PUBLIC

SAP BusinessObjects Business Intelligence Platform 4.3

Sizing Guide

Version 1.0

Date: June 9, 2023

Contents

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

1 Introduction

This document discusses the sizing of the SAP Business Intelligence 4 suite of services, hereafter referred to as SAP BI 4 and is intended to help you ensure your users can access their decision-making information in a timely manner.

A BI system is used for searching information interactively and most users would consider "the Internet broken" if their request takes more than a couple of seconds to respond. They have expectations and often a business need to have a responsive system available to them.

Sizing BI is very different compared to sizing of other types of Enterprise software. Retrieving 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.

Estimating the right amount of capacity in these four aspects for your system is key to success and depends on several variables:

- The number of users, the types of users, their usage…
- The number of used BI tools
- The number of data sources
- The deployment options…

1.1 Who Should Use This Document?

You should read this document if you are:

- Planning to deploy SAP BI 4 software.
- Optimizing or tuning an existing SAP BI 4 deployment.
- Working on a Sizing strategy, especially in conjunction with sizing tools like the SAP Quick Sizer Tool.

1.2 Components and Platforms in Scope

In this document, the guidance to estimate hardware focuses on a landscape containing:

- Web Intelligence
- Crystal Reports 2020
- Crystal Reports for Enterprise
- Analysis, edition for OLAP

General sizing tips are also provided for:

- Mobile BI
- Live Office
- SAP BI 4 Platform services: Promotion Management, Search, Data Federation
- SAP BW

It is assumed that you are already familiar with core concepts from the SAP Business Intelligence Platform Administrator Guide (SAP BI 4 Admininstration Guide). In this guide, you can find conceptual information and technical details on many topics intentionally not covered in this document.

1.3 Disclaimer

This document demonstrates how someone might perform sizing of a BI 4 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 guaranteed successful deployment.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

Customers, partners, or consultants might 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 large deployments or when your sizing requirements need to be extremely precise, we recommend considering 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.

2 Before You Start

2.1 SAP BI 4 Machine Requirements

To know how much computing power is required for a deployment, you need to know the computer's performance rating. This is measured in SAPS (SAP Application Performance Standard), a hardware-independent unit of measurement 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 processed business order line items per hour.

The SAPS rating of the system's intended hardware allows you to determine the required amount of hardware to meet a system's processing requirements. The SAPS rating takes into account CPU processing power as well as the computer's I/O capabilities. Computers can have different SAPS processing capabilities even with the same type of CPU.

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. Actual cores refer to full cores in a CPU; hyper-threaded cores should not be counted for sizing purposes.

For SAP BI 4, the minimum hardware for evaluation purposes is 8 000 SAPS of processing power and 16 GB of memory. This is defined in the SAP BI 4.3 Product Availability Matrix. You can determine the hardware requirements for deployments of any kind by following the remainder of this guide. If you are using virtualized hardware, can you dedicate CPU and memory to your BI system? If this is not the case, the SAPS number and your sizing calculations will not be accurate.

2.2 Network Bandwidth Recommendations

The time to login to SAP BI 4.3 BI Launch pad and to open a Web Intelligence document depends on the available network bandwidth since it includes the time to compute the document and transfer it over the network.

To assess this bandwidth impact, some tests have been run with Google Chrome in two configurations:

- **Without cache:** To make sure that the browser cache has no effect, the browser cache is emptied, Chrome is restarted and used in **Incognito** mode.
- **With cache:** Workflows are run several times before the test measurements and Chrome is started in **Standard** mode (not in **Incognito** mode).
- The Chrome version is the one supported by SAP BI 4.3.

The following chart shows the test results with a 21 MB Web Intelligence document containing a cross-table with 1 500 columns and 100 rows.

- The X-axis represents the available bandwidth, in Mbps.
- The Y-axis represents the time to execute the actions, in seconds.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

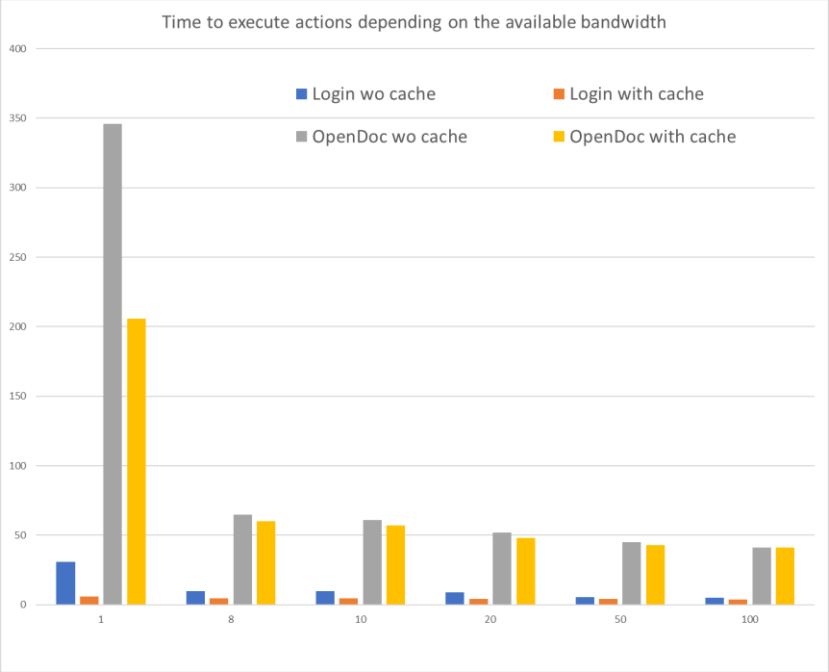

Based on these results, we strongly recommend to have, at least, 8Mbps of available network bandwidth on your client computer/laptop. If the network bandwidth is good enough, the cache has almost no impact on the final execution time.

2.3 Sizing Tools

You can use the SAP Quick Sizing Tool and the SAP BI 4.3 Configuration Wizard to start your sizing project. Since these tools do not cover all deployment scenarios nor all data sources, refer to the remainder of this document for more information to accurately size your deployment.

2.3.1 SAP Quick Sizer Tool

The SAP Quick Sizer Tool 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. Fill in the online questionnaire, an up-to-date survey 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.

The 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.

The Quick Sizer integrates the functionalities of the previous "BI4 Resource Usage Estimator" (see blog).

2.3.2 SAP BI 4.3 Configuration Wizard

Starting SAP BusinessObjects BI 4.3, the System Configuration Wizard helps you defining how many Adaptative Processing Services (APS) you need (§3.6.2, *Adaptive Processing Service*,) based on your deployment size.

To start the System Configuration Wizard, open the CMC and under **Manage**, select **System Configuration Wizard**. For more information, see the SAP BI 4 Administration Guide.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

2.4 Users

2.4.1 Active Users, Active Concurrent Users

Many sizing exercises begin with a total number of users, for example, a system sized for 10 000 **Users**. The next step is to determine how many of those users are logged-in to the system at the same time, based on your users' usage. If you cannot provide guidance on this number, it is 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 are concurrently generating load. This is another opportunity to take your users' 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 SAP BI 4 sizing tools, specify your inputs in terms of Active Concurrent Users.**

In our example, we went from 10 000 users to 100 active concurrent users, which translates to 1% concurrency.

The typical concurrency recommendations depend on the deployment size 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. In such case, use the Quick Sizer and specify more users as **Expert Users** (§2.4.2, *User Class Definitions*) to account for the increased use of the system.

2.4.2 User Class Definitions

Customers often underestimate the amount of use that a user will make of BI system. Class definitions are used to identify how much the system is used and how much load is placed on it. Knowing the mix of user classes is important since it directly affects the performance and resources required by the system. Three classes of users can be identified:

- **Information Consumers**: The least active users, they 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 users, they 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.

2.5 SAP BI Architecture

The SAP BI Platform is a set of base services that enable BI services to run in a distributed deployment. Each node in a SAP BI 4 cluster runs a Server Intelligence Agent (SIA) that proposes selected services on that node. Nodes in a cluster communicate over a service bus. The SAP BI Platform is designed using a 64-bit architecture to ensure scalability to the largest customers.

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 SAP BI 4 Administration Guide for a complete list of services and how to configure them).

2.5.1 Application Tier

This tier is made of the Web Application Services, by default Tomcat, that hosts the web applications to which the user can connect to access the SAP BI 4 system over the web:

- SAP BI Launch Pad: The users portal.
- Central Management Console (CMC): The administrators portal.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

- REST Web Services. Since SAP BI 4.3, these REST Web Services are deployed by default in the Tomcat server rather than in the Web Application Container Server (WACS) to improve scalability and remove the Single Point of Failure when deployed only in WACS.

2.5.2 Intelligence Tier

This tier, also referred as “Management Tier”, is made of:

- The Central Management Server (CMS) is a combination of service plus database for storing and retrieving this information. The database used by the CMS for its repository is sometimes called the Intelligence Tier database. It contains all the objects and configuration information that control and secure the platform.
- Core services, such as the File Repository Server (FRS).

2.5.3 Processing Tier

The Processing Tier contains the BI Services and the additional services that support them.

- Adaptive Job Service: This server is used to perform most scheduling actions. It hosts several services like the Authentication Update Scheduling Service, Users and Groups Import Scheduling Service or the Web Intelligence Scheduling Service.
- The Adaptive Processing Service (APS) is a host that runs many BI services.
  - 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….

    These services are Java-based and the APS hosts also one JVM to run these services. More than one APS may run on a given node and several nodes may run APSes in a BI landscape.
- Crystal Reports Processing Service
- Crystal Reports Cache Service
- Web Intelligence Processing Service
- Non-APS Services: Not all BI services run inside an APS. Those that do not are run separately by the SIA.

The BI Services are intended to process the SAP BI Tools, for example:

- Crystal Reports, available in two versions:
  - Crystal Reports 2020 is the original Crystal Reports product line.
  - Crystal Reports for Enterprise, which is deprecated (§2.5.4, *Deprecated Areas*).
- Web Intelligence, that provides analysis and reporting capabilities to the suite. It can query universes created with Universe Design Tool (UNV) or with Information Design Tool (UNX) as well as SAP BW and SAP HANA. The Visualization Service (also known as CVOM) provides data visualizations (charting, graphics) services, the DSL Bridge houses query generation, metadata definition and retrieval on 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.

2.5.4 Deprecated Areas

As announced in the SAP Analytics BI Statement of Direction, some products and components will no longer be supported in the future:

- Universe Design Tool and the universes created with it,
- Multi-source universes and associated connectivity,
- SAP Lumira,

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

- SAP Crystal Reports for Enterprise,
- SAP BusinessObjects Analysis, edition for OLAP,
- SAP BusinessObjects Live Office,
- AIX and Solaris platforms.

It is recommended to stop using these products and components and to start investigating alternative solutions to replace them.

3 Deployment and Sizing Methodologies

3.1 Sizing Steps and Methodology

Sizing a SAP BI 4 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 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.
- The types of data sources used influence the load and needs of the system.

These user requirements are used to define the required amount of processing and 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.

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 must be considered.

For each tier, especially the Processing Tier, the processing power and memory requirements can be determined by Quick Sizer and as suggested by System 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 (§2.1, *SAP BI 4 Machine Requirements*).

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 and the amount of memory available to the machine.

You cannot assume any SAPS rating of the machines. You must compute them, for example by visiting the website: SAP SD Standard Application Benchmark Results**.**

The Quick Sizer is used 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.

3.2 Pre-Sizing Checklist

The goal of the sizing exercise is to calculate the peak load placed on the system. To proceed with the sizing steps, you need to know the information listed below. 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 (see *User Class Definitions*, §2.4.2) of each type of BI tool.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

- **Usages**: Common user workflows are also important to know. Will the documents be scheduled to run at night and only be viewed during the day or will they be refreshed by users? If so, how frequently? If you expect users to open five documents and refresh them at the same time, that is five times the load of one user.
- **Data Sources**: What types of data sources are used: direct-access SQL databases, UNV Universes, UNX Universes, SAP BW, SAP BW on SAP HANA or SAP HANA? It's important to know which data sources will be used for most BI processing so 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 have minor 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.
- **Web Intelligence** Documents:

  In *Reading* 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 *Design* mode, the complexity of the document structure (for example: auto-fit setting, page breaks, sorts, conditional formatting, styles, sections, etc.) impacts the amount of system resources required by Web Intelligence Processing Tier when a report is generated.
- **SAPS Rating**: What types of machines will you be deploying to? You may build 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.

  For the sizing exercise, the number of **SAPS per core** needs to be determined (§2.1, *SAP BI 4 Machine Requirements*). 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.**

3.3 Quick Sizer Setup

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 reflect the sizes and types of documents the system will process.

For more information, see the following:

- Quick Sizer Tool
- Using Quick Sizer for Sizing SAP BusinessObjects Business Intelligence Suite Applications

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. thinkt.)	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

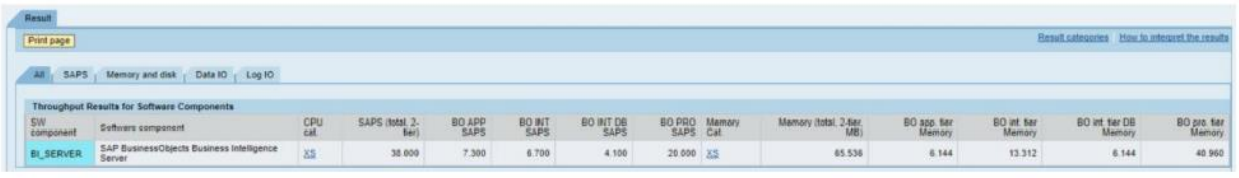
Result

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	38.800	7.300	6.700	4.100	20.000	XS	65.536	6.144	13.312	6.144	40.960

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

3.4 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.

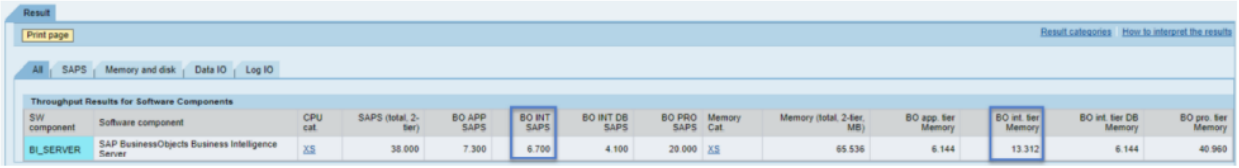

3.4.1 Central Management Server

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 many 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 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 does not offer low latency horizontal scalability, you will need to ensure to follow your DB vendor's guidance on how to size and scale vertically, especially for deployments with thousands of active concurrent users.

3.4.2 Central Management Server Database

The CMS database's performance is particularly important to the system's overall performance. It is strongly recommended that your CMS database be sized and tuned for high performance operation by a qualified DBA. For larger deployments, the Intelligence database should be installed on a dedicated machine.

Be sure to follow your database vendor's guidance on 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 free memory to do that processing.

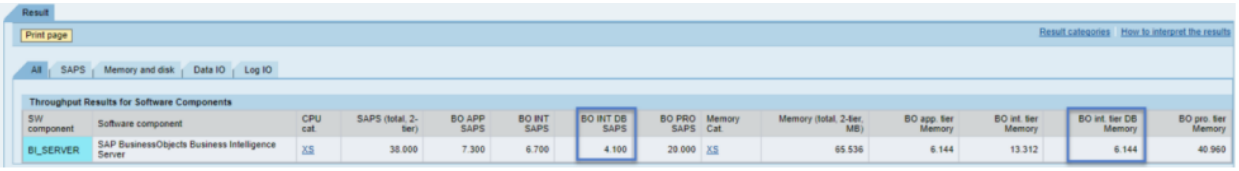

The CMS database performance significantly impacts the system's performance. The following tuning tips have been defined for the SAP Sybase databases, but they may apply to other databases as well.

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

   Data cache allows efficiently reducing IO. A sufficient 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 servers do 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 may generate deadlocks when, for example, several records are stored together on the same page. If one record is updated while another is being read, one of these two sessions must wait until another is finished. Setting Lock Granularity to use row level locking can avoid such contention issues.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

4. **Parallel Processing**

   Parallel Processing is important to set so the database can fully utilize all the CPU cores available on the machine.

We also recommend reading the KBA 2420789.

3.4.3 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 becomes 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 are critical for the operation of the FRS.

3.5 Application Tier

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

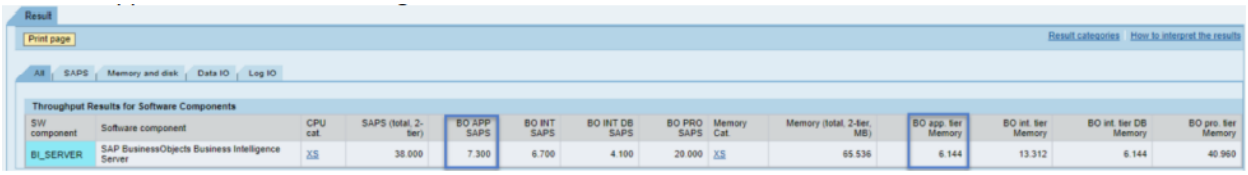
Result

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

It is recommended to follow guidelines as per Quick Sizer to allocate memory for Application Tier processing request.

As you scale out your application tier, you should add a web additional application server instance for every 500-additional active concurrent users. For better performance, it is also recommended that you configure at least 8 GB for heap size and 900 maximum threads for each instance. This is a minimum. Refer to 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 SAP 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 does not 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.

3.6 Processing Tier

3.6.1 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) used by the Connection Server running 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

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

- `<installdir>\SAP BusinessObjects\SAP BusinessObjects Enterprise XI 4.0\dataAccess\connectionServer` on Windows.

Set the `-Xms` and `-Xmx` options as shown below to set memory usage to 1GB at startup and 2GB maximum.

```
<JavaVM>
    <!-- The default JVM configuration can be overriden 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>
```

3.6.2 Adaptive Processing Service

The default installation of SAP BusinessObjects Business Intelligence Platform installs one Adaptive Processing Server (APS) per host system. The Adaptive Processing Server is a generic server that hosts services responsible for processing requests from a variety of sources.

3.6.2.1 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 to group these services 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 number of SAPS of 1000 should be used.

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

3.6.2.2 Platform Search Service

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 frequently content is changed in your system.

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

3.6.2.3 Data Federation Service

Multi-source universes created in Information Design Tool have been announced as deprecated (§2.5.4, *Deprecated Areas*). If you still use them, it is recommended to create an APS to run the Data Federation (§3.8.2, *How to Create a New APS*). In large deployments, it is easier to monitor and configure the Data Federation Services.

For medium deployments, allocate 1 GB of RAM. For larger deployments, allocate 3 GB of RAM (§3.8.4, *How to Change the Memory Setting of an APS*).

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

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

For processing calculations, a nominal number of SAPS of 1 000 should be used.

3.6.2.4 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.

3.6.3 Adaptive Job Server

The Adaptive Job Server does not require splitting or configuration as described above for the APS since the Adaptive Job Server runs services as separate dedicated processes.

Sometimes it can be necessary to increase the Xmx dedicated to the Adaptive Job Server, especially for workflows like Promotion Management. You can find more details in the KBA 2286419 and 2398299.

3.7 Sizing Analysis and Scale-out

Scaling-out your deployment is often necessary when the Quick Sizer recommends using more than one machine to answer your BI requirements. Additionally, some analysis and thought are 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.

3.7.1 Scale-out Machines

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 2 000 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 if the machine's processing power or memory limit has not exceeded the workload assigned to it. If it fits, you can move on to the next workload. If it does not, 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 dedicated to that workload should be created on the machine for those services. A machine may have other services defined on it, possibly because of the standard installation. Any services not intended for use on a machine should be stopped, disabled, and removed. 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 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.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

Similar recommendations also apply for the Application Tier hosting Web Application Server and Processing Tier's DSL Bridge Service.

3.7.2 Scale-out Memory Expectations

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

The following table lists the minimum memory requirements for some services and can be used to compute the memory requirements of a cluster node. These values are only minimum values since you must consider other services that may run on the node.

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

For Web Intelligence, the 64-bit architecture was released in SAP BI 4.3, and there is no need to set memory requirements.

3.8 APS Configuration

The Adaptive Processing Server (APS) is the application that hosts most of the BI services. It is strongly recommended to not alter the original APS defined by the installer. If that APS is not needed, it should be left intact, stopped and disabled.

To get better system scalability, you can 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 during the installation, making at least some amount of APS configuration necessary on every machine. Configuring APSes is an essential part of sizing SAP BI 4.

In SAP BI 4.3, one APS is defined during installation. The rest of this section shows the basics of APS configuration.

3.8.1 APS Service Groupings

When establishing your APSes, the following groupings are recommended for redundancy and fail-over with minimal machine dependencies.

The services in these first groups are related.

Web Intelligence APS	WebI DSL Bridge APS	Visualization APS
Web Intelligence Monitoring Service Document Recovery Service Rebean Service Custom Data Access Service Excel Data Access Service	DSL Bridge Service Web Intelligence Monitoring Service	Visualization Service Web Intelligence Monitoring Service

DF APS	Connectivity APS	Materialization APS

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

Data Federation Service	Adaptive Connectivity Service	Set Materialization Service

When duplicating services to reduce or eliminate machine dependencies, a second list of services that are mostly unrelated and have no dependencies on the first group can be defined.

Primary APS	**Promotion Management APS**	**Analysis APS**
Insight to Action Service	Visual Difference Service	Multi-Dimensional Analysis Service
Publishing Post Processing Service	Promotion Management Git Service	BEx Web Application Service
Publishing Service	Promotion Management Service	
Security Token Service	Version Management Service	
Translation Service		
BI Commentary Service		
BI Mobile Service (OCA)		
Analytics Hub Service		

Monitoring APS	**Search APS**	**Auditing APS**
Monitoring Service	Platform Search Service	Client Auditing Proxy Service

3.8.2 How to Create a New APS

1. In the CMC, go to the Servers page.
2. Right-click **Service Categories** and select **New** > **New Server** in the contextual menu.

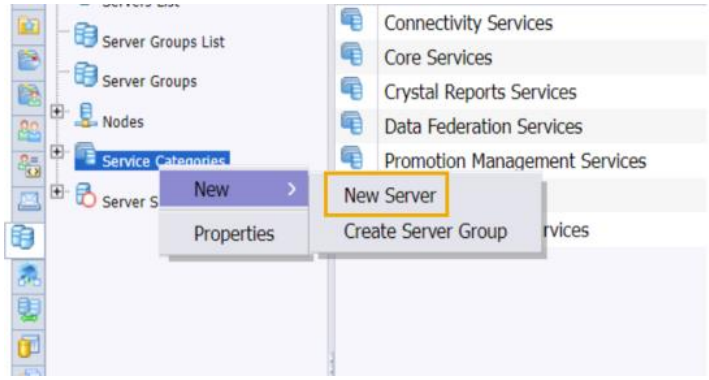

3. In the **Create New Server** dialog box, choose the most appropriate Service Category.
4. Select the first of potentially many services from the Select Service list, then press **Next**.
5. To add more services to the APS, select them in the Available Additional Services list and click “>” to add them to the Selected Services list. Choose **Next**.
6. Give the APS a unique name.
7. Choose **Create**.
8. At this point, APS is created but is not yet enabled or started. You can enable and start it by right clicking the new service in the CMC and selecting the appropriate options.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

3.8.3 How to Permanently Disable an APS

1. In the CMC, go to the **Services** listing
2. Right-click on the server (APS) that you wish to modify.
3. Choose **Disable Server**.
4. Right-click on the server (APS) that you wish to modify.
5. Choose Properties
6. Find the option **Automatically start this server when the Server Intelligence Agent starts**.
7. Uncheck the option.

3.8.4 How to Change the Memory Setting of an APS

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

3.8.5 How to Change the Services Running in an APS

1. In the CMC, go to the **Services page**.
2. Right-click on the server (APS) that you wish to modify and stop it.
3. Choose **Select Services**.

See the SAP BI 4 Administration Guide for in-depth information on configuring APSes.

3.9 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 does not run in a steady state, and is influenced by user workload.

During initial installation and testing, monitoring of the system is particularly important since the best sizing exercise cannot fully know how a system will behave and perform. Simulating user load is highly recommended, using products such as Apache JMeter or Micro Focus LoadRunner. Start with a small number of simulated users, bring the workload up slowly, and see how the system reacts. Proceeding gradually will uncover any issues 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 tools needed to do this.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

3.9.1 Scaling-out SAP BI 4 Servers

A particularly 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 SAP BI 4 Administration Guide for guidance on creating clustered deployments.

SAP BI 4.3 can be installed as a full installation on each server and any unnecessary services be turned off. This simplifies the installation procedure and allows future changes in the deployment. For example, if the role of a machine changes, all the SAP BI 4 components are there to be enabled and disabled, as necessary.

3.9.2 Building the System

It is especially important to methodically «build out» the system.

1. 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.
2. 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
3. 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 SAP BI Platform. A dedicated CMS database running on dedicated hardware is always recommended. Work with your DBA to ensure it 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 BI 4 landscape, it is recommended to methodically start each of your Server Intelligence Agent (SIA) nodes and ensure all services are correctly started before starting another node.

3.9.3 Monitoring SAP BI 4

The SAP BI 4 suite has several built-in monitoring capabilities that can be used to measure the system's performance as you build it up and continue into production.

Usage information for your SAP BI 4 deployment can be obtained by using audit reports. This can be very helpful to determine which BI Tools, documents, etc., are used, and how many users are using them. You can download a set of reports that are helpful for reporting on the SAP BI audit database here.

3.9.3.1 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:

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

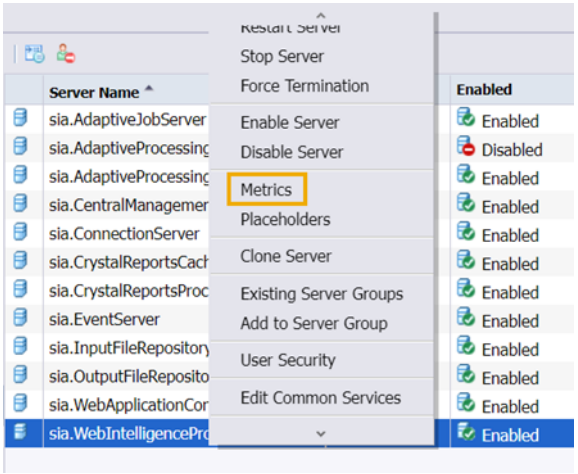

Some of the information shown includes:

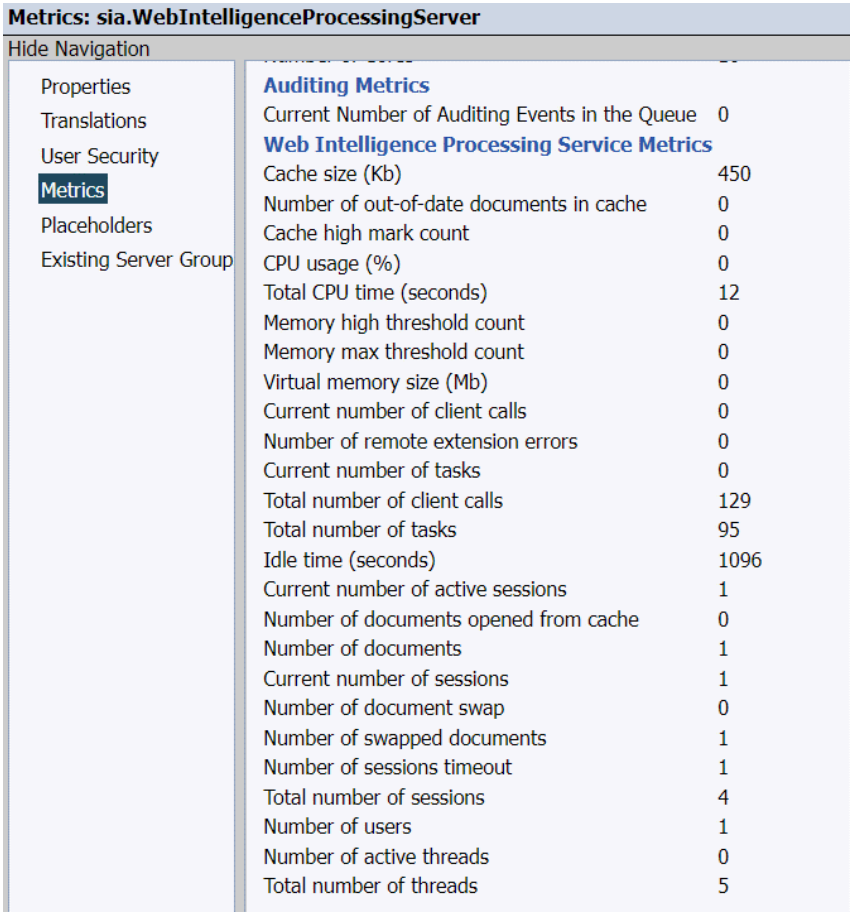

3.9.3.2 Configure Report Processing Time Alerts

The Monitoring features of SAP BI 4 allow you to receive alerts when specific reports take too long to process. To configure an alert that run a Web Intelligence or Crystal Reports document periodically and be alerted if it takes too long to run:

1. Log in the CMC and open in the **Manage** section, click **BI Admin Studio**.
2. 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 SAP BI 4 Administration Guide for more information on creating and configuring probes.
3. 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

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

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:

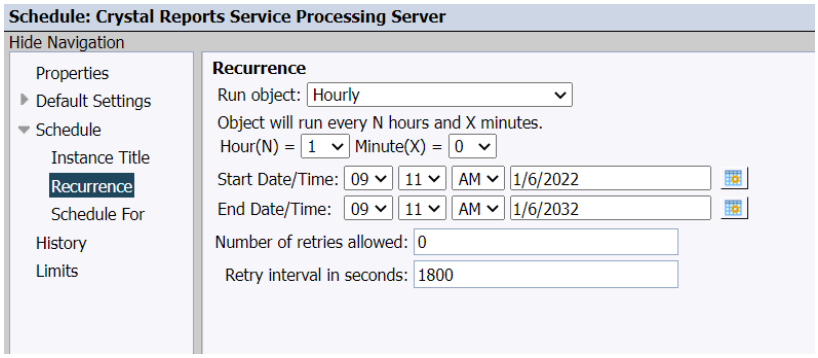

4. In the Watchlist, create a new Watch that has rules based on the execution time of the probe. Set the Caution to 5 000 (milliseconds) to receive an alert when the report takes longer than 5 seconds. This is the most time a user expects a report to require processing.
5. Set the Danger Rule to be some larger amount of time, such as 10 seconds.

See the SAP BI 4 Administration Guide for more information, including how to create 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.

3.9.4 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% to handle peaks in the range of 80% of CPU usage. Analyzing the historical usage of your deployment can help you determine whether 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.

3.9.5 BI Tool Simulation Workflows

Creating simulated user workflows is particularly 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 use.

Web Intelligence

1. Log into the BI Platform through SAP BI Launch Pad page.
2. Navigate to the folder that contains the Web Intelligence document.
3. Open the document.
4. Click **Refresh** to run the queries.
5. Answer prompts if any.
6. Navigate to a report.
7. Change input controls if any.
8. Start Drilling Mode, Drill Down, and End Drill.
9. Close the document and log out.

Crystal Reports

1. Use an OpenDocument URL to access a Crystal Reports document.
2. Enter your credentials to log into the BI Platform.
3. View the report in HTML viewer.
4. Enter the Dynamic Cascading Parameters for live data and view first page.
5. Go to the second page.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

6. Go to the page in the middle of the report.
7. Close the document and log out.

Analysis, Edition for OLAP

1. Log into the BI Platform through SAP BI Launch Pad page.
2. Navigate to the folder that contains the workspace.
3. Open the workspace.
4. Swap axes.
5. Print the analysis.
6. Close the workspace and log out.

3.10 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 do not 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 BI Platform web site for the latest information on SAP BI 4, especially the Upgrade area if that applies to you.
- To learn more about getting the most out of your infrastructure, see Sizing and Deploying SAP BusinessObjects BI 4.x Platform and Add-Ons.
- 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.

3.11 Resizing

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

To be informed accurately about the usage of the system, you may also take advantage of SAP BI 4.3 auditing features. A set of auditing reports is available for download here. These reports will help you determine how your BI services are being used.

Monitoring probe reports and CMS database performance is also recommended (§3.9.3, *Monitoring SAP BI 4*) to give you alerts when processing time is longer than expected.

4 SAP BI Client Tuning

The Quick Sizer Tool 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 with regards to individual BI tools, choose the option of changing the result level to «Line results + inputs».

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

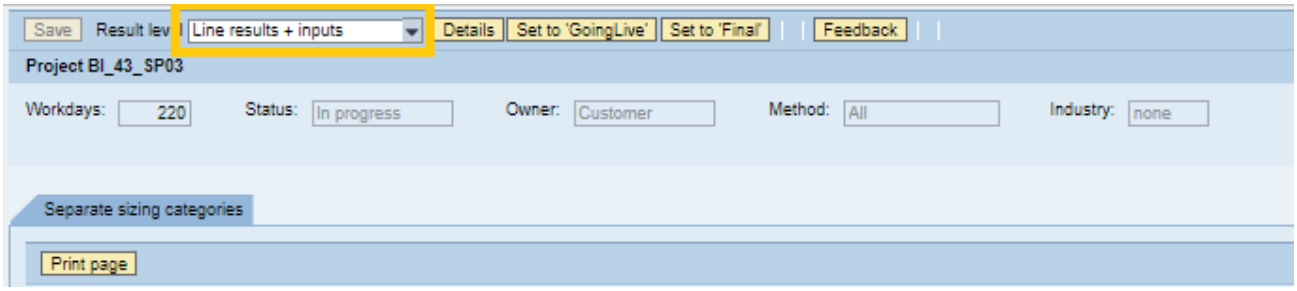

To do tool-specific sizing using the Quick Sizer, the numbers entered into the Quick Sizer are changed as each BI tool is analyzed. To get the starting independent sizing for a certain BI tool, set the user numbers for the other tools to zero.

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

4.1 SAP Web Intelligence

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

4.1.1 Instances

The Web Intelligence Processing Server is a multi-threaded process that scales both:

- Horizontally, by running Web Intelligence Processing Server on multiple machines.
- Vertically within a single machine, where multiple instances of the process may be required.

Although the Operating System allocates 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.

At least, two Web Intelligence Processing Server instances are needed for failover purpose whatever the deployment size is. More instances should be added only if either the data volume and document complexity is particularly large or the number of concurrent user requests increase. But too many Web Intelligence Processing Server instances may cause unnecessary overhead. The maximum number of instances should not exceed the number of CPU cores as this would certainly contribute towards an unnecessary overhead. It is important to keep the same settings for all Web Intelligence Processing Server instances. As a best practice, deploy and tune the first Web Intelligence Processing Server on the first machine, and then duplicate it or use the server *templates*. To do this, refer to the Business Intelligence Platform CMC Help and go to the section "**To set a configuration template".**

4.1.2 Maximum Connections

By default, a single Web Intelligence Processing server can handle up to 200 simultaneous document sessions.

It is different to the number of users or even user sessions and can be used to prevent too much memory or CPU cycles consumption due to the size and complexity of the documents. The appropriate value of this parameter depends on many factors like the request complexity or the dataset size.

It is possible to increase this number by modifying **Maximum Connections** parameter in the **Web Intelligence Core Service** properties:

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

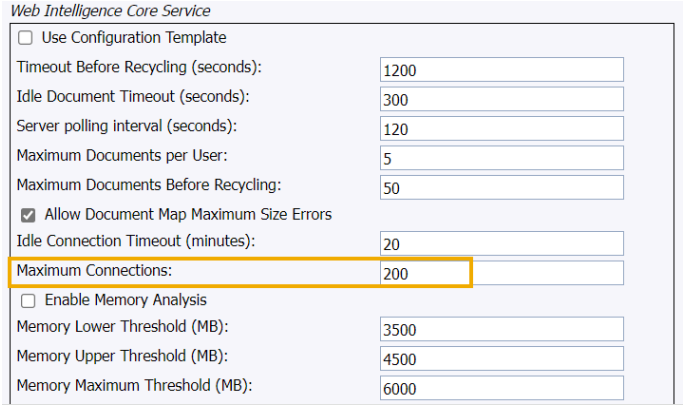

If your users encounter some «server busy» error messages, but your memory usage has not exhausted the machine, increase this **Maximum Connections** parameter until the possible errors disappear. Sometimes, it is necessary to limit server availability and thus reduce this parameter, to prevent the Web Intelligence Processing Server from becoming overloaded or consuming a huge amount of memory. The total number of Web Intelligence document sessions that can be maintained in the entire cluster is the number of instances multiplied by the number of connections.

To determine the right sizing, you need to know your maximum number of concurrent users (logged-in users plus concurrent scheduling tasks). This number is essentially the maximum number of concurrent connections you will need. Let's say you never have more than 100 concurrent users, but at peak times you may also have 25 scheduled tasks. Keep in mind that most users will only be doing one process at a time, but some will be, for example, refreshing one report while editing another simultaneously. A scheduled task will always have only one connection. So, you can safely say that you only need 200 concurrent connections (with some leeway).

4.1.3 Memory

By default, the Web Intelligence Processing server is set to use as much memory as it requires, and in most cases this setting is suitable. We strongly recommend keeping **Memory Analysis** disabled by default. Doing so, the process will use as much memory as needed, without generating an error related to memory overflow.

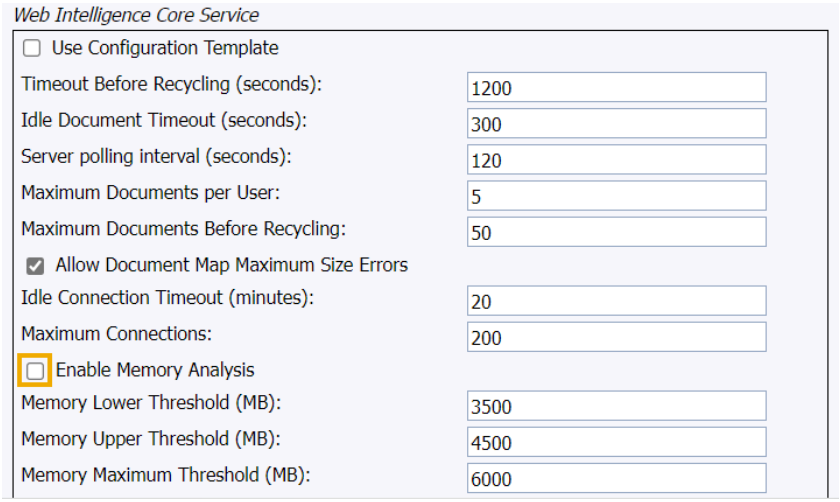

But this setting might stay enabled if you have upgraded from an earlier version.

- If you install SAP BI 4.2 or SAP BI 4.3 from scratch, this parameter is unchecked.
- If you update from SAP BI 4.0 or SAP BI 4.1, this parameter stays in its initial state: checked if it was checked and unchecked if it was unchecked.

4.1.4 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,

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

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 SAP BI 4 Administration Guide for details on Clustering.

For SAP BW, SAP HANA, and SAP Sybase sizing, you can go to the SAP Quick Sizer and find their respective administration guides:

- 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

4.1.5 Parallel Refresh

Parallel refresh has been introduced in SAP BI 4.2 and has improved refresh time of Web Intelligence reports containing several data providers.

Each connection, supported for the parallel refresh, has its own **Maximum parallel queries** limit (refer the SAP BI 4 Administration Guide for the default limits and supported connections).

Each Web Intelligence report also has its maximum number of data providers that can be refreshed in parallel:

Maximum Parallel Queries per document.

We recommend you to monitor the CPU for both SAP BI servers and databases. If you notice overload on your servers, then decrease gradually the parameters mentioned above until your deployment becomes stable:

- If you notice overload on database side, then decrease the **Maximum parallel queries** for the connection in Information Design Tool or Universe Design Tool.
- If your server consumes more than 60% of the CPU on average, then decrease the Maximum **Parallel Queries per document** in the CMC, in the Web Intelligence application's **Properties** page.

4.1.6 Query Stripping

When you run a query, you can improve the query time by selecting the Query Stripping option. With this option, objects that are not used in the document are no longer retrieved by the query (in *Design* mode, the corresponding objects are displayed in bold in the **Main** > **Objects** tab).

If you need to use this object, you need to run the query to fill the dataset.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

Although this option can reduce query time, it is even better to remove unused objects from the query, especially if you do not modify your document over time. This avoids the query stripping computation impact.

4.1.7 SAP HANA Data Sources

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

4.1.8 Disk space of your system

Before SAP BI 4.2, temporary data files were compressed when performing the first document on a report. Since SAP BI 4.2, this temporary file compression is no longer executed; which has greatly improved the performance of the first action performed in your report. However, since the file is no longer compressed when you save it, you need to allocate more space on your disk.

We recommend you multiply the space that you were using in the previous version by 10 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.

4.1.9 Comments

Since SAP BI 4.3, you can add comments to your documents. We recommend you change the default connection to the comment database from ODBC to JDBC to avoid performance degradation when opening a document or in the **Prompts** dialog box (see SAP Note 2346055).

4.1.10 Decimal Type for High Precision 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 setting the «Decimal» format option only for a measure where you need high precision.

4.1.11 Sizing for SAP Web Intelligence Upgrade

If you have already been running SAP Web Intelligence previous version and you want to know sizing for upgrade, we recommend keeping the same number of SAPS and adding an extra of 20% of memory on processing tier side for the use of new features. This percentage is based on average assumptions and must be adapted to your specific deployment and performance requirements.

4.1.12 DSL Bridge

The DSL Bridge is a service running in the default Adaptive Processing Server. It is 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.

4.1.13 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".

4.2 SAP Crystal Reports

The Quick Sizer allocates the Crystal Reports Cache Server to the Processing Tier. In general, it is recommended to run the Crystal Reports Cache Server 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.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

4.2.1 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:

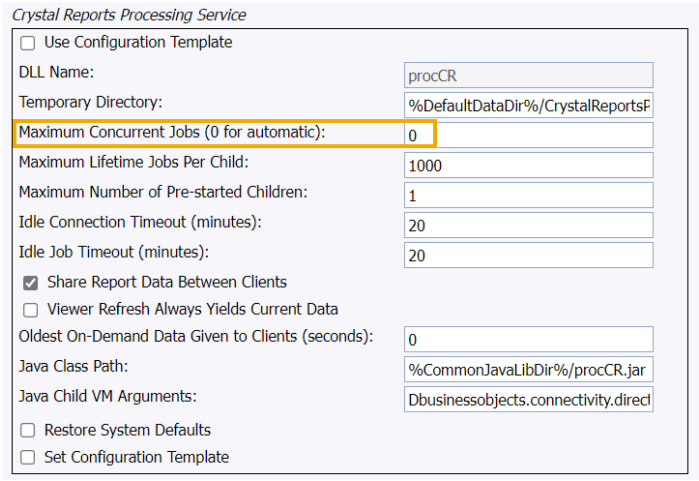

4.2.2 Increasing Memory Allocation for Improved Response Time

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.

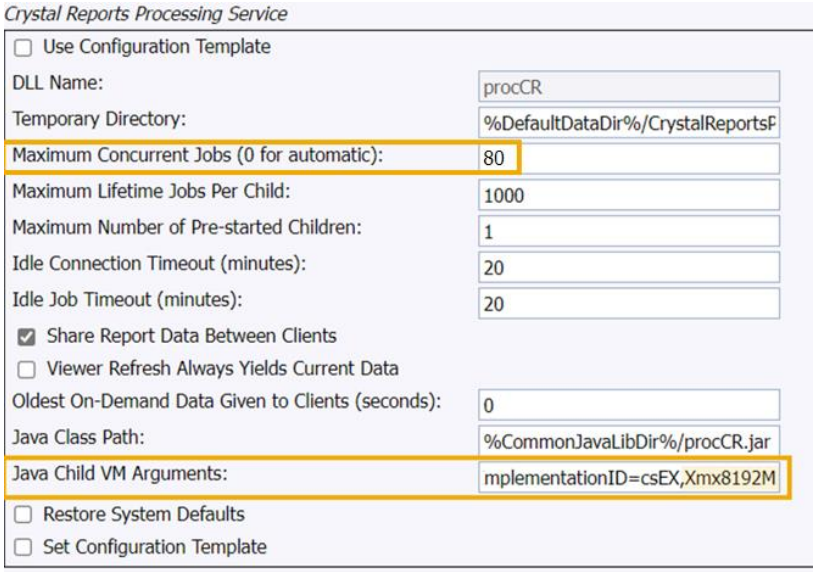

You may want to increase the number of concurrent jobs if your goal is to increase the number of reports that each child process can deliver. As shown here, the number of concurrent jobs has been set at 80.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

4.2.3 Crystal Reports 2020

To size Crystal Reports 2020 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 2020 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 2020 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 2 GB per core.

- In your sizing document, add in the SAPS and memory for Crystal Reports 2020.

4.2.4 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.

Once you know the number of cores needed, you can calculate the memory requirement by multiplying by 1.5 GB per core.

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

4.2.5 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 is not recommended to have more than one Crystal Reports Processing Service on any given machine. As you scale-out your deployment, it is better to add instances to other machines that do not currently have one.

Required memory is calculated as follows:

Crystal Reports 2020:

#of children = # of logical cores * 2.5

- Each child can use up to 800 MB of RAM
- 2 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

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

Data from the database is written to the disk in the temporary directory during processing. This disk's speed will affect processing performance.

Each Crystal Reports child process (as noted above) can process several report processing jobs at once. A job can potentially be shared among multiple user requests depending on sharing criteria. Crystal Reports 2020 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.

4.2.6 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 is recommended that you deploy them 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.

4.3 SAP Analysis, Edition for OLAP

It is recommended to create a dedicated APS (§ 3.8, *APS Configuration*) to run the Multi-Dimensional Analysis Service (MDAS) used for Analysis, edition for OLAP. In large deployments, it is easier to monitor and configure this APS.

Memory

It is recommended to configure the APS running the MDAS with a minimum of 4GB of memory (§ 3.8.4, *How to Change the Memory Setting of an APS*).

Sessions

The MDAS **Maximum Client** Sessions setting is often set to a small number by default. It is recommended to set it 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 and deploy them on separate machines.

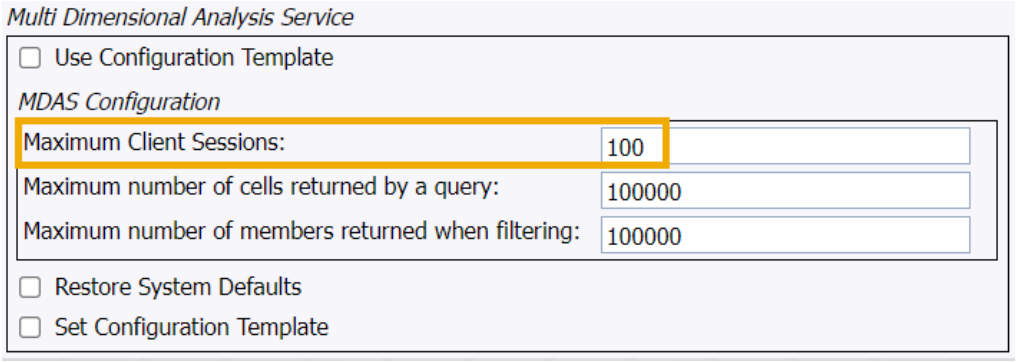

To size Analysis edition for OLAP service independently, set the user numbers in the Quick Sizer appropriately and set the user numbers for all other tools to zero.

4.4 SAP Live Office

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

If a document contains multiple BI report parts, refreshing that SAP Live Office document is equivalent to the same number of users refreshing these report parts on the BI system at the same time.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

For example, if a SAP Live Office document has three report parts and it gets refreshed, that is 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.

4.5 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 to ensure any increase in processing required by the SAP BusinessObjects Mobile server is accounted for.

Sizing information for Mobile can be obtained on the SAP Help Portal.

4.6 SAP BusinessObjects Design Studio

Sizing information for Design Studio can be obtained on the SAP Help Portal.

4.7 SAP Lumira

Sizing information for SAP Lumira 2.x can be obtained on the SAP Help Portal.

5 SAP BW Considerations and Recommendations

When using an SAP BW data source, some best practices can be used to improve performance.

5.1 SAP BW Configuration

- Apply the latest fixes available from SAP by applying the latest support packs to your SAP BW system.
- Size your SAP BW for the anticipated query load expected from the SAP BI 4 system.
- Have a BW expert tune the system generally and BEx queries specifically for your SAP BI 4 usage.

5.2 SAP Bex Query Performance

Check the performance of SAP 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 requires in the front-end and in the SAP BW analytic engine.

5.3 Navigational Attributes

A large number of navigational attributes defined in the underlying InfoProvider may impact the overall performance. This can lead the SAP BI 4 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.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

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») 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.

5.4 SAP BW Specific Tuning and Configuration

For SAP BW, you may set specific properties:

- For the Read Mode, select the **Query to Read When You Navigate or Expand** Option. This **Read Mode** option improves performance in situations where there are many Restricted Key Figures and other calculations in the Query Definition.
- Select the option **Use Selection of Structure Elements** to ensure the structure elements are sent to the database for processing

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.

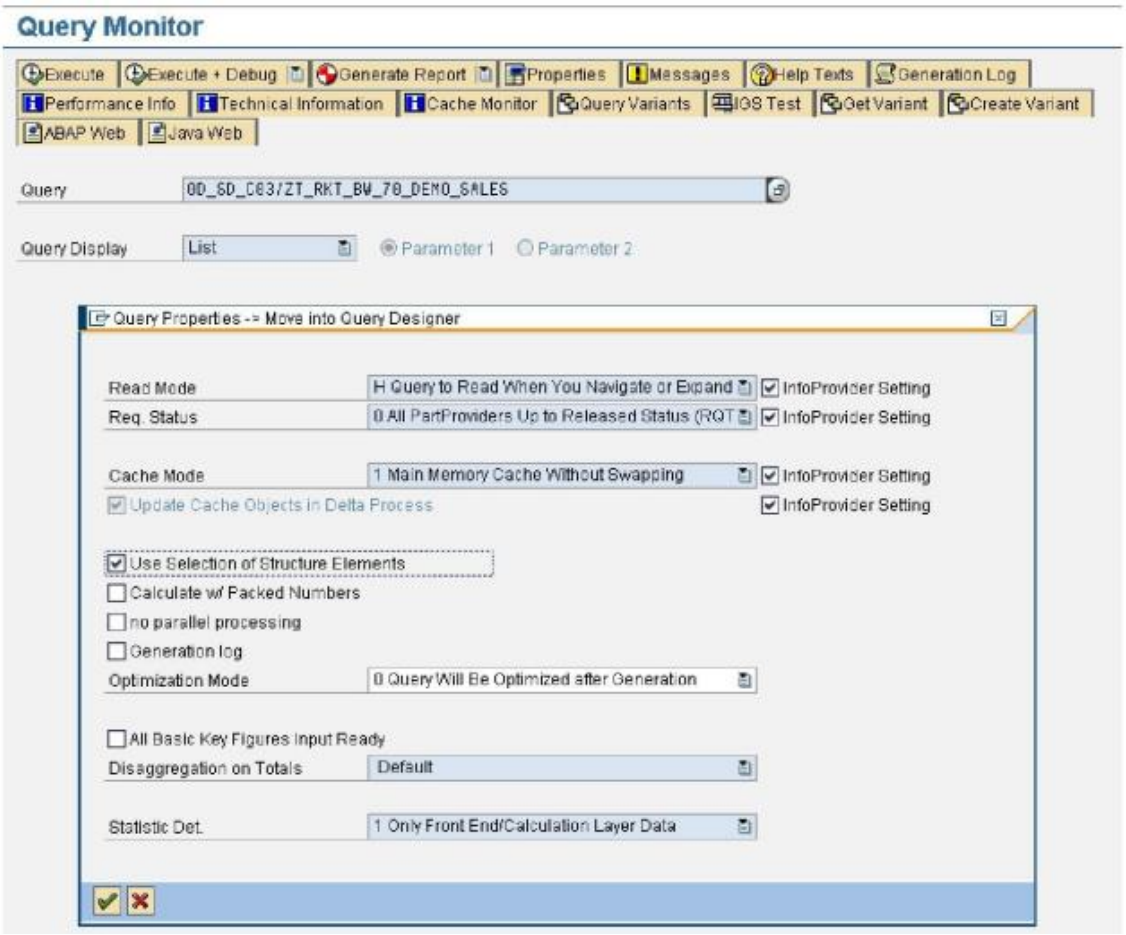

6 Best Practices for Scheduling and Publishing

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 do not have any human «think time» at all. For the number of active concurrent users, specify the number of concurrent jobs that may be running.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

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.

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 personalization that requires 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 SAP BI audit database. You can download the reports and learn more here.

6.1 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.

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.

6.2 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.

6.3 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 SAP BI 4 Administration Guide.

7 Sizing Example

This section describes how to apply the sizing methodology specified above to real customer requirements. It does not demonstrate every aspect of sizing. For your sizing, be sure to follow all the steps outlined in the previous sections.

In the customer scenario, all user counts are considered active, concurrent users. These users do not use the SAP BI Launch Pad Search heavily.

Web Intelligence

- 150 Business users, 10 expert users
- Data source: SAP BW

Crystal Reports 2020

- 80 Business users, 20 expert users
- Data source: Relational Database

Analysis, edition for OLAP

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

- 15 expert users
- Data source: Relational Database

The system is used internally only, so the web server doesn't need extra network security required for exposure on the Internet.

7.1 Quick Sizer Setup

Enter the user numbers into the SAP BI 4 Quick-Sizing (§2.3.1, *SAP Quick Sizer Tool*,) as shown below. Note that the figures in the screenshot below are based on a sample calculation. Use the Quick Sizer tool to get actual resource requirements.

Input Screen

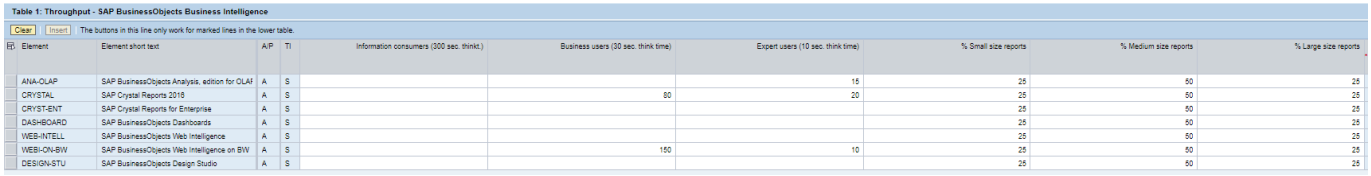

Output Screen

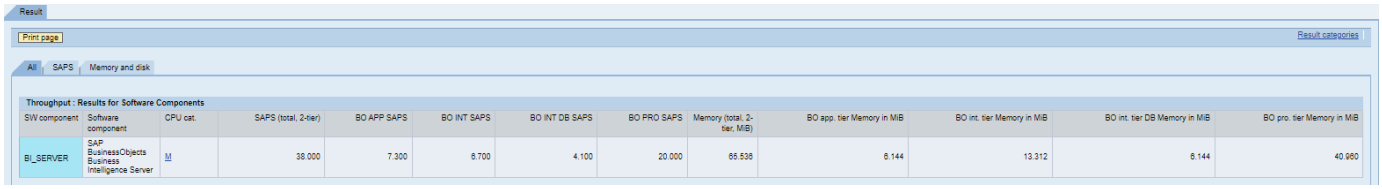

7.2 Intelligence Tier Database

The Intelligence database (SAP BI Platform repository database) must have processing headroom of 4 100 SAPS and 6.1 GB of memory available to process CMS queries. Add these values to the sizing worksheet.

7.3 Intelligence Tier

The CMS itself requires 6 700 SAPS and 13.3 GB of memory to function properly. Add these values to the sizing worksheet.

7.4 Application Tier

The web server is the main portion of the processing that occurs on the Application Tier. In this deployment, it needs 7 300 SAPS and 6.1 GB of memory. Add these values to the sizing worksheet.

7.5 Processing Tier

This tier includes all the BI tools. The internal splits of the SAPS/Memory among the individual BI tools can be obtained by changing the results type to «Line results + inputs» (as seen in the screenshot below).

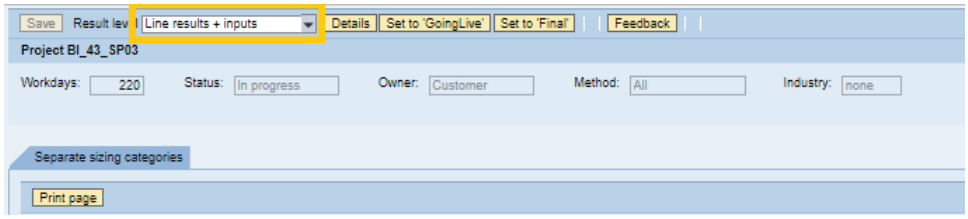

Analysis, Edition for OLAP

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

To determine how many processing and memory resources are needed for the deployment of Analysis OLAP, put only those numbers into the Quick Sizer and record the Processing Tier numbers. Since Analysis OLAP queries directly the data connectivity, no special sizing is necessary.

Crystal Reports 2020

To determine how many processing and memory resources are needed to support the deployment's Crystal Reports 2020 needs, put only those numbers into the Quick Sizer and record the Processing Tier numbers. Since Crystal Reports is used straight forward data connectivity, no special sizing is necessary. The memory required for Crystal Reports 2020 is calculated according to the processing needed.

Web Intelligence

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 SAP BW data: use the **SAP BusinessObjects Web Intelligence on BW** section

To determine how much processing and memory resources are needed to support the deployment's Web Intelligence needs, put only 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.

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. 1 000 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 1 000 SAPS of processing power. You can refer to the SAP Note 2387494 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 SAP BI 4 Administration 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:

	SAPS	Memory (GB)	Cores
Intelligence Database	4 100	6.1	2
Intelligence Tier (CMS)	6 700	13.3	3
Application Tier (Tomcat)	7 300	6.1	3
Analysis OLAP	3 528	9.5	2
Crystal Reports 2020	4 508	19.6	2
Web Intelligence on BW Processing Tier	9 132	10.8	4
Promotion Management Services	1 000	0.75	0.5
Platform Search	1 000	0.5	0.5
OS	2 000	2	1
Total	**39 268**	**68.6**	**18**

Machine SAPS	25 120
Number of cores	12
SAPS/Core	2 093

Sizing Analysis and Scale-out

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

For this deployment, the customer uses 12-core machines with 2 093 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, considering 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 hosts the CMS services and 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 hosts 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

Machine #1: CMS SAPS Memory Cores	SAPS	Memory	Cores
OS and Primary APS	2 000	2	1
Intelligence DB	4 100	6.1	2
Intelligence Tier (CMS)	6 700	13.3	3
Promotion Management Services	1 000	0.75	0.5
Total	**1 3800**	**22.15**	**6.5**

Machine #2: Web Application	SAPS	Memory	Cores
OS and Primary APS	2 000	2	1
Application Tier (Web Server)	7 300	6.1	3
Platform Search	1 000	0.5	0.5
Total	**10 300**	**8.6**	**4.5**

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

Machine #3: AOLAP, Crystal Reports	SAPS	Memory	Cores
OS and Primary APS	2 000	2	1
Analysis, Edition for OLAP	3 528	9.5	2
Crystal Reports 2020	4 508	19.6	3
Total	**10 036**	**31.1**	**6**

The Web Intelligence sizing for this customer requires 4 cores for the processing services. With the core used for the Operating System, these 5 cores are less than the 12 cores of a single machine, so two machines will not be specified for Web Intelligence processing.

Machine #4: Web Intelligence	SAPS	Memory	Cores

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.

OS and Primary APS	2 000	2	1
Web Intelligence	9 132	10.8	4
Total	**11 132**	**12.8**	**5**

In general, spreading the processing and memory for the Web Intelligence services (processing server, DSL Bridge and Data Visualization) evenly across the two machines provides processing headroom for the future.

APS Configuration

APSes need to be defined and allocated on four of the five machines, as follows (§3.8, *APS Configuration*):

- Machine #1: CMS
  - Primary APS
  - Monitoring APS
  - Auditing APS
  - Promotion Management APS
- Machine #3: Analysis, Edition for OLAP, Crystal Reports
  - Primary APS
  - Analysis APS
- Machine #4: Web Intelligence
  - Primary APS
  - WebI APS
  - Visualization APS
  - WebI DSL Bridge APS
  - Connectivity APS

Deployment

When deploying SAP BI 4 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 consider redundancy or virtualization. Those are important topics that would be part of any Enterprise software deployment. For further details on deploying SAP BI 4, see the SAP BI 4 Administration Guide.

© 2023 SAP SE or an SAP affiliate company. All rights reserved. See Legal Notice on www.sap.com/legal-notice for use terms, disclaimers, disclosures, or restrictions related to SAP Materials for general audiences.