# PTC Windchill

# PDMLink Version R8.0

## *Architecture and Sizing Guide*

## *for*

## *HP Enterprise Servers*

*Version 1.7*
*July 26, 2006*

**Michael Finocchiaro,**

**Donald Mize**

**and Stephen Fleischman**

**High Performance Computing Division**

**Hewlett-Packard Company**

# Notice

This manual and any examples contained herein are provided "as is" and are subject to change without notice. Hewlett-Packard Company makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard Co. shall not be liable for any errors or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual or the examples herein.

Microsoft is a registered trademark of Microsoft, Incorporated.

Pro/ENGINEER, Pro/INTRALINK, Pro/MECHANICA, Pro/DESKTOP, all other Pro/series product names, Expert Machinist, Parametric Technology Corporation, PTC, Windchill, Computervision, CADDS, MEDUSA, Optegra, InPart, DesignSuite, CDRS, ICEM, and other PTC product names are trademarks or registered trademarks of Parametric Technology Corporation or its subsidiaries in the United States and in other countries. Names of products from other vendors mentioned in this document may be trademarks or registered trademarks of their respective owners. Oracle is a registered trademark of Oracle Corporation and Windows and Windows NT are registered trademarks of Microsoft Corporation.

**Hewlett-Packard Company**

**Hewlett-Packard Technical Computing Division**

**3000 Waterview Parkway**

**Richardson, Texas 75080, U.S.A.**

# Table of Contents


NOTICE....2

TABLE OF CONTENTS....3

TABLE OF FIGURES....5

TABLE OF TABLES....5

EXECUTIVE SUMMARY....6

EXECUTIVE SUMMARY....6

INTRODUCTION....7

SOFTWARE VERSIONS....7

**Determining Required Patches....7**

THE WINDCHILL PDMLINK BENCHMARK....8

THE WINDCHILL PDMLINK SOLUTION....8

**Software Components and Configuration Notes....8**

ARCHITECTURAL CONSIDERATIONS....13

**PA-RISC or IA-64....13**

**3-T or not 3T....14**

**Choosing Disks....17**

**Network Bandwidth....17**

**Configuring for High Availability....18**
Mirror/UX....18
MC/ServiceGuard....18

**Planning a Backup Strategy....19**
Omniback/Data Protector....19
Oracle Archiving and Hot Backup....19

**Partitions on hp High-end Servers....19**

About Partitions ..... 19
Partitioning and Multiple Windchill instances ..... 20

**SIZING THE WINDCHILL SERVER ..... 21**

**Overview ..... 21**

**Mass Storage for the Windchill Server ..... 22**

**Network Performance for the Windchill Server ..... 23**

**Sizing the Windchill Server for MC/ServiceGuard ..... 23**

**SIZING THE WINDCHILL ORACLE RDBMS SERVER ..... 23**

**Overview ..... 23**

**Disk I/O Performance ..... 24**

**Data Protection and High Availability ..... 24**

**Sizing the Oracle Database Server for MC/ServiceGuard ..... 24**

**Network Performance ..... 25**

**CPU ..... 25**

**Memory ..... 26**

**ADDITIONAL SERVER TYPES ..... 26**

**CAD Adapter ..... 26**

**Content Cache (Replication) Server ..... 26**

**Development Systems ..... 27**

**Testing Server ..... 27**

**PERFORMANCE TOOLS ..... 27**

**WHERE TO GET MORE INFORMATION ..... 28**

**APPENDIX A: SAMPLE CONFIGURATIONS ..... 28**

**Sample Windchill Server Configurations for 3-Tier Implementations ..... 28**
100 User Windchill 3-Tier, Using rp3440 ..... 28
100 User Windchill 3-Tier Using rx1620 ..... 29
500 User Windchill 3-Tier Using rp3440 ..... 29
500 User Windchill 3-Tier Using rx4640 ..... 29
1000 User Windchill 3-Tier Using rp4440 ..... 29

1000 User Windchill 3-Tier Using rx7620 ....29
2000 User Windchill 3-Tier Using rp7420 ....29
2000 User Windchill 3-Tier Using rx8620 ....29

**Sample Oracle Database Server Configurations in a 3-Tier Implementation....30**
100 User Oracle Configuration Using rp3410 ....30
100 User Oracle Configuration Using rx1620 ....30
500 User Oracle Configuration Using rp3440 ....30
500 User Oracle Configuration Using rx1620 ....30
1000 User Oracle Configuration Using rp3440 ....30
1000 User Oracle Configuration Using rx4640 ....30
2000 User Oracle Configuration Using rp4440 ....31
2000 User Oracle Configuration Using rx4640 ....31
2000 User Oracle Configuration Using rp8420 ....31
2000 User Oracle Configuration Using rx8620 ....31

**CAD Agents....31**
CAD Agent Server, Using C-Class, c8000 Workstation ....31
CAD Agent using xw8200 Visualize Workstation ....32

**Content Replication Servers ....32**
Replication (Content Cache) Server using the rp3410....32

**APPENDIX B: PERFORMANCE TOOLS....32**
Overview ....32
GlancePlus ....33
GlancePlus Pak 2000 ....33
OpenView Performance Agent (previously known as MeasureWare) ....33
OpenView SMART Plug-Ins, (SPI's) for Databases....33
OpenView Performance manager (Previously known as PerfView) ....33
OpenView Performance Monitor (Previously known as PerfView Monitor)....34
OpenView Performance Planner (Previously known as PerfView Planner) ....35

# Table of Figures

Figure 1 - Windchill Product Architecture ....10
Figure 2 Web Architecture ....13
Figure 3 - Windchill 3-Tier Platform Example....16
Figure 4 - Windchill Server Clusters ....17

# Table of Tables

Table 1 - HP-UX Machine Configurations....14
Table 2: Comparison of 3-Tier, Load-Balanced cluster and Monolithic Architectures....15
Table 3 - Disk Configurations ....17
Table 4 - Standard Network Transfer Rates ....18
Table 5 - Availability of Virtualization on HP-UX Servers ....20
Table 6 - Partitioning and Windchill Sizing ....21
Table 7 Windchill Server CPU and Memory Requirements....22
Table 8 Database server CPU and memory requirements: non-CAD users. ....25
Table 9 Database server CPU and memory requirements: CAD users.....25

# Executive Summary

During the past year Hewlett-Packard Company and PTC have jointly participated in activities to understand and characterize system performance with PTC Windchill PDMLink software running on Hewlett-Packard computer systems products. This was accomplished so that we could help ensure our customers' success by providing advice on system configuration and sizing.

The results of our efforts yield significant benefits in two areas. The first is performance data for PTC to use in making software performance improvements and subsequently a higher performing product with improved quality for customers. The second is that it provides system performance information so that we can advise customers regarding the configuration and performance of Windchill implementations on HP servers and workstations. These results allow us to provide sizing information and guidelines in such documents as this.

This document, “Windchill Sizing Guide for hp Computer Systems”, is one of a series. As changes occur in software, such as a new release of Windchill or a new release of the HP-UX operating system, this document will be updated to reflect the most recent performance improvements. As new hardware technology or improvements become available, they will be tested allowing HP and PTC to provide you with the most recent information regarding the performance of the Windchill product line on Hewlett-Packard computer systems.

# Introduction

The purpose of this document is to give an overview of the Windchill Solution and to provide general sizing information with example system configurations based on our most recent benchmark results.

The recent benchmarking of PTC's Windchill software on Hewlett-Packard computers has resulted in the generation of a number of technical documents. These documents provide detailed information on the architecture, tuning, and installation of Windchill on Hewlett-Packard server systems. The purpose of the documents is to provide the necessary information that enables the selection and implementation of appropriately sized systems that meet performance requirements. It should be noted that acquiring the appropriate hardware and software is only one of many steps necessary to implement an enterprise Windchill solution.

# Software Versions

The sizing information presented in this document applies to Hewlett-Packard computer systems with the following software versions.

HP JDK 1.4.2_10

HP-UX Version 11.23 – see the HP Sizing Guide for Windchill PDMLink 8 for appropriate kernel parameters

Oracle 9.2.0.7 (9iR2) 64 bit for HP-UX 11i

PTC Windchill PDMLink R8.0 with M010

HP Apache Web Server 2.0.55.02 and mod_jk/1.2.10 or 1.2.15

Tomcat 5.0 (PTC) or Tomcat 5.5.9 (HP) Servlet Engine[1]

RetrievalWare 8.0

British Telecom Aphelion LDAP Directory Server 2003

SilkPerformer 7.3.0 plus HotFix

## *Determining Required Patches*

There are several patches required for a successful implementation of Windchill. Current patches can be acquired either individually or collectively as part of a patch bundle. Hewlett-Packard distributes its patch bundles for the HP-UX operating system as part of a quarterly release called the HP Extension Software CD-ROM that is included in the media updates for software support contracts. Current patch bundles are included in HP-UX systems that have the HP-UX operating system installed at the factory. For systems that are installed at customer sites it is highly recommended that the currently available HP Extension Software CD-ROM be installed as a part of the installation process or during a maintenance period. As an added convenience you may contact your local Hewlett-Packard or Parametric Technology support person and inquire about the "HP Software for Windchill" CD which contains the latest patches.

Patch numbers may be changed or be made obsolete by newer patch releases. To determine the recommended patches and other support information relating to Oracle or Windchill software, go to the PTC and Oracle web sites.

[1] PDMLink 8 at this printing only supports Tomcat 5.0 that is delivered with the PTC Windchill PDMLink CDs. Version 5.5.9, delivered with HP's Web Server version 2.0.55.02 on HP-UX 11.11 and 11.23 contains Tomcat 5.5.9. This version was proven to be more scalable than 5.0 in benchmarking due to its superior load-balancing capabilities. However, this is only an issue for customer sites with > 400 active users.

For patches relating to HP-UX 11i go to http://www.hp.com and select support or contact the local Hewlett-Packard sales representative.

# The Windchill PDMLink Benchmark

Workload simulation and benchmarking of the Windchill application took place recently at the Hewlett-Packard Americas Solutions Center performance laboratories in Cupertino, California in partnership with PTC. The tested systems were a combination of Hewlett-Packard rx7620, rx8620, and rp8420 HP-UX servers with VA7100, and EVA3000 SAN disk arrays.

Workload generation was accomplished by using HP Proliant Xeon based servers to run a variety of workloads representing the actions of a range of users, which could number up to a thousand. Workload definitions and scripts were provided by PTC and implemented using Segue's SilkPerformer load testing tool.

One of the goals of this benchmark was to simulate 'real world' users. That is, users who are active, connected, and performing tasks, with think-times and delays, as they would in their normal everyday working environment. The configurations described in Appendix B are based on the results of this benchmark.

SilkPerformer by Segue was used to control the tests, to capture response times for each client transaction and to provide statistics from each test run. Hewlett-Packard's performance tools were used to capture a wealth of system resource utilization information for the systems under test.

This combination of tools allowed the test team to capture a broad set of statistics and system performance information from which we developed system configurations, with the appropriate amount of system processing power and memory, that will provide the best overall performance for your workload requirements.

# The Windchill PDMLink Solution

The Windchill architecture is a three-tier application architecture designed and optimized for the deployment of business information applications. The client tier is the presentation layer of the architecture and uses third-party web browsers executing a combination of HTML, JavaScript, and Java applets to accomplish user tasks. The next tier, or middle tier, is the application tier and provides the business logic to support the necessary transaction processing. This tier is made up of third-party web servers, such as Apache, a servlet/jsp container such as Tomcat and the Java-based Windchill method servers. The third tier uses an Oracle 9i Database Management System (DBMS) to store structured and unstructured data. Other data to consider for mass storage include the external file vaults (which must be physically mounted or located on the machine with the Windchill application server), the LDAP directory, and (optional) the RetrievalWare indexes.

An enterprise wide collaborative product data managemen system solution using Windchill and Hewlett-Packard computer systems consists of several hardware and software components, each of which requires special consideration. Presented within this document is a discussion of the various components along with the information necessary to determine the amount of system resources, CPU, memory, disk, required for a particular implementation of Windchill based on our performance test results.

The Windchill architecture is rather complex so here are some brief descriptions of the components and some graphics to help you understand the different pieces and how they fit together.

## *Software Components and Configuration Notes*

- **PDMLink R8.0**
  This is the basic product for PDM from PTC. It consists of a java codebase where all the compiled java classes are stored as well as .jsp files (for presentation of data). Windchill runtime consists of a Server Manager, which manages the client connections and the various Windchill caches, one or more Method Servers, which manage the business objects and talk with the database and the Java RMI clients, and an optional Background Method Server, which manages the various Windchill queues. The primary

configuration file for Windchill is called wt.properties. Typically Windchill is launched with a windchill user (e.g. wcadmin

- **Info*Engine**
  Info*Engine is built-in to Windchill but can run stand-alone as well. It is the tool used by Windchill for communication between applications. For example, through Info*Engine, Windchill can talk to the LDAP server, toSAP, another Windchill, Pro/INTRALINK, another database, or a mainframe application via small programs called adaptlets in the Info*Engine server. Its schema is stored in the Aphelion LDAP Directory as discussed below. As it is built-in to Windchill, it is automatically launched when Windchill is launched.

- **File Vault**
  Windchill offers the option of storing content (CAD files, documents, etc) into either the Oracle database using Binary Large Object Blocks (or BLOBS) or onto an external file system (or vault) with only metadata stored in Oracle.

  Primarily this choice depends on the type of data that is being stored. If the data is primarily documents, (e.g., PowerPoint and Word documents, etc), from an administrative point of view, BLOBs are probably a better solution. However, if you are storing large CAD data files, a file vault is highly recommended. Things to keep in mind for a vaulted configuration:

  1. Always locate the vault as close if possible to the Windchill server if not on the Windchill server itself. Note that the file system where the vault lives must be physically attached to or mounted to the Windchill application server(s).
  2. Remember to include the vault in your backup procedure, cf. the Backup section later this report.
  3. Think about Windchill File Replication for getting the data as close as possible to its consumers, cf. the Windchill Administrator's Guide.. This is called content replication.

- **Optional Windchill modules**
  Windchill has many optional modules that include: managing sourcing, Sourcing Factor, product planning, Product Planning Factor, and integration with ERP systems, Windchill Production Factor. See the graphic below for examples.

  Additional Info*Engine adapters can also be run either on the Windchill server or on separate servers. If using TIBCO or another EAI (Enterprise Application Integration) package, it is recommended to run this on a separate Proliant, MS Windows based server.

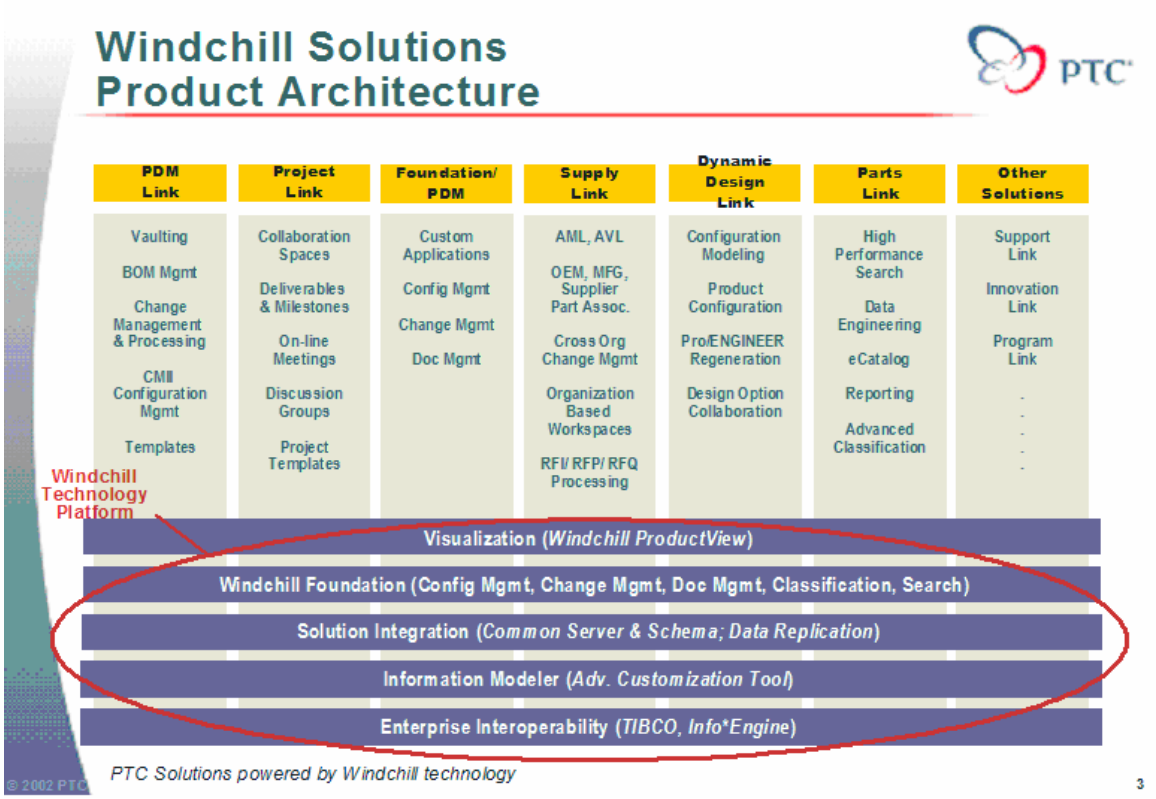


**Figure 1 - Windchill Product Architecture**

- **Windchill Visualization Services**
  Windchill Visualization Services provides Windchill with 3D representations of the CAD data stored in Windchill and enables on-line and off-line collaborative modification of CAD data in Windchill. The Windchill Engineering Factor provides adapters to standard CAD tools, such as: CADDS, CATIA, Pro/E, Unigraphics, for conversion of CAD data into thumbnails and the ProductView format. ProductView is then used for on-line collaborative design. In the modules listed above, PDMLink, ProjectLink and Dynamic Design Link include Visualization Services by default.

  The Workgroup Managers are large consumers of resources and can require a graphics card. Typically, several UNIX or Wintel workstations are installed on the same network as the Windchill server and CAD agents run for each CAD application to be processed (e.g. CADDS, Pro/E, CATIA, etc).

  **Note:** Thumbnail generation creates DerivedImage objects in the BLOBS. For better performance, these can be replicated into the file vault. Please use the Replication Manager to set up rules for replicating all derived objects to external file vault or see HP or PTC for help.

- **Oracle DBMS 9iR2**
  Oracle is the database for Windchill objects and schema. Version 9.2.x is supported. It is accessed from Windchill via a Java Database Connection Adapter (JDBC) and requires a listener for communication between clients such as the Windchill Method Server and the database. At runtime, Oracle runs executables for the database and for the listener. The database parameters are configured in a init.ora files or binary spfiles, whereas the listener is configured in listener.ora and tnsnames.ora. The tool for starting the database is sqlplus for Oracle 9.x. Typically, the database has a database administrator user (e.g. oracle) who starts both the listener and the database.

Oracle 9i comes in two versions, Standard and Enterprise, where the Enterprise Edition costs roughly 2x the Standard Edition. Windchill does not use features from the Enterprise Edition so you can safely use the Standard Edition. However, if you are deploying a large (>500 active users) Windchill installation and need more than 2 CPUs for the Oracle server, Enterprise Edition is obligatory.

Oracle 9i RAC is another Oracle option but at this time, it has not been tested in HP / PTC benchmarking. It could be an interesting alternative when the database has more than 3 Tb of data or in the case of the integration of more than active 200 WildFire seats.

- **Java 1.4.2_10 or higher**
  Java is the language in which Windchill is written. Java code (.java files) are compiled into class files (.class) and optionally bundled into Java archives (.jar files) which are then run in a Java Virtual Machine (JVM) for execution. The JVM is machine- and platform-dependant.and hp develops its own JDK (Java Development Kit which includes the Java source and the Java compiler (javac) as well as other tools). To run a Java program, we specify the Java program (java) and the .class or .jar file to be run.

- **Apache 2.0.55 or higher**
  Apache is a web server from Apache.org. The Apache server on HP-UX has been specially ported and optimized. It is fully supported by hp and highly scalable. It consists of several httpd executables which intercept and route HTTP requests coming from clients. It talks to the Tomcat servlet engine via a connector (mod_jk.so) using an Apache Java protocol (ajp13) . Apache also manages the user authentication through an LDAP server (in this case Aphelion described below) via an Apache module (auth_ldap.so). Apache configuration is done in config files (.conf) the first of which is called httpd.conf. It is launched as root and its child processes run under the www user.

  HP-UX11i comes with Apache and Tomcat pre-installed. This is the recommended configuration as it is the preferred solution from PTC, requires a minimum of setup and is free. It is highly recommended to use the Apache/Tomcat that comes with HP-UX and to regularly update the Apache with the latest security patches (http://www.hp.com/go/webservers).

  Apache can be run on separate servers from the Windchill server if desired. This would allow for less expensive web-farms front-ending a larger Windchill server. In this variation of the load-balanced cluster design, it is required to copy all static HTML and JSP to the web-server tier. This includes a full copy of the Windchill codebase to all Apache webservers on the front-end.

  The mod_deflate module in Apache can dramatically decrease bandwidth requirements and is highly recommended for deployment. Note that the libz.so module is required on the server. In June 2006, this library will be distributed free with the Internet Express Package for HP-UX. Otherwise, it is available via shareware.

- **Tomcat 5.0 or higher**
  Tomcat is a Java program which manages servlets and Java server pages (jsp). Servlets are Java programs which run on a server and intercept and process client requests. Requests to Windchill are passed from Apache through a HTTPGatewayServlet running in Tomcat to the Windchill Method Server. The Windchill portal page is written in jsp and compiled by Tomcat. Tomcat is configured by the server.xml file. Web applications which run under Tomcat are configured in web.xml files. At runtime, Tomcat runs under a JVM which is launched with the wttomcat_start (Windchill operation). It is typically run as the windchill user (e.g., wtadmin).

  For configurations of over 400 active users, it is suggested to use multiple Tomcats. The ideal configuration for Tomcat load-balancing is version 1.2.15 of mod_jk.so and Tomcat 5.5. At publication time, HP will be releasing mod_jk/1.2.15 in late August 2006 and Tomcat 5.5 support from PTC will probably come with R9 of Windchill. For configurations with more than 400 active users, please contact HP or PTC for help in this area.

- **Aphelion Directory 2003**
  Aphelion Directory 2003 is an LDAP server. It contains both the user information and schema information for Info*Engine. It is a third-party application from Syntegra and a required part of the installation process. At runtime, it consists of lde processes which listen on port 389 for ldap protocol requests (ldap://…). Typically, it is launched as root at boot time.

  Aphelion Directory Server is relatively light-weight but if there are several projects to manage on separate machines, it is suggested to place it on a separate server. It is not recommended to rely on Aphelion for a global corporate LDAP solution. Typically, iPlanet Directory Server or another third-party application will be used. In this case, there needs to be a connection to the corporate LDAP from the Apache/Windchill server. Only Info*Engine-specific information will be stored in Aphelion and Windchill will be configured to use the corporate LDAP for authentication. It has been demonstrated that Aphelion behaves better when it is limited to 1 CPU.

- **RetrievalWare 8.0 (optional)**
  RetrievalWare is an optional index/search server. It enables full-text searches of the entire Windchill database and file vaults for easy location of parts and documents. It is a third-party application from Convera. It consists of JBOSS Java based application server with administrative and indexing capabilities. It is connected to Windchill through the JBoss application. Its data is stored in a windchill_indexes directory and in the Oracle database. This data can take a long time to generate so it should be backed-up along with the Oracle database and the file vaults. Typically, it is launched either as the windchill user (e.g., wtadmin) or as a separate rware user (depending on the level of security required on the system).

  RetrievalWare requires quite a bit of resources and thus it is recommended that it is run on the Oracle database server or on its own 2-processor machine. Note also that the generation of indexes is single-threaded and can take an excessive amount of time. It is highly recommended that the generation the RetrievalWare indexes be made part of the backup strategy. Below is a diagram showing all the components that are mentioned above showing the various connections between them.

  The diagram below depicts all of the components that are mentioned above showing the various connections between them.

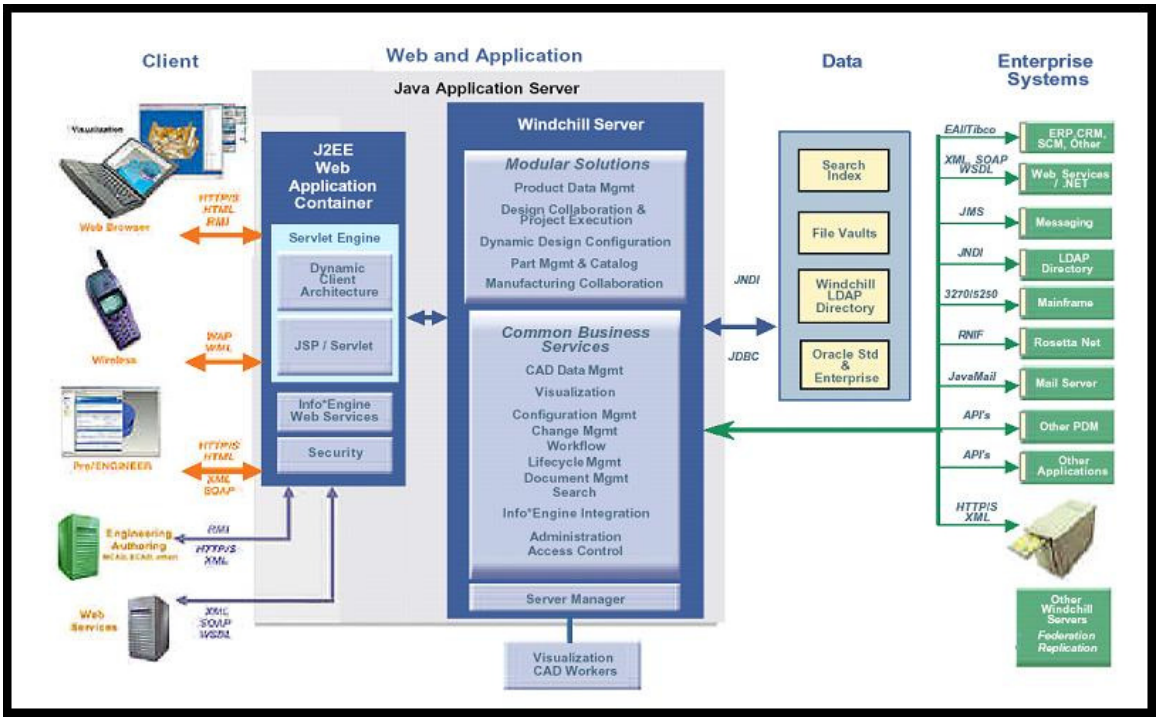


**Figure 2 Web Architecture**

# Architectural Considerations

## *PA-RISC or IA-64*

Starting with PDMLink R8, users now have a choice of two HP-UX based architectures: PA-RISC and Intel Itanium processors. The benchmarking in March 2006 demonstrated that the two platforms exhibit similar core-for-core scalability and performance. Since the long-term strategy for HP is on the Integrity architecture, it is recommended, where possible, to deploy on Intel Itanium processors.

Some considerations for Itanium:

- Not all applications run on Itanium today. RetrievalWare and Aphelion will need to be run on a separate Windows or PA machines.
- Some Windchill modules are still pending full support (ESI adapters, etc) so please check with PTC or HP support if you are deploying more than a basic PDMLink or ProjectLink.

The PA-RISC machines start with 'rp' in the product name; the IA-64 machines start with 'rx'. Please see the table below for details on the various models and their capacities.

| *Server Name* | *Max CPU's* | *PCI Slots* | *Max Memory* | *Internal Storage (max)* |
|---|---|---|---|---|
| hp Integrity rx1620 | 1-2 Itanium® 2 1.3 or 1.6 GHz | 2 | 16 Gb | 600 Gb |
| hp 9000 rp3410 | 1-2 PA 8900 800 MHz | 2 | 6 Gb | 900 Gb |
| hp Integrity rx2620 | 1-2 Itanium® 2 1.3 or 1.6 GHz | 4 | 32 Gb | 900 Gb |
| hp 9000 rp3440 | 1-4 PA-RISC 8900 800 MHz or 1.0 GHz | 4 | 32 Gb | 900 Gb |
| hp Integrity rx4640 | 1-8 Itanium® 2 1.5 or 1.6 GHz | 6 | 128 Gb | 600 Gb |

| | | | | |
|---|---|---|---|---|
| hp 9000 rp4410 | 1-4 PA-RISC 8900 800 MHz or 1.0 GHz | 6 | 128 Gb | 600 Gb |
| hp 9000 rp4440 | 2-8 PA-RISC 8900 800 MHz or 1.0 GHz | 6 | 128 Gb | 600 Gb |
| hp Integrity rx7640 | 2-16 Itanium® 2 1.6 GHz | 15 | 128 Gb | External Storage |
| hp 9000 rp7420 | 2-16 PA-RISC 8900 1.0 or 1.1 GHz | 15 | 2-128 Gb | 1.2 Tb |
| hp Integrity rx8640 | 2-32 Itanium® 2 1.6 GHz | 16 | 256 Gb | External Storage |
| hp 9000 rp8420 | 2-32 PA-RISC 8900 1.0 or 1.1 GHz | 16 or 32 | 2-256 Gb | 1.2 Tb or 2.4 Tb |
| hp Integrity Superdome 32-way | 2-32 Itanium® 2 1.6 GHz | 48 | 256 Gb | External Storage |
| Superdome 32-way | 4-32 PA-RISC 8900 1.1GHz | 48 | 256 Gb | External Storage |
| hp Integrity Superdome 64-way | 2-64 Itanium® 2 1.6 GHz | 96 | 512 Gb | External Storage |
| Superdome 64-way | 4-64 PA-RISC 8900 1.1GHz | 48-96 | 512 Gb | External Storage |
| hp Integrity Superdome 128-way | 6-128 Itanium® 2 1.6 GHz | 192 | 1 Tb | External Storage |
| Superdome 128-way | 12-128 PA-RISC 8900 1.1GHz | 96-192 | 1 Tb | External Storage |

**Table 1 - HP-UX Machine Configurations**

## *3-T or not 3T*

Modern web applications are usually deployed on a multi-layer or 3-Tier architecture. What this means is that the application layer (Web and Java servers) are not placed on the same machine as the backend or database layer. In the case of Windchill on HP-UX, there are several possibilities:

- Monolithic configuration: all the components on the same machine – or –
- 3-Tier configuration with Windchill, Tomcat, and Apache on one machine and Oracle 9i and RetrievalWare on another machine.
- Load-Balanced Cluster: Load-balancing switch forwarding requests from clients to Windchill servers (see Figure 3 below) with the Oracle server residing in the third (data) layer

The following table gives some basic differences between these choices:

| Feature | 3-Tier Architecture | Monolithic Architecture | Load-Balanced Cluster |
|---|---|---|---|
| Scalability | Full CPU's available to Windchill;<br>Very scalable | CPU's shared between Windchill and Oracle; Less scalable | Separate Tomcat's for better JSP performance; easy to add new servers to Windchill/App Server tier for future expansion |
| Investment | Requires at least two machines | Requires only one machine. Oracle licensing costs may be adversely affected unless Oracle recognized system partitioning is used. | Requires load-balancer and minimum of three machines (two for Windchill, one for Oracle); can use less expensive machines for Application Tier |
| Vaulting | Vault must stay on Windchill machine so disk space must be available to both machines | Requires enough disk for both database and vault | Vault must be NFS mounted from data layer to all machines in load-balanced cluster |
| Security | Protection of the database as it is isolated on a different machine | Database vulnerable to system failure | Normally, all data protected upstream by reverse-proxy server and isolated from Application Layer |
| High Availability | Can be integrated in an MC/ServiceGuard cluster | No H/A possible except if a standby (active-passive) | Some H/A provided by load-balancing switch (no failover |

| | | | |
|---|---|---|---|
| | | server is used in the cluster. | on Windchill). Data backend can be H/A clustered for NFS and Oracle |
| Performance | Measured 30-50% better than monolithic depending on workload | Fewer CPU cycles are available to Windchill for MethodServer but hard to say whether performance is better or worse than other cases | Probably better depending on amount of JSP and DCA manipulated |
| Robustness | Having the two systems separated gives more resource to the front-end and thus the 8:30AM syndrome can be more easily handled | No built-in protection against a sudden burst of activity on Windchill | Robust and scalable because new servers can be added to cluster as needs expand |
| Manageability | Increased amount of work due to two systems. RetrievalWare more complicated to install. | Easy to manage since there is only one machine. | Management of load-balancing switch as well as machines in cluster plus perhaps an H/A cluster for data backend |

**Table 2: Comparison of 3-Tier, Load-Balanced cluster and Monolithic Architectures**

Please note that in sizing systems in a highly available configuration using MC/Service Guard, you should plan for the worst-case fail over scenario, i.e., size the Windchill system based on the monolithic sizing data and not on the 3-tier data. This is because if a system failover occurs and if you have sized for the monolithic case, the impact on the users will be as little as possible and the performance should still be sustainable.

As a general rule, as the number of users increases, you will get better performance from a 3-Tier architecture. Also the benefits of High Availability and Security need to be considered; please see the High Availability section later in this report. For a small development site, a monolithic configuration is probably adequate. For a medium-to-large deployment, a 3-Tier solution is a better answer. The figure below depicts a typical three-tier architecture.

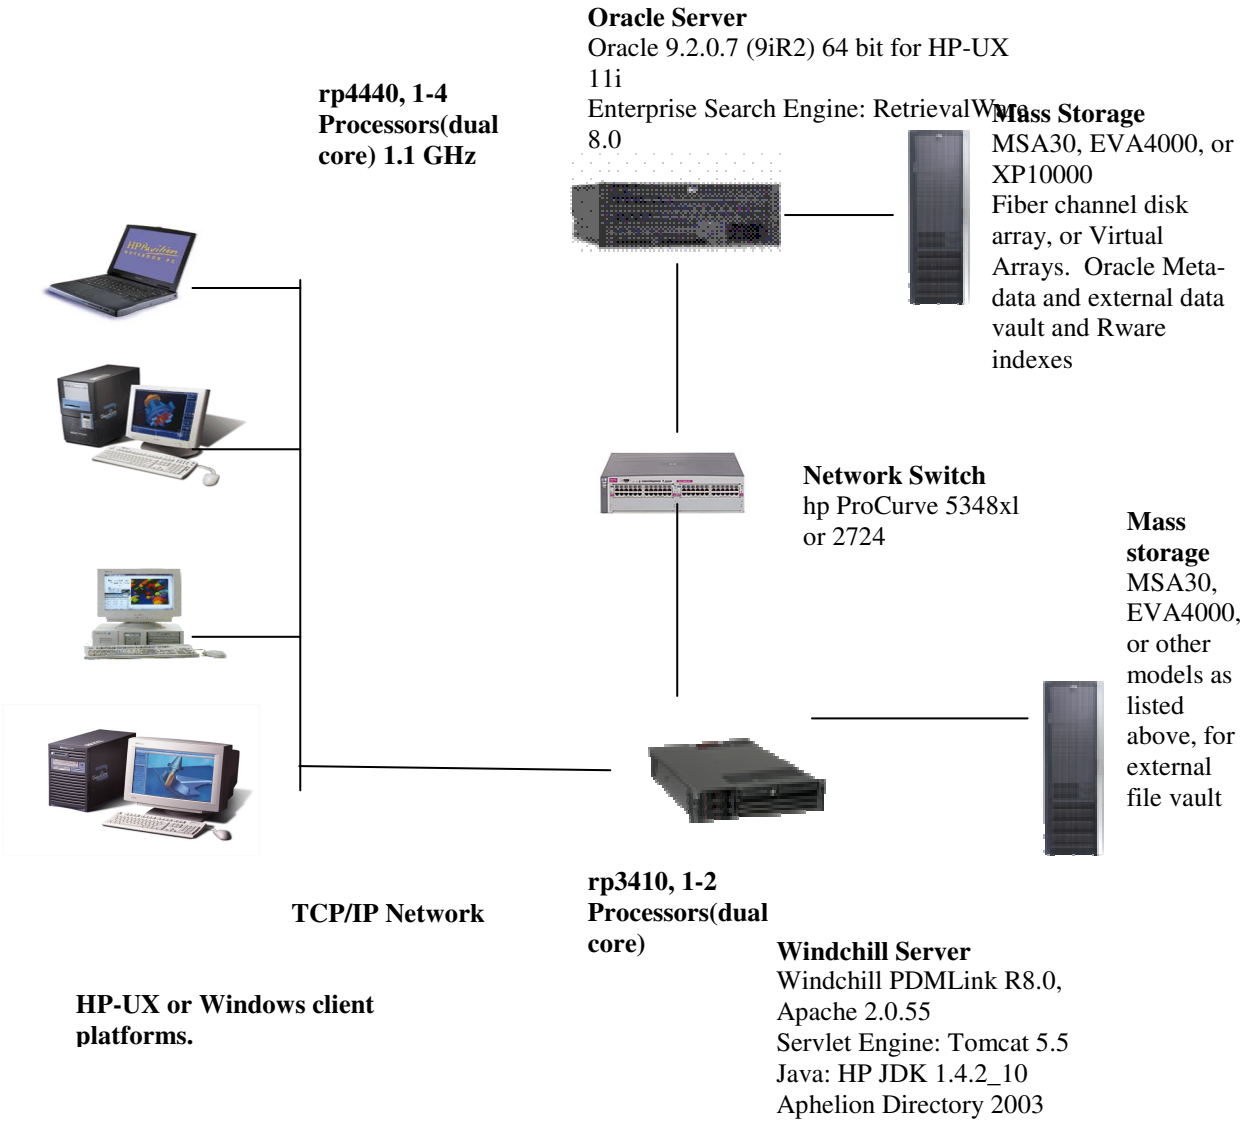


**Figure 3 - Windchill 3-Tier Platform Example**

Another option to consider is clustering, using two or more machines at the application layer along with a load balancing mechanism. Windchill has some built-in load balancing capabilities in the Method Server. To accomplish load balancing at the web-server level, the first "port of entry" for use requests, an additional IP load balancing switch must be incorporated into the system for balancing and distributing the incoming traffic. This may be in the form of a hardware appliance such as the Cisco Catalyst and Foundry Networks ServerIron switches or software such as the Apache reverse proxy capability in mod_rewrite. At this time, we do not have sizing numbers available specific to this architecture but it is expected that the 100 users/CPU recommendation will still hold if not improve. See Figure 4 below for an example.

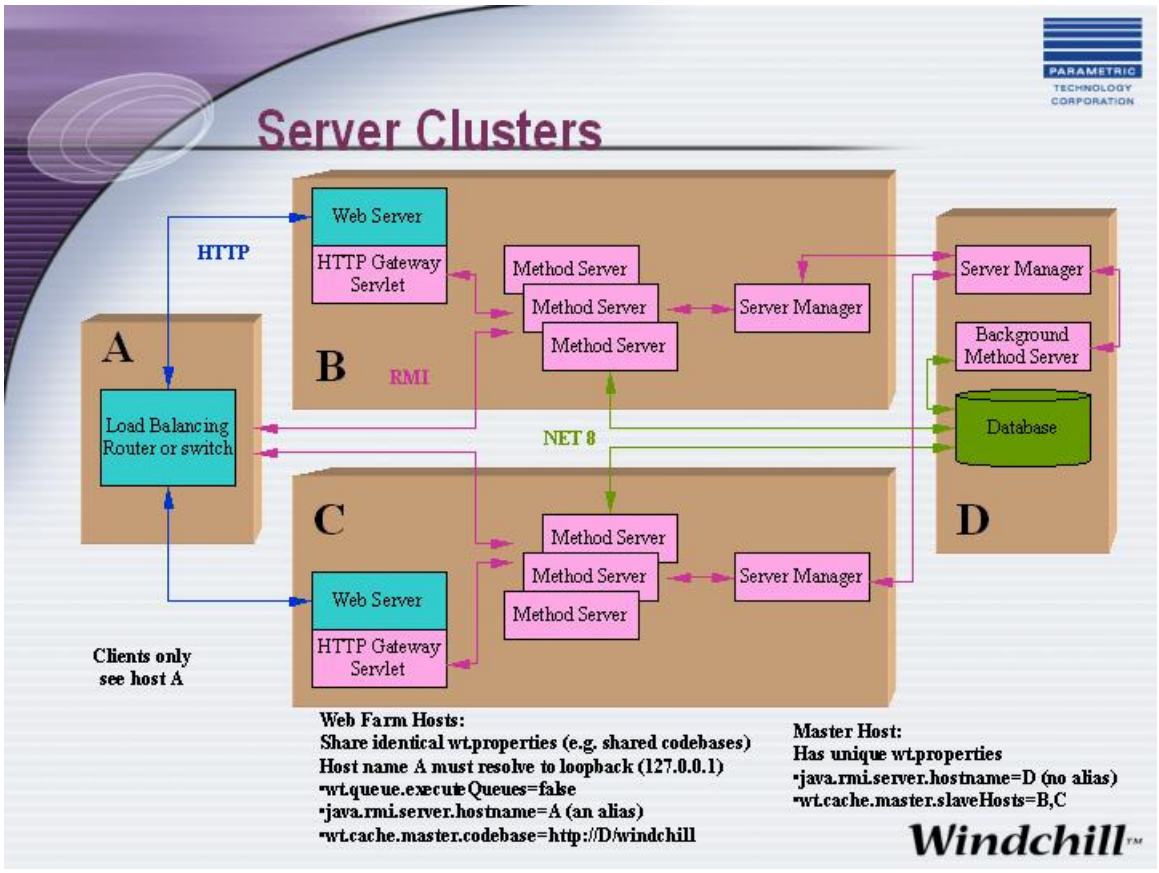


**Figure 4 - Windchill Server Clusters**

## Choosing Disks

HP can provide technical consulting to the customer to assist in choosing disk configurations appropriate to their architectures. What follows is a quick overview of some of the options to consider. Also note that the key measurements used to determine performance of disks are IOPs/second and throughput in MB/second.

| Requirement | MSA1000 | EVA4000 | XP10000 |
|---|---|---|---|
| Built-in Disk array management and security | | X | X |
| Fiber Channel connect | | X | X |
| Rack-Mount | X | | |
| Mixed environment of Windows and HP-UX | X | X | X |
| Performance but limited security | X | | |
| JBOD and Mirror/UX | X | | |
| Peak Performance Random(Mb/s) | 200 | 360 | 2100 |
| Peak Performance Random (1000 IO/s) | 30K | 200K | 500K |
| Max Capacity (Tb) | 24 | 72 | 332 |

**Table 3 - Disk Configurations**

## Network Bandwidth

In a 3-tier architecture, it is recommended that the Oracle and Windchill servers be connected into the same 1000Mbit switch for better reliability and throughput. It is also recommended to use multiple aggregated network interfaces to reduce contention for installations of more than 800 users.

The theoretical limits of network technologies are listed in the table below.

| Network Speed | Network Data Rate | Packet Rate for Small Packets in pps | Packet Rate for Large Packets in pps | Theoretical KB rate assuming 1500 byte packets | Saturation Limit in Kb/sec | Saturation Limit in packets/sec |
|---|---|---|---|---|---|---|
| ISDN 64k | 64K | 163 | 7 | 11K | 5K | 2 |
| ISDN 128k | 128K | 326 | 14 | 21K | 10K | 4 |
| T1 | 1.5M | 3255 | 136 | 200K | 100K | 42 |
| 10Base-T | 10M | 19530 | 820 | 1.2M | 600K | 250 |
| 100Base-T | 100M | 195300 | 8200 | 12M | 6M | 2500 |
| 1000Base-T | 1000M | 1953000 | 82300 | 122M | 60M | 25000 |

**Table 4 - Standard Network Transfer Rates**

The benchmarking in March 2006 demonstrated approximately 800 bytes/sec (with mod_deflate) or 3kb/sec (without mod_deflate) and thus in most cases a 64K ISDN line would be sufficient. However, for CAD data integration or document check-in for documents over 1Mb, the bandwidth is on average twice this amount and increases with model or document size. Content replication should be used where possible but no measurements of this have been performed to date. Note also that RMI should be avoided for any clients on lines under 10Base-T where possible.

# *Configuring for High Availability*

## Mirror/UX

Mirror/UX offers some rudimentary LVM RAID1 mirroring capability that is interesting if you do not have an AutoRAID or other disk array, like EVA4000 or XP10000. It is a very fast mechanism and is very reliable. During the scalability testing, a 17 Gb mirror was demonstrated to be synchronized with an offline copy in about 40 seconds.

## MC/ServiceGuard

MC/ServiceGuard is an optional software product that enables the creation of multi-node server clusters designed to ensure application availability. MC/ServiceGuard uses the concept of an application 'package' to define the service(s) running on a node within the cluster. In the event of a hardware or application failure, the package(s) will be migrated to a surviving node. Package definition files and associated scripts have been created for Windchill, Oracle, and Netscape Enterprise Server. These scripts are available from the www.hp.com/go/pdm site in the Tech Support area. Additionally, a Windchill & MC/ServiceGuard configuration white paper may be found on the same web site.

MC/ServiceGuard does *not* require dedicated, 'hot-standby' hardware to achieve High Availability. All nodes in a cluster may be active and running separate packages. When a failure is detected, the cluster's workload may be redistributed across the remaining nodes.

Because of the cluster IP addressing involved in the MC/ServiceGuard configuration, there are several things to take into account in the configuration - please see HP or PTC for assistance.

# Planning a Backup Strategy

## Omniback/Data Protector

Omniback/Data Protector can be implemented easily to perform offline backups. Its interface is self-explanatory and the only thing to remember is to shutdown the database and Windchill before starting the backup. If you are using a file vault, do not forget to backup the database and the vault at the same time. The one drawback is that if you have a failure, your only protection is this complete backup. This means that you will lose everything up to the last backup and will have to restore the entire backup image to get back online.
HP offers a full line of backup solutions and tape libraries. See http://www.hp.com/go/storage for details.

## Oracle Archiving and Hot Backup

If you are exclusively using BLOBS for data storage and if you enable archive logs in Oracle, you can perform a backup online without impacting the connected users. Oracle will continue to write to the redo logs while Omniback backs up the database files. Once the backup is complete, Oracle rolls forward the transactions stored in the redo logs and will continue working normally. If a content vault is used, it needs to be integrated into the overall strategy. One approach to vault backup is to mirror the vault to begin with and then prior to the backup the mirror is split. One copy continues online getting new transaction changes, while the split copy is backed-up with the database. The old copy is then synchronized and brought up-to-date with the new copy.

# Partitions on hp High-end Servers

## About Partitions

*Partitioning* on HP servers refers to the ability to run multiple instances of the HP-UX operating system on a single server. Partitioning is accomplished by allocating different subsets of the server's hardware resources to different *partitions*, each of which can run a single instance of HP-UX, in isolation from other partitions.

There are two main types of partitions that can be created on hp servers: hardware partitions (*nPars*) – electrically isolated machines running separate copies of HP-UX, and virtual partitions (*vPars*) – software-isolated running copies of HP-UX. Depending on the server model, one or both types of partitions may be supported. The current version of hp Partition Manager can create, delete, and administer only *hardware* partitions. For this reason, the unqualified term *partition* refers specifically to hardware partitions wherever it is used in the Partition Manager program and this help system. The terms *virtual partition* and *vPar* are used when referring to virtual partitions. The hardware nPars must be configured along cell boundaries, where a cell is a unit of, typically, 4 processors that share specific memory interconnects and are the building blocks of most of the mid to high-end servers HPUX based servers. vPars do not share this limitation but do not perform as well as nPars.

For additional sources of information about partitioning of hp servers, hp Partition Manager documentation which can be found here: http://info.fc.hp.com/projects/parmgr/11.11/help/C/overview.html.

There are several other types of partitions based in software. The first, vPars, was discussed above. Another is through processor sets or PSETs where processes can be dedicated to specific processors or groups of processors. Starting with HP-UX 11.23 on Itanium architecture, virtual machines are also now available. For more information, see the URL http://h71028.www7.hp.com/enterprise/cache/258348-0-0-0-121.html.

The preferred method of running multiple copies of operating systems (HP-UX, Windows or LINIX) on the hp Integrity platform is via Integrity Virtual Machines. For more information, please see the following URL: http://h71028.www7.hp.com/enterprise/cache/262803-0-0-0-121.html

The last partitioning possibility is using Instant Capacity or iCAP. This functionality allows customers to pay only for the processing power they require. Please see the following URL for further information: http://h71028.www7.hp.com/enterprise/cache/308083-0-0-225-121.html.

The following table summarizes the hp server solutions and their support of the various partitioning techniques mentioned above.

| | Processor Type | Max CPU Cores | Max RAM | nPar | vPar | PRM/PSET | HP Virtual Machine |
|---|---|---|---|---|---|---|---|
| **Superdome Integrity** | Itanium | 64 | 1 TB | × | × | × | × |
| **Superdome 9000** | PA-RISC 8900 | 128 | 1 TB | × | × | × | |
| **rx86x0** | Itanium | 16 | 128 GB | × | × | × | × |
| **rp84x0** | PA-RISC 8900 | 32 | 256 GB | × | × | × | |
| **rx76x0** | Itanium | 8 | 64 GB | × | × | × | × |
| **rp74x0** | PA-RISC 8900 | 16 | 128 GB | × | × | × | |
| **rx4640** | Itanium | 4 | 128 GB | | | × | × |
| **rp5470** | PA-RISC 8700 | 4 | 16 GB | | × | × | |
| **rp4440** | PA-RISC 8900 | 8 | 128 GB | | | × | |
| **rx2630** | Itanium | 2 | 32 GB | | | × | × |
| **rp3440** | PA-RISC 8900 | 4 | 32 GB | | | × | |
| **rx1620** | Itanium | 2 | 16 GB | | | × | × |
| **rp3410** | PA-RISC 8900 | 2 | 6 GB | | | × | |

**Table 5 - Availability of Virtualization on HP-UX Servers**

## Partitioning and Multiple Windchill instances

A complete Windchill environment may require separate Windchill instances for various distinct projects as well as Production, Pre-Production, Training, Integration and Development usage environment. To consolidate cost and system management, customers can use the partitioning described above to put multiple Windchill instances on the same machine divided by vPars, nPars or Virtual Machines. The memory and CPU requirements below are

minimums for each individual Windchill instance regardless of the targeted usage environment (Production, Pre-Production, Training, etc.).

One has the choice of splitting the implementation into partitions for each project or each use case. For example, if we have Project A and B for which we want Production, Pre-Production and Training/Integration/Development environments, we can use an rp8400 divided into 4 partitions (nPars or vPars)[2] of 4-CPUs each with Project A and Project B Production environments on the first partition, Pre-Production environments on the second partition and Training/Integration/Development environments on the third partition, and the Oracle database on the fourth partition. In this case, we probably want 8 GB RAM on each partition - enough to account for both Windchill projects depending on their final user counts.

Please see your local HP Services Consulting and Integration (HPS C&I) consultant for more information and sizing assistance.

Also note that it is recommended that a system contains at least 4 GB per CPU. The guidelines below are specifically for calculating how much memory/CPU is required for multiple Windchill instances assuming this minimum.

| | 3-Tier Windchill Partition | | 3-Tier Oracle Partition | |
|---|---|---|---|---|
| **User Load** | **Number of CPUs** | **Memory in GB** | **Number of CPUs** | **Memory in GB** |
| **1-100** | **2** | **8** | **1** | **4** |
| **100-500** | **6** | **24** | **2** | **8** |
| **500-1000** | **10** | **40** | **4** | **16** |
| **per 100 CAD** | **4** | **16** | **1** | **1** |

**Table 6 - Partitioning and Windchill Sizing**

# Sizing the Windchill Server

## *Overview*

The Windchill server is the middle tier in a three-tier architecture that provides authorized users access to product related information simply by using a Web browser from any secure location. It is the Windchill server that extends the access to the product data management system to the enterprise. This extension enhances communication and enables collaboration on product development related issues throughout the development cycle.

The Windchill server consists of a third-party web server such as Apache, a servlet container such as Tomcat and the Windchill method server processes. The web server handles the HTTP or HTTPS requests from web browser clients and the Windchill servlets, jsps and method servers provide the business and application logic that supports the requested transactions.

[2] Generally if using partitions that stay on hardware cell boundaries it is better to use npars.

The Windchill server can be implemented in either a monolithic configuration, in which all the server processes and Oracle execute on the same hardware platform, or in a three-tier configuration in which Oracle is executing on a separate system as discussed in the Windchill solution section. As a general rule smaller applications such as RetrievalWare or Info*Engine should reside on the Oracle Server in a three-tiered implementation or on the Windchill server in a monolithic or single server implementation. Our benchmark results indicate that the difference between these two alternative implementations is in the memory utilization. Therefore the sizing information presented here will apply to both.

General sizing information regarding CPU and memory for 3-tier implementations is provided below as well as in the example configurations in appendix B.

There are other requirements in addition to the hardware configuration in realizing a high performance Windchill server. There is the need to monitor system resource utilization using performance tools such as GlancePlus and OpenView Performance Agent as discussed in appendix A. There is also necessary Java Virtual Machine (JVM) configuration and tuning, a subject that is beyond the scope and intent of this document. It is recommended that you contact your local Hewlett-Packard Consulting organization to acquire the technical expertise necessary to properly tune the Windchill server.

A rule of thumb for sizing CPU for the Windchill server:

- **Number of Method Servers:** 1 Method Server for the first 400 users, after that 2 per 400 additional users.
- **3-Tier platform: CPU cores (PA-8900):** For up to 1000 active users, roughly 100 non-CAD users per cpu or core rounded up to the nearest integer with a minimum of 2 CPUs. Each additional 100 CAD users require 1 more CPU core.. Also on cell based machines it is recommended to use locality domains and PSETs one per cell to maximize CPU to memory performance.
- **3-Tier platform: Memory rough rule of thumb:** 4GB per CPU(core).

| System | User Load | 3-Tier CPU(core) | 3-Tier Memory GB |
|---|---|---|---|
| **rp3440 or rx1620** | 100 | 2 | 8 |
| **rp4440 or rx4640** | 500 | 6 | 24 |
| **rp7420 or rx7640** | 1000 | 10 | 40 |
| **rp8420 or rx8640** | 2000 | 20 | 80 |

**Table 7 Windchill Server CPU and Memory Requirements**

## *Mass Storage for the Windchill Server*

Mass storage for the Windchill server will consist of storage for the the local vault. For more information regarding the mass storage for the Oracle database please read the section titled “Disk IO Performance” on page 24. The storage requirements for the local vault are dependent on the amount data required to be available on-line. More detail on disk configuration can be found in the “Windchill Architecture and Planning Guide” available internally from PTC or HP. The ratio of disk reads to writes is approximately 40%.

## *Network Performance for the Windchill Server*

Some guidelines on network performance:

- It is advisable to have a 1000 BaseT (1 GigE) connection between the Oracle DB and the Windchill application server.
- For configurations of greater than 800 active users, multiple aggregated 1000Base-T network cards should be considered.
- Average bandwidth consumption per active user was ~3 kbps without mod_deflate and 700 bps with mod_deflate.
- WAN bandwidth can be reduced by using content cache/replication servers since much of the bandwidth is consumed by uploading/downloading of file content. We have not investigated this yet, but commercial WAN accelerators can be used such as the HP StorageWorks Enterprise File Services WAN Accelerator.
- Mod_deflate makes an impressive difference both on the amount of data transferred and the response time of the application. It is highly suggested for implemation.
- For cost and performance competitive network switches, we recommend the hp ProCurve Networking Switches such as the 5300XL series. A low-cost alternative is the hp ProCurve 2724 unmanaged switch.

## *Sizing the Windchill Server for MC/ServiceGuard*

When considering a high availability implementation with MC/ServiceGuard in a three tier configuration we recommend that both servers, Windchill and Oracle, be sized using the monolithic Windchill server configuration numbers. This will provide the Windchill server and the Oracle server with the additional resources necessary to assume the increased workload of Windchill or Oracle in the case of a fail-over event and minimize any performance degradation that would be experienced by the user community.

For more details on MC/ServiceGuard, please point your browser to http://www.hp.com/go/ha

# Sizing the Windchill Oracle RDBMS Server

## *Overview*

The Oracle Database server for Windchill can be a major weakness of a Windchill installation. It is important that careful consideration be given to such issues as disk and IO performance, network performance, and high availability as well as CPU and memory. General guidelines or hints are given within this document to provide a starting point for a new system implementation, but it does not provide fine tuning information. Database and system tuning is an on-going endeavor and will normally require the insights of a trained database administrator using Oracle Enterprise Manager as well as system performance tuning tools such as GlancePlus with the OpenView Agent and PerfView Analyzer. These performance tools provide a way of measuring Oracle performance and how much of system resources, such as memory and CPU, are being used. The OpenView Agent continuously collects numerous system metrics and stores them in log files for later viewing with PerfView Analyzer. This provides a way of reviewing the history of system performance and provides an opportunity to recognize trends in resource consumption. Rather than provide a historical record of performance, the GlancePlus tool gives a detailed look at how system resources are currently being used by various system processes or applications. The detailed information provided by these performance tools allow you to recognize and overcome performance inhibitors as well as to help you ensure that adequate system resources are available as the load on the system increases over time.

## Disk I/O Performance

For optimal disk I/O performance, it is important to spread the disk I/O throughput across as many separate I/O channels as possible. There should be as a minimum five disk spindles or mechanisms to accommodate the parallel I/O of the database system. For each database transaction there are corresponding accesses to logfiles - for system integrity; tables; rollback segments, and indices. Any or all of these may be accessed in the course of each database transaction. If there is only one disk, or spindle, there is a potential bottleneck in which much of the read-write speed of the disk channel is consumed by the additional disk accesses. To maximize the disk throughput and therefore complete as many database transactions as possible in a given period of time, it is necessary to spread out the I/O across multiple spindles and multiple interfaces as much as possible. Additionally, it is also important not to exceed the recommended maximum number of disk drives per SCSI or Fibre Channel interface.

Too many disks on one interface can result in attempting to read or write more data than the interface can handle, thereby exceeding the maximum data throughput capability of the interface. This usually results in some of the disks being heavily used while others are starved for data. The number of disks on an interface is dependent on the model and type of disk as well as the type of interface. For a “Fast and Wide” SCSI interface which has fifteen addresses, the number of disk drives should not exceed ten to maintain good performance. For example the AutoRAID, model 12 disk arrays, should be limited to 4 per "Fast and Wide" SCSI interface so not to exceed the maximum data throughput of the interface. Note that this recommendation would be different if you were just adding disk space with low usage where data access was infrequent. In this case, you could feel comfortable about using all the available bus addresses since data throughput and bandwidth would probably not be issues.

- The Database Server has heavy I/O rate of >3000 IOPS for 600 users and more. This requires a minimum of 2Gb Fiber Channel storage cards preferably using multiple aggregated cards.
- The ratio of disk reads to writes has been measured between 60 and 80% and up to 90% for high concurrency (>800 users).

## Data Protection and High Availability

To provide both data protection and highly available data, Hewlett-Packard offers a number of different solutions. These solutions are in the form of various product offerings, which provide varying degrees of hardware redundancy for specific failure points such as media failure or interface failure or disk controller failure. In addition, hot swapable disk drives allow for the replacement of media without having to interrupt system service allowing for highly available data to exist on a system.

It is necessary to investigate and choose the level of data availability and data protection for your particular implementation. Whether you choose to use mirrored independent disks, high availability RAID 5 disk arrays, an MC/ServiceGuard configuration should be considered. An MC/ServiceGuard implementation will provide for an alternate system to automatically take over in the event of a system failure thereby providing a very high level of system availability. The type of protection and availability you choose will depend on your need for data availability and whether the data have to be available 24 hours a day 7 days a week, or 8 hours a day, allowing for system maintenance and backup during the off periods.

## Sizing the Oracle Database Server for MC/ServiceGuard

When considering a high availability implementation with MC/ServiceGuard in a three tier configuration we recommend that both servers, Windchill and Oracle, be sized using the monolithic Windchill server configuration numbers. This will provide the Oracle server with the additional resources necessary to assume the increased workload of Windchill in the case of a fail-over event and minimize any performance degradation that would be experienced by the user community. For more information about MC/ServiceGuard, please see http://www.hp.com/go/ha.

## Network Performance

Network performance of the database server is another area of consideration. The database server is the top tier of any Windchill implementation and is therefore critical to the performance of the overall system implementation. As the number of network clients increases the network traffic also increases and a 100baseT LAN segment can quickly become a bottleneck to system performance. It is important therefore to maximize the network bandwidth to the server by using as fast a network interface as available. A 1000BaseT interface is recommended to maximize throughput into and out of the database server. Whichever interface is chosen it would be wise to monitor the network loading to make sure that the utilization is below the bandwidth limit of the interface otherwise a network bottleneck condition exists and database performance will suffer.

For cost and performance competitive network switches, we recommend the hp ProCurve Networking Switches such as the 5300XL series. A low-cost alternative is the hp ProCurve 2724 unmanaged switch.

## CPU

The CPU and memory requirements for the database server are highly dependent on concurrency and usage as well as local policies and practices, which differ from one installation to another. Additional memory can be added to allow data tables to remain resident in RAM, reducing the disk I/O and dramatically improving the database performance. Our recent benchmarks using the PA-8900-based and Itanium-based machines reveal the high performance of the new processor technology in these systems. Each processor in these systems will support several hundred clients. Wildfire users generate additional load so there is a separate sizing matrix below to account for them. Because of variations in workloads from site to site and wanting to leave additional CPU resource for future growth it is recommended that you use the following table or review the configurations found in Appendix C.

### Database Server CPU Configuration

### Non-CAD Users

| System Type | User Load | CPUs(cores) | Memory (GB) |
|---|---|---|---|
| **rp3410 or rx1620** | 100 | **1** | **4** |
| **rp3440 or rx1620** | 500 | **2** | **8** |
| **rp3440 or rx4640** | 1000 | **4** | **16** |
| **rp4440 or rx4640** | 2000 | **8** | **24** |

**Table 8 Database server CPU and memory requirements: non-CAD users.**

### CAD Users

| System Type | User Load | CPUs(cores) | Memory (GB) |
|---|---|---|---|
| **rp3410 or rx1620** | 25 | **1** | **4** |
| **rp3410 or rx1620** | 50 | **2** | **4** |
| **rp3440 or rx1620** | 100 | **2** | **8** |
| **rp4440 or rx1620** | 200 | **8** | **8** |
| **rp8420 or rx8640** | 500 | **20** | **12** |

**Table 9 Database server CPU and memory requirements: CAD users.**

These table values are based on our benchmark results and take into account periods of peak usage. For example, at the beginning of a work period many people may be logging into the Windchill system over a short span of time.

These CPU values will provide great performance and room for future growth. It is recommended that you use a method of monitoring the system resources, CPU and memory, with a tool such as the OpenView Agent that is a part of the GlancePlus Pak 2000 performance product.

### *Memory*

Too little memory in the database server can dramatically reduce the performance and cause an otherwise high performance server to have lackluster performance. The database server needs memory for the operating system, operating system buffer cache for disk I/O, a connection process database connection, and a System Global Area, SGA, to hold and sort frequently accessed data. If data being requested is not resident in memory then disk accesses are required to bring in the data. The results from SQL queries must also be sorted and if there is not enough memory to do this, then sorts are done by repeated multiple disk accesses. Both the sorting and retrieving of information by making unnecessary disk accesses are the root cause of poor database performance.

Ensuring good performance requires that there be enough memory to satisfy as high a percentage as possible of the SQL queries and have room for sorting without having to make unnecessary disk accesses. Some systems are able to satisfy 95% of the SQL queries from the buffer cache in the System Global Area, SGA, in memory without having to do disk accesses. This would be equivalent of having a buffer cache[3] hit ratio of 95%. One of the primary goals of database and system tuning is to adjust the database and system parameters to achieve the highest buffer cache hit ratio possible.

Based on our benchmark results you should install the amount of memory per the following table or review the configurations in Appendix B.

In addition to the several sample configurations in Appendix B we recommend that you contact your Hewlett-Packard sales representative to discuss your particular system requirements.

# Additional Server Types

There are several types of servers that are frequently found in PDMLink hardware architectures: The CAD Agent Server, the Content Cache Server, Development systems and Test servers.

### *CAD Adapter*

For collaborative engineering, sharing of graphic visual images, or creating thumbnails via a CAD Agent by using either PTC's ProductView or the Windchill Graphics Server, we recommend the Hewlett-Packard xw6200 or xw8200 Visualize workstation with the ATI Fire GL graphics adapter, 2GB of RAM and 18GB of disk or a c8000 HP-UX workstation with FireGL graphics, 8GB RAM and 18GB disk.

### *Content Cache (Replication) Server*

In today's design environment, designer engineers and managers, even within a single organization, may be located throughout the world with access to greatly varying network access bandwidth. The trend for dealing with the response time issues that arise from this situation is to centralize the storage of the light-weight metadata while distributing the storage of the larger content files with replication, also called content cache, servers. The number of these servers involved will depend on the distribution of workers, available network bandwidth, degree of file sharing required and content replication strategy that is chosen. In typical situations we recommend the hp 9000

[3] Note that we are refering to Oracle database buffer cache here and not OS system buffer cache.

rp3410 with 2 cpus; 8 GB RAM; 2-18 GB internal disks and external storage such as the MSA30 or EVA4000 – sized to meet the needs of the local site.

Reference architectures for both the CAD Adapter and the Content Cache Server are given in Appendix B.

### *Development Systems*

Key challenges after a successful implementation of Windchill are to maintain high levels of availability and performance as systems are moved into production. A great way to help meet these challenges is to establish the necessary test and development environments away from the production or deployment servers. This test and development environment will provide the resources for a number of activities that are conducted in support of a production Windchill implementation.

One activity is to be able to develop new data models and custom applications to solve implementation issues prior to deployment in a production environment. Another activity is to test models and tools against a copy of a production database. In a collaborative engineering environment there is the need to share information such as text and graphic visual images. There are a number of different server types that can be used to support these different activities.

A development workstation is used to develop new data models and custom applications prior to deployment in a production environment. For this purpose a small Hewlett-Packard xw6200 or an xw8200 Visualize workstation with 1 GB of RAM and an internal 18 GB disk drive will suffice. This will provide you with the hardware resources you need to carry out this custom development work.

### *Testing Server*

A more robust solution for testing against an Oracle test database while also doing development work on new data models would be either an Intel based Hewlett-Packard Proliant dl380 running Microsoft's Windows XP or Windows 2003 server or a choice of HP-UX based server systems such as an rp3410 or an rp4440 server. Whichever solution you choose will require memory and disk storage, 4 GB of memory and an 18GB disk will be sufficient.

A training environment would require just 4-CPUs so an rp4440 or one 4-CPU Superdome partition would be sufficient.

## Performance Tools

Variations in system loading can produce unexpected results in system performance. System resources such as CPU, memory and the way memory is partitioned by the operating system are fixed and finite. This doesn't mean that they can't be adjusted or reconfigured but typically this requires the kernel parameters to be adjusted and the system rebooted with a new kernel. While system resources are fixed on a running system the environment, which consists of the user community and a demand load, is quite variable. An increasing number of users, increasing demand on file system storage, increasing network traffic, as well as variations in usage profiles contribute to an ever-increasing demand on computer system resources. As demand increases, system response can actually decrease as a result of reaching the limit of these system resources resulting in a dramatic degradation of system performance.

Performance barriers, often referred to as bottlenecks, can be due to limitations in available memory, network performance, available CPU, or limitations in I/O data throughput. To identify and deal with system performance problems and to be able to predict system performance requires the use of special tools to monitor and capture system performance information.

Performance tools are a collection of software programs used to capture system performance information while the system is under load. The resulting performance information can be used to characterize system performance under a variety of conditions. When a computer system is first implemented, baseline performance information can be obtained to determine how much of available system resources are being consumed during the course of normal

system operation. Performance data can also be captured as system load increases over time, providing a means of predicting future performance with changes in system load.

Most of the performance tools are used in our testing efforts and are described in Appendix B. There are additional tools for automated monitoring of multiple systems and for systems management through a single workstation console. All of these tools are a part of the OpenView suite of system and network management tools. For more detailed information about the OpenView suite of tools beyond what is written here, please contact your local Hewlett-Packard sales representative or visit the Hewlett-Packard web site at http://www.hp.com/.

# Where to get more information

HP Global Home Page: http://www.hp.com

PTC Global Home Page: http://www.ptc.com

HP Servers: http://www.hp.com/go/servers

HP Storage: http://www.hp.com/go/storage

HP Workstations: http://www.hp.com/go/workstations

HP High Availability Solutions: http://www.hp.com/go/ha

PLM on HP: http://www.hp.com/go/plm

# APPENDIX A: Sample Configurations

## *Sample Windchill Server Configurations for 3-Tier Implementations*

**Note:** Mass storage for the local file vault is not included in the following configurations. The required amount and type of storage should be determined and added to your configuration as appropriate.

The following sample configurations each consist of the necessary memory and processors to meet the user workload as indicated. They also specify racking, internal mass storage devices, DVD or CDROM for loading of software, and a digital audiotape device for administrative use. These configurations should be viewed as examples and should be modified as necessary to meet your specific requirements. For example you may wish to modify the racking and add the appropriate number and type of disk interfaces and external mass storage devices to meet your data storage requirements.

### 100 User Windchill 3-Tier, Using rp3440

HP rp3440 2 cores, PA-8900 800MHz
8 GB RAM
2-36GB Internal disks
DVD-CDROM drive
MirrorDisk UX and GlancePlus Pak 2000

### 100 User Windchill 3-Tier Using rx1620

HP rx1620 2-CPU, Itanium® 2 1.3GHz
8 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 500 User Windchill 3-Tier Using rp3440

HP rp3440 6 cores, PA-8900 1.0 MHz
24 GB RAM
2-36GB Internal disks
IDVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 500 User Windchill 3-Tier Using rx4640

HP rx4640 6-CPU, Itanium® 2 1.5 GHz
24 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 1000 User Windchill 3-Tier Using rp4440

HP rp4440 10 cores, PA-8900 1.0 GHz
40 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 1000 User Windchill 3-Tier Using rx7620

HP rx7620 10-CPU, Itanium® 2 1.6 GHz
40 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Windchill 3-Tier Using rp7420

HP rp7420 20 cores, PA-8900 1.1 GHz
80 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Windchill 3-Tier Using rx8620

HP rx8620 6-CPU, Itanium® 2 1.6 GHz
80 GB RAM
2-36GB Internal disks
DVD-CDROM Drive
MirrorDisk UX and GlancePlus Pak 2000

## *Sample Oracle Database Server Configurations in a 3-Tier Implementation.*

### 100 User Oracle Configuration Using rp3410

HP rp3410 1-core, PA-8900 800Mhz
4GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000with 10-36GB
MirrorDisk UX and GlancePlus Pak 2000

### 100 User Oracle Configuration Using rx1620

HP rx1620 1-CPU, Itanium ® 2 1.3 GHz
4GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB
MirrorDisk UX and GlancePlus Pak 2000

### 500 User Oracle Configuration Using rp3440

HP rp3440 2-core, PA-8900 1.0 GHz
8GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB
MirrorDisk UX and GlancePlus Pak 2000

### 500 User Oracle Configuration Using rx1620HP rx1620 2-CPU, Itanium ® 2 1.6 GHz

8GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB
MirrorDisk UX and GlancePlus Pak 2000

### 1000 User Oracle Configuration Using rp3440

HP rp3440 4-cores, PA-8900 1.0 GHz
16GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

### 1000 User Oracle Configuration Using rx4640

HP rx4640 4-CPU, , Itanium ® 2 1.6 GHz
16GB RAM

2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Oracle Configuration Using rp4440

HP rp4440 8-cores, PA-8900 1.0 GHz
24GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Oracle Configuration Using rx4640

HP rx4640 8-CPU, Itanium ® 2 1.6 GHz
24GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Oracle Configuration Using rp8420

HP rp8420 8-cores, PA-8900 1.1 GHz
24GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

### 2000 User Oracle Configuration Using rx8620

HP rx8620 8-CPU, Itanium ® 2 1.6 GHz
24GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 10-36GB disks
MirrorDisk UX and GlancePlus Pak 2000

## *CAD Agents*

### CAD Agent Server, Using C-Class, c8000 Workstation

HP c8000 2-core, PA-8900 1.1Ghz
2GB RAM
1-36GB Internal disk
DVD-CDROM drive

MirrorDisk UX and GlancePlus Pak 2000

### CAD Agent using xw8200 Visualize Workstation

HP xw8200 2-CPU 3.6GHz
1GB RAM
1-36GB Internal disks
DVD-CDROM Drive
Monitor

## *Content Replication Servers*

### Replication (Content Cache) Server using the rp3410

HP rp3410 2-CPU, PA-8900 800Mhz
2 GB RAM
2-36GB Internal disks
DVD-CDROM drive
High Availability Storage Unit, MSA30 or EVA4000 with 4-15GB disks
MirrorDisk UX and GlancePlus Pak 2000

# APPENDIX B: Performance Tools

### Overview

Whether the information systems environment includes a single system or a large data center, the fundamental goal is the same: To maintain satisfactory service for users and applications while minimizing costs and optimizing use of computing resources. The amount invested in time and tools to achieve this goal depends on factors in the environment: the business importance of applications, management's requirements, and the capability of the IS staff.

The most common requirement in performance management is to respond to and correct problems when they affect user productivity and application throughput. Solving these problems may require changing priorities for running processes or programs or rescheduling non-critical workloads. To resolve these problems, you must be able to monitor current workloads, resource utilization levels, and user activity levels to have the information you need to take corrective action. Since you cannot always see problems when they occur, many times it is helpful to see historical workload information for identifying a problem after the fact.

Beyond resolving current problems, the next most common performance management challenge is understanding and reporting on the nature of system workloads, user service levels, and the consumption of system resources such as CPU, memory, disk and network bandwidth., by your applications. Also important is seeing where workloads or data and files can be re-balanced, or how system configurations can be modified to improve system responsiveness and throughput. To effectively perform these tasks, you need easy access to useful historical workload information.

Beyond problem solving and workload reporting, data center capacity planners often have to forecast future application service levels and resource use. Historical performance data combined with statistical analysis and modeling tools give you the information you need to predict the effects of implementing new applications or adding new users and resources.

Managers of larger networks of systems and high-end business systems must identify and resolve potential performance problems before they seriously affect service levels and user productivity. Also, data center managers implementing "lights out" operations need to automate as many management tasks as possible. In both cases, performance alarms are required that can notify a central system operator when resource utilization and service

levels have exceeded certain thresholds, and mechanisms are required that automatically execute corrective action when needed.

The following section describes each of the performance tools.

## GlancePlus

GlancePlus software is a real-time performance monitoring and diagnostic tool. It provides immediate performance information about your system. It lets you examine system activities, identify and resolve performance bottlenecks, and tune your system.

There are numerous system metrics, such as CPU, memory, system table resource information, and LAN resources that can be inspected. This tool is invoked when it is necessary to view system resource information.

## GlancePlus Pak 2000

GlancePlus Pak 2000 is a software bundle that includes both GlancePlus and the OpenView Agent described below. This software bundle should run on every Hewlett-Packard server to provide performance analysis capability and to provide historical performance data to the OpenView Performance Analyzer.

## OpenView Performance Agent (previously known as MeasureWare)

The OpenView Performance Agent runs continuously and collects comprehensive information on system activity to provide historical data for analysis. This information is then summarized, time stamped, and logged. The agent provides data and alarms to PerfView for analysis. The agent will capture global system metrics as well as application and process specific metrics. Global, application and process specific CPU, memory, and network utilization are just a few of the available metrics.

## OpenView SMART Plug-Ins, (SPI's) for Databases

These are available for Oracle, Informix, Sybase, and Microsoft SQL Server.

The SPI database module provides key database metrics to the OpenView Agent. A few of the many metrics are as follows: Active transactions, alerts, archive free space, cache hit ratio, call rate, long table scan ratio.

The metrics are summarized, logged, time stamped and alarmed on database metrics in the OpenView Performance Agent.

The database module automatically integrates with OpenView Performance Analyzer and provides expertise on what actions should be taken when alarm conditions occur. The SPI database modules run on the database server system with the OpenView Performance Agent.

## OpenView Performance manager (Previously known as PerfView)

The OpenView Performance Manager is a graphical user interface for viewing historical data managed by the OpenView Performance Agent. The OpenView Performance Manager graphically displays and reports on system, network, database and application metrics from OpenView Performance Agent. The OpenView Performance Manager is designed to give you a variety of views of the metrics that are important to understanding how resources are being used to support applications running in a distributed environment. The OpenView Performance Manager allows performance data from multiple OpenView Performance Agents to be viewed simultaneously on a single graph. Since each OpenView Performance Agent can monitor data from various data sources such as database management systems and networks, you can use OpenView Performance Manager to analyze key resource usage for applications running either on a single machine or highly distributed throughout the enterprise. The OpenView Performance Manager typically runs on either a HP-UX or PC NT system management workstation console.

## OpenView Performance Monitor (Previously known as PerfView Monitor)

The OpenView Performance Monitor monitors OpenView Agent alarm messages to let you manage the performance of your distributed system from a single point. The OpenView Performance Monitor uses the concept of management-by-exception to dramatically reduce the amount of human and system overhead required for effective performance management. You define specific performance or resource utilization alarm conditions using the OpenView Agent installed on each managed node in the network.

OpenView Agent alarm conditions can consist of thresholds and time duration elements for single or combinations of multiple measurements. For example, a simple alarm condition could be defined to exist when total CPU utilization is greater than 50 percent of capacity for 20 minutes or longer.

When an alarm condition occurs on a managed node, the OpenView Agent on that node sends an alarm to the OpenView Performance Monitor central console, notifying operators of a potential problem that can be corrected before it affects user productivity. If no alarms are present, you can be assured that all system resources are operating within the pre-defined acceptable parameters.

The OpenView Performance Monitor typically runs on either a HP-UX or PC XP system management workstation console.

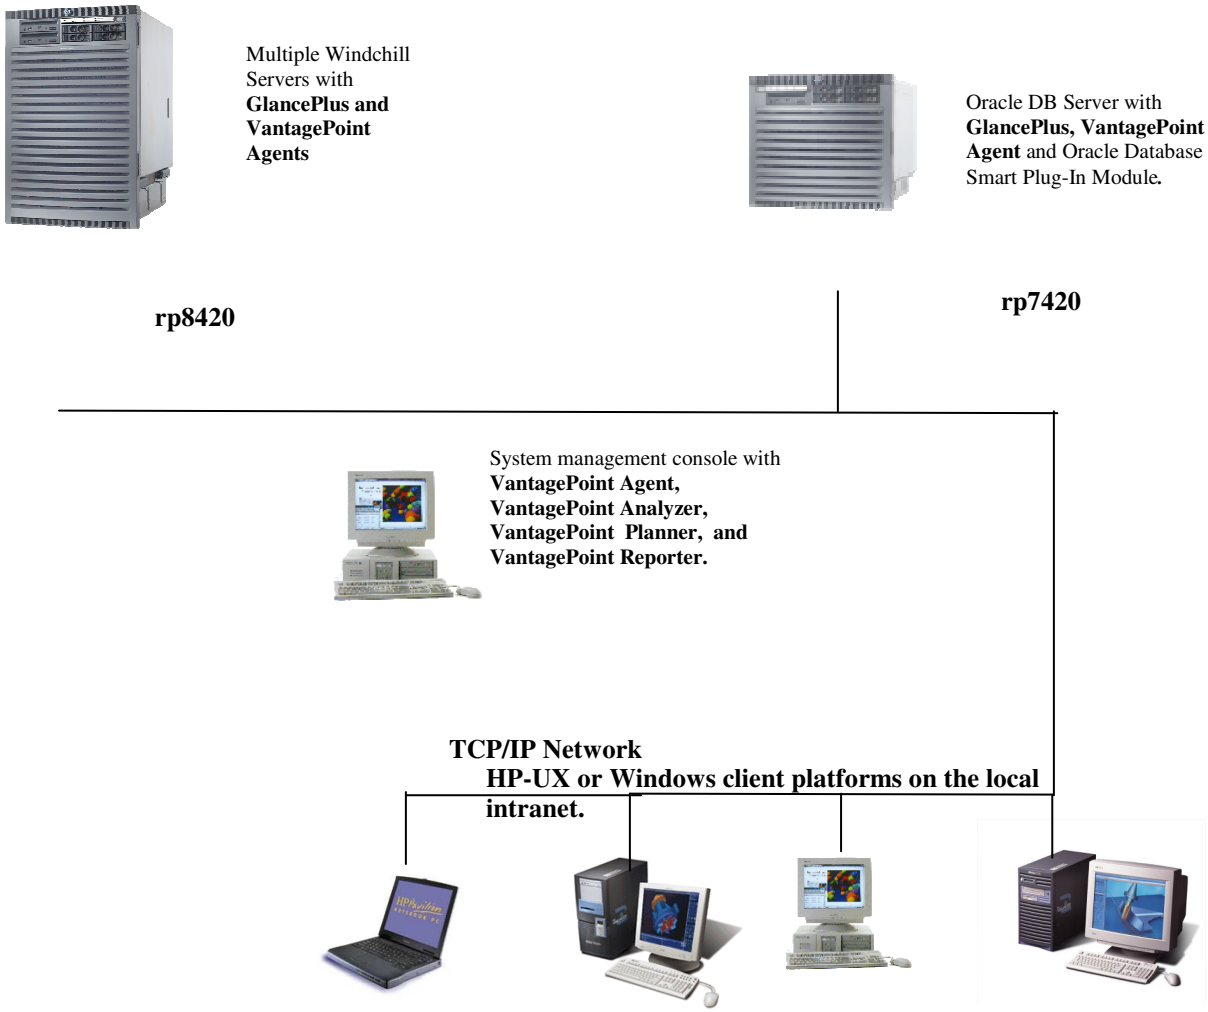


## OpenView Performance Planner (Previously known as PerfView Planner)

The addition of OpenView Performance Planner software enables you to project resource utilization and performance trends. Information can be displayed in a variety of easily understood graphical and tabular formats. This information helps you to predict future resource utilization and performance levels.

OpenView Performance Planner typically runs on either HP-UX or PC XP system management workstation consoles.