



Sizing Sun Ray™ Servers Running Windows Applications with SunPCi™ IIpro Coprocessor Cards

By Don DeVitt - Enterprise Engineering

Sun BluePrints™ OnLine - November 2001



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 816-2964-10
Revision 1.0, 10/11/01
Edition: November 2001

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun BluePrints, Java, Sun Enterprise, Solstice DiskSuite, SunPCi, Sun Ray, Sun StorEdge, Sun Quad FastEthernet, Sun StorEdge SRC/P Intelligent SCSI RAID Controller, Appliance Link Protocol, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

MetaFrame is a trademark or registered trademark of Citrix Systems, Inc in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc., in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Sun BluePrints, Java, Sun Enterprise, Solstice DiskSuite, SunPCi, Sun Ray, Sun StorEdge, Sun Quad FastEthernet, Sun StorEdge SRC/P Intelligent SCSI RAID Controller, Appliance Link Protocol, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

MetaFrame est une marque de fabrique ou une marque déposée de Citrix Systems, Inc aux Etats-Unis et à d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please
Recycle



Adobe PostScript

Sizing Sun Ray™ Servers Running Windows Applications with SunPCi™ IIpro Coprocessor Cards

Finding the lowest possible cost of ownership to support enterprise applications has continued to spawn new solutions for the thin client market. To reduce administration costs and consolidate the hardware requirements, an application server can be used to simplify application administration. Application servers run the application for end users using the resources of the server. End user, attached over the network, perform all I/O (input and output) on thin clients or appliances that have the minimal hardware complexity, and little or no administration overhead.

Sun's Sun Ray™ appliances were developed specifically to address this need to lower the Total Cost of Ownership (TCO). Sun Ray appliances are stateless, zero administration clients that process only keyboard, and mouse input, as well as screen and sound output, leaving all of the application processing and storage to the server. The Hot Desking Technology of the Sun Ray appliance allows user sessions to be completely mobile, allowing users to move from one Sun Ray appliance to another without bringing down their logged in sessions.

Supporting Applications on Thin Clients

The Sun Ray appliance is ideal for supporting Solaris™ Operating Environment (Solaris OE) (SPARC™ Platform Edition) based applications with minimal end user system administration. In the past, to support wintel (Windows running on Intel compatible processors) applications, it was required that a separate Windows NT, or Windows 2000, server be used to execute any Windows based applications that end users required. The Citrix MetaFrame™ product allows windows applications to run on a wintel server and by using the ICA (Independent Computing Architecture)

protocol, send the display, mouse, and keyboard I/O over the network to a remote client. With the Citrix MetaFrame client software running on the Sun Ray server, the Sun Ray support server could in turn take this information and channel the I/O to the Sun Ray appliance. Many customers have found the Citrix MetaFrame solution to be an excellent solution for controlling the environment in which windows based applications run.

While this solution does work well, it does require that two separate servers be maintained separately. Until now there was no way for a Sun server to support the execution of wintel applications from within the same Sun Enterprise™ server allowing for a full consolidation of the solution.

This article focuses on an integrated solution that combines several hardware and software products into one total solution. Specifically, we will narrow our discussion to developing a methodology for sizing a server that supports end users running both Solaris OE based applications and wintel applications. Individual details for configuring and administering each of the software components, that make up the solution, will not be dealt with in this article. There are many documents (see the *References* section) that cover these topics in detail. Instead we will focus on the details of integrating these components on one server.

A Solution Based on SunPCi™ IIpro, Citrix MetaFrame, and Sun Ray Products

A recent release of the SunPCi™ IIpro card and software now supports all the components necessary for supporting wintel applications on the same Sun Ray server. The latest release of this product supports running the server versions of Windows 2000 and Windows NT. In addition, the product also supports the installation of more than one card into one Sun Enterprise server allowing the solution to scale within the same server. A Sun Ray server can now support both SPARC applications and wintel applications without the need of additional servers.

The products that make up this consolidated server (other than the Sun Enterprise server, and the Solaris OE) are listed here with a brief description:

- **Sun Ray appliance** – Sun Ray appliances are stateless, zero administration clients that process only keyboard input and screen output, leaving all of the application processing and storage to the server.
- **SunPCi IIpro coprocessor card** – This PCI card can be added to a Sun workstation or a Sun Enterprise server and allow support for Microsoft operating systems. It has a 733 Mhz. Celeron Pentium processor, and dedicated ports for 10/100 Ethernet, USB, VGA, and sound I/O. Disk support is supported via the Solaris OE.

- **Citrix MetaFrame software** – This product consists of software that has both a server and client component that will allow users of a Solaris OE (SPARC Platform Edition) system to run wintel based applications running on another server supported by a Microsoft operating system.
- **Solaris™ PC NetLink software** – Solaris PC NetLink software delivers native Windows NT network services—which include directory, authentication, and file-and-print services—on Solaris servers.

It is assumed the reader will have a basic knowledge of each of these products. You may wish to refer to the *References* section at the end of the article if you are not already familiar with the each of these individual products.

The Architecture of a Full Solution

As mentioned above, Sun Ray solutions have been supporting both Solaris OE and wintel applications for some time. What is new, is that the SunPCi IIpro PCI coprocessor card can now support the execution of wintel applications for multiple users. This is now possible because the SunPCi IIpro coprocessor card and software can now support the server versions of Microsoft operating systems as well as the installation of multiple cards in one Sun Enterprise server. When enough hardware resources are allocated for each software component, the complete solution can support many users with more than acceptable performance.

Hardware Configuration

Before discussing the sizing requirements we need to look closer at the hardware and software configurations. FIGURE 1 is a diagram that shows how the various hardware components work together in an example configuration.

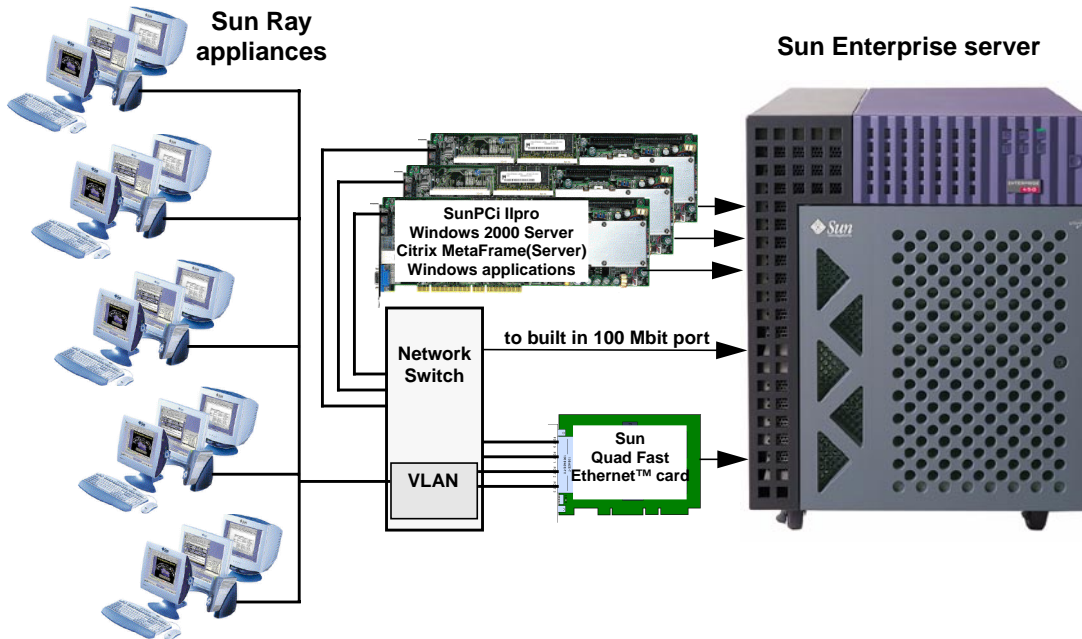


FIGURE 1 Sun Ray and SunPCI IIpro Server Hardware Solution: Hardware Component Interconnect

In this diagram, three SunPCI IIpro coprocessor PCI cards are installed into a Sun Enterprise 450 server. The SunPCI IIpro card requires a long PCI card slot and will require an additional spare PCI slot if the parallel or serial port will be used. When the SunPCI IIpro card is used as an application server, as it is in this article, you may use the parallel and serial ports which require a second slot to allow installation of a backplate that supports the serial and parallel port connectors.

Also installed into the server are Sun Quad FastEthernet™ PCI cards that supply four 100 Mbit Ethernet ports per card. This gives the system a total of five 100 Mbit ports, four on the Sun Quad FastEthernet cards and an additional port built into the Sun Enterprise 450 server.

The number of UltraSPARC™ II processors required, and the required memory will also be determined in the sizing exercise.

Each SunPCI Iipro coprocessor card has a dedicated 100 Mbit Ethernet port. All the bandwidth needed to support the Citrix MetaFrame product, and the end user's file access to their own home directories, will be supported by these ports. These ports are attached to a network switch on the same subnet as three of the 100 Mbit ports of a Sun Quad FastEthernet card. These connections allow the software running on the Windows 2000 operating system, on each of the SunPCI Iipro cards, to communicate directly with the software running on the Sun Enterprise 450 server.

The operating system, running on the SunPCI Iipro coprocessor card, will perform all disk operation through drivers that will use the PCI bus connection to communicate with the SunPCI Iipro software that supports the coprocessor card. This software will primarily support only the virtual disk drives used to support the Windows operating system and the files needed to support the applications used by the users.

In the preceding diagram the Sun Ray appliance network traffic is handled by a private network environment. The network that supports the Sun Ray appliance connections is typically implemented as a VLAN defined within the switch environment. The SunRay appliance communicate the screen, keyboard, and mouse I/O over the network by use of Appliance Link Protocol™ packets. These packets are in turn transported over the network by way of the UDP and TCP/IP protocols.

See the *References* section for a URL that displays specific information on Sun Ray networking.

Software Configuration

Now that we have reviewed the hardware configuration let us take a look at the software needed to support one user. FIGURE 2 is a chart that shows the flow of information through all the software components, from the Sun Ray appliance, to and from the applications the users are running.

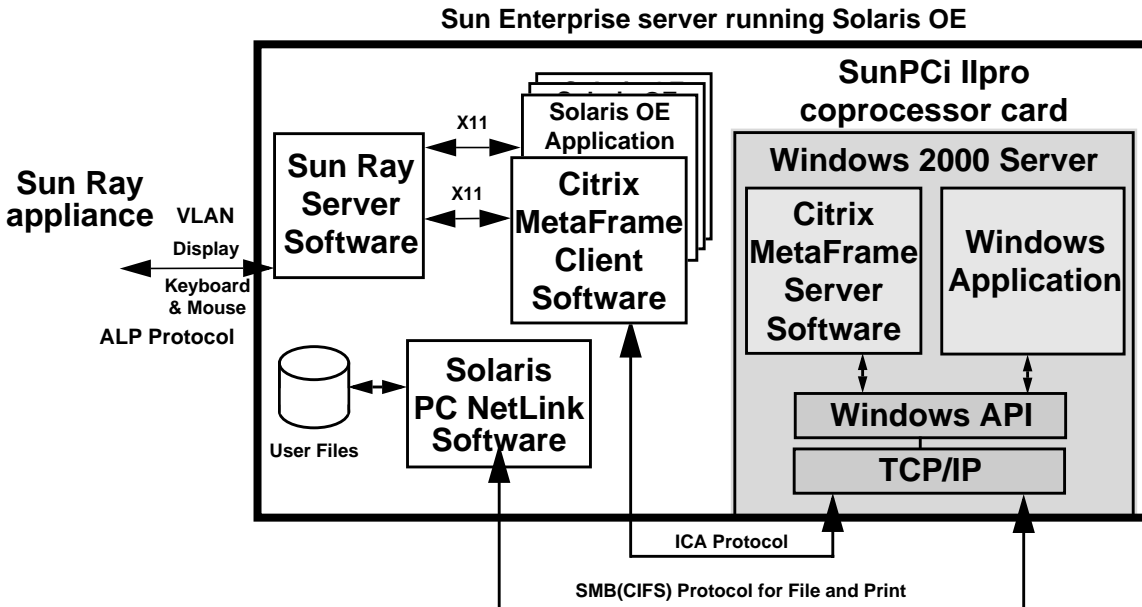


Figure #2 Software component interconnect

FIGURE 2 Sun Ray/Citrix MetaFrame/SunPCi Ilpro Software For One User

From left to right, we first see the Sun Ray appliance itself. The display, keyboard, mouse, and sound are sent to and from the Sun Ray server software via Appliance Link Protocol. This protocol uses UDP & TCP/IP to communicate to the server. The purpose of the Sun Ray server software is to translate the I/O of an X11 environment, to that of the Sun Ray appliance. For the purposes of this article we will be focusing on the I/O of the Citrix MetaFrame product. Users can simultaneously interact with any X11 based Solaris OE applications, running on the SPARC hardware.

The Citrix MetaFrame client software traditionally runs on Solaris OE workstations where users would interact with the software on the screen, keyboard, and mouse directly hooked up to the workstation. In our configuration, multiple copies of the Citrix MetaFrame client software will be running to support multiple users, each on their own Sun Ray appliance.

The Citrix MetaFrame client software talks directly to the Citrix MetaFrame server software which runs on the SunPCi Iipro card under the Windows 2000 operating system. To establish the connection between the two software components, a network interface is used on the Sun Enterprise server that is connected via a network switch to the private interface on the SunPCi Iipro card. The protocol that used between the two Citrix MetaFrame software components is a Citrix developed ICA® protocol that runs on top of the TCP/IP transport protocol.

The software running on the SunPCi Iipro card is the Windows 2000 operating system, Citrix MetaFrame server, and the applications each user will be running. The I/O for each application is intercepted by the Citrix MetaFrame server software and then sent to the Citrix MetaFrame client software. Each user running applications on the Windows 2000 server will spawn an instance of the Citrix MetaFrame server software.

Benefits of the Architecture

Using this all-in-one server approach to support an end user community offers many benefits. An abbreviated summary of these benefits is listed below:

- Benefits Sun Ray appliance:
 - Total Cost of Ownership (TCO) is reduced by 28%-58% (The Tolly Group TCO Model, 11/01/00).
 - Hot Desking Technology between Sun Ray appliances allows users to move from one Sun Ray appliance to another without logging off or logging back in by merely inserting their smart card.
 - Sun Ray appliances are completely stateless. If one breaks down, it can be quickly replaced with another.
 - All the system administration is performed on the server.
 - Most of the hardware investment is focused on the server where it benefits everyone.
- Benefits of the SunPCi Iipro coprocessor:
 - Once the first SunPCi Iipro card is configured, additional SunPCi Iipro coprocessor cards can be quickly cloned by copying the virtual C: drive (a Solaris OE file) of the first coprocessor to configure the remaining cards.
 - Like the Sun Ray appliance, the SunPCi Iipro card is stateless. If one card were to fail another card could be configured to quickly take it's place.
 - The Solaris OE software that supports the SunPCi Iipro coprocessor takes full advantage of the Solaris OE and the Sun Enterprise server to maximize performance and availability.
- Benefits of Solaris PC NetLink software:

- Allows the Sun Enterprise server to be a Primary (or Backup) Domain controller or a member server of a NT domain.
- Supplies fully compatible protocols for the SunPCi Iipro (and other Windows based PCs) to use.
- Sun Enterprise server and Solaris OE offer an extremely stable environment that all software components will be able to work together in.

Please refer to the individual product web pages for a more complete list of benefits.

Sizing the Total Solution

Unfortunately there is no one tool that can conveniently size the total solution. There are however, several tools available for sizing the individual components used in our solution. Sun Enterprise servers and the Solaris OE have been proven to scale well when handling multiple loads. By sizing each of the components separately, and insuring there are no fundamental resource conflicts, we can determine a baseline configuration that will act as a starting point for further consideration.

Goal

The general goal for this specific sizing exercise is to determine the requirements of a Sun Enterprise server needed to support a specific number of users. The goal here is to size a server that supports the following software components:

- **Sun Ray appliance software** - This includes all the software needed to support the Sun Ray appliance themselves.
- **SunPCi Iipro card** - These software components are needed to support the Solaris OE side of the coprocessor card. Primarily the resources consumed here will be supporting the access to the virtual drives used by the SunPCi Iipro cards to Windows 2000 operating system. No resources will be consumed to handle network traffic because the coprocessor have their own private 10/100 Ethernet port.
- **Solaris PC NetLink software** - These software components are needed to support the users data files. As the applications running on the SunPCi Iipro cards access user data files (spreadsheet, word processing documents, etc.), Solaris PC NetLink software will be used to support the CIFS (SMB) protocol, and Windows 2000 will use to access the files.

- **Citrix MetaFrame client software** - This software component supports the client side of the ICA protocol. The server side of the Citrix MetaFrame product redirects display, keyboard, and mouse input and output through the ICA protocol to the client portion of the software. There it is quickly converted to the X11 protocol which the Sun Ray software then display on the Sun Ray client.

During this specific sizing exercise, we wish to determine the following:

- Number of UltraSPARC II processors needed to support a specific number of users at an acceptable level of performance.
- Number of megabytes of memory in the Sun Enterprise server needed to support a specific number of users. This includes all the memory to support all software components).
- Number of SunPCi Iipro cards needed to support a specific number of end users.
- Number of megabytes of memory each SunPCi Iipro card should have.
- Number and type of network interfaces needed to support the configuration.
- Disk subsystem requirements of the server.

Assumptions

The number of variables associated with sizing a server with multiple software components can be nontrivial. You must decide at the beginning to fix certain variables or the task may become too difficult to complete in a reasonable period of time.

In the case of our server, one variable we will fix as much as possible will be the load placed on the server by each user. To determine the limits of a server, several benchmarks have been developed that attempt to simulate the load of various user communities. These benchmarks start by placing a simulated load of a particular type of end user on the server. The benchmark then proceeds to build the load, user by user, until a resource on the server becomes overloaded. The latency experienced by each end user will eventually exceed a reasonable upper limit which users find acceptable and the point where this occurs determines an upper limit for that Server configuration.

Defining the Load

Many of the benchmarks that measure latency of user application environments focus on a “knowledge worker” running Windows based desktop productivity applications. The description on the knowledge worker load can be found in the paper titled “*Windows 2000 Terminal Services Capacity and Scaling*” (Please refer to the *References* section for details on how to acquire this paper). Briefly, the knowledge

worker uses a variety of applications to add value to information in an office setting. This end user typically would run standard office productivity applications such as spreadsheets, word processors, presentation packages, and web browsers.

Additional loads have been used to simulate other kinds of users. One such load is the “data entry worker” load. This load attempts to define the load that a clerical or manufacturing worker may place on a server. This kind of load is considerably lighter than a knowledge worker because only one or two custom applications need to be supported.

While the knowledge worker will be the focus of this article, it should be kept in mind that a server can be expected to support typically 5 to 10 times as many data entry workers as knowledge workers.

In the real world, every end user places a different load on the server. Once the sizing data, based on the knowledge worker load, has been determined, you will need to determine if the end users of any target server you are planning will produce a load that is more or less that defined by the benchmarks used in the sizing exercise. If the load is greater, or the duty cycle of the end user is more frequent, more resources will need to be allocated per user. If the load is less, or the duty cycle of the end user is more infrequent, then less resources may be appropriate.

The system planner should always plan for the load of a server based on expected growth for at least 2-3 years. While the initial configuration can be designed to just meet today’s requirements, the server should be able to grow to meet the expected demand. Alternatively, additional servers can be used to support additional growth with the additional system administration that will be required to maintain it.

Assumption About the Target Server

The Sun Enterprise server products cover a vast spectrum of capabilities with respect to performance and RAS (reliability, availability, and serviceability) options. In this article we focus on the Sun Enterprise 450 server with four 400 Mhz UltraSPARC II processors. There are now many Sun Enterprise servers with more and faster processors. As long as the server meets the requirements of all the software components, it should scale up or down based on the number and type of processors.

It should be noted that server memory and SunPCi IIpro coprocessor card memory, have been determined to be a critical factors in supporting Sun Ray clients. If the server, or the SunPCi IIpro coprocessor card, runs out of physical memory for a significant period of time, the performance experienced by the end users will diminish quickly.

Baseline Configurations

Any baseline configuration determined by a sizing exercise is only a starting point for determining the final configuration of your server. It is always prudent to make resource related decisions that will configure the server to have some surplus resources. This will allow the server to have additional headroom, allowing the server to handle peak periods more gracefully, and allow for growth.

The papers listed below are used as references for our sizing exercise.

- *Supporting Microsoft Windows 2000 Server Applications from Sun Enterprise Servers* – From this paper you can determine the number of users each SunPCi II pro cards will support. This paper also determines the number of SPARC processors, and memory needed to support each SunPCi II pro card. <http://www.sun.com/blueprints/0601/windows.pdf>
- *Sizing Sun Ray Enterprise Servers* – This paper gives server sizing information for supporting Sun Ray appliances with a variety of load types. It is accompanied by a sizing tool which can be used to size the Sun Ray and Citrix MetaFrame XP (ICA client) server requirements. <http://www.sun.com/products/sunray1/whitepapers/sizing.pdf>
- *Solaris PC NetLink: Performance Sizing and Deployment* – Chapter 5 of this Sun BluePrints™ book covers sizing a Solaris PC NetLink server. <http://www.sun.com/books/catalog/devitt/index.html>
- *Solaris PC NetLink Sizing Tool* – This Sizing spreadsheet accompanies the information of chapter 5 of the Sun BluePrints book. The tool allows you to determine a baseline configuration of a server after supplying user load information. <http://www.sun.com/blueprints/tools/index.html>

Sizing the SunPCi IIpro Environment

The first paper in the list above (*Supporting Microsoft Windows 2000 Server Applications from Sun Enterprise Servers*) specifically deals with supporting a community of users running Windows 2000 applications via the SunPCi IIpro coprocessor card. It focuses on sizing a Sun Enterprise server to support SunPCi IIpro cards supporting users via Microsoft's Windows 2000 Terminal Server software. The function performed by Microsoft's Windows 2000 Terminal Server is essentially equivalent to that of the Citrix's MetaFrame product as far as the server side is concerned. Both products redirect display, keyboard, and mouse I/O from the Windows 2000 or Windows NT server to clients running over the net. Their loads are very close to identical.

In this Windows 2000 paper, it was shown that each SunPCi IIpro card can support 30 knowledge worker loads when the coprocessor card was configured with 512 MB of memory. It was also determined in that paper that each 400 Mhz UltraSPARC II processor was capable of supporting 3 SunPCi IIpro cards and 512 MB of Solaris OE memory was needed to support the 3 cards. Remember the 512 MB also includes the

memory needed to support the Solaris OE kernel. We will continue to focus on supporting 3 SunPCi IIpro cards throughout the rest of the exercise. 3 SunPCi IIpro cards, each supporting 30 users, allows the Sun Enterprise 450 server to support 90 knowledge worker loads.

Sizing the Sun Ray and Citrix MetaFrame Environments

The second paper from the preceding list (*Sizing Sun Ray Enterprise Servers*) is accompanied by a sizing tool that can size two of the software components at the same time. This Java™ programming language based tool allows you to define the load you want and includes support for sizing the server when both the Sun Ray software and Citrix MetaFrame (ICA) client software are used simultaneously.

For the sizing exercise of a knowledge worker, the load defined below was used:

- Java Applications(20%)
- Xbaseline(20%)
- CDE Utilites(10%)
- ICA client (MetaFrame) (50%)

When 90 users, with a 50% duty cycle, is entered into the tool for this user load, the resources in TABLE 1 were given to define two types of loads.

TABLE 1 Resources to Define Two Types of Loads

Resource	Average Load	Peak Load
400 Mhz Processors	2	2
7200 RPM drives	2	3
Memory	1296	2511
Swap Space	3172	2700
100 Mbit/Sec. Ethernet	2	2

Here we see two 400 Mhz CPUs will be needed to support the Sun Ray software and the Citrix MetaFrame client software. The memory requirement during the peak is 2511 MB. Keep in mind that this memory resource requirement also includes the memory requirement for the Solaris OE. We will deal with the fact that all the sizing tools include the memory overhead of the Solaris OE in the final analysis.

Sizing the Solaris PC NetLink Environment

The last component to be sized is Solaris PC NetLink software. It too has a Sizing tool in the form of a spreadsheet. This sizing tool was developed by running the Ziff Davis Netbench benchmark against many server configuration while changing resources one at a time. For a full description of the benchmark and how it was used to develop the sizing tool please refer to the References section. This benchmark also attempts to place a knowledge worker load of file operations on a file server.

When the 90 users are entered into the sizing tool, and peak period of 1 hour is chosen with a 40 MB file load, the tool recommends one 400 Mhz UltraSPARC II processor, and 458 MB of physical memory to be used. The Solaris PC NetLink software is designed to handle thousands of users. Supporting 90 users requires only a fraction of one processor. The tool is rounding up to the next whole processor. To determine what fraction of a processor the load actually places on the server, we can artificially step up the number of users until the tool recommends two 400 Mhz processors. Just below that point is how many users one full processor can support. The point where the tool recommends two processors is 336 users. To determine the fraction of one processor we will need to support 90 users is $90/336$ or .267 processors.

Again, this tool incorporates the memory overhead of supporting the Solaris OE.

Total Sizing Requirements

TABLE 2 summarizes the hardware requirements for supporting the total solution:

TABLE 2 Sizing Requirements

Software Component	400 Mhz UltraSPARC II processors	Sun Enterprise server physical memory (MB)	Disk Subsystem considerations	Network requirements and considerations
SunPCi IIpro compressors cards (3)	1	512	Software RAID 0+1, Hardware RAID 5 3x5 Gigabyte files	3 private ports on SunPCi IIpro cards (Not part of total)
Sun Ray software + Citrix MetaFrame XP (ICA) client software	2	2511-100= 2411	3 (7200 RPM) drives	Private network 2x100 Mbit Citrix ICA 1 x 100 Mbit
Solaris PC NetLink 1.2 software	.267	458 - 100 = 358	User files= 90x100Mbytes /user=9 gigabytes (RAID)	1x100 Mbit
Total Minimum requirements	3.267	3281		4 x 100 Mbit

SPARC CPU and Memory

In TABLE 2 above the CPU and memory recommendation was a summation of each individual sizing tool used for each of the products. 100 MB of memory was subtracted from the Sun Ray and Citrix software row and the Solaris PC NetLink software row because they included support for the Solaris OE which was already included in the SunPCi II pro sizing.

The performance of this solution under full load is highly dependent on having enough physical memory to support all the software components. As soon as either the SunPCi IIpro coprocessor card or the Sun Enterprise server run out of physical memory, the disk swapping will quickly cause performance problems.

The Solaris OE will utilize all spare physical memory as file system read cache. When the applications are properly configured, this read cache can play a significant role in keeping performance at a peak level. Having additional memory available will always improve performance.

To support three SunPCi Iipro cards we have a recommended baseline memory configuration of 3281 MB. Rounding up to the next valid memory configuration allows 4 GB for a Sun Enterprise 450 server. This additional memory will be put to good use by Solaris OE file caching, and to allow for some growth in the user load.

Similarly, the recommended CPU baseline configuration recommends 3.267 UltraSPARC II 400 Mhz processors. Rounding up to the next whole processor allows 4 CPUs. The additional CPU resource allows the server to have additional headroom to support the peak periods better. For the Sun Enterprise 450 server, choosing 480 Mhz would also give us additional head room.

SunPCi Iipro Memory

The SunPCi Iipro card is capable of being configured with up to 1 GB of memory. The SunPCi Iipro paper, referred to in the *References* section, showed that the SunPCi Iipro cards with 512 MB of memory are capable of supporting 30 knowledge worker loads while the same card configured with one GB of memory was capable of supporting 40 users based on benchmark results. If the real user load places a higher demand for memory than the benchmark, it may require you to use SunPCi Iipro cards that have 1 GB installed rather than 512 MB.

When ordering multiple SunPCi Iipro cards, it would be wise to order at least one SunPCi Iipro card with 1 GB of memory. This would allow you to later test to see if the additional memory allows for higher performance. Use Windows 2000 performance tools to monitor free memory so you know when the SunPCi Iipro card is approaching memory starvation.

In addition, you might wish to have at least one or two SunPCi Iipro cards configured with the additional memory (1 GB) so you can compare how SunPCi Iipro cards perform with additional memory headroom for supporting your user community. If, after an initial pilot deployment, it is determined that more users can be supported, then raising the user limit to 30 or even 40 may be possible in your particular user environment.

The SunPCi Iipro Disk Subsystem

One Solaris OE disk subsystem that needs to be sized and configured, is the disk subsystem needed to support the SunPCi Iipro virtual disks. This disk subsystem needs to contain the files needed to support the virtual drives (C: and D: Windows 2000/NT drives). Virtual drives are supported by SunPCi Iipro coprocessor support software by way of large files.

A primary benefit for using the SunPCi Iipro card, in the area of simplified system administration, is the implementation of Windows disk subsystems by way of Solaris OE files. These files can be easily copied to backup the Windows files system (while SunPCi is shutdown), or to create the template for another SunPCi Iipro card.

Using a “Server” Installation of the Application

When supporting a community of users all running the same applications, there are techniques that should be used to maximize the efficiency of the disk subsystem. One way to minimize the number of files that need to be accessed by the environment is to use a “Server” installation of the applications you wish to support. Microsoft Office, and other applications, have the ability to use a special server install that copies all the executables, libraries, and templates into one location that all users can use. Configuration files and other files that will change during normal use of the application will be placed into the user home directory. Once the application has been installed into a common location, all users can benefit from accessing the files. The Windows based operating system will need to open far fewer files to support the applications. This in turn will require considerably fewer disk requests from the Solaris OE. If there is a reasonable amount of unused physical memory, there is a high probability that the file data will be in disk cache resulting in extremely fast performance. This scenario is ideal for use with Solaris OE disk cache.

To support the Windows 2000 or Windows NT operating system, and to support the common “Server” installation of the application, five GB of disk was chosen as a recommended disk size. In your environment you may find that even three GB or even less can be allocated to support the SunPCi Iipro card virtual drives. It is important to avoid running out of disk space which may require reinstallation of some applications. After you have a working environment running on one SunPCI Iipro card, it is a simple matter to shutdown the card and make a copy of the file supporting the virtual drive. The copy can then be used as a template to configure other SunPCi Iipro cards.

To maximize performance and availability, you will want to place these files on redundant fast RAID disk subsystems. If software RAID (Veritas or Solstice DiskSuite™ software) are to be used, then RAID 0+1 would be best. If hardware RAID (such as Sun StorEdge™ A1000 array, Sun StorEdge T3 Array or the Sun StorEdge SRC/P Intelligent SCSI RAID Controller™ disk subsystem) is used then RAID 5 would be best.

If possible, placing each virtual drive file on a separate disk subsystem would be optimal to maximize the highest performance.

User Home Directories

A disk subsystem that needs to be sized separately is the disk subsystem needed to support the user home directories. This disk subsystem will be accessed via the applications running on the SunPCi Iipro coprocessor card via Solaris PC NetLink software. A full discussion on disk subsystems of Solaris PC NetLink software can be found in the Sun BluePrints book *Solaris PC NetLink Performance, Sizing, and Deployment*. In summary:

- Allocate enough disk space for each users data files, private executables and configuration files. In the sizing exercise presented here, 100 MB was allocated for each user.
- The disk subsystem should be redundant.
- The disk subsystem should be fast enough to not be the bottleneck for file access via Solaris PC NetLink software. A RAID 0+1 or Hardware RAID 5 should be used.

90 users × 100 MB per User = 9 GB. If user data files is all that needs to be supported, one RAID 1 volume made up of two 18 GB drives would be enough storage to support user data files. While this configuration would meet minimum requirements, it would be better if this user community was supported by more drives to minimize any possible contention for the same disk subsystem.

Network Considerations

Each SunPCi Iipro card has it's own private network port. Through this port all of the Citrix MetaFrame traffic (ICA protocol) and the Solaris PC NetLink file and print traffic (SMB/CIFS protocol) will flow. Benchmarks show that the combined traffic will rarely exceed 15-20% the bandwidth for these ports. All of the bandwidth for these cards will be redirected back into the Sun Enterprise server. One 100 Mbit connection should handle the throughput for all three SunPCi Iipro cards.

To allow for network redundancy and to allow for more headroom during peak periods, it would be best to allocate an additional 100 Mbit port on the server to support the traffic to and from the SunPCi Iipro coprocessor cards. The two ports could be configured as a trunked pair if your network switch supports combining the 2 ports together to improve throughput, simplify administration, and allow for a redundant path to the server.

The Sun Ray appliances will be connected to the server via the recommended 2 × 100 Mbit ports. The Appliance Link Protocol packets used by the Sun Ray clients is fully IP-compliant using both the UDP and TCP/IP transport protocols. While not absolutely necessary, it is recommended to use a switch VLAN to implement the network between the Sun Ray clients and the Sun Ray server.

So far there are requirements for 4 100Mbit ports. For a typical Sun Enterprise server this would suggest using one Sun Quad FastEthernet PCI card to support the four ports. All Sun Enterprise servers also have a 100 Mbit port built into the server. This port makes our total port count of five 100 Mbit ports, or one more than the baseline configuration requires. To put the port to good use, it would be best if the built in port and one port from the Sun Quad FastEthernet card be assigned to support the normal network traffic from the SunPCi Iipro card. Assigning the remaining three ports from the Sun Quad FastEthernet PCI card to the private network needed for the Sun Ray appliances would give that network more headroom and an opportunity for growth.

Deployment

It is rare when a real world load matches the load produced by a benchmark exactly. Some times the load is greater than what the benchmark places on the server, sometimes it is less. For this reason even though the benchmark shows that 30 knowledge workers can be supported by the card, it would be prudent to plan to support 25 users in an initial deployment or to use a pilot program to determine what the real load will be. This is especially true with user communities that are running more demanding programs than simple office productivity applications.

If possible, it would be best to assign users to be supported by each SunPCi Iipro card in groups of 5 or 10 until the real load of the community is better understood. After monitoring the utilization of the SunPCi Iipro card, by way of the Windows 2000 performance tools, you will begin to understand how the complete solution will be able to support your particular load.

Sizing Summary

While we have focused on an Sun Enterprise 450 server as a target server, any Sun Enterprise server is capable of supporting the SunPCi Iipro coprocessor card and has the ability to support the number and type of processors, memory, disk subsystem, and network that will be acceptable.

If a server is needed to support more users (90) than was presented here, this article can be used as a starting point to guide you in sizing the larger server. While a full sizing exercise is always recommended, sizing rules of thumb can be inferred from the sizing exercise presented here. By simply dividing the number of users that will be supported into a recommended baseline configuration, we get .0363 UltraSPARC II (400Mhz) processors per user, and 36.45 MB per user. This applies for knowledge worker loads only and is a recommended minimum baseline configuration. Additional resources allocated to the server will always allow the server to handle peak loading periods more gracefully.

From previous SunPCi Iipro research (see the SunPCi Iipro documents listed in the *References* section below) we see that each SunPCi Iipro card is capable of approximately 30 users, depending on the load and the amount of memory installed in the SunPCi Iipro coprocessor card.

Using the standard rule of thumb of creating a swap partition that is at least as big as the amount of physical memory would apply to this solution as well. As mentioned earlier, it is important that the active software components do not run out of physical memory and require the consistent use of swap under normal peak loading conditions.

For this solution, it would be dangerous to develop simple rules of thumb for the disk and network that would not over simplify the configuration. Please refer back to the network and disk subsystem sections above for information regarding the specific disk subsystem you should consider.

Proof of Concept Study

A proof of concept study was performed on a server utilizing all the software components discussed. In this study, the Citrix Server Test Kit (CSTK) was used to place a load on a server where Sun Ray sessions were running. Once benchmark scripts were started on a Sun Ray appliance, the session was detached using the Hot Desk Technology. The session continued to run in the background. Additional sessions were logged in and additional benchmark scripts were started.

The amount of CPU and memory resources consumed by a specific number of Sun Ray sessions running through Citrix MetaFrame XP came close to what was predicted by the sizing guide. It also showed that once the SunPCi IIpro coprocessor card ran out of physical memory, it was forced to swap out user applications which caused performance to suffer. Again, it is important that both the Sun Enterprise server and the SunPCi IIpro card consist of adequate memory to avoid swapping.

Different Loads

In this paper a knowledge worker was the target of the sizing exercise. Your particular load will ALWAYS be different. If you generally agree that the end users in your environment come close to the profile described in the papers listed in the *References* section, then using the recommended baseline configuration presented here is a good start. If you have a variety of users doing different things, then you may wish to be conservative assigning users to servers, slowly adding the number of users until you think you have reached a comfortable maximum.

If the load of your user community is significantly different than the load used here, it would be best to use this article as a methodology for sizing only. It would be prudent to go back to the individual sizing tools and perform the same sizing exercise with the new load.

Other Interoperability Solutions

The Citrix MetaFrame XP product was used in this article as the focus of the sizing exercise. There are other interoperability products that may suit your requirements better. Some of these products are SCO Tarantella, HOB HOBLINK, and GraphOn Bridges for Windows. A paper found on the Sun Ray products white papers web site is an excellent source of information on these products. It can be found at:
<http://www.sun.com/products/sunray1/whitepapers/interop.pdf>

Faster SPARC Processors

Sun Enterprise servers with UltraSPARC II 400 Mhz processors are the focus of this article. The primary reason for this is because all of the sizing tools (as of 6/01) used to size the server, include 400 Mhz UltraSPARC II processors as part of the recommended configurations. For the server solution, it is safe to assume that the faster, more capable, 480 Mhz/8 MB UltraSPARC II or UltraSPARC III processors will handle higher levels of performance. Most systems engineers find scaling the CPU solution up by the Mhz speed increase is a reasonable starting point for scaling up the servers capability for sizing. Also, make sure the Sun Enterprise server and the version of the Solaris OE have been qualified to work with all the individual software and hardware components.

Supporting Applications Running on Solaris OE vs. Wintel Environment

With the solution that was presented here, there are many situations where end users can be supported by applications running directly on the Solaris OE or the wintel environment. This sizing exercise focused primarily on supporting applications running on the SunPCi Iipro platform. If it is decided to support more applications directly on the Solaris OE, additional CPU and memory resources may be required to support these applications. Careful attention should be made to monitor the applications users are running on both environments to avoid situations where resources are consumed excessively.

Availability of SunPCi Iipro Coprocessor Cards

When supporting a large number of users with one server, it is important to maintain a high level of availability. Sun Enterprise servers have a high degree of availability. The SunPCi Iipro card relies heavily on the Sun Enterprise server it is installed into. When the Sun Enterprise server is configured with an appropriate RAID disk subsystem, the Solaris OE software that supports the SunPCi Iipro coprocessor card will take full advantage of it.

The coprocessor card itself has little redundancy built directly into the card. If one of the components fails, it typically won't bring down the Solaris OE or the other SunPCi Iipro cards running in the system, but it will no longer function for the users that use the card to run wintel applications. If you wish to avoid bringing down the Sun Enterprise server (which will mean down time for all users) to replace the failed SunPCi Iipro card, you may wish to install a spare SunPCi Iipro card in the Sun Enterprise server, so that you can quickly replace the functionality of the failed card without bringing down the server.

The fact that the SunPCi Iipro cards are almost completely stateless (with the exception of the Ethernet (MAC) address of the private 100 Mbit Ethernet port) any SunPCi Iipro coprocessor can be booted via the same virtual C: drive (a Solaris OE file) of the failed card. Only the users of the failed card will notice a short down time while the new SunPCi Iipro card is booted up to replace the failed card. After the Windows 2000 operating system does a disk check, the new SunPCi Iipro card will quickly take the place of the failed card. This unique feature of the architecture enables any SunPCi Iipro card to take the place of any other SunPCi Iipro card by simply redirecting the SunPCi Iipro card to boot from another disk image.

Implementing the Windows file system within a Solaris OE file also makes restoring the Windows 2000 system drive trivial. By simply redirecting the SunPCi Iipro software to a copy of the file that is used to implement the Windows 2000 boot disk, you can quickly recover situations that may corrupt the drive.

Conclusion

This article has focused on sizing a solution to support a community of users running wintel applications, on stateless Sun Ray appliances. The new SunPCi Iipro card enables Sun Enterprise servers to support the server versions of Microsoft operating systems. Using the SunPCi Iipro coprocessor in coordination with Citrix's MetaFrame XP product, a full solution can be offered to Sun Ray appliances on one server. This article has shown how the sizing tools, used to size a server for each software component, can be used together to size a total solution. Specific sizing "rules of thumb" for a "knowledge worker" load have been developed.

References

For additional information on the various components that make up this solution, please refer to the references below:

SunPCi Iipro PCI Coprocessor Card

Official web site: <http://www.sun.com/desktop/products/sunpci>

Documentation: <http://docs.sun.com/> Search for SunPCi

For information on scaling a SunPCi IIpro card for use with Windows terminal server refer to the Sun BluePrints OnLine article *Supporting Microsoft Windows 2000 Server Applications from Sun Enterprise Servers*. It can be found at:
<http://www.sun.com/blueprints/0601/windows.pdf>

Sun Ray appliance

Web site: <http://www.sun.com/sunray/product/>

Documentation:

<http://www.sun.com/products/sunray1/whitepapers/interop.pdf>

Sizing Paper:

<http://www.sun.com/products/sunray1/whitepapers/sizing.pdf>

Many additional reference solution and other technical papers on Sun Ray products:

<http://www.sun.com/products/sunray1/whitepapers/index.html>

Ethernet network switch information:

<http://www.sun.com/products/sunray1/switches/>

Solaris PC NetLink software

Sun BluePrints book: *Solaris PC Netlink: Performance Sizing and Deployment*

References to where this Sun BluePrints book can be found are at:

<http://www.sun.com/books/catalog/devitt/index.html>

Solaris PC NetLink Sizing Tool – This tool can be downloaded from the URL

<http://www.sun.com/blueprints/tools/index.html>

The Netbench benchmarks used to size the Solaris PC NetLink server components can be found at:

<http://www.etestinglabs.com/benchmarks/netbench/netbench.asp>

Citrix MetaFrame XP

For information on Citrix MetaFrame product please refer to:

<http://www.citrix.com/>

For Windows Terminal Server Sizing information see *Windows 2000 Terminal Services Capacity and Scaling*. It can be found at:

<http://www.microsoft.com/WINDOWS2000/techinfo/administration/terminal/tscaling.asp>

The Citrix Server Test Kit (CSTK) can be found at:

http://apps.citrix.com/cdn/SDK/cstk_sdk.asp

Acknowledgments

The author would also like to thank Craig Bender, Terry Mayer, and Michael Pratt (Sun Professional Services) for configuring the server and running the benchmarks that was the inspiration and proof of concept for this article.

Author's Bio: Don DeVitt

Don is currently a Senior Staff Engineer the PC interoperability specialist and a performance specialist within the Enterprise Engineering group. Don has been on the development teams of almost every software and hardware PC interoperability product that Sun Microsystems has produced over the past 15 years. Don is now working on Reference Configurations.

Don started his career as an electrical engineer and worked in the Automated Test industry (Teradyne Inc.), and PC operating system market (Digital Research from CP/M fame) before coming to Sun.