Eigner PLM 5.1

Hardware Sizing Recommendations for Eigner PLM

Copyrights and Trademarks

Copyright © 1995-2003 Agile Software Corporation. All rights reserved.

You shall not create any derivative works of this publication nor shall any part of this publication be copied, reproduced, distributed, published, licensed, sold, stored in a retrieval system or transmitted in any form or by any means: electronic, mechanical, photocopying, or otherwise, without the prior written consent of Agile Software Corporation, 6373 San Ignacio Avenue, San Jose, California 95119-1200 U.S.A.; Telephone 408.284.4000, Facsimile 408.284.4002, or <http://www.agile.com/>.

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to ensure its accuracy, Agile Software Corporation assumes no liability resulting from errors or omissions in this document or from the use of the information contained herein. Agile Software Corporation reserves the right to make changes in the product design without reservation and without notification to its users.

Agile Software is a registered trademark and Agile, Agile Product Collaboration, Agile Product Cost Management, Agile Product Service & Improvement, Agile Program Execution, Agile Product Interchange, AgileMD, and the Agile Logo are trademarks of Agile Software Corporation in the U.S. and/or other countries. Guaranteed Business Results is a service mark of Agile Software Corporation. All other brands or product names are trademarks or registered trademarks of their respective holders.

Java and Solaris are registered trademarks of Sun Corporation.

Microsoft, Microsoft Windows, Microsoft Word, Microsoft Excel, Internet Explorer and SQL Server are registered trademarks of Microsoft Corporation.

Oracle and Oracle8i are registered trademarks of Oracle Corporation.

NOTICE OF RESTRICTED RIGHTS:
The Software is a “commercial item,” as that term is defined at 48 C.F.R. 2.101 (OCT 1995), consisting of “commercial computer software” and “commercial computer software documentation” as such terms are used in 48 C.F.R. 12.212 (SEPT 1995) and when provided to the U. S. Government, is provided (a) for acquisition by or on behalf of civilian agencies, consistent with the policy set forth in 48 C.F.R. 12.212; or (b) for acquisition by or on behalf of units of the Department of Defense, consistent with the policies set forth in 48 C.F.R. 227.7202-1 (JUN 1995) and 227.7202-4 (JUN 1995).

Part Number:

September 20, 2004

CONTENTS

Chapter 1 Introduction 1

Chapter 2 Characteristics for Hardware Sizing 2

Chapter 3 Overall Process Model of Eigner PLM 3

Chapter 4 Sizing Services of Eigner PLM 5

Sizing the Client 5
Hard disks: 5
CPU: 5
Memory: 5
Network: 5
Sizing Database Servers 5
CPU: 6
Oracle Memory: 6
SGA: 6
SQL-Server memory: 7
General 7
Hard disks for Database: 7
RAID and DB: 8
Disk Space for database 8
Network: 9
Sizing Application Servers 12
Hard disks: 12
CPU: 12
Memory: 12
Network: 13
Batchjobs: 13
Sizing Business Services 13
Hard disks: 13
CPU: 13
Memory: 13
Network: 13
Sizing Presentation Services 14
Hard disks: 14
CPU: 14
Memory: 14
Network: 14
Sizing FMS Servers 14
Hard disks: 15

CPU: 15
Memory: 15
Network: 15
General Considerations, License Manager (FELICS) etc... 15
Minimum Requirements for Test and Customizing 16
Configure Oracle SGA 16
Oracle init parameters 18
UNIX Kernel Parameters set for Oracle 19
Shared Memory 20
Semaphores 20
Security limits 20

Chapter 5 Examples 21

Example 1 21
Example 2 24
Example 3: Comparison one or two Server installation 25
A single Eigner PLM Server configuration 25
A double Eigner PLM Server configuration 26

Chapter 1
Introduction

This manual is an extract from Emails written to PSO, pre sales and sales guys. It is intended to help you and sizing an Eigner PLM installation. After reading the document you will know the different programs and services running Eigner PLM. You will find information to optimize the service performance and security. You will have a comprehensive look from the administrator's viewpoint on Eigner PLM.

Often the data volume is underestimated and the customer only sees the needs for the initial use of Eigner PLM. Try to use scalable hardware as in future the customer will have more users on the system and probably the next Eigner PLM version will need more memory.

When sizing Windows machines, please keep in mind that you get most stable installations when splitting functions across several physically distinct server machines. Primarily the PLM server should not be domain controller, e-mail, print and file server in addition. The fastest solution for Eigner PLM is when database, axalant-Server and WebClient are on one machine.

A Windows machine should not run out of physical memory. The server gets sever problems, when this happens. Especially the Oracle database will crash if the instance tries to allocate additional memory and no physical memory is available.

The amount of main memory limits the maximum number of concurrent users. Other parameters like CPU, File I/O and network define performance but main memory especially for Windows limits the number of served users.

Chapter 2

Characteristics for Hardware Sizing

- ❑ Number of concurrent users?
- ❑ How many concurrent users are using the WebClient?
- ❑ Is Workflow in use?
- ❑ How CPU consuming is the PLM installation?
  - Which modules are in use?
  - 2D or 3D CAD, charge from special interfaces?
  - Size of structures and bill of material?
  - Use frequency and utilization profile (fixed work schedule -> peak load)
  - How much customizing will be made?
  - Special features: database replication, PDW (Project Driven Working), ..
- ❑ Which database software will be used?
- ❑ Which Server OS will be used?
- ❑ Which quantity structure of data is processed?
  (Number of items, documents or dump size)
- ❑ Is the database stored on single hard disk or on a storage system? Which RAID level is used? Determine write and read performance. Is the storage system used for different services?
- ❑ Are there different sites, locations? How are the WAN connections?
  - Transfer rate
  - Transfer speed (Router, Frame relay, VPN,..)
  - Connection costs (leased line, dial in line, flat rate)

Chapter 3

Overall Process Model of Eigner PLM

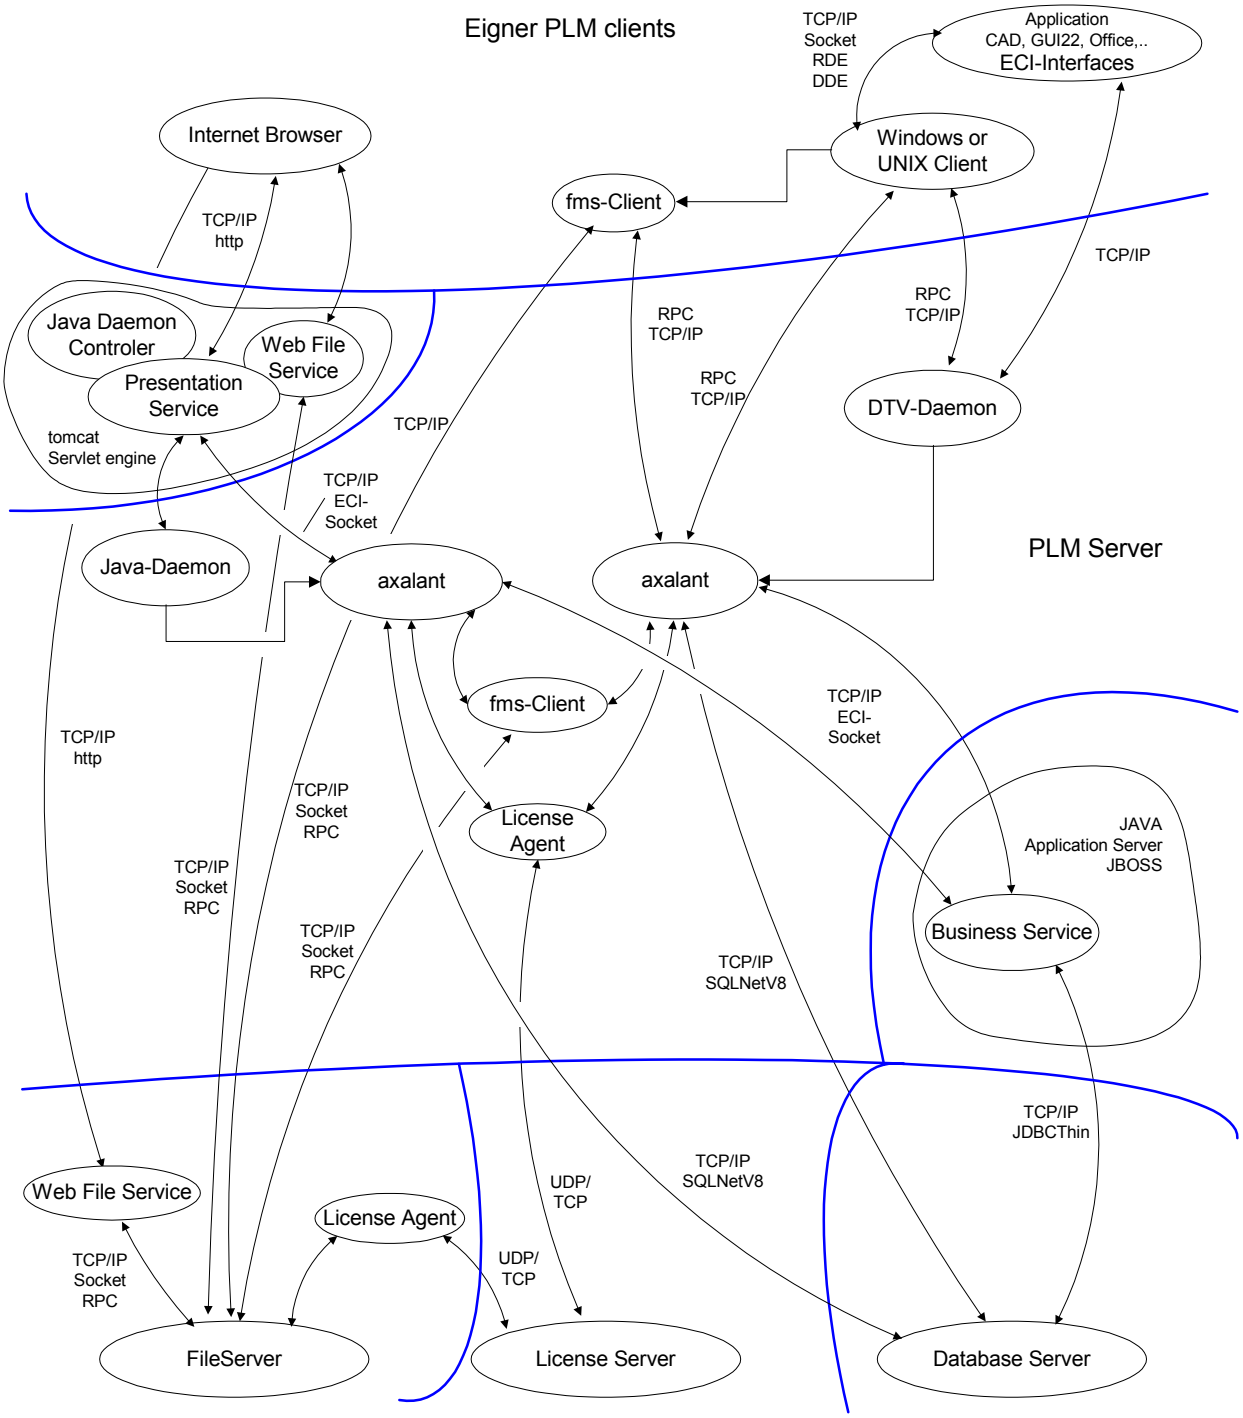

Database Server: Machine running Oracle or SQL-Server; storage of the meta data.

PLM-Server: Machine running the axalant processes. These processes perform the major business logic. One axalant process serves one user.

Presentation Service: The Web Presentation Service gives users access to PLM functionality through their web browsers. It manages web sessions for users and creates the DHTML they use to interact with the PLM Server. Based on a servlet engine (tomcat) it provides a user interface for Eigner PLM. The Frontend (Internet Browser) only show the WebClient, the servlet engine performs the GUI logic.

Business Service: Carry out the business logic for the Workflow, including Watchdog, Permission Manager and Notifier. The Java Application Server (Jboss) provides this service. It is connected via ECI-Socket to the axalant process and on the other side via JDBC to the database.

Web File Service: Provide an http interface for the FileServer. The WebClient (Internet browser) requests the Web Fileservice for a special file. The Web File Service checks out the file from the FileServer and transfer it to the browser.

Java Daemon Controller: Admin tool for the Java Daemon

License Server: Hold the license database, lock and release licenses

License Agent: Requests licenses from the license server for a client process.

FileServer: Store the files onto local disks and release files to the fms-client

Fms-client: Communicate the axalant process and perform tasks managed by the axalant process. Transfer files from local disk to the FileServer and back.

Windows Client: Native Windows Client (DataView Client)

UNIX Client: Native UNIX Client based on qt-libs

DataView Daemon: Starts axalant processes for the Windows and UNIX client

JAVA Daemon: Starts axalant processes for the Presentation Service

Internet Browser: Used as WebClient front end (Iexplorer, Netscape or Mozilla)

Chapter 4

Sizing Services of Eigner PLM

Sizing the Client

Hard disks:

Around 180 MB disk space necessary. The Windows Client often loads bitmaps. A Local installation is recommended to reduce network load and improve performance. You can set the registry string DATAVIEW_CACH to on.

CPU:

No special requirements for the DataView or UNIX Client see *Eigner PLM 5.0 Prerequisites Guide.*

The PLM Web Client needs a fast Client machine (>500 MHz) especially if Netscape or Mozilla is used as Internet Browser.

Memory:

DataView Client: 12-20 MB

Unix Client 18-30 MB

WebClient: See browser requirements (ca. 20 –30 MB)

Network:

DataView or UNIX Clients need an average of about 2,5 up to 4 Kbytes (10 to 32 kBit) network load only for Dataview Clients. That is an 64kBit ISDN Line can support 2 up to 3 or 4 connected EignerPLM Users depending how frequent and intense they are using the system and they are not transferring any additional data e.g. Documents or Drawings.

The WebClient: A 256 kBit connection will seldom by under full load. A 128 kBit connection reacts slower and 64 kbit connection is possible but the performance decrease is obvious.

20 times more information is send to the WebClient (browser) than received back.

Add Network load according to expected documents and drawings etc. shared via WAN.

Existing WAN connections indicate already existing network load!

File transfer uses the total capacity of a network connection. If a network line is under heavy load the elapsed time of the IP-packages is long. This will slow down the connection. File transfers should be fast (< 10 sec) because during this time the connection is slow. The network connection must have enough spared bandwidth.

Sizing Database Servers

Small databases (dump size) are always fast. The more data is in the slower the application performs and the more resources are needed on the DB server. If a special application or module

performs well can only be evaluated with the amount of data used in production. In general, the DB-instance uses less CPU than the axalant server process (30:70 or 40:60).

Small: - 20000 rows in one table (article, documents,..)
Medium: 20000 – 50000 rows
Big: 50000 – 100000 rows
Large: 100000 – 300000 rows

The database performance is defined by four parameters:

1. Number and performance of CPU
2. Main memory used for the database
3. Disk i/o
4. Client – Server connection speed (Network, IPC)

CPU:

One CPU >=1 GHz class should be able to support at least 50 active concurrent users on the database. The database scales very good over multiple CPUs.

Oracle Memory:

For Oracle the memory consumption is determined by the init-parameters which can be defined dynamically in the init-file (init<SID>.ora). Oracle uses shared memory which is called System Global Area (SGA). This memory is allocated at startup time of the database and grows to a certain extent during work. Each user process has its own memory (UGA) and uses the SGA in addition (shared memory).

Unix process listing shows the sum of both values. You see all the Oracle processes with big memory values, but there is only one SGA and all Oracle processes use the same memory.

Because of the low memory prices it is recommended to equip the machines with more physical memory than necessary to increase database performance.

Add at least 50 MB for OS. Unix boxes always use free memory to buffer file I/O. This improves database performance but it is better to use the Oracle buffer cache instead of the system file cache. Oracle redo log information is written with “write through cache” and is not buffered by the system cache.

SGA:

The SGA consist of three mayor parts:

- shared pool
- block buffer
- redo log buffer

The shared pool size depends on the number of users (12 –70 MB).

The block buffer is the memory buffer for the database blocks read and writes to disk. Us as much block buffer as you can. You need at least one tenth of the dump size.

The redo log buffer is used for the redo log information. The buffer must be written to disk if a commit is issued (3*64 kByte – 10MB). Higher values will not speed up the database performance. Big transactions (import, table copy) will benefit from a large block buffer.

- 30 MB for test installations
- 100 MB up to 50-70 users
- 200 – 1000 MB for 100 – 300 users

The dump size and the number of users define the size of the SGA.

Server memory for Oracle 8.1.7:

- 6 MB per connected user
- SGA size
- Add at least 50 MB for Operating system

Each connect (Oracle dedicated server process) allocate 1-2 MB memory. All open cursors (open masks) allocate additional memory to hold the database result (output of SQL-statement). If the user retrieves lots of data with a high limit (1000 rows) the Oracle server process can grow up to 25 MB. So calculate an average of 6 MB per connect.

SQL-Server memory:

The memory consumption of the SQL-Server is defined by control button.

Select the server name in the SQL-Server enterprise manager and select properties. Then you can adjust the memory. You have to reduce this after installation, because SQL-Server will allocate the whole main memory.

Server memory for SQL-Server 2000:

- 30-40 MB for Test installations,
- 80 MB is a minimum for productive use
- add at least 50 MB for Operating system

Use a base size of 30 MB, add 2 MB per connect and at least 1/10 of the dump size.

General

The more memory is allocated for the database the less disk i/o is necessary. Sorts, read and writes are buffered in the database memory. But log-information (redo) cannot be buffered. Write performance on the redo log or the log file must be permanent good.

Hard disks for Database:

A productive database installation need at least two disks. Write intensive parts of the database (rbs, redo log, temp) and system swap or page file have to be on separate disks. The database software vendors recommend using different disks for the database and the operating system to exclude any impact on the database. Each service (file service, swap, ..) which uses disk i/o can affect the database performance.

As I/O is most critical to the database be prepared to use 4 to 6 physically separate disks or an equivalent performing controller base RAID shelf (RAID 0/1) exclusively for the Database. Add a separate disk for OS. Most customers uses RAID 5, which do not provide a good write performance. This is critical for the write active parts of the database (rbs, temp, redo). The rbs and temp can be buffered by the main memory, but redo cannot be buffered. Enough constant write performance must be available. Most OS have a minimum read size of 64 kByte they are

reading from disk. A RAID 5 system reads at least one stripe. Determine the stripe size (mostly 128 kByte) and adapt the database read portion:

```
64kByte <= db_block_size * db_file_multiblock_read_count

db_block_size * db_file_multiblock_read_count <= stripe size
```

Most database uses a read size of 128 kByte.

1. disks: Only test installations
2. disks: 1-30 users
3. disks: 20 – 100 users
4. disks: 70 – 150 users
5. disks: 120 - 200 users
6. disks: 150 – 300 users
7. disks: 200 – 400 users

RAID and DB:

RAID 0 is stripping without parity. This is a fast RAID level used in movie industry for cutters. If one disk fails the data is probably lost. Because of low reliability this level is not used for database files.

RAID 1 is mirroring. Two disks are combined each disk stores the same data. If one disk fails the other is still available. The write performance is nearly that of a single disk. The read performance can reach that of two disks for index read and for sequential read it is identical to a single disk. System swap and redo logs must use different raid sets. Use multiple raid sets for a database server. You can store write intensive parts of the database on RAID 1 and others on RAID 5.

RAID 10 or 0+1 are mirrored strip sets. This is optimal for the database. The RAID 0 level provide high performance and the RAID 1 level the security. This is an expensive RAID level the capacity of the file system is half of the disk size.

RAID 5 the information is distributed over a set of disks. Each stripe contains parity information. If one disk fails the controller can reconstruct the information stored on the disk with the parity. The bad write performance of this level slows down the database but read performance is good. The hardware vendors use large file buffers to compensate the disadvantages of this RAID level. Determine the write performance of the system (test the system, ask for technical details). Estimate how other services, uses the RAID shelf, have an influence on the database performance (other DBs, PLM FileService, File Server, …).

The PLMServer, PresentationService and BusinessService need their software at startup time and have to write log files. These services do not have much file I/O.

Disk Space for database

The database server needs enough free disk space for:

- ❑ Database exports and backups
- ❑ Db log (Oracle redo logs, SQL-Server log file)
- ❑ Copies of the database dump for production, training, development, testing, upgrade,..

- Case of emergency: complete image backup (db files), db recovery

We conclude from the amount of items to the dump size

- 50 000 -> 200 to 500 MB.
- 100000 -> 1GB
- 200000 –300000 -> 2 GB

In the following we calculate with 1GB. The test, training and development dumps are held in a separate database. It is advised to use a second database instance for these environments, maybe on a second server.

The productive Oracle database is run in archive log mode. The backup strategy need a backup of all database files (media backup) and the archive logs produced since the backup. In addition Oracle exports are performed daily (logical backup). The media backup is needed to restore damaged datafiles. The logical backup is needed if a user makes a error and drops or modifies tables or rows. If a complete user dump is imported the database produce archive logs three times the size of the database dump. The archive logs contain all changes on the datafiles. The dump contains object definition and data. The database datafiles need space for the indexes in addition. The objects (tables, indexes) partially allocate more storage space than needed. Approximately 3 times more storage space is needed as the dump size.

For the archive log we provide disk space six times the size of the dump. Thought we have enough reserve size. The database will hold if disk space for the log is out.

It is advised to have the last database backup on the server machine, in order to reduce the recovery time. If the database raises a media error (defect of file), one of the first recover steps is to perform a media backup (copy of all database files). If the local stored database backup is not valid we need to restore one from the backup system (tape library). We calculate with the size of three complete database backups to have enough disk space for the recovery process. Do not use the free space for file storage. In the case of emergency you are not able to recover your database in time.

Disk space for Oracle database

	factor		single size		total
dump size			1		
Datenfiles	3	*	1	=	3
local db backups	3	*	6	=	18
local dumps	3	*	1	=	3
Archive log	6	*	1	=	6
Oracle 8i Software					1,5
OS System					1
Swap					2,5
total					35

[Gbyte]

The database server needs at least 35 GB disk space.

Network:

100 MB or 1Gbit to App server. Essential is the speed of the connection, not the throughput. 100 Mbit LAN only define the throughput. If the connection is heavily loaded the elapsed time for each IP packet is vhigh and the connection slow. The ping command sends a packet to the remote

machine and the remote machine sends acknowledgment back. The command print out the cycle time.

```
mars 164% ping venus
PING venus (194.120.231.201): 56 data bytes
64 bytes from 194.120.231.201: icmp_seq=0 ttl=255 time=3.495 ms
64 bytes from 194.120.231.201: icmp_seq=1 ttl=255 time=4.211 ms
64 bytes from 194.120.231.201: icmp_seq=2 ttl=255 time=3.107 ms
64 bytes from 194.120.231.201: icmp_seq=3 ttl=255 time=3.132 ms
64 bytes from 194.120.231.201: icmp_seq=4 ttl=255 time=3.179 ms
64 bytes from 194.120.231.201: icmp_seq=5 ttl=255 time=3.231 ms
64 bytes from 194.120.231.201: icmp_seq=6 ttl=255 time=3.150 ms
64 bytes from 194.120.231.201: icmp_seq=7 ttl=255 time=3.115 ms

----venus PING Statistics----
8 packets transmitted, 8 packets received, 0.0% packet loss
round-trip min/avg/max = 3.107/3.328/4.211 ms
```

If database and application are on the same machine the connection is faster than on separate machines. The application runs faster with an IPC (inter process communication) connection instead of a Oracle Net8 (TCP/IP). IPC is only possible on the database machine.

If Data is shared via replication between databases the network load will be at least 128 to 256 kBit. During upgrade or initial synchronization / Installation a considerable bigger network load is to be expected (>= 512 to 1024 kBit). For database replication 25% more CPU power, more memory (>50 MB) is needed. Additional Oracle server processes (ca. 10 connects) perform the replication duty.

If you split database and application for a SQL-Server installation the performance decrease is much higher than with Oracle. Use a one-server installation, when possible.

Example Server:

Simple database server:

- 1 * CPU
- 512 Mbytes Memory
- at least 2 disks (Oracle control files on different disks, separation of OS swap and db log)

A more solid machine for production:

- USP (Uninterruptible Power System) or USV (non; interruptable power supply)
- Server tower box with redundant power supply and cooling
- SCSI-Adapter with 2, 3 or 4 hard disks (more disk I/O -> better performance)
- External DAT-tape
- 1-2*CPU
- 768 Mbytes RAM or more

The following CAL Client Access License is the old license definition for the use of a server. Only if Microsoft services are used it is necessary to purchase a server OS (Windows NT/2000 Server). It is possible to run Oracle on a Windows NT Professional. Oracle 8i and 9i are certified for Windows Professional versions but Microsoft assumes that if you use a machine as server you will

have installed a Windows Server version. So for Eigner PLM Server we require a Windows Server version.

The Professional version limits the network connections to 10 concurrent (Microsoft Services). It is possible to install SQL-Server on Windows Professional. You need special installation CDs (SQL-Server Personal) and there is a limit of 5 parallel connects. The SQL-Server 2000 Personal is used for development purpose. For a productive installation at the customer side we use SQL-Server 2000 Enterprise and a Windows Server version with the amount of client licenses necessary.

Oracle does not have an active license mechanism (license software or license key).

Microsoft

Client Access License (CAL)

The information in this article applies to:

Microsoft Windows NT Advanced Server
Microsoft Windows NT Server versions 3.5, 3.51
Microsoft File and Print Services for NetWare version 3.51
Microsoft Windows 2000 Server
Microsoft Windows 2000 Advanced Server

SUMMARY

This article describes issues related to client access license.

MORE INFORMATION

There are two different licensing models that customers can choose, either a per server or per seat. The per server licensing (concurrent licensing) will cause Windows NT to implement a concurrency check to only allow the number of concurrent file/print users as configured in license manager. Per seat licensing does not implement the license manager check, though the customer must still purchase a client access license (CAL) for such clients.

The following uses require a CAL:

* Server Message Block (SMB) file & print (Windows NT Server)
* NCP file & print (FPNW)
* Remote Access Service (RAS)
* Services for Macintosh (SFM)

Examples of SMB connections are: connecting to a Windows NT server file share or printer share through Windows for Workgroups 3.11, a Microsoft Network client, a LAN manager client, or a Windows 95 client. A NetWare client access FPNW requires a separate CAL.

The following applications *do not* require a CAL:

* http
* gopher
* ftp

Additional query words: CAL

Sizing Application Servers

Hard disks:

>512MB. Provide enough Disk Space for Eigner PLM Programs, and several environments including loader files and some database dumps. I have seen growing things to 3GB in really big projects.

Temporary not used memory is paged out. Many of the Eigner PLM processes will be inactive, because the user works on different software (CAD, office). Or memory allocated for big reports is not currently used. The PLM-Server needs enough swap space or page file for the inactive memory pages.

3 * RAM: RAM < 500 MByte

2 * RAM: 500 < RAM < 2000

1 * RAM: RAM > 2000

Use the OS manual to define the swap space. Some OS have maximum sizes for the swap space. Compared to Windows, Unix will have more swap space but less main memory.

CPU:

One CPU 1 GHz class should be able to support up to 50 active concurrent users. This can differ a lot depending on how our customers use the application. However, most response time problems come from a slow DB, DB-connection, poor customizing and programming.

Memory:

CADIM/EDB 13 MB

Axalant2000: 16 MB

Eigner PLM without Workflow: 18 MB

Eigner PLM with Workflow enabled: 23 MB

Physical memory for each user. For frequent use of complex functions like BOM (structure size) copy etc. requirements are considerably bigger.

The axalant process starts up with an initial size of 10 MB then with little work grow to 13-14 MB. Intensive work especially BOM and structure will allocate 20 and more MB.

Add at least 50 MB for OS.

The maximum users for SQL-Server are around 100. Please ask PSO for their SQL-Server experience. For more than 50 users I would recommend Oracle as database software. I also would prefer UNIX as the server operating system for bigger installations (100 users). But this is a decision made by the system administration.

We recommend a maximum of

- 100 concurrent users for Oracle and Eigner PLM on one Windows2000 Server and
- 80-90 concurrent users for SQL-Server and Eigner PLM on one Windows2000 Server

To reach this values it must be a Windows 2000 Server Installation with more than 2GB main memory (when OS is installed) and the registry must be modified. A standard Windows 2000 Server Installation without modifications can server 50-60 users. This restriction does not exist on

Windows NT 4.0. Please contact support it you want to upgrade from Windows NT to Windows 2000.

Network:

The axalant process holds one database connection and general one client connection. The database connection has 10 times the throughput and much more round trips thus a fast network connection especially to the DB server is necessary (no WAN).

Batchjobs:

Batch jobs are often necessary for plotting or file conversion (XML-> PDF, CATIA->STEP). To perform the conversion the CAD software and Eigner PLM must be installed on the machine. The OS is then identical to the CAD-machines. Identify the machines for the batch jobs:

- Dedicated server
- On each CAD-Workstation
- ON the PLM Server

If the CAD Workstations are used as PLM-Servers, this will reduce the memory and CPU consumption on the central servers. This is often used on UNIX boxes.

Sizing Business Services

The Business Service is installed in conjunction with the PLM Server (Application Server). The Business Service is activated, when Workflow is enabled. It is possible to run the Business Service on a separate machine.

Hard disks:

Eigner PLM Server installation. No extra disk space

CPU:

Depends on the Workflow activities. The Jboss only works on one CPU. No additional CPU power necessary to calculate.

Memory:

Windows:

```
30 MB + ConcurrentUser * 2 MB
```

Unix:

```
70 MB + ConcurrentUser * 2 MB
```

The Java Application Server is configured with a maximum of 300 MB (java –Xmx300m). This value sets the maximum size of the JAVA virtual machine and can be modified. If you use the above formula please add sufficient reserves (min size 128 MB).

Network:

The Business Service opens a ECI connection to the PLM server process and a JDBC connection to the database.

Sizing Presentation Services

The Presentation Service should run with the PLM Server on the same machine. It is possible to run the Presentation Service on a separate machine but not advised.

Hard disks:

Eigner PLM installation plus 10 MB per environment

CPU:

Needs three times more CPU power per connect than the axalant server process. The Tomcat servlet engine only works on a single CPU. More Presentation Service environments (tomcat) can make use of multi CPU servers. The connection TCP port is then different.

Memory:

20 MByte + n * 6 MByte

n= number of users (connects)
6 MB per connect
20 MB as base size

The minimum size for the Presentation Service despite the calculation is 128 MB. The JAVA virtual machine starts with a maximum memory of 256 MB (java -Xmx256m). Raise the value if necessary.

1. Example

12 users
20 MB + 12*6MB = 92 MB
=> recommended size is 128 MB

2. Example

30 users
20 MB + 30*6MB = 200 MB
=> recommended size is 200 MB

Network:

The Presentation Service needs a fast connection to the PLM Server. It is recommended, that the axalant processes for the Presentation Service are running on the same machine.

Sizing FMS Servers

There are no special requirements for the Eigner PLM FileServer. The machine should not swap and enough free memory should be available for file buffer. The throughput of the FileServer is determined by the network connection. In normal case the possible disk I/O is higher than the throughput of the network interface. For the FileServer identical OS versions are used as the PLM Server. It is not necessary having the same OS as the client. You can install on mixed OS.

- Clients are Windows or UNIX or both
- FileServers are Windows or UNIX or both

Hard disks:

The FileServer installation needs 1 MB of disk space. Estimate the necessary disk space for the stored documents.

- How much disk space will be necessary the next month and years?
- What is the concept to raise disk space and backup volume?
- Is it possible to add new disks to the system?

The FileServer stores the files in vaults. An electronic vault is a directory and its contents. One directory is limited by the partition size. The maximal capacity of an electronic vault is the size of the partition. If the vault is running out of disk space you can create a new vault on a second partition or you have to shift the vault to a bigger partition. The directory must be moved to the bigger partition with preserved file permissions (Windows!!) and the vault definition must be changed in the PLM System.

Eigner only support local file systems. File systems in particular NFS commit the file write in a state where the file is still in the file cache and not completely written on the remote server. If the FileServer machine crashes in this situation the file is corrupt. If Storage systems (SAN,..) are utilized you have to exclude this issue. This is the same security reason, why a database only uses local file systems.

Often a RAID5 shelf is used for security and capacity reasons.

CPU:

No special requirements. CPU is needed for the TCP/IP stack and File I/O

Memory:

256 MB for a dedicated server.

The process themselves allocates 1 MB. Each connection uses main memory to buffer the transmission. So enough free memory should be available.

If the WebClient is used, it is advised to install the Presentation Service, with include the WebFileService on this machine. Therefore at least 256 MB are needed. Without WebFileService 128 MB are ok.

Network:

A fast connection to the PLM client reduce time to store and load files. The Network is the bottleneck of the FileServer. The disk system is generally faster than the network and CPU load is not high on the server. If clients are in different network segments you can use more than one network card or use multiple FileServer. But the vault definition only has one hostname (IP-address) and can only be reached over one network card.

You can combine the FileService with the PLM Server, Business Service and the Presentation Service. The FileService needs primarily disk I/O, while the other services needs CPU. There will be a load if big files are checked in or out and network load is high during transmission.

General Considerations, License Manager (FELICS) etc…

License Agents have to be installed on each Application Server machine. If you use DFM (Distributed File Management) a Felics Agent is necessary on the FileServer machine.

Oracle offers a cluster solution (Real Application Cluster) were several servers are working on the same database. But general the database exists once and only one server is working with the database (Database Server). This is the single point of failure and you are advised to install the Felics License Service on this machine.

You can install as much PLM-Servers as you want. Each axalant-Server process allocates licenses via the Felics agent on the central Felics license server and it has no influence if the license request comes from one or from several servers. The number of concurrent users and modules used concurrently are limited through the License server.

In conjunction with the axalant-Server a database client must be installed. The Oracle Client is license free. The Oracle Client is a subset of the Oracle server installation. The Oracle Client has all components necessary to connect and administrate a Oracle Server. If a customer purchase a Oracle Server on UNIX he will get the CDs for a Oracle Client on Windows in addition, to manage Oracle server from his Windows PC.

Oracle does not have a license mechanism like a license key or license server. The Oracle parameter “processes” in the init.ora limit the number of Oracle processes. The value must be higher than the number of concurrent users plus the Oracle background processes (8). Raise this parameter if you add Eigner PLM licenses.

There is an option to run axalant AppServer and FMS on the same machine as FMS requires mainly I/O and axalant mostly requires CPU. You will have to consider that they share Network throughput on this machine.

Eigner PLM can be integrated to most of currently available IP-Based fail over solutions. Sizing for this will differ according the needs of the fail over software, usually requires a second identical machine sharing APP or DB disks via Dual SCSI to a separate Disk (usually RAID) shelf. All PLM services are installed and run on both machines, except the database, which is managed by the cluster software.

Minimum Requirements for Test and Customizing

Avoid hardware resource problems (Memory/CPU) at your customer’s site for test and small customizing environments using the following minimum layout:

CPU >=800 MHz class
>= 192 MB for Database machine
>= 256 MB for PLM Server machine
>= 256 MB for Presentation Service machine
>= 512 MB when running all processes on one machine

Do NOT use this layout for any productive installation!

Configure Oracle SGA

The Oracle init file (init.ora) describes the memory allocation of the Oracle instance and their user processes (dedicated servers). The settings modify the behavior of special tasks.

Sizing Oracle SGA

concurrent users		80			input
dump size		300	MB		parameter
init.ora:					
db_files	=	32			
db_block_size	=	4096			
processes	=	110			
db_file_multiblock_read_count	=	32			
				[MB]	
db_block_buffers	=	19200		75,00	SGA
shared_pool_size	=	40000000		38,15	
large_pool_size	=	614400		0,59	
java_pool_size	=	0		0,00	
log_buffer	=	393216		0,38	
				114,11	
dml_locks	=	640			
log_checkpoint_interval	=	80000			
log_checkpoint_timeout	=	0			
compatible	=	8.1.7			
sort_area_size	=	1048576			
sort_area_retained_size	=	65536			
open_cursors	=	60			
session_cached_cursors	=	30			
optimizer_mode	=	rule			
UNIX Kernel configuration					
Oracle processes		110			
axalantrt processes		210			
HP-UX Kernel Parameter	/usr/conf/master.d/core-hpux				
SHMMAX	=	1073741824	>		119650816
SHMMIN	=	1			
SHMMNI	=	100			
SHMSEG	=	32	HP-UX depend		
MAXUSERS	=	256			
MAXUPRC	=	210			
SEMMNS	=	230	nproc		2068
SEMMNI	=	110	nfile		4016
max_thread_proc	=	1024	nkthread		3635
maxfiles	=	256	ncallout		3651
Solaris Kernel Parameters	/etc/system				
set shmsys:shminfo_shmmax	=	4294967295	>		119650816
set shmsys:shminfo_shmmin	=	1			
set shmsys:shminfo_shmmni	=	100			
set shmsys:shminfo_shmseg	=	32			
set semsys:seminfo_semmns	=	230			
set semsys:seminfo_semmni	=	110			
set semsys:seminfo_semmsl	=	120			
set semsys:seminfo_semvmx	=	32767			
set semsys:seminfo_semopm	=	100			
Linux Kernel Parameters	/proc/sys/kernel				
SHMMAX	=	RAM / 2	>		119650816
SHMMIN	=	1			
SHMMNI	=	100			
SHMSEG	=	32			
SEMMNS	=	230			
SEMMNI	=	110			
SEMOPM	=	100			
SEMVMX	=	32767			

Oracle init parameters

- **db_files:** maximum number of datafiles; should be higher than the number of current datafiles; each file entry need memory space in each oracle process
- **db_block_size:** multiple of OS block size
  - 2048: test and development systems
  - 4096: small and medium size systems
  - 8192: medium and big systems

  An 8kByte block size provides the best File I/O. SQL-Server uses general 8kByte block size. The block size cannot be changed after database creation.
- **processes:** Maximum number of Oracle processes (connects+ background processes+ replication processes). About 8 background processes perform Oracle base functions. If replication is used there are around 10 dedicated server processes. For most UNIX systems you have to adapt the amount of semaphores.
- **db_file_multiblock_read_count:** Number or database blocks read in one step. The OS reads a minimum of 64 Kbytes. The database should read the same amount or a multiple of that. RAID 5 systems read at least one stripe and the database should read the identical size.
- **db_block_buffers:** Number of database blocks held in the main memory. The RAM-read at least 1000 times faster than a disk read. The more database blocks are in the main memory the less disk I/O is needed and the faster the database performs. Raise the block buffer as much as possible. The machine should not run out of physical memory and start to swap. This will decrease database performance. In the sizing description the buffer is dimensioned 1/10 of the dump. This is the minimum value. The SGA calculation uses ¼ of the dump for better performance.
- **shared_pool_size:** Memory (in bytes) used to buffer system tablespace (Object definition, stored procedures, users and their rights,..), hold parsed statements
- **large_pool_size:** Used for parallel execution and backup. Memory mainly used for multithreaded server. Do not change value.
- **java_pool_size:** Set to zero to disable Java Server (64 Kbytes are still allocated)
- **log_buffer:** Buffer for the redo log information (in bytes). Buffer must be written to the disk at each commit. Write signal comes if buffer is filled up to 1/3. Use at least 3*64 kByte. Larger log buffer improves performance of big transactions (imp, loader, structure explosion,..).
  - 3*64 kByte (2^16 *3): test and small
  - 6*64 kByte (2^17 *3): productive small
  - 12*64 kByte (2^18 *3): productive medium
  - 24 – 128 *64 kByte (2^19 *3 – 2^23): productive big

Use a calculator. The value should be a multiple of 2^16.

- **dml_locks:** Number of locks in the system (processes * 8)
- **log_checkpoint_interval:** maximum number of redo log blocks (OS blocks) written between two checkpoints. Raise the number that no additional checkpoint occur between two log switches or set it to zero to disable the parameter:

```
log_checkpoint_interval > redo log size / OS block
log_checkpoint_interval > 1 000 000 / 512
```

- **log_checkpoint_timeout:** Maximum time in seconds between two checkpoints. A checkpoint needs a couple of seconds to be performed (1-8sec). You can disable the parameter by setting it to zero. For security reasons insert a time value, where the database must perform a checkpoint and the database is in a consistent status.
- **compatible:** The database can behave like an earlier version. Set the value like the current version of the software, to have all new features. If you raise the value to the current version it might be not possible to set it back to an older version, because new features has been used and are now in the database
- **sort_area_size:** Specifies in bytes the maximum amount of memory Oracle will use for a sort. After the sort is complete, but before the rows are returned, Oracle releases memory down to the size specified by the SORT_AREA_RETAINED_SIZE parameter. After the last row is returned, Oracle releases the remainder of the memory. Each user process allocates this amount of memory (in bytes) to perform sorts. You can use the command: alter session set sort_area_size=20000000; to modify the parameter for your session (create index, …).
- **sort_area_retained_size:** Specifies in bytes the maximum amount of the user global area (UGA) memory retained after a sort run completes. The retained size controls the size of the read buffer, which Oracle uses to maintain a portion of the sort in memory. This memory is released back to the UGA, not to the operating system, after the last row is fetched from the sort space.
- **open_cursors:** Specifies the maximum number of open cursors a session can have at once. Each DataView mask opens one cursor. Set value to 70. Often Java applications do not close their cursors, then the value is raised to 3000 – 8000.
- **session_cached_cursors:** Lets you specify the number of session cursors to cache. Repeated parse calls of the same SQL statement cause the session cursor for that statement to be moved into the session cursor cache. Subsequent parse calls will find the cursor in the cache and need not reopen the cursor.
- **optimizer_mode:** {RULE | CHOOSE | FIRST_ROWS | ALL_ROWS}
  - **RULE** specifies that rule-based optimization is to be used. This is a secure setting. The performance of the database is stable.
  - **CHOOSE** specifies cost-based optimization for a SQL statement if the data dictionary contains statistics for at least one table accessed in the statement. If no such statistics are present, the optimizer uses the rule-based approach. This is the default setting. If the statistics are old their will be a tremendous performance decrease. You have to generate statistics regular.

```
SQL> EXECUTE DBMS_STATS.GATHER_SCHEMA_STATS('AXA');
SQL> EXECUTE DBMS_STATS.DELETE_SCHEMA_STATS('AXA');
```

UNIX Kernel Parameters set for Oracle

The Oracle installation manual describes the special settings. See also online help on Oracle installation CD.

Shared Memory

- **SHMMAX:** maximum size of a shared memory segment. This setting does not affect how much shared memory is needed or used by Oracle8i or the operating system. It is used only to indicate the maximum allowable size. This setting also does not impact operating system kernel resources. The SHMMAX setting should be large enough to hold the entire SGA in one shared memory segment. A low setting can cause creation of multiple shared memory segments which may lead to performance degradation.
- **SHMMIN:** minimum size of shared memory segment
- **SHMMNI:** maximum number of shared memory identifiers in the system
- **SHMSEG:** maximum number of shared memory segments a user process can attach

Semaphores

- **SEMMNI:** maximum number of semaphore identifiers (semaphore sets) in the system
- **SEMMSL:** maximum number of semaphores in a set, set to 10 plus the largest PROCESSES parameter of any Oracle database on the system.

  semmsl = processes + 10

- **SEMMNS:** (SEMMNI * 2) maximum number of semaphores in the system. Set to the sum of the PROCESSES parameter for each Oracle database, adding the largest one twice, then add an additional 10 for each database.

  semmns = processes *2 + 10

- **SEMOPM:** maximum number of operations per semop call
- **SEMVMX:** Defines the maximum value of a semaphore.

Security limits

- **MAX_THREAD_PROC:** Defines the maximum number of kernel threads allowed per process. You may need to increase the value if required by your application. Setting it to a default or low value may lead to an out of memory error for certain applications.
- **MAXUSERS:** maximal number of users on the system. For HP-UX other parameters are derived from this parameter like nproc, nfile, nkthread, and ncallout)

  Oracle-Connections + 64

- **MAXUPRC:** Defines the maximum number of user processes.
- **NPROC:** Defines the maximum number of processes.
- **NKTHREAD:** Defines the maximum number of kernel threads supported by the system.
- **NCALLOUT:** Defines the maximum number of pending timeouts.

Chapter 5
Examples

Example 1

Starting Def:

Users: 30 users
Database Software: SQL-Server
Server OS: Windows 2000 Server

Final extension Def:

Users: 540 users
Database Software: SQL-Server
Server OS: Windows 2000 Server

We cannot recommend the amount of 500 concurrent users for SQL-Server. We have big installations with around 100 concurrent users on SQL-Server. SQL-Server is a solution for smaller installations and has less license costs. But SQL-Server is not as fast as Oracle and especially if database and application run on separate machines there is a performance decrease. Maybe the number of users comprises all employees using the system (named users) but not concurrent. You need the number of users logged in the system at one time (concurrent users).

A separation of database and application is generally not recommended. If the axalant process is working (needs CPU), the database is waiting for requests and if the axalant process sends a SQL-statement to the database it is waiting for the result. So either the database is working or axalant With bigger installations more processes are working parallel and a two-server installation will work efficiently. To have a two or multiple server installation is a decision made by the system administration. They often want to have two identical machines for better maintenance and system backup. If one server fails, the other can do the job. Or two smaller machines are cheaper than one big machine. Having one big server with a specific CPU power is always the fastest solution (not network time, better use of CPU power).

I give you a rough estimation of the hardware needs. Again it is necessary to estimate the concurrent users. Officially the estimate two named users are equivalent to one concurrent user. In practice the ratio is 3:1 or 4:1, ask someone from PSO.

With Workflow enabled and WebClient:

Eigner PLM 5.0.1 memory consumption for SQL-Server 2000

Workflow enabled	1
concurrent users	30
WebClient users	10
machines	1
dump size	250

	base size	per connect	total
axalant-Server	20	23	710
Business Service	30	2	128
Presentation Service	20	6	128
SQL-Server	30	2	90
DB block buffer			25
System	50	0,5	65
total			1146

[Mbyte]

One small server is necessary with at least 1GB main memory, one or two CPUs. More memory would be advised. The machine must have at least 2 hard disks. The FileServer can be on the same or on a different machine. The necessary disk space for the Eigner PLM FileService must be estimated. There will be a small load for the Presentation Service (Web Client). If Web Client is more often used Dual CPU is needed.

With Workflow enabled and WebClient:

Eigner PLM 5.0.1 memory consumption for Oracle 8i

Workflow enabled	1
concurrent users	540
WebClient users	200
machines	1
OS	UNIX
dump size	1000

	base size	per connect	total
axalant-Server	20	23	12440
Business Service	70	2	1150
Presentation Service	20	6	1220
Oracle	23	6	3263
DB block buffer			100
System	50	0,5	320
total			18493

[Mbyte]

For the workload of 540 users, Oracle as the database solution and a UNIX operating system is recommended. One, two or three servers can perform the tasks. At least 4 CPUs are necessary. The database need at least 5-7 independent disks or the equivalent write and read performance. The FileService must be installed on a separate machine.

If there are 540 named users, I estimate 150 concurrent users and assume that 60 of them use the WebClient as front end.

With Workflow enabled and WebClient:

Eigner PLM 5.0.1 memory consumption for Oracle 8i

Workflow enabled	1
concurrent users	150
WebClient users	60
machines	2
OS	Windows
dump size	1000

	base size	per connect	total
axalant-Server	40	23	3490
Business Service	60	2	360
Presentation Service	20	6	380
Oracle	23	6	923
DB block buffer			100
System	100	0,5	175
total			5428

[Mbyte]

This is a two-server configuration with 2 CPUs and 3-4 GB main memory each. The FileService can be installed on one machine. The other machine must be the database server. Or the FileService uses a separate machine. The database server need at least 3-5 independent disks or the equivalent write and read performance. If workflow is disabled this will reduce the memory consumption from 23 MB to 18 MB per process. The Java App Server (Business Service) is then not necessary.

We recommend a maximum of

- 100 concurrent users for Oracle and Eigner PLM on one Windows2000 Server and
- 80-90 concurrent users for SQL-Server and Eigner PLM on one Windows2000 Server

Too reach this values it must be a Windows 2000 Server Installation with more than 2GB main memory (when OS is installed) and the registry must be modified. A standard Windows 2000 Server Installation without modifications could server 50-60 users. This restriction does not exist on Windows NT 4.0. Please contact support it you want to upgrade from Windows NT to Windows 2000.

We have installations using SQL-Server with maximum users around 100. Please ask PSO for their SQL-Server experience. For more than 70 users or database intensive work (large part structure, 3D CAD) I would recommend Oracle as database software. I also would prefer UNIX as the server operating system for bigger installations (> 100 users). But this is a decision made by the customer system administration.

Example 2

Starting Def:

Users: 50 users named
Database Software: SQL-Server
Server OS: Windows 2000 Server preferred

Second step Def:

Users: 200 users named
Database Software: SQL-Server
Server OS: Windows 2000 Server preferred

The customer will have 50 active users in the first step and 200 in the second step. But the customer does not know the 'concurrency'. Users on CAD workstation often us Eigner PLM permanent while others only start Eigner PLM for retrieval or special tasks and close the application afterwards. The customer expects 40 concurrent users. The number of Eigner PLM licenses limits the concurrent users. At least the machine must be able to server this users.

40 concurrent users without Workflow

Eigner PLM 5.0.1 memory consumption for SQL-Server 2000

Workflow enabled	0
concurrent users	40
WebClient users	0
machines	1
dump size	250

	base size	per connect	total
axalant-Server	20	18	740
Business Service	30	2	0
Presentation Service	20	6	0
SQL-Server	30	2	110
DB block buffer			25
System	50	0,5	70
total			945

[Mbyte]

A dual processor machine will perform well as application and database server. At least 1GB RAM is necessary. For SQL-Server a Windows 2000 Server installation with the 40 or more client accesses are necessary.

To run the database at least two disks are necessary. An additional disk will raise the performance.

It is advised to use a separate Eigner PLM FileServer machine or at least use a different disk for database log and FileServer vaults.

The WebClient and Workflow is unaccounted for the memory calculation. Additional memory is necessary:

- Presentation Service (Web Client) >= 128MB
- Business Service (Java App Server f. Workflow) >= 128 MB

- Axalant Server 23 MB instead of 18 MB per connect.

Example 3: Comparison one or two Server installation

A single Eigner PLM Server configuration

Boundary conditions

- 60 concurrent users
- 20 Web Client users
- with Workflow
- ca. 50 000 items -> dump size 300 MB
- Oracle 8.17

The central PLM Server run the following services:

1. database server
2. axalant-Server
3. Presentation Service
4. Business Service
5. Lizenzserver

The Eigner PLM FileServer for this one server configuration will be installed on a separate machine.

Eigner PLM 5.0.1 memory consumption for Oracle 8i

Workflow enabled	1
concurrent users	60
WebClient users	20
machines	1
OS	Windows
dump size	300

	base size	per connect	total
axalant-Server	20	23	1400
Business Service	30	2	150
Presentation Service	20	6	140
Oracle	23	6	383
DB block buffer			30
System	50	0,5	80
total			2183

[Mbyte]

Disk Space for database

We conclude from the amount of items to a dump size of 200 to 500 MB.

To include reserves we calculate 1GB. The test, training and development dumps are hold in a separate database. It is advised to use a second database instance for these environments, maybe on a second server.

Disk space for Oracle database

	factor		single size		total
dump size			1		
Datenfiles	3	*	1	=	3
local db backups	3	*	6	=	18
local dumps	3	*	1	=	3
Archive log	6	*	1	=	6
Oracle Installation					1,5
OS System					1
Swap					2,5
total					35

[Gbyte]

The database server needs at least 35 GB disk space. 3-4 disks on a SCSI-Controller are advised. Alternative a RAID System can be used.

Recommendation for the PLM-Server:

- Dual CPU machine with 2 GB RAM,
- 3-4 hard disks and SCSI Controller

Because of low memory costs additional memory slots should be used to give Oracle more memory to raise performance. The physical memory limits the number of concurrent users. To server more users additional memory is necessary. Further versions of Eigner PLM will need more memory. Plan the necessary memory extension of the machine. Additional machines can also serve additional users and raise the CPU power.

Eigner PLM FileServer recommendations:

- Memory >= 256 Mbyte,
- Hard disks with enough disk Space or RAID system (level 5).

If the WebClient is used it is advised to install the Presentation Service, with include the WebFileService on this machine. Therefore at least 256 MB are needed. Without WebFileService 128 MB are ok.

A double Eigner PLM Server configuration

Two identical servers with the one server configuration are used. During normal productive operation the two servers should run the Eigner PLM system. But in the case of emergency, one of both servers should be able to run the complete Eigner PLM system.

- Dual CPU machine with 2 GB RAM,
- 3-4 hard disks and SCSI Controller

This composition should provide higher system availability. The probability of failure for two machines is twice that of one machine. But if one machine fails the other can perform the services of the crashed one. This is the concept of a fail-safe cluster (take over cluster). But without a raid system and cluster software the administrator has to shift the services onto the second machine.

The recovery time will be shorter because it is not necessary to repair the crashed server. But this is only the case if the admin has a concept for shifting the database and the FileService. He can store the files in an external box and shift the plug from one server to the other. Or if he use a storage system, attach the file system from the second server.

He can store the database backup and all produced archive logs on the second machine. For this he needs a backup and recovery concept with must be tested. Oracle themselves offers different security concepts (Standby server, Real application cluster).

Eigner PLM Server 1:

1. Database server
2. License server
3. axalant server
4. Business Service
5. Presentation Service

Eigner PLM Server 2:

1. axalant server
2. Backup License Server
3. Eigner PLM FileServer
4. Presentation Service

The Backup License server must be purchased separately. The Backup License server is only needed if server 1 fails. If Workflow is enabled the PLM Server (axalant processes) needs a running Business Service on the same or on a separate machine.

The first machine is used for the database and the Presentation Service (WebClient). The axalant processes of the WebClient are running on the same machine, with improves performance. If the Business Service starts up it needs a running database. The Business Service establishes a connection pool and retrieves user information to create the cache for the Permission Manager. To ease administration, run the Business Service on the database machine.

The second machine runs the axalant server processes for the Windows and UNIX clients. The Eigner PLM FileServer runs on this machine because no other services uses intense file I/O. The Presentation Service on the second machine includes the WebFile Service. The WebClient uses this service to access files from the FileServer.