

How to...

Perform Initial Sizing of Enterprise Portal 6.0

Enterprise Portal 6.0



Version 1.0

August 2003

Released for SAP Customers and Partners

Copyright 2003

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

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

Microsoft®, WINDOWS®, NT®, EXCEL®, Word® and SQL Server® are registered trademarks of Microsoft Corporation.

IBM®, DB2®, OS/2®, DB2/6000®, Parallel Sysplex®, MVS/ESA®, RS/6000®, AIX®, S/390®, AS/400®, OS/390®, and OS/400® are registered trademarks of IBM Corporation.

ORACLE® is a registered trademark of ORACLE Corporation, California, USA.

INFORMIX®-OnLine for SAP is a registered trademark of Informix Software Incorporated.

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

HTML, DHTML, XML, XHTML are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Laboratory for Computer Science NE43-358, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139.

JAVA® is a registered trademark of Sun Microsystems, Inc. , 901 San Antonio Road, Palo Alto, CA 94303 USA.

JAVASCRIPT® is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, SAP Logo, mySAP.com, mySAP Marketplace, mySAP Workplace, mySAP.com Business Scenarios, mySAP.com Application Hosting, WebFlow, R/2, R/3, RIVA, ABAP, SAP Business Workflow, SAP EarlyWatch, SAP ArchiveLink, BAPI, SAPPHIRE, Management Cockpit, SEM, are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other products mentioned are trademarks or registered trademarks of their respective companies.

TABLE OF CONTENTS

1	INTRODUCTION.....	2
2	AN OVERVIEW OF SAP ENTERPRISE PORTAL	2
3	SAP ENTERPRISE PORTAL ARCHITECTURE.....	3
4	FACTORS AFFECTING SIZING	5
4.1	PORTAL USERS	5
4.2	PORTAL CONTENT.....	5
4.3	SYSTEM LANDSCAPE.....	6
5	AN INITIAL SIZING APPROACH.....	6
5.1	THINK TIME AND LEVELS OF USER ACTIVITY.....	6
5.2	PORTAL T-SHIRT CPU SIZING	7
6	CONCLUSION.....	7
7	COMMENTS AND FEEDBACK.....	8

1 Introduction

This paper provides an initial CPU sizing approach for SAP Enterprise Portal 6.0. It is based on the results of performance tests that sought to anticipate portal usage in production environments. It aims to translate the data from these tests into an initial assessment of hardware requirements.

Please note that this guide is preliminary. As the Portal platform integrates your backend systems, and draws data from a variety of information sources, its technical landscape can be quite large, complex and diverse. This fact makes a thorough sizing analysis significant in scope. Sizing takes these factors into account, and goes further, to forecast user behavior, business processes, and application components in project implementations, so that customers can plan their hardware budgets. This document reflects a first stage in the sizing analysis of the product, which is still in restricted shipment. There are currently no sizing figures for Knowledge Management, Collaboration or Unification. The publication of a fuller, more detailed treatment is planned for the period leading up to the product's unrestricted shipment.

For more on the technical infrastructure, setup and architecture of the SAP Enterprise Portal, see the road map of product documentation in the *Master Guide* at <http://service.sap.com/ep60> → *Documentation & More* → *Fundamentals*

2 An Overview of SAP Enterprise Portal

SAP Enterprise Portal addresses one of the most pressing problems facing enterprise IT organizations today – how to provide an integrated, single point of access to heterogeneous IT systems. Today's modern IT landscapes typically contain many diverse technologies acquired from different vendors or developed in-house and involve a highly distributed infrastructure based on both open and proprietary protocols and standards. As part of **SAP NetWeaver** the SAP Enterprise Portal unifies enterprise applications, information, and services from both SAP and non-SAP sources into a coherent interface.

With its role-based content and collaboration tools, SAP Enterprise Portal integrates people, and with its knowledge management of unstructured data, and its unification of structured data, it integrates information.

For its users, the portal provides a window onto information both in their organization and outside it, along with the tools to manage this knowledge, to analyze and interrelate it, and to share and collaborate on the basis of it.

The core technology components which bring this breadth and depth of information to the portal desktop make up its content, and the basic unit of this content is the **iView**. An iView is a mini-application which fetches information from any of a variety of sources, such as databases, enterprise applications, file systems, and the Web, returning up-to-the-minute information each time it is launched.

3 SAP Enterprise Portal Architecture

SAP Enterprise Portal is platform-independent, able to operate on both Unix and Windows operating systems. It installs and runs on a component of the **SAP Web Application Server**, the **SAP J2EE** engine, which is a proprietary Java application server based on the Java 2 Enterprise Edition (J2EE™) standard.

SAP Enterprise Portal consists of the Portal platform, Knowledge Management, and Collaboration.

The Portal platform contains:

- **Portal Framework:** Provides components and services for the creation and running of portal content, such as iViews, pages, worksets and roles. It also supplies a persistence layer, the **Portal Content Directory**, which runs against an Oracle or Microsoft SQL database, and user management.
- **Unification:** Enables the integration of multiple information sources of structured data.
- **Connector Framework:** Supplies the pipelines of communication between the portal and backend applications.

Knowledge Management, which you implement to enable unstructured content management, and sophisticated search and retrieval functionality over unstructured content in disparate information repositories, comprises:

- **Content Management (CM):** With CM, you can author, store, manage, and display documents.
- **Retrieval and Classification (TREX):** With TREX, you can do intelligent searches and text-mining of extensive document sets and structure these sets sensibly using document classification.

Collaboration in SAP Enterprise Portal is a set of tools designed to facilitate cooperation and communication that is both synchronous and asynchronous. It lets portal users create virtual rooms in which they can organize themselves into groups, for example, teams or projects. It also allows portal users to communicate through interactive online meetings, with instant messaging and application sharing capabilities. Further, it can integrate the email and scheduling services of Groupware.

The following illustration displays the components of SAP Enterprise Portal:

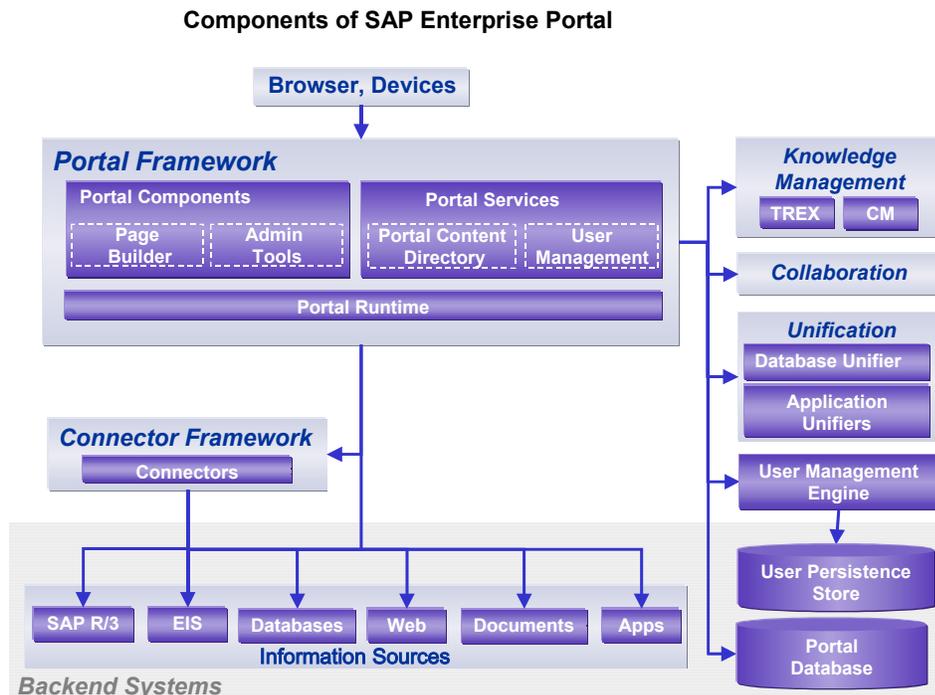


Figure 1: Components of the Enterprise Portal 6.0

4 Factors Affecting Sizing

The sizing of SAP Enterprise Portal is influenced by a wide range of factors. What follows is an overview of these factors.

4.1 Portal Users

The number and activity of users working in the portal at a given time represent key considerations in a sizing analysis.

- **Number of concurrent users and their distribution**

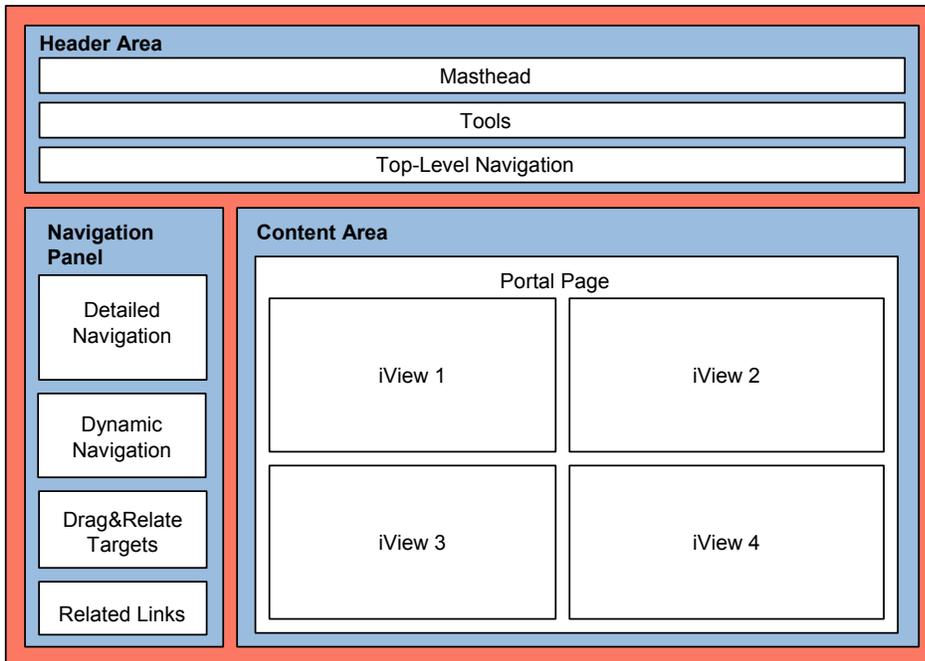
For portal sizing, the total number of possible (or named) users is misleading. For example, if you have 10,000 named users in different time zones, it is very unlikely that all 10,000 named users work in tandem. Therefore, you should only consider the number of active users when sizing. Active users are those who interact with the system in a defined period time and consume system resources.

- **User activity**

An enterprise typically has some users who work with the portal intensively, and others who work with it less intensively. For this reason, we differentiate between various types of users. For the definition of these users see 5.1.

4.2 Portal Content

The content on the Portal desktop can influence sizing numbers. The following image illustrates content on the Portal desktop:



Portal Desktop

Figure 2: Contents of the Portal Desktop

The desktop consists of multiple areas in which content is displayed:

- **Header Area**

The number of tabs in the Top-Level Navigation have an effect on performance, because these are rendered at runtime.

- **Navigation Panel**

The number of nodes in this area has a strong influence on performance, because each of these nodes are rendered at runtime.

- **Content Area:**

- The more iViews displayed in the Content Area, the longer it takes to render the information.
- Whether the portal or a backend system renders iViews can make a difference in performance. For example, the SAP Internet Transaction Server (ITS) renders a WebGUI transaction with almost no load except for parsing the link to SAP ITS and caching the content. On the other hand, the portal's rendering of Java iViews at runtime can potentially cause much more workload depending on the size and complexity of the content.

4.3 System Landscape

The system landscape also impacts sizing, and should be taken into account:

- The workload of a portal can be distributed over different portal servers. Depending on the network, it may be preferable to have a few, large portal servers in a cluster rather than many small servers.
- You need to consider the influence of network traffic between backend systems and the portal server if they are in different geographical locations. Generally, it is recommended to put the portal server in close proximity to the backend systems.

5 An Initial Sizing Approach

5.1 Think Time and Levels of User Activity

It is the task of sizing to find out how to map user activity and system interaction to necessary technical requirements. Therefore we defined pages and think time as metrics. A "page" is defined as containing four iViews. In our approach, all users perform the same actions, but at different levels of intensity, associated by their think time. The think time is that time that elapses between opening two pages.

Using the think time defined for users at three progressive levels of activity enables us to calculate the pages per hour as illustrated in the table below:

User Type	THINK TIME (SEC)	Avg. Pages Per Hour
Low	600	6
Medium	180	20
High	30	120

We assumed a distribution of user activity patterns (i.e, low, medium and high levels of activity). In our distribution, there is a large number of low intensity users, a smaller number of medium intensity users, and a still smaller number of high intensity users in the course of a day; specifically, 60% low, 34% medium, and 6% high.

5.2 Portal T-Shirt CPU Sizing

As the overall resource consumption by SAP Enterprise Portal 6.0 is planned to be equal to or better than that of SAP Enterprise Portal 5.0, we used the Quick Sizer tool (<http://service.sap.com/quicksizing>), which is based on SAP Enterprise Portal 5.0, to obtain the following initial sizing recommendations. The table also illustrates the estimated number of CPUs for various production environments in an exemplary sizing:

Production Environment	Number of Active Users	SAPS ¹ by Quick Sizer	Exemplary Sizing ²		
			Single CPU Machine	Dual CPU Machine	Quad CPU Machine
Extra Small	500	200	1	1	1
Small	1,000	400	1	1	1
Medium	5,000	2,000	4	2	1
Large	10,000	4,000	8	4	2
Extra Large	25,000	10,000	18 (not recommended)	10 (not recommended)	6

In small production environments, the extra CPUs allow a configuration to accommodate heavy usage. For medium and large production environments, additional CPUs can allow fluctuations in usage times. In extremely large production environments, extra CPUs are not needed, because the larger number of processors enables resource availability to accommodate peak times.

Note that these assumptions are continuously verified by ongoing performance tests and are subject to change.

6 Conclusion

Although the figures obtained from the Quick Sizer tool are currently based on SAP Enterprise Portal 5.0, they are an indication of the overall load expected for SAP Enterprise Portal 6.0. This means that SAP Enterprise Portal 6.0:

- Can handle an equal number of or even more users with the same amount of SAPS
- Needs the same amount or even less SAPS for the same number of users

The Quick Sizer tool is intended to support early stages of hardware acquisition, or, in other words, provide you with a first estimation of hardware requirements. The portal sizing algorithms are based on general assumptions regarding usage and click-behavior. We recommend performing a customer-specific sizing analysis. Keep in mind also that sizing is an iterative process that has to be continually verified in the course of an implementation project.

¹ SAPS is a throughput unit that describes the power of a configuration. Details can be found at <http://www.sap.com/benchmark>. 100 SAPS are defined as 2,000 fully business processed order line items per hour in the standard SD application benchmark. This throughput is achieved by processing 6,000 dialog steps (screen changes) and 2,000 postings per hour in the SD benchmark or processing 2,400 SAP transactions. For possible configurations, just check the benchmark results.

² Assuming 600 SAPS per CPU. Note that there is no fixed ratio between SAPS and CPU type.

7 Comments and Feedback

Comments and feedback are welcome. Please send them to Susanne Janssen of the Performance, Data Management & Scalability Group (susanne.janssen@sap.com), SAP AG.