

HP Storage Essentials

Software Version: 9.4.1

Storage Performance Management Guide

Document Release Date: Friday, July 08, 2011

Software Release Date: July 2011



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

© Copyright 2002-2011 Hewlett-Packard Development Company, L.P.

Trademark Notices

Java is a registered trademark of Oracle and/or its affiliates.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation.

Oracle is a registered trademark of Oracle Corporation.

UNIX® is a registered trademark of the Open Group.

Acknowledgements

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes software developed by the JDOM Project (<http://www.jdom.org/>).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.



Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users – please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at:

<http://www.hp.com/go/hpssoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

Storage Performance Management Guide.....	1
Legal Notices.....	2
Documentation Updates.....	3
Support.....	3
Contents.....	5
1 HP Storage Essentials Performance Manager.....	17
End-to-End SAN Path Topology.....	18
Flexible Access to Performance Data.....	19
Path Filtering for Performance Analysis.....	19
Performance Alerts and Notifications.....	20
Getting Started with Performance Manager.....	20
Selecting Elements to View Performance Statistics.....	20
Selecting Performance Monitoring Options.....	20
Opening the Performance Manager.....	21
General Considerations for Performance Manager.....	21
Ways to Use This Guide.....	21
Additional Resources.....	22
2 About HP Storage Essentials Performance Packs.....	23
HP Storage Essentials SRM Performance Pack Software.....	23
HP Storage Essentials Performance Pack Enterprise Software.....	24
About Performance Pack Setup and Configuration.....	24
Managing Performance Pack Licenses.....	25
About Managing Performance Pack Licenses.....	25
License Requirements.....	25
Number of Licenses.....	26
License Setup for Performance Pack Software.....	26
Reviewing the Current Usage Summary.....	27
Performing a Get Details Discovery.....	28
Importing the License.....	28

Applying the License.....	29
Next Steps.....	29
Setting Up the HP Storage Essentials SRM Performance Pack Software.....	30
Setup Steps for SRM Performance Pack for EVA.....	30
License Setup for EVA Performance Collection.....	31
About EVA License Requirements.....	31
SRM Performance Pack Supported Hardware.....	32
EVA Array-Related Software Requirements.....	32
Setting Up the HP Storage Essentials Performance Pack Enterprise Software.....	32
Setting Up the EMC Symmetrix Performance Pack.....	33
About EMC Symmetrix License Requirements.....	34
Setting Up the NetApp Performance Pack.....	34
Setting Up XP and HDS Performance Packs.....	35
Setup Steps for XP and HDS Arrays.....	35
About XP and HDS Array License Requirements.....	36
XP and HDS Array-Related Software Requirements.....	37
XP and HDS Performance Pack Enterprise Architecture.....	37
Configuring the Management Server.....	40
Create Command Device (LUN).....	41
Creating a Command LUN.....	41
Additional Related Tasks.....	42
Create the Host Proxy.....	42
Operating System Requirements.....	42
Setting Up the Host Proxy.....	42
Validate RAID Manager Library Installation.....	43
Verify that the RAID Manager Library Returns Data.....	44
3 Managing Performance Data Collection.....	45
Discovering Storage Environments for Performance Monitoring.....	45
Discovering EMC Symmetrix Storage Elements.....	46
Discovering HDS Storage Elements.....	46
Discovering HP StorageWorks EVA Arrays.....	47

Discovering HP StorageWorks XP Arrays.....	47
Discovering NetApp Storage Systems.....	48
Next Steps.....	48
Managing Performance Data Collectors.....	49
Setting Up Performance Data Collectors.....	50
Displaying the Performance Data Collectors Page.....	50
Default Collector Settings for Newly Discovered Elements.....	51
Applying Data Collector Filters.....	52
Starting Performance Data Collectors.....	53
Stopping Performance Data Collectors.....	54
Data Collector Schedules.....	54
Editing a Data Collector Schedule.....	55
Collecting Performance Data in Real Time.....	56
Viewing Data Aging Settings for Performance Statistics.....	56
Aggregation of Performance Data.....	58
Customizing Performance Data Collection.....	58
Managing XP and HDS Data Collection.....	58
About Setting Up XP Array Metric Collection.....	59
Collecting XP and HDS Array Performance Metrics.....	59
Managing EVA Data Collection.....	60
Collecting EVA Performance Metrics.....	62
EVA Configuration and Metrics Collection.....	63
EVA Performance Best Practices.....	63
Metric Collection for EVA Physical Disk Drives.....	64
EVAPerf Data Collector Considerations.....	64
4 Setting Performance Policies and Alerts.....	67
About Policy Manager.....	67
Severity Levels.....	68
Opening Policy Manager.....	68
Viewing Performance Policies.....	69
Adding Performance Policies.....	69

Modifying a Performance Policy.....	71
Deleting a Performance Policy.....	73
Setting Up E-Mail Actions for Performance Policies.....	73
Assigning an SMTP Mail Server.....	73
Setting Up E-Mail Notification for Policies.....	73
Setting Up Custom Commands for Performance Policies.....	74
Setting Up Custom Scripts.....	74
Setting Up Event Generation for Performance Policies.....	75
Setting Up Policy Event Generation.....	75
Viewing Performance Event Details.....	76
5 Viewing Performance Data.....	79
Toolbars and Menus.....	79
Creating Performance Charts.....	83
Saving Chart Settings.....	83
Opening Saved Charts.....	84
Deleting Chart Groups.....	84
Viewing Summary Charts.....	84
Setting a Custom Time Period.....	85
Viewing Trending Information for Performance.....	86
Removing Performance Data from a Graph.....	87
Interpreting Charts and Graphs.....	87
Sudden Dips in Charts in Performance Manager.....	87
Values Increase in Charts for Aggregated Drives and Aggregate Volumes.....	88
Comparing the Performance of Different Elements.....	88
6 Managing Performance in Your Storage Environment.....	91
Managing Storage System Performance.....	91
Performance Management and Metrics.....	91
Step 1: Identify the Physical Path.....	92
Step 2: Isolate the Bottleneck.....	92
Step 3: Troubleshoot the Cause.....	92
Step 4: Report on Performance Trends and Statistics.....	92

Step 5: Configure Proactive Threshold Notification using Policy Manager.....	93
Managing Application Performance.....	93
Troubleshooting Application Performance.....	93
Improving Microsoft Exchange Server Performance.....	94
Performance Baseline.....	95
Activity Spikes.....	95
Monitoring Oracle Server Performance.....	96
Step 1: Review the Database Design.....	96
Step 2: Tune the Application.....	96
Step 3: Tune the Memory.....	96
Step 4: Tune the Disk I/O.....	96
Step 5: Eliminate Database Contention.....	96
Step 6: Tune the Operating System.....	97
Monitoring Microsoft SQL Server Performance.....	97
Using Metrics to Identify Bottlenecks.....	97
Performance Management Application Metrics.....	97
Microsoft Exchange Server Performance Metrics.....	98
Oracle Server Performance Metrics.....	100
Microsoft SQL Server Performance Metrics.....	102
Sybase Performance Metrics.....	104
Informix Performance Metrics.....	104
IBM DB2 Performance Metrics.....	105
InterSystems Caché Database Performance Metrics.....	107
Managing Host Server Performance.....	108
Finding Host Performance Bottlenecks.....	108
Restrictions.....	108
Random Access Memory.....	109
Central Processing Unit.....	109
Physical Disk.....	109
Host Bus Adapter Performance Metrics.....	110
Bytes Received and Bytes Transmitted Metric.....	110

Best Practices.....	110
HBA Performance Metrics.....	110
Performance Management Host Server Metrics.....	111
Host Server Performance Metrics.....	111
Host Disk Performance Metrics.....	111
ESX Server Performance Metrics.....	112
Managing Switch Performance.....	113
Identifying Fibre Channel Performance Issues.....	113
FC Errors.....	113
CRC Errors.....	114
Link Failure.....	114
I/O Traffic.....	114
Monitoring Fabric Switch Performance.....	114
Best Practices.....	115
Fibre Channel Port Types.....	115
Performance Management Switch Metrics.....	116
Switch Performance Metrics.....	116
Managing EMC Symmetrix Performance.....	117
Collecting EMC Symmetrix Performance Statistics.....	117
Before You Collect Statistics.....	117
Best Practices.....	118
Scheduling EMC Symmetrix Performance Data Collection.....	118
Collection Frequency.....	119
Real-time Data Collection.....	119
About Counter Rollover.....	119
How Changes in Counter Values Are Calculated.....	120
Calculating the Number Base.....	120
How Rates Account for Single Rollover.....	121
Duration Values Used to Calculate Rates.....	122
Counters and KByte Conversion.....	123
EMC Symmetrix Performance Metrics.....	123

EMC Symmetrix System Performance Metrics.....	123
EMC Symmetrix Volume Performance Metrics.....	126
EMC Symmetrix Front-end Controller Performance Metrics.....	131
EMC Symmetrix Host FC Port Performance Metrics.....	133
Transient Values for EMC Symmetrix Performance Metrics.....	134
Logic for Delta Variables.....	134
Transient Values for EMC Symmetrix Systems.....	134
Transient Values for EMC Symmetrix Volumes.....	136
Transient Values for EMC Symmetrix Front-end Controllers.....	137
Transient Values for EMC Symmetrix Host FC Ports.....	138
Managing HP StorageWorks EVA Performance for Arrays.....	138
Finding EVA Bottlenecks Using Performance Metrics.....	138
I/O Per Second.....	138
Read Latency.....	139
Write Latency.....	139
Queue Depth.....	139
Monitoring EVA Performance.....	139
Performance Management EVA Metrics.....	140
Metric Table Conventions.....	141
EVA Storage System Metrics.....	141
EVA Storage System Raw Statistics.....	143
EVA Controller Metrics.....	144
EVA Controller Raw Statistics.....	145
EVA Physical Disk Metrics.....	146
EVA Physical Disk Raw Statistics.....	147
EVA Host FC Port Metrics.....	148
EVA Host FC Port Raw Statistics.....	150
EVA Volume Metrics.....	150
EVA Volume Raw Statistics.....	152
Managing XP and HDS Array Performance.....	153
XP Array Performance Tools.....	154

Finding XP and HDS Array Bottlenecks Using Performance Metrics.....	154
Best Practices.....	155
Monitoring XP Baseline Performance.....	155
Front-End Port Usage (CPU Load).....	155
Back-End ACP Pair Usage (CPU Load).....	155
Cache Utilization.....	155
Write Pending Rate.....	156
Write Response Time.....	156
Array Group Utilization.....	156
LDEV Utilization.....	156
Optimizing Your XP and HDS Array Performance.....	156
Viewing External Storage and Continuous Access Link Performance.....	157
Viewing Pool and Group Performance.....	157
Viewing Parity Group Statistics for XP and HDS Arrays.....	158
Performance Management XP and HDS Array Metrics.....	158
XP and HDS Array Storage System Metrics.....	158
Array Storage System Metrics1.....	159
XP and HDS Array Group Metrics.....	160
Best Practices.....	160
Array Group Metrics1.....	160
XP and HDS Array Volume Metrics.....	162
Best Practices.....	162
Array Volume Metrics1.....	162
XP and HDS Front-end Controller CLPR Metrics.....	164
XP and HDS Front-end Controller (CHA) Metrics.....	165
Best Practices.....	165
Controller CHA Metrics.....	165
XP Back-end Controller (DKA) Metric.....	165
Best Practices.....	166
Controller DKA Metrics.....	166
XP and HDS MPB Controller Metrics.....	166

MPB Controller Metrics 1.....	167
Managing NetApp Performance.....	168
Monitoring NetApp Performance.....	168
Using NetApp Performance Statistics to Identify Bottlenecks.....	168
Determine Latency Limits and Perform Baseline Analysis.....	168
Add Performance Policies and Alerts.....	169
Understand NetApp Performance Metrics.....	169
NetApp Performance Metrics.....	169
NetApp System Performance Metrics.....	169
NetApp Host Processor (CPU) Performance Metrics.....	173
NetApp Aggregate Performance Metrics.....	173
NetApp Front-End FC Port Performance Metrics.....	174
NetApp IP Port Performance Metrics.....	174
NetApp Logical Drive Performance Metrics.....	175
NetApp LUN Performance Metrics.....	177
NetApp Disk Drive Performance Metrics.....	178
NetApp Qtree Performance Metrics.....	178
NetApp Raw Statistics.....	179
NetApp System Raw Statistics.....	179
NetApp Aggregate Raw Statistics.....	181
NetApp Logical Drive Raw Statistics.....	181
NetApp Processor (CPU) Raw Statistics.....	182
NetApp Front-end Port Raw Statistics.....	182
NetApp IP Port Raw Statistics.....	183
NetApp LUN Raw Statistics.....	183
NetApp Disk Drive Raw Statistics.....	184
NetApp QTree Raw Statistics.....	184
7 Generating Performance Reports.....	187
Ways to Generate Performance Reports.....	187
Accessing InfoView Logon.....	187
Selecting Standard Performance Reports.....	188

List of Standard Performance Reports.....	189
Running Standard Performance Reports.....	222
Generating Performance Reports for Storage Systems.....	222
Generating Performance Reports for Physical Disk Devices.....	223
Generating Performance Reports for Controllers.....	223
Generating Performance Reports for Fibre Channel Ports.....	224
Generating Performance Reports for Top N Fibre Channel Ports by Metric.....	224
Generating Performance Reports for Storage Pools.....	225
Generating Performance Reports for Volumes.....	225
Generating Performance Reports for Top N Volumes by Metric.....	226
Generating Performance Reports for Bottom N Volumes by Metric.....	227
Generating Performance Reports for Array Groups.....	227
Generating Performance Reports for Top N Array Groups by Metric.....	228
Generating Performance Reports for Bottom N Array Groups by Metric.....	228
Generating Performance Reports for Switches.....	229
Backup Performance Forecast Reports.....	229
Forecast Calculation.....	230
Performance Forecast Report List.....	230
About Creating Simple Custom Reports.....	231
Task 1 – Select Dimension and Measure Objects.....	231
Task 2 – Select Filter Objects.....	233
Example: Host Performance Report.....	233
Creating Custom Performance Reports.....	236
Create a Web Intelligence Document.....	237
Connect to HP Storage Essentials Universe Database.....	237
Select Objects to Include in Report.....	238
Run the Report Query.....	238
Example: EVA to Host Performance Report.....	239
Creating an EVA Performance Report.....	240
Example: EVA to Host Performance Report.....	241
Task 1 – Design and Generate the Basic Report.....	241

Task 2 – Insert the First Table.....	243
Task 3 – Insert a Second Table.....	247
Creating Advanced Custom Reports.....	248
Example: Switch Port Performance Report.....	249
Creating Performance Reports with Drill Down and Trending.....	257
Example: EVA to Host Performance Trending Report with Drill Down.....	258
Task 1 – Design and Generate the Basic Report.....	258
Task 2 – Redesign the Report Format.....	260
Task 3 – Insert Drill Down.....	261
Task 4 – Add Trending Data in Charts.....	261
8 Troubleshooting Performance Issues.....	265
Performance Manager.....	265
XP Array Proxy Host is Not Accessible.....	265
Not Enough Data to Produce a Chart.....	266
Data Collection for EMC Arrays.....	266
Managing Late Data or Errors from XP and HDS Arrays.....	266
Glossary.....	269

[This page intentionally left blank]

1 HP Storage Essentials Performance Manager

HP Storage Essentials provides the key monitoring and analytical solutions that are essential to administration of critical business applications and the complex storage infrastructures supporting them. With the Performance Manager component, you have immediate visibility into the applications, hosts, switches, arrays, and path structures in your environment.

HP Storage Essentials also provides end-to-end storage area network (SAN) monitoring and reporting capabilities in a rich web-based user interface that allows you to select and display specific path performance metrics. You can view the aggregated results in easy-to-read charts and graphs, in a number of standard out-of-box reports (available using the SRM Report Optimizer), and in ad-hoc custom reports that you can develop with an optional user license. This resulting data gives you the detailed information you need for swift decision-making and efficient resolution of performance issues. As a result, you can increase system administration reaction times, troubleshoot performance bottlenecks faster, and quickly visualize the big performance picture of your array storage, hosts, applications, and SAN infrastructure.

The Performance Manager component provides a single, unified interface that enables you to view performance data appropriate for basic monitoring operations. The built-in capabilities include performance collection against managed databases, managed hosts, VMware ESX, NAS, and fabrics.

For comprehensive performance monitoring of disk arrays, HP recommends optional add-on Performance Packs—the HP Storage Essentials SRM Performance Pack for HP EVA focused on HP StorageWorks Enterprise Virtual Arrays (EVA) and the HP Storage Essentials Performance Pack Enterprise software for HP StorageWorks XP, Hitachi Data Systems (HDS), and EMC Symmetrix/DMX performance management.

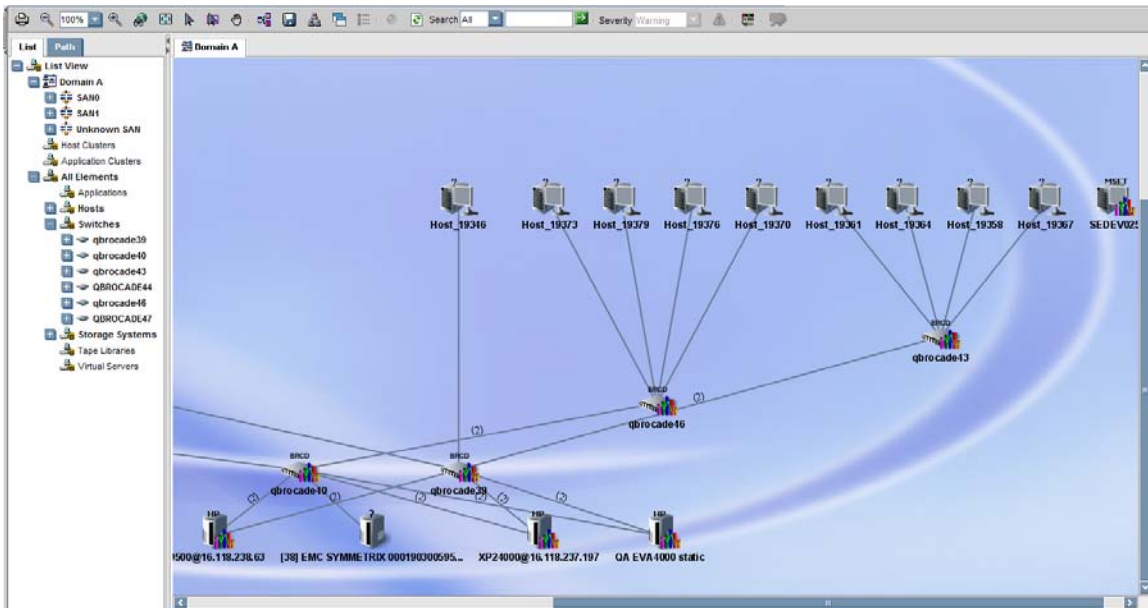
There is also an HP Storage Essentials Performance Pack Enterprise for NetApp. This Performance Pack is installed when you install the NetApp NAS Manager and does not require that you purchase a special license.

To learn more about how Performance Manager supports the day-to-day critical activities of monitoring storage systems, see the following topics about Performance Manager features:

- [End-to-End SAN Path Topology on next page](#)
- [Flexible Access to Performance Data on page 19](#)
- [Path Filtering for Performance Analysis on page 19](#)
- [Performance Alerts and Notifications on page 20](#)

End-to-End SAN Path Topology

Performance Manager maps the complete SAN path from business applications through the underlying storage fabric and SAN components to storage devices. From the host application, you can probe through the host bus adapter (HBA) into the switch and finally down to the disk to view all performance characteristics in a single screen.



You can view the performance details for each individual component or aggregate collected data for system-level analytics. Performance diagnostics are available for these path components:

- Business applications
- Host servers
- Host bus adapters
- Fabric switches
- Storage devices (SAN, DAS, and NAS)
- Storage arrays (EVA, XP, HDS, NetApp, and EMC Symmetrix)

Information about each of these performance areas is available in this guide. This guide focuses on performance metrics for disk arrays supported by the HP Storage Essentials SRM Performance Pack and Performance Pack Enterprise plug-ins for the HP Enterprise Virtual Array (EVA), HP XP, Hitachi Data Systems (HDS), and EMC Symmetrix arrays, and NetApp NAS devices.

Performance metrics are automatically collected from systems in the environment that have the appropriate CIM Extension or meet the pre-requisite for performance data collection. See the *HP Storage Essentials Installation Guide* and *User Guide* for data collection requirements.

All devices that meet the pre-requisite for performance data collection display in the HP Storage Essentials Performance Manager navigation map. You can click on any device in the map to display its performance history, or view devices in real time by selecting metrics from systems along the data path. You can also save the user-defined chart workspace and recall it at a later time for analysis of seasonal or reoccurring performance issues for a select group of systems.

HP recommends that you disable collection of unnecessary metrics on a per device basis, and as a best practice, use the sub-15 minute collection feature sparingly. Although you can set individual collectors to as little as 1 minute, these should be set only for limited duration and/or the scope of the collection.

The HP Storage Essentials SRM Performance Pack plug-in is required for performance collection of HP EVA arrays with HP Storage Essentials. The HP Storage Essentials Performance Pack Enterprise plug-in is required for performance management of HP XP or HDS arrays. For more information about HP Storage Essentials requirements for performance collection, see the *HP Storage Essentials Installation Guide* and *User Guide*.

Flexible Access to Performance Data

In addition to monitoring the statistics for each of these resources, HP Storage Essentials gives you three ways to analyze performance data using the Performance Manager:

- Real-time collection – displays tiled graphs that chart the statistics of each performance metric you select and shows you where performance is impacted along the application-storage path with a fine granularity of data points.
- Historical trend – trends performance over time using charts with customizable start and end dates, and hourly, daily, weekly, and monthly snapshots that enable you to quickly see what has changed over time.
- Trend extrapolation – leverages the trend period you selected to forecast what performance will look like in the future if the trend continues.

You are able to view performance data in real time at a fixed 15-, 20-, or 30-second interval by interactively working within the Performance Manager graphical user interface to select individual metrics in the performance charts. Performance data is otherwise collected and displayed at a default 15- or 60-minute interval and aggregated according to the weighted average for the sampling period. To customize a data sampling, you can configure the collection frequency for monitored components as needed.

You will want to take advantage of this flexible access to performance data as you complete your daily routines of verification and reporting, and as you troubleshoot performance issues.

Path Filtering for Performance Analysis

To further accelerate analysis of performance issues, Performance Manager automatically filters out any elements in the storage environment that are not relevant to the path you are monitoring. When you select an application or database instance, you see only the hosts, HBA ports, switch ports, switches, storage system ports, and storage systems, with their respective metrics, that are significant to your inquiry.

Performance Alerts and Notifications

You can set up the HP Storage Essentials Policy Manager to send notification e-mails, generate events or run custom command scripts when measured performance statistics reach thresholds or attain other conditions you've set in a policy. For example, if you are monitoring link failures on switches, you can specify in a performance policy that you want to be notified whenever the number of link failures for a specific switch exceeds your specific requirement. By setting policies, you can determine the point at which the data indicates a potential performance issue. At which point, you are notified so that you can respond immediately.

For more information about setting alerts, notifications, and events, see [Setting Performance Policies and Alerts on page 67](#).

Getting Started with Performance Manager

When you open the Performance Manager, discovered elements for your storage environment display in the List and Path tabs (left pane) and a topology view of the same elements displays in the right pane. No performance information displays until after you select an element.

Selecting Elements to View Performance Statistics

To select an element for which you want to view performance statistics, do one of the following:

- In the **List** or **Path** tab, expand the list and locate the element you want.

Click the element to select it and open the Monitoring view.

Or

- In the topology view (right pane), locate the element you want.

Click the element to select it and open the Monitoring view.

The Monitoring window panel appears showing the Monitoring tab, toolbar, performance metric options, and work windows.

Selecting Performance Monitoring Options

HP Storage Essentials collects performance data for your storage environment using any performance data collectors you have configured and enabled.

To learn how to configure performance data collectors, see [Managing Performance Data Collectors on page 49](#). For information about enabling a data collector, see [Setting Up Performance Data Collectors on page 50](#).

If data has been collected for a specific storage element, you can choose how you want to monitor this element by selecting the performance metrics you want to see and displaying the data in a chart.

HP Storage Essentials provides metrics appropriate for each type of storage element—application, host, switch, port, controller, and array. You can see which metrics are available for the element you've selected by clicking the element name in the list under the Monitoring tab. Performance monitoring options for the element appear in a panel below it. For a complete list of metrics for all storage elements, see [Managing Performance in Your Storage Environment on page 91](#).

Opening the Performance Manager

To open Performance Manager, click **Performance Manager** () on the HP Storage Essentials home page.

See Product Components in the *User Guide* for information about HP Storage Essentials and management server components.

General Considerations for Performance Manager

Check the support matrix for version requirements and the Release Notes for late-breaking documentation about Performance Manager.

Consider the following when running Performance Manager:

- If you see the message “There is not enough data to produce a chart <chart_title> at this time,” lessen the frequency option or select **All** in the **Edit Chart Property** dialog box.
- Verify that the performance collector for the element that is enabled (*Configuration > Performance*). See [Setting Up Performance Data Collectors on page 50](#) for more information about enabling performance collectors.
- Direct attached storage ports are not displayed in the storage tree in the bottom left pane.
- Performance Manager is not available to file servers.
- If one or more scheduled data points on a graph seem to be missing, an error might have occurred while collecting data at that time. Check the CIMOM log for errors for the time frame covering the sample interval of the missing data point(s).

All collectors are stopped during Get Details. This means that during Get Details, data for Performance Manager is not updated. Historical collectors, such as those available from the Configuration tab, are restarted when they are stopped during Get Details. Charts that were active in Performance Manager when Get Details was started are not restarted.

Ways to Use This Guide

This guide is intended for IT administrators, storage area network managers, and other technical personnel responsible for planning and maintaining their enterprise storage infrastructures for critical business applications. It includes detailed information about using HP Storage Essentials software to optimize performance monitoring, trending, notification, and reporting. It contains the information necessary to set up, discover, collect data, and generate reports for each component in your infrastructure.

You can find detailed information about performance metrics in the [Managing Performance in Your Storage Environment](#) section of this guide. Use the introductory sections to learn about what to look for when troubleshooting performance problems, and which metrics to monitor closely as indicators of potential problems. This section describes the performance management metrics which are collected, trended, and reported by the HP Storage Essentials Performance Manager. They contain important recommendations and best practices for monitoring the performance of your storage environment, and include metrics reference tables for each storage element in your SAN topology. This information can be useful when interpreting metrics data, analyzing performance states, and identifying performance issues for each of your SAN components. Use the metrics reference tables to understand which metrics are available for which storage element, and note any special considerations or conditions for their usage.

The [Generating Performance Reports](#) section includes both basic and advanced instructions for how to use Report Optimizer to generate reports which show the performance data gathered from your SAN environment. As you build reports, refer to the performance metric descriptions in this guide to ensure well-targeted data. For descriptions of HP Storage Essentials performance metrics, see [Managing Performance in Your Storage Environment on page 91](#)

Additional Resources

For more information about the HP Storage Essentials, see the *HP Storage Essentials Installation Guide* and *User Guide* located at:

<http://h20230.www2.hp.com/selfsolve/manuals>

You will find HP whitepapers related to performance management and monitoring at:

<http://h18006.www1.hp.com/storage/softwhitepapers.html>

For additional HP Enterprise Virtual Array (EVA) performance information, see documents located at:

<http://h71028.www7.hp.com/ERC/downloads/5983-1674EN.pdf>

<http://kb.know.hp.com/lib/KB04000/KB3903.docx>

The HP StorageWorks Sizer Tool is located at:

<http://h30144.www3.hp.com/SWDSizerWeb/default.htm>

To verify compatibility of component and software versions, see the HP Storage Essentials support matrix document on your product DVD. You can also find the support matrix document in the Documentation Center (**Help > Documentation Center**) accessible from the HP Storage Essentials user interface.

2 About HP Storage Essentials Performance Packs

HP Storage Essentials provides a basic set of performance monitoring capabilities that enable you to view and track performance metrics and to set up performance policies for your storage devices. In addition to these capabilities, you can purchase performance pack plug-ins which provide advanced monitoring features and options designed as essential components for large enterprise storage infrastructures.

HP Storage Essentials performance pack plug-ins are available in two product offerings:

- [HP Storage Essentials SRM Performance Pack Software](#)
- [HP Storage Essentials Performance Pack Enterprise Software](#)

For information about setting up and configuring performance packs, see [About Performance Pack Setup and Configuration](#) on next page.

HP Storage Essentials SRM Performance Pack Software

The HP Storage Essentials Storage Resource Management (SRM) Performance Pack Software provides advanced capabilities that enable you to visualize and analyze the end-to-end performance of the HP Enterprise Virtual Arrays (EVAs) in your storage area network (SAN).

HP Storage Essentials provides a single unified web-based interface, enabling you to monitor the performance of applications, host servers, host bus adapters (HBAs), switches, and EVA arrays. Your view is a complete end-to-end EVA and SAN view. Using topology maps, you can query individual devices and identify potential problems as they emerge. This helps to understand issues early and avoid performance bottlenecks.

The SRM Performance Pack also includes a large number of performance metric capabilities that when used with built-in reporting, give you the ability to generate reports to meet specific business and IT requirements.

The HP Storage Essentials SRM Performance Pack Software is available as a value-add plug-in to HP SRM Enterprise Edition Software for enterprise use. To accommodate larger HP EVA deployments, licensing is on a per HP EVA (array) basis. HP SRM Standard Edition Software, which targets mid-market HP storage-based SANs, includes licensing support to monitor up to three HP EVAs.

See [About Managing Performance Pack Licenses](#) on page 25 for more information about licensing.

For general information about supported hardware, see [SRM Performance Pack Supported Hardware](#) on page 32.

For a list of supported operating systems, see the *HP Storage Essentials Support Matrix*, which is available on the Documentation Center (**Help > Documentation Center**) in the HP Storage Essentials software interface.

HP Storage Essentials Performance Pack Enterprise Software

The HP Storage Essentials Performance Pack Enterprise Software is available as a licensed plug-in for the following supported storage systems:

- EMC Symmetrix Systems—see [Setting Up the EMC Symmetrix Performance Pack on page 33](#).
- Hitachi Data Systems (HDS)—see [Setting Up XP and HDS Performance Packs on page 35](#).
- HP StorageWorks XP arrays—see [Setting Up XP and HDS Performance Packs on page 35](#)
- NetApp systems—see [Setting Up the NetApp Performance Pack on page 34](#).

The HP Storage Essentials Performance Pack Enterprise Software plugs into the Performance Manager and Report Optimizer to give you end-to-end SAN visibility into the above listed storage systems. Installing the performance pack also adds an additional set of performance metrics to the basic metric set installed with the Performance Manager.

About Performance Pack Setup and Configuration

The Performance Manager is installed as part of the HP Storage Essentials software installation and provides a basic set of capabilities for monitoring performance of storage elements in your SAN. The HP Storage Essentials Storage Resource Management Performance Pack Software and the HP Storage Essentials Performance Pack Enterprise Software products enhance these capabilities by including additional performance metrics for in-depth analysis, performance policy management, and report generation. For a description of HP performance pack products, see [About HP Storage Essentials Performance Packs on previous page](#).

For step-by-step instructions on setting up a storage environment that runs performance pack software, see the following topics:

Performance Pack	For information, see...
EMC Symmetrix	Setting Up the EMC Symmetrix Performance Pack on page 33
Hitachi Data Systems (HDS)	Setting Up XP and HDS Performance Packs on page 35
HP StorageWorks Enterprise Virtual Array (EVA)	Setup Steps for SRM Performance Pack for EVA on page 30
HP StorageWorks XP	Setting Up XP and HDS Performance Packs on page 35

Performance Pack	For information, see...
NetApp	Setting Up the NetApp Performance Pack on page 34

Managing Performance Pack Licenses

This section contains the following topics:

- [About Managing Performance Pack Licenses below](#)
- [License Requirements below](#)
- [License Setup for Performance Pack Software on next page](#)

About Managing Performance Pack Licenses

How many HP Storage Essentials Performance Pack licenses you need is directly related to the performance monitoring and data collection requirements of your enterprise storage environment.

First determine which arrays and/or devices in your system require critical monitoring activity and consider implementing HP Storage Essentials Performance Pack licenses for additional performance data gathering and reporting for these arrays. If your performance monitoring includes data compliance requirements or detailed management reports, this might also impact your decision as to how many licenses you require. You'll need one license per array. For information about license requirements, see [License Requirements below](#).

NetApp NAS devices do not require a special performance pack license. The NetApp license is enabled with the NetApp NAS Manager.

Also review the metrics lists in this guide to help you determine the types of performance statistics you want to capture and how to setup and collect performance data. For more information, see [Managing Performance in Your Storage Environment on page 91](#).

For general information about HP Storage Essentials license management, see the License chapter in the *HP Storage Essentials User Guide*.

The HP Storage Essentials SRM Performance Pack and Performance Pack Enterprise Software license agreements are published in PDF format. The appropriate license file is displayed the first time you access HP Storage Essentials. You should, therefore, have the latest version of Adobe Acrobat Reader installed on the client from which you plan to perform first-time access of the HP Storage Essentials software.

License Requirements

The HP Storage Essentials Performance Pack licenses provide the ability to collect and report additional performance data. They are available for the following storage environments:

Software License	Storage Network	For information, see...
HP Storage Essentials SRM Performance Pack	HP StorageWorks Enterprise Virtual Array (EVA)	About EVA License Requirements on page 31
HP Storage Essentials Performance Pack Enterprise	EMC Symmetrix storage systems	About EMC Symmetrix License Requirements on page 34
HP StorageWorks XP arrays	Hitachi Data Systems (HDS) arrays	About XP and HDS Array License Requirements on page 36

There is no license for the HP Storage Essentials Performance Pack Enterprise Software for NetApp devices. The license is applied as part of the NAS Manager license.

For information about how to import a license file and configure settings, see [Importing the License on page 28](#).

Number of Licenses

The number of licenses required for any particular storage environment depends on how many EMC, EVA, HDS, and/or XP arrays you want to license. Each Performance Pack license allows you to monitor one array. To monitor multiple arrays, you must purchase one additional license for each array.

For information about setting up Performance Pack licenses, see [License Setup for Performance Pack Software below](#).

License Setup for Performance Pack Software

The following steps apply to either the HP Storage Essentials SRM Performance Pack or HP Storage Essentials Performance Pack Enterprise software licenses:

1. Review the Current Usage Summary table. See [Reviewing the Current Usage Summary on the facing page](#).
2. Complete a Get Details Discovery for the array you plan to license—see [Performing a Get Details Discovery on page 28](#).
3. Import the license. See [Importing the License on page 28](#).
4. Apply the license. See [Applying the License on page 29](#).

Before you import an HP Storage Essentials SRM Performance Pack or Performance Pack Enterprise software license file, you **MUST** perform a Get Details step to discover the storage arrays in your SAN for which you plan to apply the licenses.

Reviewing the Current Usage Summary

The first step in setting up a license is to review the Current Usage Summary table available in the Performance Manager. You can open this license page at any time. To view and apply Performance Pack licenses:

1. Click **Security** in the top upper right menu bar. This displays a submenu bar.
2. Click **Licenses** in the submenu.
3. Click the **Performance Licensing** tab.

A Current Usage Summary table appears showing the used and maximum numbers of managed arrays, and how many arrays can have the license capability applied.

Licensed Capacities	Used Licenses	Maximum Licensed
MAPs	415	500
Raw NAS Capacity	0.00 TB	9,999.00 TB
Managed Exchange Instances	0	1
Managed Database Instances	0	1
Managed Oracle Instances	0	
Managed SQL Server Instances	0	
Managed Sybase Instances	0	
Managed Caché Instances	0	
Managed File Server Storage	0.00 TB	9,999.00 TB
EVA Performance Pack Array-Instances *	0	1
XP, HDS Performance Pack Array-Instances *	1	2

* Use the Performance Licensing tab to apply licenses to storage systems.

From the Current Usage Summary page, you can:

- Import a license file.
- View a list of imported licenses by name.
- View properties for any licenses already imported and valid on your system.
- View cumulative license information.
- Delete licenses.
- Refresh the license usage data.

This information is important to understand prior to importing and applying a license to a storage array. For general information about managing licenses, see the *HP Storage Essentials User Guide*.

Performing a Get Details Discovery

Before you import a Performance Pack license, you must perform a Discovery- Get Details for those arrays you plan to license. This enables HP Storage Essentials to discover the arrays in your storage environment and display the correct information when you import and apply the license.

For information about the Discovery - Get Details process, see the [Discovering Storage Environments for Performance Monitoring on page 45](#). You can also find detailed Discovery information in the *HP Storage Essentials User Guide* for the array you are licensing.

For general Discovery information, see the following topics:

To Discover...	For information, see...
EMC Symmetrix	Discovering EMC Symmetrix Storage Elements on page 46
HP Enterprise Virtual Array (EVA)	Discovering HP StorageWorks EVA Arrays on page 47
HP XP	Discovering HP StorageWorks XP Arrays on page 47
HDS	Discovering HDS Storage Elements on page 46
NetApp	Discovering NetApp Storage Systems on page 48

Importing the License

When you import and configure the Performance Pack license, you can specify which arrays are to be monitored. You can make changes to your selection at a later time if you need to accommodate a new performance monitoring strategy.

To import a license file:

1. Select **Security**. This displays submenu options related to Security setup.
2. Select **Licenses** from the submenu.
3. Click the **Licenses** tab.
4. Click **Import License File**. The Import License File dialog is displayed.
5. In **Filename**, type the name of the import file or select **Browse** to navigate to the license file you want. Select the license file you want to import and click **Open** to enter it in the Filename field.
6. Click **OK** to import the file.

Click **Cancel** to cancel the import and return to the Licenses page.

After you import the license, select the array to which you want the license applied as described in the next topic.

Applying the License

After you import a Performance Pack license, you must apply the license to an array on your system. For example, if you purchased a Performance Pack license for an EVA array, then you would select an EVA array in your SAN and apply the license to it.

To apply the Performance Pack license:

1. Click the **Performance Licensing** tab and specify to which array you want the license applied by checking the checkbox beside the license name.

Licenses
Performance Licensing

Performance Pack Licenses
 EVA Total: 1 Used: 0 Available: 1
 XP, HDS Total: 2 Used: 1 Available: 1

Performance Pack licenses enable you to collect detailed statistics for a specific number of storage systems.

1. Manage licenses using the Licenses tab above
2. To license/unlicense a storage system, select/unselect the storage system check box in the table below and click Apply.
3. Start data collection for licensed storage system on the [Performance Data Collection](#) page.

Showing 1-1 out of 1 Total (1 Selected)

<input type="checkbox"/>	Name	Licensed	Serial Number	Vendor	Model	IP Address
<input checked="" type="checkbox"/>	HDS9970V@192.168.99.15	Yes	20168	Hitachi Data Systems	HDS9900V	essex.selab.usa.hp.com

Apply
Reset

2. Click **Apply**.

Next Steps

After you import the license, you can start the data collectors from the Performance Data Collection page (**Configuration > Performance > Data Collection**). This enables HP Storage Essentials to begin collecting performance data for the newly licensed array. See [Starting Performance Data Collectors](#) on page 53 for more information.

Setting Up the HP Storage Essentials SRM Performance Pack Software

You must purchase an HP Storage Essentials Storage Resource Management Performance Pack Software license for each Enterprise Virtual Array (EVA) for which you want advanced HP Storage Essentials performance monitoring, policy, and report capabilities.

Before you import an SRM Performance Pack license, you must run Discovery (see [License Setup for Performance Pack Software on page 26](#)). After you perform Discovery Get Details, HP Storage Essentials can display the discovered EVA arrays in your SAN to which you can apply your licenses.

You can view a summary of the Performance Pack licenses installed by reviewing the Current Usage Summary page (see [License Setup for Performance Pack Software on page 26](#)).

Setup Steps for SRM Performance Pack for EVA

Follow the steps in the order provided in the following table:

Step	Description	For information, see...
1	Install HP Storage Essentials software	<i>HP Storage Essentials Installation Guide</i>
2	Run Discovery - Get Details	Discovering HP StorageWorks EVA Arrays in the <i>HP Storage Essentials User Guide</i> ; also, Discovering HP StorageWorks EVA Arrays on page 47
3	<i>(Optional)</i> Review license summary usage information for your storage system	Reviewing the Current Usage Summary on page 27
4	Import the HP StorageWorks SRM Performance Pack license for the Enterprise Virtual Array to your management server	Importing the License on page 28
5	Apply the SRM Performance Pack license to the specific EVA array you want to license	Applying the License on previous page
6	Set up HP Storage Essentials data collectors	Setting Up Performance Data Collectors on page 50

Step	Description	For information, see...
7	Enable HP Storage Essentials data collectors from the Performance Data Collection page	Displaying the Performance Data Collectors Page on page 50
8	Start the EVAPerf Data Collector (EVAPerf) service	Starting Performance Data Collectors on page 53
9	Configure HP Storage Essentials for SNMP notifications	"Configuration Steps After a Finish Installation of HP Storage Essentials: Step 3 Configure HP Storage Essentials to Receive SNMP Notifications" in the <i>HP Storage Essentials Installation Guide</i>
10	(Optional) Configure Policy Manager for performance policy alerts and notifications	Setting Performance Policies and Alerts on page 67

License Setup for EVA Performance Collection

The Performance Pack plug-in is required for performance collection of HP EVA arrays with HP Storage Essentials. You can view EVA data in Performance Manager only if you are appropriately licensed.

You must complete a Get Details for EVA arrays before you import the license for the Performance Pack. After you import the license, you can start the data collectors from the Performance Data Collection page (**Configuration > Performance > Data Collection**). Although EVA arrays are displayed after you run Discovery, you must run a Get Details for the collectors to run properly.

For more information about licensing EVA arrays, see Managing Licenses in the *HP Storage Essentials User Guide*.

About EVA License Requirements

The management server restricts the number of elements it manages through its licenses. Your HP Storage Essentials SRM Performance Pack license determines how many EVA arrays you can select for performance monitoring. Each SRM Performance Pack license will let you monitor only one EVA array. To monitor multiple EVA arrays, you must purchase an SRM Performance Pack license for each EVA array.

You must keep your license up to date with the requirements of your network. If you need to monitor more arrays than your current license allows, contact HP Support for additional licenses.

As part of the license setup, you must specify which EVA arrays are to be monitored. You can make changes to your selection at a later time to accommodate new monitoring and reporting needs.

Setup for the SRM Performance Pack for EVA also require you to specify the data collectors used to implement your specific reporting needs. This is discussed in [Collecting EVA Performance Metrics on page 62](#).

SRM Performance Pack Supported Hardware

The SRM Performance Pack supports an open, enterprise heterogeneous storage (NAS, DAS, SANs, arrays) infrastructure, including fabric switches, HBAs, servers when included with HP Storage Essentials SRM Enterprise Edition Software or for HP Storage Essentials SRM Standard Edition Software, HP StorageWorks Enterprise Virtual Array, HP StorageWorks Modular Storage Array, HP ProLiant Storage Server (NAS), HP StorageWorks EML-E Series Tape Libraries; HBA and fabric switches; and Windows and Linux-based servers.

For a detailed list of supported hardware, see the HP Storage Essentials support matrix document available in the Documentation Center (**Help > Documentation Center**).

EVA Array-Related Software Requirements

The following list is a set of current best practices for EVA Performance data collection:

- If you installed the Command View EVA 9.1 or later SMI-S provider, use:
 - EVA firmware XCS 6200 or later
 - CV EVA 9.1 or later
- Command View EVA versions 9.1 and later cannot obtain cumulative metrics from the VCS 3.110 firmware. Therefore, it has no metrics to provide to HP Storage Essentials. If you cannot upgrade your VCS firmware, make sure you are running Command View 8.0.2 or 9.0.1. Do not upgrade to Command View EVA 9.1 or later. If you upgrade to Command View EVA 9.1 or later, you will no longer see your EVA metrics in HP Storage Essentials.
- A maximum of 8 EVAs per CV EVA server are supported.
- Performance data for EVAs running VCS 3.x is only point in time. Each data point for performance metrics corresponding to EVAs running VCS 3.x will only represent a one second sample period.
- VCS 4.004 is not supported for EVA5000 disk arrays.
- Management server performance provider does not support EVA VCS code earlier than 4.x.

The maximum number of vdisks supported per EVA array is 512.

You must discover the arrays for the Array Performance Pack to work.

Setting Up the HP Storage Essentials Performance Pack Enterprise Software

This section describes the steps for setting up the HP Performance Pack Enterprise software for your EMC Symmetrix, NetApp, XP or HDS storage environment. See the following topics:

- [Setting Up the EMC Symmetrix Performance Pack below](#)
- [Setting Up the NetApp Performance Pack on next page](#)
- [Setting Up XP and HDS Performance Packs on page 35](#)

Setting Up the EMC Symmetrix Performance Pack

You must perform configuration tasks on the management server to set up the HP Storage Essentials Performance Pack Enterprise software for EMC Symmetrix DMX storage in your environment. The following table outlines the administrative tasks required to set up the Performance Pack.

Step	Description	For information, see...
1	Install HP Storage Essentials software	<i>HP Storage Essentials Installation Guide</i> or <i>Hitachi Storage Services Manager Installation Guide</i>
2	Run Discovery (Get Details)	Discovering HP StorageWorks EVA Arrays in the <i>HP Storage Essentials User Guide</i>
3	(Optional) View license summary usage information for your storage system	Reviewing the Current Usage Summary on page 27
4	Import the HP Performance Pack Enterprise license for EMC Symmetrix to your management server	Importing the License on page 28
5	Apply the HP Performance Pack Enterprise license for EMC Symmetrix	Applying the License on page 29
6	Set up performance data collectors	Setting Up Performance Data Collectors on page 50
7	Start HP Storage Essentials data collection from the Performance Data Collection page	Displaying the Performance Data Collectors Page on page 50
8	Configure HP Storage Essentials for SNMP notifications	"Configuration Steps After a Finish Installation of HP Storage Essentials: Step 3 Configure HP Storage Essentials to Receive SNMP Notifications" in the <i>HP Storage Essentials Installation Guide</i>

Step	Description	For information, see...
9	(Optional) Configure Policy Manager for performance policy alerts and notifications	Setting Performance Policies and Alerts on page 67

About EMC Symmetrix License Requirements

To use all of the features of the management server, such as provisioning, with an EMC Symmetrix storage system, you must have licenses for the following products:

- BASE
- DeltaMark
- SERVER
- DevMasking
- Config Manager
- Mapping (SOLUTION_4)

If error 214 is present in the discovery log or cimom.log during discovery, the SymAPI server is not licensed for remote connections. You must acquire and install the license before discovery can occur.

Setting Up the NetApp Performance Pack

There is no Performance Pack Enterprise license required for NetApp devices. The license is installed when you install the NetApp NAS Manager.

Step	Description	For information, see...
1	Install HP Storage Essentials software	<i>HP Storage Essentials Installation Guide</i> or <i>Hitachi Storage Services Manager Installation Guide</i>
2	Run Discovery (Get Details)	"Discovering HP StorageWorks EVA Arrays" in the <i>HP Storage Essentials User Guide</i>
3	Set up performance data collectors	Setting Up Performance Data Collectors on page 50
4	Start HP Storage Essentials data collection from the Performance Data Collection page	Starting Performance Data Collectors on page 53

Step	Description	For information, see...
5	Configure HP Storage Essentials for SNMP notifications	"Configuration Steps After a Finish Installation of HP Storage Essentials: Step 3 Configure HP Storage Essentials to Receive SNMP Notifications" in the <i>HP Storage Essentials Installation Guide</i>
6	(Optional) Configure Policy Manager for performance policy alerts and notifications	Setting Performance Policies and Alerts on page 67

Setting Up XP and HDS Performance Packs

You must perform configuration tasks on the management server to set up the XP and/or HDS Performance Pack Enterprise software.

Performance Manager performance statistics are collected in-band from a command device. Therefore, to collect performance data, you must set up at least one host in your SAN with a command device from the array and install a CIM Extension. The CIM Extension automatically detects the command device. After you run Discovery, view the array's properties, then click the **Edit Proxy Host** button to select the host that will be used to collect performance data.

Setup Steps for XP and HDS Arrays

You must install the RAID Manager Library and CIM Extension software. The installation steps are included in the instructions for setting up the host proxy. Follow the steps in the order shown in the following table:

Step	Description	For information, see...
1	Install HP Storage Essentials software	<i>HP Storage Essentials Installation Guide</i> or <i>Hitachi Storage Services Manager Installation Guide</i>
2	Run Discovery and their host proxies to discover the arrays you plan to license	For XP arrays: See "Discovering HP StorageWorks XP Arrays" in the <i>HP Storage Essentials User Guide</i> For HDS arrays: See "Discovering HDS Storage Systems" in the <i>HP Storage Essentials User Guide</i>
3	(Optional) Review license summary usage information for your storage system	Reviewing the Current Usage Summary on page 27
4	Import the HP Performance Pack Enterprise software license onto your management server	Importing the License on page 28

Step	Description	For information, see...
5	Create a command device (LUN)	Create Command Device (LUN) on page 41
6	Run Discovery Get Details to get all host and array information.	Discovering HP StorageWorks XP Arrays on page 47 and Discovering HDS Storage Elements on page 46
7	Enable the license	Applying the License on page 29
8	Designate the host proxy	Create the Host Proxy on page 42
9	Verify that the RAID Manager Library (RMLIB API) is returning data through the Command LUN	Validate RAID Manager Library Installation on page 43
10	Set up HP Storage Essentials performance data collectors	Setting Up Performance Data Collectors on page 50
11	Start HP Storage Essentials data collection from the Performance Data Collection page	Displaying the Performance Data Collectors Page on page 50
12	Configure HP Storage Essentials for SNMP notifications	"Configuration Steps After a Finish Installation of HP Storage Essentials: Step 3 Configure HP Storage Essentials to Receive SNMP Notifications" in the <i>HP Storage Essentials Installation Guide</i>
13	(Optional) Configure Policy Manager for performance policy alerts and notifications	Setting Performance Policies and Alerts on page 67

About XP and HDS Array License Requirements

Your Performance Pack Enterprise license for XP and HDS arrays determines how many XP and/or HDS arrays you can select for performance monitoring. Additional licenses can be purchased if you need to monitor more XP and HDS arrays than your current license allows.

As part of the license setup, you must specify which XP or HDS arrays you want monitored. You can make changes to your selection at a later time to accommodate new monitoring and reporting needs. For information about license setup, see [License Setup for Performance Pack Software on page 26](#)

XP and HDS Array-Related Software Requirements

The Performance Pack Enterprise software for XP and HDS arrays requires the RAID Manager Library on the host proxy, as well as Command View Advanced Edition or the built-in provider.

Host Proxy requirements:

- CIM extension—The host proxy for the XP array must have a CIM extension installed and running on the proxy. The Operating System for the host proxy server is limited: Windows, Linux, HP-UX, and others. See the Support Matrix for a complete list.
- RAID Manager Lib XP (RMLIBXP)—This is not included in the CIM extension installation package. It must be installed separately on the host proxy server. For information about locating and installing RAID Manager Lib XP, refer to your XP Product documentation.

See the HP Storage Essentials support matrix for a detailed listing of required software and their versions, including the required version software of your XP or HDS array. The support matrix document is available from the HP Storage Essentials Documentation Center (**Help > Documentation Center**).

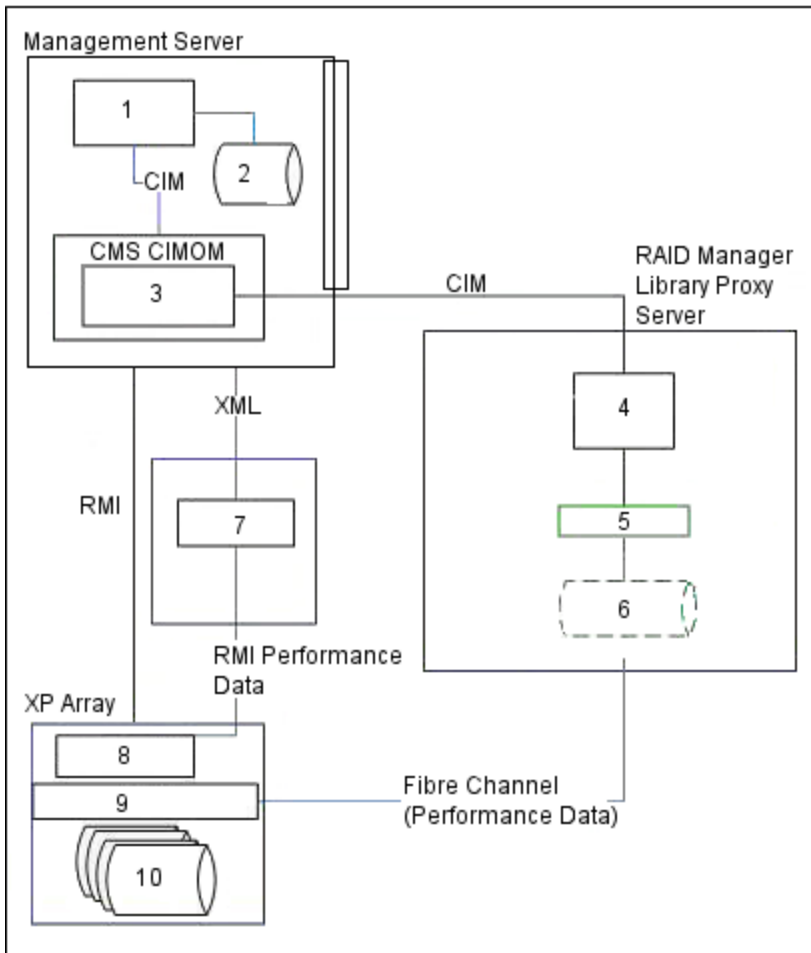
XP and HDS Performance Pack Enterprise Architecture

The Performance Pack Enterprise collection paths for the XP and HDS arrays are very similar. The main difference is that the management server uses the Command View Advanced Edition to collect data from the XP array. Likewise, the management server uses Hitachi Device Manager (HDvM) and the built-in HDS provider to collect data for the HDS array.

The management server uses a proxy host that is connected to the array to gather performance data. You cannot use the same proxy host for both XP and HDS arrays. The proxy host can be used either for multiple XP or HDS arrays, but not for both types of arrays. This host proxy can run the Windows, Linux or HP-UX operating system, and it must have a CIM Extension installed on it.

The RAID Manager Library is installed on the host proxy to get the performance data from the array using a special command LUN.

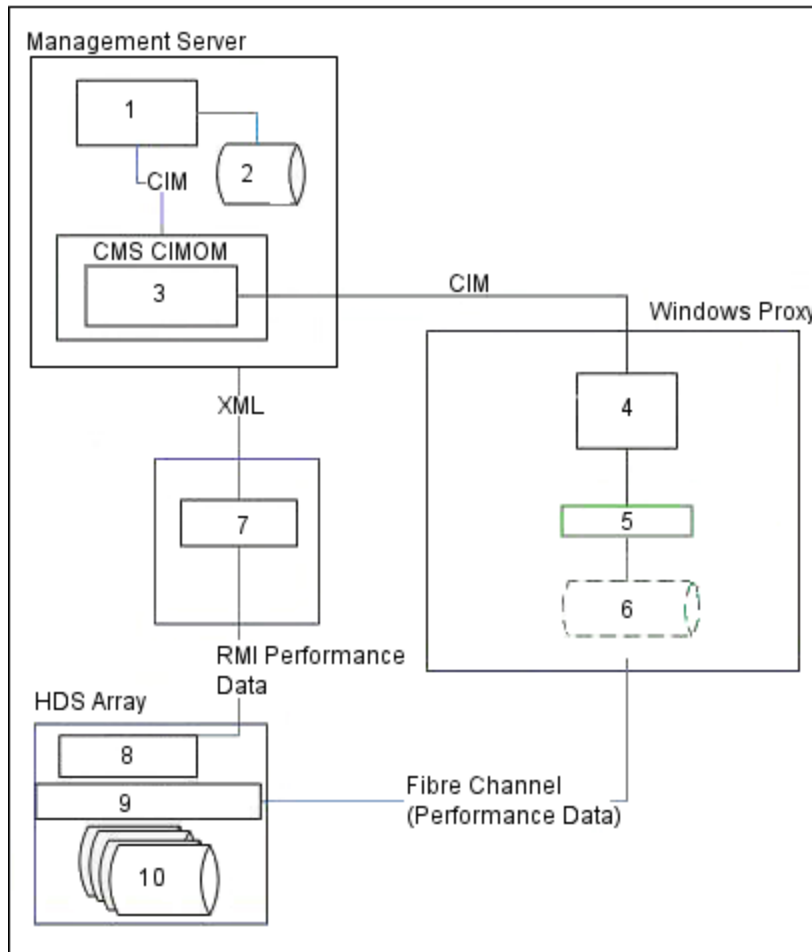
The following figure shows the collection paths for the HP Performance Pack Enterprise software for HP StorageWorks XP arrays.



Legend:

1	Management Server
2	HP Storage Essentials Oracle Database
3	XP Provider
4	HP Storage Essentials Host Agent/CIM Extension
5	RAID Manager Library
6	Command LUN
7	Command View Advanced Edition
8	Service Processor (SVP)
9	Controllers
10	Data

This diagram shows the collection paths for the HP Performance Pack Enterprise software for Hitachi Digital Systems (HDS) arrays.



Legend:

1	Management Server
2	HP Storage Essentials Oracle Database
3	HDS Provider
4	HP Storage Essentials Host Agent/CIM Extension
5	RAID Manager Library
6	Command LUN
7	Hitachi Device Manager (HDvM)
8	Service Processor (SVP)

9	Controllers
10	Data

Refer to [Create the Host Proxy on page 42](#) for information about setting up the host proxy.

The array group, containing the command device LUN that was chosen to gather statistics on the host proxy, will have increased input/output traffic because this command device LUN generates additional input/output traffic as a result of its gathering statistics. Statistical data is returned to the management server where it is stored in the database.

Configuring the Management Server

To configure the management server for the XP and HDS Array Performance Pack:

1. Install a license on your management server with XP and HDS array performance licensing enabled, as described in [License Setup for Performance Pack Software on page 26](#).
2. Discover the array:
 - XP arrays as described in [Discovering HP StorageWorks XP Arrays on page 47](#) for more information.
 - HDS arrays as described in [Discovering HDS Storage Elements on page 46](#).
3. Discover the host proxy by entering the DNS/IP information and appropriate credentials for the CIM extension running on the host proxy.
4. (*Optional*) Use the Test Button corresponding with the host connected to the XP or HDS array that you want to use as the host proxy. The Test button validates the installation of the RAID Manager Library (RMLIB API) and the creation of the command LUN. If a command LUN is available, the first available command LUN is displayed. See [Validate RAID Manager Library Installation on page 43](#).
5. Run a Get Details to get all host and array information.
6. Enable the license for the XP array or HDS array, as described in [License Setup for Performance Pack Software on page 26](#).
7. Go to the Properties page for the XP or HDS array you have licensed for performance statistics. The easiest way is directly from the **Licensing** tab screen. Click the link for the array under the name field, and it will take you directly to the Navigation page for the array. Then, click the **Properties** tab.

Storage System Properties Screen with Proxy Host Field



8. To designate the proxy host that will be used to gather statistics for an array, click **Edit**

Proxy Host. The following representative screen appears.

9. Select the host proxy that was set up, as described in [Create the Host Proxy on next page](#). There is a filter button to narrow down the selections listed. If your host proxy is not in the list, it means you have not run a successful Get Details to create the connection between the host and the array.

Create Command Device (LUN)

You must create a command LUN (command device) on the SLPR0 partition using the HP StorageWorks XP Remote Web Console (RWC) or Hitach Storage Navigator and present it to the port for which the host proxy server has access. The command device can be any device that is accessible to the host.

The command LUN should be from the first SLPR0 partition of the XP or HDS array in the case of RAID600-based or RAID500-based XP array models (which support SLPR partitioning). The SLPR0 command device LUN provides visibility to the entire array regardless of its array partitioning.

Make sure that no data exists on the volume that you select as the command device. Any data that resides on the volume becomes unavailable to the host. The volume that you designate as the command device is used only by the disk array and is blocked from the user. Also, make sure no file system has been mounted and no data is stored there.

Creating a Command LUN

To create a command LUN:

1. Launch one of the following user interfaces using administrator privileges:
 - For XP arrays: Remote Web Console (RWC)
 - For HDS arrays: Hitachi Storage Navigator
2. From the RWC or Hitachi Storage Navigator, select **GO > Lun Manager > LUN Path and Security**. A list of LDEVs appears.
3. Right-click the LDEV that you want to convert into a command device.
4. Select **Enable\Disable** from the pop-up menu.
5. Click **Apply** to save your changes and enable the selected LDEV as a command device.

Caution: Do not mount any file systems on this command LUN.

Additional Related Tasks

Creating a Command LUN might also require that you perform these tasks:

- Zone the SAN switches between the host proxy and the XP or HDS storage array port to open up a path.
- Create a host security group by allowing the command LUN on the XP or HDS port to be exposed to the HBA WWN on the RMLIB Proxy server.

Create the Host Proxy

HP Performance Pack Enterprise software for XP and HDS arrays requires a proxy host for access to both historical and real-time performance data. However, you cannot use the same proxy host for both an XP and an HDS array configured on your storage system. One proxy host can be used *either* for multiple XP arrays *or* multiple HDS arrays, but not for both.

If you use Performance Advisor software to collect XP or HDS array information, HP recommends that you use the same proxy host for Performance Advisor as the proxy host you use for the management server. Both the management server and Performance Advisor use a similar host proxy configuration, and they both use the RAID Manager Library (RMLIB API) and a command LUN.

Operating System Requirements

For XP arrays, you can set up a host proxy server that runs on the Microsoft Windows, Linux or HP-UX operating systems. For Hitachi Data Systems (HDS) arrays, you must set up a host proxy server that runs on Microsoft Windows.

Setting Up the Host Proxy

The following instructions assume that you completed steps 1-3 as described in [About Performance Pack Setup and Configuration on page 24](#).

To set up the host proxy:

1. Verify that the command LUN is accessible to the host bus adaptor (HBA) on the host proxy by using the native HBA tool set.

If you have Performance Advisor and you already installed the RAID Manager Library (RMLIB API), skip Step 2. If you are not sure where RMLIB API is installed, look at the Performance Advisor configuration to see where the agents for Performance Advisor are installed.

2. Install the RAID Manager Library (RMLIB API) (if not already installed). The RAID Manager Library can be obtained as follows:
 - HP StorageWorks XP array: Obtain the RAID Manager Library (RMLIB API) from the array firmware CD. If you do not have the RAID Manager Library software, contact HP services for assistance.

- Hitachi Data Systems (HDS) array: Contact HDS support for the RAID Manager Library for HDS storage systems.
3. Create a command LUN (LUN:0) if you have not already done so. See [Create Command Device \(LUN\) on page 41](#)
 4. Install the CIM Extension on the same host proxy server that has RMLIB API and LUN:0. The following operating systems are supported:
 - For XP arrays: Install the CIM Extension on a host proxy server running Microsoft Windows, Linux, or HP-UX.
 - For HDS arrays: Install the CIM Extension on a host proxy server running Microsoft Windows.

This is the same CIM Extension that HP Storage Essentials uses to manage and discover other hosts. No additional configuration is needed.

5. (*Optional*): Verify that the RAID Manager Library (RMLIB API) is returning data through the Command LUN by using the management server tool called **arrayScan**. See [Verify that the RAID Manager Library Returns Data on next page](#).

Validate RAID Manager Library Installation

You can validate that the RAID Manager Library (RMLIB API) is properly installed and returning data by using the Test Button that corresponds to the host proxy connected to the XP or HDS array.

The Test button validates the installation of the RAID Manager Library (RMLIB API) and the creation of the command LUN. If a command LUN is available, the first available command LUN is displayed.

Output Example

```
Name: Performance Monitoring Proxy Host Command Luns available:
\\.\PHYSICALDRIVE0 : "HP ", "OPEN-V-CM ", Rev"5001"
( Serial# 10118, RAID600or500,LDKC0, SLPR0, CLPR0, RG1-1, LDEV
00:30,
CU 0, RAID5, Port2A, PortWWN:10000000C93F0D68,
NodeWWN:20000000C93F0D68 )
...1 Array Cmd Dev Lun device paths found including any SLPR0 ones
just shown.
Model :Raid-Manager/LIB-XP/WindowsNT
VerandRev:01.12.04
```

The example shows a required SLPR0 command LUN. If the RAID Manager Library is installed, the version is displayed.

Verify that the RAID Manager Library Returns Data

Use the arrayScan tool on the management server to verify that the RAID Manager Library (RMLIB API) is returning data through the Command LUN. This tool is located in the `<CIM_extension_installation_directory>\tools` directory on the host proxy server.

The `./` prefix for arrayScan is only required on non-Windows operating systems.

To obtain information about the arrayScan tool, such as information about additional parameters, use the `-help` or `-?` parameter. For example, typing `arrayScan -?` displays Help information about the tool.

You can also use the Test button for the management server to verify whether the RMLIB API is running properly. See [Validate RAID Manager Library Installation on previous page](#)

Output Example

When the arrayScan tool is used with no parameters, it returns the selected command LUN used to get statistics. Here is an example of the output from the arrayScan tool:

```
arrayScan build date: May 21 2009:16:24:19
Return string...
\\.\PHYSICALDRIVE4 : "HP ", "OPEN-V-CM ", Rev"5001"
( Serial# 10118, RAID600or500,LDKC0, SLPR0, CLPR0, RG1-1, LDEV
00:1E,
CU 0, RAID5 , Port1A, PortWWN:10000000C95C763F,
NodeWWN:20000000C95C763F )
...1 Array Cmd Dev Lun device paths found including any SLPR0 ones
just shown.
...Return string.
Return string length: 293 (0 percent of current max 14680064 bytes).
Largest line length: 116
```

3 Managing Performance Data Collection

The HP Storage Essentials software uses performance data collectors to gather data for Performance Manager. Each time you query a SAN element for performance information, Performance Manager displays the performance data collected.

This section describes performance data collection and how to manage settings and schedules for collectors. It includes the following topics:

- [Discovering Storage Environments for Performance Monitoring below](#)
- [Managing Performance Data Collectors on page 49](#)
- [Managing XP and HDS Data Collection on page 58](#)
- [Managing EVA Data Collection on page 60](#)

Discovering Storage Environments for Performance Monitoring

Before you can collect performance data and begin monitoring the storage elements in your environment, you must run Discovery - Get Details. HP Storage Essentials can then collect performance data for any elements that are discovered.

For proper Discovery of storage elements, carefully follow instructions provided in the *HP Storage Essentials User Guide*.

The information in the following topics provide a Discovery overview of only those storage environments supported by the HP Storage Essentials Performance Pack software. The following are the supported environments and references to detailed information in the *HP Storage Essentials User Guide*.

Storage Environment	Performance Pack	See...
EMC Symmetrix	HP Storage Essentials Performance Pack Enterprise	Discovering EMC Solutions Enabler
Hitachi Data Systems	HP Storage Essentials Performance Pack Enterprise	Discovering HDS Storage Systems
HP StorageWorks Enterprise Virtual Array	HP Storage Essentials SRM Performance Pack	Discovering HP StorageWorks EVA Arrays

Storage Environment	Performance Pack	See...
HP StorageWorks XP	HP Storage Essentials Performance Pack Enterprise	Discovering HP StorageWorks XP Arrays
NetApp	HP Storage Essentials Performance Pack Enterprise	Discovering NetApp NAS Devices

Discovering EMC Symmetrix Storage Elements

EMC Symmetrix arrays are managed by an EMC Solutions Enabler server. To allow the management server to discover the Solutions Enabler and the EMC Symmetrix storage systems it manages, you must edit the `nethost` file (if you set up this file). For details, see the EMC documentation.

Additional administrative tasks required to discover and collect data from EMC Symmetrix arrays using an EMC Solutions Enabler server are as follows:

- **Open port 2707:** Make sure that port 2707 is open between the HP Storage Essentials management server and the EMC Solutions Enabler server. HP Storage Essentials communicates with EMC Solutions Enabler's service/daemon `storsrvd`, which listens on port 2707.
- **Create VCM volume:** You must create and configure a VCM volume on the storage system. You must also configure the VCM database on the EMC Solutions Enabler host. See the *EMC Solutions Enabler Symmetrix CLI Command Reference* for details.
- **Set up only one subnet:** Limit the management server to a single subnet. If your management server is on two or more subnets, discovering a storage array through EMC Solutions Enabler might not work. Limiting the management server to a single subnet allows EMC Solutions Enabler to respond correctly.

For detailed information about discovering EMC Symmetrix systems, see *Discovering EMC Solutions Enabler* in the *HP Storage Essentials User Guide*.

Discovering HDS Storage Elements

To communicate with an HDS storage system and obtain performance data, the management server must be able to communicate with the HiCommand Device Manager that manages it.

Administrative tasks required to discover and collect data from HDS storage systems are as follows:

- **Provide access information for HiCommand Device Manager:** HiCommand Device Manager is required for the management server to communicate with an HDS storage system. You must provide the IP address, user name and password for the server running HiCommand Device Manager to the management server. Do not point to the disk array for the storage system.

- **Open port 2001:** To obtain information about HDS storage systems, the management server must be able to access the port HiCommand Device Manager uses to listen. By default, HiCommand Device Manager listens on port 2001, and the management server assumes this configuration at discovery time. If HiCommand Device Manager uses a different port, specify this other port when you discover HiCommand Device Manager.
- **Set up a secure connection (optional):** The management server communicates with HiCommand Device Manager through a nonsecure connection. If you want the management server to communicate with HiCommand Device Manager through a secure sockets layer (SSL) connection, you must modify an internal property or use HTTPS when you discover HiCommand Device Manager.

For more information about discovering HDS storage systems, see *Discovering HDS Storage Systems* in the *HP Storage Essentials User Guide*.

Discovering HP StorageWorks EVA Arrays

Discovery of HP StorageWorks Enterprise Virtual Array storage systems is performed using HP StorageWorks Command View (CV) EVA. The management server supports the following Command View (CV) EVA array discovery options:

- Command View (EVA) 9.x and its SMI-S provider
- Command View (EVA) 8.x, or 9.0.x and the built-in EVA provider

The following administrative tasks are required for discovering EVA arrays by HP Storage Essentials:

- **Install CV on server not running HP Storage Essentials:** Before you can discover an HP StorageWorks EVA array, you must install the Command View (CV) software on a server that is not running HP Storage Essentials.
- **Discover both active and passive CV EVA servers:** To discover an EVA array, the Command View (CV) server that is actively managing the EVA must be discovered by the HP Storage Essentials management server. Moreover, EVA arrays can only be discovered if they are being actively managed by the CV EVA server. The EVA will not be discovered if only the CV EVA server that is passively managing the array is discovered. To continue collecting EVA data when an EVA fails over to the passive CV EVA server, both the active and passive CV EVA servers must be discovered by HP Storage Essentials.

For more detailed information about discovering EVA arrays, see *Discovering HP StorageWorks EVA Arrays* in the *HP Storage Essentials User Guide*.

Discovering HP StorageWorks XP Arrays

The HP Storage Essentials management server can discover HP StorageWorks XP arrays in your SAN environment using:

- Command View XP Advanced Edition
- Built-in XP Provider

Each of these methods requires the installation of the software required for that method. Information about using each method to discover XP arrays is described in the *HP Storage Essentials User Guide*.

Depending on the method you choose, you might be required to perform these additional administrative tasks:

- **Install Command View XP on server not running HP Storage Essentials:** Before you can discover an XP array, you must install the Command View (CV) XP software on a server that is not running HP Storage Essentials.
- **Install additional software:** Depending on the method you choose, you might need to install the Command View XP Advanced Edition software. The Command View Advanced Edition software must be installed on a server that is not running HP Storage Essentials.

For more detailed information, see *Discovering HP StorageWorks XP Arrays* in the *HP Storage Essentials User Guide*.

For instructions, see the HP StorageWorks XP Disk Array and Command View XP documentation for your version of Command View XP at <http://www.hp.com/go/hpsim/providers>.

Discovering NetApp Storage Systems

To discover a NetApp storage system and obtain performance data, the management server must be able to communicate with the NetApp Network-Attached Storage (NAS) device, also called the NetApp Fabric-Attached Storage (FAS) or NetApp filer. This requires that you provide log-on information for the device to the HP Storage Essentials management server.

Administrative tasks for discovering and collecting data from NetApp storage systems are as follows:

- **Provide access information for the device:** You must provide the IP address, user name and password for the NetApp NAS device to the management server.
- **Contact individual filers over a secure connection (optional):** Configuration of NetApp discovery allows you to contact individual filers through HTTPS, rather than being contacted through an all or nothing approach. As such, you can discover an individual NetApp device using a secure socket layer (SSL) by entering a complete URL in the Step 1 Discovery address field. This enables you to contact the device using port 443, the default NetApp SSL admin port.

For detailed information about discovering NetApp storage devices, see *Discovering NetApp NAS Devices* in the *HP Storage Essentials User Guide*. The topic also includes information about discovering NetApp virtual filers.

Next Steps

The Discovery process is important because performance data can only be collected from successfully discovered storage elements for which you have enabled data collection. After you complete the administrative tasks required to ensure successful discovery of elements in your SAN environment, you can run the HP Storage Essentials Discovery process. Detailed instructions for running Discovery are provided in the *HP Storage Essentials User Guide*.

After the Discovery completes successfully, you can optionally import the HP Storage Essentials Performance Pack licenses you require. The Performance Manager installs with a basic set of performance metrics which allow you to monitor general performance aspects of your system. The Performance Packs enhance your monitoring capabilities by providing additional performance metrics that allow you to chart and graph data at more detailed activity and trend levels. See [License Requirements on page 25](#).

When you are ready, you can next set up and start HP Storage Essentials Data Collection from the Performance Data Collection page. See [Setting Up Performance Data Collectors on next page](#).

Managing Performance Data Collectors

When a new element is discovered, collectors are created so that performance statistics and report data can be gathered periodically. The relevant element type collector setting is automatically applied when a new element is discovered. You can modify the performance data collector settings for discovered elements by navigating to the Performance Data Collection and Default Collector Settings for Newly Discovered Elements pages.

Information about discovered storage elements and the data collectors for these elements is available on the Performance Data Collectors page. For information about viewing this page, see [Displaying the Performance Data Collectors Page on next page](#)

The Performance Data Collectors page lists the collectors that the management server uses to gather performance data. The following information is available for each data collector:

Property	Description
Element	Name of the element as it appears in your storage environment.
Element Type	Type of element managed by the collector, such as Application, Host, or Switch.
Collector Type	Type of collector. The Collector Type value is correlated with the Element Type value.
Enabled	Indicates if collector is enabled to collect data.
Intervals (Minutes)	Time interval in minutes between scheduled data collection activity.
Next Scheduled	Day and time (format:yyyy-mm-dd hh:mm) of next data collection.
Edit	Click this button to edit the data collector properties.
Action	Click Stop to disable the collector process. Click Start to enable collection activity. For more information, see Starting Performance Data Collectors on page 53 and Stopping Performance Data Collectors on page 54 .

You can decide whether data should be collected by starting or stopping the collector schedule. If a collector is gathering data when you stop it, that collector will complete the collection. When you stop a collector, you are suspending any future scheduled collection processes for that collector, but not the currently active process. The Enabled column reflects whether you have enabled the collector schedule or not.

If you are not sure whether you want to start or stop a collector schedule, click **+Statistics** to see the list of performance statistics that the collector is designed to collect. Only after statistics are collected, can you then review the data using the Performance Manager.

Setting Up Performance Data Collectors

Performance data collectors are created when HP Storage Essentials discovers new elements. These performance data collectors gather performance statistics and report data that you can view in the Performance Manager.

When reviewing which data collectors are being used to collect performance data, you have the ability to start and stop a data collector and edit its property settings.

HP strongly recommends that you edit data collection properties only with instruction from HP Support.

All tasks related to viewing and editing performance data collectors are available on the Performance Data Collectors page. See [Displaying the Performance Data Collectors Page](#) below for instructions.

The Performance Data Collectors page lists all data collectors for elements in your storage environment and enables you to:

- Start data collectors—see [Starting Performance Data Collectors on page 53](#).
- Stop data collectors—see [Stopping Performance Data Collectors on page 54](#).
- Edit data collectors—see [Editing a Data Collector Schedule on page 55](#).
- Set data collector filters—see [Applying Data Collector Filters on page 52](#).
- Edit default collector settings for newly discovered elements—see [Default Collector Settings for Newly Discovered Elements on the facing page](#).
- View the properties of a discovered storage element by clicking the element name. This action displays the System Manager Navigation page, which shows the element's properties listed in a table.
- View a list of metrics actively being measured for each element by clicking **Statistics**. For a description of the metrics shown, see the performance metrics tables in [Managing Performance in Your Storage Environment on page 91](#).

Displaying the Performance Data Collectors Page

To open the Performance Data Collectors page:

1. From the Performance Manager top right menu, select **Configuration**. This displays a submenu bar.
2. Select **Performance** from the submenu.

The Performance Data Collectors page appears.

If you are navigating from the Data Aging tab, click the **Data Collection** tab to open the Performance Data Collectors page.

Element	Element Type	Collector Type	Enabled	Interval (Minutes)	Next Scheduled Run	Edit	Action
2001000DEC948501	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:15		Stop
2001000DEC87A0C1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:30		Stop
2003000DEC87A0C1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:33		Stop
2064000DEC948501	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:18		Stop
2064000DEC87A0C1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:41		Stop
20C8000DEC87A0C1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:03		Stop
2137000DEC948501	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 15:22		Stop
brocadempr1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 14:56		Stop
brocadempr1	Switch	Switch Port Collector + Statistics	✓	60	2010-09-29 14:48		Stop

Default Collector Settings for Newly Discovered Elements

The management server collects many different types of data. Instead of using a single large process, data is gathered using many different collectors. You can decide whether all or some of these collector schedules should start or stop when a new element is discovered.

To review the list of collectors that are available for each element type, access the Discovery page by selecting **Configuration > Discovery** and clicking the **Collector Settings** tab.

The Default Collector Settings page provides a comprehensive listing of what collectors are available for each element type and what category each collector is classified as.

To help you find your collectors quickly, this page offers a set of filters. The filter area is collapsed by default. To expand the filter area, click the + (plus) symbol. The following filters are supported:

- **Element Type**—Retrieves the list of collectors designed to support the specified element types.

- **Collector Category**—Specifies whether you want to see a listing of all the collectors, or only collectors designed to gather performance statistics, report data, or capacity data.

To apply the filter settings, click **Filter** to refresh the content of the page. To restore the filters to their default settings, click **Reset**, and refresh the collector page.

You can modify the default collector settings for all future discovery elements by clicking the **Edit** button, or by selecting one or more collectors using the checkboxes and clicking **Edit Selected**.

The following properties are available when you modify default collector settings:

- **Enable selected collectors for newly discovered elements**—Each collector is designed to gather data for a specific type of elements. If you select this checkbox, the management server will start the selected collectors automatically whenever it discovers a system matching the element type supported by a selected collector. This means that the start time of the selected collectors for the newly discovered system is set to the same time as when the system is discovered.

If you do not select this checkbox, the collectors are still available to the discovered elements, but the collector schedules are stopped by default. You can start any of the collectors after a system is discovered using one of the Data Collector Configuration pages; for example, **Configuration > Report** or **Configuration > Performance** or **Configuration > Backup**.

- **Repeat Interval**—Specifies how often the collector should run on a recurring basis when the management server starts a collector schedule for an element type.

Applying Data Collector Filters

The Performance Data Collector page lists a set of data filters to help you find your collectors quickly.

Filter	Description
Element Type	Use this filter to see the collectors for a specific element type or for all elements. When you select a specific element type, the Collector Type filter is adjusted to show you the collector types that are meaningful to the element type you selected.
Collector State	Use this filter to see all schedules or only the schedules that have been started (Enabled) or stopped (Disabled).
Collector Type	Use this filter to see all the collectors for a specific collector type or all collector types. This filter works in conjunction with the Element Type filter.
Element Name Contains	Use this filter to retrieve all the elements whose name contains the specified string.

The Filter settings on the Performance Data Collection page are located above the data collector table. The filter area is collapsed by default.

To select the settings to use as filters:

1. Expand the **Filter** area by clicking the + plus symbol.
2. In **Element Type**, use the drop-down menu to select one of the following options:
 - **All** – Filters for all collectors for all element types listed below.
 - **Application** – Filters and displays collectors for application performance data.
 - **Host** – Filters and displays collectors for host performance data.
 - **Switch** – Filters and displays collectors for switch performance data.
 - **Storage System** – Filters and displays collectors for storage system performance data.
 - **Tape Library** – Filters and displays collectors for tape library performance data.
 - **Fabric** – Filters and displays collectors for fabric performance data.
3. In **Collector State**, use the drop-down menu to select one of the following:
 - **Enabled** – Turns on data collection schedules for the collector you select.
 - **Disabled** – Turns off data collection schedules for the collector you select.
4. In **Collector Type**, the default is ALL. Collector Type works in conjunction with the Element Type field. If you selected "ALL" as the Element Type filter, all collector types are included in the display filter. If you selected a specific element type, for example "Switch," all the collectors for the specified element type are selected for display.
5. In **Element Name Contains**, type the string you want used to filter for specific element names.
6. Click **Filter** to apply your filter settings and refresh the Performance Data Collector table. Click **Reset** to restore filter settings to their default values.

For more information about the Data Collectors tab, see [Default Collector Settings for Newly Discovered Elements on page 51](#).


Starting Performance Data Collectors

You start performance data collection from the Performance Data Collectors page in the Performance Manager. See [Setting Up Performance Data Collectors on page 50](#).

To start a performance data collector:

1. Open the Performance Data Collectors page.
2. Select the element(s) for which you want to start collecting performance data by checking the checkbox beside the element name.
3. Click **Start Selected** for multiple collectors or click **Action** for single collectors. The Start Schedule(s) page appears.

You might prefer to start and run the collectors only for specific needs, rather than running all of them all the time; otherwise it might affect overall system response.

4. In **Next Schedule Run**, enter a date and time. You can click **Select Date**  to use a calendar to select the entry.
5. In **Repeat Interval**, indicate the time interval to use between scheduled collection of performance data. Use the pull-down menu to indicate intervals in Minutes(s), Hour(s), or Day(s) and type in the appropriate number for the time unit you select.

As a rule of thumb, set intervals to reasonable time periods. Do not set intervals that are too small. Make sure a scheduled data collection is able to complete before the next one begins. If collections are scheduled too closely together, system performance, memory usage, and database size might be impacted on the management server.

Verify the date/time specified for the first data collection. By default the first data collection is set up to 1 hour from current time. To increase the start time for the data collectors, set the start date/time to a few minutes in the future rather than the default hour.

6. Check the **Spread Start Time** checkbox if you are editing more than one performance data collector. Selecting this option enables the management server to adjust the start time of the selected collectors and spread the load across the repeat interval you specified in the previous step. Staggering the starting times for multiple collectors minimizes the impact on system performance. The most frequent data collection interval is 3 minutes.
7. Click **OK** to start data collection for the selected elements. Click **Cancel** to discard any entries and return to the Performance Data Collectors page.

Stopping Performance Data Collectors

When you stop a collector, the management server stops gathering information for which the collector is responsible. For example, if a performance collector is not running, its corresponding statistics are no longer receiving information to display. One of the following occurs:

- If there was originally no information gathered for a statistic, no data appears for that statistic.
- If information was previously gathered for a statistic, old data appears for that statistic.

To stop a collector:

1. Access the page for performance collectors (**Configuration > Performance > Data Collection**).
2. Click the **Stop** button corresponding to the collector you want to stop.

The collector stops gathering information for its corresponding reports. To stop more than one collector at once, select more than one collector and click **Stop Collectors**.

Data Collector Schedules


In addition to changing the collector schedule after an element has been discovered, you can change the default collector schedule, determine when you want the collector to run, and whether it should be started or not for future discovered elements.

To see the default collector schedule settings:


- Click **Default Collector Settings** above the Filter area.
Or
- Click **Configuration**, select **Discovery** from the submenu, and click the **Collector Settings** tab.

Editing a Data Collector Schedule

There are several ways to modify a collector schedule:

- To modify one collector at a time—Locate the collector you want to edit in the collector table, and click the **Edit** () button in that collector's row.
- To modify multiple collectors with identical changes—Locate the collectors you want to edit in the collector table, and select the collector using the checkbox in the first column of the table. You can select one or more collectors this way. Click the **Edit Selected** button above the collector table to modify the selected collector's schedules at the same time.
- To modify all of the collectors displayed on the same page of the collector table, click the checkbox in the first column header, and then click the **Edit Selected** button above the table.
- To modify all of the collectors displayed on all of the pages, click the **Select All Pages** link above the collector table, and then click the **Edit Selected** button above the table.

To edit a data collector schedule:

1. Open the Edit Schedule(s) dialog using one of the methods described above.
2. In **Next Schedule Run**, type a date and time (format: yyyy-mm-dd hh:mm) or click the calendar icon  to select a date and time for when the collector should run.

If you change the date in the box to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2010-11-31, the software assumes the date is 2010-12-01.

3. In **Repeat Interval**, specify how often the collector should run on a recurring basis. Use the pull-down menu to select an interval unit of Minute(s), Hour(s), or Day(s), and enter an appropriate number for the unit. For example, select "day(s)" and type "1" to indicate that you want the activity to occur daily.

Although it is possible to set a small performance collector interval, doing so can affect performance, memory usage, and database size on the management server.

4. Check the **Spread Start Time** checkbox, if you are editing more than one collector and you want the management server to automatically adjust the start time of the selected collectors and spread the load of the collectors across the specified repeat interval.
5. Click **OK** or **Cancel**.

Collecting Performance Data in Real Time

You can collect real-time performance statistics for your storage devices using the Monitoring capabilities of the Performance Manager. The data collected is immediately visible in charts that you can create by specifying which device and metric you wish to monitor.

To view collected performance data in real time:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Click the element for which you want to view performance statistics.
3. Under the Monitoring tab in the lower-left pane, select the element again.
4. In the lower-left pane on the Monitoring tab, select the performance metric for which you want real-time statistics and drag the metric into the chart pane to the right.

Real-time performance collection is scheduled in 20-second intervals. Performance Manager displays the last collected data statistics for the storage device you select and refreshes the data as it is collected.

Viewing Data Aging Settings for Performance Statistics

The Performance Manager Data Aging settings enable you to control data aging for both reports and performance statistics.

Data Aging includes settings for the following activities:

- **Garbage Collection (GC)**—Determines how long a set of data is preserved before it is permanently removed from the database. See step 5d below. Set the 'preserve' value to indicate the time period for which you want the data kept.
- **Data Rollup**—Controls how often a set of data is summarized. For example, how frequently hourly data is rolled into the daily table. See step 5c below.

Note: Do not modify the data on the Data Aging page unless instructed by Customer Support. Changing the settings incorrectly can adversely affect the HP Storage Essentials management server.










Data Collection Data Aging

Data Aging for Reports and Performance Statistics:

Data Aging includes Data Rollup and Garbage Collection. Data Rollup controls how often a set of data is summarized. For example, hourly data is rolled into the daily table periodically. Garbage Collection refers to how long a set of data is preserved before it is permanently removed from the database.

The settings on this page control data aging for both reports and performance statistics.



Important: Do not modify any entries unless instructed by customer support. Changing them incorrectly can adversely affect HP Storage Essentials Enterprise Edition.

Short Description	Next Scheduled Run	Interval	Preserve	Enabled	Edit
Raw GC	2010-09-30 04:45	1 day	1 day	<input checked="" type="checkbox"/>	
Hourly GC	2010-09-30 05:00	1 day	1 week	<input checked="" type="checkbox"/>	
Daily GC	2010-09-30 05:15	1 day	30 days	<input checked="" type="checkbox"/>	
Weekly GC	2010-10-03 05:45	1 week	12 weeks	<input checked="" type="checkbox"/>	
Monthly GC	2010-10-09 06:00	31 days	52 weeks	<input checked="" type="checkbox"/>	
Hourly Rollup	2010-09-29 11:10	1 hour	0 minutes	<input checked="" type="checkbox"/>	
Daily Rollup	2010-09-30 03:00	1 day	0 minutes	<input checked="" type="checkbox"/>	
Weekly Rollup	2010-10-03 04:00	1 week	0 minutes	<input checked="" type="checkbox"/>	
Monthly Rollup	2010-10-09 04:30	31 days	0 minutes	<input checked="" type="checkbox"/>	

To view and modify performance data aging settings:

1. Select **Configuration** in the upper right menu bar. This displays a submenu.
2. Select **Performance** in the submenu.
3. Click the **Data Aging** tab.

Perform the following steps *only* if Customer Support instructs you to modify one of the entries on the page.

4. Locate the Garbage Collection (GC) or Rollup description you want. Click the **Edit**  button in the far right column.
5. In the Edit GC and Rollup dialog:
 - a. To enable active garbage collection or data rollup, check the **Enable** checkbox. Uncheck the box to suspend activity.
If you are not allowed to disable the activity, the Enable option is unavailable.
 - b. In **Next Schedule Run**, type a date and time (format: yyyy-mm-dd hh:mm) or click the calendar icon  to select a date and set a time.
If you change the date in the box to a date that does not exist in a month, the software automatically calculates the date to the next month. For example, if you enter 2003-11-31, the software assumes the date is 2010-12-01.
 - c. In **Repeat Interval**, use the pull-down menu to select an interval unit of second(s), minute(s), hour(s), day(s) or week(s), and then enter the number appropriate for the unit. For example, select "day(s)" and type "1" to indicate that you want the activity to occur daily.
 - d. (*Garbage Collection only*) In **Preserve**, identify a period of time for which the data will be preserved prior to removal from the database. Use the pull-down menu to select the unit of time—second(s), minute(s), hour(s), day(s) or week(s) and enter a number that

is appropriate for the unit you select. For example, select "week(s)" and type "4" to indicate that you want the data preserved for 4 weeks.

6. Click **OK**.

Aggregation of Performance Data

Performance data collected in real time is immediately visible in charts from the Performance Manager interface. Raw performance data is stored in the database for later viewing up until the preservation interval you've specified. At which point it is rolled up and averaged into hourly data points. All other performance data is collected at 15 or 60 minute intervals by default.

Raw performance data is aggregated over time using the HP Storage Essentials daily collection and roll-up processes to save disk space on the management server. After one day, the data is aggregated to one hour intervals. After a week, it is aggregated to one day intervals, and so forth, up to monthly and yearly aggregations. Values are aggregated according to the weighted average for the sampling period, and a storage administrator can configure the collection frequency for the SAN components within the monitored system.

Customizing Performance Data Collection

Follow these best practices to ensure that you are getting optimum performance data with the least impact on your system's performance.

- Disable collection of metrics on systems that are not business critical.
- Set longer collection times for less important systems (for example, once a day, or 4 times/day).
- Disable collection (where applicable) of specific metrics which are not useful for your reporting or corporate policies.
- Leave the default of 60- or 15-minute collection intervals (where applicable).
- Set collection intervals to as low as 5 minutes for key systems only.

When you use the real-time collection feature of Performance Manager to detail a specific frequency of collection (which is fixed and not user-selectable), follow these guidelines based on device class:

Element Type	Frequency
Application, Host, or Switch	15, 20, or 30 seconds
Storage arrays (with Performance Pack plug-ins)	20 seconds

Managing XP and HDS Data Collection

You can configure and enable collectors for the XP and/or HDS array(s) you want monitored. For an overview of how to enable, edit, start, and stop a collector, see [Setting Up Performance Data Collectors](#) on page 50.

Tip: Pay particular attention to the date/time specified for the first data collection. By default this will be up to 1 hour from current time. To start the data collectors more quickly, set the start date/time to a few minutes in the future rather than the default hour.

The data collectors included in the Performance Pack Enterprise software for XP and HDS arrays are:

- Array Group Collector
- Backend Controller Collector (only if built-in provider used)
- Frontend Controller Collector
- Storage System Collector
- Volume Collector

You can also select one of the following port collectors based on the type of array:

Storage System Port Observer: For pre-RAID500 arrays (XP128K, XP512K, XP1024K). Requires switch to be discovered.

Host Fibre Port Collector: For all other arrays

For more information about configuring and enabling performance collectors, see [Managing Performance Data Collection](#) on page 45.

About Setting Up XP Array Metric Collection

XP performance data collection works by collecting data in-band from a command device. You need to setup at least one host in the SAN with a command device from the array, with a CIM Extension installed. The CIM Extension automatically detects the command device. After running Discovery Get All Element Details, view the array's properties and click **Edit Proxy Host** to select the host you want used to collect performance data.

Collecting XP and HDS Array Performance Metrics

Before you begin to collect performance metrics, read [Viewing Performance Data on page 79](#) in this guide. It describes the Performance Manager and the data collection methods for HP XP and HDS arrays available for your system.

Performance metrics for XP and HDS arrays include global resources such as cache and shared memory, front-end controllers, XP back-end controllers, ports, storage pools and their constituent volumes, Thin Provisioning/Snapshot Pools and V-VOLs, and the array groups.

The data collection process might take a considerable amount of time to complete and display in the Performance Manager. In certain cases, this might be as long as a couple of hours. Performance Manager requires two points to plot the first data point, and this affects the collection time interval along with the size of your storage array, the collector setup, and other related factors.

The minimum collection interval that can be set for the XP and HDS performance data collectors is one minute. However, collecting at frequencies less than five minutes, increases the amount of data in the database. HP recommends that you use an interval between 15 minutes and one hour. Optionally, set the collection interval lower while analyzing a problem, then restore the collector interval to the default. The interval for real time collection is always 20 seconds.

Non-performance data can be collected from the array using either the CV SMI-S provider or the built-in XP provider. However, you are required to use the built-in provider to collect back-end controller statistics. You will only see front-end controllers in the Performance Manager display when you connect to the array in Step 1 Discovery using CV SMI-S provider. Be sure to use the built-in provider in Step 1 Discovery to collect Back-end Controller statistics. See the *HP Storage Essentials User Guide* for instructions about how to do this.

When collecting performance statistics for XP logical array groups, be aware that a parity group is divided into several logical RAID groups. Thus, for example, parity group X-X can contain logical array groups X-X-1, X-X-2, and so forth. The collecting of performance statistics for logical array groups is only done at the parity group level. Because the Performance Manager for XP array statistics does not show the parity group, all the statistics are gathered at the first logical RAID group only (in our example, X-X-1), and no statistics are gathered at X-X-2 and the other subsequent logical RAID groups.

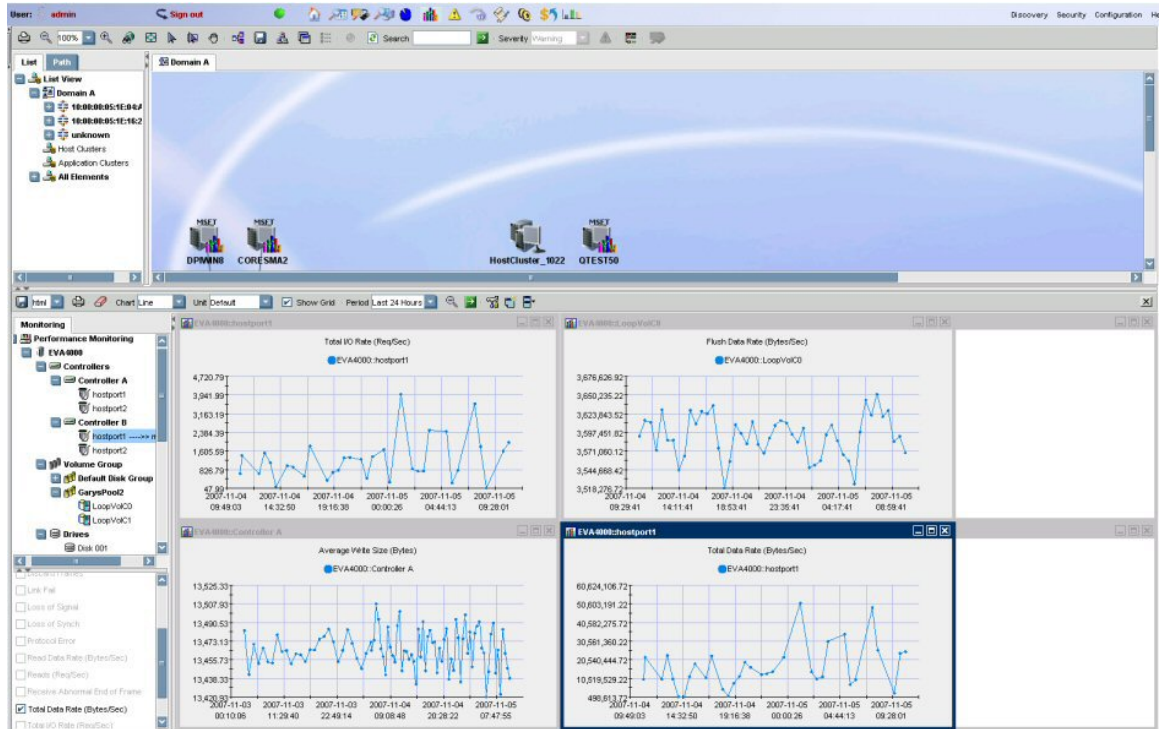
Managing EVA Data Collection

After applying the SRM Performance Pack for Enterprise Virtual Array (EVA) license to your target arrays, you must select which corresponding array-specific performance data collectors you want to enable. Each EVA Storage System-specific row in the list of performance data collectors represents a distinct set of performance metrics for the corresponding EVA array.

For information about enabling, starting, stopping, and editing collectors, see [Setting Up Performance Data Collectors on page 50](#).

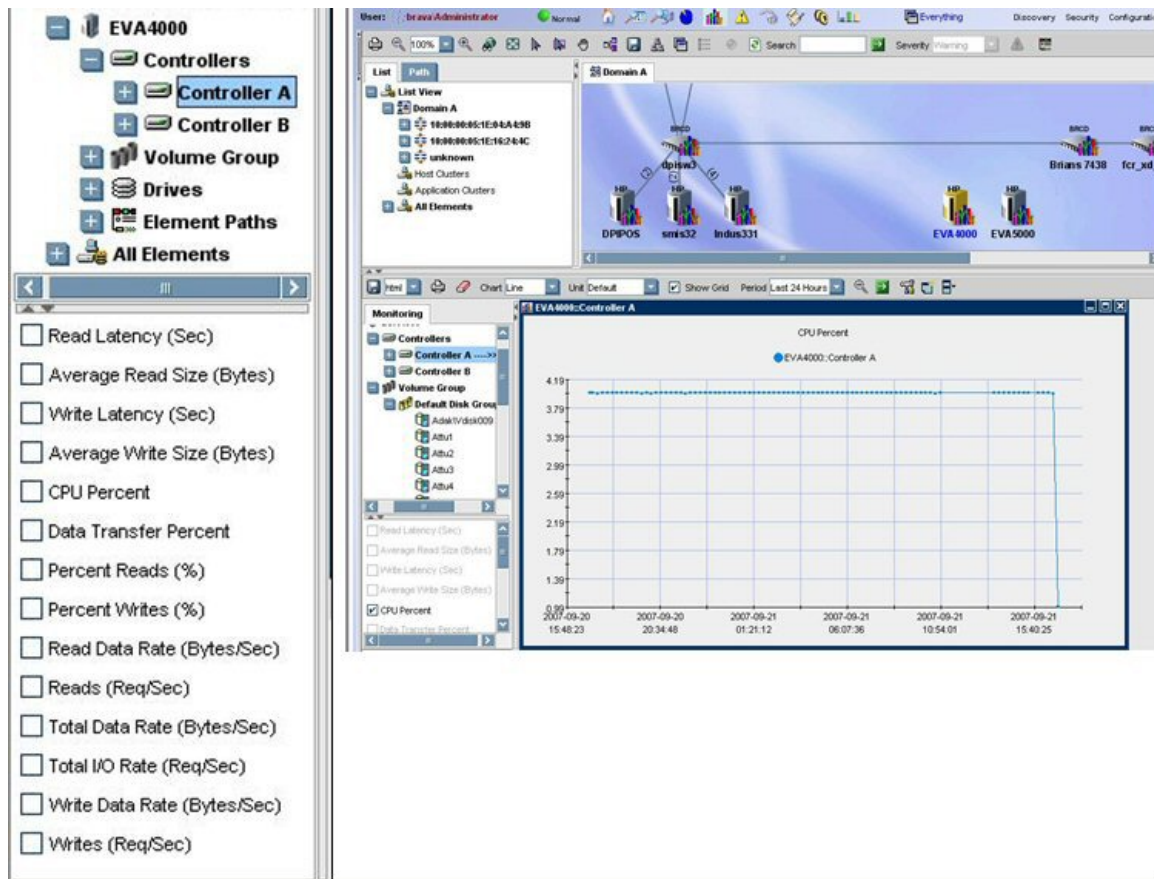
When you start a collector, you are prompted to specify a sample interval and start time for the collector. You can edit a collector's properties after a collector is started by selecting the icon for the collector. Be aware that the type, quantity, and frequency of data collectors started can affect system performance. For example, EVA physical disk metrics collection is I/O intensive, so it should be avoided unless needed.

The following example shows EVA metrics in the Monitoring tab of the Performance Manager.



The list of sub-elements and metrics shown in the Monitoring tab vary depending upon the device type. The selections determine the performance data retrieved and available for analysis.

For example, the following example shows information about data rates associated with the highlighted EVA array host port.



Collecting EVA Performance Metrics

This section describes recommendations and best practices for collecting performance metrics for HP StorageWorks Enterprise Virtual Array (EVA) systems in your environment.

It includes these topics:

- [EVA Configuration and Metrics Collection on the facing page](#)
- [EVA Performance Best Practices on the facing page](#)
- [Metric Collection for EVA Physical Disk Drives on page 64](#)

As a first step, read [Managing EVA Data Collection on page 60](#) section in this guide to understand data collection for HP EVA and which features are available for your system depending upon installed products and licensing.

No performance metrics are gathered by HP Storage Essentials during Get Details (GAED). Running the Get Details process takes time, and therefore should not be run when you need to monitor array performance or collect performance statistics.

EVAPerf must be running before you can obtain performance statistics for EVA arrays.

EVA Configuration and Metrics Collection

Any configuration changes to your HP EVA array (for example, addition or deletion of mapped volumes) require a Get Details before changes can be reflected in the Performance Monitoring tree, and in the actual performance data collected. Also, changes in the number of elements (for example, disks or volumes) can impact the aggregated performance metric data.

If you make changes to your selection of EVA arrays later, you must ensure the selected new EVA arrays have been discovered using the required Command View EVA version. For more information regarding discovery procedures, refer to the *HP Storage Essentials User Guide*.

Command View (CV) EVA supports a standby configuration whereby an array can be managed by multiple CV EVA stations simultaneously, with only one of them actively managing a specific array at a time. Discovery will include both the EVA array being actively managed and the array not being actively managed (Standby CV station). In such cases, there are limitations to the information that can be gathered, based upon whether both the managed and the unmanaged (standby) arrays are discovered by the management server. If only the unmanaged array is discovered, only the top level array information can be collected.

When collecting metrics for EVA arrays, consider the following:

- EVA performance data *cannot* be collected for EVA arrays that have controller passwords that have been set using the Operator Control Panel (OCP).
- Only EVA volumes that are mapped to a host have associated performance data.
- The management server will not be able to gather host port metrics if the EVA controller has two switches attached to it with the same DomainID, and if the same switch port number is used in each switch.

EVA Performance Best Practices

To optimize performance of your HP Enterprise Virtual Array (EVA), follow these best practices:

- Establish a baseline.
- Track IOPS rates and throughput.
- Track error rates, average queue depth and response times.
- To maximize single array performance: Fill the EVA with as many disk drives as possible.
- To maximize application performance: Consider a performance analysis before increasing the total disk count above 168.
- Configure as few disk groups as possible.
- Use 15K-rpm disks to provide the highest performance.
- To reduce the utilization of the mirror port, migrate the high I/O Intensive Vdisk from Vraid 5 to Vraid 1.
- Create disk groups using equal capacity disks.

- Attempt to optimize performance when disks of unequal capacity are present: Unless there is a specific plan to manage the differences in access density of the disk groups, it is preferable to combine disks with different capacities in the same disk group.
- Always leave read caching enabled.
- Attempt to balance the workload demand evenly across both controllers on an EVA.
- Although both controllers in an EVA can access all physical disks, a Vdisk is assigned to be managed by only one controller at a time. Consequently, a single Vdisk can only use the resources of a single controller, such as cache, data path bandwidth, Fibre Channel ports, or processor cycles. Under heavy workload demand, it is possible that only one of the controllers in an EVA pair will reach one of its resource limits if the demand is not equally distributed across both controllers. Distributing the workload evenly across both controllers ensures the maximum total performance of the array.
- Enable only the collection metrics that will actually be used. Enabling all collectors, or more than are necessary, can result in decreased Command View EVA responsiveness and intermittent performance data collection failures.

Metric Collection for EVA Physical Disk Drives

HP Enterprise Virtual Array (EVA) physical disk metrics collection is I/O intensive. Therefore, HP recommends that administrators not enable the collectors for EVA physical disk drive metrics unless required for troubleshooting.

Because there is a large amount of data associated with disk drive metrics, collection of these metrics can impact the responsiveness of Command View (CV) EVA. Collection of EVA disk drive metrics should only be done when instructed by HP service personnel or when troubleshooting a specific issue that requires examining performance data for one or more disk drives.

EVAPerf Data Collector Considerations

When running the EVAPerf Data Collector service, consider the following:

- EVAPerf Data Collector service can sometimes crash, especially in larger environments. This is less likely with EVAPerf 7.0.1 or later. Restarting the service will enable performance data collection to succeed. Consider changing the service startup parameters to **automatic** and **restart on failure**.
- If the total number of vdisks configured on the EVA exceeds 512 using EVAPerf versions prior to v7.0.1, some mapped vdisks might not have any metrics associated with them.
- Occasionally, reporting might not occur for a physical disk because the EVAPerf Data Collector returns an unknown identifier for it.
- Vdisk and physical disk metric data viewed in the EVAPerf CLI on the Command View EVA server might appear inconsistent with the management server's data. This is due to an EVAPerf CLI situation in which only point-in-time samples are returned for sample intervals greater than 30 seconds (this is for all VCS versions).

- Socket connections between the management server and the EVAPerf Data Collector service on port 860 are not immediately closed after stopping the EVA performance collectors in the management server, or when un-licensing EVA arrays for performance data collection in the management server. However, the connection will eventually time out and be closed after the default idle time for the connection expires. In these situations, although the connections are still open, no I/O occurs over the connection.
- The minimum collection interval that can be set for the EVA performance data collectors is 1 minute. However collecting at frequencies less than 5 minutes, especially if all collectors are enabled, might result in decreased stability and failure to collect metric data. The collection interval for real-time metrics is 20 seconds and cannot be changed. Frequent collection of large amounts of performance data, combined with regular I/O, might impact CV EVA responsiveness. However, this does not impact the data I/O of the EVA array.
- Only Vdisks that are mapped to a host have performance metrics associated with them. Only grouped physical disks have performance metrics associated with them.
- The default idle time-out is 2 hours and is not configurable.
- When using the internal SMI-S provider, if no performance data is collected for a period of 2 hours, the connection to the EVAPerf is closed due to an idle time-out. This will cause the EVA performance provider to be reset for the affected EVA array resulting in a failure to display a data point at the next scheduled collection interval. The failure to display the data point is because 2 consecutive data points are required after a reset in order to derive the rate data between 2 samples. Therefore, if EVA array performance collectors are scheduled at a frequency greater than 2 hours and the start times for that EVA array collectors are not staggered (that is, resulting in 2 hours periods with no collections for that EVA array), then no data will ever be displayed for that EVA array.

4 Setting Performance Policies and Alerts

Performance policies monitor the performance of your storage arrays, devices, and volumes based on specific conditions or performance thresholds that you set. Using the HP Storage Essentials Policy Manager you can set up policies to trigger alerts and events whenever these specified conditions or thresholds are reached.

For information about performance policies, see these topics:

- [Viewing Performance Policies on page 69](#)
- [Adding Performance Policies on page 69](#)
- [Modifying a Performance Policy on page 71](#)
- [Deleting a Performance Policy on page 73](#)

For information about alerts, notifications, and events, see these topics:

- [Setting Up E-Mail Actions for Performance Policies on page 73](#)
- [Setting Up Custom Commands for Performance Policies on page 74](#)
- [Setting Up Event Generation for Performance Policies on page 75](#)

About Policy Manager

The HP Storage Essentials Policy Manager manages the performance policies for your storage environment. Create performance policies to monitor performance statistics, such as utilization and I/O rates, and critical system events that might impact throughput and availability. Policy Manager sends notification e-mails, generates events or runs custom command scripts either when events occur or when measured performance indicates potential issues.

Policies for storage management only work if you have the Performance Pack license installed. For this reason, as a pre-requisite to using Policy Manager performance policies, you must install an HP Storage Essentials Performance Pack license for the specific system you want to monitor. (For information about license requirements and set-up instructions, see [License Setup for Performance Pack Software on page 26](#).) Performance Pack licenses are available for EMC Symmetrix, HP EVA and HP XP arrays, and NetApp NAS devices.

Only after you install a Performance Pack license, can you then set up alerts that notify you, for example, when specific performance thresholds are exceeded or specific events occur. For information about setting alerts, see [Setting Performance Policies and Alerts above](#).

Policies and Data Collection

After you install the Performance Pack licenses for your system, you can set up performance policies to help you monitor performance events and alert you to any throughput or availability problems. You can also review the data collectors created by HP Storage Essentials for those storage elements for which you have set policies. (See [Displaying the Performance Data Collectors Page on page 50](#))

Configure your data collectors to collect specific performance statistics that can be correlated to the acceptable performance thresholds you set in policies. Be aware that policies are not triggered when a collector is running. If you configured a policy that alerts you when a host has a critical event and the collector is running when that critical event occurs, you will not be notified of the event.

If you do not want a data collector to run during critical event periods, you can stop the data collector or reschedule data collection for another time period. See [Editing a Data Collector Schedule on page 55](#) and [Stopping Performance Data Collectors on page 54](#).

Severity Levels

Understanding the severity levels used by HP Storage Essentials is important when setting up policy events, viewing policy settings, and analyzing storage element data.


The severity levels for a storage element are set by the vendor or manufacturer of the device. To understand what a severity level represents, it is best to view the description of the event that occurred. The descriptions below provide general meanings for each level. You should view the event for an exact description.

The following severity levels are available when you create or modify a performance policy:

- **Unknown** – Severity level is not known. The event does not fall into the other categories; further information cannot be obtained from it.
- **Informational** – Indicates the event provides informational data only. An example of an informational event is a progress report event for a firmware download operation currently in progress.
- **Warning** – Indicates the event provides important warning information to an administrator. An example of a warning is one or more new physical fabric objects (device port, switch, or fabric) have been discovered.
- **Minor** – Indicates the event provides information about a minor problem for an administrator. An example of a minor event is a physical fabric object (switch port or fabric) has changed state.
- **Major** – Indicates the event provides important information about a major problem for an administrator. An example of a major event is one or more physical fabric objects (device port, switch, or fabric) have disappeared.
- **Critical** – Indicates the event is highly impactful and requires immediate attention. An example of a critical event is Brocade switches that have a failed firmware download and the failure reason code for each respective switch.

Opening Policy Manager

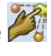
Policy Manager is a software component of HP Storage Essentials. When you launch the HP Storage Essentials software, you will find a link to Policy Manager in the Navigation panel on the left.

You will also find the Policy Manager icon  displayed in the lower half of the right window. Click this icon to open Policy Manager.

If you are already working in HP Storage Essentials, navigate to Policy Manager by clicking **Policy Manager** in the navigation panel on the far left.

Viewing Performance Policies

To view performance policies:

1. Open the Policy Manager by clicking the  icon (see [Opening Policy Manager on previous page](#)).

If you are already working within the Storage Essentials user interface, click the **Policy Manager** icon in the navigation panel on the left or if available, in the toolbar menu at the top.

2. Expand **Policies** by clicking the **+** plus symbol.
3. Locate **Performance Policies** and expand the tree by clicking the **+** plus symbol.
4. Performance Policies are organized by element type. To locate a performance policy in the tree, first drill down to the element type—Applications, Hosts, Storage Systems, Switches—and then to the specific element for which you want to view a policy.
5. To view a policy, click the element name in the performance policy tree. The policy is displayed in the Performance Policies page in the right pane.

If no policy is defined for an element, you can click **Add** to create one. See [Adding Performance Policies below](#).

You can sort on policies by name, description, and action. The Action column indicates whether the resulting policy action is an e-mail notification, generated event, or command script.

From the Performance Policies page, you can also:

- Test the policy to ensure it triggers successfully
- Modify the policy properties
- Delete the policy

See the *HP Storage Essentials User Guide* for more information about managing HP Storage Essentials policies.

Adding Performance Policies

You can create a performance policy that results in event generation, e-mail notification, or running a custom command when a discovered element is being overused. For example, you can configure Policy Manager so you receive an e-mail message when the Cache Writes Pending value for a Cache Logical Partition (CLPR) approaches 70%.

Caution: Policies for storage systems, including performance policies, only work if you have the performance license pack installed. See [License Setup for Performance Pack Software on page 26](#).

To add a performance policy:

1. Complete the steps in [Viewing Performance Policies on previous page](#).
2. Click **Add**.

Because you must select a policy template first, the Policy Template tab is open.

3. Create a template for the performance policy by selecting one of the performance metrics listed. For example, if you create a performance policy for a switch in an XP storage environment, you might select the Aggregated Port Bytes Transmitted metric or the Aggregated Port Bytes Received metric.

In step 9, you will select a performance threshold (condition) for this metric which, when reached, triggers a policy action, such as a notification e-mail that alerts you to the occurrence.

4. When ready, click **Next** or click the **Policy Properties** tab.
5. In **Name**, accept the default policy name (the element/device name plus the template/metric name) or type a new policy name.

To distinguish the policies from one another, include the condition of the policy in its name, for example, Array Groups Average Write Size >35 bytes. You could also include the name or type of devices the policy is monitoring.

6. In **Description**, accept the default text or enter a new description.
7. In **Re-arm period**, accept the default interval or enter a new time period in minutes. The re-arm period specifies the amount of time after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the policy can trigger. If you set the re-arm period to zero, the policy triggers each time conditions for the policy are met.

For example, assume you set the collector for the element to run every 2 minutes. When the collector runs, the condition for the policy is met, and an event is generated. But Event Manager is now being flooded by events. You want to be notified when the condition for the policy is met, but you want to be notified every hour instead of every 2 minutes. You would then enter 60 in the Re-arm Period field, and the policy would run every hour.

8. In **Condition**, set the threshold at which you want the policy triggered. Refer to the HP Storage Essentials metric tables to help you understand when a condition is appropriate for the policy template you selected. See [Managing Performance in Your Storage Environment on page 91](#).

Use the **Comparison Operator** drop-down menu to select one of the following:

- = — "equal to"
- != — "not equal to"
- > — "greater than"

- **>=** — "greater than or equal to"
- **<** — "lesser than"
- **<=** — "lesser than or equal to"

The field to the right of the Comparison Operator shows the units of measure for the performance metric you selected as the template. For example, if you selected the Aggregated Port Bytes Received metric for a switch policy, the unit of measure is MBytes/Sec.


Enter the numeric value for the performance threshold that you want applied to the policy. For example, if you specified the Cache Writes Pending metric when you created the policy template, you might now indicate that you want the policy to be triggered when the metric value becomes greater than (>) 60 percent (%).

9. Select the **Actions** to occur when the policy triggers. You can choose:
 - **Send E-mail** – You must set up an e-mail server to select this action. (To enable an e-mail server go to **Configuration > E-mail Server**.) The **New E-mail Action** dialog is displayed. In the comma-separated list of e-mail addresses field, type the email address to which you want the notification sent. Click **OK**. For more information, see [Setting Up E-Mail Actions for Performance Policies on page 73](#).
 - **Generate Event** – The **New Event Action** dialog is displayed. In **Event severity**, use the drop-down menu to select a severity level and then click **OK**. See [Severity Levels on page 68](#) for descriptions of levels. For more information, see [Setting Up Event Generation for Performance Policies on page 75](#).
 - **Execute Custom Command** – In the **New Custom Command** dialog, type the command you want to execute on the HP Storage Essentials management server. For more information, see [Setting Up Custom Commands for Performance Policies on page 74](#).
10. Click **Finish**.

Modifying a Performance Policy

You can edit an existing policy. You might want to do this if the threshold requirements for the performance of a device change, or if you want to receive an email notification when performance for an element depreciates.

To modify a performance policy:

1. Complete the steps in [Viewing Performance Policies on page 69](#).
2. Click the **Edit** () button that corresponds to the policy you want to modify.

You can change the following properties for the policy:

- a. In **Name**, accept the default policy name (which is the element/device name plus the template/metric name) or type a new policy name.

To distinguish the policies from one another, include the condition of the policy in its name, for example, Array Groups Average Write Size >35 bytes. You might also want to include the name or type of devices the policy is monitoring.

- b. In **Description**, accept the default text or enter a new description.
- c. The **Active** field indicates whether the policy is enabled or not. If the policy is active, the box is checked. To de-activate the policy, uncheck the checkbox.
- d. In **Re-arm period**, accept the default interval or enter a new time period in minutes. The re-arm period specifies the amount of time after the policy executes before it can execute again. The re-arm period is useful for limiting the number of times the policy can trigger. If you set the re-arm period to zero, the policy triggers each time conditions for the policy are met.

For example, if you set the collector for a port to run every 60 minutes, the policy would run every hour.

- e. In **Condition**, set the threshold at which you want the policy triggered for the performance metric. Refer to the HP Storage Essentials metric tables in [Managing Performance in Your Storage Environment on page 91](#) to understand which conditions are appropriate. This field will be related to what you understand to be normal performance parameters for the storage device.
- f. Use the **Comparison Operator** drop-down menu to select one of the following:
 - o = — "equal to "
 - o != — "not equal to"
 - o > — "greater than"
 - o >= — "greater than or equal to"
 - o < — "lesser than"
 - o <= — "lesser than or equal to".

The field to the right of the Comparison Operator shows the units of measure for the performance metric you selected as the template. For example, if you selected the Aggregated Port Bytes Received metric for a switch policy, the unit of measure is MBytes/Sec.

Enter the numeric value for the performance threshold you want applied to the policy. For example, if you specified the Cache Writes Pending metric when you created the policy template, you could now specify the policy to be triggered when the metrics value becomes greater than (>) 60 percent (%).


- 3. Select the **Actions** to occur when the policy triggers. You can choose:
 - **Send E-mail** – You must set up an e-mail server to select this action. (To enable an e-mail server go to **Configuration > E-mail Server**.) The **New E-mail Action** dialog is displayed. In the comma-separated list of e-mail addresses field, type the email address to which you want the notification sent. Click **OK**. For more information, see [Setting Up](#)

E-Mail Actions for Performance Policies on the facing page.

- **Generate Event** – The **New Event Action** dialog is displayed. In **Event severity**, use the drop-down menu to select a severity level and then click **OK**. See [Severity Levels on page 68](#) for descriptions of levels.
 - **Execute Custom Command** – In the **New Custom Command** dialog, type the command you want to execute on the HP Storage Essentials management server. For more information, see [Setting Up Custom Commands for Performance Policies on next page](#).
4. Click **Finish**.

Deleting a Performance Policy

To delete a performance policy:

1. Complete the steps in [Viewing Performance Policies on page 69](#).
2. Select the policy you want to delete from the list by checking the box next to the policy name.
3. Click **Delete** ().
4. You are asked to confirm the deletion of the policy. To continue with the deletion, click **OK**.

Setting Up E-Mail Actions for Performance Policies

You can configure Policy Manager to provide an e-mail notification when conditions specified for the policy are met, such as exceeded performance thresholds.


Assigning an SMTP Mail Server

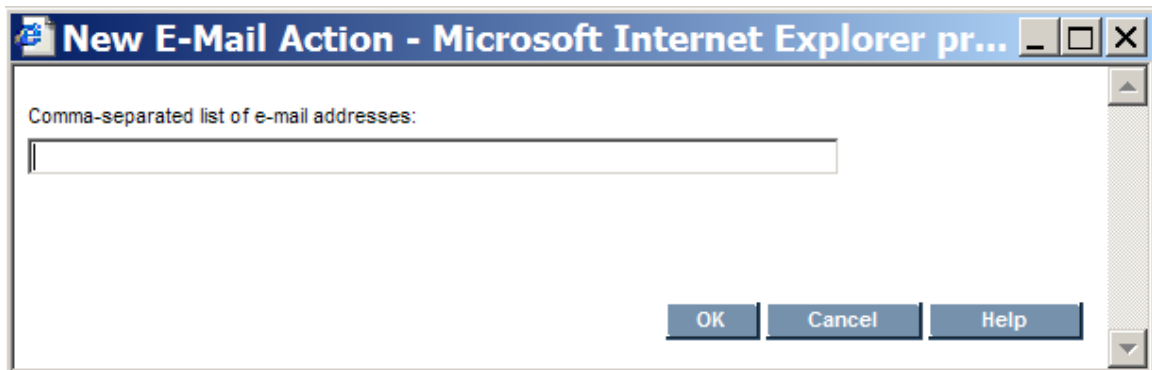
If you have not already done so, you must set up and assign an SMTP server from which the management server can send its e-mail notifications. Information about how to set up an SMTP mail server is provided in the *HP Storage Essentials User Guide*.

Setting Up E-Mail Notification for Policies

The Policy Manager sends an e-mail when the condition you specify in the policy is fulfilled. The software verifies that the address is entered in the correct format.

To set up e-mail notification for a policy:

1. Open the Policy Manager as described in [Opening Policy Manager on page 68](#).
2. Click the **Edit** () button corresponding to the policy you want to modify.
3. Click **Send E-mail**.
4. In the New E-Mail Action dialog, enter a comma-separated list of e-mail addresses to which you want a notification sent.



5. Click **OK**.

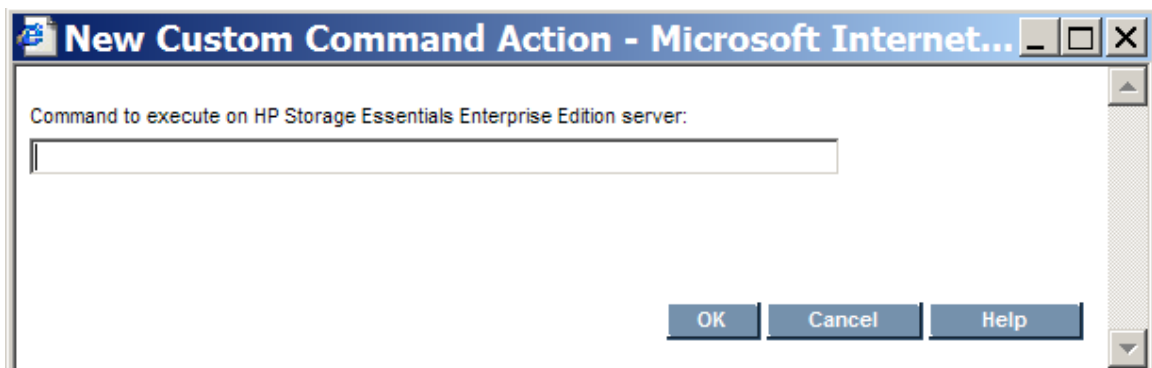
Setting Up Custom Commands for Performance Policies

You can configure Policy Manager to run a custom command on the management server when an element reaches the condition you assign in the policy. The Policy Manager executes a remote script on the management server when the condition that you specify is fulfilled.

Setting Up Custom Scripts

To set up a custom script for a policy:

1. Open Policy Manager as described in [Opening Policy Manager](#) on page 68.
2. Click the **Edit** (✎) button corresponding to the policy you want to modify.
3. Click **Execute Custom Command**.



4. Enter a command that will execute the script in the box, and click **OK**.

The software assumes you are in the `%JBOSS4_DIST%\server\appiq\remotescripts` directory on the management server when the script is executed. You can use environment variables in your script, such as `POLICY_NAME` and `POLICY_DESCRIPTION`, where `POLICY_NAME` provides the policy name and `POLICY_DESCRIPTION` provides a description of the policy.

Prefix the command with “start” if the custom command triggers a user interface component, for example, Microsoft Internet Explorer or a command prompt window. For example, if you want the custom command to open a command prompt window and list the contents in the directory, you would prefix the command as follows:

```
start dir
```

5. After you enter the command, click **OK**.

Setting Up Event Generation for Performance Policies

You can configure Policy Manager to generate a performance event when an element reaches the performance metric threshold you specify for it.

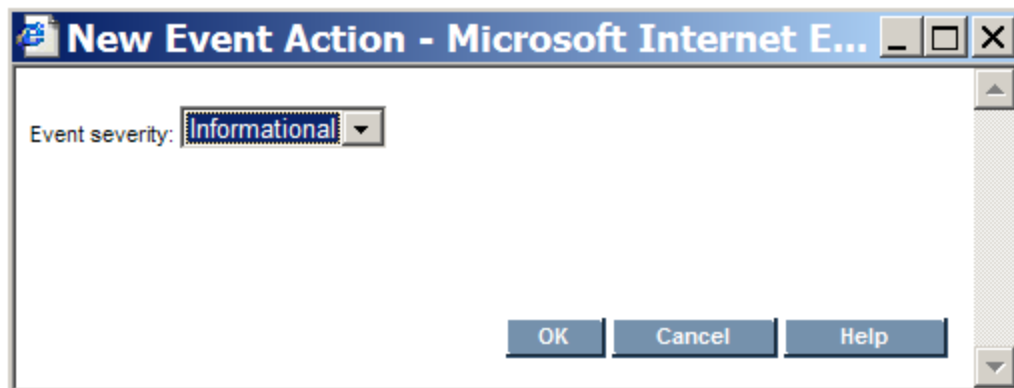
See the *HP Storage Essentials User Guide* for a list of the events supported by HP Storage Essentials.

Setting Up Policy Event Generation

The following instructions assume you are either adding or modifying a policy and therefore, have completed steps 1 - 9 in [Adding Performance Policies on page 69](#) or steps 1 - 3e in [Modifying a Performance Policy on page 71](#).

To set up event generation for a policy:

1. Open the Policy Manager. See [Opening Policy Manager on page 68](#)
2. Locate the policy you want to modify. See [Viewing Performance Policies on page 69](#)
3. Click the **Edit** (✎) button that corresponds to the policy you want to edit.
4. Click **Generate Event**.



5. In the **Event severity** field, use the drop-down menu to select the severity level at which you want a Policy Manager event generated. For descriptions of severity levels, see [Severity Levels on page 68](#).
6. Click **OK**.

Viewing Performance Event Details

The events generated by performance policies are managed by the HP Storage Essentials Event Manager. You can view event details using the event summary in the Event Details pane.

To access event details:

1. Access Event Manager by double-clicking an element displayed in the Performance Manager and then clicking the **Events** tab. This displays the events for the selected element.

To display all events, click the  Event Manager in the left navigation panel.

If you did not create performance policies for an element that generate events, no events will display for the element. See [Adding Performance Policies on page 69](#).

2. When you open the Event Manager, a summary page similar to the following example is displayed:



The screenshot shows the Event Manager interface with a table of events. The table has the following columns: ID, Severity, Time, Element, Summary Text, Event Type, Count, Cleared, and Delete. The events listed are all CIM Alert Events, with various severity levels (Warning, Error) and times. The 'Cleared' column indicates whether the event has been cleared, with 'Not Cleared' or 'Not Cleared' (in red) shown. The 'Delete' column contains a trash icon for each event.

ID	Severity	Time	Element	Summary Text	Event Type	Count	Cleared	Delete
37678	Warning	2010-10-06 11:38	000190300595 (Symm7.DMX3-6)	2010-10-06 11:33: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37647	Warning	2010-10-06 11:35	000190300595 (Symm7.DMX3-6)	2010-10-06 11:30: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37646	Warning	2010-10-06 11:35	000190300595 (Symm7.DMX3-6)	2010-10-06 11:28: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37618	Warning	2010-10-06 11:32	000190300595 (Symm7.DMX3-6)	2010-10-06 11:26: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37582	Warning	2010-10-06 11:29	000183500570 (Symm48:3830)	2010-10-06 12:00: 00.0033.00 (SEVT_ACCESS_TO_NR_DEVICE) on 000183500570 Director FA-16B	CIM Alert Event	1	Not Cleared	🗑️
37581	Warning	2010-10-06 11:29	000190300595 (Symm7.DMX3-6)	2010-10-06 11:24: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37580	Warning	2010-10-06 11:29	000190300595 (Symm7.DMX3-6)	2010-10-06 11:22: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37542	Warning	2010-10-06 11:26	000183500570 (Symm48:3830)	2010-10-06 11:58: 05.0469.00 (SEVT_FC_OPTICAL_MOD_ERROR) on 000183500570 Director FA-16B	CIM Alert Event	1	Not Cleared	🗑️
37541	Warning	2010-10-06 11:26	000190300595 (Symm7.DMX3-6)	2010-10-06 11:20: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️
37502	Warning	2010-10-06 11:23	000183500570 (Symm48:3830)	2010-10-06 11:53: 00.0033.00 (SEVT_ACCESS_TO_NR_DEVICE) on 000183500570 Director FA-16B	CIM Alert Event	1	Not Cleared	🗑️
37501	Warning	2010-10-06 11:23	000190300595 (Symm7.DMX3-6)	2010-10-06 11:18: 0x003f (SEVT_ACCESS_TO_NR_DEVICE) on 000190300595 Director FA-1D	CIM Alert Event	1	Not Cleared	🗑️

The event details summary page provides information about one or more of the following:

- **ID** – Number assigned to the event.
- **Status** – Indicates whether the event has been cleared.
- **Component** – If the event came from a component of an element, the component is listed.
- **Element** – The source of the event. An element can be a switch, host, application, fabric, or anything else on the network. If this box is blank, the event did not come from an element.
- **Severity** – Severity level. For descriptions, see [Severity Levels on page 68](#).
- **Time Reported** – Time and date the event was reported to the management server.
- **Type** – Brief label of the event.

- **Summary Text** – Explanation of the event.
- **Probable Cause** – Explanation of a probable cause.
- **Recommended Actions** – Provides recommendations.
- **Journal Entries** – Use this box to enter additional information and then click **Add Journal Entry**. This box is limited to 4,000 characters.

Events listed in Event Manager might not be attributed to the correct source until Get Details has completed.

For more information about the Event Manager and how to manage events, see the *HP Storage Essentials User Guide*.

5 Viewing Performance Data

Performance Manager provides a complete end-to-end SAN topology view which you can query to view the performance history of managed elements. The HP Storage Essentials SRM Performance Pack and Performance Pack Enterprise software licenses provide a more detailed performance view of your environment. When licensed, these optional performance packs provide additional Performance Manager data collection capabilities. Whether Performance Manager is able to display the full range of metrics described in this guide is dependent upon the number and types of licenses configured for your environment.

This section describes the Performance Manager graphical user interface toolbar and menus, and the different ways that you can view performance data. It includes the following topics:



- [Toolbars and Menus](#) below
- [Creating Performance Charts](#) on page 83
- [Saving Chart Settings](#) on page 83
- [Opening Saved Charts](#) on page 84
- [Deleting Chart Groups](#) on page 84
- [Viewing Summary Charts](#) on page 84
- [Viewing Trending Information for Performance](#) on page 86
- [Interpreting Charts and Graphs](#) on page 87
- [Comparing the Performance of Different Elements](#) on page 88





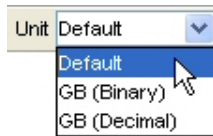

Toolbars and Menus

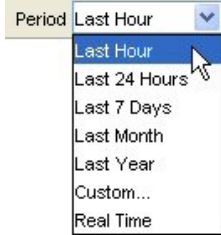




Performance Manager is an integrated component of HP Storage Essentials and shares a primary toolbar, located at the top of the upper pane, with the System Manager interface. The Performance toolbar is located in the lower pane which displays when you query a storage element for performance information.



The following table describes the tasks and actions available from the Performance toolbar:


Toolbar in Lower Pane of Performance Manager

Icon	Description
	Saves the chart in three formats: HTML, XLS, CSV. Use the drop-down menu to select the format you want.
	Opens saved charts, saves chart settings, and deletes chart groups. Use the drop-down menu to select the task you want. For additional details, see Saving Chart Settings on page 83.

Icon	Description
	Enables you to print a chart.
	Clears the chart of the elements you have selected.
	<p>Enables you to determine the type of chart displayed. Select one of the following options, and then click the  button:</p> <ul style="list-style-type: none"> • Bar – Displays each data point as a bar. The data for the different elements is displayed side by side. • Line – Displays each data point as a dot with a line connected to the previous data points. The data for the different elements for a specific point in time is displayed in the same column. • Summary – Displays a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines. See Viewing Summary Charts on page 84.
	<p>Enables you to determine the unit of measurement in the chart. Select one of the following options, and then click the  button:</p> <ul style="list-style-type: none"> • Default – Displays the data in its default unit, which is usually megabytes. • GB (Binary) – Displays the data in binary gigabytes. (The computer handles the data in binary format in which a gigabyte is equal to 1,073,741,824 bytes.) • GB (Decimal) – Displays the data in gigabytes. (A gigabyte is equal to 1,000,000,000 bytes.)


Icon	Description
	<p>Enables you to format the chart to provide data collected during the time period specified. Select the option to the left of the Period combo box. Select one of the following options, and then click the  button:</p> <ul style="list-style-type: none"> • Last Hour – Information collected during the last hour is reported. If you select Last Hour, the only frequency available is All. • Last 24 Hours – Information collected during the last 24 hours is reported. • Last 7 Days – Information collected during the last seven days is reported. • Last Month– Information collected during the last 30 days is reported. “Last Month” does not mean the previous month on the calendar. • Last Year – Information collected during the last 365 days is reported. “Last Year” does not mean the previous calendar year. • Custom... – Lets you create custom settings for the graph. • Real Time – Displays the data as it is currently being gathered by the management server. If you select Real Time, the only option for frequency is All and unit is Default. <p>To set a custom period, select Custom. See Setting a Custom Time Period on page 85 for more information.</p>
	<p>Zoom-out button allows you to return to the original data set.</p>
	<p>Applies the current period, frequency and trending information to the chart.</p>
	<p>This icon appears in the tool bar of the chart. It enables you to filter out the additional data series if the chart contains more than one series of data.</p>

Icon	Description
	<p>Enables you to modify the performance data displayed in the chart and change chart settings. When you select it, the following options are displayed:</p> <p>Performance data in the graph</p> <ul style="list-style-type: none"> • Data and Delete note pad – Displays elements on the chart (including statistics for the element) and allows you to delete some or all of them. <p>Specify graph settings</p> <p>Chart:</p> <ul style="list-style-type: none"> • Type – Allows you to select from the line, bar, or summary chart • Unit – Scales the y-axis by Gigabyte. You can select either decimal or binary value. • Line size – Affects the size of the line for line charts. • Frequency – Allows you to chart information at specified intervals. • Threshold Visible – Mandatory field to set a threshold or graph an existing threshold. • Threshold Value – Actual value of the threshold identified by a red line displayed on line and bar charts. • Show Grid – Check this field to display the Chart Grid. <p>Period:</p> <ul style="list-style-type: none"> • Period bullet – Lets you format the graph to provide data within the time period specified. • Custom – Allows you to customize a time period value by selecting start and end times from a calendar pop-up. • Trend – Specifies the number of trending data points you want to see. Trending information only applies to line and bar charts with a single statistic. The time period of the trend is determined by the frequency setting.
	<p>Creates a chart window.</p>

Icon	Description
	<p>Enables you to arrange the chart windows as follows:</p> <ul style="list-style-type: none"> • Tiled – Displays the windows within their own tile. • Cascade – Displays the windows on top of each other, with the active window on top. • Minimize All – Minimizes all windows. • Restore All – Opens all windows that have been minimized.


Creating Performance Charts

To create a performance chart for an element:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Select the element you want to monitor.
3. Under the **Monitoring** tab in the lower-left pane, select the element again. In some instances, you might need to select an element component, such as a port on a switch.
4. In the pane under the tree, select a monitoring option.
5. Use the Chart and Unit combo box to modify the chart. When you are done with your selections, click the  button in the lower pane.
6. To monitor more than one element in a chart, see [Comparing the Performance of Different Elements on page 88](#)

Saving Chart Settings


To save chart settings:

1. Create performance charts with the settings you want to save. All of the charts currently open will be saved in the chart settings group.
2. Click the  button in the lower pane, and select **Save Chart Settings** from the drop-down menu. The Save Chart Settings window opens.
3. Enter the name of the new chart group in the **Chart Group** box, or select an existing chart group to overwrite that group.
4. Click **Save**. The following settings are saved:
 - Elements
 - Observations
 - Chart type

- Unit
- Grid
- Period (custom)
- Line size
- Frequency
- Threshold
- Trend

Opening Saved Charts

To open saved charts:

1. Click the  button in the lower pane, and select **Open Charts** from the drop-down menu.
2. Select a chart group, and click **Open**. All chart observations and settings are restored, but with current data.


Custom period ranges are restored with the same period that was saved. Predefined periods (such as Last 24 hours) display data based on when the template is applied, not when it was saved.

If any charts have errors (such as a missing element), an error message is displayed.

3. Click **OK**, and the remaining charts in the saved chart group appear.

Deleting Chart Groups

To delete a chart group:

1. Click the  button in the lower pane, and select **Delete Charts** from the drop-down menu.
2. Select the chart group you want to delete, and click **Delete**. A confirmation dialog box appears.
3. Click **Yes**. The chart group is deleted.

Viewing Summary Charts

Performance Manager provides summary charts that display a single line that summarizes the values for a single statistic. Multiple statistics can be shown with multiple lines.

You cannot see more than one element at a time on a chart that has multiple elements.


The line has the following attributes:

- The vertical length of the line indicates the minimum and maximum value of the statistic within the selected data time frame.

- The green marker marks the median.
- The blue marker marks the average value.

Summary charts are not supported for real time performance display.

To view a summary chart:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Create a chart as described in [Creating Performance Charts on page 83](#).
3. Select **Summary** from the Chart combo box in the lower pane.
4. Click the  button in the lower pane for the change to take effect.

Setting a Custom Time Period

You can format the graph to provide data within a custom time period. This feature can be extremely useful for pinpointing performance changes. For example, assume you changed the firmware of a switch two weeks ago, and you want to compare the performance of that switch before and after you changed its firmware. You could create a graph that provides performance data two weeks before you changed the firmware. You could then create another graph that provides performance data two weeks after you changed the firmware.


To set a custom period:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Click the element whose performance you want to see.
3. Under the Monitoring tab in the lower-left pane, select the element again. In some instances, you might need to select an element's component, such as a port on a switch.
4. In the lower-left pane on the Monitoring tab, click a performance monitoring option.

The performance monitoring options vary according to the type of element.


5. Select **Custom** from the Period combo box in the lower pane.


The Performance Graph Editing Dialog window appears when you select the Custom option and there are no previous custom settings.


6. In the Performance Graphic Editing Dialog window, select the **Custom** option near the lower-left corner.
7. Click the calendar  button to the right of the Start box.
8. Enter the time in the time box. Make sure the time resembles a 24-hour clock, for example 22:00 for 10 p.m.
9. Click the date.


The date is highlighted in pink.

You can navigate the calendar as follows:

 - Displays the same month in the previous year

 - Displays the previous month

 - Displays the next month

 - Displays the same month in the following year

10. When you are done, click **Set**.

The start time and date are displayed in the Start box.

11. To set the end date, click the calendar  button to the right of the End box. Repeat steps 8 through 10.

The ending time and date are displayed in the End box.

12. Click **OK**.


13. Click the  button.

Viewing Trending Information for Performance

You can configure Performance Manager to display trending information in the charts you generate in the Monitoring tab. Trending charts are useful as indicators of an element's future performance based on the element's past performance data. When you request trending information, make sure that there are at least two sets of data collected within the time period you specify for the charting. This is a required condition for generating trending charts and reports.

A storage element's performance can drastically change in the future. Generated data trends are based on certain assumptions related to historical performance statistics, and as such, should not be considered fact. Trending information should be viewed only as an indicator of potential future performance states for a storage element.

To view trending information:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Click the element whose performance you want to see.
3. Under the Monitoring tab in the lower-left pane, select the element again. You might need to expand the path to locate an element. For example, for switches, you might need to expand down to select the element's port.
4. In the lower-left pane on the Monitoring tab, click the performance metric for which you want trending statistics. The performance metrics shown vary according to the type of element you select.
5. Click the **Edit Chart Property** () button in the lower pane.
6. In the Performance Graph Editing Dialog window, enter a number in the **Trend** box.

The number corresponds to the number of frequency intervals for which the trending information will be provided. For example, if you enter 5 in the Trend box, the chart provides trending information for five frequency intervals, such as five weeks (if weeks was selected from the Frequency menu). If the frequency is set to **All**, the trend interval is an hour.

7. Click **OK**.




The trending information is displayed in the chart.

If there is not enough data to display, Performance Manager does not display the chart. For example, if you selected the weekly option from Frequency menu and you only have two days of data, a chart is not displayed, regardless of the value in the Trend box. Performance Manager does not display a chart if there is not enough data, and the trending number is ignored.

See also [Creating Performance Reports with Drill Down and Trending](#) on page 257.

Removing Performance Data from a Graph

To remove multiple data from a performance graph:

1. Access Performance Manager as described in [Opening the Performance Manager](#) on page 21.
2. In Performance Manager, click the graph you want to modify in the bottom-right pane.
3. Click the **Edit Chart Property** () button in the lower pane.
4. In the Performance Graph Editing Dialog window, click the **Delete** () button corresponding with the data you want to remove from the graph.
5. When you are asked if you want to remove the data, click **Yes**.
6. Click **OK**.
7. Click the  button in the lower pane.

Interpreting Charts and Graphs

The following topics interpret certain aspects of performance that might appear in the charts generated for your system:

- [Sudden Dips in Charts in Performance Manager](#) below
- [Values Increase in Charts for Aggregated Drives and Aggregate Volumes](#) on next page

Sudden Dips in Charts in Performance Manager

In Performance Manager on the Monitoring tab, charts that display data gathered by certain volume and drive counters display their charts with the results increasing to the maximum value and then decreasing rapidly to a very low number that starts rising again.

The sudden dip in the charts occur because the counters holding the values displayed in the charts revert back to zero after reaching the upper limit, and then they continue to go up again. The counters do not usually display zero in a chart, because they are quickly gathering data again.

The following charts for individual drives are impacted: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests.

The following charts for individual volumes are impacted: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs.

Values Increase in Charts for Aggregated Drives and Aggregate Volumes

Values in charts for aggregate drives and aggregate volumes continue to rise smoothly in Performance Manager and on the Monitoring tab. The only time the values in the charts come back down to the values returned by the array is when the service for the management server restarts, and then the charts display the values returned by the array. The values in the charts continue to rise until the service for the management server is stopped.

The values in the charts continue to rise, because the values do not revert back to zero when an individual drive or volume counter that makes up the aggregation contains the maximum number of integers. Instead, the values continue to rise smoothly.

The following charts for aggregate drives are affected: ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, Unrecovered Errors, Recovered Errors, Timeouts, Retried Requests. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

The following charts for aggregate volumes are affected: Bytes Read, Bytes Read Large, Bytes Written, Bytes Written Large, ReadIOs Large, WriteIOs Large, ReadIOs, WriteIOs, TotalIOs, Bytes Transferred, ReadHitIOs. Charts for aggregated drives are also provided for arrays, controllers, and volume groups.

Comparing the Performance of Different Elements

Use Performance Manager to compare the performance of different elements.

The following example describes how to compare the physical memory used on different hosts:

1. Access Performance Manager as described in [Opening the Performance Manager on page 21](#).
2. Click the element in the topology whose performance you want to see.
This example assumes that you clicked a host named Cortez.
3. Select a component of the element. This example assumes that you selected Cortez under the **Application Path** node.
4. Select the graph monitoring option you want to see on the Monitoring tab, located in the lower pane. For example, you might select **Physical Memory Used (%)**.
5. Scroll to the bottom of the storage tree in the lower-left pane.
6. Expand the **All Elements** node by clicking the (+) symbol next to the node name.

7. Expand the node for the host you want to compare. For example, to compare a host name HostA against Cortez, you would expand the following nodes in order: **All Elements > Hosts > HostA > Element Paths > Application Path.**
8. Select the same type of component you selected previously for Cortez.
9. Select the same type of graph monitoring option in the left pane.

For example, if you selected “Physical Memory Used (%)” for the first element, you must select the same option for the other elements. You cannot select “Physical Memory Used (%)” for the first element and then “Processor Utilization (%)” for the second element because they measure different types of data.

Performance Manager displays information for the different elements in the same graph.

10. Repeat steps 7 through 9 for each host you want to compare against Cortez.

6 Managing Performance in Your Storage Environment

HP Storage Essentials Performance Manager provides visibility into complex heterogeneous storage environments and enables you to discover, trend, and manage the performance of your entire SAN, from individual devices to large storage arrays. The following topics drill down to the performance management details, techniques and best practices for optimizing your setup of Performance Manager and the hundreds of performance data metrics available for routine collection and reporting. It contains the following sections:

- [Managing Application Performance on page 93](#)
- [Managing Host Server Performance on page 108](#)
- [Managing Switch Performance on page 113](#)
- [Managing EMC Symmetrix Performance on page 117](#)
- [Managing HP StorageWorks EVA Performance for Arrays on page 138](#)
- [Managing XP and HDS Array Performance on page 153](#)
- [Managing NetApp Performance on page 168](#)

Managing Storage System Performance

To ensure availability and robust performance of business applications, storage administrators must routinely monitor and manage the arrays and storage area network (SAN) (based on the array) that support those applications. This begins with proactive planning and discovery designed to anticipate problems and enable quick reaction to application or business service slow down. Performance Manager supports this activity by giving an administrator visibility at several points along the data path: applications, host servers, switches, arrays, and array elements such as volumes, disks, controllers, and storage subsystem Performance Manager metrics for each of these path points enable you to collect the statistics, perform trending, and issue reports that support the complex decision-making process that gets the performance results your company needs.

Performance Management and Metrics

HP recommends the following performance management methodology as an approach to utilizing Performance Manager metrics to their greatest advantage when you find a performance issue. This section describes this approach using these four easy steps:

Step 1: Identify the Physical Path

Identify the physical path along which the performance issue exists. HP Storage Essentials discovers the storage environment and helps visualize the topology and paths in the storage environment. It depicts the path from application—to server—to virtual elements—to volumes—to HBA—to switches—to array. This automatic filtering helps you identify quickly the systems along a path where the performance problem exists.

Step 2: Isolate the Bottleneck

Isolate the elements on the path that cause the performance bottleneck. With the path along which the problem exists displayed, you can drill down to view the details of every port and infrastructure element involved in the data path and isolate the element that is the performance bottleneck. The isolated element can now undergo intense troubleshooting based on deep inquiry using performance metrics.

Step 3: Troubleshoot the Cause

Troubleshoot the element to understand the root cause. This guide details the important metrics to look at for each isolated element in the path. For example, if you suspect that a host server's workload is too high, and the CPU is underpowered for the load, refer to the Host Metrics section to determine which metrics will give you the information you need to confirm this.

Use metrics also to proactively measure and present indisputable data related to storage performance that gives you the ability to establish baselines and compare over time such statistics as percentage of usage, I/O throughput, available storage space, and workload trends.

Step 4: Report on Performance Trends and Statistics

The Performance Manager GUI provides standard reporting capabilities (including trending) that allow you to view application, host, switch and array metrics. Look at trends and generate reports as a pre-step to making informed decisions. Use history-based trends and dynamic reporting to collect performance data and build evidence which points to sustainable solutions. A set of out-of-box performance reports can be leveraged for decision-making purposes and future planning.

You can also use the HP Storage Essentials Report Optimizer for viewing host, switch, and array metrics. Report Optimizer user licenses can be purchased for ad-hoc reporting against metrics exposed in those cases where out-of-box reports are not sufficient. For a list of performance reports in the Report Optimizer Report Pack, see [List of Standard Performance Reports on page 189](#).

Step 5: Configure Proactive Threshold Notification using Policy Manager

After you establish a baseline, leverage the HP Storage Essentials Policy Manager by setting up performance policies which provide the early warnings needed to address future performance issues before they happen. The Policy Manager provides policy templates for many of the collected performance metrics. Use performance policies to define thresholds and resulting actions (such as generating events, emails or running scripts) when any given threshold is exceeded. For information about using the HP Storage Essentials Policy Manager and setting up performance policies, see [Setting Performance Policies and Alerts on page 67](#).

Managing Application Performance

This section describes the HP Storage Essentials performance management metrics for the applications running in your storage environment. These performance metrics are gathered by data collectors and made available for reporting. It also includes techniques and best practices for managing performance on your applications. For more information see the following topics:

- [Troubleshooting Application Performance below](#)
- [Improving Microsoft Exchange Server Performance on next page](#)
- [Monitoring Oracle Server Performance on page 96](#)
- [Monitoring Microsoft SQL Server Performance on page 97](#)
- [Performance Management Application Metrics on page 97](#)

Troubleshooting Application Performance

Follow these steps to monitor and/or troubleshoot host application performance using the HP Storage Essentials Performance Manager. The instructions assume that you launched the Performance Manager and have a displayed view of the storage path topology for your system. To begin, navigate to the host server that runs your application.

Step	Description	Considerations
1	<p>Check the host processor utilization and physical/virtual memory consumption by selecting to view these metrics:</p> <ul style="list-style-type: none"> • Free Physical Memory (KBytes) • Free Virtual Memory (KBytes) • Processor Utilization (%) 	<p>If your processor is running near 100 percent and virtual memory consumption is greater than physical memory, the problem might be that the host does not have sufficient resources to process the I/O for the application.</p>

Step	Description	Considerations
2	Select the application path where the application is running, and expand its members.	
3	Check the logical device (LDEV) read/write response time. Note if the LDEV has an extremely high response time (above 20 ms) or zero response.	If the response time is higher than expected, this can indicate a resource contention problem within the array. If there is zero response, check for a link failure.
4	Chart the adapter, switch, and array port Bytes Received and Bytes Transmitted metrics on the same graph to verify that traffic is passing through to the array without problems.	If you see large differences between the host bus adapter (HBA), switch and array port, the problem could be a bad SFP or fiber link.
5	Chart the array port and LDEV Total Data Rate.	The array port should have at least the same data rate as the LDEV. If the array port data rate is lower than LDEV that could indicate you have too much write pending in cache.
6	Select the array and chart its cache write pending percentage.	The array will begin to hold back I/O if the write pending is at or above 70 percent. Some of the causes for excessive write pending could be too much I/O from other devices in the Cache Logical Partition (CLPR), replication processing, and external storage (unless you have cache write-through mode enabled).

Improving Microsoft Exchange Server Performance

There are multiple methods for improving Microsoft Exchange Server performance. Techniques range from placing databases and transaction logs in an optimal location to using the performance metrics in this document to locate and remove Microsoft Windows bottlenecks.

Refer to the Microsoft TechNet library (<http://technet.microsoft.com/en-us/library/default.aspx>) for information about tuning your Microsoft Exchange Server. Use the information in the following sections of this guide to learn more about using HP Storage Essentials performance metrics to identify areas for improvement.

Performance Baseline

As a part of your initial data collection, create a baseline for your Microsoft Exchange servers. This baseline assumes that there are no serious problems with the servers when you take the readings. The baseline values tell you what is normal for your Microsoft Exchange servers and provide reliable values against which you can compare future measurements.

Baseline values are not valid forever however. Any time you perform a major action on your Microsoft Exchange server—such as installing a service pack or adding new hardware—the next reading should be compared to the original baseline value to determine the impact of the change. After this comparison is made, the most recent reading then becomes the new baseline.

Even if you do not make any major changes to the Microsoft Exchange server, HP recommends that you update your baseline values every four to six months by making your most recent reading the new baseline. In this way, you always have a semi-current (and realistic) baseline from which to work.

Activity Spikes

Spikes in activity are normal, even for a well-functioning server. If you are only collecting data summaries every 15 minutes, it is possible that some samples will occur during big spikes of activity, and make the Microsoft Exchange server look busier than it actually is.

To get a more accurate reading of the data, take a data collection 15 minutes and then again before 15 minutes after an activity spike. Look at those performance metrics which give you average measurements for the data over that specific time period. Not that this does not show you the spikes/transients, because these will have been averaged into the final measurement. If you want to collect more detailed statistics at a finer level of granularity, collecting data down to 5-minute intervals will better reflect actual performance during a given period. You can also collect data in real time for a typical 20-second collection by working interactively with the HP Storage Essentials interface.

When you look at the collected data for a week and if the averages seem high, it is possible that the data was collected during periods of high activity. Common advice is not to worry about it. The purpose of data collection is to look for long-term trends so that you can forecast your future hardware needs.

HP Storage Essentials provides roll up of data from raw (sub-hourly) data, to hourly, daily, weekly, and so forth up to yearly aggregation of data. Because the raw (sub hourly) is only preserved for one day before being rolled up and summarized, any spikes (where they were apparent during smaller interval sampling) would become less apparent and averaged into the larger pool of data with each roll-up. If weekly views of your server performance is a critical piece of your reporting, you can set preservation of RAW to two weeks, and set the collection interval to, for example, 5 minutes, and get a detailed view of the previous week for the elements selected.

HP recommends that you do not set RAW preservation longer than the defaults listed in the GUI (1 day). If you need to adjust the default intervals, do so only with the assistance of HP Support.

Monitoring Oracle Server Performance

You can use Oracle application performance metrics in HP Storage Essentials in association with other storage-related performance metrics to identify potential bottlenecks. The metrics available in HP Storage Essentials have been carefully chosen to provide this capability and are not a replacement for native tools.

Consider the following recommended steps when tuning Oracle applications for improved performance. To prevent unanticipated tuning side effects, please maintain the order in which steps are listed in this section. For example, it is not a good practice to increase the buffer cache if you can reduce I/O by rewriting a SQL statement.

Step 1. Review the Database Design

A well-designed database schema is the first step to solid application performance. If you are seeing poor system performance, this is usually the result of poor database design. If it is still possible, you might consider database design improvements that include the following:

- You should generally normalize to the 3NF. Selective de-normalization can provide valuable performance improvements.
- Always keep the "data access path" in mind when you are designing.
- Look at proper data partitioning, data replication, and aggregation tables for decision support systems.

Step 2: Tune the Application

Experience shows that approximately 80% of all Oracle system performance problems are resolved by coding optimal SQL. You can also enhance performance by scheduling batch tasks and other large workloads after peak working hours.

Step 3: Tune the Memory

Make sure to properly size your database buffers (for example, `shared_pool`, buffer cache, log buffer) by looking at performance metrics related to wait events, buffer hit ratios, system swapping, paging, and so forth. You might also want to pin large objects into memory to prevent frequent reloads.

Step 4: Tune the Disk I/O

To provide maximum disk subsystem throughput, properly size and place database files. Also review performance metrics for frequent disk sorts, full table scans, missing indexes, row chaining, data fragmentation, and so forth.

Step 5: Eliminate Database Contention

Application contention for database resources can slow application response time significantly. Carefully review database locks, latches, and wait events and eliminate them wherever possible.

Step 6: Tune the Operating System

Monitor and tune the operating system, CPU, and I/O and memory utilization. For information about best practices for doing this, read the related Oracle FAQ related to your specific operating system at http://www.oraFAQ.com/wiki/Operating_system.

Monitoring Microsoft SQL Server Performance

Microsoft provides information about tuning Microsoft SQL Servers for improved performance in the Microsoft TechNet library (<http://technet.microsoft.com/en-us/library/default.aspx>). Use this guide to understand how to use Storage Essential performance metrics to help you identify bottlenecks and spot potential problem areas before they affect your system.

Using Metrics to Identify Bottlenecks

Routine monitoring of your Microsoft SQL Server applications for performance bottlenecks is important. One bottleneck can manifest itself as other bottlenecks, and after you remove one, you might find the SQL Server is constrained by a completely different bottleneck. For example, on an SQL Server with disk constraints (a low RPM drive) and a small amount of available memory (causing excessive paging), you will find high processor utilization because the SQL Server processor is spending a large amount of time paging to disk. You can consider adding more memory to reduce the amount of paging, but then the disk bottleneck might also become the primary constraint.

Disk bottlenecks tend to manifest themselves in high processor utilization situations because the processor is busy managing disk I/O. Memory bottlenecks manifest themselves as high processor and disk utilization. This results because the processor is busy managing virtual memory and disk I/O. As a best practice, when examining bottlenecks, begin by examining memory bottlenecks, then disk bottlenecks and finally processor bottlenecks.

Performance Management Application Metrics

HP Storage Essentials provides metrics for the following applications:

- [Microsoft Exchange Server Performance Metrics on next page](#)
- [Oracle Server Performance Metrics on page 100](#)
- [Microsoft SQL Server Performance Metrics on page 102](#)
- [Sybase Performance Metrics on page 104](#)
- [IBM DB2 Performance Metrics on page 105](#)
- [Informix Performance Metrics on page 104](#)
- [InterSystems Caché Database Performance Metrics on page 107](#)

Microsoft Exchange Server Performance Metrics

This table lists the metrics used by HP Storage Essentials to measure performance for the Microsoft Exchange Server. An asterisk (*) marks a metric that does not collect historical data. All other metrics use collected historical data to provide statistics.

Metric	Description	Units	Use to...
Application Server			
Exchange Services	Services Running/ Not Running (Site Replication, Routing Engine, POP3, Exchange Management, MTA Stacks, Info Store, IMAP4, Events)	Count	Show whether or not a critical Microsoft Exchange service is running.
Final Destination currently unreachable Queue Size*	Contains the number of messages that cannot reach their intended destination.	Count	Set alerts. Possible reasons why final destination is not reachable are: <ul style="list-style-type: none"> • No route • Microsoft Exchange connector unreachable • No remote delivery queue available
Messages awaiting directory lookup Queue Size*	Messages awaiting directory lookup queue size. These are messages sent to a Microsoft Exchange 2000 or 2003 server.	Count	Set alerts or troubleshoot Event ID 9035, 6004, 9003, 9004. Possible causes: Message Archiving, Insufficient Permissions, or global catalog server issues.
Messages to be routed Queue Size*	Number of messages to be routed	Count	Set alerts or troubleshoot message routing issues in Microsoft Exchange.
Pre-submission Queue Size*	Pre-submission queue size	Count	Measure the number of messages waiting in the pre-submission queue. Possible causes are: <ul style="list-style-type: none"> • Third-party applications • SMTP events • DNS

Metric	Description	Units	Use to...
SMTP Local Delivery Queue Size*	Number of messages in the local queue	Count	Set alerts. Large numbers of messages indicate possible disk I/O issues or dismounted stores.
SMTP Server Queue Summary*	Summary of local and remote queues	Count	Set alerts and notifications. Helps pinpoint disk, DNS, or network I/O issues.
Storage Groups			
Log Size (MBytes)	Storage group log file size	MB	Help manage log file sizes in Microsoft Exchange.
Storage Group Size	Size of the Storage Group in MB. Contains one or more mailbox or public folders, transaction logs, and system files depending on the version.	MB	Provide a facility to track the Microsoft Exchange group utilization from a higher level than only mailbox or folders sizes.
Message Stores and Public Folders			
*Active Client Logins	Number of logons active (issued by any MAPI requests) within a specific time interval	Count	Track the number of sessions/users accessing the Microsoft Exchange server at a given time.
*Average Delivery Time	Rate at which messages are delivered to all recipients	Msg/s	Measure long delivery times that could be caused by deferred delivery queues or stalled messages.
Mail Related Objects Count	Public folder mail related objects count.	Count	Count the number of mail messages, appointments, meeting requests, tasks, task requests, contacts, remote mail items, and notes.
*Receive Queue Size	Number of messages in the mailbox store's receive queue	Count	Track MTA queues, IS queues, and SMTP queues. Possible issues are: <ul style="list-style-type: none"> • Related service • Looping or corrupt message • Source server issue

Metric	Description	Units	Use to...
*Send Queue Size1	Number of deferred delivery or submission messages.	Count	Track MTA queues, IS queues, and SMTP queues. Possible issues are: <ul style="list-style-type: none"> • Related service • Looping or corrupt message • Source server issue
Store Size	Mailbox store size	MB	Set alerts. Helps alert administrators when a store size is reaching its limits.

*Not supported with Microsoft 2007.

Oracle Server Performance Metrics

The following table lists the HP Storage Essentials Performance Manager metrics available for monitoring Oracle performance and describes what they count as well as considerations for using them.

Metric	Description	Units	Use to...
Buffer Hit Ratio	Percentage (%) of requests against data block buffer	%	Measure the effectiveness of the Oracle data block buffer. Used with a database that has an undersized DB Cache size where the “working set” of frequently-referenced data has not been cached.
Dictionary Hit Ratio	Ratio of logical reads to physical disk reads	%	Monitor ratio balance between logical and physical disk reads. As the hit ratio approaches 100 percent, more data blocks are found in memory, resulting in fewer disk I/Os and faster overall database performance.
File Read Percent	Percentage (%) of sequential reads.	%	Measure the performance and efficiency of indexes.
File Total I/O Percent	Percentage (%) of sequential reads and writes	%	Capture overall sequential database I/Os.
File Write Percent	Percentage (%) of sequential writes	%	Identify other bottlenecks like redo waits, bind variables, bulk operations, and index contention.

Metric	Description	Units	Use to...
In Memory Sort Ratio	Percentage (%) of sorts (from ORDER by clauses or index building) that are done to disk compared to in-memory	%	Differentiate and identify the number of disk sorts performed in TEMP tablespace versus those performed in-memory (RAM sorts).
Library Cache Hit Ratio (%)	Monitors the percentage (%) of entries in the library cache that were parsed more than once (reloads) over the lifetime of the instance.	%	Set the shared_pool_size large enough to prevent excessive re-parsing of SQL.
Parse CPU to Total CPU Ratio	Percentage (%) of reentrant SQL statements compared to re-parsed SQL statements; also referred to as the Parse to Execute ratio	%	Show the number of unique incoming SQL statements or that SQL statements are NOT reentrant.
Redo Buffer Allocation Retries	Total number of retries needed to allocate space in the redo buffer	%	Monitor the buffer allocation retries over a period of time while the application is running. If the Redo Buffer Allocation Retries value is continuously increasing, then increase the LOG_BUFFER value.
System Event Time Waited	Provides wait state details for ongoing Oracle transactions.	ms	Analyze time-based events, and system-wide and session wait events.
Tablespace Read Percent	Percentage (%) of tablespace reads over an interval	%	Monitor tablespace usage and growth.
Tablespace Total I/O Percent	Percentage (%) of total tablespace I/O over an interval	%	Monitor tablespace usage and growth.
Tablespace Write Percent	Percentage (%) of tablespace writes over an interval	%	Monitor tablespace usage and growth.

Microsoft SQL Server Performance Metrics

This table describes the performance metrics available for Microsoft SQL Server 2005 and 2008. For a list of supported software versions for Microsoft SQL Server, see the HP Storage Essentials support matrix document located in the Documentation Center (**Help > Documentation Center**).

Metric	Description	Units	Use to...
Blocked Processes	Number of blocked processes in your Microsoft SQL Server database	Count	Flag the database administrator (DBA) that a number of blocked processes are present. The DBA can then use a SELECT statement to identify the offending processes.
Buffer Cache Hit Rate	Percentage (%) of pages found in the buffer cache. The % is calculated as the total number of successful cache hits divided by the total number of requested cache lookups.	%	Monitor the buffer cache hit ratio. After a system maintains a steady state of operation, this metric should achieve rates of 90 percent or higher. The buffer cache hit ratio can be increased by increasing the amount of memory available to SQL Server.
CPU Usage Percentage	Percentage (%) of the total available CPU time that Microsoft SQL Server uses during the current interval.	%	Identify potential performance issues. A high CPU usage rate can indicate performance bottlenecks. Options for resolving the bottleneck are to add multiple CPUs, allocate resources more efficiently, identify resource intensive applications, and reduce workloads, or upgrade CPUs.
Cache Memory	Total amount of dynamic memory the server is using for the dynamic SQL cache.	KB	Determine if the current cache memory is meeting the needs of the SQL Server application.
Dead Locks	Number of lock requests per second that resulted in a deadlock. A deadlock is a situation where two or more competing actions are waiting for the other to finish.	Req/s	Generate an alert for a deadlock condition. You can then identify the SPIDs and the resources that are involved in a deadlock using the database logs.

Metric	Description	Units	Use to...
Lock Requests	Number of new locks and lock conversions per second requested from the lock manager.	Req/s	Identify issues related to data retrieval. A high number of lock requests with low request rate is an indicator that SQL Server must do table scans when retrieving data.
Lock Time Outs	Number of lock requests per second that timed out, including internal requests for NOWAIT locks.	Req/s	Identify areas of lock contention and congestion in the database.
Lock Waits	Lock wait time is the time that a process spends waiting for another process to release a lock.	Wait/s	Generate an alert to the database administrator when the number exceeds pre-set thresholds; also indicates database congestion.
Memory Usage Percentage	Number of pages in the procedure cache that are currently allocated to a process.	%	Monitor allocation of memory to processes. A negative number indicates that a process is freeing memory allocated by another process.
Physical I/O Percentage	Percentage of elapsed time that the disk drive was busy servicing read or write requests	%	Identify disk I/O problems in the SQL server instance. A value greater than 50 percent may indicate an I/O bottleneck.
Plan Cache Hits Ratio	Ratio between plan cache hits and lookups. The plan contains stored procedures, ad hoc and prepared Transact-SQL statements, and triggers.	%	Monitor the database plan cache performance.
Target Server Memory	Amount of dynamic memory the server can consume	KB	Manage database server memory.
Total Server Memory	Committed memory from the buffer pool. This is NOT the total memory used by the SQL Server.	KB	Manage database server memory.
Transactions	Total number of database transactions	Count	Monitor and manage load balancing.

Metric	Description	Units	Use to...
User Connections	Total number of database user connections	Count	Identify the number of users who are accessing the database.

Sybase Performance Metrics

This table lists the performance metrics available for Sybase Adaptive Server Enterprise systems.

Metric	Description	Units	Use to...
CPU Usage Percentage	Percentage of the total available CPU time that Sybase instance uses during the current interval.	%	Monitor the CPU usage of the Sybase instance. A high CPU usage rate can indicate performance bottlenecks. If you identify an issue, some options are to add multiple CPUs, allocate resources more efficiently, identify resource intensive applications, reduce workload, or upgrade CPUs.
Memory Usage Percentage(%)	Number of pages in the procedure cache that are currently allocated to a process	%	Monitor memory allocation to active processes. A negative number indicates that the process is freeing memory allocated by another process.
Physical I/O Percentage	Percentage of elapsed time that the disk drive was busy servicing read or write requests.	%	Identify disk I/O problems in the Sybase instance. A value greater than 50 percent may indicate an I/O bottleneck.

Informix Performance Metrics

The Read Cached (%) metric is an important database performance indicator. As the hit ratio approaches 100%, this indicates that more data blocks are being found in memory, which in turn results in fewer disk I/Os and faster overall database performance. Use Chunk Read and Chunk Write to monitor chunk usage, and DBSpace Reads and DBSpace Writes to monitor I/O requests to the database.

The following table describes the performance metrics for Informix.

Metric	Description	Units
Buffer Read	Number of reads from the buffer cache by the database server	count
Buffer Write	Number of writes to the buffer cache by the database server.	count
Chunk Read	Total number of reads from the chunk. Use in combination with Chunk Write metric to monitor chunk usage.	count
Chunk Write	Total number of writes to the chunk. Use in combination with Chunk Read metric to monitor chunk usage.	count
DBSpace Reads	Total number of read calls that involve this dbspace. Use in combination with DBSpace Writes metric to monitor dbspace usage.	count
DBSpace Writes	Total number of write calls that involve this dbspace. Use in combination with DBSpace Reads metric to monitor dbspace usage.	count
Disk Reads	Total number of read operations from disk by the database server	count
Disk Writes	Total number of write operations to disk by the database server	count
Page Reads	Number of pages read from disk by the database server	count
Page Writes	Number of pages transferred to disk by the database server	count
Read Cached	Percentage of all read operations that are read from the buffer cache without requiring a disk read by the database server, calculated as follows: $100 \times ((\text{buffer_reads} - \text{disk_reads}) / (\text{buffer_reads}))$ As the hit ratio approaches 100%, more data blocks are found in memory. This results in fewer disk I/Os and faster overall database performance.	%
Write Cached	Percentage of all write operations that are buffer writes by the database server, calculated as follows: $100 \times ((\text{buffer_writes} - \text{disk_writes}) / (\text{buffer_writes}))$	%

IBM DB2 Performance Metrics

This table lists the performance metrics tracked for IBM DB2 databases.

Metric	Description	Units	Use to...
Instance Private Sort Memory	Private sort memory is calculated using the formula $(db2.sort_heap_allocated / sheapthres) \times 100$. In this instance, <code>sort_heap_allocated</code> is the system monitor element, and <code>sheapthres</code> is a DBM configuration parameter.	%	Track private sort memory utilization. Used to check that there is sufficient heap space to perform sorting and that sorts do not overflow.
Instance Monitor Heap Utilization	Utilization is calculated using the formula $(db2.pool_cur_size / db2.pool_max_size) \times 100$ for the Memory Pool Identifier <code>SQLM_HEAP_MONITOR</code> .	%	Track the consumption of the monitor heap memory, If this percentage reaches the maximum 100%, monitor operations might fail.
Database Catalog Cache Hit Ratio (%)	Hit ratio indicates (as a percentage %) how well the catalog cache is working to avoid actual accesses to the catalog on disk	%	Monitor the ratio between catalog cache and physical disk reads. As the catalog cache hit ratio approaches 100 percent, more catalog blocks are found in memory, resulting in fewer disk I/Os and faster overall database performance.
Database Lock List Utilization	There is one lock list per database, and it contains the locks held by all applications concurrently connected to the database. The indicator is calculated using the formula $(db.lock_list_in_use / (locklist \times 4096)) \times 100$.	%	Track the amount of lock list memory that is being used.
Database Deadlocks	Number of lock requests that resulted in a deadlock. A deadlock is a situation where two or more competing actions are waiting for the other to finish.	Lock requests	Track deadlock conditions. The DBA will identify the SPIDs and the resources that are involved in a deadlock using the database logs.

Metric	Description	Units	Use to...
Database Shared Sort Memory Utilization	Shared sort memory is calculated using the formula $(db.sort_shrheap_allocated / sheapthres_shr) \times 100$. In this instance, sheapthres_shr is a database configuration parameter.	%	Track shared sort memory utilization for the database. Can be used to determine an appropriate value for the shared sort memory threshold.

InterSystems Caché Database Performance Metrics

The following table lists the performance metrics in HP Storage Essentials for InterSystems Caché databases.

Metric	Description	Units
Application Errors	Number of application errors logged Not supported for Caché 5.0.	Count
Global Kills	Number of global kills since startup Only supported for Caché 5.0.	Count
Cache Efficiency	Most recently measured cache efficiency. This is measured as Global references / (physical reads + writes).	GlobalRef/IO
Global Refs	Number of Global references since system startup.	Count
Global Sets	Number of Global Sets and Kills since system startup	Count
Globals Per Second	Most recently measured number of Global references per second Not supported for Caché 5.0.	GlobalRef/s
Disk Reads	Number of physical block read operations since system startup	Count
Disk Writes	Number of physical block write operations since system startup	Count
Journal Entries	Number of entries written to the system journal	Count
Logical Requests	Number of logical block requests since system startup	Count

Metric	Description	Units
Routine Loads	Number of routine loads since system startup Only supported for Caché 5.0.	Count
Routine Refs	Number of routine loads and saves since system startup. Not supported for Caché 5.0.	Count
Routine Saves	Number of routine saves since system startup. Only supported for Caché 5.0.	Count

Managing Host Server Performance

This section describes the HP Storage Essentials performance management metrics for the hosts and servers in your storage environment that are collected by data collectors and available for inclusion in reports. For more information see the following topics:

- [Finding Host Performance Bottlenecks below](#)
- [Host Bus Adapter Performance Metrics on page 110](#)
- [Performance Management Host Server Metrics on page 111](#)

Finding Host Performance Bottlenecks

A performance bottleneck is identified by an impaired or stopped flow of data which is measured, for example, in transaction response times, number of I/O requests, I/O sequences of reads and writes, and so forth. In general, bottlenecks are characterized by poor throughput, high data or resource contention, and/or inefficient handling of heavy workload demands. Bottlenecks can occur within the internal resources of the server itself, at some identifiable point along the communications network to which the server is attached, or within the storage fabric utilized by the server. Finding bottlenecks and resolving the issues that cause them is essential to keeping your servers and systems healthy and continually available to users.

HP Storage Essentials does not automatically identify hotspots or bottlenecks. The HP Storage Essentials metrics, however, can help you find bottlenecks by providing the critical iterative statistics that allow you and your managers to see performance levels on your servers and systems over time. Understanding how to configure performance metrics and interpret the data that HP Storage Essentials provides is a first step. The metrics for each of the most common bottleneck areas are described below.

Restrictions

For the Performance Manager to collect and display performance metrics, at least two data points for the hosts are needed. If only one data point is available, no performance metrics are presented.

Not all metrics are supported on all operating systems or for all storage system elements. Refer to the *HP Storage Essentials User Guide* for information about supported operating systems and metrics.

Host metrics for SAN-attached volumes are not available in the Performance Manager. SAN volume metrics are not collected from hosts. For volumes from volume manager groups (logical volumes), metrics are presented for only the first volume on that disk, although it might appear as if the metrics for all of the volumes are presented.

Random Access Memory

Use the Free Physical Memory and Virtual Memory Used metrics to measure performance related to the server random access memory (RAM). The Free Physical Memory counter displays the amount of physical RAM available to the Microsoft Windows operating system. The Virtual Memory Used counter measures "hard" page faults which are the times when data had to be swapped between the hard drive's virtual memory and the physical memory on the host.

If the Free Physical Memory counter drops below 10 percent of the RAM, while the Pages/sec counter increases significantly, this might indicate that the server does not have enough RAM to support the software programs that are currently running on the machine.

Central Processing Unit

Use the Processor Utilization % Time metric to track usage levels of the central processing unit (CPU).

The Processor Utilization % Time counter tells you what percent of CPU capacity is being used at any specific point in time. It is common for software application launches and other events to create utilization spikes between 90 and 100 percent. However, if this counter consistently measures over 80 percent, it is possible that your processor lacks the ability to handle your system workload. Because high Processor Utilization % Time metrics for a server typically result in slow response times and weak performance of applications, you might want to consider upgrading your CPU. See [Host Server Performance Metrics on page 111](#).

Physical Disk

To understand the metrics for physical disk usage, also called direct-attached storage or local disk, by your server, use the Disk Utilization % Time metric.

The Disk Utilization % Time counter displays the amount of time the hard drive spends reading or writing data. A disk that runs 40 to 50 percent of the time might need to be replaced or upgraded. Another solution is to consider adding additional storage disks.

If you have an HP EVA or HP XP Performance Pack, see [Collecting EVA Performance Metrics on page 62](#) or [Collecting XP and HDS Array Performance Metrics on page 59](#). You can gather information from volumes at the array level, but not at the host level. See [Restrictions on previous page](#) and [Host Disk Performance Metrics on page 111](#).

Host Bus Adapter Performance Metrics

The host bus adapter (HBA) performance metrics displayed in Storage Essentials are collected from the fabric switch ports, then inverted and mapped against SAN facing host/array ports and stored along with the host/array object in Storage Essentials. This provides a uniform method for performance collection across all OS/HBA combinations in the environment. If metrics are not available for a given HBA in Performance Manager, you might consider checking the status of the fabric switches in context, as they are the source of these performance metrics.

Bytes Received and Bytes Transmitted Metric

The HBA Bytes Received and Bytes Transmitted metrics can help determine transmission loads for load balancing, multi-pathing optimization and identifying the need for increased bandwidth. The goal is to determine total IOPS over an interval. If total IOPS is approaching the limitations of the HBA boundaries over a sustained period of time, then it may be beneficial to provide an additional path or to upgrade the bandwidth.

Best Practices

Follow these recommended HBA performance best practices:

- Leave the FC ports at auto-negotiate.
- Avoid putting more HBAs on a server than the bus throughput can support.
- Turn on the I/O coalesce parameter in high-performance environments.
- Install the latest HBA firmware and driver.
- Use multipathing software that supports both load balancing and path failover.

HBA Performance Metrics

The following metrics measure the performance of host bus adapters (HBA).

Metric	Description	Units	Use to...
Bytes Received	Inbound traffic in megabytes per second at the HBA port.	MB/s	Measure network traffic for load balancing, multi-path optimization, and network performance.
Bytes Transmitted	Outbound traffic in megabytes per second at the HBA port	MB/s	Measure network traffic for load balancing, multi-path optimization and network performance.
CRC Errors	Cyclic Redundancy Check errors over a period	Errors	Isolate Cyclic Redundancy Check (CRC) errors on a specific initiator or between devices.
Link Failures	Link Failures over a period	Failures	Isolate connection failures and the effect on performance.

HBA metrics might not be available for all operating systems or all HBAs.

Performance Management Host Server Metrics

By selecting a SAN host server component, you can view processor and memory utilization, host bus adapter (HBA) status, transmission metrics, and application path metrics. The application path maps from the mount point all the way to the array LDEV. Along the path you can choose any device to view its performance.

HP Storage Essentials provides the following metrics for host server performance management, and these are described in this guide:

- [Host Server Performance Metrics below](#)
- [Host Disk Performance Metrics below](#)
- [ESX Server Performance Metrics on next page](#)

Host Server Performance Metrics

Server performance metrics are collected by the CIM Extension on managed hosts. Performance metrics are not available for inferred or unmanaged servers.

You can use the following metrics to measure host performance.

Metric	Description	Units	Use to...
Free Physical Memory	Amount of physical memory available	KB	Measure available main memory for additional processes and threads.
Free Virtual Memory	Amount of virtual (paged) memory available	KB	Measure physical memory optimization and availability over time.
Physical Memory Used	Percentage (%) of main memory used by processes running on the host	%	Track memory utilization trends.
Processor Utilization	Total CPU utilization percentage (%) for all processes running on the host	%	Identify CPU bottlenecks.
Virtual Memory Used	Percentage (%) of paging file in use	%	Determine if a system has insufficient physical memory.

Server performance metrics might not be available for all supported operating systems.

Host Disk Performance Metrics

You can use the following metrics to measure disk performance on a server.

Metric	Description	Units	Use to...
Disk Read	Average time in seconds to read data from disk	KB/s	Compare read times for a given application (for example, read compared to writes).
Disk Total	Total read and write requests in seconds	KB/s	Test maximum throughput.
Disk Utilization	Based on the IRP (I/O request packets) round trip times the Average Disk Sec/Transfer. Indicates how busy a physical disk is over time.	%	Determine the average disk utilization for a given application or known number of processes. Utilization indicates how busy a disk is.
Disk Write	Average time in seconds to write data to disk	KB/s	Compare write times for a given application (for example, writes compared to reads).

Disk performance metrics might not be available for all supported operating systems.

ESX Server Performance Metrics

HP Storage Essentials tracks the following performance matrix for the ESX Server. These metrics are collected agentlessly; that is, the CIM Extension is not required on the ESX Server to view the HP Storage Essentials performance metrics.

Metric	Description	Units	Use to...
Free Physical Memory	Amount of physical memory available	KB	Measure available main memory for additional processes and threads.
Physical Memory Used	Physical memory being consumed by all processes.	%	Indicate physical memory optimization and availability over a period of time.
Processor Utilization	Total CPU utilization (%) for all processes running on the host	%	Identify CPU bottlenecks.
Disk Read	Average time in seconds to read data from disk	KB/s	Compare read times for Virtual Machines and their applications.
Disk Total	Total read/writes in seconds	KB/s	Test maximum throughput
Disk Write	Average time in seconds to write data to the disk	KB/s	Compare write times for Virtual Machines and their applications.

Metric	Description	Units	Use to...
Bytes Transmitted*	Inbound traffic in megabytes per second through the HBA.	MB/s	Measure network traffic for load balancing, multi-path optimization, and network performance.
Bytes Received*	Outbound traffic in megabytes per second through the HBA	MB/s	Measure network traffic for load balancing, multi-path optimization, and network performance.
CRC Errors*	Cyclic Redundancy Check (CRC) errors over a period of time	Errors	Isolate Cyclic Redundancy Check (CRC) errors on a specific initiator or between devices.
Link Failures*	Link Failures over a period of time	Failures	Isolate connection failures and effect on performance.

*Information displayed for the metric is discovered from the fabric switches, and not from the ESX server.

Managing Switch Performance

This section describes how to identify fibre channel performance issues and how to monitor your switch devices. It also describes in detail the HP Storage Essentials performance management metrics for the fabric switches running in your storage environment. These performance metrics are gathered by data collectors and available for reporting. For more information, see the following topics:

- [Identifying Fibre Channel Performance Issues below](#)
- [Monitoring Fabric Switch Performance on next page](#)

Identifying Fibre Channel Performance Issues

Fibre channel performance issues can be identified by performing a Cyclic Redundancy Check (CRC). CRC is a method of data integrity assurance across a transmission link. On the transmitting end, a mathematical computation is performed on the bitstream, and the result is added to the data frame. The process is reversed on the receiving end. If the two results do not match, a CRC error is generated, resulting in retransmission of the frame to maintain data integrity.

FC Errors

The following are examples of Fibre Channel (FC) errors that may be observed as a result of the above issue. This issue may not be the only reason that such errors occur. Other types of Fibre Channel errors could also potentially occur:

- Disconnects to Fibre Channel-attached storage during high I/O traffic.
- During high I/O load, CRC errors occur in conjunction with a Microsoft Windows error message: device not accessible.
- ProLiant BL20p G3 server blades stop responding with "Link failure," "loss of sync" or "loss of signal" errors logged at the switch when the HBA performs a link reset.
- Multiple path failures in multi-path environments.

Note: Because the Performance Manager does not support host events, you will need to use native tools, such as logs, event viewer, and so forth, to see host events.

CRC Errors

Brief Cyclic Redundancy Check (CRC) errors in HP Storage Essentials are a normal occurrence when an HBA is first powered on or off, or when cables are attached or detached. Excessive CRC errors during data transfers can cause performance degradation but do not compromise data integrity.

Link Failure

Link failure is the result of a loss of signal, loss of synchronization, or NOS primitive received. A link failure indicates that a link is actually "broken" for a period of time. It can possibly be due to a faulty connector, media interface adapter (MIA), or cable. The recovery for this type error is disruptive. This error is surfaced to the application using the SAN device that encountered this link failure. This causes the system to run degraded until the link recovery is complete. These errors should be monitored closely as they typically affect multiple SAN devices.

I/O Traffic

I/O traffic results have different meanings for different operating systems. The Linux and UNIX operating systems bundle small block I/O into large 128 KB block requests, and performance at the upper end of the I/O block spectrum is an important concern. Microsoft Windows, on the other hand, defaults to a maximum I/O block of 64 KB and does not bundle small requests into larger ones.

Monitoring Fabric Switch Performance

Switch performance best practices should focus on the establishment of baselines. Use Aggregate Port and Port I/O metrics to establish typical IOPS rates and throughput rates as well as common error rates, average queue depths, and response times. Monitoring SAN switch and overall SAN performance primarily involves three metrics: IOPS (I/O operations per seconds), bandwidth, and latency.

Measuring IOPS and bandwidth can tell you how much work or activity is taking place in the SAN. Measuring latency tells you how effectively the SAN is doing its work, as well as whether the SAN is meeting its service objectives. By using switches and HBAs to view error rates, you can pinpoint the source of SAN performance problems. Error rates can include loss of signal or synchronization, re-transmissions, link failure, or invalid CRC.

Best Practices

Follow these best practices to optimize switch performance:

- Keep the highest performing directors at the core of the SAN.
- Connect storage devices and the highest performing applications to the core.
- Benchmark the performance on oversubscribed ports.
- Leave the Fibre Channel (FC) ports at auto-negotiate for host and storage connections.

Fibre Channel Port Types

Understanding FC port types can help you to identify ports along a storage path on which you are monitoring or collecting performance metrics. The following tables describe the different types of Fibre Channel ports:

Node Ports	Description
N_ port	Port on the node (such as, host or storage device) used with both FC-P2P or FC-SW topologies; also known as Node port.
NL_ port	Port on the node used with an FC-AL topology; also known as Node Loop port.
F_ port	Port on the switch that connects to a node point-to-point (for example, connects to an N_ port); also known as Fabric port. An F_ port is not loop capable
FL_ port	Port on the switch that connects to an FC-AL loop (such as, to NL_ ports); also known as Fabric Loop port.
E_ port	Connection between two fibre channel switches. Also known as an Expansion port. When E_ ports between two switches form a link, that link is referred to as an inter-switch link (ISL).
EX_ port	Connection between a fibre channel router and a fibre channel switch. On the side of the switch it looks like a normal E_ port, but on the side of the router it is an EX_ port.
TE_ port	Cisco addition to Fibre Channel, now adopted as a standard. It is an extended ISL or EISL. The TE_ port provides not only standard E_ port functions but allows for routing of multiple VSANs (Virtual SANs). This is accomplished by modifying the standard Fibre Channel frame (vsan tagging) upon ingress/egress of the VSAN environment. The TE_ port is also known as Trunking E_ port.

General Ports		Description
Auto		Auto or auto-sensing port found in Cisco switches, can automatically become an E_, TE_, F_, or FL_port as needed.
Fx_port		Generic port that can become an F_port (when connected to a N_port) or a FL_port (when connected to an NL_port). Found only on Cisco devices where over-subscription is a factor.
G_port		G_port or generic port on a switch that can operate as an E_port or F_port. The G_port is found on Brocade and McData switches.
L_port		Loose term used for any arbitrated loop port, NL_port or FL_port. L_port is also known as Loop port.
U_port		Loose term used for any arbitrated port. U_port is also known as Universal port and is found only on Brocade switches.

Performance Management Switch Metrics

By selecting a SAN fabric switch, you can view the aggregated bytes transmitted and received, and the individual switch port metrics including Cyclical Redundancy Check (CRC) errors and link failures. Additionally, you can navigate to other switches, hosts, and arrays connected to the switch and view their metrics.

See [Switch Performance Metrics below](#) for a list of HP Storage Essentials performance management switch metrics.

Switch Performance Metrics

The following metrics are used to monitor switch performance.

Metric	Description	Units	Common Use
Aggregated Port Bytes Received	Sum of bytes received for all ports in a switch over an interval	MB/s	Measure inbound traffic for all ports on the switch.
Aggregated Port Bytes Transmitted	Sum of bytes transmitted for all ports in a switch over an interval	MB/s	Measure outbound traffic for all ports on the switch.
Bytes Received	Number of bytes received over a given interval	MB/s	Measure inbound traffic for specific ports on the switch.
Bytes Transmitted	Number of bytes transmitted over a given interval	MB/s	Measure outbound traffic for specific ports on the switch.

Metric	Description	Units	Common Use
CRC Errors	Number of Cyclic Redundancy Check errors over a period of time	Errors	Isolate CRC errors on a specific initiator or between devices
Link Failures	Number of link Failures over a period of time	Failures	Isolate connection failures and the effect on performance

Managing EMC Symmetrix Performance

This section describes the HP Storage Essentials performance metrics for the EMC Symmetrix devices running in your storage environment. It includes information about scheduling and collecting performance data (see [Collecting EMC Symmetrix Performance Statistics](#) below), and explains how performance metrics are derived (see [About Counter Rollover](#) on page 119).

For a list of performance management metrics, see [EMC Symmetrix Performance Metrics](#) on page 123 and [Transient Values for EMC Symmetrix Performance Metrics](#) on page 134.

Collecting EMC Symmetrix Performance Statistics

If you purchased the HP Storage Essentials Performance Pack Enterprise for EMC Symmetrix license, HP Storage Essentials uses the internal Symmetrix SMI-S provider to request and receive Symmetrix performance data. This data is sent and received through remote communications to the **storsrvd** daemon running on the EMC Solutions Enabler server that HP Storage Essentials is currently using to communicate with the corresponding Symmetrix. HP Storage Essentials also requires the **storsrvd** daemon to be running for regular (non-performance) data collection.

All EMC Symmetrix statistical performance data collection, both scheduled and real time, is triggered using the HP Storage Essentials Performance Manager user interface. In the 9.4.0 version of HP Storage Essentials, the requests for Symmetrix performance data are serviced by the HP Storage Essentials internal Symmetrix SMI-S provider (not by the EMC Symmetrix SMI-S provider bundled with EMC Solutions Enabler).

Caution: It is important that the EMC Solutions Enabler **storstpd** daemon is NOT started or running while the HP Storage Essentials Performance Manager is using the same Solutions Enabler server to collect performance data. HP Storage Essentials will likely encounter problems collecting Symmetrix performance data if the **storstpd** daemon is running. See [Scheduling EMC Symmetrix Performance Data Collection](#) on next page.

Before You Collect Statistics

Check the HP Storage Essentials support matrix (available from the Documentation Center located in the top-level directory of the StorageEssentialsDVD) for the minimum EMC Solutions Enabler version requirements. Attempts to collect Symmetrix performance data using an unsupported Solutions Enabler server can result in errors.

Best Practices

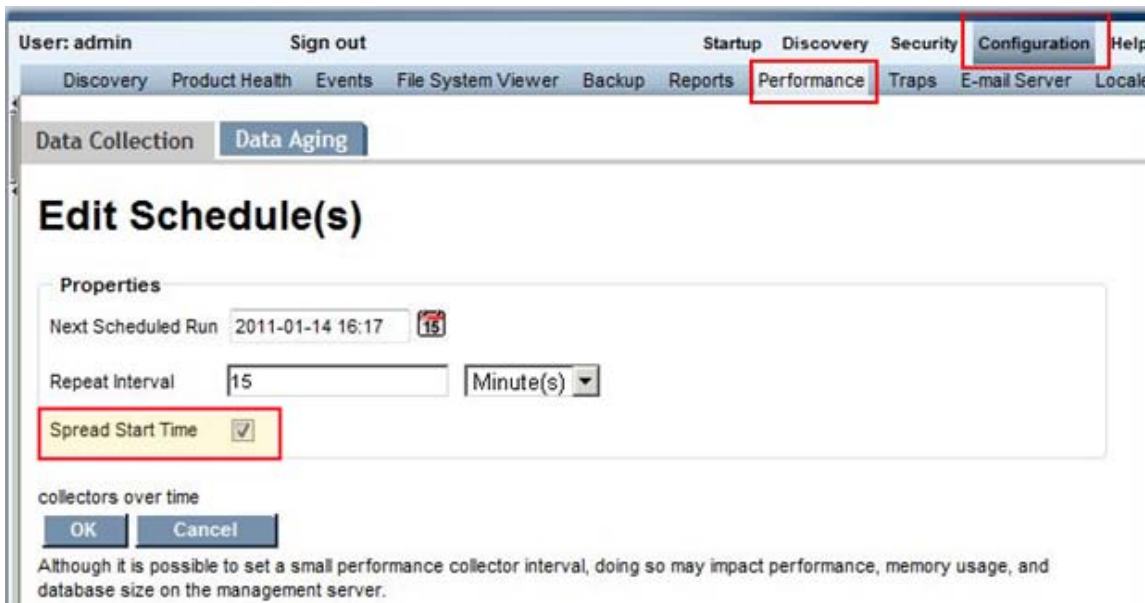
Keep in mind the following recommended best practices when collecting performance data:

- Do not schedule collection using an unsupported version of the EMC Solutions Enabler server. This will result in errors.
- When possible, schedule collection frequencies between 15 and 30 minutes. Scheduled collection frequencies greater than 2 hours will not be serviced.
- HP does NOT recommend that you schedule a collection frequency of less than 5 minutes.
- Check the HP Storage Essentials cimom and appstom logs when expected data is missing or errors are suspected.

Scheduling EMC Symmetrix Performance Data Collection

Due to limitations in the communications interface to the **storsrvd** daemon, only one request for data at a time can be made to a single Solutions Enabler server. For this reason, it is important that you schedule collection of Symmetrix performance data using the HP Storage Essentials Spread Start Time option (default) on the Data Collection - Edit Schedule page (**Configuration > Performance > Data Collection**).

This figure shows the Spread Start Time option on the Edit Schedules page on the Data Collection tab.



The amount of time to complete a single request can vary according to the number of volume statistics requested for a single Symmetrix and other factors. If the scheduled frequency (referred to as Interval) of collections (referred to as Runs) for Symmetrix Performance Data Collectors that are targeted at the same Solutions Enabler server is too short, collections could start backing up or erroring out.

Collection Frequency

The recommended frequency for Symmetrix data collection is between 15 and 30 minutes.

Data collection errors could begin surfacing for frequencies of less than 5 minutes, especially in larger environments. With the exception of the few statistics described as “point in time” in [EMC Symmetrix Performance Metrics on page 123](#), each EMC Symmetrix statistic value is averaged over the sample period between two successive collections. For example, if the scheduled collection frequency is 15 minutes, then a Read Data Rate value of 1,500 KB/s means that there was an average of 1,500 KB transferred per second between the time at which this statistic was collected and the previous collection ten minutes ago. Due to this averaging, scheduled frequencies of more than 1 hour tend to flatten the data in the line graphs, diminishing the granularity and value of the data represented by minimizing visual spikes and dips in the data rate between scheduled collections.

Although the scheduling interface allows input of frequencies for 2 hours or greater, the Performance Manager collection engine will fail to produce and store statistics for intervals greater than 2 hours, in which case, errors will be logged.

Real-time Data Collection

Real-time data collection in the Performance Manager is intended for short-term inspection of a few targeted statistics, and is not affected by, and is independent of, the previously mentioned schedules. Real-Time collectors are fixed at 20 second intervals. If expected Symmetrix performance data is missing or if errors are encountered while collecting real-time data, check both the cimom and the appstorm logs.

About Counter Rollover

When you calculate derived statistics, it is important to account for *counter rollover*. Counter rollover occurs where an accumulating performance counter in a computer system (for example, `sequential_read_ios`) reaches its maximum value and then restarts the count from zero.

A useful analogy is the odometer in your automobile. A counter rollover occurs when the odometer reaches its maximum value, rolls over to zero, and begins counting again. As a hypothetical example, if the maximum odometer count were 999,999 miles, when the car reaches 1,000,000 miles, the odometer would show 0 (zero) miles. If you calculated the miles traveled between mileage points, then you would calculate the delta between two readings taking into account the rollover. For example, if the maximum limit for the odometer is 999,999 miles, and the last reading was 900,000 and the current reading is 100,000, the change in the value since the last reading would be 200,000 miles.

In HP Storage Essentials, where possible, variables are derived by using the deltas or summations of the individual parts. This facilitates a more accurate calculation of the change in a counter value (for example, `deltaReadIOs`) in regards to handling a single rollover condition since the last time the counter value was read. You will find more information about this in the following topics, which describe how to account for counter rollover when calculating derived performance statistics:

- [How Changes in Counter Values Are Calculated below](#)
- [How Rates Account for Single Rollover on the facing page](#)
- [Duration Values Used to Calculate Rates on page 122](#)
- [Counters and KByte Conversion on page 123](#)

How Changes in Counter Values Are Calculated

Typically, a change in value is calculated as `current - previous`; that is, the current value minus the previous value. In the case of a counter rollover, however, this formula produces an incorrect value. For example, if the current value is 100,000 and the previous value was 900,000, `current - previous` equals "- 800,000," which (due to counter rollover) is incorrect.

Therefore, when the current value is less than the previous value, an assumption is made that the counter has rolled over, and only rolled over one time. For this reason, when a rollover is detected, the change is calculated using the following formula:

```
delta = (1 + max_counter_value) - previous_value + current_value
```

The notation `(1 + max_counter_value)` can be expressed more simply as b^n , where b is number base and n is the number of counter digits, referred to as `num_counter_bits`.

If we use the odometer example described in [About Counter Rollover on previous page](#), the number base (given in decimal digits) is 10, and the number of digits is 6. Therefore, b^n is 10^6 and equals 1,000,000. This is the same value as `1 + max_counter_value`.

Here is how we derive the delta value:

```
b^n = 10^6 = 1,000,00
```

```
previous value = 900,000
```

```
current value = 100,000
```

```
1,000,000 - 900,000, + 100,000 = 200,000
```

This formula accounts for a single rollover condition and produces the correct delta value of 200,000 miles.

In the previous example, b^n is the same as `max_counter_value + 1`.

Calculating the Number Base

The number base for accumulating performance counters on computer systems is binary. As such, $b=2$ and `num_counter_bits` is usually 64 or 32.

In most cases, the `(1 + max_counter_value)` is calculated as follows:

```
b^n = 2^64 = 18,446,744,073,709,551,616
```

```
b^n = 2^32 = 4,294,967,296
```

For the Symmetrix performance metrics calculations, keep in mind the following.

The `num_counter_bits` is assumed to be 64 and, thus, $(1 + \text{max_counter_value}) = 2^{64} = 18,446,744,073,709,551,616$.

Any other value for `num_counter_bits` is explicitly noted in this guide.

Any variable prefixed with “delta” implies the variable’s value was set using the following logic:

```
if current_value < previous_value
delta = (1 + max_counter_value) - previous_value + current_value
else
delta = current - last
```

How Rates Account for Single Rollover

The following example uses the Total Read Rate metric to describe how rates account for single rollover.

The main variable used to calculate the Total Read Rate metric is `deltaTotalReadIOs`. When calculating the value for `deltaTotalReadIOs`, the deltas for both `ReadIOsRandom` and `ReadIOsSequential` are first calculated, where a single rollover in each can be accounted for. Next, the delta in `TotalReadIOs` is set equal to the sum of the individual parts. The calculations used in these steps are described in detail below.

Note: When calculating counter rollover, there is an assumption that a counter has not rolled over more than once.

Here are the steps for how a typical rate such as Total Read Rate is calculated:

1. Calculate delta and account for single rollover for random I/O (if applicable):

```
deltaReadIOsRandom = current_ReadIOsRand - previous_ReadIOsRandom
```

2. Calculate delta and account for single rollover for sequential I/O (if applicable):

```
deltaReadIOsSequential = current_ReadIOsSeq - previous_ReadIOsSeq
```

3. Calculate total read I/Os using previously derived delta values in steps 1 and 2:

```
deltaTotalReadIOs = deltaReadIOsRandom + deltaReadIOsSequential
```

4. Calculate duration between collection points:

```
duration_in_seconds = current_collection_time_counter - previous_
collection_time_counter
```

5. Calculate total read rate using previously derived values in steps 3 and 4.

```
TotalReadRate = deltaTotalReadIOs / duration_in_seconds
```

Example

The following example assumes that the Read IO counters have a maximum value of 999,999. When they reach the maximum value, the counter rolls over to zero (0). The values are returned as follows:

```
previous_collection_time_counter = 600,000
current_collection_time_counter = 600,100
previous_ReadIOsRandom = 200,000
current_ReadIOsRandom = 800,000
previous_ReadIOsSeq = 900,000
current_ReadIOsSeq = 200,000
```

Using the steps, this is how the Total Read Rate is calculated using the above values:

1. Calculate delta value for random read I/Os with no rollover:

```
deltaReadIOsRandom: 800,000 - 200,000 = 600,000
```

2. Calculate the delta value for sequential read I/Os with rollover:

```
deltaReadIOsSequential (rollover): (999,999 + 1) - 900,000 +
200,000 = 300,000
```

3. Calculate the total read I/Os using the values derived in steps 1 and 2:

```
deltaTotalReadIOs = 600,000 + 300,000 = 900,000
```

4. Calculate the value for duration:

```
duration_in_seconds: 600,100 - 600,000 = 100
```

5. Calculate total read rate using values derived in steps 3 and 4:

```
TotalReadRate = 900,000 IOs / 100 seconds = 900 ReadIOs / second
```

Duration Values Used to Calculate Rates

Duration is the elapsed time between the previous statistic collection and the current statistic collection. It is typically measured in seconds and used to calculate rates (such as ReadIOs per second).

The `num_sym_timeslices` counter is the Symmetrix time stamp in half-second units and represents the time at which the performance data statistics were collected and stored in the array. This counter provides for a more accurate calculation of duration than simply using the difference in system timestamps between the last collection and the current collection.

Where available, the EMC Symmetrix `num_sym_timeslices` counter is used to calculate the duration in seconds:

```
duration = deltaNumHalfSecTimeSlices / 2
```

For statistics that do not include the `num_sym_timeslices` counter, the system timestamp is used.

Counters and KByte Conversion

Most EMC Symmetrix counters that represent KBytes have been converted from "512 Byte" blocks to KBytes by dividing by 2 because there are two blocks per KByte.

Therefore, the max value, or point of rollover, for such counters in the database is half of the normal 64-bit max value used for the actual number of blocks returned to the counter by the Symmetrix array. This means that the num_counter_bits value is reduced by one.

To help you understand to which counters the above applies, the value of num_counter_bits is explicitly noted as either 63 or 31 in the metric descriptions in [EMC Symmetrix Performance Metrics](#) below.

EMC Symmetrix Performance Metrics

The following tables describe the HP Storage Essentials performance metrics collected for EMC Symmetrix arrays. For descriptions of the transient statistics shown in the formulas, see [Transient Values for EMC Symmetrix Performance Metrics](#) on page 134.

EMC Symmetrix System Performance Metrics

EMC Symmetrix system metrics are stored in the EMCSYMSTORAGESYSTEMSTATS and MVCS_EMCSYMSSTATSVW database tables.

Metric data that is rolled up into hourly, daily, weekly, and monthly data points is stored in the corresponding EMCSYMSTORAGESYSTEMSTATS rollup tables and have a suffix of _H, _D, _W, and _M respectively. For example, hourly collected data is stored in EMCSYMSTORAGESYSTEMSTATS_H and rolled up into daily data statistics stored in the EMCSYMSTORAGESYSTEMSTATS_D table.

The following table describes the performance metrics provided by HP Storage Essentials for EMC Symmetrix storage systems.

Metric	Description	Unit	Formula
Average Read Size	Average read size.	Bytes	$\text{AvgReadSize} = (\text{deltaKBytesRead} \times 1024) / \text{deltaTotalReadIOsRandomAndSeq}$
Average Write Size	Average write size.	Bytes	$\text{AvgWriteSize} = (\text{deltaKBytesWritten} \times 1024) / \text{deltaTotalWriteIOsRandomAndSeq}$
Deferred Write Rate	Rate of deferred write requests. A deferred write is a write hit. A deferred write occurs when the I/O write operations are staged in cache and will be written to disk at a later time.	Req/s	$\text{DeferredWriteRate} = \text{deltaEMCDeferredWriteIOs} / \text{duration}$

Metric	Description	Unit	Formula
Delayed DFW Rate	Delayed DFW request rate. A delayed deferred fast write (DFW) is a write-miss. A delayed DFW occurs when the I/O write operations are delayed because the system or device write-pending limit was reached and the cache had to de-stage slots to the disks before the writes could be written to cache.	Req/s	$\text{DelayedDfwRate} = \text{deltaEMCDelayedDFWIOs} / \text{duration}$
Max Pending Flush Limit	Maximum number of write-pending slots for the entire Symmetrix. System write-pending limit is equal to 80% of the available cache slots. Symmetrix write-pending limit is not simply a sum of all Symmetrix device write-pending slots. It depends on other factors such as cache size and the Symmetrix configuration. System property. This is a point-in-time value captured at the time the statistics are taken.	Bytes	$\text{MaxPendingFlushLimit} = \text{EMCMaxKBPendingFlush} \times 1024$
Pending Flush	Number of tracks in cache that are waiting to be de-staged to disk and cannot be overwritten. This is a point-in-time value captured at the time the statistics are taken.	Bytes	$\text{PendingFlush} = \text{EMCKBPendingFlush} \times 1024$
Pending Format	Number of format pending tracks. This count can be less than the last-taken statistic; it is a point-in-time value captured at the time the statistics are taken.	Bytes	$\text{PendingFormat} = \text{EMCKBPendingFormat} \times 1024$
Percent Hits	Ratio of total hits (random and sequential) to total I/Os (random and sequential).	%	$\text{PctHitIOs} = 100 \times (\text{deltaTotalHitIOsRandomAndSeq} / \text{deltaTotalIOsRandomAndSeq})$
Percent Read Hits	Read cache hit ratio (percentage of read hits). ²	%	$\text{PctReadHitIOs} = 100 \times (\text{deltaReadHitIOsTotal} / \text{deltaTotalReadIOsRandomAndSeq})$

Metric	Description	Unit	Formula
Percent Reads	Ratio of read I/Os to total I/Os.	%	$\text{PctReadIOs} = 100 \times (\text{deltaTotalReadIOsRandomAndSeq} / \text{deltaTotalIOsRandomAndSeq})$
Percent Reads Seq ¹	Sequential read rate. (percentage of sequential reads to TotalIOs including SequentialReads). ²	%	$\text{PctSeqReadIOs} = 100 \times (\text{deltaReadIOsSeq} / \text{deltaTotalReadIOsRandomAndSeq})$
Percent Write Hit I/Os	Write cache hit ratio (percentage of write hits). ²	%	$\text{PctWriteHitIOs} = 100 \times (\text{deltaWriteHitIOsTotal} / \text{deltaTotalWriteIOsRandomAndSeq})$
Percent Writes	Ratio of write I/Os to total I/Os.	%	$\text{PctWriteIOs} = 100 \times (\text{deltaTotalWriteIOsRandomAndSeq} / \text{deltaTotalIOsRandomAndSeq})$
Prefetch Data Rate	Rate of prefetched Bytes per second.	Bytes/s	$\text{PrefetchRate} = (\text{deltaEMCKBPrefetched} \times 1024) / \text{duration}$
Read Data Rate	Read throughput rate (Bytes per second).	Bytes/s	$\text{ReadDataRate} = (\text{deltaKBytesRead} \times 1024) / \text{duration}$
Read Hits	Read cache hit rate.	Req/s	$\text{ReadHitRate} = \text{deltaReadHitIOsTotal} / \text{duration}$
Read Hits Seq	Rate of read cache hits per second (sequential hits only).	Req/s	$\text{SeqReadHitRate} = \text{deltaReadHitIOsSeq} / \text{duration}$
Read Rate Random	Random read cache request rate (requests per second).	Req/s	$\text{ReadRate} = \text{deltaReadIOs} / \text{duration}$
Read Rate Seq	Sequential read rate. ²	Req/s	$\text{SeqReadRate} = \text{deltaReadIOsSeq} / \text{duration}$
Read Rate Total	Read request rate that includes both random and sequential reads.	Req/s	$\text{ReadRateTotal} = \text{deltaTotalReadIOsRandomAndSeq} / \text{duration}$
Total Data Rate	Total bytes read and written per second.	Bytes/s	$\text{TotalDataRate} = (\text{deltaKBytesTransferred} \times 1024) / \text{duration}$
Total I/O Rate	I/O rate which includes random and sequential reads and writes.	Req/s	$\text{TotalIORate} = \text{deltaTotalIOsRandomAndSeq} / \text{duration}$
Write Data Rate	Write throughput rate (Bytes per second).	Bytes/s	$\text{WriteDataRate} = (\text{deltaKBytesWritten} \times 1024) / \text{duration}$

Metric	Description	Unit	Formula
Write Flush Data Rate	Number of tracks (expressed in Bytes) written per sec from cache to disks.	Bytes/s	$\text{WriteFlushRate} = (\text{deltaEMCWriteKBytesFlushed} \times 1024) / \text{duration}$
Write Hits	Write cache hit rate.	Req/s	$\text{WriteHitRate} = \text{deltaWriteHitIOsTotal} / \text{duration}$
Write Hits Seq	Rate of write cache hits per second (sequential hits only).	Req/s	$\text{SeqWriteHitRate} = \text{deltaWriteHitIOsSeq} / \text{duration}$
Write Rate Total	Write cache request rate (requests per second) that includes both random and sequential writes.	Req/s	$\text{WriteRate} = \text{deltaTotalWriteIOsRandomAndSeq} / \text{duration}$
Write Rate Seq	Write cache request rate (requests per second) and includes only sequential writes.	Req/s	$\text{SeqWriteRate} = \text{deltaWriteIOsSeq} / \text{duration}$

¹ The SYM CLI is calculating setPctSeqReadIOs as ratio of ReadSeq to Total Reads. The EMC Perf Metrics glossary defines this as ReadsSeq to TotalIOs.

² This metric does NOT apply to Symm3 data.

EMC Symmetrix Volume Performance Metrics

EMC Symmetrix volume metrics are stored in the EMCSYMBOLUMESTATS and MVCS_EMCSYMBOLUMESTATSVW database tables.

Metric data that is rolled up into hourly, daily, weekly, and monthly data points is stored in the corresponding EMCSYMBOLUMESTATS rollup tables having a suffix of _H, _D, _W, and _M respectively. For example, the EMCSYMBOLUMESTATS_W table contains the rolled-up weekly performance data statistics for EMC Symmetrix volumes.

The following table describes the HP Storage Essentials performance metrics for volumes connected to EMC Symmetrix storage systems.

Metric	Description	Unit	Formula
Average I/O Size	Average size of an I/O operation performed by the Symmetrix device.	Bytes	$\text{AvgIOSize} = (\text{deltaKBytesTransferred} \times 1024) / \text{deltaTotalIOsRandomAndSeq}$
Average Read Size	Average size of a read I/O operation performed by the Symmetrix device.	Bytes	$\text{AvgReadSize} = (\text{deltaKBytesRead} \times 1024) / \text{deltaTotalReadIOsRandomAndSeq}$

Metric	Description	Unit	Formula
Average Write Size	Average size of a write I/O operation performed by the Symmetrix device.	Bytes	$AvgWriteSize = (\text{deltaKBytesWritten} \times 1024) / \text{deltaTotalWriteIOsRandomAndSeq}$
Max Write Pending Threshold	Maximum number of write-pending slots available (expressed in Bytes) for the Symmetrix device. ¹	Bytes	$MaxWritePendingThreshold = \text{current_EMCMaxKBPendingFlush} \times 1024$
Pending Flush ²	Number of cache slots (expressed in Bytes) that were write pending for the logical volume at a point in time. This number changes according to the cache de-stage activity rate and the number of writes. A write is pending when it has been written to cache but has not yet been written to the disk.	Bytes	$PendingFlush = \text{current_EMCKBPendingFlush} \times 1024$
Percent Hits	Percentage of I/O cache hit operations performed by the Symmetrix device that were immediately satisfied by cache.	%	$PctHitIOs = 100 \times (\text{deltaTotalHitIOsRandomAndSeq} / \text{deltaTotalIOsRandomAndSeq})$
Percent Misses	Percentage of read and write operations performed by the Symmetrix device that were misses.	%	$PctMissIOs = 100 - PctHitIOs$
Percent Read Hits Random	Ratio of read hit I/Os to Total I/Os.	%	$PctReadHitIOsRandom = 100 \times (\text{deltaReadHitIOs} / \text{deltaTotalIOsRandomAndSeq})$
Percent Read Hits Total	Percentage of read cache hit I/Os (including both random and sequential) operations performed by the Symmetrix device.	%	$PctReadHitIOsTotal = 100 \times (\text{deltaReadHitIOsTotalRandomAndSeq} / \text{deltaTotalReadIOsRandomAndSeq})$

Metric	Description	Unit	Formula
Percent Reads	Percentage of read I/O operations performed by the Symmetrix device.	%	$PctReadIOs = 100 \times (\Delta TotalReadIOsRandomAndSeq / \Delta TotalIOsRandomAndSeq)$
Percent Read Misses Random	Ratio of read cache miss I/Os to Total I/Os.	%	$PctReadMissIOsRandom = 100 \times (\Delta ReadMissIOsRandom / \Delta TotalIOsRandomAndSeq)$
Percent Read Misses Total	Percentage of read cache miss I/O operations performed by the Symmetrix device.	%	$PctReadMissIOsTotal = 100 \times (\Delta ReadMissIOsTotalRandomAndSeq / \Delta TotalReadIOsRandomAndSeq)$
Percent Write Hits	Percentage of cache write hit I/O operations performed by the Symmetrix device.	%	$PctWriteHitIOs = 100 \times (\Delta WriteHitIOsTotalRandomAndSeq / \Delta TotalWriteIOsRandomAndSeq)$
Percent Writes	Percentage of total write I/O operations performed by the Symmetrix device.	%	$PctWriteIOs = 100 \times (\Delta TotalWriteIOsRandomAndSeq / \Delta TotalIOsRandomAndSeq)$
Percent Write Misses	Percentage of write I/O operations performed by the Symmetrix device that were write misses.	%	$PctWriteMissIOs = 100 \times (\Delta WriteMissIOsTotalRandomAndSeq / \Delta TotalWriteIOsRandomAndSeq)$
Read Data Rate	Number of Bytes read by the Symmetrix device each second.	Bytes	$ReadDataRate = (\Delta KBytesRead \times 1024) / \text{duration}$
Read Hit Rate Random	Number of random read hit I/O operations performed each second by the Symmetrix device. The read hits per sec metric for the Symmetrix device statistic does not include sequential read hits. In contrast, the Read Hit Rate Total metric includes random and sequential read hits per second.	Req/s	$ReadHitRateRandom = \Delta ReadHitIOs / \text{duration}$

Metric	Description	Unit	Formula
Read Hit Rate Seq	Number of sequential read hit requests performed (per second) by the Symmetrix device.	Req/s	$\text{SeqReadHitRate} = \text{deltaReadHitIOsSeq} / \text{duration}$
Read Hit Rate Total	Total number of read hit operations (random and sequential) performed each second by the Symmetrix device.	Req/s	$\text{ReadHitRateTotal} = \text{deltaReadHitIOsTotalRandomAndSeq} / \text{duration}$
Read Rate Random	Number of I/O operations performed each second by the Symmetrix device that were random reads. This Random Reads per sec metric for the Symmetrix device statistic does not include sequential reads. In contrast, the Read Rate Total metric includes random and sequential read hits per second.	Req/s	$\text{ReadRateRandom} = \text{deltaReadIOs} / \text{duration}$
Read Rate Seq	Number of sequential read I/O operations performed each second by the Symmetrix device.	Req/s	$\text{SeqReadRate} = \text{deltaReadIOsSeq} / \text{duration}$
Read Rate Total	Read request rate including both random and sequential read operations performed each second by the Symmetrix device.	Req/s	$\text{ReadRateTotal} = \text{deltaTotalReadIOsRandomAndSeq} / \text{duration}$
Sampled Average Read Time	Completion time of a read as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os.	MS	$\text{SampledAvgReadTimeMs} = \text{curr.getEMCSampledReadsTime}(), \text{curr.getEMCSampledReads}(), \text{null}$

Metric	Description	Unit	Formula
Sampled Average Write Time	Completion time of a write as measured by the host director. Measurements are taken for a sample set of approximately 30% of the I/Os.	MS	$\text{SampledAvgWriteTimeMs} = \text{current_EMCSampledWritesTime} / \text{current_EMCSampledWrites}$
Total Data Rate	Total Bytes read and written per second.	Bytes/s	$\text{TotalDataRate} = \text{deltaKBytesTransferred} \times 1024 / \text{duration}$
Total Hit Rate	Total number of I/O operations (random and sequential) performed each second by the Symmetrix device that were immediately satisfied by cache.	Req/s	$\text{TotalHitRate} = \text{readHitRateTotalRandomAndSeq} + \text{writeHitRateTotalRandomAndSeq}$
Total I/O Rate	Total number of read I/O and write I/O operations (random and sequential) performed each second by the Symmetrix device.	Req/s	$\text{TotalIORate} = \text{readRateTotalRandomAndSeq} + \text{writeRateTotalRandomAndSeq}$
Total I/O Rate Random	Number of I/O operations performed each second by the Symmetrix device, including writes and random reads. In contrast, the Total IO Rate metric includes writes, random reads, and sequential reads.	Req/s	$\text{TotalIORateRandom} = \text{deltaTotalIOsRandom} / \text{duration}$
Total Miss Rate	Total number of I/O operations (random and sequential) performed each second by the Symmetrix device that were NOT immediately satisfied by cache	Req/s	$\text{TotalMissRate} = \text{TotalIORate} - \text{TotalHitRate}$
Write Data Rate	Number of Bytes written by the Symmetrix device each second.	Bytes/s	$\text{WriteDataRate} = (\text{deltaKBytesWritten} * 1024) / \text{duration}$

Metric	Description	Unit	Formula
Write Hit Rate Total	Number of write hit operations (random and sequential) performed each second by the Symmetrix device.	Req/s	$\text{WriteHitRate} = \frac{\text{deltaWriteHitIOsTotalRandomAndSeq}}{\text{duration}}$
Write Hit Rate Seq	Number of sequential write hits that occurred each second for the Symmetrix device.	Req/s	$\text{SeqWriteHitRate} = \frac{\text{deltaWriteHitIOsSeq}}{\text{duration}}$
Write Miss Rate	Number of write misses that occurred for the Symmetrix device each second.	Req/s	$\text{WriteMissRate} = \frac{\text{deltaWriteMissIOsTotalRandomAndSeq}}{\text{duration}}$
Write Rate Random	Write cache request rate (requests per second) that includes only random writes by the Symmetrix device.	Req/s	$\text{WriteRate} = \frac{\text{deltaWriteIOs}}{\text{duration}}$
Write Rate Seq	Number of sequential write I/O operations performed each second by the Symmetrix device.	Req/s	$\text{SeqWriteRate} = \frac{\text{deltaWriteIOsSeq}}{\text{duration}}$
Write Rate Total	Write cache request rate (requests per second) including both random and sequential I/Os performed for the Symmetrix device.	Req/s	$\text{WriteRateTotal} = \frac{\text{deltaTotalWriteIOsRandomAndSeq}}{\text{duration}}$

¹ Max Write Pending Threshold is not a static number. It depends on Symmetrix activity. Each Symmetrix device is assigned a limit of write-pending slots that can dynamically change between a base value and a value three times the base (the maximum value). When the Max Write Pending Threshold reaches three times the base value, writes to the device are delayed so that the cache can destage, which frees the cache slots. As cache slots are freed, the writes resume. While the write-pending limit is reached, disk directors operate in a priority destage write mode. This gives write data higher priority than normal. During the delay, writes to this Symmetrix device are counted as write misses.

² Certain HP Storage Essentials metrics, such as Pending Flush, are shown in bytes, whereas their corresponding EMC metric are shown in tracks. The corresponding EMC metric for Pending Flush is EMC Num WP Tracks. The Tracks-to-KB conversion factor is based on the Enginuity version of the EMC Symmetrix array as follows: for Enginuity < 5771: Tracks-to-KB = 32; for Enginuity >= 5771: Tracks-to-KB = 64 Num Tracks = (number of Bytes / 1024) / Tracks-to-KB.

EMC Symmetrix Front-end Controller Performance Metrics

EMC Symmetrix front-end controller metrics are stored in the EMCSYMFECROLLERSTATS and MVCS_EMCSYMFECNTSTATSVW database tables.

Metric data that is rolled up into hourly, daily, weekly, and monthly data points is stored in the corresponding EMCSYMFECROLLERSTATS roll-up tables having a suffix of _H, _D, _W, and _M respectively. For example, monthly rolled up data is stored in the EMCSYMFECROLLERSTATS_M database table.

The following table describes the performance metrics provided by HP Storage Essentials for EMC Symmetrix front-end controllers.

Metric	Description	Unit	Formula
Device Write Pending Event Rate	Number of times each second that the write-pending limit for a specific Symmetrix device was reached. When the limit is reached, additional write I/O operations are deferred while waiting for data in cache to be destaged to the disk.	Events/s	$\text{DeviceWritePendingEventRate} = \frac{\text{deltaEMCDeviceFlushPendingEvents}}{\text{duration}}$
Percent Hits	Percentage of requests performed by the host director and immediately satisfied by cache.	%	$\text{PctHitIOs} = 100 \times \left(\frac{\text{deltaEmcTotalHitIOs}}{\text{deltaTotalIOs}} \right)$
Percent Reads	Percentage of read requests performed by the host director.	%	$\text{PctReadIOs} = 100 \times \left(\frac{\text{deltaReadIOs}}{\text{deltaTotalIOs}} \right)$
Percent Writes	Percentage of write requests performed by the host director over the sample interval.	%	$\text{PctWriteIOs} = 100 \times \left(\frac{\text{deltaWriteIOs}}{\text{deltaTotalIOs}} \right)$
Read Rate	Number of random read requests performed each second by Symmetrix host director.	Req/s	$\text{ReadRate} = \frac{\text{deltaReadIOs}}{\text{duration}}$
Slot Collision Rate	Number of slot collisions each second. A slot collision occurs when two or more directors try to access the same cache slot and the slot happens to be locked for an update operation by one of the directors.	Collisions/s	$\text{SlotCollisionRate} = \frac{\text{deltaEMCSlotCollisions}}{\text{duration}}$

Metric	Description	Unit	Formula
System Write Pending Event Rate	Number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.	Events/s	$\text{SystemWritePendingEventRate} = \frac{\text{deltaEMCSystemFlushPendingEvents}}{\text{duration}}$
Total Data Rate	Number of Bytes transferred through the Symmetrix Director each second.	Bytes/s	$\text{TotalDataRate} = \frac{(\text{deltaKBytesTransferred} \times 1024)}{\text{duration}}$
Total Hit Rate	Number of read and write requests performed each second by the host director that was immediately satisfied by cache.	Req/s	$\text{TotalHitRate} = \frac{\text{deltaEMCTotalHitIOs}}{\text{duration}}$
Total I/O Rate	Number of I/O operations performed each second by the Symmetrix host director. This metric represents activity between the Symmetrix device and the host or SAN device.	Req/s	$\text{TotalIORate} = \frac{\text{deltaTotalIOs}}{\text{duration}}$
Write Rate	Number of write requests performed each second by the host directors.	Req/s	$\text{WriteRate} = \frac{\text{deltaWriteIOs}}{\text{duration}}$

EMC Symmetrix Host FC Port Performance Metrics

EMC Symmetrix Front-End Controller metrics are stored in the EMCSYMHOSTFCPORTSTATS and MVCS_EMCSYMHOSTFCPTSTATSVW database tables.

Metric data that is rolled up into hourly, daily, weekly, and monthly data points is stored in the corresponding EMCSYMHOSTFCPORTSTATS roll-up tables having a suffix of _H, _D, _W, and _M respectively. For example the daily roll-up performance statistics are stored in the EMCSYMHOSTFCPORTSTATS_D database table.

The following table describes the performance metrics provided by HP Storage Essentials for EMC Symmetrix host Fibre Channel (FC) ports.

Metric	Description	Unit	Formula
Average I/O Size	Average number of Bytes transferred through the Symmetrix host port per I/O operation.	Bytes/s	$AvgIOSize = (\text{deltaKBytesTransferred} \times 1024) / \text{deltaTotalIOs}$
Total Data Rate ¹	Number of Bytes transferred through the Symmetrix host port each second.	Bytes/s	$TotalDataRate = (\text{deltaKBytesTransferred} \times 1024) / \text{duration}$
Total I/O Rate	Number of I/O operations performed each second by the Symmetrix host port. This metric represents activity between the Symmetrix device and the host or SAN device.	Req/s	$TotalIORate = \text{deltaTotalIOs} / \text{duration}$

¹ For this metric the `num_counter_bits` value equals 31.

Transient Values for EMC Symmetrix Performance Metrics

The following transient variables are used to calculate rates for the EMC Symmetrix performance metrics. These variables are pre-fixed with "delta" and use the current–previous logic described in this section.

Logic for Delta Variables

Any variable prefixed with “delta” that is calculated as `current–previous` implies that the variable value was set using the following logic:

```
if current_value < previous_value
delta = (1 + max_counter_value) - previous_value + current_value
else
delta = current - last
```

In this instance, $(1 + \text{max_counter_value})$ is equal to $2^{\text{num_counter_bits}}$. The `num_counter_bits` value is 64 unless stated otherwise.

In all formulas, the term "duration" is defined as follows:

```
duration = deltaNumHalfSecTimeSlices / 2
```

Transient Values for EMC Symmetrix Systems

The following delta variables are used for EMC Symmetrix system performance rates. The Formula column shows the equation used to derive the value of the variable on the left.

Variable Name	Formula
deltaEMCDelayedDFWIOs	currentEMCDelayedDFWIOs - previousEMCDelayedDFWIOs
deltaEMCDeferredWriteIOs	currentEMCDeferredWriteIOs - previousEMCDeferredWriteIOs
deltaEMCKBPrefetched	current_EMCKBPrefetched - previous_EMCKBPrefetched
deltaEMCWriteKBytesFlushed	currentEMCWriteKBytesFlushed - previousEMCWriteKBytesFlushed
deltaKBytesRead*	current_KBytesRead - previous_KBytesRead
deltaKBytesTransferred	deltaKBytesRead + deltaKBytesWritten
deltaKBytesWritten*	current_KBytesWritten - previous_KBytesWritten
deltaReadHitIOs	current_ReadHitIOs - previous_ReadHitIOs
deltaReadHitIOsSeq	current_EMCSequentialReadHitIOs - previous_EMCSequentialReadHitIOs
deltaReadHitIOsTotal	deltaReadHitIOs + deltaReadHitIOsSeq
deltaReadIOs	current_ReadIOs - previous_ReadIOs
deltaReadIOsSeq	current_EMCSequentialReadIOs - previous_EMCSequentialReadIOs
deltaTotalHitIOsRandomAndSeq	deltaReadHitIOsTotal + deltaWriteHitIOsTotal
deltaTotalIOsRandom**	current_TotalIOs - previous_TotalIOs
deltaTotalIOsRandomAndSeq	deltaTotalReadIOsRandomAndSeq + deltaTotalWriteIOsRandomAndSeq
deltaTotalReadIOsRandomAndSeq	deltaReadIOs + deltaReadIOsSeq
deltaTotalWriteIOsRandomAndSeq	deltaWriteIOs + deltaWriteIOsSeq
deltaWriteHitIOs	current_WriteHitIOs - previous_WriteHitIOs
deltaWriteHitIOsSeq	current_EMCSequentialWriteHitIOs - previous_EMCSequentialWriteHitIOs
deltaWriteHitIOsTotal	deltaWriteHitIOs + deltaWriteHitIOsSeq
deltaWriteIOs	current_WriteIOs - previous_WriteIOs
deltaWriteIOsSeq	current_EMCSequentialWriteIOs - previous_EMCSequentialWriteIOs

*The value of num_counter_bits is 63 for calculating this delta value. Blocks are converted to KB by dividing by 2, therefore use 63 counter bits for value when roll-over occurs. See Logic for Delta Variables above.

**The total_io counter holds the sum of random reads plus random writes.

Transient Values for EMC Symmetrix Volumes

The following table lists the delta variables used for EMC Symmetrix volume performance rates. The Formula column shows the equation used to derive the value of the variable on the left.

Variable Name	Formula
deltaKBytesRead*	current_KBytesRead - previous_KBytesRead
deltaKBytesTransferred	deltaKBytesRead + deltaKBytesWritten
deltaKBytesWritten*	current_KBytesWritten - previous_KBytesWritten
deltaReadHitIOs	current_ReadHitIOs - previous_ReadHitIOs
deltaReadHitIOsSeq	current_EMCTSequentialReadHitIOs - previous_EMCTSequentialReadHitIOs
deltaReadHitIOsTotalRandomAndSeq	deltaReadHitIOs + deltaReadHitIOsSeq
deltaReadIOs	current_ReadIOs - previous_ReadIOs
deltaReadIOsSeq	current_EMCTSequentialReadIOs - previous_EMCTSequentialReadIOs
deltaReadMissIOsRandom	deltaReadIOs - deltaReadHitIOs
deltaReadMissIOsSeq	deltaReadIOsSeq - deltaReadHitIOsSeq
deltaReadMissIOsTotalRandomAndSeq	deltaReadMissIOsRandom + deltaReadMissIOsSeq
deltaTotalHitIOsRandomAndSeq	deltaReadHitIOsTotalRandomAndSeq + deltaWriteHitIOsTotalRandomAndSeq
deltaTotalIOsRandom**	current_TotalIOs - previous_TotalIOs
deltaTotalIOsRandomAndSeq	deltaTotalReadIOsRandomAndSeq + deltaTotalWriteIOsRandomAndSeq
deltaTotalReadIOsRandomAndSeq	deltaReadIOs + deltaReadIOsSeq
deltaTotalWriteIOsRandomAndSeq	deltaWriteIOs + deltaWriteIOsSeq
deltaWriteHitIOs	current_WriteHitIOs - previous_WriteHitIOs

Variable Name	Formula
deltaWriteHitIOsSeq	current_EMCTotalHitIOs - previous_EMCTotalHitIOs
deltaWriteHitIOsTotalRandomAndSeq	deltaWriteHitIOs + deltaWriteHitIOsSeq
deltaWriteIOs	current_WriteIOs - previous_WriteIOs
deltaWriteIOsSeq	current_EMCTotalWriteIOs - previous_EMCTotalWriteIOs
deltaWriteMissIOsRandom	deltaWriteIOs - deltaWriteHitIOs
deltaWriteMissIOsSeq	deltaWriteIOsSeq - deltaWriteHitIOsSeq
deltaWriteMissIOsTotalRandomAndSeq	deltaWriteMissIOsRandom - deltaWriteMissIOsSeq

*The value of num_counter_bits is 63 for calculating this delta value. Blocks are converted to KB by dividing by 2; therefore, use 63 counter bits for value when roll-over occurs. See [Logic for Delta Variables on page 134](#).

**The total_io counter holds the sum of random reads plus random writes.

Transient Values for EMC Symmetrix Front-end Controllers

The delta variables are used for EMC Symmetrix front-end controller performance rates. The Formula column shows the equation used to derive the value of the variable on the left.

Variable Name	Formula
deltaKBytesTransferred	current_KBytesTransferred - previous_KBytesTransferred
deltaReadIOs	current_ReadIOs - previous_ReadIOs
deltaWriteIOs	current_WriteIOs - previous_WriteIOs
deltaTotalIOs	current_TotalIOs - previous_TotalIOs
deltaEMCTotalHitIOs	current_EMCTotalHitIOs - previous_EMCTotalHitIOs
deltaEMCSlotCollisions	current_EMCSlotCollisions - previous_EMCSlotCollisions
deltaEMCSystemFlushPendingEvents	current_EMCSystemFlushPendingEvents - previous_EMCSystemFlushPendingEvents
deltaEMCDeviceFlushPendingEvents	current_EMCDeviceFlushPendingEvents - previous_EMCDeviceFlushPendingEvents

Transient Values for EMC Symmetrix Host FC Ports

The following delta variables are used for EMC Symmetrix host Fibre Channel (FC) port performance rates. The Formula column shows the equation used to derive the value of the variable on the left. The num_counter_bits for calculating the KBytesTransferred for host FC ports is 31.

Variable Name	Formula
deltaKBytesTransferred	current_KBytesTransferred – previous_KBytesTransferred
deltaTotalIOs	current_TotalIOs – previous_TotalIOs

Managing HP StorageWorks EVA Performance for Arrays

This section describes the HP Storage Essentials performance management metrics for the HP StorageWorks Enterprise Virtual Arrays (EVAs) running in your storage environment. The following topics include techniques and best practices for using Performance Manager with the SRM Performance Pack software to manage your EVA arrays.

- [Finding EVA Bottlenecks Using Performance Metrics below](#)
- [Monitoring EVA Performance on the facing page](#)
- [Performance Management EVA Metrics on page 140](#)

Finding EVA Bottlenecks Using Performance Metrics

HP Storage Essentials metrics for the HP EVA array system can help you find bottlenecks by enabling you to see the performance throughput (I/O) rates and response times (latency). Bottlenecks are characterized by unexpected latencies, which may be due to heavy data workloads and resource contention. Identifying bottlenecks early can be an ongoing challenge for storage administrators. The following sections describe a few of the indicators and the HP Storage Essentials metrics that provide the information you need to manage the performance of your array.

I/O Per Second

To determine transmission loads (for managing load balancing and increased bandwidth needs), determine the total I/O per seconds (IOPS) over a time interval. An IOPS is $1000 \text{ ms} / (\text{CommandTime} + \text{SeekTime} + \text{SettledownTime} + \text{DataTransferTime})$.

Read Latency

Chart the total array MB/s, total array requests, and FP port latencies as a baseline for normal performance of your EVA array. When looking for bottlenecks, check the average time it takes to complete a read request (from initiation to information receipt) from the physical disks (for all volumes in the EVA array) by looking at the Total Volume Avg Read Miss Latency (Sec). If the latency rate increases (rule of thumb: 30 ms or more), and there is no correlated increase in IOPS or MB/s, you should investigate further. Although it is acceptable for certain environments to run at >30 ms, you should base what you consider to be an acceptable latency rate on your specific environment design and workload profile.

Write Latency

Check the average time it takes to complete a write request (from initiation to receipt of write completion) for all volumes by selecting the Total Volume Avg Write Latency (sec). If the latency is 5 ms or more, this indicates a potential bottleneck in the system and you should investigate further.

Write latency may also increase due to increased write activity to the controller. As the total write activity to the controller increases, flush rates will become higher and steadier. Check your flush rates using the Total Volume Flush Data Rate (Bytes/Sec) metric. When flush rates are consistently equal to the aggregate write rates, the maximum write rate to that controller has probably been achieved. At that point, write latencies might increase significantly.

Remote data replication using HP Continuous Access EVA might also be the cause of write latencies, especially if Sync replication is enabled. Hosts which access their LUNs through the replication ports can also be negatively affected.

Queue Depth

Disk queue depth can indicate whether the workload asked of the physical disk is greater than the intended design. In such cases, you might consider adding more physical disks.

FP queue depths can be caused by array latency, as well as not balancing hosts for fair access. For example, a host that has a queue depth of 128 can starve a host that has a queue depth of 8.

Monitoring EVA Performance

To optimize performance of your HP Enterprise Virtual Array (EVA), follow these best practices:

- Establish a performance baseline. Collect metrics that represent normal performance conditions for your system. You can then use this baseline to compare future metrics and analyze issues.
- Track IOPS rates and throughput.
- Track error rates, average queue depth and response times.
- To maximize single array performance: Fill the EVA with as many disk drives as possible.

- To maximize application performance, you should perform a performance analysis before adding any significant application load to an existing EVA.
- Configure as few Disk Groups as possible.
- When creating Disk Groups, as a best practice use equal capacity disks.
- Where different physical disk sizes are used in a Disk Group, set up workloads so that the larger physical disks house a proportionately higher amount of the data and that workloads do not exceed design capacities.
- To provide the highest performance, use 15K-RPM disks. This is a guideline for performance management and not a remedy for availability.
- For high I/O write workloads, to reduce the utilization of the mirror port, migrate the high I/O intensive Vdisk from VRAID 5 to VRAID 1.
- The only implication of having different physical disk sizes in a disk group is that the larger physical disks will house a proportionally higher amount of the data. If physical disks are not asked to do more than they are designed to do, then no harm is done.
- Always leave read caching enabled.
- Attempt to balance the workload demand evenly across both controllers on an EVA.
- Although both controllers in an EVA can access all physical disks, a Vdisk is assigned to be managed by only one controller at a time. Because the Vdisks can be accessed by either controller, all FP ports can be used. ALUA aware drivers will use the owning controller's FP ports for the read data and all available ports for the write access.
- Enable a common subset of metrics for standard collection. These might include CPU utilization (%), total I/O rate (MB/s), total requests per second, and FP read and write latencies. If you create a baseline that includes these metrics, you can then quickly determine 1) if the latency increases and 2) if there is an increased workload that is responsible for it. Consider placing these metrics into a basic EVA template.
- Enabling all collectors, or more than are necessary, can result in decreased Command View EVA responsiveness and intermittent performance data collection failure.

Performance Management EVA Metrics

HP Storage Essentials enables deep visibility into the performance levels of HP StorageWorks Enterprise Virtual Arrays (EVA) storage arrays by providing metrics for the following EVA storage infrastructure components:

- [EVA Storage System Metrics on the facing page](#)
- [EVA Controller Metrics on page 144](#)
- [EVA Physical Disk Metrics on page 146](#)
- [EVA Host FC Port Metrics on page 148](#)
- [EVA Volume Metrics on page 150](#)

Click the metric set above that interests you to view a table that describes the statistics collected for that array component.

Metric Table Conventions

You will find the following conventions used in EVA performance statistics tables listed above:

- To simplify formulas, the value Time represents the difference in seconds between the most recent two StatisticTime values returned from the SMI-S provider.
- Items marked with an asterisk (*) indicate they are duplicates of SMI-S statistics.
- All statistic names used in shown formulas are listed in a separate "raw statistics" table for each array component. See references provided in the introduction to each table.

EVA Storage System Metrics

The metrics track performance for the HP StorageWorks Enterprise Virtual Array (EVA). Storage system metrics are stored in the HPEVASTORAGESYSTEMSTATS database table. For descriptions of statistics in formulas, see [EVA Storage System Raw Statistics on page 143](#) and the [EVA Volume Raw Statistics on page 152](#).

Metric	Description	Units	Formula
Total Data Rate*	Rate data can be transmitted between devices for the entire storage system.	Bytes/s	$(\Delta \text{KBytesTransferred} \times 1024) / \Delta \text{Time}$
Total I/O Rate*	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the entire storage system	Req/s	$\Delta \text{TotalIOs} / \Delta \text{Time}$
Total Volume Avg Read Hit Latency	Average time to complete a read request (from initiation to information receipt) from the array cache memory for all volumes in the array.	Sec	$(\Delta \text{ReadHitLatency} / 1000) / \Delta \text{ReadHitIOs}$
Total Volume Avg Read Miss Latency	Average time to complete a read request (from initiation to information receipt) from the physical disks for all volumes	Sec	$(\Delta \text{ReadMissLatency} / 1000) / \Delta \text{ReadMissIOs}$
Total Volume Avg Read Size	Average data read size for all volumes	Bytes	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{ReadIOs}$

Metric	Description	Units	Formula
Total Volume Avg Write Latency	Average time to complete a write request (from initiation to receipt of write completion) for all volumes	Sec	$(\Delta \text{KBytes Transferred} \times 1024) / \Delta \text{Time}$
Total Volume Avg Write Size	Average write size for all volumes	Bytes	$(\Delta \text{KBytes Written} \times 1024) / \Delta \text{WritelOs}$
Total Volume Data Rate	Rate data can be transmitted between devices for all volumes	Bytes/s	$(\Delta \text{KBytes Transferred} \times 1024) / \Delta \text{Time}$
Total Volume Flush Data Rate	Rate at which data is written to physical disks in array	Bytes/s	$(\Delta \text{FlushKBytes} \times 1024) / \Delta \text{Time}$
Total Volume Flush Rate	Aggregate of all flush counters: mirror flush, cache flush, host writes to snapshots and snapclones	Bytes/s	$\Delta \text{FlushRequests} / \Delta \text{Time}$
Total Volume I/O Rate	Average number of I/O operations per second for both sequential and non-sequential read and write operations for all volumes	Req/s	$\Delta \text{TotalIOs} / \Delta \text{Time}$
Total Volume Mirror Data Rate	Rate at which data travels across the mirror port to complete read and write requests to all virtual disks	Bytes/s	$(\Delta \text{MirrorKBytes} \times 1024) / \Delta \text{Time}$
Total Volume Pct Read IOs	Percentage (%) of read I/O operations per second for both sequential and non-sequential reads for all volumes	%	$100 \times (\Delta \text{ReadIOs} / \Delta \text{TotalIOs})$
Total Volume Pct Write IOs	Percentage (%) of write I/O operations per second for both sequential and non-sequential writes for all volumes	%	$100 \times (\Delta \text{WritelOs} / \Delta \text{TotalIOs})$
Total Volume Prefetch Data Rate	Rate data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected	Bytes/s	$(\Delta \text{PrefetchKBytes} \times 1024) / \Delta \text{Time}$

Metric	Description	Units	Formula
Total Volume Read Data Rate	Rate data is read from the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$
Total Volume Read Hit Data Rate	Rate at which data is read from the array cache memory because of read hit requests	Bytes/s	$(\Delta \text{ReadHitKBytes} \times 1024) / \Delta \text{Time}$
Total Volume Read Hit Rate	Number of read requests per second completed from the array cache memory	Req/s	$\Delta \text{ReadHitIOs} / \Delta \text{Time}$
Total Volume Read Miss Data Rate	Rate at which data is read from physical disks because the data was not present in the array cache memory	Bytes/s	$(\Delta \text{ReadMissKBytes} \times 1024) / \Delta \text{Time}$
Total Volume Read Miss Rate	Number of read requests (per second) that were not available from cache memory and therefore were completed from the physical disks instead	Req/s	$\Delta \text{ReadMissRequests} / \Delta \text{Time}$
Total Volume Read Rate	Number of read requests per second completed from a virtual disk that were sent to all hosts	Req/s	$\Delta \text{ReadIOs} / \Delta \text{StatisticTime}$
Total Volume Write Data Rate	Rate at which data is written to the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	$\Delta \text{KBytesWritten} \times 1024 / \Delta \text{Time}$
Total Volume Write	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	$\Delta \text{WriteIOs} / \Delta \text{Time}$

*This metric is a duplicate of an SMI-S statistic.

EVA Storage System Raw Statistics

HP Storage Essentials collects the following storage system array statistics and stores them in the HPEVASTORAGESYSTEMSTATS database table.

Raw Statistic	Type	Units
Total I/Os*	uint64	Count

Raw Statistic	Type	Units
Kbytes Transferred*	uint64	KB

*This metric is a duplicate of an SMI-S statistic.

EVA Controller Metrics

For EVA storage systems in which HP Storage Essentials supports performance data collection, the following performance metrics are monitored for EVA controllers. See the [EVA Controller Raw Statistics on the facing page](#) for descriptions of the statistics shown in the formulas.

Metric	Description	Units	Formula
Average Read Size*	Amount of data read (per second) from physical disk	Bytes	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{ReadIOs}$
Average Write Size*	Amount of data written (per second) to physical disks	Bytes	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{WriteIOs}$
CPU Percent**	Percentage (%) of time that the central processing unit on the controller is active. A completely idle controller shows 0%. A controller saturated with activity shows 100%.	%	$100 \times (\Delta \text{CpuBusyCounter} / \Delta \text{StatisticsTime})$
Data Transfer Percent**	Similar to % Processor Time except that it does not include time for internal processes not related to host-initiated data transfers	%	$100 \times (\Delta \text{DataTxCounter} / \Delta \text{StatisticsTime})$
Percent Reads*	Percentage (%) of CPU time dedicated to reads	%	$100 \times (\Delta \text{ReadIOs} / \Delta \text{TotalIOs})$
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	$100 \times (\Delta \text{WriteIOs} / \Delta \text{TotalIOs})$
Read Data Rate*	Rate at which data is read from the controller by all disks	Bytes/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$
Read Latency	Average time it takes to complete a read request (from initiation to receipt of write completion) through the controller	Sec	$(\Delta \text{ReadLatency} / 1000) / \Delta \text{ReadIOs}$
Reads*	Rate at which data is read from each host port	Req/s	$\Delta \text{ReadIOs} / \Delta \text{Time}$

Metric	Description	Units	Formula
Total Data Rate*	Rate at which data can be transmitted between devices for the controller	Bytes/s	$(\Delta \text{KBytesTransferred} \times 1024) / \Delta \text{Time}$
Total I/O Rate*	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the controller	Req/s	$\Delta \text{TotalIOs} / \Delta \text{Time}$
Write Data Rate*	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{Time}$
Write Latency	Average time it takes to complete a write request (from initiation to receipt of write completion)	Sec	$(\Delta \text{WriteLatency} / 1000) / \Delta \text{WriteIOs}$
Writes*	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	$\Delta \text{WriteIOs} / \Delta \text{Time}$

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

EVA Controller Raw Statistics

The following values are collected by HP Storage Essentials and, except for the CpuCounter and DataTxCounter, are aggregated by the provider from the underlying port metrics. All values are stored in the HPEVACONTROLLERSTATS database table.

Raw Statistic	Type	Units
CpuCounter	uint64	Ms
DataTxCounter	uint64	Ms
Total I/Os*	uint64	Count
Kbytes Transferred*	uint64	KB
Read I/Os*	uint64	Count
KBytesRead*	uint64	KB
Write I/Os*	uint64	Count
KBytesWritten*	uint64	KB
Read Latency	uint32	ms
Write Latency	uint32	ms

*This metric is a duplicate of an SMI-S statistic.

EVA Physical Disk Metrics

The following metrics track performance statistics for EVA physical disks. These values are stored in the HPEVADISKSTATS database table. For descriptions of the statistics shown in formulas, see [EVA Physical Disk Raw Statistics on the facing page](#).

Metric	Description	Units	Formula
Average Read Size*	Amount of data read from physical disk	Bytes	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{ReadIOs}$
Average Write Size*	Amount of data written to physical disk	Bytes	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{WriteIOs}$
Drive Latency metric**	Average time to complete read/write requests from the physical disk drive	Sec	$(\Delta \text{DriveLatency} / 1000) / \Delta \text{TotalIOs}$
Percent Reads*	Percentage (%) of CPU time dedicated to reads	%	$100 \times (\Delta \text{ReadIOs} / \Delta \text{TotalIOs})$
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	$100 \times (\Delta \text{WriteIOs} / \Delta \text{TotalIOs})$
Queue**	Average number of outstanding requests against the physical disk	Req	$\Delta \text{DriveQueueDepth} / \Delta \text{Statistic Time}$
Read Data Rate***	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$
Read Latency**	Average time to complete a read request (from initiation to information receipt) from the array volume	Sec	$(\Delta \text{ReadLatency} / 1000) / \Delta \text{ReadIOs}$
Reads***	Rate at which data is read from each host port	Req/s	$\Delta \text{ReadIOs} / \Delta \text{Time}$
Total Data Rate*	Rate at which data can be transmitted between devices for the host port	Bytes/s	$(\Delta \text{KBytesTransferred} \times 1024) / \Delta \text{Time}$
Total I/O Rate*	Average number of I/O operations (requests per second) for both sequential and non-sequential reads and writes for the host port	Req/s	$\Delta \text{TotalIOs} / \Delta \text{Time}$

Metric	Description	Units	Formula
Write Data Rate***	Rate at which data is written to the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{Time}$
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	Sec	$(\Delta \text{WriteLatency} / 1000) / \Delta \text{WriteIOs}$
Writes***	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

***This metric is collected by both HP Storage Essentials and EVAPerf and is a duplicate of an SMI-S statistic.

EVA Physical Disk Raw Statistics

The following values are collected by HP Storage Essentials and stored in the HPEVADISKSTATS database table.

Raw Statistic	Type	Units
TotalIOs*	uint64	Count
KBytesTransferred*	uint64	KB
ReadIOs*	uint64	Count
KBytesRead*	uint64	KB
WriteIOs*	uint64	count
KBytesWritten*	uint64	KB
DriveQueueDepth	uint64	Count
DriveLatency	uint64	ms
ReadLatency	uint64	ms
WriteLatency	uint64	ms

*This metric is a duplicate of an SMI-S statistic.

EVA Host FC Port Metrics

For EVA storage systems in which HP Storage Essentials supports performance data collection, the following items are monitored for EVA Host FC Ports. See [EVA Host FC Port Raw Statistics](#) on page 150 for descriptions of the statistics shown in the formulas.

Metric	Description	Units	Formula
Average Read Size	Amount of data read (per second) from physical disks	Bytes	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{ReadIOs}$
Average Write Size	Amount of data written (per second) to physical disks	Bytes	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{WriteIOs}$
Bad CRC	Number of bad CRC errors. Indicates that the Cyclic Redundancy Check (CRC) which compares a data stream against a stored checksum, has found the data stream changed and therefore no longer reliable. Use to help the transmitter detect errors in the frame that are caused by bad writes, bad media, damaged links/hardware, excessive link errors, and transfer rates.	Count	–
Bad Receive Characters	Number of bad receive characters in the bit stream. Use to determine the number of bad frames associated with the Bad CRC metric above.	Count	–
Discard Frames	Number of frames discarded due to Bad CRCs. Frames are the basic unit of communication between two N_ports, and are composed of a starting delimiter, header, payload, CRC, and end delimiter.	Count	–
Link Fail	Number of link failures. Use to find issues with the fiber optic cable or transceiver or the SAN infrastructure.	Count	–
Loss of Signal	Number of times the receiver reports loss of signal. Indicator that fiber optic signal no longer exists. Use to assist in troubleshooting signal loss.	Count	–

Metric	Description	Units	Formula
Loss of Sync	Number of times the receiver logic reports loss of sync has timed-out. Use to determine the number of times an intermittent loss of synchronization in communication signals was received by an enclosure connected to a Fibre Channel (FC) loop.	count	–
Percent Reads	Percentage (%) of CPU time dedicated to reads.	%	$100 \times (\Delta \text{ReadIOs} / \Delta \text{TotalIOs})$
Percent Writes	Percentage (%) of CPU time dedicated to writes	%	$100 \times (\Delta \text{WriteIOs} / \Delta \text{TotalIOs})$
Protocol Error	Number of errors in the protocol between the channel and the control unit. Use to differentiate between protocol errors and link errors.	count	–
Queue	Average number of outstanding host requests against all virtual disks accessed through this host port	Count	$\Delta \text{QDepth} / \Delta \text{Time}$
Read Data Rate	Rate at which data is read from the controller by all disks	Bytes/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$
Read Latency**	Average time to complete a read request (from initiation to receipt of write completion) through the controller	Sec	$(\Delta \text{ReadLatency} / 1000) / \Delta \text{ReadIOs}$
Reads**	Rate at which data is read from each host port	Req/s	$\Delta \text{ReadIOs} / \Delta \text{Time}$
Receive Abnormal End of Frame	Number of times a bad frame was detected during data transmission	Count	–
Total Data Rate*	Rate in which data can be transmitted between devices for the host port	Bytes/s	$(\Delta \text{KBytesTransferred} \times 1024) / \Delta \text{Time}$
Total I/O Rate*	Average number of I/O operations as requests per second for both sequential and non-sequential reads and writes for the host port	Req/s	$\Delta \text{TotalIOs} / \Delta \text{Time}$
Write Data Rate**	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{Time}$

Metric	Description	Units	Formula
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	Sec	$(\Delta \text{WriteLatency} / 1000) / \Delta \text{WritelOs}$
Writes**	Number of write requests per second completed to a virtual disk that were received from all hosts	Req/s	$\Delta \text{WritelOs} / \Delta \text{Time}$

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and EVAPerf performance statistic.

EVA Host FC Port Raw Statistics

The following values are collected by HP Storage Essentials and stored in the HPEVAHOSTFCPORTSTATS database table.

Raw Statistic	Type	Units
TotalIOs*	uint64	Count
KbytesTransferred*	uint64	KB
ReadIOs	uint64	count
KBytesRead	uint64	KB
ReadLatency	uint32	ms
WritelOs	uint64	Count
KBytesWritten	uint64	KB
WriteLatency	uint32	ms
FlushRequests	uint64	Count
QDepth	uint64	Count

*This metric is a duplicate of an SMI-S statistic.

EVA Volume Metrics

The following metrics track performance for HP StorageWorks Enterprise Virtual Array (EVA) volumes. See [EVA Volume Raw Statistics](#) on page 152 for descriptions of the statistics shown in the formulas.

Metric	Description	Units	Formula
Average Read Size*	Amount of data read (per second) from physical disks	Bytes	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{ReadIOs}$

Metric	Description	Units	Formula
Average Write Size*	Amount of data written (per second) to physical disks	Bytes	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{WriteIOs}$
Flush Data Rate	Rate at which data is written to a physical disk for the associated virtual disk	Bytes/s	$(\Delta \text{FlushKBytes} \times 1024) / \Delta \text{Time}$
Flushes	Number of cache flush requests per second	Req/s	$\Delta \text{FlushRequests} / \Delta \text{Time}$
Mirror Data Rate	Rate at which data travels across the mirror port to complete read and write requests for the associated virtual disk	Bytes/s	$(\Delta \text{MirrorKBytes} \times 1024) / \Delta \text{Time}$
Percent Reads*	Percentage (%) of CPU time dedicated to reads	%	$100 \times (\Delta \text{ReadIOs} / \Delta \text{TotalIOs})$
Percent Writes*	Percentage (%) of CPU time dedicated to writes	%	$100 \times (\Delta \text{WriteIOs} / \Delta \text{TotalIOs})$
Prefetch Data Rate	Rate at which data is read from the physical disk to cache in anticipation of subsequent reads when a sequential stream is detected	Bytes/s	$(\Delta \text{PrefetchKBytes} \times 1024) / \Delta \text{Time}$
Read Data Rate*	Rate at which data is read from the virtual disk by all hosts, including transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesRead} \times 1024) / \Delta \text{Time}$
Read Hit Data Rate**	Rate at which data is read from the array cache memory because of read hit requests	Bytes/s	$(\Delta \text{ReadHitKBytes} \times 1024) / \Delta \text{Time}$
Read Hit Latency**	Average time to complete a read request (from initiation to information receipt) from the array volume	Sec	$(\Delta \text{ReadHitLatency} / 1000) / \Delta \text{ReadHitIOs}$
Read Hits***	Number of read requests (per second) completed from the array cache memory	Req/s	$\Delta \text{ReadHitIOs} / \Delta \text{Time}$
Read Miss Data Rate**	Rate at which data is read from physical disks because the data was not present in the array cache memory	Bytes/s	$(\Delta \text{ReadMissKBytes} \times 1024) / \Delta \text{Time}$

Metric	Description	Units	Formula
Read Miss Latency**	Average time it takes to complete a read request (from initiation to information receipt) from the physical disks for all volumes	Sec	$(\Delta \text{ReadMissLatency} / 1000) / \Delta \text{ReadMissIOs}$
Read Misses**	Number of read requests (per second) that failed to complete from the array cache memory and were completed from physical disks instead	Req/s	$\Delta \text{ReadMissRequests} / \Delta \text{Time}$
Reads*	Rate at which data is read from each host host port	Req/s	$\Delta \text{ReadIOs} / \Delta \text{Time}$
Total Data Rate*	Rate at which data can be transmitted between devices for the host port	Bytes/s	$(\Delta \text{KBytesTransferred} \times 1024) / \Delta \text{Time}$
Total I/O Rate***	Average number of I/O operations in requests per second for both sequential and non-sequential reads and writes for the hostport	Req /s	$\Delta \text{TotalIOs} / \Delta \text{Time}$
Write Data Rate***	Rate at which data is written to the virtual disk by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{KBytesWritten} \times 1024) / \Delta \text{Time}$
Write Latency**	Average time to complete a write request (from initiation to receipt of write completion)	Sec	$(\Delta \text{WriteLatency} / 1000) / \Delta \text{WriteIOs}$
Writes***	Number of write requests per second completed to a virtual disk that were received from all hosts	Req /s	$\Delta \text{WriteIOs} / \Delta \text{Time}$

*This metric is a duplicate of an SMI-S statistic.

**This metric is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

***This metric is a duplicate of an SMI-S statistic and is collected as both an HP Storage Essentials and an EVAPerf performance statistic.

EVA Volume Raw Statistics

The following values are collected by HP Storage Essentials and stored in the HPEVAVOLUMESTATS database table.

Raw Statistic	Type	Units
Total I/Os*	uint64	Count
Kbytes Transferred*	uint64	KB

Raw Statistic	Type	Units
Read I/Os*	uint64	Count
Read Hit I/Os*	uint64	Count
KBytesRead*	uint64	KB
Write I/Os*	uint64	Count
KBytesWritten*	uint64	KB
Read Hit KBytes	uint32	KB
Read Hit Latency	uint32	Ms
Read Miss Requests	uint64	Count
Read Miss Kbytes	uint64	KB
Read Miss Latency	uint32	ms
Write Latency	uint32	ms
Flush Requests	uint64	Count
Flush Kbytes	uint64	KB
Mirror KBytes	uint64	KB
Prefetch KBytes	uint64	KB

*This metric is a duplicate of an SMI-S statistic.

Managing XP and HDS Array Performance

HP Storage Essentials provides over a hundred performance management metrics for the XP and HDS arrays in your storage environment. A basic set of metrics is available with installation of the HP Storage Essentials Performance Manager, and a larger set installs with the HP Storage Essentials Performance Pack Enterprise software. Metrics are gathered by data collectors, and available for inclusion in reports and performance policies.

For more information see the following topics:

- [Finding XP and HDS Array Bottlenecks Using Performance Metrics on next page](#)
- [Monitoring XP Baseline Performance on page 155](#)
- [Optimizing Your XP and HDS Array Performance on page 156](#)
- [Viewing External Storage and Continuous Access Link Performance on page 157](#)
- [Viewing Pool and Group Performance on page 157](#)

XP Array Performance Tools

HP Storage Essentials is the only performance tool for XP arrays that displays a full SAN topology, showing you all devices along the connection path from application to disk volume. Using the Performance Manager you can generate performance charts for each component in the path for which data has been collected and generate reports which show performance statistics.

To fully use Performance Manager capabilities, use the tool in collaboration with other HP Storage Essentials components. Use the Discovery features of HP Storage Essentials to discover each element in your SAN environment. You can then display these discovered elements in a SAN topology path that clearly shows the connections from hosts to arrays. Select the performance metrics you want to monitor using Performance Manager and when you are ready, generate reports for hosts, switches, and arrays using the Report Optimizer.

Depending on your IT practices, you could also be using one of the following HP StorageWorks XP performance tools:

- **XP Command View AE** enables you to configure common tasks for hosts, switches, and arrays; manages multiple XP disk arrays; uses the Remote Web Console to manage one XP disk array at a time.
- **XP Performance Advisor** enables you to monitor hosts, applications running on the hosts, and switches and arrays connected to the hosts; focuses on performance management of XP disk arrays; provides reporting, performance estimates, and a history database.

You will find certain instructions for the Remote Web Console feature of XP Command View in this guide. For the most part, instructions are specific to HP Storage Essentials Performance Manager.

Finding XP and HDS Array Bottlenecks Using Performance Metrics

The following sections describe considerations when examining performance metrics for HP XP and HDS arrays.

When determining the cause of a performance bottleneck, it is important to observe how an array utilizes its memory and cache. Selecting an XP or HDS array element allows you to view shared memory and cache usage statistics for the array. The cache statistics are itemized per Cache Logical Partition (CLPR).

To troubleshoot and isolate factors that could affect your application performance:

1. Navigate to the host running your application and display performance metrics by clicking it. You can click on either the application name in the path tree or application icon in the topology map.
2. Check the host processor utilization and physical/virtual memory consumption. If your processor is running near 100% and virtual memory consumption is greater than physical memory, the problem is most likely with the host not having enough resources to process the I/O for the application.
3. Select the application path where the application is running, and expand its members.

4. Check the LDEV read/write response time. The LDEV may have an extremely high response time (above 20ms) or zero response. When the response time is higher than expected, there may be a HP RESTRICTED resource contention problem within the array. In the case of zero response, there could be a link failure.
5. Chart the HBA, switches, and array port Bytes Transmitted/Received metrics on the same graph to verify that traffic is passing through to the array without problems. If you see large differences between the HBA, switch and array port, the problem could be a bad SFP or fiber link.
6. Chart the array port and LDEV Total Data Rate. The array port should have at least the same data rate as the LDEV. If the array port data rate is lower than LDEV that could mean you have too much write pending in cache.

Best Practices

As I/O is running, view the Cache Writes Pending for the CLPR (to which your host/application LUN(s) belong) in real time. To do this, select Real Time from the Period drop down. Write Pending indicates the amount of I/O waiting to be de-staged to disk. As Cache Writes Pending approaches 70 percent, typically the array begins to hold back I/O, which in turn affects your host response time. Some of the causes for excessive write pending could be too much I/O from other devices in the CLPR, replication processing, and external storage (unless you have cache write-through mode enabled).

Monitoring XP Baseline Performance

The following are recommendations for monitoring your XP baseline performance monitoring. These are useful as guides for general thresholds in HP Storage Essentials Policy Manager as well.

See [About Policy Manager on page 67](#) for more information about setting performance policies and alerts.

Front-End Port Usage (CPU Load)

HP recommends an average load under 50 percent. At 50 percent or more, the I/O response time gets doubled. Set the Policy Manager to send Info-Msg at 60% and Alert at 70%.

Back-End ACP Pair Usage (CPU Load)

Recommendations for back-end ACP pair usage are the same as front-end port usage. Set Policy Manager to send Info-Msg at 60% and Alert at 70%. This parameter is not available in the HP Storage Essentials XP Perf Policy. Also for 99 percent of XP customers, back-end CPU is not a performance issue.

Cache Utilization

HP recommends a threshold below 70 percent. If this value is greater than 70 percent, you should check the cache characteristics of your XP. Set Policy Manager to send Alert if greater than 70 percent.

Write Pending Rate

The write pending rate is very important. The value of the Percent Write Pending Data (%) is normally stable around 30%. If the percentage reaches 50 percent, you want to be alerted immediately. Set your Policy Manager to send Alerts if it is >40% to >50%.

Write Response Time

The write response time indicated by the Write Response Time (Sec) for XP array volumes should be below 10ms (also in sync-CA environment). The baseline for this parameter is important. Set your Policy to send Alert if the response time is greater than 10 milliseconds (ms).

Array Group Utilization

HP recommends that you maintain an average load of under 50 percent. Set Policy Manager to send Info-Msg at >60% and Alert at >70%.

LDEV Utilization

The probability that one logical device (LDEV) has more than 70 percent of the load of an array group is very low. If this situation is true, you need to analyze manually. The overall significance of this metric value is low, as it is more important to look into the overall array group utilization.

Optimizing Your XP and HDS Array Performance

RAID group utilization will let you know what kind of I/O is affecting performance. Some disk types and RAID levels work faster with Sequential versus Random I/O, and/or Reads versus Writes. The storage admin should optimize the array to satisfy the customer's Service Level Agreement (SLA) while minimizing the cost of the array. The user will also need to know the origin of the I/O (that is, the volumes in use), the data path, and array resource usage.

To optimize performance on your HP XP array:

1. Chart RAID group utilization metrics and check to see what kind of I/O is most prevalent.
2. Determine which volumes are connected to the RAID group and chart their I/O type and response times.
3. Verify that the response times meet your Service Level Agreement.
4. Determine which volumes are candidates for optimization based on their I/O type and response times.
5. Check the ports to which the volumes are mapped and graph their throughput.
6. Verify that the port is not a bottleneck based on its maximum rated throughput versus its actual throughput.
7. Determine the host to which the volumes are mapped, and make sure its resources are not excessively utilized.
8. Determine which applications are using the volumes in the RAID group.

9. Locate new volumes to use as candidates for migration.
10. Check the candidate's current RAID Group utilization to verify that resources are available.

Viewing External Storage and Continuous Access Link Performance

Currently, you must use the Remote Web Console (RWC) to determine which ports are used for Continuous Access and/or External Storage connectivity. Plot the ports you suspect may have a failure, and the switches used to interconnect the arrays. The port charts will show you how well the port is being utilized, and the switch chart will show any link failures between the arrays.

To diagnose your links:

1. Navigate to front-end controllers and chart the IOPS from initiator or target Port. If you want to see performance from all arrays simultaneously you will need to launch another browser session for each array.
2. Chart the back-end processor (XP DKA) metrics for the external arrays.
3. Chart cache write pending for all arrays. Your XP or HDS disk array should have cache passthrough enabled for external storage; therefore, you will see minimum impact on the CLPR metrics assigned to the external LUNs.
4. Monitor the switches to determine the link failures. You will need to launch another browser session to view switches and arrays simultaneously.
5. Assuming there are no link failures, determine which external array connections are excessively utilized by checking the cache writes pending versus link IOPS. If cache writes pending on the external array is low and link IOPS are high, you could add another external link path to improve performance.

Viewing Pool and Group Performance

THP/Snapshot/Continuous Access Journal Pool Performance

XP Pools are composed of LDEVs from RAID Groups. To fully analyze your Pool, you must determine the RAID Groups to which the LDEVs belong. Best practice for creating a Pool is to dedicate the entire RAID Group to that Pool, which makes the performance more predictable, because no external source of IO and disk utilization will affect the Pool.

To diagnose your Pool:

1. Select Storage Pools and right click on the storage pool in question, then choose **Show Properties**. Select the pool volumes to determine the RAID Groups they belong to.
2. Chart the V-VOL performance from the storage pool.
3. Navigate to Array Groups and find the RAID Groups to which the Pool LDEVs belong. Chart the RAID Group utilization.
4. If your RAID Group utilization is high, focus your search down to I/O type. Chart Percent Reads/Writes for the RAID Group(s) and V-VOL(s) connected to the Pool.

5. Look for patterns in I/O type compared to V-VOL usage to determine which V-VOLs are most responsible for RAID Group utilization.

Viewing Parity Group Statistics for XP and HDS Arrays

Storage Essentials provides performance statistics at the Parity Group level for XP and HDS arrays. It does not provide statistics for each Logical Array Group (also called Logical RAID Groups) within the Parity Group.

Performance statistics for the Parity Group display in the information fields for the first Logical Array Group within that Parity Group. To view the performance statistics for a Parity Group, therefore, look at the first Logical Array Group statistics. The subsequent non-first Logical Array Groups show statistics of zero, which is intentional to avoid double-counting along the entire Parity Group when statistics are shown against the first Logical Array Group.

Performance Management XP and HDS Array Metrics

Performance Manager enables you to drill down deep into an XP or HDS array and view all component performance metrics. These include cache and shared memory statistics, CHA processor and port statistics for front-end controllers, and DKA processor statistics for back-end controllers, volume statistics for storage pools, and RAID group statistics for array groups.

HP Storage Essentials provides performance metrics for the following storage infrastructure components for HP XP and HDS arrays. Click a metric set that interests you in the following list to see the statistics collected for that component.

- [XP and HDS Array Storage System Metrics](#) below
- [XP and HDS Array Group Metrics](#) on page 160
- [XP and HDS Array Volume Metrics](#) on page 162
- [XP and HDS Front-end Controller CLPR Metrics](#) on page 164
- [XP and HDS Front-end Controller \(CHA\) Metrics](#) on page 165
- [XP and HDS MPB Controller Metrics](#) on page 166

Back-end controller DKA metrics are provided for HP XP arrays only. See [XP Back-end Controller \(DKA\) Metric](#) on page 165.

XP and HDS Array Storage System Metrics

Selecting an XP or HDS array element in Performance Manager allows you to view shared memory and cache usage statistics. The cache statistics are itemized per Cache Logical Partition (CLPR). It is important for the user to observe how the array is using its memory and cache in order to help determine the root cause of a performance bottleneck.

De-staging activities in a data storage system are controlled by providing a write pending list of elements, where each element is defined to store information related to a cache memory data element for which a "write to storage" is pending, and maintaining the write pending list so that de-staging of a data element can be based on the maturity of the pending write. Use the Write Pending Data (MBytes) metric in the following table to monitor this information.

Array Storage System Metrics¹

The following metrics measure the performance of HP XP or HDS arrays, including the P9500/XP storage arrays. See footnotes for exceptions.

Metric	Description	Units
CHA Cache Memory Busy Rate*	Rate at which the CHA cache is busy de-staging data to the DKA. Available for HP XP arrays only.	%
CHA Shared Memory Busy Rate*	Rate at which the CHA shared memory is busy	%
Cache Usage	Cache utilization in megabytes	MB
DKA Cache Memory Busy Rate*	Rate at which the DKA cache is busy de-staging data to the disk. Available for HP XP arrays only.	%
DKA Shared Memory Busy Rate*	Rate at which the DKA shared memory is busy. Available for HP XP arrays only.	%
Percent Cache Usage**	Cache utilization percent	%
Percent Sidefile Usage**	Percent utilization of the sidefile. A sidefile is an internal buffer that saves a copy of the data to be transmitted to a remote XP array. Use to track continuous access (CA) sidefile cache utilization and the potential impact of DR activities. Formula: $\text{SidefileUsage}/\text{CacheSize}$	%
Percent Write Pending Data**S	Percentage of pending writes based on the percentage of cache being used to buffer writes on the selected controller. Use to determine if a CLPR is needed or if attention needs to be directed toward journal parity groups. Formula: $\text{WritePendingData}/\text{CacheSize}$	%
Read Hits**	Read I/O requests per second satisfied from cache	Req/s
Sidefile Usage**	Sidefile cache utilization in megabytes	MB
Write Pending Data**	Indicator of pending writes based on cache in megabytes used to buffer writes on the selected controller.	MB

¹Performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

*Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays use MPB controllers.

**The default Cache Logical Partition (CLPR) used for obtaining performance information is CLPR 0. The default MP for the XP controller is MP 0.

XP and HDS Array Group Metrics

The Array Groups section lists all of the RAID Groups in the array. Select an array group to view its performance (for example, Utilization). You can also distinguish the I/O types and determine how the RAID group is being used to determine whether data should be migrated to another group.

Best Practices

Simultaneously graph the percentage of reads and writes, random and sequential I/O, average read and write size, and total utilization. By comparing the percentage of reads and writes, I/O types, and utilization, you can see how the RAID group is being used. If there is excessive random access, it could affect your sequential performance. You might, therefore, move your LUNs used for random I/O to another RAID group; for example, by using Auto LUN or Business Copy.

Array Group Metrics¹

The following performance metrics are available for HP XP and HDS array groups and includes the P9500/XP storage arrays (see footnote for exceptions). The size of I/Os being issued affects disk latency; that is, large I/O sizes typically result in slightly higher latency.

Metric	Description	Units
Average Read Size	Average number of reads in bytes to the array group.	Bytes
Average Write Size	Average number of writes in bytes to the array group	Bytes
Percent Read Hits	Percentage (%) of cache reads for the array group	%
Percent Read Hits Random	Percentage (%) of random read I/O's that were de-staged from cache	%
Percent Read Hits Seq	Percentage (%) of sequential read I/O's that were de-staged from cache	%

Metric	Description	Units
Percent Reads	Percentage (%) of reads from cache	%
Percent Writes	Percentage (%) of writes from cache	%
Read Data Rate	Rate data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Read Data Rate Random	Rate at which random data is read from the array group by all hosts	Bytes/s
Read Data Rate Seq	Rate at which sequential data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Read Hits	Cache read hits in requests per second	Req/s
Read Response Time	Time required to complete a read I/O in seconds	Sec
Reads	Number of read I/O's in requests per second	Req/s
Total Data Rate	Rate at which data can be transmitted between devices for the selected array group	Req/s
Total I/O Rate	Total number of read or write operations taking place per second for the selected array group	Req/s
Utilization	Percentage of time disks in the array group are busy	%
Write Data Rate	Rate at which data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Write Data Rate Random	Rate at which random data is written to the array group by all hosts	Bytes/s
Write Data Rate Seq	Rate at which sequential data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s
Write Response Time	Time required to complete a write I/O in seconds	Sec
Writes	Number of write I/O's in requests per second	Req/s

¹Performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

XP and HDS Array Volume Metrics

Volume performance is monitored by selecting the storage pool containing the volume. The pool shows all volumes (LDEVs, LUSEs, and V-VOLs) associated with it. Select individual volumes to view their performance metrics. When you display volume metrics, you can view its response time to the host and data rate.

Best Practices

When diagnosing a performance bottleneck, first graph the application's volume response time and data rates. If response time is high, and data rate is low, then there might be a problem with any of the stages in between (switch, CHA, DKA, and so forth). If response time is low/normal and data rate is high, storage is not the source of your application's performance problem. If the application is using more than one LUN, you can graph metrics that have the same units on the same graph, but not in real time.

Array Volume Metrics¹

The following metrics are available for HP XP and HDS array volumes, including the P9500/XP storage array volumes (see footnotes for exceptions). To simplify formulas, the value Time represents the difference in seconds between the most recent two StatisticTime values returned from the provider.

Metric	Description	Units	Formulas
Average Read Size	Average number of reads in bytes to the array group	Bytes	$(\Delta \text{RandomReadData} + \Delta \text{SequentialReadData}) / (\Delta \text{RandomReadIOs} + \Delta \text{SequentialReadIOs})$
Average Write Size	Average number of writes in bytes to the array group	Bytes	$(\Delta \text{RandomWriteData} + \Delta \text{SequentialWriteData}) / (\Delta \text{RandomWriteIOs} + \Delta \text{SequentialWriteIOs})$
Percent Read Hits	Percentage (%) of cache reads for the array group	%	$100 \times (\Delta \text{RandomReadHitIOs} + \Delta \text{SequentialReadHitIOs}) / (\Delta \text{RandomReadIOs} + \Delta \text{SequentialReadIOs})$
Percent Read Hits Random	Percentage (%) of random read I/O's that were de-staged from cache	%	$100 \times \Delta \text{RandomReadHitIOs} / \Delta \text{RandomReadIOs}$

Metric	Description	Units	Formulas
Percent Read Hits Seq	Percentage (%) of sequential read I/O's that were de-staged from cache	%	$100 \times \Delta \text{SequentialReadHitIOs} / \Delta \text{SequentialReadIOs}$
Percent Reads	Percentage (%) of reads from cache	%	$100 \times (\Delta \text{RandomReadIOs} + \Delta \text{SequentialReadIOs}) / (\Delta \text{RandomReadIOs} + \Delta \text{RandomWriteIOs} + \Delta \text{SequentialReadIOs} + \Delta \text{SequentialWriteIOs})$
Percent Writes	Percentage (%) of writes from cache	%	$100 \times (\Delta \text{RandomWriteIOs} + \Delta \text{SequentialWriteIOs}) / (\Delta \text{RandomReadIOs} + \Delta \text{RandomWriteIOs} + \Delta \text{SequentialReadIOs} + \Delta \text{SequentialWriteIOs})$
Read Data Rate	Rate at which data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{RandomReadData} + \Delta \text{SequentialReadData}) / \Delta \text{StatisticTime}$
Read Data Rate Random	Rate at which random data is read from the array group by all hosts	Bytes/s	$\Delta \text{RandomReadData} / \Delta \text{StatisticTime}$
Read Data Rate Seq	Rate at which sequential data is read from the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	$\Delta \text{SequentialReadData} / \Delta \text{StatisticTime}$
Read Hits	Cache read hits in requests per second	Req/s	$(\Delta \text{RandomReadHitIOs} + \Delta \text{SequentialReadHitIOs}) / \Delta \text{StatisticTime}$
Read Response	Time required to complete a read I/O in seconds	Sec	$\Delta \text{ReadResponseTimeCounter} / (\Delta \text{RandomReadIOs} + \Delta \text{SequentialReadIOs})$
Reads	Number of read I/O's in requests per second	Req/s	$(\Delta \text{RandomReadIOs} + \Delta \text{SequentialReadIOs}) / \Delta \text{StatisticTime}$

Metric	Description	Units	Formulas
Total Data Rate	Rate in which data can be transmitted between devices for the selected array group	Req/s	$(\Delta \text{RandomReadData} + \Delta \text{RandomWriteData} + \Delta \text{SequentialReadData} + \Delta \text{SequentialWriteData}) / \Delta \text{StatisticTime}$
Total I/O Rate	Total number of read or write operations taking place per second for the selected array group	Req/s	$(\Delta \text{RandomReadIOs} + \Delta \text{RandomWriteIOs} + \Delta \text{SequentialReadIOs} + \Delta \text{SequentialWriteIOs}) / \Delta \text{StatisticTime}$
Utilization	Percentage (%) of time that disks in the array group are busy	%	$100 \times \Delta \text{ActiveTime fields} / \Delta \text{StatisticTime}$
Write Data Rate	Rate at which data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	$(\Delta \text{RandomWriteData} + \Delta \text{SequentialWriteData}) / \Delta \text{StatisticTime}$
Write Data Rate Random	Rate at which random data is written to the array group by all hosts	Bytes/s	$\Delta \text{RandomWriteData} / \Delta \text{StatisticTime}$
Write Data Rate Seq	Rate at which sequential data is written to the array group by all hosts and includes transfers from the source array to the destination array	Bytes/s	$\Delta \text{SequentialWriteData} / \Delta \text{StatisticTime}$
Write Response Time	Time required to complete a write I/O in seconds	Sec	$\Delta \text{WriteResponseTimeCounter} / (\Delta \text{RandomWriteIOs} + \Delta \text{SequentialWriteIOs})$
Writes	Number of write I/O's in requests per second	Req/s	$(\Delta \text{RandomWriteIOs} + \Delta \text{SequentialWriteIOs}) / \Delta \text{StatisticTime}$

¹Performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

XP and HDS Front-end Controller CLPR Metrics

The following Cache Logical Partition (CLPR) metrics are available for the front-end controller.

Note: The MPB (Micro Processor Blade) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. These front-end controller CLPR metrics are therefore not used to measure the performance for XP P9500 arrays.

Metric	Description	Units	Formulas
Total Data Rate	Rate at which data can be transmitted between devices for the selected CLPR	Bytes\s	$\Delta \text{TotalDataTransferred} / \Delta \text{StatisticTime}$
Total I/O Rate	Total number of read or write operations taking place per second for the selected front-end controller CLPR	Req\s	$\Delta \text{TotalIOs} / \Delta \text{StatisticTime}$

XP and HDS Front-end Controller (CHA) Metrics

Front-end controllers route I/O to/from hosts to their disk cache slots. The front-end view shows all channel adapters (CHA) and their constituent ports. Selecting the CHA allows you to view the processor (MP) utilization.

Note: The MPB (Micro Processor Blade) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. This front-end controller (CHA) metric is therefore not used to measure performance for XP P9500 arrays.

Best Practices

Graph CHA MP Utilization and Port data rate simultaneously to determine if the port is the bottleneck. If MP utilization is less than 90%, and the port data rate is near its maximum, then you know that data is being de-staged to disk normally assuming the host(s) is running at its peak. If not, check the DKA and RAID Group utilization.

Controller CHA Metrics

The following performance metrics are available for the front-end controller.

Metric	Description	Units	Formula
Processor Utilization*	Processor utilization rate on the selected CHA controller	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$

*Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays utilize MPB controllers.

XP Back-end Controller (DKA) Metric

Back-end controllers route I/O from cache slots to the disk. The back-end view for XP arrays shows all the disk controller adapters (DKA). Selecting the DKA allows you to view processor utilization (MP).

The Hitachi Data Systems (HDS) Discovery process does not expose back-end controllers for HDS arrays. For this reason, there is no back-end controller (DKA) metric for HDS.

Note: The MPB (Micro Processor Blade) controllers for XP P9500 storage arrays have a separate set of metrics to measure performance. This back-end controller (DKA) metric is therefore not used to measure performance for XP P9500 arrays.

Best Practices

You can graph DKA MP Utilization and RAID Groups IOPS to determine if the controller is acting as a bottleneck. If MP utilization is around 90% and the RAID Group data rate is much lower than expected for the RAID level and disk type, there might be a background process such as replication that is over-utilizing the DKA. Typically, you want your back-up disks on a different RAID Group than your production disks, and the back-up disks should be connected to a different DKA.

Controller DKA Metrics

The following performance metrics are available for the XP back-end controller.

Metric	Description	Units	Formula
Processor Utilization*	Processor utilization rate on the selected DKA controller	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$

*Metric not available for XP and HDS RAID700 arrays. XP and HDS RAID700 arrays use MPB controllers.

XP and HDS MPB Controller Metrics

The following table describes the performance metrics collected for MPB (Micro Processor Blade) controllers for P9500/XP storage arrays (see footnote for exception). MBP processor statistics are obtained by HP Storage Essentials CIM extensions.

MPB Controller Metrics¹

Metric	Description	Units	Formula
Back End Utilization	Processor utilization for the back-end processes (back-end activities for target I/O requests) as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Mainframe External Initiator Utilization	Processor utilization for the mainframe external initiator processes as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Mainframe Target Utilization	Processor utilization rate for the mainframe target processes (front-end activities for processing mainframe I/O requests) as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Open External Initiator Utilization	Processor utilization rate for the open external initiator processes (external storage access activities) as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Open Initiator Utilization	Processor utilization rate for the open initiator processes (continuous access replication activities) as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Open Target Utilization	Processor utilization rate for open target processes (front-end activities) as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
Processor Utilization	Processor utilization rate on the selected MPB controller. This rate is the sum of the other seven metrics listed in this table.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$
System Utilization	Processor utilization rate of the array system processes as a percentage of total processor time.	%	$\Delta \text{BusyTimeCounter} / \Delta \text{ElapsedTimeCounter}$

¹The performance metrics in this table are not available for XP P9500 storage arrays discovered using Command View Advanced Edition (CVAE). For more information, see the Licensing chapter in the *Installation Guide*.

Managing NetApp Performance

HP Storage Essentials Performance Pack Enterprise software for NetApp systems provides more than 80 performance metrics that you can use to monitor your NetApp storage devices. A complete list of performance metrics along with descriptions and derived formulas is provided in this section. See these topics for information:

- [Monitoring NetApp Performance below](#)
- [NetApp Performance Metrics on the facing page](#)
- [NetApp Raw Statistics on page 179](#)

Monitoring NetApp Performance

HP Storage Essentials collects performance statistics derived from the NetApp devices in your storage environment and makes this data available to you in multiple ways for your monitoring and assessment activities. For example, you can use the Performance Manager to view data in real-time or chart historical data, and you can use Report Optimizer to generate reports and graphs for comparative analysis. Additionally, NetApp also provides information about monitoring NetApp devices for improved performance. You can find this information on the NetApp website at <http://partners.netapp.com/go/techontap/matl/monitor-troubleshoot.html>.

Using NetApp Performance Statistics to Identify Bottlenecks

HP Storage Essentials can help you identify bottlenecks and isolate potential problem areas before they affect your NetApp system. Two strong indicators of potential bottlenecks are throughput (I/O rates) and latency levels.

- Throughput is measured by a raw count of bytes or data packets transmitted and received along SAN connection paths; that is, along a storage system's Ethernet or target Fibre Channel interfaces. When a storage device performs at expected levels, without delays or queue backups, throughput statistics can be used as a measurement of baseline performance.
- Latency is the measured lapse in time between a data request send and its receipt (also called a round-trip loop). Latency can be measured for reads and writes: read and write hits and misses, reads and writes for a specific interface or protocol, and so forth. A low latency rate indicates optimum transmission of data.

Determine Latency Limits and Perform Baseline Analysis

Each application has acceptable limits for latency. Beyond these limits increasing degradation of performance becomes noticeable; user tasks take longer to execute and response times slow to unacceptable levels. Typically, this occurs over time as the number of users and applications increase, and heavier workloads are added to the existing storage system. Latency therefore is a strong indicator of how well your storage system handles workloads and where potential bottlenecks may develop over time.

However, what may be acceptable for one application or device, may be totally unacceptable for another. To determine the acceptable latency limits for your specific environment, establish a performance baseline that represents normal workloads and usage. Monitor your storage performance by comparing current, hourly, and daily performance data against your baseline statistics. If you expect heavier usage requirements at different times of the day, allow for this in establishing your performance baselines.

Add Performance Policies and Alerts

As an aid, set up performance policies with threshold alerts which advise you when latency limits are being exceeded. Like baselines, performance policies are a standard best practice for monitoring, evaluating, and verifying performance quality over time. See the *HP Storage Essentials Storage Performance Management Guide* for more information about setting up performance policies.

Understand NetApp Performance Metrics

You will find descriptions of the NetApp performance metrics available in HP Storage Essentials in [NetApp Performance Metrics](#) below and [NetApp Raw Statistics](#) on page 179. These metrics are used in charting, reporting, policy creation, and real-time analysis of the performance of your NetApp devices.

NetApp Performance Metrics

The following tables list the HP Storage Essentials performance metrics for NetApp storage systems.

NetApp System Performance Metrics

The following performance metrics are provided by HP Storage Essentials for NetApp storage systems, and include iSCSI, CIFS, DAFS, FCP, HTTP, and NFSv3 operations.

Metric	Description	Unit	Formula
Cache			
Buffer Cache Hits Count	Buffer Cache or system memory read cache hits per second. Use to determine if access latency is contributing to performance issues.	Req/s	$\Delta \text{Buf_Load_Cnt} / \Delta \text{Time}$
Buffer Cache Misses Count	Buffer cache miss count per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. Use to determine PAM deployment requirements and configuration.	Req/s	$\Delta \text{Buff_Miss_Cnt} / \Delta \text{Time}$

Metric	Description	Unit	Formula
Inode Cache Hits Count	Number of hits for inodes read from disk that are cached and subsequently accessed. Use to determine cause of increase in file system performance	Req/s	$\Delta \text{Inode_Cache_Hit} / \Delta \text{Time}$
Inode Cache Misses Count	Number of inode cache misses per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. The inode cache behaves in the same fashion. Use to determine if the inode cache needs to be increased.	Req/s	$\Delta \text{Inode_Cache_Miss} / \Delta \text{Time}$
Name Cache Hits	Number of name cache hits per second (rate). Use to determine frequency of name cache hits. The name cache improves file lookup in the file system.	Req/s	$\Delta \text{Name_Cache_Hit} / \Delta \text{Time}$
Name Cache Misses	Number of name cache misses per second (rate). A cache miss is simply data that is not in the cache. The result is that the system must fetch the data from the disk. The name cache behaves in the same fashion. Use to determine if the name cache needs to be increased.	Req/s	$\Delta \text{Name_Cache_Miss} / \Delta \text{Time}$
Common Internet File System (CIFS) I/O Operations			
CIFS Latency	Average latency for Common Internet File System (CIFS) operations in milliseconds.	ms	$\Delta \text{CIFS_Latency} / \Delta \text{CIFS_Latency_Base}$
CIFS Operations	Number of Common Internet File System (CIFS) operations per second.	Req/s	$\Delta \text{CIFS_Ops} / \Delta \text{Time}$
Direct Access File System (DAFS) I/O Operations			
DAFS Operations	Number of Direct Access File System (DAFS) operations per second.	Req/s	$\Delta \text{DAFS_Ops} / \Delta \text{Time}$
Fibre Channel Protocol (FCP) I/O Operations			
FCP Operations	Number of Fibre Channel Protocol (FCP) operations per second.	Req/s	$\Delta \text{FCP_Ops} / \Delta \text{Time}$

Metric	Description	Unit	Formula
FCP Read Data	FCP bytes read per second.	Bytes/s	$\Delta \text{FCP_Read_Data} / \Delta \text{Time}$
FCP Read Latency	Average latency for read operations observed over all LUNs in the system accessed over FCP in milliseconds.	ms	$\Delta \text{FCP_Read_Latency} / \Delta \text{FCP_Read_Ops}$
FCP Read Operations	Total number of read operations per second observed over all the LUNs in the system accessed over FCP.	Req/s	$\Delta \text{FCP_Read_Ops} / \Delta \text{Time}$
	FCP bytes written per second.	Bytes/s	$\Delta \text{FCP_Write_Data} / \Delta \text{Time}$
FCP Write Latency	Average latency for write operations observed over all LUNs in the system accessed over FCP in milliseconds.	ms	$\Delta \text{FCP_Write_Latency} / \Delta \text{FCP_Write_Ops}$
FCP Write Operations	Total number of write operations per second observed over all the LUNs in the system accessed over FCP.	Req/s	$\Delta \text{FCP_Write_Ops} / \Delta \text{Time}$
Hypertext Transfer Protocol (HTTP) I/O Operations			
HTTP Operations	Number of Hypertext Transfer Protocol (HTTP) operations per second.	Req/s	$\Delta \text{HTTP_Ops} / \Delta \text{Time}$
Internet Small Computer System Interface (iSCSI) I/O Operations			
iSCSI Operations	Number of Internet Small Computer System Interface (iSCSI) operations per second.	Req/s	$\Delta \text{iSCSI_Ops} / \Delta \text{Time}$

Metric	Description	Unit	Formula
ISCSI Read Data	iSCSI bytes read per second.	Bytes/s	$\Delta \text{ISCSI_Read_Data} / \Delta \text{Time}$
ISCSI Read Latency	Average latency of read operations observed over all LUNs in the system accessed over iSCSI in milliseconds.	ms	$\Delta \text{ISCSI_Read_Latency} / \Delta \text{ISCI_Read_Ops}$
ISCSI Read Operations	Total number of read operations per second observed over all the LUNs in the system accessed over iSCSI.	Req/s	$\Delta \text{ISCSI_Read_Ops} / \Delta \text{Time}$
ISCSI Write Data	iSCSI bytes written per second.	Bytes/s	$\Delta \text{ISCSI_Write_Data} / \Delta \text{Time}$
ISCSI Write Latency	Average latency of write operations observed over all LUNs in the system accessed over iSCSI in milliseconds.	ms	$\Delta \text{ISCSI_Write_Latency} / \Delta \text{ISCI_Write_Ops}$
ISCSI Write Operations	Total number of write operations per second observed overall the LUNs in the system accessed over iSCSI.	Req/s	$\Delta \text{ISCSI_Write_Ops} / \Delta \text{Time}$
Network File System (NFS) Operations			
NFSV3 Average Operations Latency	Average latency of the NFS v3 operations in milliseconds.	ms	$\Delta \text{NFSv3_Avg_Op_Latency} / \Delta \text{NFSv3_Avg_Op_Latency_Base}$
NFS Operations	Number of Network File System (NFS) operations per second.	Req/s	$\Delta \text{NFS_Ops} / \Delta \text{Time}$

Metric	Description	Unit	Formula
NFSV3 Read Latency	Average latency for NFS v3 read operations in milliseconds.	ms	$\Delta \text{NFSv3_Read_Latency} / \Delta \text{NFSv3_Avg_Read_Latency_Base}$
NFSV3 Write Latency	Average latency for NFS v3 write operations in milliseconds.	ms	$\Delta \text{NFSv3_Write_Latency} / \Delta \text{NFSv3_Avg_Write_Latency_Base}$

NetApp Host Processor (CPU) Performance Metrics

The following performance metrics are provided for NetApp host processors (CPU).

NetApp does not expose memory utilization statistics because all memory that is not used for initialization is allocated to the file system for caching. In other words, all memory is used all of the time. How it is allocated is what really matters. Processor utilization and cache hit ratio are better indicators of how utilized a filer is. If the cache hit ratio is low and CPU utilization is high, the filer might be reaching maximum utilization.

Metric	Description	Unit	Formula
Processor Utilization	Total CPU utilization (%) by all the processes running on the filer. Indicates the percentage (%) of time that the processor is active. A completely idle processor shows 0%. A processor saturated with activity shows 100%. Use to identify CPU bottlenecks.	%	$100 \times (\Delta \text{Processor_Busy} / \Delta \text{Processor_Elapsed_Time})$

NetApp Aggregate Performance Metrics

The following metric is provided for measuring NetApp aggregate transfer rates.

Metric	Description	Unit	Formula
Total Transfers	Total number of transfers per second serviced by the aggregate.	Req/s	$\Delta \text{Total_Transfers} / \Delta \text{Time}$

Metric	Description	Unit	Formula
User Read Block	Number of blocks read per second on the aggregate.	Blocks/s	$\Delta \text{User_Read_Blocks} / \Delta \text{Time}$
User Reads	Number of user reads per second to the aggregate.	Req/s	$\Delta \text{User_Reads} / \Delta \text{Time}$
User Write Blocks	Number of blocks written per second to the aggregate.	Blocks/s	$\Delta \text{User_Write_Blocks} / \Delta \text{Time}$
User Writes	Number of user writes per second to the aggregate.	Req/s	$\Delta \text{User_Writes} / \Delta \text{Time}$

NetApp Front-End FC Port Performance Metrics

The following front-end Fibre Channel (FC) port metrics measure performance state using HP Storage Essentials.

Metric	Description	Unit	Formula
Other Operations	Number of other operations per second.	Req/s	$\Delta \text{Other_Ops} / \Delta \text{Time}$
Queue Full Responses	SCSI queue full responses per second.	Req/s	$\Delta \text{Queue_Full} / \Delta \text{Time}$
Read Data	Number of bytes read from the filer per second.	Bytes/s	$\Delta \text{Read_Data} / \Delta \text{Time}$
Read Operations	Number of read operations per second.	Req/s	$\Delta \text{Read_Ops} / \Delta \text{Time}$
Write Data	Number of bytes written to the filer per second.	Bytes/s	$\Delta \text{Write_Data} / \Delta \text{Time}$
Write Operations	Number of write operations per second.	Req/s	$\Delta \text{Write_Ops} / \Delta \text{Time}$

NetApp IP Port Performance Metrics

The following metrics are available to measure NetApp IP port performance.

Metric	Description	Unit	Formula
Bytes Received	Inbound traffic in megabytes per second through the filer network interface controller (NIC). Use to assess network traffic for load balancing, multi-path optimization, and network performance.	MB/s	$(\Delta \text{Recv_Data} \times 1024) / \Delta \text{Time}$

Metric	Description	Unit	Formula
Bytes Transmitted	Outbound traffic in megabytes per second through the filer network interface controller (NIC). Use to assess network traffic for load balancing, multi-path optimization, and network performance.	MB/s	$(\Delta \text{Send_Data} \times 1024) / \Delta \text{Time}$
Packets Received	Inbound traffic in packets per second through the filer network interface controller (NIC). Network packets contain data headers, address source and destination, payload and CRC fields. Use to measure network traffic for load balancing, multi-path optimization, and network performance. Contrary to bytes received testing, packet testing is actually a better test because inbound packets have either arrived or not. Byte testing does not indicate whether or not a packet transmission completed.	Pkt/s	$\Delta \text{Recv_Packets} / \Delta \text{Time}$
Packets Transmitted	Outbound traffic in packets per second through the filer network interface controller (NIC). Network packets contain data headers, address source and destination, payload and CRC fields. Use to measure network traffic for load balancing, multi-pathing optimization and network performance. Contrary to bytes transmitted testing, packet testing is actually a better test because outbound packets have been either sent or not. Byte testing does not indicate whether or not a packet transmission completed.	Pkt/s	$\Delta \text{Send_Packets} / \Delta \text{Time}$
Receive Errors	Errors per second while receiving packets.	Req/s	$\Delta \text{Recv_Errors} / \Delta \text{Time}$
Send Errors	Errors per second while sending packets.	Req/s	$\Delta \text{Send_Errors} / \Delta \text{Time}$

NetApp Logical Drive Performance Metrics

The following NetApp logical drive performance metrics are available. The metrics for NetApp file system and volume statistics are combined.

Metric	Description	Unit	Formula
Average Latency	Average latency in milliseconds for all operations on the volume.	ms	$\Delta \text{Avg_Latency} / \Delta \text{Total_Ops}$
Other Latency	Average latency time for other writes to the volume in milliseconds.	ms	$\Delta \text{Other_Latency} / \Delta \text{Other_Ops}$
Other Operations	Number of other operations per second to the volume.	Req/s	$\Delta \text{Other_Ops} / \Delta \text{Time}$
Read Data	Bytes read per second from the volume.	Bytes/s	$\Delta \text{Read_Data} / \Delta \text{Time}$
Read Latency	Average latency time for reads to the volume in milliseconds.	ms	$\Delta \text{Read_Latency} / \Delta \text{Read_Ops}$
Read Operations	Number of reads per second to the volume.	Req/s	$\Delta \text{Read_Ops} / \Delta \text{Time}$
Reserved Inodes	Provides the count of reserved inodes in a file system. Use to count the reserved inodes. The first 10 inodes on a file system are special inodes. Inodes 7-10 are reserved and usually not used.	counter	–
Total Inodes	Total number of inodes. Inodes are file system data structures or metadata used to store basic file data like ownership and file permissions. Use to view the inode limit and determine if more are needed.	counter	–
Total Operations	Total number of operations per second serviced by the volume.	Req/s	$\Delta \text{Total_Ops} / \Delta \text{Time}$

Metric	Description	Unit	Formula
Used Inodes	Total number of inodes that are currently used. Use to show the number of inodes in use. Use to alert the admin when inode utilization is approaching total inodes available.	counter	–
Write Data	Bytes written per second to the volume.	Bytes/s	$\Delta \text{Write_Data} / \Delta \text{Time}$
Write Latency	Average latency time for writes to the volume in milliseconds.	ms	$\Delta \text{Write_Latency} / \Delta \text{Write_Ops}$
Write Operations	Number of writes per second to the volume.	Req/s	$\Delta \text{Write_Ops} / \Delta \text{Time}$

NetApp LUN Performance Metrics

The following metrics measure NetApp LUN performance.

Metric	Description	Unit	Formula
Average Latency	Average latency in milliseconds for all operations on the LUN	ms	$\Delta \text{Avg_Latency} / \Delta \text{Total_Ops}$
Other Operations	Number of other operations per second	Req/s	$\Delta \text{Other_Ops} / \Delta \text{Time}$
Queue Full Responses	Number of queue full responses per second	Req/s	$\Delta \text{Queue_Full} / \Delta \text{Time}$
Read Data	Number of bytes read per second	Bytes/s	$\Delta \text{Read_Data} / \Delta \text{Time}$
Read Operations	Number of read operations per second	Req/s	$\Delta \text{Read_Ops} / \Delta \text{Time}$
Total Operations	Total number of operations on the LUN per second	Req/s	$\Delta \text{Total_Ops} / \Delta \text{Time}$
Write Data	Number of bytes written per second	Bytes/s	$\Delta \text{Write_Data} / \Delta \text{Time}$
Write Operations	Number of write operations per second	Req/s	$\Delta \text{Write_Ops} / \Delta \text{Time}$

NetApp Disk Drive Performance Metrics

The following performance measurements are available for NetApp disk drives.

Metric	Description	Unit	Formula
Disk Busy	Percentage of time there was at least one outstanding request to the disk	%	$100 \times (\Delta \text{Disk_Busy} / \Delta \text{Disk_Busy_Base})$
Total Transfers	Total number of disk operations involving data transfer initiated per second	Req/s	$\Delta \text{Total_Transfers} / \Delta \text{Time}$
User Read Blocks	Number of blocks transferred for user read operations per second	Blocks/s	$\Delta \text{User_Read_Blocks} / \Delta \text{Time}$
User Read Latency	Average latency per block in milliseconds for user read operations	ms	$\Delta \text{User_Read_Latency} / \Delta \text{User_Read_Blocks}$
User Reads	Number of disk read operations initiated each second for retrieving data or metadata associated with user request	Req/s	$\Delta \text{User_Reads} / \Delta \text{Time}$
User Write Blocks	Number of blocks transferred for user write operations per second	Blocks/s	$\Delta \text{User_Write_Blocks} / \Delta \text{Time}$
User Write Latency	Average latency per block in milliseconds for user write operations	ms	$\Delta \text{User_Write_Latency} / \Delta \text{User_Write_Blocks}$
User Writes	Number of disk write operations initiated each second for retrieving data or metadata associated with user requests	Req/s	$\Delta \text{User_Writes} / \Delta \text{Time}$

NetApp Qtree Performance Metrics

The following performance metrics are provided for the NetApp QTree.

Metric	Description	Unit	Formula
CIFS Operations	Number of Common Internet File System (CIFS) operations per second to the qtree	Req/s	$\Delta \text{CIFS_Ops} / \Delta \text{Time}$

Metric	Description	Unit	Formula
Internal Operations	Number of internal operations generated by activities such as snapmirror and backup per second to the qtree	Req/s	$\Delta \text{Internal_Ops} / \Delta \text{Time}$
NFS Operations	Number of NFS operations per second to the qtree	Req/s	$\Delta \text{NFS_Ops} / \Delta \text{Time}$

NetApp Raw Statistics

The following raw values calculate the performance metric data provided in the Performance Manager for NetApp system devices. All raw metrics are counters derived from the NetApp device.

For a list of the Performance Manager metrics that use these raw statistics, see the [NetApp Performance Metrics](#) on page 169.

NetApp System Raw Statistics

The following values are derived from the NetApp device and are stored in the NAS_SYSTEM_STATS database table. This database contains the iSCSI, CIFS, FCP, NFSv3 and general system statistics.

Raw Statistic	Description	Unit
buf_load_cnt	Number of buffer cache or system memory read cache hits.	Hit count
buff_miss_cnt	Number of buffer cache misses.	Miss count
cifs_latency	Average latency for Common Internet File System (CIFS) operations.	Time count in milliseconds
cifs_latency_base	Total observed CIFS operations to be used as a base counter for CIFS average latency calculation.	Time count in milliseconds
cifs_ops	Number of Common Internet File System (CIFS) operations.	Request count
dafs_ops	Number of Direct Access File System (DAFS) operations.	Request count
fcp_ops	Number of Fibre Channel Protocol (FCP) operations.	Request count
fcp_read_data	Number of FCP bytes read.	Byte count
fcp_read_latency	Average latency for read operations observed over all LUNs in the system accessed over FCP.	Time count in milliseconds

Raw Statistic	Description	Unit
fcg_read_ops	Total number of read operations observed over all the LUNS in the system accessed over FCP.	Request count
fcg_write_data	Number of FCP bytes written.	Byte count
fcg_write_latency	Average latency for write operations observed over all LUNs in the system accessed over FCP.	Time count in milliseconds
fcg_write_ops	Total number of write operations observed over all the LUNs in the system accessed over FCP.	Request count
http_ops	Number of Hypertext Transfer Protocol (HTTP) operations.	Request count
inode_cache_hit	Number of hits for inodes read from disk that are cached.	Hit count
inode_cache_miss	Number of inode cache misses.	Miss count
iscsi_ops	Number of Internet Small Computer System Interface (iSCSI) operations.	Request count
iscsi_read_data	iSCSI bytes read.	Byte count
iscsi_read_latency	Average latency of read operations observed over all LUNs in the system accessed over iSCSI.	Time count in milliseconds
iscsi_read_ops	Total number of read operations per second observed over all the LUNs in the system accessed over iSCSI.	Request count
iscsi_write_data	Number of iSCSI bytes written.	Byte count
iscsi_write_latency	Average latency of write operations observed over all LUNs in the system accessed over iSCSI.	Time count in milliseconds
iscsi_write_ops	Total number of write operations observed on all LUNs in the system accessed over iSCSI.	Request count
name_cache_hit	Number of name cache hits.	Hit count
name_cache_miss	Number of name cache misses.	Miss count

Raw Statistic	Description	Unit
nfsv3_avg_op_latency	Average latency of the NFS v3 operations.	Time count in milliseconds
nfsv3_avg_op_latency_base	Array of select NFS v3 operation counts for latency calculation.	Time count in milliseconds
nfs_ops	Number of Network File System (NFS) operations.	Request count
nfsv3_read_latency	Average latency for NFS v3 read operations.	Time count in milliseconds
nfsv3_write_latency	Average latency for NFS v3 write operations.	Time count in milliseconds

NetApp Aggregate Raw Statistics

The following values are derived from the NetApp device and are stored in the NAS_AGGREGATE_STATS database table.

Raw Statistic	Description	Unit
total_transfers	Number of transfers.	Request count
user_read_blocks	Number of blocks read.	Data block count
user_reads	Number of user reads.	Request count
user_write_blocks	Number of blocks written.	Data block count
user_writes	Number of user writes.	Request count

NetApp Logical Drive Raw Statistics

The following values are derived from the NetApp device and are stored in the NAS_FILESYSTEM_STATS database table. This database contains the volume and file system statistics.

Raw Statistic	Description	Unit
avg_latency	Average latency for all operations on the volume.	Time count in milliseconds
other_latency	Latency time for other writes to the volume.	Time count in milliseconds
other_ops	Number of other operations to the volume.	Request count

Raw Statistic	Description	Unit
read_data	Number of bytes read from the volume.	Byte count
read_latency	Average latency time for reads to the volume.	Time count in milliseconds
read_ops	Number of read operations to the volume.	Request count
total_ops	Total number of operations serviced by the volume.	Request count
write_data	Number of bytes written to the volume.	Byte count
write_latency	Average latency time for writes to the volume.	Time count in milliseconds
write_ops	Number of write operations to the volume.	Request count
wv_fsinfo_inos_reserve	Reserved inodes in the volume.	Inode count
wv_fsinfo_inos_total	Total number of inodes in the volume.	Inode count
wv_fsinfo_inos_used	Used inodes in the volume.	Inode count

NetApp Processor (CPU) Raw Statistics

The following values are derived from the NetApp device and are stored in the HOSTCPUSTAT database table.

Metric	Description	Unit
processor_busy	Time CPU is being utilized. Uses the wall-clock time since boot and is used for calculating processor utilization.	Time count in seconds
processor_elapsed_time	Time CPU is running, including busy and idle periods.	Time count in seconds

NetApp Front-end Port Raw Statistics

The following values are derived from the NetApp device and are stored in the NAS_FCPORT_STATS database table.

Raw Statistic	Description	Unit
other_ops	Number of other operations to the volume.	Request count
read_data	Number of bytes read from the filer.	Byte count
read_ops	Number of read operations to the volume.	Request count

Raw Statistic	Description	Unit
write_data	Number of bytes written to the filer.	Byte count
write_ops	Number of write operations to the volume.	Request count
queue_full	Number of responses that SCSI queue is full.	Response count

NetApp IP Port Raw Statistics

The values for the following statistics are derived from the NetApp device and are stored in the IP_PORTSTATS database table.

Raw Statistic	Description	Unit
rcv_data	Amount of inbound data in bytes received through the filer network interface controller (NIC).	Byte count
rcv_errors	Number of errors which occurred while receiving packets.	Error count
rcv_packets	Number of inbound packets received through the filer network interface controller (NIC).	Packet count
send_data	Amount of outbound data in bytes sent through the filer network interface controller (NIC).	Byte count
send_errors	Number of errors which occurred while sending packets.	Error count
send_packets	Number of outbound packets sent through the filer network interface controller (NIC).	Packet count

NetApp LUN Raw Statistics

The following values for statistics are derived from the NetApp device and are stored in the NAS_LUN_STATS database table.

Raw Statistic	Description	Unit
avg_latency	Average latency time count for all operations on the LUN.	Time count in milliseconds
other_ops	Number of other operations.	Request count
queue_full	Number of queue full responses.	Response count
read_data	Number of bytes read.	Byte count

Raw Statistic	Description	Unit
read_ops	Number of read operations from the LUN.	Request count
total_ops	Number of total operations on the LUN.	Request count
write_data	Number of bytes written to the LUN.	Byte count
write_ops	Number of write operations to the LUN.	Request count

NetApp Disk Drive Raw Statistics

The following values for statistics are derived from the NetApp device and are stored in the NAS_DISKDRIVE_STATS database table.

Raw Statistic	Description	Unit
disk_busy	Time count for when there was at least one outstanding request to the disk.	Time count
disk_busy_base	Time base for disk_busy calculation.	Time count
total_transfers	Total number of disk operations involving data transfers.	Request count
user_read_blocks	Number of blocks transferred for user read operations.	Block count
user_read_latency	Latency for user read operations.	Time count in microseconds
user_reads	Number of disk read operations associated with a user request.	Request count
user_write_blocks	Number of blocks transferred for user write operations.	Block count
user_write_latency	Latency for user write operations.	Time count in microseconds
user_writes	Number of disk write operations associated with user requests.	Request count

NetApp QTree Raw Statistics

The following values are derived from the NetApp device and stored in the NAS_QTREE_STATS database table.

Raw Statistic	Description	Unit
cifs_ops	Number of Common Internet File System (CIFS) operations to the qtree.	Request count
Internal_ ops	Number of internal operations generated by activities (such as snapmirror and backup) to the qtree.	Request count
nfs_ops	Number of NFS operations to the qtree.	Request count

7 Generating Performance Reports

This section describes the Report Optimizer and different ways to generate performance reports that meet your administrative and management needs. It includes the following topics:

- [Ways to Generate Performance Reports below](#)
- [Accessing InfoView Logon below](#)
- [Selecting Standard Performance Reports on next page](#)
- [List of Standard Performance Reports on page 189](#)
- [Running Standard Performance Reports on page 222](#)
- [About Creating Simple Custom Reports on page 231](#)
- [Creating Custom Performance Reports on page 236](#)
- [Creating an EVA Performance Report on page 240](#)
- [Creating Advanced Custom Reports on page 248](#)
- [Creating Performance Reports with Drill Down and Trending on page 257](#)

For more information about the HP Storage Essentials Report Optimizer see *Creating Reports in a Complex Environment* available from the HP Storage Essentials Documentation Center (**Help > Documentation Center**).

Ways to Generate Performance Reports

You can generate performance reports in two ways. Both methods are described in this guide.

- Select a Report Optimizer standard performance report (see [Selecting Standard Performance Reports on next page](#)).
- Generate a custom performance report (see [Creating Custom Performance Reports on page 236](#)).

Accessing InfoView Logon

InfoView enables you to access and apply approximately 3000 HP Storage Essentials database objects and associated query filters. Use InfoView to navigate Report Optimizer and apply objects to create HP Storage Essentials Custom Reports.

To access InfoView:

1. From a supported browser, go to `http://<report_optimizer_server_name_or_IP_address>:8080/InfoViewApp/logon.jsp`.
2. Log on with the following credentials:

Username: Administrator

Password: (blank)

Log On to InfoView | [Help](#)

Enter your user information and click Log On.
(If you are unsure of your account information, contact your system administrator.)

User Name:

Password:

Selecting Standard Performance Reports

SRM Report Optimizer provides a set of standard performance reports that you can use. To access the standard performance reports:

1. Open Report Optimizer.
2. Click **Document List**.
3. Expand **Public Folders > Report Packs > Standard Reports > Performance Reports**.
4. Under Performance Reports, select the type of performance report you want:
 - NetApp
 - Storage System - EMC Symmetrix
 - Storage System - EVA
 - Storage System - XP
 - Storage System - HDS
 - Switch Performance

A list of the standard performance reports provided for that area appears in the right pane.

See [List of Standard Performance Reports on the facing page](#) for a list and description of available reports.

List of Standard Performance Reports

The following performance reports provided in the Storage Resource Management (SRM) Report Optimizer. These are standard reports that require you to enable data collection for the metrics gathered in the reports, as well as specify the collection period and storage objects for which you want data. All standard reports are Web Intelligence report documents. All report data is historical data (unless noted otherwise) and dependent on data collection schedules. For information about setting up performance data collection for standard reports, see [Setting Up Performance Data Collectors on page 50](#).

A report can be more than one page in length. Look at the page indicator in the upper right corner of the report screen to see the number of pages. Click the next and previous arrows to move between pages and make your report selection.

Use the following table to learn about performance reports and to help you determine which reports are appropriate for your reporting needs.

Name	Description	Performance Data
Performance		
Performance - NetApp*		
NetApp FC Port Total Request Rate	<p>Displays the front-end Fibre Channel (FC) port request rates for the selected NetApp ports.</p> <p>Port statistics are plotted on graphs for the given time period with the controllers for a single NetApp on a single graph.</p>	<p>Data: NetApp System Name, Front End Port Name, Collection Time</p> <p>Metric: Total Requests/Second</p>
NetApp FC Port Total Throughput	<p>Displays the front-end Fibre Channel (FC) port throughput for the selected NetApp ports.</p> <p>Port statistics are graphed for the specified time period, and controllers for a single NetApp are plotted on a single graph.</p>	<p>Data: NetApp System Name, Front End Port Name, Collection Time</p> <p>Metric: Total Throughput Rate (Bytes/Sec)</p>

Name	Description	Performance Data
NetApp IP Port Total Request Rate	<p>Displays the front-end Ethernet IP port request rate for the selected NetApp ports.</p> <p>Port statistics are graphed for the specified time period with the ports for a single NetApp on a single graph.</p>	<p>Data: NetApp System Name, Front End Port Name, Collection Time</p> <p>Metric: Total Requests/Second</p>
NetApp IP Port Total Throughput	<p>Displays the front-end Ethernet IP port throughput for the selected NetApp ports.</p> <p>Port statistics are graphed for the specified time period with the ports for a single NetApp on a single graph.</p>	<p>Data: NetApp System Name, Front End Port Name, Collection Time</p> <p>Metric: Total Throughput Rate (KBytes/Sec)</p>
NetApp System Latency and Operations	<p>Displays the NetApp system latency and operations per second (req/s) for each protocol for the NetApps selected.</p> <p>Statistics are graphed for the specified time period, and protocols for a single NetApp are plotted on a single graph.</p>	<p>Data: NetApp System Name, Collection Time</p> <p>Metrics: CIFS Latency, FCP Write Latency, FCP Read Latency, iSCSI Read Latency, iSCSI Write Latency, NFS Read latency, NFS Write latency, CIFS Operations/Second, FCP Read Operations/Second, FCP Write Operations/Second, iSCSI Read Operations/Second, iSCSI Write Operations/Second, NFS Read Operations/Second, NFS Write Operations/Second</p>

Name	Description	Performance Data
NetApp Top N Aggregate Total Request Rate	<p>Displays the top 20 NetApp aggregates with the largest average request rate for the selected NetApps. The number (N) of displayed aggregates varies if multiple aggregates have the same calculated average total request rate.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum(([Total Transfers (ops/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Aggregate]))/(Sum([Duration (ms)] ForEach ([Collection Time];[Statistics Type];[NAS + Aggregate])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Aggregate Name, Collection Time</p> <p>Metrics: Total OPS, Read OPS, Write OPS</p>
NetApp Top N Disk Utilization	<p>Displays the top 20 NetApp disks with the highest average utilization for the selected NetApps. The number (N) of displayed disks varies if multiple disks have the same calculated average, such as '100', and this average falls within the top 20 calculated.</p> <p>The average is calculated using the following formula:</p> <pre>'Sum(([Disk Busy]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Disk]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Disk]))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Disk Name, Collection Time</p> <p>Metrics: Utilization Rate, Disk Size, Disk Type</p>

Name	Description	Performance Data
NetApp Top N Ethernet Port Total Request Rate	<p>Displays the top 10 Ethernet ports with the largest average request rate for the selected NetApps. The number (N) of displayed ports varies if multiple ports have the same calculated average total request rate.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum(([Packets Total (packets/sec)]* [Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Port]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Controller, Port, Port Speed, Collection Time</p> <p>Metrics: Packets Received, Packets Transmitted, Total Packets</p>
NetApp Top N Ethernet Port Total Throughput	<p>Displays the top 10 Ethernet ports with the largest average data throughput rate for the selected NetApps. The number (N) of displayed ports varies if multiple ports have the same calculated average total throughput.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum(([Bytes Total (KBytes/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Port]))/(Sum ([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Controller, Port, Port Speed, Collection Time</p> <p>Metrics: Bytes Received, Bytes Transmitted, Total Bytes</p>

Name	Description	Performance Data
NetApp Top N FC Port Total Request Rate	<p>Displays the top 10 NetApp front-end Fibre Channel ports with the largest average request rate for the selected NetApps. The number (N) of displayed ports varies if multiple ports have the same calculated average total operations per second.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total Operations (ops/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Site + NAS + Port]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Site + NAS + Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Controller, Port, Port Speed, Collection Time</p> <p>Metrics: Read Operations (ops/sec), Write Operations (ops/sec), Other Operations (ops/sec), Total Operations (ops/sec)</p>
NetApp Top N FC Port Total Throughput	<p>Displays the top 10 NetApp front-end Fibre Channel ports with the largest average throughput for the selected NetApps. The number (N) of displayed ports varies if multiple ports have the same calculated throughput.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum(([Total Data (KBytes/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Port]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Controller, Port, Port Speed, Collection Time</p> <p>Metrics: Read Bytes, Write Bytes, Total Bytes</p>

Name	Description	Performance Data
NetApp Top N LUN Latency	<p>Displays the top 20 NetApp LUNs with the longest average latency for the selected NetApps. The number (N) of displayed LUNs varies if multiple volumes have the same calculated average latency.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Average Latency (ms)]*[Duration(ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Volume]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Volume Name, Collection Time</p> <p>Metrics: Average Latency, Read Latency, Write Latency, Other Latency</p>
NetApp Top N LUN Total Requests	<p>Displays the top 20 NetApp LUNs with the highest average rate of total operations per second for the selected NetApps. The number (N) of displayed LUNs varies if multiple volumes have the same calculated average total requests per second.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total Operations (ops/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS +LUN]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + LUN])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, LUN Name, Collection Time</p> <p>Metrics: Total OPS, Read OPS, Write OPS, Other OPS</p>

Name	Description	Performance Data
NetApp Top N LUN Total Throughput	<p>Displays the top 20 NetApp LUNs with the highest average rate of total bytes per second for the selected NetApps. The number (N) of displayed LUNs varies if multiple LUNs have the same calculated average total throughput.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total Data (KBytes/sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + LUN]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + LUN])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, LUN Name, Collection Time</p> <p>Metrics: Total KBytes/sec, Read KBytes/sec, Write KBytes/sec</p>
NetApp Top N Volume Latency	<p>Displays the top 20 NetApp volumes with the longest average latency for the selected NetApps. The number (N) of displayed volumes varies if multiple volumes have the same calculated average latency.</p> <p>The average is calculated using this formula:</p> <pre>'Sum([Average Latency (ms)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[NAS + Volume]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[NAS + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Volume Name, Collection Time</p> <p>Metrics: Average Latency, Read Latency, Write Latency, Other Latency</p>

Name	Description	Performance Data
NetApp Top N Volume Total Request Rate	<p>Displays the top 20 NetApp volumes with the highest average request rate for the selected NetApps. The number (N) of volumes displayed in the report varies if multiple volumes have the same calculated average total request rate.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total Operations (ops/sec)]*[Duration (ms)])ForEach([Collection Time];[Statistics Type];[NAS + Volume]))/(Sum([Duration(ms)]ForEach([Collection Time];[Statistics Type];[NAS + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Volume Name, Collection Time</p> <p>Metrics: Total OPS, Read OPS, Write OPS, Other OPS</p>
NetApp Top N Volume Total Throughput	<p>Displays the top 20 NetApp volumes with the highest average throughput for the NetApps you select. The number (N) of volumes displayed varies if multiple volumes have the same calculated average total throughput.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total Data (KBytes/sec)]*[Duration (ms)])ForEach([Collection Time];[Statistics Type];[NAS + Volume]))/(Sum([Duration (ms)]ForEach([Collection Time];[Statistics Type];[NAS + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: NetApp System Name, Volume Name, Collection Time</p> <p>Metrics: Total KBytes/sec, Read KBytes/sec, Write Kbytes/sec</p>
Performance - Storage System - EMC Symmetrix		

Name	Description	Performance Data
Symmetrix Bottom N Front End Controller Hit Rate	<p>Displays the bottom 10 front-end EMC Symmetrix controllers with the lowest I/O hit percentage for cache memory for the selected storage system. The number (N) of displayed controllers varies if multiple controllers have the same calculated average cache hit rate.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([% Hit I/Os]* [Duration (ms)])ForEach ([Array+Controller];[Collection Time];[Statistics Type]))/(Sum([Duration (ms)] ForEach([Array+Controller];[Collection Time];[Statistics Type])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Storage System Name, Front End Controller Name, Collection Time</p> <p>Metric: Total I/O (Req/Sec)</p> <p>The bottom controller averages are graphed for the specified time period with the lowest average listed first. Controllers that have an average I/O request rate of less than .01 req/s are ranked lower to keep them out of the report or listed last.</p>

Name	Description	Performance Data
Symmetrix Bottom N Volume Cache Hit Rate	<p>Displays the bottom 20 EMC Symmetrix volumes that have the lowest average cache hit rate (% Hit I/O metric) for the selected array(s). N is the number (1-20) of volumes to display in the report.</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total IO (Req/Sec)]*[Duration (ms)]) ForEach([Collection Time]; [Statistics Type];[Site + Storage + Volume]))/(Sum([Duration (ms)] ForEach ([Collection Time];[Statistics Type];[Site + Storage + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Symmetrix System Name, Volume Name, Collection Time</p> <p>Metrics: % Hit I/Os, % Miss I/Os, % Read Hit I/Os, % Read Miss I/Os, % Write Hit I/Os, % Write Miss I/Os</p> <p>The bottom volume averages are graphed for the specified time period with the lowest average listed first. Volumes with an average request rate of zero with a hit % rate (no I/O) are ranked below a 100% hit rate.</p>
Symmetrix Cache Usage	<p>Displays the EMC Symmetrix cache utilization statistics for the selected storage array(s). Utilization statistics are graphed for the time period you specify.</p>	<p>Data: Storage System Name</p> <p>Metrics: Read Hit %, Write Hit %, Total Hit %, Max Pending Flush Limit, Pending Flush, Total Permacache Slots, Used Permacache Slots, Delayed Fast Write Rate</p>
Symmetrix Front End Controller Total Request Rate	<p>Displays the front-end EMC Symmetrix controller total request rate (req/s) for the array(s) and controller(s) you select.</p> <p>The report graphs the controller statistics for time period you specify with the controllers for a single array on a single graph.</p>	<p>Data: Storage System Name, Front End Controller Name, Collection Time</p> <p>Metrics: Total I/O Rate (Req/Sec)</p>

Name	Description	Performance Data
Symmetrix Front End Port Total Request Rate	<p>Displays the EMC Symmetrix front-end port request rate (req/s) for the selected array port(s).</p> <p>Port statistics are graphed for the specified time period with the ports for a single array on a single graph.</p>	<p>Data: Storage System Name, Front End Port Name, Collection Time</p> <p>Metrics: Total I/O (Req/Sec)</p>
Symmetrix Front End Port Total Throughput	<p>Displays the EMC Symmetrix front-end port total data throughput (KB/s) for the selected array port(s).</p> <p>Port statistics are graphed for the specified time period with the ports for a single array on a single graph.</p>	<p>Data: Storage System Name, Front End Port Name, Collection Time</p> <p>Metric: Total Data Rate (KBytes/Sec)</p>
Symmetrix Top N Front End Controller Total Request Rate	<p>Displays the top 10 EMC Symmetrix controllers for the selected storage array(s) based on the average I/Os per second (IOPS).</p> <p>The average is calculated using this formula:</p> <pre>'Sum(([Total I/O (Req/Sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Array+Controller])) / (Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Array+Controller])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Storage System Name, Front End Controller Name, Collection Time</p> <p>Metrics: Total IOPS, Read IOPS, Write IOPS, Total Cache Hit IOPS</p> <p>Top controller averages are graphed for the specified time period with the highest average listed first.</p>

Name	Description	Performance Data
Symmetrix Top N Front End Port Total Request Rate	<p>Displays the top 20 EMC Symmetrix ports that have the highest average number of I/Os per second (IOPS) over a selected time period for the selected array(s).</p> <p>The average is calculated using this formula:</p> <pre>'=Sum([Total I/O (Req/Sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Array+Port])/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Array+Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Storage System Name, Front End Port Name, Collection Time</p> <p>Metric: Total I/O (Req/Sec)</p> <p>Top 20 port averages are graphed for the specified time period with the highest average listed first.</p>
Symmetrix Top N Front End Port Total Throughput	<p>Displays the top 20 EMC Symmetrix front-end ports with the highest average total data throughput (KB/s) for the selected array(s). The number (N) of displayed volumes varies if multiple volumes have the same calculated average total throughput.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum([Total Data Rate (Bytes/Sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Array+Port])/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Array+Port])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Storage System Name, Front End Port Name, Collection Time</p> <p>Metrics: Total Data Rate (KBytes/Sec)</p> <p>Top 20 port averages are graphed for the specified time period with the highest average listed first.</p>

Name	Description	Performance Data
Symmetrix Top N Volume Request Rate	<p>Displays the top 20 EMC Symmetrix volumes with the highest average request rate (req/s) for the selected array(s). The number (N) of volumes displayed in the report varies if multiple volumes have the same calculated average total request rate.</p> <p>The average is calculated using this formula:</p> <pre>'=Sum(([Total IO (Req/Sec)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Site + Storage + Volume]))/(Sum([Duration (ms)] ForEach ([Collection Time]; [Statistics Type];[Site + Storage + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric.</p>	<p>Data: Symmetrix System Name, Volume Name, Collection Time</p> <p>Metrics: Total req/sec, Read req/sec, Write req./sec., Total Hit req/sec, Total Miss req/sec</p> <p>Top volume averages are graphed with the highest average listed first.</p>
Symmetrix Top N Volume Response Time	<p>Displays the top 20 EMC Symmetrix volumes that have the highest average response times for the array(s) you select. The number (N) of volumes displayed in the report varies if multiple volumes have the same calculated average total response time.</p> <p>Response times are measured from the host directors on approximately 30% of the operations. The average is then calculated using this formula:</p> <pre>'Sum([Max AVG (ms)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Site + Storage + Volume]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Site + Storage + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric. The maximum is calculated by using the greater of the measured read or write sampled average times for the collection time.</p>	<p>Data: Symmetrix System Name, Volume Name, Collection Time</p> <p>Metrics: Sampled AVG Read Time (ms), Sampled AVG Write Time (ms)</p>

Name	Description	Performance Data
Symmetrix Top N Volume Throughput	<p>Displays the top 20 EMC Symmetrix volumes with the highest average throughput rates (KB/s) for the array(s) you select. The number (N) of volumes displayed in the report varies if multiple volumes have the same calculated average throughput.</p> <p>Response times are measured from the host directors on approximately 30% of the operations. The average is then calculated using this formula:</p> <pre>'Sum(([Max AVG (ms)]*[Duration (ms)]) ForEach([Collection Time];[Statistics Type];[Site + Storage + Volume]))/(Sum([Duration (ms)] ForEach([Collection Time];[Statistics Type];[Site + Storage + Volume])))'</pre> <p>where duration is the duration of the collection time for the metric. The maximum is calculated by using the greater of the measured read or write sampled average times for the collection time.</p>	<p>Data: Symmetrix System Name, Volume Name, Collection Time</p> <p>Metrics: Sampled AVG Read Time (ms), Sampled AVG Write Time (ms)</p> <p>Top volume averages are graphed with the highest average listed first.</p>
Symmetrix Volume Cache Hit Rate	<p>Displays the EMC Symmetrix volume cache hit rate for the selected storage volume(s).</p> <p>The graph lists volumes from those with the highest to lowest average cache hit rate. Volumes with an average request/sec rate of 0 (no I/O) and a hit rate of 0% are listed after volumes with a hit rate of 100%.</p>	<p>Data: Symmetrix System Name, Volume Name, Volume Pool Name, Raid Type, Collection Time</p> <p>Metrics: % Hit I/O, % Read Hit I/O - Random, % Read Hit I/O - Total, % Write Hit I/O</p>

Name	Description	Performance Data
Symmetrix Volume Request Rate	Displays the EMC Symmetrix volume request rate (req/s) for the selected storage volume(s). The graph lists volumes from those with the highest to the lowest average request rate.	Data: Symmetrix System Name, Volume Name, Volume Pool Name, Raid Type, Collection Time Metrics: Total I/O Req/Sec, Total Miss I/O Req/Sec, Read I/O - Total Req/Sec, Read I/O - Random Req/Sec, Read Miss I/O Req/Sec, Write I/O Req/Sec, Write Miss I/O Req/Sec
Symmetrix Volume Response Times	Displays the EMC Symmetrix volume average response times for the selected volume(s). Response times are measured from the host directors on about 30% of the operations for the selected storage volumes. The graph lists volumes from those with the highest to lowest average response times.	Data: Symmetrix System Name, Volume Name, Volume Pool Name, Raid Type, Collection Time Metrics: Sampled AVG Read Time (ms), Sampled AVG Write Time (ms)
Symmetrix Volume Throughput	Displays the EMC Symmetrix throughput rates for the storage volume(s) you select. Volumes are listed from those with the highest to lowest average throughput rate.	Data: Symmetrix System Name, Volume Name, Volume Pool Name, Raid Type, Collection Time Metrics: Total Data (KBytes/sec), Read Data (KBytes/sec), Write Data (KBytes/sec)
Performance - Storage System - EVA		

Name	Description	Performance Data
EVA Physical Disk Average Drive Latency	Displays the average drive latency for a physical disk over a time period that you specify.	Data: Storage System Name, Disk Drive Name, Collection Time Metric: Drive Latency
EVA Physical Disk Average Queue Depth	Displays the average queue depth for a physical disk drive in the collection time range selected.	Data: Storage System Name, Disk Drive Name, Collection Time Metric: Queue Depth
EVA Physical Disk Latency	Displays the read and write latencies for a physical disk drive in the collection time range selected.	Data: Storage System Name, Disk Drive Name, Collection Time Metrics: Average Read Latency, Average Write Latency
EVA Physical Disk Request Rate	Displays the I/O request rates for a physical disk drive.	Data: Storage System Name, Disk Drive Name, Collection Time Metrics: Total I/O, Read I/O, Write I/O, % Read I/O, % Write I/O
EVA Physical Disk Throughput	Displays the throughputs for a physical disk drive.	Data: Storage System Name, Disk Drive Name, Collection Time Metrics: Total (MB/s), Read (MB/s), Write (MB/s)

Name	Description	Performance Data
EVA Storage Controller CPU Utilization	Displays the EVA controller utilization over a selected time period.	Data: Storage System Name, Storage Processor Name Metric: % CPU
EVA Storage Controller Data Transfer	Displays the EVA controller data transfer over a selected period of time.	Data: Storage System Name, Storage Processor Name Metric: % Data Transfer
EVA Storage Controller Latency	Displays the EVA controller average read and average write latency over a selected time period.	Data: Storage System Name, Storage Processor Name Metric: Average Read Latency, Average Write Latency
EVA Storage Controller Request Rate	Displays the I/O request rates for a storage controller over a selected time period.	Data: Storage System Name, Storage Processor Name Metrics: Total I/O, Read I/O, Write I/O, % Read I/O, % Write I/O
EVA Storage Controller Throughput	Displays the throughputs for a storage controller over a selected time period.	Data: Storage System Name, Storage Processor Name Metric: Read (MB/s), Write (MB/s)

Name	Description	Performance Data
EVA Storage FC Port Average Queue Depth	Displays the average queue depth statistics of storage Fibre Channel (FC) ports over a selected time period.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metrics: Average Queue Depth
EVA Storage FC Port Latency	Lists the storage Fibre Channel (FC) port latency statistics over a selected time period.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metrics: Average Read Latency, Average Write Latency
EVA Storage FC Port Request Rate	Displays the I/O request rates for a storage Fibre Channel (FC) port over a selected time period.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metrics: Total I/O, Read I/O, Write I/O, % Read I/O, % Write I/O
EVA Storage FC Port Throughput	Lists the throughput statistics of storage Fibre Channel (FC) port over a selected time period.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metrics: Read (MB/s), Write (MB/s)

Name	Description	Performance Data
EVA Storage Pool Latency	Displays the EVA storage pool average read miss, average read hit, and average write latency statistics over a selected time period.	Data: Storage System Name, Storage Pool Name, Collection Time Metrics: Average Read Hit Latency, Average Read Miss Latency, Average Write Latency
EVA Storage Pool Request Rate	Displays the I/O request rates of an EVA storage pool over a selected period of time.	Data: Storage System Name, Storage Pool Name, Collection Time Metrics: Total I/O, Read I/O, Read Hit I/O, Read Miss I/O, Write I/O, % Read I/O, % Write I/O
EVA Storage Pool Throughput	Displays the EVA throughput statistics of a storage pool over a selected time period.	Data: Storage System Name, Storage Pool Name, Collection Time Metric: Read Hit MB/s, Read Miss MB/s, Write MB/s
EVA Storage System Request Rate	Displays the EVA storage system request (I/O) rate over a selected time period.	Data: Storage System Name, Collection Time Metrics: Total I/O, Read I/O, Write I/O, % Read I/O, % Write I/O
EVA Storage System Throughput	Displays the EVA throughputs for a storage system over a selected time period.	Data: Storage System Name, Collection Time Metrics: Read MB/s, Write MB/s

Name	Description	Performance Data
EVA Storage Volume Latency	Displays the EVA storage volume average read miss, average read hit, and average write latency over a selected time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Average Read Hit Latency, Average Read Miss Latency, Average Write Latency
EVA Storage Volume Request Rate	Displays the EVA storage volume request rate statistics over a selected time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Total I/O, Read I/O, Read Hit I/O, Read Miss I/O, Write I/O, % Read I/O, % Write I/O
EVA Storage Volume Throughput	Displays the EVA storage volume throughput over a specified period of time.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Read Miss MB/s, Read Hit MB/s, Write MB/s

Name	Description	Performance Data
EVA Top N Volumes by Activity	Displays the top N storage volumes for each storage pool based on the average of the Total I/O per storage volume for a specified period of time.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: Read Response, % Utilization, Total I/O, Read Hit I/O, Read Miss I/O, Write I/O
Performance - Storage System - HDS		
HDS Array Group Cache	Displays the HDS array group cache performance statistics over a selected time period.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Write I/O (Req/Sec), % Read Hit, % Sequential Read Hit, % Random Read Hit
HDS Array Group Request Rate	Displays the HDS array group request (I/O) rate statistics over a selected period of time.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Read Hit I/O (Req/Sec), Write I/O (Req/Sec), % Read I/O, % Write I/O

Name	Description	Performance Data
HDS Array Group Response Time	Displays the read and write response times for each HDS storage array group over a selected period of time.	Columns: Storage System Name, Array Group Name, Collection Time Metrics: Read Response Time, Write Response Time
HDS Array Group Throughput	Displays the throughput (MB/s) for each array group on HDS arrays over time.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Read (MB/s), Sequential Read (MB/s), Random Read (MB/s), Write (MB/s), Sequential Write (MB/s), Random Write (MB/s)
HDS Bottom N Array Groups by Random Read Hit Percent	Displays the bottom N HDS array groups by read hit (random) percent rate for each storage system based on the average of the % Random Read Hit statistic per array group over a selected time period. N is the number of array groups to display in the report and is specified prior to generating the report.	Data: Storage System Name, Collection Time Metrics: Array Group name, % Utilization, Total I/O, Read Hit I/O, % Random Read Hit
HDS Bottom N Array Groups by Sequential Read Hit Percent	Displays the bottom N HDS array groups by read hit (sequential) percent rate for each storage system based on the average of % Sequential Read Hit statistic per array group over a selected time period. N is the number of array groups to display in the report.	Data: Storage System Name, Collection Time Metrics: Array Group name, % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit

Name	Description	Performance Data
HDS Bottom N Volumes by Random Read Hit Percent	Displays the bottom N HDS volumes by read hit (random) percent rate for each storage pool based on the average of the % Random Read Hit statistic per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: % Utilization, Total I/O (Req/Sec), Read Hits I/O (Req/Sec), % Random Read Hit
HDS Bottom N Volumes by Sequential Read Hit Percent	Displays the bottom N HDS volumes by read hit (sequential) percent rate for each storage pool based on the average of % Sequential Read Hit statistics per storage volume over a selected period of time. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: % Utilization, Total I/O (Req/Sec), Read Hits I/O (Req/Sec), % Sequential Read Hit
HDS Storage Controller Utilization	Displays the HDS storage controller utilization percent rate based on Utilization statistics over a specified period of time.	Data: Storage System Name, Storage Controller Name, Collection Time Metrics: % MP#0 Utilization, % MP#1 Utilization, % MP#2 Utilization, % MP#3 Utilization, % MP#4 Utilization, % MP#5 Utilization, % MP#6 Utilization, % MP#7 Utilization

Name	Description	Performance Data
HDS Storage FC Port Total Request Rate	Displays the storage Fibre Channel (FC) port total I/O request rate statistics for FC ports on an HDS array over a specified time period.	Data: Storage System Name, Storage Controller Name, FC Port Name, Collection Time Metric: Total I/O
HDS Storage FC Port Total Throughput	Displays the storage Fibre Channel (FC) port total throughput of each HDS FC port over a selected time period.	Data: Storage System Name, Storage Controller Name, FC Port Name, Collection Time Metric: Total MB/s
HDS Storage Volume Cache	Displays the cache hit details for each HDS storage volume over a selected time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time, Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Write I/O (Req/Sec), % Read Hit , % Sequential Read Hit , % Random Read Hit
HDS Storage Volume Request Rate	Displays the request (I/O) rate for each HDS storage volume over a selected time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Read Hit I/O (Req/Sec), Write I/O (Req/Sec), % Read I/O, % Write I/O

Name	Description	Performance Data
HDS Storage Volume Response Time	Displays the response time for each HDS storage volume over a selected period of time.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Read Response Time, Write Response Time
HDS Storage Volume Throughput	Displays the throughput for each HDS storage volume based on the Read Response Time and Write Response Time statistics for the storage volumes over a specified time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Read Response Time, Write Response Time
HDS Top N Array Groups by Read Response Time	Displays the top N HDS array groups based on the average of Read Response Time per array group over a selected time period. N is the number of array groups to display in the report and is specified prior to generating the report.	Data: Array Group Name Metrics: Read Response Time, % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit, % Random Read Hit
HDS Top N Array Groups by Utilization	Displays the top N HDS array groups based on the average of the % Utilization statistic per array group over a selected time period. N is the number of array groups to display in the report.	Data: Array Group Name Metrics: % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit, % Random Read Hit

Name	Description	Performance Data
HDS Top N Array Groups by Write Response Time	Displays the top N HDS array groups based on the average of Write Response Time statistics per array group over a selected time period. N is the number of array groups to display in the report.	Data: Array Group Name Metrics: Write Response Time, % Utilization, Total I/O
HDS Top N Storage FC Port Total Request Rate	Displays the top N FC ports for each storage controller on HDS arrays based on the average of Total I/O per FC port over a selected time period. N is the number of ports to display in the report.	Data: Storage System Name, Storage Controller Name, FC Port Name Metric: Total I/O
HDS Top N FC Port Total Throughput	Displays the top N Fibre Channel (FC) ports on HDS arrays based on total throughput per FC Port over a selected time period. N is the number of ports to display in the report.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metric: Total MB/s
HDS Top N Volumes by Read Response Time	Displays the top N storage volumes for each HDS storage pool based on the average of Read Response Time statistics per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: Read Response Time, % Utilization, Total I/O, Read I/O, % Read Hit I/O, % Sequential Read Hit, % Random Read Hit

Name	Description	Performance Data
HDS Top N Volumes by Utilization	Displays the top N storage volumes for each HDS storage pool based on the average of the % Utilization statistic per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: % Utilization, Total I/O, Read I/O, % Read Hit I/O, % Sequential Read Hit, % Random Read Hit
HDS Top N Volumes by Write Response Time	Displays the top N storage volumes for each HDS storage pool, based on the average of Write Response Time statistics per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: Write Response Time, % Utilization, Total I/O
Storage System - XP		
XP Array Group Cache	Displays the XP array group cache performance statistics over a selected time period.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Write I/O (Req/Sec), % Read Hit, % Sequential Read Hit, % Random Read Hit

Name	Description	Performance Data
XP Array Group Request Rate	Displays the XP array group I/O request rate statistics over a selected time period.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Read Hit I/O (Req/Sec), Write I/O (Req/Sec), % Read I/O, % Write I/O
XP Array Group Response Time	Displays the read and write response times for each XP storage array group over a selected period of time.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Read Response Time, Write Response Time
XP Array Group Throughput	Displays the throughput (MB/s) for each array group on XP arrays over a selected period of time.	Data: Storage System Name, Array Group Name, Collection Time Metrics: Read (MB/s), Sequential Read (MB/s), Random Read (MB/s), Write (MB/s), Sequential Write (MB/s), Random Write (MB/s)
XP Bottom N Array Groups by Random Read Hit Percent	Displays the bottom N XP array groups by read hit (random) percent rate for each storage system based on the average of the % Random Read Hit statistic per array group over a selected time period. N is the number of array groups to display in the report and is specified prior to generating the report.	Data: Storage System Name, Collection Time Metrics: Array Group name, % Utilization, Total I/O, Read Hit I/O, % Random Read Hit

Name	Description	Performance Data
XP Bottom N Array Groups by Sequential Read Hit Percent	Displays the bottom N XP array groups by read hit (sequential) percent rate for each storage system based on the average of % Sequential Read Hit statistic per array group over a selected time period. N is the number of array groups to display in the report.	Data: Storage System Name, Collection Time Metrics: Array Group Name, % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit
XP Bottom N Volumes by Random Read Hit Percent	Displays the bottom N XP volumes by read hit (random) percent rate for each storage pool based on the average of the % Random Read Hit statistic per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time. Metrics: % Utilization, Total I/O (Req/Sec), Read Hits I/O (Req/Sec), % Random Read Hit
XP Bottom N Volumes by Sequential Read Hit Percent	Displays the bottom N XP volumes by read hit (sequential) percent rate for each storage pool based on the average of % Sequential Read Hit statistics per storage volume over a selected period of time. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: % Utilization, Total I/O (Req/Sec), Read Hits I/O (Req/Sec), % Sequential Read Hit

Name	Description	Performance Data
XP Storage Controller Utilization	Displays the XP storage controller utilization percent rate based on Utilization statistics over a specified period of time.	Data: Storage System Name, Storage Controller Name, Collection Time Metrics: % MP#0 Utilization, % MP#1 Utilization, % MP#2 Utilization, % MP#3 Utilization, % MP#4 Utilization, % MP#5 Utilization, % MP#6 Utilization, % MP#7 Utilization
XP Storage FC Port Total Request Rate	Displays the storage Fibre Channel (FC) port total I/O request rate statistics for FC ports on an XP array over a specified time period.	Data: Storage System Name, Storage Controller Name, FC Port Name, Collection Time Metric: Total I/O
XP Storage FC Port Total Throughput	Displays the storage Fibre Channel (FC) port total throughput of each XP FC port over a selected time period.	Data: Storage System Name, Storage Controller Name, FC Port Name, Collection Time Metric: Total MB/s

Name	Description	Performance Data
XP Storage Volume Cache	Displays the cache hit details for each XP storage volume over a selected time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time, Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Write I/O (Req/Sec), % Read Hit , % Sequential Read Hit, % Random Read Hit
XP Storage Volume Request Rate	Displays the request (I/O) rate over a selected time period for each XP storage volume.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Total I/O (Req/Sec), Read I/O (Req/Sec), Read Hit I/O (Req/Sec), Write I/O (Req/Sec), % Read I/O, % Write I/O
XP Storage Volume Response Time	Displays the response time for each XP storage volume over a selected period of time.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Read Response Time, Write Response Time

Name	Description	Performance Data
XP Storage Volume Throughput	Displays the throughput for each XP storage volume based on the Read Response Time and Write Response Time statistics for the storage volumes over a specified time period.	Data: Storage System Name, Storage Pool Name, Storage Volume Name, Collection Time Metrics: Data:
XP Top N Array Groups by Read Response Time	Displays the top N XP array groups based on the average of Read Response Time per array group over a selected time period. N is the number of array groups to display in the report and is specified prior to generating the report.	Data: Array Group Name Metrics: Read Response Time, % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit, % Random Read Hit
XP Top N Array Groups by Utilization	Displays the top N XP array groups based on the average of the % Utilization statistic per array group over a selected time period. N is the number of array groups to display in the report.	Data: Array Group Name Metric: % Utilization, Total I/O, Read Hit I/O, % Sequential Read Hit, % Random Read Hit
XP Top N Array Groups by Write Response Time	Displays the top N XP array groups based on the average of Write Response Time statistics per array group over a selected time period. N is the number of array groups to display in the report.	Data: Array Group Name Metrics: Write Response Time, % Utilization, Total I/O
XP Top N Storage FC Port Total Request Rate	Displays the top N FC ports for each storage controller on XP arrays based on the average of Total I/O per FC port over a selected time period. N is the number of ports to display in the report.	Data: Storage System Name, Storage Controller Name, FC Port Name Metric: Total I/O

Name	Description	Performance Data
XP Top N FC Port Total Throughput	Displays the top N Fibre Channel (FC) ports for XP arrays based on total throughput per FC port over a selected time period. N is the number of ports to display in the report.	Data: Storage System Name, Storage Processor Name, FC Port Name, Collection Time Metric: Total MB/s
XP Top N Volumes by Read Response Time	Displays the top N storage volumes for each XP storage pool based on the average of Read Response Time statistics per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: Read Response Time, % Utilization, Total I/O, Read I/O, % Read Hit I/O, % Sequential Read Hit, % Random Read Hit
XP Top N Volumes by Utilization	Displays the top N storage volumes for each XP storage pool, based on the average of the % Utilization statistic per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: % Utilization, Total I/O, Read I/O, % Read Hit I/O, % Sequential Read Hit, % Random Read Hit
XP Top N Volume by Write Response Time	Displays the top N storage volumes for each XP storage pool, based on the average of Write Response Time statistics per storage volume over a selected time period. N is the number of volumes to display in the report.	Data: Storage System Name, Storage Pool Name, Storage Volume Name Metrics: Write Response Time, % Utilization, Total I/O

Name	Description	Performance Data
Switch Performance		
Switch Port Input and Output Performance	Lists the switch port I/O performance statistics over a selected time period.	User selects the specific performance metrics and switch ports for which they want performance I/O statistics. Report columns and metrics are not predetermined.

*NetApp performance reports are limited to *physical* filers. Do not select virtual filers, because there is no data to return for these devices.

Running Standard Performance Reports

HP Storage Essentials provides a large number of pre-designed performance reports that allow you to view and monitor performance of specific devices and storage elements in your environment. For a list of Storage Essentials standard reports, see [List of Standard Performance Reports on page 189](#).

The first time you run a report, data is not displayed until you click the **Refresh Data** button. Many reports prompt you to filter report data by prompting you to specify a start and end date.

The generated report displays the historical data collected during the period you specify. However, based on the report statistics selected for a report, the report roll-up schedules, and garbage collection, you might find data not charting for the entire time period you've selected. Granularity of data varies depending on the data rollup settings selected. You can find the settings for Garbage Collection and Data Rollup on the Data Aging page (**Configuration > Reports > Data Aging** tab). See [Viewing Data Aging Settings for Performance Statistics on page 56](#) for more information.

Note: You can click the **Refresh All** button (upper right corner) to display the Prompts dialog, which enables you to enter the begin and end dates for the reported data, as well as other information to be displayed in the report. Click **Run Query** to start report generation.

Generating Performance Reports for Storage Systems

The following standard storage system reports are provided by Report Optimizer:

- EVA Storage System Throughput
- EVA Storage System Request Rate

To run standard performance reports for storage systems:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance data and click the > button. Do not select All because this results in a blank report.
5. Click **Run Query**.

Generating Performance Reports for Physical Disk Devices

The following standard reports are provided for physical disk devices on your EVA arrays:

- EVA Physical Disk Throughput
- EVA Physical Disk Request Rate
- EVA Physical Disk Latency
- EVA Physical Disk Average Queue Depth
- EVA Physical Disk Average Drive Latency

To run standard performance reports for physical disk devices:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance data and click the > button. Do not select All because this results in a blank report.
5. Select the disk drives to include in the report and click the > button. Do not select All because this results in a blank report.
6. Click **Run Query**.

Generating Performance Reports for Controllers

The following standard reports are provided for the controllers on your EVA array:

- EVA Storage Controller Throughput
- EVA Storage Controller Request Rate
- EVA Storage Controller Latency
- EVA Storage Controller Data Transfer
- EVA Storage Controller CPU Utilization

To run standard performance reports for controller devices:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance data and click the > button. Do not select All because this results in a blank report.
5. Select the storage controllers for which you would like to view performance statistics, and click the > button.
6. Click **Run Query**.

Generating Performance Reports for Fibre Channel Ports

The following standard reports are provided for Fibre Channel (FC) ports:

- EVA Storage FC Port Request Rate
- XP Storage FC Port Total Request Rate
- HDS Storage FC Port Total Request Rate
- EVA Storage FC Port Throughput
- XP Storage FC Port Total Throughput
- HDS Storage FC Port Total Throughput
- EVA Storage FC Port Latency
- EVA Storage FC Port Average Queue Depth

To run standard reports for Fibre Channel ports:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage controllers to view in the report and click the > button. Do not select All because this results in a blank report.
6. Select the FC ports for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
7. Click **Run Query**.

Generating Performance Reports for Top N Fibre Channel Ports by Metric

The following standard reports display a specified number of Fibre Channel Ports for a performance metric, such as highest total request rate for ports:

- EVA - *No standard report*
- XP Top N FC Port Total Throughput
- XP Top N Storage FC Port Total Request Rate
- HDS Top N Storage FC Port Total Request Rate

To run standard reports for top N Fibre Channel ports by metric:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage controllers to view in the report and click the > button. Do not select All because this results in a blank report.
6. Enter a number to indicate how many ports you want displayed in the report. The report selects those ports with the highest statistics for the performance metric and displays only the number (top downward) that you specify.
7. Click **Run Query**.

Generating Performance Reports for Storage Pools

The following standard reports are provided for EVA storage pools:

- EVA Storage Pool Throughput
- EVA Storage Pool Request Rate
- EVA Storage Pool Latency

To run standard performance reports for storage pools:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage pools for which you would like to view performance statistics, and click the > button. Do not select All because this results in a blank report.
6. Click **Run Query**.

Generating Performance Reports for Volumes

The following standard reports are provided for volumes:

- EVA Storage Volume Latency
- XP Storage Volume Cache
- HDS Storage Volume Cache
- EVA Storage Volume Throughput
- XP Storage Volume Request Rate
- HDS Storage Volume Request Rate
- EVA Storage Volume Request Rate
- XP Storage Volume Response Time
- HDS Storage Volume Response Time

To run standard performance reports for volumes:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage pools for which you would like to view performance statistics, and click the > button. Do not select All because this results in a blank report.
6. Select the storage volumes you want to view in the report and click the > button. Do not select All because this results in a blank report.
7. Click **Run Query**.

Generating Performance Reports for Top N Volumes by Metric

The following standard reports are provided for top N volumes by metric:

- EVA Top N Volumes by Activity
- XP Top N Volumes by Read Response Time
- HDS Top N Volumes by Read Response Time
- XP Top N Volumes by Utilization
- HDS Top N Volumes by Utilization
- XP Top N Volumes by Write Response Time
- HDS Top N Volumes by Write Response Time

To run standard performance reports for top N volumes by metric:

1. Select a start date.
2. Select an end date.

3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage pools for which you would like to view performance statistics, and click the > button. Do not select All because this results in a blank report.
6. Enter the number of volumes (top down) to display in the report.
7. Click **Run Query**.

Generating Performance Reports for Bottom N Volumes by Metric

The following standard reports are provided for bottom N volumes by activity:

- EVA - *No standard report*
- XP Bottom N Volumes by Random Read Hit Percent
- HDS Bottom N Volumes by Random Read Hit Percent
- XP Bottom N Volumes by Sequential Read Hit Percent
- HDS Bottom N Volumes by Sequential Read Hit Percent

To run standard performance reports for top N volumes by metric:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the storage pools for which you would like to view performance statistics, and click the > button. Do not select All because this results in a blank report.
6. Enter the number of volumes (bottom up) to display in the report.
7. Click **Run Query**.

Generating Performance Reports for Array Groups

The following standard reports are provided for array groups:

- XP Array Group Cache
- HDS Array Group Cache
- XP Array Group Request Rate
- HDS Array Group Request Rate
- XP Array Group Response Time
- HDS Array Group Response Time

To run standard performance reports for storage pools:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Select the array group you want to view in the report and click the > button. Do not select All because this results in a blank report.
6. Click **Run Query**.

Generating Performance Reports for Top N Array Groups by Metric

The following standard reports are provided for top N array groups by activity:

- EVA - *No standard report*
- XP Top N Array Groups by Read Response Time
- HDS Top N Array Groups by Read Response Time
- XP Top N Array Groups by Utilization Definition
- HDS Top N Array Groups by Utilization Definition
- XP Top N Array Groups by Write Response Time
- HDS Top N Array Groups by Write Response Time

To run standard performance reports for top N array groups by metric:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Enter the number of array groups (top down) that you want displayed in the report.
6. Click **Run Query**.

Generating Performance Reports for Bottom N Array Groups by Metric

The following standard reports are provided for bottom N array groups:

- EVA - *No standard report*
- XP Bottom N Array Groups by Random Read Hit Percent
- HDS Bottom N Array Groups by Random Read Hit Percent

- XP Bottom N Array Groups by Sequential Read Hit Percent
- HDS Bottom N Array Groups by Sequential Read Hit Percent

To run standard performance reports for bottom N array groups by metric:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the storage systems for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Enter the number of array groups (bottom up) to display in the report.
6. Click **Run Query**.

Generating Performance Reports for Switches

The following standard report is provided for switch ports:

- Switch Port Input and Output Performance

To run a standard performance report for switches:

1. Select a start date.
2. Select an end date.
3. Select a statistic type. The available options are: Raw, Hourly, Daily, Weekly, and Monthly.
4. Select the switches for which you want to view performance statistics and click the > button. Do not select All because this results in a blank report.
5. Click **Run Query**.

Backup Performance Forecast Reports

Forecast reports use historical data points to calculate data trends for some pre-determined period of time in the future.

Report Optimizer provides two Backup performance forecast reports:

- Client Performance Forecast
- Library Performance Forecast

Both reports display data for a future period of fifteen days.

Forecast Calculation

Trending is calculated based on a linear regression algorithm. Reports require a minimum of two data entries for calculation. If there is only one data point, the Forecast is the same as History (based on historical data only). If there is no change in historical values (or the slope calculated is zero), the graphed trend shows the average of the historical values (a straight line). If the trend goes below zero (negative values), the Forecast is restricted to zero. If there is no history for the last seven days, only Forecast data is shown in the report graph.

Forecasted data is calculated using the following formula:

$$\text{Projected Total} = ((\langle \text{Future_Month} \rangle - \text{to_date}('01\text{-JAN-1970}', 'dd\text{-mm-} \\ \text{yyyy}')) - [\text{Intercept}]) / [\text{Slope}]$$

In this instance:

- Intercept = ((Sum of Y) - (Slope x (Sum of X))) / N
- Slope = Numerator / Denominator ((Sum of Y) - (Slope x (Sum of X))) / N
- N = count(collectiontime)
- Y = collectiontime - to_date('01-JAN-1970', 'dd-mm-yyy')
- X = totalGB
- Numerator = (N x Sum(X x Y)) - ((Sum of X) x (Sum of Y))
- Denominator = (N x Sum of (X x X)) - ((Sum of X) x (Sum of X))

Performance Forecast Report List


The following table lists the Backup throughput performance reports provided as standard reports in Report Optimizer with forecast trending information. To access these reports select **Document List > Public Folders > Report Pack > Standard Reports**.

The first time you open a report, only the report title displays without data. To see report data, click **Refresh All** in the upper toolbar. You are then prompted to enter parameters that determine which data to use to populate the report.

The following table shows where to find a specific forecast trending report in Standard Reports. For example, if the header shows "Backup - Forecast," expand the Backup node in the navigation tree and then the Forecast node. You should see the list of Backup forecast reports.

Report Name	Description
Backup - Forecast	
Client Performance Forecast	Displays future trends in Backup throughput (I/O) forecasted for the next fifteen (15) days.
Library Performance Forecast	Displays future trends in library throughput (I/O) forecasted for the next fifteen (15) days.

The following figure shows the Client Performance Forecast report. The chart shows 7 days of historical data and 15 days of forecasted data.



For information about generating custom trending reports and charts, see [Creating Performance Reports with Drill Down and Trending on page 257](#) and [Viewing Trending Information for Performance on page 86](#)



About Creating Simple Custom Reports


Creating a simple custom report using Report Optimizer requires the following tasks:


- Task 1 – Select Dimension and Measure objects you want included in the report. These are described below.
- Task 2 – Select Filter objects you want as conditions for filtering the report results.

The following example shows how to accomplish the tasks.

Task 1 – Select Dimension and Measure Objects

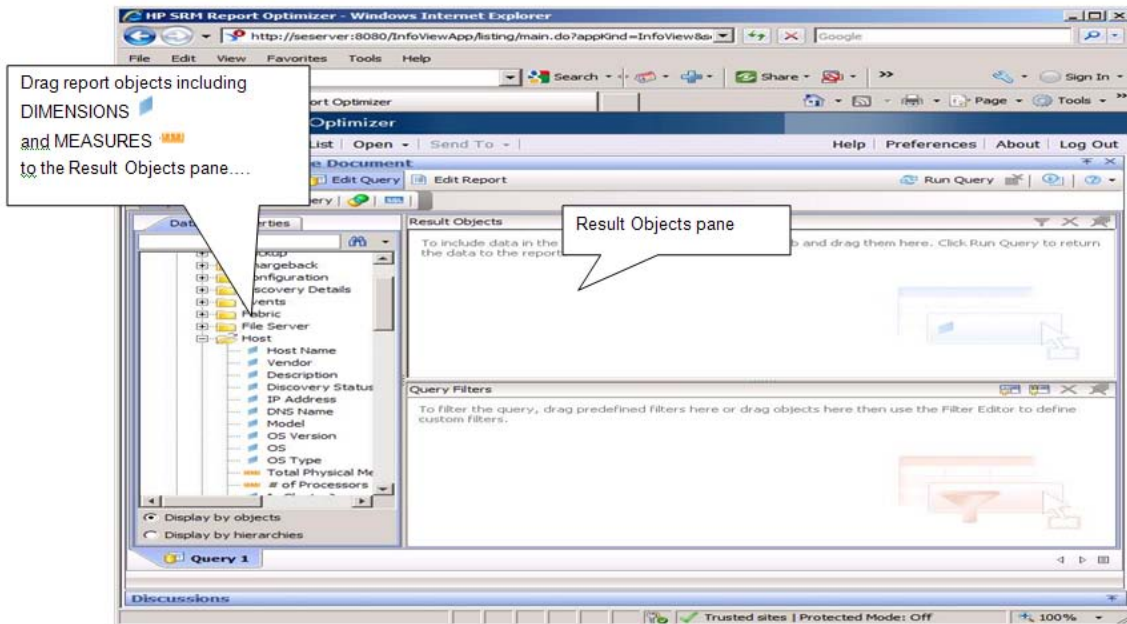
HP Storage Essentials provides roughly 3000 database objects in the form of Dimensions () and Measures () that you can use to construct your report. After you create a new Web Intelligence Document (see [Creating Custom Performance Reports on page 236](#), you can find the database objects listed in the left pane of the New Web Intelligence Document window.

A Dimension object is any character-based object that is indicated by the symbol . For performance reports, you will be interested in Dimension objects such as Storage System Name, FC Port Name, Disk Drive Name, and so on. When selected, these objects are displayed as column headers in your report. These objects have properties, which you will be asked to specify for each object as appropriate. For example, if you select Storage System Name, you also must specify which storage system to display. Report Optimizer prompts you for this information.


A Measure object is any number-based object indicated by the symbol . For performance reports, performance metrics are Measure objects; for example, Total I/O Rate, % Utilization, Write Data Rate, Read Rate, and so on. HP Storage Essentials provides one Measure object for each performance metric supported. To create a custom performance report, you select any Measure object you want included in a report. For example, to create a performance report for your host bus adapters (HBA) that shows the number of link failures, you would include the Link Failures object in the Results Objects list. A report generates, showing the Link Failure performance data for each HBA.

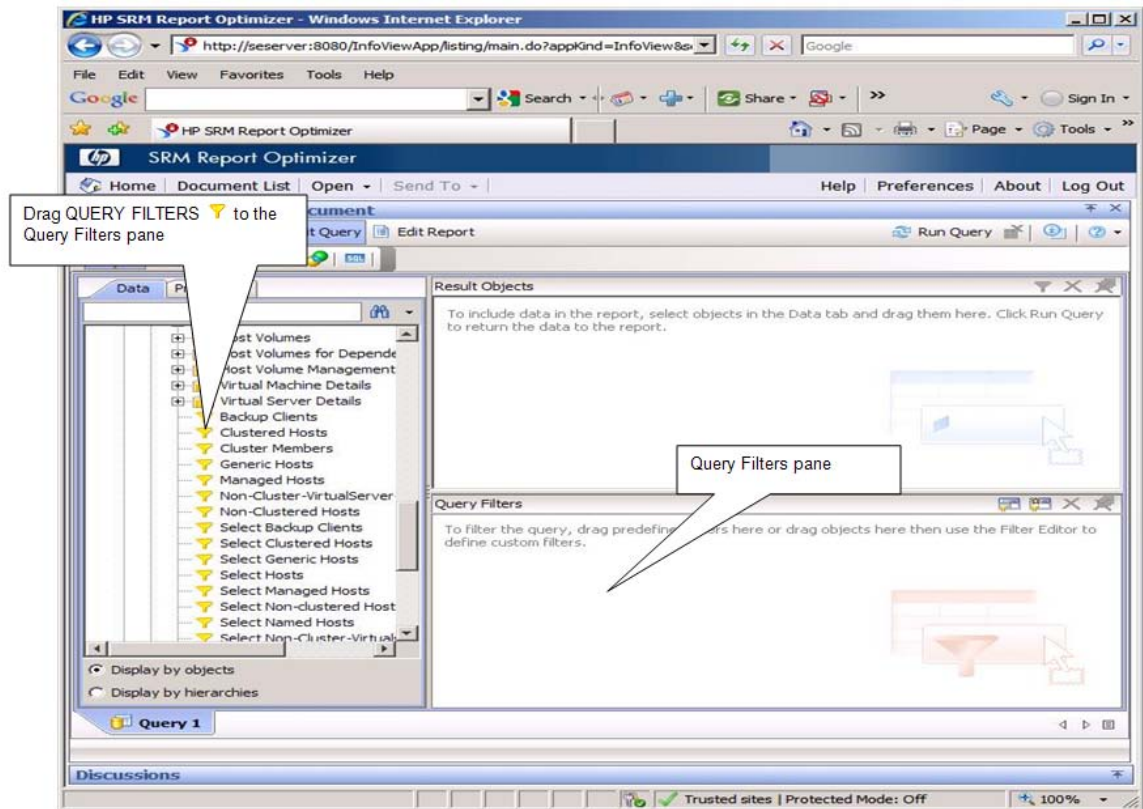
To use Measure objects effectively, make sure that data is being actively collected for that metric. Check your data collectors and collection schedules to ensure that you are gathering data for the time period you need.

To design a performance report, drag a Dimension object or a Measure from the Data tab into the Results Object pane on the right. This builds an SQL query that is sent to the Storage Essentials database to retrieve a result, which then displays in your report.



Task 2 – Select Filter Objects

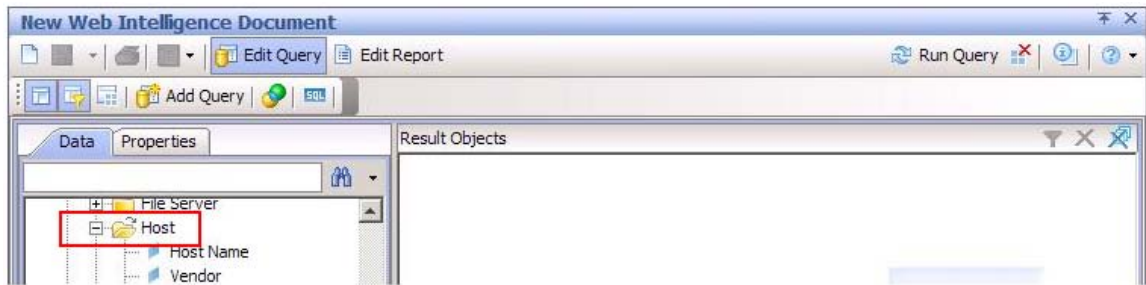
A Query Filter is an object with a symbol  that is used to extract certain results based on a condition. To add filters, drag a Query Filter object from the Data tab on the left to the Query Filter pane on the bottom right. Use only filters from the categories (host, switch, storage system, and so on) from which you want to retrieve performance data.



Example: Host Performance Report

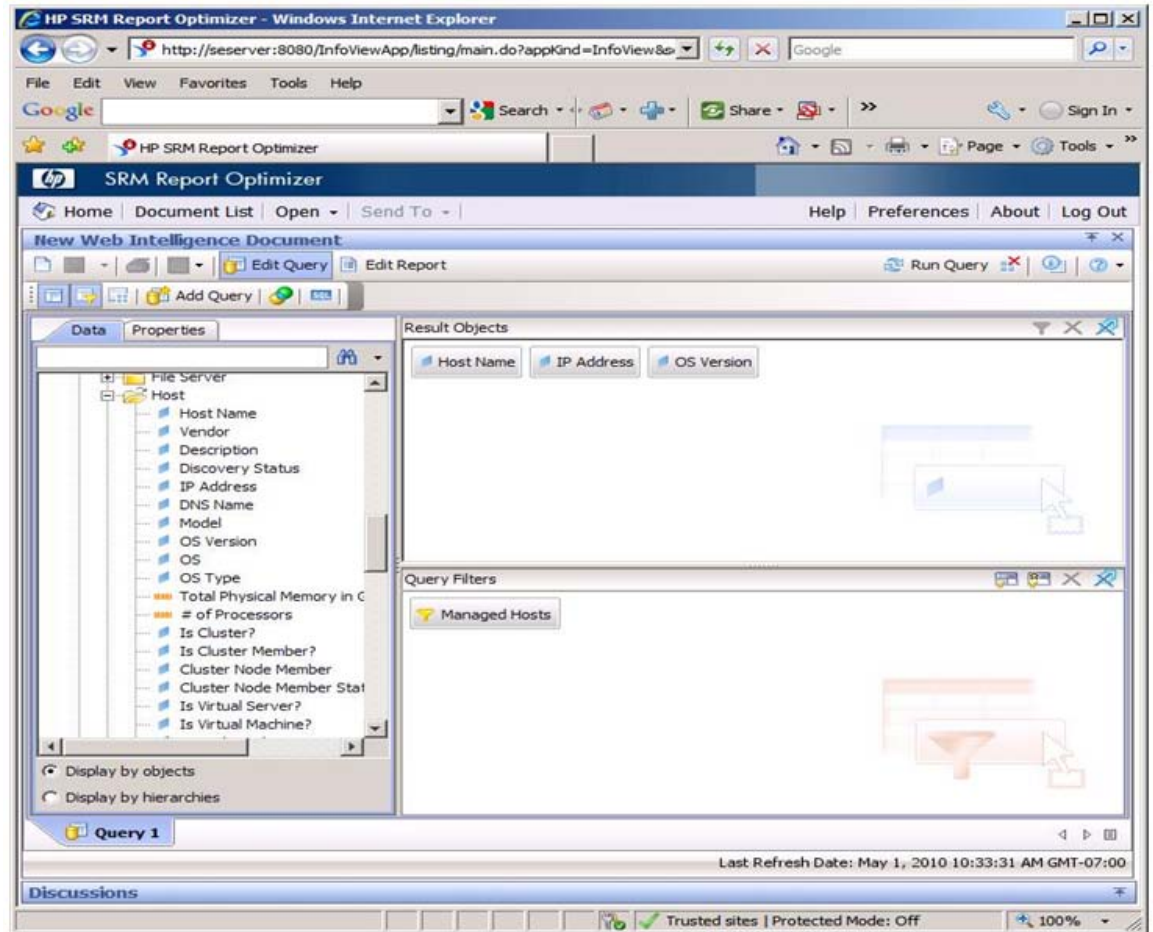
The following example shows how to create a basic HOST report that displays Host Name, IP Address, and OS Version, and is then filtered on Managed Hosts.

1. Select the Host category from the Data tab on the left and expand the tree to display the Host objects.

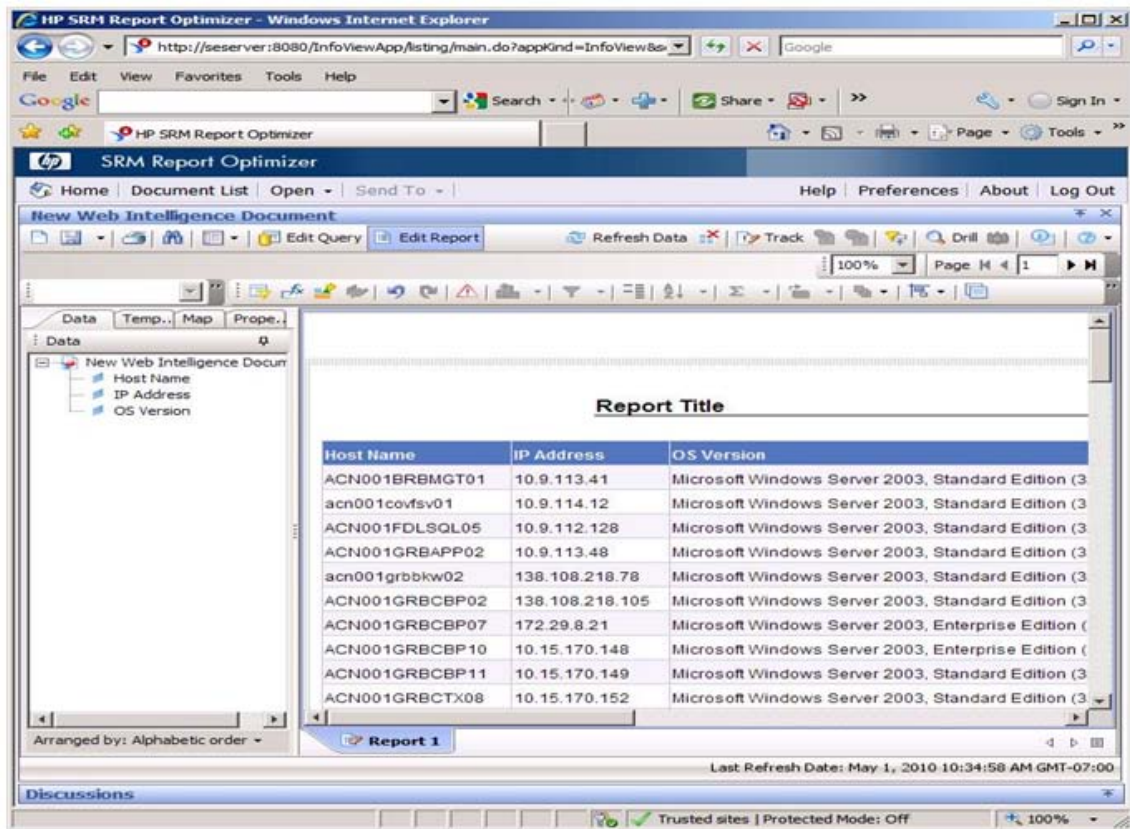


2. Select the objects for host name, IP address, and OS version to include in your report:
 - a. Drag or Double-click **Host name** from the Data tab on the left to the Results Objects pane on the right.
 - b. Drag or Double-click **IP Address** to the Results Object pane.
 - c. Drag or Double-click **OS version** to the Results Object pane.
3. Drag or Double-click **Managed Hosts** from the Data tab on the left to the Query Filters pane on the right.

The Web Intelligence Document should look like the following example:



4. When ready, click **Run Query** on the toolbar at the upper right.
5. The Web Intelligence Document page displays the data from the results of the report query. You can edit the report by clicking **Edit Report** in the upper toolbar. The Data tab on the left displays the Result Objects selected for the report.



The report is filtered on Managed Hosts. Only the hosts that Storage Essentials is managing using an agent (CIM Extension) or a VMware Virtual Center/ ESX server are displayed.

To go back and edit the query for this report, click **Edit Query**.

Creating Custom Performance Reports

Before you begin to create a custom performance report, verify that performance data collectors exist for each performance metric you want displayed in the report. Make sure these are enabled and actively collecting data (check collection schedules to see when data was last collected). To see the most recent data (that is, data which reflects statistics since the last collection), do a Refresh Data. Be sure to wait for the refresh to complete before generating your report.

Make sure that Storage Essentials is actively collecting the performance data that you want to display in your reports. To review data collectors and schedules, open the Performance Manager and go to **Configuration > Performance > Data Collection**. For more information, see [Displaying the Performance Data Collectors Page on page 50](#).

When you are ready, perform these four simple steps to create a custom performance report. Instructions for each step follow.

Step 1 —Create an empty Web Intelligence Document that you will use to format your web performance report.

Step 2—Connect to the HP Storage Essentials Universe database using the Report Connector. This enables you to view the objects in the database that you can select for your reports.

Step 3—After connecting to the database, open the Performance folder and select the objects you want to include in your report.

Step 4—Run the report query using the performance objects and filters selected in Step 3.

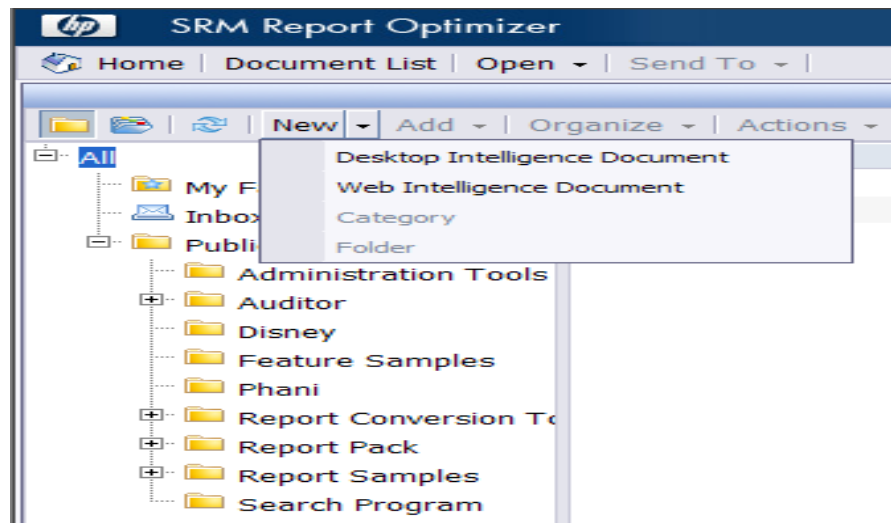
The following instructions assume that you opened the Report Optimizer InfoView and are ready to create a report. See [Accessing InfoView Logon on page 187](#) for instructions about how to open the SRM Report Optimizer.

Create a Web Intelligence Document

To create a Web Intelligence Document:

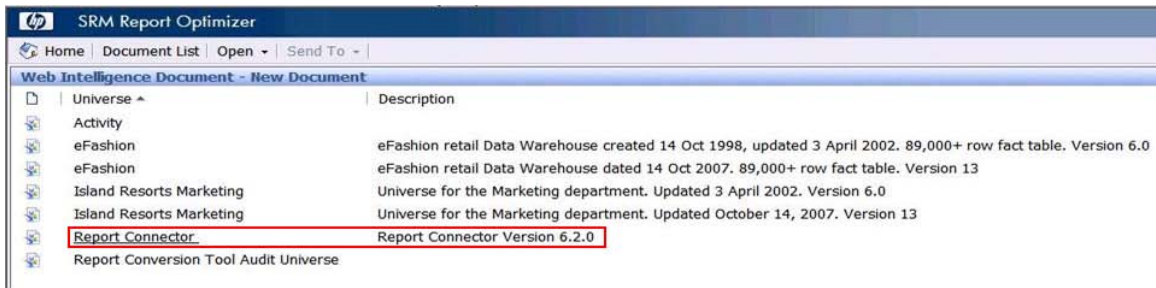
1. From Report Optimizer InfoView, under Navigate (left list) click **Document List**.
2. In the menu bar, click **New** and select **Web Intelligence Document** from the drop-down menu.

You can also navigate to Performance Reports in the Document List tree (**Public Folders > Report Packs > Standard Reports > Performance Reports**), right-click **Performance Reports**, and select **New > Web Intelligence Document**.



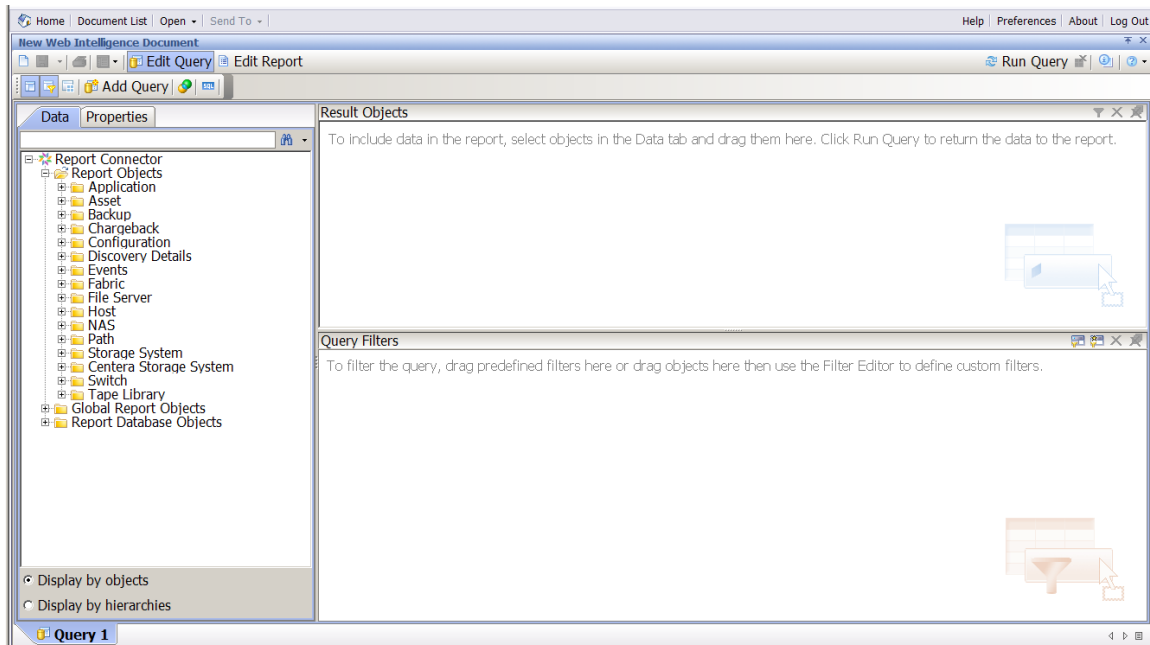
Connect to HP Storage Essentials Universe Database

With the New Web Intelligence Document page open, connect to the HP Storage Essentials Universe database by clicking **Report Connector**. The Report Connector connects you to the HP Storage Essentials Universe (database). The Report Connector should be version 6.2.0 or higher.



Select Objects to Include in Report

After you connect to the HP Storage Essentials Universe database, the SRM Report Optimizer displays a list of business objects from which you can create a query to generate your performance report. For an overview of how to create custom performance reports using these objects, see [About Creating Simple Custom Reports](#) on page 231.



When selecting to view information for devices in a storage system, as a best practice, include the Storage System Name object as well as the device category name, such as FC Port Name or Disk Drive Name. If you decide to filter data based on a specific selection of devices (such as the Select Disk Drives filter), Report Optimizer prompts you to specify those devices for which you want to display performance data.

Run the Report Query

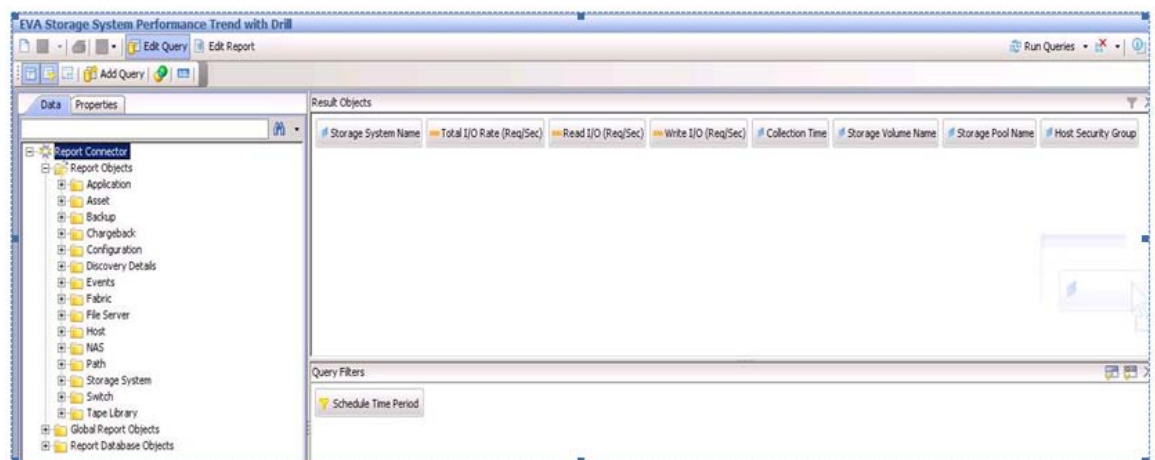
When you are ready, click **Run Query** to generate the performance report.

Example: EVA to Host Performance Report

The following example shows EVA to Host performance statistics that relate to I/O per sec performance. The report will display the storage system, volume, and storage pool names, along with host security groups. A trend chart for each host will be produced.

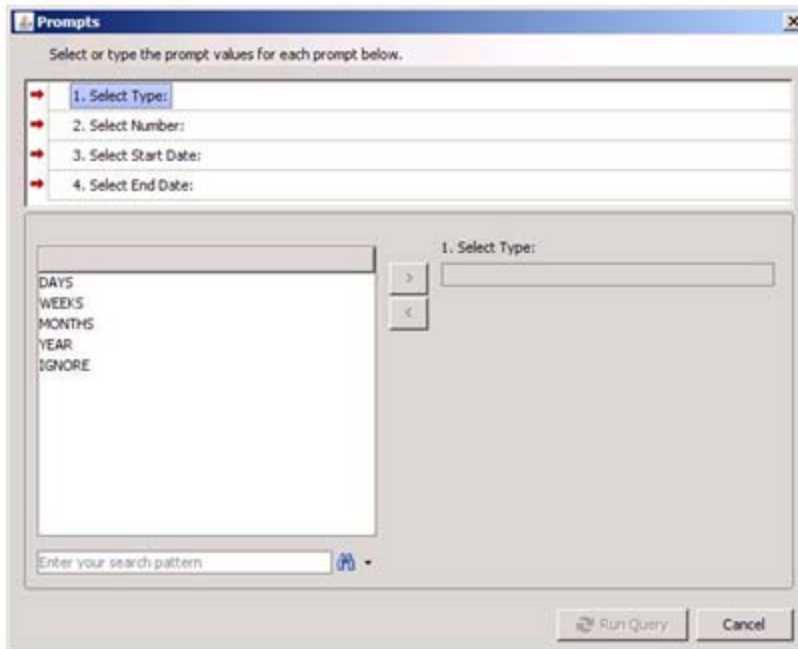
1. Log on to InfoView. See [Accessing InfoView Logon on page 187](#).
2. Create a new report by clicking **New Web Intelligence Document**. See [Create a Web Intelligence Document on page 237](#).
3. Expand the **Storage System** category.
4. Expand **Performance Statistics > EVA > EVA Storage Volume Statistics - Historical**.
5. Drag or double-click the following data objects to the Result Objects pane:
 - **Storage System Name**
 - **Total I/O Rate (Req/Sec)**
 - **Read I/O (Req/Sec)**
 - **Write I/O (Req/Sec)**
 - **Collection Time**
 - **Storage Volume Name**
 - **Storage Pool Name**
 - **Host Security Group**
6. Scroll down to the bottom of the EVA Storage Volume Statistics - Historical category and drag or double-click **Schedule Time Period**.

The resulting Web Intelligence Document should look like the following example:



7. When ready, click **Run Query** in the toolbar in the upper right.

- Enter prompt values in the Prompts dialog.



The performance report appears.

Storage System No	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	Collection Time	Storage Volume No	Storage Pool Name	Host Security Group
ctb_mia_eva_01	2.32	2.32	0	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLDEP
ctb_mia_eva_01	2.41	2.41	0	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLDEP
ctb_mia_eva_01	0.63	0.63	0	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLDEP
ctb_mia_eva_01	1.16	1.16	0	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLDEP
ctb_mia_eva_01	57.74	35.66	22.08	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLNXP
ctb_mia_eva_01	67.74	35.66	32.08	5/10/10	CTX_MIA_HPUXFT	Default Disk Group	CTX_MIA_FTLNXP

Creating an EVA Performance Report

Creating trending information for an EVA-to-host performance report requires the following tasks:

- Task 1 – Design the report by specifying the Dimension and Measure objects in the Results Object pane and the Filter objects in the Query Filter pane. When ready, run the report query. Depending on your type of report, you may also need to select the query context, answer prompts, and run the report query again.
- Task 2 – Redesign the report format by removing the generated table and reformatting the report structure by inserting a new table using report templates and the Formula Editor.
- Task 3 – Insert additional tables to the report as required.

The following example shows how to accomplish the tasks.

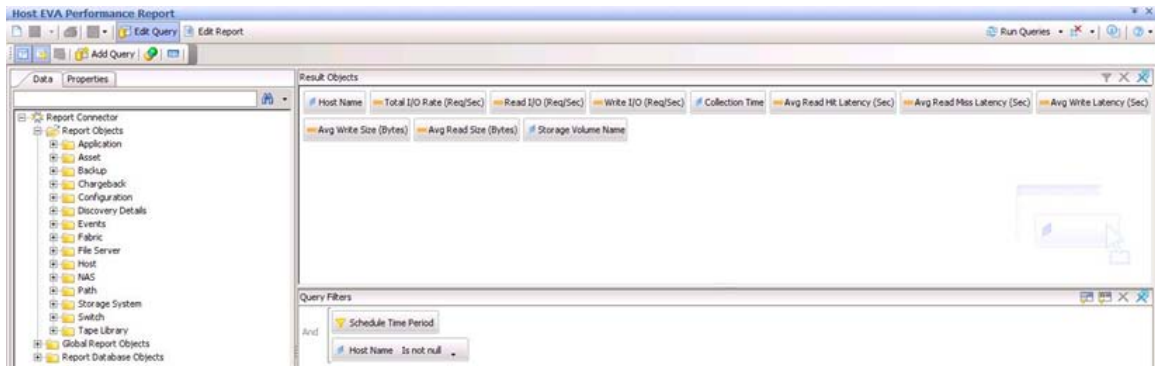
Example: EVA to Host Performance Report

This example shows you how to create a report that displays EVA to host performance statistics. The example captures the collected data for I/Os per second, latency, and the data rate. It places this information into a trending chart for each host.

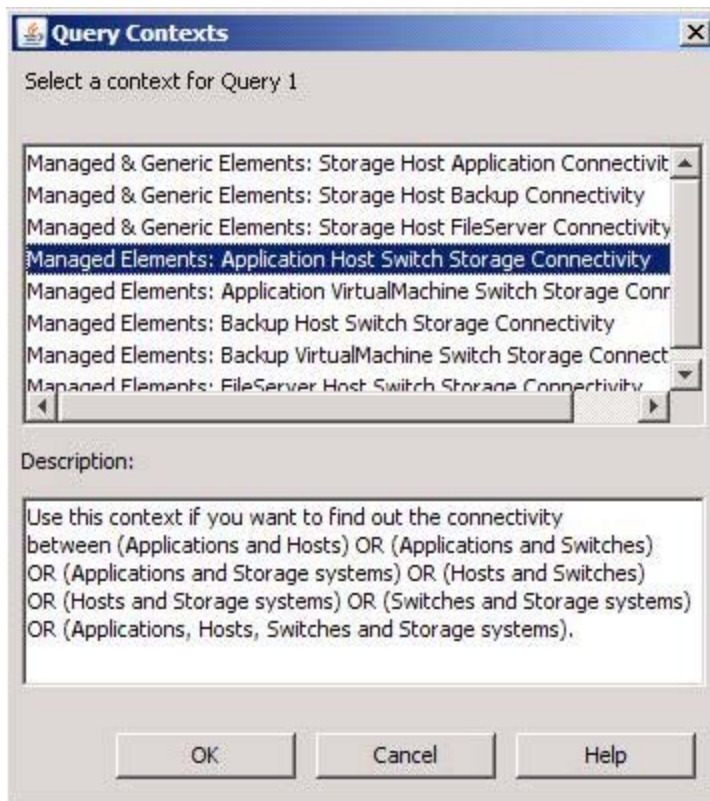
The following instructions assume that you logged on to Report Optimizer, created a Web Intelligence Document, and connected to the HP Storage Essentials Universe database. See [Creating Custom Performance Reports on page 236](#) for instructions.

Task 1 – Design and Generate the Basic Report

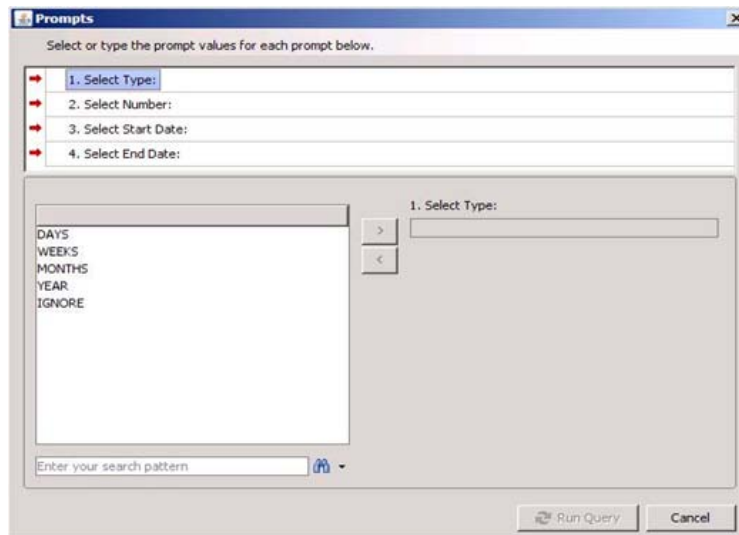
1. With the Data tab selected in the left pane, navigate to the **Host** folder and expand the object tree to display all Host objects.
2. Drag or double-click **Host Name** into the Results Object pane on the right.
3. Navigate to the Storage System folder and expand the object tree to display objects.
4. Drag or double-click **Storage System Name** into the Results Object pane.
5. Navigate to the Storage System folder and expand along this tree path: **Performance Statistics > EVA > EVA Storage Volume Statistics - Historical**.
6. Drag or double-click the following Measure objects into the Results Object pane on the right.
 - **Total I/O Rate (Req/Sec)**
 - **Read I/O Rate (Req/Sec)**
 - **Read I/O Rate (Req/Sec)**
 - **Avg Read Hit Latency (Sec)**
 - **Avg Read Miss Latency (Sec)**
 - **Avg Write Latency (Sec)**
 - **Avg Read Size (Bytes)**
 - **Avg Write Size (Bytes)**
7. Drag or double-click the following Dimension object into the Results Object pane on the right:
Collection Time
8. Scroll to the bottom of the Storage Volume Statistics - Historical folder and drag the following object to the Query Filters pane:
Schedule Time Period
9. Return to the Host folder and drag **Host Name** to the Query Filters pane.



10. Click **Run Query** in the upper right toolbar.
11. Select **Query Contexts**. The Query Contexts dialog opens.
You must select the query context because the Host Name provider and the Storage System provider are co-located on the same report query.
12. Select **Managed Elements: Application Host Switch Storage Connectivity** and click **OK**.



13. When the Prompts dialog appears, enter answers to prompts as appropriate.



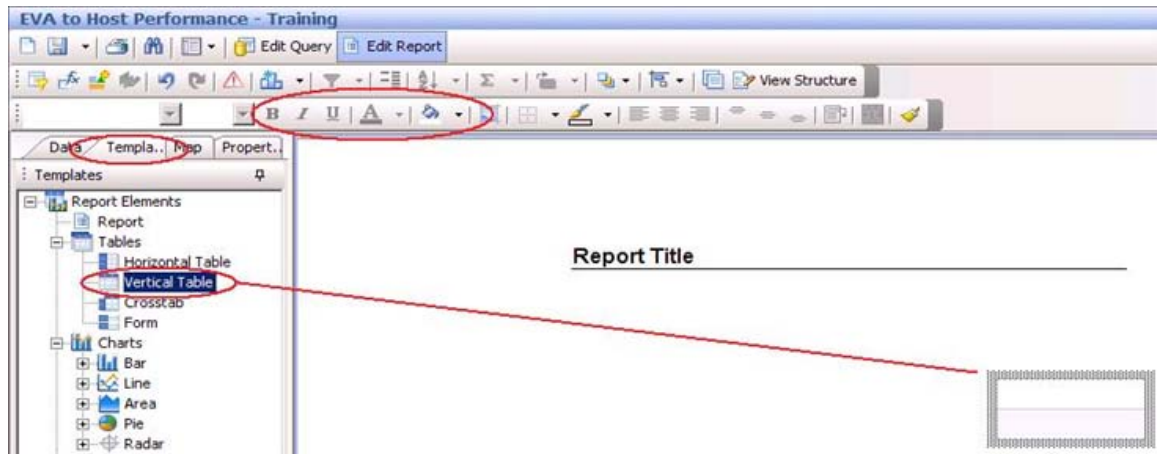
14. Click **Run Query**. The report similar to the following example is displayed.

Host Name	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	Collection Time	Avg Read Hit Latency	Avg Read Miss Latency	Avg Write Latency
CTX-MIA-EVACV01	0	0	0	0:5:6/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:7/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:8/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:9/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:10/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:11/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:12/10	0	0	0
CTX-MIA-EVACV01	0	0	0	0:5:13/10	0	0	0

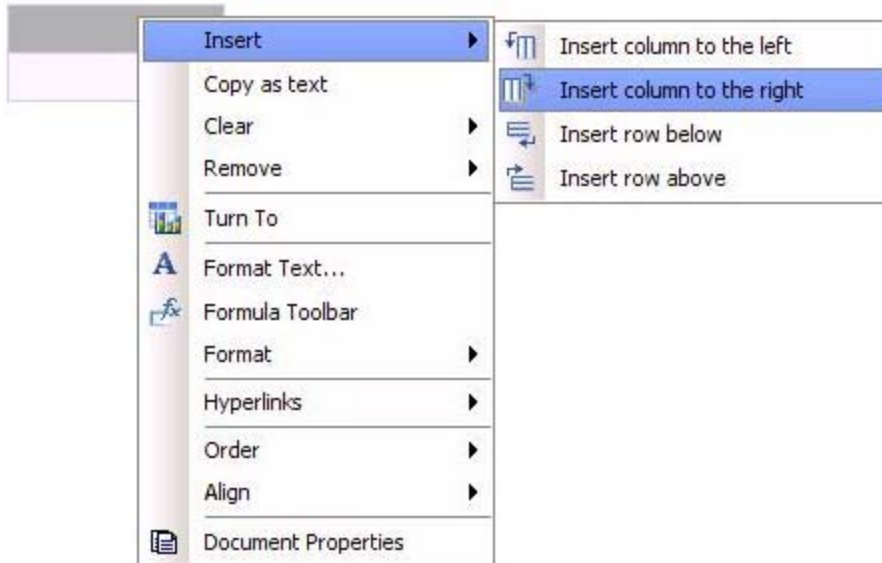
15. Select the table, right-click and select **Remove > Remove Table**. This removes the table format from the report and gives you a clean palette on which you can redesign the presentation of the data in this report.

Task 2 – Insert the First Table

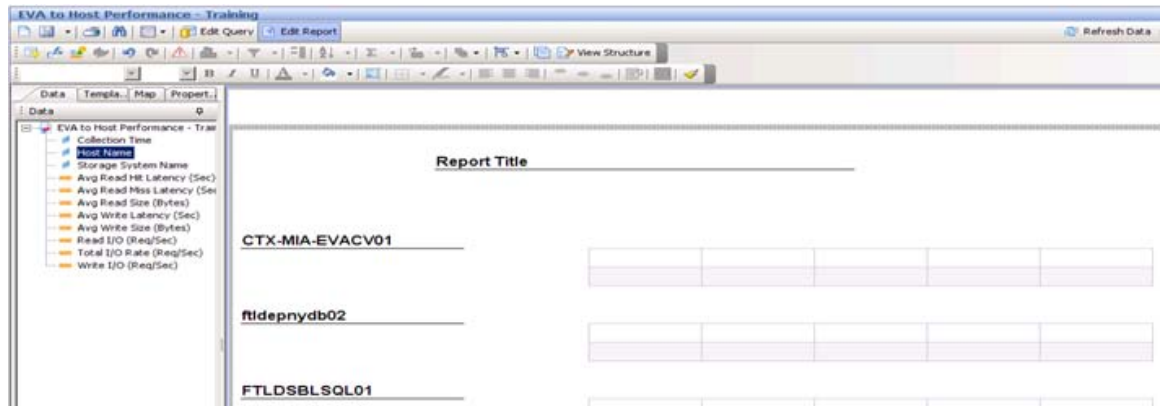
1. Click the **Templates** tab.
2. Drag the **Vertical Table** object onto the report. Change the background color of the table to white and the foreground color to black, if you choose.



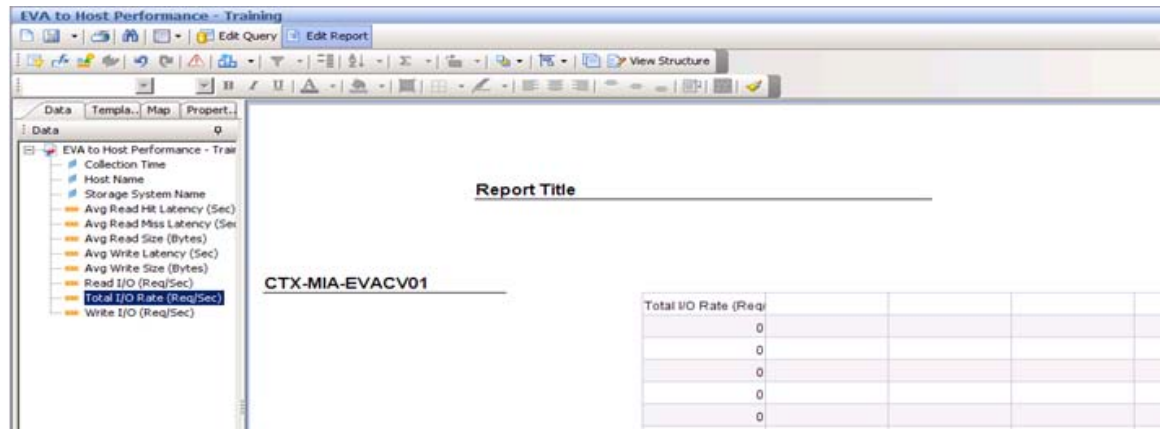
3. Right-click the table object and select **Insert > Insert column to the right** and add five columns.




4. Drag the **Host Name** object onto the report.



5. Drag the **Total I/O Rate (Req/Sec)** object into the first column.



6. Insert formula into first column.
 - a. Double-click the value in the first cell.
 - b. Press **Ctrl-Enter** to display the Formula Editor.
 - c. Insert the “Average” function as follows: **=average([Total I/O Rate (Req/Sec)])**



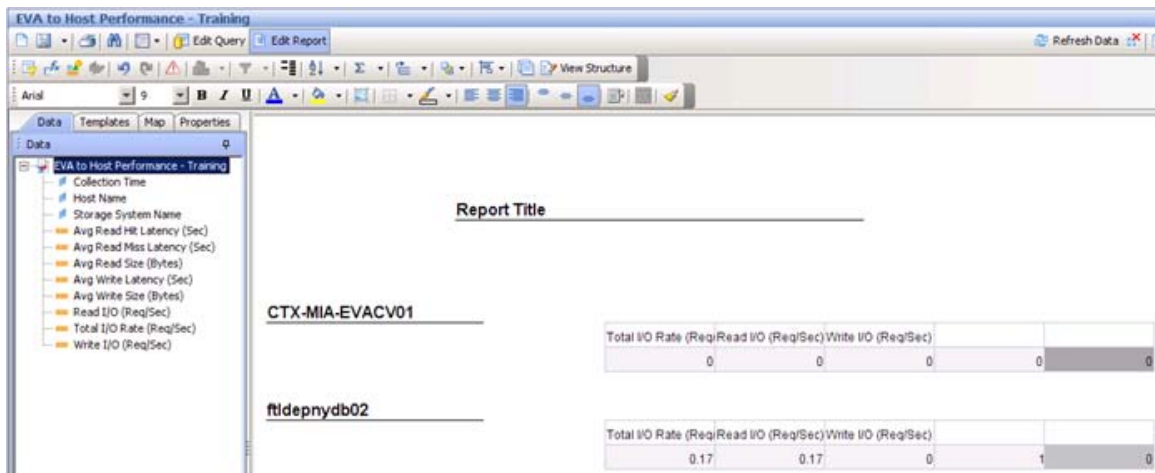
- d. Click **OK**.
7. Drag the **Read I/O (Req/Sec)** object to the second column in the table. Modify the value in the cell as described in Step 21 above by inserting the “Average” function.
8. Drag the **Write I/O (Req/Sec)** object to the third column in the table. Modify the value in the cell as described in Step 21 above by inserting the “Average” function.

9. Drag the **% Reads** object into the fourth column. Modify the value in the cell by inserting the following function:


```
=if iserror(Average([Read I/O (Req/Sec)]) / Average([Total I/O Rate (Req/Sec)])) then 0 else Average([Read I/O (Req/Sec)]) / Average([Total I/O Rate (Req/Sec)])
```

10. Drag the **% Write** object into the fifth column. Modify the value in the cell by inserting the following function:

```
=If IsError(Average([Write I/O (Req/Sec)]) / Average([Total I/O Rate (Req/Sec)])) Then 0 Else Average([Write I/O (Req/Sec)]) / Average([Total I/O Rate (Req/Sec)])
```




11. Select **View Structure** from the toolbar.
12. Click the **Templates** tab and navigate to the **Charts > Line** folder and expand.
13. Insert a **Vertical Mixed** chart below the host name.
14. Drag the **Total I/O Rate (Req/Sec)** object to the “place Measure objects here” section of the chart, and then follow these steps:
- Double-click the **Total I/O Rate (Req/Sec)** object on the chart.
 - Press **Ctrl-Enter** to display the Formula Editor.
 - Insert the “Average” function as follows: **=average([Total I/O Rate (Req/Sec)])**



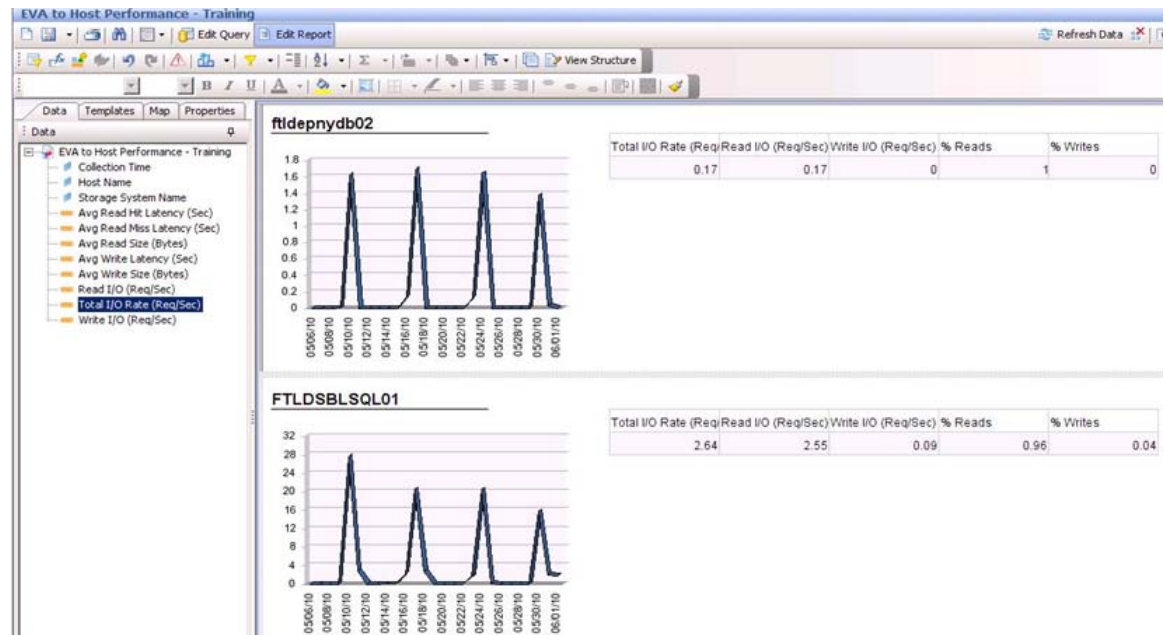
- Click **OK**.
15. Drag the **Collection Time** object to the “place Dimension objects here” at the bottom of the

chart, and then follow these steps:

- a. Double-click the **Collection Time** object on the chart.
- b. Press **Ctrl-Enter** to display the Formula Editor.
- c. Insert the "Formatdate" function as follows: **=average([Total I/O Rate (Req/Sec)])**



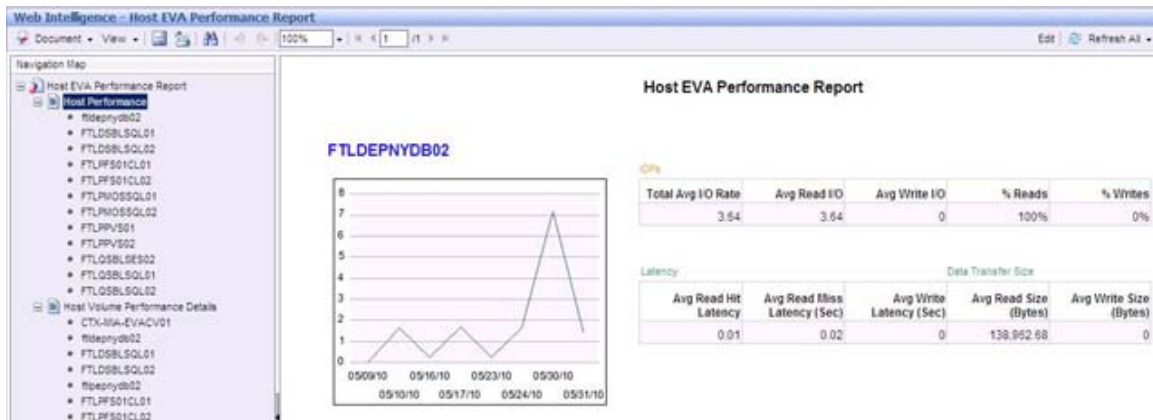
- d. Click **OK**.
16. Drag the **Host Name** object to the "place Dimension objects here" at the top-right of the chart.
17. Select **View Results** from the toolbar. The report should look similar to the following figure:



Task 3 – Insert a Second Table

1. Selecting **Templates** tab.
2. Drag the **Vertical Table** object to the report.
3. Drag the following Measure objects to the report—one to each column—and insert the "Average" function: **=average([<object_name>])** into each column.
 - **Avg Read Miss Latency (Sec)**
 - **Avg Write Latency (Sec)**

- **Avg Read Size (Bytes)**
 - **Avg Write Size (Bytes)**
4. Select **View Results** to format the report.
 5. Perform these steps to adjust chart properties:
 - Select the chart and click the **Properties** tab.
 - In the Properties tab, expand **Display** and de-select **3D Look**.
 - Expand **Data** and de-select **Show data markers**.
 6. Perform these steps to modify text and headers in the report:
 - Remove the underline from the Host Name and change the text color as desired.
 - Select each table header and select **right-justify** from the toolbar.
 - Add headers to each table section. Modify header colors as desired.
 - Double-click the **Report Title** object and type the title for the report: Host EVA Performance Report.
 - Move the **Report Title** object to the report header section and remove the underline.
 7. Select **View Results** in the toolbar. The report should look similar to the following example:



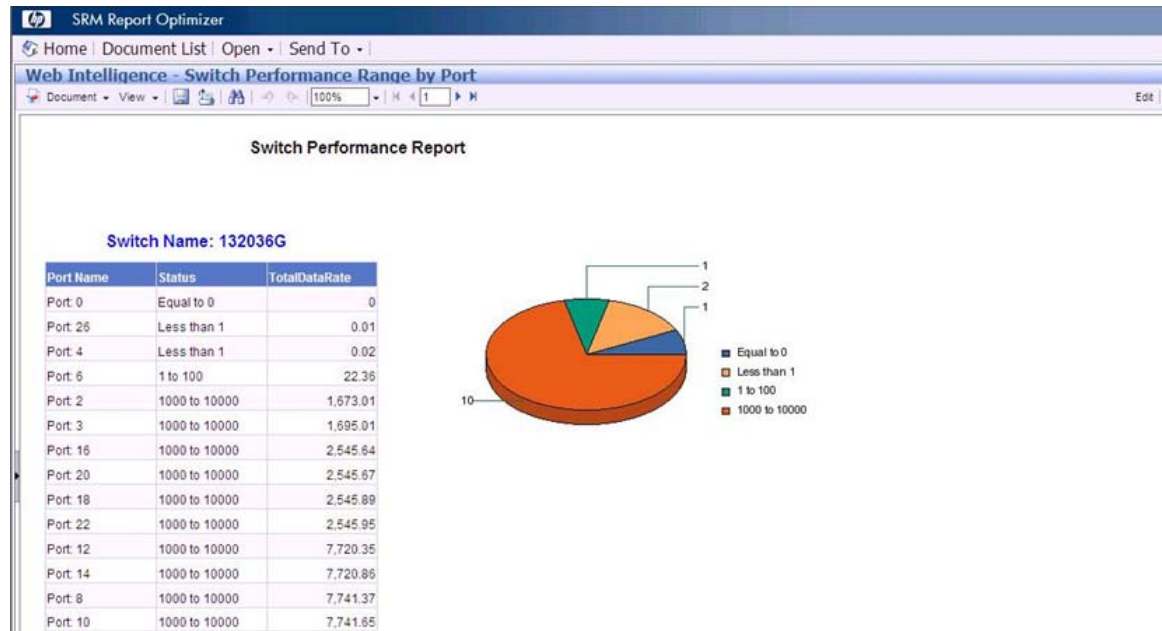
Creating Advanced Custom Reports

SRM Report Optimizer gives you the ability to create reports that show more than raw data statistics. Performance reports can show comparisons in performance metric data, illustrate information visually, include drill down capabilities, and display performance trends. By creating advanced custom reports, you improve the presentation of your reports and include only the data you need to meet your audience's specific interests.

The following example shows how to create a switch port performance report that breaks out the data by port and illustrates total data rates in a pie chart. The techniques you'll learn in this example can easily be applied to other report instances, whether you are creating a host performance report or pulling together comparative statistics on controllers.

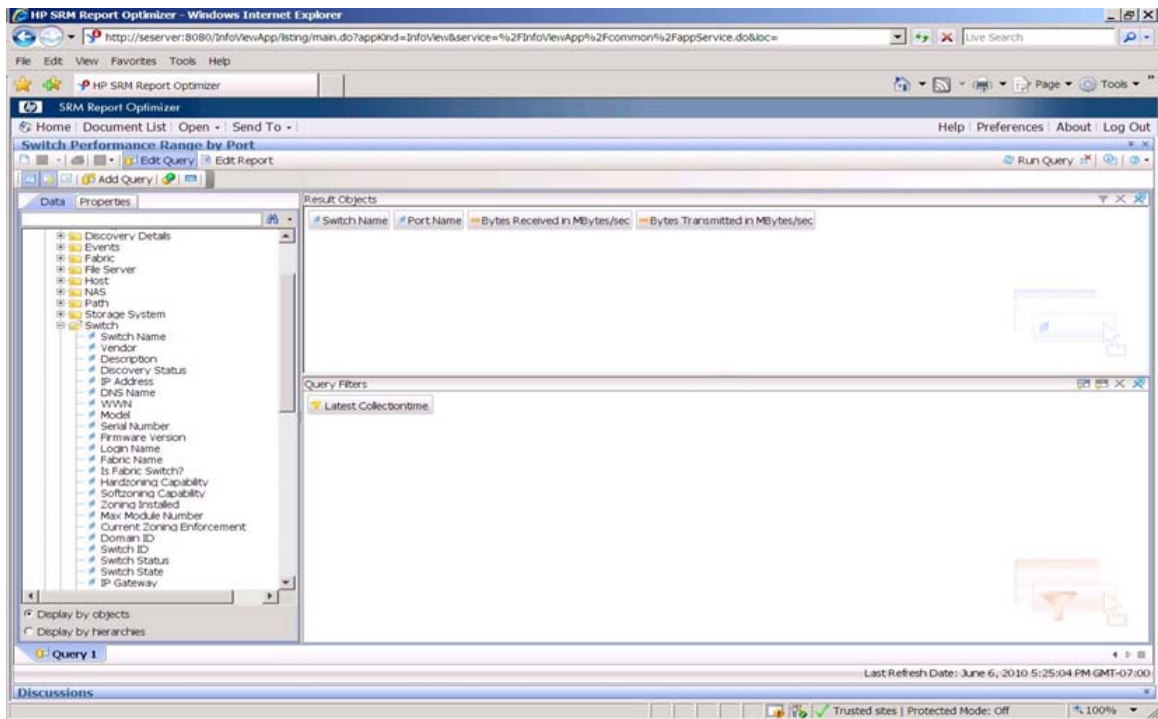
Example: Switch Port Performance Report

This example creates a Switch Port Performance Report that shows switch port performance ranges for total Mbytes transmitted and received. The categories selected are Total Bytes equal to 0, less than 1K, 1K to 100K, 100K to 1000K, 1000K to 10000K, and 10000K to 50000K. We'll also display the count for these in a pie chart. The end report should look something like this:

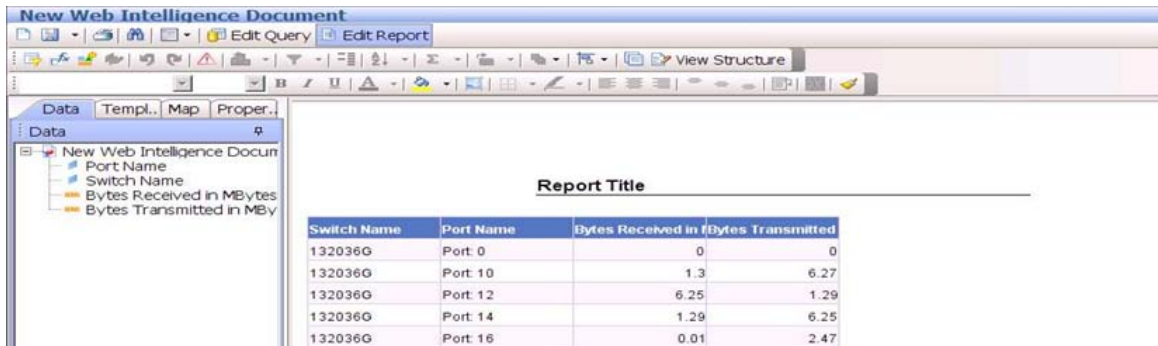


The following instructions assume that you logged on to Report Optimizer, created a Web Intelligence Document, and connected to the HP Storage Essentials Universe database. See [Creating Custom Performance Reports on page 236](#) for instructions.

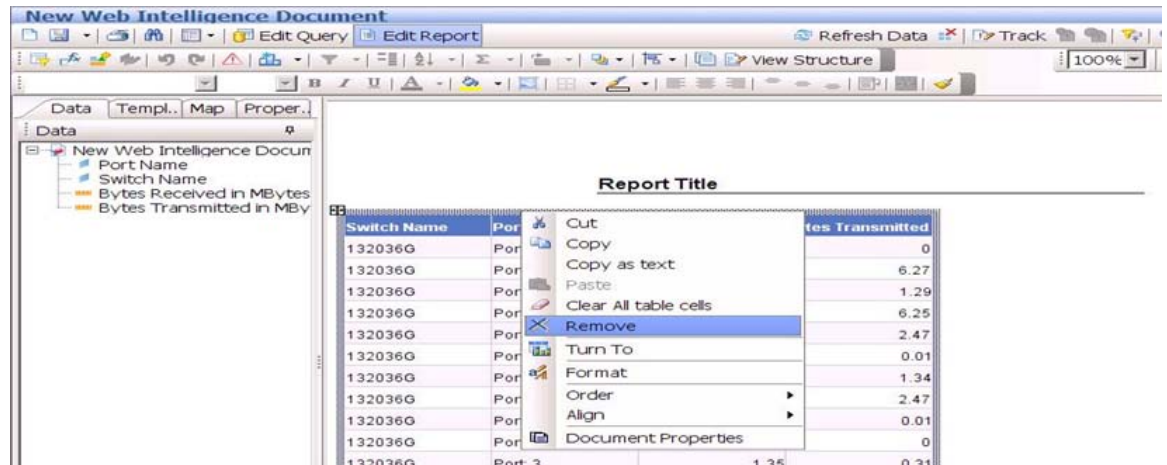
1. With the **Data** tab open, expand the object tree and navigate to the **Switch** folder.
2. Select and expand the **Switch** folder and drag (or double-click) the **Switch Name** object onto the Result Objects pane.
3. Select and expand the **Switch Ports** folder and drag (or double-click) the **Port Name** object onto the Result Objects pane.
4. Select the **Switch Port Performance Statistics - Historical** folder and drag (or double-click) the **Bytes Received in MBytes/sec** object onto the Result Objects pane.
5. Drag (or double-click) the **Bytes Transmitted in MBytes/sec** object onto the Result Objects pane.
6. Scroll to the bottom of the **Switch** folder and drag (or double-click) **Latest Collection Time** into the Query Filters pane.



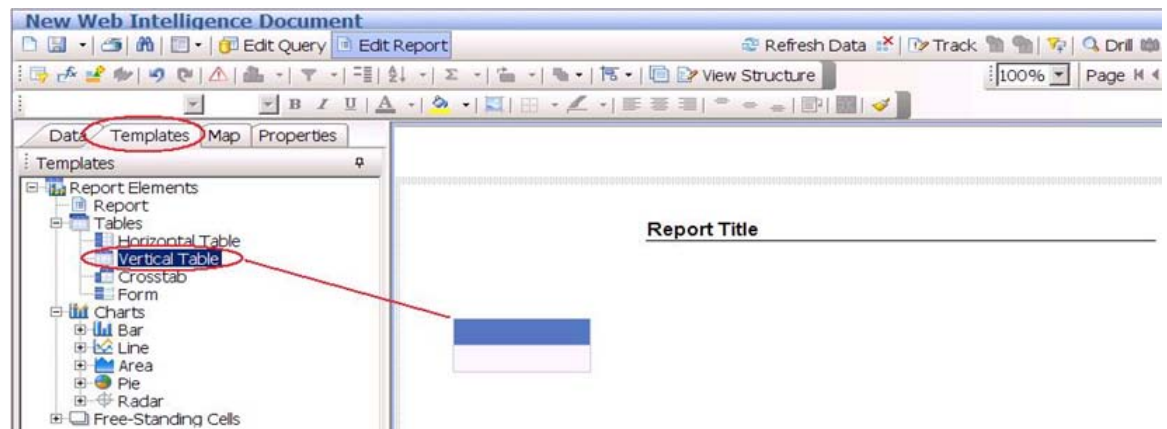
- When ready, click **Run Query** in the upper right toolbar. Your report should look similar to this:



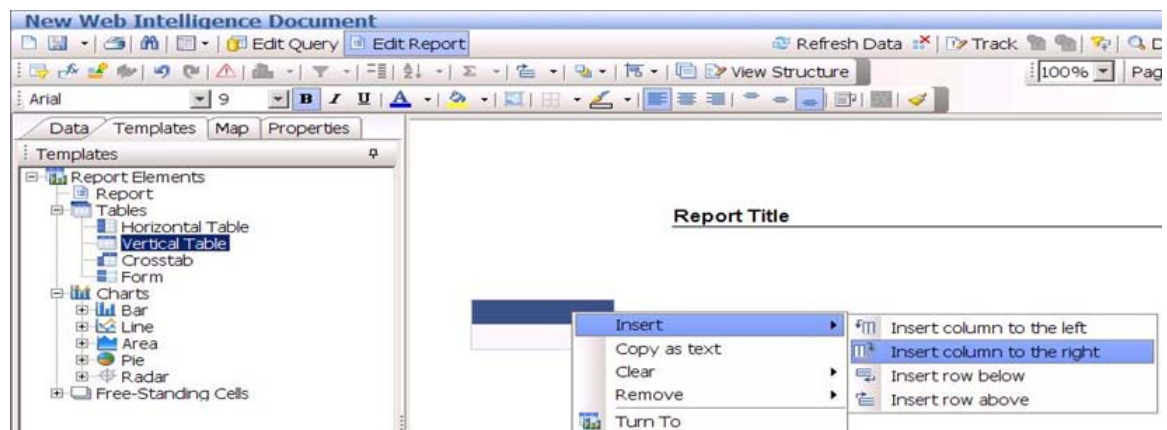
- We want to redesign the table presentation, so we first need to remove the existing table. To do this, select the report table, and right-click to select **Remove > Remove Table** from the drop-down menu



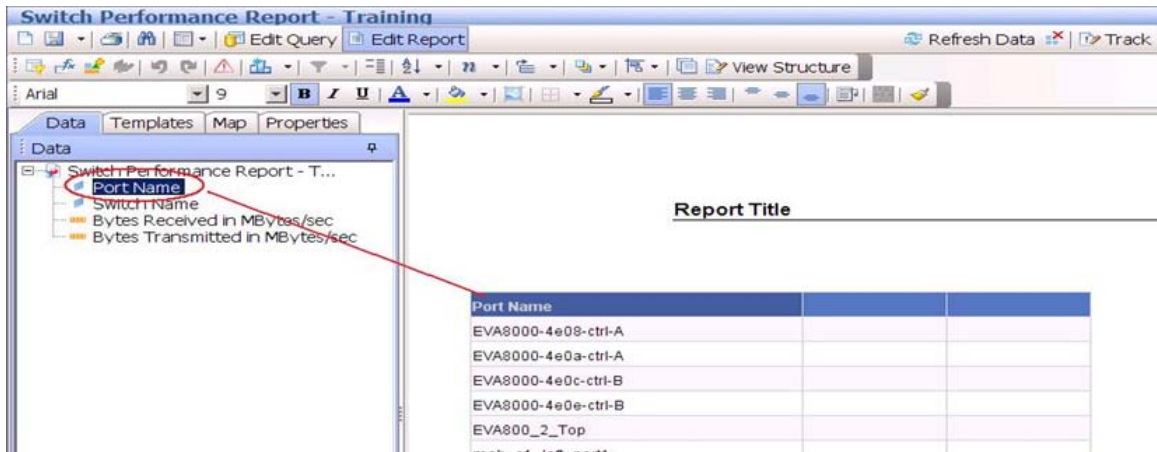
9. Select the Templates tab, then drag the **Vertical Table** object over to the report pane.



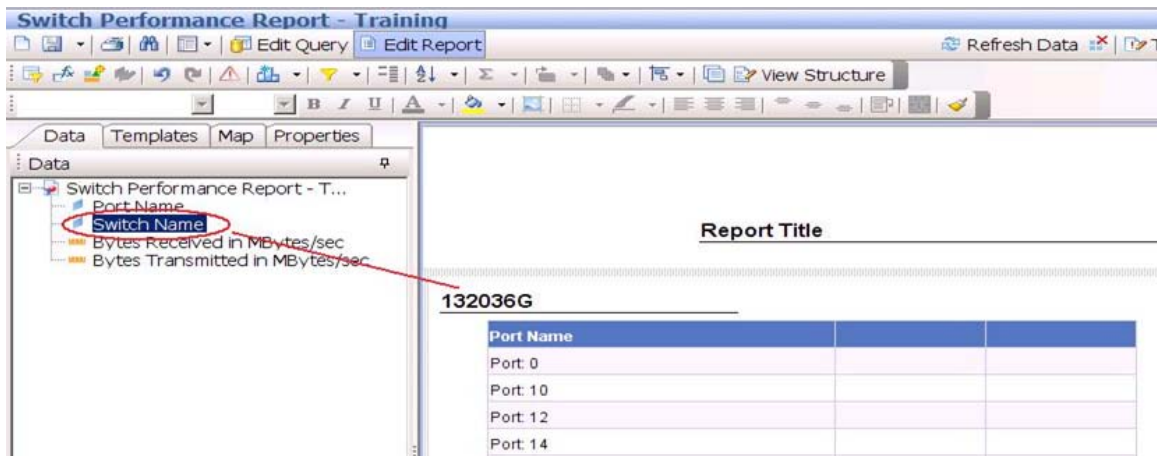
10. Select the Vertical Table table object and right-click to select **Insert > Insert column to the right** from the drop-down menu. Add two more columns for a total of three.



11. Drag the **Port Name** object to the first column.




12. To show report data by switch name, drag the **Switch Name** object to the report pane and position it above the table as follows:



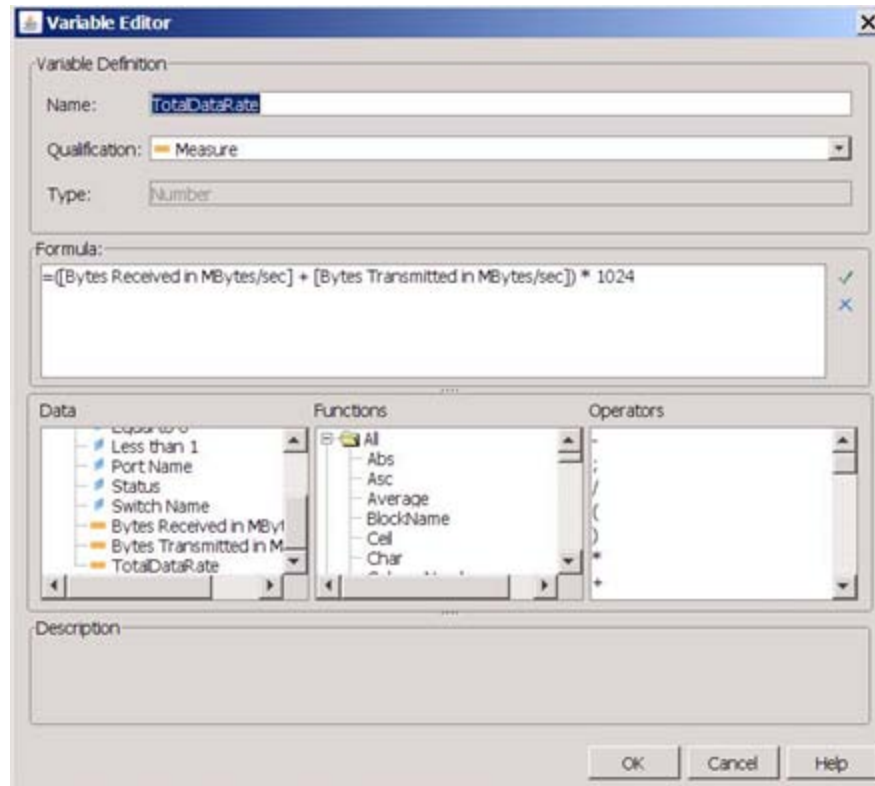
13. The next step is to create the following variables for your report.

Variable Name	Type	Formula
TotalDataRate	Measure	$=(\text{Bytes Received in MBytes/sec} + \text{Bytes Transmitted in MBytes/sec}) * 1024$
Equal to 0	Dimension	$=\text{If}[\text{TotalDataRate}] = 0 \text{ Then "Equal to 0"}$
Less than 1	Dimension	$=\text{If}[\text{TotalDataRate}] > 0 \text{ And } [\text{TotalDataRate}] < 1 \text{ Then "Less than 1"}$
1 to 100	Dimension	$=\text{If}[\text{TotalDataRate}] > 1 \text{ And } [\text{TotalDataRate}] < 100 \text{ Then "1 to 100"}$
100 to 1000	Dimension	$=\text{If}[\text{TotalDataRate}] > 100 \text{ And } [\text{TotalDataRate}] < 1000 \text{ Then "100 to 1000"}$
1000 to 10000	Dimension	$=\text{If}[\text{TotalDataRate}] > 1000 \text{ And } [\text{TotalDataRate}] < 10000 \text{ Then "1000 to 10000"}$
10000 to 50000	Dimension	$=\text{If}[\text{TotalDataRate}] > 10000 \text{ And } [\text{TotalDataRate}] < 50000 \text{ Then "10000 to 50000"}$
Status	Dimension	$=\text{If}[\text{TotalDataRate}] = 0 \text{ Then "Equal to 0" Elseif } [\text{TotalDataRate}] < 1 \text{ And } [\text{TotalDataRate}] > 0 \text{ Then "Less than 1" Elseif } [\text{TotalDataRate}] > 1 \text{ And } [\text{TotalDataRate}] < 100 \text{ Then "1 to 100" Elseif } [\text{TotalDataRate}] > 100 \text{ And } [\text{TotalDataRate}] < 1000 \text{ Then "100 to 1000" Elseif } [\text{TotalDataRate}] > 1000 \text{ And } [\text{TotalDataRate}] < 10000 \text{ Then "1000 to 10000" Elseif } [\text{TotalDataRate}] > 10000 \text{ And } [\text{TotalDataRate}] < 50000 \text{ Then "10000 to 50000" Elseif } [\text{TotalDataRate}] > 50000 \text{ Then "Greater than 50K"}$


From the toolbar, open the Variable Editor by clicking the icon.



14. Using the Variable Editor dialog, create the MEASURE variable with the formula $=[\text{Bytes Received in MBytes/sec}] + [\text{Bytes Transmitted in MBytes/sec}] * 1024$. This variable gives the switch throughput as the total number of MBytes received and transmitted. When done, click **OK**.

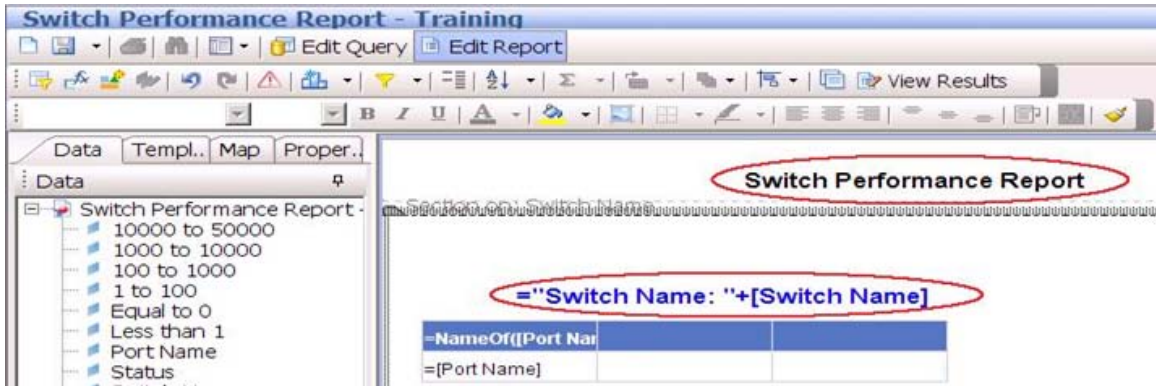


15. The next step is to format the table, fill in the table data, and create the pie chart. To begin, select **View Structure** in the toolbar.

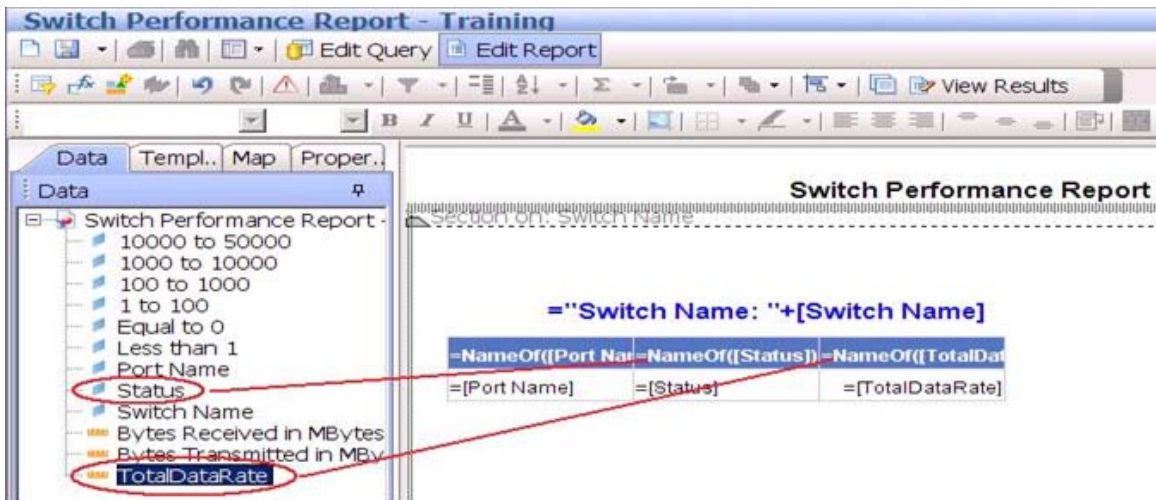


16. Move the report title object to the header space and change the name of the report to "Switch Performance Report".
17. Double click the [Switch Name] object and insert "Switch Name: "+" between the "=" and the

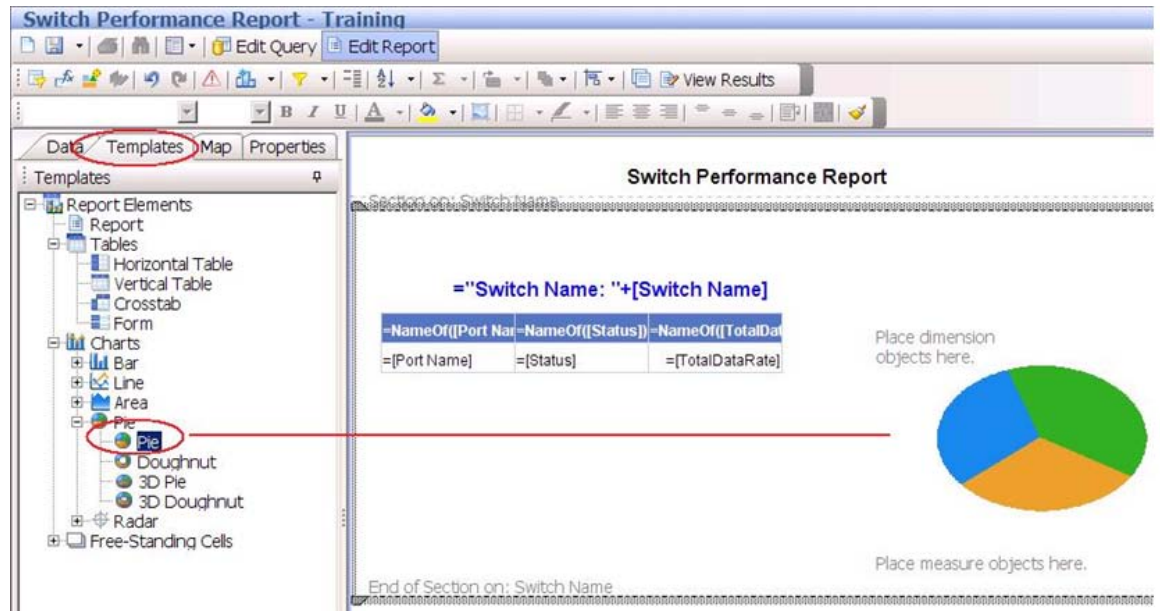
[Switch Name] object. Center both the report title and the switch name object.



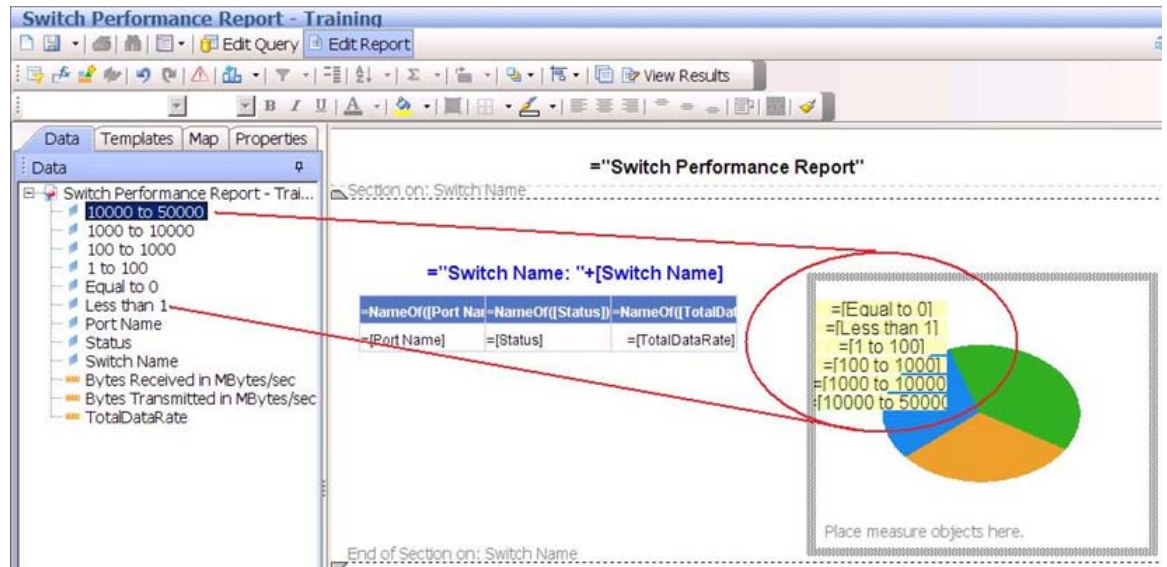
18. Drag the **Status** object to column two and drag the **TotalDataRate** object to column three.



19. Click the **Templates** tab and expand **Charts > Pie**.
20. Select **Pie** and drag the Pie Chart to the right side of the report.

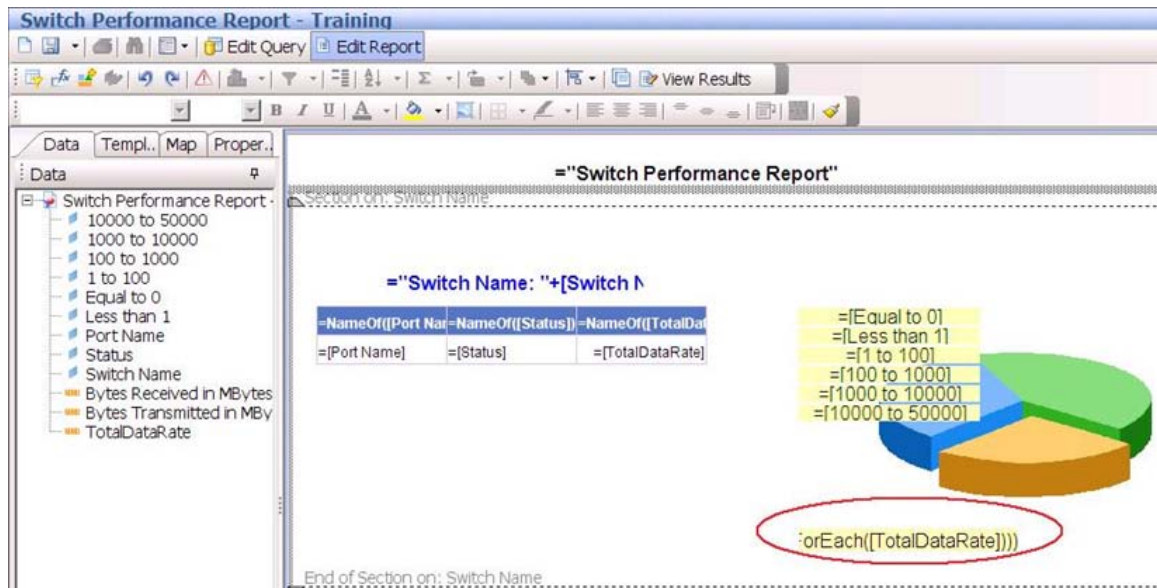


21. Drag the following Dimension objects to the “Place dimension objects here” section of the Pie Chart. The objects *must* be entered in the exact order below:



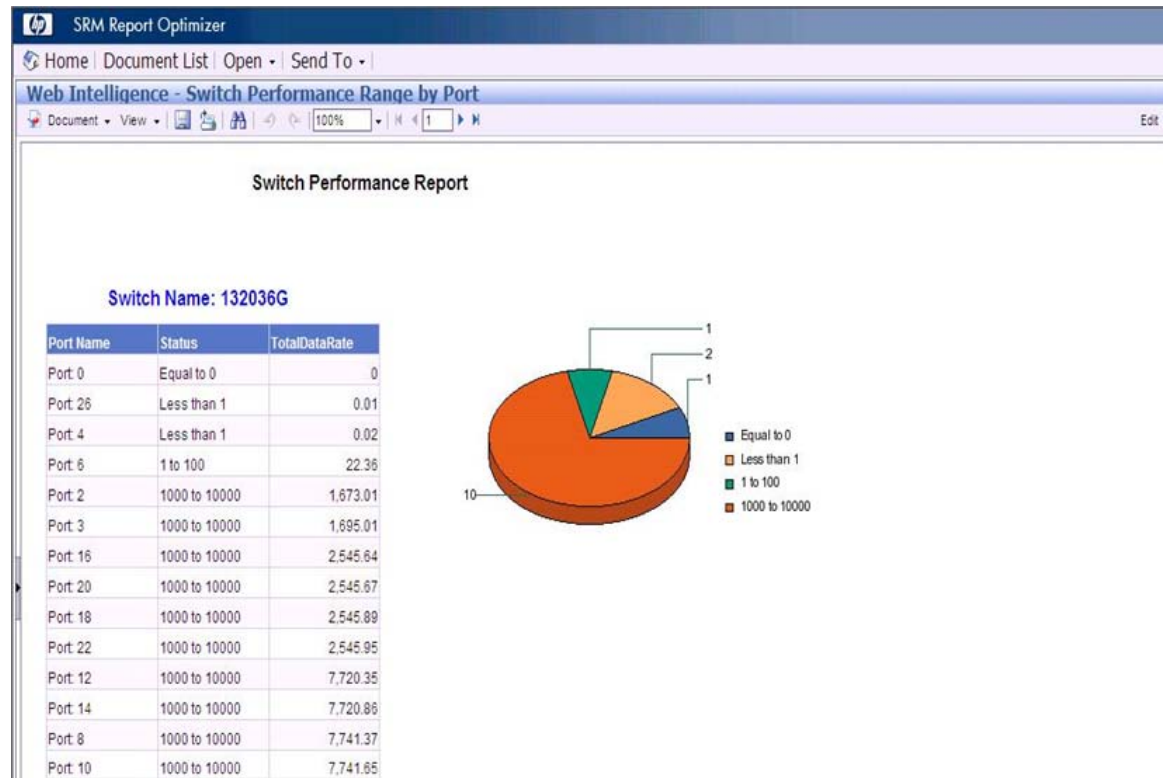
22. Add the following MEASURE object formula to the “Place measure objects here” section of the Pie Chart:

```
=Sum(Count([Equal to 0] ForEach([TotalDataRate]))) + Sum(Count([Less than 1] ForEach([TotalDataRate]))) + Sum(Count([1 to 100] ForEach([TotalDataRate]))) + Sum(Count([100 to 1000] ForEach([TotalDataRate]))) + Sum(Count([1000 to 10000] ForEach([TotalDataRate]))) + Sum(Count([10000 to 50000] ForEach([TotalDataRate])))
```



To get this formula into the Pie Chart, drag over any Measure object from the Data tab, and replace that object with the newly created formula.

23. Select **View Results** from the toolbar. This displays the Switch Performance Report as shown below:



Creating Performance Reports with Drill Down and Trending

After you set up a performance report and run the report query to generate the report, you can format that information in trending graphs. Trending information assists you in understanding the past and current performance levels of a storage application, host, or device on your system.

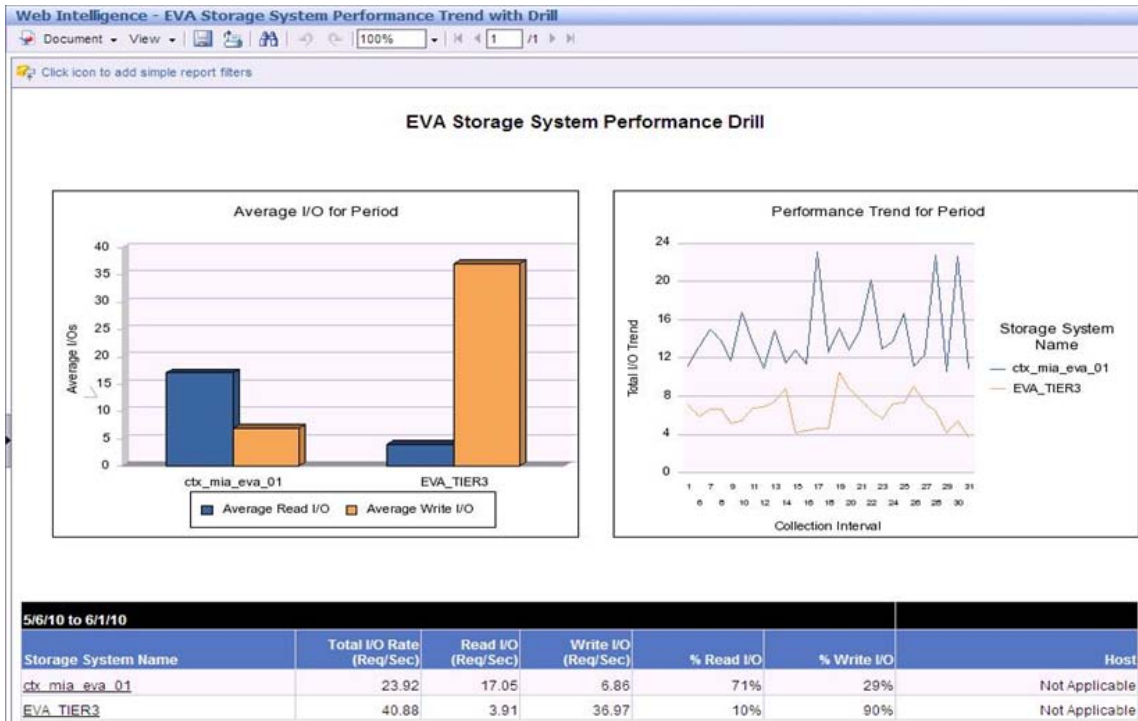
Creating a trending graph requires the following tasks:

- Task 1 – Design the report by specifying the Dimension and Measure objects in the Results Object pane and then the Filter objects in the Query Filter pane. Run the report query to see the basic report. Depending on your type of report, you might also need to select the query context, answer prompts, and run the report query again.
- Task 2 – Redesign the report format by removing columns and adding new ones. Insert formulas as needed into column cells.
- Task 3 – Insert Drill Filter functions wherever you want drill down capability.
- Task 4 – Create variables for trending data. Insert line and bar charts that display how this data trends over a period of time.

The following example shows how to accomplish the tasks.

Example: EVA to Host Performance Trending Report with Drill Down

This example explains how to create a report that displays EVA-to-host performance statistics. The example captures the collected data for I/Os per second, latency, and the data rate. It places this information into a trending chart for each host.



The following instructions assume that you logged on to Report Optimizer, created a Web Intelligence Document, and connected to the HP Storage Essentials Universe database. See [Creating Custom Performance Reports on page 236](#) for instructions.

Task 1 – Design and Generate the Basic Report

1. With the Data tab selected in the left pane, navigate to the **Storage System** folder and expand **Performance Statistics > EVA > EVA Storage Volume Statistics - Historical**.
2. Drag or double-click the following objects into the Results Object pane on the right.
 - **Storage System Name**
 - **Storage Volume Name**
 - **Storage Pool Name**
 - **Host Security Group**
 - **Total I/O Rate (Req/Sec)**
 - **Read I/O Rate (Req/Sec)**

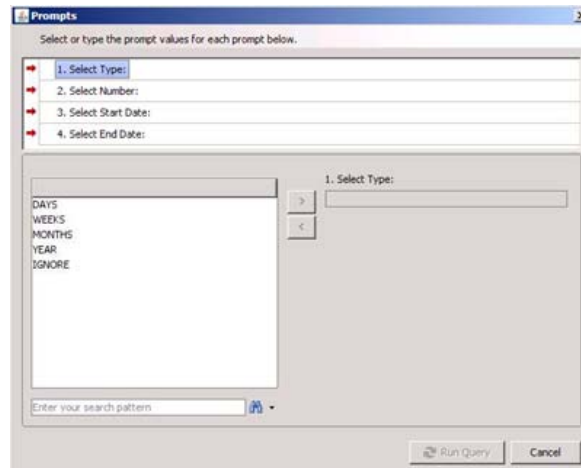
- **Write I/O Rate (Req/Sec)**
 - **Collection Time**
3. Scroll to the bottom of the Storage Volume Statistics - Historical folder and drag the following object to the Query Filters pane.:

Schedule Time Period

4. The resulting Web Intelligence page should look similar to the following figure.:



5. Click **Run Query** (upper right in toolbar).
6. The Prompts dialog appears. Enter answers to prompts as appropriate.



7. Click **Run Query**. The report is displayed.

Storage System Name	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	Collection Time	Storage Volume Name	Storage Pool Name	Host Security Group
ctx_mia_eva_01	2.32	2.32	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	2.41	2.41	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	0.63	0.63	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	1.16	1.16	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	57.74	35.66	22.08	5/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	57.74	35.66	22.08	5/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP

Task 2 – Redesign the Report Format

1. Select the **Storage Volume Name** column, right-click, and select **Remove > Remove Column**.

Storage System Name	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	Collection Time	Storage Volume Name	Storage Pool Name	Host Security Group
ctx_mia_eva_01	2.32	2.32	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	2.41	2.41	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	0.63	0.63	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	1.16	1.16	0.51	0/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	57.74	35.66	22.08	5/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP
ctx_mia_eva_01	57.74	35.66	22.08	5/10/10	CTX_MIA_HPUMFT	Default Disk Group	CTX_MIA_FTLDEP

2. Remove the **Storage Pool Name** and **Collection Time** columns. Right-click the column and select **Remove > Remove Column**.
3. Double-click a cell in the **Total I/O Rate** column and insert the following “Average” function:
`=average([Total I/O Rate (Req/Sec)])`

Formula Editor

Formula: `=average([Total I/O Rate (Req/Sec)])`

4. Double-click a cell in the **Read I/O** column and insert the following “Average” function:
`=average [Read I/O (Req/Sec)]`.
5. Double-click a cell in the **Write I/O** column and insert the following “Average” function:
`=average([Write I/O (Req/Sec)])`.
6. Right-click the table and select **Insert > Insert column to the right** and add two new

columns. These columns contain the % Read and % Write statistics.

- Select the first new column and press **Ctrl-Enter** to display the Formula Editor. Modify the value in the first new column by inserting the following function: `=average([Read I/O (Req/Sec)]) / average([Total I/O Rate (Req/Sec)])`. Type "%Read" in the header for this column.
- Select the second new column and press **Ctrl-Enter** to display the Formula Editor. Modify the value in the second new by inserting the following function: `=average([Write I/O (Req/Sec)]) / average([Total I/O Rate (Req/Sec)])`. Type "%Write" in the header for this column.

Task 3 – Insert Drill Down

- Double-click a cell in the Host Security Group column and insert the following function: `=if DrillFilters()=[Storage System Name] + " - " + [Storage Pool Name] then [Host Security Group] else "Not Applicable"`.

This formula displays the host security group name if the user drills down to the volume level. Otherwise, it displays "Not Applicable"

- Click **Drill** in the upper-right toolbar.



- Your report should look like the following:

Report Title


Storage System Name	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	% Reads	% Writes	Host Security Group
ctx_mia_eva_01	23.22	16.58	6.64	0.71	0.29	Not Applicable
EVA_TIER3	17.65	1.67	15.98	0.09	0.91	Not Applicable

- Perform the following clean-up steps on the report title:
 - Double-click the Report Title object and modify the title to read "EVA Storage System Performance Trend with Drill".
 - Remove the underline in the title.
 - Move the Report Title to the report header section.

Task 4 – Add Trending Data in Charts

- Select **View Structure** from the toolbar.
- Click the **Templates** tab. In the object tree (left pane), navigate to the **Charts > Bar** folder and expand.
- Insert a **Vertical Grouped** bar chart above the table block.

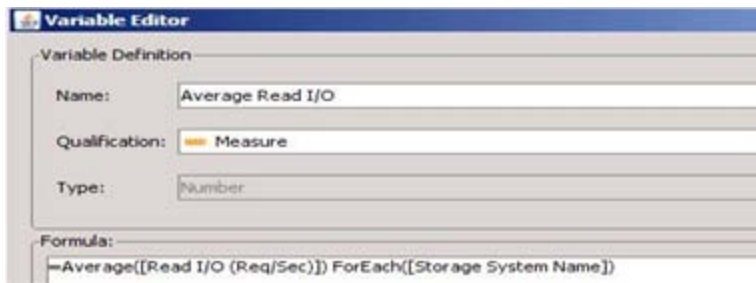
4. Navigate to the **Charts > Line** folder and insert a **Vertical Mixed** line chart to the right of the bar chart and above the table block.
5. Click **Variables Editor** in the toolbar.



In the Variable Editor, create the following variables:

Variable Name	Qualification	Formula
Period	Dimension	=DayNumberOfMonth([Collection Time])
Average Read I/O	Measure	=Average([Read I/O (Req/Sec)]) ForEach([StorageSystemName])
Average Write I/O	Measure	=Average([Write I/O (Req/Sec)]) ForEach([StorageSystemName])
Total I/O Trend	Measure	=Average([Total I/O (Req/Sec)]) ForEach([StorageSystemName])

The following example shows how to create the Average Read I/O variable:

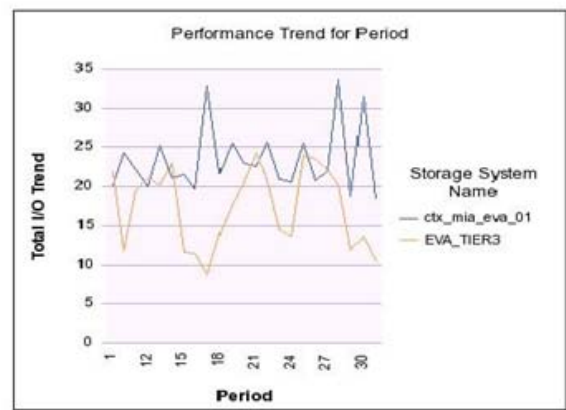



6. Drag the **Average Read I/O** object to the “Place measure objects here” section of the Bar Chart.
7. Drag the **Average Write I/O** object to the same location and place it under the Average Read I/O object.
8. Drag the **Storage System Name** object to the “Place dimension objects here” in the section of the Bar Chart.
9. Drag the **Total I/O Trend** object to the “Place measure objects here” section of the Line Chart.

10. Drag the **Period** object to the “Place dimension objects here” in the bottom section of the Line Chart.
11. Drag the **Storage System Name** object to the “Place dimension objects here (optional)” top-right section of the Line Chart.
12. Select **View Results**.
13. (Optional) Change the properties of the charts.

The resulting report should look like the following:

EVA Storage System Performance Trend with Drill



Storage System Name	Total I/O Rate (Req/Sec)	Read I/O (Req/Sec)	Write I/O (Req/Sec)	% Reads	% Writes	Host Security Group
ctx_mia_eva_01	23.22	16.58	6.64	0.71	0.29	Not Applicable
EVA_TIER3	17.65	1.67	15.98	0.09	0.91	Not Applicable

8 Troubleshooting Performance Issues

This Troubleshooting section describes known issues related to monitoring storage performance in a system managed by HP Storage Essentials..

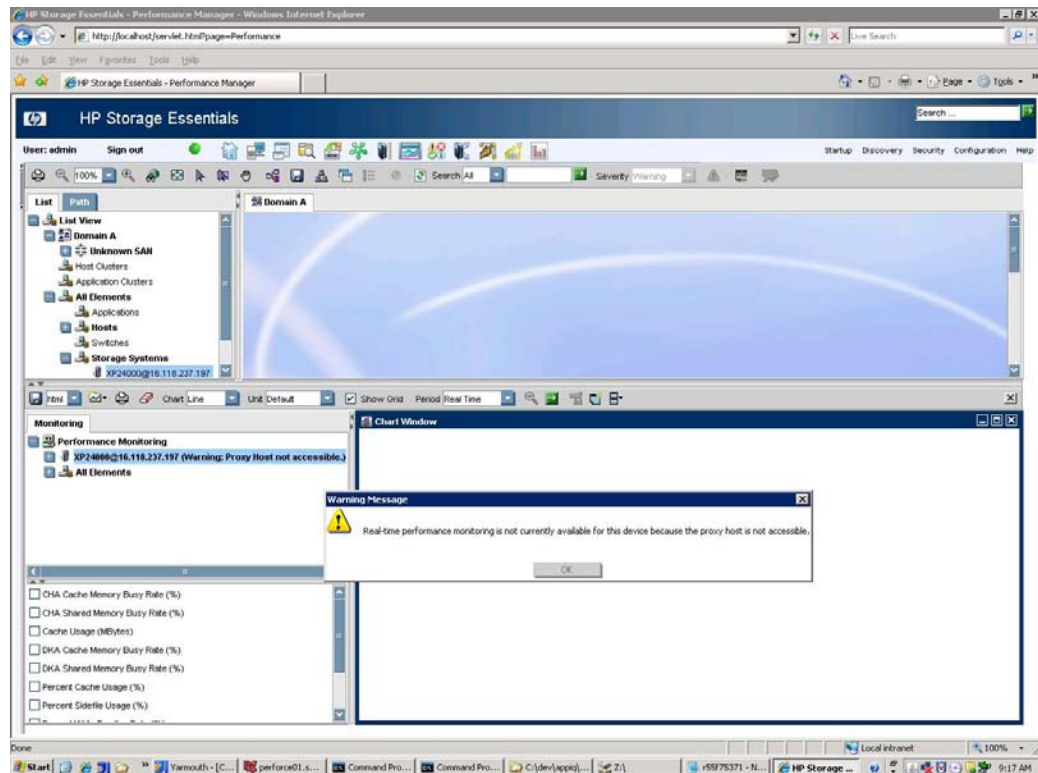
Performance Manager

XP Array Proxy Host is Not Accessible

HP Performance Pack Enterprise software for XP arrays requires a proxy host. You can only view real-time or historical performance data when the proxy host is available.

The Performance Manager displays the following message if you attempt to query real-time performance statistics for an XP array whose proxy host is not accessible:

"Real-time performance monitoring is not currently available for this device because the proxy host is not accessible."



Click **OK** to close the message.

For information about setting up a proxy host, see [Create the Host Proxy on page 42](#) [Setting Up XP and HDS Performance Packs on page 35](#)

You cannot use the same proxy host for both an XP and an HDS array configured on your storage system. One proxy host can be used *either* for multiple XP arrays *or* multiple HDS arrays, but not for both.

Not Enough Data to Produce a Chart

When you attempt to create a chart, if you see the message:

```
"There is not enough data to produce a chart [chart_title] at this time."
```

Verify that there is active performance data collection for the storage element or device that you are charting. Active data collectors are listed on the Performance Data Collectors page (**Configuration > Performance > Data Collection**). Verify that the performance collector for the storage element is enabled. See [Managing Performance Data Collection on page 45](#) for more information about enabling performance collectors.

If data collection is enabled, you might consider decreasing the collection frequency. Also consider editing the chart properties by selecting All in the Edit Chart Property dialog box.

Data Collection for EMC Arrays

The EMC Solutions Enabler STP Daemon (storstpd) is an EMC daemon that collects performance data counters at regular intervals. HP Storage Essentials does not depend on this daemon for storage performance data collection or for the derivation of any performance statistics. If you are running this daemon while collecting performance data using Storage Essentials, be sure to monitor the daemon to avoid heavy CPU usage which might cause issues with Storage Essentials data collection. If the storstpd daemon seems to be interfering with Storage Essentials' ability to successfully collect statistical performance data, then it may be necessary to stop the storstpd daemon and set the daemon or service to not start automatically.

Managing Late Data or Errors from XP and HDS Arrays

If you are performing real-time data collection, and the element is not returning the information in time, you will see a message resembling the following displayed in red:

```
Data is late or an error occurred...
```

The message means the software cannot obtain the information in a timely manner because one or more of the following occurred:

- The CV XP SMI-S provider went down.
- The built-in XP or HDS provider went down.
- Someone logged into the built-in provider for the array on the Service Processor (SVP) and changed the "View" mode to "Modify."

When this happens, in most cases, statistics will continue to be collected because they are collected using the host proxy, not through CV XP SMI-S provider or the built-in provider.

However, if Step 3 Discovery is attempted when CV XP SMI-S provider or the built-in provider is down, then statistics will stop being collected. In this case, the error that would appear in the cimom.log is the following:

```
HDS_MISSING_CONNECTION=Connection information is missing for Storage Array with Serial Number {0}. Try rediscovering the access points in Discovery Setup.
```

To resolve the error, restart the CV XP SMI-S provider or the built-in provider and then redo Step 1 Discovery for the array.

You also might want to look in the appstorm and cimom logs around the time the error occurred to see if there are any error messages that might explain why the management server failed to collect the data. The product provides substantial error handling and logging for failed XP and HDS performance data collections.

Glossary

This glossary contains a list of commonly used terms that relate to storage system performance and/or conditions relevant to describing performance metrics.

Disk drive service time—Time period equal to the Sum of Seek + Latency + Transfer Time.

Disk latency—Secondary operation that occurs after the “seek”, which is the time it takes for the data to reach the read/write head of a disk drive.

Exchange—A single exchange contains one or more frame sequences. It is comparable to a SCSI I/O process, and is the mechanism for coordinating the exchange of information between two communicating N_Ports in a single operation.

FC port—Port defined by the Fibre Channel (FC) protocol. See the FC port types listed in the Fabric Switch section of this guide.

Frame—Smallest indivisible unit of user data that is sent on the Fibre Channel link. Frames can be variable in length, up to a maximum of 2148 bytes long. Frame size depends on implementation, not hardware or software. Each frame contains a four-byte Start of Frame delimiter, a 24-byte header, up to 2112 bytes of FC-4 payload consisting of zero to 64 bytes of optional headers and zero to 2048 bytes of ULP data, a four-byte CRC, and a four-byte End of Frame delimiter.

Frame sequence—A set of one or more related frames. For example, a large file transfer would be accomplished in a sequence consisting of multiple frames.

IOPS—I/O per second. IOPS is calculated using the following formula:
$$\text{IOPS} = 1000\text{ms} / (\text{CommandTime} + \text{SeekTime} + \text{SettledownTime} + \text{DataTransferTime})$$

Latency—Time it takes a given network packet to travel from source to destination and back; also referred to as the *round-trip time*.

Response time—Number of I/Os processed per second by the storage system. The number of I/Os per second (IOPS) is another metric used to define the performance of a storage system. This describes how efficiently the storage controller processes transactions. Generally, OLTP applications, which have heavy on-line transaction workloads, are most concerned with I/Os per seconds (IOPS).

Seek—Initial operation a disk performs to place the read head on the right track of a disk drive.

Storage performance—Amount of work completed per unit of time, which refers to the number of I/O requests (work) completed by the storage system per second. The two measurements that are typically used to define the performance of a storage system are *throughput* and *latency*.

Throughput—Volume of requests that can be processed by the storage system per unit of time. For storage performance, this is usually measured in megabytes per second ((MB/s). Throughput can be described as how effectively the storage controller is able to supply the application with a lot of data. Generally, batch jobs and scientific applications such as seismic evaluation are most affected by throughput.

Transfer time—Time it takes for data to be read from or written to the host after seek and latency.

Workload—Specific I/O sequence of reads and writes, cache hits and misses, and sequential or random activity that an application demands of a storage system.