

Agile Product Lifecycle Management

Capacity Planning Guide

Release 9.3.6

E71149-17

April 2023

Copyright © 2010, 2023, Oracle and/or its affiliates. All rights reserved.

Primary Author: Oracle Corporation

Contributing Author: Edlyn Sammanasu

Contributor:

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	vii
Audience.....	vii
Documentation Accessibility	vii
Related Documents	vii
Conventions.....	vii
 1 Requirements	
Application Server.....	1-1
Load Balancer.....	1-1
Clients.....	1-1
Database Server.....	1-2
Agile File Manager	1-2
LDAP Directory Server	1-2
AutoVue for Agile PLM.....	1-2
Agile PLM System Overview	1-2
Additional Agile PLM Software Components	1-3
Software Requirements.....	1-4
Hardware Requirements	1-8
 2 Capacity Planning	
Agile PLM Application Server.....	2-1
Java Heap Size Recommendations	2-1
Hardware Sizing.....	2-2
Hardware Sizing for Exalogic/Exadata.....	2-3
Hardware Sizing for Linux	2-3
Hardware Sizing for Windows.....	2-4
Hardware Sizing for Solaris.....	2-4
Hardware Sizing for AIX	2-5
Database Server Sizing	2-6
Oracle Database Sizing.....	2-6
Small Database	2-7
Medium Database	2-7
Large Database	2-7
Extra-large Database.....	2-8
CPU and Memory	2-8

Storage	2-8
Network Card.....	2-8
Sizing Summary	2-8
Hardware Resource Plan for Database Models	2-8
Disk Space and Tablespace Configurations.....	2-8
One-Disk	2-9
Two-Disk Configuration.....	2-9
Four-Disk Configuration.....	2-10
Eight-Disk Configuration.....	2-11
Twelve-Disk Configuration.....	2-12
Hardware Load Balancer and Proxy Web Server	2-13
CPU	2-13
Network Card	2-13
Distributed File Management	2-13
AutoVue for Agile PLM.....	2-14
Events	2-14

3 Breakpoint Testing Results

Reports Scalability Data	3-1
<i>Heap Consumption in Reports</i>	<i>3-1</i>
<i>Protecting the Application Server.....</i>	<i>3-1</i>
<i>Example: Configuration Scenario</i>	<i>3-1</i>
Solution Results on Windows and Linux	3-2
Product Collaboration	3-2
Change Object.....	3-2
Item Object	3-3
Product Quality Management.....	3-4
PSR Object	3-4
QCR Object.....	3-4
Product Portfolio Management	3-4
Product Governance & Compliance	3-5
Product Governance & Compliance Nightly Rollup.....	3-7
Product Cost Management.....	3-7
Distributed File Manager	3-9
Export.....	3-9
Solution Results on Solaris	3-10
Product Collaboration	3-10
Change Object.....	3-10
Item Object	3-11
Product Quality Management.....	3-12
PSR Object	3-12
QCR Object.....	3-12
Product Portfolio Management	3-12
Product Cost Management.....	3-14
Distributed File Manager	3-16
Export.....	3-16
Solution Results on AIX.....	3-17

Product Collaboration	3-17
Change Object.....	3-17
Item Object.....	3-18
Product Quality Management.....	3-18
PSR Object	3-19
QCR Object.....	3-19
Distributed File Manager	3-19
Export.....	3-20
 4 Client Requirements	
Client Requirements	4-1
 5 Directory Server	
Overview	5-1
System Architecture.....	5-1
Implementation Details.....	5-2
Configuration	5-3
Directory Service Connection Parameters	5-3
Multiple Directory Server Support.....	5-3
Schedule Synchronization.....	5-3
On-Demand Synchronization	5-3
Utilities.....	5-4
Frequently Asked Questions (FAQ)	5-4
What happens when I delete a user from Agile?.....	5-4
What happens when I delete a user from the directory server?.....	5-4
Does the Agile server allow login ID (user ID) changes?.....	5-4
My directory server provides a feature called "Activate/Inactivate". How does this relate to "Enable/Disable" within Agile?	5-4
I want to create a user in the directory server and log in to Agile immediately. How can I do this?	5-4
Can I still create a user from within Agile? How does it reflect in the directory server?	5-5
 6 File Management Server and Viewer	
Agile File Management Server.....	6-1
File Vaults	6-1
Distributed File Manager	6-2
Using AutoVue for Agile PLM in a Distributed Environment.....	6-2
How Distributed File Manager Works.....	6-2
Do I need distributed file management?	6-3
Can I synchronize my distributed File Managers?.....	6-3
Do I need to back up all of my distributed File Managers?.....	6-3
Do I need an AutoVue for Agile PLM for every distributed File Manager?	6-4
Should I install Anti-virus software on my file vault?.....	6-4
File Management Security.....	6-4
Server Security	6-4
Client Access Security	6-4

AutoVue for Agile PLM.....	6-4
Caching Viewer	6-5
AutoVue for Agile PLM Security	6-5

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
AutoVue for Agile PLM	7-4
Network Bandwidth Requirements	7-4

8 Performance Tips

Tuning Memory for Java Applets.....	8-1
Running PPM Gantt Charts.....	8-1
Optimum JVM Parameters on Solaris.....	8-1
Small System Configuration.....	8-2
Medium System Configuration	8-3
Large System Configuration	8-4
Extra Large System Configuration	8-5

Preface

Agile PLM is a comprehensive enterprise PLM solution for managing your product value chain.

Audience

This document is intended for administrators and users of the Agile PLM products.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

Oracle's Agile PLM documentation set includes Adobe® Acrobat PDF files. The Oracle Technology Network (OTN) Web site <http://www.oracle.com/technetwork/documentation/agile-085940.html> contains the latest versions of the 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 Agile PLM Documentation folder available on your network from which you can access the Agile PLM documentation (PDF) files.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.

Convention	Meaning
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Requirements

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 AutoVue for Agile PLM.

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. As the System Configuration Overview figure 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.

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 can be deployed in the Demilitarized Zone (DMZ) in order to proxy requests from outside the corporate firewall to the application server in the Safe Zone.

A load balancer is required if you are installing application servers in a cluster and/or if you have multiple primary file manager instances. The load balancer is required in order to enforce session persistence for both the clustered application server instances and the redundant file manager instances. For the clustered application server, the load balancer must enforce persistence using a cookie-insert mechanism where the load balancer injects its own cookie into the request. The application passes this cookie to the file manager during file operations so that any request from the file manager to the application server on behalf of the end-user is also persisted properly.

Clients

Agile PLM 9.3.6 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 T3 for Oracle 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. Check ["Software Requirements"](#) on page 1-4 for details about which database software is supported.

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

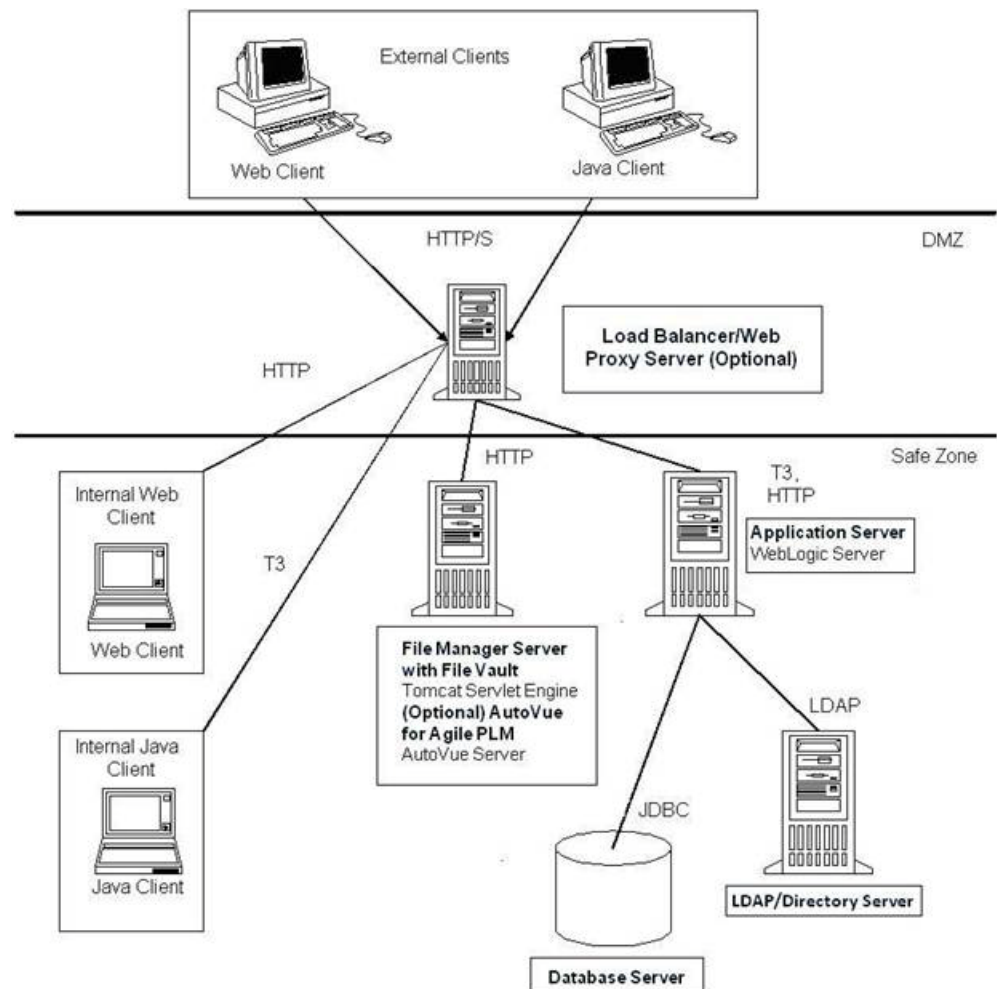
In an effort to better support the industry standard authentication schemes, Agile PLM 9.3.6 supports Lightweight Directory Access Protocol (LDAP) based authentication. LDAP support enables you to integrate Agile with existing directory servers so user accounts can be managed in one place. Integrating 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.

AutoVue for Agile PLM

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

Agile PLM System Overview

The following diagram provides an overall depiction of the Agile PLM System Overview.

Figure 1–1 Agile PLM System Overview

Additional Agile PLM Software Components

Agile PLM 9.3.6 requires new versions of OPLA, AutoVue, EC MCAD, EC ECAD, PLM Mobile, D2R AIA PIP for EBS, D2R AIA PIP for SAP, and a patch to enable RMW 9.3.4. These additional software components will release after the initial release of Agile PLM 9.3.6. The list below defines the minimum versions of these software components:

- ⌚ OPLA 3.6
- ⌚ AutoVue 21.0.1 or AutoVue 21.0.2 for Agile PLM
- ⌚ EC MCAD 3.6
- ⌚ EC ECAD 3.6
- ⌚ PLM Mobile 4
- ⌚ D2R AIA PIP for EBS 3.6

- ? D2R AIA PIP for SAP 3.6
- ? RMW 9.3.4 patch

Software Requirements

The Agile PLM 9.3.6 Software Requirements table 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.

The table below lists the supported application servers.

Agile Components	Operating Systems	Certified Software
Database Server (Standalone or RAC)	<p>Oracle Solaris 11.2+ (SPARC64, x86-64)</p> <p>Oracle Linux 8 (requires Database 19c or later, FMW 12.2.1.4 and Agile PLM 9.3.6 RUP 13 or later), 7, 6 (x86-64)</p> <p>Red Hat Enterprise Linux 8 (requires Database 19c or later, FMW 12.2.1.4 and Agile PLM 9.3.6 RUP 13 or later), 7, 6 (x86-64)</p> <p>Red Hat Enterprise Linux 7, 6 (x86-64)</p> <p>SUSE Linux 12 (x86-64)</p> <p>Microsoft Windows Server 2022 and 2019 (with Agile PLM 9.3.6 Installer patch 31040722), 2016 (with Agile PLM 9.3.6 Installer patch 27871474), 2012R2, 2012</p> <p>IBM AIX 7.2, 7.1 (POWER 64-bit)</p> <p>HP-UX 11.31 (Itanium64)</p> <p><i>Note: Oracle Exadata Certified</i></p> <p>Cloud Infrastructure:</p> <p>Oracle Cloud Infrastructure</p> <p><i>Also supported, but not tested with other infrastructure offerings.</i></p> <p>Virtualization:</p> <p>Oracle VM 3 with Oracle Linux, Oracle Solaris and Microsoft Windows guest Operating Systems listed above.</p> <p><i>Refer to Oracle Support Note "Certified Software on Oracle VM"</i></p> <p>Oracle Solaris Containers/Zones</p> <p>IBM AIX LPARs</p> <p>Microsoft Hyper-V (See Support Note 1563794.1)</p>	<p>(RUP 10 and later) Oracle Database Server 19c (version 19.3.0.0.0+)</p> <p>(RUP 8 and later) Oracle Database Server 18c (version 18.0.0.0+)</p> <p>(RUP 1 and later) Oracle Database Server 12cR2 (version 12.2.0.1+)</p> <p>Oracle Database Server 12cR1 (version 12.1.0.1+)</p> <p><i>Note: Enterprise Edition, Standard Edition, Standard Edition One, Standard Edition Two certified requires Database.</i></p>

Agile Components	Operating Systems	Certified Software
Application Server	Same as Database Server (see above) <i>Note: Oracle Exalogic certified.</i>	(RUP 13 and later) Oracle FMW 12c Infrastructure version 12.2.1.4 - requires 64-bit JDK 1.8.0_211 or later (fresh install of Agile PLM 9.3.6 on FMW 12.2.1.4 supported with installer patch 31040722) (RUP 6 and later) Oracle FMW 12c Infrastructure version 12.2.1.3 - requires 64-bit JDK 1.8.0_181 or later (fresh install of Agile PLM 9.3.6 on FMW 12.2.1.3 supported with installer patch 27871474). (RUP 5 and earlier) Oracle FMW 12c Infrastructure version 12.2.1.1 - requires 64-bit JDK 1.8.0_172 or later. <i>Note: WebLogic Suite, WebLogic Server Enterprise Edition, WebLogic Server Standard Edition certified.</i>
File Manager (DFM)	Same as Application Server (see above)	Note: Apache Tomcat used for File Manager is now updated to a recent version with every RUP. See the RUP ReadMe for a table listing the Apache Tomcat version delivered with each RUP. Note: Uses JDK provided during installation of Agile PLM.
AutoVue for Agile PLM (AutoVue Server) <i>(Note: Refer to Oracle AutoVue documentation for latest system requirements.)</i>	See AutoVue documentation for detailed information on operating systems support.	(RUP 8 and later) Oracle AutoVue 21.0.2.8 for Agile PLM (RUP 8 and later) Oracle AutoVue 21.0.2.7 for Agile PLM (RUP 8 and later) Oracle AutoVue 21.0.2.6 for Agile PLM (RUP 11 and later) Oracle AutoVue 21.0.2.4 for Agile PLM (RUP 11 and later) Oracle AutoVue 21.0.2.2 for Agile PLM (RUP 8 and later) Oracle AutoVue 21.0.2.0 for Agile PLM (RUP 7 and earlier) Oracle AutoVue 21.0.1.5 for Agile PLM (patch 27998404) (RUP 7 and earlier) Oracle AutoVue 21.0.1.0 for Agile PLM.
Web Server (Proxy)	Oracle Solaris 11 (SPARC64, x86-64) Oracle Linux 8, 7, 6 Red Hat Enterprise Linux 8, 7, 6 SUSE Linux 12	Apache HTTP Server 2.x Oracle HTTP Server (OHS) 12c
	Microsoft Windows Server 2022, 2019, 2016, 2012 R2, 2012	IIS (bundled with Windows) Oracle HTTP Server (OHS) 12c

Agile Components	Operating Systems	Certified Software
Web Client	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit)	Google Chrome 46 and later. Microsoft Edge Microsoft Internet Explorer 11 (no longer supported with RUP 19 and later due to uptake of 3rd-party library that no longer supports IE11) Mozilla Firefox ESR 38 and later Java SE 8
Web Client	Oracle Linux 8, 7, 6	Mozilla Firefox ESR 38 and later Java SE 8
Web Client	Apple macOS 10.12 (Sierra) and later	(RUP 9 and later) Apple Safari 13, 12, 11 Apple Safari 10, 9 and 8 Java SE 8
Java Client	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit)	Google Chrome 46 and later Microsoft Edge Microsoft Internet Explorer 11 (no longer supported with RUP 19 and later due to uptake of 3rd-party library that no longer supports IE11) Mozilla Firefox ESR 38 and later Java SE 8
Java Client	Oracle Linux 8, 7, 6	Mozilla Firefox ESR 38 and later Java SE 8
SDK Client	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit) Oracle Linux 8, 7, 6	Java SE 8 JDK
SDK Client	Apple macOS 10.12 (Sierra) and later	Java SE 8 JDK
EC Services	Same as Application Server (see above)	
Excel Integration for Import/Export	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit)	(RUP 10 and later) Microsoft Excel 365, 2021, 2019, 2016, 2013, 2010, 2007 and 2003
Excel Integration for PG&C	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit)	(RUP 10 and later) Microsoft Excel 365, 2021, 2019, 2016, 2013, 2010, 2007 and 2003
Project Integration for PPM	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit)	Microsoft Project 2021, 2019, 2016, 2013 and 2010
Directory Servers (LDAP)	Same as Application Server (see above)	Oracle Internet Directory (OID) 11g (version 11.1.1.7+) Oracle Virtual Directory (OVD) 11g (version 11.1.1.7+) Oracle Unified Directory (OUD) (version 12.2.1.4+)

Agile Components	Operating Systems	Certified Software
Directory Servers (LDAP)	Oracle Solaris	Oracle Directory Server Enterprise Edition (ODSEE)
Directory Servers (LDAP)	Microsoft Windows Server	Microsoft Active Directory (bundled with Windows) and AD-LDS
SSO - SAML2	Oracle Solaris 11 (SPARC64, x86-64) Oracle Linux 8, 7, 6 Red Hat Enterprise Linux 8, 7, 6 SUSE Linux 12 Microsoft Windows Server 2022, 2019, 2016, 2012 R2, 2012	(RUP 9 and later) Oracle IDCS as IdP (should also work with other SAML2 IdPs)
SSO - Oracle Access Manager	See OAM documentation for supported operating systems.	Oracle Access Manager (OAM) 11gR2 (version 11.1.2.2+), 11g (version 11.1.1.7+)
SSO-NTLM (Windows)	Microsoft Windows 11, 10, 8.1, 7 (64-bit, 32-bit) <i>Note: Not supported with Google Chrome browser.</i>	Microsoft Windows Server 2022, 2019, 2016, 2012 R2, 2012 (bundled IIS Web Server)
Enterprise Manager	See OEM documentation for supported operating systems.	Oracle Enterprise Manager (OEM) 12c (version 12.1.0.1+)
User Productivity Kit (UPK)		Oracle UPK 11.1
ACS	(same as Application Server)	Oracle WebLogic JMS (WLS 12.2.x)
Reporting	(same as Application Server)	Oracle BI Publisher 11g (version 11.1.1.6+).
Scripting	(same as Application Server)	Groovy 2.4.5

Hardware Requirements

The table below shows the minimum hardware needed to deploy an Agile PLM 9.3.6 system.

Agile Server	Hardware Requirements
Application Server	Dual CPU Intel Xeon 3 GHz and above OR Dual CPU Ultra Sparc IV+ 1.50 GHz and above (or equivalent) 2 MB L2 cache 4 GB RAM
Database Server	Dual CPU Intel Xeon 1.8 GHz OR Dual CPU Ultra Sparc-IIIi 1.1GHz (or equivalent) 512 KB L2 Cache 2 GB RAM (minimum), 4GB (recommended)
File Manager and AutoVue for Agile PLM	Dual CPU Intel Xeon OR Dual CPU Ultra Sparc IV+ (or equivalent) 512 KB L2 cache 2 GB RAM

Agile Server	Hardware Requirements
Web Server	Intel Xeon 2.8 GHz OR Ultra Sparc III (or equivalent) 512 KB L2 Cache 1 GB RAM

Note: CPU utilization for Agile PLM tends to be relatively high for the application server component and can be relatively light for other components (especially for file manager and database, AutoVue CPU utilization depends heavily on usage including number and size of files being rendered). Therefore, scalability is often focused on the application server component. Also for this reason, the application server is most sensitive to CPU clock speed (the faster the core clockspeed, the better the throughput and scalability). Thus, in general, if you deploy on a slower core you can expect the application to scale somewhat worse and conversely, if you deploy on a faster core you can expect the application to scale somewhat better.

The following table shows the typical hardware recommendations for production deployments:

Database Server	at least 8 cores and 8 GB physical memory per instance
Application Server	at least 4 cores and 4 GB physical memory for each application server (JVM) instance running Agile PLM
File Manager	at least 1 core and 1 GB physical memory for each Agile PLM File Manager instance
AutoVue Server	at least 1 core and 1 GB physical memory for each AutoVue Server instance (increase as usage dictates)
Web Server	at least 1 core and 1 GB physical memory for each instance (increase as usage dictates)

Note: When running more than one application server (JVM) against a single Agile PLM schema, the application server instances must be clustered. An application server cluster typically has at least two application server instances (JVMs) running the Agile PLM application. Clustering introduces some measurable overhead (estimated at approximately 10%) for things like cache synchronization over JMS, and so on and thus the application does not scale completely linearly as application server instances are added. This is especially noticeable when going from one application server instance (standalone) to two or more application server instances (clustered).

Capacity Planning

This chapter helps you to plan and gauge server capacity.

Agile PLM Application Server

To determine the application server capacity, the average Transactions per second (TPS) the server can support for a given Agile solution must be determined. For each solution, 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 solution.

In the first phase, tests were conducted on individual solutions to determine the TPS. A single, 2 CPU Dual core application server supported an average three second response time.

The TPS for the Agile PLM solutions is as follows:

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

Java Heap Size Recommendations

The table below lists the default heap sizes configured by the Agile PLM installer.

Note: The heap size recommendations below apply to JVMs running the core Agile PLM application. There is no need to alter the heap settings for the Agile PLM File Manager or the WebLogic Admin Server components.

Recommended Settings (Installer Defaults)

Configuration	Heap Size (min=max)	MaxPermSize	NewSize	MaxNewSize
UNIX/Linux (64-bit JDK 7)	3072m	512m	1300m	1300m
Windows (64-bit JDK 7)	3072m	512m	1300m	1300m

Note: For the Agile PLM application server component that runs in WebLogic Server, the JVM heap settings can be found in the shell script, `setUserOverrides`, that is located in the directory `AGILE_HOME/agileDomain/bin`. On Windows, the heap settings are also present in the service install scripts and you must uninstall and re-install the service to put a change into effect. Always back up the original file before making changes.

Note: The benchmark tests presented below were performed on 64-bit operating systems using 64-bit JVMs with a 3 GB (min=max) heap size. Note that when configuring heap size, larger is not necessarily better. The best heap configuration is arguably the smallest heap size that can accomplish the task without heap thrashing. Therefore, for most deployments, a 3 GB (3072m) initial heap size is recommended (assumes 64-bit JDK).

Hardware Sizing

Agile conducts extensive load tests to determine scalability for individual product components, and for combinations of modules. Agile uses HP Load Runner 9.1 to simulate virtual user load for the benchmark tests.

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

- ? Average user load
- ? Peak user load
- ? User distribution across different modules, if more than one module is implemented
- ? Network configuration
 - Latency
 - Bandwidth

The goal of hardware sizing is to balance hardware costs with user response times. To effectively accomplish this, you must 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 you can configure the system and tune it 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 must 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.

Hardware Sizing for Exalogic/Exadata

Agile PLM on Oracle Exalogic (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for Exa based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Total Number of Users
Intel Xeon E5-2690 @ 2.90GHz	1	4	400
Intel Xeon E5-2690 @ 2.90GHz	2	8	720
Intel Xeon E5-2690 @ 2.90GHz	Add 1 Server with 4 cores for every 360 users.		720+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
Intel Xeon E5-2690 @ 2.90GHz	Oracle Exalogic X3-2	Intel Xeon Processor, 2 CPU - 16 Cores, 2.9 GHz, 32 GB RAM

Hardware Sizing for Linux

Agile PLM on Oracle Linux (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for Linux based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Total Number of Users
Intel Xeon E5-2690 @ 2.90GHz	1	4	280
Intel Xeon E5-2690 @ 2.90GHz	2	8	500
Intel Xeon E5-2690 @ 2.90GHz	Add 1 Server with 4 cores for every 250 users.		500+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
Intel Xeon E5-2690 @ 2.90GHz	Sun Server X3-2L	Intel Xeon Processor, 2 CPU - 16 Cores, 2.9 GHz, 32 GB RAM

Hardware Sizing for Windows

Agile PLM on Microsoft Windows Server (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for Windows based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Total Number of Users
Intel Xeon E5-2690 @ 2.90GHz	1	4	280
Intel Xeon E5-2690 @ 2.90GHz	2	8	500
Intel Xeon E5-2690 @ 2.90GHz	Add 1 Server with 4 cores for every 250 users.		500+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
Intel Xeon E5-2690 @ 2.90GHz	Sun Server X3-2L	Intel Xeon Processor, 2 CPU - 16 Cores, 2.9 GHz, 32 GB RAM

Hardware Sizing for Solaris

Agile PLM on Sun Solaris x86-64 (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for Solaris based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Total Number of Users
Intel Xeon E5-2690 @ 2.90GHz	1	4	280
Intel Xeon E5-2690 @ 2.90GHz	2	8	500
Intel Xeon E5-2690 @ 2.90GHz	Add 1 Server with 4 cores for every 250 users.		500+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
Intel Xeon E5-2690 @ 2.90GHz	Sun Server X3-2L	Intel Xeon Processor, 2 CPU - 16 Cores, 2.9 GHz, 32 GB RAM

Agile PLM on Sun Solaris SPARC64 (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for Solaris based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Number of Users*
SPARC T7 processor @ 4.13GHz	1	4	280
SPARC T7 processor @ 4.13GHz	2	8	510
SPARC T7 processor @ 4.13GHz	Add 1 Server with 4 cores for every 255 users.		510+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Note: 4 cores = 32 virtual processors (4*8 threads per core)

*Note: These results were achieved using the following JVM parameters:

```
-ms3072M -mx3072M -XX:PermSize=512M -XX:MaxPermSize=512M -XX:NewSize=1300M
-XX:MaxNewSize=1300M -XX:+UseCompressedOops -XX:+AlwaysPreTouch -XX:+UseTLAB
-XX:+AggressiveOpts -XX:SurvivorRatio=6 -XX:TargetSurvivorRatio=90 -XX:+UseParallelGC
-XX:ParallelGCThreads=16 -XX:+UseLargePages -XX:LargePageSizeInBytes=256M
```

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
SPARC T7 processor @ 4.13GHz	Oracle SPARC T7-1	SPARC T7-1 Processor, 2 CPU - 32 Cores (8 threads per core totals to 256 virtual processors), 4.13 GHz, 512 GB RAM

Hardware Sizing for AIX

Agile PLM on IBM AIX (64-bit OS with 64-bit JVM)

The following table shows the recommended hardware sizing for AIX based on user load.

Processor Type and Speed	Number of Application Servers (JVMs)	Total Number of Cores (four cores per JVM)	Total Number of Users
IBM POWER7 @ 3.86 GHz	1	4	320

IBM POWER7 @ 3.86 GHz	2	8	570
IBM POWER7 @ 3.86 GHz	Add 1 Server with 4 cores for every 285 users.		570+

Note: To support multiple application servers, clustering must be implemented, which adds an additional 10% load on each server.

Server Configuration by Processor Type

Processor Type and Speed	Server Model	Server Details
IBM POWER7 @ 3.86 GHz	IBM Power 780	4 CPU - 16 Cores, 32 GB RAM

Database Server Sizing

For production environments, it is recommended to run the database server on dedicated hardware. Database hardware sizing depends 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, and to 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 must 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 must be estimated by monitoring database growth over the first few months of normal operation to predict future disk size needs.

The following tables show the Agile PLM 9.3.6 Database Sizing Matrix for Oracle.

Database Sizing Model:

? Small

Agile DB Configuration	Number of Users	CPU	RAM in GB	Disks
D	1000	8	8	9
C	500	4	4	4
B	250	4	2	4
A	100	2	1	4

? Medium

Agile DB Configuration	Number of Users	CPU	RAM in GB	Disks
D	1000	12	12	9
C	500	8	8	9

B	250	4	4	4
A	100	2	2	4

? Large

Agile DB Configuration	Number of Users	CPU	RAM in GB	
D	1000	16	16	13
C	500	8	8	11
B	250	4	4	9

? Extra Large

Agile DB Configuration	Number of Users	CPU	RAM in GB	Disks
D	1000	24	24	15
C	500	12	12	13

The following table shows the 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 9.3.6 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 9.3.6 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 9.3.6 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 9.3.6 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 the sizing tables, the Agile PLM 9.3.6 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 9.3.6 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. You can use the Agile PLM 9.3.6 database utility optimize up to 15 separate disks.

Hardware Resource Plan for Database Models

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

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 can result in disk I/O contention when the storage device is a single physical disk. As both database size and usage increase, performance can decline significantly. A one-disk configuration is best for a demonstration, preproduction, and testing environment or where the database files are stored on a RAID array, SAN, or other storage subsystem with built-in striping and mirroring. The configuration can be implemented as shown in the table below.

The following table shows a one-disk configuration for OFA implementation.

Disk	Resource
Disk 1	ORACLE_HOME
	SYSTEM
	TOOL
	UNDO
	TEMP
	USERS
	INDX
	AGILE_DATA1
	AGILE_DATA2
	AGILE_DATA3
	AGILE_DATA4
	AGILE_DATA5
	AGILE_INDX1
	AGILE_INDX2
	AGILE_INDX3
	AGILE_INDX4
	AGILE_INDX5
	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.

The following table shows a two-disk configuration for OFA implementation.

Disk	Resource
Disk 1	ORACLE_HOME
	SYSTEM
	TOOL
	UNDO
	AGILE_DATA1
	AGILE_DATA2
	AGILE_DATA3
	AGILE_DATA4
	AGILE_DATA5
	LOG1
	LOG2
Disk 2	TEMP
	USERS
	INDX
	AGILE_INDX1
	AGILE_INDX2
	AGILE_INDX3
	AGILE_INDX4
	AGILE_INDX5
	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 the table below 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.

The following table shows a four-disk configuration for OFA implementation.

Disk	Resource
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
	AGILE_DATA5
	Controlfile03
Disk 4	AGILE_INDX1
	AGILE_INDX2
	AGILE_INDX3
	AGILE_INDX4
	AGILE_INDX5

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 the table below.

The following table shows an eight-disk configuration for OFA implementation.

Disk	Resource
Disk 1	ORACLE_HOME
	SYSTEM
	TOOL
	UNDO
	LOG1/2/3/4
	Controlfile01

Disk	Resource
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
	AGILE_DATA5
Disk 7	AGILE_INDX1
	AGILE_INDX2
Disk 8	AGILE_INDX3
	AGILE_INDX4
	AGILE_INDX5

Twelve-Disk Configuration

Further separating the AGILE_DATA and AGILE_INDX tablespaces, twelve-disk configurations can be implemented as shown in the table below. 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.

The following table shows a twelve-disk configuration for OFA implementation.

Disk	Resource
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	Resource
Disk 8	AGILE_DATA4
	AGILE_DATA5
Disk 9	AGILE_INDX1
Disk 10	AGILE_INDX2
Disk 11	AGILE_INDX3
Disk12	AGILE_INDX4
	AGILE_INDX5

Hardware Load Balancer and Proxy Web Server

A load balancer or proxy web server is deployed to direct requests to the application server(s). When external users need access to Agile, at least one of these 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.

The following table shows the 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.

AutoVue for Agile PLM

AutoVue for Agile PLM performance is a function of the number of files being viewed concurrently and the average file size being viewed. For the latest requirements, refer to the Oracle AutoVue OTN documentation site.

Events

The Events feature allows custom business logic, which is implemented as Java code or Groovy script, to be invoked as part of PLM actions. This feature is a powerful capability that enables you to do validation, auto-populate attributes, and automate dependent tasks.

Enabling events, processing event subscriptions, and executing the handler with custom logic triggers additional processing by the server as part of the PLM action, which consequently impacts throughput and response times. The actual impact depends on the number of events enabled, the number of handlers registered for the enabled events, and the amount of computing done by the handler.

Internal tests done with a simple handler enabled for every action in the system reduced overall throughput by 2.5% and slowed response times for individual actions up to 10%. These numbers, however, are for illustration purposes only. In reality, you will most likely enable only a subset of the events, but the handlers are likely to be more compute intensive than the test scenario. You should consider the additional computational load from events and adjust hardware sizing accordingly.

Breakpoint Testing Results

This chapter includes scalability data and breakpoint testing results.

Reports Scalability Data

This section provides results for reports scalability data.

Heap Consumption in Reports

When calculating resource needs, assume that reports will consume one-third of application server resources (Heap and CPU). For example, if the application server heap size is set to 1280, then the available heap on that application server for reports should be estimated as 426MB (1280m/3 = 426m).

Note: Depending on the report, 1024 MB of available heap can typically handle up to 300K rows.

Total Heap Size (in megabytes)	Available Heap Size for Reports	Maximum rows based on available heap
1024	341	100,000
1280	427	125,000
2048	683	200,000
3072	1024	300,000

Protecting the Application Server

Agile PLM now provides two configuration parameters to limit report result set sizes and avoid out-of-memory errors. The configuration parameters are:

- report.maxStandardBOMReportResults - set in the agile.properties file, this parameter sets the maximum number of BOM reports.
- MaxReportResults - set in the Java Client, this parameter sets the maximum number of all non-BOM reports.

Example: Configuration Scenario

Scenario: The environment has a clustered deployment with 4 managed servers each having heap size of 3072m (assumes 64-bit JDK).

Considering the above scenario, the following table provides suggestions for the value of the report configuration parameters, depending on the desired number of reports to be executed concurrently.

Total Heap Size	Maximum Result Set Per Node	Desired Concurrent Reports Per Node	Suggested Value For Above Parameters
3072m	300K rows	3	100K
3072m			
3072m	300K rows	5	60K
3072m	300K rows	10	30K
3072m	300K rows	12	25K
3072m	300K rows	15	20K

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, and Product Governance & Compliance solutions, 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 Oracle Enterprise Linux 4 or RedHat Linux version 4

Note: The response times for all of the tables are measured in seconds and are for single user only.

Product Collaboration

The following are breakpoint testing results for the Product Collaboration solution.

Change Object

The following table lists the results for Copy and Paste Affected Items.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	40	10%	10%
2000	240	25%	10%
3000	570	25%	10%
5000	Timed out	25%	12%

The following table lists the results for Save As Change.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	17	10%	5%
2500	57	15-25%	5%
5000	153	25%	5%
7500	268	25%	5%

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
10000	450	25%	5%

The following table lists the results for Change Status to Release (with Preliminary Items).

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	27	5%	10%
2500	70	5%	12%
5000	140	5%	12%
7500	214	5%	12%
10000	273	5%	12%

The following table lists the results for Change Status to Pending.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	42	5%	10%
2500	142	5%	10%
5000	305	25%	10%
7500	453	25%	10%

The following table lists the results for Change Status to Released. Each item has at least 3 pending changes.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	25	<5%	12%
2500	73	<5%	12%
5000	150	5%	10%
7500	220	<5%	12%

Item Object

The following table lists the results for Copy and Paste BOM.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	15	<5%	<5%
2500	70	15-20%	5%
5000	150	15-20%	5%

The following table lists the results for Expand BOM-All Levels.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
3596	22 minutes	75%	NA

Each page load took 5 seconds.

Product Quality Management

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

- ? PSR - 23908 items
- ? QCR - 23525 items

PSR Object

The following table lists the results for Copy and Paste Affected Items (Newly Imported Items).

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	52	10%	5%
2500	403	10-15%	5%
3500	720	10-15%	5%
5000	Timed out after 23 minutes	10-15%	5%

The following table lists the results for Copy and Paste Affected Items (Existing Items).

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	302	10%	5%
2500	Timed out	10-15%	5%

QCR Object

The following table lists the results for Copy and Paste Affected Items.

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
2500	360	10-15%	5%
5000	Timed out after 22 min	10-15%	5%

Product Portfolio Management

The following table lists the results for several PPM scenarios.

Scenarios	1500 Object Tree (Response time in seconds)	3000 Object Tree (Response time in seconds)
Save As Project	30	50
Save As Template	18	35
Create Project From Template	38	95
Add 10 Team Members	21	42
Reschedule Root Project	13	30
Reschedule Child	11	24
Delegate Owner	15	27
Substitute Resource	11	24
Change Status	18	30
Expand All (Web Client)	160	NA

Tree Details:

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

Product Governance & Compliance

The following tables list the results for several PG&C scenarios.

Note: Each declaration has 2 specs. Example: Dataset 1000=Declaration with 500 parts + 2 specs.

Table 3–1 Object: Declaration Tab Name: Items

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Items	12	25	35	15%	5%
View Items Tab	5	5	5	NA	NA
Calculate Compliance	4	4	4	NA	NA
Release the Declaration	40	72	150	NA	12%
Import Items (no BOS)	NA	NA	NA	NA	NA

Table 3–2 Object: Declaration Tab Name: Mfr Parts

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Mfr Parts	12	21	50	10-15%	5%
View Mfr Parts Tab	5	5	5	NA	NA
Calculate Compliance	4	4	4	NA	NA
Release the Declaration	40	75	90	NA	12%
Import Mfr Parts (no BOS)	NA	NA	NA	NA	NA

Table 3–3 Object: Declaration Tab Name: Part Groups

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Part Groups	12	22	50	10-15%	5%
View Part Groups Tab	5	5	5	NA	NA
Calculate Compliance	4	4	4	NA	NA
Release the Declaration	42	72	150	NA	12%
Import Part Groups (no BOS)	NA	NA	NA	NA	NA

Table 3–4 Object: Substance Groups Tab Name: Substances

Scenario	50 Substances	100 Substances
Search & Add Substances	6	8
View Substances tab	3	3
Import Substances	6	11

Table 3–5 Object: Specification Tab Name: Substances

Scenario	50 Substances	100 Substances
Search & Add Substances and/or Groups	6	8
View Substances tab	3	3
Import Substance Groups	5	9

Product Governance & Compliance Nightly Rollup

Table 3–6 Stats to Rollup 109462 BOMs with 1 Spec

Total Time	Average App Server CPU	Average Database CPU
14 Hours 30 Minutes	15-20%	12%

Note: If there are more specifications, the rollups need to perform more processing. This also has an impact on performance.

Product Cost Management

The following table lists the results for several PCM scenarios.

Action	Project with 2500 Items	Project with 5000 Items	BOM with 5000 Items	Project with 10000 Items
Display project item tab	4	7	4	9
Search and Add Items into project from PC item master	41	65	22	NA
Update item from PC item master	31	48	70	117
Add/modify Partners (10)	25	45	28	73
Import item from external source	22	67	NA	160
Export item to external source	8	14	14	25
Project Save As	22	110	300	420
Quantity rollup	5	12	13	25
Display project AML tab	5	7	5	6
Display item BOM	5	7	4	9
Display of project analysis	5	6	6	8
Opening Responses Look Up	5	38	100	270
Project response price lookup	12	16	57	200
Set as best	21	45+55 (45 sec to open page + 55 to set as best)	80+95	110+138
Publish Quote History, Authoring	35+170	124+622	97+490	420 + Waited for 25 minutes

Action	Project with 2500 Items	Project with 5000 Items	BOM with 5000 Items	Project with 10000 Items
Publish Published Price Redlining	Out of memory - Major perf. Improvement made in 9368			
Items RFQ Creation	30	153	110	306
Open RFQ	15	50	45	95
Display of RFQ Responses	15	21	10	24
Response Status Tab	3	8	6	7
Opening RFQ Response Look Up	12	38	105	260
RFQ Response Look Up	9	17	120	215
RFQ Response Export	18	56	57	107
RFQ Response Import	35	120	80	270
RFQ Response status display	3	8	6	7
Edit/Modify Responses	14	53	35	180
Supplier Response Export	19	45	35	87
Supplier Response Import	37	104	73	220
Submit Responses	19	53	35	104
Costed BOM Comparison (Buyer) Report	9	20	7	14
Delete Project	4	4	2	2
Close RFQ	14	45	2	90
Close RFQ (after Lock RFQ)	15	50	39	90
Delete RFQ	35	40	35	85
Lock RFQ	15	55	45	110
Analysis tab export	32	93	90	180
Data Sets: <ul style="list-style-type: none"> Project with 2500 items, each having one AML Project with 5000 items, each having two AMLs Project with 10000 items, each having two AMLs Five-level BOM with a total of 5000 items and 2 AMLs for each leaf component. The top-level assembly contains 10 immediate children with each child having 5 immediate children. 				

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.

The following tables list the results for file upload and download.

Note: Response time is measured in seconds.

Table 3–7 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 1 GB	96	NA	<5%	5-10%
Download - 1 GB	155	NA	25%	5-10%
Upload - 2 GB	Crashes	NA	NA	NA
Download - 2 GB	357	NA	10%	5%

The following table lists the results for advanced file upload and download.

Table 3–8 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 5GB	435	435	<5%	5%
Download -5GB	153	152	15%	10%

Export

Export was tested using the Web Client.

The following table lists the results for Export scenarios.

Table 3–9 Item/BOM Export Scenario

Scenario	Data Set	Response Time	JVM Heap	CPU Usage - App Server	CPU Usage - DB Server	Comments
Export CSV	4.8 Million Rows	4 h 30 mins	1200 MB	20% continuous	<5% continuous	6 level deep BOM Output file size: 3 GB
Export PDX/aXML	4.8 Million Rows	25 mins	1200 MB	20% continuous	<5% continuous	Same BOM as above. Duplicated rows won't be exported to PDX and aXML formats.

Table 3–10 *Changes Export Scenario*

Scenario	Data Set	Response Time	JVM Heap	CPU Usage - App Server	CPU Usage - DB Server	Comments
Export CSV	3000 rows	3 h	1200 MB	20 %	<5%	NA

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 solutions, 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, 16 GB RAM, 1.5 GHz, 32 MB L2 Cache
- Database Server: v890, IV+ Processor, 8 CPU - 16 Core, 16 GB RAM, 1.5 GHz, 32 MB L2 Cache

Note: The response times for all of the tables are measured in seconds and are for single user only.

Product Collaboration

This section provides breakpoint testing results for Product Collaboration.

Change Object

The following table lists the results for Copy and Paste Affected Items on Solaris.

Table 3–11 *Results for Copy and Paste Affected Items on Solaris*

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	37	12%	<5%
2000	346	13%	<5%
3000	Timed out	14%	<5%

The following table lists the results for Save as Change on Solaris.

Table 3–12 *Results for Save as Change on Solaris*

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	15	6%	1%
2500	45	15%	1%
5000	102	13%	1%
7500	210	15%	1%

The following table lists the results for Change Status to Released (with Preliminary Items) on Solaris.

Table 3–13 Results for Change Status to Released (with Preliminary Items) on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	7	7%	<5%
2500	20	12%	<5%
5000	44	10%	<5%
7500	62	13%	<5%

The following table lists the results for Change Status to Pending on Solaris.

Table 3–14 Results for Change Status to Pending on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	26	6%	<5%
2500	88	6%	<5%
5000	194	16%	<5%
7500	300	16%	<5%

The following table lists the results for Change Status to Released (each item has at least 3 pending changes) on Solaris.

Table 3–15 Results for Change Status to Released (each item has at least 3 pending changes) on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	9	4%	<5%
2500	22	6%	<5%
5000	46	11%	<5%

Item Object

The following table lists the results for Copy and Paste BOM on Solaris.

Table 3–16 Results for Copy and Paste BOM on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	11	12%	1%
2500	47	10%	1%
5000	110	12%	2%

The following table lists the results for Expand BOM (All Levels) on Solaris.

Table 3–17 Results for Expand BOM (All Levels) on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
3596	20	50%	NA

Product Quality Management

This section provides breakpoint testing results for Product Quality Management.

PSR Object

The following table lists the results for Copy and Paste Affected Items (newly imported items) on Solaris.

Table 3–18 Results for Copy and Paste Affected Items (newly imported items) on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	51	9%	1%
2500	451	10%	1%
3500	748	10%	1%
5000	Timed out	10%	1%

The following table lists the results for Copy and Paste Affected Items (existing items) on Solaris.

Table 3–19 Results for Copy and Paste Affected Items (existing items) on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	261	4%	1%
2500	Timed out	10%	1%

QCR Object

The following table lists the results for Copy and Paste Affected Items on Solaris.

Table 3–20 Results for Copy and Paste Affected Items on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
2500	390	9%	1%
5000	Timed Out	NA	NA

Product Portfolio Management

The following table lists the results of several PPM scenarios on Solaris.

Note: Response times are in seconds.

Table 3–21 PPM Scenarios on Solaris

Scenarios	1200 Object Tree	2400 Object Tree
Save As Program	36	55
Save As Template	25	40
Create Project From Template	42	99

Table 3–21 (Cont.) PPM Scenarios on Solaris

Scenarios	1200 Object Tree	2400 Object Tree
Add 10 Team Members	29	43
Reschedule Root Project	17	34
Reschedule Child	15	27
Delegate Owner	18	29
Substitute Resource	14	28
Change Status	22	39
Expand All	160	Out of memory

Tree Details:

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

The following tables list the results for several PG&C scenarios.

Note: Each declaration has 2 specs. Example: Dataset 1000=Declaration with 500 parts + 2 specs.

Table 3–22 Object: Declaration Tab Name: Items

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Items	14	26	39	15%	5%
View Items Tab	4	4	4	NA	NA
Calculate Compliance	5	4	4	NA	NA
Release the Declaration	49	70	157	NA	12%
Import Items (no BOS)	NA	NA	NA	NA	NA

Table 3–23 Object: Declaration Tab Name: Mfr Parts

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Mfr Parts	12	22	52	10-15%	5%
View Mfr Parts Tab	4	6	6	NA	NA
Calculate Compliance	5	4	4	NA	NA
Release the Declaration	48	77	93	NA	12%

Table 3–23 (Cont.) Object: Declaration Tab Name: Mfr Parts

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Import Mfr Parts (no BOS)	NA	NA	NA	NA	NA

Table 3–24 Object: Declaration Tab Name: Part Groups

Scenario	250 Compositions	500 Compositions	1000 Compositions	App Server CPU Usage	DB Server CPU Usage
Copy & Paste Part Groups	12	20	52	10-15%	5%
View Part Groups Tab	4	6	6	NA	NA
Calculate Compliance	4	4	4	NA	NA
Release the Declaration	49	75	156	NA	12%
Import Part Groups (no BOS)	NA	NA	NA	NA	NA

Table 3–25 Object: Substance Groups Tab Name: Substances

Scenario	50 Substances	100 Substances
Search & Add Substances	8	13
View Substances tab	3	3
Import Substances	8	14

Table 3–26 Object: Specification Tab Name: Substances

Scenario	50 Substances	100 Substances
Search & Add Substances and/or Groups	7	11
View Substances tab	3	3
Import Substance Groups	7	11

Product Cost Management

The following table lists the results for several PCM scenarios.

Note: Response times are in seconds.

Action	Project with 5000 Items	Project with 10000 Items
Display project item tab	6	5
Search and Add Items into project from PC item master	57	NA

Action	Project with 5000 Items	Project with 10000 Items
Update item from PC item master	45	194
Add/modify Partners (10)	32	60
Import item from external source	64	160
Export item to external source	13	29
Project Save As	210	545
Quantity rollup	17	25
Display project AML tab	4	6
Display item BOM	6	5
Display of project analysis	6	5
Opening Responses Look Up	14	30
Project response price lookup	14	26
Set as best	37+47	80+105
Publish Quote History, Authoring	47+45	163+75
Publish Published Price Redlining	Out of memory - Major perf. Improvement made in 9368	
Items RFQ Creation	45	110
Open RFQ	10	110
Display of RFQ Responses	4	6
Response Status Tab	4	6
Opening RFQ Response Look Up	10	16
RFQ Response Look Up	55	98
RFQ Response status display	4	6
Edit/Modify Responses	150	550
Supplier Response Export	135	480
Supplier Response Import	90	180
Submit Responses	88	510
Costed BOM Comparison (Buyer) Report	25	45
Delete Project	2	2
Close RFQ	270	690
Delete RFQ	30	33
Lock RFQ	291	690
Data Sets: ? Project with 5000 items, each having two AMLs ? Project with 10000 items, each having two AMLs		

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.

The following table lists the results for basic file upload and download.

Table 3–27 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 1 GB	93	NA	<5%	5%
Download - 1 GB	160	NA	25%	5%
Upload - 2 GB	Crashes	NA	NA	NA
Download - 2 GB	349	NA	10%	5%

The following table lists the results for advanced file upload and download.

Table 3–28 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 5GB	441	429	<5%	5%
Download -5GB	161	157	15%	5%

Export

Export was tested using the Web Client.

The following table lists the results for Export scenarios.

Table 3–29 Items/BOM Export Scenario

Scenario	Data Set	Response Time	JVM Heap	CPU Usage - App Server	CPU Usage - DB Server	Comments
Export CSV	4.8 Million Rows	4 h 30 mins	1200 MB	20% continuous	<5% continuous	6 level deep BOM Output file size: 3 GB
Export PDX/aXML L	4.8 Million Rows	25 mins	1200 MB	20% continuous	<5% continuous	Same BOM as above. Duplicated rows won't be exported to PDX and aXML formats.

Table 3–30 Changes Export Scenario

Scenario	Data Set	Response Time	JVM Heap	CPU Usage - App Server	CPU Usage - DB Server	Comments
Export CSV	3000 Rows	3 h	1200 MB	20%	<5%	NA

Solution Results on AIX

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

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

- Application Server: PowerPC_Power6, P6 Processor, 4.7 GHz, 4 Cores, 16 GB RAM
- Database Server: PowerPC_Power6, P6 Processor, 4.7 GHz, 4 Cores, 16 GB RAM

Note: The response times for all of the tables are measured in seconds and are for single user only.

Product Collaboration

This section provides breakpoint testing results for Product Collaboration.

Change Object

The following table lists the results for Copy and Paste Affected Items on AIX.

Table 3–31 Results for Copy and Paste Affected Items on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	32	22%	5%
2000	255	25%	5%
3000	700	25%	5%
5000	Timed out after 15 min	25%	5%

The following table lists the results for Save as Change on AIX.

Table 3–32 Results for Save as Change on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	12	15%	5%
2500	41	15-25%	5%
5000	87	15-25%	5%
7500	168	15-25%	5%

The following table lists the results for Change Status to Released (with Preliminary Items) on AIX.

Table 3–33 Results for Change Status to Released (with Preliminary Items) on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	12	10%	5%
2500	20	10%	5%
5000	41	25%	5%

Table 3–33 (Cont.) Results for Change Status to Released (with Preliminary Items) on

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
7500	70	25%	12%

The following table lists the results for Change Status to Pending on AIX.

Table 3–34 Results for Change Status to Pending on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	31	10%	6%
2500	85	25%	6%
5000	185	25%	6%
7500	300	25%	6%

The following table lists the results for Change Status to Released (each item has at least 3 pending changes) on AIX.

Table 3–35 Results for Change Status to Released on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	10	10%	10-15%
2500	21	10%	10-15%
5000	43	10%	10-15%

Item Object

The following table lists the results for Copy and Paste BOM on AIX.

Table 3–36 Results for Copy and Paste BOM on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
500	11	15%	<5%
2500	50	15-25%	<5%
5000	107	15-25%	5-10%

The following table lists the results for Expand BOM (All Levels) on AIX.

Table 3–37 Results for Expand BOM (All Levels) on AIX

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
3596	18	10%	<5%

Product Quality Management

This section provides breakpoint testing results for Product Quality Management.

PSR Object

The following table lists the results for Copy and Paste Affected Items (Newly Imported Items).

Table 3–38 Results of Copy and Paste Affected Items (Newly Imported Items)

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	48	10%	5%
2500	355	10%	5%
3500	710	10%	5%
5000	Timed out after 15 minutes	10%	5%

The following table lists the results for Copy and Paste Affected Items (Existing Items).

Table 3–39 Results for Copy and Paste Affected Items (Existing Items)

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
750	292	5%	2%
2500	Timed out after 15 minutes	NA	NA

QCR Object

The following table lists the results for Copy and Paste Affected Items on Solaris.

Table 3–40 Results for Copy and Paste Affected Items on Solaris

Number of Affected Items	Response Time (in seconds)	CPU Usage - Application Server	CPU Usage - Database
2500	320	10%	5%
5000	Timed Out after 15 minutes	NA	NA

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.

The following table lists the results for basic file upload and download.

Table 3–41 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 1 GB	45	NA	11%	10%
Download - 1 GB	10	NA	15%	5%
Upload - 2 GB	86	NA	11%	5%
Download - 2 GB	30	NA	50%	5%

The following table lists the results for advanced file upload and download.

Table 3–42 Basic File Upload and Download

Scenario	Response Time - Web Client	Response Time - Java Client	CPU Usage - Client	CPU Usage - DFM Server
Upload - 7 GB	205	255	10%	12%
Download - 7 GB	160	156	15-25%	20%

Export

Export was tested using the Web Client.

The following table lists the results for Export scenarios.

Table 3–43 Changes Export Scenario

Scenario	Data Set	Response Time	CPU Usage - App Server	CPU Usage - DB Server
Export - CSV	3500	30	<10%	<2%
Export - PDX	3500	35	<10%	<2%

Client Requirements

This chapter describes client-side requirements.

Client Requirements

Clients require desktop class computers supported by the browsers. 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 Requirements](#)" on page 7-4 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. Therefore, 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 can use the desktop client for some 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 can reformat the table by itself.

The Java client connects directly to the application server using the application server specific native protocol, such as T3 with Oracle WebLogic.

Directory Server

In an effort to support industry standard authentication schemes and enable central management of user information, Agile PLM 9.3.6 supports integration with the industry leading directory server and LDAP-based authentication for the Agile solution suite.

Overview

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.

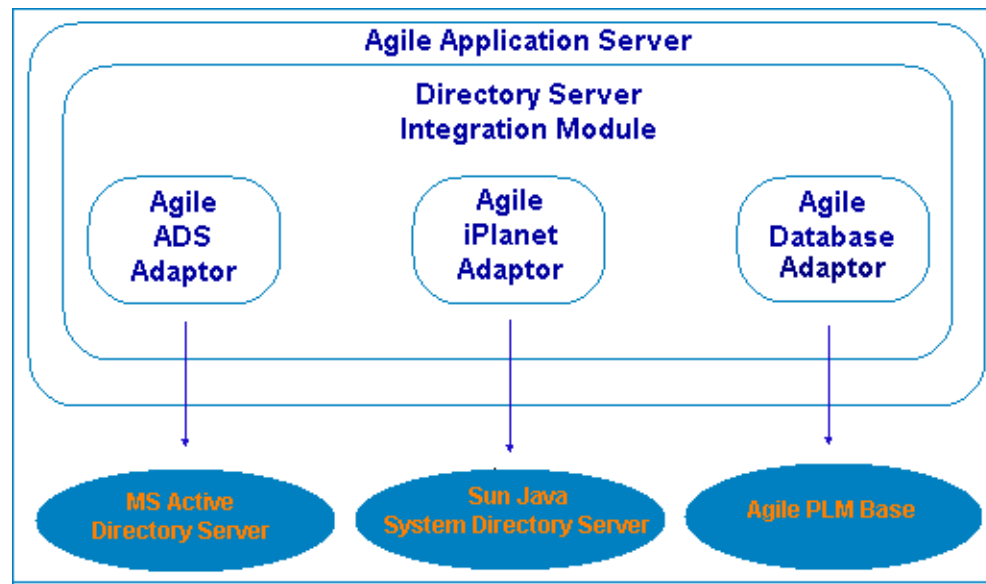
This chapter describes the architecture of the Agile PLM 9.3.6 Directory Server Integration Module, implementation details, configuration, and contains an FAQ section to help you gain a better understanding of how this solution works.

System Architecture

The figure below illustrates the high-level architecture of the Agile PLM 9.3.6 Directory Server Integration Module. As shown in the diagram, all components in the Agile solution use the Directory Server Integration Module platform component of the core Agile Application Server. 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 System Architecture



The Directory Service Integration Module is an interface that describes the contract between the module and the Agile PLM 9.3.6 platform. The interface does not make any assumptions about how the implementation provides these services. This interface driven design de-couples the Agile PLM 9.3.6 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, and so on, 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 9.3.6. 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 follow 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:

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 9.3.6 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, and the password, are managed by the source only and displayed in Agile as read-only attributes.

You can configure the Directory Server Integration Module 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. 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 those accounts cannot be used for authentication.

You can configure one directory server during the Agile installation. You can configure additional directory servers manually after installation. Agile provides two scripts to enable configuration after installation:

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

Frequently Asked Questions (FAQ)

The following is a list of frequently asked questions and their corresponding answers.

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. You must 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. 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.

File Management Server and Viewer

This chapter describes the file management server and vault, and the file viewer.

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 file name/Agile ID mapping information is maintained in the Agile database. When a user requests a file (Get, View, Checkout, and so on), 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, AutoVue for Agile PLM, and any other deployed File Managers. File Managers can be clustered in one location, if required.

The features of the Agile PLM 9.3.6 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 enables 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 AutoVue for Agile PLM in a Distributed Environment

If AutoVue for Agile PLM is used, an AutoVue for Agile PLM should also be installed locally with each distributed file manager. The local AutoVue for Agile PLM can be installed on the same system 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.

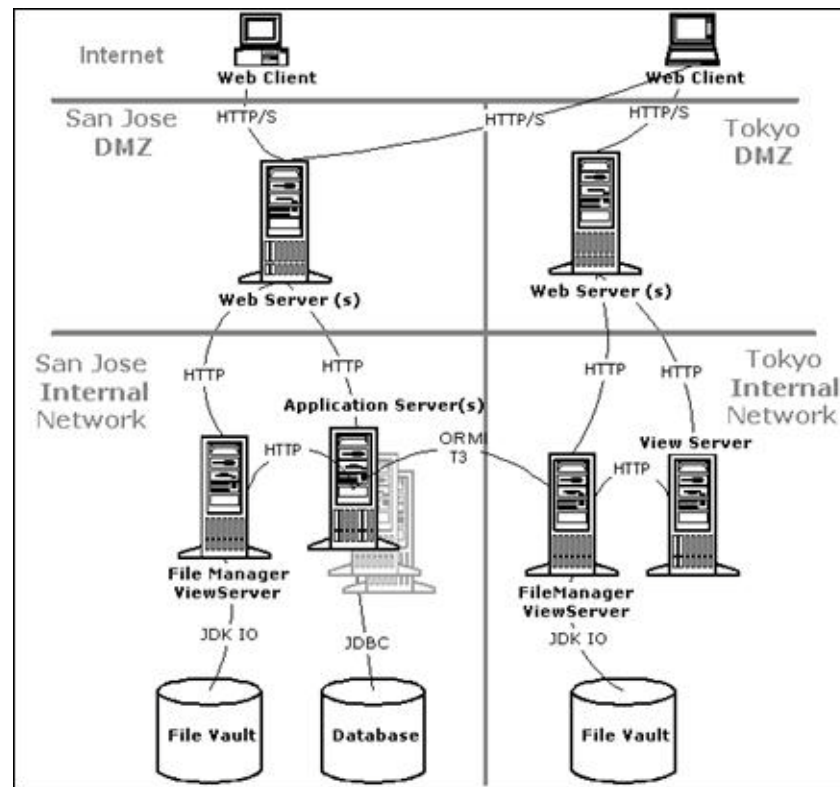
Agile recommends deploying the File Manager and AutoVue for Agile PLM on a single dual CPU system with the following minimum specifications:

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:

1. A user from the Tokyo site selects a file attachment to view.
2. 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 AutoVue for Agile PLM or directly, depending on whether the viewer is configured for the given file type.
3. Another user selects the same file attachment.
4. Because the file now resides in the Tokyo File Manager vault, it is served directly to the user.

Figure 6-1 Sample Distributed File Manager Configuration

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 usually necessary 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 back up all file vaults.

Do I need an AutoVue for Agile PLM for every distributed File Manager?

Yes. If you decide to use AutoVue for Agile PLM, then it is recommended that you install a viewer at every location where you have a File Manager installed. By installing AutoVue for Agile PLM 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.

AutoVue for Agile PLM

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

The AutoVue 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 AutoVue for Agile PLM on the same system or separate systems. No additional configuration is required except that specified for the Agile Web client and the caching viewer. The VueServlet is used to access AutoVue for Agile PLM across firewalls from external clients, such as the Agile Web client, on standard HTTP/HTTPS ports.

Caching Viewer

AutoVue for Agile PLM converts files from their native format to a streaming file format that is used to render 2D/3D images. These files can be cached and reused to improve viewer performance.

There are two types of streaming file caching; on-demand and offline. On-demand caching occurs in the background when a user views a file in AutoVue for Agile PLM. 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 viewer. AutoVue for Agile PLM should not be used for offline caching as performance can be impacted. Therefore, the caching viewer should not be installed on the same system as AutoVue for Agile PLM. Also, offline caching is only supported on the primary File Manager. Distributed file managers installed in remote locations cannot use offline caching.

AutoVue for Agile PLM Security

The primary security consideration pertaining to AutoVue for Agile PLM surrounds the native product data files that it caches when processing a view request. Deploying AutoVue for Agile PLM 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 AutoVue for Agile PLM to stay behind the firewall along with the other main components is a significant security benefit.

Communication and Data Flow

This chapter describes communication and data flow configurations used by the Agile system.

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.

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	7001 or 23791
Java Client	Distributed File Manager	HTTP	80/443
Portal Server	Application Server	T3	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
AutoVue for Agile PLM	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	NA	NA
File Manager	File Vault Mapped Drive	NA	NA
Application Server	Database Server	JDBC	1521
Application Server	LDAP Server	LDAP	389
SDK Client Application	Application Server	T3 or HTTP	7001 or 7777 or 23791

Client Communication

The Agile application adapts to your current network infrastructure and utilizes the available bandwidth to transmit data through 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 the application server-specific protocol. Oracle 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 must 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

The figure below 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, before completing any requested transactions. Additionally, client-side proxy software may be configured as a transparent or a non-transparent proxy.

Figure 7-1 Typical Network Configuration

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. 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 viewer. The application server connects to the file vault, database server, and optional LDAP directory server using the protocols.

Directory Server

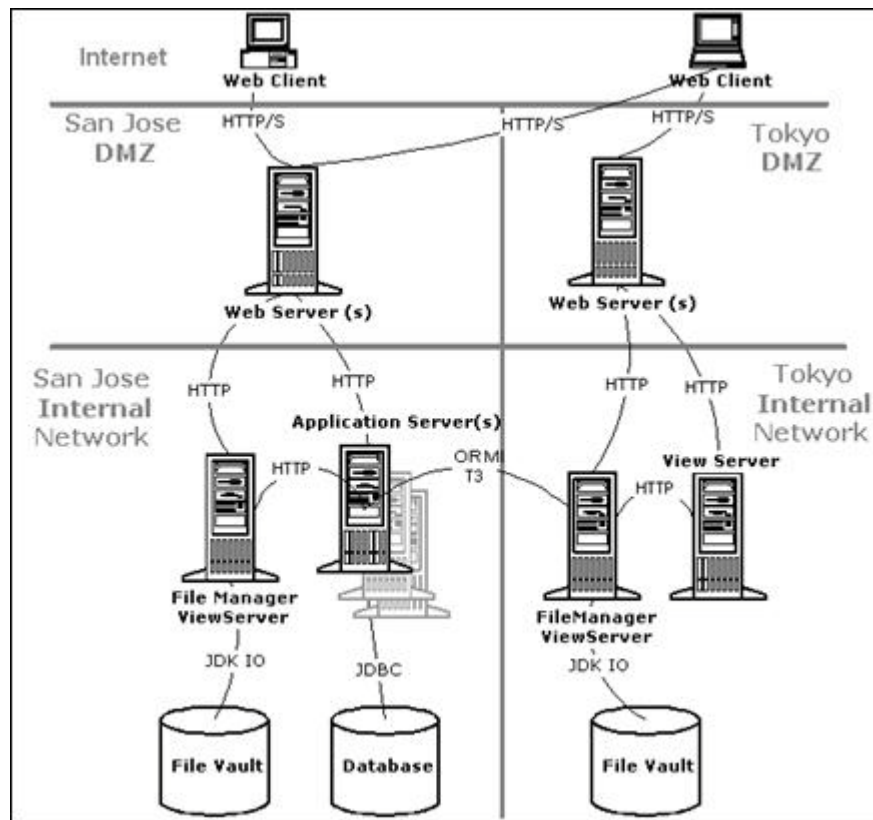
Directory server communication occurs using LDAP protocol. If a directory server is being used to manage accounts in Agile, then it is also used for authentication. Every time login occurs, the directory server is asked to authenticate the user. The application server also connects to the directory server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

Distributed File Management

As the figure below 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 redlining performed through AutoVue for Agile PLM. 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 7–2 Distributed File Management Server Example

AutoVue for Agile PLM

Viewing a file through the AutoVue for Agile PLM is handled differently than getting the file. If you view a file instead of getting it, the request is sent to AutoVue for Agile PLM to render the file. The viewer requests the file from the file server. The AutoVue for Agile PLM and file server communicate using the HTTP protocol. The file server gets the file from the vault and sends it back to the viewer. AutoVue for Agile PLM 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 9.3.6 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.

Performance Tips

This chapter describes recommendations and tips for optimal performance.

Tuning Memory for Java Applets

The Agile Web Client uses Java applets for advanced functionality. Examples include the Gantt Chart and AutoVue for Agile PLM. 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, AutoVue for Agile PLM, 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 WebLogic*.

Running PPM Gantt Charts

When using the PPM Gantt Chart on a project containing over 100 lines, make sure the Java Heap Size is 256MB.

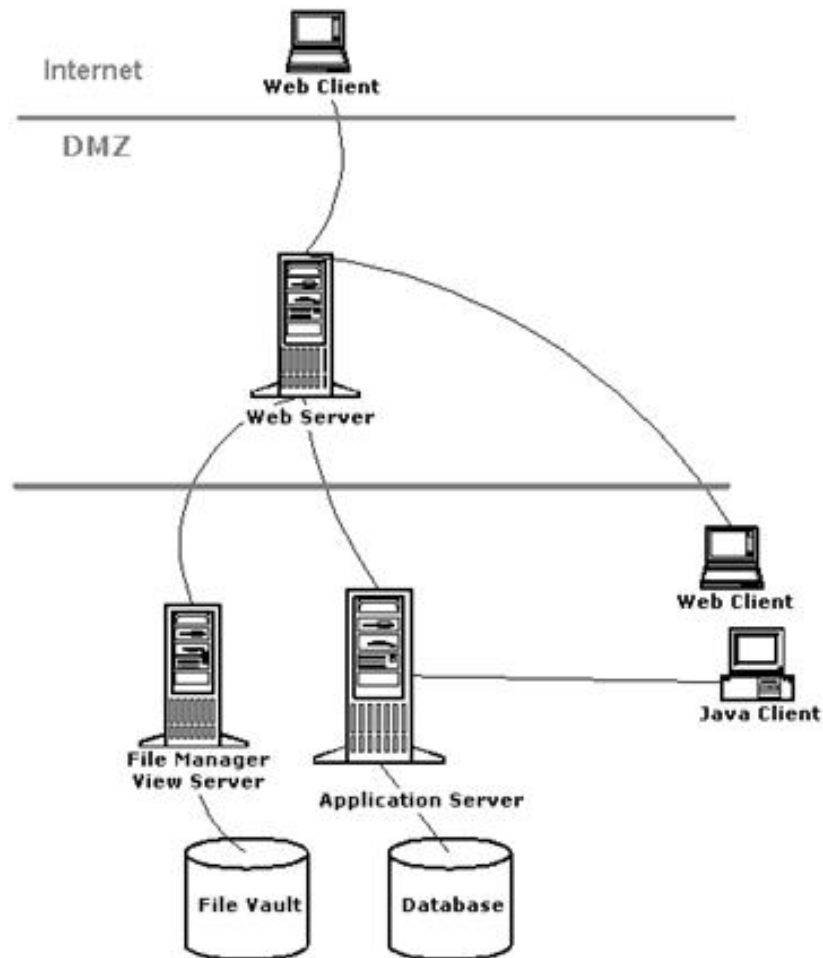
Optimum JVM Parameters on Solaris

After running several tests on Solaris systems 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
```

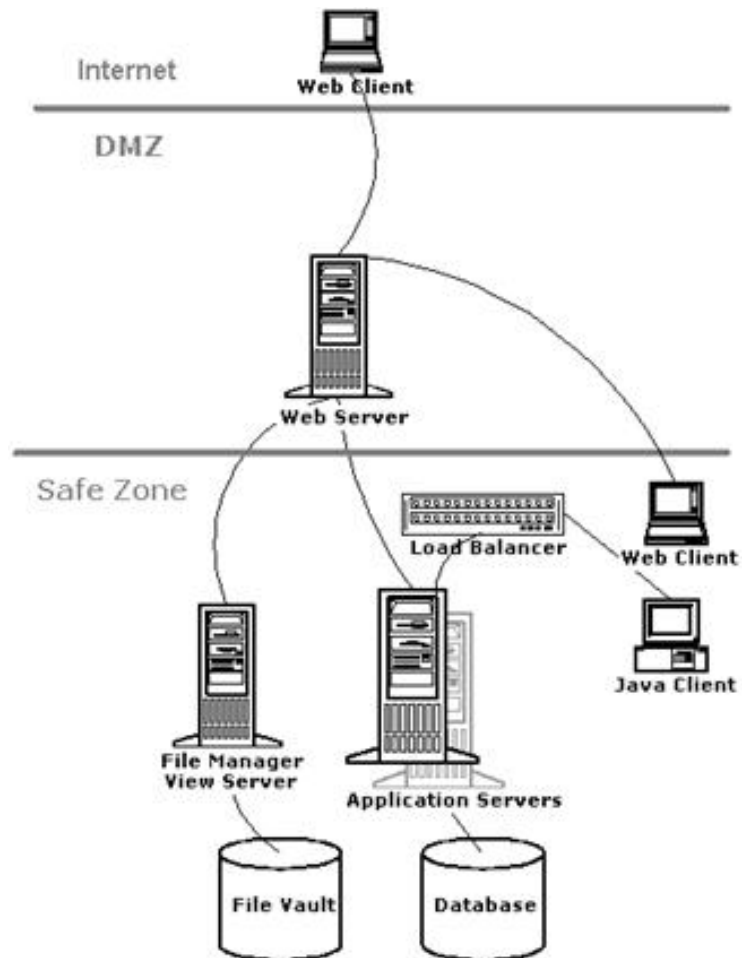
Small System Configuration

Figure 8–1 Example of a Small System Configuration



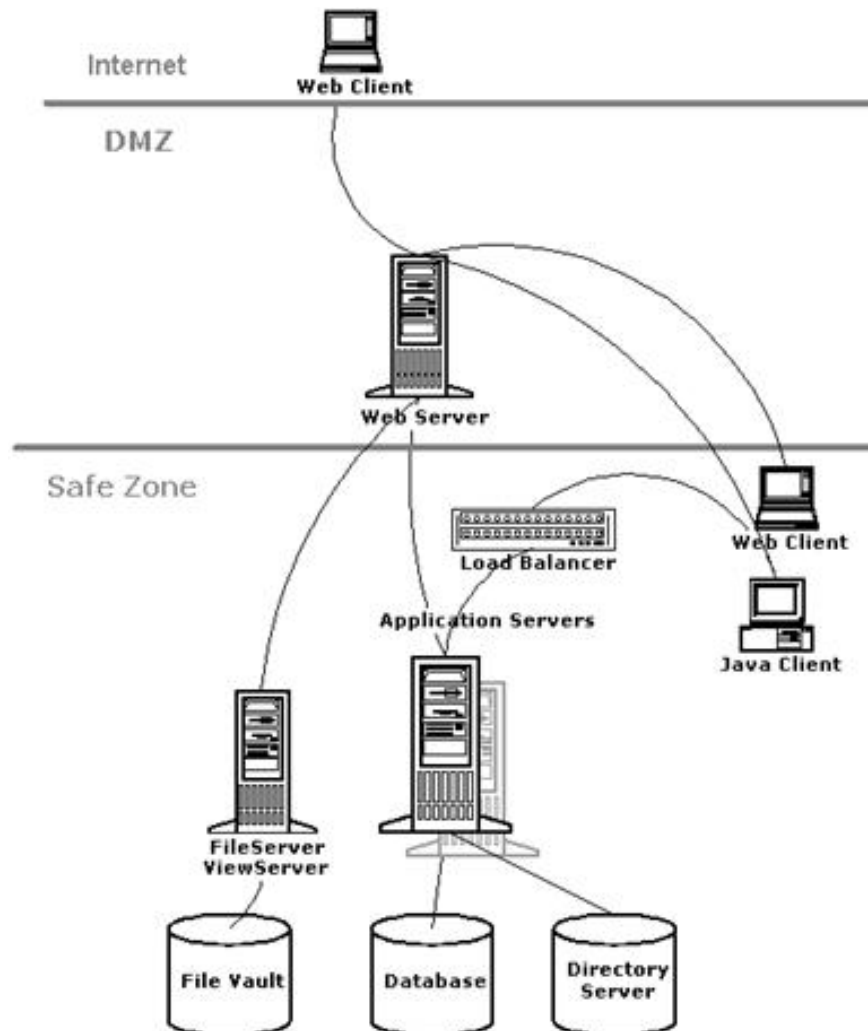
Medium System Configuration

Figure 8–2 Example of Medium System Configuration



Large System Configuration

Figure 8–3 Example of Large System Configuration



Extra Large System Configuration

Figure 8-4 Example of Extra Large System Configuration

