



Hewlett Packard
Enterprise

HPE Operations Bridge Reporter

Software Version: 10.01

Windows® and Linux operating systems

Performance, Sizing, and Tuning Guide

Document Release Date: June 2016

Software Release Date: June 2016

Legal Notices

Warranty

The only warranties for Hewlett-Packard Development Company, L.P. 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. HPE 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 HPE 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 Notice

© Copyright 2016 Hewlett Packard Enterprise Development LP

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.

Documentation Updates

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

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

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

<https://softwaresupport.hp.com>

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

<https://hpp12.passport.hp.com/hppcf/createuser.do>

Or click the **the Register** link at the top of the HP Software Support page.

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

Support

Visit the HP Software Support Online web site at: **<https://softwaresupport.hp.com>**

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:

<https://hpp12.passport.hp.com/hppcf/createuser.do>

To find more information about access levels, go to:

<https://softwaresupport.hp.com/web/softwaresupport/access-levels>

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is

<http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Chapter 1: Introduction	5
Chapter 2: Sizing Approach	6
Sizing the Deployment	6
Calculating Content Load	6
System Content	6
Network Content (by NPS integration)	7
RUM/BPM Content	7
OMi Content	7
Service Health Content	7
Throughput Calculation and Custom Content in OBR	8
Retention Period	8
Deployment Size	9
Hardware and Software Configuration	12
Chapter 3: General Recommendations and Best Practices	15
Hardware and Software	15
Operating System	15
Software	16
HPE Operations Bridge Reporter Server	16
Data Extraction	16
Data Processing	18
Vertica Database	20
Configuring and Monitoring Virtual Machines	22
SAP BusinessObjects	23
Adding Additional Web Intelligence (WebI) Servers	24
Chapter 4: Benchmark	27
Benchmark Scenario 1	27
Benchmark Scenario 2	29
Benchmark Scenario 3	31
Send Documentation Feedback	33

Chapter 1: Introduction

HPE Operations Bridge Reporter (OBR) is a cross-domain performance reporting solution. HPE OBR uses SAP BusinessObjects Enterprise for all its business intelligence and reporting needs. HPE OBR uses the HPE Vertica database for storing performance metrics for long periods. In addition to SAP BusinessObjects and HPE Vertica, OBR 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 OBR in your environment and to modify various applications, databases, and operating system parameters to achieve optimal performance.

[Chapter - Sizing Approach](#) provides the guidelines to determine the size of the deployment and the hardware and software required for various deployments.

[Chapter - General Recommendations and Best Practices](#) provides general guidelines and best practices to obtain optimal performance from the HPE OBR application, HPE Vertica database, and the operating system.

[Chapter - Benchmark](#) provides details of various performance benchmark tests conducted on HPE OBR. You can use the results of these tests to choose a system configuration for specific OBR loads. Note that these tests were conducted in a controlled environment and should only be used as an indication of the capacity of the system.

Chapter 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 guidelines to calculate the load for some of the out-of-the-box HPE OBR content. The load is computed based on certain assumptions and approximations. So, while choosing hardware include enough headroom to handle the actual load.

Note: If you are using additional content packs that are not part of HPE OBR media and/or custom content, ensure that you factor the additional throughput incurred from those content. See corresponding content documentation for guidance.

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 ($n/w\ if$). OBR 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 $\sim (t * 12)$ records per hour

Network Content (by NPS integration)

The size of the environment for network content (by NPS integration) 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*. OBR extracts hourly summarized data from the network data source, so the throughput requirement is computed as,

Total number of CIs = $n + n/w \text{ if}$

Throughput requirement is $\sim (n + n/w \text{ if}) * 1 \text{ if per hour}$

Note: Network content calculations mentioned above does not factor NNMi direct load.

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 \sim applications(a) + transactions(t) + locations(l)

Throughput requirement is $\sim (\text{RUM MAX EPS} + \text{BPM MAX EPS}) * 60 * 60$

OMi Content

In the case of OMi content, the size of the environment is determined by the OMi Events Per Second (EPS). Refer to the *BSM Administration Guide* for details on calculating MAX EPS for your environment.

Throughput requirement is $\sim \text{EPS} * 60 * 60 * 138.889$

Service Health Content

In the case of Service Health content, the size of the environment is determined by the number of KPIs configured for logging (k) and HIs configured for logging (h). Refer to the *BSM Administration Guide* for details on calculating MAX EPS for your environment.

Totals number of CIs $\sim k + h$

Throughput requirement is $\sim (k + (h * 139)) * 12$

Throughput Calculation and Custom Content in OBR

Apart from the content packs mentioned above, you need to calculate the throughput for other out-of-the-box content packs and custom content packs. Throughput for each content pack is calculated based on the number of records added to database per hour, which is calculated from the CSV files collected per hour.

For example, a content pack in which there are 2 rate tables, which for a small deployment, collects CSV files with 1000 and 2000 records respectively for each table per hour. The throughput for the content pack would be 3000 summing up the number of records pumped in by each CSV, without considering any custom procedure or aggregation.

For OMi, we notice that even though the number of events generated per hour and the number of HI and KPI instances logged per hour is not very high. There are many other parameters to be considered, especially in the custom procedures section for this content pack, which is why the throughput seems to be much higher even though the number of lines in the CSV file is relatively less. This is because we have come up on a suitable factor based on record generation estimates that we finalized up on after analyzing the custom procedures and aggregates in the content pack.

For each content pack, there are aggregates and custom procedures which create more records and this also has to be factored in as part of the throughput (records per hour). While preparing for a HPE OBR setup with multiple content packs, that may or may not include other out-of-the-box content packs and custom content packs not mentioned above, calculate the throughput based on these parameters. Add the calculated throughput to the additional throughput section of the sizing calculator to get more accurate and realistic configuration requirements.

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 the following Table 1. You should plan for more disk space on the Vertica database server 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 OBR, the deployment size is categorized as small, medium, large, and enterprise based on the number of CIs collected from data sources. Small, medium, large and enterprise deployment corresponds to 500, 5000, 20000, 40000 nodes respectively. Total number of CIs and throughput requirement for these deployments is shown in the following tables:

Table 2: Total CIs and Throughput requirement in Operations Bridge Deployment (with APM and Network)

Deployment Size	System Nodes	Network Nodes	Network Interfaces	Application (RUM + BPM)	RUM 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

Table 3: Total CIs and Throughput requirement in Operations Bridge Deployment (Systems and Events)

Deployment Size	System Nodes	OMi Events (per sec)	Total Number of CIs	Throughput Requirement (records/hour)
Medium	5,000	3	~160K	~3,400K (3.4 million)
Large	20,000	6	~640K	~10,700K (10.7 million)

Table 4: 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

Deployment Size	System Nodes	Network Nodes	Network Interfaces	Total Number of CIs	Throughput Requirement (records/hour)
Large	20,000	20,000	70,000	~730K	~8,000K (8 million)

Table 5: Total CIs and Throughput requirement in Enterprise Deployment (System)

Deployment Size	System Nodes	Total Number of CIs	Throughput Requirement (records/hour)
Enterprise	40,000	~1280K	~16,000K (16 million)

For High Availability environment, the sizing for the shared and local storage can be calculated as follows:

- Shared storage for SAP BusinessObjects : *<space allocated for disk as mentioned in this guide>*
- Storage for HPE Vertica: *<dbpace allocated for a 3-node cluster as mentioned in this guide>*
- Local storage, CPU, and memory must be of the same specification on both the OBR servers in high availability environment. Similarly, both the BusinessObjects high availability servers (if OBR and BusinessObjects components are installed separately) must have identical configuration.
- In case of Vertica clustering, each of the nodes in the Vertica cluster must have identical configuration.

Table 6: CI Distribution Details

Data Source/Content		Small	Medium	Large	Enterprise
Agent	System Node	500	5,000	20,000	40,000
	File System	5,000	50,000	200,000	400,000
	Disk	5,000	50,000	200,000	400,000
	Network	2,500	25,000	100,000	200,000
	CPU	3,000	30,000	120,000	240,000

Data Source/Content		Small	Medium	Large	Enterprise
BPM	Applications	20	50	1,000	1,500
	Transactions	100	500	5,000	7,500
	Locations	10	50	1,000	1,500
	Trx-Loc Combinations	500	5,000	200,000	300,000
	Max EPS	1	10	220	300
RUM	Applications	5	20	100	150
	Transactions	150	500	5,000	7,500
	End User groups	100	500	10,000	15,000
	Locations	50	500	10,000	15,000
	Servers	5	15	100	150
	Events	10	50	100	150
	Trx-Loc Combinations	2,000	25,000	200,000	300,000
	Max EPS	100	300	900	1,500
NNM iSPI Performance for Metrics	Polled addresses	5,000	10,000	20,000	30,000
	Polled interfaces	10,000	50,000	70,000	100,000
OMi	OMi EPS	1	3	5	10

Above calculations include only the content that contributes the largest load to OBR. 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 OBR.

Hardware and Software Configuration

Table 7, Table 8, Table 9, and Table 10 show the minimum configuration based on benchmark tests.

Note: This section lists the minimum hardware and software you must provision for HPE Operations Bridge Reporter. Ensure that you supply more hardware resources (CPU, RAM, and Disk Space) than the minimum requirements for optimal performance of OBR.

Table 7: Hardware and Software Configuration for a Single-System Deployment

Managed Environment Size		System 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 **
Small	4	24 CPU Cores	96	700 GB	600 GB
Medium	3	28 CPU Cores	128	1.7 TB	700 GB

** This column captures disk space requirement for software and run-time data.

Table 8: Hardware Configuration for a Distributed Deployment

Managed Environment Size		OBR System Configuration			Vertica System Configuration 3-Node Cluster (per node)		
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 (in GB)	Disk Space
Medium	4	16 CPU Cores	64	700 GB	10 CPU Cores	120	1.6 TB
Large***	4	24 CPU	96	1.2	24	288	4.2 TB

Managed Environment Size		OBR System Configuration			Vertica System Configuration 3-Node Cluster (per node)		
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 (in GB)	Disk Space
		Cores		TB	CPU Cores		
Enterprise***	4	48 CPU Cores	128	1.9 TB	24 CPU Cores	288	7.8 TB

*** For large and enterprise deployments of OA/SiS collections, additional remote collectors should be deployed on separate systems.

Note: SAP BusinessObjects requires a minimum of 16 GB RAM and 16 GB disk space.

The OBR collector component is tested for a maximum of 10,000 nodes (~320K CIs).

Table 9: Collector Configuration

Multiple Remote Collectors are recommended in the cases of agent-based collections like Operations Agent (OA) collection. An additional Remote Collector of the same configuration as specified below should be added for every 10,000 nodes in case of OA collection.

In the case of SiteScope collection, factors such as the number of instances monitored and the polling frequency are crucial, and these have to be considered when deciding the number of additional Remote Collectors required.

Deployment Size(Number of Nodes)	System Configuration (per remote collector)			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	250	2500	6

Note:

While factoring in for sizing calculations, we have considered the HPE Vertica recommendations as suggested at the below given links:

- <https://my.vertica.com/documentation/vertica/configuring-hardware-and-virtual-machines/>
- http://my.vertica.com/docs/Hardware/HPE_Vertica_HW_planning.pdf

Chapter 3: General Recommendations and Best Practices

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

Hardware and Software

Processor

You can deploy OBR 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 latest Gen 9 processors are recommended.

Disk for OBR Server

Disk performance is important for high scale environments that are medium tier or higher. For OBR server, it is recommended to use RAID 5+0 (50) 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.

For information about disk space recommendations for HPE Vertica, see ["Vertica Database" on page 20](#).

Operating System

Linux

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

Windows

OBR establishes a 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).

Software

To see the list of supported operating systems, refer to the *HPE Operations Bridge Reporter Support Matrix*.

HPE Operations Bridge Reporter Server

Operations Bridge Reporter (OBR) implements an Extract, Transform, and Load (ETL) layer to collect, transform, and load data into its data warehouse. The collector component in OBR communicates with data sources and extracts data. The data warehouse is implemented in a Vertica column store database. OBR allows you to deploy the reporting server (SAP BusinessObjects), collector and Vertica components on separate systems. Based on the size of the 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 OBR application are:

Data Extraction

Initial Data Collection

OBR 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 11: Initial history collection period.

Table 11: Initial history collection period

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

Table Type	Initial History Collection period
OMi (Events and KPIs)	15 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 in your collector system. The number of HP Performance Agents polled for data concurrently is controlled by the number of threads configured in the OBR 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 OBR will not be exceeded 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. For deployments with OA/SiS collections having over 10,000 agents, it is recommended to run remote collectors only and disable the 'local' collector.

A large volume of data is extracted from the Profile and Network databases. 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 OBR.

After the collection is complete, set the thread count to the default value.

Missing Data Collection

If OBR is down for some period for maintenance or other reasons or if data source is unreachable for some period, OBR collects the missing data from the data sources. If the collection stops for some reason, `collector.maxHistory` parameter defined in the file `{PMDB_HOME}/data/config` determines the maximum amount of historical data that may be collected by OBR from the HP Performance Agents. 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 OBR from the BSM Profile and Network databases. The default value is 15 days (360 hrs). If OBR 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, set the socket read timeout to a lower value for agent communication by executing the following 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

OBR 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 OBR 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 to avoid expensive dimension updates of all Content Packs. Performance of OBR 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 OBR Administration console. Increasing the collection interval results in latency increase.

Data Retention Period for the Collector

The OBR 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 OBR Processes

Content Packs installed in OBR 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 OBR. 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: OBR Stream



All data movement within OBR is controlled through a data processing framework. This framework allows the administrator to control the total number of orchestration services that runs at any given time. It is also possible to control the number of processes per step type. If the OBR system has limited resources or is consuming very high CPU resources, providing a limit for the total number of OBR data processes and limiting per step type process can help reduce resource utilization. This, however, can slow down the movement of data into OBR. Similarly if there is high latency with data movement then you can increase the limits of OBR processes depending on the hardware resources available to OBR.

To limit the number of OBR data processes, see the *Managing data processes* section in the *HPE OBR Online Help for Administrators* section . To limit the number of process per step type, execute the following command:

```
abcAdminUtil -setResourceCount -resourceType <resource type> -value  
<pool count>
```

where,

<resource type> is type of the step. Example: COLLECT_PROC, TRANSFORM_PROC, RECONCILE_PROC, STAGE_PROC, LOAD_PROC, AGGREGATE_PROC, and EXEC_PROC_PROC

<pool count> is the limit on the number of processes of type <resource type>.

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

Step Type	Default Process Limit
EXEC_PROC_PROC	20

Each data movement step that is processed in OBR has a maximum time limit. By default, this limit is set to 60 minutes. 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 OBR folders affects performance of disk operations. OBR components move the files to failed folder if it encounters errors while processing the data in the file. These files contain data rejected by OBR'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, {PMDB_HOME}/stage/failed_to_load and {PMDB_HOME}/stage/failed_to_reconcile folders.

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 OBR. 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.


For more information, see *HPE Operations Bridge Reporter Troubleshooting Guide*.

Vertica Database

Traditional OLTP databases store data row-wise, which is the preferred mechanism for transaction processing. HPE Vertica stores data by column which is suitable for queries that extract few fields from a table. HPE Vertica performance is generally limited by the CPU and Memory.

HPE Vertica supports scale up and scale out as per your need. For scale up, more hardware resources such as CPU and RAM need to be added. For Scale out, HPE Vertica clustering requires minimum of 3 nodes, and all the nodes must have the same hardware configuration.

Clustering lets you scale out your database cluster easily by adding more hardware.



Columns are duplicated across cluster nodes. If one machine goes down, you still have a copy:

- Data warehouse log-based recovery is impractical. Instead, stores enough projections for K-safety

New cluster node queries existing nodes for the data it needs:

- Rebuilds missing objects from other nodes, another benefit of multiple sort orders

The following hardware configurations provide optimal performance for your HPE Vertica database.

Processor

See [HPE Vertica Hardware Planning Guide](#) for information about processor requirements for Vertica.

Memory

HPE Vertica requires a minimum of 8 GB of memory per physical CPU core in each server. However, in high-performance applications, you should run 12-16 GB of memory per physical core. The memory should be at least DDR3-1600 (preferably DDR4-2133), and should be appropriately distributed across all memory channels in the server.

Storage

HPE Vertica requires a minimum read/write speed of 40 MB/s per physical core of the CPU. However, for best performance, you should have 60–80 MB/s per physical core. Each node should have 1–9 TB of storage post RAID. In a production setting, HPE Vertica recommends RAID 10. RAID 50 can be a viable alternative.

Due to the heavy compression/encoding that HPE Vertica performs, you do not need to use solid-state drives (SSDs). To satisfy HPE Vertica requirements, a RAID array of more, less expensive hard disk drives (HDDs) works just as well as a RAID array of fewer SSDs.

HPE Vertica recommends that you use enterprise grade direct attached storage (DAS) instead of storage area network (SAN) or Network Attached Storage (NAS). Doing so usually results in faster data retrieval.

Note: If you intend to use RAID 50 for your data partition, keep a spare node in every rack. This allows for manual failover of a Vertica node in the case of a drive failure. (Recovering a Vertica node is faster than rebuilding a RAID 50. To keep node recovery times at an acceptable rate, never put more than 10 TB compressed data on any node.)

Network

HPE Vertica recommends 10G networking over 1G networking in almost every situation.

Configuring and Monitoring Virtual Machines

HPE Vertica recommends that the virtual servers reach at least these performance goals:

Networking

- 100 MB/s of UDP network traffic per node on the private network (as measured by vnetperf)
- 20 MB/s per core of TCP network traffic on the private network (as measured by vnetperf)
- Independent public network

I/O:

- Measured by vioperf concurrently on all HPE Vertica nodes:
- 25 MB/s per core of write
- 20+20 MB/s per core of rewrite
- 40 MB/s per core of read
- 150 seeks per second of latency (SkipRead)
- Thick provisioned disk, or pass-through storage

Note: HPE Vertica does not support VMware Vmotion and Logical Volume Manager (LVM) on any drive where database (catalog and data) files are stored.

For best performance

- Disable CPU scaling on the physical hosts.
- Configure the disk blocks to align with the blocks that ESX creates. Unaligned blocks may cause reduced I/O performance during high load.

All virtual machines in a virtualized HPE Vertica cluster must be configured with the same specifications. HPE Vertica recommends that you configure your virtual machine as follows:

- One socket per virtual machine and 4 GB of memory per core in that socket
- Configure all volumes attached to each virtual machine as:
 - Thick Provisioned Eager Zeroed
 - Independent
 - Persistent

SAP BusinessObjects

OBR 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 6 GB. For further sizing of reports based on the number of users, you can distribute the Adaptive Processing Servers (APS). The confirmation options are available under OBR Central Management Console (CMC) home under **System Configuration Wizard**.

The load on the APS 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 200. If the load on the server per request is high, then you may see the Server Busy error message. Consider reducing this parameter and add one more Web Intelligence Server to support additional connection requests. While adding additional servers, always add one Web Intelligence Processing Server per CPU core per machine.

To summarize, 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 4 GB limit.

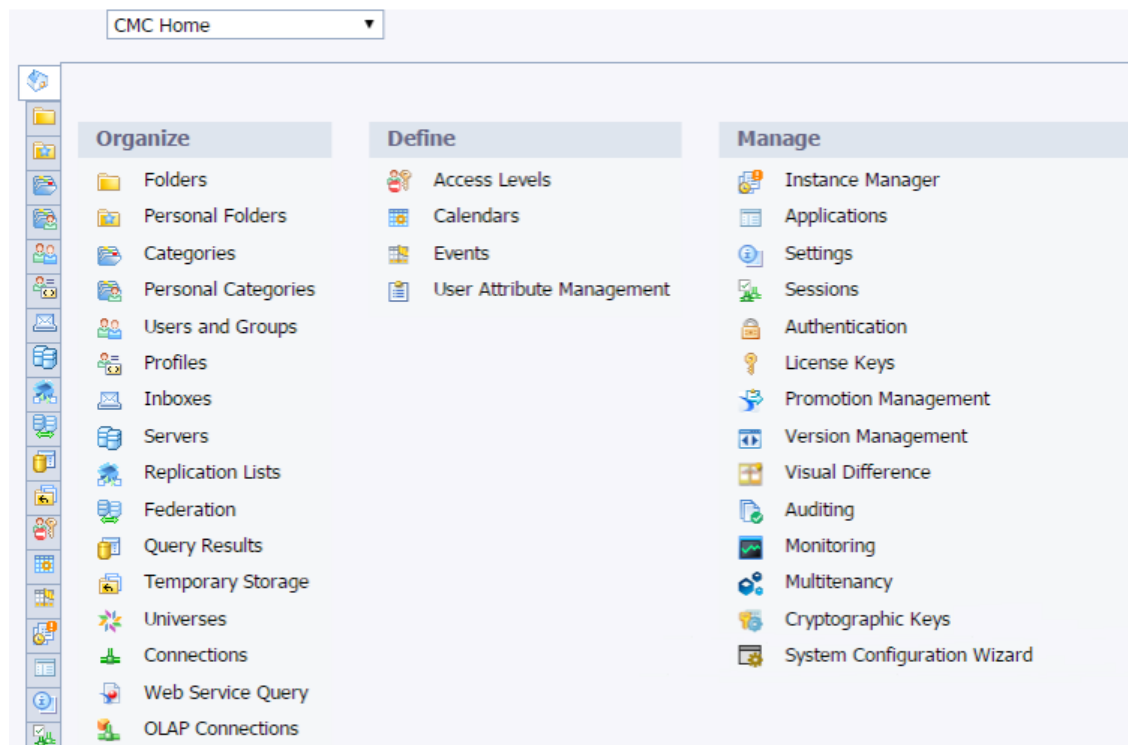
Adding Additional Web Intelligence (WebI) Servers

Perform the following steps to add additional WebI Servers:

1. Log in to Central Management Console (https://<OBRSysTem_FQDN>:8443/BOE/CMC), where, <OBRSysTem_FQDN> is the fully qualified domain name of the system where SAP BusinessObjects is installed.


The **Central Management Console** home page is displayed.

Central Management Console




2. Click on **Servers**  and then click **Servers List**.

The list of servers are displayed. You will find the **OBR.WebIntelligenceProcessingServer**.



Server Name ^	State	Enabled
OBR.AdaptiveJobServer	Running	Enabled
OBR.AdaptiveProcessingServer	Running	Enabled
OBR.CentralManagementServer	Running	Enabled
OBR.ConnectionServer	Running	Enabled
OBR.DashboardsCacheServer	Running	Enabled
OBR.DashboardsProcessingServer	Running	Enabled
OBR.EventServer	Running	Enabled
OBR.InputFileRepository	Running	Enabled
OBR.OutputFileRepository	Running	Enabled
OBR.WebApplicationContainerServer	Running	Enabled
OBR.WebIntelligenceProcessingServer	Running	Enabled

- Right-click on **OBR.WebIntelligenceProcessingServer** and then click **Clone Server**.



Server Name ^	State
OBR.AdaptiveJobServer	Running
OBR.AdaptiveProcessingServer	Running
OBR.CentralManagementServer	Running
OBR.ConnectionServer	Running
OBR.DashboardsCacheServer	Running
OBR.DashboardsProcessingServer	Running
OBR.EventServer	Running
OBR.InputFileRepository	Running
OBR.OutputFileRepository	Running
OBR.WebApplicationContainerServer	Running
OBR.WebIntelligenceProcessingServer	Running

- Properties
- Start Server
- Restart Server
- Stop Server
- Force Termination
- Enable Server
- Disable Server
- Metrics
- Placeholders
- Clone Server**
- Existing Server Groups
- Add to Server Group
- User Security
- Edit Common Services
- New >
- Translations
- Tools >
- Delete

The **Clone Server** dialog box is displayed.

- Enter the **New Server Name** as **OBR.WebIntelligenceProcessingServer1**.
The Clone to Node is selected by default as **OBR**.


Clone Server: OBR.WebIntelligenceProcessingServer

New Server Name:

Clone to Node:

An additional WebI Server is created to balance the load when the reports with data more than 4 GB is running simultaneously.

Manage ▾ Actions ▾



Server Name ^	State
OBR.AdaptiveJobServer	Running
OBR.AdaptiveProcessingServer	Running
OBR.CentralManagementServer	Running
OBR.ConnectionServer	Running
OBR.DashboardsCacheServer	Running
OBR.DashboardsProcessingServer	Running
OBR.EventServer	Running
OBR.InputFileRepository	Running
OBR.OutputFileRepository	Running
OBR.WebApplicationContainerServer	Running
OBR.WebIntelligenceProcessingServer	Running
OBR.WebIntelligenceProcessingServer1	Running

5. Click **OK**.
6. To enable the server, right-click on new **OBR.WebintelligenceProcessingServer1** and click **Enable Server**.
7. To start the server, right-click on **OBR.WebintelligenceProcessing server1** and click **Start Server**.

Chapter 4: Benchmark

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

Note: The benchmark tests were performed with HPE OBR database as the only database on a Vertica instance.

Benchmark Scenario 1

OBR and Vertica are installed with all Content Packs on different systems. This test was carried out using System content with a load of 5,000 hosts and all content with medium load(3.8 million throughput).

Windows

For Windows operating system, benchmark test was performed with OBR server on the Windows and Vertica on Linux.

Deployment Name	Hardware Configuration - Medium (Distributed)
OBR (Medium Deployment - OBR Server)	OBR and Vertica installed on different systems
	Model: HP ProLiant DL385 G8
	CPU: 8 (AMD Opteron 6386SE @2.80 GHz)
	RAM: 16 GB & Virtual Memory: 24 GB
	HDD size(preferably of type RAID 5) : 500 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 2012 R2

Linux

Deployment Name	Hardware Configuration - Medium (Distributed)
OBR (Deployment - Remote DB) – Vertica single node	OBR and Vertica installed on different systems
	Model: HP ProLiant DL385 G8
	CPU: 8 (AMD Opteron 6386SE @2.80 GHz)
	RAM: 32 GB & Virtual Memory: 34 GB
	HDD size(preferably of type RAID 5) : 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: RHEL 6.5

Benchmark Scenario 2

OBR and Vertica installed on different systems, with Vertica as a 3 node cluster. This test was carried out using System content with a load of 20,000 hosts and all content with large load (10.8 million throughput).

Windows

For Windows operating system, benchmark test was performed with OBR server on the Windows and Vertica on Linux.

Deployment Name	Hardware Configuration - Large (Distributed)
OBR (Deployment - OBR Server)	OBR and Vertica installed on different systems
	Model: HP ProLiant DL380 G9
	CPU: 16 (Intel Xeon CPU ES2630 @2.40 GHz)
	RAM: 32 GB & Virtual Memory: 47 GB
	HDD size(preferably of type RAID 5) : 650 GB
	Storage Type: HP Flexible Smart Array P440ar
	Drive Type: SAS
	Rotational Speed: 15K RPM
	Transfer Speed PHY 1:3 GBPS
	Disk Cache Battery: 1 GB
	OS: Windows 2012 R2

Linux

Deployment Name	Hardware Configuration - Large (Distributed)
OBR (Deployment - Remote DB) – Vertica single node	OBR and Vertica installed on different systems
	Model: HP ProLiant DL380 G9
	CPU: 16 (Intel Xeon CPU ES2670 @2.60 GHz)
	RAM: 64 GB & Virtual Memory: 34 GB
	HDD size(preferably of type RAID 5) : 850 GB
	Storage Type: HP Flexible Smart Array P440ar Drive Type: SAS Rotational Speed: 15K RPM Transfer Speed PHY 1:3 GBPS Disk Cache Battery: 1 GB OS: RHEL 6.6

Benchmark Scenario 3

OBR and Vertica are installed with SM Content Pack on the same system. It is then deployed in an HPOM environment. This test was carried out using System content with a load of 40,000 hosts and SM content with enterprise load (15.4 million throughput). Benchmark test was performed on the Windows operating system (for OBR server) and Linux (3 node Vertica cluster).

Windows

For Windows operating system, benchmark test was performed with OBR server on the Windows and Vertica on Linux (Vertica 3-node cluster).

Deployment Name	Hardware Configuration - Enterprise (Distributed)
OBR (Deployment - OBR Server) (specs mentioned are per node)	OBR and Vertica installed on different systems, with Vertica as a 3-node cluster
	Model: HP ProLiant DL385 G7
	CPU: 24 (AMD Opteron CPU 6180SE @2.50 GHz)
	RAM: 96 GB & Virtual Memory: 108 GB
	HDD size(preferably of type RAID 5) : 1.64 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 2012

Linux

Deployment Name	Hardware Configuration - Enterprise (Distributed)
OBR (Deployment - Remote DB)	OBR and Vertica installed on different systems, with Vertica as a 3-node cluster

Deployment Name	Hardware Configuration - Enterprise (Distributed)
– Vertica 3 node cluster (specs mentioned are per node)	
	Model: HP ProLiant DL380 G87
	CPU: 32 (Intel Xeon CPU ES2670 @2.60 GHz)
	RAM: 252 GB & Virtual Memory: 343 GB
	HDD size(preferably of type RAID 5) : 3.6 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: RHEL 6.5

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Performance, Sizing, and Tuning Guide (Operations Bridge Reporter 10.01)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to docfeedback@hpe.com.

We appreciate your feedback!