HP Service Health Reporter

for the Windows® and Linux operating systems

Software Version: 9.30

Performance, Sizing, and Tuning Guide

Document Release Date: March 2014

Software Release Date: July 2013

Legal Notices

Warranty

Restricted Rights Legend

Copyright Notices

Trademark Notices

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

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.

The key objective of this guide is to provide steps to arrive at sizing of the hardware needed to deploy the SHR in your environment and how to modify various applications, databases, and operating system parameters to achieve optimal performance, and also to record the results of various performance tests carried out on the product within HP lab.

Chapter 2 provides the guidelines to determine the size of the deployment and the hardware and software requirements for different deployments.

Chapter 3 provides general guidelines and best practices to obtain optimal performance from the SHR application, the Sybase IQ database, and the operating system.

Chapter 4 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 Sizing Approach

The objective of sizing is to estimate the system resource required to ensure the deployed system meets the performance objectives.

Sizing the Deployment

The factors that affect hardware sizing are,

- The content you will deploy and the load for each of the content pack.
- The retention period for each of the content

Calculating Content Load

This section provides the guidelines to calculate the load for some of the out of box SHR content. The load is computed based on certain assumptions and approximations. So, **while choosing hardware you should include enough headroom to handle the actual load**.

System content

The size of the environment for system content is determined by the total number of physical and virtual nodes (n), average number of file systems per node (fs), average number of disks per node (disk), average number of CPUs per node (cpu), and average number of network interfaces per node (n/w if). SHR extracts 5 minute summarized data for System content, so total number of records extracted per hour per CI is 60/5 = 12. Hence throughput requirement is computed as,

> Total number of CIs (t) = n + n * (fs + disk + cpu+ n/w if)
>
> Throughput requirement is ~ (t * 12) records per hour

Network Content

The size of the environment for network content is determined by the number of performance polled network nodes (n) and performance polled interfaces (n/w if) in your deployment by Network Node Manager iSPI Performance for Metrics. SHR extracts hourly summarized data from the network data source, so the throughput requirement is computed as,

> Total number of CIs = n + n/w if
>
> Throughput requirement is ~ (n + n/w if) * 1 if per hour

RUM/BPM Content

In the case of RUM/BPM content, the size of the environment is determined by the number of transactions (t), applications(a),locations(l) and MAX EPS . Refer to the BSM Administration guide for details on calculating MAX EPS for your environment.

Totals number of CIs ~= applications(a) + transactions(t) + locations(l) Throughput requirement is ~ (RUM MAX EPS + BPM MAX EPS) * 60 * 60

Retention Period

You should determine the retention period for each of the content. Out-of-the-box retention period for different summary tables is shown in Table 1. You should plan for more disk space, if you increase the retention period.

Table 1: Out-of-the-box Retention Period

Table Type	Default Retention (Days)
Raw	90
Hourly	365
Daily	1,825

Deployment Size

In SHR, the deployment size is categorized as small, medium and large based on the number of CIs collected from data sources. Small, medium and large deployment corresponds to 500, 5000 and 20000 nodes respectively. Total number of CIs and throughput requirement for these deployments is shown in Table 2.

Table 2 : Total CIs and Throughput requirement in SaOB (Service and Operations Bridge) Deployment

Deployment Size	System Nodes	Network Nodes	Network Interfaces	Application (RUM + BPM)	Event Rate	Total Number of CIs	Throughput Requirement (records/hour)
Small	500	5,000	10,000	~100/sec	10/sec	**~30K**	**~600K**
Medium	5,000	10,000	50,000	~300/sec	20/sec	**~220K**	**~3,200K**

Total CIs and Throughput requirement in HPOM Deployment

Deployment Size	System Nodes	Network Nodes	Network Interfaces	Total Number of CIs	Throughput Requirement (records/hour)
Small	500	5,000	10,000	**~30K**	**~200K**
Medium	5,000	10,000	50,000	**~220K**	**~2,000K**
Large	20,000	20,000	70,000	**~730K**	**~8,000K (8 million)**

Total number of CIs and records/hour in Table 3 are calculated based on entries in Table 2. Each of the system nodes is assumed to have **10 file systems, 10 disks, 5 network i/f, and 6 CPUs**.

Table 3: CI Distribution Details

Data Source/Content		Small	Medium	Large
Agent	System Node	500	5,000	20,000
	File System	5,000	50,000	200,000
	Disk	5,000	50,000	200,000
	Network	2,500	25,000	100,000
	CPU	3,000	30,000	120,000
BPM	Applications	20	50	1,000
	Transactions	100	500	5,000
	Locations	10	50	1,000
	Trx-Loc Combinations	500	5,000	200,000
	Max EPS	1	10	220
RUM	Applications	5	20	100
	Transactions	150	500	5,000
	End User groups	100	500	10,000
	Locations	50	500	10,000
	Servers	5	15	100
	Events	10	50	100
	Trx-Loc Combinations	2,000	25,000	200,000
	Max EPS	100	300	900
NNM iSPI Performance for Metrics	Polled addresses	5,000	10,000	20,000
	Polled interfaces	10,000	50,000	70,000

Above calculations include only the content that contributes the largest load to SHR. You should provide enough headroom for other content like KPI, HI, and so on.

NOTE: You should also do a similar exercise for the custom content you deploy in SHR.

Hardware and Software Configuration

Table 4, Table 5, Table 6, and Table 7 show the **minimum** configuration based on benchmark tests.

Note: You must provision these minimum hardware and software requirements for *HP Service Health Reporter*. Ensure that you supply more hardware resources (CPU, RAM, and Disk Space) than the minimum requirements for healthy performance of SHR.

Table 4: Hardware and Software Configuration for a Single-System Deployment

Managed Environment Size		**System Configuration**				**Sybase IQ Configuration**			
Deployment Type	**No. of Content Packs**	**CPU(64 bit) x-86-64**	**RAM (in GB)**	**Disk Space for DB**	**Disk space for s/w *** **	**iqmc (in GB)**	**iqtc (in GB)**	**main dbspace (in GB)**	**temp dbspace (in GB)**
Small*	3	4 CPU Cores	8	400 GB	100 GB	1.7	1.7	49	49
Medium	6	8 CPU Cores	16	800 GB	200 GB	3.5	3.5	98	98
Medium	All	8 CPU Cores	24	1.6 TB	400 GB	5.5	5.5	98	98
Large**	All	24 CPU Cores	64	4.5 TB	0.5 TB	24	24	192	192

Table 5: Hardware Configuration for a Dual-System Deployment

Managed Environment Size		**SHR System Configuration**			**Sybase System Configuration**		
Deployment Type	**No. of Content Packs**	**CPU(64 bit) x-86-64**	**RAM (in GB)**	**Disk space for s/w *** **	**CPU(64 bit) x-86-64**	**RAM (inGB)**	**Disk Space**
Medium	All	8 CPU Cores	16	400 GB	8 CPU Cores	16	1.6 TB
Large**	All	16 CPU Cores	32	0.5 TB	16	32	4.5 TB

Table 6: Sybase IQ Configuration for a Dual-System Deployment

Managed Environment Size		**Sybase IQ Configuration**			
Deployment Type	**No. of Content Packs**	**iqmc (in GB)**	**iqtc (in GB)**	**main DB space (in GB)**	**Temporary DB space(in GB)**
Medium	All	7.0	7.0	98	98
Large**	All	15.0	15.0	192	192

*For systems with 4 CPUs, add the following entry in `{SYBASE}/IQ-15_4/scripts/pmdbconfig.cfg`:
-iqgovern 50

** For large deployment you should deploy collectors on separate systems. In the benchmark tests, collectors were deployed on two separate systems each collecting data for 10000 nodes each.

*** This column captures disk space requirement for software and runtime data

The SHR collector component is tested for a maximum of 10,000 nodes (~320K CIs). Table 7 shows the **minimum** configuration of the collector.

Table 7: Collector Configuration

Deployment Size(Number of Nodes)	**System Configuration**			**Collector Configuration**	
	CPU(64 bit) x-86-64	**RAM (in GB)**	**Disk Space (in GB)**	**Threads**	**Max Heap Size(in GB)**
10,000	4 CPU Cores	8	300	2500	6

3 General Recommendations and Best Practices

This section provides the guidelines and best practices for better performance of SHR.

Hardware and Software

Processor

You can deploy SHR on systems with Intel 64-bit (x86-64) or AMD 64-bit (AMD64) processors. It is recommended to use Intel processors.

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

Disk

Disk performance is important for high scale environments that are medium tier or higher. It is recommended to use RAID 1+0 (10) with battery-backed write cache on disks of 15,000 rpm or high performance SAN storage. Disk configurations that do not meet this level of performance are not adequate.

Software

Refer to SHR support matrix to see supported Operating System.

If you are using virtual machines, it is recommended to use VMware ESXi 5.0 or later minor version. Virtual environment must meet the *x86-64 or AMD64 hardware requirements*.

Operating System

The Linux kernel provides a system to limit the number of file descriptors and other resources on a per-process basis. SHR uses sockets and file-system files extensively, so SHR service start scripts sets this limit to 65,536.

SHR establishes connection to various data sources to collect monitoring data. When a connection is established the client side of the connection uses a port number. The ephemeral port range configured on a Windows system limits the maximum number of connections from one system to another. You should increase this range to approximately 60,000 by executing the steps mentioned in http://support.microsoft.com/kb/319502

You should configure **virtual memory to at least twice the physical memory** (that is, twice the size of the RAM).

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 collector component in SHR communicates with data sources and extracts data. The data warehouse is implemented in a Sybase IQ column store database. SHR allows you to deploy the collector and Sybase IQ components on separate systems. Based on the size of your deployment, you can deploy the collector component on multiple systems. This deployment enables you to distribute the load of central server. You can also choose to deploy the collector close to the data sources to reduce network bandwidth usage.

Some of the best practices for tuning SHR application are:

Data Extraction

Initial Data Collection

SHR collectors provide the ability to collect historical data when it starts collecting from a particular data source. The default settings for various data sources are shown in Table 8.

Table 8: Initial history collection period

Table Type	Initial History Collection period
Agents	15 Days
BSM Profile DB & Network DB	15 Days
OMi (HIs and KPIs)	7 Days

These default settings may be changed to get additional historical data. However, increase in duration affects the RAM usage and increases the time taken for completing the 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 controlled by the number of threads configured in the SHR collection. The **org.quartz.threadPool.threadCount** parameter in the file `{PMDB_HOME}/config/ramscheduler.properties` 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 5,000 hosts and 15 days of initial history collection, the recommended thread count is 50 for initial history collection.

Large volume of data is extracted from the Profile database and Network database. If more than 15 days of data is required, modify the **dbcollector.initHistory** parameter in the file `{PMDB_HOME}/data/config.prp`. If more historical data is required, set the thread count in the file `{PMDB_HOME}/config/ramscheduler.properties` 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, set the thread count to the default value.

Missing Data Collection

If SHR is down for some period for maintenance or other reasons or if data source is unreachable for some period, SHR collects the missing data from the data sources. `collector.maxHistory` parameter defined in the file `{PMDB_HOME}/data/config` determines the maximum amount of historical data that may be collected by SHR from the HP Performance Agents if collection stops for an agent for some reason. The default value is set to 15 days (360hrs). For profile and network database collection, `dbcollector.maxHistory` parameter determines the maximum amount of historical data that may be collected by SHR from the BSM Profile database and network database. The default value is 15 days (360 hrs). If SHR collects missing data for a number of data sources, you may reduce the `org.quartz.threadPool.threadCount` value as in the case of initial history collection.

Agent Response Timeout

If the agents in your environment do not respond after connection is established, you will see socket read connection timeout errors in logs. This slows down the data collection from other data sources. To overcome this issue, you can set the socket read timeout to a lower value for agent communication by executing the commands,

```
ovconfchg -ns bbc.cb -set RESPONSE_TIMEOUT <timeout in secs>
ovc -restart
```

However, if you set this to very low value, then socket connection closes before agent responds and this result in loss of data.

Collection Interval

SHR uses BSM Run-time Service Model (RTSM), HP Operations Management (HPOM), or VMware vCenter as its topology source. The default collection interval 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. You should set this parameter value depending on the frequency with which the topology sources are updated. If RTSM or HPOM is updated at a lower frequency, you may increase the collection interval. This helps in avoiding expensive dimension updates of all Content Packs. Performance of SHR decreases if you reduce the collection interval.

The collection interval for data collection from HP Performance Agent, Profile Database, and Network database is set to one hour by default. You can change this parameter from SHR Administration console. Increasing the collection interval, results in latency increase.

Data Retention Period for the Collector

The SHR server pulls (copies if collector coexists with server) data from collector and archives it in the {PMDB_HOME}/extract/archive folder on the collector system. You can configure the retention period for the archive folder using the archivefilecleanup.job.freq and archive.retention.period parameters in the {PMDB_HOME}/config/collection.properties file.

The archivefilecleanup.job.freq parameter indicates the frequency of the cleanup job in minutes and archive.retention.period indicates the retention period in hours.

Data Processing

Number of SHR Processes

Content Packs installed in SHR deploy data processing streams as shown in Figure 1, 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. To keep the performance overhead of idle Content Packs low, it is recommended to install only those Content Packs that have data sources configured.

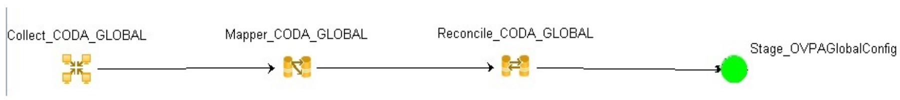

Figure 1: SHR Stream

All data movement within SHR is controlled through a data processing framework. This framework allows the administrator to control the total number of SHR processes that runs at any given time. It is also possible to control the number of processes per step type. 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 and limiting per step type process can help reduce resource utilization. This, however, can slow down the movement of data into SHR. Similarly if there is high latency with data movement then you can increase the limits of SHR processes depending on the HW resources available to SHR.

To limit the number of SHR data processes, see the *Online Help for Administrators* section "Managing data processes." To limit the number of process per step type, execute the command,

```
abcAdminUtil -setResourceCount -resourceType <type> -value <value>
where,
        <type> : Type of the step, e.g COLLECT_PROC, TRANSFORM_PROC, RECONCILE_PROC, STAGE_PROC, LOAD_PROC, AGGREGATE_PROC, EXEC_PROC_PROC.
         <value> : Limit on the number of process of <type>. e.g 40.
```

Default values set for each of these steps are listed in the table below:

Step Type	Default Process Limit
COLLECT_PROC	20
TRANSFORM_PROC	20
RECONCILE_PROC	20
STAGE_PROC	20
LOAD_PROC	30
AGGREGATE_PROC	20
EXEC_PROC_PROC	20

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 until the data processing is complete.

Disk space usage

Increase in number of files in SHR folders affects performance of disk operations. SHR components move the files to failed folder if it encounters errors while processing the data in the file. These files contain data rejected by SHR's ETL layer and may need to be corrected manually. Accumulation of files in these folders can increase disk space usage and may affect other disk operations. You are required to manually process the data in the `{PMDB_HOME}/stage/failed_to_transform`, `{PMDB_HOME}/stage/failed_to_stage` and `{PMDB_HOME}/stage/failed_to_load` folders as defined in SHR guides.

After the data is loaded into the stage tables, the collected data is archived as CSV files in the folder `{PMDB_HOME}/stage/archive`. These files are deleted periodically by SHR. An increase in the number of files increases disk space usage and may affect other disk operations.

Increasing log file size from the default settings uses more disk space. Before increasing log file size, make sure you have adequate disk space.

Control load on Sybase IQ

SHR pre-calculates the summary data to reduce the time taken to query large sets of data in reports. The summarization processes are modeled as steps in SHR stream and executes at the background. The aggregation functions used in the summarization process includes average, maximum, minimum, count, 90^{th} percentile, 95^{th} percentile, linear forecast etc. The summarizations calculated by out of the box content packs are defined in the file `{PMDB_HOME}/config/aggregate_config.xml.` The aggregations that are not used by out of box reports are disabled in this file. If you do not require some of the pre-summarization you can turn it off at per aggregation per metric level in this file to reduce the load on Sybase IQ. If you modify the out of box settings in this file you should execute following command to deploy the change,

```
aggrgen regenerateall=true
```

Sybase IQ is down for some period

If there are too many files accumulated in the `{PMDB_HOME}/stage` or the `{PMDB_HOME}/collect` folder, decrease the thread count of the collector to reduce data inflow into SHR until the backlog is cleared. This situation can occur if Sybase IQ was down or not accessible, or steps in streams failed to execute for a period of time while data collection was running.

Sybase IQ Database

Traditional OLTP databases store data row-wise fashion which is the preferred mechanism for transaction processing. Sybase IQ stores data by column which is suitable for queries that extract few fields from a table. Sybase IQ performance is generally limited by the CPU, Memory and storage available to Sybase IQ process. An increase in CPU drives more memory and disk usage, so all aspects of the system should be taken into account while scaling up for Sybase IQ.

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

- **iqgovern**: Sybase IQ calculates the value for this parameter based on the system configuration and there is no need to modify this parameter except when Sybase IQ is deployed on a lower configuration system. If Sybase IQ is deployed on 4 CPU, 8 GB RAM system, add the entry iqgovern=50 in the file {SYBASE}\IQ-15_4\scripts\pmdbconfig.cfg.
- **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 prevents 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. SHR sets iqmc=1.7 GB and iqtc=1.7 GB for small deployment, iqmc=3.5 GB and iqtc=3.5 GB for medium deployment and iqmc=7GB and iqtc=7GB for large deployment. You can increase the value of the buffer cache for better database performance depending on the physical memory available on the system.
- **main dbspace**: SHR creates main dbspace 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 certain threshold value. The threshold value is configured using dbspace.max.percentage parameter in config.prp file. The default value of this parameter is 85 percent. It is recommended that an initial file size of higher volume be set instead of depending on SHR IM Service to add the file. Multiple smaller sets of data files will degrade the performance. Sybase IQ performs best if one large file compared to multiple smaller files.
- **temp dbspace**: The SHR IM Service does 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).

To add database files by using Sybase Central:

1. Open Sybase Central:
    a. On Windows, click **Start** -> **Programs** -> **Sybase** -> **Sybase IQ 15.4** -> **Sybase Central v6.1 Edition**
    b. On Linux, run **/opt/HP/BSM/Sybase/shared/sybcentral610/scjview**.
2. In the right pane, double-click **Sybase IQ 15**.
3. On the Connections -> Connect with Sybase IQ 15...
4. In the **Connect** dialog box, in the **Identification** tab, type the user credentials.
5. In the **Database** tab, select the database you want to connect to, and then click **OK**.
6. In the **Contents** tab, double-click **Dbspaces**. You can create a new dbspace file by clicking the **Create a dbspace** option in the left pane.

To add database files by using dbisql:

1. Open interactive SQL:
    a. On Windows, click **Start** -> **Programs** -> **Sybase** -> **Sybase IQ 15.4** -> **Interactive SQL**.
    b. On Linux, run the following command:

        /opt/HP/BSM/Sybase/IQ-15_4/bin64/dbisql
2. In the **Connect** dialog box, in the **Identification** tab, type the user credentials.
3. In 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:
    On Windows:
    ALTER DBSPACE pmdb_user_main ADD FILE pmdb_user_main02

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

- To improve performance, it is recommended to relocate the following Sybase IQ database files to different physical drives prior to the start of data collection:
    - **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.

Sybase IQ performs better if you increase CPU and memory resource on the system.

SAP BusinessObjects

SHR reports are web intelligence documents. The Web Intelligence (WebI) Report Server in SAP Business Objects is responsible for generating Web Intelligence documents. The maximum memory available to the WebI server is only 2 GB as it is a 32-bit process. To overcome this limitation, you should estimate the load on the server and deploy the required number of WebI servers.

The load on the WebI Server depends on the number of simultaneous connections to the server and the complexity and size of the report documents accessed. If the server is not configured correctly, while accessing reports you may get errors like “Web Intelligence Server is busy” and “Server reached maximum number of simultaneous connections.”

Following are some of the steps you can do to avoid these errors:

- While accessing reports, default values for prompts can result in several thousands of records being loaded to the Web Intelligence server. You should specify appropriate values for prompts to avoid high load on the server. For example, when opening reports you should specify values for business service or node group prompts such that the number of nodes fetched from database is not more than 1000 to 2000.

- The default value for the maximum number of connections is 50. If the load on the server per request is high, then you may see the Server Busy error message. Consider reducing this parameter and instead add one more Web Intelligence Server to support additional connection requests. While adding additional servers, you should not break the golden rule "One Web Intelligence Processing Server per CPU core per machine".

In summary, the objective is to arrive at the number of WebI servers and maximum connections per server so that the server allows all users to connect and open report documents without reaching the 2 GB limit.

4 Benchmark

This chapter describes the benchmark test scenarios and the test methodology used for performance tests.

The following table lists the performance benchmark scenarios:

Benchmark Scenario	System	Topology Source	Deployment Size	Content Packs (Out of the box)
1	All components on single system	HPOM	Medium **(~2,000K records/hour)**	All
2	SHR and Sybase IQ on separate system	RTSM	Medium **(~3,300K records/hour)**	All
3	All components on single system	RTSM	Medium **(~3,300K records/hour)**	All
4	All components on single system	HPOM	10,000 Nodes for System content and Medium load for other content **(~4,000K records/hour)**	All
5	SHR server and Sybase IQ on same system with collectors on two separate systems	HPOM	Large for system and network, medium for other content **(~8,000K records/hour)**	All

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out using the configuration as documented in *Chapter 2 Deployment Size* section.
- The latency is measured as the time taken for the source system data to be available in various SHR tables.
- 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.

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 testing was carried out on a medium (5000 hosts) deployment on Linux and Windows.

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 in file {SYBASE}\IQ-15_4/scripts/pmdbconfig.cfg

Benchmark Scenario 2

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 (5000 hosts) deployment on the Windows operating system.

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

To Achieve the Results

Increase the Sybase IQ Main/Temp Cache on remote database system to 12.28GB in file `{SYBASE}/IQ-15_4/scripts/pmdbconfig.cfg`

Benchmark Scenario 3

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 (5000 hosts) deployment on both Windows and Linux.

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

To Achieve the Results

Increase the Sybase IQ Main/Temp Cache to 5.5 GB. The `pmdbconfig.cfg` file is located on Windows in `%SYBASE%\IQ-15_4\scripts` and on Linux in `$SYBASE/IQ-15_4/scripts`.

Benchmark Scenario 4

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 using System content with a load of 10,000 hosts and other content with medium load. Benchmark test was performed on the Windows operating system.

Hardware Configuration

Deployment Name	HPOM
SHR (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

To Achieve the Results

1 Increase the Sybase IQ Main and Temp Cache to 11 GB each in the file `{SYBASE}/IQ-15_4/scripts/pmdbconfig.`

2 Increase the Collection JVM Memory (Xmx) to 6 GB.

On Windows:

a Run the following command:
CollectionServiceCreation.bat -remove "C:\HP-SHR\" "C:\HP-SHR\"

b Modify -Xmx to `–Xmx6144m` set in JVM_ARGS in **CollectionServiceCreation.bat**.

c Run the following command:
CollectionServiceCreation.bat -install "C:\HP-SHR\" "C:\HP-SHR\"

d Create dependent services:
sc config HP_PMDB_Platform_Collection depend= HP_PMDB_Platform_IM/HP_PMDB_Platform_Message_Broker/HP_PMDB_Platform_Sybase

On Linux:

a Stop the collection service by running the following command:
service HP_PMDB_Platform_Collection stop

b Set –Xmx in JVM_ARGS to –Xmx6144m in hpbsm_pmdb_collector_start.sh.

c Start the collection service:
service HP_PMDB_Platform_Collection start

Benchmark Scenario 5

SHR and Sybase IQ are installed on the same system with all Content Packs supported in OM deployment scenario. The SHR collector component is installed on two separate systems. This testing was carried out for a large (20000 hosts) deployment on both Windows and Linux.

Test Methodology

The following test methodology was used to perform the tests:

- The test was carried out on a live environment with 20000 UNIX and Microsoft Windows hosts running the HP Operations agent or HP Performance Agent.
- The latency is measured as the time taken for the source system data to be available in various SHR tables.
- 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

SHR Components	HPOM
SHR Server	SHR and Sybase IQ installed on the same system
	Model: HP ProLiant DL580 G5
	CPU: 24 (Intel Xeon CPU X7350 @2.93 GHz)
	RAM: 64GB & Virtual Memory: 128 GB
	HDD size (preferably disk array with RAID5): 5 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
SHR collector	The SHR collector is installed on Linux and Windows with the following configuration:
	Model: HP ProLiant DL580 G5
	CPU: 4 (Intel Xeon CPU X7350 @2.93 GHz)
	RAM: 8GB
	HDD size: 300 GB

To Achieve the Results

1. Increase the Sybase IQ Main and Temp Cache to 24 GB each; location of file is `{SYBASE}\IQ-15_4\scripts\pmdbconfig.cfg`
2. Increase the Collection JVM Memory (Xmx) to 6 GB (Default value 4GB).
3. On Windows, follow these steps to increase the maximum memory of collection JVM:

   a Run the following command:
   CollectionServiceCreation.bat -remove "C:\HP-SHR\" "C:\HP-SHR\"

   b Set JVM_ARGS=`-Xmx6144m` in CollectionServiceCreation.bat.

   c Run the following command:
   CollectionServiceCreation.bat -install "C:\HP-SHR\" "C:\HP-SHR\".

   d Create dependent services:
   sc config HP_PMDB_Platform_Collection depend= HP_PMDB_Platform_IM/HP_PMDB_Platform_Message_Broker/HP_PMDB_Platform_Sybase

   On Linux, follow these steps:

   a Stop the collection service by running the following command:
   service HP_PMDB_Platform_Collection stop

   b Set –Xmx in JVM_ARGS to –Xmx6144m in hpbsm_pmdb_collector_start.sh.

   c Start collection service:
   service HP_PMDB_Platform_Collection start