HP Service Health Reporter

for the Windows® operating system

Software Version: 9.20

Performance and Configuration Guide

Document Release Date: December 2012

Software Release Date: October 2012

Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notices

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.
Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.
UNIX® is a registered trademark of The Open Group.
Java is a registered trademark of Oracle and/or its affiliates.
Intel® and Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

Acknowledgements

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

This product includes software developed by Andy Clark.

This product includes asm software which is Copyright (c) 2000-2005 INRIA, France Telecom

All rights reserved.

This product includes jquery.sparkline.js software which is Copyright (c) 2007-2009, Adolfo

Marinucci All rights reserved.

Support

Visit the HP Software Support web site at:

www.hp.com/go/hpsoftwaresupport

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

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

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

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

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

To find more information about access levels, go to:

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

Disclaimer

The performance numbers provided in this document are obtained in a controlled test environment and therefore, may not apply to a customer production environment. Consult with HP before using the performance results and hardware recommendations provided in this document.

Contents

1 Introduction 6

2 Hardware and Software Requirements 7

Hardware 7
Processor 7
Disk Space 7
Virtual Memory 7
Software 7
Operating System 7
Virtualization 7

3 General Recommendations and Best Practices 9

HP Service Health Reporter Application 9

4 Sybase IQ Database 11

5 Operating System Bottlenecks 13

6 Benchmark 14

Benchmark Scenario 1 16

Test Methodology 16
Hardware Configuration 16
To Achieve the Results 16
Test Result Summary 17

Benchmark Scenario 2 18

Test Methodology 18
Hardware Configuration 18
To Achieve the Results 18
Test Result Summary 18

Benchmark Scenario 3 20

Test Methodology 20
Hardware Configuration 20
To Achieve the Results 20
Test Result Summary 21

Benchmark Scenario 4 22
Test Methodology 22
Hardware Configuration 22
To Achieve the Results 22
Test Result Summary 22

1 Introduction

HP Service Health Reporter (SHR) is a cross-domain performance reporting solution. SHR uses SAP BusinessObjects Enterprise for all its business intelligence and reporting needs. SHR uses the Sybase IQ database for storing performance metrics for long periods. In addition to SAP BusinessObjects and Sybase IQ, SHR consists of several collectors that gather performance metrics from various data sources. Therefore, hardware sizing for SHR proves to be challenging.

One solution is to conduct performance tests on the product. The benchmark test results are based on the hardware on which the tests are conducted.

The key objective of this guide is to provide steps to modify various applications, databases, and operating system parameters, and also to record the results of various performance tests carried out on the product.

Chapter 2 provides general guidelines and best practices to obtain optimal performance from the SHR application, the Sybase IQ database, and the operating system.

Chapter 3 provides details of various performance benchmark tests conducted on SHR. You can use the results of these tests to choose a system configuration for specific SHR loads. These tests were conducted in a controlled environment and should only be used as an indication of the capacity of the system. **Do not replicate the results directly in your environment**.

2 Hardware and Software Requirements

Hardware

Processor

- Intel 64-bit (x86-64) or AMD 64-bit (AMD64):
  - Supported for Windows only.
  - For Intel 64-bit (x86-64), the following Xeon processor families are recommended: Penryn, Nehalem, Westmere, and Sandy Bridge.
  - For AMD 64-bit (AMD64), the following Opteron processor families are recommended: Istanbul, Lisbon, and Valencia.

Disk Space

Before allocating disk space for SHR 9.20, consider the following:

- The recommendations in the table are minimum disk space required based on average test environment at HP (considering the default retention time). More complex environments may require more disk space.
- Disk performance is important for high scale environments that are Medium tier or higher. HP recommends RAID 1+0 (10) with battery-backed write cache on disks of 15,000 rpm or better. Disk configurations that do not meet this level of performance are not adequate.
- Increasing log file size from the default settings uses more disk space. Before increasing log file size, make sure you have adequate disk space.

Virtual Memory

Recommended size is at least twice physical memory (i.e., 2 multiplied by RAM).

Software

Operating System

Windows Server 2008 R2 x 64 Enterprise Editions with Service Pack 1.

Virtualization

- VMware ESXi 5.0 or later minor version:
  - Supported for Windows only.
  - Virtual environment must meet the x86-64 or AMD64 hardware requirements listed above.

The recommendations listed in this support matrix apply to SHR 9.20 running under the default settings.

The distribution of CIs is categorized as follows for various deployment sizes:

	Performance Target for SHR (Distribution of CIs)			
Agent		**Small**	**Medium**	**Large**
	Host	500	5000	20000
BPM	Applications	20	50	1000
	Transactions	100	500	5000
RUM	Applications	5	20	100
	Transactions	150	500	5000
NNM iSPI Performance for Metrics	Nodes	250-3000	3000-8000	8000-18000
	Interfaces	120000	400000	900000

The following table describes different tiers of managed node environments and the hardware requirements. The values stated here are approximate and reflect levels tested by HP. For a complex environment, you may need to increase the Java heap size and upgrade hardware as indicated. Managed environments larger than these tiers are not supported without additional HP approval.

- The requirements mentioned are out of the box SHR 9.20 Default Retention Periods

Table Type	**Default Retention(Days)**
Raw	90
Hourly	365
Daily	1825

- Default initial history collection for different sources SHR 9.20 Collects

	Init History Collection
Agents	15 Days
BSM Profile DB	15 Days
HP OM (Events)	15 Days
OMi (HIs and KPIs)	7 Days

Single System Managed Environment Size		**Recommended Hardware System Requirements**				**Collection Thread/Heap**		**Sybase IQ**					
Deployment Type	**No. of CP's**	**CPU (64-bit) x86-64 AMD64**	**RAM**	**Disk space for Product installation**	**Disk Space for Database**	**No .of Collection Threads**	**Recommended Collection Java heap size**	**gm**	**iqmc**	**iqtc**	**iqgovern**	**Main dbspace**	**Temp dbspace**
Small* **	3	4 CPU cores	8 GB	3 GB	500 GB	50	4 GB(-Xmx	150	1.7	1.7	50	49GB	49GB
Medium	6	8 CPU cores	16 GB	3 GB	1 TB	50	4 GB(-Xmx)	150	3.5	3.5	26	98GB	98GB
Medium	All	8 CPU cores	24 GB	3 GB	1 TB	50	4 GB(-Xmx)	150	5.5	5.5	26	98GB	98GB
Large	All	16 CPU cores	32 GB	3 GB	2 TB	200	8 GB(-Xmx)	150	8.5	8.5	32	192GB	192GB

*** For Small Deployment (4 CPU x 8 GB RAM), add the following entry -iqgovern 50 in `%SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg`.

3 General Recommendations and Best Practices

HP Service Health Reporter Application

SHR implements an Extract, Transform, and Load (ETL) layer to collect, transform, and load data into its data warehouse. The data warehouse is implemented in a Sybase IQ column store database. Several parameters affect the performance of this ETL layer. You can apply the following recommendations to the ETL layer for optimal performance:

- Sybase IQ is CPU-intensive. Higher the number of CPU allocated better the performance.
- Content Packs installed in SHR deploy data processing streams to audit and control the data flow. These streams consist of steps that implement various ETL tasks and also control the sequence of execution of these tasks. Each Content Pack deploys one or more streams in SHR. These streams are launched periodically and each step launches a process that runs the specified task. The SHR 9.20 Content Packs deploy more than 140 steps. To keep the performance overhead of idle Content Packs low, it is recommended to install only those Content Packs that have data sources configured.
- The SHR collectors provide the ability to collect historical data. The default settings for various data sources are:
    - HP Performance Agent: 15 days.
    - Business Service Management (BSM) Profile database: 15 Days.
    - BSM Service and Operations Bridge (OMi) database (Health Indicators (HIs) and Key Performance Indicators (KPIs)): 7 days.
    - HP Operations Manager (HPOM) events: 15 Days.

These default settings may be changed to get additional historical data. However, performance can decline with increase in duration and number of agents. It primarily affects the RAM usage and increases the corresponding time taken for completing this operation.

To collect additional historical data from the HP Performance Agents, increase the **collector.initHistory** parameter in the `config.prp` file that is present in the `%PMDB_HOME%/data` folder. The number of HP Performance Agents polled for data concurrently is governed by the number of threads configured in the SHR collection. The **org.quartz.threadPool.threadCount** parameter in the `%PMDB_HOME%/config/ramscheduler.properties` file identifies the maximum number of threads that may be generated and therefore, the maximum number of HP Performance Agents that may be polled simultaneously. If the requested historical data is huge, decrease the number of threads. This ensures the memory requirement of SHR does not exceed and result in an `OutOfMemory` error. With 5000 hosts and 15 days of initial history collection, the recommended thread count is 50.

The Profile database is another source for large volumes of data. The default value is 15 days. If more than 15 days of data is required, modify the **dbcollector.initHistory** parameter in the `%PMDB_HOME%/data/config.prp` file. If more historical data is required, set the thread count in the `%PMDB_HOME%/config/jdbcscheduler.properties` file to a very low value. This slows down the HP Performance Agent collection but allows the Profile

database data to be collected, which can increase the heap memory consumption of SHR. After the collection is complete, the thread count may be reset to the default value.

A **MaxHistoryTime** parameter is defined in the `%PMDB_HOME%/data/config.prp` file for HP Performance Agent collection, in conjunction with **InitHistoryTime**. This parameter determines the maximum amount of historical data that may be collected by SHR from the HP Performance Agents. **collector.maxHistory** is an important parameter if the SHR collection restarts after going down for several days. By default the value is set to 2 days (**48hrs**)

A **dbcollector.maxHistory** parameter is defined in the `%PMDB_HOME%/data/config.prp` file for Profile database data collection, in conjunction with **dbcollector.initHistory**. This parameter determines the maximum amount of historical data that may be collected by SHR from the Profile database. **dbcollector.maxHistory** is important if the SHR collection restarts after going down for several days. By default the value is set to 6 hrs.

After the data is loaded into the stage tables, the collected data is archived as CSV files in the `%PMDB_HOME%/stage/archive` folder. These files are deleted periodically. An increase in the number of files will affect the disk I/O and the overall performance of the SHR system.

You are required to manually process the data in the `%PMDB_HOME%/stage/failed_to_transform`, `%PMDB_HOME%/stage/failed_to_stage` and `%PMDB_HOMEstage/failed_to_load` folders. Accumulation of files in these folders can adversely affect the disk I/O. These files contain data rejected by SHR's ETL layer and may need manual processing before reprocessing is attempted.

If there are too many files accumulated in the `%PMDB_HOME%/stage/` (Agent collection files) or `%PMDB_HOME%/collect` (DB Collection files) folder, decrease the thread count to reduce data inflow into SHR till the backlog is cleared. This situation can occur if Sybase IQ is down for a period of time while data collection was running.

SHR uses BSM Run-time Service Model (RTSM), HP Operations Management (HPOM), or VMware vCenter as its topology source. However, performance tests for SHR are done considering only the RTSM and HPOM topology sources. The default synchronization frequency for topology sources is set to 24 hours. This is the recommended minimum period. However, this value can be changed through the SHR Administration Console. The ideal value for this frequency depends on the frequency with which the topology sources are updated. If RTSM or HPOM is updated at a much lower frequency, you should reduce the SHR synchronization frequency accordingly. This helps avoid expensive dimension updates of all Content Packs. Performance of SHR decreases with increasing frequency of synchronization.

The frequency of data collection from HP Performance Agent is set to one hour by default. You can change this parameter for each host in the Administration Console.

All data movement within SHR is controlled through a data processing framework. This framework allows the administrator to control the number of SHR processes generated at any given time. If the SHR system has limited resources or is consuming very high CPU resources, providing a limit for the total number of SHR data processes can help reduce resource utilization. This, however, can slow down the movement of data into SHR. For more information about configuring the resource usage of the data processes, see the *Online Help for Administrators* section "Managing data processes."

Each data movement step that is processed in SHR has a maximum time limit. This limit is set to 60 minutes by default. In certain cases where a large amount of data is being processed, steps like pre-aggregation and forecasting might exceed this limit. This causes the data processing stream to display an error state. In such cases, you must wait till the stream completes.

4 Sybase IQ Database

You can apply the following recommendations to tune the Sybase IQ database for optimal performance:

- Retention period for SHR are as follows:
    - Raw (as-collected) data – 90 days
    - Hourly aggregated data – 365 days
    - Daily aggregated data – 1825 days

You can change the retention period. However, increasing the retention for raw data increases the size of the raw data table in Sybase IQ.

The following Sybase IQ startup parameters, located in the `%SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg` file, can be tuned for better SHR performance. If you make any changes to the following parameters, you must restart the Sybase IQ database:

- **gm :** This parameter limits the total number of concurrent user connections to the Sybase IQ server; by default, SHR sets this parameter to 150. If you have installed only one or two SHR Content Packs, you can set this parameter to a lower value for better performance. Note that Sybase IQ allocates memory for both *active* and *idle* connections and a lower value of gm helps avoid overheads.
- **iqmc** and **iqtc:** Sybase IQ uses main and temporary buffer caches for database operations. The data is stored in one of the two caches whenever it is in memory. For better Sybase IQ performance, consider adding more data files to the dbspace manually, preferably from a different disk, after the post-install configuration phase. This will increase the I/O rate and evenly distribute data in the database files, improving the overall database performance. You can add additional default, SHR sets iqmc=3.5 GB and iqtc=3.5 GB for small deployment, iqmc=5.5 GB and iqtc=5.5 GB for medium deployment and iqmc=8.5GB and iqtc=8.5 GB for enterprise deployment. You can increase the value of the buffer cache for better database performance depending on the total amount of physical memory available on the system

SHR creates database files and temporary dbspace files in the same directory (disk). The SHR Internal Monitoring (IM) Service extends the **pmdb_user_main** database size automatically by adding new files when the database space usage crosses 85 percent (The threshold value is controlled through config.prp - dbspace.max.percentage). It is recommended that an initial file size of higher volume be set instead of allowing the SHR IM Service to add the file.

SHR IM Service adds new files to the database in the same directory (same disk) as **pmdb.db** and will not extend the temporary **dbspace**. For better Sybase IQ performance, consider adding more data files to the **dbspace** manually, preferably from a different disk, after the post-install configuration phase. This will increase the I/O rate

and evenly distribute data in the database files, which will improve the overall database performance. You can add additional files to a **dbspace** either by using Sybase Central or from Interactive SQL Java (dbisql).

SHR IM Service automatically adds the data file to the existing **dbspace** based on the deployment selection during the post-install configuration phase. Multiple smaller sets of data files will degrade the performance. Sybase IQ performs best if one big data file is added to the **dbspace** prior to the SHR IM Service creation.

To add database files by using Sybase Central:

1. Click **Start** -> **Programs** -> **Sybase** -> **Sybase IQ 15.4** -> **Sybase Central v6.1 Edition**. Sybase Central opens.
2. On the right pane, double-click **Sybase IQ 15**.
3. On the Connections -> Connect with Sybase IQ 15...
4. In the **Connect** dialog box, on the **Identification** tab, type the user credentials.
5. On the **Database** tab, select the database you want to connect to and then click **OK**.
6. On the **Contents** tab, double-click **Dbspaces**. You can create a new dbspace file by clicking the **Create a dbspace** option on the left pane.

To add database files by using dbisql:

1. Click **Start** -> **Programs** -> **Sybase** -> **Sybase IQ 15.4** -> **Interactive SQL**. Interactive SQL opens.
2. In the **Connect** dialog box, on the **Identification** tab, type the user credentials.
3. On the **Database** tab, select the database you want to connect to and then click **OK**.
4. Use the ALTER DBSPACE command to add a file:

   ALTER DBSPACE <*dbspace name***> ADD FILE <***logical name***> '<***complete file path***>' SIZE <***size***>**
   Example: **ALTER DBSPACE pmdb_user_main ADD FILE pmdb_user_main02**

   'C:\dbfile\pmdb_user_main02.iq' SIZE 20GB

If the following Sybase IQ database files are relocated to different drives prior to the start of data collection, it can enhance the performance:

- **Catalog Store** (for example, pmdb.db) - After the database is created, this file cannot be moved.
- **IQ Store or IQ_SYSTEM_MAIN** (for example, pmdb.iq) - After the database is created, this file cannot be moved.
- **IQ Temporary store or IQ_SYSTEM_TEMP** (for example, pmdb.iqtmp) -This file can be relocated post database creation.
- **IQ message log or IQ_SYSTEM_MSG** (for example, pmdb.iqmsg) - This file can be relocated post database creation.
- **Catalog Store transaction log** (for example, pmdb.log) - After the database is created, this file cannot be moved.

User main or PMDB_USER_MAIN (for example, pmdb_user_main(x).iq) - During the creation of the database, you can specify a different location.

5 Operating System Bottlenecks

To identify system-level resource bottlenecks, it is recommended that you monitor the performance of the operating system. Install HP Performance Agent on the SHR server for monitoring system performance. The following are a few rules to identify the operating system-level resource bottlenecks, using the metrics extracted from HP Performance Agent:

- Consistent high global CPU utilization (GBL_CPU_TOTAL_UTIL > 90%), and significant Run Queue or Load Average (GBL_PRI_QUEUE or GBL_RUN_QUEUE > 3)
- High physical memory utilization (GBL_MEM_UTIL > 95%), and significant pageout rate (GBL_MEM_PAGEOUT_RATE > 1), or any deactivations (GBL_MEM_SWAPOUT_RATE > 0), or Vhand process consistently active (vhand's ROC_CPU_TOTAL_UTIL > 5%).
- Consistent high utilization on at least one disk device (GBL_DISK_UTIL_PEAK or highest BYDSK_UTIL > 50%).
- Significant queuing lengths (GBL_DISK_SUBSYSTEM_QUEUE > 3 or any BYDSK_REQUEST_QUEUE > 1).

Processes or threads blocked on I/O wait reasons (PROC_STOP_REASON = CACHE, DISK, IO).

6 Benchmark

This chapter describes the benchmark test scenarios, the test methodology that is used, and the results of the tests. Since it is difficult to perform tests on different platforms and SHR configurations, a representative set of SHR configurations were used to perform these tests.

Deployment Size categorized based on the CIs distributed as below:

	Performance Target for SHR (Distribution of CIs)			
		Small	**Medium**	**Large**
Agent	Host	500	5000	20000
	File System	1500	15000	60000
	Disk	1500	15000	60000
	Network	1000	10000	40000
	CPU	1000	10000	40000
BPM	Applications	20	50	1000
	Transactions	100	500	5000
	Locations	10	50	1000
	Trx-Loc Combinations	500	5000	200000
	Max EPS	1	10	220
RUM	Applications	5	20	100
	Transactions	150	500	5000
	End User groups	100	500	10000
	Locations	50	500	10000
	Servers	5	15	100
	Events	10	50	100
	Trx-Loc Combinations	2000	25000	200000
	Max EPS	100	300	900
NNM iSPI Performance for Metrics	Nodes	250-3000	3000-8000	8000-18000
	Interfaces	120000	400000	900000
	Polled addresses	5000	10000	20000
	Polled interfaces	10000	50000	70000
	Custom polled objects	30000	50000	75000
	Polled node components	40000	60000	80000

The following table lists the performance benchmark scenarios:

Test Scenario	SHR and SybaseIQ System	Topology Source	Deployment Size	Content Packs (Out of the box)
1	Single System	HPOM	Enterprise	All
2	Single System	HPOM	Medium	All
3	Two System	RTSM	Medium	All
4	Single System	RTSM	Medium	All

Benchmark Scenario 1

SHR and Sybase IQ are installed with all Content Packs on the same system. It is then deployed in an HPOM environment. This test was carried out in an enterprise (supports 10000 hosts) deployment.

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out on a live environment with 10000 live UNIX and Microsoft Windows hosts running HP Operations agent or HP Performance Agent.
- The average time taken to collect was measured.
- The average time taken by various steps of the data processes was measured.
- The CPU, memory, and disk I/O utilization of the SHR system was collected at various periods during the test.

Hardware Configuration

Deployment Name	HPOM
SHR (Enterprise Deployment - Standalone)	SHR and Sybase IQ installed on the same system
	Model: HP ProLiant DL580 G5
	CPU: 16 (Intel Xeon CPU X7350 @2.93 GHz)
	RAM: 32 GB & Virtual Memory: 64 GB
	HDD size (preferably disk array with RAID5): 2 TB Storage Type: P6000 EVA Storage Systems Drive Type: SAS Rotational Speed: 10K RPM Transfer Speed PHY 1:3 GBPS Disk Cache Battery: 1 GB
	OS: Windows 2008 R2 SP1

To Achieve the Results

1. Increase the Aggregate JVM memory (Xmx) to 128M (Default value 32M) location of file `%pmdb_home%\config\startup\aggregate.ini`.
2. Increase the Sybase IQ Main/Temp Cache to 8.5 GB location of file `%SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg`.
3. Increase the Collection JVM Memory (Xmx) to 8 GB (Default value 4GB) location of file `%pmdb_home%\bin\CollectionServiceCreation.bat`.

   Recreation of the **CollectionService** is required. In the command prompt, type:

   a. `CollectionServiceCreation.bat -remove "C:\HP-SHR\" "C:\HP-SHR\"`.

   b. Set JVM_ARGS=`-Xmx8192m` in **CollectionServiceCreation.bat**.

   c. `CollectionServiceCreation.bat -install "C:\HP-SHR\" "C:\HP-SHR\"`.

d. Dependent service creation command: `C:\>sc config HP_PMDB_Platform_Collection depend= HP_PMDB_Platform_IM/HP_PMDB_Platform_Message_Broker/HP_PMDB_Platform_Sybase`.

4. Increase the number of collection threads to 200 (Default value 50) location of file `%pmdb_home%\config\jdbcscheduler.properties)`.

Test Result Summary

The following table lists the test results for SHR in the HPOM deployment scenario with all Content Packs installed.

Latency:

	Content Packs		
	Network	**System**	**Application**
Daily Table	6Hrs 30Min		
Hourly Table	3Hrs 30Min	3Hrs	4Hrs 30Min
Rate Table	2Hrs 40Min	2Hrs 30Min	3Hrs 10Min

HPOM Enterprise Deployment Scenario with All Content Packs (Live 10000 Nodes)

Metric Name	**Performance Metrics**	
HPOM Topology Collection	Average 95th Percentile	6 Hours 00 Minutes 6 Hours 02 Minutes
Data Collection (Performance Agent/Operations agent)	Average 95th Percentile	26 Minute 20 Second 30 Minute 33 Second
Data Collection (DB)	Average 95th Percentile	0 Hours 55 Minutes 3 Hours 41 Minutes
Reconciliation	Average 95th Percentile	31 Seconds 1 Minute 46 Seconds
Mapper	Average 95th Percentile	10 Seconds 35 Seconds
Stage	Average 95th Percentile	24 Seconds 56 Seconds
Data Load	Average 95th Percentile	50 Seconds 2 Minutes 48 Seconds
Data Aggregation (Aggregate Data Store)	Average 95th Percentile	39 Seconds 2 Minutes 47 Seconds
SQLExecutor (Custom Procedure for Data movement)	Average 95th Percentile	07 Seconds 24 Seconds
Database Maintenance	Average 95th Percentile	4 Minutes 07 Seconds 8 Minutes 19 Seconds

Benchmark Scenario 2

SHR and Sybase IQ are installed with all Content Packs on the same system. It is then deployed in an HPOM environment. This testing was carried out on a medium (supports 5000 hosts) deployment.

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out on a live environment with 5000 live UNIX and Microsoft Windows hosts running HP Operations agent or HP Performance Agent.
- The average time taken to collect was measured.
- The average time taken by various steps of the data processes was measured.
- The CPU, memory, and disk I/O utilization of the SHR system was collected at various periods during the test.

Hardware Configuration

Deployment Name	HPOM
SHR (Medium Deployment - Standalone)	SHR and Sybase IQ on the same system
	Model: HP ProLiant DL380p Gen8
	CPU: 8 (Intel Xeon CPU E5-26900 @2.9 GHz)
	RAM: 24 GB
	Virtual Memory: 48 GB
	HDD size (preferably disk array with RAID5): 1 TB Storage Type: P6000 EVA Storage Systems Drive Type: SAS Rotational Speed: 10K RPM Transfer Speed PHY 1:3 GBPS Disk Cache Battery: 1 GB

To Achieve the Results

Increase the Sybase IQ Main/Temp Cache to 5.5GB location of file %SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg

Test Result Summary

The following table lists the test results for SHR in the HPOM deployment scenario with all Content Packs installed.

Latency:

	Content Packs		
	Network	System	Application
Daily Table	6Hrs 30Min		
Hourly Table	4Hrs 30Min	3Hrs	3Hrs 30Min
Rate Table	4Hrs	2Hrs 30Min	3Hrs

HPOM Medium Deployment Scenario with All Content Packs (Live 5000 Nodes)

Metric Name	Performance Metrics	
HPOM Topology Collection	Average 95th Percentile	2 Hours 47 Minutes 2 Hours 49 Minutes
Data Collection (Performance Agent/Operations agent)	Average 95th Percentile	26 Minutes 20 Seconds 30 Minutes 33 Seconds
Data Collection (DB)	Average 95th Percentile	3 Seconds 13 Seconds
Reconciliation	Average 95th Percentile	12 Seconds 24 Seconds
Mapper	Average 95th Percentile	4 Seconds 10 Seconds
Stage	Average 95th Percentile	17 Seconds 36 Seconds
Data Load	Average 95th Percentile	39 Seconds 57 Seconds
Data Aggregation (Aggregate Data Store)	Average 95th Percentile	15 Seconds 54 Seconds
SQLExecutor (Custom Procedure for Data movement)	Average 95th Percentile	7 Seconds 23 Seconds
Database Maintenance	Average 95th Percentile	52 Seconds 1 Minute 42 Seconds

Benchmark Scenario 3

SHR and Sybase IQ are installed on different systems and all Content Packs are installed. It is then deployed in an RTSM environment. This testing was carried out on a medium (supports 5000 hosts) deployment.

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out on a live environment with 5000 live UNIX and Microsoft Windows hosts running HP Operations agent or HP Performance Agent.
- The average time taken to collect was measured.
- The average time taken by various steps of the data processes was measured.
- The CPU, memory, and disk I/O utilization of the SHR system was collected at various periods during the test

Hardware Configuration

Deployment Name	RTSM
SHR (Medium Deployment – Remote DB Box)	SHR and Sybase IQ on different systems
	Model: ProLiant DL385 G7
	CPU: 8 (AMD Opteron 6174 @2.2 GHz)
	RAM: 16 GB & Virtual Memory: 32 GB
	HDD size (preferably disk array with RAID5): 750 GB Storage Type: P6000 EVA Storage Systems Drive Type: SAS Rotational Speed: 10K RPM Transfer Speed PHY 1: 3 GBPS Disk Cache Battery: 1 GB
	OS: Windows 2008 R2 SP1
SHR (Medium Deployment – SHR Box)	SHR and Sybase IQ on different systems
	Model: ProLiant DL385 G7
	CPU: 8 (AMD Opteron 6174 @2.2 GHz)
	RAM: 16 GB & Virtual Memory: 32 GB
	HDD size (preferably disk array with RAID5): 250 GB Storage Type: P6000 EVA Storage Systems Drive Type: SAS Rotational Speed: 10K RPM Transfer Speed PHY 1: 3 GBPS Disk Cache Battery: 1 GB
	OS: Windows 2008 R2 SP1

To Achieve the Results

Increase the Sybase IQ Main/Temp Cache on remote database box to 12.28GB location of file `%SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg`

Test Result Summary

The following table lists the test results for SHR in the RTSM deployment scenario with all Content Packs installed.

Latency:

	Content Packs		
	Network	System	Application
Daily Table	22Hrs 30Min		
Hourly Table	4Hrs 30Min	3Hrs	4Hrs 30Min
Rate Table	4Hrs	2Hrs 30Min	4Hrs

RTSM Medium Deployment Scenario with All Content Packs (Live 5000 Nodes)

Metric Name	Performance Metrics	
RTSM Topology Collection	Average 95th Percentile	2 Minutes 49 Seconds 3 Minutes 10 Seconds
Data Collection (Performance Agent/Operations agent)	Average 95th Percentile	23 Minutes 26 Seconds 49 Minutes 32 Seconds
Data Collection (DB)	Average 95th Percentile	18 Seconds 11 Seconds
Reconciliation	Average 95th Percentile	1 Minute 39 Seconds 9 Minutes 25 Seconds
Mapper	Average 95th Percentile	1 Minute 39 Seconds 9 Minutes 26 Seconds
Stage	Average 95th Percentile	3 Minutes 27 Seconds 14 Minutes 56 Seconds
Data Load	Average 95th Percentile	1 Minute 43 Seconds 3 Minutes 13 Seconds
Data Aggregation (Aggregate Data Store)	Average 95th Percentile	44 Seconds 3 Minutes 1 Second
SQLExecutor (Custom Procedure for Data movement)	Average 95th Percentile	31 Seconds 3 Minutes 5 Seconds
Database Maintenance	Average 95th Percentile	1 Minute 0 Second 1 Minute 1 Second

Benchmark Scenario 4

SHR and Sybase IQ are installed on the same system with all Content Packs installed. It is then deployed in an RTSM environment. This testing was carried out on a Medium (supports 5000 hosts) deployment.

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out on a live environment with 5000 live UNIX and Microsoft Windows hosts running HP Operations agent or HP Performance Agent.
- The average time taken to collect was measured.
- The average time taken by various steps of the data processes was measured.
- The CPU, memory, and disk I/O utilization of the SHR system was collected at various periods during the test

Hardware Configuration

Deployment Name	**RTSM**
SHR (Medium Deployment - Standalone)	SHR and Sybase IQ on the same system
	Model: ProLiant DL380 G7
	CPU: 16 (Intel Xeon X5650 @2.67GHz)
	RAM: 24 GB & Virtual Memory: 48 GB
	HDD size (preferably disk array with RAID5): 1 TB Storage Type: P6000 EVA Storage Systems Drive Type: SAS Rotational Speed: 10K RPM Transfer Speed PHY 1: 3 GBPS Disk Cache Battery: 1 GB
	OS: Windows 2008 R2 SP1

To Achieve the Results

Increase the Sybase IQ Main/Temp Cache to 5.5 GB location of file `%SYBASE%\IQ-15_4\scripts\pmdbconfig.cfg`

Test Result Summary

The following table lists the test results for SHR in the RTSM deployment scenario with all Content Packs installed.

Latency:	Content Packs		
	Network	System	Application
Daily Table	24Hrs 30Min		
Hourly Table	4Hrs 30Min	2Hrs 30Min	3Hrs 20Min
Rate Table	4Hrs	2Hrs	3Hrs

RTSM Medium Deployment Scenario with All Content Packs (Live 5000 Nodes)

Metric Name	Performance Metrics	
RTSM Topology Collection	Average 95th Percentile	1 Minute 55 Seconds 2 Minutes 24 Seconds
Data Collection (Performance Agent/Operations agent)	Average 95th Percentile	20 Minutes 4 Seconds 28 Minutes 17 Seconds
Data Collection (DB)	Average 95th Percentile	12 Seconds 27 Seconds
Reconciliation	Average 95th Percentile	37 Seconds 41 Seconds
Mapper	Average 95th Percentile	27 Seconds 19 Seconds
Stage	Average 95th Percentile	1 Minute 50 Seconds 2 Minutes 2 Seconds
Data Load	Average 95th Percentile	3 Minutes 30 Seconds 1 Minute 53 Seconds
Data Aggregation (Aggregate Data Store)	Average 95th Percentile	2 Minutes 29 Seconds 4 Minutes 4 Seconds
SQLExecutor (Custom Procedure for Data movement)	Average 95th Percentile	2 Minutes 7 Seconds 8 Minutes 4 Seconds
Database Maintenance	Average 95th Percentile	1 Minute 16 Seconds 2 Minutes 23 Seconds