PTC Windchill 5.1 on Sun Server Sizing Guide

Technical White Paper

November 2001

Printed in the United States of America.
901 San Antonio Road, Palo Alto, California 94303 U.S.A

Please Recycle

Contents

Introduction 1

In Spring 2001, tests of PTC's Windchill Foundation 5.1 on a Sun Fire™ 4800 server were conducted at the Enterprise Technology Center in Burlington, MA by Sun Microsystem's Market Development Engineering staff. Assistance was provided by Segue support staff, PTC Windchill performance engineers, and Taylor Consulting, Inc.

The results of these tests were used to compile this sizing guide, which was written as an aid for PTC Application Engineers, PTC Global Services Consultants, Sun System Engineers, and other professional service organizations involved in sizing configurations, planning, and installing Windchill application servers.

The History of Windchill Sizing Tests

This was the fourth set of tests, the goal of which was to assess and improve the performance and scalability of Windchill 5.1 on Sun Fire servers as measured by the number of concurrent users. The previous test was performed with a earlier version of Windchill and should not be compared with the results from this test. This test, along with the previous three tests, resulted in feedback to PTC Research and Development, enabling improved performance and scalability of Windchill.

It is important to note that the test scripts used during this test were unique and written by PTC to test core Windchill functions on the Sun platform. These scripts were different from those used during previous PTC tests with other vendors, and were not part of a standardized benchmark. Comparing Sun and HP, or any other vendor's results, would therefore be irrelevant.

1

What is Windchill?

Windchill Foundation 5.1 offers a federated, Web-based environment that focuses on the creation, management, and evolution of product and process information. Its core capabilities include inter-enterprise document management, structure management, life cycle management, work flow management, and system administration.

Windchill Foundation is a member of the Windchill Collaboration Factor! suite of products, which creates an environment where businesses can share product and process information, regardless of format or where it resides. Windchill Foundation links together, or federates, the dispersed systems, applications, and processes of a company and it's supply chains. With this federated, n-tiered, Web-based framework, companies have the flexibility needed for business-to-business collaboration, and can take advantage of emerging Internet technology to increase their competitive advantage.

One of the key aspects of an N-tier architecture is the use of diagonal scaling — a combination of horizontal and vertical scaling, each of which is described below.

Horizontal Configurations

Horizontal scaling enables rapid application change and simplified change management. By using multiple servers, changes can be deployed incrementally, enabling applications or services to be replicated quickly to multiple servers in a controlled manner. Using a second system, or system domain, also increases tuning flexibility. The database system can be tuned for faster Oracle performance and the Windchill server can be tuned for network application performance. Figure 1 illustrates an example of a horizontal Windchill application architecture.

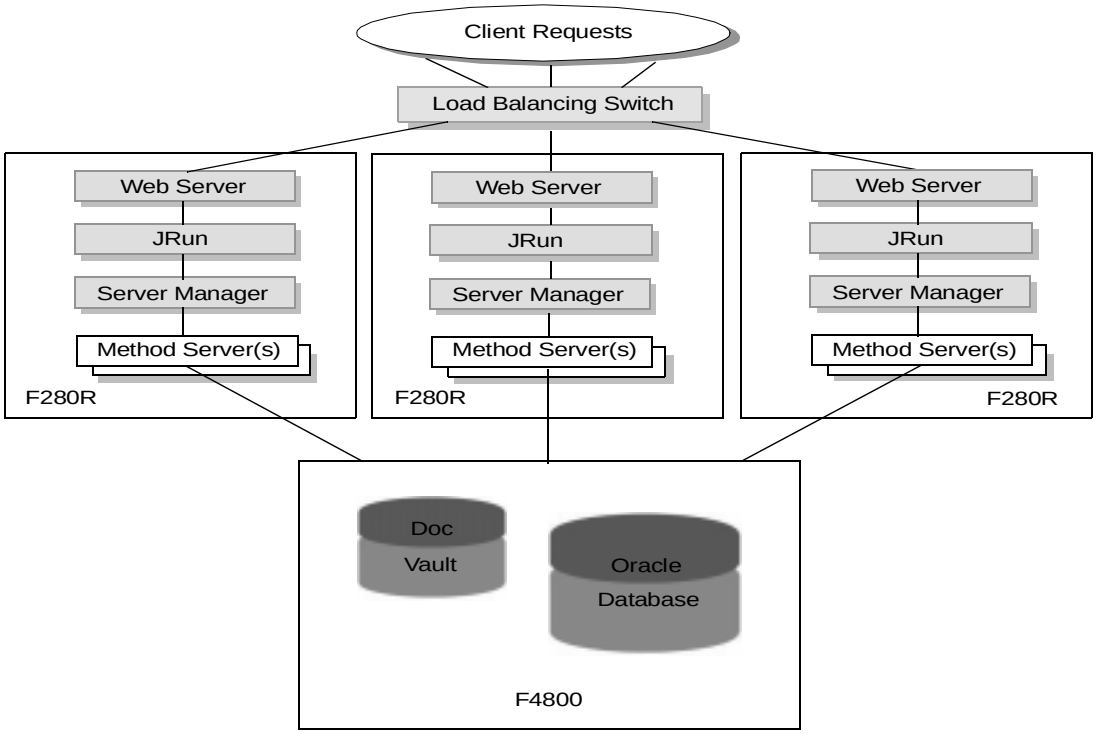

Figure 1 Horizontal Windchill Application Architecture using Sun Fire 280R and Sun Fire 4800 servers

Vertical Configurations

Vertical scaling is characterized by slowly changing applications, and is often necessary for data-intensive services such as databases, video servers, mail stores, transaction processing, search indexes, and directories. Services are scaled within the system — resources such as CPUs, memory, and storage can be incrementally added to the server over time to increase scalability. Figure 2 depicts a vertical Windchill application architecture.

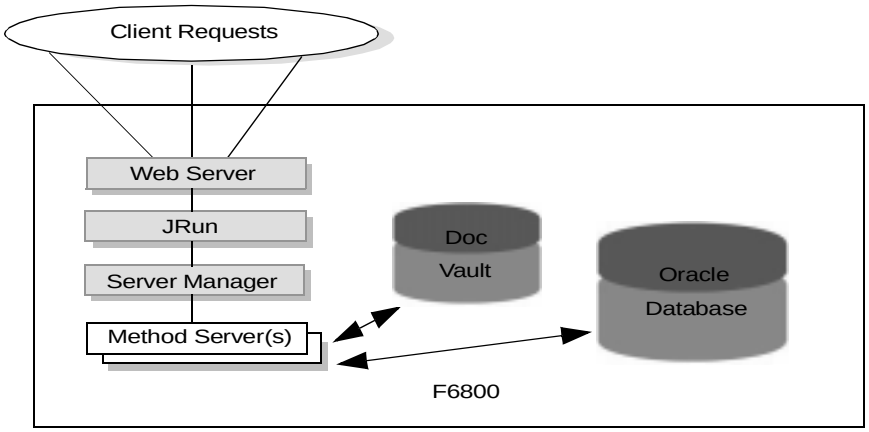

Figure 2 Vertical Windchill Application Architecture using a single Sun Fire 6800 server

By combining horizontal and vertical scaling, companies can deliver more consistent access to applications and services while at the same time improving system performance. This combination can be achieved through multiple servers, or by employing the powerful domain functionality of Sun's Enterprise™ 10K or Sun Fire servers.

Windchill delivers information management through simple, everyday Web mechanisms within standard Web browsers. Windchill Web pages provide access to Java™ applets, search engines, URLs, e-mail, upload/download services, integration with plug-in and helper applications, and the ability to hyperlink and navigate information across supplier and partner systems.

Windchill Foundation uses standard Web technology such as: Web servers, browsers, Java and JavaBeans™, HTML, HTTP, XML, Web security models, and search engines.

For more information see the Sun white paper, *Deploying a Collaborative Product Commerce (CPC) Solution with PTC Windchill.*

Other Software Components

- *Solaris™ 8 Operating Environment*: This series of tests was performed using the Solaris 8 Operating Environment and application software that was certified for it at the time.Previous Windchill tests were executed on earlier versions of the Solaris Operating Environment.
- *Oracle 8.1.7:* The Oracle 8 database stores all information needed for Windchill to manage enterprise product data.
- *iPlanet™ Web Server 4.1:* The iPlanet Web Server Enterprise Edition provides Web and application services.
- *iPlanet™ Directory Server:* In all configurations, LDAP Directory Services should be provided through the customer's existing infrastructure to centralize user management. If Directory Services are not available, sizing for a Directory Server should be performed using iPlanet guidelines appropriately adjusted for the expected number of users. (See *http://docs.iplanet.com*)
- *Macromedia JRUN Servlet Engine 3.0 SP2:* Macromedia Inc. supplies software products designed for rapidly and cost effectively building e-business applications. The JRun server is a powerful, high-performance Java™ 2 Platform, Enterprise Edition 2 (J2EE™) application server for Windchill applications.
- *RetrievalWare Search Engine 6.7.1:* Convera RetrievalWare gives users the ability to quickly extract knowledge from large collections of information — in any form, from any source — using simple natural language queries. It provides an intelligent interface that seamlessly connects users to the full spectrum of enterprise knowledge assets — including scanned paper documents, real time information streams (such as e-mail and news feeds), word processor documents, legacy databases and multimedia libraries.
- *Segue Load Generator/Load Balancer:* Segue's SilkPerformer is a multiple PC based load generation software system. It employs a script driven console and multiple agents to simulate large numbers of users driving a large Web-based server system. Segue is typically used to either profile servers or application software residing on servers. To simulate a managed workload of 1,000 to 1,500 users, 20 PCs were required: 18 with 600 MHz CPUs and 128 MB of memory, and two with 200 MHz CPUs and one gigabyte memory.

  In production environments, a load balancer can be employed to balance incoming client requests across the Web and Windchill servers in horizontal configurations.

1

Definition of Terms Used in this Guide

Several terms particular to Windchill are used in this guide and are defined below:

- *Named Users:* Every user that is defined in the Windchill User Database, and has the *ability* to login to a Windchill system, is considered a *Named User*, regardless of whether they are logged in and active (concurrent), or not. There are generally three to ten times more named users than concurrent users, depending on customer use.
- *Concurrent Users*: Users that are *currently* logged into the Windchill system and have the ability to execute any transaction that is available to them as a user. Concurrent users can be in either think mode or executing a transaction.
- *Concurrent Transactions:* A small subset of the Concurrent Users that are actually *executing* a transaction at the time of the sampling. This number helps understand exactly how busy the system is at any given moment.
- *Enterprise Search and Local Search:* An Enterprise search is any search that uses the RetrievalWare search engine. Local searches are those that only search local metadata *do not* use RetrievalWare.
- *Document/Part Author (Author)*: A type of user in Windchill that is responsible for creating or modifying content in the Windchill database. The main difference between a Document and Part Author is the type of transaction and the content of those transactions that are performed.
- *Consumer*: A type of user in Windchill that does not modify content but instead is a simple reader or consumer of that information. This type of user is a lighter weight user than an Author and does not constitute a large load on the system. It also represents a larger portion of the user community than the Authors do.
- *Vertical Configuration:* A hardware configuration in which all Windchill subsystems are installed on a single server.
- *Servlet Engine:* A Java runtime program, such as JRun or Tomcat, which manages and executes JavaServer™ Pages and Java servlets.
- *BLOB*: An acronym for ***b****inary* ***l****arge* ***ob****ject*, a collection of binary data stored as a single entity in the Oracle database.
- *Vault*: A storage technique whereby the data resides in the UNIX® file system, and only a reference to the data is stored in the Oracle database.

Layout of this Guide

Chapter 2 of this guide presents the CPU, memory, and I/O sizing guidelines which were created using the results from the fourth iteration of Windchill sizing tests performed on Sun Fire servers.

Chapter 3 provides eight reference configurations from pilot and development/test systems, to highly available productions systems.

The details of the test and the test results are contained in *Appendix A*, and should be referenced before using the guide to identify factors in a potential customer installation that might affect the results of the guide. Software settings for the various applications are also supplied in *Appendix A* to help in installation planning.

To give a sense of how customers are currently implementing Windchill systems, *Appendix B* illustrates three real life customer configurations.

Appendix C provides some general sizing guidelines as well as information on where to go for more detailed information.

1

CPU, Memory, and I/O Sizing Guidelines 2

The objectives of this iteration of tests were to measure CPU utilization and performance, based on user loads, in a simulated Windchill Foundation 5.1 environment, and to formulate hardware sizing guidelines. PTC engineers defined a randomly generated workload with think times, composed of 39 different tasks typical in a Windchill Foundation implementation (for more information on the specifics of the tests, see Appendix A). During the tests, Segue SilkPerformer was employed to simulate workloads from 1,000 to 1,500 concurrent users in 100 user increments. All server-side Windchill components, including database queries and a six gigabyte sample production database, were running on one Sun Fire 4800 server.

The Sun and PTC engineers ran a series of test in which they increased the number of users from 1,000 to 1,500 on a 12 CPU Sun Fire 4800 server. Tests were then run with 1,000 users on seven, eight, and nine CPU systems. For each specific Windchill task, the engineers measured response times, overall data throughput, transaction concurrency, and CPU utilization in order to formulate the number of users supported per CPU. The results exhibited the outstanding linear scalability and excellent performance of the Sun Fire servers.

CPU/Memory Sizing Guide

The engineers used the results from the test to create a guide for determining the number of CPUs needed in a Windchill installation, based on the number of concurrent users expected on the system. It is important to note that *concurrent users* are users that are logged in to the system and performing transactions, as

opposed to *named users*, which are simply *all* of the possible users defined in the Windchill User Database. The guide, which can be used for both vertical and horizontal configurations, is contained in Figure 3.

Windchill Application Server CPU/Memory Guide										
Concurrent Users	E220R E250 450Mz or 480MHz	F280R 750MHz	E420R E450 450Mz or 480MHz	E3500 400MHz	E4500 E5500 400MHz	V880 F3800 750MHz	E6500 400MHz	F4800 F4810 750MHz	E10K 400MHz	F6800 750MHz
100	1	1	1	1	1					
200	2	1	2	2	2		2			
300	2	2	2	3	3	2	3	2		2
400		2	3	4	4	2	4	2		2
500		2	4	5	5	2	5	2		2
600				6	6	4	6	4		4
700				7	7	4	7	4		4
800				8	8	4	8	4		4
900					9	4	9	4		4
1,000					10	6	10	6		6
1,250					12	6	12	6		6
1,500						8	14	8		8
1,750						8	16	8	16	8
2,000							18	10	18	10
2,250							20	10	20	10
2,600								12	24	12
3,500									32	16
5,400									48	24

Application Tier: (1) 750 MHz CPU for every 225 Users
Application Tier: (1) 450 MHz CPU for every 140 Users
Application Tier: (1) 400 MHz CPU for every 112 Users
Web server and JRun: Add (1) CPU for up to 2,000 Users
Oracle: 1 CPU for every 4 CPU's in Application Tier
RetrievalWare (less than 1,200 users): Add 1 CPU for every 4 CPUs in Application Tier
RetrievalWare (more than 1,200 users): Add 1 CPU for every CPU in Application Tier
Memory: 1 GB of memory per CPU up to 16 CPUs, and 0.5 GB of memory per CPU over 16

Figure 3 Windchill Application Server CPU/Memory Guide

With the final test results, the team was able to make several conclusions. One was that vertical configurations are possible, i.e., Windchill, Oracle, Web server, servlet engine, search engine and directory server can all run on one physical system. A second was that the Sun Fire 4800 server should be configured to support no more than 1,500 concurrent users. And, since most of the Windchill 5.1 subsystems exhibited linear growth, they concluded that further scaling would be possible on larger Sun Fire servers. Information on the results that lead to these conclusions can be found in Appendix A.

Generating Server Configurations Using the Guide

Configuring a baseline Windchill system with the CPU/Memory Guide is a fairly simple task, however, it must be understood that these are general, baseline configurations only, and that the guide itself was a product of simulated user loads. Detailed information on the test and user workload specifics are located in Appendix A, and should be referenced to determine differences in customer environments that may affect the results of the guide. It is highly recommended that customers work with Sun and PTC Professional Services to design and plan, exact, detailed configurations and installations.

Note: Oracle sizing information was based on a relatively small test database. Customers should consult an Oracle professional to properly size the Oracle database for production environments.

To configure a general, baseline system, first determine the number of concurrent users for the system. Second, decide on a vertical verses horizontal configuration (see the Introduction for more detailed descriptions), and third, use the guide to choose the server(s) needed.

For example, a 1,200 concurrent user vertical configuration (Web server, Windchill server, and database all on one system) could be configured in the following manner:

- *Start with six 750 MHz CPUs for the Windchill Application server*
- *Add one CPU for the Web server and JRun (less than 2,000 users)*
- *Add two CPUs for Oracle (one for every four in the Application server)*
- *Add one CPU for the RetrievalWare search engine*
- *Add ten gigabytes of memory (one GB for every CPU up to 16)*

For a total of ten 750 MHz CPUs and ten gigabytes of memory. A Sun Fire 6800 server would easily fill these requirements and accommodate future growth.

Using the number above, a horizontal configuration might contain:

- *Four Sun Fire 280R Windchill Application and Web/JRun servers (two CPUs, two GB RAM each). This is a total of eight CPUs, six for the Windchill Application, and two for the Web server and JRun, spread across four servers.*
- *One Sun Fire 4800 Oracle/RetrievalWare server (four CPUs, four GB RAM)*

Growing this configuration to 1,750 users is achieved by adding a Sun Fire 280R Windchill server, and three more CPUs and gigabytes of memory to the Sun Fire 4800 server (one for Oracle and two for RetrievalWare). This configuration would also need a load balancing switch or other load balancing software to

distribute incoming client requests across the Windchill Application and Web/JRun servers. Logical examples of both of these configurations were given if Figures 1 and 2 in the Introduction section of this paper.

Memory Sizing and Allocation

As a general rule, the test results indicated that at least one gigabyte of memory should be included for each CPU configured. An important part of optimizing Windchill performance is balancing all application memory requirements with physical memory. Both the Windchill method server and the Oracle database server can efficiently cache persistent data in memory to minimize disk reads, increasing transaction throughput, and therefore require large amounts of memory.

In a vertical configuration, the available memory needs to be balanced across the system to avoid unnecessary disk paging. Table 1 provides a memory allocation model for vertical configurations.

Application	Percent of Available Memory
Windchill	31%
Oracle	31%
RetrievalWare	6.25%
Solaris Operating Environment	12.5%
JRun Servlet Engine	6.25%
Web Server	12.5%

Table 1 Memory Allocation

For horizontal configurations in which Oracle and Windchill are on separate servers or system domains, the percentages for both Oracle and Windchill should be increased to approximately 60 percent of the available memory.

I/O Sizing

An important part of configuring a Windchill system is determining the amount of disk space needed. Space must be allocated for system and application software, however, it is a small and constant amount compared with the very large and changing requirements of the Oracle database and RetrievalWare indexes.

Oracle Database Sizing

The PTC install scripts automatically configure database disk space, for either BLOBs or external file vaults, depending on whether the installer chooses a *small, medium, or large* configuration, as follows:

- *Small/Demo*: BLOBs=1.2 GB, Vault=1.3 GB
- *Medium*: BLOBs=2 GB, Vault=2 GB
- *Large*: BLOBs=4 GB, Vault=4.2 GB

These are conservative, initial database sizes, and take into account the desire of most installations to limit the amount of pre-allocated disk space.

The following PTC formula can be used to calculate a more realistic approximation of the target database size when all the data has been loaded:

Load Point + Content Data + Meta Data = Total Database Size

The *Load Point* is a constant of 600 megabytes, and represents the amount of disk space required for Oracle home and Oracle executables.

The value for *Content Data* is applicable for either BLOBs or external file vaults, and is calculated with the formula:

*Content = X*M + Y*N + (X + Y) / 50 G*

Where:

*X = Number of parts * average number of iterations*

M = Average size of part content in gigabytes

*Y = Number of documents or change objects * average number of iterations*

N = Average size of document content in gigabytes

The formula for obtaining *Meta Data* is:

*Meta Data = (X + Y) * .04C*

Where:

*X = Number of parts * average number of iterations*

*Y = Number of documents or change objects * average number of iterations*

C = Average size of document content in gigabytes

Database Sizing Example

This example uses a site that will contain 200,000 Windchill business objects after an initial legacy load. The characteristics for the data in this example are:

- 50,000 parts - with an average of two iterations per part
- 25,000 documents - with an average of four iterations per document
- 300 KB for document associated content
- 1 MB for part associated content

Following the formula, first calculate the Content Data:

*((Number of part * Number of iterations) * Average part content) + (Number of documents * Number of iterations) * Average document content) + Total Windchill Business Objects / 50 gigabytes*

Or:

*(100,000 * .300) + (100,00 * 1) + 200,000/50 GB = 134,000 GB*

Next, calculate Meta Data:

*(Number of parts * Number of iterations) + (Number of documents * Number of iterations) * .04 (Average size of all content)*

Or:

*Meta Data = (200,000) * .04 (.65)*

Finally, calculate the database size using the initial formula:

Load Point + Content Data + Meta Data = Total Database Size

Or:

.600 GB = 134 GB + 5.2 GB = 140 gigabytes

Therefore, in this example, the size of the database would be 140 gigabytes.

RetrievalWare Disk Sizing

It is extremely important to allocate disk space for the RetrievalWare search engine data indexes when calculating the total disk space required for a Windchill installation. These indexes are capable of requiring as much disk space as the actual content that was indexed. Depending on the configuration chosen for indexing data, the disk space needed by the indexes may be ten to 20 percent of the original content. Disk space should also be allocated for some future amount of data, i.e., one years anticipated growth.

2

Reference Configurations 3

The following examples were constructed by taking a cell from the *Windchill Application Server CPU/Memory Guide* for the number of users expected in an installation and adding CPUs per the guide to form an appropriate configuration. These are base configurations and because each customer installation will differ, do not include disk storage other than the system default.

Pilot and Development/Test Servers

1-200 Concurrent Users

The Sun Enterprise 220R server configuration in Figure 4 can support up to 200 concurrent users, has limited scalability, and requires rack mounting.

Figure 4 1-200 concurrent user Pilot or Development/Test configuration using a Sun Enterprise 220R server

1-200 Concurrent Users, Scalability to 500 Concurrent Users, Rack Mount

The Sun Enterprise 420R server configuration in Figure 5 can support up to 500 concurrent users with the addition of two 450 MHz CPUs and two gigabytes of memory.

Figure 5 1-200 concurrent user Pilot or Development/Test configuration using a Sun Enterprise 420R server, scalable to 500 concurrent users

1-200 Concurrent Users, Scalability to 500 Concurrent Users, Desk Side

The Sun Enterprise 450 server in Figure 6 can support up to 500 concurrent users with the addition of two 480 MHz CPUs and two gigabytes of memory.

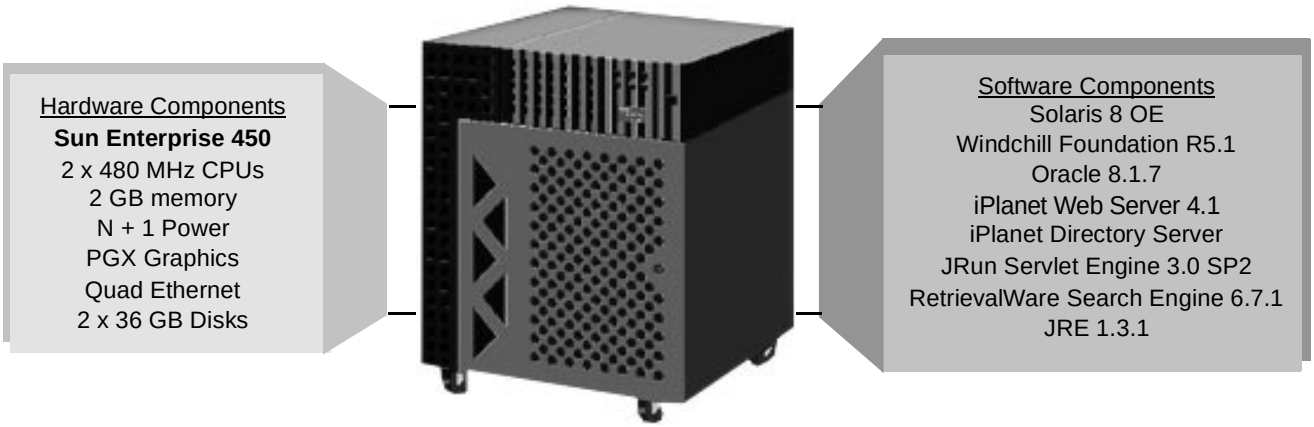

Figure 6 1-200 concurrent user Pilot or Development/Test configuration using a Sun Enterprise 450 server, scalable to 500 concurrent users

1-300 Concurrent Users

The Sun Fire 280R server in Figure 7 supports up to 300 concurrent users with two UltraSPARC™-III 750 MHz CPUs and two gigabytes of memory, and requires rack mounting.

Figure 7 1-300 concurrent user Pilot or Development/Test configuration using a Sun Fire 280R server

Production Servers

1-500 Concurrent Users, Scalability to 1,250 Concurrent Users

The Sun Fire 3800 server in Figure 8 can support up to 1,250 concurrent users with the addition of four 750 MHz CPUs and four gigabytes of memory and is rackmount or rack-ready. The Sun Fire 3800 server is Sun's entry level midrange server, offering unprecedented availability, performance, and flexibility, along with mainframe features such hardware redundancy and support for two Dynamic System Domains.

Figure 8 1-500 concurrent user Production configuration using a Sun Fire 3800 server, scalable to 1,250 concurrent users

1-500 Concurrent Users, Scalability to 2,250 Concurrent Users

The Data Center Class Sun Fire 4800 server in Figure 9 can support up to 2,250 concurrent users with the addition of eight 750 MHz CPUs and eight gigabytes of memory. The Sun Fire 4800 server is a versatile server with exceptional value and scalability for companies requiring an affordable, highly available business server with tremendous computing power, the ability to scale system performance and capacity as their needs grow, and the versatility available with two Dynamic System Domains.

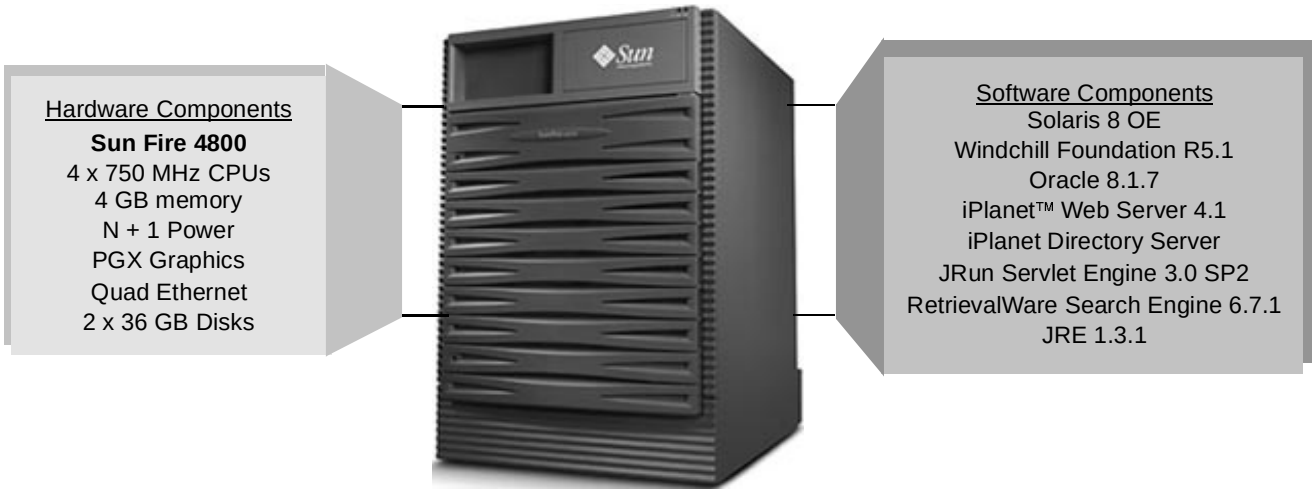

Figure 9 1-500 concurrent user Production configuration using a Sun Fire 4800 server, scalable to 2,250 concurrent users

High Availability Production Servers

1,250 Concurrent Users, Scalability to 5,400 Concurrent Users

The highly available solution in Figure 10 uses two Sun Fire 6800 servers in a cluster configuration to provide application failover and scalability through multi-domain support. This configuration can support up to 5,400 concurrent users with the addition of 12 CPUs and 12 gigabytes of memory. Support for up to four domains in the Sun Fire 6800 server increases flexibility and availability, and the expansion boards can be used to fine tune the servers by dynamically reallocating CPUs and memory as needed.

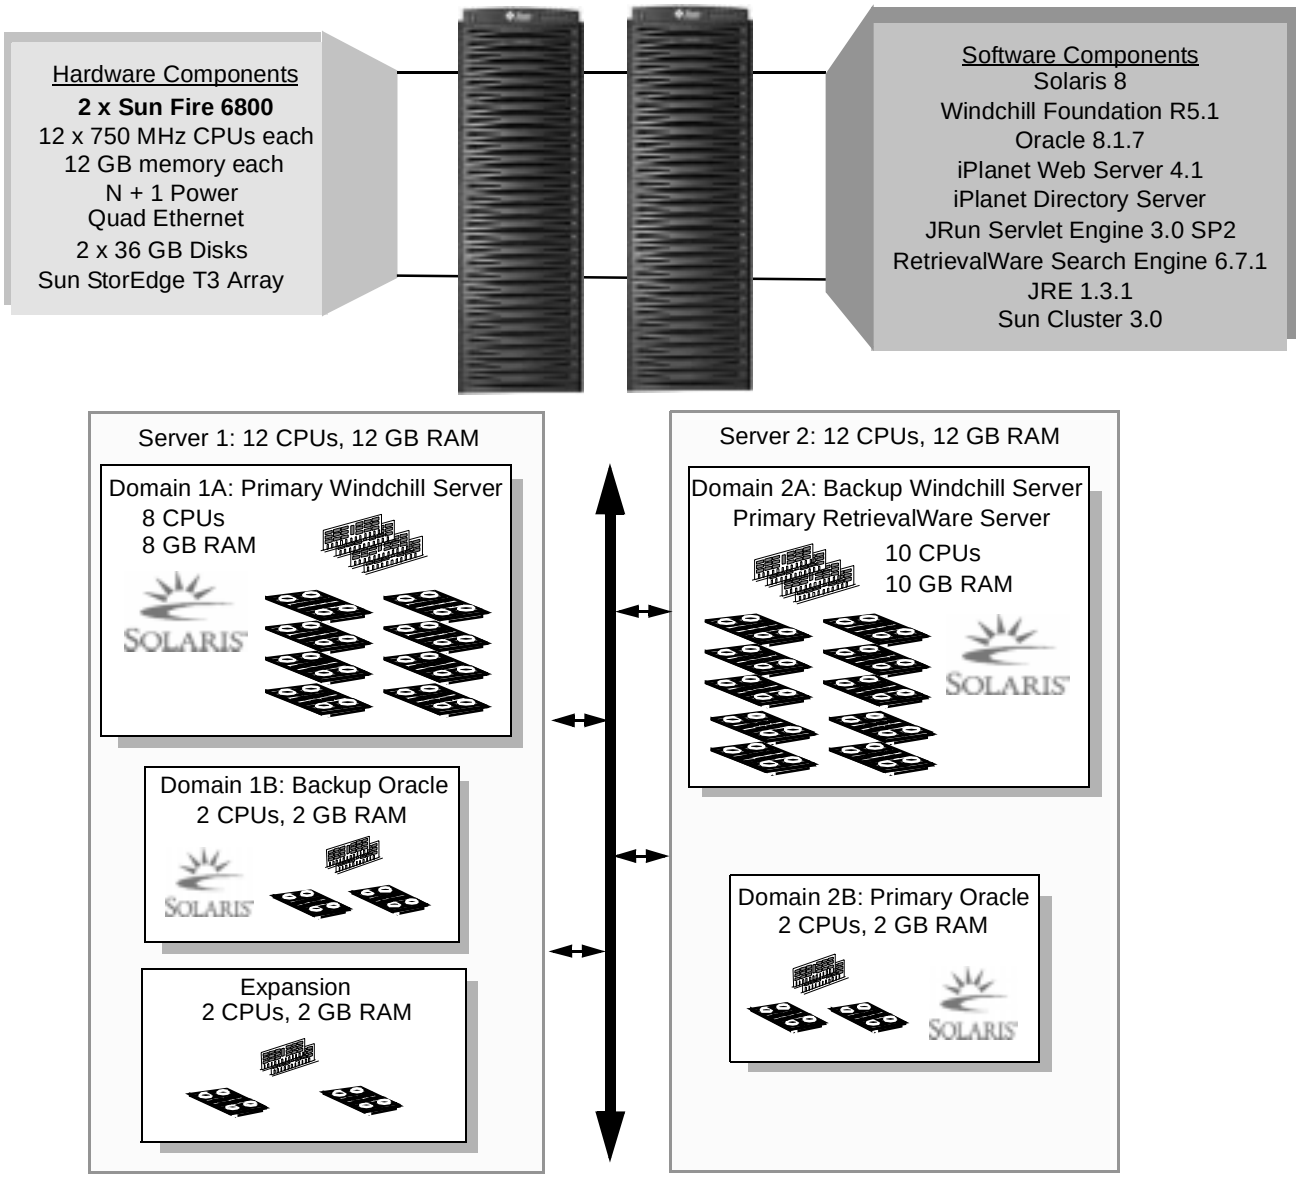

Figure 10 1,250 concurrent user Highly Available Production configuration using two Sun Fire 6800 servers in a clustered environment, scalable to 5,400 concurrent users

Horizontal Configuration — 2,000 Concurrent Users

The horizontal configuration in Figure 11 employs three Sun Enterprise 420R servers, two Sun Fire 3800 servers, and a hardware load balancing switch, providing application load balancing and failover in an n-tier solution. The horizontal nature of the Windchill Application and Web/JRun servers enables the increased flexibility of rapid application change and simplified change management, and allows the administrator to tune for network and application performance.

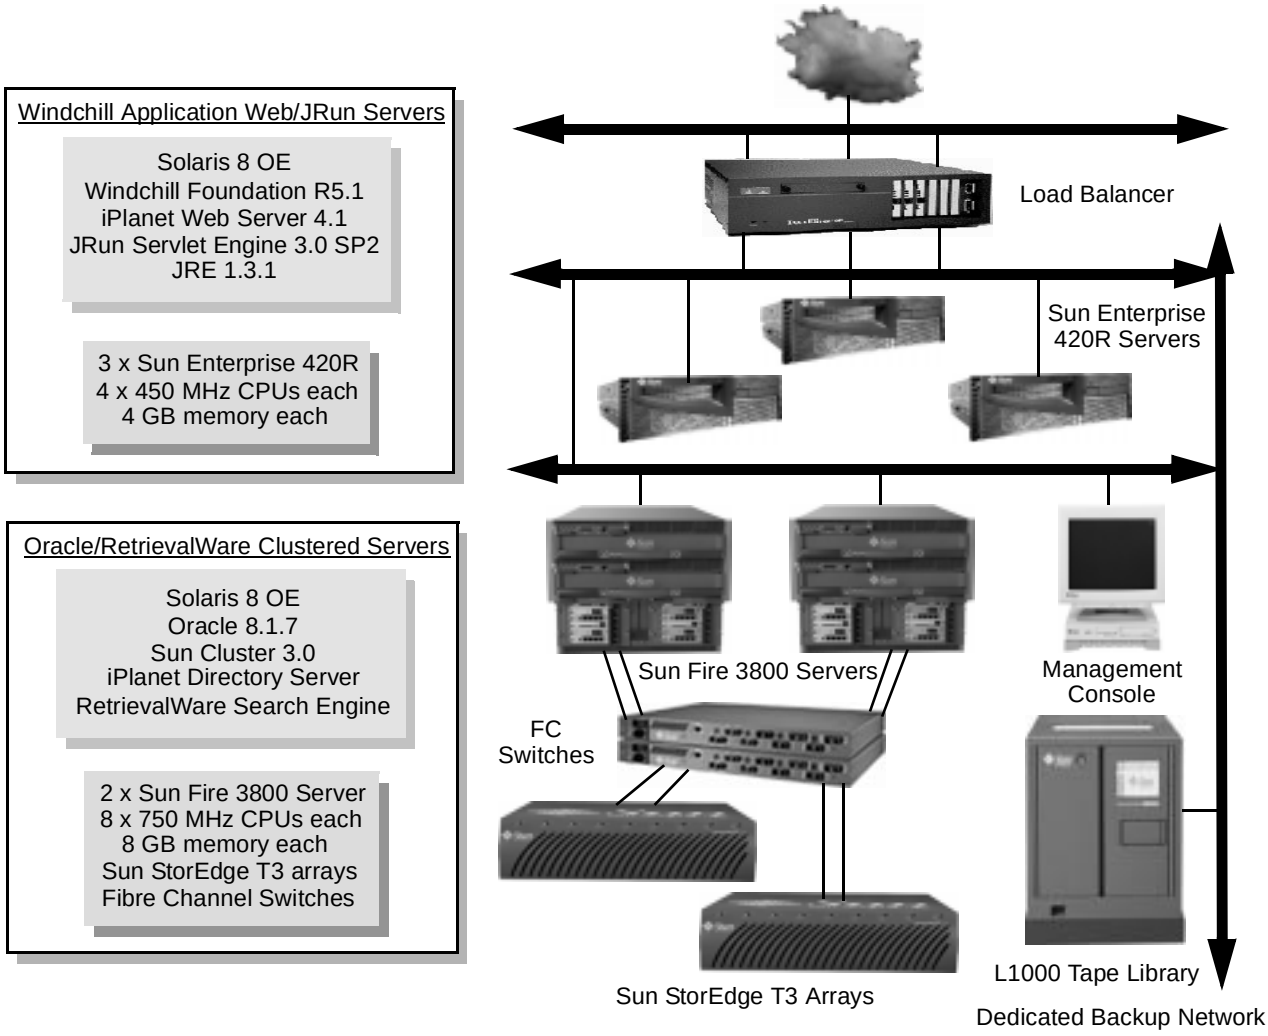

Figure 11 2,000 concurrent user, Horizontal configuration

Placing the Oracle database and RetrievalWare search engine on a separate, highly available cluster of systems allows this portion of the configuration to scale vertically — resources such as CPUs, memory, and storage can be incrementally added to the server over time to increase scalability — enables the administrator to tune the system for database and transaction performance, and

provides exceptional availability for the critical data stored on the system. This configuration also includes a tape backup solution with a dedicated backup network, illustrating a complete data center environment.

Appendix A: Details of the Test A

This chapter describes the characteristics of the tests and Windchill workload on a Sun Fire 4800 server. The characteristics of actual Windchill installations will differ unless the workload applied by the actual users is similar to the test workload. Reviewing the data in this chapter may help to determine the extent to which these test results may be applied to potential installations. Software settings for the various applications are also supplied to help in installation planning.

Test Environment

These tests were conducted on an UltraSPARC-III based Sun Fire 4800 Midframe server, with twelve 750 MHz CPUs and 24 gigabytes of memory. The data was stored on two Sun StorEdge™ T3 arrays and one Sun StorEdge A5200 array. The network was switched TCP/IP Ethernet, with a gigabit connection from the Sun Fire 4800 server to the switch, and 100 Mbps connections from the Segue PCs to the switch.

User Scripts

The scripts for this iteration of testing were adapted from the previous study conducted by Sun and PTC and did not include customizations such as vaults, vertical applications, administrative functions, visualization, or Info*Engine. The simulation of the client side load was performed with Segue SilkPerfomer, which supports the recording and playback of HTTP traffic and incorporates user

"think time" — the typical delays between client transactions. The script contained two phases in which the users logged in during phase one, and executed a series of transactions in phase two.

Test Results

A four stage test was performed on each configuration. The first stage started by simulating 60 *concurrent connections* (concurrent users waiting for a transaction to complete). During the second stage, a number of users (1,000 - 1,500) began to log in and became "concurrent" users (phase one of the user scripts). When all users were logged in, the third stage began, in which each user waited a random time period of up to 15 minutes before executing its first transaction. The fourth stage was the "measured test", a one hour period in which the transactions were measured. The initial delay of user transaction executions was used to keep the user concurrency constant for the one hour duration of the test.

In an attempt to provide mid-range test results, the team first tested 1,000 users on a 12 CPU system. The Sun Fire 4800 server ran effortlessly, as was indicated by a low number of concurrent connections and a 50 percent CPU utilization rate.

The next test simulated 1,500 users on a 12 CPU system. The number of concurrent connections was higher in this test as compared to the 1,000 user, 12 CPU test, however, average response time was highly acceptable, and the CPU utilization rate was about 85 percent, with RetrievalWare processes accounting for 40 percent of the usage. These test results indicated that the 12 CPU Sun Fire 4800 server was comfortably capable of handling a 1,500 user workload.

Subsequent tests started with 1,000 users on a seven CPU system. In this test, the number of concurrent connections was higher than the 1,000 user, 12 CPU test, and the CPU utilization rate was over 90 percent. Average response time was acceptable and Bill of Material (BOM) response times met the criteria for a well tuned Windchill system (average not to exceed 30 seconds). Additional tests were performed on eight and nine CPU systems with expected results.

With the final test results, the team was able to make several conclusions, which helped in the creation of the CPU/Memory Guide. One was that vertical configurations are possible, i.e., Windchill, Oracle, Web server, servlet engine, search engine and directory server can all run on one physical system. A second was that the Sun Fire 4800 server should be configured to support no more than 1,500 users. And, since most of the Windchill 5.1 subsystems exhibited linear growth, they concluded that further scaling would be possible on larger Sun Fire servers.

Software Settings and Information

The software settings detailed below are those used in the performance tests and are for a vertical configuration only. They are provided as a guide to help in planning and installing Windchill servers. Please reference the *PTC Windchill Performance Tuning Guide* for instructions on modifying files for horizontal configurations.

Oracle 8.1.7

The total size of the Oracle database for this test was less than six gigabytes and the System Global Area (SGA) was set to 1.3 gigabytes. The user table space, BLOBs and Oracle redo log files were stored on Sun StorEdge T3 arrays. Oracle rollback segments, as well as other UNIX and Oracle data sets were stored on a Sun StorEdge A5200 array.

This series of tests used BLOBs, which is Windchill's default data storage method. Other PTC test results have indicated that vault storage — a non-default storage mechanism — can be substantially faster than BLOB's, but was not tested due to time constraints. The high performance T3 arrays did not exhibit any significant load with the BLOBs as compared to the redo and rollback functions, thus indicating that the use of vaults to increase performance may no longer be necessary when T3 arrays are used.

The PTC Windchill Performance Tuning Guide states that significant performance improvements had been measured when placing BLOB data files on file systems mounted in direct I/O mode. Tests were conducted with direct I/O, however, the increased CPU time required to support it outweighed the possible benefit of reduced latency. Therefore, file system I/O was used in the final test runs. To further enhance the performance of the Oracle database, the PTC Windchill Performance Tuning Guide also recommends the use of raw disk devices administered with a volume management tool such as Veritas.

Solaris Operating Environment

The following changes were made to the /etc/system file:

set shmsys:shminfo_shmmax=0xffffffff

set semsys:seminfo_semmns=600

set symsys:seminfo_semmsl=200

set rlim_fd_max=8192

set rlim_fd_cur=1024

set tcp:tcp_conn_hash_size=65536

set fastscan=65536

set maxpgio=65536

Solaris 7 OE only:

set priority_paging=1

Windchill 5.1

PTC SPR numbers 870543 and 870544 performance updates were installed on Windchill for the tests, however, installers should contact PTC for current updates.

The two main configuration files for the Windchill application are:

WindchillHome/codebase/wt.properties

WindchillHome/codebase/dbs/db.properties

These files are located in the *codebase* subdirectory and can be edited by hand or through the Windchill application tool. The *wt.properties* file is modified for all Windchill specific customizations, site-specific settings, and some performance related settings. The *db.properties* file contains all database related settings such as passwords and simultaneous connections and is located in the *codebase/db* subdirectory.

Method Servers

Multiple method servers may be started on a single server to distribute and balance the load across multiple operating system processes. The optimal number of method servers is affected by the number of CPUs, location of the database, search engine, and scalability of the Java™ Virtual Machine (JVM), and should be determined experimentally for each installation. As a general rule, PTC suggests an initial configuration of one method server for every two or three CPUs. However, if Oracle and all the Windchill applications are running on the same server, it is necessary to allocate at least one CPU to execute other applications. The number of method servers is set with:

wt.manager.monitor.start.MethodServer=8

For the best performance on multiple CPU systems, a separate background method server should be configured as follows:

wt.manager.cmd.BackgroundMethodServer=$(wt.manager.cmd.MethodServer)

wt.method.serviceName=BackgroundMethodServer

wt.queue.executeQueues=true

wt.method.minPort=3000

MaxRedirects is the number of times the server manager will try another server before moving on to the next if the current method server is too busy, and should be set to the number of method servers minus one. In this example:

wt.method.loadbalance.maxRedirects=7

In order to determine the optimal number of method servers, the amount of CPU time the method servers collect should be monitored. Load balancing is affected if one server collects an unusually large amount, indicating that there are too many or too few method servers. The first of the next two entries causes a ServerLoadException to be thrown when the maximum number of concurrently active contexts is reached, before moving on to the next server. The second entry defines the maximum.

wt.manager.ServerSelector.MethodServer=wt.manager.BalancedServerSelect

wt.method.loadbalance.activeContext=8

The method servers write log entries in a single file by default, making it difficult to separate entries from multiple method servers. The use of multiple log files is helpful in determining performance problems, such as why a single method server has accumulated a disproportionate amount of CPU time. To write a separate log file for each method server, add the following lines to the end of *wt.properties*:

wt.manager.cmd.MethodServerLauncher=wt.manager.MethodServerLauncher

wt.method.log.file=$(wt.logs.dir)\ \ MethodServer{0}.log

wt.queue.queueMonitorHelper.verbose=true

wt.queue.waitStateChange=300

The frequency of log file entries is crucial in monitoring system performance and determining if the configuration is adequate for a specific load, and can be controlled with the property:

wt.method.methodSummaryInterval

Java Specific Settings

The *wt.properties* file contains reasonable configuration instructions. The following lines provide the various Java parameters and Windchill application startup options that were used during the test. The commented section is an excerpt from the *wt.properties* file.

```
#Note: The cmd and classpath variables below assume that your codebase contains
#all the class files necessary to run the server manager and method server.
#If you with to depend on your PATH and CLASSPATH environment variables,
#comment out the current wt.java.cmd and wt.java.classpath and uncomment
#the ones currently commented out.
wt.jdk1.1.8.workAround=false
#wt.jdk=/usr/java
#wt.java.cmd=java
#wt.java.cmd=/usr/java/bin/java
wt.jdk=/local/java/j2sdk1_3_1
wt.java.cmd=/local/java/j2sdk1_3_1/bin/java
#wt.java.classpath=$(wt.home)/codebase:$(java.class.path)
wt.java.classpath=$(wt.home)/codebase:$(wt.jdk)/jre/lib/rt.jar
```

To launch the Server Manager with Java command line options, use the appropriate lines from below:

```
#JSDK 1.3.1 execute line
wt.manager.cmd.ServerManager=$(wt.java.cmd) -server -classpath
$(wt.java.classpath) -Xms256m -Xmx256m -Xnoclassgc -noverify
wt.manager.ServerManagerMain
#JDK 1.2.2 execute line
wt.manager.cmd.ServerManager=$(wt.java.cmd) -classpath $(tw.java.classpath)
-Xms256m -Xmx256m -Xnoclassgc -noverify wt.manager.ServerManagerMain
```

To launch the Method Server with Java command line options, use the appropriate lines from below:

#JSDK 1.3.1 execute line

wt.manager.cmd.MethodServer=$(wt.java.cmd) -server -classpath $(wt.java.classpath) -Djava.protocol.handler.pkgs=HTTPClient -Xms640m -Xmx640m -XX:NewSize=128m -XX:MaxNewSize=128m -noverify wt.method.MethodServerMain

#JDK 1.2.2 execute line

wt.manager.cmd.MethodServer=$(wt.java.cmd) -classpath $(wt.java.classpath) -Djava.protocol.handler.pkgs=HTTP Client -Xms640 -Xmx640 -noverify wt.method.MethodServerMain

Miscellaneous Settings

Errors were periodically encountered during test runs. The error rate decreased after increasing the *wt.properties* file values of *clientTimeout* to 1,200 and *maxSockets* to 2,000.

wt.manager.rmi.clientTimeout=1200

wt.manager.rmi.maxSockets=1200

wt.method.rmi.clientTimeout=1200

wt.method.rmi.maxSocket=2000

Changes to the db.properties File

The following modifications to the *db.properties* file relate to how Windchill connects to the Oracle database and the resources consumed by that connection.

wt.pom.maxDbConnections=8

wt.pom.feebackInterval=200

wt.pom.statementCacheSize=300

wt.pom.cachedStatementReuseLimit=9999999

iPlanet Web Server 4.1

The iPlanet Web Server was installed as per Windchill instructions and required only one change to the *magnus.conf* file:

RqThrottle 4096

The *RqThrottle* parameter specifies the maximum number of simultaneous transactions the Web server can manage, and has a default value of 512. Increasing this value has the effect of throttling the server, resulting in minimized transaction latencies. Although this parameter affects multiple virtual servers, it does not function as a load-balancer.

The *obj.conf* file must be modified if using JRun to provide a link to the JRun servlet in order to bypass iPlanet's internal servlet engine. An entry to include the Windchill authorization gateway also needs to be added. Changes to this file are standard and are provided in detail in the JRun and Windchill 5.1 documentation, however, they may be unnecessary in future releases.

Caution: If Apache/Tomcat is used, check the iPlanet *obj.conf* file for the existence of a log level parameter. A setting of 'debug' has been reported to cause performance problems.

iPlanet Directory Server

The standard installation of the iPlanet Directory Server was sufficient for these tests, however, consistent use of the root name across all software components in the Windchill installation is necessary to avoid domain name problems.

JRUN Servlet Engine 3.0 SP2

The stated goal of this round of tests was to use the servlet engine included with the iPlanet Web Server, however, it was not yet certified with Windchill 5.1. At the time, Windchill supported JRun and Apache Software Foundation's Tomcat, an open source software. Tomcat was used in the initial tests, but was abandoned because of poor performance and because of customer concerns associated with support of open source products. Therefore, JRun was used during the tests.

The settings for general java properties, Windchill CLASSPATH settings, port numbers, log level reporting, etc. are located in *./servers/admin/local.properties* and *./lib/global.proerties.* Necessary modifications to these files are discussed in the Windchill documentation, and in the comments within the files.

In order to ensure the completion of all user transactions, regardless of time, the time out value was increased with the following parameter:

jcp.endpoint.main.timeout

Logging parameters can be set in the *global.properties* file to reflect preferences for log levels, log file names, and number of files saved.

Note: JRun uses a single, out of process Java Runtime Engine (JRE). With larger numbers of users, multiple JREs may be needed. PTC intends to increase the number of JavaServer Pages in future releases of Windchill, which will induce a heavier load on JRun, causing it to become a bottleneck. For this reason, Sun recommends using iPlanet Web Server 6.0 with Windchill 6.0.

RetrievalWare Search Engine 6.7.1

Changes to the *./windchill_indexes/config/rware.cgf* file are performed via a shell script from RetrievalWare. The installation for these tests followed standard procedure as documented in the Windchill installation Guide, however, some problems were encountered with the installation scripts, which were remedied by changing the default permissions.

Two steps are required to enable the RetrievalWare search engine. The first is to create the indexes, and the second is to load and configure the indexes for use. Problems were encountered with Windchill's bulk index load tool, forcing the test team to use a copy of the indexes from Windchill. New installations should expect to work closely with Windchill to acquire the latest scripts and patches until Windchill 6.0 is released.

Types of Users and Tasks

User Categories

The Windchill Performance Team defined two broad categories of users — Authors and Consumers — which represent the basic types of tasks performed in Windchill production environments.

- ***Authors:*** Users who *create and modify* any type of information that is part of a Windchill installation. They may be engineers or office personnel that use a variety of systems (from CAD/CAM/CAE to documentation systems) to create and modify business objects. The business objects most often represent documents, parts, and changes. These users typically search the database or scan cabinets/folders for business objects, and may checkout an object for modification or retrieve it to initiate an analysis application on the retrieved data set. Others will create, analyze, and approve changed objects. Team leads generally interact with their Worklist several times a day to analyze and/or approve changes and assign tasks. They will often remain in the application without interacting with the system for long time periods, averaging thirty-minutes between transactions. In large installations, authors comprise approximately 10 percent of the total number of users.

- ***Consumers:*** Users who *access* any type of information that is part of a Windchill installation. They are almost always the larger community compared to authors. Typical consumer tasks include: retrieving and viewing meta-data, performing data searches, and retrieving and viewing actual product data.

A ratio of 85 percent consumer to 15 percent author was used in this iteration of tests. Actual customer installation ratios of consumers to authors will vary from small installations with a higher ratio of authors, to large installations, which consist primarily of consumers.

Tasks and Repetitions

Table 2 summarizes the types of author and consumer tasks and number of times each task was expected to execute during the test cycle.

Authors	Repetitions per Shift	Consumers	Repetitions per Shift
Local Search	5	Enterprise Search	7
Expand Folder Structure	5	Local Search	3
Expand Product Structure	2	View Document	3
View Document	4	View BOM Report	2
Create Part Revision	1	Expand Product Structure	1
Create Document/Part	1	Expand Folder Structure	1
Checkout Document/Part	2		
Checkin Document/Part	2		
Perform Work Item	3		
Create Change Issue	1		
Add Subparts	1		

Table 2 Tasks and Repetitions

Transaction Weighting

Table 3 contains the Windchill transactions used in the tests to simulate real life user loads. These transactions, and their weights (the percentage of times they would be executed compared to the other transactions), are representative of typical Windchill environments, and should provide a baseline to use during the design of individual configurations.

Weight	Transaction	Weight	Transaction
0.03%	ADDUSES_PART_FORM	3.09%	ENTERPRISE_SEARCH_PAGE
0.55%	BOM_ASSEMBLY_100	1.15%	EXPAND_LEVEL1
0.10%	BOM_ASSEMBLY_200	1.15%	EXPAND_LEVEL2
2.27%	BROWSE_CABINETS_FORM	1.15%	EXPAND_LEVEL3
0.16%	CHANGE_REQUEST_ACTION	1.15%	EXPAND_LEVEL4
0.16%	CHANGE_REQUEST_PAGE	2.19%	ITERATION_HISTORY_FORM
2.19%	CHANGES_FORM	6.88%	LOCAL_SEARCH_FORM
2.27%	CHECKED_OUT_FOLDER_FORM	1.94%	PART_OBJPROPS_FORM
0.07%	CHECKIN_DOC_FORM	1.94%	PART_STRUCTURE_FORM
0.07%	CHECKIN_PART_FORM	1.94%	PART_TYPE_FORM
0.29%	CHECKOUT_DOC_FORM	24.97%	PERSONAL_CABINET_FORM
0.17%	CHECKOUT_PART_FORM	0.03%	REVISE_PART_FORM
0.04%	CREATE_DOCUMENT_PAGE	2.28%	SELECT_CABINET_FORM
0.03%	CREATE_PART_FORM	2.28%	SELECT_FOLDER_FORM
2.19%	DEPENDENCY_VIEW_FORM	14.78%	VIEW_DOC_FORM
4.54%	DO_DOCUMENT_SEARCH_FORM	0.27%	WORKITEM_FORM
1.94%	DO_PART_SEARCH_FORM	0.27%	WORKLIST_ACTION_FORM1
7.39%	DOC_OBJPROPS_FORM	0.27%	WORKLIST_ACTION_FORM2
4.54%	DOCUMENT_TYPE_FORM	0.27%	WORKLIST_FORM
3.00%	ENTERPRISE_SEARCH_ACTION		

Table 3 Transactions and Weights

Test Data

The data used in the tests was programmatically generated and was modeled by Windchill performance engineers to appropriately stress all aspects of system performance, and is described in Table 4.

Database Field	Value	Database Field	Value
Users	5,000	Part Reference Link	468,116
Groups	67	Part Describe Link	468,116
Projects	53	Iterations per Part	5
Shared Cabinets	1,000[1]	Revisions per Part	1
Document Masters	117,032	IBAs per Part Iteration	14
Document Iterations	117,032	Change Requests	722
Part Masters	29,258	File Vaults	2
Part Iteration	146,287	Assemblies	242
Part Usage Link	29,016		

1. Only Shared Cabinets 1-300 are readable by users 1-5,000.

Table 4 Test Database Data

Appendix B: Real Life Customer Configurations

B

There are many customers who are currently using horizontal configurations of Sun servers in Windchill environments. The following three examples detail real life customer installations supporting 2,000, 2,500, and 3,500 concurrent users.

2,000 User Customer Configuration

This large manufacturer of medical devices supports 2,000 users with two clustered Sun Enterprise 4500 servers, and three Sun Enterprise 450 servers, using an *Application Farm* approach, as shown in Figure 12.

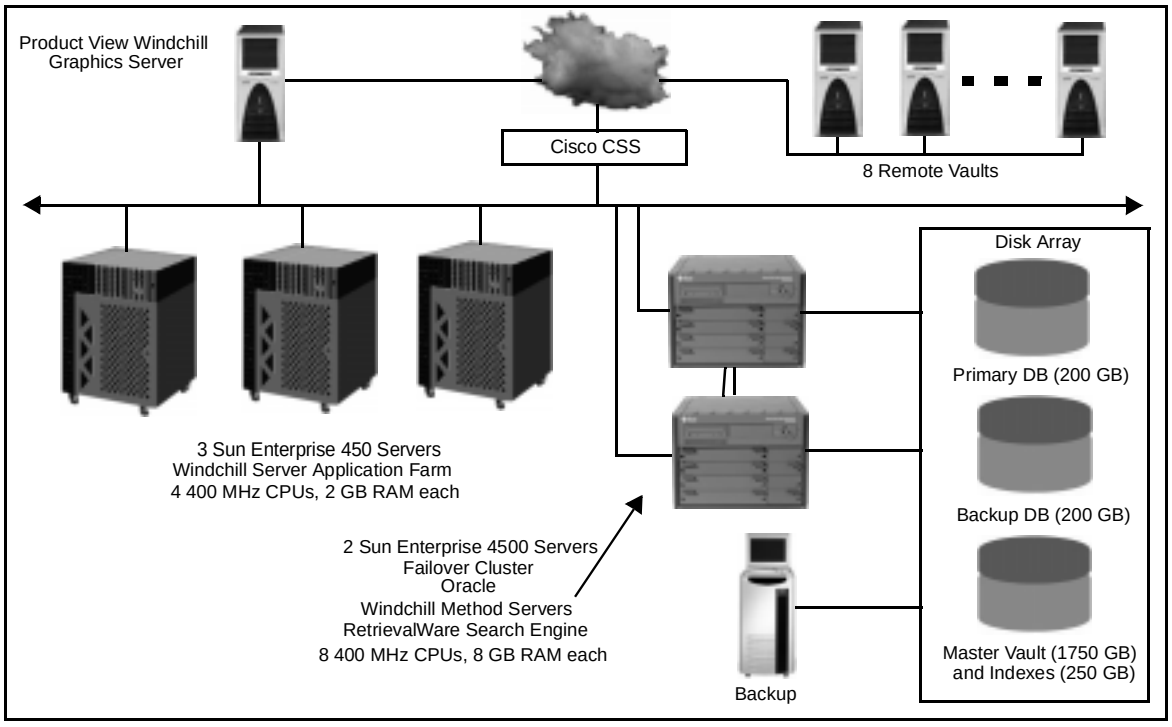

Figure 12 Real Customer 2,000 User Configuration

2,500 User Customer Configuration

The Windchill environment of a large data storage manufacturer supporting 2,500 users is illustrated in Figure 13. This configurations uses two Sun Enterprise 450 servers for the Windchill application servers, and two clustered Sun Enterprise 4500 servers for the database servers.

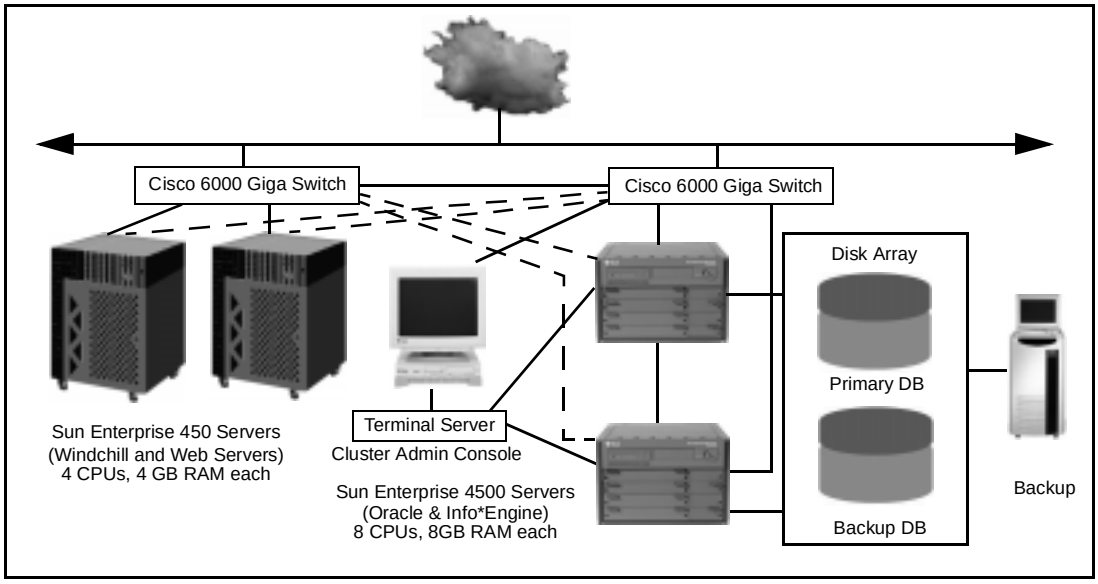

Figure 13 Real Customer 2,500 User Configuration

3,500 User Customer Configuration

A leading International electronics and electrical engineering company is currently supporting 3,500 Windchill users on Sun Enterprise servers. The configuration, depicted in Figure 14, consists of a pair of clustered Sun Enterprise 6500 servers running Oracle and NFS services, and three Sun Enterprise 4500 servers, two for Windchill and Web services, and one for development and test.

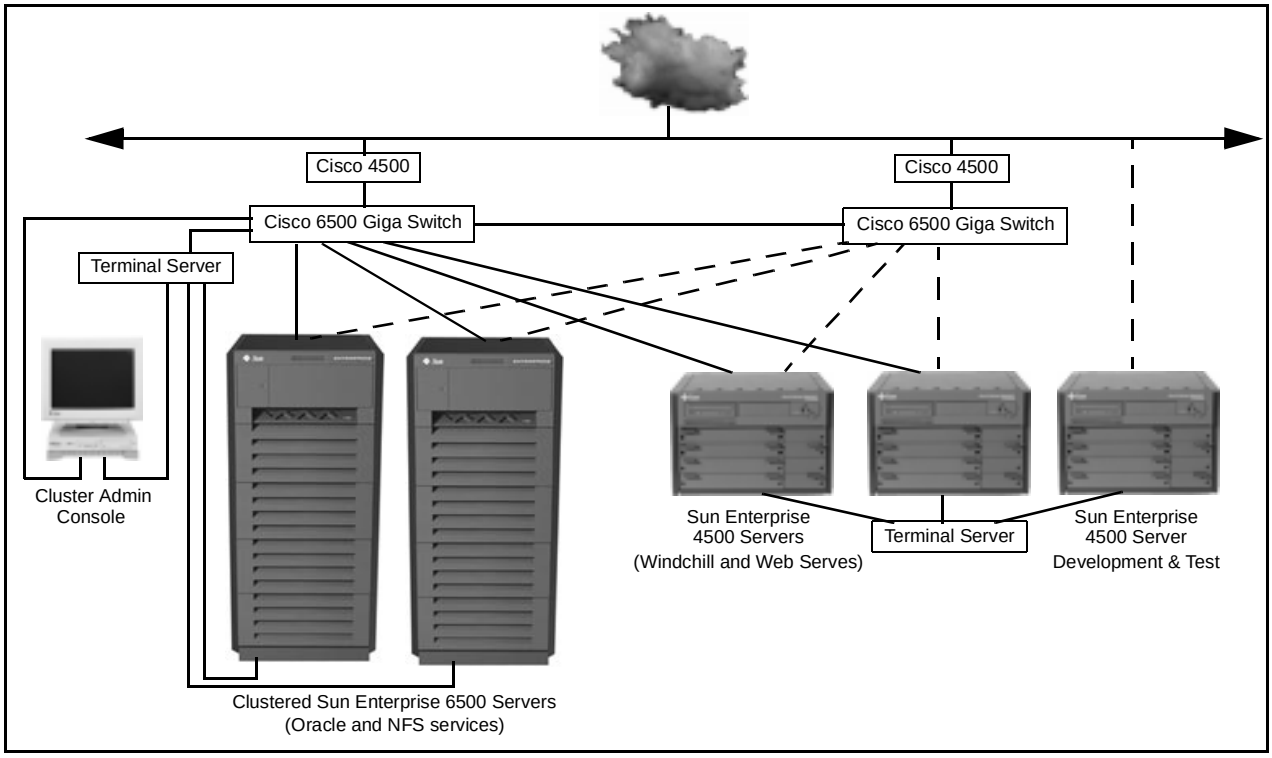

Figure 14 Real Customer 3,500 User Configuration

Appendix C: General Sizing Guidelines

C

Database Sizing

The size and content of the database in an implementation will impact overall performance. Larger databases can be more resource-intensive in both I/O and CPU utilization. A well-tuned, indexed database will scale better than one that has not been tuned. Windchill customers are expected to employ or contract RDBMS personnel who are capable of monitoring and tuning database performance.

Network Sizing

As the number of network clients increases, the network traffic will also increase and cause possible bottlenecks on slower network connections. All Sun servers come standard with fast ethernet, so network bottlenecks should not be an immediate issue. However, properly sizing the network is still an important task. The PTC Windchill Performance Tuning Guide, version 6.0 contains a section on Windchill Network Sizing Guidelines, which should be consulted in order to adequately size each segment of the network for optimal performance and to avoid possible bottlenecks.

Disk Sizing

One way to increase performance in a system is to carefully plan the I/O disk layout. Since disk access has the greatest latency, the initial task should be to keep what has been retrieved from disk in memory. The performance of

retrieving the data from disk can be increased by spreading the data evenly across multiple disks, and by ensuring there are enough disk controllers to transfer the data into memory.

Disk failure is the most common cause of system failure, and although the mean time between failure of individual disks is very high, the ever-increasing number of disks in large databases results in more frequent disk failures. Protecting the data from disk failure, and spreading the data across multiple disks with adequate controllers for performance can be accomplished by using hardware RAID 5 disk arrays such as the Sun StorEdge T3 array.

The Sun StorEdge T3 array's flexible architecture and dual-hosting capabilities make it an ideal storage medium for general applications and highly available cluster configurations. To protect against system failures in a Sun™ Cluster 3.0 07/01 environment, groups of arrays are connected to dual, redundant, failover, eight-port Sun StorEdge network Fibre Channel (FC) switches. Capacity ranges from 163 gigabytes in a single unit (using 18 GB disks), up to 2.6 terabytes in a rack of two four-unit partner groups (using 36 GB disks).

Availability Considerations

Sun offers many high availability solutions ranging from redundant system components to clustered failover servers. The level of availability needed will depend on the customers uptime requirements and budget. Implementing a system that delivers high availability means including strategies to minimize downtime, as well as techniques to enable recovery from disasters, such as deploying backup and disaster recovery procedures. This topic is discussed in more detail in the Sun white paper, *Sun StorEdge Products and Sun Cluster.*

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Internet World Wide Web Home page at: *http://www.sun.com*.

Look for the these and other Sun technology white papers:

Deploying a Collaborative Product Commerce (CPC) Solution with PTC Windchill

Sun Solutions for Collaborative Product Commerce (CPC) — A White Paper

Sun StorEdge Products and Sun Cluster.

PTC references and links:

http://www.ptc.com

Windchill Foundation R5.1 Scalability Test Definition

PTC Windchill Performance Tuning Guide, Releases 5.0 and 6.0

Other:

Oracle Essentials, Oracle 8 and Oracle 8i, by Rick Greenwald, Robert Stackowiak, and Jonathan Stern, O'Rielly & Associates

Sales Offices

Africa (North, West and Central): +9714-3366333
Argentina: +5411-4317-5600
Australia: +61-2-9844-5000
Austria: +43-1-60563-0
Belgium: +32-2-704-8000
Brazil: +55-11-5187-2100
Canada: +905-477-6745
Chile: +56-2-3724500
Colombia: +571-629-2323
Commonwealth of Independent States: +7-502-935-8411
Czech Republic: +420-2-3300-9311
Denmark: +45 4556 5000
Egypt +202-570-9442
Estonia: +372-6-308-900
Finland: +358-9-525-561
France: +33-01-30-67-50-00
Germany: +49-89-46008-0
Greece: +30-1-618-8111
Hungary: +36-1-202-4415
Iceland: +354-563-3010
India: +91-80-5599595
Ireland: +353-1-8055-666
Israel: +972-9-9710500
Italy: +39-039-60551
Japan: +81-3-5717-5000
Kazakhstan: +7-3272-466774
Korea: +822-2193-5114
Latvia: +371-750-3700
Lithuania: +370-729-8468
Luxembourg: +352-49 11 33 1
Malaysia: +603-264-9988
Mexico: +52-5-258-6100
The Netherlands: +00-31-33-45-15-000
New Zealand: +64-4-499-2344
Norway: +47 23 36 96 00
People's Republic of China:
- Beijing: +86-10-6803-5588
- Chengdu: +86-28-619-9333
- Guangzhou: +86-20-8755-5900
- Hong Kong: +852-2202-6688
- Shanghai: +86-21-6466-1228

Poland: +48-22-8747800
Portugal: +351-21-4134000
Russia: +7-502-935-8411
Singapore: +65-438-1888
Slovak Republic: +421-7-4342 94 85
South Africa: +2711-805-4305
Spain: +34-91-596-9900
Sweden: +46-8-631-10-00
Switzerland:
- German: 41-1-908-90-00
- French: 41-22-999-0444

Taiwan: +886-2-2514-0567
Thailand: +662-636-1555
Turkey: +90-212-335-22-00
United Arab Emirates: +9714-3366333
United Kingdom: +44-1-276-20444
United States: +1-800-555-9SUN OR +1-650-960-1300
Venezuela: +58-2-905-3800
Worldwide Headquarters:

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
Phone: 650-960-1300 or 800-555-9SUN
Internet: www.sun.com

Printed in USA