



Capacity Planning and Deployment Guide

Release 9.2.1

Part No. E11130-01

Make sure you check for updates to this manual at the
[Oracle Technology Network \(OTN\) Web site](#)

Copyrights and Trademarks

Copyright © 1995, 2007, Oracle. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle and Agile are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

September 13, 2007

REVISIONS

Revision	Date	Pages Affected	Description
A	06/23/06	All	Initial release of manual.
B	07/10/06	9-1	Added JVM parameters for Solaris.
C	07/17/06	6-4	Changed caching view server description.
D	07/19/06	3-2	Changed the name of Table 3-3. Added Table 3-5, Change Status to Released (Each Item has at least 3 Pending Changes.
		9-1	Removed “Using Compression Filters.”
E	07/24/06	Preface	Corrected Training Website URL.
F	07/26/06	1-4	Added Linux as a supported platform for Apache 2.0.55.
G	08/09/06	1-4	Corrected Windows 2003 version number.

CONTENTS

Chapter 1 Overview

Overview	1-1
Application Server	1-1
Load Balancer	1-1
Clients	1-1
Database Server	1-1
Agile File Manager	1-2
LDAP Directory Server	1-2
View Server	1-2
Software Requirements	1-4
Hardware Requirements	1-5

Chapter 2 Capacity Planning

Hardware Sizing	2-1
Agile Application Server	2-2
Hardware Sizing for Windows and Linux	2-3
Hardware Sizing for Solaris	2-4
Database Server Sizing	2-5
Oracle Database Sizing	2-5
Hardware Resource Plan for Database Models	2-7
Disk Space and Tablespace Configurations	2-7
Hardware Load Balancer and Proxy Web Server	2-10
CPU	2-10
Network Card	2-10
Distributed File Management	2-10
View Server	2-10

Chapter 3 Breakpoint Testing Results

Solution Results on Windows and Linux	3-1
Product Collaboration	3-2
Product Quality Management	3-3
Distributed File Manager	3-5
Export	3-5
Product Cost Management	3-6
Solution Results on Solaris	3-7
Product Collaboration	3-8
Product Quality Management	3-8
Distributed File Manager	3-9
Product Cost Management	3-11
Portlets Performance	3-12

Chapter 4 Client Requirements

Client Requirements	4-1
---------------------------	-----

Chapter 5 Directory Server

Overview	5-1
System Architecture	5-1
Implementation Details	5-2
Configuration	5-2
Utilities	5-3
Frequently Asked Questions (FAQ)	5-3

Chapter 6 File Management Server and View Server

Agile File Management Server	6-1
File Vaults	6-1
Distributed File Manager	6-2
How Distributed File Manager Works	6-2
File Management Security	6-4
View Server	6-4
View Server Security	6-5

Chapter 7 Communication and Data Flow

Communication and Data Flow	7-1
Client Communication	7-2
Proxy Web Server	7-2
Typical Network Configurations	7-2
Application Server	7-3
Directory Server	7-3
Distributed File Management	7-3
View Server	7-4
Network Bandwidth Requirements	7-4

Chapter 8 Product Interchange

About Product Interchange	8-1
Product Interchange Requirements	8-1

Chapter 9 Performance Tips

Tuning Memory for Java Applets	9-1
Optimum JVM Parameters on Solaris	9-1

Appendix A System Configuration Examples

Small System Configuration	A-1
Medium System Configuration	A-2
Large System Configuration	A-2
Extra Large System Configuration	A-3

PREFACE

The Agile documentation set includes Adobe® Acrobat™ PDF files. The [Oracle Technology Network \(OTN\) Web site](#) contains the latest versions of the Oracle|Agile PLM PDF files. You can view or download these manuals from the Web site, or you can ask your Agile administrator if there is an Oracle|Agile Documentation folder available on your network from which you can access the Oracle|Agile documentation (PDF) files.

To read the PDF files, you must use the free Adobe Acrobat Reader™ version 7.0 or later. This program can be downloaded from the www.adobe.com.

The [Oracle Technology Network \(OTN\) Web site](#) can be accessed through **Help > Manuals** in both the Agile Web Client and the Agile Java Client. If applicable, earlier versions of Oracle|Agile PLM documentation can be found on the www.agile.com/support.

If you need additional assistance or information, please contact support@agile.com or phone (408) 284-3900 for assistance.

Before calling Agile Support about a problem with an Oracle|Agile PLM manual, please have ready the full part number, which is located on the title page.

Readme

Any last-minute information about Oracle|Agile PLM can be found in the Readme file on the [Oracle Technology Network \(OTN\) Web site](#).

Agile Training Aids

Go to the [Agile Training Web page](#) for more information on Agile Training offerings.

CHAPTER 1

Overview

Overview

The Agile Solution has an N-tier, J2EE architecture. These tiers are typically deployed across several servers in a production environment: Application Server, File Manager, Web Server, Database Server, and View Server.

Application Server

The Agile Application Server is the center of the Agile system, the base for the PLM platform, where all common services and business logic reside for the entire solution. The Agile Application Server runs on industry-leading J2EE application servers. See Table 1-1 for the list of supported application servers. As Figure 1-1 illustrates, all client servers and users connect to the Application Server either directly or indirectly. The application server connects to the components in a persistence layer where product content is stored, the database server, the file manager, and the LDAP server.

Load Balancer

The hardware load balancer brokers client communications without compromising the security of your internal network. Clients communicate through the load balancer with the application server. There are no Agile Software components running on the hardware load balancer. They are usually deployed in the Demilitarized Zone (DMZ) where it proxies requests from outside the corporate firewall to the application server in the Safe Zone. A load balancer is necessary if you are installing Oracle Application Servers in a cluster.

Clients

Agile PLM 9.2.1 includes two clients, a Web client and a Java client. The Web client is a thin HTML client that uses firewall-friendly protocols (HTTP/S). The Java client is a Java-based client that can use application server-specific protocols, such as ORMI for Oracle Application Server and T3 for BEA WebLogic, to connect to the server. Each client has its own strengths and weaknesses from a functional, architectural, and performance standpoint. This document compares the clients from an architectural and performance standpoint.

Database Server

The Agile Database Server persists or stores all product content and system settings. Agile's database server runs on Oracle 9i or 10g.

Agile File Manager

The Agile File Manager stores all documents, drawings, and other files within the Agile system. Due to the geographically dispersed nature of the global enterprise, multiple Agile File Managers can be deployed in a distributed configuration for efficient distribution of product content. Agile File Manager is made up of two main components: the file server and the file vault. The file vault represents the file system where the actual files reside. The file vault can be located on the application server or a dedicated storage system.


LDAP Directory Server

In an effort to better support the industry standard authentication schemes, Agile PLM 9.2.1 supports Lightweight Directory Access Protocol (LDAP) based authentication. LDAP support enables you to integrate Agile with existing directory servers so user accounts can be managed in one place. Integrating with LDAP is optional. Users can be managed within Agile without a directory server. There are no Agile Software components deployed on the Directory Server.

View Server

The View Server allows users to view and mark up documents and drawings in a supported Web browser. One of the advantages of using the viewer is that users can view files stored in Agile without having the native application that created the file installed on their desktop machines. Another advantage of the View Server is that it decreases bandwidth requirements and improves response time. Instead of sending large native files to the user, the View Server streams a smaller image file to the View Server applet on the client.

Figure 1-1: System Configuration Overview



Software Requirements

Table 1-1 shows the operating system and other software supported or required for each Agile component. The application server includes all server components for AIS, SDK, Reports, Import, Export, Agile Content Service, LDAP, and all solutions.

Table 1-1: Agile PLM 9.2.1 Software Requirements

Agile Components	Third-Party Software Requirements
Server Operating Systems	Windows 2003 SP 1, 32-bit Edition Solaris 9 and 10, SPARC only Red Hat Linux 4, All Editions
Client Operating Systems	Windows 2000 SP4 Pro Windows XP SP2 Pro Apple OS X 10.4 Solaris 10
Database Servers	Oracle 9i R2 (9.2.0.7) (Windows and Solaris only) or Oracle 10g R2 (10.2.0.1.0) Standard Edition for up to 4 CPUs or Enterprise Edition for more than 4 CPUs
Application Server	Oracle Application Server 10g Java Edition v10.1.2.0.2 or BEA WebLogic 8.1 SP5
Web Server	IIS 6.0 for Windows 2003 Apache 2.0.55 for Solaris and Linux
LDAP Server	Microsoft Active Directory Server or Sun Java System Directory Server
Portal Server	IBM WebSphere Portal Enable 5.1.0.2 with supported database
Portal Database	See IBM's website for supported databases and operating systems. For Agile PLM 9.2.1, the portal database is certified with Oracle 9i on Windows 2003.
File Manager	TomCat 5.0.30
View Server	Cimmetry Jvue 19 on Solaris 9 only (19SP2 for Windows) Not supported on Linux or Solaris 10
Java Client	JRE 5.0 and supported web browsers for Windows and Solaris
Web and Portal Client	Internet Explorer 6.0 SP2 and Firefox 1.5 for Windows Firefox 1.5 for Unix operating systems Safari 2.0 and JRE 5.0 Release 2 for Apple operating systems JRE 5.0 is required for applets.
Express Client	JRE 5.0
ChangeCAST Client	JRE 5.0 on Windows only
SDK Clients	Same JRE as the Application Server
Dataload Client	Same JRE as the Application Server Microsoft Access
Microsoft Excel for Product Interchange	Microsoft Excel 2000
Microsoft Excel Integration for PG&C	Microsoft Excel 2000 SP3, XP, 2003 on supported Microsoft Windows client operating systems
Microsoft Project Integration for PPM	Microsoft Project 2002 and 2003 on supported Microsoft Windows client operating systems

Table 1-2: Supported Platforms for Agile 9.2.1 Product Interchange Server

Platforms	Database Server	Database Server	Web Server
	Oracle 9i 9.2.0.1	Oracle 10g 10.2.0.1.0	Tomcat 5.0.30
Windows 2003 (Intel Pentium IV or higher)	X	X	X
Solaris (Ultra Sparc III or higher)	X	X	Not Supported

Hardware Requirements

Table 1-3 shows the minimum hardware needed to deploy an Agile PLM 9.2.1 system.

Table 1-3: Agile PLM 9.2.1 Hardware Requirements

Agile Server	Hardware Requirements
Application Server	Dual CPU Pentium Xeon, 3 GHz and above OR Dual CPU Ultra Sparc III, 900 MHz and above 2 MB L2 cache 2 GB RAM
Database Server	Dual CPU Intel Xeon 2.8 GHz OR Dual CPU Ultra Sparc III 512 KB L2 Cache 1 GB RAM
File Manager and View Server	Dual CPU Intel Xeon 2.8 GHz OR Dual CPU Ultra Sparc III 512 KB L2 cache 2 GB RAM
Web Server	Intel Xeon 2.8 GHz OR Ultra Sparc III 512 KB L2 Cache 1 GB RAM

Table 1-4 shows the minimum hardware needed to deploy Agile Product Interchange.

Table 1-4: Agile 9.2.1 Product Interchange Server Hardware Requirements

Agile 9.2.1 Product Interchange Server	Hardware Requirements
Application Server	Dual CPU Intel Xeon, 3.0 GHz and above OR Dual CPU Ultra Sparc III, 900 MHz and above 2 MB L2 Cache 2 GB RAM 2x36MB SCSI 10K RPM HDD
Database Server	Dual CPU Intel Xeon 2.8 GHz and above OR Dual CPU Ultra Sparc III, 900 MHz 512 KB L2 Cache 1 GB RAM 2x36MB SCSI 10K RPM HDD

CHAPTER 2

Capacity Planning

Hardware Sizing

Agile conducts extensive load tests to determine scalability for individual product components, as well as for combinations of modules. Agile uses Silk Performer 6.0 by Segue Software to simulate virtual user load for the benchmark tests.

To determine the required hardware for a given implementation, a number of factors must be considered. These factors are:

- 1 Average user load
- 2 Peak user load
- 3 User distribution across different modules, if more than one module is implemented
- 4 Network configuration
 - a Latency
 - b Bandwidth

The goal of hardware sizing is to balance hardware costs with user response times. To do this effectively, you need to accurately estimate and plan for both peak and average system load. Peak load is the load on the system during peak times. For example, users may access the system heavily between 9:00 AM and 10:00 AM, then again between 1:00 PM and 2:00 PM. Average load is determined by calculating load during all periods and averaging it.

If the peak load occurs on a regular basis, such as, daily or weekly, it would be ideal to configure and tune systems to meet the peak load requirements. Those users who access the application during non-peak times would experience better response times than the peak-time users. If peak load times are infrequent or do not deviate much from average load and higher response times during peak usage is acceptable, then the system can be configured and tuned to average load. This leads to a decrease in hardware investment at the cost of higher response times during infrequently high server load.

Another major factor that needs to be considered for hardware sizing is the average wait time between actions or clicks for a given user. The average wait time can vary from one second to 15 seconds to several minutes, depending on how the user uses the system. The user spends time on analyzing or reading data received between transactions and performing other tasks such as reading email, using the telephone, and chatting with a colleague. All of these actions contribute to the average wait time between actions performed in the Agile system.

The Transaction Processing Performance Council (<http://www.tpc.org>) that publishes the benchmarks for different applications, recommends a wait time of 7 to 70 seconds between subsequent actions. For sizing calculations, the average wait time must be considered. The lower the average wait time, the smaller the number of users the server can support.

Agile Application Server

Agile PLM 9.2.1 supports Oracle Application Server 10g and BEA WebLogic 8.1 as the application server on Windows, Solaris, and Linux operating systems.

To determine the application server capacity, the average Transactions per second (TPS) the server can support for a given Agile module must be determined. For each module, business scenarios were identified that users with different roles would perform daily. Based on these scenarios and the user distribution, the workload is designed per module.

In the first phase, tests were conducted on individual modules to determine the TPS. A single, dual-CPU application server supported an average three second response time. This TPS became the base calculation.

The TPS for the Agile PLM modules is as follows:

Module	Transactions per Second (TPS)
Product Collaboration (PC)	19
Product Cost Management (PCM)	12
Product Portfolio Management (PPM)	16
Product Quality Management (PQM)	19
Product Governance and Compliance (PGC)	17

With Product Collaboration, the test was conducted for five and seven second wait times with a test duration of 10 minutes. The calculation and results are as follows:

TPS=19

Total duration=60 x 10 = 600 seconds

Total transactions expected=600 x 19 = 11400

Wait time of five seconds and an average response time of three seconds, the user time = 3+5 = 8 seconds

Total transactions one user can perform during 10 minutes of test time = 600/8 = 75

Total number of users that can be loaded = 11400/75 = 152

Performing the tests with 160 users with a wait time of five seconds between clicks yielded the following results:

Total transactions: 11970

Average response time: 3.03 seconds

With a seven second wait time, the results are as follows:

Wait time of seven seconds and an average response time of three seconds, the user time = 3+7 = 10 seconds

Total transactions one user can perform during 10 minutes of test time = 600/10 = 60

Total number of users that can be loaded = 11400/60 = 190

Performing the tests with 190 users with a wait time of seven seconds between clicks yielded the following results:

Total transactions: 11754

Average response time: 3.15 seconds

In order to get a uniform scale across the modules listed in the previous table, a normalization factor is calculated with Product Collaboration as the base module.

Module	Transactions per Second (TPS)	Normalization Factor
Product Collaboration (PC)	19	1
Product Cost Management (PCM)	12	0.63
Product Portfolio Management (PPM)	16	0.84
Product Service and Improvement (PSI)	19	1
Product Governance and Compliance (PGC)	17	1

For example, if you have multiple modules and the users are distributed across the modules as follows:

PC: 100 users

PPM: 40 users

PCM: 30 users

Normalizing the user load based on the total users = $100/1 + 40/.84 + 30/.63 = 195.24$

Hardware Sizing for Windows and Linux

Table 2-1: Recommended hardware sizing for Windows and Linux based on user load

Number of Servers*	User Load	
	Low Wait Time**	High Wait Time***
1	160	250
2	290	450
3	420	675
4	550	900
* Application Server Details: Dual CPU Pentium Xeon, 3GHz and above, 2MB L2 cache, 2GB RAM ** Assuming an average wait time of 5 seconds between clicks, the number of users supported. *** Assuming an average wait time of 10 seconds between clicks, the number of users supported.		

To support multiple application servers, clustering must be implemented which adds an additional 10% load on the cluster.

Hardware Sizing for Solaris

Table 2-2: Recommended hardware sizing for Solaris based on user load

Number of Users	Processor Type	Number of Servers	Number of Containers* per Server	Number of CPUs
150	T1	1	1	1 x 1 = 1
	IIIi	1	1	1 x 1 = 1
300	T1	1	2	1 x 1 = 1
	IIIi-	2	1	2 x 2 = 4
500	IIIi	3	1	3 x 2 = 6
	IV+	1	1	1 x 4 = 4
750	IIIi	5	1	5 x 2 = 10
	IV+	1	2	2 x 4 = 8
1000	IIIi	7	1	7 x 2 = 14
	IV+ (v490)	2	2	2 x 4 = 8
	IV+ (v890)	1	3	1 x 8 = 8
1500	IV+ (v490)	2	2	2 x 4 = 8
	IV+ (v890)	1	4	1 x 8 = 8
1500+	IV+ (v490)	2 servers for the first 1500 users + 1 server for every additional 750 users		
	IV+ (v890)	1 server for the first 1500 users + 1 server for every additional 1500 users		

* Containers are only applicable for Solaris 10

Table 2-3: Server Configurations with Processor Type

Processor Type	Server Model	Server Details
T1	T2000	UltraSparc T1 Processor, 1 CPU - 8 Core, 8 GB RAM, 1.2 GHz
IIIi	v240	UltraSparc IIIi Processor, 2 CPU, 8 GB RAM, 1.5 GHz
IV+ (v490)	v490	UltraSparc IV+ Processor, 4 CPU - 8 Core, 16 GB RAM, 1.5 GHz, L2 32 MB
IV+ (v890)	v890	UltraSparc IV+ Processor, 8 CPU - 8Core, 16 GB RAM, 1.5 GHz, L2 32 MB

T1 processors are slower than IIIi processors, but there are cost advantages with T1-based servers. For customers with up to 300 users, T1-based servers are recommended with the caveat that if there are more than 150 users, containers must be configured. Benchmark tests of different Agile components have shown that T1 processors are 50% slower than IIIi processors running the same load. Agile does not recommend setting up more than 2 containers on T1-based servers, as internal tests have shown that no additional benefit is achieved. The cost of a T1-based server (T2000) is approximately the same as the IIIi-based v240 server. So, for 300 users, a T2000 is more cost-effective.

IV+-based servers perform and scale better than IIIi processors, but are more expensive. For 750 users, 5 IIIi-based v240 servers are required compared to 1 IV+-based v490 server. But, on a cost basis, 5 v240 servers cost less than the single v490 server. Also, having multiple servers would be advantageous should one of the servers fail. With the v490 server, there is only a single point of failure. The advantage of a v490 server would be the possible future expansion to a few thousand users on your system. A similar case can be made when comparing v490 and v890 servers.

Database Server Sizing

For production environments, it is recommended to run the database server on dedicated hardware. Database hardware sizing is dependent on both concurrent usage and the amount of data or size of the database. The best measure of database size is schema dump file size and estimated monthly incremental increases. Exporting the Agile schema at periodic intervals and analyzing its size helps you determine if a larger database sizing model is needed to better manage database growth, as well as minimize ongoing database maintenance and tuning.

Oracle Database Sizing

For existing Agile customers, getting the initial dump file size as a baseline is easy. For new customers, the dump file size needs to be estimated. If there is an existing database, use the Oracle Export Utility to verify the dump file size. If there is no existing database to reference, the size of the database needs to be estimated by monitoring database growth over the first few months of normal operation to predict future disk size needs.

Table 2-4: Agile PLM 9.2.1 Database Sizing Matrix for Oracle





Agile DB Configuration	Number of Users	Database Server Hardware											
		CPU	RAM	# of disk	CPU	RAM	# of disk	CPU	RAM	# of disk	CPU	RAM	# of disk
D	1000	8	8GB	9	12	12GB	9	16	16GB	13	24	24GB	15
C	500	4	4GB	4	8	8GB	9	8	8GB	11	12	12GB	13
B	250	4	2GB	4	4	4GB	4	4	4GB	9	Extra Large		
A	100	2	1GB	4	2	2GB	4	Large					
Database Sizing Model		Small			Medium								
													

Table 2-5: Oracle Database Sizing Model

Size	Initial Dump File Size (MB)	Monthly Increment (MB)
Small	<1024	<50
Medium (Regular)	<5120	<200
Large	<16384	<400
Extra Large	<38912	<1000

Each database sizing model requires an initial database configuration for deployment. For scalability and concurrency support, you need additional hardware resources, such as CPU, RAM, and number of disks.

Small Database

The Agile PLM 9.2.1 small database sizing model can be used in a demo or test environment with the minimum hardware requirements. In a production environment, a small database (default) requires the settings of configuration A as an initial configuration. Configurations B, C, and D can be used for scalability and the addition of more concurrent users.

Medium Database

The Agile PLM 9.2.1 medium database sizing model requires configuration A as an initial configuration with additional RAM. Configurations B, C, and D can be used for scalability and the addition of more concurrent users.

Large Database

The Agile PLM 9.2.1 large database sizing model requires configuration B as an initial configuration. Configurations C and D can be used for scalability and the addition of more concurrent users.

Extra-large Database

The Agile PLM 9.2.1 extra-large database sizing model requires configuration C as an initial configuration. Configuration D can be used for scalability and the addition of more concurrent users.

CPU and Memory

As you can see from Table 2-4, the Agile PLM 9.2.1 database, CPU, and memory requirements are roughly the same as with previous versions of Agile. With further improvement on bind variables and optimization of SQL, memory resource should be primarily used for the DB cache (or buffer), which is directly proportional to the amount of data. For servers running Windows, the minimum recommended CPU is an Intel 2.8 GHz Xeon with 512 KB L2 cache.

Storage

It is recommended to start with a 4-disk configuration. The starting disk space requirement, 4x18GB=72GB, may seem quite large when comparing it to the size of the initial dump file, but considering the storage needs of the Agile PLM 9.2.1 features, including full text search, this may actually be on the low side.

Network Card

The database should have a 100Mbps network card.

Sizing Summary

RAM sizing is directly related to database size. Storage sizing should be based on the number of spindles. The Agile PLM 9.2.1 database utility can be used to optimize up to 15 separate disks.

Hardware Resource Plan for Database Models

The following table lists recommended hardware resources for different database size models.

Table 2-6: Recommended hardware resources for different database sizes

Database Size	CPU	RAM	Disks *
Demo	1	512 Mb	1
Small	2	1 Gb	4
Medium	2	2 Gb	4
Large	4	4 Gb	8
Extra-Large	12	8 Gb	12

* Each disk has 18 Gb disk space.

Disk Space and Tablespace Configurations

While the proper sizing of extents minimizes dynamic extensions in the same segments, disk I/O contention within the same logical tablespace or physical data file can also be harmful.

You can improve disk I/O performance for multiple disk configurations by spreading the I/O burden across multiple disk devices. The following sections describe the use of multiple disks for the Oracle database server. It is always advisable to use more disks.

One-Disk

A one-disk configuration is best for a demonstration, preproduction, and testing environment. This configuration results in disk I/O contention. In addition, as both usage and database size increase, performance significantly declines. The one-disk configuration is intended for demo database applications only, and the configuration can be implemented as shown in Table 2-1.

Table 2-1: One-disk configuration for OFA implementation

Disk	Oracle_Home	Tablespaces	Redo Logfiles
Disk 1	ORACLE_HOME	SYSTEM TOOL UNDO TEMP USERS INDX AGILE_DATA1 AGILE_DATA2 AGILE_DATA3 AGILE_DATA4 AGILE_INDX1 AGILE_INDX2 AGILE_INDX3 AGILE_INDX4	LOG1 LOG2 LOG3 LOG4

There is no beneficial gain from OFA for the one-disk configuration from the perspective of disk I/O contention. There should be no significant impact on a current production database if you implement the default Oracle settings with a one-disk configuration.

Two-Disk Configuration

A two-disk configuration is best for a small database. To eliminate potential I/O contention, AGILE_DATA and AGILE_INDX data files are on separate disks. As usage and database size increases, performance declines.

Table 2-2: Two-disk configuration for OFA implementation

Disk	Oracle_Home	Tablespaces	Redo Logfiles
Disk 1	ORACLE_HOME	SYSTEM TOOL UNDO AGILE_DATA1 AGILE_DATA2 AGILE_DATA3 AGILE_DATA4	LOG1 LOG2
Disk 2		TEMP USERS INDX AGILE_INDX1 AGILE_INDX2 AGILE_INDX3 AGILE_INDX4	LOG3 LOG4

Four-Disk Configuration

A four-disk configuration is best for an enterprise-level implementation of Agile. A four-disk configuration spreads the various data files, control files, and redo log files across multiple disk devices.

- ❑ The three control files can be mirrored onto three different disks for best recovery protection.
- ❑ All potential I/O demand-intensive data files can be distributed onto their own separate disk. Redo log files are completely isolated from the rest of the data files, as the log files can cause significant I/O contention during transactions if they are sharing disks with other data files. The UNDO data file is separated from the schema data files and log files as well, so I/O contention can be minimized.
- ❑ The Agile schema tablespaces can be isolated from the rest of the SYSTEM, TEMP, TOOL, and UNDO data files.

The four-disk configuration shown in Table 2-3 is recommended. For production database sites, the four-disk configuration represents the minimum requirements for an OFA implementation and provides the minimum hardware configuration for performance tuning.

Table 2-3: Four-disk configuration for OFA implementation

Disks	Oracle_Home	Tablespaces	Redo Logfiles	Control files
Disk 1	ORACLE_HOME	SYSTEM TOOL UNDO	LOG1/2/3/4	Controlfile01
Disk 2		TEMP USERS INDX	Archive log file	Controlfile02
Disk 3		AGILE_DATA1 AGILE_DATA2 AGILE_DATA3 AGILE_DATA4		Controlfile03
Disk 4		AGILE_INDX1 AGILE_INDX2 AGILE_INDX3 AGILE_INDX4		

Eight-Disk Configuration

In addition to the advantages associated with a four-disk configuration, an eight-disk configuration supports an enterprise-level implementation of Agile by further spreading various data files and redo log files across multiple disk devices.

Application schema can get additional performance gains in terms of I/O load spread by further separating the AGILE_DATA1, AGILE_DATA2, AGILE_DATA3, and AGILE_DATA4 and AGILE_INDX1, AGILE_INDX2, AGILE_INDX3 data files because of potential I/O contention between the AGILE_DATA data files and AGILE_INDX data files. A complete separation of potential large datafiles in its own disk spindle should help I/O contention as physical disk I/O is inevitable, due to the share size of data, as shown in Table 2-4.

Table 2-4: Eight-disk configuration for OFA implementation

Disks	Oracle_Home	Tablespaces	Redo Logfiles	Control files
Disk 1	ORACLE_HOME	SYSTEM TOOL UNDO	LOG1/2/3/4	Controlfile01
Disk 2		TEMP USERS INDX	Archive log file	Controlfile02
Disk 3		AGILE_DATA1		Controlfile03
Disk 4		AGILE_DATA2		
Disk 5		AGILE_DATA3		
Disk 6		AGILE_DATA4		
Disk 7		AGILE_INDX1 AGILE_INDX2		
Disk 8		AGILE_INDX3 AGILE_INDX4		

Twelve-Disk Configuration

Further separating the AGILE_DATA and AGILE_INDX tablespaces, twelve-disk configurations can be implemented as shown in Table 2-5. This results in complete independent spindles for AGILE_DATA1, AGILE_DATA2, AGILE_DATA3, and AGILE_DATA4 and AGILE1_INDX, AGILE_INDX2, AGILE_INDX3, and AGILE_INDX4.

Table 2-5: Twelve-disk configuration for OFA implementation

Disks	Oracle_Home	Tablespaces	Redo Logfiles	Control files
Disk 1	ORACLE_HOME	SYSTEM TOOL	LOG1/2/3/4	Controlfile01
Disk 2		USERS INDX	Archive log file	Controlfile02
Disk 3		UNDO		Controlfile03
Disk 4		TEMP		
Disk 5		AGILE_DATA1		
Disk 6		AGILE_DATA2		
Disk 7		AGILE_DATA3		
Disk 8		AGILE_DATA4		
Disk 9		AGILE_INDX1		
Disk 10		AGILE_INDX2		
Disk 11		AGILE_INDX3		
Disk 12		AGILE_INDX4		

Hardware Load Balancer and Proxy Web Server

A load balancer or proxy web server is deployed to protect the application server. When external users need access to Agile, this device is deployed in the DMZ. The load balancer or proxy web server does not need to be installed in the DMZ if Agile is only accessed internally from within the corporate firewall.

Note Load balancers can be used with the Java client and the Web client. Proxy web servers can only be used with the Web client.

Much like the application server, the dominant factor in determining hardware sizing for the proxy web servers is concurrent usage. Use the following table to determine the hardware needed for the web server tier.

Table 2-1: Web Server Sizing Matrix

Peak Logged In Users	Number of Servers	Number of CPUs	Memory (GB)
100	1	1	1
250	1	1	1
500	1	2	2
1000	2	2	2

CPU

For servers running Windows or Linux, the minimum recommended CPU is an Intel 2.8 GHz Xeon with 512 KB L2 cache.

Network Card

Each proxy web server should have a 100Mbps network card for up to 500 peak logged in users. For more than 500 logged in users, each server should have a 1Gbps network card.

Distributed File Management

The performance of Distributed File Management is a function of how many files are being downloaded or uploaded concurrently along with how large the files are. A site handling up to 100 logged in users requires a server with two CPUs running the processors previously mentioned and 2GB of RAM. File vault storage size is a function of the expected amount of data to be stored there.

View Server

View Server performance is a function of how many files are being viewed concurrently and the average file size being viewed. One server with dual CPUs as specified, and 2GB of memory is required for a server handling up to 100 users. Memory requirements are a function of file size. For example, 50 people each viewing a 1MB file requires 50MB of free memory.

Breakpoint Testing Results

Solution Results on Windows and Linux

The following tables show the results of breakpoint testing on the Product Collaboration, Product Quality Management, Product Portfolio Management, Product Cost Management modules, and the Distributed File Manager and Export components. The testing was initially performed with Agile PLM running on Windows 2003 systems, but follow-up testing resulted in the same findings on RedHat Linux version 4 systems.

Note The response time for all of the tables is measured in seconds.

Product Collaboration

Change Object

Table 3-1: Search and Add Affected Items

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
500	75	10%	5%
2000	660	20-25%	5%
3000	Timed out after 16 minutes	25%	5%

Table 3-2: Save As Change

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
750	24	10-15%	10%
2500	96	20-25%	10-15%
5000	298	20-25%	<5%
7500	765	20-25%	<5%
10000	Timed out after 15 minutes	20-25%	<5%

Table 3-3: Change Status to Released (With Preliminary Items)

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
750	105	15-20%	5%
2500	430	20-25%	5-10%
5000	1140	20-25%	5-10%
7500	Timed out after 17 minutes	20-25%	5-10%

Table 3-4: Change Status to Pending

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
750	335	20-25%	5-10%
2500	1265	20-25%	5-10%
5000	Timed out	20-25%	5-10%

Table 3-5: Change Status to Released (Each Item has at least 3 Pending Changes)

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
750	460	20-25%	5-10%
1500	900	20-25%	5-10%
>1500	Timed out	20-25%	5-10%

Item Object

Table 3-6: Search and Add BOM

Number of Items	Response Time	CPU Usage	
		Application Server	Database
500	23	<5%	<5%
2500	261	20-25%	<5%
5000	705	20-25%	<5%
7000	920	20-25%	<5%
7500	Timed out	20-25%	<5%

Table 3-7: Expand BOM - All levels

Number of Items	Response Time	CPU Usage	
		Application Server	Database
>7000 rows	8 minutes	20% for the first 25 seconds	Clean
Client browser consumed a large amount of memory that placed Internet Explorer in a non-responsive state. On some machines, incremental rendering did not work. In this case, the browser hung when the Expand All button was clicked which led to 100% CPU usage on the client machine.			

Product Quality Management

During the breakpoint testing for Product Quality Management, the table sizes used were as follows:

- ❑ PSR - 23908 items
- ❑ QCR - 23525 items

PSR Object

Table 3-8: Search and Add Affected Items

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
2500	610	15-20%	<5%
3500	Timed out after 16 minutes (added 3,384 affected items)	10-15%	<5%

Table 3-9: Change Status to Released

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
10000	1 minute	<10%	<5%

QCR Object

Table 3-10: Search and Add Affected Items

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
2500	Waited for 25 minutes (added only 1248 affected items)		

Table 3-11: Change Status to Released

Number of Affected Items	Response Time	CPU Usage	
		Application Server	Database
10000	1 minute	<10%	<5%

Product Portfolio Management

Table 3-12: PPM Scenarios

Scenarios	1500 Object Tree	3000 Object Tree	6000 Object Tree
Save As Program	20	32	71
Save As Template	11	23	50
Create Program From Template	22	45	96
Add 10 Team Members	14	26	50
Reschedule Root Program	7	12	25
Reschedule Child	5	7	14
Delegate Owner	3	3	4
Substitute Resource	13	23	45
Change Status	6	11	22
View Gantt Chart	27		
Add Dependencies in Gantt			

Tree Details: 6000 objects, 3 levels

- 10 Team Members at each level
- 10 Links

- 10 Deliverables
- 10 External Dependencies

Distributed File Manager

The breakpoint testing for distributed file manager tested file upload and download on the distributed file management server. Basic file upload and download was tested on both supported application servers without a web proxy. Advanced file upload and download was tested using the Java Client on both supported application servers with a web proxy.

Table 3-13: Basic File Upload and Download

Scenario	Break Point	Response Time	CPU Usage			Issue/ Exception
			Client	Application Server	DFM Server	
Upload	> 2GB	5 minutes, 5 seconds	<10%	<3%	8%	MIME limitation
Download	> 3GB	11 minutes, 5 seconds	<15%	<3%	<5%	Stopped responding after downloading a 3GB file

Table 3-14: Advanced File Upload and Download

Scenario	Break Point	Response Time	CPU Usage		
			Client	Application Server	DFM Server
Upload	Tested up to 80GB	170 minutes	<10%	<3%	8%
Download	Tested up to 80GB	175 minutes	<15%	<3%	<5%

Export

Export was tested using the Web Client on Oracle Application Server.

Table 3-15: Export scenarios

Scenario	Break Point	Response Time	CPU Usage		Issue/Exception
			Application Server	Database	
Export - CSV	>16000	15 minutes	20%	<5%	Timed out after 15 minutes
Export - PDX	16000	1 minute	20%	<5%	

BOM Structure:

- ❑ Size of BOM: 28,000
- ❑ Total Levels: 5
- ❑ Assembly + SubAssembly + Items: 3,600
- ❑ Manufacturer Parts: 24,400

Product Cost Management

Action	Project with 2500 Items	Project with 5000 Items	Project with 10000 Items
Display project item tab	2	2	2
Display project expanded item tab	-	-	-
Import item from PC item master	55	200	220
Update item from PC item master	50	75	236
Add/modify Partners	9	11	33
Import item from external source	39	47	193
Export item to external source	22	40	113
Quantity rollup	50	-	Application Error after 9 minutes
Display project AML tab	2	2	2
Display item BOM	2	2	2
Display of project analysis	2	3	3
Opening Responses Look Up	58	138	340
Project response price lookup	11	7	221
Set as best	1	3	3
Publish Quote History, Authoring	60+62	132+110	347+273
Publish Published Price Redlining	60+58	132+105	347+266
Items RFQ Creation	16	16	35
Add Partners (2)	9	11	33
Open RFQ	33	40	128
Display of RFQ Responses	2	2	2
Response Status Tab	3	3	3
Opening RFQ Response Look Up	89	136	Application Error after 9 minutes
RFQ Response Price Look Up	4	7	Application Error
RFQ Response export	235	340	Application Error after 8 minutes
RFQ Response import	-	-	-
RFQ Response status display	2	2	2
Edit/Modify Responses	2	2	2
Supplier Response Export	75	110	420
Supplier Response Import	58	72	232
Submit Responses	55	-	-
View Affected Items	-	2	-

Action	Project with 2500 Items	Project with 5000 Items	Project with 10000 Items
Add Affected Items		440	
Redline updates		2	
Release PCO			
Costed BOM Comparison (Buyer) Report	3	9	23
Costed BOM Comparison (Buyer) Report Export	64	148	369
Save as	29	38	192
Delete	1	1	1
Close RFQ	50	55	267
Delete RFQ	3	4	15
Lock RFQ	55	60	279

Solution Results on Solaris

The following tables show the results of breakpoint testing on the Product Collaboration, Product Quality Management, Product Portfolio Management, Product Cost Management modules, and the Distributed File Manager component.

Breakpoint testing on the Solaris platform was performed on the following configuration:

- ❑ Application Server: v490, IV Processor, 4 CPU - 8 Core, 16GB RAM, 1.35 GHz, 8MB L2 Cache
- ❑ Database Server: v890, IV Processor, 8 CPU, 16GB RAM, 1.35 GHz, 8MB L2 Cache

Note The response time for all of the tables is measured in seconds.

Product Collaboration

Change Object

Table 3-16: Search and Add Affected Items

Number of Affected Items	Response Time
2500	865
4500	Waited for 35 minutes

Table 3-17: Save As Change

Number of Affected Items	Response Time
2500	65
10000	800
13000	1440

Table 3-18: Change Status to Released

Number of Affected Items	Response Time
2500	63
7500	330
13000	Database errors occurred after 15 minutes

Table 3-19: Change Status to Pending

Number of Affected Items	Response Time
2500	200
7500	855

Item Object

Table 3-20: Search and Add BOM

Number of Items	Response Time
2500	100
5000	200
10000	513
13000	840

Product Quality Management

PSR Object

Table 3-21: Search and Add Affected Items

Number of Affected Items	Response Time
2500	730
4500	Timed out after 15 minutes

QCR Object

Table 3-22: Search and Add Affected Items

Number of Affected Items	Response Time
2500	Waited for 25 minutes (added only 1243 affected items)

Product Portfolio Management

Table 3-23: PPM Scenarios

Scenarios	3000 Object Tree	6000 Object Tree
Save As Program	36	105
Save As Template	22	80
Create Program From Template	42	163
Add 10 Team Members	22	80
Reschedule Root Program	11	35
Reschedule Child	4	10
Delegate Owner	3	8
Substitute Resource	4	55
Change Status	6	17
Gantt Chart Load	18	Out of memory

Tree Details: 6000 objects, 3 levels

- 10 Team Members at each level
- 10 Links
- 10 Deliverables
- 10 External Dependencies

Distributed File Manager

The breakpoint testing for distributed file manager tested file upload and download on the distributed file management server. Basic file upload and download was tested on both supported application servers without a web proxy. Advanced file upload and download was tested using the Java Client on both supported application servers with a web proxy.

Table 3-24: Basic File Upload and Download

Scenario	Break Point	Response Time	CPU Usage			Issue/ Exception
			Client	Application Server	DFM Server	
Upload	> 2GB	5 minutes, 5 seconds	<10%	<3%	8%	MIME limitation
Download	> 3GB	11 minutes, 5 seconds	<15%	<3%	<5%	Stopped responding after downloading a 3GB file

Table 3-25: Advanced File Upload and Download

Scenario	Break Point	Response Time	CPU Usage		
			Client	Application Server	DFM Server
Upload	Tested up to 20GB	170 minutes	<10%	<3%	8%
Download	Tested up to 20GB	175 minutes	<15%	<3%	<5%

Product Cost Management

Price Change Order Object

Table 3-26: Search and Add Prices

Number of Prices	Response Time
1000	890
1500	Timed out

Table 3-27: Release PCO

Number of Prices	Response Time
1000	78

Table 3-28: Export Prices

Number of Prices	Response Time
100	80
1000	510
1500	Out of memory after 11 minutes

PCM Scenarios

Action	Response Time for 5000 Objects	Response Time for 10000 Objects
Display project item tab	2	2
Import item from PC item master	77	225
Update from PC item master	97	318
Add/modify Partners	30	60
Import item from external source	60	156
Export item to external source	52	144
Quantity rollup	27	51
Display project AML tab	2	2
Display item BOM	2	2
Display of project analysis	2	3
Opening responses lookup	2	2
Project response price lookup	7	7
Set as best	3	3
Publish Quote History, Authoring	45+41	160+72
Publish Published Price Redlining	40+46	100+75
Items RFQ Creation	32	95
Add Partners	30	60
Open RFQ	2	118
RFQ Responses Tab	2	2

Action	Response Time for 5000 Objects	Response Time for 10000 Objects
Response Status Tab	2	4
Opening RFQ Response lookup	2	10
RFQ Response price lookup	42	16
Export RFQ Response Lines	Waited for 30 minutes	Not Responding
Import RFQ Response Lines	-	-
RFQ Response status display	2	2
Edit/Modify Responses	400	3
Supplier response export	62	120
Supplier response import	74	164
Submit responses	3	4
Costed BOM Comparison (Buyer) Report	15	37
Costed BOM Comparison (Buyer) Report Export	-	-
Save as	28	44
Delete project	2	2
Close RFQ	260	668
Delete RFQ	26	30
Lock RFQ	280	687

Portlets Performance

Based on the same configurations of Agile PLM on both Windows and Solaris, the performance of portlets is three times slower.

Client Requirements

Client Requirements

Clients require desktop class computers supported by the browsers listed in Table 1-1. For clients running the Windows operating system, the minimum recommended CPU is a Pentium 4 with 100MB of available RAM. Client monitor resolution should meet or exceed 1024x768 and should support at least 256 colors.

For effective operation, Agile recommends at least 250Kb/sec transfer rate for each Java client and 1000Kb/sec transfer rate for each Web client. The Web client requires more bandwidth than the Java client because it is a thin client. Thin clients depend on the server not only for data, but also formatting information, which increases the total amount of information transferred. See “Network Bandwidth” in Chapter 6 for information on how to determine how much bandwidth is needed for a given system.

Although the two clients are expected to perform similarly under ideal network conditions and because the Java client is more sensitive to network latency, performance of the Java client and Web client is expected to diverge as network latency increases. For this reason, Agile recommends using the Java client within corporate firewalls with low network latency.

A Java client user is expected to consume slightly less application server resources than a Web client because the Java client is able to use the desktop client for some of the work. Sorting data tables is a good example of how the Java client differs from the Web client. The Web client must ask the server to resort the table after selecting a new sort order. The Java client is able to reformat the table by itself.

The Java client connects directly to the application server using application server specific native protocols, such as ORMI with Oracle Application Server or T3 with BEA WebLogic.

CHAPTER 5

Directory Server

Overview

In an effort to support industry standard authentication schemes and enable central management of user information, Agile PLM 9.2.1 supports integration with industry leading directory server and LDAP-based authentication for the Agile solution suite. Directory server integration enables you to seamlessly integrate the Agile solution suite with your existing directory servers, so users can be managed in one place. Agile PLM 9.2.1 supports Microsoft Active Directory Server and Sun Java System Directory Server.


This chapter describes the architecture of the Agile PLM 9.2.1 Directory Server Integration Module, implementation details, configuration, and contains a FAQ section to help you gain a better understanding of how this module works.

System Architecture

Figure 5-1 illustrates the high-level architecture of the Agile PLM 9.2.1 Directory Server Integration Module. As shown in the diagram, the Directory Server Integration Module is a platform component of the core Agile Application Server that is used by all components in the Agile solution. This module provides three main services to the Agile solution:

- 1 Authentication
- 2 Obtaining up-to-date user listings (synchronize users)
- 3 Synchronizing user profiles

Figure 5-1: LDAP Architecture



The Directory Service Integration Module is an interface that describes the contract between the module and the Agile PLM 9.2.1 platform. The interface does not make any assumptions as to how the implementation provides these services. This interface driven design de-couples the Agile PLM 9.2.1 platform from user management and authentication, allowing new implementations to be easily added in the future for expanded support.

User information that is managed through the Directory Server Integration Module must be maintained by the source. If a directory server is selected to manage user accounts, then user accounts must be managed through the directory server, not through Agile.

New users must be added to the directory server and specific user attributes such as user ID, password, email, phone, etc., must be managed through the directory server. Only if the Agile database is used, all user information can be managed through the Agile Java Client. There is one exception to this rule; all supplier users must be managed through Agile, regardless of whether directory server integration is chosen.

Implementation Details

Three implementations of this interface are provided in Agile PLM 9.2.1. These implementations are depicted as *Adaptors* in Figure 5-1. An adaptor provides an implementation of the interface that is specific to a particular external service, typically, a directory server. For example, the ADS Adaptor provides an implementation specifically designed to work with Microsoft Active Directory Server. The Directory Server Integration Module enables you to outsource, from Agile, the functionality of managing user accounts and authenticating users through their choice of directory server.

Directory servers vary greatly in terms of the features they offer and the information they provide. Therefore, the Agile Directory Server Integration Module makes minimum assumptions about the services offered by the directory servers and complies with industry standards. At a minimum, the directory server must comply with the LDAP standards and support attributes mandated by InetOrgPerson schema from the LDAP standard. The following table shows these attributes, as reflected in the Agile User Profile:

Table 5-1: InetOrgPerson Schema for Agile User

LDAP Attribute	Agile Database Column	Example
sn	AGILEUSER.LAST_NAME	Smith
givenName	AGILEUSER.FIRST_NAME	Joe
title	AGILEUSER.TITLE	Manager, Product Development
uid	AGILEUSER.LOGINID	Jsmith
mail	AGILEUSER.EMAIL	jsmith@company.com
telephoneNumber	AGILEUSER.WORK_PHONE	+1 408 555 1862
facsimileTelephoneNumber	AGILEUSER.FAX	+1 408 555 1992
mobile	AGILEUSER.MOBILE	+1 408 555 1941

The Agile PLM 9.2.1 database still contains the full list of users and other vital user information needed by the application. However, if a directory server is used, the previously listed attributes, as well as the password, are managed by the source only and displayed in Agile as read-only attributes.

The Directory Server Integration Module can be configured at install time or later. Configuration settings are found in the Agile Java Client. The following section discusses each configuration option and how it affects the system.

Configuration

The Directory Server Integration Module is intended to be flexible, yet simple to use. The module provides the following configuration parameters to control system behavior.

Directory Service Connection Parameters

Connection parameters include the hostname, port, protocol, account name and filter. The account name is used to connect to the directory server during synchronization; therefore, it must have the appropriate privileges. The filter is used to select only a subset of the users defined in the directory server as Agile users.

Multiple Directory Server Support

It is possible to define multiple sets of connection parameters to configure integration with multiple directory servers. This may be useful if you have users in multiple domains that need access to Agile or if you have backup directory servers to provide fail-over support. It may be necessary to configure a separate directory server to manage Agile users who are not employees. If a backup or secondary directory server is configured, the authentication module tries the backup server if access to the primary server fails.

Schedule Synchronization

The time interval to synchronize users with Agile is set in the Task Configuration node on the Admin tab of the Agile Java Client. During synchronization, any newly created users are added to Agile and any modified user attributes are synchronized. Note that all newly created users are disabled in the Agile database until they are enabled through the Agile Java Client.

On-Demand Synchronization

In addition to scheduling synchronization, it is possible to synchronize user account information on demand through the Agile Java Client.

Utilities

The upgrade script, `migrateUsersToDB`, is a command line script used to migrate all users from a directory server to the Agile database. The script applies the same rules that are applied when synchronizing. User records that are not matched remain active in the database. They are not deleted or disabled, but note that those accounts can not be used for authentication.

One directory server can be configured during the Agile installation. Additional directory servers can be configured manually after installation. Agile provides two scripts to enable configuration after installation:

- 1 `encryptpwd`—`Ldapconfig.xml` needs an encrypted password for the directory server administrator user. This script generates it based on the existing administrator password.
- 2 `checkLDAPConfig`—Use for checking LDAP configurations. All errors should be fixed, if encountered.

Frequently Asked Questions (FAQ)

What happens when I delete a user from Agile?

The user is not automatically deleted from the directory server. On the other hand, the user is not able to log in to Agile. Within Agile, the user appears on the **Deleted Users** page and can be undeleted from there.

What happens when I delete a user from the directory server?

The user is not automatically deleted from Agile. On the other hand, the user is not able to log in to Agile. When synchronizing user profiles, currently only updates and creates are considered.

Does the Agile server allow login ID (user ID) changes?

Agile allows login ID changes only for integrations with directory servers. To do this, change the user ID in the directory server and synchronize the user within the Agile server.

Note Supplier user's user ID cannot be modified because they are only managed in Agile.

My directory server provides a feature called “Activate/Inactivate.” How does this relate to “Enable/Disable” within Agile?

This does not, in any way, affect the enable/disable functionality within Agile. If a user has been “Inactivated” in the directory server or disabled in Agile, they cannot log in to the Agile system.

I want to create a user in the directory server and log in to Agile immediately. How can I do this?

This can be done through on-demand synchronization. On-demand synchronization immediately synchronizes with the directory server. Note that newly created users are disabled by default. You must assign proper roles and privileges, then enable the user before they can log in.

Can I still create a user from within Agile? How does it reflect in the directory server?

If you choose to integrate with a directory server, only supplier users can be created in the Agile database. All other users must be managed through the directory server.

CHAPTER 6

File Management Server and View Server

Agile File Management Server

There are two main components to the Agile File Manager: the file server and the file vault. When a file is added to Agile, it is assigned an internal Agile identifier (ID) number by the file server and added to the file vault. It is not stored in the file vault under its original file name. The filename/Agile ID mapping information is maintained in the Agile database. When a user requests a file (Get, View, Checkout, etc.), that request is routed to the file server which looks up the file's ID, retrieves the file from the file vault, and sends it to the user.

File Manager uses standard HTTP protocol to communicate with the Agile Application Server, View Server, and any other deployed File Managers. File Managers can be clustered in one location, if required.

The features of the Agile PLM 9.2.1 File Manager are:

- ❑ The File Manager is not bundled as a part of the Application server, but it is deployed as a separate process.
- ❑ File Managers can be deployed at any location with no Internet domain restrictions.
- ❑ File Managers are firewall friendly with all communication using the HTTP protocol.
- ❑ Improvements have been made to LAN/WAN performance.
- ❑ Support for 70+ File Managers
- ❑ A logging feature has been added to help with troubleshooting.
- ❑ Support for large file upload and download.

There can be any number of File Managers deployed, based upon the different geographical locations in which the user base resides. Typically, the File Manager that resides along with the Application Server is designated as the Primary File Manager. All File Managers are equal, except for the following additional transactions that are limited to the Primary File Manager:

- ❑ Uploading files by the Microsoft Project Sync Integration component during publishing.
- ❑ Downloading files for Full Text Search indexing

File Vaults

The file vault contains all file attachments stored in Agile. The File Manager supports two kinds of file vaults:

- 1 Standard Vault
- 2 Custom Vault

Custom vaults simplify the initial setup allowing you to attach a disk containing terabytes of data as a vault without actually uploading the files. This setup is performed using Agile FileLoad to set up the necessary database entries to avoid uploading the actual data.

Each File Manager supports cascading of multiple vaults. The primary vault contains all new uploaded files and redlines while the secondary vault contains the older files. When locating files, each vault is searched in a cascaded manner as configured in the Java client.

Distributed File Manager

Due to the geographically dispersed nature of the global enterprise, multiple Agile File Managers can be deployed in a distributed configuration for efficient distribution of product content. A Distributed File Manager configuration allows you to manage files efficiently at remote locations. Deploying a distributed file manager reduces download time by placing Agile files closer to where they are needed, allowing users to configure which file manager to use.

Using The Agile View Server in a Distributed Environment

If the Agile Viewer is used, an Agile View Server should also be installed locally with each distributed file manager. The local View Server can be installed on the same machine as the file manager. If local users are accessing Agile from outside the firewall, a proxy is recommended in the distributed configuration.

How Distributed File Manager Works

Agile File Managers have a peer-to-peer relationship. When a user requests a file, the request is directed to their configured file manager. If the file is found, it is served to the user. If the file is not found at that location, the file manager obtains the list of servers that have the file, then tries to retrieve the file from the closest File Manager. The file is then saved to the local vault and served to the user.

Agile recommends deploying the File Manager and View Server on a single two CPU machine with the following minimum specifications:

Processor: Intel Xeon, 2.8 Ghz, 512KB L2 cache, 2GB RAM

Disk Space: 50MB of available disk space

An example of a distributed file management scenario might be as follows: the Agile Application Server, Primary File Manager, and View Server are installed in San Jose, California and a File Manager is installed in Tokyo, Japan. A user from the Tokyo site selects a file attachment to view. The File Manager in Tokyo determines that the file is located in San Jose and downloads the file, copies it to the file vault in Tokyo, and serves the file to the user through the View Server or directly, depending on whether the View Server is configured for the given file type. Another user selects the same file attachment. Because the file now resides in the Tokyo File Manager vault, it is served directly to the user.

Do I need distributed file management?

If your company has several sites or locations, then you can consider using distributed file management. In general, the following criteria can help you to determine if you need additional file managers:

- ❑ Multiple remote locations—If you have multiple remote sites connected over a WAN, clients at remote locations who need to store and retrieve files from the main file server can experience large time delays.
- ❑ Reducing inter-network traffic—If your network is divided into subnets to reduce network traffic, then you should consider putting a distributed file manager in each subnet that contains Agile users to reduce your intersubnet network traffic.

Can I synchronize my distributed File Managers?

Yes. While it is not necessary in most cases to synchronize the distributed file management servers, there are some customers whose business processes warrant the need to synchronize. If you decide that synchronization is necessary, there are several utilities you can use (Agile does not recommend a specific utility.):

- ❑ Robocopy Utility—delivered with the Microsoft NT and Win2K Resource Kits. Visit www.microsoft.com for more information.
- ❑ XXCopy—visit www.xxcopy.com for more information.

Do I need to back up all of my distributed File Managers?

Yes, Agile recommends that you backup all file vaults.

Do I need a View Server for every distributed File Manager?

Yes. If you decide to use the View Server, then it is recommended that you install a View Server at every location where you have a File Manager installed. By installing a View Server with each File Manager, you reduce network traffic.

Should I install Anti-virus software on my file vault?

Yes, you are responsible for protecting all Agile systems and files managed in Agile from viruses.

File Management Security

There are two security considerations pertaining to the security of the File Manager: the server and client access.

Server Security

Content in the Agile file vault must be protected from deletion or modification by unauthorized users. Agile recommends allowing access only to administrators of Agile. System users who access files through the clients do not need Add and Read privileges to the Agile file vault or file directory because the Agile File Manager retrieves the files for the user, not individual users.

Client Access Security

Whether you are accessing files from the Java client or from the Web client, your files are secure. When the client is run from inside your firewall, the files are transferred behind the firewall, which secures your files from outside intervention. When the clients are accessing files from outside the firewall, Secure Sockets Layer (SSL) communications protocol is supported.

View Server

The Agile View Server is an optional component that allows users to view and mark up documents, drawings, and CAD files in supported Web browsers. By using the Agile Viewer, you can view files without having the native application that created the file installed on your desktop machines. Another advantage of the Agile Viewer is that it can decrease system bandwidth requirements. Instead of sending rather large native files to the user, the View Server sends a smaller image file that can be viewed using the View Server applet on the client. A Web browser downloads the applet from the File Manager and stores it in the browser cache. The next time the Agile Viewer is launched, the applet is loaded from the cache, unless a new version is found on the File Manager.

The Agile View servlet (VueServlet) is packaged and deployed as part of the File Manager installation. The JvueProxy Servlet is also deployed on the File Manager and tunnels the requests to the Agile Application Server using HTTP protocol. Because of this communication, each File Manager should have an associated View Server on the same machine or separate machines. No additional configuration is required except that specified for the Agile Web client and the caching View Server. The VueServlet is used to access Agile View Server across firewalls from external clients, such as the Agile Web client, on standard HTTP/HTTPS ports.

Caching View Server

The Agile View Server converts files from their native format to a metafile format (CMF). The metafile format is used to render the 2D/3D images from the cache. These files are approximately 20% the size of the source files which helps to improve View Server performance.

There are two types of metafile caching; on-demand and offline. On-demand caching occurs in the background when a user views a file in the Agile Viewer. Offline caching occurs without any user intervention through a caching utility located on the File Manager.

During offline caching, the files are translated on a dedicated caching view server. The Agile View Server should not be used for offline caching as performance can be impacted. So, the caching view server should not be installed on the same machine as the Agile View Server. Also, offline metafile caching is only supported on the primary File Manager. Distributed file managers installed in remote locations cannot use offline caching.

View Server Security

The primary security consideration pertaining to the View Server surrounds the native product data files that it caches when processing a view request. Deploying the View Server behind the firewall protects the native files and ensures that the only traffic going into the DMZ is streamed data for secure transmission. The ability of the View Server to stay behind the firewall along with the other main components is a significant security benefit.

CHAPTER 7

Communication and Data Flow

Communication and Data Flow

The following table summarizes the protocol and ports used for communication between each component in the Agile system. Keep in mind that ports are usually configurable. The ports listed are either industry standard ports or default ports that are configured out of the box.

Table 7-1: Protocols and default ports

Client	Server	Protocol	Default Port
Web Client	Web Server	HTTP/S	80/443
Web Client	File Manager	HTTP/S	80/443
Web Client	Application Server	HTTP	7001 or 8888
Java Client	Application Server	T3 or ORMI	7001 or 23791
Java Client	Distributed File Manager	HTTP	80/443
Portal Server	Application Server	T3 or ORMI	7001 or 23791
Portal Client	Portal Server	HTTP	9000
Portal Client	Web Server	HTTP/S	80/443
View Applet	Web Server	HTTP/S	80/443
View Server	File Manager	TCP	5099
Distributed File Manager	Application Server	HTTP	80
Distributed File Manager	Distributed File Manager	HTTP	80
Distributed File Manager	File Vault Mapped Drive	N/A	
File Manager	File Vault Mapped Drive	N/A	
Application Server	Database Server	JDBC	1521
Application Server	LDAP Server	LDAP	389
SDK Client Application	Application Server	T3 or ORMI or HTTP	7001 or 7777 or 23791
Product Interchange Web Server	PLM Application Server	HTTP/S	80/443
Product Interchange Web Client	Product Interchange Web Server	HTTP/S	8080/443
Product Interchange Web Server	Database Server	JDBC	1521

Client Communication

The Agile application adapts to your current network infrastructure and utilizes the available bandwidth to transmit data via TCP protocol. Typical traffic varies based on the number of concurrent users and the activity performed. Agile's architecture is very network friendly. The metadata or database traffic is fairly light and the amount of physical data rows transmitted across the network for the web client is limited and controlled to a set amount. For instance, if your query requested 1000 rows worth of data, the search result may only display the first 50 rows. Of these rows, only a portion of the attribute data is transferred for your review. You can then drill into the appropriate record to display the detailed information. At that time, another much smaller query is issued based on an index to efficiently retrieve the data. Attachments are not transferred until viewed or opened. Once requested, the file exists as part of the local client cache to be opened or viewed locally.

For security reasons, it is recommended to use SSL for all communication with clients outside the firewall. HTTPS protocol is supported and typically runs on port 443. HTTPS is a firewall-friendly protocol using SSL so the data is transferred securely.

When communicating to clients inside the firewall, there are other viable protocol options. SSL increases communication overhead, therefore, using HTTP protocol instead of HTTPS for internal clients results in faster response time. In addition to using HTTP, the Java client can connect directly to the application servers using application server specific protocols. Oracle Application Server uses ORMI, while BEA WebLogic uses T3. Using native application server-specific protocols is recommended because it is more efficient than using HTTP.

Proxy Web Server

The proxy web server forwards traffic from the clients to the application server. HTTP protocol is used between the web server and the application server. Encryption is supported only between IIS and the client's browser, not between the proxy web server and the application server, so do not put a certificate on the application server.

The proxy web server is sensitive to source or destination IP Network Address Translation (NAT). You need to bind the Application Server instance to the DNS name and use the same DNS name both inside and outside of the firewall. The proxy server is not sensitive to source Port Address Translation (PAT).

Typical Network Configurations

Figure 7-1 depicts a typical network configuration. The right side of the figure shows the internal LAN/WAN environment protected by a firewall. Then you have a DMZ area where web servers and a proxy cache or server-side proxy server exist. An additional firewall protects the DMZ from the Internet and other company environments. The proxy cache server is typically placed in front of the web servers and provides caching activities to improve performance for the clients behind the LAN/WAN firewall. The client in this case is beyond another corporate firewall and may have a client-side proxy server functioning as part of the firewall. A client-side proxy server brokers communication between the Internet and client browsers behind the firewall, protecting the clients from direct communication and possible threats. Client-side proxy software may require authentication, often username and password verification, prior to completing any requested transactions. Additionally, client-side proxy software may be configured as a transparent or a non-transparent proxy.

The initial installation of the Java client requires the transfer of a JAR file. Proxy configurations need to be considered when deploying the Java client to supply chain partners.

Application Server

There are two components on the application server that clients can connect to, the web container and the EJB container. Clients connect to the web container using HTTP protocol. Clients connect to the EJB container, using application server specific protocols as shown in Table 7-1. All components in the Agile system connect in one of these two ways including web clients, Java client, portal server, SDK clients, distributed file management, proxy web server, and view server. The application server connects to the file vault, database server, and optional LDAP directory server using the protocols also shown in Table 7-1.

Directory Server

Directory server communication occurs using LDAP protocol. If a directory server is being used to manage accounts in Agile, then it is also used for authentication. Every time login occurs, the directory server is asked to authenticate the user. The application server also connects to the directory server for synchronization. Synchronization can be user and/or schedule driven.


Distributed File Management

As Figure 7-2 depicts, a distributed file management server fits into the system configuration much like the application server. In fact, it is a local application server with only file operation capabilities. All users connect to the main application server through the proxy web servers, but when a remote user requests a file, they are pointed to their configured distributed file manager. A distributed proxy web server can be deployed in proxy calls between users and the file manager.

An Agile distributed file manager may connect to the Agile Application Server or another file manager if a request file does not exist in the local vault. The file manager connects to the Agile Application Server for other file operations such as deleting and redlines performed through the View Server. File manager to file manager communication occurs over application specific protocols and configured ports.

With remote centers, all attachment viewing, opening, and updating of documents is accomplished locally within the remote center. The only time a file attachment is transferred across the WAN is if the file version resides on another remote file server. At that time, the file is copied down to the local file server and then remains local when accessed again.

Figure 7-1: Sample Distributed File Management configuration



View Server

Viewing a file through the View Server is handled differently than getting the file. If you view a file instead of getting it, the request is sent to the view server to render the file. The view server requests the file from the file server. The view server and file server communicate using the HTTP protocol. The file server gets the file from the vault and sends it back to the view server. The view server then streams an image of the file across the network to the viewer applet that runs in the user's browser over HTTP/S. Using the viewer is very network friendly because it uses firewall friendly protocols and the streamed image is smaller than the native file. For example, a 2MB native file may typically be compressed to approximately 100KB.

Network Bandwidth Requirements

The data transfer of the Web client for Agile PLM 9.2.1 averages about 140KB per interaction. With a low of approximately 50KB when selecting a subclass on the Item Creation window to 900KB on the Quick Search Results window with 250 rows per page displayed. The amount of data transferred is proportional to the data displayed. For example, on the Item BOM page with the default of 10 rows per page setting and the default fields, the data transfer increases by 3.8KB. Similarly, on the Change Object-Affected Items Tab, each row above the default 10 rows adds an additional 6.8KB to the data transfer. Agile recommends using the default of 10 rows per page, unless your business process requires you to change it. Also, the first time the Web client is accessed, all of the static content like images and style sheets are downloaded. This causes a higher data transfer rate per page. After the initial file load until the web browser cache is emptied or there is an update to the application on the server, the average data transfer to and from the Web client for a page request is in the range of 100 to 120KB.

Over a WAN, the network latency and bandwidth are the major factors affecting the Web client performance. On a 2Mbps T1 line, the effective bandwidth is approximately ~225Kbps compared to ~10Mbps on a 100Mbps Ethernet LAN. So, the higher bandwidth would be required on a WAN to support the number of concurrent users expected on the WAN. Conversely, the response time increases if the bandwidth is not sufficient to support all of the concurrent requests over the WAN.

For the Java client, the data transfer depends on what transactions are performed and the amount of data that results from the transaction. For a business object, loading the data is less than that for a corresponding transaction on the Web client, except during a search. For a search that results in 5000 records, data for all 5000 records is transferred which produces 1.5MB of data on the Java client; whereas, the same search on the Web client would result in a data transfer of 150 to 600KB, depending on a rows per page setting of 10 to 250. The reason for this difference is the data for all 5000 records is transferred to the Java client at once while in the Web client, pagination controls the number of records the server sends at a time.

A 100Mbps network will slow down with about 100 concurrent hits, not counting any other network traffic. Therefore, it is highly recommended that all Agile servers are connected to a Gigabit LAN.

CHAPTER 8

Product Interchange

About Product Interchange

The Agile Product Interchange solution enables users to identify any structural or parts related issues in BOMs/ AMLs and resolve these issues in a systematic fashion.

The **Product Translation** function performs automated conversion of a wide range of text-based BOM/AML formats into user-defined format. The supported input formats include:

- MS-Excel, MS-Word, Text, HTML, PDX formats
- Single/Multiple Levels, Single/Multiple Files
- Delimited/fixed width formats, Complex rows/columns, Multi-line fields, Wrap-around fields

It also performs comprehensive ‘rules-driven’ validation of the BOM/AML structure and provides an interactive framework for correcting errors.

The **Product Cleansing** function performs manufacturer name aliasing by converting manufacturer name variants into standard ‘corporate’ manufacturer names. It also performs part number aliasing by mapping ‘dirty’ incoming part numbers (manufacturer or internal) to parts numbers previously ‘cleaned’ in the organization. Over a period of time this knowledge base of manufacturer name aliases and part number aliases leads to more accurate parts information through out the PLM processes.

The **Product Validation** function enables users to confirm the accuracy of part information by matching against reference databases or parts catalogs. This also enables users to find alternate replacement parts where issues are detected.

The **Product Enrichment** function provides a framework for adding a breadth of attributes to the individual part / item records to make the product information more complete and usable throughout the product lifecycle. The additional product attributes can be sourced from one or more reference sources.

Product Interchange Requirements

The Agile Product Interchange application server is deployed on Apache Tomcat 5.0.30. The Product Interchange database can be collocated with the Agile PLM database instance or can be installed on a separate database. It is recommended to deploy the Product Interchange Web Server and Agile PLM Application Server on separate physical servers, specifically if the PLM Application Server has two CPUs. Agile File Manager and the Agile Product Interchange Server should not be installed on the same physical server.

An Agile Integration Services (AIS) license is required to use Agile Product Interchange. To use the Product Interchange Dataload utility, Microsoft Excel should be installed on the Product Interchange Web Server.

Typically, the Product Interchange Web Server should reside in the Safe Zone behind the corporate firewall and collocated along with the Agile PLM Application and Database Servers.

The following diagram shows a typical deployment. Note that in this diagram, only clients from within the Firewall or Intranet would be able to access the Product Interchange server.

In order to provide external access, the Product Interchange Application Server would be required to be deployed in the DMZ and port 1521 (or the port where the Product Interchange database instance is configured to listen) need to be opened in the Firewall between the DMZ and the Intranet or Safe Zone.

Another possible scenario would be to install the Product Interchange Web Server and Product Interchange Database on the same physical server to completely separate the Agile PLM and Product Interchange installations. In this case, there would be additional Oracle License requirements. For this deployment, the Product Interchange server should be located in the Intranet/Safe Zone. For external access, a separate Proxy server would need to be set up in the DMZ.

CHAPTER 9

Performance Tips

Tuning Memory for Java Applets

The Agile Web Client uses Java applets for advanced functionality. Examples include the Gantt Chart and the Agile Viewer. These applets use the Java Plug-in to run inside your browser.

The amount of memory an applet requires depends on the content it attempts to load. If you experience memory problems while running the Gantt Chart, the Agile Viewer, or other Java applets, you should increase the amount of memory available to Java applets. For information on configuring Java applet runtime parameters using the Java Control Panel, see *Installing Agile PLM with Oracle Application Server* or *Installing Agile PLM with BEA WebLogic*.

Optimum JVM Parameters on Solaris

After running several tests on Solaris machines with various amounts of user distribution loads for the components, the following parameters have been identified to improve out-of-box system performance:

```
-server -Xms3g -Xmx3g -XX:NewSize=1300m -XXMaxNewSize=1300m -XX:SurvivorRatio=3 -  
XX:TargetSurvivorRatio=90 -XX:MaxPermSize=256m -XX:PermSize=100m -XX:+DisableExplicitGC -  
XX:+UseMPSS
```



APPENDIX A

System Configuration Examples


This appendix shows diagrams of typical system configurations. It contains the following sections:

- ❑ *Small System Configuration*
- ❑ *Medium System Configuration*
- ❑ *Large System Configuration*
- ❑ *Extra Large System Configuration*


Small System Configuration



Medium System Configuration



Large System Configuration



Extra Large System Configuration

