

SAP BusinessObjects Business Intelligence
Document Date: 12-15-2017

BI 4.2 Sizing Guide

© Copyright 2017 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, OS/2, Parallel Sysplex, MVS/ESA, AIX, S/390, AS/400, OS/390, OS/400, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere, Netfinity, Tivoli, and

Informix are trademarks or registered trademarks of IBM Corporation in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun

Microsystems, Inc., used under license for technology invented and implemented by Netscape.

MaxDB is a trademark of MySQL AB, Sweden.

SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver, SAP BusinessObjects and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any.

Nothing herein should be construed as constituting an additional warranty.

Disclaimer

Some components of this product are based on Java™. Any code change in these components may cause unpredictable and severe malfunctions and is therefore expressly prohibited, as is any decompilation of these components.

SAP Library document classification:
CUSTOMERS & PARTNERS

Documentation in the SAP Service Marketplace

You can find this documentation at the following address: <http://service.sap.com/sizing>

Table of Contents

1. Document History	6
2. Who should use this document?	7
3. What You Need to Know	8
4. Pre-Sizing Checklist	9
5. Post-Sizing Checklist	10
6. BI4 Machine Requirements	11
7. Introduction	12
8. Components and Platforms in Scope	13
9. Disclaimer	14
10. Resizing	15
11. Sizing Tools	16
11.1 BI4.2 Configuration Wizard	16
12. BI 4 Architecture	17
13. User Class Definitions	18
13.1 Users, Active Users, Active Concurrent Users	18
14. Conceptual tiers and services	19
15. Deployment and Sizing Methodology	20
15.1 Sizing Steps and Methodology	20
15.2 Processing and Memory Requirements Calculations	20
15.3 Prerequisites	21
15.4 Step 1: Quick Sizer Setup	22
15.5 Step 2: Intelligence DB Tier	22
15.6 Step 3: Intelligence Tier	23
15.7 Step 4: Application Tier	23
15.8 Step 5: Processing Tier	23
15.8.1 Web Intelligence	24
15.8.2 Analysis OLAP	24
15.8.3 Crystal Reports 2016	25
15.8.4 Crystal Reports for Enterprise	26
15.8.5 Dashboard Design	26
15.8.6 Adaptive Processing Service.....	27
15.9 Step 6: Sizing Analysis and Scale-out	28
15.10 Step 7: APS Configuration	29
15.10.1 APS Service Groupings	30
15.10.2 Basic APS Configuration Instructions	31
15.10.3 Adaptive Job Server	32
15.11 Step 8: Deployment and Monitoring	32
15.11.1 Scaling-out BI4 Servers.....	32
15.11.2 Building the System.....	33

15.11.3 Monitoring BI4	33
15.11.4 Using Audit Reports	35
15.11.5 Monitoring CPU and Memory Usage	35
15.11.6 BI Tool Simulation Workflows	35
16. Scheduling and Publishing	36
16.1 Scheduling and Publishing Best Practices for Performance	37
16.1.1 Using Dedicated Machines	37
16.1.2 Scaling	37
16.1.3 I/O	37
17. Service Tuning	38
17.1 Web Intelligence	38
17.1.1 Parallel refresh	41
17.1.2 SAP HANA data sources	42
17.1.3 Disk space of your system	42
17.1.4 BI Commentary feature	42
17.1.5 Decimal Type for big number	42
17.1.6 Sizing for SAP BusinessObjects Web Intelligence Upgrade	42
17.1.7 DSL Bridge	42
17.2 Central Management Server (CMS)	42
17.3 Crystal Reports Cache Service	43
17.4 File Repository Service (FRS)	43
17.5 Connection Server	43
17.6 Web Application Server	44
17.7 Scale-out Memory Expectations	44
17.8 Crystal Reports	45
17.9 Dashboards	46
17.10 Analysis OLAP	47
18. SAP BW Considerations and Recommendations	48
18.1 BW Configuration	48
18.2 BEx Query Performance	48
18.3 Navigational Attributes	48
18.4 BW Specific Tuning and Configuration	48
18.4.1 Set Specific Properties	48
18.5 Additional Content	49
19. CMS Database Tuning	50
20. Virtualization: Service Level and Performance	51
21. Desktop Virtualization	52
22. Sizing Example	53
Step 1: Quick Sizer Setup	53
Step 2: Intelligence Tier DB	54
Step 3: Intelligence Tier	54
Step 4: Application Tier	54
Step 5: Processing Tier	54
23. Sizing References	59

23.1 Sizing SAP BusinessObjects Live Office	59
23.2 Sizing SAP BusinessObjects Explorer	60
23.3 Sizing SAP BusinessObjects Mobile	61
23.4 Sizing SAP BusinessObjects Design Studio	61
23.5 Sizing SAP Lumira 2.0.....	61
24. Appendix.....	62
24.1 BI4 P&R Benchmark Testing Details	62
24.1.1 Test report documents and data characteristics	62
24.1.2 Report data description	65
24.1.3 CMS repository data description	66
24.1.4 Test hardware specifications	67

1. Document History

Date	Details
June, 2013	Sizing methodology reworked to be easier to follow as well as more comprehensive.
September, 2013	Additional tuning suggestions added (Apache Split, BI Pattern Books). BI 4.1 sizing added.
December, 2013	<p>Added (restored) Analysis OLAP tuning, enhanced Scheduling and Publishing guidance, Resizing guidance, added audit database reporting guidance to the Monitoring section.</p> <p>Added sizing checklists.</p>
February, 2014	<p>Added Connection Server tuning: JVM memory settings for JDBC connectivity including HANA.</p> <p>Updated Web Intelligence suggested maximum connections, CMS sizing guidance, importance of CMS database tuning. Added a CMS Database Tuning section.</p> <p>Updated note about multiple Promotion Management instances: supported starting BI 4.1 SP3.</p> <p>Added a link to Design Studio sizing. Updated link to the PAM for BI 4.1.</p> <p>Updated Apache Split links to new content for BI 4.1 and Apache 2.4.</p>
June, 2017	<p>Updated the sizing guide w.r.t BI 4.2 SP4.</p> <p>Included Web Intelligence sizing updates and Quick Sizer tool.</p>
July, 2017	Updated Web Intelligence sizing and APS services.
Dec, 2017	Updated guide with Web Intelligence changes and re-structured the document.

2. Who should use this document?

This document provides recommendations and best practices to help you deploy and scale the various services in SAP BusinessObjects BI 4.

We recommend consulting this guide if you are:

- Working on a [Sizing](#) strategy, especially in conjunction with the SAP Quick Sizer.
- Planning to deploy SAP BusinessObjects BI 4 software.
- Optimizing or tuning an existing SAP BusinessObjects BI 4 deployment.

For all deployments, it is necessary to consult a sizing expert in order to validate sizing exercises after initial planning using the SAP Quick Sizer.

Sometimes customers, partners, or consultants develop their own applications, strongly modify SAP's out-of-the-box solutions, or implement complex integrations involving multiple systems. In these cases, when sizing for the large deployments or when your sizing requirements need to be extremely precise, you should consider an "expert sizing". Expert sizing is a hands-on exercise performed by or in collaboration with an SAP Professional, where customer-specific data is analyzed and used to achieve a sizing result with greater precision. The main objective is to determine the resource consumption of customized content, applications, and usage patterns by taking comprehensive measurements. For more information, see [Sizing Quick Sizer Tool](#).

3. What You Need to Know

To Size your deployment effectively and accurately, you need to know the following, which are explained further in this document:

- The SAP Application Performance Standard (SAPS) is a hardware-independent unit that describes the performance of a system configuration in the SAP environment. It is derived from the Sales and Distribution (SD) Benchmark, where 100 SAPS is defined as the computing power to handle 2,000 fully business processed order line items per hour. (For more information about SAPS, see **Measuring in SAPS** on <http://www.sap.com/benchmark>).

The SAPS rating of your hardware is the performance rating of the hardware you are going to use for your deployment. You also need to determine the number of actual* cores in the system and in turn determine the SAPS-per-core value.

- If you are using virtualized hardware, you need to know if the IT department will allow CPU and memory reservations. Without reservations, the SAPS number will not be accurate and your sizing calculations will be incorrect.
- The number of users of each BI Tool that you intend to use. Within that context, the types of each user (consumer, business user, expert). This is explained further in this document.
- The relative document sizes for each BI Tool.
- The data sources the BI Tools will access.

*'Actual' cores refer to 'full' cores in a CPU. Hyper-threaded cores should not be counted for the purposes of sizing.

4. Pre-Sizing Checklist

These are the things you need to do before you start your sizing exercise:

- How many Active Concurrent users will you have for each BI tool? See “Users, Active Users, Active Concurrent Users” in this document for more information.
- What types of users will be using each tool? Information Consumers (low workload), Business Users (medium workload) or Experts (high workload)? See “User Class Definitions” in this document for more information.
- Do you know how users will use each BI tool? Will they refresh documents rapidly? Will they schedule them at off-peak time? Will they open three documents at once? Will they build dashboards that consume multiple reports? How users use the tools affects the workload they produce.
- What types of machines will you be deploying to? What is the SAPS rating of each machine? See “BI4 Machine Requirements” in this document for more information.
- What types of data sources will your users access? See “Data Sources” in this document for more information.
- Will you be deploying to a virtualized environment? See “Virtualization: Service Level and Performance” in this document and <http://www.sap.com/bivirtualization> for more information.
- Do not assume you are finished sizing if you only fill in the Quick Sizer (which could be accessed using <http://service.sap.com/quicksizer>). These tools do not cover all deployment scenarios nor all data sources. You must follow the Deployment and Sizing Methodology section below to accurately size your deployment.

5. Post-Sizing Checklist

Here are the things you should consider once you have followed the sizing steps in this guide:

- Do a sanity check on your deployment landscape. Are there too many I/O-intensive services running on the same machine? Is it configured to allow for peak usage? I.e., do you think it will run comfortably (around 60% average utilization) so that peak times don't cause issues?
- Are the other parts of your IT infrastructure ready and sized for the load the BI system will place on them? Has your BW system been resized and patched? Are the relational databases tuned and ready for the load?
- Check the [BI Pattern Books](#) for best practices in how to configure your deployment.
- Check <https://www.sap.com/products/bi-platform.html> for the latest information on BI4, especially the [Upgrade](#) area if that applies to you.
- To learn more about getting the most out of your infrastructure, see the resources available at <http://www.sap.com/bisizing>.
- To ensure good performance in a virtualized environment, see "Virtualization: Service Level and Performance" in this document and <http://www.sap.com/bivirtualization> for more information.
- Install auditing reports so you can monitor your system. A set of auditing reports is available for download [here](#).
- Schedule a re-sizing in six months. See the Resizing section in this document for more information.

6. BI4 Machine Requirements

For BI4, the minimum hardware for evaluation purposes is 8000 SAPS of processing power and 16 GB of memory. This is defined in the [BI 4.2 PAM](#), found here: <http://scn.sap.com/docs/DOC-70167>

You can determine the hardware requirements for deployments of any kind by following the remainder of this guide.

7. Introduction

This document discusses the Sizing of the Business Intelligence 4 suite of services, hereafter referred to as BI4. The BI4 Suite is designed to serve customers from small business to SAP's largest customers. The number of users, types of users, usage patterns, number of BI tools included in the suite, number of data sources supported by the suite as well as the deployment options supported by the suite all factor into a series of variables that affect the successful deployment of the BI suite. No configuration fits all customers. The purpose of this document is to help guide you through the Sizing Exercise for BI4.

Sizing BI is very different compared to sizing of other types of Enterprise software. Deriving BI is a very resource intensive task. The act of extracting information from a potentially large amount of data requires adequate amounts of processing power and exercises all the important subsystems of a computer: CPU, memory, disk and network. Having the right amount of capacity of these four aspects of your system is key to success. BI can also be very bursty, since the load can rely a lot on the schedule of its users.

A BI system is used much like Google is: searching for information interactively. If Google required a number of seconds to respond, most users would consider "the Internet broken". They have an expectation and often a business need to have a responsive system available to them. The purpose of this document is to help ensure your users can access BI decision-making information in a timely manner.

No tool or document can replace human judgment. So while this document attempts to cover as many aspects of the Sizing Exercise as possible, you must always consult sizing experts at SAP.

Make sure you have the latest version of this document and other BI4 Sizing resources by visiting <http://www.sap.com/bisizing>

8. Components and Platforms in Scope

The guidance in order to estimate hardware, i.e. referenced in this document are based on the following:

- Analysis OLAP
- Crystal Reports 2016
- Crystal Reports for Enterprise
- Dashboard Designer
- Web Intelligence

General sizing tips are also provided for:

- Explorer
- Mobile BI
- Live Office
- BI4 Platform services: Promotion Management, Search, Data Federation
- SAP BW
- Web Intelligence

It is assumed that readers of this document are already familiar with core concepts from the [Business Intelligence Platform Administrator Guide](#) (BI4 Admin Guide). In the BI4 Admin Guide, you can find conceptual information and technical details on a wide number of topics which are intentionally not covered in this guide.

9. Disclaimer

This document demonstrates how someone might perform sizing of a BI4 system. The methodology and sample walkthroughs offered here are examples of the tasks and thinking involved. The performance and functioning of an actual system may vary for many reasons. The examples offered here should not be considered as a guarantee of success of a particular deployment.

You should work with your SAP account representative and ensure you get the advice of a sizing expert when making final decisions with regards to a BI4 deployment.

10. Resizing

It is recommended that you resize your environment every six months or sooner if you have had significant changes in the scope of your deployment. It can be difficult to anticipate how your users will make use of the BI environment at the outset of a project. Resizing is recommended to ensure your systems are configured for how your users are using the system.

In order to be informed accurately about the usage of the system, it is recommended that you take advantage of the auditing features and auditing reports available for BusinessObjects BI4. A set of auditing reports is available for download [here](#). These reports will help you determine how your BI services are being used.

Monitoring and probe reports are also recommended. Probe reports can give you alerts when time to run expectations are longer than expected. See the [BI4 Admin Guide](#) for more information on creating probes. CMS database performance should also be monitored.

Monitoring and recording of the basic CPU and memory usage of the machines in your deployment is also recommended. The goal is to have an average CPU usage of 60% in order to handle peaks in the range of 80%. Analyzing the historical usage of your deployment can help you determine if more resources are needed for a particular node. Similarly, if memory is frequently being fully used, you may be experiencing reduced performance and need to add more.

11. Sizing Tools

Quick Sizer Tool

Quick Sizer is a Web-based tool designed to make the sizing of SAP Business Suite easier and faster. It has been developed by SAP in close cooperation with all platform partners and is free of cost.

With Quick Sizer you can translate business requirements into technical requirements. Simply fill in the online questionnaire, an up-to-date survey that is based on business-oriented figures. The results you obtain can help you select an economically balanced system that matches your company's business goals. This is especially useful for initial budget planning.

Quick Sizer calculates CPU, disk, memory and I/O resource categories based on reports' size usage and the number of users working with the different SAP solutions in a hardware and database independent format.

11.1 BI4.2 Configuration Wizard

Starting with SAP BusinessObjects BI 4.2, the Configuration Wizard helps the sizing exercise by defining a number of APSEs based on the deployment size chosen. This document explains how to perform a custom sizing exercise. The Configuration Wizard's APSEs can be used as a time saver when it comes to defining the APSEs for your deployment. For more information, see the [BI4 Admin Guide](#).

12. BI 4 Architecture

The BI4 architecture can be a complex set of services and technologies. The following is a representation of the system from a technology perspective:

Find the link to the latest BI 4 Architecture diagram [here](#).

For the purposes of sizing BI4, the essential aspects of the BI4 architecture are as follows:

The BI4 Platform: This is a set of base services that enable BI services to run in a distributed deployment. Each node in a BI4 cluster runs a Server Intelligence Agent (SIA) which runs selected services on that node. Nodes in a cluster communicate over a service bus.

The BI4 Platform and suite of services are designed using a 64-bit architecture to ensure scalability to the largest customers.

The CMS: This is where all the objects and configuration information that control and secure the platform are located. The CMS is a combination of service plus database for storing and retrieving this information.

Basic Services: Core services such as the File Repository (FRS) need to run somewhere in the cluster.

APS: The Application Processing Service (APS) is a service host. A large number of the BI services run inside APSes. There can be more than one APS running on a given node, and there can be many nodes running APSes in a BI landscape. Services running in the APS are Java-based and the APS controls the JVM that the services run in. There is one JVM per APS. Non-APS Services: Not all BI services run inside an APS. Those that don't are run separately by the SIA.

BI Services: The BI services in the system have specific architectural considerations:

Crystal Reports: Crystal Reports services are available in two versions: Crystal Reports for Enterprise is designed to have great Universe and SAP connectivity. It is a newer version of Crystal Reports. It uses data access technology common to the other tools in the suite.

Crystal Reports 2016 is the original Crystal Reports product line, maintained and available for backward compatibility.

Web Intelligence (aka WebI): WebI provides analysis and reporting capabilities to the suite. It provides compatibility with old Universes (UNV), new Universes (UNX) as well as SAP BW and SAP HANA. The Visualization Service (aka CVOM) provides data visualizations (charting, graphics) services to WebI, the DSL Bridge house the query generation, metadata definition and retrieval on BusinessObjects universes and data retrieval for SAP BW. The data federation service handles multiple data provider sources into the universe by the SAP BusinessObjects Data Federator component and the BI Commentary service lets you add and manage comments in your Web Intelligence reports.

13. User Class Definitions

Three user classes are defined below. Class definitions are used to identify how much the system will be used and how much load will be placed on it. Information Consumers use the system sparsely and Expert Users use the system heavily. Knowing the mix of user classes is important since it directly affects the performance and resources required by the system.

Customers often underestimate the amount of use that a user will make of BI system.

Information Consumers - The least active of all the user types. Information consumers spend an average of 300 seconds (5 minutes) idle in between navigation steps. These users typically view predefined and static content and perform relatively little drilling and filtering on their own.

Business Users - These users perform some moderate amount of drilling and filtering on their own. Business users spend an average of 30 seconds idle in between navigation steps.

Expert Users - The most active of all the user types. Expert users spend an average of 10 seconds idle in between navigation steps. These users are much more likely to perform resource-intensive operations in the system including ad-hoc analysis and customization of reports, retrieving a large number of rows, and heavy client-side filtering.

13.1 Users, Active Users, Active Concurrent Users

Many sizing exercises begin with a total number of users, for example: the customer needs a system sized for 10,000 **users**. The next step is to determine how many of those users will be logged-in to the system at the same time. The best way is to work with the customer to understand their usage patterns. If a customer can't provide guidance on this number, it's common to estimate 10% as a minimum. We call these **active users**.

In this example, there would be 1,000 active users. Because users can spend much of their time being idle after logging-in, you also need to determine how many of the active users in the system will be concurrently generating load. This is another opportunity to take customer usage patterns into account, but if there is no information available on this, the most common estimate is 10% as a minimum. We call these users who are logged-in and concurrently generating load, **active concurrent users**. In our example, there would be 100 active concurrent users.



When you use the BI 4 sizing tools, specify your inputs in terms of Active Concurrent Users.

In the example above, we went from 10,000 users to 100 active concurrent users, which translate to 1% concurrency. It's important to understand that the typical concurrency recommendations vary depending on the size of deployment and other factors including the kind of work users perform.

If you expect or experience typical concurrency higher than these nominal ratios, you should expect a heavier load and should compensate accordingly. For guidance in such a case, you may want to use the Quick Sizer and specify more users as 'expert users' to account for the increased use of the system.

14. Conceptual tiers and services

The SAP BI platform can be thought of as a set of conceptual tiers. A subset of each list of services is provided here for convenience. Refer to the [BI4 Admin Guide](#) for the complete list.

Application Tier

- Web Application Service (typically Tomcat)
 - CMC – Central Management Console
 - BI launch pad

Intelligence Tier (Referred to in the [BI4 Admin Guide](#) as “Management Tier”)

- CMS*¹
- FRS

Processing Tier

- Adaptive Job Service
- Adaptive Processing Service (APS)
 - Analysis OLAP and Multi-Dimensional Analysis Service (MDAS)
 - DSL Bridge Service
 - Visualization Service (CVOM)
 - BI Commentary Service, Client Auditing Service, Monitoring Service, Web Intelligence Monitoring Service, Document Recovery Service....
- Crystal Reports Processing Service
- Crystal Reports Cache Service
- Dashboard Cache Service
- Web Intelligence Processing Service

*¹ The database which the CMS uses for its repository may sometimes be referred to as the “Intelligence Tier DB

15. Deployment and Sizing Methodology

Sizing a BI4 deployment requires a reasonable degree of planning so that calculations and predictions can be made about the needs of the system. The number of users and the needs of those users can be used to predict load on the system. The types of data sources used have an influence on the load and needs of the system. The needs of the users include the BI services that they need to use. Some users use the services a lot, some use them a little. Some reports are scheduled to be processed at night and viewed during the day. Some reports need to be refreshed when viewed, which causes more load as the number of users increases.

Once the user requirements are defined, the system can be defined that will achieve the required amount of processing. The final step in the sizing is to apply it to the hardware landscape. The deployment hardware can range from many small machines to one large machine. The Sizing Exercise includes the allocation of BI services to the nodes in the system, considering the CPU, memory, disk and network capabilities of nodes to be used in the construction of the system landscape.

15.1 Sizing Steps and Methodology

The basic approach to sizing is working through the requirements of all the tiers in the deployment, accounting for the processing and memory requirements of each subsystem. Most subsystems and services have special considerations that need to be taken into account.

Once the processing and memory requirements are determined, the next major step is to fit the services to the hardware landscape. This requires knowing the SAPS per core rating of the computer to be used as well as the amount of memory available to the machine or machines.



You must determine the SAPS rating of the machines that will make up your deployment. You cannot assume any SAPS rating from this or any other document.

You can determine the SAPS rating of machines by visiting the website: [SAP SD Standard Application Benchmark Results](#)

The Quick sizer is used as a tool to help calculate the processing load to be handled by each tier in your system. Depending on the BI tool under consideration, additional calculations may need to be made with corresponding changes and/or additions to the sizing.

15.2 Processing and Memory Requirements Calculations

For each tier, especially the Processing Tier, the processing power and memory requirements can be determined by Quick Sizer and as suggested by Configuration wizard in BI systems. This is the core of the Sizing Exercise. Once these numbers are determined, scale-out to deployment hardware can be done.

The processing power is calculated and specified in a processing unit called SAPS. This is a processor and computer independent unit that describes the processing power of a CPU in each computer. SAPS is an SAP-derived processing rating that allows you to calculate the power required by a deployment. It takes into account the I/O processing capabilities of the computer in addition to its raw computation abilities. Computers can have different SAPS processing capabilities even with the same type of CPU.

15.3 Prerequisites

The goal of the sizing exercise is to calculate the peak load that will be placed on the system. In order to proceed with the sizing steps, you need to know the following information listed here. Gathering this data accurately is the most important part of the sizing exercise since all the sizing calculations derive from this information.

Users: How many Information Consumers, Business Users and Expert Users of each type of BI tool. An average user workflow is also important to know. If you expect users to open five BI documents and refresh them all at the same time, that's five times the load of one user. That needs to be accounted for since the system needs to handle that load.

It is very important to know if the common workflow of the users is going to include refreshing reports and if so, how frequently. Will the reports be scheduled to run at night and only be viewed during the day or will they be refreshed by users every hour? That's an important part of the load prediction and thus the sizing estimate.

Data Sources: What types of data sources will be used: direct-access SQL databases, UNV Universes, UNX Universes, SAP BW, BW on HANA or HANA direct? It can be a mix. It's important to know which types of data sources will be used for the majority of BI processing so that peak load can be predicted. Some customers expect their mix of data sources to change over time. This is an important thing to consider.

Document Size: The relative size of each document is important to know. Will most documents be small and cause little impact on the system or will most be large and require a lot of processing? This should be determined for each BI tool expected to be used.

Document Complexity for Web Intelligence:

In viewing mode, two additional factors impact SPECIFIC system resource usage:

- data volume: the larger the data volume, the more systems resources are required by Web Intelligence Processing Tier.
- query complexity: the more complex the query (for example, queries with multiple filters and sorts), the more systems resources are required by Web Intelligence Processing Tier.

In report editing mode, the following factor also impacts system resource usage:

- complexity of the document structure (for example, documents that include the auto-fit setting, page breaks, sorts, alerters, styles, sections, expand/collapse etc.) increases the amount of system resources required by Web Intelligence Processing Tier when a report is generated.

SAPS Rating: In order to know how much computing resources are needed for a deployment, you need to know the computer's performance rating. This is measured in SAPS. You may be building the system using one computer or many. You may have the computer(s) specified in advance or not until the initial sizing has been calculated.

The SAPS rating of the system's intended hardware allows you to determine the required amount of hardware to meet the system's processing requirements. The SAPS rating takes into account CPU processing power as well as the computer's I/O capabilities. To learn the exact amount of SAPS produced by your hardware, check the [SAP SD Standard Application Benchmark Results](#) website.

For the sizing exercise, the amount of **SAPS per core** needs to be determined. The benchmark shows the SAPS rating for the computer. You then need to divide by the number of actual cores of the CPU.

Note: Hyper-threaded cores should not be used for this calculation.

15.4 Step 1: Quick Sizer Setup

With BI 4.2 sizing onwards, we are enhancing the capabilities of “BI4 Resource Usage Estimator”, to estimate the hardware requirements, based on customer needs.

With additional components and functionalities being introduced in consequent releases of BI 4.2, we have decided to upgrade the “BI4 Resource Usage Estimator” and move it into “Quick Sizer”. For more details, refer this [blog](#).

For more information, see the following:

- [Quick Sizer Tool](#)
- [Using Quick Sizer for Sizing SAP BusinessObjects Business Intelligence Suite Applications](#)

The Quick Sizer should be initialized with the number of users of each type for each type of BI tool. The report size sliders should also be set to reflect the sizes and types of documents that will be processed by the system.

It might look like this:

Input screen

Table 1: Throughput - SAP BusinessObjects Business Intelligence

Clear Insert The buttons in this line only work for marked lines in the lower table

Element	Element short text	A/P	TI	Information consumers (300 sec. think.)	Business users (30 sec. think. time)	Expert users (10 sec. think. time)	% Small size reports	% Medium size reports	% Large size reports
ANA-OLAP	SAP BusinessObjects Analysis, edition for OLAP	A	S			15	25	50	25
CRYSTAL	SAP Crystal Reports 2016	A	S		80	20	25	50	25
CRYST-ENT	SAP Crystal Reports for Enterprise	A	S						
DASHBOARD	SAP BusinessObjects Dashboards	A	S						
WEB-INTELL	SAP BusinessObjects Web Intelligence	A	S						
WEBI-ON-BW	SAP BusinessObjects Web Intelligence on BW	A	S		150	10	25	50	25
DESIGN-STU	SAP BusinessObjects Design Studio	A	S						
LUMIRA	SAP Lumira	A	S						
LUMIRA-BW-LIVE	SAP Lumira for BW Live	A	S						
LUMIRA-HANA-LIVE	SAP Lumira for HANA Live	A	S						

Output screen

Result

Print page Result categories How to interpret the results

All SAPS Memory and disk Data IO Log IO

SW component	Software component	CPU cat.	SAPS (total, 2-ter.)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory cat.	Memory (total, 2-ter. MB)	BO app. tier Memory	BO int. tier Memory	BO int. tier DB Memory	BO pro. tier Memory
BI_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

15.5 Step 2: Intelligence DB Tier

The Intelligence database is the CMS database. For larger deployments, the Intelligence database should be installed on a dedicated machine.

Be sure to follow your database vendor’s guidance regarding how to size and scale the deployment of the database. Be aware that your database vendor may have strict recommendations regarding deployment to a virtual machine.

If you intend to use an existing database that will be shared with other workloads, you must account for the processing and memory requirements shown in the Quick Sizer. I.e., it should have free headroom to handle the given amount of processing (SAPS) and have the given amount of memory free to do that processing.

SW component	Software component	CPU cat	SAPS (total, 2-tier)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory Cat	Memory (total, 2-tier, MB)	BO app. tier Memory	BO int. tier Memory	BO int. tier DB Memory	BO pro. tier Memory
BI_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

15.6 Step 3: Intelligence Tier

The workload of the Intelligence Tier is primarily the CMS. There are other services included in the Intelligence Tier such as the FRS and Promotion Management.

SW component	Software component	CPU cat	SAPS (total, 2-tier)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory Cat	Memory (total, 2-tier, MB)	BO app. tier Memory	BO int. tier Memory	BO int. tier DB Memory	BO pro. tier Memory
BI_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

15.7 Step 4: Application Tier

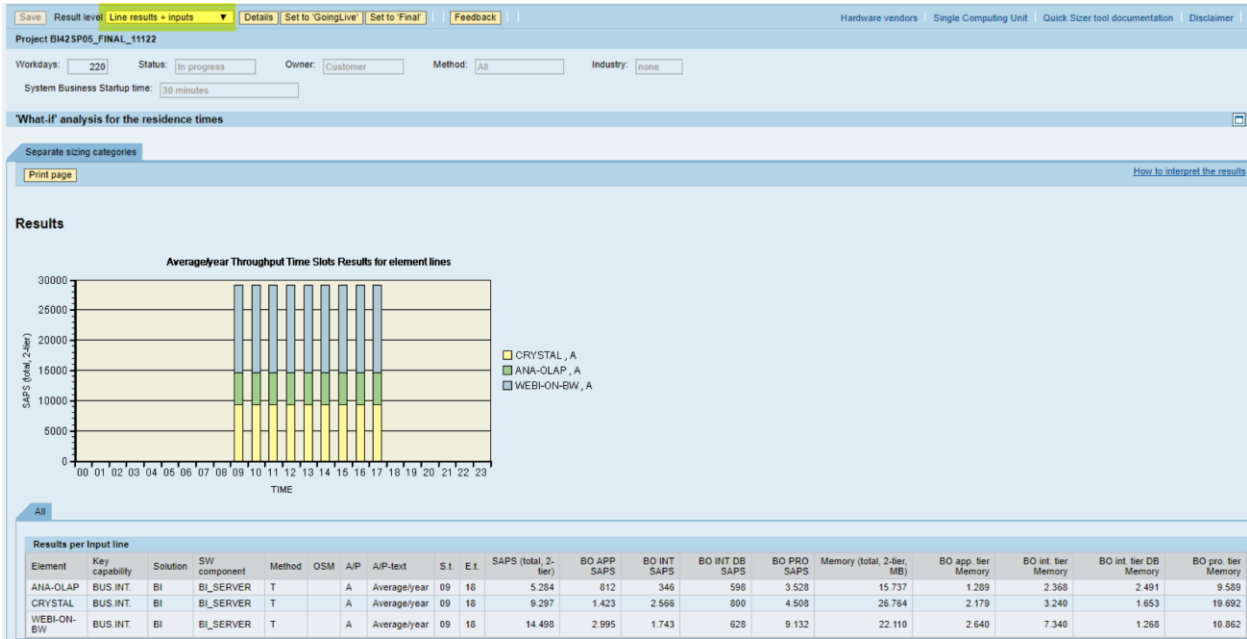
The Application Tier is the Web Application Server processing. Almost all interaction with the system is through the Web Application Server, including mobile clients.

SW component	Software component	CPU cat	SAPS (total, 2-tier)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory Cat	Memory (total, 2-tier, MB)	BO app. tier Memory	BO int. tier Memory	BO int. tier DB Memory	BO pro. tier Memory
BI_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

It is recommended to follow guidelines as per Quick Sizer for to allocate memory for Application Tier processing request. For every 500 active concurrent users, it is recommended that an additional Web Application Server instance be created.

15.8 Step 5: Processing Tier

Quick sizer provides the total memory and SAPS recommendations for Processing Tier, which is a resultant of the number entered into the Quick Sizer for multiple BI tools. However, to understand the usage w.r.t. individual BI tools, choose the option of changing the result level to "Line results + inputs" (as seen in the below screenshot).



SW component	Software component	CPU cat	SAPS (total, 2-tier)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory (total, 2-tier, MB)	BO app. tier Memory	BO int. tier Memory	BO int. tier DB Memory	BO pro. tier Memory	
BL_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

In order to do tool-specific sizing using the Quick Sizer, the numbers entered into the Quick Sizer are changed as each BI tool is analyzed. To get the starting independent sizing for a certain BI tool, set the user numbers for the other tools to zero.

Note: The Quick Sizer allocates the Crystal Reports Cache Server to the Processing Tier. In general, it is recommended that the Crystal Reports Cache Server be run on the same machine as the Crystal Reports Processing Servers. In a scale-out scenario, a Crystal Reports Cache Server instance should be allocated to run on each of the machines running Crystal Reports Processing Servers. It can be disabled on nodes that are not running Crystal Reports Processing Servers.

15.8.1 Web Intelligence

To size Web Intelligence independently, set the user numbers appropriately and set the user numbers for all other tools to zero.

15.8.1.1 Data Visualization

The Data Visualization (CVOM) service is used by Web Intelligence. Memory allocation for CVOM needs to be done, based on "how the charts are created".

15.8.2 Analysis OLAP

To size Analysis OLAP (aka MDAS) independently, set the user numbers in the Quick Sizer appropriately and set the user numbers for all other tools to zero.

It is recommended that an APS be created for the purpose of running the Analysis OLAP service. This allows the Analysis services to be monitored and configured easily.

Note: the default limit for the number of concurrent users of the Analysis OLAP service is 15. This should be changed as you deploy the Analysis OLAP service to one or more APS instances.

Multi Dimensional Analysis Service

<input type="checkbox"/> Use Configuration Template	
<i>MDAS Configuration</i>	
Maximum Client Sessions:	100
Maximum number of cells returned by a query:	100000
Maximum number of members returned when filtering:	100000
<input type="checkbox"/> Restore System Defaults	
<input type="checkbox"/> Set Configuration Template	

Note: Analysis OLAP is also referred to as the 'Multi-Dimensional Analysis Service' or MDAS.

See the [BI4 Admin Guide](#) for more information on changing Analysis OLAP settings.

15.8.3 Crystal Reports 2016

To size Crystal Reports 2016 independently, set the user numbers appropriately and set the user numbers for all other tools to zero.

If your deployment will be to one large machine, you can record these requirements in your sizing document and move to the next tool. Proceed here if you will be scaling out your deployment.

The Crystal Reports 2016 services required resources are determined by the processing requirements.

The number of cores is determined by the SAPS shown in the Quick Sizer for the Processing Tier (while having set only the Crystal Reports 2016 user numbers). To determine the required CPU cores, divide the SAPS needed by your SAPS per core calculation (see the Prerequisites section above for more information). You should round up when you encounter fractional cores.

Once you know the number of cores needed, you can calculate the memory requirement by multiplying by 1.25 GB per core.

- In your sizing document, add in the SAPS and memory for Crystal Reports 2016.

15.8.3.1 Crystal Reports Processing Service

When sizing for workflows involving a small number of users viewing Crystal Reports, you may notice high values for recommended memory. You may also notice that for a larger number of users the required memory is almost constant (it does not increase much when more users are added). This is the expected behavior and it is due to the internal functionality of the Crystal Reports processing servers and engine, which are optimized to offer smoother overall performance, especially for a higher number of users over a longer period of time. The optimal allocated memory is based on an algorithm that uses the number of CPUs on the machine rather than the number of users in the system. Because of this, it's not recommended to have more than one Crystal Reports Processing Service on any given machine. As you scale-out your deployment, it's better to add instances to other machines that don't

current have one.

Required memory is calculated as follows:

Crystal Reports 2016:

of children = # of logical cores * 2.5 Each
child can use up to 500 MB of RAM
➤ 1.25 GB per core

Crystal Reports for Enterprise:

of children = # of logical cores ÷ 2
Each child can use up to 3 GB of RAM
➤ 1.5 GB per core

Because of the volume and frequency of communication between the Processing Service and the Cache Service, it's recommended that you deploy in pairs. That is, on each machine where there is a Processing Service, also deploy a Cache Service and vice versa.

Data from the database is written to disk in the temporary directory during processing. The speed of this disk will have an effect on processing performance.

Note: Each Crystal Reports child process (as noted above) can process a number of report processing jobs at once. A job can potentially be shared among multiple user requests depending on sharing criteria. Crystal Reports 2016 child processes default to 40 jobs each while Crystal Reports for Enterprise child processes default to 20 jobs each. The number of jobs per child determines how many child processes are created. The number of users served by a child process can be greater than the job limit if job sharing is possible. Sharing can occur when a job's report processing criteria match, such as database security and parameter prompts. It is the role of the cache server to control job sharing.

15.8.4 Crystal Reports for Enterprise

To size Crystal Reports for Enterprise independently, set the user numbers appropriately and set the user numbers for all other tools to zero.

The Crystal Reports for Enterprise services required resources are determined by the processing requirements.

The number of cores is determined by the SAPS shown in the Quick Sizer for the Processing Tier (while having set only the Crystal Reports for Enterprise user numbers). To determine the required CPU cores, divide the SAPS needed by your SAPS per core calculation (see the Prerequisites section above for more information). You should round up when you encounter fractional cores.

15.8.5 Dashboard Design

To size Dashboard Design independently, set the user numbers appropriately and set the user numbers for all other tools to zero.

15.8.5.1 Dashboard Cache Service, Dashboard Processing Service

Because the technology used for the Crystal Reports and Dashboard Cache and Processing Services are largely the same, the scalability characteristics and recommendations are also the same. Refer to the [17.3 Crystal Reports Cache Service](#) and [15.8.3.1 Crystal Reports Processing Service](#) and use the same information for Dashboards.

Required memory is calculated as follows:

of children = # of logical cores ÷ 2 Each
child may use up to 1.5 GB of RAM

15.8.6 Adaptive Processing Service

15.8.6.1 *Promotion Management Services*

'Promotion Management' refers to the following group of services:

- Visual Difference Service
- Promotion Management ClearCase Service
- Promotion Management Service
- Trace Log Service
- Version Management Service

It is recommended that these services be grouped into a Promotion Management-specific APS for medium and larger deployments. For medium to large sized deployments, this APS should be allocated at least 750 MB of memory. For very large deployments, this APS should be allocated at least 1GB of memory.

For processing calculations, a nominal amount of SAPS of 1000 should be used.

- In your sizing document, add in the SAPS and memory for Promotion Management.

15.8.6.2 *Platform Search Services*

Sizing of the Search Service depends on your company's intended use of search. It can be turned off, used lightly or used heavily. If you intend to use search, the amount of resources to allocate to it depends on how much and how frequent content is changed in your system.

For a system where content doesn't change frequently during the day, you can schedule search to index at night. It is recommended that a separate APS be created just for the Search Service. This will allow for easy monitoring and configuration of resources for search.

15.8.6.3 *Data Federation Service*

If you intend to use multi-source Universes, you need to size for the Data Federation service.

It is recommended that an APS be created for the purpose of running the Data Federation, especially for large deployments. This allows the data federation services to be monitored and configured easily.

For medium deployments, allocate 1 GB of RAM. For larger deployments, allocate 3 GB of RAM.

- In your sizing document, add in the SAPS and memory for the Data Federation service.

For processing calculations, a nominal amount of SAPS of 1000 should be used.

For more thorough sizing of the Data Federation service, see the Data Federation Service Sizing and Tuning Companion Guide, located on [SAP Service Marketplace \(SMP\)](#).

15.8.6.4 *Adaptive Connectivity Service*

For medium and larger deployments, the Adaptive Connectivity service needs sizing. For medium deployments, 500 MB of memory is recommended. For large deployments, 1 GB of memory is recommended.

It is recommended that an APS be created for the purpose of running the Adaptive Connectivity services, especially

for large deployments. This allows the data federation services to be monitored and configured easily.

- In your sizing document, add in the SAPS and memory for the Adaptive Connectivity service if needed.

15.9 Step 6: Sizing Analysis and Scale-out

Scaling-out your deployment is often necessary, where the size of the deployment is larger than a single machine that is expected to run the BI system. Additionally, some analysis and thought is needed to selectively place services on machines in a cluster so that optimal performance and reliability is achieved. Once you have the basic sizing as calculated in the above sections, you should have a list of all the services with SAPS and memory requirements for each. The next step is to ensure the hardware machine used in the deployment can handle the load.

Before assigning services to machines in the system, you need to determine what the constraints of the system are. There will very likely be processing power and memory limits on each machine. Some machines may be different.

For each of the tiers: Intelligence DB tier, Intelligence tier and Application tier and **for each** of the BI tools to be used in your deployment, do the following for each of the workloads:



1. Start by assigning the workload to a new machine in the system.
2. Add in memory to account for the operating system and basic BI infrastructure components such as the SIA. 2GB of memory and 2000 SAPS are recommended for these functions. (More than 2 GB of memory may be needed depending on the hardware platform and operating system used. Consult your machine vendor for more information.)
3. Check to see if either of the processing power of the machine or the memory limit has not exceeded for the workload that has been assigned to it. If it fits, you can move on to the next workload. If it doesn't fit, you need to add additional machines until the full processing power and/or memory for the workload have been accommodated.

For workloads made up of APS-based services, an APS specific to that workload should be created on the machine for those services. A machine may have other services defined on it, possibly as a result of the standard installation. Any services not intended for use on a machine should be disabled. See the APS Configuration section below for more information.

In general, Services should be distributed so the expected load is distributed across the machines within the cluster, so one individual machine is not unnecessarily loaded more than another. As an example: let's say your users are using an equal processing demand for Web Intelligence and Crystal Reports and you have determined that you need 2 Processing Tier machines. One option would be to place Web Intelligence on one machine and Crystal Reports on the other as these are I/O intensive. Another option would be to place both Web Intelligence and Crystal Reports on both the machines, as the load balancing of the system would generally distribute the load. However, it is probably best not to place Web Intelligence and Crystal Reports on the same machine, and all other services on the other. There are no strict rules regarding the above. Similar recommendations hold good for the Application Tier hosting Web Application Server and Processing Tier's DSL Bridge Service

15.10 Step 7: APS Configuration

The Adaptive Processing Server (APS) is the application that hosts a lot of the BI services. To get better system Scalability, we split the APS into multiple instances, which enables better resource usage by isolating specific tasks into java container. Where, each java container would have its own java heap setting. Multiple APSes may be defined on multiple nodes within a deployment. In almost all cases more than one APS will be needed in the system, both for management and maintenance of the running services. Often, nodes in the system need to run only a small subset of the available services in the default APS created at install time, making at least some amount of APS configuration necessary on every machine. Configuring APSes is an essential part of sizing BI4.

In BI 4.2, one APS is defined at installation. The rest of this section shows the basics of APS configuration.

15.10.1 APS Service Groupings

When establishing your APSes, the following APS definitions are recommended:

Core APS
Insight to Action Service
Publishing Post Processing Service
Publishing Service
Security Token Service
Translation Service
BI Commentary Service
BI Mobile Service (OCA)

Monitoring APS
Monitoring Service
Client Auditing Proxy Service

Search APS
Platform Search Service

Auditing APS
Client Auditing Proxy Service

Promotion Management APS
Visual Difference Service
Promotion Management ClearCase Service
Promotion Management Service
Version Management Service

Analysis APS
Multi-Dimensional Analysis Service
BEx Web Application Service

Web Intelligence APS
Web Intelligence Monitoring Service
Document Recovery Service
Rebean Service
Custom Data Access Service
Excel Data Access Service

Visualization APS
Visualization Service
Web Intelligence Monitoring Service

WebI DSL Bridge APS
DSL Bridge Service
Web Intelligence Monitoring Service

DF APS
Data Federation Service

Connectivity APS
Adaptive Connectivity Service

Materialization APS
Set Materialization Service

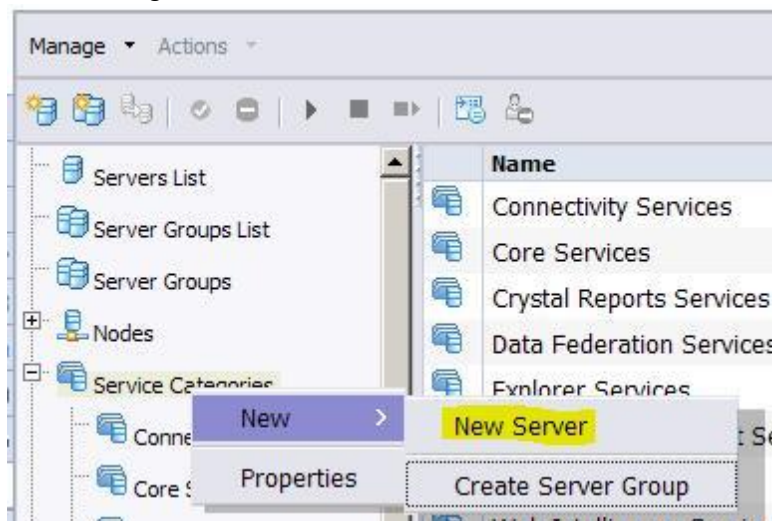
Note: The color-coded groupings shown above can be used when deploying APSes for redundancy and fail-over with minimal machine dependencies. The red group (right-hand column) services are related while the blue group (left-hand column) are mostly unrelated but have no dependencies on the red group. In cases where you want to duplicate services and also reduce or eliminate machine dependencies, this partitioning can be used.

15.10.2 Basic APS Configuration Instructions

It is strongly recommended that you not alter the original APS defined by the installer. If that APS is not needed, it should be left intact, stopped and disabled.

How to Create a New APS

- In the CMC, go to the Services listing.
- On the Manage menu, choose New and then New Server.



- On the Create New Server dialog box, choose the most appropriate Service Category.
- Select the first of potentially many services from the Select Service list, then press Next.
- To add more services to the APS, select them in the Available Additional Services list press > to add them to the Selected Services list. Choose Next when done.
- Give the APS a unique name.
- Choose Create.

At this point the APS will exist but will not be enabled or started. You can enable and start your new APS by right-clicking on the new service in the CMC and selecting the appropriate options.

How to Permanently Disable an APS

- In the CMC, go to the Services listing
- Right-click on the server (APS) that you wish to modify
- Choose Disable Server
- Right-click on the server (APS) that you wish to modify
- Choose Properties
- Find the option: 'Automatically start this server when the Server Intelligence Agent starts'.
- Uncheck the option.

How to Change the Memory Setting of an APS

- In the CMC, go to the Services listing

- Right-click on the server (APS) that you wish to modify.
- Choose Properties.
- Locate the Command Line Parameters.
- Scroll to the right until you locate **-Xmx1G** (or similar - it may have a different number).
- Change the number after the -Xmx. For example, to set the APS to use 6GB, change the command line to read -Xmx6G.

How to Change the services running in an APS

- In the CMC, go to the Services listing.
- Right-click on the server (APS) that you wish to modify.
- Choose Select Services.

See the [BI4 Admin Guide](#) for in-depth information on configuring APSes, on the [SAP Help Portal](#).

15.10.3 Adaptive Job Server

Note that the Adaptive Job Server does not require “splitting” or configuration as described above for the APS. This is because the Adaptive Job Server runs services as separate dedicated processes.

15.11 Step 8: Deployment and Monitoring

An important part of BI4 deployment is monitoring the operation of the various services to make sure they are running at a reasonable level of utilization. For BI systems, an average utilization goal is 60% of a node’s resources. This goal is prescribed because of the bursty nature of BI, which doesn’t run in a steady state but is very influenced by user workload.

During initial installation and testing, monitoring of the system is very important since the best sizing exercise can’t fully know how a system will behave and perform. Simulating user load is highly recommended, using products such as Apache JMeter or HP LoadRunner. Start with a small number of simulated users and bring the workload up slowly, watching how the system reacts. Proceeding gradually will uncover any trouble spots that can be dealt with before progressing to a full workload.

It is critical to monitor all aspects of resource usage in the system:

- Disk
- Network
- CPU
- Memory

This needs to be done across the entire BI4 landscape, including the CMS database and reporting databases. The monitoring services in BI4 provide a subset of the needed tools to do this.

15.11.1 Scaling-out BI4 Servers

A very important aspect of sizing that is inferred by the scenarios above is the creation and deployment of dedicated BI4 servers. The goal is to have dedicated machines running just the services that are necessary for its role. This is a common requirement of a scale-out deployment. See the [BI4 Admin Guide](#) for guidance on creating clustered deployments.

The Sizing Guide suggests that BI4 be installed as a full installation on each server and then any unnecessary services be turned off. This allows for a simpler installation procedure and also allows for future changes to the deployment. For example, if the role of a machine changes, all the BI4 components are there to be enabled and disabled as necessary.

15.11.2 Building the System

It is very important to methodically “build out” the system.

- a. Start with a smaller landscape, using a smaller number of users to gain confidence in your configuration. Gradually increase the load in increments of 50 to 200 users, only adding services/servers as necessary.
- b. Carefully monitor and analyze the performance and resource usage across the entire landscape, including the CMS repository database, the Web application tier, and any other SAP BI 4 platform servers involved in the test to identify bottlenecks in either the underlying infrastructure or server
- c. Take the appropriate actions to resolve problem areas as you find them. For example, adding another CMS host if CMS CPU utilization is above 80%. Testing gradually and solving issues as they are found is important, since progressing immediately to a full workload can make identifying the root cause of issues challenging.

The CMS database is key to the overall performance and scalability of the BI platform. A dedicated CMS database running on dedicated hardware is always recommended. Work with your DBA to ensure the CMS database server is correctly sized, configured, and monitored by referring to database vendor materials on sizing.

The underlying infrastructure, including machines and network, is critical to the overall performance and scalability of the BI platform; work with your infrastructure administrators to ensure the environment is correctly sized, configured, and monitored.

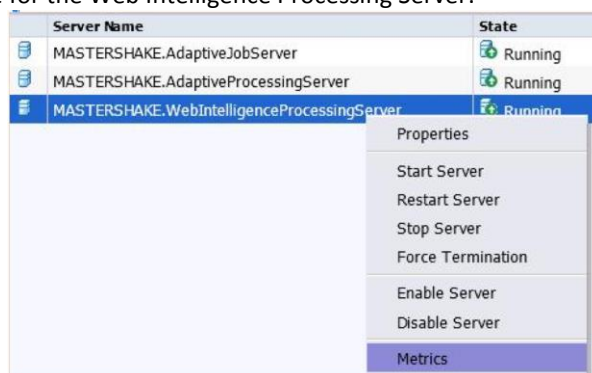
When starting your SAP BI4 landscape, it’s recommended to methodically start each of your Server Intelligence Agent (SIA) nodes and ensure all services are correctly started before starting another node.

15.11.3 Monitoring BI4

The BI4 suite has a number of built-in monitoring capabilities that can be used to measure the performance of the system as you build it up and continue into production.

15.11.3.1 Check the Processing Server Metrics

It is recommended that you frequently check the Processing Server Metrics for the BI processing services that your BI deployment uses, especially in the setup phase of your project. Metrics are maintained for each processing service such as Average Processing Time, Maximum Processing Time, and Minimum Processing Time. Here is an example for the Web Intelligence Processing Server:



Some of the information shown includes:

Metrics: MOD36DC1511.WebIntelligenceProcessingServer	
Hide Navigation	
Properties	Auditing Metrics
Translations	Current Number of Auditing Events in the Queue 0
User Security	Web Intelligence Processing Service Metrics
Metrics	Cache size (Kb) 0
Placeholders	Number of out-of-date documents in cache 0
Existing Server Group	Cache high mark count 0
	CPU usage (%) 0
	Total CPU time (seconds) 98
	Memory high threshold count 0
	Memory max threshold count 0
	Virtual memory size (Mb) 0
	Current number of client calls 0
	Number of remote extension errors 0
	Current number of tasks 0
	Total number of client calls 0
	Total number of tasks 0
	Idle time (seconds) 233458
	Current number of active sessions 0
	Number of documents opened from cache 0
	Number of documents 0
	Current number of sessions 0
	Number of document swap 0
	Number of swapped documents 0
	Number of sessions timeout 0
	Total number of sessions 0
	Number of users 0
	Number of active threads 0
	Total number of threads 5

15.11.3.2 Configure Report Processing Time Alerts

The Monitoring features of BI4 allow you to receive alerts when specific reports take too long to process. You can configure probes that run a Web Intelligence or Crystal Reports document periodically. You can be alerted if any of the documents take too long to run.



Step 1: In the Monitoring area of the CMC, open the properties for the Crystal Reports Service (Processing Server) to configure the probe for Crystal Reports. Open Interactive Analysis to configure the probe for Web Intelligence. You should already have documents created that you would like to monitor the performance of. See the [BI4 Admin Guide](#) for more information on creating and configuring probes.

While the system is in testing and validation phases, report processing probes should be run frequently (every few minutes) in order to effectively monitor the performance of the system.

Once a system is placed into production, it is suggested that you run report probes hourly so you can be alerted if system performance degrades.

You can change the schedule for a report probe as shown here:

Schedule: Crystal Reports Service (Processing Server)	
Hide Navigation	
Properties	Recurrence
▶ Default Settings	Run object: Hourly
▼ Schedule	Object will run every N hours and X minutes.
Instance Title	Hour(N) = 1 Minute(X) = 0
Recurrence	Start Date/Time: 03 19 PM 5/10/2013
Schedule For	End Date/Time: 03 19 PM 5/10/2023
History	Number of retries allowed: 0
Limits	Retry interval in seconds: 1800

reserved.

Step 2: In the Watchlist, create a new Watch that has rules based on the execution time of the probe. Set the Caution to 5000 (milliseconds) to have a caution alert occur when the report takes longer than 5 seconds. This is the most time a user expects a report to require to process. Set the Danger Rule to be some larger amount of time, such as 10 seconds.

See the [BI4 Admin Guide](#) for more information, including information on creating CMS Database alerts. The performance of the CMS is vital to the overall operation of the BusinessObjects environment and setting probes to monitor its performance is recommended.

15.11.4 Using Audit Reports

Usage information for your BI4 deployment can be obtained by using audit reports. This can be very helping in determining how many users are using certain BI tools, certain documents, etc. You can download a set of reports that are helpful for reporting on the BusinessObjects audit database [here](#).

15.11.5 Monitoring CPU and Memory Usage

Monitoring and recording of the basic CPU and memory usage of the machines in your deployment is also recommended. The goal is to have an average CPU usage of 60% in order to handle peaks in the range of 80%. Analyzing the historical usage of your deployment can help you determine if more resources are needed for a particular node. Similarly, if memory is frequently being fully used, you may be experiencing reduced performance and need to add more.

15.11.6 BI Tool Simulation Workflows

Creating simulated user workflows are very important in the testing of a new deployment. The following workflows are starting points for your user simulation. Be sure your simulations make use of the features that most users will use.

Crystal Reports	Dashboards	Web Intelligence	Analysis OLAP
1. Logon to OpenDocument web application.	1. Open BI launch pad page	1. Open BI launch pad page and login.	1. Open BI launch pad page.
2. View the report in HTML viewer.	2. Logon to BI launch pad	2. Navigate to the folder that contains the Web Intelligence report.	2. Logon to BI launch pad.
3. Enter the Dynamic Cascading Parameters for live data and view first page.	3. Navigate to Document List Tab	3. Open the report.	3. Navigate to Document List.
4. Go to the 2 nd page.	4. Expand Public Folder.	4. Open Advanced Prompt and Run Query for retrieve live data.	4. Open Saved Workspace.
5. Go to the page in the middle of the report.	5. Open the folder containing the Dashboard.	5. Navigate to Sub Report > Nation.	5. Swap Axes.
6. Close the document.	6. View the Dashboard.	6. Start Drilling Mode, Drill Down, and End Drill.	6. Print Analysis.
7. Log Off.	7. Refresh the dashboard.	7. Close Document.	7. Close Workspace.
	8. Close the Dashboard.	8. Log Off.	8. Log Off.
	9. Log Off.		

16. Scheduling and Publishing

Sizing for Scheduling and Publishing is very different than sizing for interactive users. You typically have a time budget to work with but are not as concerned with response time. However, if you want to have your jobs run in a time window, sizing is needed.

When considering a starting point for sizing, specify the user count in the Quick Sizer as Expert users. This most closely matches the demand made by the job schedulers, which don't have any human 'think time' at all. For the number of active concurrent users, specify the number of concurrent jobs that may be running.

When sizing for scheduling and publishing, you need to consider the time of day that those operations are done with respect to the workload of your interactive users. If they overlap, the amount of workload from scheduling and/or publishing needs to be added to the workload of your users. If the schedules do not overlap, for example if scheduling is done at night and your users only use the system during your day, then you size for the larger of the two scenarios.

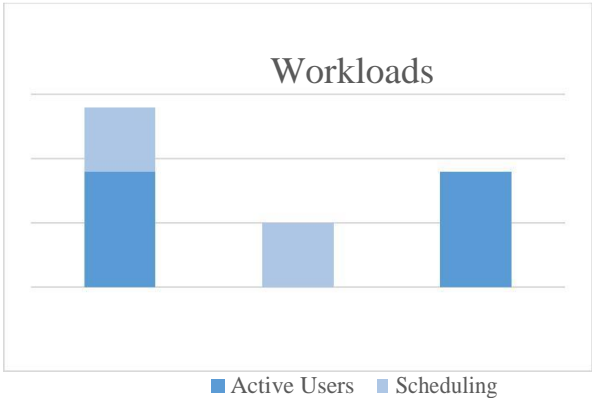


Figure 1: Workloads: Combined or separate

For Publishing, it is especially important to understand how many concurrent personalization jobs will be running as part of each publication. For example, if you need to serve three regions in a Publishing job and thus have three separate database queries, that would be equivalent to three active concurrent users. If you are publishing with personalizations that require a query-per-recipient, you need to determine the number of queries that might be able to be processed at once and use that number as the active concurrent user count.

If you are unsure about how much time is being spent running scheduled reports, you can deploy the auditing reports available from SAP. These reports help you to report on auditing information stored in the BusinessObjects audit database. You can download the reports and learn more [here](#).

16.1 Scheduling and Publishing Best Practices for Performance

16.1.1 Using Dedicated Machines

If both CPU and memory for the Adaptive Processing Server (APS) are heavily utilized during job processing, move that APS to its own machine.

Be sure BI4.2 SP02 or later is installed. This feature pack introduced better scaling and CPU utilization for scheduling and publishing.

For Publishing, isolate the Publishing Service and the Publication Post Processing Service on dedicated APS instances on dedicated machines. This will allow post processing to occur while a subsequent publishing job is processed.

16.1.2 Scaling

Horizontally scaling out the Publishing Service and Publication Post Processing Services across multiple APSes will enable more publication instances to proceed concurrently.

For publications with many recipients, vertically scale the APS on machines that have more CPU cores and memory. This will enable the Publishing Service to concurrently process more recipients and the APS to generate more jobs.

16.1.3 I/O

Scheduling and especially Publishing are very I/O intensive. A machine with fast I/O for disk and/or networking to a SAN is important for the location of the FRS folders.

For more guidance on ensuring scheduling and publishing jobs are configured to run as efficiently as possible, see the [BI4 Admin Guide](#).

17. Service Tuning

Configuration optimizations for the services listed below may improve their performance.

17.1 Web Intelligence

Instances

The Web Intelligence Processing Server is a multi-threaded process and scales both horizontally, by running the process on multiple machines, and vertically within a single machine. When scaling vertically multiple instances of the process may be required.

Although the Operating System will allocate a single Web Intelligence Processing Server process to execute on more than one CPU core at any one time, there are very small parts of the process that require exclusive lock over all threads within that process. The exclusive lock forms a tiny fraction of the overall process. This characteristic, amongst others, means that more than one instance of the process may be required when scaling vertically. Typically, this requirement is only when the load on the process is particularly significant.

Data volume and document complexity dramatically impact the performance of the Web Intelligence Processing Server. With simple documents, smaller data volume and less complex document structures, the greater the number of concurrent users a single process can support. Once either the documents become large, in terms of data volume and complexity, or the number of concurrent user requests is great, then multiple instances of the process may be required to gain the highest scalability. If additional instances are not added, then the scalability will not be as optimal.

We need a minimum of 2 WRS instances for failover purpose whatever the deployment size is. Only if either the data volume and document complexity is particularly large or the concurrent user requests are great should more instances be added. Too many Web Intelligence Processing Server instances may cause unnecessary overhead. The very maximum number of instances should not exceed the number of CPU cores as this would certainly contribute towards an unnecessary overhead. It's important the settings are the same and consistent so as best practice, deploy and tune the first Web Intelligence Processing Server on the first machine, and then duplicate it or use the server 'templates'.

Maximum Connections

By default, a single Web Intelligence Processing server supports up to 200 concurrent document sessions. This number can be increased or decreased by modifying the "Maximum Connections" parameter, in the server properties, as follows:

Web Intelligence Core Service

<input type="checkbox"/> Use Configuration Template	
Timeout Before Recycling (seconds):	1200
Idle Document Timeout (seconds):	300
Server polling interval (seconds):	120
Maximum Documents per User:	5
Maximum Documents Before Recycling:	50
<input checked="" type="checkbox"/> Allow Document Map Maximum Size Errors	
Idle Connection Timeout (minutes):	20
Maximum Connections:	200
<input type="checkbox"/> Enable Memory Analysis	
Memory Lower Threshold (MB):	3500
Memory Upper Threshold (MB):	4500
Memory Maximum Threshold (MB):	6000
<input type="checkbox"/> Enable APS Service Monitoring	

The 'maximum connections' isn't the number of users or users' sessions, but instead the number of document sessions. Determining the maximum connections setting is dependent on many factors, however in all cases the setting should only be used to limit the process and protect it from consuming too much memory or CPU cycles due to the size and complexity of the documents and the intensity of the requests.

In general, the 'maximum connections' should be set high, to prevent users from experiencing 'server busy' error messages. If your users experience 'server busy' errors, but your memory usage has not consumed the machine, then keep increasing the maximum connections until the errors are no longer experienced. However, there are times its necessary to limit the servers' availability and then the 'maximum connections' should be reduced, if there is a need to prevent any one Web Intelligence Processing Server becoming heavily loaded or consume huge amounts of memory. The total number of Web Intelligence document sessions that can be maintained across the entire cluster is the number of instances multiplied by the number of connections.

Memory

By default, the Web Intelligence Processing server is set to utilize as much memory as it requires and for the clear majority of cases this setting is suitable. However, if either the available memory on the machine is limited, or the documents are particularly large, in terms of data volume and complexity, when the memory usage can be considerable, it may be required to 'Enable Memory Analysis'. This feature will reclaim memory when memory thresholds are breached. For each memory threshold 'Lower', 'Upper' and 'Maximum', different mechanisms are enabled that recover memory and limit memory use. This means you should not set these values to be the same as each other, but rather keep the ratios between them similar to the defaults.

'Enable Memory Analysis' is now unchecked by default when you perform a new installation from BI 4.2 onwards. If you have updated from an earlier version this setting may be enabled. We now recommend this is disabled as

your default setting, enabling the process to use as much memory as it needs and without the need to reclaim memory or generate any related error due to memory limits becoming breached.

If you need to limit the memory allocated to the Web Intelligence Processing server, you can:

- Reactivate this parameter by checking the option in the properties of the Web Intelligence Processing Server and then restarting the server.

Web Intelligence Core Service

<input type="checkbox"/> Use Configuration Template	
Timeout Before Recycling (seconds):	<input type="text" value="1200"/>
Idle Document Timeout (seconds):	<input type="text" value="300"/>
Server polling interval (seconds):	<input type="text" value="120"/>
Maximum Documents per User:	<input type="text" value="5"/>
Maximum Documents Before Recycling:	<input type="text" value="50"/>
<input checked="" type="checkbox"/> Allow Document Map Maximum Size Errors	
Idle Connection Timeout (minutes):	<input type="text" value="20"/>
Maximum Connections:	<input type="text" value="200"/>
<input type="checkbox"/> Enable Memory Analysis	
Memory Lower Threshold (MB):	<input type="text" value="3500"/>
Memory Upper Threshold (MB):	<input type="text" value="4500"/>
Memory Maximum Threshold (MB):	<input type="text" value="6000"/>

Increase/decrease the maximum connection for each Web Intelligence Processing server and monitor your system to avoid overhead of your system's memory.

Web Intelligence Core Service

<input type="checkbox"/> Use Configuration Template	
Timeout Before Recycling (seconds):	<input type="text" value="1200"/>
Idle Document Timeout (seconds):	<input type="text" value="300"/>
Server polling interval (seconds):	<input type="text" value="120"/>
Maximum Documents per User:	<input type="text" value="5"/>
Maximum Documents Before Recycling:	<input type="text" value="50"/>
<input checked="" type="checkbox"/> Allow Document Map Maximum Size Errors	
Idle Connection Timeout (minutes):	<input type="text" value="20"/>
Maximum Connections:	<input type="text" value="200"/>
<input type="checkbox"/> Enable Memory Analysis	
Memory Lower Threshold (MB):	<input type="text" value="3500"/>
Memory Upper Threshold (MB):	<input type="text" value="4500"/>
Memory Maximum Threshold (MB):	<input type="text" value="6000"/>
<input type="checkbox"/> Enable ADF Service Monitoring	

Tuning recommendations

It is recommended to perform a load test to simulate end user activity utilizing Web Intelligence to determine the optimal configuration. The simulation should mimic the user's workflow in terms of events (opening, navigating, drilling, refreshing etc.), their concurrent requests with suitable think times and the documents they will typically use.

The initial load test should start with 1 or 2 instances of Web Intelligence, on the machines assigned to run Web Intelligence content, with all default settings for 'maximum connections' and memory settings. Start with a load of a single user and record the response time. Then progressively add more users ensuring the response time is not degraded. Should the response time degrade, then:

1. Fine tune dependent Adaptive Processing Servers, such as 'DSL Bridge Service' and 'Visualization Service' by reviewing garbage collection logs and making appropriate changes such as increasing the 'Xmx' Java memory setting or increasing the number of instances of those services.
2. Fine tune Web Intelligence by increasing the number of instances. Typically, only for very heavy loads may you find the need to also increase or decrease the 'max connections', enable memory analysis and alter the memory threshold limits.

Do this, until you identify the maximum scalability. You will then have identified the correct number of Web Intelligence instances for your users, the documents they are processing and for your hardware setup.

When sizing and tuning Web Intelligence, it's important to consider the data source(s) being used. The performance of the end-user operations that directly interact with the data sources such as refreshing data of the report contents, creation of documents, and Data filtering by Drill out of scope are impacted by the data source's server performance. We recommend that the databases are well sized and tuned by the administrator.

Finally, check and fine-tune the end-to-end landscape from the web server to the database servers. Refer to the [BI4 Admin](#) guide for details on Clustering.

For SAP BW, SAP Hana and SAP Sybase sizing you can go to the SAP Quick Sizer and also find the respective administration guide here:

- SAP BW administration guide
- SAP HANA administration guide
- SAP Sybase administration guide

For others, refer to the 3rd party database administrators.

To optimize your performance, review the following Web Intelligence best practices:

<http://scn.sap.com/docs/DOC-58571>

<http://scn.sap.com/docs/DOC-58532>

A lot of new features have been introduced in the SAP Business Objects 4.2, which offers a lot of new possibilities to analyze your data. You will find in the following sections tips and tricks on the best way to use those new functionalities.

17.1.1 Parallel refresh

Parallel refresh has been introduced in 4.2, which has a lot improved refresh time of Web Intelligence reports containing multi data providers.

Each connection, supported for the parallel refresh, has its own "maximum parallel queries" limit (refer the [BI4 Admin guide](#) for the default limits and supported connections).

Each Web Intelligence report has also its maximum number of data providers that can be refreshed in parallel: "Maximum Parallel Queries per document".

We recommend you to monitor CPU on both BusinessObjects Enterprise Servers side and Database side.

If you notice overload on your servers then decrease gradually the above parameters until your deployment becomes stable:

- If you notice overload on database side, you can decrease gradually the "maximum parallel queries" per connection from Information Design Tool and Universe Design Tool
- If your BOE server consume more than 60% of the CPU in average, then you can decrease gradually the « Maximum Parallel Queries per document » from the CMC.

17.1.2 SAP HANA data sources

HANA database SP10 includes optimization on Web Intelligence queries. So, we recommend you to get the SAP HANA database SP10 version to get better Web Intelligence report refresh-response time.

17.1.3 Disk space of your system

In Previous versions, we were compressing temporary data file when performing the first action on a report, which was taking a lot of time.

Since 4.2, this compression of the temporary file is not done any more. The performance has been highly improved for the first action performed in your report but you need to allocate more space on your disk.

We recommend you multiply by 10, the space that you were using in the previous version, for the temporary files located on the following folders:

```
<INSTALL_DIR>\SAP BusinessObjects\SAP BusinessObjects Enterprise XI 4.0\Data\<Server  
Name>\<SIA_Name>.\WebIntelligenceProcessingServer_Name>\sessions
```

Each Web Intelligence Processing Server has its own “sessions” folder where the temporary files are stored.

17.1.4 BI Commentary feature

You can insert comments in your reports from BI 4.2 onwards.

We recommend you change the default connection to the comment database from odbc to jdbc to avoid open document or get prompt box actions performance degradation. (see SAP Note: [2346055](#)).

17.1.5 Decimal Type for big number

Your performance may be impacted if you use many measures of the type “Decimal” ([Web Intelligence End User Guide](#)). If you notice performance degradation, we recommend you to set the “Decimal” format option only for a measure where you need high precision.

17.1.6 Sizing for SAP BusinessObjects Web Intelligence Upgrade

If you have already been running SAP BusinessObjects Web Intelligence previous 4.1 version and you want to know sizing for upgrade, we recommend you to keep the same number of SAPS as in 4.1 and add an extra of 20% of memory on processing tier side for the use of 4.2 new features (such as comments, geomap, etc.)

This percentage is based on the assumptions provided in the Web Intelligence Appendix and so need to be adapted to your specific deployment and performance requirements.

17.1.7 DSL Bridge

It's highly recommended to have one or more dedicated APS instances running only the DSL Bridge Service. More DSL Bridge service instances can be created, if necessary, to support the number of active concurrent users required.

17.2 Central Management Server (CMS)

The Central Management Server (CMS) can scale horizontally and vertically. The CMS provides server, user, session management, and security (access rights and authentication) management.

A single CMS instance can service a large number active concurrent users, depending on the resources of the machine on which an instance runs. 500 active concurrent users per instance is a conservative number useful for the purposes of planning your deployment. The recommended method for adding CMS instances to your deployment is to monitor the CPU and memory consumption of each CMS instance with a target goal of around 60% utilization. Adjust the number of CMSes while under representative user load to achieve approximately 60% utilization. Too many CMSes can add unnecessary synchronization traffic between instances. Too few CMSes will hinder performance of the deployment.

Note that all CMS instances share the same repository database. If your repository DB vendor doesn't offer low-latency horizontal scalability, you'll need to ensure you follow your DB vendor's guidance on how to size and scale vertically, especially for deployments with thousands of active concurrent users.

The performance of the CMS database is very important to the overall performance of the system. It is strongly recommended that your CMS database be sized and tuned for high performance operation by a qualified DBA. See the section, CMS Database Tuning for more information.

17.3 Crystal Reports Cache Service

The Crystal Reports Cache Service scales both up and out and can support 400 active concurrent users per instance. You will very likely not need more than 1 instance on any given machine, since the Cache Service is never enough of a bottleneck to make this necessary.

Because of the volume and frequency of communication between the Processing Service and the Cache Service, it's recommended that you deploy in pairs. That is, on each machine where there is a Cache Service, also deploy a Processing service and vice versa. This can yield better performance than having the Cache Service and Processing Service on different machines.

17.4 File Repository Service (FRS)

While your deployment can have multiple Input and Output FRS instances, the first FRS pair (input + output) to register with the CMS become the only active FRS pair. Other FRS instances are considered passive backups. Although all FRS instances run simultaneously, only the active FRS pair handles requests. If an active FRS fails, a passive FRS is changed to active status. When the previously active FRS becomes operational again, it is registered as a passive FRS with the CMS.

Good disk and I/O performance is critical for the operation of the FRS.

17.5 Connection Server

For data connectivity scenarios that use the Connection Server (commonly Web Intelligence and Crystal Reports), the recommended database connectivity is ODBC or native middleware, where available.

For Java-based (JDBC) middleware, including JDBC drivers for HANA, additional performance configuration may be required for the Java Virtual Machine (JVM) that is created by the Connection Server that hosts the JDBC driver. For larger user loads, the default JVM memory settings may be inadequate. The memory allocated to the JVM created by the Connection Server can be increased by editing the **cs.cfg** file, located in `.../installdir/sap_bobj/enterprise_xi40/dataAccess/connectionServer` on Unix and Linux and here on Windows: `C:\Program Files (x86)\SAP BusinessObjects\SAP BusinessObjects Enterprise XI 4.0\dataAccess\connectionServer`.

Set the `-Xms` and `-Xmx` settings as shown below. Here they are being set to 1GB at startup and 2GB maximum memory usage.

```

<JavaVM>
  <!-- The default JVM configuration can be overridden here -->
  <!-- Use an absolute path for the JVM -->
  <!--
  <LibraryName JNIVersion="JNI_VERSION_1_4">ABSOLUTE_PATH/jvm.dll</LibraryName>
  -->
  <Options>
    <Option>-Xrs</Option>
    <Option>-Xms1024m</Option>
    <Option>-Xmx2048m</Option>
  </Options>
</JavaVM>

```

17.6 Web Application Server

As you scale out your application tier, you should add an additional application server instance for every 500- additional active concurrent users. For better performance, it's also recommended that for each instance you configure at least 8 GB for heap size and 900 maximum threads. This is a minimum. Refer the [Quick Sizer](#) values to see if more memory is necessary for your deployment. These recommendations are based on Tomcat; capacity for other Web Application Server vendors may vary.

When deploying your Web Application Server, it is more important to have a low latency connection to the other BI 4 services than it is between the Web Application Server and the Web client. As such, you should optimize the network connection between the servers where possible.

For deployments that involve sizable numbers of users, mobile users or remote users, it is recommended that the Web tier be split to deliver static content by a separate Web application server. Static content (HTML pages, images, documents, JavaScript and Cascading Style Sheets) – content that doesn't change – can be delivered by a Web server dedicated to that task. The Apache Web server is typically used for this.

You can learn more about splitting the Web tier by consulting the *Web Application Deployment Guide* found in the 'Installation and Upgrade' section on the [SAP Help Portal](#). To learn more about deploying multiple instances of Tomcat, see the [SAP Business Intelligence Platform Pattern Books](#).

17.7 Scale-out Memory Expectations

As you scale out your deployment, you will need to know how much memory to allocate to each node in the cluster of machines. The amount of memory required will depend on which services each node is running.

The following table lists the expected memory requirements for these services. You should use these values when calculating the memory requirements of a cluster node.

Service	Expected Memory Requirements
Crystal Reports Cache Service	250 MB
Dashboard Cache Service	250 MB
CMS	1 GB
SIA	500 MB
Visualization Service (CVOM)	2 GB

See the Business Intelligence Platform Administrator Guide (BI4 Admin Guide) for a full list of services and how to configure each service.

17.8 Crystal Reports

Controlling Job Creation

To ensure the number of job processing child processes created by the Crystal Reports Processing Service doesn't exceed your needs, it is recommended that you configure the Maximum Concurrent Jobs entry to match the number of users you expect to support. Open the Properties of the CrystalReportsProcessingServer and locate the Crystal Reports Processing Service section as shown here:



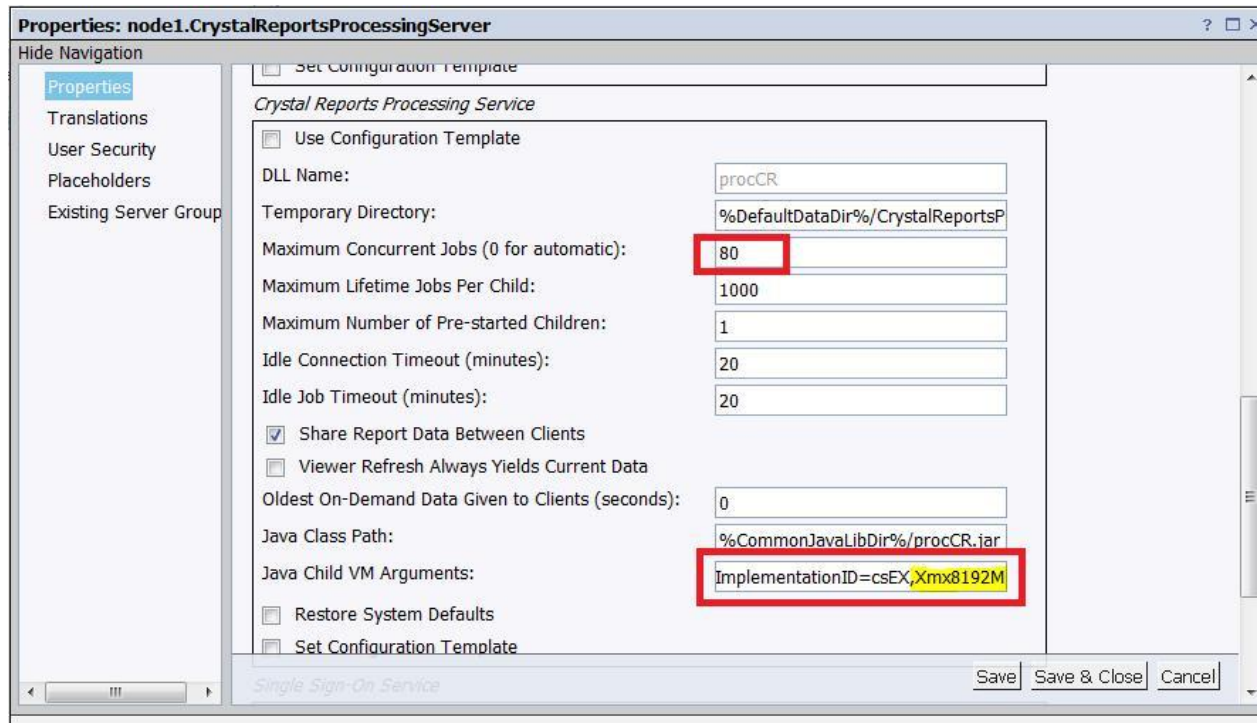
Crystal Reports Processing Service

<input type="checkbox"/> Use Configuration Template	
DLL Name:	procCR
Temporary Directory:	%DefaultDataDir%/CrystalReportsP
Maximum Concurrent Jobs (0 for automatic):	0
Maximum Lifetime Jobs Per Child:	1000
Maximum Number of Pre-started Children:	1
Idle Connection Timeout (minutes):	20
Idle Job Timeout (minutes):	20
<input checked="" type="checkbox"/> Share Report Data Between Clients	
<input type="checkbox"/> Viewer Refresh Always Yields Current Data	
Oldest On-Demand Data Given to Clients (seconds):	0
Java Class Path:	%CommonJavaLibDir%/procCR.jar
Java Child VM Arguments:	Dbusinessobjects.connectivity.direct
<input type="checkbox"/> Restore System Defaults	
<input type="checkbox"/> Set Configuration Template	

Increasing Memory Allocation for Improved Response Time

Depending on your reporting needs and the design of your reports, it may be possible to improve the processing time of reports by allocating more memory for Crystal Reports to work with. This is done by increasing the amount of memory allocated to each Crystal Reports child process. You may want to do this in the case where you have extra processing and memory headroom on a machine beyond your essential sizing requirements. Allowing Crystal Reports to more quickly process reports has obvious responsiveness benefits and allows the system to be better able to handle spikes in demand.

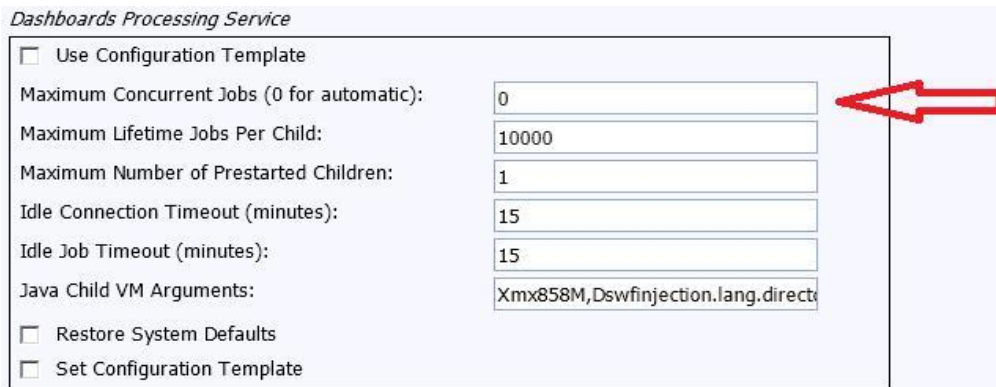
Each Crystal Reports processing child process has a default memory allocation of approximately 2GB. You can increase this by setting the Java memory parameter in the Java Child VM Argument list, as shown below in the settings for the CrystalReportsProcessingServer in the CMC. In this example, the memory is changed to 8192MB using the parameter setting Xmx8192M.



You may want to increase the number of concurrent jobs if your goal is to increase the number of reports that each child process can deliver. As shown here, the number of concurrent jobs has been set to 80.

17.9 Dashboards

To ensure the number of job processing child processes created by the Dashboards Processing Service doesn't exceed your needs, it is recommended that you configure the Maximum Concurrent Jobs entry to match the number of users that you expect to support.



17.10 Analysis OLAP

Analysis OLAP (Multi-Dimensional Analysis Service or MDAS) works the best when run in its own APS. See the section Step 7: APS Configuration, “Step 7: APS Configuration”.

Memory

Once in its own APS, it is recommended that MDAS be configured with a minimum of 4GB of memory. The section Step 7: APS Configuration outlines the steps for configuring the memory settings of an APS.

Sessions

The MDAS “Maximum Client Session” setting is often set to a low number by default. It is recommended that it be set in a range of 45 to 100 depending on the expected load.

Each MDAS can support up to 100 client sessions. To support more sessions, create additional APS instances that host an MDAS service. It is recommended that additional MDAS APSes be deployed on separate machines.

Multi Dimensional Analysis Service

Use Configuration Template

MDAS Configuration

Maximum Client Sessions:	100
Maximum number of cells returned by a query:	100000
Maximum number of members returned when filtering:	100000

Restore System Defaults

Set Configuration Template

18. SAP BW Considerations and Recommendations

When using BI4 to report on SAP BW data, the following recommendations are made to ensure the best performance possible:

18.1 BW Configuration

- Apply the latest fixes available from SAP by applying the latest support packs to your BW system
- It is strongly suggested that your BW be sized for the anticipated query load expected from the BI4 system.
- Have a BW expert tune the system generally and BEx queries specifically for your BI4 usage

18.2 BEx Query Performance

Check the performance of BEx queries. See the [BEx Query Runtime Statistics](#) found on the SAP Help Portal.

You can determine how much time the execution of certain user actions require in the front-end and in the SAP BW analytic engine.

18.3 Navigational Attributes

A large number of navigational attributes defined in the underlying InfoProvider may impact the overall performance. This can lead the BI4 tools to generate a lot of "Crossjoin" operators, causing extra unused data to be returned in queries. Revisit your BEx queries to be sure they include only the information necessary for BI.

Characteristic attributes can be converted into navigational attributes. They can be selected in the query in exactly the same way as the characteristics for an InfoCube. In this case, a new edge/dimension is added to the InfoCube. During the data selection for the query, the data manager connects the InfoProvider and the master data table ('join') in order to fill the Query.

From a pure performance point of view, you should model an object on a characteristic rather than on a navigational attribute.

18.4 BW Specific Tuning and Configuration

18.4.1 Set Specific Properties

- Check the option: Use Selection of Structure Elements
- Choose the Read Mode option: Query to Read When You Navigate or Expand

Checking the *Read Mode* property helps improve performance in situations where there are many Restricted Key Figures and other calculations in the Query Definition.

If analysis of the query performance indicates very high EVENTID 3200 times and/or the FEMS number is very high, try enabling this property and check performance.

Checking the *Use Selection of Structure Elements* option ensures the structure elements are sent to the database for processing.

You can see these settings shown below:

Query Monitor

The screenshot displays the SAP Query Monitor interface. At the top, there is a toolbar with icons for 'Execute', 'Execute + Debug', 'Generate Report', 'Properties', 'Messages', 'Help Texts', and 'Generation Log'. Below the toolbar, there are tabs for 'Performance Info', 'Technical Information', 'Cache Monitor', 'Query Variants', 'IGS Test', 'Get Variant', and 'Create Variant'. At the bottom of the toolbar, there are buttons for 'ABAP Web' and 'Java Web'. The main area shows the 'Query' field with the value '00_SD_C03/ZT_RKT_BW_70_DEMO_SALES'. Below the query field, there is a 'Query Display' dropdown set to 'List' and radio buttons for 'Parameter 1' and 'Parameter 2'. A dialog box titled 'Query Properties -> Move into Query Designer' is open, showing various settings:

- Read Mode: H Query to Read When You Navigate or Expand (checked) InfoProvider Setting
- Req. Status: 0 All PartProviders Up to Released Status (RQT) (checked) InfoProvider Setting
- Cache Mode: 1 Main Memory Cache Without Swapping (checked) InfoProvider Setting
- Update Cache Objects in Delta Process (checked) InfoProvider Setting
- Use Selection of Structure Elements
- Calculate w/ Packed Numbers
- no parallel processing
- Generation log
- Optimization Mode: 0 Query Will Be Optimized after Generation
- All Basic Key Figures Input Ready
- Disaggregation on Totals: Default
- Statistic Det.: 1 Only Front End/Calculation Layer Data

At the bottom of the dialog box, there are 'OK' and 'Cancel' buttons.



18.5 Additional Content

More BW performance tuning information can be found here:

- [Performance Optimizing SAP BusinessObjects Reports Based Upon SAP BW Using BICS Connectivity](#)
- [Business Intelligence Performance Tuning](#)
- [Performance & Scalability Best Practices and Tuning Recommendations for SAPBO XI 3.1 on BW](#)

19. CMS Database Tuning

The CMS database performance has a significant impact on the performance of the system. The following Sybase database tuning tips are offered. These tips may apply to other database servers as well.

1. **Data Cache [buffer pool]**

Data cache allows efficiently reducing IO. Having a sufficient size of data cache can speed up queries.

2. **Procedure Cache**

Configure the Procedure Cache so frequently used database procedures are cached. This allows the database to not need to reload procedures from disk as often.

3. **Lock Granularity**

The CMS does a lot of reading and writing from/to the database. By default, the granularity is set to page level, which is too high. Page level locking promotes contention by the CMS when, for example, several records are stored together in a same page. One of them may be being updated but another one is being read. One of these two sessions has to wait until other has finished.

Setting Lock Granularity to use row level locking can avoid contention and deadlock issues.

4. **Parallel Processing:**

This is important to set so Sybase can fully utilize all the CPU cores available on the machine.

20. Virtualization: Service Level and Performance

BI4 works very well in virtualized environments. Testing has shown it performs equally well on virtualized hardware as it does on dedicated hardware. However, in order to achieve equal performance, the hardware it runs on must be dedicated to the virtual machines that the BI services run in.



SAP strongly recommends that BI virtual machines have reservations for the memory and CPUs assigned to them.

BI is a bursty application workload. It can be very I/O intensive when heavily loaded by users. When experiencing maximum load as defined by your sizing, all parts of the system can go from moderately loaded to very busy momentarily. If the BI hardware is being shared with other workloads in a virtualized environment, the response to bursts in activity can result in slow-downs between dependent BI services and ultimately end user response time.

It is recommended that you build your system to **target 60% average utilization** so that bursts in activity can be handled. Performance of both physical and virtual systems degrades when system utilization is greater than 80%.

If you are unable to secure memory and CPU reservations for your virtualized deployment, it is recommended that additional scale-out be done to provide adequate BI processing resources for peak periods. The amount of additional scaling out will depend on your IT infrastructure and performance policies.

The goal of Sizing BI4 is to have a deployment that can handle the workload needed and provide good response time for end users. If you deploy to a shared virtualized environment, **extra care is needed to ensure your sizing exercise results in a system which has a Service Level and responsiveness that can be maintained.**

For more in-depth information on BI and Virtualization, see the *SAP BI 4 Virtualization* site at <http://www.sap.com/bivirtualization>.

21. Desktop Virtualization

When running BI4 desktop applications are deployed to virtualized desktop environments such as Citrix and Windows Server Remote Desktop Services, sizing is required to determine how many concurrent users can be supported for commonly used BI applications.

The following table shows the memory footprint of applications that you may run in a virtualized desktop environment. You can use the memory footprint information to calculate the amount of memory required to service the number of users required.

Product	Memory Footprint per User
Web Intelligence Rich Client	1.5 GB
Crystal Reports for Enterprise	1.5 GB
Crystal Reports 2016	150 MB
Interactive Design Tool	500 MB
Universe Designer	500 MB

22. Sizing Example

The following is an example customer requirement and walkthrough of how to apply the sizing methodology specified above.

Note: this sample doesn't demonstrate every aspect of sizing. For your sizing, be sure to follow all the steps outlined in the section [15. Deployment and Sizing Methodology](#).

For this example, the following are the requirements set out by the customer. All user counts are considered active-concurrent users.

- Web Intelligence:
 - 150 Business users, 10 expert users
 - Data source: BW
- Crystal Reports 2016:
 - 80 Business users, 20 expert users
 - Data source: Relational database
- Analysis OLAP:
 - 15 expert users
 - Data source: Relational Database
- The customer has indicated they won't be using search heavily.
- The customer has indicated the system will be used internally only, so the Web server doesn't need extra network security, etc. required for exposure on the Internet.

The rest of this section demonstrates how to follow the methodology described above in the sizing of this customer scenario.

Step 1: Quick Sizer Setup

Enter the user numbers into the BI4 Quick-Sizing <http://service.sap.com/quicksizing> as shown below. Note that figures in the below screenshot are based on a sample calculation. Use the Quick Sizer tool to get actual resource requirements.

For more information, see [11. Sizing Tools](#) and [Quick Sizer Tool](#).

Input screen

Table 1: Throughput - SAP BusinessObjects Business Intelligence

Clear Insert The buttons in this line only work for marked lines in the lower table.

Element	Element short text	A/P	TI	Information consumers (300 sec. think.)	Business users (30 sec. think time)	Expert users (10 sec. think time)	% Small size reports	% Medium size reports	% Large size reports
ANA-OLAP	SAP BusinessObjects Analysis, edition for OLAP	A	S			15	25	50	25
CRYSTAL	SAP Crystal Reports 2016	A	S		80	20	25	50	25
CRYST-ENT	SAP Crystal Reports for Enterprise	A	S						
DASHBOARD	SAP BusinessObjects Dashboards	A	S						
WEB-INTELL	SAP BusinessObjects Web Intelligence	A	S						
WEBI-ON-BW	SAP BusinessObjects Web Intelligence on BW	A	S		150	10	25	50	25
DESIGN-STU	SAP BusinessObjects Design Studio	A	S						
LUMIRA	SAP Lumira	A	S						
LUMIRA-BW-LIVE	SAP Lumira for BW Live	A	S						
LUMIRA-HANA-LIVE	SAP Lumira for HANA Live	A	S						

Output screen

Result

Print page Result categories How to interpret the results

All SAPS Memory and disk Data IO Log IO

Throughput Results for Software Components

SW component	Software component	CPU cat.	SAPS (total, 2-tier)	BO APP SAPS	BO INT SAPS	BO INT DB SAPS	BO PRO SAPS	Memory Cat.	Memory (total, 2-tier, MB)	BO app. tier Memory	BO int tier Memory	BO int tier DB Memory	BO pro. tier Memory
BI_SERVER	SAP BusinessObjects Business Intelligence Server	XS	38.000	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

Step 2: Intelligence Tier DB

The Intelligence DB (the CMS database) must have processing headroom of 4100 SAPS and 6.1 GB of memory available for processing CMS queries. Add these values to the sizing worksheet.

Step 3: Intelligence Tier

The CMS itself requires 6700 SAPS and 13.3 GB of memory to function properly. Add these values to the sizing worksheet.

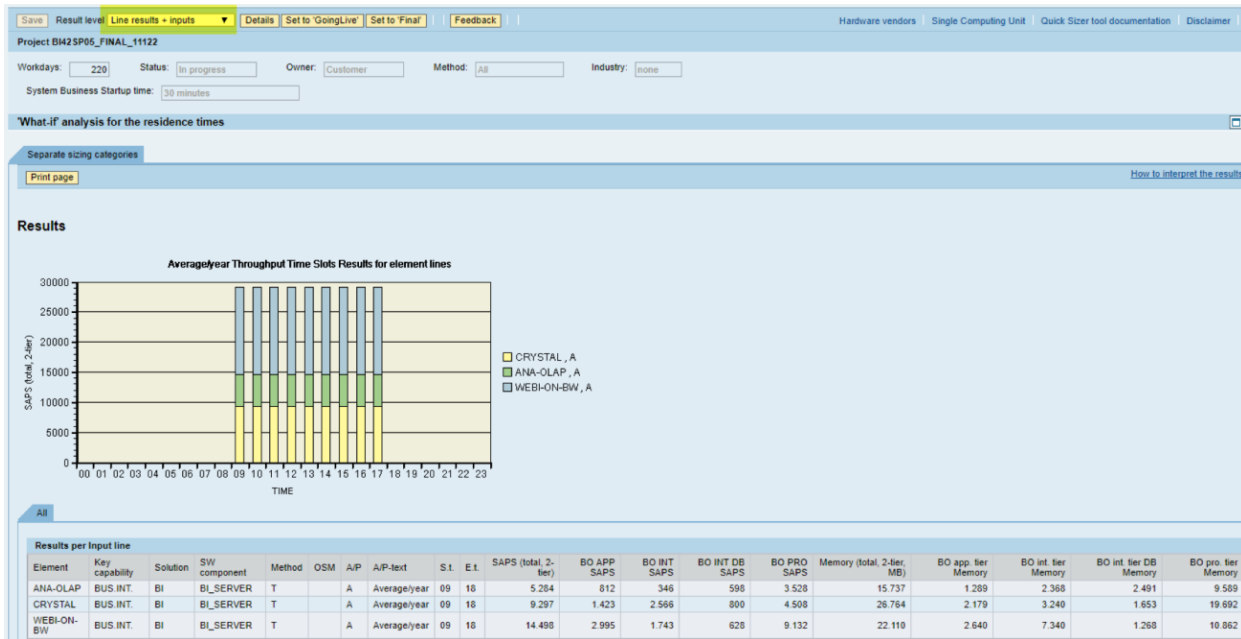
Step 4: Application Tier

The Web server is the main portion of the processing that occurs on the Application Tier. In this deployment, it needs 7300 SAPS and 6.1 GB of memory. Add these values to the sizing worksheet.

Step 5: Processing Tier

This is the tier that includes all the BI tools.

The internal splits of the SAPS/Memory w.r.t. individual BI tools, can be obtained by changing the results type to "Line results + inputs" (as seen in the screenshot below).



Analysis OLAP:

To determine how much processing and memory resources are needed to support the deployment’s Analysis OLAP needs, put just those numbers into the Quick Sizer and record the Processing Tier numbers. Since Analysis OLAP will be using straight forward data connectivity, no special sizing is necessary.

Crystal Reports 2016

To determine how much processing and memory resources are needed to support the deployment’s Crystal Reports 2016 needs, put just those numbers into the Quick Sizer and record the Processing Tier numbers. Since Crystal Reports will be using straight forward data connectivity, no special sizing is necessary. The memory required for Crystal Reports 2016 is calculated according to the processing needed.

Web Intelligence

To determine how much processing and memory resources are needed to support the deployment’s Web Intelligence needs, put just those numbers into the Quick Sizer, in the “SAP BusinessObjects Web Intelligence” section (WebI on BW), and record the Processing Tier numbers. Add these values to the sizing worksheet.

You can perform Web Intelligence sizing using the Quick Sizer tool using two options:

- for relational data: use the “SAP BusinessObjects Web Intelligence” section
- for BW data: use the “SAP BusinessObjects Web Intelligence on BW” section

The Web Intelligence performance and memory numbers can go into the sizing worksheet as shown. Add that to the worksheet.

Promotion Management Services

The Promotion Management services for this size of deployment, a medium size deployment, requires 750 MB of memory. 1000 SAPS are adequate for the processing needed.

Platform Search Services

Platform Search can be sized for such a deployment, considering they won't be using search heavily, at 500 MB of memory and 1000 SAPS of processing power. You can refer this [note](#) for more details.

Note: Considering the customer wishes to make light usage of Search, it would be wise to configure the Search Services to work in a limited fashion. See the [BI4 Admin Guide](#) for more details for configuring Search Services.

Add these values to the sizing worksheet.

At this point, your sizing worksheet might look like this:

Sizing					
	SAPS	Memory (GB)	Cores		
				Machine SAPS	25120
Intelligence DB	4100	6.1	2	Number of cores	12
Intelligence Tier (CMS)	6700	13.3	3	SAPS/Core	2093
Application Tier (Tomcat)	7300	6.1	3		
Analysis OLAP	3528	9.5	2		
Crystal Report 2016	4508	19.6	2		
"Web Intelligence on BW" Processing Tier	9132	10.8	4		
Promotion Management Services	1000	0.75	0.5		
Platform Search	1000	0.5	0.5		
OS	2000	2	1		
Total	39268	68.6	18		

Sizing Analysis and Scale-out

For this deployment, recall that the customer has 12 core machines with 2093 SAPS per core. In the spreadsheet shown above, the number of cores needed for each service has been calculated by dividing the SAPS for each service by the SAPS per core of the specified processor.

Now that we know the number of cores needed, we can focus on distributing the services across machines, taking into account the number of cores per machine. It is assumed that memory is relatively cheap compared to additional machines, so some machines will have requirements for more memory than others.

Remember that for each new machine, an operating system needs to be accounted for as well as the basic communication infrastructure of the BI system.

Distributing the Intelligence and Application Tiers

For this deployment, the relatively small processing load required by the CMS and Web server could allow it to be deployed to the same machine. However, the focus of performance for user interaction is primarily the Web application server and the CMS, so it is not suggested to run them on the same machine.

Machine #1 will house the CMS services and also Promotion Management. Promotion Management runs infrequently and makes use of the CMS so running them on the same machine for this scenario makes sense.

Machine #2 will host the Web application server. At this point the processing usage is low but this leaves room in the future for additional load by mobile services, etc. Since the processing load on this machine is relatively low, Platform Search can be allocated to it as well.

Distributing the BI Processing Services

<u>Machine #1: CMS SAPS Memory Cores</u>	<u>SAPS</u>	<u>Memory</u>	<u>Cores</u>
<u>OS and BI Core</u>	<u>2000</u>	<u>2</u>	<u>1</u>
<u>Intelligence DB</u>	<u>4100</u>	<u>6.1</u>	<u>2</u>
<u>Intelligence Tier (CMS)</u>	<u>6700</u>	<u>13.3</u>	<u>3</u>
<u>Promotion Management Services</u>	<u>1000</u>	<u>0.75</u>	<u>0.5</u>
-	<u>13800</u>	<u>22.15</u>	<u>6.5</u>
-	-	-	-
<u>Machine #2: Web App</u>	<u>SAPS</u>	<u>Memory</u>	<u>Cores</u>
<u>OS and BI Core</u>	<u>2000</u>	<u>2</u>	<u>1</u>
<u>Application Tier (Web Server)</u>	<u>7300</u>	<u>6.1</u>	<u>3</u>
<u>Platform Search</u>	<u>1000</u>	<u>0.5</u>	<u>0.5</u>
-	<u>10300</u>	<u>8.6</u>	<u>4.5</u>

For this deployment, the Analysis OLAP and Crystal Reports processing fit on one machine.

<u>Machine #3: AOLAP, CR</u>	<u>SAPS</u>	<u>Memory</u>	<u>Cores</u>
<u>OS and BI Core</u>	<u>2000</u>	<u>2</u>	<u>1</u>
<u>Analysis OLAP</u>	<u>3528</u>	<u>9.5</u>	<u>2</u>
<u>Crystal Reports 2016</u>	<u>4508</u>	<u>19.6</u>	<u>3</u>
-	<u>10036</u>	<u>31.1</u>	<u>6</u>

The Web Intelligence sizing for this customer requires 4 cores for the processing services. 5 cores does not exceed a single machine's 12, so two machines will not be specified for WebI processing.

<u>Machine #4, WebI</u>	<u>SAPS</u>	<u>Memory</u>	<u>Cores</u>
<u>OS and BI Core</u>	<u>2000</u>	<u>2</u>	<u>1</u>
<u>Web Intelligence</u>	<u>9132</u>	<u>10.8</u>	<u>4</u>
-	<u>11132</u>	<u>12.8</u>	<u>5</u>

Note: In general, spreading the processing and memory for the Web Intelligence services (processing server, DSL Bridge and Data Visualization) evenly across the two machines will provide processing headroom for the future, which is desirable.

APS Configuration

APSEs need to be created on four of the five machines. Following the 15.10 Step 7: **APS Configuration** guidance, the following APSEs should be defined and allocated to the machines as follows:

Machine #1: CMS

- Core APS
- Monitoring APS
- Auditing APS
- Promotion Management APS

Machine #3: AOLAP, CR

- Analysis APS

Machine #4: WebI

- WebI APS
- Visualization APS
- WebI DSL Bridge APS
- Connectivity APS

Deployment

When deploying BI4 to these machines, all services should be installed and then selectively disabled once the cluster is created. This will aid future expansion and patching, so that services can be turned back on where and when they are needed. For example, patching of machines can be done in parallel if a CMS is running on each node in the cluster. For day-to-day processing, having a CMS running on each node is not necessary. However, it is relatively easy to bring up a CMS instance on each node if it is already installed.

Redundancy and Virtualization

This sample deployment doesn't take into consideration redundancy or virtualization. Those are important topics that would be part of any Enterprise software deployment. For further details on deploying BI4, see the [BI4 Admin Guide](#).

23. Sizing References

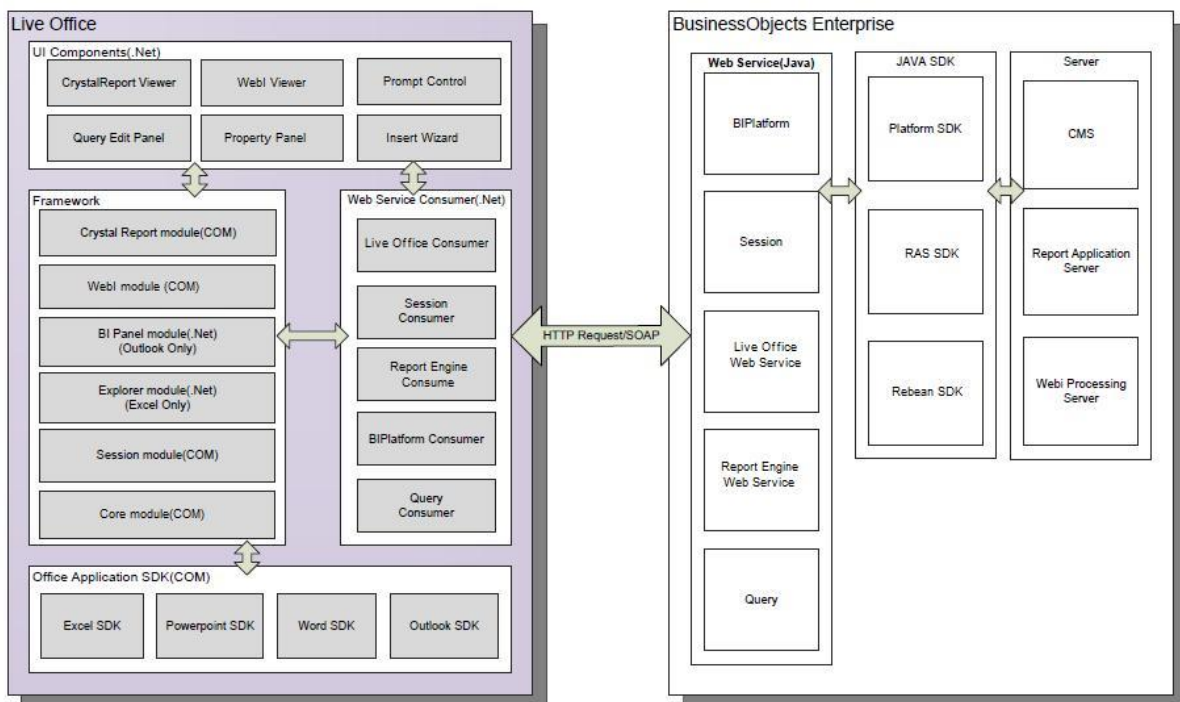
23.1 Sizing SAP BusinessObjects Live Office

SAP Live Office uses the Web Intelligence and Crystal Reports engines depending on which report parts are inserted in your documents. When performing Sizing of BI4, Live Office usage counts as usage of these BI services, including data sources, etc.

It is important to realize that if a document contains multiple BI report parts, refreshing that Office document is the equivalent of the number of users interacting with the BI system all at the same time.

For example, if an Office document has three report parts and it gets refreshed, that's the same as three active concurrent users. Be sure to take this into consideration when accounting for the number of active concurrent users of the BI system.

Here is a more detailed diagram of the Live Office architecture:



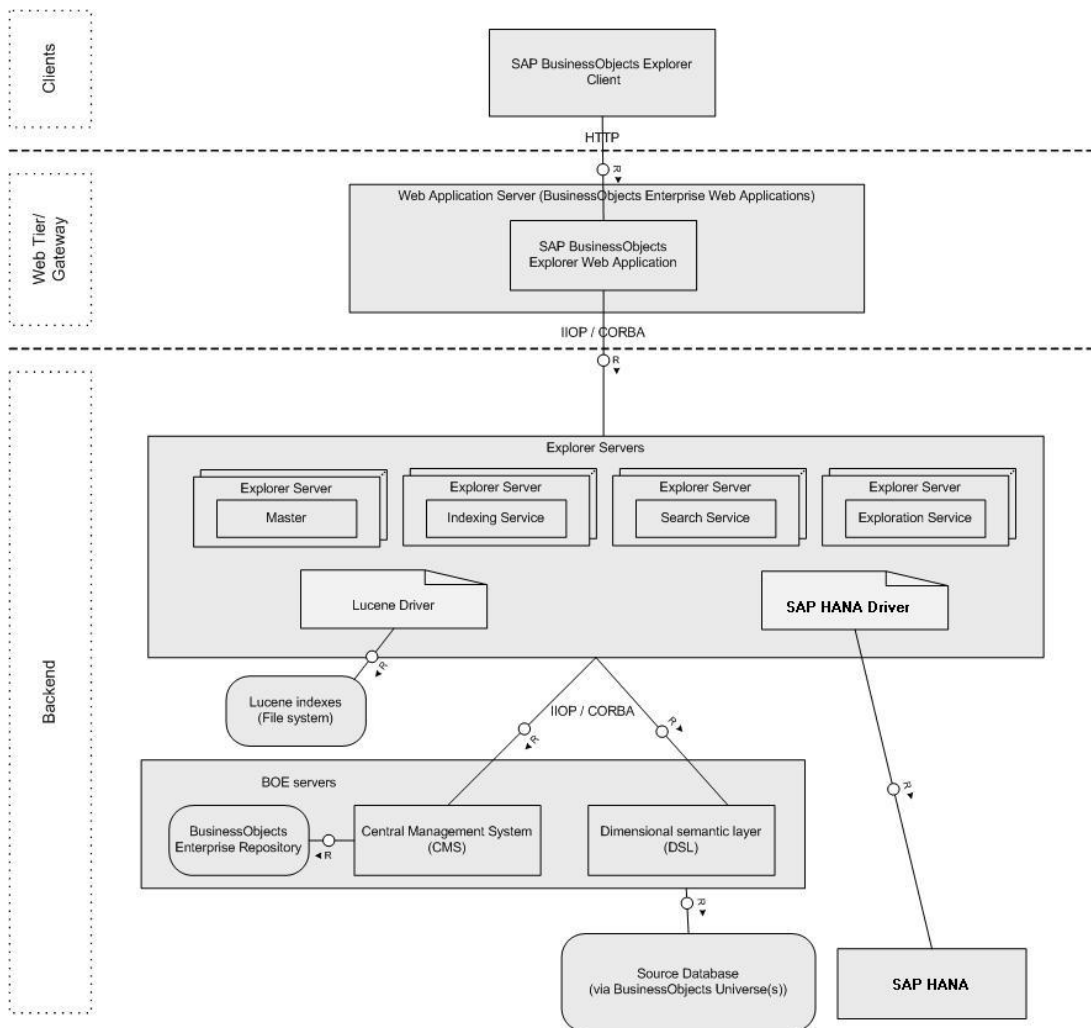
For more information on SAP Live Office, see the [Live Office](#) website.

23.2 Sizing SAP BusinessObjects Explorer

SAP BusinessObjects Explorer has two deployment types:

Explorer	Data Providers	Components
Explorer 4.2: Provides fast processing of small to large data sets	<ul style="list-style-type: none"> • BusinessObjects Universes (UNX) • Excel files 	<ul style="list-style-type: none"> • SAP BusinessObjects Explorer • SAP BusinessObjects Enterprise
Accelerated Version: Provides processing of very large data sets	<ul style="list-style-type: none"> • NetWeaver BWA indexes enabled for Explorer • SAP HANA in-memory computing engine • Excel files 	<ul style="list-style-type: none"> • SAP BusinessObjects Explorer • SAP BusinessObjects Enterprise • SAP NetWeaver Business Warehouse (BWA) • SAP HANA

The following is the Explorer architecture:



For information on Sizing Explorer, see the [SAP BusinessObjects Explorer Sizing Guide](#) in the Performance and Scalability area of SAP Service Market Place.

23.3 Sizing SAP BusinessObjects Mobile

SAP BusinessObjects Mobile solution allows end users to access SAP BI documents (Web intelligence, Crystal Reports, etc.) through their mobile devices. It is suited for ad hoc query, reporting, and analysis.

SAP BusinessObjects Mobile backend services run within the Web application server (typically Tomcat) in the BI4 environment. The Mobile server provides alternate rendering and workflows for BI in ways that are optimized for mobile devices.

When sizing for deployments that include Mobile BI, follow the sizing methodology outlined in this document. Be certain to consider the load all users will put on the system, including mobile users. The load created by a mobile user needs to be considered in the same way as a regular desktop user: Information Consumers, Business Users and Experts need to be identified and their load accounted for in the sizing exercise.

Depending on the portion of mobile users vs. desktop users, you may want to increase the memory and CPU resources allocated to the Web application server in order to ensure any increase in processing required by the SAP BusinessObjects Mobile server is accounted for.

[Sizing information for Mobi](#) can be obtained on the SAP Service Marketplace.

23.4 Sizing SAP BusinessObjects Design Studio

[Sizing information for Design Studio](#) can be obtained on the SAP Service Marketplace.

23.5 Sizing SAP Lumira 2.0

[Sizing information for SAP Lumira 2.0](#) can be obtained on the SAP Help Portal.

24. Appendix

24.1 BI4 P&R Benchmark Testing Details

The following information describes the P&R test environment used to generate the Sizing benchmark data.

24.1.1 Test report documents and data characteristics

The following table is a summary of the data volume returned for the various documents in the sizing tests. For more information about the complexity and design of the documents see the document-specific sections further down. There is a separate section later, which describes not only the data volume returned as the result of a query, but also the total size and structure of the data as it existed in the reporting database.

Product	Small	Medium	Large
Crystal Reports (for Enterprise and 2016)	1,000 rows	10,000 rows	100,000 rows
Dashboards	25 rows	250 rows	2,500 rows
Web Intelligence	2,000 rows	20,000 rows	200,000 rows
Analysis, edition for OLAP	35,000 cells	89,500 cells	375,000 cells

Crystal Reports for Enterprise and Crystal Reports 2016

Information about the composition and complexity of Crystal Reports documents used in the sizing tests is summarized below.



Fields in query result set: 23
 Tables joined: 4
 Record selection clauses: 4
 Parameters/Prompts: 4
 Interactive filter clauses: 0

	Simple Elements					Complex Elements			Total
	Total	Field	Text	Total	Picture	Total	Chart	Crosstab	
Report Header	4	3	1	0	0	0	0	0	4
Page Header	6	1	4	0	1	0	0	0	6
Group Header 1	7	1	3	3	0	2	2	0	9
Group Header 2	7	1	3	3	0	1	1	0	8
Group Header 3	3	1	1	1	0	0	0	0	3
Group Header 4	5	3	1	1	0	1	0	1	6
Body	3	2	1	0	0	0	0	0	3
Group Footer 4	6	2	3	1	0	0	0	0	6
Group Footer 3	2	0	1	1	0	0	0	0	2
Group Footer 2	2	0	1	1	0	0	0	0	2

Group Footer 1	3	3	0	0	0	0	0	0	3
Report Footer	15	0	8	7	0	0	0	0	15
Page Footer	3	1	2	0	0	0	0	0	3
Total	66	18	29	18	1	4	3	1	70

Web Intelligence

For “SAP BusinessObjects Web Intelligence” (aka WebI on Relational data), the Sizing estimation by Quick Sizer is based on the assumptions provided in this section. A different assumption results in a different guideline. The only difference between definition of users is the click speed:

User definition	Information user	Business user	Expert user
Think time between each action in seconds	300 sec	30 sec	10 sec

Document T-shirt definition for Relational Data is the following:

Reports have been created with four data providers to see the impact on the parallel refresh new feature as follows:

Web Intelligence Report Size	Small	Medium	Large
Composition of each data provider in the report	The report contains 4 data providers (DP): 2DP: 500 rows – 2000 Cells 1DP: 500 rows – 5000 Cells 1 DP: 500 rows – 5500 Cells In total : 4DP: 2 000 rows – 14500 cells	The report contains 4 data providers (DP): 2DP: 5 000 rows-20 000 cells 1DP: 5 000 rows-50 000 cells 1DP: 5 000 rows-55 000 cells In total: 4DP: 20 000 rows- 145 000 cells	The report contains 4 data providers (DP): 2DP: 50 000 rows-200 000 cells 1DP: 50 000 rows-500 000 cells 1DP: 50 000 rows-550 000 cells In total: 4DP: 200 000 rows-1 450 000 cells

The reports include charts and tables and use the Quick Display mode.

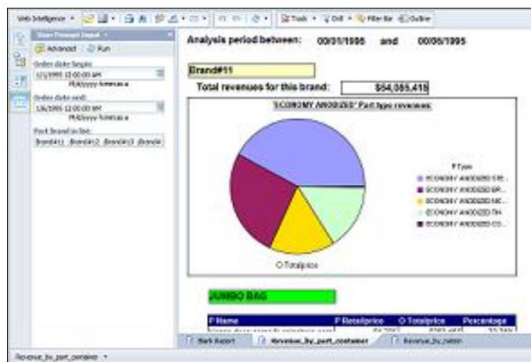
The data provider volume refers to the size of resulting data, retrieved from the reporting database by the query. It is stored in the Web Intelligence document. In this example, the data source on the RDBMS side consists of 3 million rows.

The following Relational Scenario has been used for the sizing: Login – Open Document - Navigation - Refresh – Drill - Export to PDF – Close and Logout.

Data-source type used: UNX, single source

Web Intelligence

The Web Intelligence documents used in the sizing tests have been summarized below, with respect to their composition and complexity.



Fields in query result set: 9

Tables joined: 8

Record selection clauses: 2

Parameters/Prompts: 3

Report filter clauses: 0

	Simple Elements		Complex Elements			Total
	Total	Cell	Total	Chart	Table	
Blank Report						
Page Header	0	0	0	0	0	0
Body	5	5	0	0	0	5
Page Footer	0	0	0	0	0	0
Revenue_by_part...						
Page Header	0	0	0	0	0	0
Body	1	1	0	0	0	1
Section 1	7	7	1	1	0	8
Section 2	1	1	1	0	1	2
Page Footer	0	0	0	0	0	0
Revenue_by_nation...						
Page Header	0	0	0	0	0	0
Body	1	1	0	0	0	1
Section 1	0		2	1	1	2
Page Footer	0	0	0	0	0	0
Total	15	15	4	2	2	19

Specific parameters used to get the Quick sizer results:

BI platform setting - Auditing on. Platform Search Indexing has been stopped.

Two Web Intelligence Processing server with following parameter changed – Maximum Connection 500

1 Custom APS with BI Commentary service – connection to BI commentary database in jdbc to get better performance on some specific actions.

1 Custom APS with DSL Bridge service and Web Intelligence Monitoring Service.

1 Custom APS with Client Auditing Proxy Service, Monitoring Service, Visualization Service, Web Intelligence Monitoring Service, Document Recovery Service.

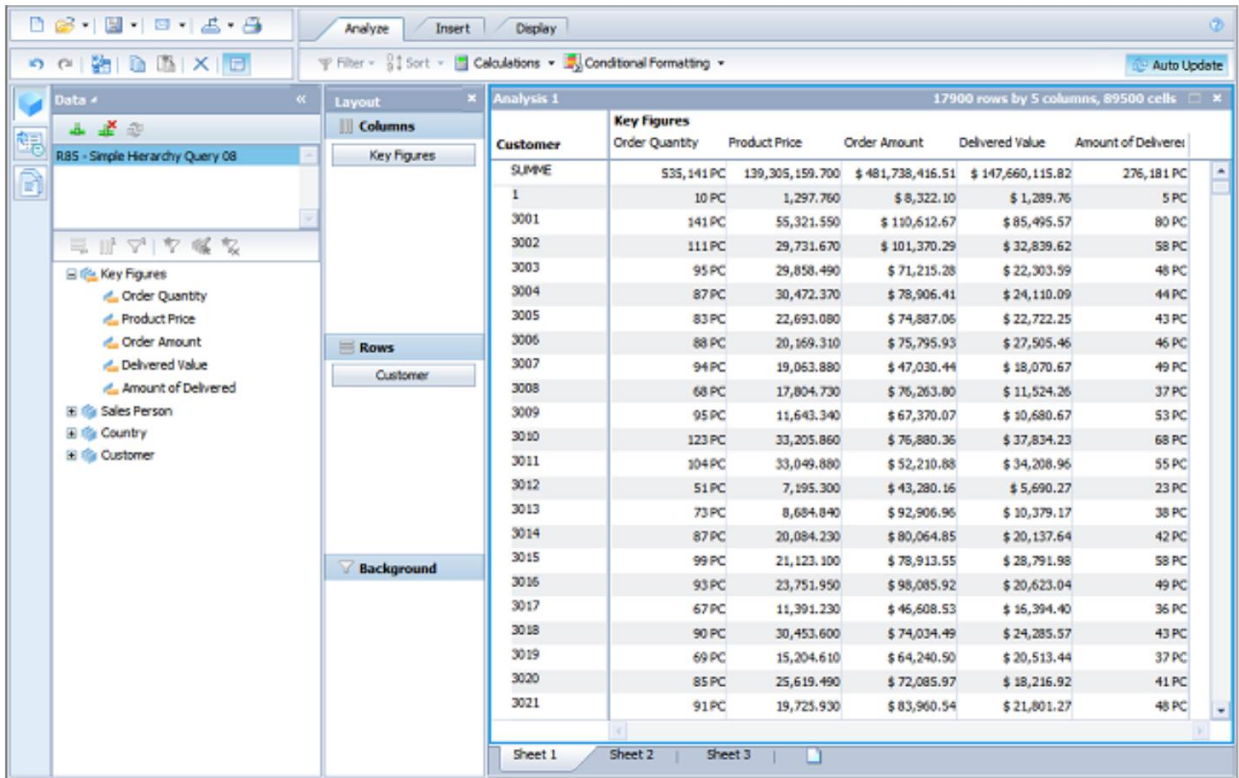
Web Server (Tomcat8) — Max Thread 400

Cache enabled (enabled by default).

Parallel DP Refresh enabled (BI 4.2 new feature, enabled by default).

Analysis, edition for OLAP

The information below represents the typical Analysis Workspace used in the sizing tests.



Dashboards

Dashboards like the following were used in the sizing tests. There were 2 queries issued every time the dashboard was refreshed.

C Custkey	C N	O O	O O	O Totalprice
5	Cust	33	F	94427.03
1	Cust	71	O	162977.69
2	Cust	97	F	64261.51
7	Cust	99	F	82237.18
2	Cust	132	F	107117.34
1	Cust	164	F	186807.98
1	Cust	224	F	137011.1

C Custkey	C Name
1	Customer#000000001
2	Customer#000000002
3	Customer#000000003
4	Customer#000000004
5	Customer#000000005
6	Customer#000000006
7	Customer#000000007

C Address
IvHzIApeRb ot,c,E
XSTf4,NCwDVaWNe6tEgwfmRchLXak
MG9kdTD2WBHm
XxV5JslAGtn
KvpyuHCplrB84WgAiGV6sYpZq7Tj
sKZzOCsnMD7mp4Xd0YrBvx,LREYKUWAH yv
TrGe5oaZNoVePvLI5kRvXRfkasDTea

C Phone
25-989-741-2988
23-768-687-3665
11-719-748-3364
14-128-190-5944
13-750-942-6364
30-114-968-4951
28-190-982-9759

24.1.2 Report data description

This table represents total data set sizes. Other tables earlier in this document show the data volume returned from a query which is always a subset of the entire dataset. Crystal Reports for Enterprise, Crystal Reports 2016, and Dashboards tests were executed against DS2 and DS3 on a relational database. Web Intelligence tests were executed against DS1, also on a relational database.

Data Set	LINEITEM	ORDERS	PARTSUPP	PART	CUSTOMER	SUPPLIER	NATION	REGION
DS1	367.9 MB	81.5 MB	55.75 MB	12.9 MB	11.6 MB	0.75 MB	0.01 MB	< 0.01 MB
	3M rows	750K rows	400K rows	100K rows	75K rows	5K rows	25 rows	5 rows
DS2	72.75 MB	18.3 MB	13.7 MB	3.1 MB	2.9 MB	0.2 MB	0.01 MB	< 0.01 MB
	400K rows	100K rows	53K rows	13K rows	10K rows	650 rows	25 rows	5 rows
DS3	4.8 MB	1.1 MB	0.75 MB	0.2 MB	0.2 MB	0.02 MB	0.01 MB	< 0.01 MB
	40K rows	10K rows	5K rows	1K rows	1K rows	65 rows	25 rows	5 rows

Analysis, edition for OLAP was tested against an SAP BW data source. Information about the dataset used in the sizing tests is below. The average BEx query used had 4 characteristics, 5 hierarchies, 3 key figures, and 0 variables.

Dimension	Hierarchy / Characteristic	Members	Levels
Material	Material	100K	1
Material	BMK_HIER	100K	6
Material	MAT_N2_HIGH_LARGE	60K	11
Material	MAT_N6_LINK_NODES	60K	11
Sold-to Party	Sold-to Party	100K	4
Sold-to Party	CUST_N1_FLAT_LARGE	100K	4
Sold-to Party	CUST_N5_UNBALANCED	100K	8
Sales Organization	Sales Organization	1K	1
Sales Organization	SALESORG_N3_SMALL	500	4
Calendar Day	Calendar Day	365	1
Calendar Day	QuaDay	365	2
Calendar Day	QuaMonday	365	3
Country of Sold-to Party	Country of Sold-to Party	21	1
Calendar Year/Month	Calendar Year/Month	12	1
Product Hierarchy	Product Hierarchy	10	1

24.1.3 CMS repository data description

Info Object Type	# of Info Objects
Folders	37
Documents - WebI	268
Documents – Crystal Reports	22
Documents - Others	58
Inboxes	5005
Personal Categories	5005
Users	5006
Groups	61
Deployment Files	3048
Others	732
Application Objects	927
Total	8020

24.1.4 Test hardware specifications

Machine	CPU (SAPS)	CPU Arc h	RAM (GB)	Disk W (MB/s)	Disk R (MB/s)	Network
M1 - BusinessObjects BI 4	23,000	x64	48	852	1170	10 Gbit
M2 - Reporting DB (Relational)	23,000	x64	48	863	1120	10 Gbit
M3 - Reporting DB (SAP BW 1)	23,000	x64	48	955	1252	10 Gbit
M4 - Reporting DB (SAP BW 2)	23,000	x64	24	891	1192	1 Gbit
M5 - CMS Repository DB	23,000	x64	48	974	1219	10 Gbit
M6 - Load Generator	12,000	x64	32	25	616	1 Gbit