
- ◆ Troubleshooting the ProblemA-3
- ◆ Before Calling the Customer Support CenterA-4
- ◆ Documenting the Problem.....A-5
- ◆ Reporting a New ProblemA-6
- ◆ Sending Problem Documentation.....A-7

Overview of Detecting and Resolving Problems

EMC software products are supported directly by the EMC Customer Support Center in the United States.

EMC uses the following process to resolve customer problems with its software products (Figure A-1).




Figure A-1 Problem Detection and Resolution Process

Troubleshooting the Problem

Please perform the relevant diagnostic steps before you contact the EMC Customer Support Center:

1. Read the documentation carefully.
2. Reconstruct the events leading up to the problem and describe them in writing.
3. Run some test cases to reproduce the problem.

If you encounter a problem that requires technical programming or analysis, call the nearest EMC office or contact the EMC Customer Support Center at one of the following numbers:

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

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

Worldwide: **(508) 497-7901**

Please do not request a specific support representative unless one has already been assigned to your particular system problem.

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

<http://powerlink.EMC.com>

Before Calling the Customer Support Center

Have the following information available before calling the Customer Support Center or your support representative (if one has been assigned to you):

- ☐ Your company name
- ☐ Your name
- ☐ Your phone number
- ☐ For an existing problem, the problem tracking system ID, if one was previously assigned to the problem by a support representative
- ☐ For an MVS problem, the JESLOG, SYSPRINT, all STDOUT DD members of the server job output and similar output for the client, and the relevant portion of the SYSLOG

Documenting the Problem

If the EMC Customer Support Center requests information regarding the problem, please document it completely, making sure to include the following information:

- ☐ Your company name and address
- ☐ Your name
- ☐ Your telephone number
- ☐ The importance of the problem, so that it can be assigned a priority level

To expedite the processing of your support request, you can photocopy this list and include it with the package.

Reporting a New Problem

For a new problem, please provide the following information:

- ☐ Release level of the software that you are running
- ☐ Software installation parameters
- ☐ Host type on which you are running
- ☐ Operating system you are running and its release number
- ☐ Functions of the software that you are running
- ☐ Whether you can reproduce the problem
- ☐ Previous occurrences of the problem
- ☐ Whether the software has ever worked correctly
- ☐ Time period that the software did work properly
- ☐ Conditions under which the software worked properly
- ☐ Changes to your system between the time the software worked properly and the problem began
- ☐ Exact sequence of events that led to the system error
- ☐ Message numbers and complete text of any messages that the system produced
- ☐ Log file dated near the time the error occurred
- ☐ Results from tests that you have run
- ☐ Other related system output
- ☐ Other information that may help solve the problem

Sending Problem Documentation

Use one of the following methods to send documentation of the problem to the EMC Customer Support Center:

- ◆ E-mail
- ◆ FTP
- ◆ U.S. mail to the following address:

EMC Customer Support Center
45 South Street
Hopkinton, MA 01748-9103

If the problem was assigned a number or a specific support representative, please include that information in the address as well.

A

- advanced performance characteristics 3-38
- Agent, about 1-3
- alignment, data 4-14
- Analyzer
 - chart, introduced 3-9
 - menu option 3-3
 - right-click menu options 3-3
 - starting 2-2, 3-2
- archive data, displaying 3-9, 3-56
- archive file
 - charts 2-3, 3-7
 - polling interval for 3-8
- arguments, spaces in 5-5

B

- bottlenecks 4-6
 - eliminating 4-7

C

- cache
 - page size 4-11
 - performance 4-10, 4-11
- characteristics, performance 3-38
- chart
 - Analyzer, introduced 3-9
 - archive file and real time 3-7
 - data, exporting 3-48
 - display, customizing 3-50
 - IO Detail 3-28
 - menu
 - IO Size Distribution 3-30

- Performance Detail 3-21
- resizing, zooming 3-45
- CLI
 - Java, basic commands 5-2
 - Java, spaces in arguments 5-5
- command line format, spaces in 5-5
- comma-separated values (csv) 3-48
- component objects, selecting 3-15, 3-24
- csv (comma-separated values) 3-48
- Customer Support Center A-3, A-7
- customizing chart display 3-50

D

- data, alignment on stripe and performance 4-14
- device menu
 - Performance Detail 3-24
 - Performance Summary 3-16
- dirty pages
 - counter 3-40
 - meaning 4-12
- disk
 - factors in performance 4-9
 - menu 3-5
- displaying archive data 3-9, 3-56
- domain, introduced 1-2

E

- environments, Manager 1-2
- exporting chart data 3-48

F

- forced flush rate 4-11

H

high watermark flush rate 4-12
 hit ratio
 read cache 4-13
 write cache 4-12
 host performance issues 4-14

I

IO Size Distribution chart 3-28

J

jpeg (.jpg) images 3-30, 3-48

L

light loads 4-5
 load, light 4-5
 locality of reference 4-5
 low watermark flush rate 4-12
 LUN (logical unit)
 bottleneck, identifying 3-52
 menu 3-5
 operating system device names 3-2
 performance (Performance Survey chart)
 3-11
 Survey chart 3-52

M

Manager
 environments 1-2
 starting 2-2, 3-2
 menu
 for storage-system components 3-4, 3-5, 3-6
 IO Size Distribution chart 3-30
 option, Analyzer 3-3
 Performance Summary 3-16
 Performance Survey 3-13
 metaLUN, menu 3-5
 metrics, sorting 3-18

N

navicli.jar CLI command 5-2
 Navisphere Manager, see Manager

O

operating system device names of storage-system
 LUNs 3-2

P

page
 dirty 4-12
 size, cache 4-11
 performance
 characteristics
 advanced 3-38
 basic 3-38
 data, understanding, Chapter 3
 problem, identifying 4-6
 survey chart 3-11
 Performance Detail chart 3-21
 Performance Summary chart 3-14
 Performance Survey chart 3-11
 menu 3-13
 pictures from charts (exporting) 3-48
 polling interval 3-8
 printing graphic displays 3-48

Q

queue length characteristic
 counter 3-41
 using 4-4

R

RAID factors in performance 4-9
 read cache
 counters 3-43
 performance 4-10
 real-time charts 3-7
 resizing charts 3-45
 right-click menu options 3-3

S

server
 performance issues 4-14
 throttle 4-14
 sorting based on latest or maximum values 3-50
 sorting by performance characteristics 3-11
 sorting metrics 3-18

SP (storage processor) menu 3-4, 3-5, 3-6
starting Analyzer 2-2, 3-2
stripe size factors and performance 4-9

T

thresholds
 LUN Survey chart 3-14
 useful values 3-52
throttle, server 4-14
tool tips 3-26, 3-38

U

utilization characteristic
 counter 3-39
 using 4-4, 4-6, 4-8

W

watermark flush rates 4-12
workload, light 4-5
workload, understanding 4-2
write cache
 counters, using 4-11
 performance 4-10
 size 4-10

Z

zooming charts 3-45

