SAP BusinessObjects

SAP BusinessObjects Explorer Sizing Guide

- SAP BusinessObjects Enterprise 4.0 Support Package 03

2012-07-16

SAP

2012-07-16

Contents

About this guide

This document provides an overview of the input and output parameters you need to measure in order to size your SAP BusinessObjects Explorer system.

In this guide SAP HANA and SAP BWA data sources are referenced as Explorer Accelerated version.

1.1 Deployment senarios

Sizing information is provided for the following types of Explorer deployment:

Explorer	Data providers	Components
Explorer 4.0: Provides fast processing of small to large data sets.	• BusinessObjects Universes .UNX (relational) • Excel flat files .	• SAP BusinessObjects Explorer • SAP BusinessObjects Enterprise
Accelerated Version: Provides accelerated processing of very large data sets.	• NetWeaver BWA indexes enabled for Explorer • SAP In Memory Computing Engine (SAP HANA) • Excel flat files	• SAP BusinessObjects Explorer • SAP BusinessObjects Enterprise • SAP NetWeaver Business Warehouse Accelerator (BWA) • SAP In Memory Computing Engine (SAP HANA)

1.2 Definition of "user"

These sizing guidelines are calculated based on the average load represented by active users. Active users are those users who generate any activities on Explorer in a given time period of one hour. When applying these sizing guidelines to your deployment, it is necessary to calculate the number of active users for which you need to size your system.

Note:
Active users are not to be confused with the notion of "named users."

1.3 Definition of SAPS

The target output of the sizing algorithm provided in this document is a hardware vendor independent measure referred to as SAPS (SAP Application Performance Standard). Please refer to your hardware vendor to obtain a suitable hardware configuration for the calculated resource requirements.

Related documentation

The following SAP documentation provides information for SAP BusinessObjects Explorer 4.0:

Information	Documentation	Location
List of known issues and workarounds.	SAP BusinessObjects 4.0 Release Notes	http://service.sap.com/releasenotes
Supported platforms and third party software.	Product Availablity Matrix	SAP Service Marketplace: http://service.sap.com/pam In the "Search" field, type: Explorer 4.0
List of the new features introduced with the latest release.	What's New in SAP BusinessObjects 4.0	SAP Help Portal: http://help.sap.com
Architecture and technical landscape of the SAP BusinessObjects Business Intelligence platform 4.0, as well as links to required documentation and SAP notes.	• *Business Intelligence Platform Installation Guide* • *Business Intelligence Platform Administrator Guide* • *Business Intelligence Platform Web Application Deployment Guide* • *Business Intelligence Platform Upgrade Guide*	
Error messages explained	SAP BusinessObjects 4.0 Error Message Guide	
Explorer installation tasks	*SAP BusinessObjects Explorer Installation Guide*	
Explorer server administration tasks	*SAP BusinessObjects Explorer Administrator's Guide*	
End-user information on creating, managing and exploring data using the Explorer application interface.	*SAP BusinessObjects Explorer Online Help* PDF	
	SAP BusinessObjects Explorer Online Help Online Help	Log into the application then click **Help**.

Architecture overview

3.1 Explorer architecture overview

The architecture of Explorer is structured into the following layers:

Layer	Description
Clients	SAP BusinessObjects Explorer client
Web tier/gateway	Web servers and Web Application Servers
Back end	Explorer servers and the BusinessObjects Enterprise servers

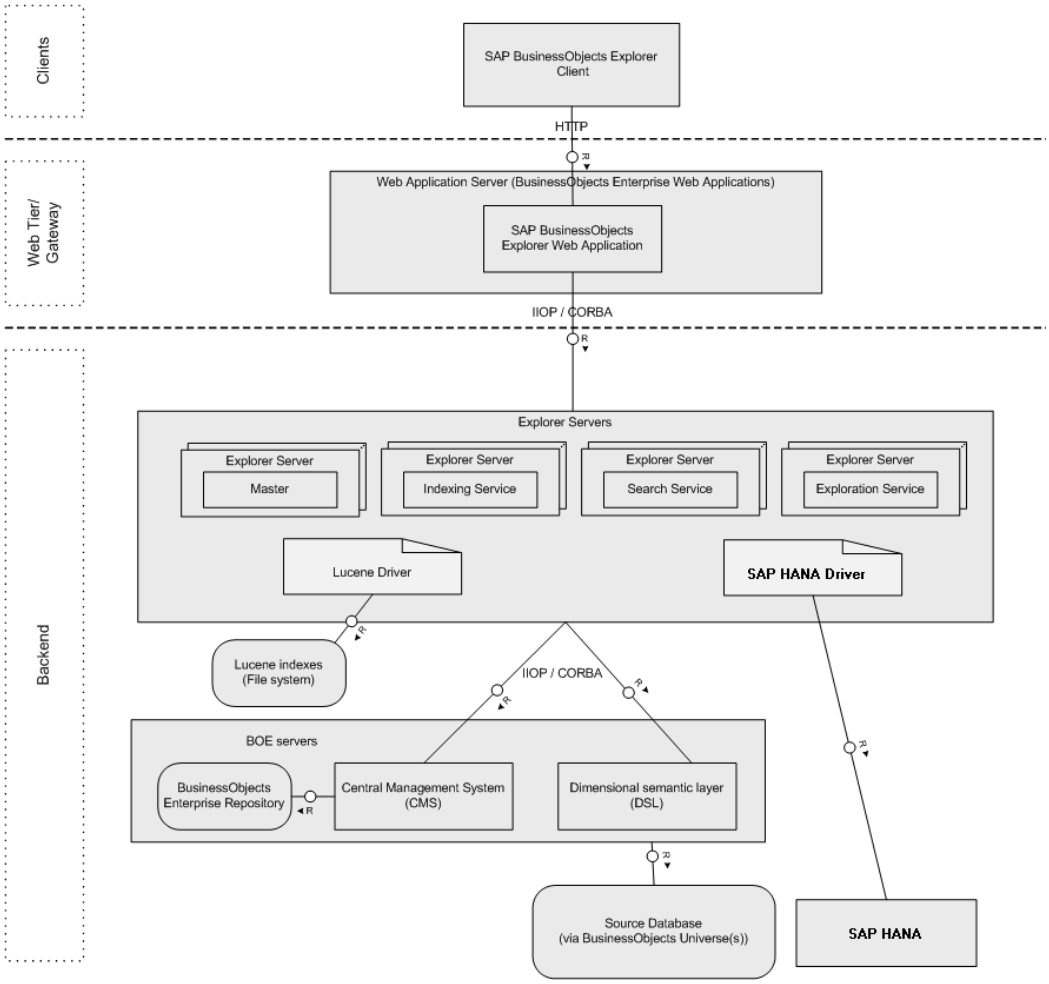

3.2 Input parameters affecting client, web tier, and back end layers

The main capabilities of Explorer that impact performance on each layer of the Explorer architecture are:

Deployment layer	Parameters affecting performance
Client	• The size of the Information Space(s) opened by the end user • The number of exploration actions (clicks) that the end user performs within the same user session increases the amount of memory consumed by the browser.

Deployment layer	Parameters affecting performance
Web tier and gateway	The number of active user sessions and the intensity of user activity per active session impacts the load on the Explorer Web Application Server.
back end	• Using Search on the Home tab - stresses the Search services • Filtering and drill analysis - stresses the Exploration services • Indexing - stresses the Indexing servers

3.3 Other factors that impact performance

Other factors that impact performance are:

- The speed and number of dedicated CPUs,or CPU cores.
- The speed of the storage system. This influences the performance of the Search and Exploration servers when SAP BusinessObjects Explorer is deployed just with SAP BusinessObjects Enterprise.
- The speed of the network connection between the following:
  - The component layers of your Explorer system
  - Your Explorer system and your BusinessObjects Enterprise system.

During performance and scalability tests it was observed that a machine at 100% CPU had the following network consummation on 1 GB network interface:

Component	Peak %	Minimum %
Gateway	25	10
Explorer/BusinessObjects Enterprise	3	1

3.4 Minimum hardware requirements

The Product Availability Matrix (PAM) contains information on client requirements for BusinessObjects Enterprise deployments.

The minimum and recomended hardware requirements for a BusinessObjects Explorer deployment are as follows:

Table 3-4: Hardware requirements for services

Component		Minimum requirement	Recomended requirement
Explorer service	CPU/Core Server	1 CPU/Core	2 CPU/Core
	Physical memory	2 GB	4 GB
	Disk size	70 GB fast disk storage	140 GB fast disk storage

Table 3-5: Hardware requirements for servers

Component		Minimum requirement	Recomended requirement
Explorer Server	CPU/Core Server	4 CPU/Core. By default, Explorer has four services. (4890 SAPS level).	8 CPU/Core (11240 SAPS level).
	Physical memory	8 GB	16 GB
	Disk size	70 GB fast disk storage	250 GB fast disk storage
	Network interface speed	1 Gbps	10 Gbps

Table 3-6: Hardware requirements for clients

Component		Minimum requirement	Recomended requirement
Explorer Client	CPU/Core Server	1 CPU/Core	2 CPU/Core
	Physical memory	1 GB RAM free.	2 GB RAM free.
	Network interface speed	1 Gbps	10 Gbps

3.5 Minimum required services

The minimum required services for an Explorer deployment are the following:

Functional domain	Minimum deployment
Web server/gateway	1 instance
BusinessObjects Enterprise	1 instance of: • CMS • APS • Input and output FRS • Connection server • Connection server 32 • WebIntelligence server
Explorer	1 instance of Master, Indexing, Search, and Exploration servers.

3.6 Functional domains

The sizing information provided in this document is organized by three functional domains that represent the three most common usage types, or type of interaction performed by users:

Functional domains	Usage
Exploration	Navigate, filter and drill values in an Explorer Information Space.
Index generation	Three types of index generation can occur: • Generate an index for a new or reconfigured Information Space. In the case of Explorer 4.0 the third-party software component Lucene generates data and metadata.
Search	Two types of search action can occur: • On the Explorer Home page, search across Information Spaces for the Information Spaces that include matches to the search string. • Search within an Explorer Information Space for a subset of data.

Exploration sizing

4.1 Exploration sizing

This section provides information to help you size your Explorer system so that acceptable performance is achieved when users perform exploration actions such as navigation, filter and drill, while interacting with Explorer Information Spaces.

4.2 Sizing parameters for exploration

The sizing guideline below shows the number of SAPS and memory needed for small, medium, and large sized information spaces for both gateway and Explorer. Explorer includes Master, Index, Search, and Exploration Services.

Information spaces sizes are defined as follows:

- Small: Lucene Information spaces (small = 100.000 cells)
- Medium: Lucene Information spaces (medium = 1.000.000 cells)
- Large: Lucene Information spaces (large = 10.000.000 cells)
- SAP HANA information spaces (large 10.000.000 cells)

For larger volume information spaces, it is recommended to add supplementary Exploration server instances. For example, the table Large considers 1 Exploration server for every 100 active users.

Contact your hardware vendor or SAP if you plan for more active users.

Table 4-1: Small information space sizing guideline

Number of active users	SAPS		Memory (GB)				
	Gate way	Explorer	Gate way	Explo ration	Mas ter	Search	Index
50	600	1400	2	2	2	2	2
100	1200	2800	2	2	2	2	2

Number of active users	SAPS		Memory (GB)				
200	2400	5600	2	4	2	2	2
300	3600	8400	2	4	2	2	2

Note:
For exploration, no disk or memory sizing is necessary for the Explorer component.

Table 4-2: Medium information space sizing guideline

Number of active users	SAPS		Memory (GB)				
	Gate way	Explor er	Gate way	Explo ration	Master	Search	Index
50	600	3100	2	2	2	2	2
100	1200	6200	2	2	2	2	2
200	2400	12400	2	4	2	2	2
300	3600	18600	2	4	2	2	2

Table 4-3: Large information space sizing guideline

Number of active users	SAPS		Memory (GB)				
	Gate way	Explor er	Gate way	Explo ration	Mas ter	Search	Index
50	600	9550	2	4	2	2	4
100	1200	19100	2	4	2	2	4
200	2400	38200	2	8	2	2	4
300	3600	57300	2	12	2	2	4

Table 4-4: SAP HANA Large information space sizing guideline

Number of active users	SAPS			Memory (GB)					
	Gateway	Explorer	HANA	Gateway	Exploration	Master	Search	Index	HANA
50	600	1700	3100	2	2	2	2	4	See SAP HANA Sizing guide
100	1200	3400	6200	2	2	2	2	4	
200	2400	6800	12400	2	4	2	2	4	
300	3600	10200	18600	2	4	2	2	4	

Note:
The numbers are based on the assumptions described previously and are subject to vary when the assumptions are different. HANA's resource consumption could have a specific impact depending on the structure of the cube on which an information space is built. Refer to the HANA Sizing Guide for the HANA specific information.

4.3 Exploration input types

There are three exploration input (user-interaction) types performed when users have a session open in Explorer:

- Open: open an Information Space
- Search: search inside an already opened Information Space for a specific sub-set of data.
- Navigate: navigate the data in an Information Space. This action is also described as "drill" or "filter".

4.4 Load per component during exploration

This section describes the load for each functional domain by component.

4.4.1 Open

Load for the Open action during exploration:

Component	Load
Explorer	• If the user uses Search to retrieve and open an Information Space, this launches a Search command and puts load on the Search services. • If the user clicks on a bookmark to a specific exploration view of an Information Space, this launches multiple filter commands to retrieve that specific exploration view, which puts load on the Exploration services. • Retrieving the data to the Information Space launches the command to connect to the specified data provider, for example the specific BusinessObjects universe or BWA index, and puts load on the Exploration services. • Generating the chart launches the visualization command and the best chart recommendation command and puts load on the Exploration services.
SAP HANA	For Explorer Accelerated Version, it is the initial exploration of an Information Space that causes the complete dataset of an InfoCube to be scanned. Scanning the complete dataset is the most expensive type of navigation.

4.4.2 Search

Load for the Search action during exploration:

Component	Load
Explorer	Searching values within an Information Space puts load on the Explorer Exploration services. **Note:** Using search on the Home page puts load on the Explorer Search services.
SAP HANA	For Explorer Accelerated Version, searching on the Home page or in an Information Space puts load on SAP HANA.

4.4.3 Navigate

Load for the Navigate action during exploration:

Component	Load
Explorer	Selecting a value from one of the facets of the Information Space as a filter value or to drill down into the next level of details puts load on the Explorer Exploration services. **Note:** Although there are other types of navigation actions, for example export of data, or switching to table views, we have restricted our sizing considerations to these most common navigation steps.
SAP HANA	Naviation in accelerated version has the same cost as navigation in a small Lucene information space. Refer to the SAP HANA documentation for sizing and performance information Hana Docupeidia

4.5 Exploration user profiles

We have defined three usage profiles based on the typical usage patterns of users performing exploration actions. Each profile is associated with specific load patterns, in order to help customers size their CPU:

User profile	Load patterns during exploration
Information consumer	The Information Consumer is a business user who uses Explorer occasionally to obtain very specific information via a well defined search path. These users access a limited number of Information Spaces via bookmarks to specific exploration views. Their business need is to view the updated values on familiar data sets. The number of navigations (clicks) for this type of user is low; we assume an average total of 10 clicks per hour. The predominant type of navigation performed by this user is adding filters and drilling down on values. (On a typical user base, 70% active users fit this profile.)

User profile	Load patterns during exploration
Executive user	The Executive User is a business user who uses Explorer frequently to perform basic analysis. These users access a small number of distinct Information Spaces via Search. Their business need is to obtain a personalized view of the overall data set. For this reason, these users perform significantly more navigations per hour than the Information Consumer; we assume an average total of 150 clicks per hour. The predominant type of navigation action performed by this user is searching for values, applying filters and performing drill analysis. (On a typical user base, 25% active users fit this profile).
Expert user	The Expert User is a business user who uses Explorer daily to perform advanced analysis. These users access a high number of distinct Information Spaces via Search. Their business need is to provide deep analytical insight into a wide range of corporate data and to share their analyses with business collaborators. The number of navigations (clicks) for this type of user is high; we assume an average total of 300 clicks per hour. The predominant type of navigation action performed by this user is searching for Information Spaces and values, applying multiple filters and performing extensive drill analysis. (On a typical user base, 5% active users fit this profile).

Note:
These profiles are provided as an example and may differ from the specific usage patterns across your deployment.

4.6 User interaction and activity level by user profile

The following matrix is used to define the share of each of the interaction types, along with the click frequencies, for each of the user profiles:

User profile	Click frequency (clicks/h)	Exploration and Search	Open Information Space
Information consumer	10	70	10
Executive user	150	50	20
Expert user	300	30	30

Experience with standard Explorer users has shown that an average distribution of business users according to the above user scenarios is approximately 70% : 25% : 5% (Information Consumers : Executive Users : Expert Users).

Note:
The distribution of users into user scenarios as well as the distribution of navigation step types within the user scenarios can be changed at any time, should more specific information from customers be available.

4.7 Sizing example: SAPS by information space size

It is assumed that there are 100 active users interacting with Information Spaces of variable size and complexity, based on BusinessObjects universes (Lucene) and BWA indexes. There is an average think time of 30 seconds between each navigation action. The performance requirement is a response time of less than 5 seconds per click.

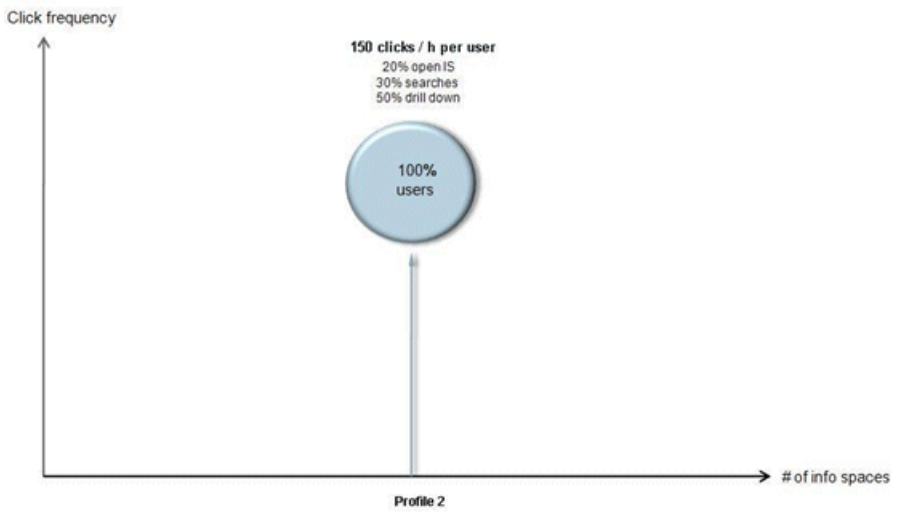

This chart describes the SAPS necessary for the Explorer backend, Explorer web tier/gateway for the following example scenario: 150 clicks / h per user : 20% open Information Spaces; 30% search in Information Spaces; 50% drill down.

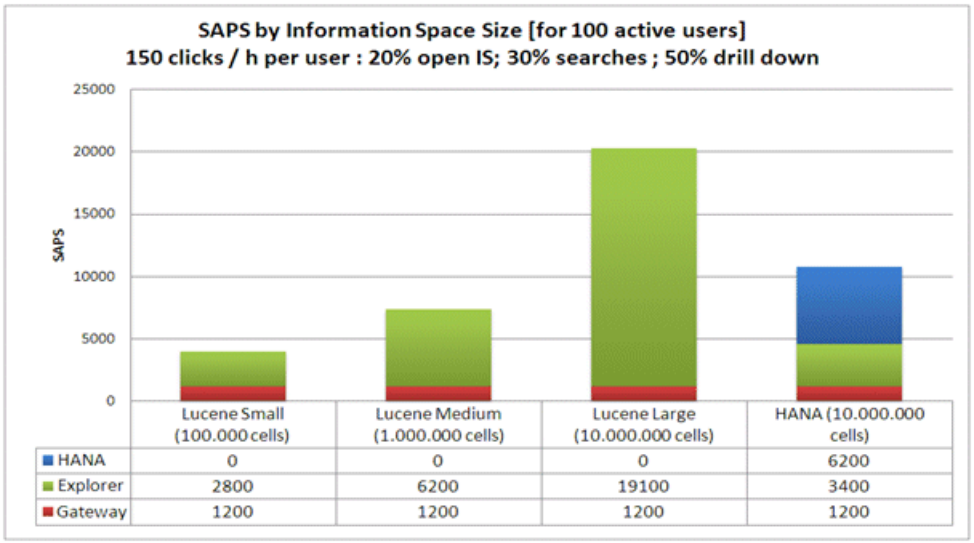

4.8 Sizing example: SAPS by user profile

It is assumed that there are 100 active users interacting with Information Spaces of variable size and complexity, based on BusinessObjects universes and BWA indexes.

There is an average think time of 30 seconds between each navigation action or click. The performance requirement is a response time of less than 5 seconds per click.

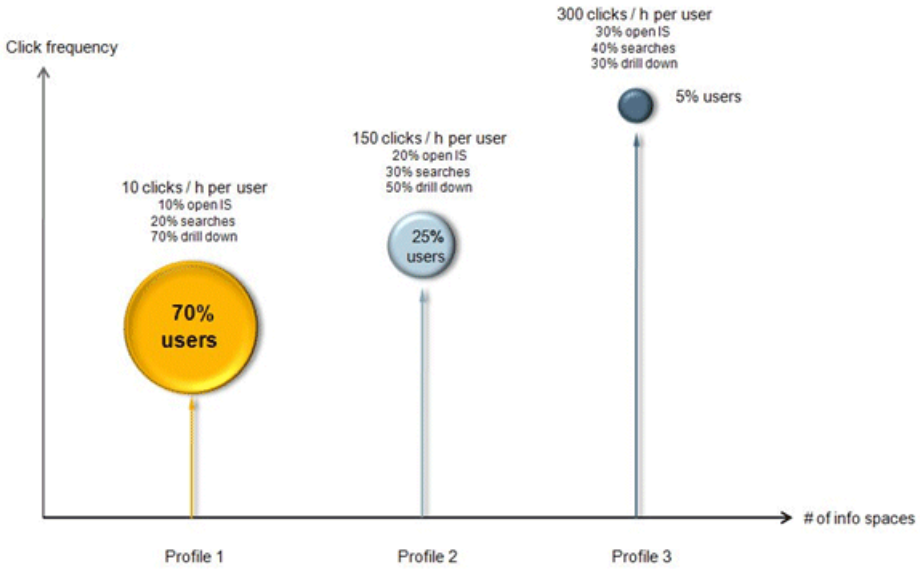

This example is based on the following scenarios:

- Scenario 1: 70% of users are Information Consumers/
- Scenario 2: 25% of users are Executive Users.
- Scenario 3: 5% are Expert Users.

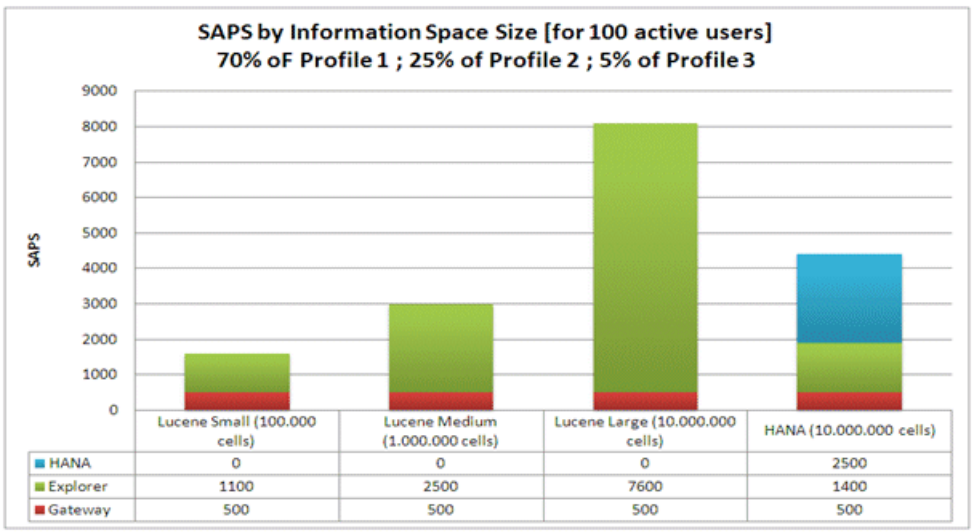

4.9 Sizing example: Exploration of Excel spreadsheets

Uploading an Excel spreadsheet as a data provider for an Explorer Information Space causes an index to be generated before the Information Space can be explored.

The SAPS required for index generation or exploration of an Excel data source stored in the BusinessObjects Enterprise CMC is equal to the SAPS required for universe data sources (for the same data volume). For example: SAPS (Universe 100.000 cells) = SAPS (Excel 100.000 cells) for index generation or exploration.

Note:
The maximum number of uploads that can be done concurrently is 30 by default. If more than 30 uploads are performed concurrently, then the requests are queued. If you wish to increase this value you can either modify the parameter using the CMC (for full information, refer to the SAP BusinessObjects Explorer Administrator's Guide) or add additional Indexing services to your deployment.

Once an index has been generated, it is stored for duration specified in the polestar.service.properties file. Once the defined duration has expired, the bookmark is removed from the disk. The storage duration can be modified by changing the bookmark validity parameter using the Manage > Applications menu in the CMC. For full information, refer to the SAP BusinessObjects Explorer Administrator's Guide.

Example:

Example of SAPS for Excel upload:

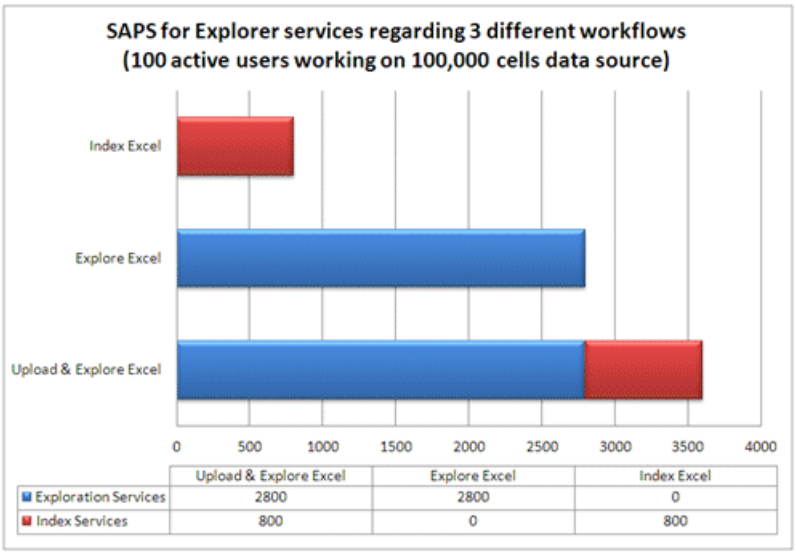

4.10 Sizing of exploration views

When you plan to deploy Exploration Views, it is recommended that you multiply by 2.5 times the SAPS indicated for the Exploration Sizing.

Index generation sizing

5.1 Index generation sizing

This section provides information to help you size your Explorer system so that acceptable performance is achieved when users perform actions to generate indexes. These indexes are consumed by Explorer Information Spaces.

5.2 Sizing guidelines for index generation

The following table gives sizing guidelines for index generation on a universe based data source.

Table 5-1: Guideline: Index generation with Explorer

No.cells	Data vol (MB)	SAPS (per job)		Memory (GB)				Disk (MB/index)	Estimated elapse time *(Minutes)
		Index server	Other Explorer services (sum)	Index server	Exploration	Search	Master		
100,000	2	5	<10	2	2	2	2	42	1
1,000,000	20	25	<10	2	2	2	2	180	2
10,000,000	200	250	<50	4	4	2	2	150	7

* Based on a 4890 SAPS machine.

5.2.1 Index generation input types

Indexes are created by the following components:

Generate an index for a new Information Space:

For this Explorer version	
Explorer 4.0 x	Data and metadata are indexed with the third-party software component Lucene.
Explorer Accelerated version	Data and metadata is indexed using SAP HANA technolgy.

Explorer index generation requests are queued up until an Indexing service is available.

The following input parameters, regarding data complexity, impact sizing for index generation:

- The number of facets (or dimensions, characteristics)
- The number of distinct values in each facet
- The number of measures (or key figures)
- The number of records in the fact table.

Note:
The upload Excel action also initiates an indexation. In this specific case indexations are processed in parallel instead of being queued.

5.2.2 Load on Explorer during index generation

- Load on Explorer:

Sizing domain	Description
CPU sizing	CPU sizing of the Explorer component for indexing stresses the Indexing services, as well as having a smaller impact on the Exploration and Search services during the final phase of indexation when the index is replicated by these two services. Only one Indexing job can be processed per Indexing service at any given time, so Indexing requests are queued up until an Indexing service is available. The complexity of the Information Spaces being indexed also impacts the load here. For example, the number of facets (dimensions or characteristics), the number of hierarchies and their structure, or the number of distinct values .

Sizing domain	Description
Disk sizing	Disk sizing is proportional in relation to the size of the source data being indexed. The following files are generated to disk during index generation: • one principal fact table • as many pre-aggregated fact tables as there are facets • Lucene generates an index of the facet values for each facet in the Information Space • each distinct value in each facet is mapped in memory during the duration of the index generation.

Search sizing

6.1 Search Sizing

This section provides information to help you size your Explorer system so that acceptable performance is achieved when users perform a search on the Explorer Home page.

6.2 Constant and variable sizing for search

The following table describes which index generation sizing parameters are variable.

These parameters require sizing according to:

- Input parameters described in the section Search Input Types
- Fixed index generation sizing parameters which it is necessary to meet the recommended requirements provided by your SAP consultant:

Input parameters	Explorer	SAP HANA
Memory	Variable (minimum 2GB per service)	Variable (depends on data volume)
CPU	Variable (quantity of Information Spaces and data volume)	Variable (quantity of Information Spaces and data volume)
Disk usage	Fixed (no data is created on the disk during search)	Fixed (no data is created on the disk during search)

6.3 Search input types

Search actions performed on the Explorer Home page generate a search across all of the Information Spaces stored on the Explorer system.

Note: In an exploration search step a subset of the data of an Information Space is searched for a given search term. Searching values within an Information Space puts load on the Explorer Exploration services. Information about exploration search is included in the Exploration Sizing section of this document.

6.3.1 Load by component during search

Load by component during search:

Load on	Description
Explorer	When you use Search on the Explorer Home tab to retrieve and open an Information Space, this launches a Search command and puts load on the Search services.
SAP HANA	When SAP HANA is accessed and used by Explorer, the Search on the Explorer Home tab puts load on the Explorer Search services for the searching of meta data of the Explorer Information Spaces, for example name, description, key values, and facets. This search also puts load on the SAP HANA component for the searching of the content that is assigned to the respective meta data, for example key values and facets defined for the Explorer Information Space.

Tuning techniques

7.1 Tuning techniques

The following tuning techniques allow you to optimize performance without modifying your CPU or hardware:

- Tuning the web tier by opting for a specific web application server or by clustering multiple web application servers.
- Adjusting your deployment architecture according to the usage patterns in your deployment, for example by adding Exploration services or Search services, to ensure optimum scalability.
- Increasing virtual memory on the backend servers, by configuring the JVM heap size.
- Increasing the number of concurrent requests possible on the web application server by configuring the number of the Corba threads.

Note:
The guidelines provided here need to be adapted to your specific deployment and performance requirements.

7.1.1 Tuning the web tier gateway

Depending on the load at the Web Tier (or gateway), Explorer performance can be improved by setting various JVM tuning parameters and/or clustering multiple Web Application Servers. Multiple application servers can also extend the scalability of Explorer deployments.

Consult your Web Application Server vendor for more information.

7.1.2 Configure Tomcat memory pool

Set the Apache Tomcat memory pool sizes as follows:

- Initial memory pool: 2048
- Maximum memory pool: 2048

7.1.3 Increase number of threads on Apache Tomcat

It is recommended to set the maximum number of threads higher than the number of concurrent active users. For Tomcat 6, you can define the level in the following file:

```
<Tomcat6_install>Tomcat6\conf\server.xml
```

The following example shows that it sets 500 by editing the “maxThreads=500”. By default the value is 200.

```
<Executor name="tomcatThreadPool" namePrefix="catalina-exec-" maxThreads="500"
minSpareThreads="4" />

<Connector executor="tomcatThreadPool" port="8080" protocol="HTTP/1.1" ……
/>
```

For other Web App Servers refer to the corresponding Administration Guide.

7.2 Tuning the Explorer and web application JVM

BusinessObjects Explorer 4x services run on SAP JVM 64-bit.

To ensure optimum performance, it is recommended you tune some of the settings on the Explorer services.

BusinessObjects Explorer on SAP JVM 64-bit allows treating larger data volumes. It also extends the scalability of a single Explorer service to handle as many concurrent users compared to the multiple instances of Explorer 3.x that were required on 32-bit JVM.

JVM options can be added in the Explorer service command line parameters from the CMC.

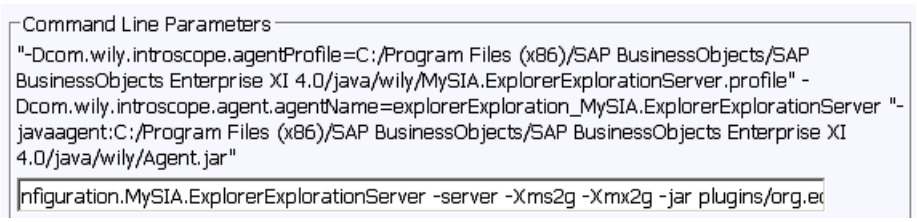

JVM options vary depending on the different JVM vendors. Consult the product documentation for your JVM vendor.

7.2.1 Increasing Java virtual memory on the Explorer services

The amount of virtual memory required by the Explorer services depends on the size of the Information Spaces being explored and indexed across your deployment. You can increase the amount of virtual memory available on each server by changing the JVM heap size value as necessary.

It is advisable to set identical values for both minimum and max heap sizes. For example :

Server	Initial memory size	Maximum memory size
Master	-Xms2g	-Xmx2g
Search	-Xms2g	-Xmx2g
Exploration	-Xms4g	-Xmx4g
Index	-Xms4g	-Xmx4g

7.2.2 Tuning JVM options

There are numerous JVM options available to optimize your Explorer services and Tomcat memory usage. Consult the list provided by the JVM vendors. Some typical tuning options are the following:

- -XX:+UseConcMarkSweepGC
- -Xmn<size>
- -Xss<size>
- -XX:+UseParNewGC
- -XX:ParallelGCThreads=<value>
- -XX:+AggressiveHeap
- -XX:+UseParallelGC
- -XX:ParallelGCThreads=<value>
- -XX:SurvivorRatio=<value>
- -XX:MaxTenuringThreshold=<value>
- -XX:+UseParallelOldGC

7.2.3 Understanding heap activities by GClog

Understanding GClog usage is one of the most effecient ways to analyze the best JVP options.

Below is a typical way of the garbage collection log (gclog) for SAP JVM. You can trace gclog by adding the following parameters along with the JVM Options. Consult JVM vendors for syntax.

-Xloggc:<file> -XX:+PrintGCDetails -XX:+PrintGCTimeStamps

Where <file> is the location and filename of the log.

Example : -Xloggc:C:/gclogs/explorerExplorationServer.log

There are a number of tools available to analyze GClog, for example GC Viewer.

Note:
Intense gclog tracing may have an impact on performance in a production environment.

7.2.4 Monitoring and trouble shooting tools

You can use the following methods to monitor and analyze Java usage:

Technique	Description
SAP JVM Profiler	SAP JVM Profiler is a Eclipse plug-in that allows profiling of your Explorer implementation. It shows detailed SAP JVM memory consumption and helps identify root cases to possible response time degradation issues. As Explorer runs on SAP JVM, there are no supplementary parameters on Explorer services. SAP JVM Profiler is available for download from SAP Customer support .
Wily Introscope	BusinessObjects Explorer 4x is facilitated to deploy Wily's Introscope. It also traces and profiles on the heap activities and the response times. It integrates the Java services, and C++ programs such as the BusinessObjects Enterprise back-end, allowing a consolidated view for monitoring the entire BusinessObjects Enterprise and Explorer deployment. Refer to the BusinessObjects Administration Guide.

7.2.5 Physical memory and JVM heap size

When swapping memory to hard disk, the JVM heap size value you define should always be lower than the amount of physical memory available on the server. Having a low amount of physical memory and configuring large values for the heap size of each server results in the swapping of memory to the hard disk. For example, if there is 2 GB of RAM, it is not efficient to provide a heap size of 1024 MB for each Explorer server. SAP BusinessObjects Explorer functions correctly but memory swapping occurs, therefore having an impact on performance.

Best practices

8.1 Deploying multiple exploration servers for improved Information Space exploration

If the main activity of your user population is exploration, then it is recommended you deploy SAP BusinessObjects Explorer in a cluster with additional Explorer servers to ensure maximum performance when users navigate Information Spaces.

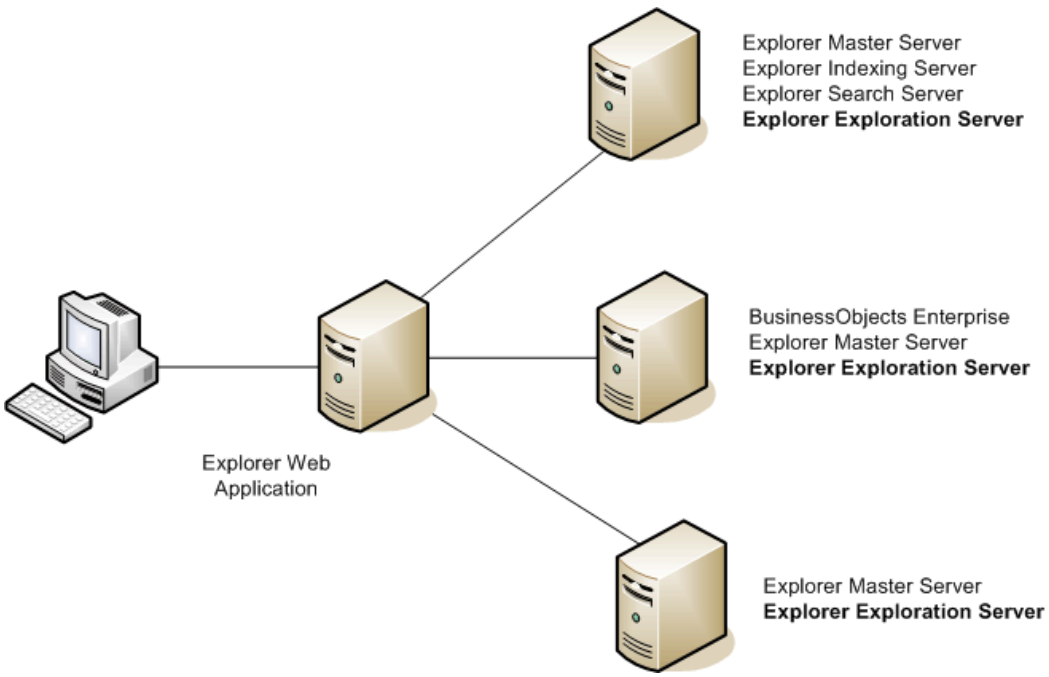

Deploying a high-end machine to the cluster improves the performance and lowers any server constraints.

8.2 Deploying multiple index servers for improved indexing

The indexing of Information Spaces is dependent on the following:

- The number of Explorer services deployed and how they are deployed.
- The hardware (CPU, memory, hard disk) used for Explorer hosts.
- The Java Virtual Machine heap.

If your aim is to improve indexing performance, it is recommended you put one installation of all four Explorer services (Master, Indexing, Search, and Exploration) on the host where BusinessObjects Enterprise is installed, and additional Explorer Indexing Servers on separate hosts, ensuring they are directed to the BusinessObjects Enterprise installation. The indexing load is shared across all the indexing services.

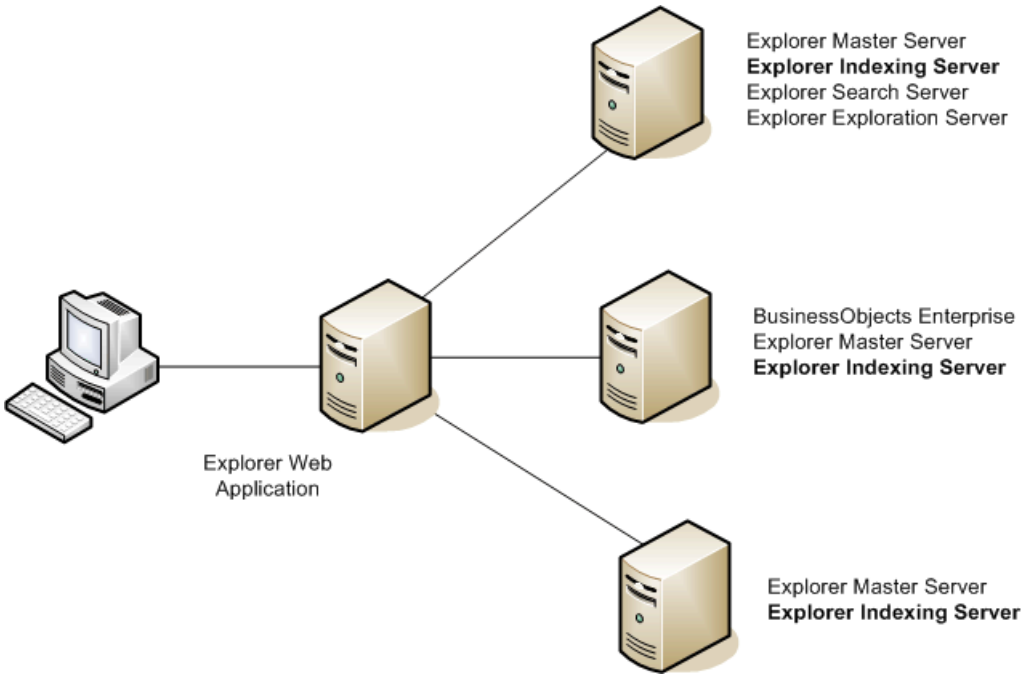

The number of servers required is dependent on the number of users expected to use Explorer. For example, if you expect a high number of users indexing the large Information Spaces at the same time (an extreme scenario), then an additional host is required. Indexing many Information Spaces has an impact on explorers while they are exploring. It is recommended you schedule Information Spaces for indexing when there is less activity, such as over night.

Note:
It is good practice to install an Explorer Search service and an Explorer Exploration service on the additional hosts. This allows those machines to share the load of searching and exploration when there is no indexing in progress.

8.3 Best practices

Before creating Information Spaces, gather the information requirements of your end users by asking the following questions:

- What exactly is the business need of the Information Space?

If you know what the Information Space is going to be used for, then you can simply identify the related data source objects. For example, the business need is for knowing the sales revenue last year for all of your European stores. You could select the Sales Revenue measure, the Country, City, and Store dimensions, and finally, the Last Year filter.

- How many users are expected to access and explore the Information Space?

If you know that the Information Space is for several users, select only necessary objects. If you select too many objects that can have little use for the user, exploration and indexing can be impacted. It can also cause confusion to users.

- What are the security expectations?

Ensure that you select objects that are only meant to be in the Information Space.

- Is a single Information Space the best option?

Several small Information Spaces can often be better than a single Information Space.

- What is the best data provider to use?

Depending on the business need and user demand, choose a source data system and data provider that is the most efficient and most accurate.

- What is the context of the Information Space?

While choosing your data source objects, ensure that you know if any contexts are required. A context makes certain that the Information Space represents the desired perspective. For example: Sales or Reservations.

- If my Information Space is created on a BusinessObjects universe, what filters can be applied so that only data of interest is retrieved?

By using filters, only the data necessary for a specific information need is included into the Information Space. For example, by including a filter called "Last Year,", only data from the previous year is retrieved into the Information Space when users explore it.

Note:
Filters are created at the data provider level when the BusinessObjects universe or BWA index is designed.

- Is the definition you want valid?

Validate the definition of your Information Space before indexing, by clicking the Validate button when you have selected the objects and filters you want to include.

8.4 Personalization

Personalization enables power users or administrators to create a "Reference" Information Space that maps the Explorer user login per end user to the values they are allowed to view for a specific object or facet. For example, regional managers can only view results for the cost centers in their area of responsibility.

The Information Space that contains the business data, for example, the cost center information, can be associated with the Reference Information Space, so that when users open the business data Information Space, they only see the values they are authorized to view.

8.5 Scheduled indexing

Performance during indexing is dependent upon:

- Information space size data volume: number of dimensions, number of distinct values in a dimension.
- server hardware resource.

If users only access SAP BusinessObjects Explorer during working hours, you can schedule the indexing over night, so that users are not impacted by indexing

If you have medium sized information spaces and concurrent user access is not expected, then a single high-end machine is considered to be efficient. However, if you have many users indexing and exploring large information spaces constantly, ensure the following:

- Explorer is deployed in a cluster with additional machines each having extra Explorer services.
- Sufficient system resource on each machine.
- Fast disk on each host is advantageous.

8.6 Auditing services

Auditing is enabled by default. You may wish to adjust the audit level or stop it depending on your needs as an intensive audit tracing may degrade performance. For more information on BOE auditing refer to the SAP BusinessObjects Enterprise Administrator's Guide.

8.7 Controlling Platform Search Crawling

BusinessObjects Enterprise 4x offers a platform search feature from BI launch pad. To facilitate this feature, the Platform Search Crawling (or indexing) on the CMS repository contents is by default active. In certain deployment practices, for example when BI launch pad is not in use, you can schedule the crawling at night when there are fewer end user activities. Highly frequent crawling may degrade performance. Refer to SAP BusinessObjects Enterprise Administrator's Guide.

Index