




Supporting Microsoft Windows 2000 Server Applications from Sun Enterprise™ Servers

By Don DeVitt - Enterprise Engineering

Sun BluePrints™ OnLine - June 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.: 816-1636-10
Revision 01, 06/05/01
Edition: June 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, Sun Enterprise, SunPCi, UltraSPARC, Sun Ray, Solaris PC NetLink, Sun StorEdge, Solstice DiskSuite, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

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, Sun Enterprise, SunPCi, UltraSPARC, Sun Ray, Solaris PC NetLink, Sun StorEdge, Solstice DiskSuite, 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. 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

Supporting Microsoft Windows 2000 Server Applications from Sun Enterprise™ Servers

Microsoft Windows 2000 and NT server consolidations

The quest to reduce the total cost of ownership of deploying and supporting productivity applications has spawned several trends by which organizations solve application management problems. Some organizations have decided to solve their PC application management problem by using Microsoft's Terminal Server software running on Windows 2000 servers for executing applications displayed on a variety of devices attached to the network.

Another trend that reduces the total cost of ownership, in the network servers segment, is replacing several slower, less stable servers with a single faster, more highly available server. This paper addresses a Sun solution that unifies both of these trends.

Wintel (the term "wintel" means an Intel compatible processor running a Microsoft Windows operating system) servers can only be consolidated to a limited degree, using wintel hardware as the target of the consolidation. Servers utilizing Sun hardware and the Solaris™ Operating Environment (Solaris OE) have earned the reputation of having High Availability (HA) and have a very high level of scalability. This makes them ideal target servers for Windows NT & Windows 2000 server consolidations, as long as the target server can support the required software. Sun Enterprise servers, based on SPARC™ processors and running the Solaris OE, are not capable of running wintel software directly. This is an issue that has blocked some consolidations in the past. The latest release of the Sun SunPCI™ IIpro coprocessor card has changed this.

This article explores using multiple SunPCi IIpro PCI cards, running in Sun Enterprise servers, to support services that previously could only be supported on Windows NT or Windows 2000 Servers. Specifically, this article will focus on the task of running wintel applications remotely for a variety of clients. While this article discusses the use of the SunPCi IIpro card to support Microsoft Windows 2000 Terminal Server, it helps illustrate that many Windows 2000 and Windows NT consolidation problems can be potentially solved by using similar Sun Enterprise server configurations, with SunPCi IIpro cards supporting the PC specific applications.

SunPCi IIpro coprocessor PCI card

The SunPCi IIpro coprocessor PCI card is a Sun product that, until recently, was dedicated to supporting wintel based applications on Sun's PCI based workstations. Previously, it supported only the non-server versions of Microsoft Windows operating systems. As of 2001, Sun has created a new version of the SunPCi IIpro software (2.2.1) that allows up to six SunPCi IIpro cards to be installed in Sun's workgroup server line (Sun Enterprise 220, 250, 280, 420, 450 servers), in addition to the Sun workstation systems. Another key feature of this new software is the following server versions of Microsoft Operating Systems are now supported: Windows NT 4.0, Windows NT Terminal Server Edition, and Windows 2000 Server (including Terminal Services). The new SunPCi IIpro card, which allows for up to one Gbyte of RAM and a 733 MHz Celeron Processor, coupled with the new software, now allows Sun servers to support wintel server applications at current server performance levels.

An abbreviated specification for the SunPCi IIpro PCI coprocessor card is:

- 733-MHz Intel Celeron processor
- Up to one Gbyte of memory
- Ports - High speed USB, built in HD-15 VGA, SoundBlaster compatible mike in and line out, 10/100 ethernet, (serial DB9 (COM1), parallel DB25 (LPT1:)) are available via a separate backplate
- Disk, floppy, and CD ROM devices are supported by the native Solaris OE.

A layout of the SunPCi IIpro card is shown in Figure 1.




FIGURE 1 SunPCi IIpro coprocessor card

For a full description of the SunPCi IIpro card, see:
<http://www.sun.com/desktop/products/SunPCI/>.

When installed into a Sun PCI based workstation, the SunPCi IIpro card can support a full Microsoft Windows 98 or Windows 2000 Professional environment on the Sun desktop. A previous Sun BluePrints™ Online article highlights many of the administrative capabilities of an earlier version of the SunPCi coprocessor card. See: <http://www.sun.com/software/solutions/blueprints/0899/pc2.pdf>.

The most recent release of the SunPCI IIpro software supports the installation of multiple SunPCi IIpro cards within one UltraSPARC II based Sun Enterprise Servers with PCI buses (support for UltraSPARC III based servers will soon follow). The latest software also supports Microsoft Windows 2000 (Server Version), as well as Windows 2000 Professional (Workstation Version). Utilizing the on-board 10/100 ethernet port, supporting multiple Windows 2000 server applications from within the same Sun Enterprise server is now possible.

Multiple SunPCi IIpro coprocessor cards in a single Sun Enterprise server

Before discussing a specific configuration, there are several points that should be covered:

Power and cooling

Most current Sun workgroup servers (such as Sun Enterprise 450 and 250 server) have multiple PCI buses which support the SunPCi IIpro PCI card. While any PCI slot is able to support a SunPCi IIpro card, close attention should be paid to the

power and cooling specification of the server. It is important to not exceed these specifications when adding multiple cards to one server. In the case of an Enterprise 450, the maximum number of Sun PCi Iipro cards that can be installed is six. More than six, will exceed the power and cooling specification.

In addition, the server should have enough backup power in case one power supply fails. If SunPCi Iipro cards are added to the server, the power requirements necessary to keep the system running during a power supply failure may be exceeded. In the example Sun Enterprise 450 server presented in this paper, three power supplies are needed, instead of two, in order to keep the system running if a power supply fails. Before adding SunPCi Iipro cards to a server, the power supply requirements should be reviewed to make sure enough power is available even if one power supply fails.

PCI bus considerations

The Sun Enterprise 450 server supports considerable PCI bus bandwidth by way of 10 PCI Slots, supported by six PCI buses and implemented by three PCI controller chips each controlling two PCI buses. While the PCI bus traffic going through a SunPCi Iipro card can at times be significant, it is unlikely to saturate a PCI bus. Like any high bandwidth PCI card, it would be best if it were installed into a PCI slot that is attached to a PCI bus shared with as few high throughput PCI cards as possible. The following chart shows what PCI slots are available in a Sun Enterprise 450 server. It should be used as a guide to avoid a potential bottleneck by placing too many high bandwidth PCI cards on one PCI bus. Placing a SunPCi Iipro card on

TABLE 0-1 Sun Enterprise 450 PCI slots

PCI Slot Number	PCI Slot type	Chip & Bus	/device special file
1	64 Bit/33Mhz 5Volt	Chip C Bus 2	/pci@6,4000
2	64 Bit/33Mhz 5Volt	Chip C Bus 2	/pci@6,4000
3	64 Bit/33Mhz 5Volt	Chip C Bus 2	/pci@6,4000
4	64 Bit/66Mhz 3.3 Volt	Chip C Bus 1	/pci@6,2000
5	64 Bit/66Mhz 3.3 Volt	Chip A Bus 1	/pci@1f,2000
6	64 Bit/66Mhz 3.3 Volt	Chip B Bus 1	/pci@4,2000
7	64 Bit/33Mhz 5Volt	Chip B Bus 2	/pci@4,4000
8	32 Bit/33Mhz 5Volt	Chip B Bus 2	/pci@4,4000
9	32 Bit/33Mhz 5Volt	Chip B Bus 2	/pci@4,4000
10	32 Bit/33Mhz 5Volt	Chip A Bus 2	/pci@1f,4000

each PCI bus would be the best way to distribute the load evenly. The “Chip & Bus” column of the chart shows which slots are attached to which PCI buses. Note that Sun Enterprise 450 server PCI buses are supported by an integrated circuit that supports two PCI buses per chip. “Chip A” has a “Bus 1” and a “Bus 2,” as does “Chip B” and “Chip C”.

With six PCI buses, PCI bandwidth or bus contention is rarely an issue with the Sun Enterprise 450 (and is generally true for all Sun servers with PCI buses). No PCI slot shares its PCI bus with more than two other cards. It is best not to move pre-installed PCI cards, such as SCSI, unless absolutely necessary.

The “/device special file” column in the previous table can be used to determine if the Solaris OE has correctly configured a specific SunPCi Iipro card into the system. The command “ls -go /dev|grep sunpci” can be used to see how many SunPCi Iipro cards have been installed in the server. The following sample output shows what is seen on a server with three SunPCi Iipro installed into an Sun Enterprise 450 serve. Each SunPCi Iipro requires three drivers (sunpci2drv, dummy, and volif).

```
testsys 75 =>ls -go|grep sunpci
lrwxrwxrwx 1 48 May 7 10:19 sunpci2drv0 -> ../devices/pci@1f,4000/pci108e,7063@4:sunpci2drv
lrwxrwxrwx 1 43 May 7 10:19 sunpci2drv1 -> ../devices/pci@1f,4000/pci108e,7063@4:dummy
lrwxrwxrwx 1 43 May 7 10:19 sunpci2drv2 -> ../devices/pci@1f,4000/pci108e,7063@4:volif
lrwxrwxrwx 1 48 May 7 10:19 sunpci2drv3 -> ../devices/pci@1f,2000/pci108e,7063@1:sunpci2drv
lrwxrwxrwx 1 43 May 7 10:19 sunpci2drv4 -> ../devices/pci@1f,2000/pci108e,7063@1:dummy
lrwxrwxrwx 1 43 May 7 10:19 sunpci2drv5 -> ../devices/pci@1f,2000/pci108e,7063@1:volif
lrwxrwxrwx 1 47 May 7 10:19 sunpci2drv6 -> ../devices/pci@4,4000/pci108e,7063@2:sunpci2drv
lrwxrwxrwx 1 42 May 7 10:19 sunpci2drv7 -> ../devices/pci@4,4000/pci108e,7063@2:dummy
lrwxrwxrwx 1 42 May 7 10:19 sunpci2drv8 -> ../devices/pci@4,4000/pci108e,7063@2:volif
lrwxrwxrwx 1 47 May 7 10:19 sunpci2drv9 -> ../devices/pci@4,4000/pci108e,7063@3:sunpci2drv
```

Applications that require multiple processors to support one instance of the application

While every Windows 2000 Server application that doesn’t require special hardware support should run on the SunPCi Iipro card, it is not the best choice for running applications that require multiple intel processors to support a single instance of the application software to maintain an acceptable level of performance. If the Windows 2000 server application requires more than one CPU to support acceptable performance, it is better to find the equivalent SPARC based application to run on the native SPARC based Solaris OE. Example applications that fall into this category are web services and messaging. There are many mainstream products that run on the SPARC platform that can support these services in a optimized and scalable way.

Supporting Microsoft 2000 Terminal Server clients

Microsoft 2000 Terminal Servers allow PC clients and UNIX® workstation users to run Microsoft Windows 2000 based applications running on a Microsoft Windows 2000 server. The application runs on the server and the display information is sent back to the PC client or the workstation, where it is presented to the end user. Windows 2000 now supplies the server and client code necessary to support PC clients. To support Solaris OE based workstations, or SunRay™ appliances, additional software will be required to support running wintel based applications from the Sun environment. One such product is the Metaframe product from Citrix.

Servers running Microsoft Terminal Servers (or the Citrix Metaframe product) are good candidates for consolidation by a SunPCi Iipro card running in Sun servers. One SunPCi Iipro card, running a Windows 2000 Server, can support many users. Once the maximum number of users has been reached, based on application latency, new users can then be assigned to additional SunPCi Iipro cards installed on the same Sun Enterprise server.

SunPCi Iipro software requirement for an X window.

Wintel based servers typically do not support headless (no display) operation. Likewise, the current Solaris OE based software supporting the SunPCi Iipro card requires an X environment to display a view of what is happening on the console.

In normal operation of a Windows 2000 Terminal Server, the server display simply has a login screen when no one is accessing the server from the console. This requirement means that a X server (the display) be available to accept the display information at all times.

If a graphics card is installed on the server where the SunPCi Iipro card is installed, it should be used to display the multiple SunPCi Iipro console displays. If no graphic card is installed on the server, then another X environment (perhaps on the administrator workstation) needs to be ready to accept the output of the SunPCi Iipro software. This will require that the appropriate `xhost <host>` command be issued to allow the X window information to be displayed from the server.

Starting the SunPCi Iipro cards after a server reboot

The SunPCi Iipro card installation procedure does not currently set up the SunPCi Iipro software to start automatically after the server it is installed in reboots. By default, after the Solaris OE boots up, the system administrator needs to start the SunPCi Iipro software manually. Alternatively, a script can be launched via a control file in the `/etc` directory. The X window dependency mentioned above, as well as

the location of the SunPCi IIpro virtual drive files, will require altering the script to suit a particular configuration. What follows is sample script that could be used to start a six card SunPCi IIpro environment:

```
setenv DISPLAY sysadmin:0
/opt/SUNWspci2/bin/sunpci -c Card1 -p /sunpcildrive/inifile/ntserver1.ini&
/opt/SUNWspci2/bin/sunpci -c Card2 -p /sunpci2drive/inifile/ntserver2.ini&
/opt/SUNWspci2/bin/sunpci -c Card3 -p /sunpci3drive/inifile/ntserver3.ini&
/opt/SUNWspci2/bin/sunpci -c Card4 -p /sunpci4drive/inifile/ntserver4.ini&
/opt/SUNWspci2/bin/sunpci -c Card5 -p /sunpci5drive/inifile/ntserver5.ini&
/opt/SUNWspci2/bin/sunpci -c Card6 -p /sunpci6drive/inifile/ntserver6.ini&
```

Each `ntserver(X).ini` file, referred to in the previous example, are SunPCi IIpro “.ini” files that control, among other things, which virtual drive file should be used by the SunPCi IIpro software as the “C:” drive when the card boots. A sample “.ini” file (`ntserver1.ini`) is shown:

```
[Drives]
A drive=/dev/rdiskette
C drive=/sunpcildrive/C.roboserver1
CD=/vol/dev/aliases/cdrom0

[Display]
OldText

[Disk32]
Enabled=Yes
```

Again, the script and “.ini” file will need to be changed to suit individual server configuration and other environmental factors.

Scaling and sizing a SunPCi IIpro coprocessor card running Microsoft 2000 Terminal server

Setting up a Windows 2000 Terminal server on a SunPCi IIpro card, without knowing how many users each card is capable of supporting, makes it difficult to size a server to support a particular user community. A benchmark for determining how many users can be supported by a Windows 2000 Terminal server is available

on the Microsoft web site. See:

<http://www.microsoft.com/windows2000/techinfo/administration/terminal/tscaling.asp>.

The benchmark is designed to run scripts on PC clients that simulate three different types of loads. Specifically the loads the benchmark is designed to produce are: “Structured Task Worker” (heavy load); “Knowledge Worker” (medium load); and “Data Entry Worker” (light load). To test the load capacity of a SunPCi Iipro card, the “knowledge worker” load of the benchmark was executed by 50 PCs to determine how many users each card could support when used in parallel with five additional cards. Please see the benchmark documentation for a full explanation of “knowledge worker” load.

The benchmark is designed to direct a load against a server, starting with one user load. The load is increased by five users for each iteration of the benchmark. The benchmark proceeds by increasing the load to 50 users, and measuring the time it takes to complete each iteration of the benchmark. A threshold is defined as being 10% slower than a standard run of the benchmark for one user. Once the measured benchmark times exceed the defined threshold, the number of user scripts running at that time determine the upper limit of users that the environment can be expected to support for that load. This is the limit a particular Windows 2000 Terminal server can support for this work load.

This benchmark was run against six SunPCi Iipro cards on an Sun Enterprise 450 server to determine how many users each card could support. The following configuration was used to benchmark the server:

Server:

1 Sun Enterprise 450 server - 4x400Mhz CPUs, 4 GByte RAM, 8 x 18 GByte drives, 1 Gigabit ethernet

6 SunPCi Iipro boards (each with 512Mbyte RAM) software V 2.2.1

Solaris OE 2.6

Clients:

60 x 300Mhz PCs running Windows NT Workstation 4.0

Figure 2 illustrates the benchmark configuration.




FIGURE 2 Benchmark configuration for Windows 200

When the benchmark was used to direct a load of 50 users to each SunPCI Ilpro card, the following latency chart (Figure 3) was the result.

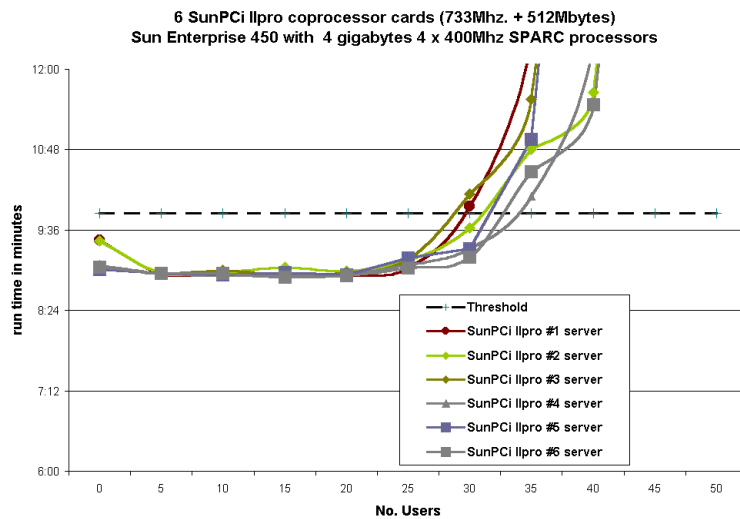


FIGURE 3 Runtime performance vs. PC Client

In this chart, the time it takes to run each iteration of the benchmark is plotted on the vertical axis. The benchmark starts with just one user. The time needed to execute one iteration of the benchmark for one user is used to define the baseline upper limit of acceptable performance. As the benchmark proceeds, more and more simulated users are run against each SunPCI Ilpro card. The same load is placed on all six

cards at the same time. Eventually, the load exceeds the throughput of one of the subsystems used by the SunPCi Iipro card, and the time it takes to finish an iteration of the test exceeds the baseline.

The number of users that each card can support, based on the “knowledge worker” load of the benchmark, is 30 users. This figure is reached by multiplying the six cards in the system, times 30 users per card, giving an upper limit of 180 users for this configuration.

Hardware limits to performance

The previous benchmark shows that one or more hardware components eventually reached their limit. The SunPCi Iipro environment, coupled with the Sun Enterprise 450 server running Solaris 2.6 OE, has a multitude of components that could have caused the benchmark to reach the user limit at 30. SunPCi Iipro card CPU, memory, Sun Enterprise 450 processors, the disk subsystem, or the network could have caused a bottleneck.

Network

The SunPCi Iipro coprocessor card has a private 10/100 Ethernet port. All network traffic is channeled to these private ports.

Older versions of the SunPCi Iipro card did not have their own 10/100 Ethernet port. Instead, network traffic was channeled through the Solaris OE on the host system to create a virtual network connection. The new SunPCi Iipro software will continue to support this network access when there is only one SunPCi Iipro card installed in a system, which is typically the case on workstations. If there are more than one card installed, the SunPCi Iipro software will disallow this virtual network access. Only the 10/100 Ethernet port on the SunPCi Iipro coprocessor card itself will be used to access the network.

If you are attempting to use the SunPCi Iipro card in a server capacity, and you are using only one card in the system, make sure you use the physical ethernet connection, as the virtual network connection is less efficient and places an additional load onto the SPARC processors.

After running the benchmark, the statistics from the network switch were checked to see if the 100Mbit Ethernet connection approached saturation during the test. This benchmark load never caused the port utilization to exceed 10%. Clearly, the network interface was not a limiting factor in supporting this kind of load.

Disk subsystem

The SunPCi Iipro card utilizes the Solaris OE files to contain the disk image for the system, and boot disk for the Windows 2000 Server Operating system. This virtual disk image is treated no differently than any other file on the UNIX file system, and, as such, can be moved, backed up, copied, etc. This disk image file can range in size from one MByte to eight GBytes in size, and is created and defined by the system administrator upon initial use of the card. Two virtual disk image files (C: Drive and D: Drive) can be assigned to a SunPCi Iipro card, along with the CD/DVD-ROM located on the Sun Enterprise 450 server. This virtual disk image file looks like a unformatted hard drive to the operating system that will run the SunPCi Iipro card. When the operating system is loaded, the system administrator creates the file system used by the OS, such as FAT16 (Win98, Win ME, NT4), FAT32 (Win98, Win ME, Win2K), NTFS4 (NT4), and NTFS5 (Win2K).

The file that is used to support the virtual disk subsystem in a real world environment should always be placed on a redundant RAID environment. If a hardware RAID exists, either by way of a Sun SRC/P PCI RAID card, or by way of a Sun StorEdge™ A1000 disk subsystem or a Sun StorEdge T3 Array disk subsystem, then a RAID 5 volume would be an acceptable solution to eliminate the vulnerability to a single disk failure and still have good performance. If software RAID (Veritas or Solstice DiskSuite™ software) is available, then a RAID 0+1 volume should be used to accomplish the fast, redundant RAID environment.

Supporting a Windows 2000 disk subsystem by way of a Solaris OE based UFS file has many advantages when compared to a stand alone Windows 2000 server. The first advantage is that it becomes trivial to create duplicate Windows 2000 Server environments. By simply copying (assuming the appropriate licensing exists) the file that is used to support the disk image, it can be made available to another SunPCi Iipro card software to create a duplicate Windows 2000 Server environment. After the copy operation is completed, an initial boot of the SunPCi Iipro card will require changing network settings to allow the SunPCi Iipro card to have a unique network name and IP address.

Backups are relatively simple to perform, as only one Solaris OE file needs to be backed up. For the Windows Terminal Server environment, described in this paper, users should be allocated space on a separate volume to maintain their private files. This volume could be accessible via PC Clients by using Solaris™ PC NetLink software.

For the previous benchmark, six 2-GByte C:\ drives were created, one for each SunPCi Iipro card. The disk image files were spread across six different 18 Gbyte SCSI disks. During the benchmark, a moderate to low amount of disk activity was seen as the SunPCi Iipro card software accessed these files. The key to making sure the disk performance does not bottleneck this kind of server environment is the amount of memory of the Solaris OE environment.

Solaris OE memory and the disk subsystem

The Solaris OE will use all unused physical memory for the purpose of disk read cache. As files are opened and read, unused physical memory is used to keep copies of the read data available for possible future access. This fact is of particular significance when it comes to supporting the SunPCi IIpro software. When the SunPCi IIpro card software initializes for the first time, it opens the file that supports the disk image of the Windows 2000 Server environment. This file is not closed until the SunPCi IIpro software exists. This scenario will receive maximum benefit from the Solaris OE read cache. As the Windows 2000 Terminal Server environment settles into supporting users, the disk data and executables needed to support the users environments has a high probability of settling into read cache, thus making disk performance appear to be extremely fast. The more memory a Solaris OE has available, the faster a SunPCi IIpro disk subsystem is likely to perform.

To fully optimize a specific set of applications running individually across multiple SunPCi IIpro environments, it is best to do a “server installation” of the application. Many Windows 2000 applications support running applications from a server where the files on the server are installed in a Read Only mode. Because they will not change, each SunPCi IIpro environment can access the same files. The Solaris OE read cache would be better utilized in this environment, because there would be fewer files to cache.

Several benchmarks were executed to determine the maximum number of SunPCi IIpro cards standard Solaris OE memory configurations can support.

Additional benchmarks were executed to determine how many SPARC processors were needed to support the SunPCi IIpro cards in the server. While the chart below shows one SPARC processor as the recommended minimum for one and three SunPci IIpro cards, for purposes of redundancy, it would be prudent to have an additional processor in the system.

The results are summarized in the following table:

TABLE 1

Solaris OE memory configuration	Recommended Maximum number of SunPCi IIpro cards	Recommended Minimum number of SPARC processors
256 Mbytes	1	1
512 Mbytes	3	1
1 Gbyte (or more)	6	2

Version of Solaris OE to use with the SunPCi Ilpro card.

While the SunPCi Ilpro card is supported on many versions of the Solaris OE, there are several compelling reasons for installing the card on a server running a current version of Solaris 8 OE. By doing so, it will automatically pick up fixes and performance enhancements added to the Solaris OE over time. Two specific examples (there are many more) are the new memory manager of Solaris 8 OE, and the enhancements to file system that were first shipped in Solaris 8 OE (1/01, also known as update 3). These two changes to the Solaris OE allow the SunPCi Ilpro software to work at peak efficiency.

Running other services on the server concurrently

It should be emphasized that the recommended resources needed to support a particular configuration of SunPCi Ilpro cards does not in any way include the resources needed to support other services (email, file and print, web, etc.) that may be running on the server. Separate sizing exercises should be performed for each of these services. The resources needed to support each service should then be combined to determine the final server configuration. Special attention should be made to insure there is little or no conflict for the same resources from the various services running concurrently on the server.

SunPCi Ilpro coprocessor memory

Another key component in supporting a Windows 2000 Server environment is the SunPCi Ilpro on card memory available to support the application it is asked to execute. This is especially true in the case of supporting a Windows 2000 Terminal server environment, where the server is actually executing the applications each user is executing. As long as all the users applications can be kept in physical memory, performance of the CPU is maximized. As soon as the memory is exhausted, the environment will start to swap memory, causing performance to suffer dramatically.

Two benchmarks were run to show how a lack of memory may have caused the previous benchmarks to eventually exceed the prescribed performance threshold. By installing an additional 512 Mbytes of memory in one SunPCi Iipro card, and running the benchmark again, the run times of a 512 Mbyte card versus a one Gbyte card were charted. The results are shown in Figure 4.




FIGURE 4 Run time performance of one SunPCi Iipro (512 Mbyte vs. 1 Gbyte).

The chart above shows that by adding memory to one SunPCi Iipro card, the benchmark continues to have acceptable performance for up to 40 users. Memory was at least one of the limiting factors in the first benchmark.

High Availability considerations

There is no redundancy built into each SunPCi Iipro card. If any critical component of the card fails, the card will not function. Unlike most stand alone Windows 2000 server however, each SunPCi Iipro card is almost completely stateless. Only the MAC (ethernet) address of the 10/100 ethernet port on the card differentiates one SunPCi Iipro card from another. This allows any idle SunPCi Iipro card installed in the system to be used to quickly restore any service caused by a failed SunPCi Iipro card. The system administrator can boot the new card with the old virtual disk image file, and the Windows 2000 server environment will quickly restore itself. Installing a spare SunPCi Iipro card and leaving it idle is the best way to reduce the down time of the services supported by a SunPCi Iipro card.

Conclusion

This paper has shown that the latest SunPCi Iipro hardware and software can be configured to support both the workstation and server versions of Microsoft operating systems. Each SunPCi Iipro card installed into a Sun Enterprise 450 server has been shown to support approximately 30 clients in a Windows 2000 Terminal Server environment. Because the Windows 2000 Terminal Server environment exercises all aspects of a Window 2000 server, it represents a worst case scenario for supporting Windows 2000 server applications. By supporting multiple SunPCi Iipro cards in one system, Sun Enterprise servers can now support a multitude of server environments that could not be supported before.

References

For additional information on SunPCi Iipro card, refer to the web site:

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

For additional information on Windows 2000 Terminal Server, please refer to the Microsoft web site.

For information on Citrix Metaframe Product please refer to

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

Solaris SunPCi Iipro documentation is available on

<http://docs.sun.com>

Acknowledgments

The author would like to thank James Freed (SunPCi Engineering Team) for performing the benchmarks used in this paper. The author would also like to thank Craig Bender (Sun Professional Services) for his part in writing this paper.

Author's Bio: Don DeVitt

Don is currently a Senior Staff Engineer and PC interoperability specialist within the Enterprise Engineering group, and a member of the Solaris PC NetLink engineering team, where he has focused on performance-related issues. 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 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.