

HP 3PAR System Reporter 2.9 Software User's Guide

Abstract

This guide provides the information you need to install and use HP 3PAR System Reporter to monitor performance, create charge back reports, and plan storage resources for HP 3PAR Storage Systems, and is intended for system and storage administrators who monitor and direct system configurations and resource allocation for HP 3PAR Storage Systems.



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

The information contained herein is subject to change without notice. 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.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft®, Windows®, Windows® XP, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java and Oracle are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Perl DBD, MySQL Driver Copyright & License

This module is Copyright (c) 2003 Rudolf Lippan; Large Portions Copyright (c) 1997-2003 Jochen Wiedmann, with code portions Copyright (c) 1994-1997 their original authors. This module is released under the same license as Perl itself. See the Perl README for details.

MySQL TCL Copyright & License

Copyright (c) 1994, 1995 Hakan Soderstrom, Enskede, Sweden and Tom Poindexter, Denver, Colorado Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice and this permission notice appear in all copies of the software and related documentation.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the names of the copyright owners nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OraTcl Copyright and License

Copyright (c) 2000-2005 Todd M. Helfter

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the name of Todd Helfter may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Todd Helfter.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

IN NO EVENT SHALL TODD HELFTER BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

MySQL Client Library Copyright & License

Copyright © 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

SQLite Copyright & License

The original author of SQLite has dedicated the code to the public domain. Anyone is free to copy, modify, publish, use, compile, sell, or distribute the original SQLite code, either in source code form or as a compiled binary, for any purpose, commercial or noncommercial, and by any means.

Apache HTTP Server Copyright & License

Licensed under the Apache License Version 2.0, January 2004, <http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
 - c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and

wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

TWAPI Copyright & License

Copyright (c) 2003-2008, Ashok P. Nadkarni All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of the copyright holder and any other contributors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

GNU General Public License

PREAMBLE

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The ``Program'', below, refers to any such program or work, and a ``work based on the Program'' means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term ``modification''.) Each licensee is addressed as ``you''.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

Accompany it with a written offer, valid for at least three years, to give any third-party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions

for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

Documentation

For the latest version of this document, go to <http://www.hp.com/go/3par/>, navigate to your product page, click **Support for your product**, and then click **Manuals**.

Contents

1	Introduction	18
	User Interfaces	18
	Related Documentation	18
	Typographical Conventions	19
	Advisories	19
2	Overview and Features	20
	About System Reporter	20
	Configuring System Reporter Policy Settings	20
	Choosing a Method for Accessing Reports	21
	Creating Customized Reports	21
	Formatting Reports for Presentation	23
	Customizing the Reporting Format	23
	Generating Email Alerts	23
	Accessing Report Data using Web Queries	23
	Using the Database Schema to Build Custom Reports	24
	Analyzing and Optimizing Storage Utilization	24
3	Installation, Configuration, Update, and Removal	25
	Before You Begin	25
	System Reporter CD Contents	25
	System Requirements	26
	Using the System Reporter Sizing Spreadsheet	27
	System Reporter Components	28
	Installing the System Reporter Components	30
	Required Components	30
	Optional Components	30
	Installing the Web Server	31
	Installing the Database Server	32
	Choosing the Appropriate Database	32
	SQLite	32
	Microsoft SQL	33
	MySQL	33
	Oracle	33
	Installing and Setting Up the Microsoft SQL Server Database (Optional)	34
	System Reporter Host Configuration	34
	Installing and Setting Up the MySQL Database (Optional)	38
	Installing MySQL Server	38
	Configuring the MySQL Database	39
	Setting the max_allowed_packet parameter	40
	Creating MySQL Users and Schema	40
	Setting Up the Oracle Database (Optional)	41
	Creating an Oracle database	41
	Creating the Oracle Users for System Reporter	42
	Creating the Oracle Net Service	42
	Installing or Upgrading the System Reporter Tools	43
	Installing or Upgrading System Reporter Tools on Windows	43
	Installing System Reporter Tools on Linux	46
	Adding Storage Systems	47
	Setting File Permissions (Optional)	48
	Verifying Installation	48
	Installing the Excel Client (Optional)	48

About Apache HTTP Server Authorization and Access Control.....	49
Removing the System Reporter Components.....	50
Removing the System Reporter Tools on Windows.....	50
Removing the System Reporter Tools on Linux.....	51
Removing the MySQL Database (Optional).....	51
4 Getting Started.....	52
Accessing the Main System Reporter Interface.....	52
System Reporter Main Menu Window.....	52
Getting Help.....	54
Using the Object Selection Controls.....	54
Simple Selection List.....	54
Filtered Multi-Selection Control.....	54
Accessing the Optional Excel Client.....	55
Starting the System Reporter Excel Client.....	55
Setting the Security Level for Excel.....	56
Setting the Security for Excel 2003.....	56
Setting the Security Level for Excel 2007.....	57
Connecting to the Web Server from the Excel Client.....	59
Changing the Excel Client Server and Resetting the Workbook.....	61
Changing the Web Server.....	61
Resetting the Workbook.....	61
Accessing the Excel Client Query Log.....	61
Deleting Excel Sheets.....	61
Deleting Menu Sheets.....	61
Deleting Data Sheets.....	62
Deleting the Query Log.....	62
Saving and Exiting the Excel Client.....	62
5 Accessing and Creating Reports.....	64
Accessing Quick Reports.....	103
Accessing Scheduled Reports.....	65
Creating Custom Reports.....	66
Building the Report Menu.....	67
Sample Resolution.....	67
Report.....	68
Select Systems.....	68
Select Domains.....	69
Generating the Report.....	69
Controls Available by Report.....	72
Report Selection.....	76
Standard Report Types.....	76
Select Summary Columns.....	81
Compare.....	81
Max. Number.....	82
Select Peak.....	82
Group By.....	82
Order By.....	83
Low Bucket.....	84
High Bucket.....	84
Sum End Buckets.....	84
Time Selection.....	84
Reload After (mins).....	84
Begin/At Time.....	84
End Time.....	85
Count.....	85

Object Selection.....	85
Select PDIDs.....	85
Include Chunklets.....	85
Select Cage IDs.....	86
Select Disk Types.....	86
Select Disk Speed.....	86
Include PDs in States.....	86
Select Usr CPGs.....	87
Select Snp CPGs.....	87
Select CPGs.....	87
Select LDs.....	88
Select RAID Types.....	88
Select LD Usage.....	88
Select LUN.....	88
Select VVs.....	89
Select Prov Types.....	89
Select VV Types.....	89
Select Port Types.....	90
Select Port Rates.....	90
Select Ports (n:s:p).....	90
Select Hosts.....	91
Select Nodes.....	91
Select To Nodes.....	91
Select Queues.....	91
Format Selection (Web interface only).....	91
Report Information.....	92
Chart Lib.....	92
Chart Types.....	92
Chart X Pixels.....	96
Chart Y Pixels.....	96
Time Labels.....	96
Legend Position.....	96
Time Format.....	96
Select Space Unit.....	97
Working with Excel Charts.....	97
Choosing a Chart Type.....	97
Refreshing Charts.....	98
Retaining Charts.....	99
Exporting Executable Reports.....	100
Customizing Excel Reports.....	100
Adding Custom Columns to a Data Sheet.....	101
Creating a Custom Chart.....	101
Saving Customized Settings.....	102
6 Interpreting Report Metrics.....	103
Summary Report.....	103
Space Reports.....	106
PD Space Reports.....	106
CPG Space Reports.....	106
LD Space Reports.....	107
VV Space Reports.....	107
Performance Reports.....	108
Common Performance Metrics.....	108
VV Cache Performance Reports.....	109
Node Cache Performance Reports.....	109

CPU Performance Reports.....	110
Link Performance Reports.....	110
Adaptive Optimization Reports.....	111
7 Configuring the System Reporter Policy Settings.....	112
Accessing the System Reporter Policies.....	112
Configuring Sampling Policies.....	112
Compacting Databases.....	113
Editing Sampling Policies.....	114
Adding, Removing and Re-Configuring HP 3PAR Storage Systems.....	115
Adding a Storage Server.....	115
Verifying Addition of a Storage Server.....	115
Removing a Storage Server.....	115
Re-Configuring a Storage Server Sampling.....	116
Configuring Rules for Email Alerts.....	116
Adding an Alert Rule.....	117
Changing an Alert Rule.....	117
Removing an Alert Rule.....	118
Alert Rule Parameters.....	118
Data Table.....	118
Resolution.....	119
System.....	119
Metric.....	119
Direction.....	121
Limit Value.....	121
Limit Count.....	121
Condition.....	121
Condition Value.....	121
Recipient.....	121
Example of an Alert Rule.....	121
Scheduling Reports.....	122
Adding a Scheduled Report.....	123
Changing a Scheduled Report.....	123
Removing a Scheduled Report.....	124
Configuring Adaptive Optimization.....	124
About File Permissions.....	124
Modifying the Apache HTTP Server Configuration Files.....	125
8 Web Queries.....	126
About Web Queries.....	126
Data Format Version Report.....	190
About Reports and Options.....	127
Report Options.....	128
category.....	128
example:.....	128
compare.....	128
values:.....	128
example:.....	130
comparesel.....	130
values:.....	130
table.....	131
values:.....	131
example:.....	131
refresh.....	131
values:.....	131
example:.....	132

charttab.....	132
values:.....	132
example:.....	132
charttype.....	132
values:.....	132
example:.....	133
graphx.....	133
value:.....	133
example:.....	133
graphy.....	133
value:.....	133
example:.....	133
graphlegpos.....	134
value:.....	134
example:.....	134
tableformat.....	134
value:.....	134
example:.....	134
alllabels.....	134
values:.....	134
example:.....	135
timeform.....	135
values:.....	135
example:.....	135
begintsecs.....	135
value:.....	135
example:.....	136
endtsecs.....	136
value:.....	136
example:.....	136
groupby.....	136
values:.....	136
example:.....	138
orderby.....	138
values:.....	138
example:.....	140
histbegin.....	140
values:.....	140
example:.....	141
ldspaceunit.....	141
values:.....	141
examples:.....	141
pdspaceunit.....	141
values:.....	141
examples:.....	142
vvspaceunit.....	142
values:.....	142
examples:.....	142
selcageid.....	142
value:.....	143
examples:.....	143
selchunks.....	143
values:.....	143
example:.....	144
selcpg.....	144

value:.....	144
example:.....	144
selsrcpg.....	144
selsnpcpg.....	145
seldisktype.....	145
value:.....	145
example:.....	145
seldiskspeed.....	145
value:.....	145
example:.....	145
seldomain.....	146
value:.....	146
example:.....	146
selhost.....	146
value:.....	146
example:.....	146
selld.....	147
value:.....	147
example:.....	147
sellduse.....	147
values:.....	147
example:.....	148
selnode.....	148
value:.....	148
example:.....	148
selnsp.....	148
value:.....	148
example:.....	149
selpdid.....	149
value:.....	149
example:.....	149
selpdstate.....	149
values:.....	149
example:.....	150
selporttype.....	150
values:.....	150
example:.....	150
selraidtype.....	150
values:.....	150
example:.....	151
selsys.....	151
values:.....	151
example:.....	151
selvv.....	151
values:.....	151
example:.....	152
selvtype.....	152
values:.....	152
selprovtype.....	152
values:.....	152
Report Types.....	153
Summary Report.....	153
summary.....	153
example:.....	153
Space Reports.....	153

pd_space_time.....	153
example:.....	153
pd_space_group.....	154
example:.....	154
ld_space_time.....	154
example:.....	154
ld_space_group.....	154
example:.....	154
vv_space_time.....	154
example:.....	155
vv_space_group.....	155
example:.....	155
Performance Reports.....	155
pd_perf_time.....	155
example:.....	155
pd_perf_group.....	155
example:.....	156
pd_svt_hist.....	156
example:.....	156
pd_svt_histvstime.....	156
example:.....	156
port_perf_time.....	156
example:.....	156
port_perf_group.....	156
example:.....	157
port_svt_hist.....	157
example:.....	157
port_svt_histvstime.....	157
example:.....	157
ld_perf_time.....	157
example:.....	157
ld_perf_group.....	157
example:.....	158
ld_svt_hist.....	158
example:.....	158
ld_svt_histvstime.....	158
example:.....	158
vlun_perf_time.....	158
example:.....	158
vlun_perf_group.....	158
example:.....	159
vlun_svt_hist.....	159
example:.....	159
vlun_svt_histvstime.....	159
example:.....	159
vv_cmp_time.....	159
example:.....	159
vv_cmp_group.....	159
example:.....	159
cpu_perf_time.....	160
example:.....	160
cpu_perf_group.....	160
example:.....	160
Adaptive Optimization Reports.....	160
regiodensity.....	160

regiondensity example.....	160
regmoves.....	160
regmoves example.....	161
vttier.....	161
vttier example.....	161
Options Available by Report.....	161
Objects Present Reports.....	167
objectpres= cageid.....	168
example:.....	168
output:.....	169
objectpres= cpg_name.....	169
example:.....	169
output:.....	169
objectpres= diskspeed.....	170
example:.....	170
output:.....	170
objectpres= disktype.....	170
example:.....	170
output:.....	170
objectpres= domain_name.....	171
example:.....	171
output:.....	171
objectpres= host_name.....	171
example:.....	171
output:.....	171
objectpres= ld_name.....	172
example:.....	172
output:.....	172
objectpres=node.....	172
example:.....	172
output:.....	173
objectpres=pdid.....	173
example:.....	173
output:.....	173
objectpres=port.....	174
example:.....	174
output:.....	174
objectpres= system.....	174
example:.....	174
output:.....	174
objectpres= vv_name.....	175
example:.....	175
output:.....	175
objectpres= vv_name_nosnap.....	175
example:.....	175
output:.....	176
objecctpres= porttype.....	176
example:.....	176
output:.....	176
sample_times.....	177
example:.....	177
output:.....	177
9 Using Adaptive Optimization.....	178
Overview.....	178

Selecting an HP 3PAR Storage System for Adaptive Optimization.....	179
Configuring Adaptive Optimization.....	179
Tier Definition (CPG Name and GiB).....	180
Schedule.....	182
Measurement Hours.....	182
Adaptive Optimization Mode.....	183
Adaptive Optimization Active.....	183
Adding an Adaptive Optimization Configuration.....	183
Changing an Adaptive Optimization Configuration.....	184
Removing an Adaptive Optimization Configuration.....	184
Adaptive Optimization Reports.....	184
Region IO Density Report.....	184
VV Tiers Report.....	186
Adaptive Optimization Space Moved Report.....	187
The showvvcpg CLI extension.....	188
10 Database Schema.....	190
Using The Database Schema.....	190
Table Versions and Naming Conventions.....	190
Administrative Tables.....	190
The policy Table.....	190
The sampleinserv Table.....	190
The system Table.....	191
The alertconfig Table.....	192
The scheduled_report Table.....	193
The scheduled_report_log Table.....	193
The adprovconfig Table.....	193
The adprov_log Table.....	194
The ldrgmoves Table.....	194
The recentsample Table.....	195
The Sample Data and Inventory Tables.....	196
The pdspace Table.....	196
The pdspaceinv Table.....	197
The ldspace Table.....	198
The ldspaceinv Table.....	198
The vvspace Table.....	199
The vvspaceinv Table.....	201
The cpgspace Table.....	201
The cpgspaceinv Table.....	203
Common Performance Table Columns.....	203
Calculating Common Performance Metrics.....	205
Service Time Histogram Buckets.....	206
IO Size Histogram Buckets.....	206
The statpd Table.....	206
The statpdinv Table.....	207
The statld Table.....	207
The statldinv Table.....	208
The statport Table.....	208
The statportinv Table.....	209
The statvlun Table.....	209
The statvluninv Table.....	210
The statvvcmp Table.....	210
The statvvcmpinv Table.....	211
The statcpu Table.....	211
The statcpuinv Table.....	212

The statlink Table.....	212
The statlinkinv Table.....	213
The statldrg Table.....	213
The statldrg_timeinv Table.....	214
The perf_timeinv Table.....	214
The space_timeinv Table.....	214
11 Troubleshooting.....	215
Troubleshooting Overview.....	215
Troubleshooting the Sampling Components.....	215
Sampling Component Log Files.....	215
Installation Configuration.....	216
Sampling Errors.....	216
Errors incurred by the insertsample process.....	216
Errors incurred by the CLI sampler process.....	217
Sampler Service/Daemon Errors.....	217
Troubleshooting Adaptive Optimization Components.....	217
Troubleshooting the Web Server Components.....	217
Troubleshooting the Excel Client.....	217
Extracting Data from the Database with sysbck.....	218
Examples of Using sysbck for Various Databases.....	219
Dumping the Entire Database.....	222
Dumping the Entire SQLite Database.....	222
Dumping the Entire MySQL Database.....	222
Index.....	223

1 Introduction

This guide provides the information you need to install and use HP 3PAR System Reporter to monitor performance, create charge back reports, and plan storage resources for HP 3PAR Storage Systems. An optionally licensed component called Adaptive Optimization is also included with System Reporter that can be used to monitor and automatically optimize utilization of storage resources. If you do not have a licence for Adaptive Optimization, the monitoring and system reporting features are still accessible but the optimization capabilities are inactivated.

User Interfaces

NOTE: The InServ Storage Server has been rebranded as HP 3PAR Storage System. There are instances in this document where screenshots and/or menu items and command output refer to the HP 3PAR Storage System as InServ or InServ Storage Server.

The HP 3PAR InForm® Operating System (InForm OS) has two user interfaces: the HP 3PAR InForm Command Line Interface (CLI) and the HP 3PAR InForm Management Console (IMC). You can use the IMC or the InForm CLI to monitor storage servers and access statistics. From the IMC or InForm CLI, you also can export table data to a .CSV (Comma-Separated Variable) file and then use Microsoft® Excel® or similar software to import that file. However, to generate reports directly in Microsoft Excel and to access historical information, you must use the system reporting tools described in this guide.

Two user interfaces are offered as part of the HP 3PAR Storage System: a Web-based interface and an Excel client.

- The Web-based interface enables you to monitor HP 3PAR Storage Systems and access statistics using a standard Web browser. Configuration of the System Reporter database sampling policies must be done via a Web browser.
- The Excel client enables you to monitor HP 3PAR Storage Systems and access statistics using Microsoft® Excel®. This option also offers additional customization options not offered by the Web-based interface.

Related Documentation

The following documents provide information related to HP 3PAR Storage Systems, the InForm® Operating System, the InForm® Management Console and InForm® CLI:

For information about...	Read the...
Storage system hardware specifications, site-planning, environmental considerations, power requirements, cabling, network configurations and third-party rack implementations	<i>HP 3PAR Storage System Physical Planning Manual</i>
Identifying storage server components and fixing problems	<i>HP 3PAR InForm OS Message and Operator's Guide</i>
Using the HP 3PAR InForm OS Command Line Interface (CLI)	<i>HP 3PAR InForm OS Command Line Interface Reference</i>
Configuring and managing HP 3PAR Storage Systems and using the InForm Management Console (IMC)	<i>HP 3PAR InForm OS Management Console Online Help</i>

Typographical Conventions

This guide employs the following typographical conventions:

Typeface	Meaning	Example
ABCDabcd	Used for dialog elements such as titles, button labels, and other .	When prompted, click Finish to complete the installation.
ABCDabcd	Used for paths, filenames, and screen output.	Open the file <code>\gui\windows\setup.exe</code>

Advisories

To avoid unexpected configuration or operational problems, be sure to observe the notes and cautions in this guide.

 **WARNING!** Warnings alert you to actions that can cause injury to people or irreversible damage to data or the operating system.

 **CAUTION:** Cautions alert you to actions that can cause damage to equipment, software, or data.

NOTE: Notes are reminders, tips, or suggestions that supplement the procedures included in this guide.

2 Overview and Features

This chapter provides a brief overview of the features and capabilities offered by System Reporter.

About System Reporter

System Reporter monitors performance and the usage of storage resources and allows you to generate charts and graphs that report useful statistics for planning and configuring the operation of HP 3PAR Storage Systems.

System Reporter provides the following main features:

- Convenient access to configuration options for selecting systems to include for reporting, specifying sampling parameters, scheduling reports, and generating alerts.
- Extensive selection of reports for obtaining performance and storage utilization statistics on selected objects (i.e., hosts, ports, nodes, physical disks, virtual volumes, etc.).
- Quick access to pre-defined reports that contain useful statistics for most common types of installations.
- Scheduling of reports with predefined parameters that are initiated at predetermined times and can then be accessed when needed.
- Customization of reports using the standard web interface, (or through the provided Excel client), that provide specifically selected and formatted reports for specified systems.
- Options for choosing the time and duration for the collection of reporting statistics which can be initiated at a specific time, collected over a period of time, and/or compared between a range of times.
- Options for viewing and comparing report information in a variety of formats through a selection of charts and tables.
- Alerts that can be configured to send email notifications to a specified address when certain reporting criterion are met.
- Support for customized formatting of report data via the Comma Separated Values (CSV) file formatting standard for inclusion in web or Excel applications.
- Support for creating custom web or Excel reports using well documented Web queries.
- Access to the database schema for direct queries to the reporting data stored in the database used by System Reporter.
- Inclusion of an optionally licensed component called Adaptive Optimization that can analyze access rates for sub-volume level regions and move regions between tiers of storage for more optimal and efficient storage utilization.

Configuring System Reporter Policy Settings

Although you can begin using System Reporter immediately after it has been installed to quickly access reports, you may want to make some adjustments or set up some of the capabilities for your particular operating environment by selecting the appropriate configuration options through the Policy Settings.

The following configuration options are available from the Policy Settings:

- **Sampling Policies** – Specifies the report sampling policies that determine sampling intervals, how long reports are stored, when to perform database compaction, etc.
- **InServ Systems** – Selects the HP 3PAR Storage Systems to include for sampling, report generation and optional optimization. The systems are generally added during the installation, but you can access this option to add or remove systems as needed.

- **Alert Rules** – Used to add, modify and remove alert rules and to specify the parameters for generating email reports.
- **Scheduled Reports** – Used to add, modify, and remove scheduled reports for generation at a specified time.
- **Adaptive Optimization** – Configures the optionally licensed Adaptive Optimization components that are used to monitor and optimize storage utilization.

See for details on how to configure the policy settings for your site.

Choosing a Method for Accessing Reports

There are three basic methods for accessing reports through System Reporter:

Quick Reports – Allows you to immediate access a variety of pre-defined reports (created through CGI programs) that are deemed useful for most installations. The reports are made available through a menu tree that you expand and collapse to select the systems and options of interest. For instance, one of the reports provides a comparison of the peak hourly system CPU utilization for the past seven days for all systems while another compares the utilization of the most active physical disks for a given system.

Scheduled Reports – Allows you to view reports that were created at scheduled times based on pre-selected parameters and stored in a reserved directory structure. You can either view the reports via the web interface provided through System Reporter or you can copy the report subdirectories to another area and view them there. Maintaining a report directory for scheduled reports allows an administrator to limit the users who have access to each report directory. One example where this may be useful is when multiple departments share an array. In this case, the administrator can schedule various reports specifically designed for each department, place them in different report directories, and then grant each department access to the relevant scheduled reports in their own report directory.

Custom Reports – Custom reports are more complex to generate because they allow for a great deal of customization, but they are also very useful. You can select the resolution, the type of report, and then specify which systems and domains are to be monitored for the reporting information that is collected. You also have the choice of specifying how the information is to be formatted for presentation through a selection of tables and graphs.

Creating Customized Reports

System Reporter provides a number of options for building customized reports as follows:

- Report Selection. Allows you to select from one of the reports shown in [Table 1 \(page 21\)](#):

Table 1 Available Reports

Report Name	Description	Report Metrics Cross Reference
Summary	Displays summary information.	“Accessing Quick Reports” (page 103)
PD Space	Displays physical disk space usage.	“PD Space Reports” (page 106)
CPG Space	Displays CPG disk space usage.	“PD Space Reports” (page 106)
LD Space	Displays logical disk space usage.	“LD Space Reports” (page 107)
VV Space	Displays virtual volume space usage.	“VV Space Reports” (page 107)
Port Performance	Displays port performance.	“Common Performance Metrics” (page 108)

Table 1 Available Reports *(continued)*

Report Name	Description	Report Metrics Cross Reference
VLUN Performance	Displays VLUN (volume-LUN) performance.	"Common Performance Metrics" (page 108).
LD Performance	Displays logical disk performance.	"Common Performance Metrics" (page 108)
PD Performance	Displays physical disk performance.	"Common Performance Metrics" (page 108)
VV Cache Performance	Displays VV (virtual volume) cache performance.	"VV Cache Performance Reports" (page 109)
Node Cache Performance	Displays Node Cache Performance.	"Node Cache Performance Reports" (page 109)
CPU Performance	Displays CPU performance.	"CPU Performance Reports" (page 110)
Link Performance	Displays performance of links between nodes.	"Link Performance Reports" (page 110)
Adaptive Optimization	Displays Adaptive Optimization reports.	"Adaptive Optimization Reports" (page 111)

- **Sample Resolution** – Specifies how often operational metrics are sampled for a given report:
- **Daily** – Sampling is performed once a day.
- **Hourly** – Sampling is performed on an hourly basis.
- **Hi Resolution** – Sampling is performed at much more frequent intervals as defined by the "Requested High-Res sampling interval in minutes" option defined through the Sampling Policies setting which offers the finest granularity of operational statistics.

The retention period for the sampled data in the database tables is set through the sampling policy (see "[Configuring Sampling Policies](#)" (page 112)). The data is purged based on the limit set through the policy settings. You also have the ability to set the data retention period for individual systems (see for details).

- **System and Domain Selection** – Specifies the systems and domains that are to be included in the report.
- **Scope of Report Selection** – The options for all reports (except for the Summary report), are divided into four sections:
 - **Versus Time** reports plot data over a period of time (time is along the X-axis). These reports can either plot aggregate metrics or separate metrics for up to 16 objects. If you want to compare more than 16 objects, or do more complex object selection you must use the **At Time** reports described below.
 - **At Time** reports allow you to compare a large number of objects, defined in more complex ways, but only at a specific point in time. The objects are plotted along the X-axis.
 - **Histogram At Time** reports show the distribution of service times and IO sizes at a specific point in time.
 - **Histogram vs Time** reports show the distribution of service times and IO sizes over a period of time (time is along the X-axis).
- **Time Selection** – Selects the time or range of times when the report is to be generated.

- **Object Selection** – Selects the list of objects that are to be included in the reporting data. For example, you can select the VVs, ports and hosts for a VLUN.
- **Format Selection** – Selects the format of the report including whether the report should show charts and/or tables and the type and size of charts.

See for complete details on how to build customized reports.

Formatting Reports for Presentation

When building customized reports, you can specify how the reporting data is formatted for viewing by selecting the appropriate table or type of chart. You can choose from the following chart types:

- Stacked Bars
- Staked Vertical Bars
- Stacked Area
- Bars
- Vertical Bars

In addition to choosing the type of chart that is best suited for the report being generated, you can also fine tune the presentation by selecting the position of the legend, choosing a time format and specifying how the units of storage measurement are represented (i.e., MBs, TBs, chunklets, etc.).

When using the Excel client to build reports, you have additional flexibility in the formatting of a report by modifying the datasheet and including customized columns.

See for descriptions and options available for formatting reporting data.

NOTE: Access to a currently unsupported charting package is also available in this version and offered as one of the formatting options (under **Chart Lib**) when generating a report.

To use the unsupported charting capabilities offered through this “Dynamic” charting option, you must use a browser that supports html5 canvas tags (the latest versions of Windows E9, Firefox and Google Chrome browsers currently offer this support).

Customizing the Reporting Format

Although System Reporter provides a number of built-in controls for specifying the format of a report, in some cases you may want to present the data using a different format that is not available from the standard selection. In this case, you can save the data for the report using the Comma Separated Values (CSV) format as an option that is accessible with the click of an icon next to the heading of the report. You can then open the CSV formatted file in an application (such as a spreadsheet) where it can then be charted or formatted as desired. If the browser is set up to open files with a `.csv` extension in a spreadsheet application, this can be done automatically.

Generating Email Alerts

You can configure System Reporter to send email alerts when certain metrics meet specified conditions. For example, suppose you want to receive an email alert when any VLUN has an average read service time of more than 100 ms in any high-resolution sampling interval. To do this, all you need to do is fill in a form with the specified details and then submit the query.

See for complete details on how to configure email alerts.

Accessing Report Data using Web Queries

The querying mechanism employed by System Reporter is made available for specialized customization of reports. For instance, instead of using the Excel client provided with System Reporter, you can use web queries in an Excel spreadsheet to bring various report data into your spreadsheets. As another alternative, you might like to view multiple standard reports presented

together in a single web page. You can do this by creating an html file that includes multiple iframe elements, each of which includes a standard web report query.

See “Web Queries” (page 126) for details on how to instigate web queries.

Using the Database Schema to Build Custom Reports

The databases used by System Reporter are standard relational databases capable of supporting SQL queries. The default (built-in) database is SQLite (see <http://www.sqlite.org>) but MySQL and Oracle database servers can optionally be used instead. The database schema is documented so you can write your own reports by directly querying the System Reporter database.

See for details on how to utilize the database schema.

Analyzing and Optimizing Storage Utilization

An optionally licensed component called Adaptive Optimization is included with System Reporter. Adaptive Optimization analyzes sub-volume, region level disk access rates for a given array over a scheduled period of time and then performs a data migration of regions between tiers according to a cost versus performance preference. Disk utilization is optimized by having frequently accessed data moved to the higher performance tier (e.g., RAID 1 using Solid State Disks or SSDs) while infrequently accessed data is moved to the lower cost tier (e.g., RAID 6 on Nearline disks). By using Adaptive Optimization, you can achieve a much higher degree of control over disk usage by reserving your faster and more expensive storage resources for the data that is frequently accessed and relegating your slower and less expensive drives to storing data that is only occasionally accessed.

3 Installation, Configuration, Update, and Removal

This chapter describes how to install and configure as well as remove the HP 3PAR System Reporter components.

Before You Begin

Before installing HP 3PAR System Reporter components, select a system on which to run the System Reporter sampler and Web server. This system must use Windows® Server 2003, Windows® Server 2008 or Red Hat® Enterprise Linux 5.

If necessary, perform the following setup activities on the designated Windows® or Red Hat® Linux server:

- Install the InForm OS Command Line Interface (CLI), version 2.3.1 or above.
- Install the Apache Server.
- If you are using an Oracle database, you must install the Oracle client.
- (Optional). To avoid a DHCP-assigned address changing, assign a fixed IP address to the system.

NOTE: The CLI installer executable for the InForm OS can be found on the System Reporter CD at: `Windows\CLI\setup.exe` (for Windows) or `Linux/CLI/setup.bin` (for Linux)

See the *InForm OS Management Console Online Help*, located on the System Reporter CD at: `Document/CLI\231_CLI_Administrator_Manual.pdf`, for instructions on installing the InForm CLI.

System Reporter CD Contents

InServ Storage ServerCD includes the following contents (Table 2 (page 25)):

Table 2 InServ Storage Server CD Contents

Directory	Filename	Description
/Documents/	CLI\ 231_CLI_Administrator_Manual.pdf	HP 3PAR InForm OS CLI Administrators Guide.
	SR_Sizing.xls	System Reporter server sizing spreadsheet (see “Using the System Reporter Sizing Spreadsheet” (page 27)).
	System_Reporter.pdf	HP 3PAR System Reporter Software User Guide (this document).
/Linux/	sampleloop-2.9-1.i386.rpm	System Reporter sampleloop package for Red Hat Linux.
	sysrptwebserv-2.9-1.i386.rpm	System Reporter sysrptwebserv package for Red Hat Linux. (Version 2.3.1 of the InForm CLI has only been validated for RedHat 5.2. Version 2.2.4 of the InForm CLI has been validated for RedHat 5.0 and 5.1.)
/Linux/ CLI/	setup.bin	HP 3PAR InForm OS CLI installer executable for Red Hat Linux.
/Scripts/	showvvcpg.tcl	CLI script implementing the CLI showvvcpg extension command (see “The showvvcpg CLI extension” (page 188)).
/Windows/	apache-2.2.4-win32-x86-msvc-0.9.8.exe	Apache HTTP Server installer executable for Microsoft Windows. This installer is provided only as a convenience. You may choose to download and install directly from: http://httpd.apache.org .

Table 2 InServ Storage Server CD Contents *(continued)*

Directory	Filename	Description
	installer.exe	System Reporter installer executable for Microsoft Windows.
/Windows/ CLI/	setup.exe	HP 3PAR InForm OS CLI installer executable for Microsoft Windows.

System Requirements

HP 3PAR System Reporter supports the use of the HP 3PAR InForm OS, versions 2.2.4 through 2.3.1 MU1 and beyond.

Observe the following minimum system hardware requirements:

Table 3 Minimum Hardware Requirements

System Component	Minimum Requirement
CPU	Pentium 4, 3GHz (or faster)
Memory	1 GB
Disk	20 GB free space (application and data)

For better performance, it is recommended that you use separate disks for application data and database data.

NOTE: When running System Reporter on a Virtual Machine, it is best to have the database server running on a separate VM, especially for Microsoft SQL, MySQL or Oracle.

HP 3PAR System Reporter components have the following requirements ([Table 4 \(page 26\)](#)):

Table 4 HP 3PAR Storage System Requirements

Component	Supported Operating Systems	Additional Requirements
Database	Windows® Server 2008 Red Hat Enterprise Linux 5.4 32 or 64-bit OS versions are supported except that only the 32-bit versions of Linux are supported for the Sampling service when using the Oracle 11g Database since we need to use the 32-bit Oracle 11g client which is only supported on the 32-bit Linux.	MySQL® Database Server version version 5.1 (Optional) Oracle 11g® Database (Optional) (Oracle RAC is not currently supported.)
Microsoft SQL Server Note that Microsoft SQL is not supported on Linux platforms.		Microsoft SQL Server Note that Microsoft SQL is not supported on Linux platforms.
Sampling service		Oracle 11g Client (32-bit only) when used with Oracle 11g Database
Web server		Apache HTTP Server 2.0 or later (installer provided for version 2.2.14 on Microsoft Windows)
Web Browser (one of the following)		<ul style="list-style-type: none"> Internet Explorer 7 or later Mozilla Firefox 3.0 or later (To use the dynamic charting capabilities offered as an unsupported formatting option, you must have Firefox 3.6 or later or IE 9.)
Excel client		Microsoft® Excel 2003, 2007 (not supported on Mac OS X)

Table 4 HP 3PAR Storage System Requirements (continued)

Component	Supported Operating Systems	Additional Requirements
Adaptive Optimization		InForm OS version 2.3.1 MU1 and beyond. Not supported when using SQLite database.
On-Line Help (User's Guide)		Adobe Acrobat Reader

NOTE: The optionally licensed Adaptive Optimization component is only supported by InForm OS versions 2.3.1 MU 1 and beyond. Adaptive Optimization is not supported for SQLite databases; only Oracle and MySQL are supported.

Using the System Reporter Sizing Spreadsheet

Before installing System Reporter, you need to determine the system resources that are required to adequately run System Reporter and the selected data base server. The optimal computing power (CPU cores), memory and storage space needed to accommodate the data that is sampled by System Reporter and used to generate reports depends on the number of systems and the amount of storage that is being monitored for your particular installation. An Excel spreadsheet named SR_Sizing.xls (Figure 1 (page 28)) is included with System Reporter for the purpose of determining the optimal system resources that are recommended for running System Reporter and the database server at your site. This spreadsheet is included on the installation CD.

NOTE: The System Reporter Sizing Spreadsheet is protected except for the cells in red, but you can turn off the protection to see the formulas that are used to derive a given value.

Entities in the spreadsheet are color coded. Only the values into controls or cells colored red need to be set. Cells or ranges in yellow are output values that you can use to size or configure the server that you will use to run System Reporter. Values in blue are empirically determined parameters used to drive the sizing model.

Figure 1 System Reporter Sizing Spreadsheet

The screenshot displays a spreadsheet titled 'System Reporter Sizing Spreadsheet' with columns labeled A through X. The spreadsheet is divided into several sections:

- Instructions:** Located at the top, providing guidance on how to use the spreadsheet.
- Table Row Sizes:** A table with columns for 'Table Name', 'Row Size (Bytes)', 'Number of Rows', and 'Total Size (Bytes)'. Rows include tables like 'table_row_sizes', 'table_row_sizes_index', etc.
- Sampling Parameters:** A section with red headers for 'Sample Rate' and 'Sample Size', with values for various tables.
- Database Requirements:** A section with yellow headers for 'Database Requirements', including 'Table Row Sizes', 'Table Row Sizes Index', and 'Table Row Sizes Index Index'.
- System Resources:** A section with yellow headers for 'System Resources', including 'CPU Requirements', 'Memory Requirements', and 'Disk Requirements'.

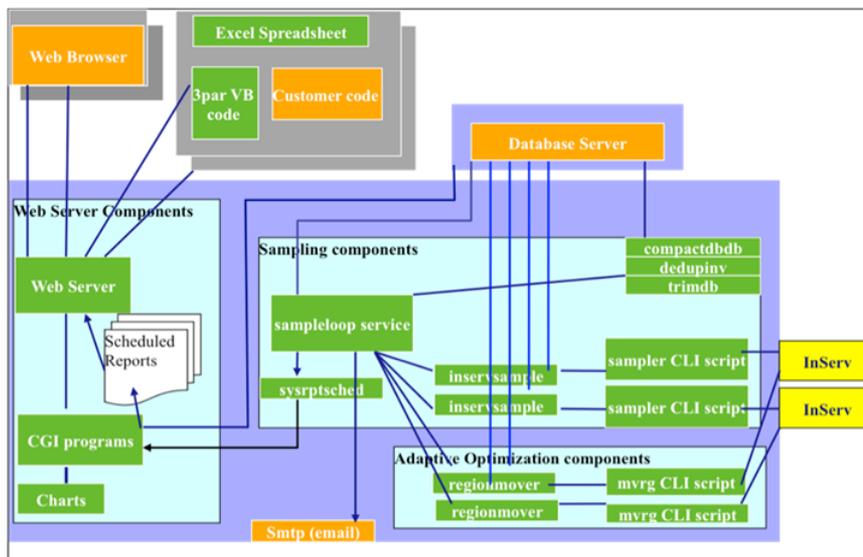
1. Click on the appropriate type of database that you are using for System Reporter at your site. You will notice that SQLite only supports one HP 3PAR Storage System for use with System Reporter. Please consider another database if you plan on using System Reporter to monitor more than a single HP 3PAR Storage System.
2. Modify the **Sampling Parameters** in red to reflect the approximate values for your installation. The values for **Table Row Sizes** in blue are approximate table row sizes (bytes) based on observations derived by HP during testing System Reporter instances and includes some extra space to accommodate typical inventory tables, indexes, etc.
3. Observe the requirements that are shown in the yellow cells to determine the suggested allocation of system resources for your System Reporter server or database server.

NOTE: The database server should have roughly the same amount of memory and CPU cores as the System Reporter server.

System Reporter Components

System Reporter includes the following components (Figure 2 (page 29)):

Figure 2 HP 3PAR Storage System and User Interface Components



- **Database.** The HP 3PAR System Reporter database is capable of supporting SQL queries. Four types of databases are supported: Microsoft SQL Server, SQLite, MySQL and Oracle (see [“Installing the Database Server”](#) (page 32)).
- **Sampler components.** The sampler components are responsible for sampling performance and space information from the HP 3PAR Storage System arrays, saving the data in the database, and removing old samples from the database. The sampler components include the following:
 - The sampling service (sampleloop) runs as a service on Windows or as a daemon on Linux. It creates and coordinates other sampler components.
 - A separate inservsample process is created for each HP 3PAR Storage System being sampled. This process in turn creates a separate CLI process that runs the sampler script which logs into the array and collects the required data. The inservsample process then inserts the data into the database. Along with the tables that contain the space and performance data, the database also includes inventory tables that allow quicker creation of report menus. The inservsample process also update the inventory tables.
 - For MySQL, Microsoft SQL Server (Windows only), and Oracle databases, the sampleloop service/daemon creates a separate process to maintain the database (since SQLite supports limited concurrency, these functions are performed by the sampleloop process). These include:
 - The trimdb process removes samples that are older than the sample retention period specified in the sampling policies (see [“Configuring Sampling Policies”](#) (page 112)).
 - The dedupinv process removes duplicate entries from the inventory tables.
 - The compactdb process periodically compacts and optimizes the database when necessary.
- **Adaptive Optimization components.** The sampleloop process also creates a regionmover process for each HP 3PAR Storage System array for which Adaptive Optimization is configured (see [“Selecting an HP 3PAR Storage System for Adaptive Optimization”](#) (page 179)). Each regionmover process analyzes region level performance data for each Adaptive Optimization configuration and generates region moves between tiers. The regionmover process creates a separate CLI process that logs into the array and executes the mvrg program that executes the data movement commands on the array.

- *Web server components.* The HP 3PAR System Reporter Web server has customized scripts that access the database to generate reports to standard browsers as clients, enabling users to access the reports and configure System Reporter without installing any additional software.
- *Excel client.* The customized HP 3PAR System Reporter Excel client has built-in Web queries to the Web server. Use of the Excel client is optional, but offers you additional options. You must have Microsoft® Excel® in order to use the HP 3PAR System Reporter Excel client.

Installing the System Reporter Components

Required Components

To begin using HP 3PAR System Reporter, you must install the following components which are supplied on the installation CD:

- CLI for InForm OS
- Web server (Apache HTTP Server)
- System Reporter tools (sampler, default SQLite database, and web server tools for generating reports)
 - *Windows*
 - *installer.exe*
 - *Linux*
 - *sysrptwebserv-2.9-1.i386.rpm*
 - *sampleloop-2.9-1.i386.rpm*

NOTE: To install and use System Reporter, you must have a valid license. If you plan use Adaptive Optimization, you must also have an additional license to take advantage of the storage optimization capabilities offered by this optional component.

Optional Components

Installation of the following components are optional:

- Excel client (optional). While the necessary Excel sheet is automatically included, Microsoft Excel is required by the machine accessing the Excel sheet from the System Reporter server.
- The Microsoft SQL, MySQL or Oracle database (optional). Except for the SQLite database which is supplied with System Reporter, any other choice of database must be obtained from the respective vendor.
- Adaptive Optimization (Offered by HP as an option and requires a separate license)

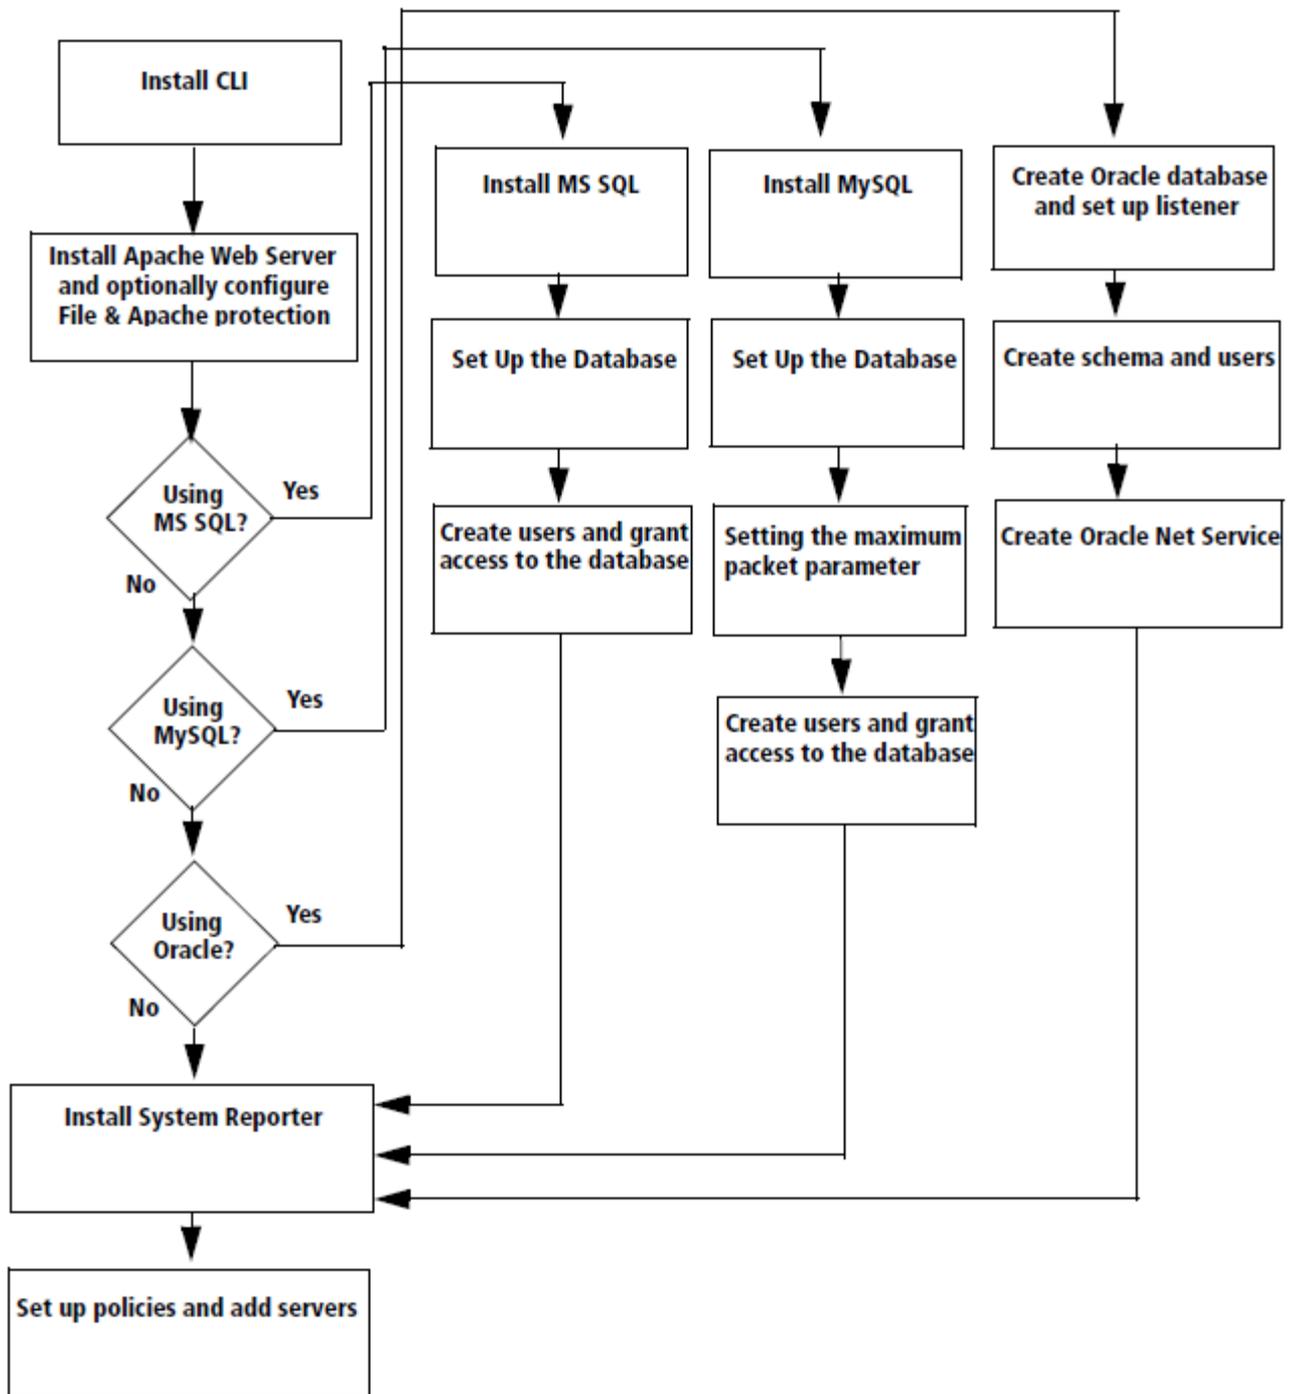
NOTE: If you are upgrading an existing System Reporter installation to the 2.9 release, you do not need to re-install the CLI, optional MySQL database or Apache web server. You only need to re-install the System Reporter Tools as described in [“Installing or Upgrading the System Reporter Tools” \(page 43\)](#). If you wish to keep the previously collected data, you must specify the same database during the installation process. The System Reporter installer for Microsoft Windows will automatically detect previous settings and offer the option to upgrade the existing installation.

System Reporter 2.9 will automatically convert the earlier release database tables into the format that it needs. If you wish to retain the tables in an earlier format, you must make a backup copy of the database before upgrading.

All table updates are logged in the Sampleloop.log. Contact your customer service representative if any tables fail to update properly.

The following flowchart ([Figure 3 \(page 31\)](#)) provides a visual representation of the installation process for a new installation that is described in detail in the following sections.

Figure 3 Installation Process Flowchart



Installing the Web Server

HP 3PAR System Reporter requires Apache HTTP Server, version 2.2 or later.

A copy of the installer executable for Windows of Apache HTTP Server, version 2.2, is located on the HP 3PAR System Reporter CD or you can download the most recent version from www.apache.org. Red Hat Linux typically includes the Apache Web Server.

NOTE: System Reporter 2.9 and beyond supports secure connections on Apache servers using the SSL protocol. For details on implementing the SSL protocol on Apache servers, please consult the following links:

http://httpd.apache.org/docs/2.2/ssl/ssl_intro.html

http://httpd.apache.org/docs/2.2/ssl/ssl_howto.html

http://httpd.apache.org/docs/2.1/ssl/ssl_faq.html

Use the **Apache HTTP Server Installation Wizard** to complete a typical installation, accepting all defaults (with possibly the exception of the Administrator's Email address).

To verify successful installation of Apache HTTP Server, open a Web browser and point it at the Web server (or **localhost** if the browser is on the same machine).

If the installation was successful, you will see a message in the Web browser that confirms that your installation was successful.

NOTE: On Linux, Apache should be configured to run as the default **apache** user and the sampleloop daemon will also be configured to run as the **apache** user. The default locations for html files (`/var/www/html`) and CGI files (`/var/www/cgi-bin`) should be used because the `sysrptwebsrv` package installs files in these directories.

Please note that System Reporter has not been tested with virtual domains in Apache.

NOTE: If SELinux (Security Enhanced Linux) is implemented at your site, you may need to either disable SELinux entirely or temporarily grant permissive access to the system to install the System Reporter tools. Please consult your Linux documentation for details on how to grant permissive access.

NOTE: It is recommended that you increase the default `Timeout` parameter to at least 360 seconds (or more for installations with very large databases) in the appropriate apache configuration file (typically `httpd.conf`) and then restart apache to let this change take effect. This is to ensure that CGI calls that take a long time to generate a report do not timeout.

On Windows, by default the `Timeout` parameter is specified in the following file:

```
C:\Program Files\Apache Software Foundation\Apache2.2\conf\
httpd-default.conf
```

On Linux, by default the `Timeout` parameter is specified in the following file:

```
/etc/httpd/conf/httpd.conf
```

Installing the Database Server

HP 3PAR System Reporter supports several database servers: Microsoft SQL, SQLite, MySQL and Oracle. Before beginning the installation, select the database based on the considerations described below and summarized in [Table 5 \(page 33\)](#).

Choosing the Appropriate Database

Before beginning, you want to choose the right database for your particular installation.

SQLite

SQLite is an embedded database that is included with System Reporter. It does not run as a separate application but rather as a module within the sampling tool and web server. Since it requires no installation or configuration, it is very simple to use but it has several limitations that make it only suitable for small database sizes and limited concurrency.

SQLite supports only limited database concurrency because it relies on a single file lock for the entire database. The sampling process needs exclusive access to the database when inserting new samples. If the database is very large or if there are a large number of concurrent web requests,

the sampler or the Web server can time out. Furthermore, database compaction (see “[Compacting Databases](#)” (page 113)) cannot be done in parallel with sample insertion so there could be a large gap in samples during database compaction.

Sampling of multiple HP 3PAR Storage Systems is not supported with an SQLite database, either the Microsoft SQL Server, MySQL or the Oracle database servers should be used instead.

For large databases (over about 1 GB) or a large number of concurrent Web users, either the MySQL®, Microsoft SQL, or the Oracle database servers should be used instead.

NOTE: The optionally licensed Adaptive Optimization feature is not supported with the SQLite database. Please consider another database if you plan on utilizing the storage optimization capabilities provided by Adaptive Optimization.

Microsoft SQL

System Reporter supports Microsoft SQL version 2008 R2 database server but it is not included and must be purchased separately. The Microsoft SQL server is not supported on Linux and does not support data generated by System Reporter when installed on a Linux platform.

MySQL

A MySQL database server can be used with System reporter but is not included with it. You must purchase it separately. MySQL is a good choice for System reporter because it supports MyISAM tables. These tables do not support transactions (which are not required for System Reporter) and are consequently somewhat higher performance than Oracle for sample insertion and deletion.

Oracle

System Reporter supports an Oracle 11g database server but it is not included with it. You must purchase it separately. Compared with MySQL, Oracle is slightly slower and more resource intensive for sample insertion and deletion. However, Oracle is also a good choice for large and active System Reporter installation.

NOTE: System Reporter does not currently support Oracle RAC implementations.

NOTE: System Reporter requires the 32-bit Oracle client and does not support the 64-bit version.

NOTE: If you choose Oracle as the database, you cannot install System Reporter on a 64-bit Linux system since it requires a 32-bit Oracle client which is not supported on 64-bit Linux. This limitation does not apply to the 64-bit Windows version since the 32-bit Windows version of Oracle Client works on 64-bit Windows platforms.

Table 5 (page 33) provides a comparison of features for the various database choices.

Table 5 Database Comparison

Comment	SQLite	Microsoft SQL	MySQL	Oracle
Included with System Reporter?	Yes	No	No	No
Sample more than one HP 3PAR Storage System?	No	Yes	Yes	Yes
Supports Adaptive Optimization?	No	Yes	Yes	Yes
Can have remote database server?	No	Yes	Yes	Yes
Scale to large database size?	No, recommend less than 1 GB	Yes	Yes	Yes

Table 5 Database Comparison (continued)

Comment	SQLite	Microsoft SQL	MySQL	Oracle
Scale to large number of concurrent users?	No	Yes	Yes	Yes
Performance for sample insertion and deletion?	Adequate for small databases	Good	High	Good

Installing and Setting Up the Microsoft SQL Server Database (Optional)

To install the Microsoft SQL Server Database, you must create the database and then set up a valid login and user. Consult the appropriate Microsoft SQL Server documentation for details.

System Reporter Host Configuration

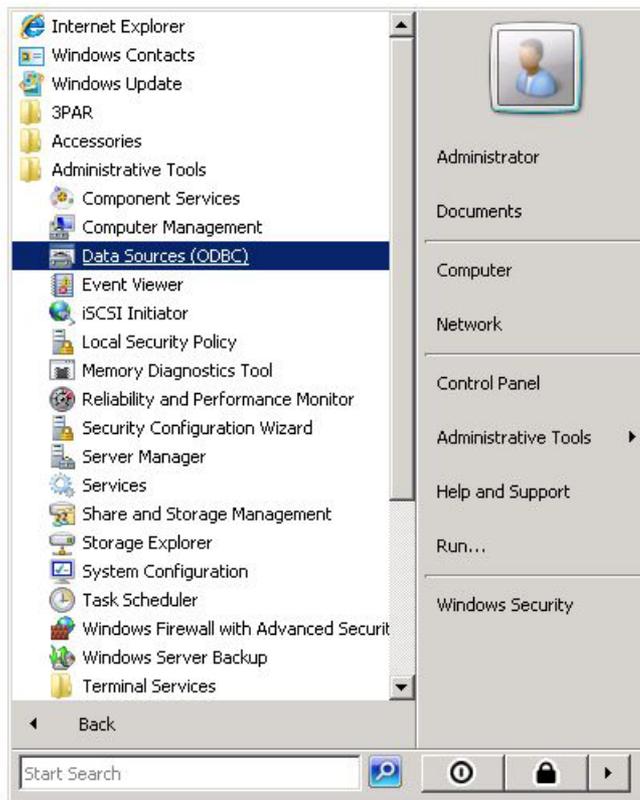
Once you have set up the Microsoft SQL database and created a login and user, you need to configure the System Reporter host to create the data source (ODBC) using the appropriate Administrative Tools [Figure 4 \(page 34\)](#).

NOTE: The screens shown in the following examples are specific for the 32 bit ODBC Administrator utility but are similar to those for the 64 bit Administrator utility.

For a 64 bit Windows Server please run the 32 bit ODBC Administration utility to create the DSN. Usually, this utility can be found in the following path:

C:\Windows\SysWOW64\odbcad32.exe

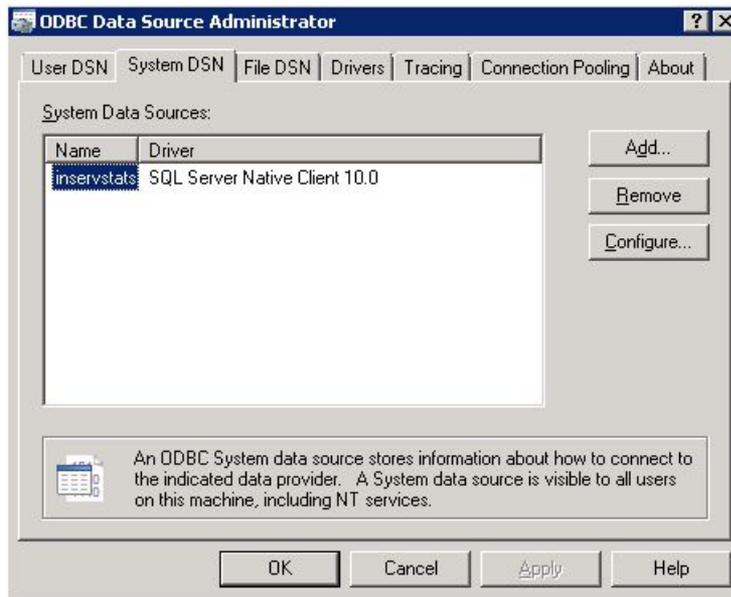
Figure 4 Selecting Data Sources (ODBC)



On the System Reporter host create a System ODBC DSN as follows:

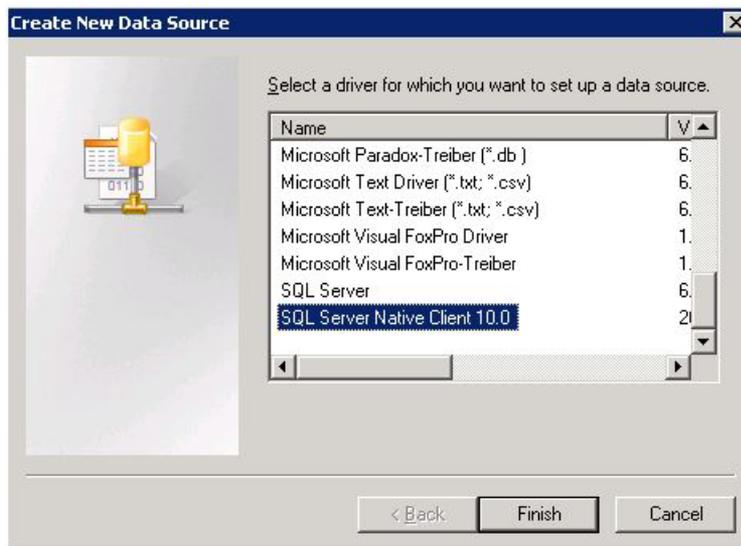
1. Select **Administrative Tools**→**Data Sources (ODBC)**.
The ODBC Data Source Administrator screen appears:

Figure 5 ODBC Data Source Administrator Screen



2. Select the **System DSN** tab.
 3. Click **Add**.
- The **Create New Data Source** screen appears:

Figure 6 Selecting the data source for the SQL driver

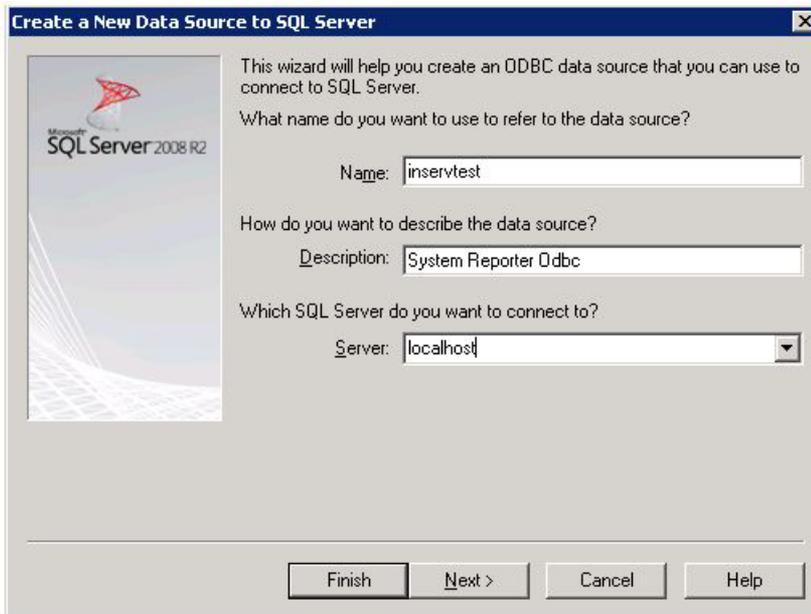


4. Choose **SQL Server Native Client 10.0**. Any SQL server ODBC driver will work. However, the SQL Server Native Client ODBC is supplied with the Microsoft SQL Server client and is considered to be the most efficient.

NOTE: Although you have the option of using the default SQL Server ODBC, if the SQL Server client is already installed on the machine, choose the latest ODBC.

5. Click **Finish**.
- The **Create a New Data Source to SQL Server** screen [Figure 7 \(page 36\)](#) appears:

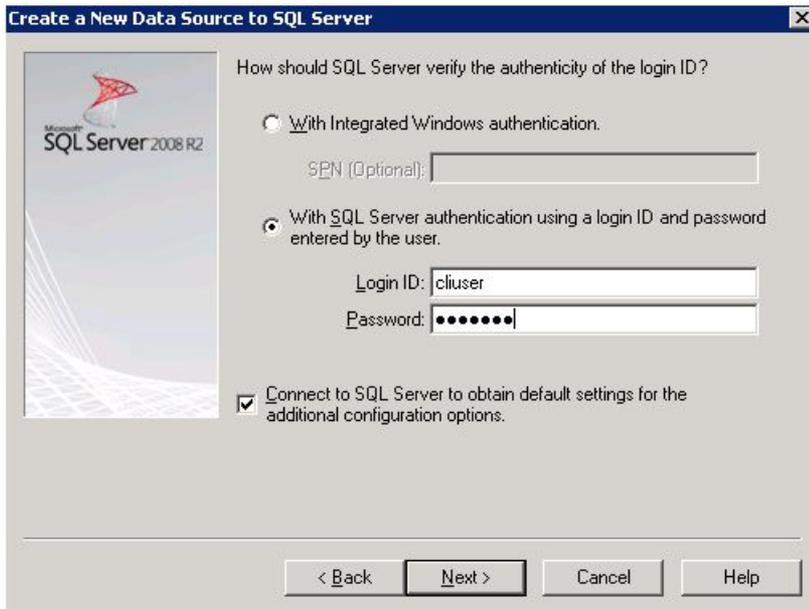
Figure 7 Create a New Data Source to SQL Server Screen



Note that the screens appear somewhat differently depending on whether you have chosen to use the default version of the SQL ODBC.

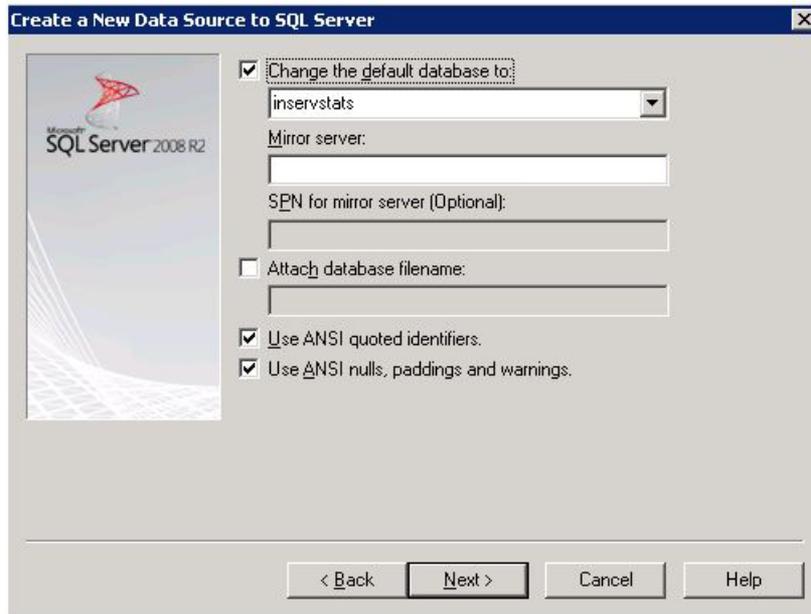
6. Click **Next**.

Figure 8 Entering the Login ID and Password



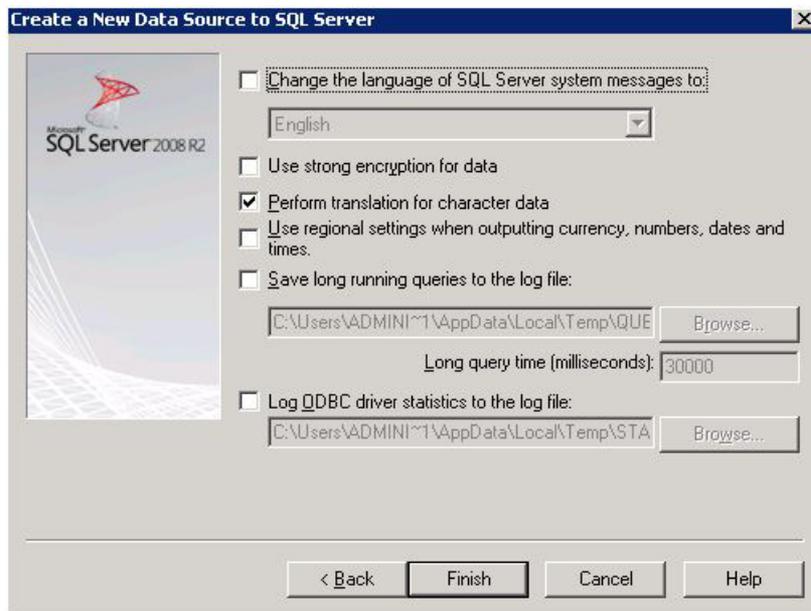
7. Enter the user's log in and password information, then click **Next**.

Figure 9 Specifying the MSSQL Database Name



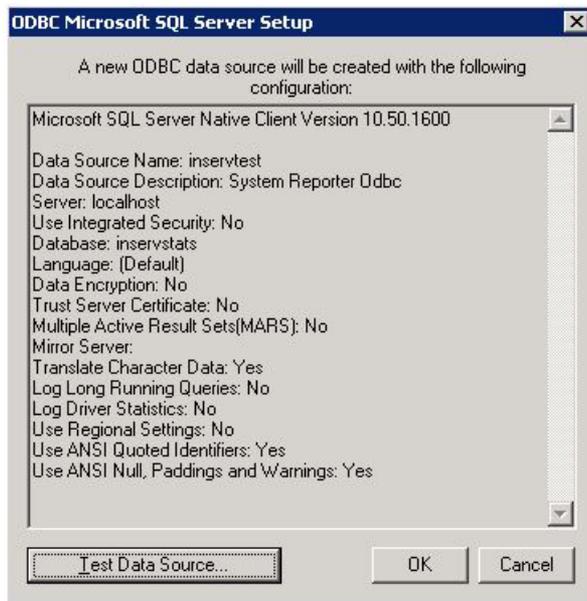
8. Check the **Change the default database** check box, and enter the name of the MSSQL database that has been created for System Reporter.
9. Click **Next**.

Figure 10 Checking the Details for the Database



10. Click **Finish**.

Figure 11 ODBC Microsoft SQL Server Setup Screen.



11. Check the details for the selected database, then click **Test Data Source**. The screen must show the following (Figure 12 (page 38)):

Figure 12 SQL Server ODBC Data Source Test Screen



If the screen does not show the appropriate information, please run the configuration again.

Installing and Setting Up the MySQL Database (Optional)

To use MySQL, you must install MySQL Database Server, version 5.1, following the instructions provided in this section. MySQL Database Server is optional third-party software and is not provided by HP.

Installing MySQL Server

To use MySQL for the HP 3PAR System Reporter database, obtain a software license and a copy of MySQL Database Server, version 5.1. See www.mysql.com for licensing and download instructions.

NOTE: Beginning with the 2.2 release, System Reporter can use a MySQL database on a remote server. The MySQL installation and configuration steps are the same for the remote server. See section for configuring System Reporter to use a remote MySQL database. This type of configuration is highly recommended.

After downloading the MySQL executable, use the **MySQL Server Setup Wizard** to set up and configure MySQL as follows:

1. Open the MySQL executable.
2. When prompted, click **Next**.
3. Choose a **Typical** installation and click **Next**.
4. Click **Install**.
5. Click **Skip Sign-Up** then click **Next**.
6. Make sure that **Configure MySQL Server now** is checked and then click **Finish**.

The **MySQL Server Instance Configuration Wizard** appears.

NOTE: When installing on a Redhat Linux server where the MySQL server is running locally on the same server, you must complete one of the following procedures:

Create a symbolic link using the following command:

```
ln -s /var/lib/mysql/mysql.sock /var/run/mysqld/mysqld.sock
```

Or, make the following change in the `/etc/my.cnf` file:

```
socket=/var/run/mysqld/mysqld.sock
```

and then restart the MySQL server.

Without one of these modifications, the sampleloop logs the following error:

```
2010-03-09 12:54:01: Could not connect to database: mysqlconnect/db
server: Can't connect to local MySQL server through
socket '/var/run/mysqld/mysqld.sock' (2)
```

Configuring the MySQL Database

Use the **MySQL Server Instance Configuration Wizard** to configure the MySQL database as follows:

1. In the **MySQL Server Instance Configuration Wizard**, click **Next** to begin configuring the database.
2. Select **Detailed Configuration** and then click **Next**.
3. Select **Server Machine** then click **Next**.
4. Select **Multifunctional Database** then click **Next**.
5. When prompted to select a drive for the warehousing data file, accept the default drive and directory and click **Next**.
6. Choose **Decision Support (DSS)/OLAP** and click **Next**.
7. Confirm that **Enable TCP/IP Networking** is selected and that the **Port Number** is **3306** and then click **Next**.
8. Confirm that **Standard Character Set** is selected and then click **Next**.
9. Confirm the following and then click **Next**:
 - **Install As Windows Service** is selected
 - the default **Service Name** is **MySQL**
 - the **Launch the MySQL Server automatically** checkbox is checked

10. Confirm that **Modify Security Settings** is selected, enter a root password, and then retype the password to confirm.

△ **CAUTION:** Be sure to record the password for later reference.

11. Make sure that **Create An Anonymous Account** is not selected and then click **Next**.
12. Click **Execute**.
13. When prompted, click **Finish**.

△ **CAUTION:** System Reporter only supports and creates MyISAM tables. Conversion of these tables to use any other type of storage engine is not supported.

Setting the `max_allowed_packet` parameter

Change the `max_allowed_packet` MySQL parameter by editing its value in the MySQL initialization file.

- On Linux, the default initialization file is `/etc/my.cnf`.
- On Windows, the default initialization file is `C:\Program Files\MySQL\MySQL Server 5.1\my.ini`.

The initialization file should include the following line:

```
max_allowed_packet=32M
```

After changing the `max_allowed_packet` parameter, you need to restart the MySQL server.

- On Windows you can do this by restarting the MySQL service.
- On Linux, you can run:

```
/etc/init.d/mysqld restart
```

△ **CAUTION:** It is necessary to use a large `max_allowed_packet` length because, when sampling a large number of systems or very large systems, the SQL INSERT command length can be quite large and the `max_allowed_packet` must be large enough to hold the entire statement. If the `max_allowed_packet` is not large enough, the sampler will get an error and sampling will stop.

Creating MySQL Users and Schema

If you are using MySQL, after installing and configuring MySQL you must create the MySQL users and schema as described in this section.

1. Log in to MySQL as root. You may need to do this on the machine on which MySQL is installed since root login is typically restricted to local logins only.

```
mysql -u root -p
```

2. Run the following command at the mysql prompt to create the database (named `inservstats` by default):

```
create database inservstats;
```

3. Run the following commands to create the two users **cliuser** and **webuser**.

```
create user cliuser identified by 'cliuserpassword';  
create user webuser identified by 'webuserpassword';
```

These users have different privilege levels that correspond with optional Apache HTTP Server privilege levels, as described in [“Installing or Upgrading the System Reporter Tools”](#) (page 43).

- **cliuser** is the default user name for the sampler and policy change user. This user can create and edit database tables and change database sampling policies (see [“Editing Sampling Policies”](#) (page 114)).
- **webuser** is the default user name for the Web reports user. This user has database read privileges only.

As part of installing the System Reporter tools (see [“Installing or Upgrading the System Reporter Tools”](#) (page 43)), the installer creates two Apache HTTP Server `config.tcl` files to store the user names and passwords for the **cliuser** and the **webuser**. See for instructions on editing these `config.tcl` files if you change the MySQL usernames and passwords at a later time.

NOTE: When creating the users, you must assign them passwords in order for System Reporter to function properly.

4. Grant the two users the privilege levels described above by running the commands below:

```
use insertstats
grant all on * to cliuser;
grant select on * to webuser;
exit;
```

5. Exit the `mysql` command line by running the following command:

```
exit;
```

Setting Up the Oracle Database (Optional)

System Reporter supports the use of MySQL® Database Server or Oracle 11g Database Server. To use Oracle, you must set up Oracle, following the instructions provided in this section. To use and install MySQL go to [“Installing and Setting Up the Microsoft SQL Server Database \(Optional\)”](#) (page 34).

Oracle is optional third-party software and is not provided by HP.

Installing the Oracle database server is beyond the scope of this guide. We assume that you already have an Oracle database server. In this section we describe how to create a database and users on that server for the System Reporter and how to set up a connection to the database from the System Reporter.

NOTE: To ensure that all environment variables are properly updated, install Oracle prior to installing the Apache Web Server. If Apache is already installed, be sure to restart it after the Oracle client is installed.

Creating an Oracle database

On the Oracle database server:

1. Start the Oracle **Database Configuration Assistant**.
2. Click **Next** at the welcome screen.
3. Select the **Create a Database** radio button and click **Next**.
4. Select the radio button for the **Data Warehouse** template and click **Next**.
5. Enter a name for the database and Oracle System Identifier (**SID**) and click **Next**.
6. Keep defaults for **Configure Enterprise Manager**, and click **Next**.
7. Select passwords for various Oracle users and click **Next**.
8. Select the desired storage mechanism and click **Next**.
9. At this point you can click **Finish** or continue to set further options.

10. After clicking **Finish**, check the settings in the confirmation screen and click **OK**.
11. Click **Exit** after the database has been created.

Creating the Oracle Users for System Reporter

Use the instructions that follow to create two users: **cliuser** and **webuser**. These users have different privilege levels that correspond with optional Apache HTTP Server privilege levels, as described in “About Apache HTTP Server Authorization and Access Control” (page 49).

- **cliuser** is the default user name for the sampler and policy change user. The System Reporter tables are created in the **cliuser** schema. This user can create and edit database tables and change database sampling policies (see “Editing Sampling Policies” (page 114)).
- **webuser** is the default user name for the Web reports user. This user has select privileges only for the tables in the **cliuser** schema.

As part of installing the System Reporter tools (see “Installing or Upgrading the System Reporter Tools” (page 43)), the installer creates two Apache HTTP Server `config.tcl` files to store the user names and passwords for the **cliuser** and the **webuser**. See for instructions on editing these `config.tcl` files if you change the Oracle usernames and passwords at a later time.

To create the users, open the **Oracle Enterprise Manager** for the database in a web browser. When you create the database in section “Creating an Oracle database” (page 41), a link to the **Oracle Enterprise Manager** URL for the database would typically have been added to the Windows programs submenu for Oracle.

1. At the login screen for the **Oracle Enterprise Manager** login as SYSTEM with the password for the SYSTEM user.
2. Click the link (or tab) for **Server**.
3. Under **Security**, click on the link for **Users**.
4. Click the **Create** button.
5. In the create user screen, enter the user name **cliuser** along with the password. Select a default and temporary tablespace, select **Unlocked** for the status.
6. Click on the link (or tab) for **Roles**.
7. The **cliuser** needs to have the **CONNECTION** and **RESOURCE** roles. The **CONNECTION** role should be there by default, otherwise add it along with the **RESOURCE** role by clicking the **Edit List** button and adding **RESOURCE** and/or **CONNECTION** and clicking **OK**. The **cliuser** should also have the **CREATE TABLE** role.
8. Click **OK** to create the **cliuser** user.

Repeat the steps above to create the **webuser** user as well. The **webuser** needs SELECT privileges for all the tables in the **cliuser** schema. System Reporter will automatically grant SELECT privileges to the **webuser** when it creates the tables.

NOTE: To improve performance, System Reporter runs the `dbms_stats` utility in Oracle. Please grant execute permission for this utility in Oracle.

NOTE: System Reporter 2.9 logs more information for debugging purposes. It uses two Oracle views `v$instance` and `v$version`. Please grant permissions to these views for the user that sets the policy settings.

Creating the Oracle Net Service

You need to create an Oracle Net Service on the host on which the System Reporter will be installed. This is used by the System Reporter sampler and the web CGI programs to connect to the Oracle database.

1. Install the Oracle Client (the Runtime version is sufficient) on the host. This should include the **Net Configuration Assistant** under the Oracle client program menu.

2. Start the **Net Configuration Assistant**.
3. Select the **Local Net Service Name Configuration** radio button and click **Next**.
4. Select the **Add** radio button and click **Next**.
5. Enter the service name for the database according to the directions and click **Next**.
6. Select **TCP** protocol and click **Next**.
7. Enter the host name for the database server along with the appropriate port number and click **Next**.
8. Select **Yes, perform a test** and click **Next**. The test may fail if the default user and password are not set. Click the **Change Login** button and enter `cliuser` and password and click **OK**. The test should succeed. Click **Next**.
9. Enter a net service name. This is the name you will need to enter when installing the System Reporter tools. Click **Next**.
10. Follow instructions to complete the net service configuration and exit the program.

Installing or Upgrading the System Reporter Tools

You have several options for installing the HP 3PAR System Reporter tools, depending on whether you choose to use the SQLite database (the default), or one of the other optional databases.

If you are using Microsoft SQL Server, MySQL or Oracle for your database, before installing the HP 3PAR System Reporter tools you must install and configure the database as described in:

- [“Installing and Setting Up the Microsoft SQL Server Database \(Optional\)”](#) (page 34)
- [“Installing and Setting Up the MySQL Database \(Optional\)”](#) (page 38)
- [“Setting Up the Oracle Database \(Optional\)”](#) (page 41)

NOTE: Before installing the System Reporter tools, make sure that you have installed Apache HTTP Server as described in 3.6.2 Optional Components. [Figure 3](#) (page 31) provides a visual representation of the installation process for a new installation that is described in detail in the following sections.

There are different procedures for installing or upgrading System Reporter tools on Windows and Linux. For Windows, refer to section [“Installing or Upgrading System Reporter Tools on Windows”](#) (page 43) and for Linux refer to section [“Installing System Reporter Tools on Linux”](#) (page 46).

Installing or Upgrading System Reporter Tools on Windows

Install the HP 3PAR System Reporter tools on Windows as follows:

1. Insert the HP 3PAR System Reporter Software CD into your CD-ROM drive and open the file.
`installer.exe`
The **System Reporter Installer** window opens.
2. If the installer detects a previous System Reporter installation, it will initialize the settings for the installation to the existing settings instead of the default settings and include an **Upgrade** button.

Click the **Upgrade** button to keep the existing settings and skip to step 10.

Click the **Next** button to step through each of the settings as described in the following steps.

NOTE: If the installer detects a previous version of System Reporter, it will pop up a window at this point. If the `compactdb` process is running, the popup will advise you to wait for the compaction to complete and perform the upgrade later. It will not provide an option to stop the previous installation since terminating a `compactdb` prematurely causes problems with the database for MySQL. If `compactdb` is not running, the popup window asks for permission to stop the previous installation before proceeding with the new installation.

3. When prompted, use the **Browse...** button to select and enter the full path to the InForm CLI executable `cli.exe` (if it is different from the default) then click **Next**.
4. When prompted, use the **Browse...** button to select and enter the full path to the location where you installed Apache HTTP Server (if it is different from the default) then click **Next**.
5. Use the **Browse...** button to select and enter the installation directory for the HP 3PAR System Reporting Tools sampler or accept the default directory (`C:/Program Files/3PAR/System Reporter`) and then click **Next**.
6. Make sure that the desired database is selected, then click **Next**.
7. This step differs based on the type of database you selected.
 - a. If you selected **SQLite** for the database, you are prompted for the installation directory for the HP 3PAR Reporting Tools SQLite database. Use the **Browse...** button to select and enter the directory or accept the default directory (`C:/Program Files/3PAR/System Reporter/inservstats`) and then click **Next**. Since SQLite databases are merely on local files no user name or password are required and you will skip step 8 and be taken directly to step 9.
 - b. If you selected **Microsoft SQL** for the database, you are prompted for the DSN name and user name password. Enter the DSN name for a remote Microsoft SQL server if different from the default database that is displayed (`inservstats`). Enter the user name password. Click **Next**.
 - c. If you selected **MySQL** for the database, you are prompted for the MySQL server host name and database name. Enter the host name or IP address for a remote MySQL server or leave it as **localhost** for local MySQL server. Enter the database name if different from the default database that is displayed (`inservstats`). Click **Next**.
 - d. If you selected **Oracle** for the database you are prompted for the host string. Enter the net service name you entered in step 9 and click **Next**.

NOTE: If you are upgrading from an earlier version of the System Reporter or reinstalling an existing 2.9 version for SQLite, and if you wish to retain the data collected from the previous installation, you must specify the same (existing) database, otherwise a new database will be created.

System Reporter 2.9 will automatically convert the earlier release database tables into the format that it needs. If you wish to retain the tables in an earlier format you must make a backup copy of the database before upgrading.

NOTE: When upgrading from an earlier version of System reporter (for MySQL only), the Sampling process may take a while to start sampling. This delay is necessary to address an issue related to case insensitiveness with default installs of MySQL.

8. When prompted, type the **webuser** password in the **Web reports password** field and the **cliuser** password in the **Sampler and policy change password** field and then click **Next**.

NOTE: See and for additional information about **cliuser** and **webuser**.

These users correspond with the Apache HTTP Server privilege levels (see [“About Apache HTTP Server Authorization and Access Control”](#) (page 49)) and are also defined in the Apache Server HTTP `config.tcl` files (see [“Modifying the Apache HTTP Server Configuration Files”](#) (page 125)).

The installer places these user names and passwords in the Apache HTTP Server `config.tcl` files and the **cliuser** user name and password in the database password file called `dbpwwfile` located in the installation directory for the HP 3PAR System Reporting Tools sampler (default file path `C:/Program Files/3PAR/System Reporter/dbpwwfile`). See for information on setting the appropriate permissions for the files.

9. Enter optional SMTP (Simple Mail Transfer Protocol) parameters to allow System Reporter to send email. If you do not wish to use the email alerts feature (see [“Configuring Rules for Email](#)

Alerts” (page 116)) you may leave all entries in this screen blank. If you want to be able to configure System Reporter to generate email alerts, you must enter the IP address or name of the SMTP server, the user name and password if the SMTP server requires authentication, and the SMTP originator name (the mailbox name that the email appears to be from). This information is stored in the `sampleloop_config.tcl` file in the System Reporter installation directory. Click **Next** to continue.

NOTE: In addition to configuring the SMTP parameters for alerts, it is also used for sending an email when the **If there is a problem with sampling, send email to this address** option is set through the Sampling Policies settings.

NOTE: Some firewalls or antivirus software may prevent the HP 3PAR System Reporter sampler service (an executable named `sampleloop.exe`) or the facility for sending email links for reports (an executable named `inserv_perf.exe`) from connecting to the SMTP server. Ensure that the security software allows these executables to connect to the SMTP server. To check if the connection is working, create an alert as described in that will always generate an alert on a high-resolution table. Then check that email was received. If email was not received, check for error information in the `sampleloop.log` file (see “[Troubleshooting the Sampling Components](#)” (page 215)).

10. Click **Next** and then, when prompted, click **Finish**.

The HP **3PAR System Reporter Installer** window closes.

11. On the Windows **Control Panel**, open **Administrative Tools** and then **Services**.
12. In the **Services** window, locate the entry for HP **3PAR System Reporter sampler** and verify that the **Status** is **Started**.

NOTE: If the **Status** is not **Started**, the installation of the HP 3PAR Reporting Tools may be unsuccessful. Check the `sampleloop.log`. If the installation was unsuccessful, remove all System Reporter components using the instructions provided in and then repeat the installation steps described in “[Installing or Upgrading the System Reporter Tools](#)” (page 43).

See Chapter 11, Troubleshooting for additional troubleshooting information.

13. Close the **Services** window.

At this point the System Reporter sampler service has been installed and started.

If this is an upgrade from a previous version of System Reporter, the sampler process will automatically convert existing tables in the database into newer version tables and create new tables necessary for version 2.8.

 **WARNING!** Do not interrupt the sampler process during a System Reporter upgrade. Stopping the sampling service while it is in the process of upgrading the database tables may leave it in a state where it cannot continue. If you encounter a problem and need to stop the sampler process, you should first consult the `sampleloop.log` to make sure that the database tables have been upgraded before attempting to stop or restart the sampler process.

NOTE: The upgrade process may take some time during which web access to the reports will not be available. Consult the `sampleloop.log` file in the installation directory to check on the progress of the table conversion.

After all the new tables are available, sampling will start automatically.

Once all the tables have been created, the main web page for system reporter should be accessible at `http://<host_name>/3par` where `<host_name>` is the name of the server where Apache and System Reporter are installed.

Proceed to section “[Adding Storage Systems](#)” (page 47).

Installing System Reporter Tools on Linux

System Reporter tools on Linux consists of two RPM packages available on the CD:

- `sampleloop` (`sampleloop-2.9-1.i386rpm`) consists of the files needed to install the `sampleloop` daemon.
- `sysrptwebsrv` (`sysrptwebsrv-2.9-1.i386.rpm`) consists of the files needed to install the System Reporter web server code.

Both these packages must be installed after logging in to a terminal session on the server as the root user. Make sure that `httpd` and `mysqld` are running before starting `sampleloop` (when using the MySQL database). Note that `mysqld` may not be running on the same server.

NOTE: The `sysrptwebsrv` requires the 32-bit version of the `gd` (version 2.0 or greater) package to be installed. If running the 64-bit version of Red Hat Enterprise Linux 5, the package requirement incorrectly appears to be met by the 64-bit version of the `gd` package. You must ensure that the 32-bit (i386) version of the `gd` package is installed.

NOTE: Although the packages are installed as root, the `sampleloop` daemon and the CGI programs executed from the Apache web server run as the `apache` user.

NOTE: If SELinux (Security Enhanced Linux) is implemented at your site, you may need to either disable it or temporarily grant permissive access to the system to install the System Reporter tools. Please consult your Linux documentation for details on how to grant permissive access.

1. If there is an existing System Reporter installation, stop the `sampleloop` daemon and remove the older packages as follows:

```
/etc/init.d/sampleloop stop
rpm --erase sampleloop
rpm --erase sysrptwebsrv
```

The `rpm --erase` command will remove the previous installation of the packages but save any edited configuration files with a `.rpmsave` suffix appended to their names.

2. Install the `sampleloop` package by mounting the System Reporter CD, changing to the directory on the CD that includes the rpm packages and running:

```
rpm --install sampleloop-2.9-1.i386.rpm
```

3. Edit the `/etc/sampleloop.conf` file following the directions in the file. If a previous version of that file was saved you may keep the previous settings and simply replace the `/etc/sampleloop.conf` file with the `/etc/sampleloop.conf.rpmsave` file.

4. For MySQL and Oracle databases, create/edit the `/etc/sampleloop_dbpwwfile` file to include a single line that has the database user name and password separated by a space.

5. Start the `sampleloop` daemon by running:

```
/etc/init.d/sampleloop start
```

Look at the `sampleloop.log` file in the `/var/log/sampleloop/` directory to see status and progress of the `sampleloop` daemon. If this is a new installation, the `sampleloop` daemon will create all the required tables and initialize the policy table. Since no systems have yet been added to the sampling list, none will be sampled.

If this is an upgrade of an existing installation, `sampleloop` will automatically convert tables to the newer versions required for System Reporter 2.9. Then it will resume sampling.

If errors are reported in the `sampleloop.log` file (for example due to incorrect database setup, or incorrect `sampleloop.conf` parameters) correct those error and restart the `sampleloop` daemon:

```
/etc/init.d/sampleloop restart
```

6. Install the `sysrptwebsrv` package by mounting the System Reporter CD, changing to the directory on the CD that includes the rpm packages and running:


```
rpm --install sysrptwebsrv-2.9-1.i386.rpm
```
7. Edit the `/var/www/cgi-bin/3par-rpts/config.tcl` and `/var/www/cgi-bin/3par-policy/config.tcl` files following the directions in those files. If previous versions of the file were saved, you may keep the previous settings and simply replace the files with the saved versions. If the configuration files are correct, you should be able to open the web page <http://hostname/3par>.

Once all the tables have been created, the main web page for system reporter should be accessible at `http://<host_name>/3par` where `<host_name>` is the name of the server where Apache and System Reporter are installed.

Adding Storage Systems

Follow these instructions for adding HP 3PAR Storage Systems to be sampled.

1. Open a Web browser and point it at `http://<host_name>/3par/` where `<host_name>` is the Web server where Apache HTTP Server is installed.

The HP **3PAR System Reporter** window appears.

2. Click on **Policy Settings** in the Extras Menu area.

The **System Reporter Policies** window appears.

NOTE: If you upgraded from a previous version of System Reporter (and you changed the location of the CLI password directory with that previous release) you must first remove, and then re-add your HP 3PAR Storage Systems.

3. Choose the **InServ Systems** tab.
4. Click **Add InServ**.
5. In the **Add InServ System to be Sampled** window, type the IP Name or address of a storage server you wish to add, your **CLI User** name, and **CLI Password**.

NOTE: If you do not plan on using the capabilities offered by the optionally licensed Adaptive Optimization feature, you can specify a Browse level access since no special privileges are required to view reports. However, if you are using Adaptive Optimization to perform system optimization, you should specify an Edit level user access. Specifying a Super level privilege is not required or recommended. For 3.1.1 systems, users can have the 3PAR_AO role instead of the edit role. This role has browse level privileges plus the ability to move regions between Common Provisioning Groups (CPGs).

6. Check the **Use SSL** box if the CLI connection should be over SSL.
7. Leave the **Skip LD Performance data** box checked if you do not need to collect performance data for Logical Disks (LDs). If you intend to configure Adaptive Optimization for the HP 3PAR Storage System, you should uncheck the box since Adaptive Optimization uses LD performance data.
8. Check the **Sample Adaptive Optimization Data** box if you want to collect performance statistics for the optional Adaptive Optimization component. When choosing this option, you must make sure that the **Skip LD Performance data** check box is unchecked, otherwise an error message is generated.
9. You have the option of specifying the amount of time to retain certain types of data or you can accept the default values:
 - Hours to keep High Res Data – Specifies the amount of time (in hours) to keep High Res data.
 - Days to keep Hourly Data – Specifies the number of days to retain hourly data.

- Days to keep Daily Data – Specifies the number of days to retain daily data.
- Days to keep Adaptive Optimization Data – Specifies the number of days to retain Adaptive Optimization data.

10. Click **Submit Query**.

NOTE: Since performance data for LDs can increase the database size substantially you should check the **Skip LD Performance data** box unless you intend to configure Adaptive Optimization for the system.

You can change the value of the **Use SSL** and **Skip LD Performance data** parameters or even disable sampling for a storage server entirely as described in section “[Re-Configuring a Storage Server Sampling](#)” (page 116).

11. Repeat step 5 to add additional storage servers and then click **Return to Policies**.

NOTE: The maximum CPG size supported when region data sampling is enabled for Adaptive Optimization is 250 TB. The maximum Adaptive Optimization configuration size is limited to 125 TB.

Setting File Permissions (Optional)

Set file permissions as desired for the password and configuration files created during installation. See *About File Permissions* for additional information.

Verifying Installation

Perform the following steps to verify that the System Reporter tools were installed correctly and that systems are being sampled successfully:

1. Wait several minutes after adding the last storage server as described in step 5??? (for SQLite configurations).
2. Look at the sampleloop.log file. On Windows this is in the log subdirectory in the directory where System Reporter is installed (C:\Program Files\3par\System Reporter\log\sampleloop.log). On Linux, open /var/log/sampleloop.log.
3. At the bottom of the log file, for each system to be sampled, look for an entry similar to the following:

```
2009-04-29 10:16:15: Adding system inserv1 into system table

2009-04-29 10:16:21: Sample inserted (1 systems) (time=2009-04-29 10:14:53) (hires,
space=0) (9 secs)
```

You can now return to the HP **3PAR System Reporter** window and execute database queries using a Web browser.

NOTE: For performance type reports, at first you may only be able to see high resolution performance information. For space type reports, it may take more than 30 minutes before you can sample high resolution data. Daily and hourly data samples may take even longer.

Installing the Excel Client (Optional)

If you plan to use the HP 3PAR System Reporter Excel client to access the HP 3PAR System Reporter database, you must install the Excel client on a local machine as described in this section.

NOTE: Use of the HP 3PAR System Reporter Excel client requires that you have Microsoft® Excel® 2003 or 2007 software already installed on your local machine. See for additional details.

The HP 3PAR System Reporter Excel client is contained in a single Microsoft® Excel® workbook with the file name `3PAR System Reporter.xls`. After downloading this file, you can save it in any convenient location.

To download the HP 3PAR System Reporter Excel client:

1. Open a browser window and point it at `http://<host_name>/3par/` where `<host_name>` is the server where Apache HTTP Server is installed (see [Before You Begin and Optional Components](#)).
2. Near the top of the page is the **Excel Client** link (see [Figure 14 \(page 53\)](#)).
3. Use the **Excel Client** link to download and save the file `3PAR System Reporter.xls` to any convenient location on your local machine.

Methods for saving this file locally will vary according to the Web browser you are using.

△ CAUTION: Depending on your Web browser, simply clicking the **Excel Client** link may cause the Excel client to open in your browser window, which will not give you access to the full Excel client functionality. Be sure to save the file on your local machine and then refer to [“Accessing the Optional Excel Client” \(page 55\)](#) for further instructions.

About Apache HTTP Server Authorization and Access Control

You can configure Apache HTTP Server to require a username and password before allowing a user to:

- read the HP 3PAR System Reporter database or view the database sampling policies.
 - make changes to database sampling policies.
-

NOTE: Viewing and editing the HP 3PAR System Reporter database sampling policies can only be done via Web browser, as described in [“Editing Sampling Policies” \(page 114\)](#).

One username and password pair can be used for both, or a different username/password pair can be used for each of these functions. Using two username/password pairs will allow you to configure System Reporter such that a user is able to read the database and view sampling policies, but is not allowed to modify System Reporter database policy settings.

Apache HTTP Server authorization and access control information is stored in the main configuration file, `httpd.conf`. You can add one or two `Directory` directives to the `httpd.conf` file to control user access to the System Reporter database as described above. Refer to the contents of the `httpd.conf` file as well as your Apache documentation for more information regarding the directives contained in this file and how to add `Directory` directives.

When adding `Directory` directives to the `httpd.conf` file, a username and password pair applied to the `cgi-bin/3par-rpts` directory will control reading of the database and the display of policy settings. A username and password pair applied to the `cgi-bin/3par-policy` directory will control modifications to the policy settings.

Within the `Directory` directive, the `AuthName` property corresponds to a field that appears in the connection dialog box that the user sees when trying to access restricted functionality (see and for examples). Depending on how you configure the Apache `httpd.conf` file, the user will see this dialog box when attempting to read the database (or database sampling policies) or when attempting to edit the sampling policies via Web browser. Note that the appearance of this dialog box may vary depending on which version of Windows you are using.

The following example provides two sample `Directory` directives that can be placed in the Apache `httpd.conf` file for the purpose of user access control. This example allows for the

distinction between users with database and sampling policy read access and those with database sampling policy edit privileges. Note that, when adding `Directory` directives to the `httpd.conf` file, you must also create a password file `{password_file}`, as described in your Apache documentation.

```
<Directory "cgi-bin/3par-rpts">
  AuthType Basic
  AuthName "3PAR System Reporter Database Query"
  AuthUserFile {password_file}
  Require user {username_read}
</Directory>
<Directory "cgi-bin/3par-policy">
  AuthType Basic
  AuthName "3PAR System Reporter Policy Update"
  AuthUserFile {password_file}
  Require user {username_edit}
</Directory>
```

In the above example:

- `{password_file}` is a password file that you must create (see the Apache documentation for instructions). Both directives can use the same password file.
- `{username_read}` is the username for a user that is allowed read access to the System Reporter database and database sampling policies
- `{username_edit}` is the username for a user that is allowed to edit database sampling policies.

Removing the System Reporter Components

This section describes how to remove the HP 3PAR System Reporter Components. This includes removing the System Reporter sampler and default database, as well as deleting the optional MySQL database (when applicable).

Removing the System Reporter Tools on Windows

To remove the System Reporter tools on Windows:

1. Go to the Windows **Control Panel**, open **Administrative Tools** and then **Services**.
2. In the **Services** window, right-click **3PAR System Reporter sampler** and then click **Stop**. Close the **Services** window.
3. From a command prompt, change to the directory where you installed the HP 3PAR System Reporter tools:

```
# cd \program files\3par\system reporter
```

4. Execute the following command:

```
# winserv uninstall "3PAR System Reporter sampler"
```

5. Exit and close the command prompt window.

```
# exit
```

6. Using **Windows Explorer**, go to the location where you installed the HP 3PAR System Reporter tools (for example, `C:\Program Files\3par\System Reporter`) and delete the entire directory.
7. Locate and remove the `inservstats.db` file if the file was in a location other than the System Reporter folder. Also remove the file `inservstats.db-journal` if present in the same folder.
8. (Optional) In the location where you installed Apache HTTP Server, locate the `htdocs` directory (for example, `C:\Program Files\Apache Group\Apache2\htdocs`) and delete the `3par` subdirectory (`C:\Program Files\Apache Group\Apache2\htdocs\3par`).
9. (Optional) Go to the location where you installed Apache HTTP Server and locate the `cgi-bin` directory (for example, `C:\Program Files\Apache Group\Apache2\cgi-bin`), then delete the subdirectories `3par-policy` and `3par-rpts`.

Removing the System Reporter Tools on Linux

1. Stop the `sampleloop` daemon by running:

```
/etc/init.d/sampleloop stop
```

2. Remove the `sampleloop` and `sysrptwebsrv` packages by running:

```
rpm --erase sampleloop
```

```
rpm --erase sysrptwebsrv
```

Although these commands remove the files associated with the packages, user-modified configuration files are saved with a `.rpm_save` suffix appended to their names. If you do not plan to re-install these packages later with similar configuration settings you may delete these configuration files.

3. If you used SQLite as the database, locate and remove the `inservstats.db` file, typically in the `/var/inservstats` directory. Also remove the file `inservstats.db-journal` if present in the same directory.

Removing the MySQL Database (Optional)

If you are using a MySQL database and you want to completely remove all HP 3PAR System Reporter components, you should also remove the `inservstats` database using the MySQL Administrator, as described in this section.

1. Log in to MySQL as `root`. You may need to do this on the machine on which MySQL is installed since `root` login is typically restricted to local logins only.

```
mysql -uroot -p
```

2. Run the following command at the `mysql` prompt to delete (drop) the database (named `inservstats` by default):

```
drop database inservstats;
```

⚠ CAUTION: Deleting the database schema will remove the `inservstats` database and cannot be undone.

3. Delete the users created by running the following commands:

```
drop user cliuser;
```

```
drop user webuser;
```

4. Exit the MySQL command line.

```
exit;
```

4 Getting Started

This chapter describes how to start using HP 3PAR System Reporter from the primary web interface or Excel client.

NOTE: For information on using a Web browser to edit database sampling policies, add or remove storage servers from the database, schedule reports or optimize storage utilization with the optional Adaptive Optimization feature, see [“Configuring the System Reporter Policy Settings” \(page 112\)](#).

Accessing the Main System Reporter Interface

You can access the HP 3PAR System Reporter main features using any standard Web browser. A Web browser is not provided on the *3PAR System Reporter* CD.

To access HP 3PAR System Reporter using a Web browser, point your browser at `http://<host_name>/3par/` where `<host_name>` is the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed (see [“Before You Begin” \(page 25\)](#)).

NOTE: If your Apache Server is configured for secure connections, be sure to precede the address with `https://` (instead of `http://`) when accessing System Reporter. For complete details on configuring the Apache Server, please consult [“About Apache HTTP Server Authorization and Access Control” \(page 49\)](#).

If the Apache HTTP Server has been configured to require a password to read the database, a connection dialog box appears ([Figure 13 \(page 52\)](#)). Type the user name and password and click **OK**.

Figure 13 Connection Dialog Box

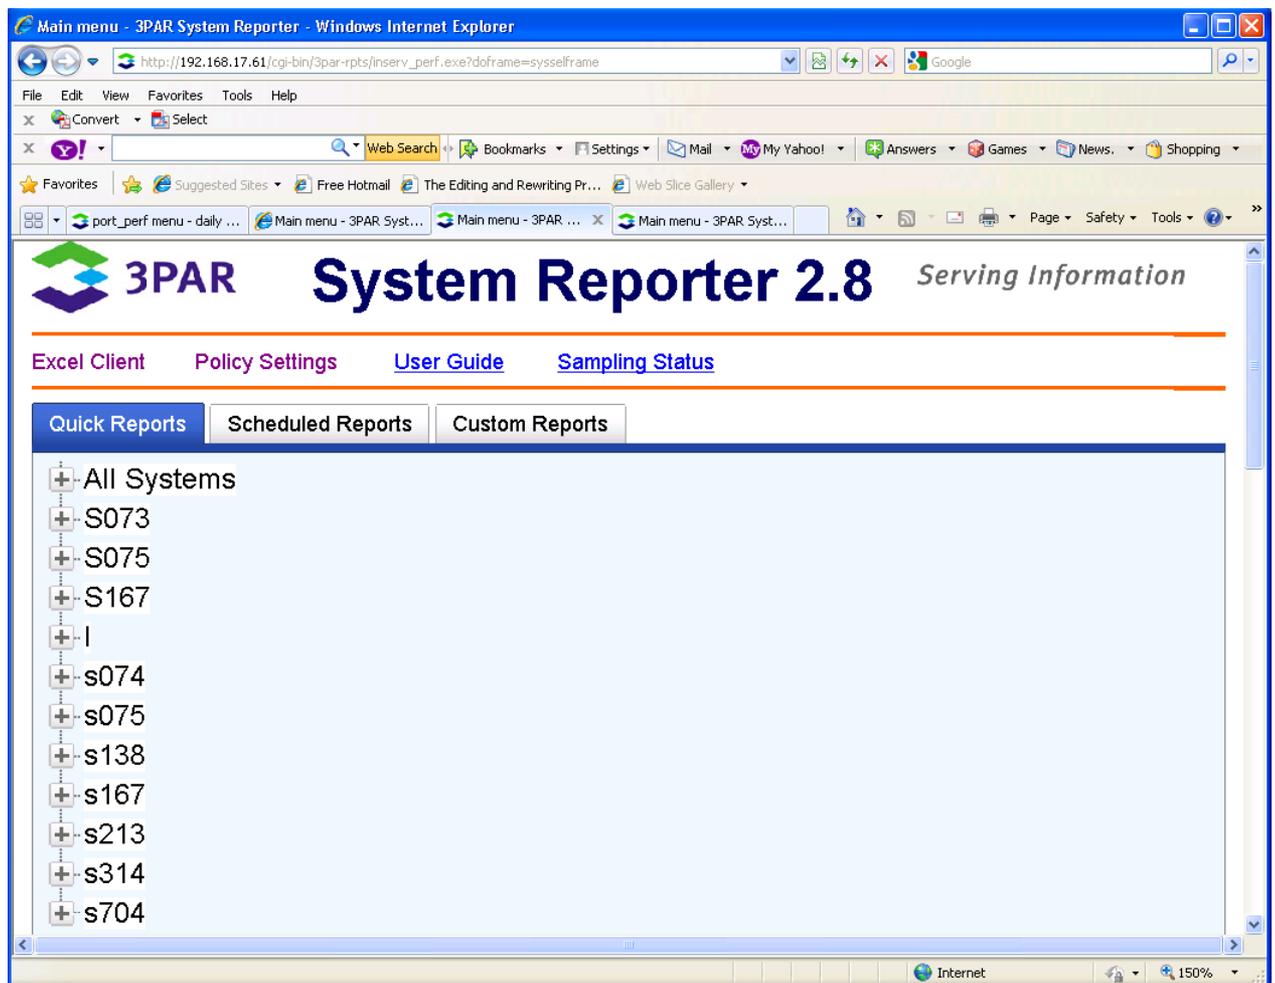


NOTE: See for more instructions on setting Apache HTTP Server to require a username and password.

System Reporter Main Menu Window

The various System Reporter menu options are accessed through the Main Menu Window.

Figure 14 The System Reporter Main Window



The **HP 3PAR System Reporter** Main Menu (Figure 14 (page 53)) is divided into the following areas:

- **Title Pane** which shows the version of System Reporter.
- **Main Menu** tabs which provide options for generating quick reports, accessing scheduled reports and creating customized reports. (See for details).
- **Extras Menu** area has links to:
 - **Excel Client** used to create and access custom reports.
 - **Policy Settings** (see “Configuring the System Reporter Policy Settings” (page 112) for complete details on configuring policy settings). The following tabs are provided from the policy settings window:
 - **Add Systems** – used to specify the HP 3PAR Storage Systems for reporting purposes.
 - **Alert Rules** – used to specify the alert rules.
 - **Schedule Report** – used to schedule reports at specified times.
 - **Adaptive Optimization** – used to configure optimization for storage utilization.

- **User Guide** (this document in PDF format).
- **Sampling Status** (the time since the most recent high-resolution sample for each system being sampled).

NOTE: Instead of opening a new window, browsers can be configured to open a new tab instead.

For Firefox, open the options window (**Tools**→**Options**), click on the **Tabs** tab, and select the option that opens new pages in a new tab.

For Internet Explorer, open the Internet Options window (**Tools**→**Internet Options**), click on the **General** tab and the **Settings** button for **Tabs**. This opens the **Tabbed Browsing Settings** window. Ensure that the **Enable Tabbed Browsing** checkbox is checked, and select the **Always open popups in a new tab option**.

Getting Help

While accessing HP 3PAR System Reporter using a Web browser, you can open the *HP 3PAR System Reporter Software User's Guide* and access online help from your browser as follows:

1. From the **Extras Menu**, click **User Guide**. The *HP 3PAR System Reporter Software User's Guide* opens in a separate browser window.
2. Click the underlined labels for various menu items to access the appropriate section of the User Guide that provides help.

NOTE: Adobe Acrobat® is required to view the *HP 3PAR System Reporter Software User's Guide*. See www.adobe.com/support/downloads for download instructions.

Using the Object Selection Controls

Menus in System Reporter often provide a way to select a subset of objects to include in a report from among a list of available objects. System Reporter provides two types of selection lists to implement this, a simple selection list and a filtered multi-selection control as described in the following sections.

Simple Selection List

When the number of available objects is fixed or expected to be small, System Reporter uses a standard selection list in which multiple items can be selected (using **CTRL+CLICK**). An example is the selection list to select from the available disk speeds (see “[Select Disk Speed](#)” (page 86)). Since the number of distinct disk speeds available is small, a simple selection list is adequate.

Filtered Multi-Selection Control

When the number of available objects is large (for example there may be thousands of VVs in a VV selection list), it is difficult to scroll through a simple selection list to find a desired object or to select specific objects in this list using **CTRL+CLICK**. To make it easier to find specific objects and select them, System Reporter provides a filtered multi-selection control. Two examples of this, for selecting systems and for selecting domains, are shown in [Figure 26](#) (page 70).

A filtered multi-selection control consists of two lists: one showing the available objects and one showing the currently selected objects. Objects can be moved between lists by either double-clicking on the object or by selecting objects and clicking on the appropriate arrow button in the control. If no specific object is selected, then all objects are considered selected.

The list of available objects can be filtered by typing an expression pattern in the text input area above the list of available objects and clicking on the **Filter** button.

CAUTION: Filtering large numbers of objects may exceed the limit for your server, in which case you should reduce the number of objects being filtered and try again.

NOTE: When using an expression pattern, you should be aware that the Web interface uses a Regular Expression syntax for filtering whereas the Excel client uses a LIKE operator syntax. For example, to search for an object with "DB" anywhere in the name, the filter text would use **DB** for the web interface and ***DB*** for the Excel client. As another example, to search for an object starting with "**DB**" and ending with "12", the filter text would be **^DB.*12** for the Web interface and **DB*12** for the Excel client. For more complex search patterns, please refer to the Regular Expression syntax for the Web and the LIKE Operator help in Excel.

Accessing the Optional Excel Client

The HP 3PAR System Reporter Excel client is supported on Microsoft® Windows® systems for creating and accessing custom reports. Microsoft® Excel® 2003 or 2007 is required but is not provided as part of HP 3PAR System Reporter.

NOTE: In order to use the HP 3PAR System Reporter Excel client, you must first install and configure the Excel client as described in "[Installing the Excel Client \(Optional\)](#)" (page 48).

Starting the System Reporter Excel Client

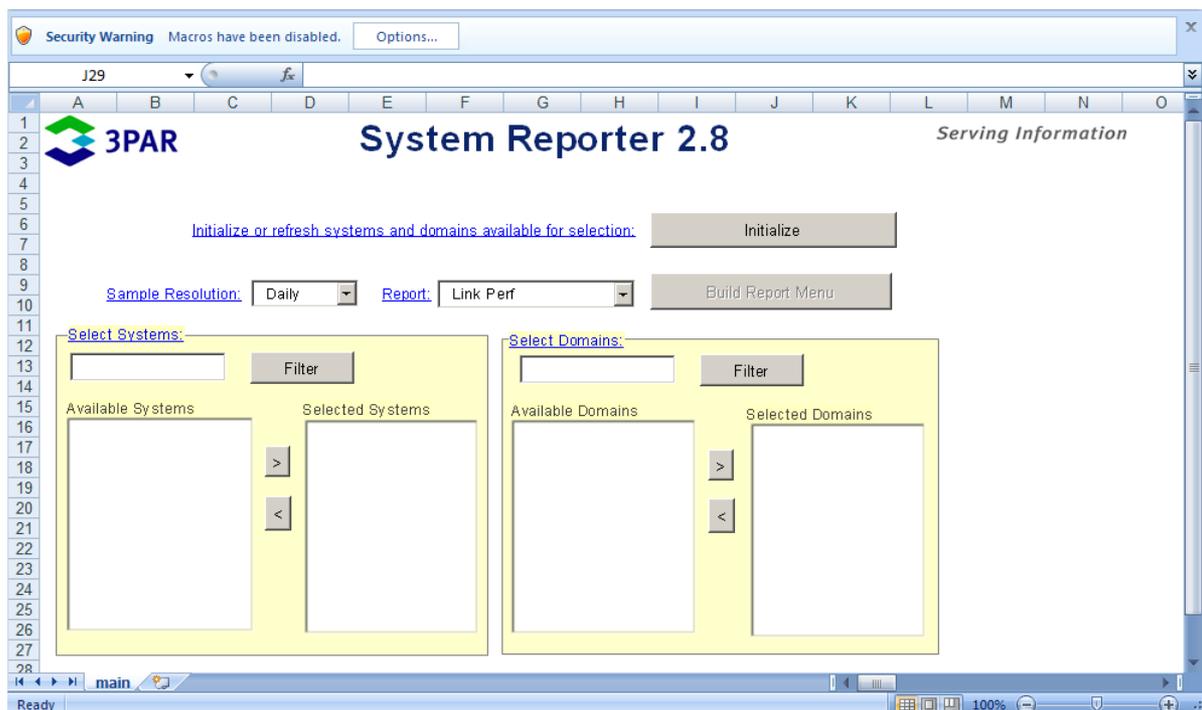
To start the HP 3PAR System Reporter Excel client, use your file manager to locate the Excel worksheet that was saved locally during the installation and access your local version.

NOTE: If you have not yet saved a version of the Excel worksheet locally, see for instructions on saving the `3PAR System Reporter.xls`

file on your local machine.

The HP 3PAR System Reporter Excel client opens in Microsoft® Excel® ("[System Reporter Excel Client](#)" (page 56)).

Figure 15 System Reporter Excel Client



Setting the Security Level for Excel

The HP 3PAR System Reporter Excel client uses macro code to implement its functionality. In order for this code to be executed, you must set the proper security settings depending on the version of Excel that you are using.

Setting the Security for Excel 2003

In most cases, when accessing the System Reporter client for Excel 2003, a Security Warning screen appears indicating that macros have been disabled (Figure 16 (page 56)).

Figure 16 Security Warning for Enabling Macros in Excel 2003



1. Click **Enable Macros** to continue launching the Excel client.
If you do not see the Security Warning, try completing steps 2 thru 7 to set the macro security level for Microsoft® Excel® 2003, before continuing.
2. On the Microsoft® Excel® **Tools** menu, point to **Macro**, then **Security**.
The **Security** dialog box appears with the **Security Level** tab enabled (Figure 17 (page 57)).

Figure 17 Setting the Security Level



3. On the **Security Level** tab, set the security level to **Medium** or **Low** (Medium is the default and suggested setting) and then click **OK**.

NOTE: The HP 3PAR System Reporter Excel client will not work if the Microsoft® Excel® security level is set to **High** or **Very High**.

4. If you changed the security level setting, close the Excel workbook and then reopen it. With a security setting of **Medium**, Excel will prompt you when the workbook is opened to specify whether or not macros should be enabled.
5. You need to select **Enable Macros** for the workbook to function.

NOTE: Please consult your company's security policy before changing any security settings.

Setting the Security Level for Excel 2007

In most cases, when accessing the System Reporter client for Excel 2007, a security warning, indicating that macros have been disabled, is shown in the upper left part of the screen. If you do not see the Security Warning, try completing steps 4 thru 9 before continuing.

To enable macros for the current session, follow these steps:

1. Click the **Options** button next to the security alert.
The **Security Alert - Macro** screen appears (Figure 18 (page 58)).

Figure 18 Macro Security Alert for Excel 2007



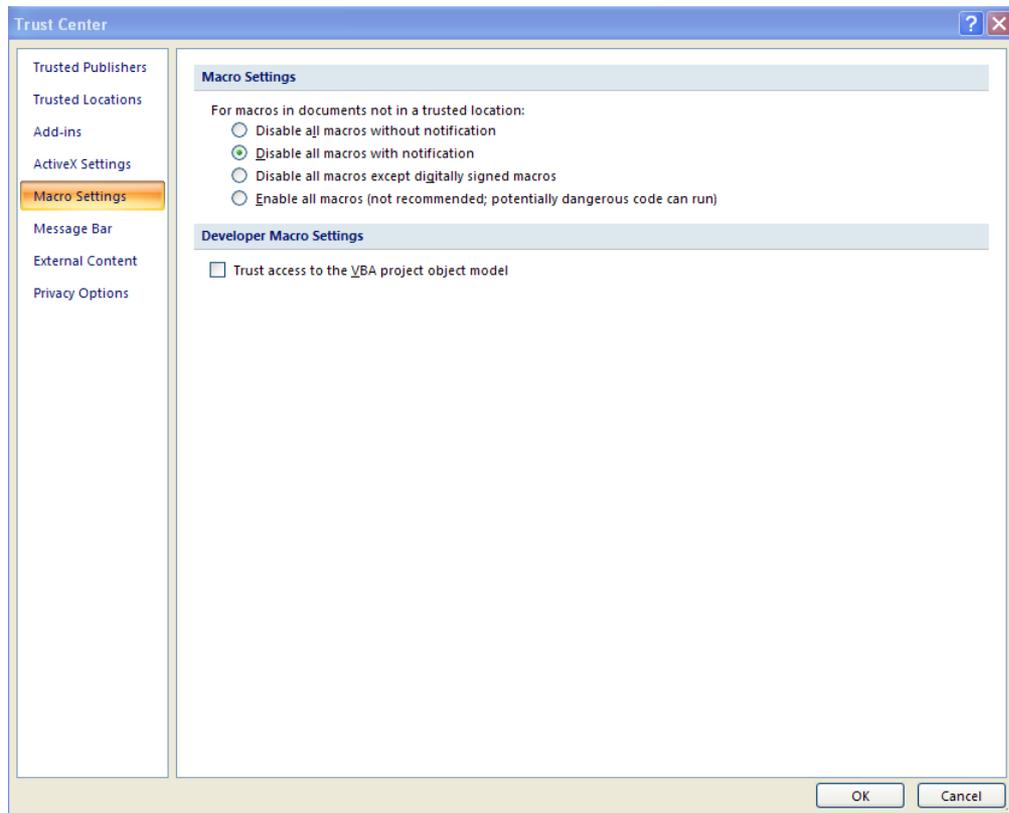
2. Choose the **Enable this content** radio button.
3. Click **OK**.

If you do not see a Security Warning when accessing the Excel client, it is likely that the main Excel Macro Settings are not set up properly and need to be changed. In this case, follow these steps:

4. Click the Windows icon in the left corner of the screen.
5. Click **Excel Options** at the bottom of the menu.
6. Click **Trust Center**.
7. Click **Trust Center Settings**.

The **Trust Center** screen appears (“Trust Center Screen” (page 59))

Figure 19 Trust Center Screen



8. If it is not already set, from the list of Macro Settings, choose **Disable all macros with notification**.

NOTE: The HP 3PAR System Reporter Excel client will not work if the Macro Setting is set to **Disable all macros without notification** or **Disable all macros except digitally signed macros**.

9. Click **OK**, then click **OK** again to return to the System Reporter Excel client main screen.
10. Close the Excel client, then launch it again from the System Reporter main menu.
11. Continue with step 1 thru 3.

NOTE: When launching the Excel 2007 client after making and saving changes, you may see a connection warning message. See *Saving and Exiting the Excel Client* for more details.

NOTE: Please consult your company's security policy before changing any security settings.

Connecting to the Web Server from the Excel Client

After opening the HP 3PAR System Reporter Excel client workbook with the correct security level setting, you must initialize or refresh the systems and domains.

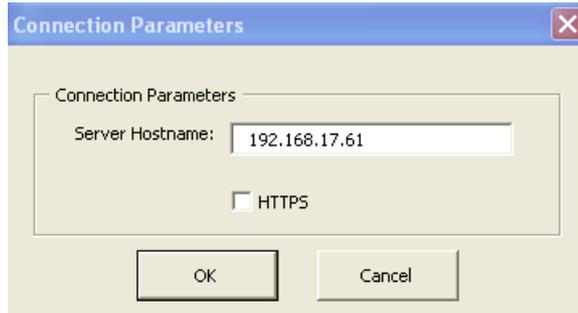
NOTE: You must connect to a Web server each time you open the Excel client workbook. The workbook does not store the server information.

To connect to a Web server from the Excel client workbook:

1. After opening the HP 3PAR System Reporter Excel client workbook with the correct security level setting, click **Initialize**.

The **Connection Parameters** dialog box appears (Figure 20 (page 60)).

Figure 20 Connection Parameters Dialog Box



2. In the **Server Hostname** box, type a host name and click **OK**.

The host name is the name of the Web server where the Apache HTTP Server and HP 3PAR System Reporter Web server scripts are installed (see “Before You Begin” (page 25)). If the server is configured for secure connections, check the HTTPS box.

NOTE: You can point the Excel client to a different Web server at any time. See *Changing the Web Server* for instructions on changing the Web server.

If the Web server has been configured to require a password to read the database, a connection dialog box appears (Figure 21 (page 60)).

Figure 21 Connection Dialog Box



3. Type your user name and password and click **OK**.

NOTE: See for more instructions on setting Apache HTTP Server to require a username and password.

The versions of the System Reporter’s Web server component and the Excel client are compared. If the data format provided by the server is incompatible with the Excel client, you must choose a compatible server before the Excel client can be used.

NOTE: If your server version is newer, but still compatible, a dialog appears stating that some features of the data provided by the server cannot be displayed by the Excel client.

Changing the Excel Client Server and Resetting the Workbook

The HP 3PAR System Reporter Excel client workbook includes a custom **3PAR** menu that is accessible from the Microsoft® Excel® menu tab called **Add-Ins**. This menu tab enables you to change the Web server, reset the workbook, and delete menu sheets.

Changing the Web Server

To change the Web server used for querying the System Reporter database:

1. Click the **Add-Ins** menu tab, choose the 3PAR pull down menu, and then select **Set server name...**
The **Connection Parameters** dialog box appears.
2. In the **Server Hostname** box, type the new server name.
3. Click **OK**.

Resetting the Workbook

To reset the HP 3PAR System Reporter Excel client workbook click the **Add-Ins** menu tab, choose the **3PAR** pull down menu and then select the **Reset all** option.

All reporting tool sheet controls are reset.

Accessing the Excel Client Query Log

The HP 3PAR System Reporter Excel client workbook uses various Web queries to initialize user interface controls and draw charts. The most recent queries are recorded on a worksheet called the **Query Log**. By default, the **Query Log** is hidden. However, you can make this sheet visible as follows:

To make the **Query Log** visible:

1. Click on the tab associated with the report for which you want to examine the query log.
2. Right click on the tab.
3. Choose the **unhide** option. The unhide dialog box appears.
4. Choose the query log.

A new **Query Log** tab appears at the bottom of the Microsoft® Excel® window.

In the query log, there is a row for each query with the time and URL. Click on a URL to bring up a Web browser and display the results of the query in table form.

When the number of entries in the query log reaches 150, the older 50 are automatically deleted, leaving the most recent 100 entries.

You can hide, unhide, or delete the **Query Log** sheet as desired. If deleted, a new hidden sheet is created for the next query. See to learn how to delete the query log and other sheets.

Deleting Excel Sheets

When using the HP 3PAR System Reporter Excel client, you can reset the entire workbook to its original state by following the instructions in [“Resetting the Workbook” \(page 61\)](#), or you can delete unwanted sheets individually as described in this section.

Deleting Menu Sheets

When creating reports with the Excel client, the menu sheets can begin to clutter up the work space after a while and you may want to delete menus that are no longer needed.

To delete menu sheets, click the **Add-Ins** menu tab, choose the **3PAR** pull down menu and then select the **Delete menu sheets** option.

Deleting Data Sheets

Data sheets remain until you manually delete them. Manually deleting a data sheet causes any embedded charts associated with that data sheet to disappear as well.

To delete a data sheet:

1. At the bottom of the Microsoft® Excel® window, click a data sheet tab to enable the sheet to be deleted.
2. Point to the name of the data sheet tab and then right-click.
3. On the shortcut menu that appears, click **Delete**.
4. In the confirmation dialog that appears, click **OK**.

The data sheet is deleted. Any embedded charts associated with that data sheet disappear as well.

Deleting the Query Log

You can also delete the query log, which is by default a hidden sheet, by making it visible and then using the method described above to delete it.

To delete the query log:

1. Make the Query Log tab visible by following the instructions in [“Accessing the Excel Client Query Log” \(page 61\)](#).
2. At the bottom of the Microsoft® Excel® window, point to the **Query Log** tab and then right-click.
3. On the shortcut menu that appears, click **Delete**.
4. In the confirmation dialog that appears, click **OK**.

The query log is deleted. The next time that you initiate a new query using the Excel workbook, a new hidden query log sheet is created.

Saving and Exiting the Excel Client

To exit the HP 3PAR System Reporter Excel client workbook, on the **File** menu, click **Close**.

Upon exiting, Microsoft Excel will prompt you to save changes. If you have added customized charts and/or data sheets, as described in [“Working with Excel Charts” \(page 97\)](#), you may want to save your changes when prompted. However, it is not necessary to save changes and in most cases it is not desirable to do so. Saving the Excel workbook does not save the Web server information and saved workbooks do not automatically update embedded charts or data sheets when they are reopened.

When opening a previously saved Excel worksheet, you may receive a security alert indicating that macros and data connection have been disabled ([Figure 22 \(page 63\)](#)).

Figure 22 Security Alerts Multiple Issues



To enable macros and the data connection, choose the **Enable this content** buttons.

NOTE: If it is ever necessary, you can obtain an entirely new copy of the HP 3PAR System Reporter Excel client workbook by following the instructions provided in “Installing the Excel Client (Optional)” (page 48).

5 Accessing and Creating Reports

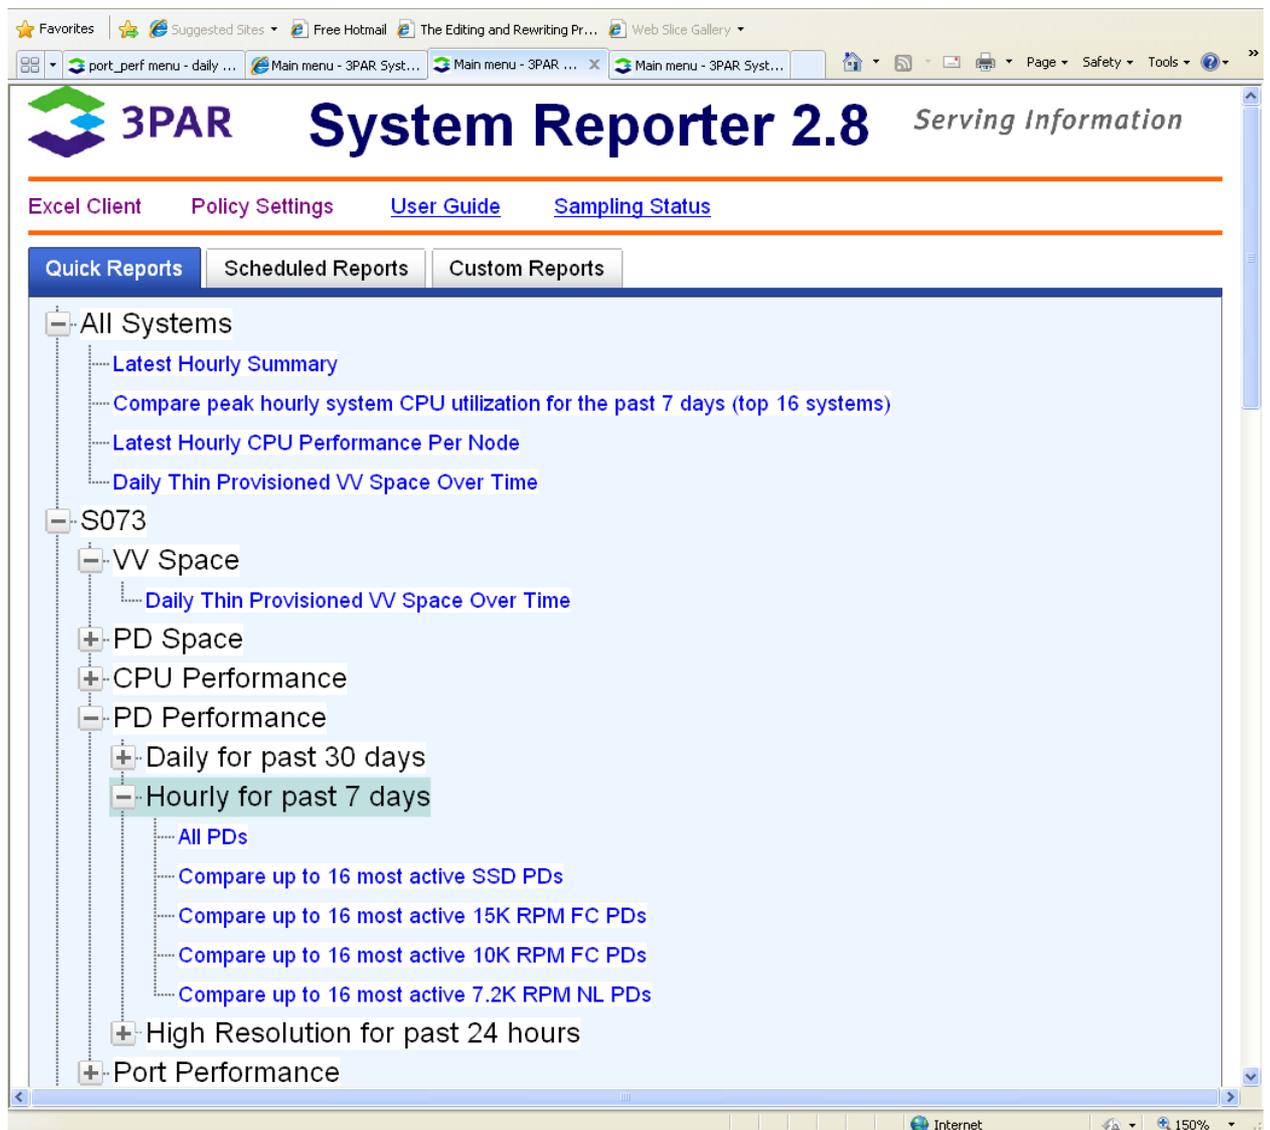
This chapter describes how to access quick reports, scheduled reports, and build custom reports that are specifically tailored for your operating environment. The menu selections and controls available for creating custom reports from either the web interface or the Excel client are functionally identical, so the information in this chapter applies to both with explanations provided for any minor variations.

NOTE: For information on using a Web browser to edit database sampling policies, adding or removing storage servers from the database, scheduling reports or optimizing storage utilization with the optional Adaptive Optimization feature, see [“Configuring the System Reporter Policy Settings” \(page 112\)](#).

Accessing Quick Reports

Quick reports provide predefined reports that you can generate quickly by choosing the **Quick Report** tab from the System Reporter Main Menu Window and then selecting one of the customized reports from the list of directories as shown in [Figure 23 \(page 65\)](#). Reports are provided for all the systems that are currently configured for System Reporter (available by expanding the **All Systems** branch) and from individual systems (available by expanding a particular system’s branch) and encompass sampling over a given period of time.

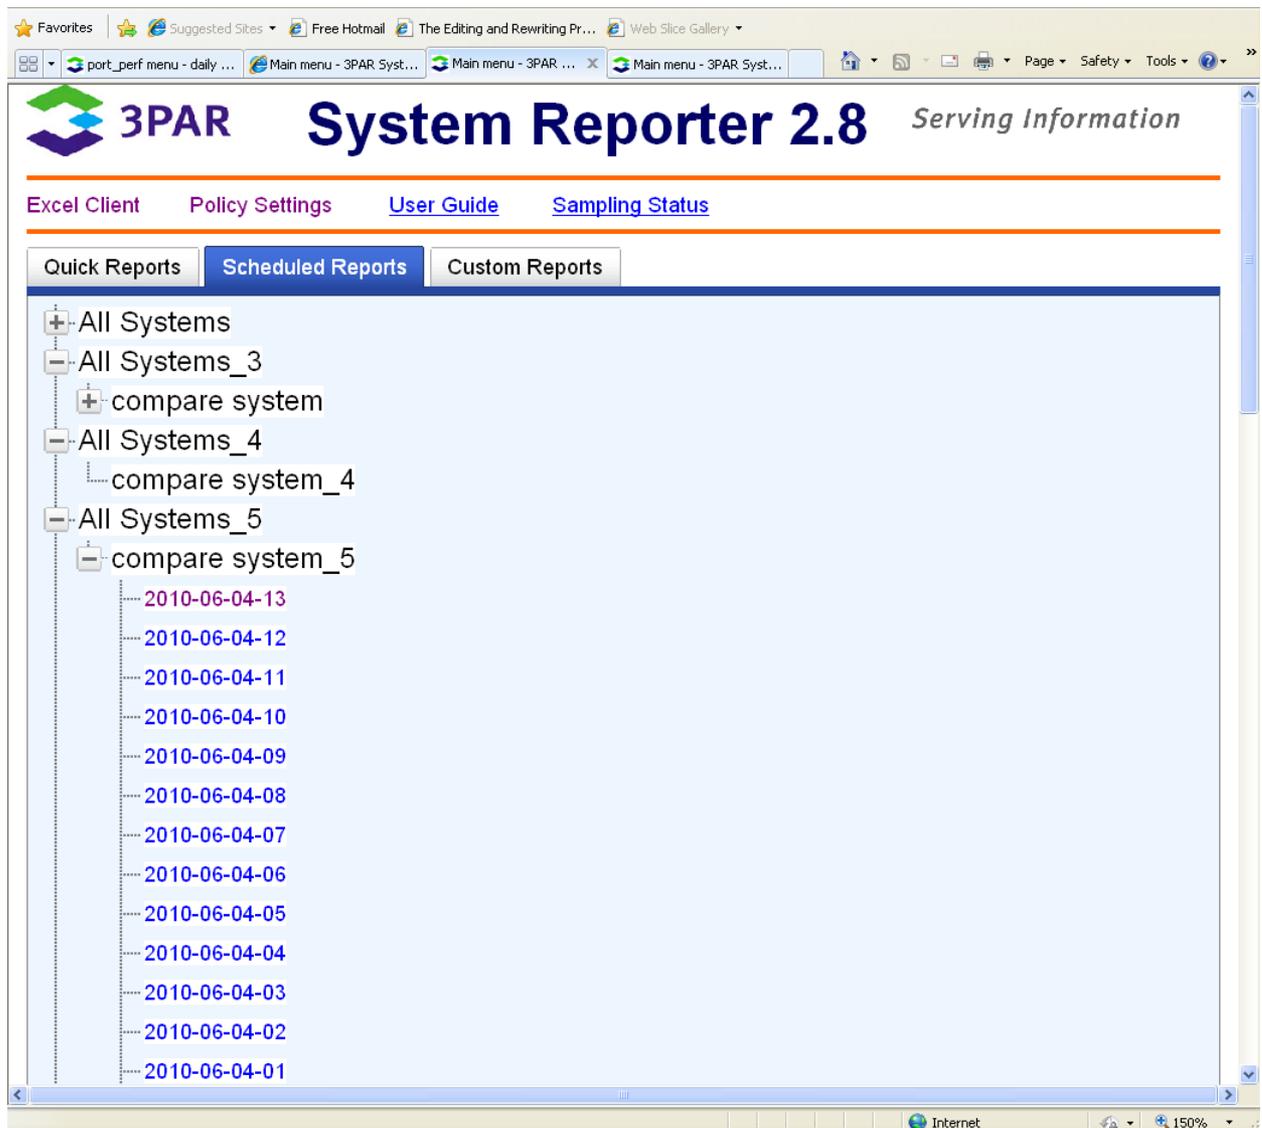
Figure 23 Quick Reports Window



Accessing Scheduled Reports

Scheduled reports are generated using pre-selected parameters at scheduled times and are stored in sub-directories whose structure is reflected in the tree shown in [Figure 24 \(page 66\)](#). Scheduled reports are stored with a name that is assigned when the report is scheduled. Each time the report runs, a new subdirectory is created for that instance of the report using the timestamp as the name. You can either view the reports via the **Scheduled Reports** tab or you can copy the report subdirectories to another area where they can be accessed by authorized users according to whatever permissions are deemed appropriate.

Figure 24 Scheduled Reports Window



NOTE: See "Scheduling Reports" (page 122) for complete details on scheduling a report.

Creating Custom Reports

You can create customized reports by selecting the **Custom Reports** tab from the Main Menu or by choosing the **Excel Client** from the Extras Menu.

NOTE: The web interface and Excel client for creating custom reports are identical except that the Excel client charting functions are accessed separately from the main interface and provide some additional flexibility.

In general, there are two main steps involved in creating a report:

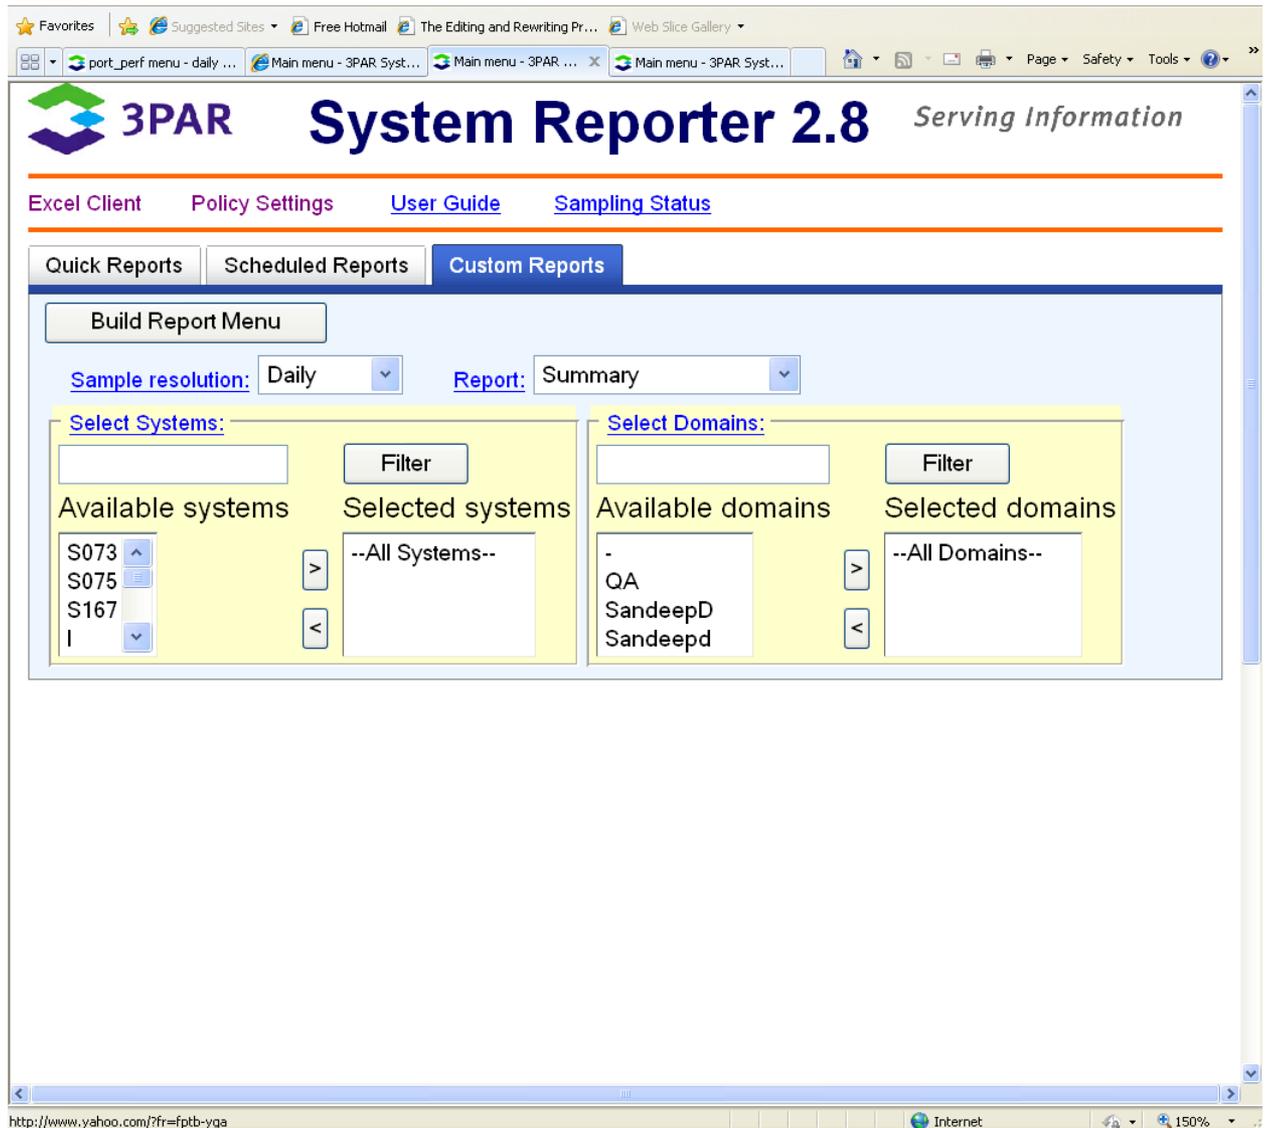
1. Choose a report, select the sample resolution, select the system(s) and/or domain(s) and click on the **Build Report Menu** button. This will build the report menu for that report in a new window (or tab).
2. Set the appropriate controls in the report menu and click on the **Generate Report** button. The report will be generated in a new window (or tab).

A detailed description of the various menu choices are described in the following sections.

Building the Report Menu

The first step in creating a report is to choose the appropriate options for the type of report that you want to generate.

Figure 25 Custom Reports Window



NOTE: The layout and placement of controls may vary slightly for the Excel client but operate in the same manner as described for the web interface.

Sample Resolution

The **Sample resolution** control enables you to select from one of three resolutions corresponding to the tables that are maintained in the database (see [Table 6 \(page 68\)](#)). The samples can be retained for a different lengths of time for each resolution (see for instructions on altering the default sampling policies).

Since a different set of objects may be present in data tables for different resolutions, the generated report menu is populated with the objects present in the tables for the selected resolution.

Table 6 Sample Resolution

Option	Purpose
Daily	Choose the Daily data table, which covers the largest time range, but has the coarsest resolution with only one sample per day.
Hourly	Choose the Hourly data table, which covers a shorter time range than the daily data table, but has one sample per hour.
High-Res	Choose the High-Resolution data table, which spans an even smaller time range, but has the finest granularity.

Report

Use the Report menu to choose from the following reports (Table 7 (page 68)):

Table 7 Reports and Uses

Reports	Purpose
Summary	Display summary information. See “Working with Excel Charts” (page 97) for the types of statistics provided by this report.
PD Space	Display physical disk space usage. See “PD Space Reports” (page 106) for the types of statistics provided by this report.
CPG Space	Display Common Provisioned Group (CPG) disk space usage. See “CPG Space Reports” (page 106) for the types of statistics provided by this report.
LD Space	Display logical disk space usage. See “LD Space Reports” (page 107) for the types of statistics provided by this report.
VV Space	Display virtual volume space usage. See “VV Space Reports” (page 107) for the types of statistics provided by this report.
Port Performance	Display port performance. See “Common Performance Metrics” (page 108) for the types of statistics provided by this report.
VLUN Performance	Display VLUN (volume-LUN) performance. See “Common Performance Metrics” (page 108) for the types of statistics provided by this report.
LD Performance	Display logical disk performance. See “Common Performance Metrics” (page 108) for the types of statistics provided by this report.
PD Performance	Display physical disk performance. See “Common Performance Metrics” (page 108) for the types of statistics provided by this report.
VV Cache Performance	Display VV (virtual volume) cache performance. See “VV Cache Performance Reports” (page 109) for the types of statistics provided by this report.
Node Cache Performance	Display Node Cache Performance. See “Node Cache Performance Reports” (page 109) for the types of statistics provided by this report.
CPU Performance	Display CPU performance. See “CPU Performance Reports” (page 110) for the types of statistics provided by this report.
Link Performance	Display performance of links between nodes. See “Link Performance Reports” (page 110) for the types of statistics provided by this report.
Adaptive Optimization	Display Adaptive Optimization reports. See “Adaptive Optimization Reports” (page 111) for the types of statistics provided by this report.

Select Systems

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more systems to include in the report.

Table 8 Select System

Option	Purpose
-All Systems-	The report will include all the systems that are included in the database.
list of specific system names	Select one or more systems to include in the report. Multiple systems can be selected.

NOTE: When selecting systems for System Reporter where there are multiple HP 3PAR Storage Systems available for reporting, consider limiting the number of systems that you select to avoid having to wait for longer report generations.

NOTE: When generating an Adaptive Optimization report, only one system can be selected at a time. See [“Using Adaptive Optimization” \(page 178\)](#) for details.

Select Domains

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more domains to be included in the report. This control is available for LD Space, VV Space, LD Performance, VLUN Performance and VV Cache Performance reports.

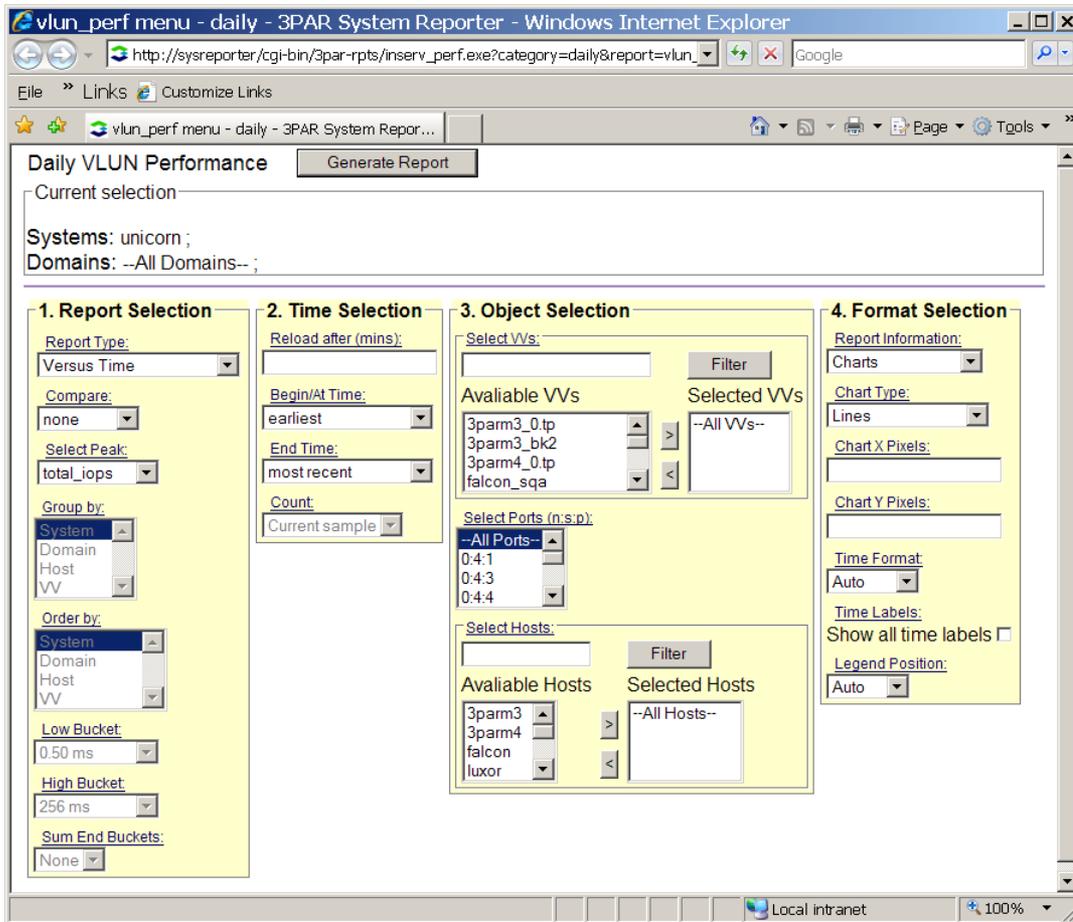
Table 9 Select Domains

Option	Purpose
-All Domains-	The report will include all the domains that are included in the database.
list of specific domain names	Select one or more domains to include in the report. Multiple domains can be selected.

Generating the Report

When you click on the **Build Report Menu** button in the Custom Reports window, a new window (or tab) is created containing a report menu for the selected report. An example of this report menu is shown in [Figure 26 \(page 70\)](#). This is the report menu for Daily VLUN Performance.

Figure 26 Report Menu for Daily VLUN Performance



The top of the report menu includes a section that lists the current selection of systems of domains. In this example, the menu includes only system named `unicorn` and all domains.

For most reports, the controls in the report menu are divided into four sections:

- **1. Report Selection:** controls the type of report that will be generated, see “[Report Selection](#)” (page 76).
- **2. Time Selection:** controls the time or range of time for which the report is generated, see “[Time Selection](#)” (page 84).
- **3. Object Selection:** controls the specific objects to be included in the report, see “[Object Selection](#)” (page 85).
- **4. Format Selection (web client only):** controls the format of the report, see “[Format Selection \(Web interface only\)](#)” (page 91).

NOTE: The Summary report does not include the Object selection, but provides options for selecting columns instead; see “[Select Summary Columns](#)” (page 81)

NOTE: The Excel client does not include a Format Selection but instead allows you to choose and modify the charts and graphs after the report is generated. See *Working with Excel Charts* for details.

NOTE: Adaptive Optimization reports have different options for the report selection and do not include the object selection section. For details on Adaptive Optimization reports, see “[Using Adaptive Optimization](#)” (page 178).

Use the controls in the menu to select the desired report parameters and then click **Generate Report**. See for a listing of the controls that are available for each report and then consult the relevant sections later in this chapter for complete details.

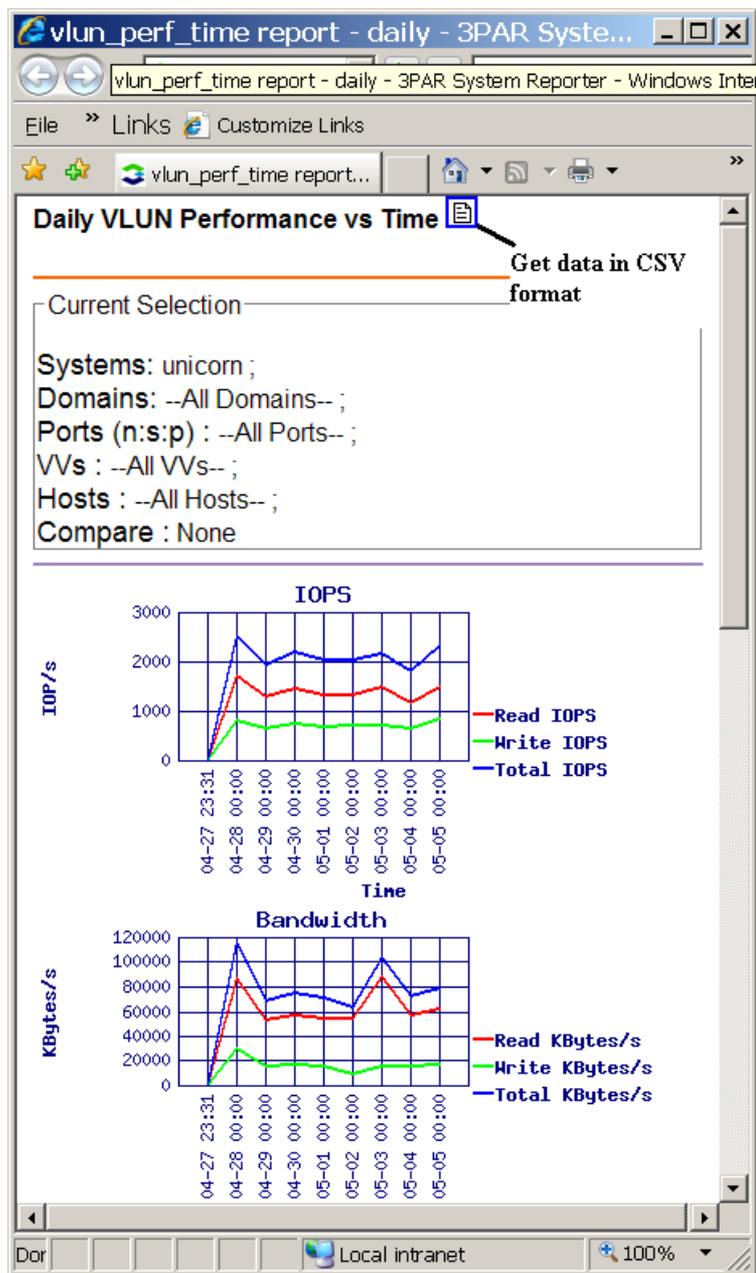
CAUTION: Filtering large numbers of objects may exceed the limit for your server, in which case you should reduce the number of objects being filtered and try again.

Once the report is generated, query results appear in one or more charts and/or tables in a new window (or tab) of the browser ([Figure 27 \(page 72\)](#)). The top of the report includes a section that describes the various selection parameters used to generate the report.

NOTE: Internet Explorer has a URL length limit of 2048 characters. The Firefox (Browser) has a 65,536 character limit. Apache produces a "413 Entity Too Large" error after 4000 characters. To avoid this type of error, break the query into a smaller size (especially for PD/LD reports).

NOTE: All the data that is shown in the results pane can be downloaded as a CSV format document (has a `.csv` extension) by clicking on a document icon that appears next to the main heading. If your browser is configured to open `.csv` files in a spreadsheet program such as Microsoft Excel this is a very convenient way of importing data into a spreadsheet.

Figure 27 Query Results



If necessary, use the scrollbox on the right to view all charts and/or tables.

Controls Available by Report

Table 10 (page 73) lists the set of controls that are valid for generating each report. The entries have the following meanings:

- A blank entry indicates that the control is not valid for report.
- A **Y** entry indicates that the control is valid for all report types.
- A **T** entry indicates that the control is valid for the report but only for **Versus Time** report types.
- An **H** entry indicates that the control is valid for the report but only for **Histogram At Time** report types.
- An **h** entry indicates that the control is valid for the report but only for **Histogram Versus Time** report types.
- An **A** entry indicates that the control is valid for the report but only for **At Time** report types.

Table 10 Controls for Each Report

Report: Section and Control	Smry	PD Space	CPG Space	LD Space	VV Space	PD Perf	Port Perf	LD Perf	VLUN Perf	VV Cache Perf	Node Cache Perf	CPU Perf	Link Perf	Adaptive Optimization
Main Menu														
"Sample Resolution" (page 67)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Select Systems" (page 68)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Select Domains" (page 69)				Y	Y			Y	Y	Y	Y			Y
Report Selection														
"Standard Report Types" (page 76)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Compare" (page 81)						T	T	T	T			T	T	
"Max. Number" (page 82)						T	T	T	T					
"Select Peak" (page 82)		Y	Y	Y	Y	Y	Y	Y	Y			Y		
"Group By" (page 82)		A	A	A	A	A	A	A	A	A	A	A	A	
"Order By" (page 83)		A	A	A	A	A	A	A	A	A	A	A	A	
"Low Bucket" (page 84)						h	h	h	h					
"High Bucket" (page 84)						h	h	h	h					
"Sum End Buckets" (page 84)						h	h	h	h					
Time Selection														
"Reload After (mins)" (page 84)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Begin/At Time" (page 84)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 10 Controls for Each Report *(continued)*

Report: Section and Control	Summary	PD Space	CPG Space	LD Space	VV Space	PD Perf	Port Perf	LD Perf	VLUN Perf	VV Cage Perf	Node Cache Perf	CPU Perf	Link Perf	Adaptive Optimization
"End Time" (page 85)		T	T	T	T	T	T	T	T	T	Y	T	T	T
"Count" (page 85)						H	H	H	H					
Object Selection														
"Select PDIDs" (page 85)		Y	Y			Y								
"Include Chunklets" (page 85)		Y	Y											
"Select Cage IDs" (page 86)		Y	Y											
"Select Disk Types" (page 86)		Y	Y	Y			Y							
"Select Disk Speed" (page 86)		Y	Y				Y							
"Include PDs in States" (page 86)		Y	Y											
"Select Usr CPGs" (page 87)					Y									
"Select Snp CPGs" (page 87)					Y									
"Select CPGs" (page 87)			Y	Y										Y
"Select LDs" (page 88)				Y				Y						
"Select RAID Types" (page 88)			Y	Y										
"Select LD Usage" (page 88)				Y										
"Select LUN" (page 88)									Y					

Table 10 Controls for Each Report *(continued)*

Report: Section and Control	Summary	PD Space	CPG Space	LD Space	VV Space	PD Perf	Port Perf	LD Perf	VLUN Perf	VV Gate Perf	Node Cache Perf	CPU Perf	Link Perf	Adaptive Optimization
"Select VVs" (page 89)					Y				Y	Y				
"Select Prov Types" (page 89)					Y									
"Select VV Types" (page 89)					Y									
"Select Hosts" (page 91)							Y							
"Select Port Types" (page 90)							Y							
"Select Port Types" (page 90)						Y	Y		Y					
"Select Port Rates" (page 90)							Y							
"Select Ports (n:s:p)" (page 90)														
"Select Hosts" (page 91)					Y				Y					
"Select Nodes" (page 91)											Y	Y	Y	
"Select To Nodes" (page 91)													Y	
"Select Queues" (page 91)													Y	
Format Selection														
"Report Information" (page 92)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Chart Types" (page 92)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Chart X Pixels" (page 96)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 10 Controls for Each Report (continued)

Report: Section and Control	Summary	PD Space	CPG Space	LD Space	VV Space	PD Perf	Port Perf	LD Perf	VLUN Perf	VV Gate Perf	Node Cache Perf	CPU Perf	Link Perf	Adaptive Optimization
"Chart Y Pixels" (page 96)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Time Labels" (page 96)		T	T	T	T	T	T	T	T	T	T	T	T	
"Legend Position" (page 96)		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
"Time Format" (page 96)		T	T	T	T	T	T	T	T	T	T	T	T	
"Select Space Unit" (page 97)		Y	Y	Y	Y									

Detailed descriptions for each of the controls described in this table are provided in the following sections.

Report Selection

For each report except for the Summary Report, you can use the **Report Type** pull-down list to choose a report type. There are two selections of report types: standard report types and adaptive optimization report types that only apply to the Adaptive Optimization feature that is available through an optional license with System Reporter.

Standard Report Types

The following standard report types are available (Table 11 (page 76)):

Table 11 Standard Report Types

Report Type	Purpose
Versus Time	<p>Display metrics versus time. The beginning of the time range is specified by the Begin/At Time control (see "Begin/At Time" (page 84)) and the end of the time range is specified by the End Time control (see "End Time" (page 85)).</p> <p>Unless the Compare control is available and set to something other than none, the Versus Time report shows aggregate metrics for all the selected objects over a range of time. See Figure 28 for an example.</p> <p>If the Compare control (see section "Compare" (page 81)) is set to a specific object, then the Versus Time report plots the metrics separately for each of up to 16 of the objects so that they can be compared. Figure 29 shows an example where the performance of systems is being compared.</p>
At Time	<p>Display metrics at the time specified by the Begin/At Time control (see "Begin/At Time" (page 84)). For this type of report, the Group By (see section "Group By" (page 82)) and Order By (see section "Order By" (page 83)) controls are enabled and can be used to compare different groupings of object type. "Group By System:Host At Time Report Example" (page 79) shows an example of an example with System and Host selected in the Group By control and IOPs/s selected in the Order By control.</p>
Histogram Versus Time	<p>Display service time histogram versus time, with a separate data series for each service time bucket between the Low Bucket (see "Low Bucket" (page 84)) and the High Bucket (see "High Bucket" (page 84)).</p>

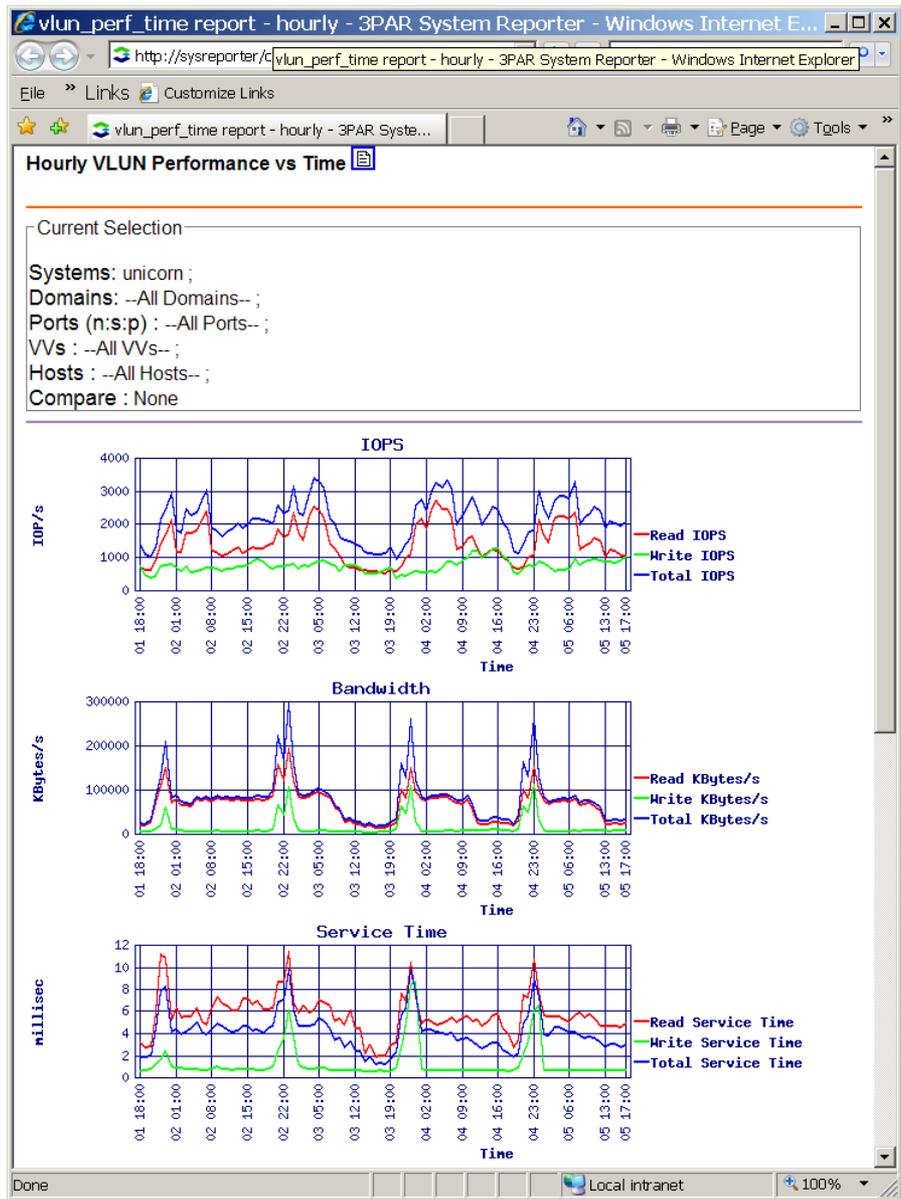
Table 11 Standard Report Types (continued)

Report Type	Purpose
	The beginning of the time range is specified by the Begin/At Time control (see “Begin/At Time” (page 84)) and the end of the time range is specified by the End Time control (see “End Time” (page 85)). See “Histogram Versus Time Report Example” (page 80).
Histogram At Time	Display histogram of metrics at the time specified by the Begin/At Time control (see “Begin/At Time” (page 84)). This control is only available for PD, LD, Port and VLUN performance reports. See “Histogram At Time Report Example” (page 81).

For each report type, controls in the report menu that do not apply to that report type are disabled.

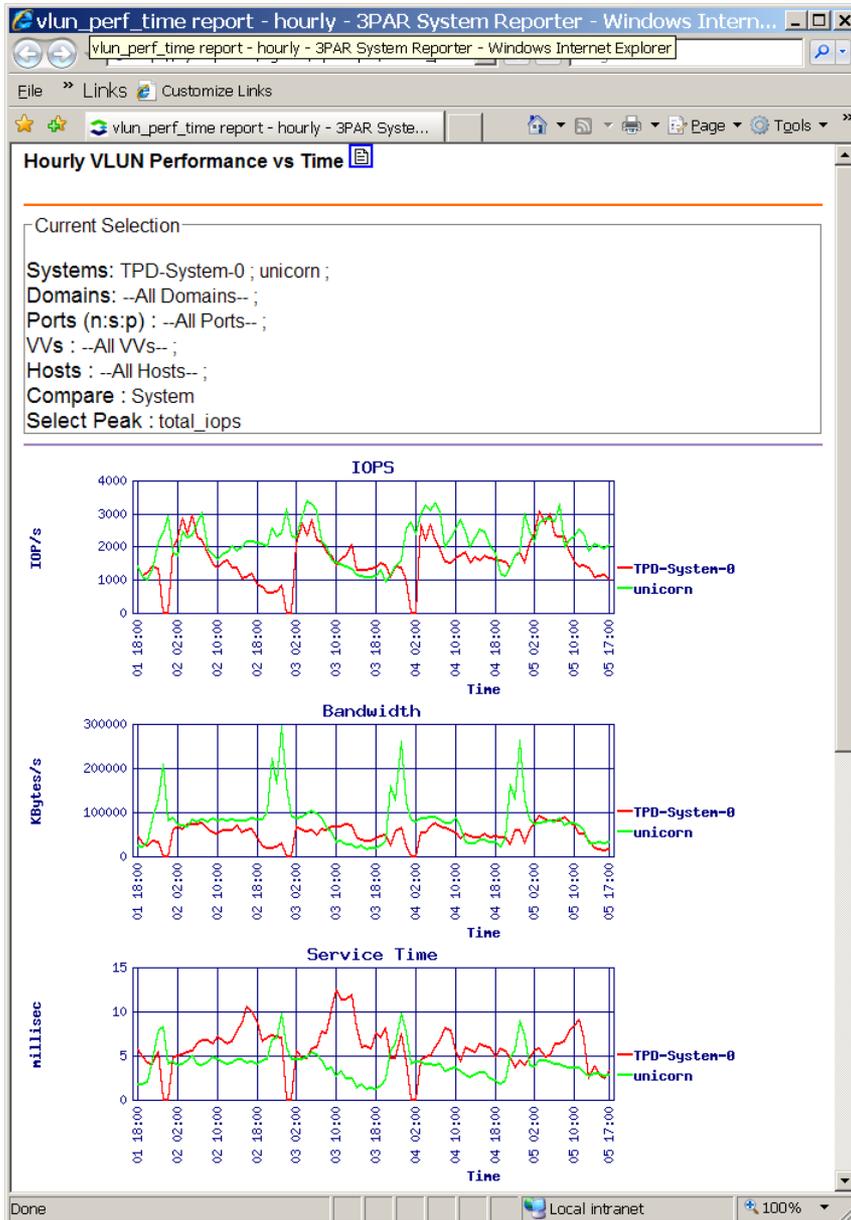
NOTE: See for complete details on Adaptive Optimization reports.

Figure 28 Versus Time Report Example



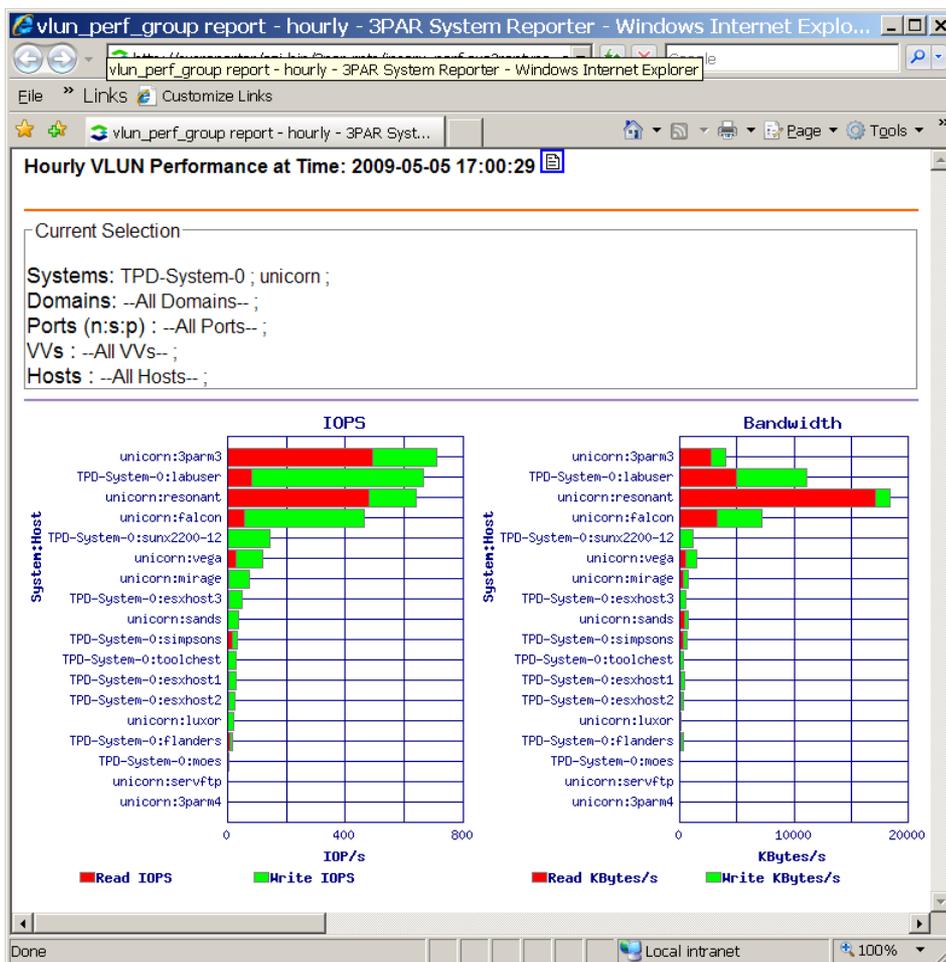
In the aggregate performance of all the selected objects (possibly in multiple systems) is shown. Typically, some components of the metric may be shown separately, for example Read and Write metrics are shown separately.

Figure 29 Versus Time with Compare=System Report Example



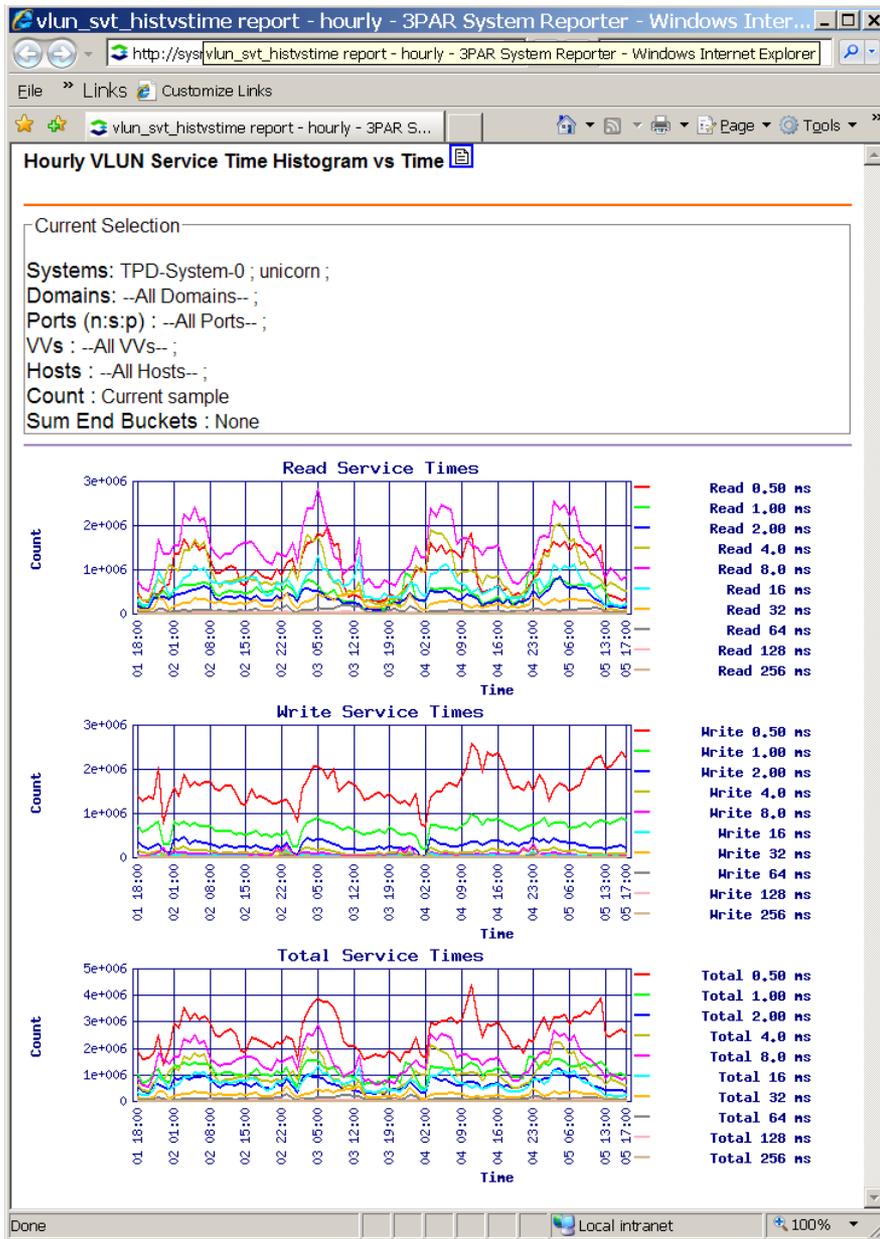
If you want to compare the performance of a particular class of objects (for example, systems), then you can use the **Compare** control (see section “Compare” (page 81)) available in some reports if the **Report Type** is **Versus Time**. Figure 29 shows an example where systems are being compared. In this kind of report, for each metric there is a single line per object being compared so it is not possible to break down components (for example Reads and Writes) of the metric.

Figure 30 Group By System:Host At Time Report Example



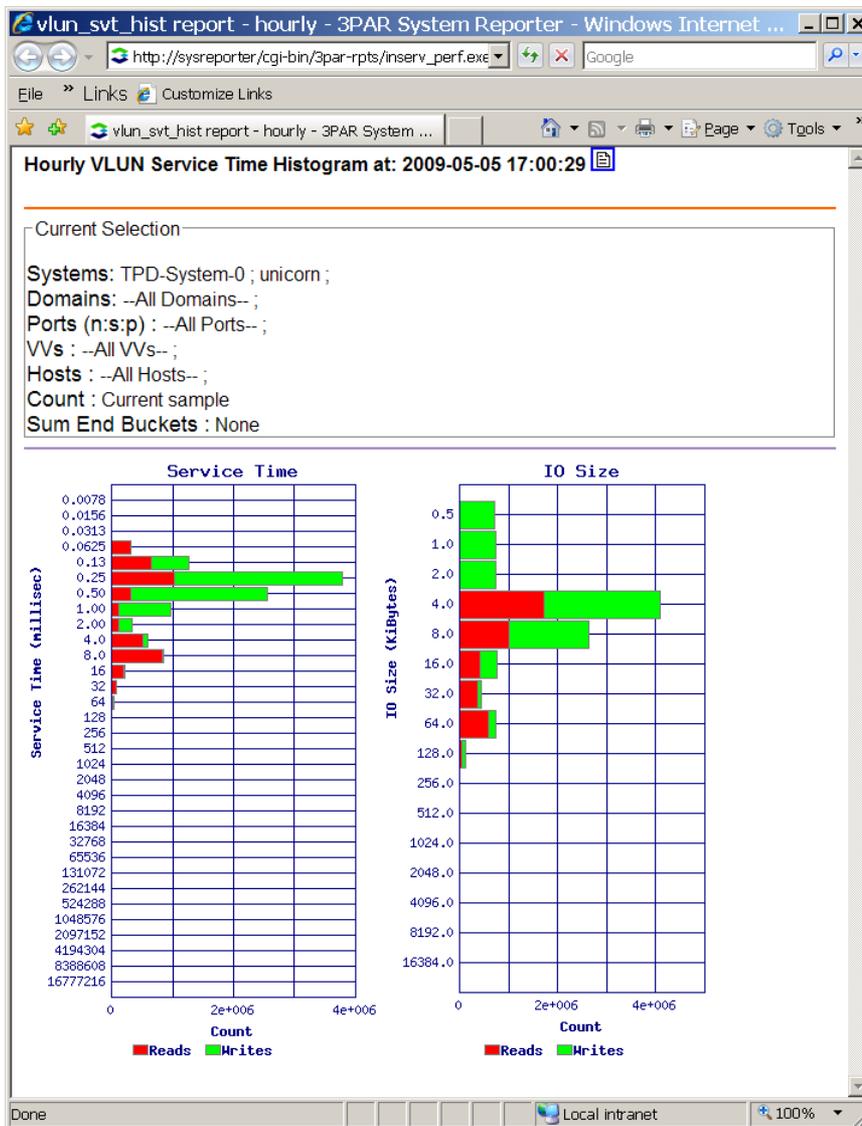
The **At Time** report can be used in conjunction with the **Group By** and **Order By** controls to do more detailed comparisons at a specific point in time. Figure 30 shows an **At Time** report with **System** and **Host** selected in the **Group By** control and **IOPs/s** selected in the **Order By** control. This allows us to easily compare the performance of hosts on different systems.

Figure 31 Histogram Versus Time Report Example



The Histogram Versus Time report provides a histogram or service times for the aggregate of the selected objects as shown in Figure 31 "Histogram Versus Time Report Example" (page 80). There is a separate line per service time bucket.

Figure 32 Histogram At Time Report Example



The Histogram At Time report provides a histogram of service times and IO sizes at a specific time for the aggregate of the selected objects as shown in Figure 32 “Histogram At Time Report Example” (page 81).

Select Summary Columns

This control is valid only for the Summary report and selects which columns are displayed in the Summary report.

Compare

This control is only valid for some **Versus Time** performance reports and the options available depend on the report. If the **none** option is selected, the aggregate performance metrics of the selected systems and objects are presented versus time. However, instead of the aggregate -- if you want to compare systems or objects -- select the appropriate **Compare** option. For example, if the **System** option is selected, the performance metrics for each selected system are plotted separately. An example of this is shown in Figure 29 “Versus Time with Compare=System Report Example” (page 78).

The number of objects that can be compared in a report is limited to 16 (the number of distinct colors available for charts). With some reports, if there are more than 16 objects present then the

Select Peak control (see section “[Max. Number](#)” (page 82)) can be used to select 16 objects to compare.

Max. Number

Specifies the maximum number of graphs that will be generated for a given report. This option is only valid when a compare is requested and when set to none, this option is ignored.

Select Peak

For some reports, when the number of objects to be compared (see section “[Compare](#)” (page 81)) exceeds 16, the **Select Peak** control can be used to select which 16 objects to compare. The **Select Peak** control is used to select a metric (for example, **total_iops**) and the 16 objects that have the highest peak in the selected metric in the selected time range are chosen.

Group By

This control is only valid for **At Time** reports and specifies the X (category) axis for the reports. The menu options available depend on the report selected and are listed with each report. One or more **Group By** options can be selected (use **CTRL+CLICK** to select multiple options) and the default if none are selected is **System**. There is an X axis value for (and/or a table row for) each unique combination of selected objects that exists in the data table. For example, if **System** and **Host** in the **Group By** options are selected for the VLUN performance report, a table row and chart X axis value exists in the report for each System:Host combination that exists in the VLUN data table as shown in Figure 30 “[Group By System:Host At Time Report Example](#)” (page 79). See for a listing of all available options.

Table 12 Group By

Option	Purpose
System	Group by HP 3PAR Storage System.
Domain	Group by domain name.
RAID Type	Group by RAID Type. See “ Select RAID Types ” (page 88) for the list of RAID types and descriptions.
Node	Group by Controller Node.
CageID	Group by Cage ID.
Mag	Group by magazine.
PDID	Group by ID of the Physical Disk.
ProvType	Group by VV provisioning type.
Disk	Group by disk position in magazine.
DiskType	Group by Disk Type.
DiskSpeed	Group by Disk Speed.
LD	Group by LD name.
LDUsage	Group by LD Usage. See “ Select LD Usage ” (page 88) for the list of LD usages and descriptions.
VV	Group by VV name.
VVType	Group by VV type.
CPG	Group by the CPG that the LD belongs to.
UsrCPG	Group by the Usr CPG to which the VV belongs.
Snpcpg	Group by the Usr CPG to which the VV belongs.

Table 12 Group By (continued)

Option	Purpose
PortNode	Group by port's node number (The n in n:s:p).
n:s:p	Group by the port (n:s:p).

Order By

This control is only valid for **At Time** reports and specifies the order of the X (category) axis for the reports. The menu options available depend on the report selected and are listed with each report. One or more **Order By** options can be selected (use **CTRL+CLICK** to select multiple options) and the default if none are selected is **System**. See for a listing of all available options.

Table 13 Order By

Option	Purpose
System	Order by HP 3PAR Storage System.
Domain	Order by domain name.
RAID Type	Order by RAID Type. See "Select RAID Types" (page 88) for the list of RAID types and descriptions.
Node	Order by Controller Node.
CageID	Order by Cage ID.
Mag	Order by magazine.
PDID	Order by ID of the Physical Disk.
ProvType	Order by VV provisioning type.
Disk	Order by disk position in magazine.
DiskType	Order by Disk Type.
DiskSpeed	Order by Disk Speed.
LD	Order by LD name.
LDUsage	Order by LD Usage. See "Select LD Usage" (page 88) for the list of LD usages and descriptions.
LDSize	Order by LD Slze.
LDFreeSize	Order by free space in LD.
LDRawSize	Order by LD Raw Slze.
VV	Order by VV name.
VVType	Order by VV type.
VirtSize	Order by VV virtual size.
TotalRsvdSize	Order by total reserved space associated with the VV.
TotalRawSize	Order by total raw space associated with the VV.
CPG	Order by the CPG that the LD belongs to.
UsrCPG	Order by the Usr CPG to which the VV belongs.
SnpCPG	Order by the Usr CPG to which the VV belongs.
PortNode	Order by port's node number (The n in n:s:p).
n:s:p	Order by the port (n:s:p).

Table 13 Order By *(continued)*

Option	Purpose
SelChunkSize	Order by total of the selected chunklets for PD space reports.
IOPs/s	Order by total IOPs.
Bandwidth	Order by total Bandwidth.
Service Time	Order by total service time.
IO Size	Order by total IO Size.
Queue Length	Order by queue length.
Busy Percentage	Order by the percentage of time that the VLUN is busy (i.e., has at least one outstanding IO operation).
Total Accesses	Order by total number of accesses for VV and node cache performance.
ReadHit%	Order by read hit% for VV and node cache performance.
WriteHit%	Order by write hit% for VV and node cache performance.
Idle	Order by Idle time for CPU performance reports.
Intr	Order by interrupts per sec for CPU performance reports.
Ctxt	Order by context switches per sec for CPU performance reports.

Low Bucket

This control is only valid for **Histogram Versus Time** reports and specifies the lowest service time bucket that will be included in the report.

High Bucket

This control is only valid for **Histogram Versus Time** reports and specifies the highest service time bucket that will be included in the report.

Sum End Buckets

This control is only valid for **Histogram Versus Time** reports and controls whether the **Low Bucket** and/or **High Bucket** should include the counts for lower or higher buckets respectively. For example, if **High Bucket** is 256 ms and **Sum End Buckets** is **High**, then the counts in the 256 ms bucket will include the sum of counts for all buckets higher than 256 ms.

Time Selection

Reload After (mins)

If you enter a number into the **Reload After (mins)** field, the browser will recompute and reload the report periodically after the specified number of minutes.

NOTE: Reloading the report is not equivalent to periodically clicking **Submit Query**. The report will be reloaded with the parameter values that were set at the time **Submit Query** was clicked.

Begin/At Time

For **Versus Time** or **Histogram Versus Time** reports this control selects the beginning of the time span of the report. For **At Time** or **Histogram At Time** reports this control selects the sample that is used in the report.

Table 14 Begin/At Time

Option	Purpose
earliest	Begin with the earliest sample available.
a list of specific time values	Begin at the specified time value. The values in the menu are populated based on the actual samples in the selected data table.
a list of relative time values	Begin at the specified time before the current time.

NOTE: Since the Summary Report requires data from both Space and Performance samples, the **Begin/At Time** control will only include times that have both space and performance samples.

End Time

The End Time control specifies the end of the time range. It is only valid for **Versus Time** or **Histogram Versus Time** reports.

Table 15 End Time

Option	Purpose
most recent	End with the most recent sample available.
a list of specific time values	End at the specified time value. The values in the menu are populated based on the actual samples in the selected data table.

Count

This control is only valid for **Histogram At Time** or **Histogram Versus Time** reports.

Table 16 Count

Option	Purpose
Current sample	Only count accesses within the current sample interval
Total	Count all accesses from the beginning (object creation or system reboot) until the current sample.

Object Selection

The type of object that can be selected and the available options depend on the type of report that is being generated as described in the following section.

Select PDIDs

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control”](#) (page 54)) used to select one or more PDIDs to include in the report. This control is available for PD Space and PD Performance reports.

Table 17 Select PDIDs

Option	Purpose
-All PDIDs-	The report will include all the PDIDs that are included in the database.
list of specific PDIDs	Select one or more PDIDs to include in the report.

Include Chunklets

This control is only valid for PD space reports and it allows you to control which chunklet types will be included in the report.

Table 18 Include Chunklets

Option	Purpose
-All Chunklets Types-	Show all the chunklet types in the report.
list of specific chunklet types	Select one or more chunklet types to include in the report. Multiple chunklet types can be selected using (CTRL+CLICK) . See Table 44 (page 106) for a list of chunklet types and their description.

Select Cage IDs

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more Cage IDs to include in the report. This control is available for PD Space reports.

Table 19 Select Cage IDs

Option	Purpose
-All Cage IDs-	The report will include all the Cages that are included in the database.
list of specific Cage IDs	Select one or more Cage IDs to include in the report. Multiple Cage IDs can be selected.

Select Disk Types

This control is available for PD Space, CPG Space, LD Space and PD Performance reports.

Table 20 Select Disk Types

Option	Purpose
-All Disk Types-	The report will include all the Disk Types that are included in the database.
list of specific Disk Types.	Select one or more Disk Types to include in the report. Multiple Disk Types can be selected using (CTRL+CLICK) . Currently available disk types are: <ul style="list-style-type: none"> • FC - Enterprise Fibre Channel disks • NL - Nearline Fibre Channel disks • SSD - Solid State Disks

Select Disk Speed

This control is available for PD Performance and PD Space reports.

Table 21 Select Disk Speed

Option	Purpose
-All Disk Speeds-	The report will include all the Disk Speeds that are included in the database.
list of specific Disk Speeds.	Select one or more Disk Speeds to include in the report. Multiple Disk Speeds can be selected using (CTRL+CLICK) . Currently available disk speeds are: <ul style="list-style-type: none"> • 10 - 10 K RPM disks • 7 - 7200 RPM disks • 15 - 15 K RPM disks • 150 - SSDs

Include PDs in States

This control is only available for PD Space reports.

Table 22 Include PDs in States

Option	Purpose
-All PD States-	Include all PD in the report.
valid	Include PDs in a valid state in the report.
Not valid	Include PDs not in a valid state in the report.

Select Usr CPGs

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more Usr CPGs to include in the report. This control is available for VV Space reports.

NOTE: For InForm OS releases prior to 2.3.1, a single CPG could be associated with a VV to provision its snapshot copy space. For thin provisioned VVs, this snapshot copy space was also used for provisioning the base VV’s space (the usr space).

Beginning with the InForm OS release 2.3.1, thin provisioned VVs can have separate spaces and corresponding separate CPGs to provision the base VV (usr CPG) and the snapshot copy space (snp CPG).

The VV space report therefore provides separate controls to select the Usr CPGs and the Snp CPGs. Since LDs can only belong to one CPG, the LD space reports only have a single Select CPG control (see [“Select CPGs” \(page 87\)](#)).

Table 23 Select Usr CPGs

Option	Purpose
-All CPGs-	The report will include all the Usr CPGs that are included in the database.
list of specific CPGs	Select one or more Snp CPGs to include in the report. Multiple CPGs can be selected.

Select Snp CPGs

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more Snp CPGs to include in the report. This control is available for VV Space reports. See note in [“Select Usr CPGs” \(page 87\)](#).

Table 24 Select Snp CPGs

Option	Purpose
-All CPGs-	The report will include all the Snp CPGs that are included in the database.
list of specific CPGs	Select one or more Snp CPGs to include in the report. Multiple CPGs can be selected.

Select CPGs

This is a filtered multi-selection control (see [“Filtered Multi-Selection Control” \(page 54\)](#)) used to select one or more CPGs to include in the report. This control is available for CPG Space, LD Space, and Adaptive Optimization Region IO Density reports.

Table 25 Select CPGs

Option	Purpose
-All CPGs-	The report will include all the CPGs that are included in the database.
list of specific CPGs	Select one or more CPGs to include in the report. Multiple CPGs can be selected.

Select LDs

This is a filtered multi-selection control (see “[Filtered Multi-Selection Control](#)” (page 54)) used to select one or more LDs to include in the report. This control is available for LD Space and LD Performance reports.

Table 26 Select LDs

Option	Purpose
-All LDs-	The report will include all the LDs that are included in the database.
list of specific LDs	Select one or more LDs to include in the report. Multiple LDs can be selected.

Select RAID Types

This control is only available for CPG and LD Space reports.

Table 27 Select RAID Types

Option	Purpose
-All RAID Types-	Include LDs of all RAID Types in the report.
RAID 0	Include RAID 0 LDs in the report.
RAID 1	Include RAID 1 LDs in the report.
RAID 5	Include RAID 5 LDs in the report.
RAID 6	Include RAID 6 LDs in the report.

Select LD Usage

This option is only available for the LD Space reports.

Table 28 Select LD Usage

Option	Purpose
-All LD Uses-	Include all LD uses.
V: Volume User Space	Include LDs mapped to VV user space.
C,SD: Copy Snapshot Data	Include LDs mapped to VV snapshot copy data space.
C,SA: Copy Snapshot Admin	Include LDs mapped to VV snapshot copy admin space.
P: Preserved data	Include LDs that are marked to hold data that cannot be flushed due to RAID set failure.
P,F: First Preserved data	Include Preserved data LDs that are also distinguished as being the first in the group of Preserved data LDs.
log: Logging	Include LDs that are marked for use in temporarily storing data during physical disk replacement operation. Typically, each Controller Node has 20GB of logging LD space.

Select LUN

This is a filtered multi-selection control (see “[Filtered Multi-Selection Control](#)” (page 54)) used to select one or more LUNs to include in the report. This control is available for VLUN Performance report.

Table 29 Select LUN Usage

Option	Purpose
-All LUNs-	The report will include all the LUNs that are included in the database.
list of specific LUNs	Select one or more LUNs to include in the report.

Select VVs

This is a filtered multi-selection control (see “[Filtered Multi-Selection Control](#)” (page 54)) used to select one or more VVs to include in the report. This control is available for VV Space, VLUN Performance and VV Cache reports.

Table 30 Select VVs

Option	Purpose
-All VVs-	The report will include all the VVs that are included in the database.
list of specific VVs	Select one or more VVs to include in the report. Multiple VVs can be selected using (CTRL+CLICK) . For VV Space reports, these only include base VVs, not snapshots. For VLUN Performance reports these include only VVs that have been exported as VLUNs.

Select Prov Types

This control is available for VV Space reports.

Table 31 Select Prov Types

Option	Purpose
-All Prov Types-	The report will include all the VV provisioning types that are included in the database.
list of specific provisioning types	Select one or more VV provisioning types to include in the report. Multiple VV provisioning types can be selected using (CTRL+CLICK) . Provisioning types include: <ul style="list-style-type: none"> • full - Fully provisioned VV. The base Usr space is fully provisioned with either with no Snp (snapshot) space or with statically allocates Snp space. • cpvv - Commonly provisioned VV. The base Usr space for the VV is fully provisioned and the snapshot space is associated with a Snp CPG. • tpvv - Thin provisioned VV. The base Usr space for the VV is provisioned from the Usr CPG. Snapshots allocate copy space from a Snp CPG if any. • tpsd - Old-style thin provisioned VV (created with an older release than 2.3.1) where both the base VV space and the snapshot copy space are allocated from the same Snp space associated with the same Snp CPG. • snp - The VV is a virtual copy (snapshot).

Select VV Types

This control is available for VV Space reports and is used to select the copy type of the VV.

Table 32 Select VV Types

Option	Purpose
-All VV Types-	The report will include all the VV Types that are included in the database.
list of specific VV Types	Select one or more VV Types to include in the report. Multiple VV Types can be selected using (CTRL+CLICK) . VV Types include: <ul style="list-style-type: none"> • base - Base volume (not a copy) • pcopy - Physical (full) copy • vcopy - Virtual copy (snapshot)

Select Port Types

This control is only available for Port Performance reports.

Table 33 Select Port Types

Option	Purpose
-All Port Types-	The report will include all the Port Types that are included in the database.
list of specific Port Types	Select one or more Port Types to include in the report. Multiple Port Types can be selected using (CTRL+CLICK) . Currently, the port types available include <ul style="list-style-type: none"> • disk - Fibre Channel disk ports • free - Free (unused) ports • host - Fibre Channel host ports • rcip - Remote Copy IP ports • rcfc - Remote Copy Fibre Channel ports • iscsi - iSCSI host ports

Select Port Rates

This is a filtered multi-selection control (see “[Filtered Multi-Selection Control](#)” (page 54)) used to select port rates to include in the report. This control is available for Port Performance reports.

Table 34 Select Port Rates

Option	Purpose
-All Port Rates	The report will include data for all port rates.
list of specific Port Rates	Select one or more port rates to include in the report.

Select Ports (n:s:p)

This control is available for PD Performance, Port Performance and VLUN Performance reports.

Table 35 Select Ports (n:s:p)

Option	Purpose
-All Ports-	The report will include all the Ports that are included in the database.
list of specific ports	Select one or more ports to include in the report. Multiple ports can be selected using (CTRL+CLICK) . The format for ports is n:s:p where <ul style="list-style-type: none"> • n is the controller node • s is the PCI slot in the controller node • p is the port number in a PCI slot

Select Hosts

This is a filtered multi-selection control (see “[Filtered Multi-Selection Control](#)” (page 54)) used to select one or more Hosts to include in the report. This control is available for VV Space and VLUN Performance reports.

Table 36 Select Hosts

Option	Purpose
-All Hosts-	The report will include all the Hosts that are included in the database.
list of specific Hosts	Select one or more Hosts to include in the report. Multiple Hosts can be selected.

Select Nodes

This control is available for CPU and Link Performance reports. For Link Performance reports it is the source node for the link.

Table 37 Select Nodes

Option	Purpose
-All Nodes-	The report will include all the Controller Nodes that are included in the database.
list of specific Nodes	Select one or more Nodes to include in the report. Multiple Nodes can be selected using (CTRL+CLICK).

Select To Nodes

This control is available for Link Performance reports and it selects the destination node for the link.

Table 38 Select To Nodes

Option	Purpose
-All Nodes-	The report will include all the Controller Nodes that are included in the database.
list of specific Nodes	Select one or more Nodes to include in the report. Multiple Nodes can be selected using (CTRL+CLICK).

Select Queues

This control is available for Link Performance reports.

Table 39 Select Queues

Option	Purpose
-All Queues-	The report will include all the queues that are included in the database.
list of specific Queues	Select one or more Queues to include in the report. Multiple Queues can be selected using (CTRL+CLICK).

Format Selection (Web interface only)

The format selection controls are used to specify how the output of a given report is formatted for presentation. The various controls for specifying the formatting options are described in the following sections.

NOTE: The Formatting options available from the web interface are not provided through the Excel client when generating a report but can be specified instead after the report is generated. For information on formatting reports when using the Excel client, see “[Working with Excel Charts](#)” (page 97).

Report Information

For each report except for the Summary Report, you can use the **Report Information** pull-down list to choose between displaying query results in the form of a chart, a table, or both.

Tables display query results in tabular format, as in the example shown in [Figure 33 \(page 92\)](#).

Figure 33 Example of a Table

Daily PD Space Usage											
Time	Normal, Used, OK Chunklets	Normal, Used, Failed Chunklets	Normal, Unused, Initialized Chunklets	Normal, Unused, Uninitialized Chunklets	Normal, Unused, Failed Chunklets	Spare, Used, OK Chunklets	Spare, Used, Failed Chunklets	Spare, Unused, Initialized Chunklets	Spare, Unused, Uninitialized Chunklets	Spare, Unused, Failed Chunklets	
2006/04/21 00:02:12	29540	0	1637	0	32	138	0	1504	0	29	
2006/04/22 00:00:41	61593	0	38211	0	92	141	0	5658	0	29	
2006/04/23 00:00:48	61593	0	38211	0	92	141	0	5658	0	29	
2006/04/24 00:00:51	61593	0	38211	0	93	141	0	5657	0	29	
2006/04/25 00:00:59	56279	0	38157	0	90	141	0	5657	0	29	
2006/04/26 00:01:01	55727	0	9977	0	90	141	0	3696	0	29	
2006/04/27 00:01:07	55727	0	9977	0	90	141	0	3696	0	29	
2006/04/28 00:04:21	56159	0	38279	0	88	141	0	5657	0	29	
2006/04/29 00:00:04	56591	0	40929	0	88	141	0	5791	0	29	
2006/04/30 00:00:07	56422	0	40916	182	88	312	0	5620	0	29	
2006/05/01 00:00:11	56374	0	41145	0	89	358	0	5574	0	29	
2006/05/02 00:03:03	56590	0	40929	0	89	142	0	5789	0	30	
2006/05/03 00:03:39	56467	0	40894	158	89	437	0	5494	0	30	
2006/05/04 00:03:48	57090	0	51909	0	81	142	0	5789	0	30	

Chart Lib

Allows you to optionally select the charting package used by System Reporter. By default, the (Static) charting capabilities are used, but you can try out the unsupported (Dynamic) option providing you are using a browser that supports html5 canvas tags (the latest versions of Windows E9, Firefox and Google Chrome browsers currently offer this support).

Chart Types

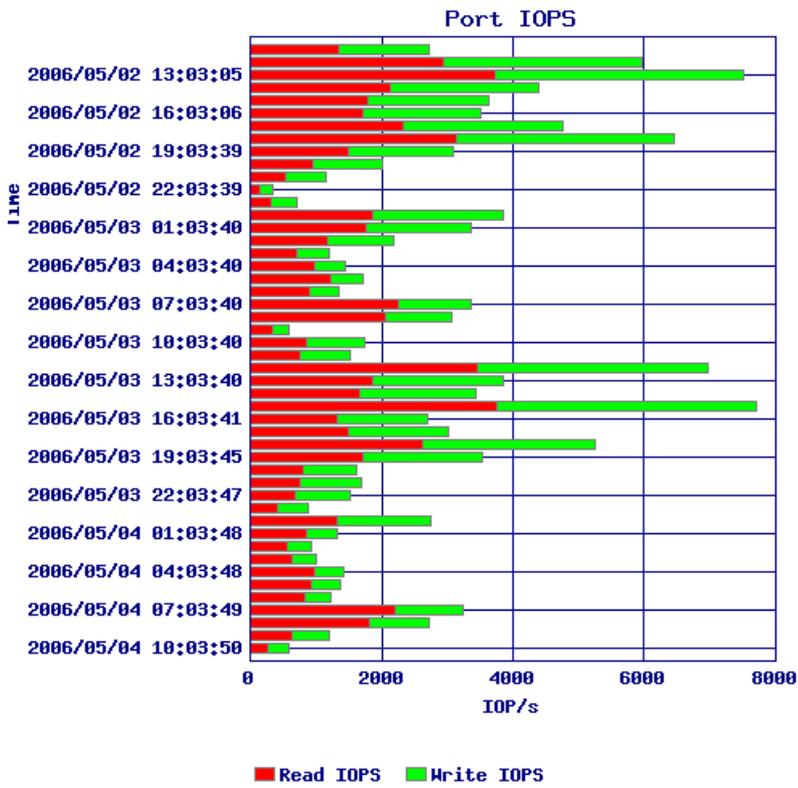
Stacked charts stack data on top of each other so you can also see totals. However, since the total height of the stacked charts (with **Stacked Bars**, **Stacked Vert Bars** and **Stacked Area** charts) is the sum of the individual segments, they are not suitable for some forms of data. For example, in performance charts, the average service time for I/O is not the sum of the average service time for reads plus the average service time for writes. Similarly, the average I/O size for I/O is not the sum of the average I/O sizes of reads plus the average I/O size of writes. Therefore, when you select stacked charts, only the total service time and total I/O sizes are shown. If you want to see the service time and I/O sizes for reads and writes separately, select **Bars**, **Vert Bars** or **Lines** from the **ChartType** list.

You can choose from the following chart types:

- **Stacked Bars**

Figure 34 Chart Using Stacked Bars

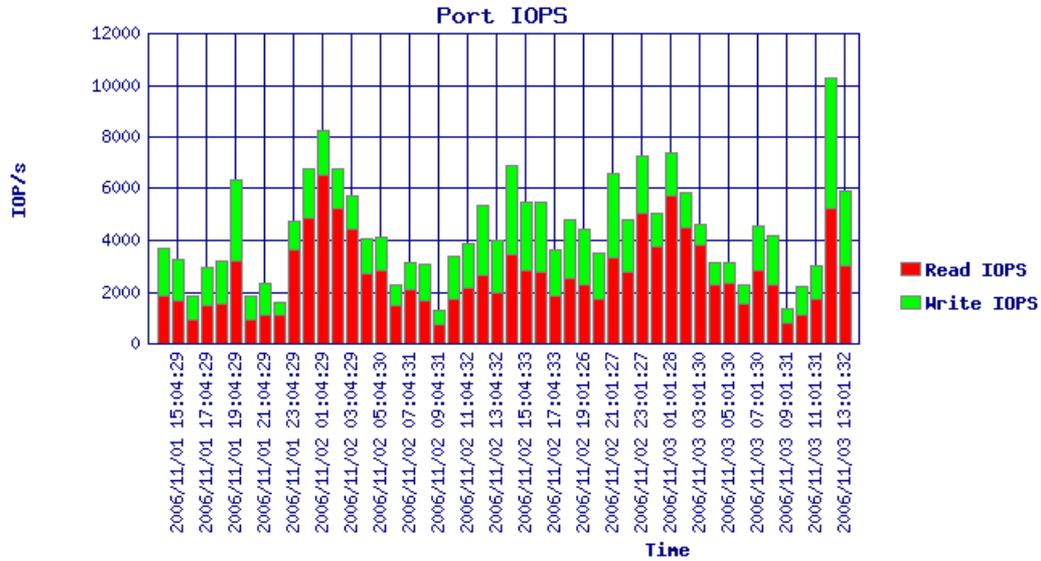
hourly Port Performance



- **Stacked Vert Bars**

Figure 35 Chart Using Stacked Vertical Bars

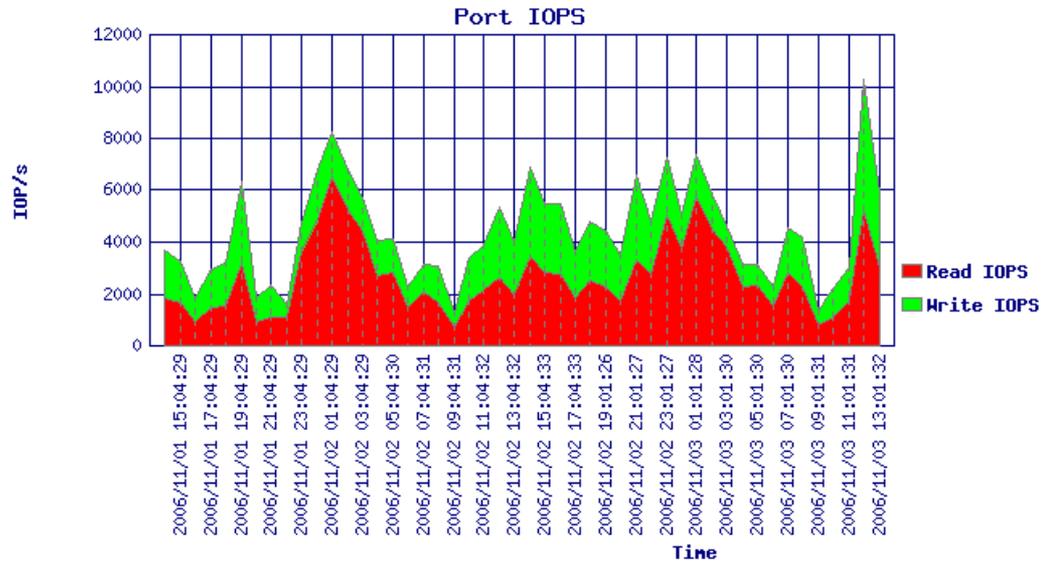
Hourly Port Performance



- Stacked Area

Figure 36 Chart Using Stacked Areas

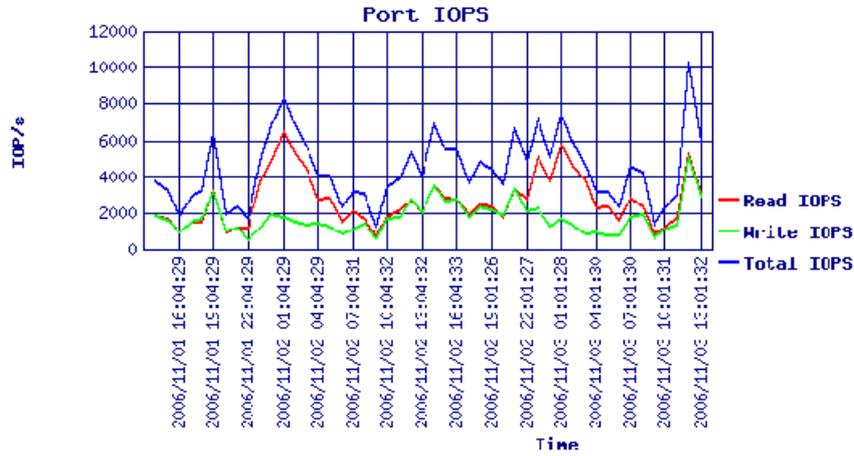
Hourly Port Performance



- Lines

Figure 37 Chart Using Lines

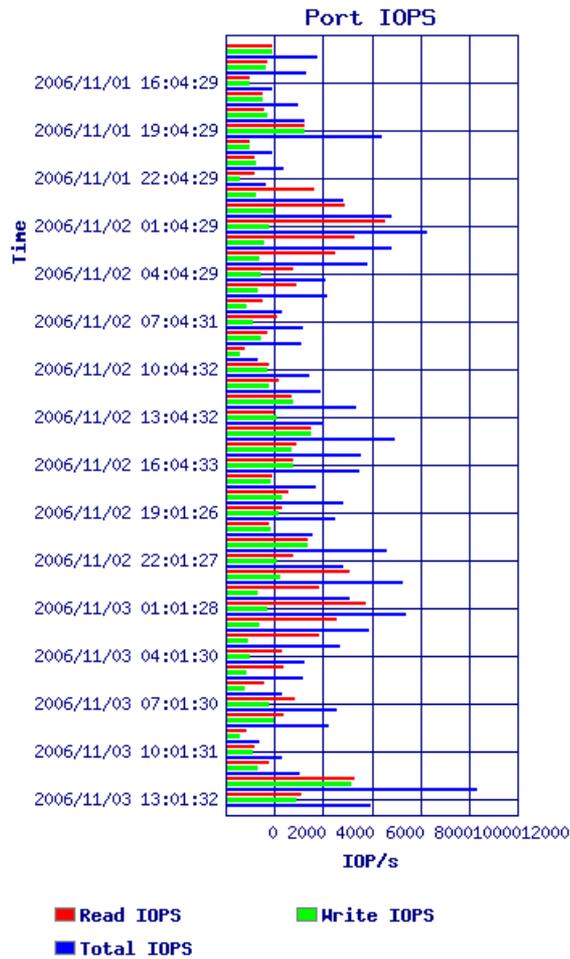
Hourly Port Performance



- Bars

Figure 38 Chart Using Bars

Hourly Port Performance



- **Vert Bars**

Figure 39 Chart Using Vert Bars

Hourly Port Performance

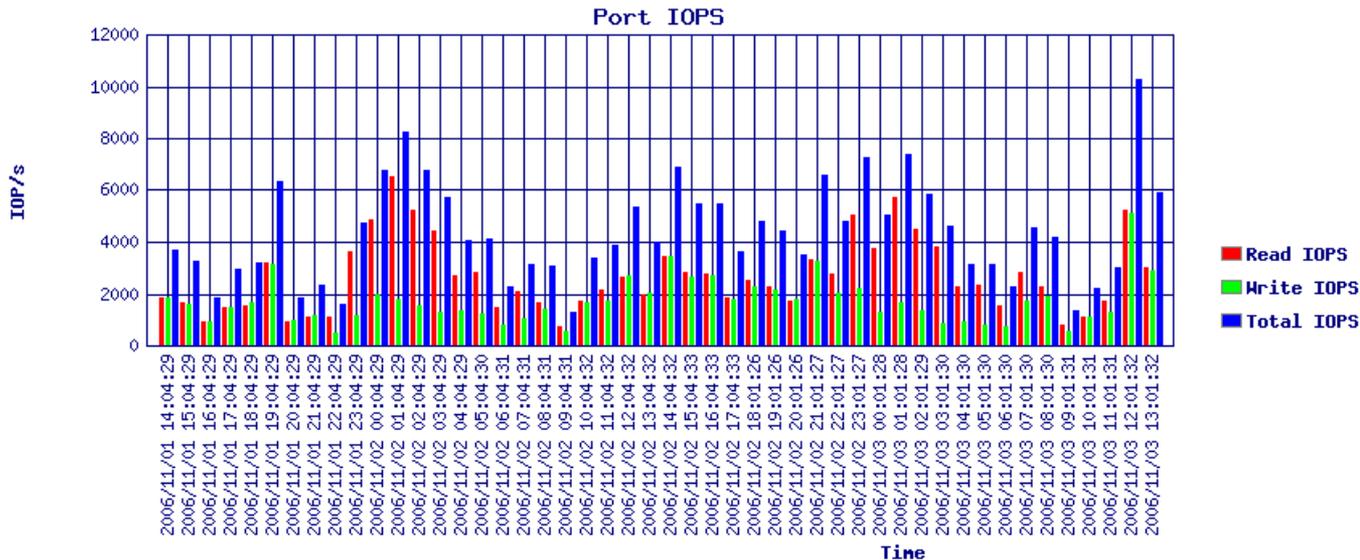


Chart X Pixels

You can explicitly specify the length in pixels of the X (category) axis of charts. If this is left blank the X axis length will be automatically chosen. If the length specified is too small, the chart may not be drawn.

Chart Y Pixels

You can explicitly specify the length in pixels of the Y (value) axis of charts. If this is left blank the Y-axis length will be automatically chosen. If the length specified is too small, the chart may not be drawn.

Time Labels

This is a checkbox that selects whether or not each point on the time axis of charts is labelled. If checked, each point on the time axis is labelled, otherwise only some points along the axis are labelled so that the length of the X (time) axis can be reduced.

Legend Position

You can choose the position of the chart legend.

Table 40 Legend Position

Option	Purpose
Auto	Legend position is chosen automatically based on the Chart Type .
Bottom	Legend is placed below the chart.
Right	Legend is placed to the right of the chart.

Time Format

This control selects how time values are labelled and is only valid when **Report Type** is **Versus Time**.

Table 41 Time Format

Option	Purpose
Auto	Time format is chosen automatically based on the sample resolution. If resolution is Daily, the format includes the month, date, hours and minutes, for example: 11-29 08:32. If the resolution is hourly, the format includes the date, hours and minutes, for example: 29 08:32. If the resolution is High-Res, the format include date, hours, minutes and seconds, for example: 29 08:32:05.
Full	Time values are labelled fully, for example: 2005-11-01 22:45:15
Relative	Time values are labelled relative to the current time, for example: 5d 07:50:03 (meaning 5 days, 7 hours, 50 mins, 3 secs ago).

Select Space Unit

This control selects the units in which storage space is represented. It is only valid for space reports.

Table 42 Select Space Unit

Option	Purpose
Chunklets	Space unit is chunklet (256 MB). This option is only available for PD Space reports.
MiB (2 ²⁰ bytes)	Space unit is 2 ²⁰ bytes ("binary" megabyte). This option is not available for PD space reports.
GB (10 ⁹ bytes)	Space unit is 10 ⁹ bytes.
GiB (2 ³⁰ bytes)	Space unit is 2 ³⁰ bytes ("binary" gigabyte).
TB (10 ¹² bytes)	Space unit is 10 ¹² bytes.
TiB (2 ⁴⁰ bytes)	Space unit is 2 ⁴⁰ bytes ("binary" terabyte).

Working with Excel Charts

When working with reports created with the Excel client, you have a variety of choices on how to format the information for presentation once the report has been generated.

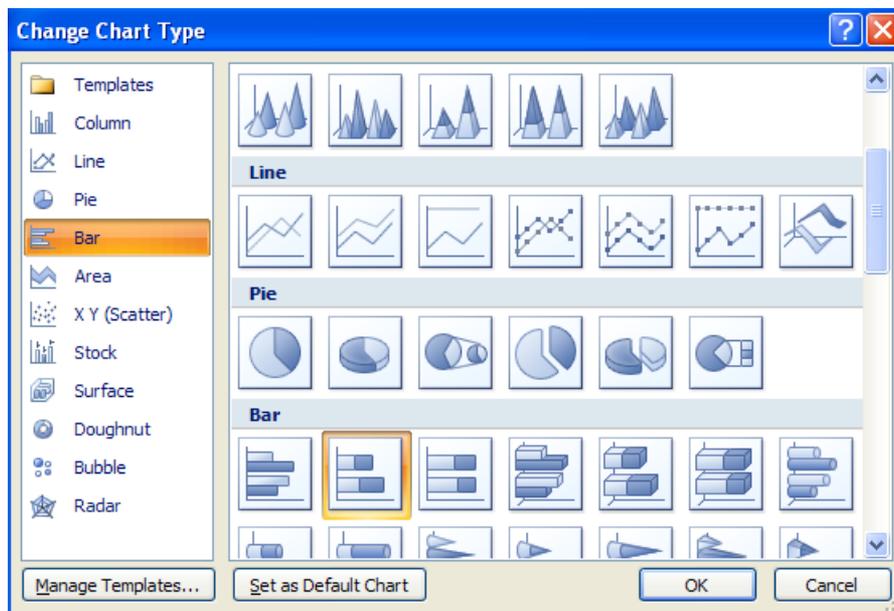
Choosing a Chart Type

If you want to change the default format for a report chart once it has been generated, you can have the data formatted using another type of chart:

1. Click the existing data that you want to reformat using another type of chart.
2. Click the **Design** tab.
3. Choose **Change Chart Type**.

The selection of available charts shown in appears.

Figure 40 Change Chart Type Menu for the Excel Client



4. Choose the appropriate selection from the list of chart types, then select on the chart that you want to use to display the report data.
5. Click **OK**.

△ **CAUTION:** The Excel Chart controls are configured for the default Chart Type. If you change the default Chart Type, you may also need to modify other chart controls for optimal viewing.

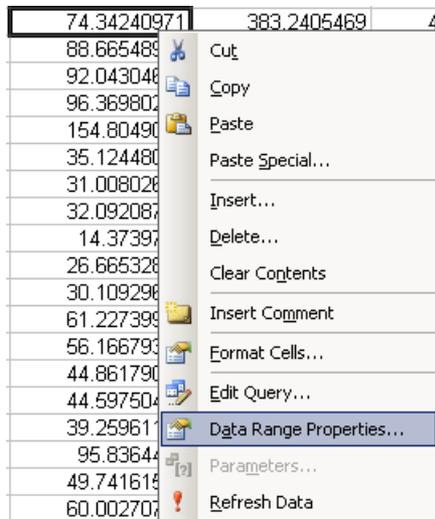
Refreshing Charts

Using a feature provided by System Reporter, you can refresh a report at any time with the latest reporting data by clicking on the **Refresh** button in the upper right side of the spread sheet. Excel allows for custom charts to be created along with the existing charts. These charts also get refreshed when the **Refresh** button is clicked.

Using another option provided through Excel, you can have the data refreshed according to a specific schedule by modifying the properties of query table from which the charting information is derived. When the table data is refreshed, any charts based on the table will also be refreshed. Setting a table to be automatically refreshed is most applicable with a table that was created using relative time selections. Here are the steps for modifying the query table to refresh the data according to a given schedule.

1. Right-click in any data cell and click **Data Range Properties** in the shortcut menu that appears.

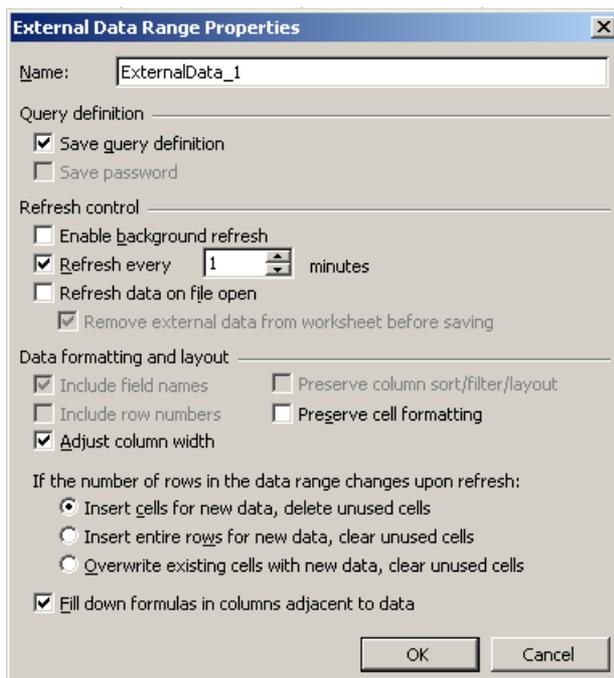
Figure 41 Data Range Properties Shortcut Menu



The **External Data Range Properties** dialog box appears.

2. Select the **Refresh every** checkbox and enter the desired value.

Figure 42 External Data Range Properties Dialog Box



3. Click **OK**.

Retaining Charts

The HP HP 3PAR System Reporter Excel client workbook is designed to reuse data sheets, redrawing charts on the reporting tool sheets with updated data or modified query parameters as required. To redraw a chart on a reporting tools sheet with updated data or modified query parameters, use the reporting tool sheet controls to make any changes and then click **Refresh**. The values on the corresponding data sheet are updated and the charts on the reporting tool sheet are redrawn.

It is also possible to retain a chart for further reference. This can be done by moving the chart and the data sheet linked to that chart to new worksheets. See the instructions that follow for a detailed description of how to complete this procedure.

To retain a chart for further reference

1. On the chart to be retained, right-click to bring up the shortcut menu.
2. On the shortcut menu, click **Location** (for Excel 2003) or **Move Chart** (for Excel 2007).

The **Chart Location** (for Excel 2003) or **Move Chart** (for Excel 2007) dialog appears.

3. In the **Chart Location** or **Move Chart** dialog box, click **As new sheet** or **New sheet** and then click **OK** to move the chart to a new sheet.

A new tab (with the specified name) appears at the bottom of the window for the new worksheet that holds the moved chart.

NOTE: Moving a chart may cause it to be resized. When a chart becomes smaller, you can adjust the axis labels to make them more readable. To adjust the axis, right-click the axis, click **Format Axis...** on the shortcut menu that appears, and then use the controls in the **Format Axis** dialog box.

4. Identify the data sheet that contains the source data for the moved chart. For example, a chart generated with the **PDIO** reporting tool sheet is linked to data on either the **PDSpace Time Data** or **PDSpace Items Data** sheet.

NOTE: If you don't know the name of the data sheet that holds the source data for a chart, you can right-click the chart and click **Source Data...**, then click the **Series** tab in the **Source Data** dialog box that appears. On the **Series** tab, look at the **Name** and **Values** boxes. The worksheet name appears in both boxes, just before the exclamation point.

5. Right-click the source data sheet tab and click **Rename**, then edit the data sheet name. Be sure to use a unique name that is not used for any other reporting tool or data sheets in the Excel client workbook.

After editing the data sheet name, you can return to the reporting tool sheet to perform additional queries without affecting the relocated chart or the data sheet holding its source data.

Exporting Executable Reports

When working with reports generated through the Excel client, you have the ability to save the contents to a file where the formatted data can be accessed. This is useful when you want to have data in a given report accessible to anyone in a directory or to email the reports. Here are the steps for saving an executable report:

1. Click on the appropriate tab at the bottom of the screen to access the report you want to save as an executable file.
2. Click on the **Windows Office Button** in the upper right corner.
3. Choose **Save As**, then select **Microsoft Office Excel Workbook** (for Excel 2007, you also have the option of saving as **Excel Macro Enabled Workbook**).
4. Assign a name to the file and click **Save**.

Customizing Excel Reports

Read this section for information on customizing the charts used to plot the results of database queries initiated with the HP 3PAR System Reporter Excel client.

To customize a chart, you must edit the data sheet associated with that chart. Customized charts typically make use of custom columns added to a data sheet.

NOTE: The Visual Basic code for the Excel client is password protected to prevent accidental modification. For customers who would like to add their own code, the password is “password”. Please observe caution when considering using another password other than the default, since subsequent access would not be possible if knowledge of the altered password were to become unavailable for some reason.

Adding Custom Columns to a Data Sheet

The following instructions illustrate how to add custom columns to a data sheet. In this example, you will add the custom columns **Non-failed spare** and **Percent spare used** to the **Dly PDSpace (1)** sheet. In the following section, “[Creating a Custom Chart](#)” (page 101), you will create a custom chart using one of these new columns.

NOTE: User defined columns can only be added as the last column in the data sheet.

To add custom columns to the **Dly PDSpace (1)** sheet:

1. On the **PD Space** reporting tool sheet, create a chart with data that spans any time range.
The **Dly PDSpace (1)** sheet is created to hold the data represented in the chart at the bottom of the **PD Space** reporting tool sheet.
2. On the **Dly PDSpace (1)** sheet, scroll to the first empty cell to the right of Data table. In cell N52, click and type a new header for **Non-failed spare**.
3. Below the new header, in cell N53, click and then add a formula to sum the columns that have values for spares that are not failed:
`=H53+J53+K53`
Press **ENTER** after typing in the formula. A calculated value should appear in cell N53.
4. Click cell N53 and drag the cell selection down so that the formula is copied to all rows of the table.
To drag the cell selection down, you must position your pointer over the lower right corner of the cell and then click and drag. Calculated values should appear in all cells of column N.
5. In cell O52, click and type a new header for **Percent spare used**.
6. Below the new header, in cell O53, add a formula to calculate the percentage of non-failed spare space that is used:
`=IF(N53>0, H53/N53*100, 0)`
Press **ENTER** after typing in the formula. A calculated value should appear in cell O53.
7. Click cell O53 and drag the cell selection down so that the formula is copied to all rows of the table.
To drag the cell selection down, you must position your pointer over the lower right corner of the cell and then click and drag. Calculated values should appear in all cells of column O.
On the **PDSpace Time Data** sheet, the newly added columns are appropriately resized and populated with values that reflect the most recent query results.

Creating a Custom Chart

After adding custom columns to a data sheet, you can use those columns to create custom charts. The following instructions illustrate how to create a custom chart from one of the custom **Dly PDSpace (1)** sheet columns you added in “[Adding Custom Columns to a Data Sheet](#)” (page 101).

NOTE: The following example uses the XY chart type, which works well for **High-Res**, **Daily**, and **Hourly** resolutions (see “[Sample Resolution](#)” (page 67)). Other chart types may not work as expected.

For Excel 2003:

To create a custom chart with the **Percent spare used** column in Excel 2003:

1. On the **Dly PDSpace (1)** sheet, click and drag to select the **Time** column (header cell and all data cells) and then **CTRL+CLICK** and drag to select the **Percent spare used** column.
2. On the Microsoft Excel 2003 main menu, click **Insert** and then **Chart** to bring up the **Chart Wizard**.
3. Use the chart wizard to create a new chart.
 - a. In **Step 1** of the wizard: under **Chart type** select **XY** and under **Chart sub-type** select **Scatter with data points connected by lines without markers** and then click **Next**.
 - b. In **Step 2** and **Step 3** of the wizard, click **Next** to accept the defaults.
 - c. In **Step 4** of the wizard, click **As new sheet** and then click **Finish**.

A new chart appears on a new workbook sheet. This chart is associated with the custom column on **PDSpace Time Data** sheet. This data sheet is tied to a query that can be modified and refreshed using the user interface controls on the **PD Space** reporting tool sheet.

For Excel 2007:

To create a custom chart with the **Percent spare used** column in Excel 2007:

1. On the **Dly PDSpace (1)** sheet, click and drag to select the **Time** column (header cell and all data cells) and then **CTRL+CLICK** and drag to select the **Percent spare used** column (header cell and all data cells).
2. On the Microsoft Excel main menu, click **Insert**→**Scatter**→**Scatter with Straight line**.
New chart will appear on the same sheet
3. Click on the Chart and from the Excel main menu select **Layout** to change desired chart properties (Chart Title, Axis Titles, Axis properties, etc.).
If you want to move the Chart to a new sheet, right click on the chart and select **Move Chart**.

Saving Customized Settings

After customizing data sheets and creating customized charts, save the customized workbook in Excel by using the **Save As...** option and saving a copy of the workbook with a new name.

This saves the entire workbook, including the customizations you made to the data sheet and any custom charts. When you close and then reopen the new workbook, the ability to create new charts using the same formulas is preserved.

NOTE: To obtain an entirely new copy of the HP 3PAR System Reporter Excel client workbook, follow the instructions provided in [“Installing the Excel Client \(Optional\)”](#) (page 48).

6 Interpreting Report Metrics

This chapter describes the metrics that are shown for each of the various reports generated through System Reporter.

Summary Report

Figure 43 (page 103) shows a typical example of the metrics that are provided by a Summary Report.

Figure 43 Summary Report Metrics

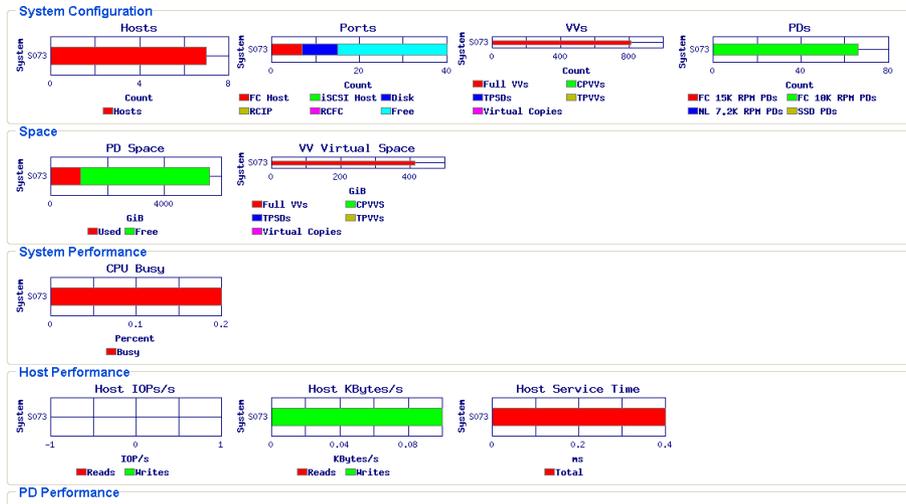


Table 43 (page 103) describes all of the metrics available for Summary Reports. Only some of the metrics described here are shown by default, the others depend on which columns are selected through the **Select Summary Columns** control (see “Select Summary Columns” (page 81)).

Table 43 Summary Report Metrics

Metric	Description
All Columns	Shows all columns.
IP Address	IP Address of the HP 3PAR Storage System.
OS Revision	The InForm OS Revision.
System Model	The HP 3PAR Storage System model.
Serial	Serial number of the HP 3PAR Storage System.
Controller Nodes	Number of Controller Nodes.
Hosts	Number of hosts.
Exported VVs	Number of distinct VVs exported as VLUNs. Shown by default.
All VVs	Shows all VVs.
Virtual Copies	Shows all Virtual Copies.
Full VVs	Shows all full VVs.
CPVVs	Shows all CPVVs.
TPSDs	Shows all TPSDs.

Table 43 Summary Report Metrics *(continued)*

Metric	Description
TPVVs	Shows all TPVVs.
Virtual Copies	Shows all virtual copies.
Host Ports	Number of Host Ports on which VLUNs are exported.
FC Host Ports	Number of host ports.
iSCSI Host Ports	Number of iSCSI host ports.
Disk Ports	Number of disk ports.
RCIP Ports	Number of RCIP ports.
RCFC Ports	Number of RCFC ports.
Free Ports	Number of free ports.
FC 15K RPM PDs	Number of FC 15K physical disks.
FC 10K RPM PDs	Number of FC 10K physical disks.
NL 7K RPM PDs	Number of 7K near line physical disks.
SSD PDs	Number of 150 SSD physical disks.
PDs	Number of PDs (physical disks).
CPU Busy%	Average percentage of time that the CPUs are busy on all controller nodes. Shown by default.
Host Read IOPs	IO read operations/sec to host.
Host Write IOPs	IO write operations/sec to host.
Host Total IOPs	Total (reads + writes) IO operations/sec to hosts. Shown by default.
Host Read KByte/s	Host reads in KByte/s.
Host Write KByte/s	Host writes in Kbyte/s.
Host Total KByte/s	Total (reads + writes) bandwidth in KBytes/s to hosts. Shown by default
Host Total Svct ms	Total (reads + writes) average service time in millisec to hosts. Shown by default.
Host Total IOSz Kbytes	Total (reads + writes) average IO size in KBytes to hosts.
Host Total Queue Length	Total (reads + writes) queue length to hosts. Shown by default.
PD Read IOPs	IO read operations/sec to PDs.
PD Write IOPs	IO write operations/sec to PDs.
PD Total IOPs	Total (reads + writes) IO operations/sec to PDs.
PD Read KByte/s	PD reads in Kbytes.
PD Write KByte/s	PD writes in Kbytes.
PD Total KByte/s	Total (reads + writes) bandwidth in KBytes/s to PDs.
PD Total Svct ms	Total (reads + writes) average service time in millisec to PDs.

Table 43 Summary Report Metrics *(continued)*

Metric	Description
PD Total IOSz KByte/s	Total (reads + writes) average IO size in KBytes to PDs.
PD Total Queue Length	Total (reads + writes) queue length to PDs.
Free PD Capacity GiB	Free PD space available. Shown by default.
Used PD Capacity GiB	Used PD space.
Total PD Capacity GiB	Total PD space available. Shown by default.
Free (FC 10,15K), (NL 7k), (SSD) RPM PDS Capacity GiB	Free space for a specified drive in GiBs.
Used (FC 10,15K),(NL 7k), (SSD) RPM PD Capacity GiB	Used space for a specified drive type in GiBs.
Total (FC 10,15K), (NL 7k), (SSD) RPM PD Capacity GiB	Total space for a specified drive type in GiBs.
VCopies Virtual Size GiB	Virtual size in GiBs for VCopies.
CPVVs Virtual Size GiB	Virtual size in GiBs for CPVVs.
Full Virtual Size GiB	Total virtual size in GiBs.
TPSDs Virtual Size GiB	Virtual size in GiBs for TPSDs.
TPVVs Virtual Size GiB	Virtual size in GiBs for TPVVs.
All VVs Virtual Size GiB	All VV virtual size in GiBs.
Admin Rsvd GiB	Reserved admin space in GiBs.
Snap Rsvd GiB	Reserved snap space in GiBs.
User Rsvd GiB	Reserved user space in GiBs.
Total Rsvd GiB	Total reserved in GiBs.
Admin Raw Rsvd GiB	Reserved raw admin space in GiBs.
Snap Raw Rsvd GiB	Reserved raw snap space in GiBs.
User Raw Rsvd GiB	Reserved user snap space in GiBs.
Total Raw Rsvd GiB	Total reserved raw space in GiBs.
LDs	Logical disks
LD Space GiB	Logical Disk space in GiBs.
RAID (0,1,5,6) LD	RAID logical disks
RAID (0,1,5,6) LD Space GiB	RAID logical disk space in GiBs.
FC RAID (0,1,5,6) LDs	Fibre Channel RAID logical disks.
FC RAID (0,1,5,6) LD Space GiB	Fibre Channel RAID logical disk space in GiBs.
NL RAID (0,1,5,6) LDs	Near Line RAID logical disks.
NL RAID (0,1,5,6) LD Space GiB	Near Line RAID logical disk space in GiBs.
SSD RAID (0,1,5,6) LDs	SSD RAID logical disks.
SSD RAID (0,1,5,6) LD Space GiB	SSD RAID logical disk space in GiBs.

Space Reports

This section describes the metrics that are shown for the Physical Disk (PD), Common Provisioning Group (CPG), Logical Disk (LD), and Virtual Volume (VV) space reports used to determine storage capacities and usage.

PD Space Reports

PD (Physical Disk) Space reports show the amount of physical disk space composed of chunklets according to the usage types shown in [Table 44 \(page 106\)](#).

Table 44 Chunklet Usage Types

Usage Type	Description
Normal, Used, OK	Chunklets allocated to logical disks and containing data that is currently accessible to the system.
Normal, Used, Failed	Chunklets allocated to logical disks and containing data but inaccessible because they are failed.
Normal, Unused, Initialized	Chunklets currently initialized and available for use by logical disks.
Normal, Unused, Uninitialized	Chunklets being cleaned.
Normal, Unused, Failed	Chunklets initialized but unallocated; these chunklets do not contain data and are failed.
Spare, Used, OK	Chunklets currently used for sparing and containing data that is accessible by the system.
Spare, Used, Failed	Chunklets used for sparing but containing data that is currently inaccessible to the system because they are failed.
Spare, Unused, Initialized	Chunklets initialized and available for use, but only as spares.
Spare, Unused, Uninitialized	Un-initialized chunklets reserved for use as spares.
Spare, Unused, Failed	Chunklets reserved as spares but unusable because they are failed.
Total	Total of all the selected chunklets. The total is only shown in non-stacked charts.

For charts, scroll down to locate the color key listing the chunklet types represented in the chart ([Figure 44 \(page 106\)](#)).

Figure 44 Chunklet Color Key

■ Normal, Used, OK	■ Normal, Used, Failed
■ Normal, Unused, Initialized	■ Normal, Unused, Uninitialized
■ Normal, Unused, Failed	■ Spare, Used, OK
■ Spare, Used, Failed	■ Spare, Unused, Initialized
■ Spare, Unused, Uninitialized	■ Spare, Unused, Failed

CPG Space Reports

CPG (Common Provisioning Group) Space reports show the amount of space according to the usage types shown in [Table 45 \(page 106\)](#).

Table 45 CPG Space Usage Types

Usage Type	Description
Admin Used Space GiB	Amount of Admin space currently used for CPG allocation.
Snap Used Space GiB	Amount of raw snap space currently used for CPG allocation.
User Used Space GiB	Amount of user space currently used for CPG allocation.
Total Used Space GiB	Total amount of space currently used for CPG allocation.

Table 45 CPG Space Usage Types *(continued)*

Usage Type	Description
Admin Free Space GiB	Amount of admin free space available for CPG allocation.
Snap Free Space GiB	Amount of raw snap space available for CPG allocation.
User Free Space GiB	Amount of user free space available for CPG allocation.
Total Free Space GiB	Total amount of free space available for CPG allocation.
Admin Space GiB	Amount of remaining Admin space available for currently allocated CPGs.
Snap Space GiB	Amount of remaining Snap space available for currently allocated CPGs.
User Space GiB	Amount of remaining User space available for currently allocated CPGs.
Total Space GiB	Total amount of remaining space available for currently allocated CPGs.
Growth Space GiB	Amount of space available for CPG growth.

LD Space Reports

Table 46 (page 107) describes the metrics that are shown for LD (Logical Disk) Space Reports.

Table 46 LD Space Report Metrics

Metric	Description
Raw Space	The raw physical disk space used by the LDs.
Used Space	The LD space that is mapped to VVs.
Free Space	The LD space that is not mapped to VVs. Note that logging and preserved data LDs cannot be mapped to VVs and will show up as free space.
Total Space	The total LD space (Used Space + Free Space). This metric is not shown in stacked charts.

VV Space Reports

Table 47 (page 107) describes the metrics that are shown for VV (Virtual Volume) Space Reports

Table 47 VV Space Report Metrics

Metric	Description
User Raw Rsvd Space	Raw physical disk space reserved for User (base volume) Space for the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Snap Raw Rsvd Space	Raw physical disk space reserved for Snap (snapshot copy) Space for the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Admin Raw Rsvd Space	Raw physical disk space reserved for Admin (snapshot metadata) Space for the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Total Raw Rsvd Space	Total (user + snap + admin) Raw physical disk space reserved for the selected VV. This is zero for Virtual Copy (snapshot) VVs.
User Used Space	User space that is used by the selected VVs. For fully provisioned VVs, all the User Rsvd Space is considered used. For thin provisioned VVs, only some part of the User Rsvd Space may be used (the rest is User Free Space). This is zero for Virtual Copy (snapshot) VVs.
User Free Space	User Rsvd Space that is not actually used. This is zero for Virtual Copy (snapshot) VVs.
User Rsvd Space	Space that is reserved for User (base volume) Space for the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Snap Used Space	Snap space that is used by the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Snap Free Space	Snap Rsvd Space that is not actually used. This is zero for Virtual Copy (snapshot) VVs.

Table 47 VV Space Report Metrics (continued)

Metric	Description
Snap Rsvd Space	Space that is reserved for Snap Space for the selected VVs (Snap Used Space + Snap Free Space). This is zero for Virtual Copy (snapshot) VVs.
Snap Vcopy Space	Snap Space used by Virtual Copy (snapshot) VVs only. This is zero for all VVs that are not Virtual Copies (snapshots). Note that this value is calculated on the HP 3PAR Storage System by running the <code>updatesnapSPACE</code> command. System Reporter does not run this command. Beginning with the InForm OS 2.3.1 release running the <code>updatesnapSPACE</code> command can be scheduled periodically on the HP 3PAR Storage System.
Admin Used Space	Admin (snapshot metadata) space that is used by the selected VVs. This is zero for Virtual Copy (snapshot) VVs.
Admin Free Space	Admin Rsvd Space that is not actually used. This is zero for Virtual Copy (snapshot) VVs.
Admin Rsvd Space	Space that is reserved for Admin Space for the selected VVs (Admin Used Space + Admin Free Space). This is zero for Virtual Copy (snapshot) VVs.
Admin Vcopy Space	Admin Space used by Virtual Copy (snapshot) VVs only. This is zero for all VVs that are not Virtual Copies (snapshots). Note that this value is calculated on the HP 3PAR Storage System by running the <code>updatesnapSPACE</code> command. System Reporter does not run this command. Beginning with the InForm OS 2.3.1 release running the <code>updatesnapSPACE</code> command can be scheduled periodically on the system.
Total Vcopy Space	Total (Snap + Admin) Vcopy space used by Virtual Copy (snapshot) VVs only. This is zero for all VVs that are not Virtual Copies (snapshots). Note that this value is calculated on the HP 3PAR Storage System by running the <code>updatesnapSPACE</code> command. System Reporter does not run this command. Beginning with the InForm OS 2.3.1 release running the <code>updatesnapSPACE</code> command can be scheduled periodically on the system.
Total Used Space	Total used space for the selected VVs (User Used Space + Snap Used Space + Admin Used Space). This is zero for Virtual Copy (snapshot) VVs.
Total Rsvd Space	Total reserved space for the selected VVs (User Rsvd Space + Snap Rsvd Space + Admin Rsvd Space). This is zero for Virtual Copy (snapshot) VVs.
Virtual Size	Virtual size of the selected VVs.

Performance Reports

This section describes the metrics for reports that provide performance data for a given system.

Common Performance Metrics

Table 48 (page 108) describes the metrics that are shown in common for (PD, Port, LD and VLUN) performance reports when selecting the **Versus Time** and **At Time** report types.

Table 48 Performance Metrics for PD, Port, LD and VLUN Performance Reports

Metric	Description
Read IOPS	Read operation per sec.
Write IOPS	Write operations per sec.
Total IOPS	Total (read + write) operations per sec.
Read KBytes/sec	Read bandwidth in KBytes/s.
Write KBytes/sec	Write bandwidth in KBytes/s.
Total KBytes/sec	Total (read + write) bandwidth in KBytes/s.
Read Svct ms	Average read service time in millisec.
Write Svct ms	Average write service time in millisec.

Table 48 Performance Metrics for PD, Port, LD and VLUN Performance Reports *(continued)*

Metric	Description
Total Svct ms	Average total (read + write) service time in millisec.
Read IOSz KBytes	Average size of read operations in KBytes.
Write IOSz KBytes	Average size of write operations in KBytes.
Total IOSz KBytes	Average size of read and write operations in KBytes.
Queue Length	Queue length at the sample time. Note that (unlike the other metrics above), the queue length is an instantaneous measure at the sample time, not an average over the sample interval. Please note that due to the way that RCFC ports process data, the Queue Length may not be a valid measure.
Avg Busy %	The percentage of time that the object is busy (i.e., has at least one outstanding IO operation). Note that this is not the same thing as the percentage of available object bandwidth that is being used. For example, a port shown as 50% busy may have much less than 50% of its bandwidth utilized. This is because the port may be idle for a substantial period of time between the request and response. Please note that for RCFC ports, the value will always be at 100% since the read requests are always waiting in the queue to be processed.

For **Histogram At Time** and **Histogram Versus Time** report types, the PD, Port, LD and VLUN Performance reports show the read counts, write counts or total counts in various service time buckets. Depending on the **Count** control (see [“Count” \(page 85\)](#)), these counts are either the accesses within the sample interval or the total count from the beginning (object creation or system reboot) until the current sample.

The CPU and VV Cache performance reports provide different metrics that are described separately in the following sections.

VV Cache Performance Reports

Table 49 (page 109) describes the metrics that are shown for VV Cache Performance Reports.

Table 49 VV Cache Performance Report Metrics

Metric	Description
Read Hits	Number of reads that hit in the cache.
Read Misses	Number of reads that miss in the cache.
Read Total	Total number of reads. Not shown in charts.
Write Hits	Number of writes for which the page is already in cache AND is dirty (i.e., has previously written data that has not yet been flushed to disk).
Write Misses	Number of writes that miss in the cache. A write is considered a miss if the page is not in the cache or if the page is not dirty in the cache (see above).
Write Total	Total number of writes. Not shown in charts.
Read Hit%	Percentage of reads (out of total reads) that hit in the cache.
Write Hit%	Percentage of write (out of total writes) that hit in the cache.
Total	Total number of accesses (reads + writes). Not shown in charts.

Node Cache Performance Reports

Table 50 (page 110) describes the metrics that are shown for Node Cache Performance Reports.

Table 50 Node Cache Performance Report Metrics

Metric	Description
Read Hits	Number of reads that hit in the cache.
Read Misses	Number of reads that miss in the cache.
Read Total	Total number of reads. Not shown in charts.
Write Hits	Number of writes for which the page is already in cache AND is dirty (i.e., has previously written data that has not yet been flushed to disk).
Write Misses	Number of writes that miss in the cache. A write is considered a miss if the page is not in the cache or if the page is not dirty in the cache (see above).
Write Total	Total number of writes. Not shown in charts.
Read Hit%	Percentage of reads (out of total reads) that hit in the cache.
Write Hit%	Percentage of write (out of total writes) that hit in the cache.
Total	Total number of accesses (reads + writes). Not shown in charts.
Lock Blocks	Number of locked blocks.
Free	Free space for the cache.
Clean	Number of clean cache pages (valid data on page).
Write1	Number of dirty pages that have been modified exactly 1 time. A page is dirty when it has been modified in cache but not written to disk.
WriteN	Number of dirty pages that have been modified more than 1 time.
WriteScheduled	Number of pages scheduled to be written to disk.
Writing	Number of pages being currently written by the flusher to disk.
Recov	During node down, number of dead node pages recovered that have to be written to disk.

CPU Performance Reports

Table 51 (page 110) describes the metrics that are shown for CPU Performance Reports.

Table 51 CPU Performance Report Metrics

Metric	Description
User%	Percent of CPU time in user-mode.
Sys%	Percent of CPU time in system mode.
Idle%	Percent of CPU time in idle.
Intr/s	Number of interrupts per second.
CtxtSw/s	Number of context switches per second.

Link Performance Reports

Table 52 (page 111) describes the metrics that are shown for Link Performance Reports.

Table 52 Link Performance Report Metrics

Metric	Description
Transfers/s	Transfers per second on the link.
Bandwidth	Bandwidth used on the link.
Transfer size	Average size of the transfers on the link.

Adaptive Optimization Reports

Table 53 (page 111) describes the metrics that are shown for Adaptive Optimization Reports.

Table 53 Adaptive Optimization Report Metrics

Metric	Description
IO Rate Density	Shows IO rate density for a given region.
Tiers 0-2	Shows movement on a given VV for a type of tier.
Tier movement	Space moved between tiers.

NOTE: Adaptive Optimization reports supports reporting for only one system at a time.

7 Configuring the System Reporter Policy Settings

This chapter describes how to edit HP 3PAR System Reporter policies, such as editing the sampling policies, adding or removing storage servers from the database, and scheduling reports. These actions can only be performed using a Web browser and cannot be performed using the System Reporter Excel client. The end of the chapter also describes how to edit the Apache HTTP Server configuration files, which hold the Apache HTTP Server usernames and passwords.

Accessing the System Reporter Policies

You can configure and modify policies by clicking on **Policy Settings** in the System Reporter Main Menu Window and selecting the appropriate tab from the System Reporter Policies window:

3PAR System Reporter Policies

Description	Value	Change
Full path to directory containing InServ CLI password files	C:/Program Files/3PAR/System Reporter/clipwdir	Change
Requested High-Res sampling interval in minutes	1	Change
Number of hours to keep High-Res samples	48	Change
Number of days to keep hourly samples	7	Change
Number of days to keep Adaptive Optimization hourly samples	7	Change
Number of days to keep daily samples	366	Change
Ratio of High-Res performance samples to High-Res space samples	6	Change
Perform compaction of the database after next sample	daily	Change
If there is a problem with sampling, send email to this address.		Change

The following options are available from the menu tabs:

- **Sampling Policies** – Configures report sampling policies.
- **InServ Systems** – Specifies the HP 3PAR Storage Systems to include for sampling and subsequent report generation and optional optimization.
- **Alert Rules** – Specifies alert rules for email reports.
- **Scheduled Reports** – Schedules reports for generation at a specified time.
- **Adaptive Optimization** – Configures the optionally licensed Adaptive Optimization components.

NOTE: For details on configuring Adaptive Optimization, see “Using Adaptive Optimization” (page 178).

Configuring Sampling Policies

You can use a Web browser to adjust certain HP 3PAR System Reporter database sampling policies such as the sampling interval for **high-res** samples.

Table 54 (page 112) shows a list of policies that you can edit and their default values.

Table 54 Sampling Policies

Policy	Default
Full path to directory containing InForm CLI password files	C:/Program Files/3par/ System Reporter/clipwdir (on Windows) /etc/clipwdir (on Linux)
Requested hi-res sampling interval in minutes.	5

Table 54 Sampling Policies *(continued)*

Policy	Default
Default number of hours to keep hi-res samples (in force by default for all systems unless specified otherwise on an individual basis)	24
Default number of days to keep hourly samples (in force by default for all systems unless specified otherwise on an individual basis)	7
Default number of days to keep Adaptive Optimization hourly samples (in force by default for all systems unless specified otherwise on an individual basis)	7
Default number of days to keep daily samples (in force by default for all systems unless specified otherwise on an individual basis)	366
Ratio of Hi-Res performance samples to Hi-Res space samples. Typically space changes are slower than performance so space sampling can be configured at a slower rate than performance sampling.	6
Perform compaction of the database after next sample (SQLite only. Sampling will be stopped and database will be unavailable for several minutes) see “Compacting Databases” (page 113)	daily
If there is a problem with sampling, send email to this address. Email is only sent if the sampling problem requires the sampler service/daemon to exit.	(blank)

Compacting Databases

When some sampling policies are changed or when the number or sizes of systems being sampled is reduced the size of the database may be reduced significantly.

If a MySQL or Oracle database is being used and there is a significant change in size, the database’s tables do not automatically shrink. A **compaction** must be performed on the database in order to recover the file space. To start a compaction on the database, change the policy labelled Perform compaction of the database after next sample to **now**, **daily** or to the specific day of the week depending on whether you want to perform the compaction now, immediately after each daily sample or immediately after the daily sample on the specified day of the week.

If the policy value is set to **now** the compaction is done once and the value is automatically reset to 0.

To achieve compaction with MySQL databases, System Reporter uses the OPTIMIZE command. For Oracle databases System Reporter uses the ALTER TABLE SHRINK SPACE command. Refer to your SQLite, MySQL or Oracle database documentation for further information about database compaction using these commands.

For MySQL and Oracle databases, compaction is done by an external process so that data sampling can continue in parallel with database compaction. However, for SQLite databases sampling is suspended until the compaction is complete.

For MS SQL, the compaction is performed using the following procedure:

```

DBCC SHRINKDATABASE
    ( database_name [ , target_percent ]
      [ , { NOTRUNCATE | TRUNCATEONLY } ]
    )

DBCC SHRINKFILE
    ( { file_name | file_id }
      { [ , target_size ]
        | [ , { EMPTYFILE | NOTRUNCATE | TRUNCATEONLY } ]
      }
    )

```

```
) }  
}
```

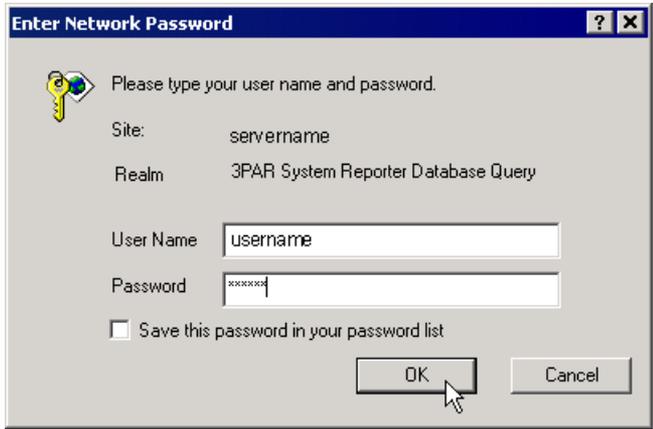
NOTE: Please note that System Reporter does not support using Microsoft SQL with Linux systems.

Editing Sampling Policies

To edit a System Reporter database sampling policy:

1. Access HP 3PAR System Reporter by pointing your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** window appears.
2. In the Extras Menu area, select **Policy Settings**.
3. If the Web server has been configured to require a password to read the database and view sampling policies, a connection dialog box appears. Type the user name and password and click **OK** (Figure 45 (page 114)).

Figure 45 Connection Dialog Box



NOTE: See for more instructions on setting Apache HTTP Server to require a username and password to read the database or view sampling policies.

4. In the **Sampling Policies** screen that appears, locate the policy to be changed and, to the right of the policy, click **Change**.
5. If the Web server has been configured to require a password to edit sampling policies, a connection dialog box appears. Type the user name and password and click **OK** (Figure 46 (page 115)).

Figure 46 Connection Dialog Box



NOTE: See for more instructions on setting Apache HTTP Server to require a username and password in order to edit sampling policies.

6. In the **New Value** box, type a new value and then click **Submit Query**.
The **Sampling Policies** window appears with the new value.
7. When finished editing policies, close the **Sampling Policies** window.

Adding, Removing and Re-Configuring HP 3PAR Storage Systems

System Reporter can sample multiple systems and place that information in the database. This section describes how to add or remove systems from the list of systems being sampled by System Reporter, and how to re-configure sampling parameters for a storage system.

- To start sampling a new storage server, add the system to the list of systems to be sampled (“Adding a Storage Server” (page 115)).
- To stop sampling a storage server, remove that system from the system table as described in “Removing a Storage Server” (page 115).
- To re-configure sampling parameters for a storage server see section “Re-Configuring a Storage Server Sampling” (page 116)

Adding a Storage Server

Refer to section for instructions on how to add a system to be sampled by System Reporter.

Verifying Addition of a Storage Server

To verify that a new storage server was successfully added to the system table refer to “Verifying Installation” (page 48).

NOTE: For performance type reports, you may only be able to see high resolution performance information at first. For space type reports, it may take more than ten minutes before you can sample high resolution data. Daily and hourly data samples may take even longer.

Removing a Storage Server

You can remove a storage server from the list of systems being sampled by System Reporter as follows.

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.

2. Click on **Policy Settings** in the Extras Menu area.
The **System Reporter Policies** window appears.
3. Choose the **InServ Systems** tab.
4. Click **Remove InServ**.
5. Choose the storage server from the drop-down list and then click **Submit Query**.
6. An alert window appears confirming that the system was removed.
7. Click **OK** to return to the **Sampling Policies** window.

NOTE: If you want to temporarily disable sampling of a storage server, you do not need to remove it from the list of systems being sampled. Instead you can disable sampling for the storage server by checking the Disable Sampling parameter for the storage server as described in “Re-Configuring a Storage Server Sampling” (page 116).

NOTE: Removing a storage server from the database table stops data collection for that storage server but does not remove the database or remove HP 3PAR System Reporter or any of its components. See for deinstallation instructions.

Re-Configuring a Storage Server Sampling

You can re-configure sampling parameters for a storage server being sampled by System Reporter as follows.

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click on **Policy Settings** in the Extras Menu area.
The **System Reporter Policies** window appears.
3. Choose the **InServ Systems** tab.
4. Locate the storage server in the table of HP 3PAR storage systems being sampled. Click on the **Change** link for that storage server. A web form titled **Change InServ Sampling Parameters** appears with the following options:
 - **IP Name or Address of system** – Shows the IP address for the selected HP 3PAR Storage System.
 - **Use SSL** – Specifies whether to use SSL for secured connections.
 - **Skip LD Performance Data** – Specifies whether to skip LD performance data.
 - **Sample Adaptive Optimization Data** – Indicates whether to sample Adaptive Optimization data.
 - **Hours to keep High Res Data** – Specifies the amount of time (in hours) to keep High Res data.
 - **Days to keep Hourly Data** – Specifies the number of days to retain hourly data.
 - **Days to keep Daily Data** – Specifies the number of days to retain daily data.
 - **Days to keep Adaptive Optimization Data** – Specifies the number of days to retain Adaptive Optimization data.
5. Check or clear the boxes for the sampling parameters.
6. Click **Submit Query**.

Configuring Rules for Email Alerts

System Reporter can be configured to evaluate numerous rules and generate email alerts to specified recipients for each rule whose conditions are met.

Adding an Alert Rule

To add an alert rule:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click on **Policy Settings** in the Extras Menu area.
The 3PAR System Reporter Policies window appears.
3. Choose the **Alert Rules** tab.
4. Click **Add Alert**.
5. Choose the **Data Table** to which the rule applies from the drop-down list.
6. Choose the **Resolution** of the samples to which the rule applies from the drop-down list. The rule will be evaluated for each sample of the chosen resolution.
7. Choose the **System** to which the rule applies from the drop-down list. Leave this blank if you wish the rule to apply to all systems.
8. Choose the Metric that the rule should calculate from the down-list. The available metrics depend on the chosen Data Table and changing the Data Table will reset the selected Metric.
9. Choose the **Direction** that determines how the Metric is compared to the Limit Value from the drop-down list. The available values are > (greater than) and < (less than).
10. Enter the **Limit Value** as a number. The metric is compared against this number.
11. Enter the **Limit Count** as an integer (zero or larger). For each sample interval, an alert email will only be generated if the Metric exceeds the Limit Value (as compared by Direction) for more than Limit Count objects.
12. Enter the **Condition** (min_read_iops, min_write_iops, min_total_iops) to indicate the type of condition that is to be monitored.
13. Enter the **Condition Value** to specify the minimum amount that is to be met for the associated condition.
14. Enter the **Recipient** email address to whom the alert for this rule should be sent.
15. Click **Submit Query**.
16. An alert window appears confirming that the alert rule was added.
17. Click **OK** to return to the **Sampling Policies** window.

Changing an Alert Rule

To change an alert rule:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click on **Policy Settings** in the Extras Menu area.
The **3PAR System Reporter Policies** window appears.
3. Choose the **Alert Rules** tab.
4. Click **Change** next to the existing alert that you want to modify.
The **Change Alert Rule** menu appears.
5. Make the necessary modifications.
6. Click **Submit Query**.
7. An alert window appears confirming that the alert rule has been changed.
8. Click **OK** to return to the **Sampling Policies** window.

Removing an Alert Rule

To remove an alert rule:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click on **Policy Settings** in the Extras Menu area.
The **3PAR System Reporter Policies** window appears.
3. Choose the **Alert Rules** tab.
4. Click **Remove Alert**.
5. Select the alert rule to remove from the drop-down list.
6. Click **Submit Query**.
7. An alert window appears confirming that the alert rule was removed.
8. Click **OK** to return to the **Sampling Policies** window.

Alert Rule Parameters

Alert rules are specified using parameters described in the following sections.

Data Table

The data table on which the rule operates. The currently supported data tables are described in [Table 55 \(page 118\)](#).

Table 55 Data Table

Data Table	Description
statport	Port performance table.
statport_8g statport_4g statport_2g statport_1g	Represents the port bandwidth utilized by the 8g port. Represents the port bandwidth utilized by the 4g port. Represents the port bandwidth utilized by the 2g port. Represents the port bandwidth utilized by the 1g port.
statport_nodesum	Represents the sum of all the ports on a node. This allows you to generate alerts when nearing the limits of IOPs/s and bandwidth of a node.
statport_disknodesum	Represents the sum of all the disk ports on a node. This allows you to generate an alert when nearing the limits of back-end (disk) IOPs/s or bandwidth for the node.
statport_hostnodesum	Represents the sum of all the ports on a node by host. This allows you to generate alerts when nearing the limits of IOPs/s and bandwidth of a node by host.
statport_syssum	Represents the sum of all the ports on an HP 3PAR Storage System. This allows you to generate alerts when nearing the limits of IOPs/s and bandwidth of the system.
statpd	Physical Disk performance table.
statpd_nl7	Restricts statpd to only the NL 7.2K RPM PDs.
statpd_fc10	Restricts statpd to only the FC 10K RPM PDs.
statpd_fc15	Restricts statpd to only the FC 15K RPM PDs.
statpd_ssd150	Restricts statpd to only the SSD PDs with speed 150.
statpd_sum	Represents the aggregation of all the PDs in the HP 3PAR Storage System.
statpd_nl7sum	Represents the aggregation of all the NL 7.2K RPM PDs in the HP 3PAR Storage System.
statpd_fc10sum	Represents the aggregation of all the FC 10K RPM PDs in the HP 3PAR Storage System.
statpd_fc15sum	Represents the aggregation of all the FC 15K RPM PDs in the HP 3PAR Storage System.

Table 55 Data Table *(continued)*

Data Table	Description
statpd_ssd150sum	Represents the aggregation of all the SSD PDs of speed 150 in the HP 3PAR Storage System.
statld	Logical Disk performance table.
statvlun	VLUN performance table.
statcpu	CPU performance table.

Resolution

The resolution of data table on which the rule operates. This also implicitly determines how often the rule is evaluated because the rule is evaluated after each sample at the specified resolution.

The available resolutions are described in [Table 56 \(page 119\)](#).

Table 56 Resolution

Resolution	Description
hires	Data table that contains the high resolution samples.
hourly	Data table that contains the hourly samples
daily	Data table that contains the daily samples

System

If **System** is specified, the alert rule is calculated only for objects in the specified system. If **System** is left blank the alert rule is calculated for all systems.

Metric

The metric that is calculated for the given data table and resolution. The metric is calculated per relevant object (not the aggregate of the objects) in the data table (ports for **statport**, PDs for **statpd**, LDs for **statld**, VLUNs for **statvlun** and nodes for **statcpu**). The supported metrics depend on the value of **Data Table**. Currently supported metrics are described in [Table 57 \(page 119\)](#).

Table 57 Metric

Metrics	Description	Data Tables
read_iops	Read IO operations/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum, statpd_ssd150sum, statld, statvlun
write_iops	Write IO operations/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum, statpd_ssd150sum, statld, statvlun
total_iops	Total IO (read + write) operations/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum,

Table 57 Metric (continued)

Metrics	Description	Data Tables
		statpd_ssd150sum, statld, statvlun
read_kbps	Read KBytes/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum, statpd_ssd150sum, statld, statvlun
write_kbps	Write KBytes/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum, statpd_ssd150sum, statld, statvlun
total_kbps	Total (read + write) KBytes/sec	statport, statport_nodesum, statport_syssum, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statpd_sum, statpd_nl7sum, statpd_fc10sum, statpd_fc15sum, statpd_ssd150sum, statld, statvlun
read_svctms	Read service time in millisecond averaged over the sample	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
write_svctms	Write service time in millisecond	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
total_svctms	Total (read + write) service time in millisecond	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
read_ioszkb	Read IO Size in KBytes	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
write_ioszkb	Write IO Size in KBytes	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
total_ioszkb	Total (read + write) IO Size in KBytes	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun
total_qlen	Total queue length	statport, statpd, statpd_nl7, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun

Table 57 Metric (continued)

Metrics	Description	Data Tables
idle_perc	Percentage idle time for the node	statcpu
total_utilperc	Total utilization per second. For instance, when the PD utilization exceeds a given high (say 90%) utilization meaning that 90% of the time there is some IO operation outstanding at the PD.	statport, statpd, statpd_n17, statpd_fc10, statpd_fc15, statpd_ssd150, statld, statvlun

Direction

Specifies whether the **Metric** should be less than (<) or greater than (>) the **Limit Value** for the alert to be generated.

Limit Value

Specifies the value that the **Metric** is compared to. This should be a number.

Limit Count

The alert will only be generated if the number of objects for which the **Metric** exceeds the **Limit Value** (in the direction specified by Direction) is greater than Limit Count in any given sample. This should be an integer.

Condition

Specifies the condition (min_read_iops, min_write_iops, min_total_iops) that should be monitored.

Condition Value

Specifies the condition value.

Recipient

Email address to which the alert email should be sent.

Example of an Alert Rule

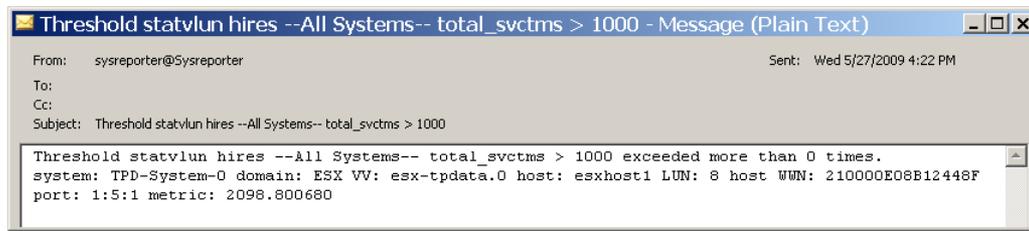
Figure 47 (page 121) shows an example of a configured alert. In this example, System Reporter will send an email alert to email@address each time it inserts a high-resolution sample in which more than 0 (Limit Count) VLUNs (including VLUNs in all systems) has a total (reads and writes included) average service time in millisecs of greater than (>) 1000 (Limit Value).

Figure 47 Example of an alert rule

Alert Rules											
Alert ID	Data Table	Resolution	System	Metric	Direction	Limit Value	Limit Count	Condition	Condition Value	Recipient	Change
1	statport	hires	--All Systems-	total_svctms	>	1000	0	min_read_iops	300	email@address	Change
Add Alert											
Remove Alert											

An example of the email alert is shown Figure 48 (page 122):

Figure 48 Example of an alert email

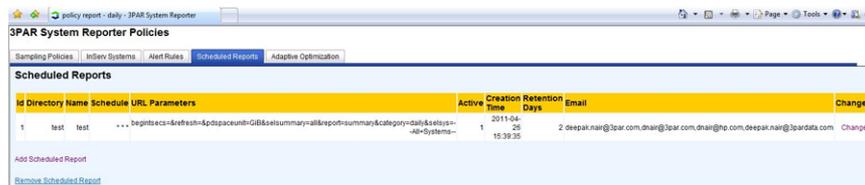


The email shows that a single VLUN instance had an average total service time of 2098.800680 millisecond during a high-resolution sampling interval.

Scheduling Reports

You can have a predefined report run according to a predetermined schedule by choosing the **Scheduled Reports** tab from the System Reporter Policies window (Figure 49 (page 122)).

Figure 49 Scheduled Reports Window



Scheduled reports offer the following advantages:

- Allows quick access to pre-configured reports on a scheduled basis instead of keeping active tabs in a web browser.
- Allows the generation of reports to take place in the background at off-peak times to minimize impact on system performance.
- Allows distribution of scheduled reports to users from a selected directory without giving access to policy configuration or system information outside a particular users authority.

When you schedule a report, a subdirectory is created with the specified name in the "scheduledreports" directory according to the following structure:

```
scheduledreports/<report directory>/<report name>/<YYYYMMDDHH>
```

Each time the report runs, a new subdirectory is created for that instance of the report with the timestamp as the name. All of the PNG image files, the .CSV file and the .html file will be placed in that subdirectory.

You have the choice of making the scheduled reports accessible through the System Reporter Main Menu Window or you can copy the directory structure to some other location with whatever permissions are appropriate for the users who access the reports. The benefit of a having a report directory structure is that you can limit the users who have access depending on the permissions that are assigned (for example by creating `.htaccess` files). One example where this may be useful is when multiple departments share an array. You can schedule various reports specific to each department, place them in different report directories, and then allow each department access only to their respective report directory. By default, the scheduled report process can run 10 reports at a time. This can be re-configured if required.

As another option while scheduling a report, you can have an email sent to a specified address that provides a link when a scheduled report is generated.

NOTE: When making scheduled reports available, the entire report directory should be published along with all the associated sub directories.

Adding a Scheduled Report

To add a scheduled report:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click **Policy Settings** in the Extras Menu area.
The **System Reporter Policies** window appears.
3. Choose the **Scheduled Reports tab**.
The **Scheduled Reports** window appears.
4. Click **Add Scheduled Report**.
5. Specify the report directory.
This is the name of the directory under the top-level directory for scheduled reports (`htdocs/3par/scheduledreports/`). You can specify multiple reports within a report directory and each of these reports will have its own sub-directory.
6. Specify the report name.
7. Specify the report schedule. Hold the **CTRL** key to pick multiple dates and times. The minimum possible granularity of a scheduled report is one hour.
8. Specify the report URL parameters (you can copy these from the address bar of the report that define the report). Do not include the `http://` portion of the URL. The maximum size of a URL should not exceed 4000 characters.
9. Choose whether or not the report should be currently active or not (`true/false`). This option allows you to turn the report generation on and off as needed.
10. Specify the number of days to retain a scheduled report. System Reporter will automatically delete reports that are older than the specified number of days. The default is seven days.
11. If appropriate, enter an email address specifying where a link should be sent when the scheduled report is generated.
12. When done, click **Submit Query**. An alert window appears confirming that the report has been scheduled.
13. Click **OK** to return to the **Sampling Policies** window.

Changing a Scheduled Report

To change a scheduled report:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click **Policy Settings** in the Extras Menu area.
The **System Reporter Policies** window appears.
3. Choose the **Scheduled Reports tab**.
The **Scheduled Reports** window appears.
4. Click **Change** next to the scheduled report that you want to change.
The **Change Scheduled Report** window appears.
5. Change the appropriate settings.
6. When done, click **Submit Query**. An alert window appears confirming that the scheduled report has been changed.
7. Click **OK** to return to the **Sampling Policies** window.

Removing a Scheduled Report

To remove a scheduled report:

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.
The **3PAR System Reporter** main window appears.
2. Click **Policy Settings** in the Extras Menu area.
The **System Reporter Policies** window appears.
3. Choose the **Scheduled Reports tab**.
The **Scheduled Reports** window appears.
4. Click **Remove Scheduled Report**.
The Select Scheduled Report to Remove window appears.
5. Choose the scheduled report to remove from the pull down list.
6. When done, click **Submit Query**. An alert window appears confirming that the scheduled report has been deleted.
7. Click **OK** to return to the **Sampling Policies** window.

Configuring Adaptive Optimization

An optionally licensed Adaptive Optimization component is included with System Reporter that can be used to monitor and optimize storage utilization.

NOTE: See for complete details on how to configure and use Adaptive Optimization.

About File Permissions

The files installed or created by the HP 3PAR System Reporter installer have default permissions that allow any user to read file contents. Use the appropriate Windows or Linux security mechanisms to limit file access as desired. The minimum level for correct operation is READ access by the SYSTEM account (the Apache HTTP Server and HP 3PAR System Reporter services run as the SYSTEM account).

For Windows Systems

If System Reporter has been configured to use MySQL, Microsoft SQL or Oracle, the following files (shown in their default locations) contain password information that you may wish to protect:

```
C:\Program Files\3par\System Reporter\dbpwfile
```

```
C:\Program Files\Apache Group\Apache2\cgi-bin\3par-policy\config.tcl
```

```
C:\Program Files\Apache Group\Apache2\cgi-bin\3par-rpts\config.tcl
```

Similarly, if configured to generate email alerts with an SMTP user and password, the following file contains the SMTP user name and password:

```
C:\Program Files\3par\System Reporter\sampleloop_config.tcl
```

The SMTP location for sending email for scheduled reports is in:

```
C:\Program Files\Apache Group\Apache2\cgi-bin\3par-rpts\config.tcl
```

For Linux Systems

If System Reporter has been configured to use MySQL, Microsoft SQL or Oracle, you can access the following files (shown in their default locations) that contain password information you may wish to protect:

```
ls -ltra /var/www/cgi-bin/3par-rpts
```

```
-r--r--r-- 1 apache apache 1148 Sep 22 2010 config.tcl
```

```
drwxr-xr-x 2 apache apache 4096 Jan 21 17:33 .
```

```
-rwxr-xr-x 1 apache apache 6717686 Jan 21 23:49 inserv_perf.exe
-rwxr-xr-x 1 apache apache 1211700 Jan 21 23:49 chart.exe
ls -ltra /var/www/cgi-bin/3par-policy
-r--r--r-- 1 apache apache 1208 Sep 22 2010 config.tcl
drwxr-xr-x 3 apache apache 4096 Jan 21 17:31 .
```

```
-rwxr-xr-x 1 apache apache 6683099 Jan 21 23:49 inserv_policy.exe
```

You must make sure that the files are owned by Apache. If necessary, you can use the `chown` command to change permissions for the files.

The SMTP location for sending email for scheduled reports is in:

```
/var/www/cgi-bin/3par-rpts/config.tcl
```

Modifying the Apache HTTP Server Configuration Files

You must modify the Apache HTTP Server `config.tcl` files located in the `3par-policy` and `3par-rpts` subdirectories whenever you change MySQL user names and passwords.

1. Go to the Apache HTTP server `cgi-bin` directory and locate the `3par-policy` subdirectory (for example: `C:\Program Files\ApacheGroup\Apache2\cgi-bin\3par-policy`).
2. In the `3par-policy` subdirectory, locate the `config.tcl` file (for example, `C:\Program Files\Apache Software Foundation\Apache2.2\cgi-bin\3par-policy\config.tcl`) and open with a text editor.
3. At the bottom of the file, alter the appropriate fields as necessary.

```
# Select the type of database (sqlite, mysql, oracle or mssql):
set Sysdb::dbtype <oracle,mysql,mssql,sqlite>

set Sysdb::dbhost <TNS Names in oracle,hostname in mysql,dsn name in mssql>
# Enter database name
set Sysdb::dbname <Schema name in oracle,db name in mysql and username in MSSQL>
# Enter database user name and password
set Sysdb::dbuser cliuser
set Sysdb::dbpasswd "cliuser"
```

4. Save the changes and close the file.
5. Go to the Apache HTTP server `cgi-bin` directory and locate the `3par-rpts` subdirectory (for example, (for example, `C:\Program Files\ApacheGroup\Apache2\cgi-bin\3par-rpts`).
6. In the `3par-rpts` subdirectory, locate the `config.tcl` file (for example, `C:\Program Files\Apache Software Foundation\Apache2.2\cgi-bin\3par-rpts\config.tcl`) and open with a text editor.
7. At the bottom of the file, alter the appropriate fields as necessary.

```
# Select the type of database (sqlite, mysql, oracle or mssql):
set Sysdb::dbtype oracle
set Sysdb::dbhost istat
# Enter database name
set Sysdb::dbname <Schema name in oracle,db name in mysql and user name in MSSQL>
# Enter database user name and password
set Sysdb::dbuser webuser
set Sysdb::dbpasswd "webuser"
```

8 Web Queries

This chapter documents the queries that the HP 3PAR System Reporter interfaces make to the System Reporter Web server.

About Web Queries

System Reporter Web queries have the following basic format:

```
http://<server_name>/cgi-bin/3par-rpts/  
inserv_perf.exe?report=<report_name>[&<option>...]
```

where:

- <server_name> is the Web server where Apache HTTP Server and the System Reporter Web scripts are installed.
- <report_name> is the name of a report described in this chapter.
- <option>... represents one or more general or report-specific options that are described in this chapter.

NOTE: In the above example and elsewhere in this chapter, the following typographical conventions are used:

- variables are denoted by angle brackets (for example, <server_name>)
- braces denote optional items (for example [&<option>])
- ellipses denote that the preceding component may be repeated (for example [&<option>...]).

The following example shows a query for Web server named localhost:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=pd_space_time&charttab=table&category=hires
```

This query requests

- a report showing physical disk space usage over time (pd_space_time).
- results posted in the form of a table (charttab=table).
- results pulled from the high-resolution database table (category=hires).

Because a specific time range is not specified, the query will default to all times present in the specified database table (hires).

NOTE: To learn more about Web queries, you can perform any System Reporter Web query by using a browser (see “Getting Started” (page 52)) and then selecting the **New window** option and then clicking **Submit Query** to open the report in a new window. In the new window, the address bar displays the Web query that you created using the browser interface.

Data Format Version Report

If the data returned by web queries is going to be processed, it is important to check that the format of the data is what is expected by the scripts or code processing it. To determine if your processing code is compatible with the server, issue the following query:

```
http://<server_name>/cgi-bin/3par-rpts/inserv_perf.exe?report=version
```

where <server_name> is the name of your server. The following is an example of the table returned:

Figure 50 System Reporter Version Report

System Reporter Version

Category	Major	Minor
Product	2	6
ReportData	4	0

This query returns a 3x3 table with the column headers **Category**, **Major**, and **Minor**. The two rows after the header contain major and minor numbers for **Product** and **ReportData**.

The **Product** version numbers are informational only and are the numbers assigned to the product when it is released. The **Product** version number is not necessarily related to the **ReportData** version. The **Product** version of the server will be the same as is displayed in the upper frame of the Web client.

The **ReportData** version numbers correspond to the format of the data as provided by the server. These numbers will change according to the following rules:

- Bug fixes do not cause a change in version number. If data in a table is being incorrectly reported, fixing the server to report it correctly does not change the version.
- Adding one or more new columns to a table without changing any of the existing columns causes the minor number to be incremented.
- Removing a column causes the major number to be incremented.
- Changing the meaning of a column causes the major number to be incremented.

Clients can accept lower version numbers if the client has been coded to work with previous versions. Clients can accept newer minor numbers when the major number is the same since the data format is compatible, but may wish to warn the user that not all features of the server's data will be represented. The client should not attempt to use data from a server with a newer major number. The simplest approach is for the client to reject a server unless the major and minor number of the server's data format is the same as the data format specification used to code the client.

The definitions of queries and the data they return given below are version 4.1. A client coded to these definitions would compare major number 4 and minor number 1 against the server's Report Data major and minor numbers.

About Reports and Options

There are five categories of reports that you can generate using System Reporter:

- summary reports (see ["Summary Report" \(page 153\)](#))
- space reports (see ["Space Reports" \(page 153\)](#))
- performance reports (see ["Performance Reports" \(page 155\)](#))
- adaptive optimization reports (see ["Adaptive Optimization Reports" \(page 160\)](#))
- objects present reports (see ["Objects Present Reports" \(page 167\)](#))

Reports are generated by sending a Web query to the System Reporter database. Each query must specify a report type from one of the three categories listed above. These report types are discussed in further detail in ["Report Types" \(page 153\)](#).

Each report type also has options that you can use to customize the report results. Not all options can be used with all report types. Refer to ["Report Options" \(page 153\)](#) to learn which options are compatible with which report types.

Report Options

Options are included in the query URL by using the following syntax:

```
<option>=<value>
```

where `<option>` and `<value>` are the options and values described below.

category

Use the `category` option to specify which database table to use when pulling results. This option can only be specified once in a given query string.

NOTE: The System Reporter maintains three database tables that each offer a different granularity of resolution and stored time range. See also [“Sample Resolution” \(page 67\)](#).

Specify one of the following:

`daily`

(Default). Specifies that results will be pulled from the daily database table. For the daily database table, the storage server is sampled once daily, shortly after midnight.

`hourly`

Specifies that results will be pulled from the hourly database table. For the hourly database table, the storage server is sampled once an hour.

`hires`

Specifies that results will be pulled from the hires database table. For the hires database table, the storage server is sampled every 5 minutes (by default) or less frequently, depending on how the database sampling policy has been configured (see [“Editing Sampling Policies” \(page 114\)](#)).

example:

The following example specifies that query results are to be pulled from the hourly database table:

```
&category=hourly
```

compare

The `compare` option allows you to compare the performance of multiple (up to 16) systems or system objects on one report. Use the `compare` option to specify an object as described in [“values:” \(page 128\)](#) below and compare that specified object against objects of the same type. This option can only be specified once in a given query string.

values:

Specify one of the following:

`none`

(Default) No comparison is performed.

`System`

Compares the set of systems specified with the query or all systems if none are specified.

`Domain`

Compares the set of domains specified with the query or all domains if none are specified.

`PortType`

Compares selected port types. This value is only available for the `port_perf_time` report.

DiskType

Compares selected disk types. This value is only available for the `pd_perf_time`, `ld_space_time` and `pd_space_time` reports.

DiskSpeed

Compares selected disk types. This value is only available for the `pd_perf_time` and `pd_space_time` reports.

RaidType

Compares selected RAID types. This value is only available for the `ld_space_time` report.

n:s:p

Compares selected port positions. For this value, an escape sequence may be required to retain spacing and special characters (for example: `n%3As%3Ap`).

LD

Compares selected logical disks. This value is only available for the `ld_perf_time` and `ld_space_time` reports.

PDID

Compares selected physical disk IDs. This value is only available for the `pd_perf_time` and `pd_space_time` reports.

Host

Compares selected hosts. This value is only available for the `vlun_perf_time` report.

VV

Compares selected volumes. This value is only available for the `vlun_perf_time` and `vv_space_time` reports.

VVType

Compares selected VV Type. This value is only available for the `vv_space_time` report.

ProvType

Compares selected VV Provisioning Type. This value is only available for the `vv_space_time` report.

LDUsage

Compares selected LD Usage. This value is only available for the `ld_space_time` report.

CPG

Compares selected CPG. This value is only available for the `ld_space_time` and `ld_perf_time` reports.

UsrCPG

Compares selected Usr CPG. This value is only available for the `vv_space_time` report.

Snpcpg

Compares selected Snp CPG. This value is only available for the `vv_space_time` report.

Node

Compares selected controller nodes. This value is only available for the `cpu_perf_time` report.

CageID

Compares selected drive cages. This value is only available for the `pd_space_time` report.

Mag

Compares selected drive magazine. This value is only available for the `pd_space_time` report.

Disk

Compares selected disk magazine position. This value is only available for the `pd_space_time` report.

example:

The following example specifies that all logical disks are compared:

```
&compare=LD
```

comparesel

When the compare option (see “compare” (page 128)) is used, up to 16 objects can be compared. However, if there are more than 16 objects selected, then the comparesel option is used to select a metric. The 16 objects with the highest peak values of the metric are chosen to be compared.

values:

Specify one of the following to choose the peak metric:

`total_iops`

(Default) The total (read + write) IOPs for the object.

`total_kbps`

The total (read + write) bandwidth (KBytes/s) for the object.

`total_svctms`

The total (read + write) service time in millisec for the object.

`total_qlen`

The total queue length for the object.

`busy_pct`

The percentage of time that the object is busy (i.e., has at least one outstanding operation).

`cpu_busy`

The percentage of CPU busy time.

`interrupts`

The CPU interrupts per second.

`ctxtsws`

The CPU context switches per second.

`total_space`

The total selected space.

`raw_space`

The raw space (including RAID overhead).

`used_space`

The used space.

`free_space`

The free space.

`TotalRsvdSize`

The total reserved size.

`TotalRawSize`

The total raw (including RAID overhead) size.

`VirtSize`

The virtual size.

table

Use the `table` option to specify a table from which data is returned when using one of the `objects_present` reports (see “[Objects Present Reports](#)” (page 167)). This option can only be specified once in a given query string.

values:

Specify one of the following:

`ldspace`

Returns values from the `ldspace` database tables.

`pdspace`

Returns values from the `pdspace` database tables.

`vvspace`

Returns values from the `vvspace` database tables.

`statpd`

Returns values from the `statpd` database tables.

`statld`

Returns values from the `statld` database tables.

`statvlun`

Returns values from the `statvlun` database tables.

`statvvcmp`

Returns values from the `statvvcmp` database tables.

`statcmp`

Returns values from the `statcmp` database tables.

`statport`

Returns values from the `statport` database tables.

`statcpu`

Returns values from the `statcpu` database tables.

`statlink`

Returns values from the `statlink` database tables.

example:

The following example returns values for the `pdspace` database tables, which are relevant to physical disk space reports (see “[Space Reports](#)” (page 153)):

```
table=pdspace
```

refresh

Use the `refresh` option to have the web browser automatically reload the report with the same options after the specified number of minutes. This option can only be specified once in a given query string.

values:

Specify an integer number of minutes. Any other value will be ignored.

example:

The following example specifies that the report is to be reloaded every 2 minutes:

```
&refresh=2
```

charttab

Use the `charttab` option to specify whether to post the results in tabular format, as charts, or both. The number of charts and/or tables used to post the results depend on the report type. This option can only be specified once in a given query string.

values:

Specify one of the following:

`chart`

Specifies that results will be displayed in the form of one or more charts. When using `charttab=chart`, you can also use the `charttype` option to specify a chart type (see “[charttype](#)” (page 132)).

`table`

Specifies that results will be displayed in tabular format.

`both`

(Default). Specifies that results will be displayed in the form of one or more charts, followed by one or more tables.

example:

The following example specifies that query results are to be displayed in the form of one or more charts:

```
&charttab=chart
```

charttype

Use the `charttype` option in conjunction with `charttab=chart` or `charttab=both` to specify the chart type to be used when results are to be displayed in chart form. This option can only be specified once in a given query string.

NOTE: Since the total height of the stacked charts (`stackedbars/stackedarea`) is the sum of the individual segments, they are not suitable for some forms of data. For example, in performance charts, the average service time for I/O is not the sum of the average service time for reads plus the average service time for writes. Similarly, the average size for I/O is not the sum of the average I/O size of reads plus the average I/O size of writes. Therefore, when you select stacked charts, only the total service times and total I/O sizes are shown. If you want to see the service times and I/O sizes for reads and writes separately, use `bars` or `lines`.

values:

Specify one of the following:

`stackedbars`

Specifies that charts will use stacked bars. Stacked charts stack data on top of each other so you can see the total as well. See [Figure 34 \(page 93\)](#) for an example of a chart using stacked bars.

`stackedvbars`

Specifies that charts will use stacked vertical bars. Stacked charts stack data on top of each other so you can see the total as well. See [Figure 35 \(page 94\)](#) for an example of a chart using stacked vertical bars.

stackedarea

Specifies that charts will use stacked areas instead of stacked bars. See [Figure 36 \(page 94\)](#) for an example of a chart using stacked areas.

bars

(Default). Specifies that bar charts will be used. See [Figure 38 \(page 95\)](#) for an example of a bar chart.

vbars

Specifies that vertical bar charts will be used. See [Figure 39 \(page 96\)](#) for an example of a bar chart

lines

Specifies that line graph charts will be used. See [Figure 37 \(page 95\)](#) for an example of a chart using lines.

example:

The following example specifies that query results are to be displayed in the form of one or more charts that use lines:

```
&charttab=chart&charttype=lines
```

graphx

Use the `graphx` option to specify the length of the chart X (category) axis in pixels. If not defined, the length is computed automatically. If too few pixels are specified, the chart may not be drawn. This option can only be specified once in a given query string.

value:

Specify an integer number of pixels.

example:

The following example specifies that the chart X axis should be 2000 pixels.

```
charttab=chart&charttype=lines&graphx=2000
```

graphy

Use the `graphy` option to specify the length of the chart Y (value) axis in pixels. If not defined, the length is computed automatically. If too few pixels are specified, the chart may not be drawn. This option can only be specified once in a given query string.

value:

Specify an integer number of pixels.

example:

The following example specifies that the chart Y axis should be 300 pixels.

```
charttab=chart&charttype=lines&graphy=300
```

graphlegpos

Use the `graphlegpos` option to specify the position of the chart legend. If not defined, the legend position is selected automatically based on the `charttype`. This option can only be specified once in a given query string.

value:

`bottom`

The legend is placed below the chart.

`right`

The legend is placed to the right of the chart.

example:

The following example specifies that the chart legend should be to the right of the chart.

```
charttab=chart&graphlegpos=right
```

tableformat

Use the `tableformat` option to specify a table format when results are to be displayed in the form of a table. If not defined, the table is formatted for viewing in a browser. This option can only be specified once in a given query string.

value:

`excel`

This option causes the results in the left column of the table to be preceded by an apostrophe so that Excel will not perceive values as numbers. See [“Objects Present Reports” \(page 167\)](#) for additional details.

example:

The following example specifies that results in the left column of the table are to be preceded by an apostrophe so that Excel will not perceive the values as a numbers:

```
tableformat=excel
```

alllabels

The `alllabels` option selects whether or not each point on the time axis of a chart is labelled.

values:

Specify one of the following:

`1`

All points on the time axis will be labelled with the time. The time axis may be longer to ensure that spacing between samples is sufficient for readable labels.

`0`

(Default). Time axis tables will be spaced at regular intervals and the spacing between samples will be reduced to make the time shorter.

example:

The following example specifies that all the time axis points should be labelled:

```
&alllabels=1
```

timeform

Use the `timeform` option to specify a time format (either `relative` or `absolute`) when using a report that requires a time range. This option can only be specified once in a given query string.

values:

Specify one of the following:

`relative`

Times used in the table(s) and/or chart(s) will be expressed relative to the present and expressed as using the format `<xx>d <hh>:<mm>:<ss>`, where this string expresses the number of days, hours, minutes and seconds ago.

`absolute`

(Default). Times used in the table(s) and/or chart(s) will be expressed in the following format:

`yyyy/mm/dd hh:mm:ss`.

`full`

The times used in the table(s) and/or chart(s) will be expressed in the following format: `yyy/mm/dd hh:mm:ss`.

`auto`

The times used in the table(s) and/or chart(s) will be expressed depending on the resolution that is selected.

example:

The following example specifies that the time formats should be expressed in terms relative to the present point in time:

```
&timeform=relative
```

begintsecs

Use the `begintsecs` option to specify a start time when using a report that requires a time range, or to indicate a specific point in time when using other report types.

If this option is not defined,

- for reports that require a time range, the earliest available time is applied.
- for reports that only require a `begintsecs` time, the most recent available time is applied.
- for `objects_present` reports (see [“Objects Present Reports” \(page 167\)](#)), the earliest available time is applied.

NOTE: Available times may vary according to the database table being used (specified with the `category` option, as described in [“category” \(page 128\)](#)).

This option can only be specified once in a given query string.

value:

`<tsecs>`

If `<tsecs>` is non-negative it represents a point in time, expressed as the number of seconds since January 1, 1970, GMT. For example, to express 2005/07/13 11:00:34 as a `<tsecs>` value, use the value 1121277634.

If `<tsecs>` is negative, it represents a point in time `<tsecs>` seconds before the current time.

NOTE: You can use the `sample_times` report (see “[sample_times](#)” (page 177)`sample_times` on page 8.78) to see `tsecs` values paired with times expressed using an easier-to-read format.

example:

The following example specifies a time of 11:00:34 on July 13, 2005:

```
&begintsecs=1121277634
```

endtsecs

Use the `endtsecs` option to specify an end time when using a report that requires a time range. Used in conjunction with the `begintsecs` option (“[begintsecs](#)” (page 135)). If not defined, then the most recent time is used. This option can only be specified once in a given query string.

value:

`<tsecs>`

If `<tsecs>` is non-negative it represents a point in time, expressed as the number of seconds since January 1, 1970, GMT. For example, to express 2005/07/13 11:00:34 as a `<tsecs>` value, use the value 1121277634.

If `<tsecs>` is negative, it represents a point in time `<tsecs>` seconds before the current time.

NOTE: You can use the `sample_times` report (see “[sample_times](#)” (page 177)`sample_times` on page 8.78) to see `tsecs` values paired with times expressed using an easier-to-read format.

example:

The following example specifies an endrange time of 11:00:34 on July 13, 2005:

```
&endtsecs=1121277634
```

groupby

Use the `groupby` option to select what is charted on the category axis (X-axis). The default for `groupby` is `System`, which means that you will see a bar (when using a `barchart`) for each system. The value is the total of the metric for the entire system. The `groupby` option can be used multiple times in the same query string, and when specified multiple times functions as an `OR` operation. If an object name exists on more than one system and if `System` is not selected as a `groupby` parameter, the resulting metric for that object will be the aggregate of the metric across all the selected systems that include the object name. This can be useful, for example if a host is connected to multiple systems and that host is assigned the same name on all those systems. In that case, selecting `Host` as the `groupby` parameter will show that hosts total performance across all systems.

values:

Specify one of the following:

`System`

(Default). When using `groupby=System`, you will not be able to see the breakdown of the metric for individual components (for example, volumes or hosts) in the system.

Domain

Group results by domain.

Host

Group results by hosts.

CPG

Group results by Common Provisioning Group (CPG) name.

UsrCPG

Group results by Common Provisioning Group (CPG) name for the User Space of VVs.

SnpCPG

Group results by Common Provisioning Group (CPG) name for the Snap Space of VVs.

VV

Group results by Virtual Volume (VV) name.

VVType

Group results by Virtual Volume (VV) Type.

ProvType

Group results by virtual volume (VV) provisioning type.

LD

Group results by Logical Disk (LD) name.

n:s:p

Group results by port. For this value, an escape sequence may be required to retain spacing and special characters (for example: n%3As%3Ap).

PortNode

Group results by port node number (The n part in n:s:p).

RaidType

Group results by RAID level (for example, RAID 0, RAID1, RAID 5 and RAID 6).

LDUsage

Group results by logical disk usage. See [Table 28 \(page 88\)](#) for a list of these usage types.

DiskType

Group results by physical disk type.

DiskSpeed

Group results by physical disk speed.

CageID

Group results by drive cage ID number.

Mag

Group results by magazine position within the drive cage (0 through 7).

Disk

Group results by disk position on a magazine (0 through 3).

PDID

Group results by physical disk ID number.

Node

Group results by controller node number (0-7, depending on the storage server model and configuration).

PortType

Group results by port type.

example:

The following example specifies that results should be grouped by system:

```
&groupby=System
```

The following example specifies that results should be grouped by system, drive magazine, and physical disk ID:

```
&groupby=System&groupby=Mag&groupby=PDID
```

orderby

Use the `orderby` option to determine how to order results. The `orderby` option can be used multiple times in the same query string. When multiple `orderby` options are used the order of the objects specified with the `orderby` option is relevant. The leftmost `orderby` is the primary sort key, the next `orderby` is the secondary key, etc.

Note that for any given report only some of the values below are valid.

values:

Specify one of the following:

System

(Default). Order results by system.

Domain

Order results by domain name.

Host

Order results by host name.

CPG

Order results by Common Provisioning Group (CPG) name.

UsrCPG

Order results by Common Provisioning Group (CPG) name for the User Space of VVs.

SnpCPG

Order results by Common Provisioning Group (CPG) name for the Snap Space of VVs.

VV

Order results by Virtual Volume (VV) name.

VVType

Order results by Virtual Volume (VV) Type.

ProvType

Order results by Virtual Volume (VV) provisioning type.

LD

Order results by Logical Disk (LD) name.

n:s:p

Order results by port. For this value, an escape sequence may be required to retain spacing and special characters (for example: n%3As%3Ap).

PortNode

Order results by port node number (The n part in n:s:p).

RaidType

Order results by RAID level (for example, RAID 0, RAID1, RAID 5 and RAID 6).

LDUsage

Order results by logical disk usage. See [Table 28 \(page 88\)](#) for a list of these usage types.

DiskType

Order results by physical disk type.

DiskSpeed

Order results by physical disk speed.

CageID

Order results by drive cage ID number.

Mag

Order results by magazine position within the drive cage (0 through 7).

Disk

Order results by disk position on a magazine (0 through 3).

PDID

Order results by physical disk ID number.

Node

Order results by controller node number (0-7, depending on the storage server model and configuration).

IOPs

Order results by I/O per second.

Bandwidth

Order results by bandwidth.

Service Time

Order results by service time. For this value, an escape sequence may be required to retain spacing (for example, "Service%20Time").

IO Size

Order results by I/O size. For this value, an escape sequence may be required to retain spacing (for example, "IO%20Size").

Queue Length

Order results by queue length. For this value, an escape sequence may be required to retain spacing (for example, "Queue%20Length").

PortType

Order results by port type.

SelChunkSize

Order results by total size of all the selected chunklets.

Idle

Order results by CPU idle percentage.

Intr

Order results by CPU interrupts/sec.

Ctxt

Order results by CPU context switches/sec.

VirtSize

Order results by VV virtual size.

TotalRsvdSize

Order results by VV reserved space size.

TotalRawSize

Order results by VV total raw reserved space size.

LDSize

Order results by total LD size.

LDFreeSize

Order results by total free LD size.

LDRawSize

Order results by total raw LD size.

TotalAccesses

Order results by total number of accesses to the VV.

ReadHit%

Order results by total read hit percentage to the VV.

WriteHit%

Order results by total write hit percentage to the VV.

example:

The following example specifies that results should be ordered by system:

```
&orderby=System
```

The following example specifies that results should be ordered by system, disk type, and physical disk ID:

```
&orderby=System&orderby=DiskType&orderby=PDID
```

histbegin

Use the `histbegin` option to specify whether the histogram should include only accesses that occurred during the sampling interval corresponding to the category, or all accesses up to the current sample. This option can only be specified once in a given query string.

values:

Specify one of the following:

0

(Default) Include only accesses in the sampling interval.

1

Include all accesses.

example:

The following example specifies that the histogram should include all accesses:

```
&histbegin=1
```

ldspaceunit

Use the `ldspaceunit` option to determine the unit of measure for logical disk space reports. This option can only be specified once in a given query string.

values:

Specify one of the following:

MiB (2²⁰ bytes)

(Default). Use megabytes (1,048,576 bytes) as the logical disk space unit. For this value, an escape sequence may be required to retain spacing and special characters (for example, "MiB%20%282%5E20%20bytes%29").

GB (10⁹ bytes)

Use gigabytes (1,000,000 bytes) as the logical disk space unit. For this value, an escape sequence may be required to retain spacing and special characters (for example, "GB%20%2810%5E9%20bytes%29").

GiB (2³⁰ bytes)

Use gibibytes (1,024 megabytes) as the logical disk space unit. For this value, an escape sequence may be required to retain spacing and special characters (for example, "GiB%20%282%5E30%20bytes%29").

examples:

The following example specifies that the logical disk space unit should be in megabytes:

```
&ldspaceunit=MiB%20%282%5E20%20bytes%29
```

The following example specifies that the logical disk space unit should be in gigabytes:

```
&ldspaceunit=GB%20%2810%5E9%20bytes%29
```

The following example specifies that the logical disk space unit should be in gibibytes:

```
&ldspaceunit=GiB%20%282%5E30%20bytes%29
```

pdspaceunit

Use the `pdspaceunit` option to determine the unit of measure for physical disk space reports. This option can only be specified once in a given query string.

values:

Specify one of the following:

Chunklets

(Default). Use the chunklet (a 256-MB block of contiguous space on a physical disk) as the physical disk space unit.

GB (10⁹ bytes)

See [“ldspaceunit” \(page 141\)](#).

GiB (2^{30} bytes)

See [“ldspaceunit” \(page 141\)](#).

examples:

The following example specifies that the physical disk space unit should be in chunklets:

```
&pdspaceunit=Chunklets
```

The following example specifies that the logical disk space unit should be in gigabytes:

```
&pdspaceunit=GB%20%2810%5E9%20bytes%29
```

The following example specifies that the logical disk space unit should be in gibibytes:

```
&pdspaceunit=GiB%20%282%5E30%20bytes%29
```

vvspaceunit

Use the `vvspaceunit` option to determine the unit of measure for virtual volume space reports. This option can only be specified once in a given query string.

values:

Specify one of the following:

MiB (2^{20} bytes)

See [“ldspaceunit” \(page 141\)](#).

GB (10^9 bytes)

See [“ldspaceunit” \(page 141\)](#).

GiB (2^{30} bytes)

(Default). See [“ldspaceunit” \(page 141\)](#).

examples:

The following example specifies that the virtual volume space unit should be in megabytes:

```
&vvspaceunit=MiB%20%282%5E20%20bytes%29
```

The following example specifies that the virtual volume space unit should be in gigabytes:

```
&vvspaceunit=GB%20%2810%5E9%20bytes%29
```

The following example specifies that the virtual volume space unit should be in gibibytes:

```
&vvspaceunit=GiB%20%282%5E30%20bytes%29
```

selcageid

Use the `selcageid` option to select a drive cage. When no value is specified, no cages are excluded based on their cage IDs.

value:

Specify one or more of the following:

--All cage IDs--

(Default). Does not exclude cages based on their cage IDs.

<cage_ID>

This value consists of a drive cage ID.

examples:

The following example shows two ways to include cages regardless of their cage IDs:

```
&selcageid=  
  
OR  
  
&selcageid=--All%20Cage%20IDs--
```

The following example includes only physical disks in cage number 1:

```
&selcageid=1
```

The following example includes only physical disks in cages number 1, 3 or 5:

```
&selcageid=1&selcageid=3&selcageid=5
```

selchunks

Use the `selchunks` option to include chunklets according to their chunklet type. If no value is specified, no chunklets are excluded based on their chunklet states.

values:

Specify one or more of the following:

NRM_USED_OK

Include only normal used chunklets that are OK.

NRM_USED_FAIL

Include only normal used chunklets that have failed.

NRM_AVAIL_CLEAN

Include only normal available (unused) chunklets that have been cleaned (initialized).

NRM_AVAIL_DRTY

Include only normal available (unused) chunklets that are still being cleaned (and are uninitialized).

NRM_AVAIL_FAIL

Include only normal available (unused) chunklets that are failed.

SPR_USED_OK

Include only spare used chunklets that are OK.

SPR_USED_FAIL

Include only spare used chunklets that have failed.

SPR_AVAIL_CLEAN

Include only spare available (unused) chunklets that have been cleaned (initialized).

SPR_AVAIL_DRTY

Include only spare available (unused) chunklets that are still being cleaned (and are uninitialized).

SPR_AVAIL_FAIL

Include only spare available (unused) chunklets that have failed.

example:

The following example shows how to include chunklets regardless of their type:

```
&selchunks=
```

The following example specifies using only normal, available (unused) chunklets that are clean and available for use:

```
&selchunkls=  
NRM_AVAIL_CLEAN
```

The following example specifies including normal, available (unused) chunklets that are clean and available for use as well as spare available (unused) chunklets that are cleaned and available for use:

```
&selchunkls=  
NRM_AVAIL_CLEAN&selchunks=SPR_AVAIL_CLEAN
```

selcpg

Use the `selcpg` option to select a common provisioning group (CPG).

value:

Specify one or more of the following:

--All CPGs--

(Default). This value does not exclude CPGs based on their names.

<CPG_name>

This value consists of a CPG name.

example:

The following example shows two ways to include all CPGs regardless of name:

```
&selcpg=OR&selcpg=--All%20CPGs--
```

The following example includes only the CPGs named CPG1 and CPG2:

```
&selcpg=CPG1&selcpg=CPG2
```

selusrpcpg

Use the `selusrpcpg` option to select a Common Provisioning Group (CPG) that is used for the user space of VVs. Usage and values are similar to `selcpg` (see “[selcpg](#)” (page 144)).

selsnpcpg

Use the `selusrpcpg` option to select a Common Provisioning Group (CPG) that is used for the snap space of VVs. Usage and values are similar to `selcpg` (see “`selcpg`” (page 144)).

seldisktype

Use the `seldisktype` option to select a disk type.

value:

Specify one or more of the following:

`--All Disk Types--`

(Default). This value does not exclude disks based on their type.

FC

This value allows you to include only Fibre Channel type disks.

NL

This value allows you to include only Near Line type disks.

example:

The following examples show two ways to include disks regardless of type:

```
&seldisktype=OR&seldisktype=--All%20Disk%20Types--
```

The following example specifies that only Fibre Channel disks are to be included:

```
&seldisktype=FC
```

seldiskspeed

Use the `seldiskspeed` option to select a disk speed.

value:

Specify one or more of the following:

`--All Disk Speeds--`

(Default). This value does not exclude disks based on their speed.

15

This value allows you to include only 15K RPM disks.

10

This value allows you to include only 10K RPM disks.

7

This value allows you to include only 7200 RPM disks.

example:

The following examples show two ways to include disks regardless of speed:

```
&seldiskspeed=OR&seldiskspeed=--All%20Disk%20Speeds--
```

The following example specifies that only 10K RPM disks are to be included:

```
&sel diskspeed=10
```

seldomain

Use the `seldomain` option to select a domain according to the domain name.

value:

Specify one or more of the following:

```
--ALL Domains--
```

(Default). Does not exclude any domain based on domain names.

```
<domain_name>
```

This value consists of a storage server domain name.

example:

The following example shows two ways to specify that no domain be excluded based on host name:

```
&seldomain=OR&seldomain=--All%20Domains--
```

The following example specifies including only VLUNs belonging to domain `domain1`:

```
&seldomain=domain1
```

The following example specifies including only VLUNs belonging to domain `domain1` or `domain2`:

```
&seldomain=domain1&seldomain=domain2
```

selhost

Use the `selhost` option to select a host according to the host name.

value:

Specify one or more of the following:

```
--ALL Hosts--
```

(Default). Does not exclude any hosts based on host names.

```
<host_name>
```

This value consists of a storage server host name.

example:

The following example shows two ways to specify that no hosts be excluded based on host name:

```
&selhost=  
OR  
&selhost=--All%20Hosts--
```

The following example specifies including only VLUNs exported to host `winhost1`:

```
&selhost=winhost1
```

The following example specifies including only VLUNs exported to host `winhost1` or `winhost2`:

```
&selhost=winhost1&selhost=winhost2
```

selld

Use the `selld` option to select logical disks according to their LD names.

value:

Specify one or more of the following:

--All LDs--

(Default). Does not exclude logical disks based on their names.

<ld_name>

This value consists of a logical disk name.

example:

The following example shows two ways to specify that LDs not be excluded based on their names:

```
&selld=  
OR  
&selld=--All%20LDs--
```

The following example includes only logical disk named `ld1`:

```
&selld=ld1
```

sellduse

Use the `sellduse` option to select logical disks of a specific usage type.

See [Table 28 \(page 88\)](#) for descriptions of each usage type.

values:

Specify one or more of the following:

V

Includes only logical disks used for volume user space.

C, SD

Includes only logical disks used for copy snapshot data.

C, SA

Includes only logical disks used for copy snapshot administration.

P

Includes only logical disks used for preserved data.

P, F

Includes only logical disks used for first preserved data.

log

Includes only logical disks used for logging.

example:

The following example includes logical disks of all usage types:

```
&sellduse=
```

The following example specifies only logging logical disks:

```
&sellduse=log
```

selnode

Use the `selnode` option to select controller nodes according to their numbers (0-7).

value:

Specify one or more of the following:

```
--All Nodes--
```

(Default). Specifies that nodes are not to be excluded based on their names.

```
<node>
```

This value consists of a controller node number (value is an integer between 0 and 7).

example:

The following example shows two ways to specify that nodes are not to be excluded based on name:

```
&selnode=  
OR  
&selnode>--All%20Nodes--
```

The following example specifies that nodes 0 and 1 will be included:

```
&selnode=0&selnode=1
```

selnsp

Use the `selnsp` option to select ports according to their location.

value:

Specify one or more of the following:

```
--All Ports--
```

(Default). Specifies that no ports are to be excluded based on location.

```
<n:s:p>
```

This value consists of a port position location expressed in `<n:s:p>` format, where `<n>` (0,1,2,3,4,5,6 or 7) is a node number (which depends on the number of nodes and the storage server model), `<s>` (0-5 for each node) is a node PCI slot numbered consecutively from left to right, and `<p>` (1-4, depending on the type of adapter) represents an adapter port numbered consecutively from top to bottom.

example:

The following example shows two ways to specify that no ports are to be excluded based on location:

```
&selnsp=  
OR  
&selnsp=--All%20Ports--
```

The following example specifies that only port 0:1:1 will be included:

```
&selnsp=0%3A1%3A1
```

selpdid

Use the `selpdid` option to select physical disks according to the physical disk ID.

value:

Specify one or more of the following:

```
--ALL PDIDs--
```

(Default). Do not exclude PDs based on their ID number.

```
<PD_ID>
```

This value consists of a physical disk ID number.

example:

The following example shows two ways to specify that PDs are not to be excluded based on their PD IDs:

```
&selpdid=  
OR  
&selpdid=--All%20PDIDs--
```

The following example includes PDs with IDs 1 through 5:

```
&selpdid=1&selpdid=2&selpdid=3&selpdid=4&selpdid=5
```

selpdstate

Use the `selpdstate` option to select physical disks according to their state (`valid` or `Not valid`).

values:

Specify one or more of the following:

```
--All PD States--
```

(Default). PDs are not excluded based on their state.

```
valid
```

Specify this value to include all valid physical disks.

```
Not valid
```

Specify this value to include only physical disks that do not have a state of `valid`. For this value, an escape sequence may be required to retain spacing (for example, `"Not%20valid"`).

example:

The following example shows two ways to specify that physical disks are not to be excluded based on their states:

```
&selpdstate=  
OR  
&selpdstate=--All%20PD%20States--
```

The following example specifies that only physical disks with a current state of `Not valid` should be used:

```
&selpdstate=Not%20valid
```

selporttype

Use the `selporttype` option to select ports according to their type.

values:

Specify one or more of the following:

--All Port Types--

(Default). All port types are selected.

disk

Specify this value to include only ports connected to disks.

free

Specify this value to include only free ports.

host

Specify this value to include only ports connected to hosts.

rcip

Specify this value to include only Remote Copy over IP ports.

rcfc

Specify this value to include only Remote Copy over Fibre Channel ports.

iscsi

Specify this value to include only iscsi ports.

example:

The following example specifies that ports connected to hosts are included:

```
&selporttype=host
```

selraidtype

Use the `selraidtype` option to select logical disks with specific RAID levels.

values:

Specify one or more of the following:

0

Specify this value to include RAID 0 logical disks.

1

Specify this value to include RAID 1 logical disks.

5

Specify this value to include RAID 5 logical disks.

6

Specify this value to include RAID 6 logical disks.

example:

The following example specifies inclusion of all logical disks regardless of RAID level:

```
&selraidtype=
```

The following example includes RAID 1 and RAID 5 logical disks only:

```
&selraidtype=1&selraidtype=5
```

selsys

Use the `selsys` option to select a particular storage system.

values:

Specify one or more of the following:

--All Systems--

(Default). Does not exclude storage servers based on their system names.

<sysname>

This value consists of a storage server system name.

example:

The following example shows two ways to specify that storage servers should not be excluded based on their names:

```
&selsys=  
OR  
&selsys=--All%20Systems--&
```

The following example includes data from InServ1:

```
&selsys=  
InServ1
```

selvv

Use the `selvv` option to specify virtual volumes according to their names.

values:

Specify one or more of the following:

--All VVs--

(Default). Does not exclude virtual volumes based on their names.

<vvname>

This value consists of a virtual volume name.

example:

The following example shows two ways to specify that volumes be included regardless of name:

```
&selvv=  
OR  
&selvv=--All%20VVs--
```

The following example includes volumes with names `vol1`, `vol2`, and `vol3` only:

```
&selvv=vol1&selvv=vol2&selvv=vol3
```

selvvtype

Use the `selvvtype` option to specify virtual volumes according to their types.

values:

Specify one or more of the following:

`--All VV Types--`

(Default). Does not exclude virtual volumes based on their types.

`base`

Select base VVs.

`vcopy`

Select Virtual Copy (snapshot) VVs.

`pcopy`

Select Full Copy (physical copy) VVs.

selprovtype

Use the `selprovtype` option to specify virtual volumes according to their provisioning types.

values:

Specify one or more of the following:

`--All VV Types--`

(Default). Does not exclude virtual volumes based on their types.

`full`

Select fully provisioned VVs with no snapshot space or statically allocated snapshot space.

`cpvv`

Select commonly provisioned VVs. These have fully provisioned User Space and Snapshot space allocated from a Snp CPG.

`tpvv`

Select thin provisioned VVs that have separate User and Snap spaces.

`tpsd`

Select old-style (created before InForm OS release 2.3.1) thin provisioned VVs that have user and snapshot data allocated from a single Snap Space.

`snp`

Select snapshot VVs.

Report Types

Report types are identified with names that typically use the following format:

```
<object>_<yaxis>_<xaxis>
```

where

- <object> is the entity being measured
- <yaxis> represents the metrics plotted on the y-axis of the charts
- <xaxis> is the measure that is on the x-axis of a chart.

For example, the `pd_perf_time` report type plots various performance metrics of physical disks on the y-axis with time on the x-axis.

In order to generate a report, you must specify a report type using the following syntax:

```
report=<report_type>
```

where <report_type> is one of the report types described below.

Report types can be grouped into the following four categories:

- summary report
- space reports
- performance reports
- objects present reports

Summary Report

summary

Shows summary information for all systems.

example:

The following example shows the entire query string for a summary report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&pdspaceunit=Chunklets&report=summary
```

Space Reports

pd_space_time

Plots the total number of chunklets of different types summed over the specified systems and PDs over a period of time (x-axis). Can show the rate of usage of raw space.

example:

The following example shows the entire query string for a `pd_space_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&endtsecs=&  
timeform=absolute&pdspaceunit=Chunklets&selsys=--All%20Systems--&  
selpdid=--All%20PDIDs--&selchunks=&selpdstate=--All%20PD%20States--&
```


example:

The following example shows the entire query string for a `vv_space_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
endtsecs=&timeform=absolute&vvspaceunit=Mib%20%282%5E20%20bytes%29&selsys=--All%20Systems--&  
selcpg=--All%20CPGs--&selvv=--All%20VVs--&  
report=vv_space_time
```

vv_space_group

Plots the virtual size as well as the raw and LD space used by the user, snap and admin (on y-axis) at a specific time, summed over (grouped by) any combination of parameters (for example, system, volume, etc.). For the LD space, it separates free and used space.

example:

The following example shows the entire query string for a `vv_space_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
vvspaceunit=Mib%20%282%5E20%20bytes%29&selsys=--All%20Systems--&  
selcpg=--All%20CPGs--&selvv=--All%20VVs--&groupby=System&orderby=System&  
report=vv_space_group
```

Performance Reports

Performance reports show information about physical disk, logical disk, port, VLUN (export), VV cache, and CPU performance.

pd_perf_time

Returns a chart and/or table of performance metrics (IOPs, bandwidth, service time, I/O size) for physical disks versus time, limited by any options selecting systems, PD IDs, or ports. With no comparison type the metrics are an aggregate of the selected objects. With a comparison type of `System`, `PDID`, or `n:s:p`, data is returned for each selected System, PD, or port respectively, being aggregated over the other object types (e.g., aggregated over selected PDs and ports when `compare=System`).

example:

The following example shows the entire query string for a `pd_perf_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
endtsecs=&timeform=absolute&selsys=--All%20Systems--&selpdid=--All%20PDIDs--&  
selnsp=--All%20Ports--&report=pd_perf_time
```

pd_perf_group

Plots the aggregate standard performance metrics (IOPs, bandwidth, service time, I/O size) of all the physical disks in the selected systems at a specific time, grouped by any combination of parameters (for example, `System`, `PDID`, `n:s:p`, etc.).

example:

The following example shows the entire query string for a `pd_perf_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selpdid=--All%20PDIDs--&selsnp=--All%20Ports--&  
groupby=System&orderby=System&report=pd_perf_group
```

pd_svt_hist

Plots the service time histogram of all the physical disks in the selected systems at a specific time.

example:

The following example shows the entire query string for a `pd_svt_hist` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selpdid=--All%20PDIDs--&selsnp=--All%20Ports--  
&histbegin=1&report=pd_svt_hist
```

pd_svt_histvstime

Plots the service time histogram of all the physical disks in the selected systems versus time.

example:

The following example shows the entire query string for a `pd_svt_histvstime` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=lines&category=daily&begintsecs=&  
selsys=--All%20Systems--&selpdid=--All%20PDIDs--&selsnp=--All%20Ports--  
&histbegin=1&report=pd_svt_histvstime
```

port_perf_time

Returns a chart and/or table of performance metrics (IOPs, bandwidth, service time, I/O size) for ports versus time, limited by any options selecting systems, port types, or ports. With no comparison type the metrics are an aggregate of the selected objects. With a comparison type of `System`, `PortType`, or `n:s:p` data is returned for each selected System, port type, or port respectively, being aggregated over the other object types (e.g., aggregated over selected port types and ports when `compare=System`).

example:

The following example shows the entire query string for a `port_perf_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin  
tsecs=&endtsecs=&timeform=absolute&selsys=--All%20Systems--&selporptype=--  
All%20Port%20Types--&selsnp=--All%20Ports--&compare=none&report=port_perf_time
```

port_perf_group

Plots the aggregate standard performance metrics (IOPs, bandwidth, service time, I/O size) of all the ports in the selected systems at a specific time, grouped by any combination of parameters (for example, `System`, `Port Type`, `n:s:p`, etc.).

example:

The following example shows the entire query string for a `port_perf_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
/inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin  
tsecs=&selsys=--All%20Systems--&selporttype=--All%20Port%20Types--&selnsp=--  
All%20Ports--&groupby=System&orderby=System&report=port_perf_group
```

port_svt_hist

Plots the service time histogram of all the ports in the selected systems at a specific time.

example:

The following example shows the entire query string for a `port_svt_hist` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin  
tsecs=&histbegin=&selsys=--All%20Systems--&selporttype=--All%20Port%20Types--  
&selnsp=--All%20Ports--&report=port_svt_hist
```

port_svt_histvstime

Plots the service time histogram of all the ports in the selected systems versus time.

example:

The following example shows the entire query string for a `port_svt_histvstime` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin  
tsecs=&histbegin=&selsys=--All%20Systems--&selporttype=--All%20Port%20Types--  
&selnsp=--All%20Ports--&report=port_svt_histvstime
```

ld_perf_time

Returns a chart and/or table of performance metrics (IOPs, bandwidth, service time, I/O size) for logical disks versus time, limited by any options selecting systems or logical disks. With no comparison type the metrics are an aggregate of the selected objects. With a comparison type of `System`, or `LD` data is returned for each selected System or logical disk respectively, being aggregated over the other object types (e.g., aggregated over selected logical disks when `compare=System`).

example:

The following example shows the entire query string for a `ld_perf_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
endtsecs=&timeform=absolute&selsys=--All%20Systems--&sellid=--All%20LDs--  
&report=pd_perf_time
```

ld_perf_group

Plots the aggregate standard performance metrics (IOPs, bandwidth, service time, I/O size) of all the logical disks in the selected systems at a specific time, grouped by any combination of parameters (for example, `System`, `LD`, etc.).

example:

The following example shows the entire query string for a `ld_perf_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=  
&selsys=--All%20Systems--&sellid=--All%20LDs--&  
groupby=System&orderby=System&report=ld_perf_group
```

ld_svt_hist

Plots the service time histogram of all the logical disks in the selected systems at a specific time.

example:

The following example shows the entire query string for a `ld_svt_hist` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=  
selsys=--All%20Systems--&sellid=--All%20LDs--&histbegin=1&report=ld_svt_hist
```

ld_svt_histvstime

Plots the service time histogram of all the logical disks in the selected systems versus time.

example:

The following example shows the entire query string for a `ld_svt_histvstime` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=  
selsys=--All%20Systems--&sellid=--All%20LDs--&histbegin=1&report=ld_svt_histvstime
```

vlun_perf_time

Returns a chart and/or table of performance metrics (IOPs, bandwidth, service time, I/O size) for VLUNs versus time, limited by any options selecting systems, hosts, VVs, or ports. With no comparison type the metrics are an aggregate of the selected objects. With a comparison type of `System`, `Host`, `VV`, or `n:s:p` data is returned for each selected System, host, virtual volume, or port respectively, being aggregated over the other object types (e.g., aggregated over selected hosts, VVs, and ports when `compare=System`).

example:

The following example shows the entire query string for a `vlun_perf_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=  
endtsecs=&timeform=absolute&selsys=--All%20Systems--&selhost=--All%20Hosts--&  
selvv=--All%20VVs--&selnsp=--All%20Ports--&report=vlun_perf_time
```

vlun_perf_group

Plots the aggregate standard performance metrics (IOPs, bandwidth, service time, I/O size) of all the VLUNs (exports) in the selected systems at a specific time, grouped by any combination of parameters (for example: `System`, `PDID`, `n:s:p`, etc.).

example:

The following example shows the entire query string for a `vlun_perf_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selhost=--All%20Hosts--&selvv=--All%20VVs--&  
selnsp=--All%20Ports--&groupby=System&orderby=System&report=vlun_perf_group
```

vlun_svt_hist

Plots the service time histogram of all the VLUNs (exports) in the selected systems at a specific time.

example:

The following example shows the entire query string for a `vlun_svt_hist` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selhost=--All%20Hosts--&selvv=--All%20VVs--&  
selnsp=--All%20Ports--&histbegin=1&report=vlun_svt_hist
```

vlun_svt_histvstime

Plots the service time histogram of all the VLUNs (exports) in the selected systems versus time.

example:

The following example shows the entire query string for a `vlun_svt_histvstime` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selhost=--All%20Hosts--&selvv=--All%20VVs--&  
selnsp=--All%20Ports--&histbegin=1&report=vlun_svt_histvstime
```

vv_cmp_time

Returns a chart and/or table of cache performance (accesses, hits, misses) for VV cache versus time, limited by any options selecting systems or VVs.

example:

The following example shows the entire query string for a `vv_cmp_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin  
tsecs=&endtsecs=&timeform=absolute&selsys=--All%20Systems--&selvv=--All%20VVs--  
&report=vv_cmp_time
```

vv_cmp_group

Plots the cache performance (accesses, hits, misses) of VV cache in the selected systems and virtual volumes at a specific time, grouped by system and/or virtual volume.

example:

The following example shows the entire query string for a `vv_cmp_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?category=daily&charttab=chart&charttype=stackedbars&refresh=&begin
```

```
tsecs=&selsys=--All%20Systems--&selvv=--All%20VVs--&groupby=System&orderby=System&report=vv_cmp_group
```

cpu_perf_time

Returns a chart and/or table of performance CPU metrics versus time, limited by any options selecting systems or nodes.

example:

The following example shows the entire query string for a `cpu_perf_time` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
endtsecs=&timeform=absolute&selsys=--All%20Systems--&selnode=--All%20Nodes--&  
report=cpu_perf_time
```

cpu_perf_group

Plots the CPU performance metrics in the selected systems and nodes at a specific time, grouped by system and/or node.

example:

The following example shows the entire query string for a `cpu_perf_group` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&category=daily&begintsecs=&  
selsys=--All%20Systems--&selnode=--All%20Nodes--&  
groupby=System&orderby=System&report=cpu_perf_group
```

Adaptive Optimization Reports

regiodensity

Returns a region I/O density report over the requested interval (from `begintsecs` to `endtsecs`). The report is a histogram where the buckets are various region I/O access rates (accesses per GB per minute). Two values are shown for each bucket:

- the total space for regions in that bucket, and
- the total accesses/min for the regions in that bucket.

regiodensity example

The following example shows the entire query string for a `regiodensity` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&begintsecs=-10800&  
selsys=3par1&report=regiodensity
```

regmoves

Returns a report showing the amount of data moved by Adaptive Optimization between various tiers over the requested interval (from `begintsecs` to `endtsecs`).

regmoves example

The following example shows the entire query string for a `regmoves` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&begintsecs=-10800&  
selsys=3par1&report=regmoves
```

vtier

Returns a report showing the amount of data in each tier for each VV space (`usr`, `snp` and `adm`) at the selected time. The amount of data is shown in both absolute terms (MiB) and as a percentage.

vtier example

The following example shows the entire query string for a `vtier` report:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?charttab=chart&charttype=stackedbars&selsys=3par1&report=vtier
```

Options Available by Report

The following figures display the options that are valid for summary and space report and lists the options that are valid for performance reports.

An entry can be:

- blank: Indicates that the option is not valid for the report.
- Y: Indicates that the option can be specified once for the report.
- M: Indicates that the option can be specified multiple times for the report.

Figure 51 Valid Options for Summary, Space and Adaptive Optimization Reports (I)

Report: Option	summary	pd_space_time	pd_space_group	cpq_space_time	cpq_space_group	ld_space_time	ld_space_group	vv_space_time	vv_space_group	regiodensity	regmoves	vvtier
category	Y	Y	Y	Y	Y	Y	Y	Y	Y			
compare		Y			Y	Y		Y				
compareselel		Y				Y		Y				
table												
refresh	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
charttab		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
chartlib	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
charttype		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
graphx		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
graphy		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
grpahlegpos		Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
tableformat	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
alllabels		Y	Y	Y	Y	Y		Y				
timeform		Y				Y		Y				
begintsecs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
endtsecs		Y	Y	Y	Y	Y		Y		Y	Y	
selMaxNumber	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
groupby			M	M	M		M		M			
orderby			M	M	M		M		M			
histbegin												

Figure 52 Valid Options for Summary, Space and Adaptive Optimization Reports (II)

Report:	summary	pd_space_time	pd_space_group	cpq_space_time	cpq_space_group	ld_space_time	ld_space_group	vv_space_time	vv_space_group	regiodensity	regmoves	vv_tier
histlbkt												
histhbkt												
histsum												
pdspaceunit		Y	Y									
ldspaceunit				Y	Y	Y	Y					
vvspaceunit								Y	Y			
selcageid		M	M									
selchunks		M	M									
selcpq				M	M	M	M					
selusrpq								M	M			
selnpcq								M	M			
seldisktype		M	M	M	M	M	M					
seldiskspeed												
selhost								M	M			
selld						M	M					
sellduse						M	M					
selnode												
seltonode												
selqueue												
selnsp												

Figure 53 Valid Options for Summary, Space and Adaptive Optimization Reports (III)

Report: Option	summary	pd_space_time	pd_space_group	qpg_space_time	qpg_space_group	ld_space_time	ld_space_group	vv_space_time	vv_space_group	regiodensity	regmoves	vvTier
selpdid		M	M									
selpdstate		M	M									
selporttype												
selprovtype								M	M			
selraidtype				M	M	M	M					
selsys		M	M	M	M	M	M	M	M	Y	Y	Y
seldomain				M	M	M	M	M	M	Y	Y	Y
selvv								M	M			
selvtype								M	M			

Figure 54 Valid Options for Performance Reports (I)

Report: Option	pd_perf_time	pd_perf_group	pd_svt_hist	pd_svt_histstime	port_perf_time	port_perf_group	port_svt_hist	port_svt_histstime	ld_perf_time	ld_perf_group	ld_svt_hist	ld_svt_histstime	vlun_perf_time	vlun_perf_group	vlun_svt_hist	vlun_svt_histstime	vv_cmp_time	vv_cmp_group	cpu_perf_time	cpu_perf_group	link_perf_time	link_perf_group	nd_cmp_time	nd_cmp_group
category	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
compare	Y				Y				Y				Y						Y		Y			
comparesel	Y				Y				Y				Y						Y		Y			
table																								
refresh	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Figure 55 Valid Options for Performance Reports (II)

Report: Option	pd_perf_time	pd_perf_group	pd_svt_hist	pd_svt_histstime	port_perf_time	port_perf_group	port_svt_hist	port_svt_histstime	ld_perf_time	ld_perf_group	ld_svt_hist	ld_svt_histstime	vln_perf_time	vln_perf_group	vln_svt_hist	vln_svt_histstime	vw_cnp_time	vw_cnp_group	cpu_perf_time	cpu_perf_group	link_perf_time	link_perf_group	nd_cnp_time	nd_cnp_group
charttab	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
charttype	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
graphx	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
graphy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
grpahlegpos	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
tableformat	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
alllabels	Y				Y				Y				Y				Y		Y		Y		Y	
timeform	Y				Y				Y				Y				Y		Y		Y		Y	
begintsecs	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
endtsecs	Y				Y			Y	Y			Y	Y			Y	Y		Y		Y		Y	
groupby		M				M				M				M				M	M	M	M	M	M	M
orderby		M				M				M				M				M	M	M	M	M	M	M
histbegin			Y	Y			Y	Y			Y	Y			Y	Y								
histlbkt				Y				Y				Y				Y								
histbkt				Y				Y				Y				Y								
histsum				Y				Y				Y				Y								
pdspaceunit																								
ldspaceunit																								
vwspaceunit																								

Figure 56 Valid Options for Performance Reports (III)

Report: Option	pd_perf_time	pd_perf_group	pd_svt_hist	pd_svt_histvstime	port_perf_time	port_perf_group	port_svt_hist	port_svt_histvstime	ld_perf_time	ld_perf_group	ld_svt_hist	ld_svt_histvstime	vlun_perf_time	vlun_perf_group	vlun_svt_hist	vlun_svt_histvstime	vv_cmp_time	vv_cmp_group	cpu_perf_time	cpu_perf_group	link_perf_time	link_perf_group	nd_cmp_time	nd_cmp_group
selcageid																								
selchunks																								
selcpg									M	M	M	M												
selusrcpg																								
selsnpcpg																								
seldisktype	M	M	M	M																				
seldiskspeed	M	M	M	M																				
selhost													M	M	M	M								
selld									M	M	M	M												
sellduse																								
selnode																			M	M	M	M	M	M
seltonode																					M	M		
selqueue																					M	M		
selnsp	M	M	M	M	M	M	M	M					M	M	M	M								
selpdid	M	M	M	M																				
selpdstate																								
selporttype					M	M	M	M																
selprovtype																								
selraidtype																								

Figure 57 Valid Options for Performance Reports (IV)

Report: Option	pd_perf_time	pd_perf_group	pd_svt_hist	pd_svt_histvstime	port_perf_time	port_perf_group	port_svt_hist	port_svt_histvstime	ld_perf_time	ld_perf_group	ld_svt_hist	ld_svt_histvstime	vlun_perf_time	vlun_perf_group	vlun_svt_hist	vlun_svt_histvstime	vv_cmp_time	vv_cmp_group	cpu_perf_time	cpu_perf_group	link_perf_time	link_perf_group	nd_cmp_time	nd_cmp_group
selsys	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M	M
seldomain									M	M	M	M	M	M	M	M	M	M						
selvv													M	M	M	M	M	M						
selvvttype																								

Objects Present Reports

Objects present reports return the objects present in a table over a specified time for use in user interface controls. For example, to allow the user to select a display of PD space usage at a particular time, you would need to know which times, systems, and PD IDs are present in a table. Sort order of columns is ascending, starting with the left column, then the next to the right, and so on.

To see more examples of objects present report query strings, visit the System Reporter Excel client query log (see [“Accessing the Excel Client Query Log”](#) (page 61)). In the Excel query log, note the use of `tableformat=Excel` in some of the query strings. Use of this option inserts an apostrophe (') before any table value that should be interpreted literally. This prevents Excel from potentially converting a value (such as a CPG name) that begins with a number into a number string.

The `objects_present` report requires the `category` (see [“category”](#) (page 128)) and `table` (see [“table”](#) (page 131)) parameters. The `category` parameter must be one of `daily`, `hourly` or `hires`. The valid table values for the report are described in [Table 58](#) (page 167). Exactly one value for the `category` and `table` parameters must be specified.

For example, the following query string produces the output displayed at the top in [Figure 58](#) (page 167):

```
http://localhost/cgi-bin/3par-rpts/
inserv_perf.exe?report=objects_present&objectpres=cpg_name&table=ldspace&
category=hourly&tableformat=excel
```

Removing `tableformat=excel` from this query string produces the output displayed at the bottom in [Figure 58](#) (page 167).

Figure 58 Using the `tableformat=excel` Option

Objects in `ldspace_hourly_1, system_1`

System	CPG
'Scott's	'---
'Scott's	'cpg
"TPD System #0	'---
"TPD System #0	'bfarm-root-cpg
"TPD System #0	'rootfs
"TPD System #0	'workspaces

Table 58 Objects Present Report table values

Report	table values
<code>objects_present&"objectpres= cageid"</code> (page 168) <i>objectpres= cageid</i> on page 8.65	pdspace
<code>objects_present&"objectpres= cpg_name"</code> (page 169) <i>objectpres= cpg_name</i> on page 8.66	ldspace vvspace statld
<code>objects_present&"objectpres= diskspeed"</code> (page 170) <i>objectpres= diskspeed</i> on page 8.67	pdspace statpd
<code>objects_present&"objectpres= disktype"</code> (page 170) <i>objectpres= disktype</i> on page 8.68	ldspace pdspace statpd

Table 58 Objects Present Report table values (continued)

Report	table values
objects_present&"objectpres= domain_name" (page 171)objectpres= domain_name on page 8.69	ldspace vvspace statld statvlun statvvcmp
objects_present&"objectpres= host_name" (page 171)objectpres= host_name on page 8.69	statvlun
objects_present&"objectpres= ld_name" (page 172)objectpres= ld_name on page 8.70	ldspace statld
objects_present&"objectpres=node" (page 172)objectpres=node on page 8.71	statcpu
objects_present&"objectpres=pdid" (page 173)objectpres=pdid on page 8.72	pdspace statpd
objects_present&"objectpres=port" (page 174)objectpres=port on page 8.73	statpd statvlun statport
objects_present&"objectpres= system" (page 174)objectpres= system on page 8.74	ldspace pdspace vvspace statcpu statpd statport statld statvlun statvvcmp
objects_present&"objectpres= vv_name" (page 175)objectpres= vv_name on page 8.75	statvlun statvvcmp vvspace
objects_present&"objectpres= vv_name_nosnap" (page 175)objectpres= vv_name_nosnap on page 8.76	vvspace
objects_present&"objecctpres= porttype" (page 176)objecctpres= porttype on page 8.77	statport
"sample_times" (page 177)sample_times on page 8.78 Note: The report is sample_times, objects_present and objectpres should not be specified.	ldspace pdspace vvspace statcpu statpd statport statld statvlun statvvcmp

objectpres= cageid

Returns a list of drive cages.

example:

The following example shows the entire query string for:

objects_present&objectpres=cageid

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=cageid&table=pdspace&  
category=daily
```

output:

Returns a row for each drive cage ID number on each system.

Output columns are:

- System: the system name
- CageID: drive cage ID numbers

The following is example output of the query:

Objects in pdspace_daily_2, system_1

System	CageID
Scott's	0
Scott's	1
Scott's	2
TPD System #0	1
TPD System #0	2
TPD System #0	3
TPD System #0	4
TPD System #0	5
TPD System #0	6

objectpres= cpg_name

Returns a list of common provisioning group (CPG) names.

example:

The following example shows the entire query string for:

objects_present&objectpres=cpg_name

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=cpg_name&table=ldspace&  
category=hares
```

output:

Returns a row for each common provisioning group on each system.

Output columns are:

- System: the system name
- CPG: common provisioning group name.

The following is example output of the query:

Objects in ldspace_hires_1, system_1

System	CPG
Scott's	---
Scott's	cpg
TPD System #0	---
TPD System #0	bfarm-root-cpg
TPD System #0	rootfs
TPD System #0	workspaces

objectpres= diskspeed

Returns a list of disk speeds.

example:

The following example shows the entire query string for:

objects_present&objectpres=diskspeed

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=diskspeed&table=statpd&  
category=daily
```

output:

Returns a row for each disk speed on each system.

Output columns are:

- System: the system name
- DiskSpeed: list of disk speed

objectpres= disktype

Returns a list of disk types.

example:

The following example shows the entire query string for:

objects_present&objectpres=disktype

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=disktype&table=pdspace&  
category=daily
```

output:

Returns a row for each disk type on each system.

Output columns are:

- System: the system name
- Disk Types: list of disk types

The following is example output of the query:

Objects in pdspace_daily_2, system_1

System	Disk Types
Scott's	FC
TPD System #0	FC

objectpres= domain_name

Returns a list of domain names.

example:

The following example shows the entire query string for:

objects_present&objectpres=disktype

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=domain_name&table=statvln&  
category=daily
```

output:

Returns a row for each domain present over the time range on each system.

Output columns are:

- System: the system name
- Domain: domain names

objectpres= host_name

Returns a list of host names.

example:

The following example shows the entire query string for:

objects_present&objectpres=disktype

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=host_name&table=statvln&  
category=daily
```

output:

Returns a row for each host present over the time range on each system.

Output columns are:

- System: the system name
- Host : host names

The following is example output of the query:

Objects in statvluu_daily_1, system_1

System	Host
Scott's	sqa-sune450-06
TPD System #0	bfar64bit
TPD System #0	maude
TPD System #0	ned
TPD System #0	rod
TPD System #0	simpsons
TPD System #0	todd
TPD System #0	toolchest

objectpres=ld_name

Returns a list of logical disk names.

example:

The following example shows the entire query string for:

objects_present&objectpres=disktype

```
http://localhost/cgi-bin/3par-rpt/  
/inserv_perf.exe?report=objects_present&objectpres=ld_name&table=ldspace&  
category=hourly
```

output:

Returns a row for each logical disk present over the time range for each system.

Output columns are:

- System: the system name
- LD: logical disk names

The following is example output of the query:

Objects in ldspace_hourly_1, system_1

System	LD
Scott's	admin.usr.0
Scott's	tp-0-sa-0.0
Scott's	tp-0-sa-0.1
Scott's	tp-0-sa-0.2
Scott's	tp-0-sd-0.0
Scott's	tp-0-sd-0.1

objectpres=node

Returns a list of controller node numbers.

example:

The following example shows the entire query string for:

objects_present&objectpres=node

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=node&table=statcpu&  
category=hourly
```

output:

Returns a row for each node present over the time range for each system.

Output columns are:

- System: the system name
- Node: controller node number

The following is example output of the query:

Objects in statcpu_hourly_1, system_1

System	Node
Scott's	0
Scott's	1
TPD System #0	0
TPD System #0	1

objectpres=pdid

Returns a list of physical disk IDs.

example:

The following example shows the entire query string for:

objects_present&objectpres=pdid:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=pdid&table=pdspace&  
category=daily
```

output:

Returns a row for each physical disk present over the time range for each system.

Output columns are:

- System: the system name
- PDID: physical disk ID numbers

The following is example output of the query:

Objects in pdspace_daily_2, system_1

System	PDID
Scott's	-1
Scott's	0
Scott's	1
Scott's	2
Scott's	3

objectpres=port

Returns a list of ports.

example:

The following example shows the entire query string for
objects_present&objectpres=port

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=port&table=statvlnu&  
category=daily
```

output:

Returns a row for each port present over the time range for each system.

Output columns are:

- System: the system name
- N:S:P: The port location specified as controller node number:PCI slot number:port number

The following is example output of the query:

Objects in statvlnu_daily_1, system_1

System	N:S:P
Scott's	1:0:2
TPD System #0	0:1:1
TPD System #0	1:1:1

objectpres= system

Returns a list of storage systems.

example:

The following example shows the entire query string for
objects_present&objectpres=system using:

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=system&table=pdspace&  
category=daily
```

output:

Returns a row for each system present over the time range for each system.

Output columns are:

- System: system name
- IP Name: system IP name

The following is example output of the query:

Objects in pdspace_daily_2, system_1

System	IP Name
Scott's	nodeg4
TPD System #0	kang

objectpres= vv_name

Returns a list of virtual volume names.

example:

The following example shows the entire query string for:

objects_present&objectpres=vv_name

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=vv_name&table=statvlnu&  
category=daily
```

output:

Returns a row for each virtual volume present over the time range for each system.

Output columns are:

- System: system name
- VV: virtual volume names

The following is example output of the query:

Objects in statvlnu_daily_1, system_1

System	VV
Scott's	tpvw
TPD System #0	asic-tapeout
TPD System #0	bart-tp-ws
TPD System #0	bfar2-root
TPD System #0	foo2
TPD System #0	generic-root

objectpres= vv_name_nosnap

Returns a list of virtual volume names, excluding snapshot volumes.

example:

The following example shows the entire query string for:

objects_present&objectpres=vv_name_nosnap

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=vv_name_nosnap&table=vvspace&  
category=daily
```

output:

Returns a row for each volume present over the time range (excluding snapshot volumes) for each system.

Output columns are:

- System: system name
- VV: virtual volume name

The following is example output of the query:

Objects in vvspace_daily_1, system_1

System	VV
Scott's	admin
Scott's	tpvv
TPD System #0	26root
TPD System #0	admin
TPD System #0	asic-tapeout
TPD System #0	bart-tp-ws

objectpres= porttype

Returns a list of port types.

example:

The following example shows the entire query string for:

objects_present&objectpres=porttype

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=objects_present&objectpres=porttype&table=statport&  
category=daily
```

output:

Returns a row for each time sample in the given table for the given category (resolution).

Output columns are:

- System: system name
- Port Types: port type

The following is example output of the query:

Objects in statport_daily_1, system_1

System	Port Types
Scott's	disk
Scott's	free
Scott's	host
Scott's	rcip
TPD System #0	disk
TPD System #0	free
TPD System #0	host
TPD System #0	rcip

sample_times

Returns valid times of database samples.

example:

The following example shows the entire query string for:

sample_times

```
http://localhost/cgi-bin/3par-rpts/  
inserv_perf.exe?report=sample_times&category=daily&table=pdspace
```

output:

Returns a row for each time sample in the given table for the given category (resolution).

Output columns are:

- time: string with format yyyy-mm-dd hh:mm:ss
- tsecs: Unix seconds, the number of seconds since 01/01/1970, GMT.

The following is example output of the query:

Times in pdspace_daily_2

Time	tsecs
2005/07/02 00:00:49	1120287649
2005/07/03 00:00:58	1120374058
2005/07/04 00:00:10	1120460410
2005/07/05 00:00:29	1120546829
2005/07/06 00:00:41	1120633241
2005/07/07 00:00:49	1120719649
2005/07/08 00:00:51	1120806051
2005/07/09 00:00:57	1120892457
2005/07/10 00:01:02	1120978862
2005/07/11 00:01:10	1121065270
2005/07/12 00:01:18	1121151678
2005/07/13 00:01:21	1121238081
2005/07/14 00:01:21	1121324481

9 Using Adaptive Optimization

This chapter provides information about the optionally licensed Adaptive Optimization feature that works in conjunction with System Reporter to optimize storage utilization across multiple tiers.

NOTE: Adaptive Optimization is only supported by InForm OS versions 2.3.1 MU 1 and beyond.

Overview

Adaptive Optimization analyzes sub-volume, region level disk access rates for a given array over a scheduled period of time and then performs a data migration of regions between tiers according to a cost versus performance preference. Disk utilization is optimized by having frequently accessed data moved to the higher performance tier (e.g., RAID 1 using Solid State Disks or SSDs) while infrequently accessed data is moved to the lower cost tier (e.g., RAID 6 on Nearline disks). By using Adaptive Optimization, you can achieve a much higher degree of control over disk usage by reserving your faster and more expensive storage resources for the data that is frequently accessed and relegating your slower and less expensive drives to storing data that is only occasionally accessed.

Multiple Adaptive Optimization configurations can be defined for an HP 3PAR Storage System array where each corresponds to a related set of hosts or applications that require optimization. You can assign either two or three storage tiers (defined through tiers numbered 0, 1 and 2) for data migration in a given configuration where tier 0 is the fastest, tier 1 is the intermediate speed, and tier 2 is the slowest. Tiers are defined by Common Provisioning Groups (CPGs) and all of the cost and performance characteristics of the tier are determined by the setting of the CPG, such as the RAID level, number of disks used, disk type and speed. You can also control the maximum space available for each tier. See [“Tier Definition \(CPG Name and GiB\)” \(page 180\)](#) section for more details.

In addition to the space available for the tiers, you can also differentiate configurations by an Adaptive Optimization mode described in [“Adaptive Optimization Mode” \(page 183\)](#).

You must also specify the schedule (the dates, weekdays and hours) when the configuration will execute (see [“Schedule” \(page 182\)](#)) as well as the number of hours of data prior to the execution time that the analysis should consider for optimization purposes (see [“Measurement Hours” \(page 182\)](#)).

The analysis for an Adaptive Optimization configuration first calculates the space available in each CPG (tier). It then calculates the access rate (IO accesses/(GiB * minute)) over the measurement interval (specified in the configuration) for all the VV regions in the CPGs for a configuration as well as the average access rates for each CPG (tier). Based on the space available in each tier, and the performance of each region in comparison with the averages for each of the tiers, the data is moved from one tier to the other as applicable using a CLI program (mvrg).

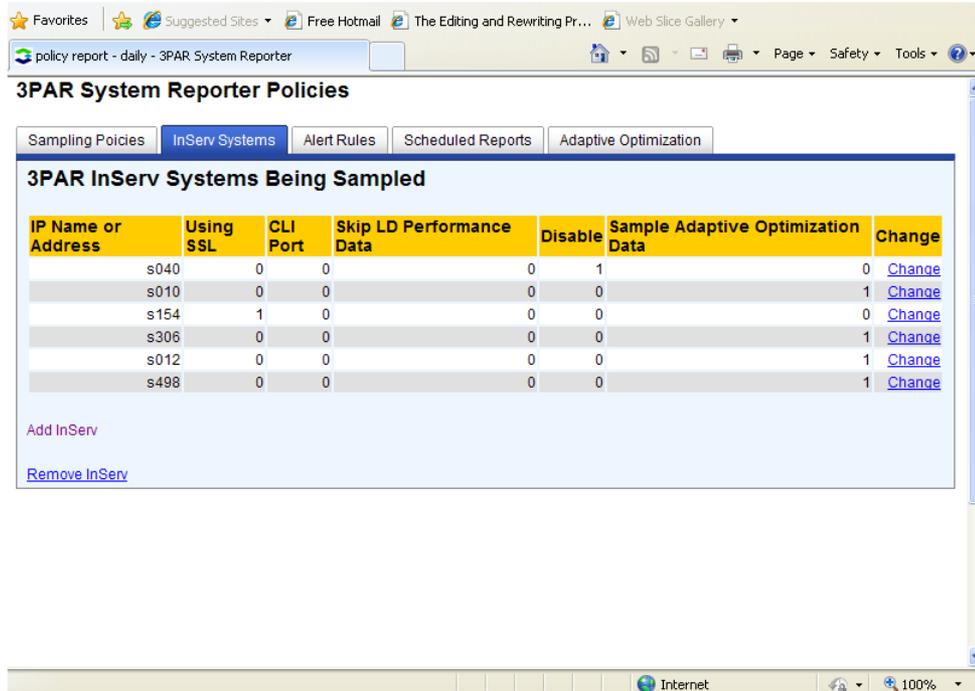
NOTE: Adaptive Optimization analyzes each configuration independent of the other configurations. Only the space available in each tier of the configuration and the relative access rates of the regions in the tiers (CPGs) of the configuration matter. Of course, there may be indirect influences from other configurations. For example, if CPGs from two configuration use the same PDs, then space used by a CPG in one configuration may cause less space to be available for the CPG in the other configuration.

System Reporter can generate a number of different reports that you can use to track storage utilization and monitor the movement of data performed by Adaptive Optimization. See [“Creating Custom Reports” \(page 66\)](#) for details on generating reports.

Selecting an HP 3PAR Storage System for Adaptive Optimization

To begin using Adaptive Optimization on an HP 3PAR Storage System array, you must either add it configured to include collection of data (region-level data) for sampling purposes (see “Adding Storage Systems” (page 47)) or change the configuration to include the data. To change the configuration, click on the **InServ Systems** tab in the System Reporter Policies window.

Figure 59 Setting Up the Storage System for Adaptive Optimization



From this menu, you must ensure that the system that you want to optimize is configured and that the **Sample Adaptive Optimization Data** parameter for the desired system has been checked and is set to **1**. In addition, the **Skip LD Performance Data** should not be set since Adaptive Optimization also looks at LD performance data.

NOTE: In order to use Adaptive Optimization to perform system optimization, you need to have Edit level privileges on the arrays where this capability is implemented.

Configuring Adaptive Optimization

To configure Adaptive Optimization,

1. Point your browser at the Web server where Apache HTTP Server and the HP 3PAR System Reporter Web server scripts are installed.

The **3PAR System Reporter** main window appears.

2. Click **Policy Settings** in the Extras Menu area.

The **System Reporter Policies** window appears.

- Click the **Adaptive Optimization** tab.
The **Adaptive Optimization Configuration** window appears (Figure 60 (page 180)).

Figure 60 Adaptive Optimization Configuration Main Window

Id	InServ	Tier 0 CPG Name	Tier 1 CPG Name	Tier 2 CPG Name	Tier 0 GiB	Tier 1 GiB	Tier 2 GiB	Schedule	Measure Hours	Mode	Active	Change
1	498	sr-per710-ssd-r1	sr-per710-fc-r5	sr-per710-nl-r6	150	800	800	** 0,12	12	Performance	1	Change
2	498	esx-ssd-r1	esx-fc-r1		100	1500	100	** 3,15	12	Balanced	1	Change
3	498	sysrpt-ssd-r1	sysrpt-fc-r1	sysrpt-nl-r6	100	400	800	** 4,16	12	Cost	1	Change
5	498	sw-dev-ssd-r1	sw-dev-fc-r6	sw-dev-nl-r6	150	100000	100000	** 1,13	12	Performance	1	Change
6	498	labuser-ssd-r1	labuser-fc-r1	labuser-nl-r1	100	100000	100000	** 2,14	12	Balanced	1	Change
4	ASIC		buildfarm-FC-r5	buildfarm-NL-r5	100	1000	1000	***	4	Balanced	1	Change

[Add Adaptive Optimization Configuration](#)
[Remove Adaptive Optimization Configuration](#)

The following information is provided for each Adaptive Optimization configuration:

- **Id** – A unique identification number assigned automatically to the configuration when it is created.
- **InServ** – Name of the HP 3PAR Storage System for which the configuration is implemented.
- **Tier # CPG Name** – CPG name associated with a given tier (0 - 2) (see “Tier Definition (CPG Name and GiB)” (page 180)).
- **Tier # GiB** – Maximum size of the storage space that can be used for the corresponding tier (see “Tier Definition (CPG Name and GiB)” (page 180)).
- **Schedule** – Schedule for the optimization process (see “Schedule” (page 182)).
- **Measure Hours** – Amount of time a given CPG is monitored for optimization purposes (see “Measurement Hours” (page 182)).
- **Mode** – Optimization mode (Cost/Performance/Balanced. See “Adaptive Optimization Mode” (page 183)).
- **Active** – Indicates whether or not optimization is currently implemented (true/false).
- **Change** – Clicking on this link brings up a screen used to change the configuration.

Tier Definition (CPG Name and GiB)

Each tier is defined by the CPG and the maximum space that the tier should use. The CPG definition includes all the parameters necessary to define a storage tier’s quality of service: the disk type, disk speed, RAID type, etc. You can specify either two or three tiers. When you specify two tiers you can choose which tier to leave out by setting its CPG to “-”.

NOTE: When you specify two tiers, for the purposes of analysis and data movement only the order of the tiers matters, not the actual tier number. For example specifying tiers 0 and 1 is equivalent to specifying 0 and 2 or 1 and 2. The lower numbered tier is considered the faster tier and the higher numbered tier is considered the slower tier.

However, you should consider whether you may want to add a tier in the future, and if so what tier number that is likely to be, and leave out that tier number for the two tier configuration. For example, if you have SSD PDs and NL PDs, you may want to use SSD PDs for the tier 0 CPG, NL PDs for the tier 2 CPG, and leave tier 1 empty for the possibility of a CPG using FC PDs.

The maximum space in GiB for the tier allows you to limit how much space in a tier is available for a given configuration. If you do not want to place a limit on the amount of space for a tier, you can enter a large value that is unlikely to be exceeded.

Adaptive Optimization looks at several limits on the CPG:

- The CPG allocation (hard) limit
- The CPG warning limit
- The AO space limit for the tier

Adaptive Optimization takes the minimum of these limits and treats that as a space limit for the CPG and attempts to keep the total LD used space in the CPG below this limit (by a margin of about two times the grow size of the CPG).

The TPVV usr space and Snapshot copy space for a CPG can be examined using the InForm CLI `showvv -cpg alloc` command. Since this space may need to grow on demand as the host writes the VV or snapshot, you should make sure that the CPG associated with the TPVV usr space and Snapshot copy space has room to grow (i.e., you should not associate TPVV usr space with a space-constrained SSD CPG).

As an example, suppose your default TPVV usr cpg is CPG1 (tier 1), and you set a tier space limit on tier 1 to 100GB. In between iterations of Adaptive Optimization, suppose the TPVVs grow and the space used in CPG1 grows to 110 GB. Also, suppose the grow size is 8GB. When the next iteration of Adaptive Optimization runs, it will try to reduce the used space in CPG1 to $100 - (2 * 8) = 84$ GB, and attempts to move 26GB from that tier to either tier 0 or tier 2.

The reason it needs two times the grow size is that the HP 3PAR Storage System does not currently allow usr space for fully provisioned VVs and Snap space or TPVV space to be on the same LDs so it needs the equivalent of the grow size for both parts. This can be inefficient, especially for SSD CPGs that are not used for automatic growth, only Adaptive Optimization movement. Therefore, it is recommended that you decrease the grow size for SSD CPGs to the smallest possible amount,

except in the case where the CPG is the default for snapshot or base TPVV allocation as noted in the following caution.

△ **CAUTION:** For some CPGs, especially those that use small PDs such as SSDs, the additional buffer space in the CPG may be too large to be left unused. To minimize the buffer space you can use the CLI to set the CPG grow size to the smallest possible value. However, if you do so, you should make sure that the CPG is not the default CPG for snapshot or base TPVV allocation (the CPGs shown in the `showvv -cpgalloc` CLI command). The default CPG for snapshot or base TPVV allocation should always be one that has sufficient room to grow.

△ **CAUTION:** You can remove a CPG from a configuration by either specifying a different CPG in its place or by setting the tier's CPG to "-" (no CPG). However, when a CPG is removed from a configuration, all the VV space mapped to LDs in that CPG remain in that CPG, Adaptive Optimization does not move them to a different CPG. And since Adaptive Optimization no longer considers that CPG in its analysis that space will remain there.

Therefore, before removing a CPG from an Adaptive Optimization configuration, you should first make sure that there is no VV space mapped to that CPG. You can do so by setting the space for that tier to 0 GiB. Adaptive Optimization will then move all the VV space out from that CPG to other tiers in that configuration. This may take a few iterations of execution of that configuration. You can use the LD space report for that CPG to see when the CPG is empty. You should also make sure that the CPG is not the default CPG for snapshot and base TPVV allocation otherwise new VV space may get mapped to that CPG.

NOTE: Please see ["Adding Storage Systems"](#) (page 47) for details on adding an HP 3PAR Storage System.

Schedule

The schedule consists of three components: dates, weekdays and hours. You can choose multiple values of each of these components to create a schedule for when each adaptive optimization configuration executes.

NOTE: Adaptive optimization configurations may be scheduled to run at intervals that are shorter than their measurement hours (see ["Measurement Hours"](#) (page 182)). However, regions that have been moved are not considered in the analysis for subsequent iterations until at least the measurement interval has expired.

△ **CAUTION:** There is single regionmover process for each HP 3PAR Storage System array, and it completes analysis and data movement for all the adaptive optimization configurations scheduled for the current hour before proceeding to the next hour. If the time taken to complete the current configurations is so long that the hour is past the execution time for the configurations scheduled next, it may skip those configurations. You should therefore ensure that configurations are either executed together or sufficiently spaced.

Measurement Hours

You can specify the number of hours that you want to consider in order to determine the activity level of regions. Control of the measurement duration is important especially for applications whose performance is not uniformly important throughout the day.

For example, consider an application whose performance is important only during the first few hours of a business day (Monday through Friday, 8 am through 11 am) and then not particularly important throughout the rest of the day, though not necessarily idle the rest of the day. If we measure the application throughout the day, the tiering will consider the performance during times that are not important and as a result the performance during the important times of the day may

not be optimal. We can achieve optimal results by scheduling adaptive optimization execution for 11 am Monday through Friday with a measurement interval of 3 or 4 hours. That way only performance measurements during the important period for the application will be considered.

NOTE: The measurement hours selected should not exceed the retention time for adaptive optimization samples set in the sampling policies (see [“Editing Sampling Policies”](#) (page 114)).

Adaptive Optimization Mode

You have the choice of three modes to define the cost versus performance bias that controls the manner in which data is moved for a given configuration as follows:

- **Performance** – The analysis parameters are biased towards moving the data to the faster tiers and away from the slower tiers.
- **Balanced** – The analysis parameters are balanced between Performance and Cost.
- **Cost** – The analysis parameters are biased towards moving data to the slower (less costly) tiers and away from the faster tiers.

You can use the mode parameter to differentiate between the resources that are available to various configurations. For example, you may set the mode for a critical application’s configuration to **Performance** and the mode parameter for a less important application to **Cost**. The critical application’s data will then get relatively more of the higher performance tier space than the less important application.

Adaptive Optimization Active

You can temporarily disable adaptive optimization by setting the Active bit to 0 (false) and re-enable it by setting Active to 1 (true). Setting the Active bit to 0 only disables analysis and data movement, it does not prevent System Reporter from sampling the region counters.

Adding an Adaptive Optimization Configuration

To add an Adaptive Optimization configuration:

1. From the Adaptive Optimization Configuration window, click **Add Adaptive Optimization Configuration**.
The **Add Adaptive Optimization Configuration** window appears.
2. Select the system for which optimization will be enabled. The systems that are available from the pull-down menu are added through the **InServ Systems** tab in the System Reporter Policies window.

NOTE: When you add a new HP 3PAR Storage System to be sampled, the system, its domains and the CPGs that are available may not appear until after the CPG space data has been sampled for that system.

3. Choose a domain.
4. Choose the CPG that is to be monitored and optimized for each tier. Up to 3 tiers may be defined.
5. Enter the amount of storage space (in GBs) to make available for the tier.
6. Schedule when the optimization is to occur. You can specify the date/day of the week and hour to mitigate impact on system performance by having the optimization run at off-peak hours. Hold the **CTRL** key to add multiple selections.
7. Enter the measurement in hours that the sampling for optimization will take place.
8. Select the appropriate mode. Cost/Balanced/Performance.
9. Choose whether the configuration is to be currently active (true/false). This switch allows you to turn the optimization off and on as required.

10. When done, click **Submit Query**. An alert window appears confirming that Adaptive Optimization configuration has been added.
11. Click **OK** to return to the **Sampling Policies** window.

Changing an Adaptive Optimization Configuration

To change an Adaptive Optimization configuration:

1. From the Adaptive Optimization Configuration window, click **Change** next to the configuration that you want to change.
The **Change Adaptive Optimization Configuration** window appears.
2. Make the necessary changes.
3. When done, click **Submit Query**. An alert window appears confirming that the Adaptive Optimization has been changed.
4. Click **OK** to return to the **Sampling Policies** window.

Removing an Adaptive Optimization Configuration

To remove an Adaptive Optimization configuration:

1. From the Adaptive Optimization Configuration window, click **Remove Adaptive Optimization**.
The **Remove Adaptive Optimization Configuration** window appears.
2. Choose the configuration to remove from the pull-down menu based on tiers that are defined.
3. When done, click **Submit Query**. An alert window appears confirming that the Adaptive Optimization has been removed.
4. Click **OK** to return to the **Sampling Policies** window.

NOTE: When removing a CPG (tier) associated with an Adaptive Optimization configuration, the VV data remains on that tier. If you want to remove a tier, you should first make sure that there is no VV data on the tier. To do so, set the space limit for that tier to 0. Adaptive Optimization will, over time (it may require multiple iterations), move the data to the other tiers. Once that CPG has no used space, you can then safely remove or change that CPG.

Adaptive Optimization Reports

System Reporter provides three types of Adaptive Optimization reports as follows:

- **Region IO Density** – Show the I/O access rate and distribution of data for each tier.
- **VV Tiers** – Shows the distribution of space in each VV among the tiers in absolute terms and also as a percentage.
- **Space Moved** – Shows the space moved over a selected time interval.

You can access these reports through the System Reporter Main Menu by clicking on the Custom Reports tab and then selecting the Adaptive Optimization report.

NOTE: See “Adaptive Optimization Reports” (page 111) for details on interpreting an Adaptive Optimization report.

Region IO Density Report

The Region IO Density report for an Adaptive Optimization configuration, shows three charts. The first two charts (Figure 61 (page 185) and Figure 62 (page 185)) show histograms with the X-axis (category) illustrating the IO access rate.

Figure 61 Region IO Density Showing Total Space

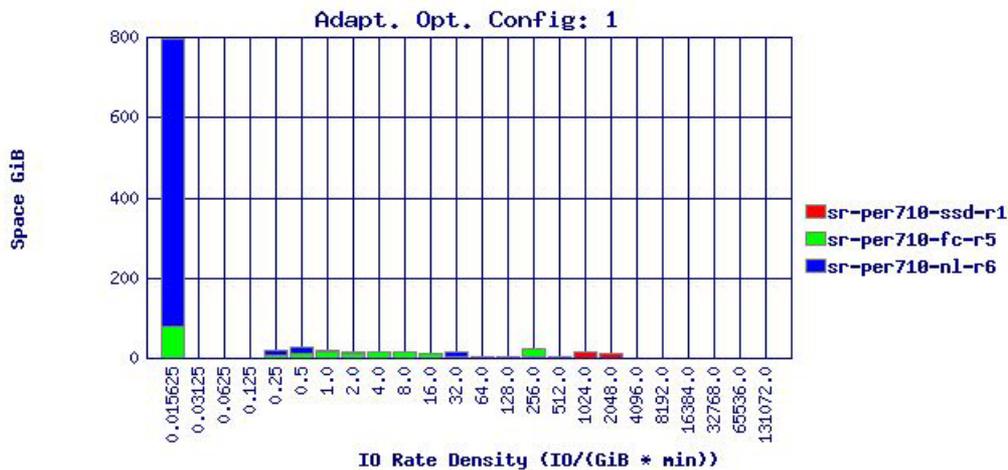
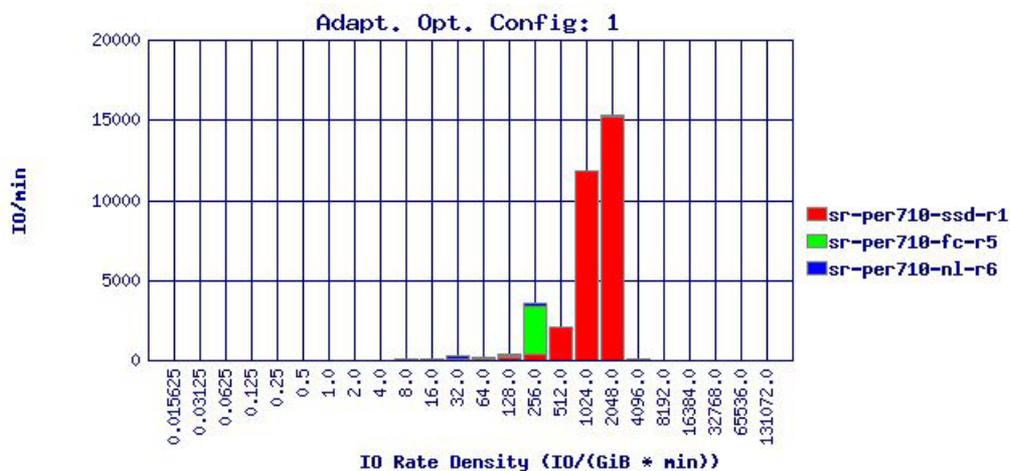


Figure 62 Region IO Density Showing Total IO Access

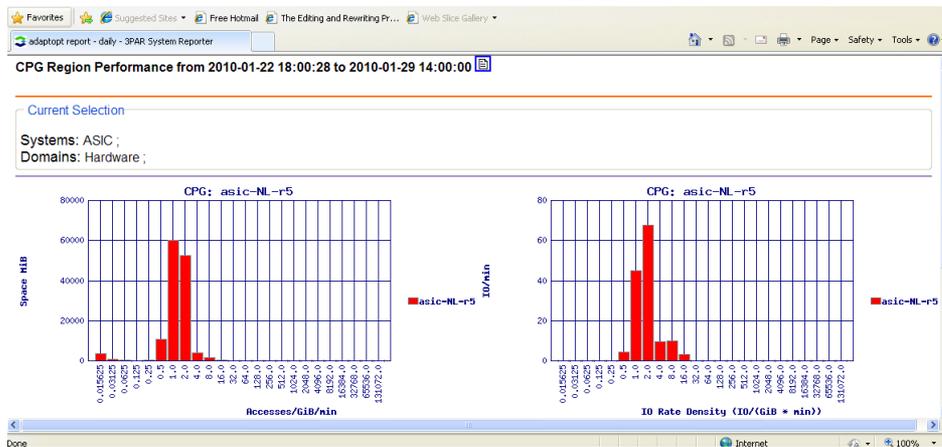


Each region is put into the appropriate histogram bucket on the X-axis based on its IO access rate during the specified time period. The Y-axis (value) for the first chart is the total space for the regions in the histogram bucket, and the Y-axis (value) for the second chart is the total IO accesses/min for the regions in the histogram bucket. The three CPGs corresponding to the three tiers on an Adaptive Optimization configuration are shown with different colors on the same chart, in this case regions in green are tier 0, regions in blue are tier 1 and regions in blue are tier 2.

You can see from first chart (Figure 61 (page 185)) that most of the space is used by tier 2 (blue) for regions that have a very small access rate. You can also see from the second chart (Figure 62 (page 185)) that most of the IO accesses are handled by the regions in tier 0 (red) even though the first chart shows that these regions occupy very little space. This case illustrates Adaptive Optimization working very well: the fastest tier handles most of the IO accesses even though it uses very little space, and conversely most of the space is on the slower tier even though it does very little IO. The region IO density report can also be useful to see if Adaptive Optimization will be useful for the space in a CPG.

Figure 63 (page 186) shows the Region IO density report for a CPG without Adaptive Optimization. The charts shown are very similar to the charts in Figure 62 (page 185) and Figure 61 (page 185) except that there is only one tier (CPG).

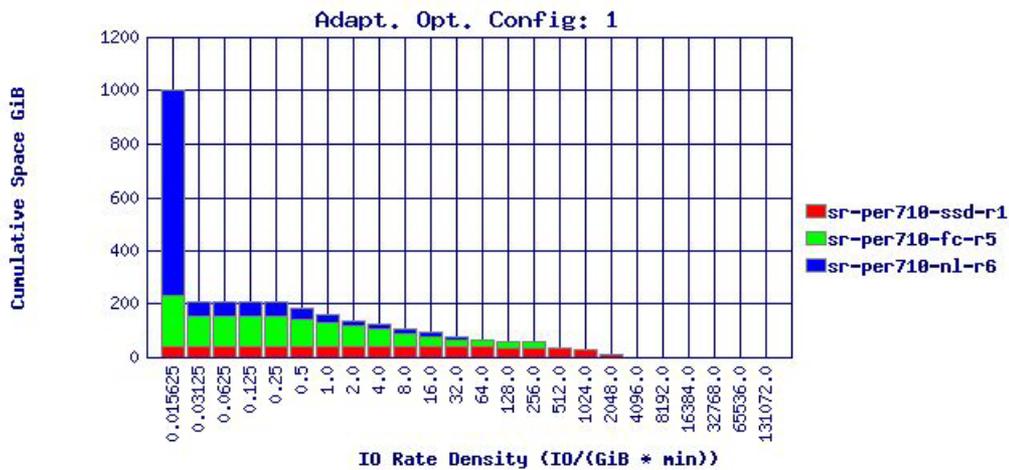
Figure 63 Region IO Density without Adaptive Optimization



You can see that much of the space is used for regions with very few IO accesses and most of the IO accesses are handled by a very small number of regions. Adaptive Optimization would work very well in this situation since the busiest regions would be moved to the fastest tier and the most idle regions would be moved to the slowest tier.

A third report includes cumulative information that can assist in identifying ideal Adaptive Optimization candidates, i.e., applications/profiles that stand to benefit based on the concentration of I/O.

Figure 64 Region IO Density Showing Cumulative Space



The Cumulative Space chart is useful for the following purposes:

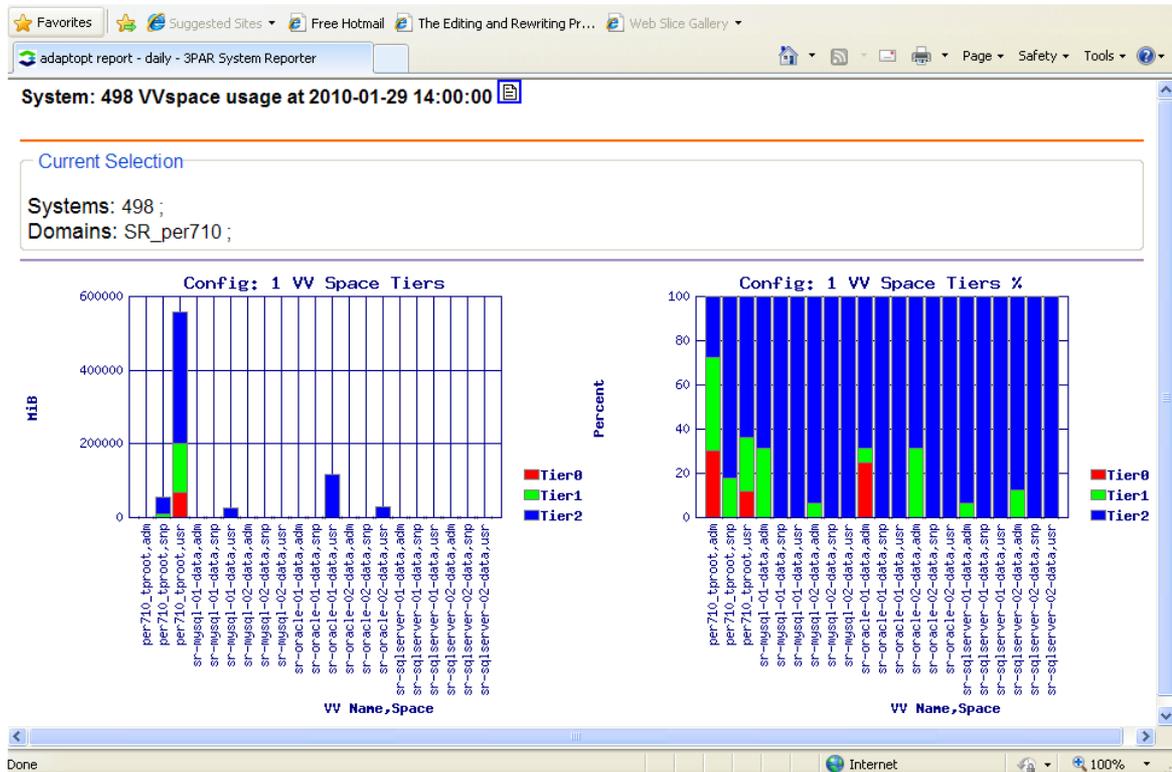
- To determine the ideal candidates for Adaptive Optimization.
- To identify what percent of capacity needs to remain or be removed from a high-performance tier and what percent of IOPS are served for that particular Adaptive Optimization profile.

NOTE: Some features provided by the cumulative reports are only implemented by the charts provided through the Excel client due to limitations of the web based charting capabilities.

VV Tiers Report

The VV Tiers report shows the distribution of space in each VV among the tiers in an Adaptive Optimization configuration, in absolute terms and also as a percentage (Figure 65 (page 187)).

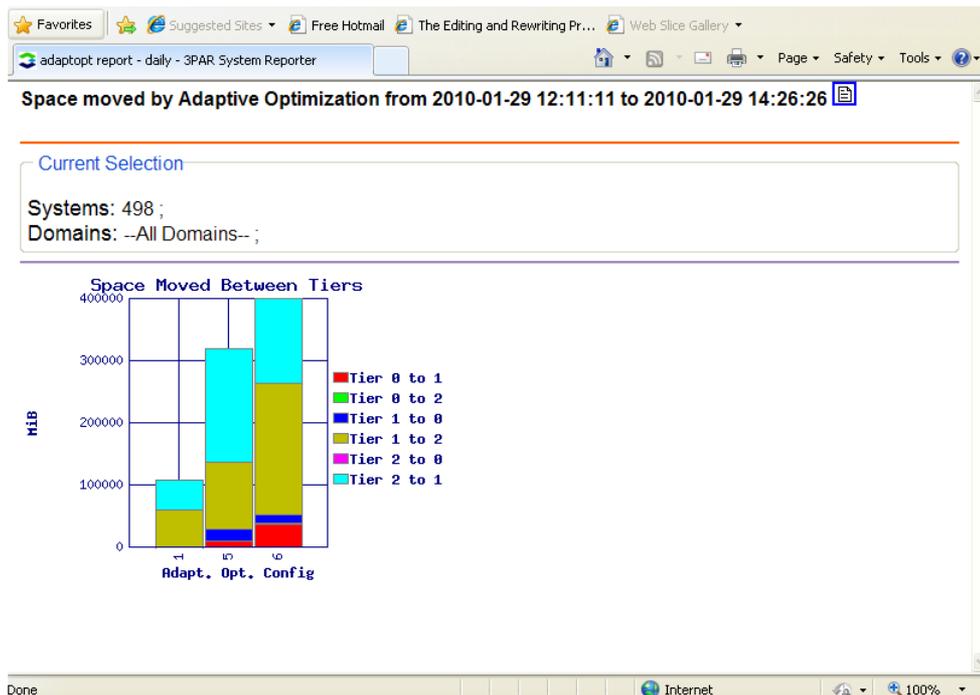
Figure 65 VV Tiers Report



Adaptive Optimization Space Moved Report

The Space Moved report shows the space moved over a selected time interval for each Adaptive Optimization configuration (Figure 66 (page 187)).

Figure 66 Space Moved Report



The space moved between tiers represents a performance overhead for Adaptive Optimization since it takes system resources and disk bandwidth to move data between tiers. The Space Moved

report shows the space moved over the selected time interval for each Adaptive Optimization configuration.

The showvvcpg CLI extension

Adaptive Optimization may move regions of a VV into different CPGs. In order to more easily see how much of each VV is in each tier, the CD for System Reporter 2.8 includes a CLI script file, `showvvcpg.tcl`, that adds a new CLI command when it is sourced from the CLI.

Once you start the CLI and connect to the HP 3PAR Storage System array, you can source the `showvvcpg.tcl`:

```
cli% source showvvcpg.tcl
```

This adds a new `showvvcpg` command to the CLI. You can get the help text for the command as follows:

```
cli% showvvcpg -h
```

```
Usage: showvvcpg [options] [<vvname|pattern> ...]
```

The [options] include

```
-domain <domain_name|pattern>[,<domain_name|pattern>...]
```

Show only VVs that are in domains with names matching one or more of the specified `domain_name` or patterns. Only VVs within domains that the user is a member of can be listed. Multiple domain names or patterns can be repeated using a comma-separated list (for example `-domain <domain_name>,<domain_name>...`).

```
-sortcol <col>[,<dir>][:<col>[,<dir>]...]
```

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

```
inc
```

Sort in increasing order (default).

```
dec
```

Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

An example of the output of the command is shown below, where we have chosen to show only the VVs in the domain ESX:

```
cli% showvvcpg -domain ESX
```

```
----- (MB) ----- -New (MB) --
```

```
Name CPG Adm Snp Usr Adm Snp Usr
```

```
esx-ui esx-fc-r1 352 - 161152 - - -
```

```
esx-ui esx-ssd-r1 160 - 11904 - - -
```

```
esx-st esx-fc-r1 384 - 35584 - - -
```

```
esx-st esx-ssd-r1 128 - 2304 - - -
```

```
esx-sysrpt esx-fc-r1 384 - 112128 - - -
```

```

esx-sysrpt esx-ssd-r1 128 - 7680 - - -
esx-sqa esx-fc-r1 384 - 44544 - - -
esx-sqa esx-ssd-r1 128 - 1536 - - -
esx-sw esx-fc-r1 224 - 249088 - - -
esx-sw esx-ssd-r1 288 - 16512 - - -
esx-rdm-mysqlserver esx-fc-r1 384 1024 44032 - - -
esx-rdm-mysqlserver esx-ssd-r1 128 - - - - -
esx-rdm-dev-oracle esx-fc-r1 384 1024 78848 - - -
esx-rdm-dev-oracle esx-ssd-r1 128 - 7168 - - -
esx-rdm-sysreporterfield esx-fc-r1 320 - 470400 - - -
esx-rdm-sysreporterfield esx-ssd-r1 448 - 896 - - -
-----
8 - 4352 2048 1243776 0 0 0
498 cli%

```

The columns have the following meanings:

- **Name** -- The VV name.
- **CPG** -- The CPG name.
- **Adm** -- There are two columns with this name. The first is the Adm (administration meta data) space in MiB, the second is the new Adm space in MiB. The new space is usually empty unless there are regions currently being moved. While it is being moved from one CPG (tier) to another, a region occupies space in both source and destination (new) CPGs. Once the region movement is complete, the space is no longer shown in the New column.
- **Snp** --The snapshot copy on write space in MiB. For the same reasons as the Adm spaces above, there is an additional Snp column for new space.
- **Usr** -- The user space in MiB. For the same reasons as the Adm spaces above, there is an additional Usr column for new space.

10 Database Schema

This chapter documents the HP 3PAR System Reporter database schema.

Using The Database Schema

The description of the database schema is provided as a convenience to those wishing to query the database directly instead of using the web queries described in “Web Queries” (page 126). However, if you wish to write your own applications that directly query the database, it is important to note the following caution:

CAUTION: The database schema may change between releases of the System Reporter and users must be prepared to revise their applications accordingly.

HP does not provide direct assistance in generating SQL queries and only users who are confident of their own database expertise to use the schema description in this chapter should attempt to use it to create applications.

Users may only issue queries that read the database, not alter it.

Table Versions and Naming Conventions

All tables are versioned and their version number appears at the end of the table names.

Administrative tables contain only the base name and version number. For example **policy_1** is the policy table and it is version 1.

The sample data tables and the inventory tables include the base name, sampling resolution and the version number. For example, **pdspace_hourly_3** contains the PD (physical disk) space data for hourly resolution samples and the table is version 3. Similarly, the **pdspaceinv_hourly_2** table contains the inventory of the distinct PDs in the **pdspace_hourly_3** table and is of revision 2. Note that the revision number of the inventory table need not be the same as the revision number of the data table whose inventory it contains.

When System Reporter is upgraded some tables may need to be converted to new version numbers. This is done automatically by the System Reporter when the sampling service starts.

Administrative Tables

The administrative tables are used to control the behavior of System Reporter and to keep track of the systems that are to be sampled.

The policy Table

The **policy_1** table stores the names and values of the sampling policy parameters. There are a fixed number of rows, one per parameter. The parameters are described in Table 54 (page 112).

Table 59 policy_1 Table

Column	Type	Comment
name	VARCHAR(32) PRIMARY KEY	Name of the policy parameter
val	VARCHAR(255)	Value of the policy parameter

The sampleinserv Table

The **sampleinserv_5** table contains a row for each of systems that the System Reporter must sample.

Table 60 sampleinserv_5 Table

Column	Type	Comment
name	VARCHAR(64) NOT NULL	IP address of the system to sample.
usesssl	INTEGER NOT NULL	Determines whether CLI connections to the InServ will use SSL. 0 - Do not use SSL 1 - Use SSL
cliport	INTEGER NOT NULL	Currently unused. In future this will be used to explicitly specify the port number to use for the CLI connection to the system.
noldperf	INTEGER NOT NULL	0 - Collect LD performance data 1 - Do not collect LD performance data.
disable	INTEGER NOT NULL	0 - Data collection for the system is enabled 1 - Data collection for the system is disabled
sample_chrg	INTEGER NOT NULL	0 - Do not collect region level data for Adaptive Optimization 1 - Collect region level data for Adaptive Optimization
HIGHRESHRS	INTEGER NOT NULL	Number of hours the hires sample should be retained for a particular system.
HOURLYDAYS	INTEGER NOT NULL	Number of days the hourly sample should be retained for a particular system.
DAILYDAYS	INTEGER NOT NULL	Number of days the daily sample should be retained for a particular system.
RGCHHOURLYDAYS	INTEGER NOT NULL	Number of days the hourly sample for region stats should be retained for a particular system.

The system Table

The **system_4** table contains a row for each system that has been sampled. A single physical system may have multiple entries in this table because any changes in the columns for that system will result in a new row in the table.

Table 61 system_4 Table

Column	Type	Comment
uniq_id	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for the system instance. All samples in the data tables have this uniq_id in the sys_uid column.
serial	INTEGER NOT NULL	Serial number of the system.
name	VARCHAR(32) NOT NULL	InServ system name, This is what gets used in most reports to identify the system.
mhz	INTEGER NOT NULL	The CPU frequency in MHz. For os_rev prior to 2.2.2, the MHz is used to calculate the times for the service time histogram buckets. For os_rev of 2.2.2 or later, the service time histogram buckets are in terms of absolute time in milliseconds rather than relative to CPU frequency. See “Service Time Histogram Buckets” (page 206) .
ip_name	VARCHAR(32) NOT NULL	IP Address of the system.
os_rev	VARCHAR(32) NOT NULL	InForm OS revision. This is needed in some cases to determine how to interpret some data. For example, see the mhz column description above.

Table 61 system_4 Table (continued)

Column	Type	Comment
system_model	VARCHAR(32) NOT NULL	System hardware model, for example: InServ S400X InServ T400
CHSIZEMB	INTEGER NOT NULL	Indicates the chunklet size.

The alertconfig Table

The **alertconfig_2** table contains a row for each alert configuration rule.

Table 62 alertconfig_2 Table

Column	Type	Comment
alertid	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each rule.
data_table	VARCHAR(32) NOT NULL	The data table on which the rule operates. Currently supported values are: statport statport_1g statport_2g statport_4g statport_8g statport_hostnodesum statport_nodesum statport_disknodesum statport_syssum statpd, statpd_nl7 stapd_fc10 statpd_fc15 statpd_ssd150 statpd_sum statpd_nl7sum statpd_fc10sum statpd_fc15sum statpd_ssd150sum statid statvlun statcpu
resolution	VARCHAR(32) NOT NULL	The sample resolution, one of hires , hourly or daily .
inserv	VARCHAR(32) NOT NULL	The system name that the alert rule is applied to. If this is the empty string (blank) then the alert rule is applied to all systems. If this is not the empty string, it should match one of the values in the name column of the system_3 table.
metric	VARCHAR(32) NOT NULL	The metric that is calculated. The supported values depend on the data_table value.
direction	VARCHAR(32) NOT NULL	How the metric is compared against the limit_val . Currently this should be either > (greater than) or < (less than).
limit_val	DOUBLE PRECISION NOT NULL	The calculated metric is compared based on direction against limit_val .
limit_count	INTEGER NOT NULL	In each sample interval if the number of objects for which the metric values exceed limit_val when compared using direction is greater than limit_count then an email alert is generated.

Table 62 alertconfig_2 Table (continued)

Column	Type	Comment
condition_code	VARCHAR(32)	Defines an additional condition that must be met for certain alerts to be generated. See “Condition” (page 121).
condition_val	DOUBLE PRECISION NOT NULL	Defines a value used in conjunction with condition_code. See “Condition Value” (page 121).
recipient	VARCHAR(64)	The email address to which alerts for this rule should be sent.

The scheduled_report Table

The **scheduled_report_2** table contains a row for each scheduled report and defines where the report should be stored, when it should be run and parameters used to generate the report.

Table 63 scheduled_report_2 Table

Column	Type	Comment
rpt_id	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each scheduled report.
rpt_dir	VARCHAR(64) NOT NULL	The top level report directory (under the <code>3par/scheduledreports/</code> directory).
rpt_name	VARCHAR(64) NOT NULL	The report name. A sub-directory of this name is created (under the <code>rpt_dir</code> directory) in which all scheduled reports are created.
rpt_sched	VARCHAR(200) NOT NULL	The schedule for this report. The format is a space-separated list of three items: the date list, the weekday list and the hour list. Each of these three lists is either a “*” or a comma-separated list of individual dates, weekdays and hours respectively.
rpt_url	VARCHAR(4000) NOT NULL	The report’s URL parameters (after the “?” in the URL).
rpt_active	INTEGER NOT NULL	0 - Report is not active. 1 - Report is active.
rpt_createtime	VARCHAR(64) NOT NULL	The time when this scheduled report was created.
rpt_retentiondays	INTEGER NOT NULL	The number of days for which created reports should be kept. Reports that are older are automatically deleted.
RPT_EMAIL	VARCHAR(1000) COMMA SEPARATED	The email address for delivery of scheduled reports.

The scheduled_report_log Table

The **scheduled_report_log_1** table is used to keep track of when each scheduled report has run.

Table 64 scheduled_report_log_1 Table

Column	Type	Comment
rpt_id	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each scheduled report.
rpt_run_start_time	VARCHAR(64) NOT NULL	The time the report was started.
rpt_run_end_time	VARCHAR(64) NOT NULL	The time the report finished executing.
rpt_run_hr	INTEGER	The hour that the report ran.
rpt_run_status	INTEGER	The status for the report.

The adprovconfig Table

The **adprovconfig_1** table has a row for each adaptive optimization configuration.

Table 65 adprovconfig_1 Table

Column	Type	Comment
adprovid	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each adaptive optimization configuration.
sysname	VARCHAR(32) NOT NULL	The system name.
t0_cpg_name	VARCHAR(32) NULL	Tier 0 CPG name. If there is no tier 0, this should be the empty string.
t1_cpg_name	VARCHAR(32) NULL	Tier 1 CPG name. If there is no tier 1, this should be the empty string.
t2_cpg_name	VARCHAR(32) NULL	Tier 2 CPG name. If there is no tier 2, this should be the empty string.
ap_sched	VARCHAR(200) NOT NULL	The schedule for when this adaptive optimization configuration should be run. The format is a space-separated list of three items: the date list, the weekday list and the hour list. Each of these three lists is either a "*" or a comma-separated list of individual dates, weekdays and hours respectively.
stats_hours	INTEGER NOT NULL	The number of hours for which the region statistics should be considered.
ao_mode	VARCHAR(32) NOT NULL	One of three modes: <ul style="list-style-type: none"> • Performance - Configuration is biased to move data towards faster tiers. • Balanced - Configuration determines appropriate distribution between Performance and Cost. • Cost - Configuration is biased to move data towards slower (less costly) tiers.
ap_active	INTEGER NOT NULL	0 - Configuration not active. 1 - Configuration is active.

The adprov_log Table

The **adprov_log_1** table is used to keep track of when each adaptive optimization configuration has been executed.

Table 66 adprov_log_1 Table

Column	Type	Comment
adprovid	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each adaptive optimization configuration.
start_tsecs	INTEGER NOT NULL	Time in secs when the adaptive optimization configuration was started. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
end_tsecs	INTEGER NOT NULL	Time in secs when the adaptive optimization configuration ended. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
status	VARCHAR(64) NOT NULL	The status can be one of the following: <ul style="list-style-type: none"> • started - Execution has started for the configuration. • done - Execution completed without error for the configuration. • error - Execution encountered an error.

The ldrgmoves Table

The **ldrgmoves_1** table is used to keep track of the LD region moves generated for each adaptive optimization configuration that has been executed.

Table 67 ldrgmoves_1 Table

Column	Type	Comment
sched_tsecs	INTEGER NOT NULL	Time in secs when the LD region move was scheduled. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance that refers to the uniq_id for the system instance in the system_3 table.
adprovid	INTEGER PRIMARY KEY (AUTO_INCREMENT)	A unique identifier for each adaptive optimization configuration.
ldid	INTEGER NOT NULL	The logical disk ID of the source LD region.
ld_name	VARCHAR(32) NOT NULL	The LD name of the source LD region.
ldoffmb	INTEGER NOT NULL	Offset in MiBs of the source LD region.
vvid	INTEGER NOT NULL	The VV ID for the source LD region.
vvspace	VARCHAR(8) NOT NULL	The VV space of the source LD region.
vvoffmb	INTEGER NOT NULL	Offset in MiBs in the VV space of the region.
src_tier	INTEGER NOT NULL	Specifies the source tier.
dst_tier	INTEGER NOT NULL	Specifies the destination tier.
dst_cpgname	VARCHAR(32) NOT NULL	Specifies the destination CPG name.
move_reason	VARCHAR(32) NOT NULL	The reason for the move, can be one of the following: <ul style="list-style-type: none"> • busy_space - Busy region moved to faster tier because of lack of space in source tier. • busiest - Busiest region moved to faster tier since it met movement criteria. • idle_space - Idle region moved to slower tier because of lack of space in source tier. • idlest - Idlest region moved to slower tier since it met movement criteria. • zero - Regions not accessed at all moved to slowest tier.
status	VARCHAR(64) NOT NULL	The status can be one of the following <ul style="list-style-type: none"> • scheduled - The LD region move has been scheduled. • issued - The CLI command to execute the LD region move has been issued. • done - The LD region move has completed. • cancelled - The LD region move has been cancelled.
done_tsecs	INTEGER NOT NULL	Time in secs when the LD region move was done. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.

The recentsample Table

The **recentsample_1** table contains a row for each sample resolution containing the time of the most recent sample for that resolution.

Beginning with the 2.4.1 release System Reporter inserts data from each system into the database separately. Consequently, a given data table at a specific time may contain data from only some of the systems since the remaining systems' data may not have been inserted yet. In order to avoid reporting on such partial samples, System Reporter updates the recentsample table with the time of the most recent complete sample for each resolution.

Table 68 recentsample_1 Table

Column	Type	Comment
resolution	VARCHAR(32) NOT NULL	The sample resolution, one of hires , hourly or daily .
tsecs	INTEGER NOT NULL	Sample time in secs of the most recent complete sample of the resolution. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.

The Sample Data and Inventory Tables

The sample data tables include the space and performance samples of various objects in the systems. For each type of sample, there are three separate tables, one for each sampling resolution:

- **hires** - High resolution samples
- **hourly** - Hourly samples
- **daily** - Daily sample.

The tables for the different resolutions are identical in their structure. In the table names, the substring <res> can be replaced by either **hires**, **hourly** or **daily**.

The database often needs to be queried to determine all the objects that are present in the table. Often, this information is used to populate the menus in a web browser or Excel client. Since the data tables can be very large, this type of query can take a long time in spite of indexes on the data tables. To improve the efficiency of these queries, a separate set of inventory tables is maintained, one for each sample data table. The inventory tables contain a row for each distinct object in the corresponding sample data table. When the sampler inserts new samples into a table, or removes older samples from a table, it also updates the corresponding inventory table.

The pdspace Table

The **pdspace_<res>_3** table holds physical disk (PD) space data samples.

Table 69 pdspace_<res>_3 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the InServ system instance refers to the uniq_id for the system instance in the system_3 table
pdid	INTEGER NOT NULL	The physical disk ID.
cageid	INTEGER NOT NULL	The cage ID that the PD is in
cageside	INTEGER NOT NULL	The side of the disk chassis that the cage is in.
mag	INTEGER NOT NULL	The magazine number in the cage that the PD is in.
disk	INTEGER NOT NULL	The disk number in the magazine.
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are FC - Fibre Channel drive. NL - Nearline drive.
diskspeed	INTEGER NOT NULL	The approximate disk speed in K RPM
state	VARCHAR(32) NOT NULL	The state of the PD.
wwn	CHAR(16) NOT NULL	The WWN of the PD.
NRM_USED_OK	INTEGER NOT NULL	Number of normal chunklets that are used in LDs.
NRM_USED_FAIL	INTEGER NOT NULL	Number of normal chunklets that were used in LDs but have failed.

Table 69 pdspace_<res>_3 Table (continued)

Column	Type	Comment
NRM_AVAIL_CLEAN	INTEGER NOT NULL	Number of normal unused chunklets that have been initialized. These are free for use in new LDs.
NRM_AVAIL_DRTY	INTEGER NOT NULL	Number of normal unused chunklets that have not yet been initialized. These will be free for use in new LDs after the system has had time to initialize them.
NRM_AVAIL_FAIL	INTEGER NOT NULL	Number of normal unused chunklets that have failed.
SPR_USED_OK	INTEGER NOT NULL	Number of spare chunklets that are used in LDs.
SPR_USED_FAIL	INTEGER NOT NULL	Number of spare chunklets that were used in LDs but have failed.
SRP_AVAIL_CLEAN	INTEGER NOT NULL	Number of spare unused chunklets that have been initialized. These are free for use in new LDs.
SRP_AVAIL_DRTY	INTEGER NOT NULL	Number of spare unused chunklets that have not yet been initialized. These will be free for use in new LDs after the system has had time to initialize them.
SRP_AVAIL_FAIL	INTEGER NOT NULL	Number of spare unused chunklets that have failed.

The pdspaceinv Table

The **pdspaceinv_<res>_2** table holds the inventory for the **pdspace_<res>_3** table.

Table 70 pdspaceinv_<res>_21Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs of the most recent sample for the PD in the corresponding data table. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
pdid	INTEGER NOT NULL	The physical disk ID.
cageid	INTEGER NOT NULL	The cage ID that the PD is in
cageside	INTEGER NOT NULL	The side of the disk chassis that the cage is in.
mag	INTEGER NOT NULL	The magazine number in the cage that the PD is in.
disk	INTEGER NOT NULL	The disk number in the magazine.
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are <ul style="list-style-type: none"> • FC - Fibre channel drive. • NL - Nearline drive.
diskspeed	INTEGER NOT NULL	The approximate disk speed in K RPM.
wwn	CHAR(16) NOT NULL	The WWN of the PD.

The Idspace Table

The **ldspace_<res>_3** table holds logical disk (LD) space data samples.

Table 71 ldspace_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the LD belongs to. If the LD is not assigned to a domain, this column contains the value "-".
ldid	INTEGER NOT NULL	The logical disk ID.
ld_name	VARCHAR(32) NOT NULL	The LD name.
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are: <ul style="list-style-type: none"> • FC - Fibre channel drive. • NL - Nearline drive.
owner	INTEGER NOT NULL	The node number of the owner of the LD.
backup_flag	INTEGER NOT NULL	The node number of the backup owner of the LD.
raid_type	INTEGER NOT NULL	The RAID type of the LD
set_size	INTEGER NOT NULL	The set size of the LD.
step_size	INTEGER NOT NULL	Step size of the LD.
row_size	INTEGER NOT NULL	Row size of the LD.
rawsizemb	INTEGER NOT NULL	Raw size of the LD in MB.
sizemb	INTEGER NOT NULL	Size of the LD in MB.
usedmb	INTEGER NOT NULL	Used MB in the ID.
cpg_name	VARCHAR(32) NOT NULL	The name of the CPG that the LD belongs to. If the LD does not belong to a CPG the name is "-".
ld_use	VARCHAR(32) NOT NULL	Describes what the LD is used for. The possible values are described in Table 28 (page 88) .
ld_state	VARCHAR(32) NOT NULL	The state of the LD.
avail	VARCHAR(32) NOT NULL	The intended availability of the LD.
cavail	VARCHAR(32) NOT NULL	The current (actual) availability of the LD.

The ldspaceinv Table

The **ldspaceinv_<res>_2** table holds the inventory for the **ldspace_<res>_2** table.

Table 72 ldspaceinv_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the LD belongs to. If the LD is not assigned to a domain, this column contains the value "-".
ldid	INTEGER NOT NULL	The logical disk ID.
ld_name	VARCHAR(32) NOT NULL	The LD name.
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are <ul style="list-style-type: none"> • FC - Fibre channel drive. • NL - Nearline drive.
owner	INTEGER NOT NULL	The node number of the owner of the LD.
backup_flag	INTEGER NOT NULL	The node number of the backup owner of the LD.
raid_type	INTEGER NOT NULL	The RAID type of the LD.
cpg_name	VARCHAR(32) NOT NULL	The name of the CPG that the LD belongs to. If the LD does not belong to a CPG the name is "-".
ld_use	VARCHAR(32) NOT NULL	Describes what the LD is used for. The possible values are described in Table 28 (page 88) .

The vvspace Table

The **vvspace_<res>_3** table holds virtual volume (VV) space data samples.

Table 73 vvspace_<res>_3 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VV belongs to. If the VV is not assigned to a domain, this column contains the value "-".
vvid	INTEGER NOT NULL	The virtual volume ID.
bsid	INTEGER NOT NULL	The ID of the base VV for the VV.
vv_name	VARCHAR(32) NOT NULL	The VV name.
wwn	CHAR(16) NOT NULL	The WWN of the VV.
snp_cpg_name	VARCHAR(32) NOT NULL	The name of the CPG, if any, that is associated with the VV's snapshot space. If the VV snapshot space is not associated with a CPG, this column is "-".
usr_cpg_name	VARCHAR(32) NOT NULL	The name of the CPG, if any, that is associated with the VV's user space. If the VV user space is not associated with a CPG, this column is "-".

Table 73 vvspace_<res>_3 Table (continued)

Column	Type	Comment
provv_type	VARCHAR(32) NOT NULL	The provisioning type of the VV (see the <code>PROV</code> column for the <code>showvv</code> CLI command in InForm OS 2.3.1 or later).
vv_type	VARCHAR(32) NOT NULL	The VV Type (see the <code>TYPE</code> column for the <code>showvv</code> CLI command for InForm OS 2.3.1 or later).
copyof	VARCHAR(32) NOT NULL	The name of the VV that this VV is a copy of.
vv_state	VARCHAR(32) NOT NULL	The state of the VV.
createtime	VARCHAR(32) NOT NULL	The creation time for the VV.
spacecalctime	VARCHAR(32) NOT NULL	Only useful for snapshots. The most recent time that <code>updatesnapspace</code> was run for the VV.
vsize_mb	DOUBLE PRECISION NOT NULL	Virtual size of the VV in MB.
adm_rsvd_mb	DOUBLE PRECISION NOT NULL	Snapshot admin space reserved for the VV in MB.
adm_rawrsvd_mb	DOUBLE PRECISION NOT NULL	Raw (physical space including RAID overhead) snapshot admin space reserved for the VV in MB.
adm_used_mb	DOUBLE PRECISION NOT NULL	Snapshot admin space used by the VV in MB.
adm_zn	DOUBLE PRECISION NOT NULL	Admin zones in the VV.
adm_free_zn	DOUBLE PRECISION NOT NULL	Free admin zones in the VV.
snp_rsvd_mb	DOUBLE PRECISION NOT NULL	Snapshot data space reserved for the VV in MB.
snp_rawrsvd_mb	DOUBLE PRECISION NOT NULL	Raw (physical space including RAID overhead) snapshot data space reserved for the VV in MB.
snp_used_mb	DOUBLE PRECISION NOT NULL	Snapshot data space used by the VV in MB.
snp_zn	DOUBLE PRECISION NOT NULL	Snapshot zones in the VV.
snp_free_zn	DOUBLE PRECISION NOT NULL	Free snapshot zones in the VV.
usr_rsvd_mb	DOUBLE PRECISION NOT NULL	User space reserved for the VV in MB.

Table 73 vvspace_<res>_3 Table (continued)

Column	Type	Comment
usr_rawsvd_mb	DOUBLE PRECISION NOT NULL	Raw (physical space including RAID overhead) user space reserved for the VV in MB.
usr_used_mb	DOUBLE PRECISION NOT NULL	User space used by the VV in MB.
usr_zn	DOUBLE PRECISION NOT NULL	User zones in the VV.
usr_free_zn	DOUBLE PRECISION NOT NULL	Free user zones in the VV.

The vvspaceinv Table

The **vvspaceinv_<res>_3** table holds the inventory for the **vvspace_<res>_3** table.

Table 74 vvspaceinv_<res>_3 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VV belongs to. If the VV is not assigned to a domain, this column contains the value "-".
vvid	INTEGER NOT NULL	The virtual volume ID.
bsid	INTEGER NOT NULL	The ID of the base VV for the VV.
vv_name	VARCHAR(32) NOT NULL	The VV name.
wwn	CHAR(16) NOT NULL	The WWN of the VV.
snp_cpg_name	VARCHAR(32) NOT NULL	The name of the CPG, if any, that is associated with the VV's snapshot space. If the VV snapshot space is not associated with a CPG, this column is "-".
usr_cpg_name	VARCHAR(32) NOT NULL	The name of the CPG, if any, that is associated with the VV's user space. If the VV user space is not associated with a CPG, this column is "-".
provv_type	VARCHAR(32) NOT NULL	The provisioning type of the VV (see the Prov column for the showvv CLI command in InForm OS 2.3.1 or later).
vv_type	VARCHAR(32) NOT NULL	The VV Type (see the Type column for the showvv CLI command for InForm OS 2.3.1 or later).

The cpgspace Table

The **cpgspace_<res>_1** table holds Common Provisioning Group (CPG) space data samples.

Table 75 cpgspace_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the CPG belongs to. If the CPG is not assigned to a domain, this column contains the value ".".
cpgid	INTEGER NOT NULL	The common provisioning group ID.
cpg_name	VARCHAR(32) NOT NULL	The common provisioning group name.
adm_total_mb	DOUBLE PRECISION NOT NULL	The total admin LD space in the CPG in MB.
adm_rawtotal_mb	DOUBLE PRECISION NOT NULL	The total raw (including RAID overhead) admin space in the CPG in MB.
adm_used_mb	DOUBLE PRECISION NOT NULL	The used admin LD space in the CPG in MB.
adm_rawused_mb	DOUBLE PRECISION NOT NULL	The raw (including RAID overhead) used admin LD space in the CPG in MB.
snp_total_mb	DOUBLE PRECISION NOT NULL	The total snapshot LD space in the CPG in MB.
snp_rawtotal_mb	DOUBLE PRECISION NOT NULL	The total raw (including RAID overhead) snapshot space in the CPG in MB.
snp_used_mb	DOUBLE PRECISION NOT NULL	The used snapshot LD space in the CPG in MB.
snp_rawused_mb	DOUBLE PRECISION NOT NULL	The raw (including RAID overhead) used snapshot LD space in the CPG in MB.
usr_total_mb	DOUBLE PRECISION NOT NULL	The total user LD space in the CPG in MB.
usr_rawtotal_mb	DOUBLE PRECISION NOT NULL	The total raw (including RAID overhead) user space in the CPG in MB.
usr_used_mb	DOUBLE PRECISION NOT NULL	The used user LD space in the CPG in MB.
usr_rawused_mb	DOUBLE PRECISION NOT NULL	The raw (including RAID overhead) used user LD space in the CPG in MB.
warn_mb	DOUBLE PRECISION NOT NULL	The warning limit of the CPG in MB.
limit_mb	DOUBLE PRECISION NOT NULL	The allocation (hard) limit of the CPG in MB.
gsize_mb	DOUBLE PRECISION NOT NULL	The growth size of the CPG in MB.
rawfree_mb	DOUBLE PRECISION NOT NULL	The raw free space available for the CPG.
usablefree_mb	DOUBLE PRECISION NOT NULL	The usable free space available for the CPG.

The cpgspaceinv Table

The **cpgspaceinv_<res>_1** table holds the inventory for the **cpgspace_<res>1** table.

Table 76 cpgspaceinv_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the CPG belongs to. If the CPG is not assigned to a domain, this column contains the value "-".
cpgid	INTEGER NOT NULL	The common provisioning group ID.
cpg_name	VARCHAR(32) NOT NULL	The common provisioning group name.

Common Performance Table Columns

The **statport**, **statvlun**, **statld** and **statpd** tables all have a common set of columns for the performance data. These columns are described in [Table 77 \(page 203\)](#).

Table 77 Common Performance Columns

Column	Type	Comment
begin_msec	DOUBLE PRECISION NOT NULL	Timestamp value indicating when the counting started (typically when the system was rebooted or the device was created). Two count samples should only be compared if the begin value is the same for the two samples. For example, the difference between rcount values of two samples represents the number of read operations in the interval between those samples only if the begin values are the same for those samples.
d_begin_msec	DOUBLE PRECISION NOT NULL	Difference between begin value from last sample for the same object.
now_msec	DOUBLE PRECISION NOT NULL	Timestamp value in microseconds. The difference in the now values of two samples should be used in calculating the time interval between two samples
d_now_msec	DOUBLE PRECISION NOT NULL	Difference between now value from last sample for the same object.
qlen	DOUBLE PRECISION NOT NULL	Queue length at the time of the sample. qlen is an instantaneous measure, not an average queue length over the sample period. qlen may not be available for all objects and its value will be negative to indicate that it is invalid.
d_qlen	DOUBLE PRECISION NOT NULL	Difference between qlen value from last sample for the same object.
busy	DOUBLE PRECISION NOT NULL	Time in microsecs that the device has been busy.
d_busy	DOUBLE PRECISION NOT NULL	Difference between busy value from last sample for the same object.
rcount	DOUBLE PRECISION NOT NULL	Total number of reads since the object was first created or restarted.
d_rcount	DOUBLE PRECISION NOT NULL	Difference between rcount value from last sample for the same object.

Table 77 Common Performance Columns *(continued)*

Column	Type	Comment
rbytes	DOUBLE PRECISION NOT NULL	Total number of bytes read since the object was first created or restarted
d_rbytes	DOUBLE PRECISION NOT NULL	Difference between rbytes value from last sample for the same object.
rerror	DOUBLE PRECISION NOT NULL	Total number of read errors since the object was first created or restarted
d_rerror	DOUBLE PRECISION NOT NULL	Difference between error value from last sample for the same object.
rdrops	DOUBLE PRECISION NOT NULL	Total number of dropped reads since the object was first created or restarted.
d_rdrops	DOUBLE PRECISION NOT NULL	Difference between rdrops value from last sample for the same object.
rticks	DOUBLE PRECISION NOT NULL	Total number of ticks (microsecs) taken by read operations since the object was first created or restarted.
d_rticks	DOUBLE PRECISION NOT NULL	Difference between rticks value from last sample for the same object.
wcount	DOUBLE PRECISION NOT NULL	Total number of writes since the object was first created or restarted
d_wcount	DOUBLE PRECISION NOT NULL	Difference between wcount value from last sample for the same object.
wbytes	DOUBLE PRECISION NOT NULL	Total number of bytes written since the object was first created or restarted
d_wbytes	DOUBLE PRECISION NOT NULL	Difference between wbytes value from last sample for the same object.
werror	DOUBLE PRECISION NOT NULL	Total number of write errors since the object was first created or restarted
d_werror	DOUBLE PRECISION NOT NULL	Difference between werror value from last sample for the same object.
wdrops	DOUBLE PRECISION NOT NULL	Total number of dropped writes since the object was first created or restarted
d_wdrops	DOUBLE PRECISION NOT NULL	Difference between wdrops value from last sample for the same object.
wticks	DOUBLE PRECISION NOT NULL	Total number of ticks (microsecs) taken by write operations since the object was first created or restarted
d_wticks	DOUBLE PRECISION NOT NULL	Difference between wticks value from last sample for the same object.
rhist0	DOUBLE PRECISION NOT NULL	Number of reads in service time bucket number 0.
d_rhist0	DOUBLE PRECISION NOT NULL	Difference between rhist0 value from last sample for the same object.
...
rhist31	DOUBLE PRECISION NOT NULL	Number of reads in service time bucket number 31.
d_rhist31	DOUBLE PRECISION NOT NULL	Difference between rhist31 value from last sample for same object.

Table 77 Common Performance Columns *(continued)*

Column	Type	Comment
whist0	DOUBLE PRECISION NOT NULL	Number of writes in service time bucket number 0.
d_whist0	DOUBLE PRECISION NOT NULL	Difference between whist0 value from last sample for the same object.
...
whist31	DOUBLE PRECISION NOT NULL	Number of writes in service time bucket number 31.
d_whist31	DOUBLE PRECISION NOT NULL	Difference between whist31 value from last sample for same object.
rsize0	DOUBLE PRECISION NOT NULL	Number of reads in IO size bucket number 0 (512 bytes).
d_rsize0	DOUBLE PRECISION NOT NULL	Difference between rsize0 value from last sample for the same object.
...
rsize15	DOUBLE PRECISION NOT NULL	Number of reads in IO size bucket number 15 (16MiBytes).
d_rsize15	DOUBLE PRECISION NOT NULL	Difference between rsize15 value from last sample for same object.
wsize0	DOUBLE PRECISION NOT NULL	Number of writes in IO size bucket number 0 (512 bytes).
d_wsize0	DOUBLE PRECISION NOT NULL	Difference between wsize0 value from last sample for the same object.
...
wsize15	DOUBLE PRECISION NOT NULL	Number of writes in IO size bucket number 16 (16MiBytest).
d_wsize15	DOUBLE PRECISION NOT NULL	Difference between wsize15 value from last sample for same object.

Although in principle it was not necessary to store the difference values (**d_rcount** for example), since these can be derived from previous samples of the same object, in practice having the difference values stored in the data tables greatly simplifies the queries and improves the performance of the queries. As described in the comment for the **begin** column in [Table 77](#) (page 203), the differences are only valid when the objects have the same **begin** time.

Calculating Common Performance Metrics

[Table 78](#) (page 205) shows SQL expressions that can be used to calculate various performance metrics from the common performance columns for any sample row.

Table 78 Common Performance Metrics for a Sample Row

Metric	SQL Expression
Read IOPs/s	CASE WHEN d_now > 0 THEN d_rcount * 1000000.0/d_now ELSE 0 END
Write IOPs/s	CASE WHEN d_now > 0 THEN d_wcount * 1000000.0/d_now ELSE 0 END
Total IOPs/s	CASE WHEN d_now > 0 THEN (d_rcount + d_wcount) * 1000000.0/d_now ELSE 0 END
Read KBytes/sec	CASE WHEN d_now > 0 THEN d_rbytes * 1000.0/d_now ELSE 0 END

Table 78 Common Performance Metrics for a Sample Row (continued)

Metric	SQL Expression
Write KBytes/sec	CASE WHEN d_now > 0 THEN d_wbytes * 1000.0/d_now ELSE 0 END
Total KBytes/sec	CASE WHEN d_now > 0 THEN (d_rbytes + d_wbytes)* 1000.0/d_now ELSE 0 END
Read service time (millisec)	CASE WHEN d_rcount > 0 THEN d_rticks / (d_rcount * 1000.0) ELSE 0 END
Write service time (millisec)	CASE WHEN d_wcount > 0 THEN d_wticks / (d_wcount * 1000.0) ELSE 0 END
Total service time (millisec)	CASE WHEN (d_rcount + d_wcount) > 0 THEN (d_rticks + d_wticks) / ((d_rcount + d_wcount) * 1000.0) ELSE 0 END
Read IO Size (KBytes)	CASE WHEN d_rcount > 0 THEN d_rbytes / (d_rcount * 1000.0) ELSE 0 END
Write IO Size (KBytes)	CASE WHEN d_wcount > 0 THEN d_wbytes / (d_wcount * 1000.0) ELSE 0 END
Total IO Size (KBytes)	CASE WHEN (d_rcount + d_wcount) > 0 THEN (d_rbytes + d_wbytes) / ((d_rcount + d_wcount) * 1000.0) ELSE 0 END
Queue length	CASE WHEN qlen > 0 THEN qlen ELSE 0 END The check for qlen > 0 is because qlen < 0 indicates that qlen is invalid for that sample.
Busy%	AVG(CASE WHEN ((d_now > 0) AND (d_busy < busy)) THEN (d_busy * 100.0/d_now) ELSE 0 END)

Service Time Histogram Buckets

The service time range corresponding to each of the rhist and whist buckets depends on the version of the InFormOS (**os_rev** column in the system table, see “The system Table” (page 191)). For **os_rev** prior to 2.2.2, the nth bucket (**rhistn** and **whistn**) correspond to service times from the previous bucket to $(2^n * 4.096 / \text{mhz})$ millisec where **mhz** is the mhz column in the **system_3** table (see “The system Table” (page 191)). This means that service time histograms from systems with different CPU frequencies could not be correctly aggregated. Beginning with the 2.2.2 release, the buckets were made independent of the CPU frequency, and the nth bucket (**rhistn** and **whistn**) corresponds to service times from the previous bucket to $(2^n - 7)$ millisec.

IO Size Histogram Buckets

There are 16 IO Size buckets correspond IO sizes 512 bytes to 16MiBytes. The nth bucket holds the count for IO Sizes from the previous bucket’s max size to $512 * 2^n$ bytes.

The statpd Table

The **statpd_<res>_4** table holds physical disk (PD) performance data samples.

Table 79 statpd_<res>_4 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
pdid	INTEGER NOT NULL	The physical disk ID.
port_n	INTEGER NOT NULL	The node that the disk’s primary port is connected to (the “n” part of n:s:p).

Table 79 statpd_<res>_4 Table (continued)

Column	Type	Comment
port_s	INTEGER NOT NULL	The PCI slot in the node that the disk's primary port is connected to (the "s" part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the disk's primary port is connected to (the "p" part of n:s:p).
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are <ul style="list-style-type: none"> • FC - Fibre channel drive. • NL - Nearline drive.
diskspeed	INTEGER NOT NULL	The approximate disk speed in K RPM.
(common performance columns)		The rest of the columns are the common performance columns described in "Common Performance Table Columns" (page 203)).

The statpdinv Table

The **statpdinv_<res>_2** table holds inventory for the statpd_<res>_3 table.

Table 80 statpdinv_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
pdid	INTEGER NOT NULL	The physical disk ID.
port_n	INTEGER NOT NULL	The node that the disk's primary port is connected to (the "n" part of n:s:p).
port_s	INTEGER NOT NULL	The PCI slot in the node that the disk's primary port is connected to (the "s" part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the disk's primary port is connected to (the "p" part of n:s:p).
disktype	VARCHAR(32) NOT NULL	The disk type. Currently, the disk types are: <ul style="list-style-type: none"> • FC - Fibre channel drive. • NL - Nearline drive.
diskspeed	INTEGER NOT NULL	The approximate disk speed in K RPM.

The statld Table

The **statld_<res>_5** table holds logical disk (LD) performance data samples.

Table 81 statld_<res>_5 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the LD belongs to. If the LD is not assigned to a domain, this column contains the value "-".
ldid	INTEGER NOT NULL	The logical disk ID.

Table 81 statld_<res>_5 Table (continued)

Column	Type	Comment
node	INTEGER NOT NULL	The node that owns the LD
cpg_name	VARCHAR(32) NOT NULL	The name of the CPG that the LD belongs to. If the LD is not in a CPG, this column contains the value "----".
(common performance columns)		The rest of the columns are the common performance columns described in "Common Performance Table Columns" (page 203)).

The statldinv Table

The **statldinv_<res>_4** table holds the inventory for the statld_<res>_5 table

Table 82 statldinv_<res>_4 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the LD belongs to. If the LD is not assigned to a domain, this column contains the value ".".
ldid	INTEGER NOT NULL	The logical disk ID.
node	INTEGER NOT NULL	The node that owns the LD.
cpg_name	VARCHAR(32) NOT NULL	The name of the CPG that the LD belongs to. If the LD is not in a CPG, this column contains the value "----".

The statport Table

The **statport_<res>_4** table holds port performance data samples.

Table 83 statport_<res>_4 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
port_n	INTEGER NOT NULL	The node that the port is connected to (the "n" part of n:s:p).
port_s	INTEGER NOT NULL	The PCI slot in the node that the port is connected to (the "s" part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the port is connected to (the "p" part of n:s:p).
porttype	VARCHAR(32) NOT NULL	The port type. Currently, the port types are: <ul style="list-style-type: none"> • disk - Fibre channel ports that disks are connected to. • host - Fibre channel ports that hosts are connected to (directly or through a fabric). • rcip - Remote Copy IP ports. • rfc - Remote Copy fibre channel ports.

Table 83 statport_<res>_4 Table (continued)

Column	Type	Comment
		<ul style="list-style-type: none"> iscsi - iSCSI host ports. free - ports that are not connected.
(common performance columns)		The rest of the columns are the common performance columns described in "Common Performance Table Columns" (page 203).

The statportinv Table

The **statportinv_<res>_4** table holds the inventory for the **statport_<res>_4** table.

Table 84 statportinv_<res>_4 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
port_n	INTEGER NOT NULL	The node that the port is connected to (the "n" part of n:s:p).
port_s	INTEGER NOT NULL	The PCI slot in the node that the port is connected to (the "s" part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the port is connected to (the "p" part of n:s:p).
porttype	VARCHAR(32) NOT NULL	The port type. Currently, the port types are: <ul style="list-style-type: none"> disk - fiber channel ports that disks are connected to. host - fiber channel ports that hosts are connected to (directly or through a fabric). rcip - Remote Copy IP ports. rcfc - Remote Copy fiber channel ports. iscsi - iSCSI host ports. free - ports that are not connected.

The statvlun Table

The **statvlun_<res>_3** table holds VLUN performance data samples.

Table 85 statvlun_<res>_3 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VLUN belongs to. If the VLUN is not assigned to a domain, this column contains the value "-".
vv_name	VARCHAR(32) NOT NULL	The VV name.
host_name	VARCHAR(32) NOT NULL	The host name.
lun	INTEGER NOT NULL	The LUN number.

Table 85 statvlun_<res>_3 Table (continued)

Column	Type	Comment
host_wwn	VARCHAR(128) NOT NULL	The host WWN (world wide name) or iscsi name.
port_n	INTEGER NOT NULL	The node that the port is connected to (the “n” part of n:s:p).
port_s	INTEGER NOT NULL	The PCI slot in the node that the port is connected to (the “s” part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the port is connected to (the “p” part of n:s:p).
(common performance columns)		The rest of the columns are the common performance columns described in “Common Performance Table Columns” (page 203).

The statvluninv Table

The **statvluninv_<res>_4** table holds the inventory for **statvlun_<res>_4** table.

Table 86 statvluninv_<res>_4 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VLUN belongs to. If the VLUN is not assigned to a domain, this column contains the value “.”.
vv_name	VARCHAR(32) NOT NULL	The VV name.
host_name	VARCHAR(32) NOT NULL	The host name.
lun	INTEGER NOT NULL	The LUN number.
host_wwn	VARCHAR(128) NOT NULL	The host WWN (world wide name) or iscsi name.
port_n	INTEGER NOT NULL	The node that the port is connected to (the “n” part of n:s:p).
port_s	INTEGER NOT NULL	The PCI slot in the node that the port is connected to (the “s” part of n:s:p).
port_p	INTEGER NOT NULL	The port number in the PCI slot in the node that the port is connected to (the “p” part of n:s:p).

The statvvcmp Table

The **statvvcmp_<res>_2** table holds VV cache performance data samples.

Table 87 statvvcmp_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VV belongs to. If the VV is not assigned to a domain, this column contains the value “.”.
vv_id	INTEGER NOT NULL	The VV ID.

Table 87 statvcmp_<res>_2 Table (continued)

Column	Type	Comment
vv_name	VARCHAR(32) NOT NULL	The VV name.
rdhit	DOUBLE PRECISION NOT NULL	The number of read cache hits for the VV.
d_rdhit	DOUBLE PRECISION NOT NULL	Difference between rdhit value from last sample for same VV.
rdmiss	DOUBLE PRECISION NOT NULL	The number of read cache misses for the VV.
d_rdmiss	DOUBLE PRECISION NOT NULL	Difference between rdmiss value from last sample for same VV.
wrhit	DOUBLE PRECISION NOT NULL	The number of write cache hits for the VV.
d_wrhit	DOUBLE PRECISION NOT NULL	Difference between wrhit value from last sample for same VV.
wrmiss	DOUBLE PRECISION NOT NULL	The number of write cache misses for the VV.
d_wrmiss	DOUBLE PRECISION NOT NULL	Difference between wrmiss value from last sample for same VV.

The statvcmpinv Table

The **statvcmpinv_<res>_2** table holds the inventory for the **statvcmp_<res>_2** table.

Table 88 statvcmpinv_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
dom_name	VARCHAR(32) NOT NULL	The domain name that the VV belongs to. If the VV is not assigned to a domain, this column contains the value "".
vv_id	INTEGER NOT NULL	The VV ID.
vv_name	VARCHAR(32) NOT NULL	The VV name.

The statcpu Table

The **statcpu_<res>_1** table holds CPU performance data samples.

Table 89 statcpu_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
node	INTEGER NOT NULL	The node number.
cpu	INTEGER NOT NULL	The CPU number in the node.
usert	DOUBLE PRECISION NOT NULL	Jiffies (0.01 sec) in user mode.

Table 89 statcpu_<res>_1 Table (continued)

Column	Type	Comment
d_user	DOUBLE PRECISION NOT NULL	Difference between user value from last sample for same CPU.
syst	DOUBLE PRECISION NOT NULL	Jiffies (0.01 sec) in system mode.
d_syst	DOUBLE PRECISION NOT NULL	Difference between syst value from last sample for same CPU.
idlet	DOUBLE PRECISION NOT NULL	Jiffies (0.01 sec) in idle mode.
d_idlet	DOUBLE PRECISION NOT NULL	Difference between idlet value from last sample for same CPU.
cputicks	DOUBLE PRECISION NOT NULL	The CPU tick count. A single count for all CPUs on a node is maintained in CPU 0, all other CPUs on the node have 0 values for this column.
d_cputicks	DOUBLE PRECISION NOT NULL	Difference between cputicks value from last sample for same CPU.
intr	DOUBLE PRECISION NOT NULL	The number of interrupts for the CPU. A single count for all CPUs on a node is maintained in CPU 0, all other CPUs on the node have 0 values for this column.
d_intr	DOUBLE PRECISION NOT NULL	Difference between intr value from last sample for same CPU.
ctxt	DOUBLE PRECISION NOT NULL	The number of context switches for the CPU. A single count for all CPUs on a node is maintained in CPU 0, all other CPUs on the node have 0 values for this column.
d_ctxt	DOUBLE PRECISION NOT NULL	Difference between ctxt value from last sample for same CPU.

The statcpuinv Table

The **statcpuinv_<res>_1** table holds the inventory for the **statcpu_<res>_1** table.

Table 90 statcpuinv_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
node	INTEGER NOT NULL	The node number.
cpu	INTEGER NOT NULL	The CPU number in the node.

The statlink Table

The **statlink_<res>_2** table holds the inter-node link performance data samples.

Table 91 statlink_<res>_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
node	INTEGER NOT NULL	The node number
node_to	INTEGER NOT NULL	The node number of the destination of the link.
queue	VARCHAR (32) NOT NULL	The name of the Queue.
begin_msec	DOUBLE PRECISION NOT NULL	Timestamp value indicating when the counting started (typically when the system was rebooted or the device was created). Two count samples should only be compared if the begin value is the same for the two samples.
d_begin_msec	DOUBLE PRECISION NOT NULL	Difference between begin value from last sample for the same object.
now_msec	DOUBLE PRECISION NOT NULL	Timestamp value in microseconds. The difference in the now values of two samples should be used in calculating the time interval between two samples.
d_now_msec	DOUBLE PRECISION NOT NULL	Difference between now value from last sample for the same object.
xfers	DOUBLE PRECISION NOT NULL	Total number of transfers on the queue since the node was started.
d_xfers	DOUBLE PRECISION NOT NULL	Difference between xfers value from last sample for the same object.
bytes	DOUBLE PRECISION NOT NULL	Total number of bytes transferred on the queue since the node was started.
d_bytes	DOUBLE PRECISION NOT NULL	Difference between bytes value from last sample for the same object.

The statlinkinv Table

The **statlinkinv_<res>_1** table holds the inventory for the **statlink_<res>_1** table.

Table 92 statlinkinv_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
node	INTEGER NOT NULL	The node number.
node_to	INTEGER NOT NULL	The node number of the destination of the link.
queue	VARCHAR (32) NOT NULL	The name of the Queue.

The statldrg Table

The **statldrg_hourly_2** table holds the LD region performance data. Only hourly samples are kept.

Table 93 statldrg_hourly_2 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.
ldid	INTEGER NOT NULL	The logical disk ID of the LD region.
ldoffmb	INTEGER NOT NULL	Offset in MiB of the LD region.
vvid	INTEGER NOT NULL	The VV ID for the LD region.
vvspace	VARCHAR(8) NOT NULL	The VV space of the LD region.
vvoffmb	INTEGER NOT NULL	Offset in MiB in the VV space of the region.
tstamp	INTEGER NOT NULL	Time stamp in secs for the start of the counts. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
accesses	WRITES DOUBLE PRECISION NOT NULL	Total number of writes on the inServ.

The statldrg_timeinv Table

The **statldrg_timeinv_1** table holds the time inventory for the **statldrg_hourly_2** table.

Table 94 statldrg_timeinv_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance refers to the uniq_id for the system instance in the system_3 table.

The perf_timeinv Table

The **perf_timeinv_<res>_1** table holds the sample times for performance samples for each system.

Table 95 perf_timeinv_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance that refers to the uniq_id for the system instance in the system_3 table.

The space_timeinv Table

The **space_timeinv_<res>_1** table holds the sample times for space samples for each system.

Table 96 space_timeinv_<res>_1 Table

Column	Type	Comment
tsecs	INTEGER NOT NULL	Sample time in secs. This is the Unix time and is the number of seconds since 0:00:00 Jan 1, 1970 GMT.
sys_uid	INTEGER NOT NULL	A unique identifier for the system instance that refers to the uniq_id for the system instance in the system_3 table.

11 Troubleshooting

This chapter describes how to troubleshoot and solve common problems that may occur while setting up or using System Reporter.

Troubleshooting Overview

The first step when troubleshooting System Reporter is to identify the component that is causing problems (see for the structure of System Reporter).

- If there are problems in sampling components see [“Troubleshooting the Sampling Components” \(page 215\)](#). Examples of the types of problems encountered include:
 - The sampler service or daemon stops.
 - The sampler fails to sample data from a particular HP 3PAR Storage System.
- If there are problems with the Adaptive Optimization components, see [“Troubleshooting Adaptive Optimization Components” \(page 217\)](#).
- If there are problems generating interactive reports in a web browser see [“Troubleshooting the Web Server Components” \(page 217\)](#).
- If there are problems using the Excel client, see [“Troubleshooting the Excel Client” \(page 217\)](#). Some problems generating reports in the Excel client may be associated with the web server components as well.
- If there are problems with the database server, see [“Extracting Data from the Database with sysbck” \(page 218\)](#) or [“Dumping the Entire Database” \(page 222\)](#).

Troubleshooting the Sampling Components

The System Reporter sampling components include several processes. Most of these processes create log files that are the first things you should look at when troubleshooting.

Sampling Component Log Files

In general, a sampling component process named `<process_name>` creates a log file named `<process_name>.log`. However, for some processes a separate instance of the process is created for each HP 3PAR Storage System and for these processes the log file is named `<process_name>_<inserv_ip>.log` where `<inserv_ip>` is the IP name or address of the system. On Windows, this is in the HP 3PAR System Reporter directory (default location for the directory is `C:\Program Files\3par\System Reporter\log`). On Linux, the log file is in the `/var/log/sampleloop/` directory.

NOTE: The `<process_name>` excludes the executable file name extensions `.exe` (Windows) or `.bin` (Linux).

For example, the main process (sampleloop) runs as a service on Windows or as a daemon on Linux and it creates a log file called `sampleloop.log`. As shown in [Figure 2 \(page 29\)](#), the sampleloop process creates an `inservsample` process for each HP 3PAR Storage System and these processes create log files such as `inservsample_192.168.1.19.log`.

At the beginning of each day (or whenever a process is restarted), a new `<process_name>.log` is created and the old logs are rolled to `<process_name>.<x>.log`, where `<x>` is 1, 2, 3 or 4. When log files are rolled, the previous `<process_name>.4.log` file is deleted and replaced with the `<process_name>.3.log` file.

When reporting a problem with sampling, please provide the entire log directory.

- △ **CAUTION:** If you get an error in a sampling component that causes the service or daemon to stop, save the entire log directory **before** restarting the service or daemon so that you do not lose the error information in the logs.

Installation Configuration

On Windows, you can see the parameters with which the sampler was installed by typing the following command in that directory (which should contain the `winserv.exe` file) as follows:

```
# winserv.exe showconfig "3PAR System Reporter sampler"
Configuration of 3PAR System Reporter sampler:

-displayname 3PAR System Reporter sampler (managed by WinServ)
-noninteractive
-binary C:\Program Files\3PAR\System Reporter\winserv.exe
-start auto
-errorcontrol normal
-depends
-loadordergroup
-user LocalSystem
-ipcmethod blind
*CommandLine sampleloop.exe -dbtype sqlite -db E:/inservstats/inservstats
-noexpand
*Description: Periodically collects statistics from 3PAR InServ Storage Servers and
inserts the data into the 3PAR System Reporter database.
```

On Linux, the parameters for the `sampleloop` daemon are all in the `/etc/sampleloop.conf` file.

Sampling Errors

When the `inservsample` process for a particular HP 3PAR Storage System `<ip_name>` encounters an error it logs the error in its `inservsample_<ip_name>.log` file and exits. The main `sampleloop` process will detect that the `inservsample` process has exited and log the error in its `sampleloop.log` file. For the next sample, the `sampleloop` process will create a new `inservsample` process to sample data from InServ `<ip_name>`. This new `inservsample` process will roll the `inservsample_<ipname>.log` file.

Since there is a separate `inservsample` process for each system, errors that occur while sampling one HP 3PAR Storage System do not generally impact sampling of other systems. However, in the case of timeout errors, the sampling interval for all the arrays will be stretched to the timeout limit since the `sampleloop` process waits until the timeout limit before it continues to the next sample.

There are two main types of errors incurred by either the `inservsample` process or the CLI sampler process as described in the following sections.

Errors incurred by the `inservsample` process

Errors incurred by the `inservsample` process are typically database errors including:

- Loss of network connectivity to the database server.
- Database server errors. One example when using MySQL as the database is if the `max_allowed_packet` parameter is not sufficiently large (see [“Setting the max_allowed_packet parameter” \(page 40\)](#)) and the `inservsample` process gets the error “MySQL server has gone away”.

Errors incurred by the CLI sampler process.

Errors incurred by the CLI sampler process may include loss of network connectivity to the HP 3PAR Storage System, a timeout because the CLI process took too long to get the data, or authentication failure on the CLI connection (the user was removed or password changed).

In both cases, the sampleloop process restarts the inservsample process so, if the error is not persistent, sampling can resume. If the error is persistent, then the inservsample process will continuously log the error and restart.

Sampler Service/Daemon Errors

Unlike HP 3PAR Storage System sampling errors, when the sampler service or daemon (sampleloop) incurs an error, it causes the service to exit. Before it exits, sampleloop sends an email with the error information if an email address is configured for reporting errors (see [“Configuring Sampling Policies”](#) (page 112)).

Generally, errors in the compactdb, dedupinv and trimdb processes cause these processes to log errors in their respective logs and exit, which in turn causes sampleloop to log the error and exit.

Troubleshooting Adaptive Optimization Components

Adaptive Optimization components include a regionmover process for each HP 3PAR Storage System for which adaptive optimization is enabled. Since the regionmover processes are created by the main sampleloop service or daemon, the `regionmover_<inserv_ip>.log` files are created in the same directory as the sampling components (see [“Sampling Component Log Files”](#) (page 215)).

Troubleshooting the Web Server Components

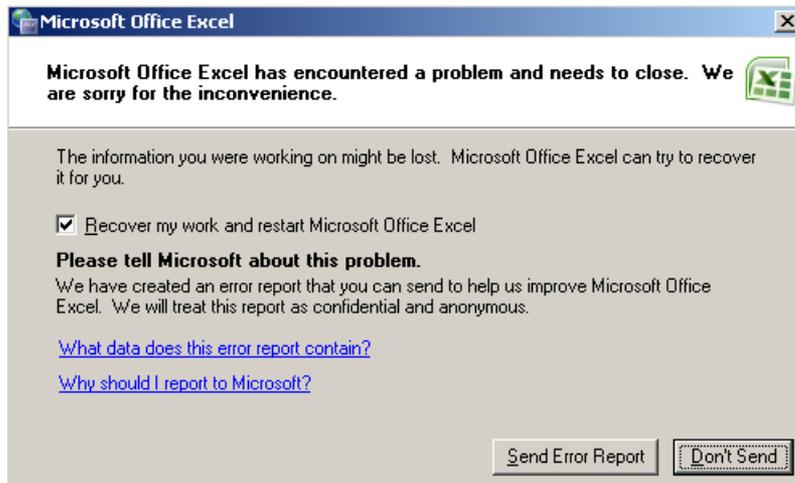
Errors generated by the Apache HTTP Server are logged in the `errors.log` file in the `logs` directory (default location on Windows is `C:\Program Files\Apache Group\Apache2\logs` and the default location on Linux is `/var/log/httpd`).

Troubleshooting the Excel Client

This section contains troubleshooting information for the System Reporter Excel client.

- If clicking on a button displays handles for resizing the button, the security level of Excel has been set too high and macro code is not being allowed to execute. Change the security level to medium or low (see [“Setting the Security Level for Excel”](#) (page 56)) and reopen the workbook.
- If you see a warning containing **"internal state was inconsistent"**, the code has detected an inconsistency and has tried to recover, but may not be able to. If errors continue, go to the **3PAR** menu and click **Reset All**. If errors persist, close and reopen the workbook. Entering and exiting Excel's design mode can cause these inconsistencies and should be avoided.
- Run-time error warnings indicate potential problems with the System Reporter code. To clear the problem, go to the **3PAR** menu and click **Reset All** or close and reopen the workbook. Using a fresh copy of the workbook is recommended.
- If Microsoft Office Excel encounters an internal problem and needs to close ([Figure 67](#) (page 218)), in most of the cases the recovered Excel copy of the workbook may not work as expected.

Figure 67 Microsoft Office Excel Internal Error



If you experience a problem with a recovered copy of the workbook (Figure 68 (page 218)), try starting with a fresh copy of System Reporter workbook.

Figure 68 Microsoft Excel Recovery Error



Extracting Data from the Database with `sysbck`

Sometimes it is necessary to extract some of the data from the System Reporter database. For example, you may want to send that data to HP to help in analyzing a performance problem. This section describes how you can extract all or some of the data from a database using the `sysbck` command into a set of files that can then be loaded into a different database.

NOTE: If necessary, when using `sysbck` to extract large amounts of data from a database, the contents may be saved in separate 2 Gb files.

NOTE: The `sysbck` command is recommended when extracting a relatively small sampling of the database. The command creates a file with SQL statements to insert each row of extracted data into a database. While this is convenient for subsequently loading the data into a database, especially if the database is of a different type than the source database, the file size can be quite large. If you want to extract the entire database or large sections of the database, it is preferable to dump the entire database as described in “[Dumping the Entire Database](#)” (page 222).

On Windows the `sysbck.exe` executable is installed in the System Reporter directory (by default `C:\Program Files\3PAR\System Reporter`), and on Linux the `sysbck.bin` executable is installed in the `/usr/bin/` directory.

The parameters for the `sysbck` command are shown below:

```
# ./sysbck.binUsage:
sysbck [parameters]

The [parameters] are
-dbtype <dbtype>  Type of database: sqlite, mysql or oracle.
```

```

-dbhost <dbhost> For mysql: name or IP address of host.
                  Default is localhost.
                  For oracle: hosts string.
                  Ignored if dbtype is sqlite.
-db <name>        Name of the database.
                  For SQLite this is the full file name, including the path,
                  of the file that contains the database.
-dbpwfile <file> Name of file whose first line is the username and
                  password for the database.
                  Ignored if dbtype is sqlite.
-dbouttype [Optional] <dbtype> Name of the outscript db type
-scripttype [Optional] <script type> Type of the script 'insert/create/both'
-backupfilename The file name along with the path to put the extract in
-begintime [Optional] begin time of the extract in '%m/%d/%YYYY %H:%M:%S' format
-endtime [Optional] end time of the extract in '%m/%d/%YYYY %H:%M:%S' format
-inservername [Optional] name of the inservs to be extracted seperated by ':'

```

NOTE: The `-inservername` parameter specifies the IP Name or Address that was assigned while configuring the HP 3PAR Storage Systems. If you are not sure, consult **Policy Settings → InServ Systems** to get this parameter.

NOTE: The `-dbouttype` parameter should not be used when extracting data for analysis by HP support.

Depending on the `-scripttype` parameter, the `sysbck` command produces one or both of the following types of files:

1. The `<backup>.str` file contains SQL code to create the database structure (tables and indexes).
2. The `<backup>.ins` contains SQL code to insert the data into the tables.

These files can be source as SQL commands to create the schema and insert the data respectively.

To limit the extracted data to one or more systems, use the `-inservername` parameter. To limit the extracted data to a particular time period, use the `-begintime` and `-endtime` parameters.

Examples of Using `sysbck` for Various Databases

Following are examples of how to use the `sysbck` command to generate output for various databases.

NOTE: The `-dbouttype` parameter is optional; by default, the extract database is MySQL.

- From MySQL to Oracle:

```

sysbck -dbtype mysql -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype
oracle -scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservername s074:s138

```

- From MySQL to MS SQL:

```

sysbck -dbtype mysql -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype mssql
-scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservername s074:s138

```

From MS SQL to SQLite:

```
sysbck -dbtype mssql -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype sqlite  
-scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime "07/13/2010  
12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From MS SQL to MySQL:

```
sysbck -dbtype mssql -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype mysql  
-scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From MySQL to MySQL:

```
sysbck -dbtype mysql -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype mysql  
-scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From Oracle to SQLite:

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype  
sqlite -scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From Oracle to MySQL:

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype  
mysql -scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From Oracle to Oracle:

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype  
oracle -scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From SQLite to SQLite:

```
sysbck -dbtype sqlite -db cliuser -dbouttype sqlite -scripttype both  
-backupfilename "c:/test/test_20100713_1322PM" -begintime "07/13/2010 12:45:00"  
-endtime "07/14/2010" -inservname s074:s138
```

- From SQLite to MySQL:

```
sysbck -dbtype sqlite -dbhost istat -db cliuser -dbpwfile dbpwfile -dbouttype  
mysql -scripttype both -backupfilename "c:/test/test_20100713_1322PM" -begintime  
"07/13/2010 12:45:00" -endtime "07/14/2010" -inservname s074:s138
```

- From SQLite to MS SQL:

```
sysbck -dbtype sqlite -db cliuser -dbouttype mssql -scripttype both -backupfilename
"c:/test/test_20100713_1322PM" -begintime "07/13/2010 12:45:00" -endtime
"07/14/2010" -inservname s074:s138
```

- From SQLite to Oracle:

```
sysbck -dbtype sqlite -db cliuser -dbouttype oracle -scripttype both
-backupfilename "c:/test/test_20100713_1322PM" -begintime "07/13/2010 12:45:00"
-endtime "07/14/2010" -inservname s074:s138
```

The following examples show how to use `sysbck` without any optional parameters (Note this will extract .ins data in the MySQL format).

- Backs up the entire database from Oracle:

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM"
```

- Backs up the entire database from MySQL:

```
sysbck -dbtype mysql -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM"
```

- Backs up the entire database from SQLite:

```
sysbck -dbtype sqlite -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM"
```

- Issuing `sysbck` without a system name (this will extract for all systems for the period July 13 2010 till July 14th 2010 00:00:00):

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM" -dbouttype mysql -begintime "07/13/2010" -endtime
"07/14/2010" -inservname s074:s138 -scripttype both
```

- Issuing `sysbck` without a begin and end date (this will extract for all data for system s138):

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM" -dbouttype mysql -inservname s138 -scripttype
both
```

- Issuing `sysbck` without a begin and end date and `dbouttype` (Note this will extract for all data for system s138; in this instance the extract will default to the MySQL extraction):

```
sysbck -dbtype oracle -dbhost istat -db cliuser -dbpwfile dbpwfile -backupfilename
"c:/test/test_20100713_1322PM" -inservname s138 -scripttype both
```

NOTE: When executing `sysbck.exe` from any other directory on Windows, ensure that the `PATH` environment variable is set to the System Reporter root directory path (by default, this path is `C:\Program Files\3par`).

Dumping the Entire Database

When you want to extract a large part of the database, it may be more efficient, both in time and disk space, to dump the entire database instead of using `sysbck`. The method for dumping the entire database depends on the database type. The process for SQLite is described in “[Dumping the Entire SQLite Database](#)” (page 222) while the process for MySQL is described in “[Dumping the Entire MySQL Database](#)” (page 222). For Oracle and MS SQL, please refer to the Oracle or Microsoft documentation.

Dumping the Entire SQLite Database

The SQLite database is just a file (by default `C:\Program Files\3par\System Reporter\inservstats\inservstats.db` on Windows and `/var/inservstats/inservstats` on Linux). However, if there are other files in the same directory created by SQLite, such as a journal file (`inservstats.db-journal`) you cannot just copy the database file to another location and open it as an SQLite database. The presence of the other files, such as the journal file, indicates that the database is not in a consistent state. The recommended way to get a copy is to follow these steps:

1. Stop the HP 3PAR System Reporter sampler service (Windows) or daemon (Linux) so that the database is not active while you are making the copy.
2. If there is no other file (such as the journal file), the database is now consistent and you can just copy the `inservstats.db` file.
3. If there is another file (such as the journal file), the database is still inconsistent. In this case, you must first open the database with SQLite which will automatically apply the journal changes. The easiest way to do this is to generate a report from the database using a web browser. The web server will open the database, the journal changes will be automatically applied, and the journal file will be automatically deleted. Then you can copy the database file.
4. After the database file has been copied, you can re-start the HP 3PAR System Reporter sampler service (Windows) or daemon (Linux).

NOTE: Ideally, when extracting data for HP support from an SQLite database, use `sysbck` without the default `dbouttype` parameter.

Dumping the Entire MySQL Database

To create a dump of a MySQL database use the `mysqldump` program as follows

```
# mysqldump --host=hostname --user=username --password=passwd --opt dbname > dbname.sql
```

This creates a text file called `dbname.sql` that contains SQL commands to create the tables and insert the data into the tables. You can simply source this file into MySQL to recreate the entire database.

NOTE: When dumping data from a large database that is to be backed up to a different location than the source, it is best to use `sysbck` as the data will be stored in multiple 2 GB files (if necessary), and is easier to manage.

Index

Symbols

3PAR System Reporter User's Guide, 54
3PAR System Reporter.xls, 49
3par-policy subdirectory, 125
3par-rpts subdirectory, 125

A

adding a system to the list of systems to be sampled, 115
Administrator's Guide, 25
Apache HTTP Server
 cgi-bin directory, 125
 config.pl files, 125
 Directory directives, 49
 httpd.conf file, 49
 installing, 31
 password, 52
 security settings, 49
 troubleshooting, 217
At Time, 76

B

bar chart (example of), 95

C

cautions, about, 19
CD contents, 25
cgi-bin directory, 125
charts
 charts v tables, 92
Chunklets
 usage types, 106
CLI, 25
CLI Reference, 18
config.pl files, 125
conventions, typographical, 19
CPG Space Reports, 21
CPU Performance Reports, 22
creating the MySQL user and schema, 40

D

daily data table, 68
daily samples, 113
data tables, 67
database
 adding or removing storage servers, 112
 length for retaining samples, 113
 sampling policies, 49
 troubleshooting, 218
database sampling policies
 hi-res samples, 112
Decision Support (DSS)/OLAP, 39
deinstallation, 25
DHCP-assigned address changing, 25
dialog elements, 19
Directory directives (Apache HTTP Server), 49

documentation, related, 18

E

Excel client
 installing, 48
Excel Client link, 49
exporting table data, 18

G

generating reports
 with Web browser, 70

H

hi-res samples, 113
hi-res sampling interval, 112
high-res data table, 68
high-res samples, 112
Histogram At Time, 77
hourly data table, 68
hourly samples, 113
httpd.conf file, 49

I

InForm CLI
 password files, 112
InForm operating environment, 18
InForm OS Command Line Interface, 25
InServ storage servers, 18
inservstats database
 installing, 32, 44
 removing, 51
 uninstalling, 51
installing the System Reporter tools, 43
interfaces, user, 18
IP address, 25

L

LD Performance Reports, 22
LD Space Reports, 21
Lines chart (example of), 94
Link Performance reports, 22

M

Microsoft® Excel®, 30
MySQL Database Server
 creating user and schema, 40
 installation and configuration, 41

N

notes, about, 19

P

PD Performance Reports, 22
PD Space Reports, 21
Physical Planning Manual, 18
policy settings, 49

Port Performance Reports, 21

R

reading of the database, 49

Red Hat Enterprise Linux, 25

Region IO Density Report, 184

related documentation, 18

remove a storage server from the system table, 115

Remove InServ, 47

Report Information control, 92

Report Type control, 76

reports

- displaying results as chart or table, 92

- results, 71

- types of, 76

requirements, 26

S

sample resolution

- resolutions, 67

Sample resolution control, 67

sampling a new storage server, 115

sampling components

- troubleshooting, 215

sampling interval, 112

sampling policies

- editing, 112

sampling service, 29

Scheduled Reports, 122

Space Moved Report, 187

special considerations, 32

Stacked Area chart (example of), 94

Stacked Bars chart (example of), 93

Stacked Vert Bars chart (example of), 93

stop sampling a storage server, 115

storage servers

- adding to database, 47

- removing from the system table, 115

- verifying addition to system table, 115

Summary Report, 21

System Reporter

- 3PAR System Reporter.xls, 49

- 3par-policy subdirectory, 125

- 3par-rpts subdirectory, 125

- components, 28

 - installation and configuration, 25

 - requirements, 26

- configuring the database, 112

- database, 29

- editing sampling policies, 112

- Excel client, 30

- System_Reporter.pdf, 25

- tools, 48

 - installing, 43

- verifying installation, 48

- Web server, 30

- window, 53

 - panes, 53

- with a Web browser, 66

system table

- removing a storage server from, 115

- verifying addition of new storage server, 115

T

tables

- exporting table data, 18

- tables and charts, 92

TCP/IP Networking, 39

troubleshooting

- verifying installation, 48

typographical conventions, 19

U

uninstalling inservstats database, 51

uninstalling System Reporter components, 25

uninstalling the System Reporter components, 50

user interfaces, 18

username and password

- Apache HTTP Server, 49

V

verifying installation, 48

Versus Time, 76

vert bar chart (example of), 96

VLUN Performance Reports, 22

VV Cache Performance Reports, 22

VV Space Reports, 21

VV Tiers Report, 186

W

Web browser

- editing sampling policies, 112

- requirements, 52

- using System Reporter with, 52

Windows® Server, 25