

SunTRI/P 4.0 Installation Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300
Fax 650 969-9131

Part No.: 805-6905-10
October 1998, Revision A

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 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.

Portions of this product may be derived from the UNIX® system, licensed from Novell, Inc., and from the Berkeley 4.3 BSD system, licensed from the University of California. UNIX is a registered trademark in the United States and in other countries and is exclusively licensed by X/Open Company Ltd. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers. 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).

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

The OPEN LOOK® and Sun™ Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox Corporation in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a nonexclusive 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.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303-4900 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie et la décompilation. Aucune partie de ce produit ou de sa documentation associée 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.

Des parties de ce produit pourront être dérivées du système UNIX® licencié par Novell, Inc. et du système Berkeley 4.3 BSD licencié par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays, et licenciée exclusivement par X/Open Company Ltd. 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.

Sun, Sun Microsystems, le logo Sun, SunTRI, OpenBoot, et Solaris sont des marques déposées ou enregistrées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC, utilisées sous licence, sont des marques déposées ou enregistrées 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.

Les utilisateurs d'interfaces graphiques OPEN LOOK® et Sun™ ont été développés de Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox Corporation 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, cette licence couvrant aussi les licenciés de Sun qui mettent en place les utilisateurs d'interfaces graphiques OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A REpondre A UNE UTILISATION PARTICULIERE OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.





Contents

- 1. SunTRI/P Adapter Introduction 1**
 - Token Ring Networks 1
 - SunTRI/P Adapter Overview 3
 - Hardware Overview 3
 - Software Overview 3
 - SunTRI/P Adapter Features 4
 - Preparing for Installation 4
 - Requirements 4
 - Adapter Handling Precautions 4
 - Connecting to the Network 5
 - Recommended Types of Cabling 5
 - Platform Compatibility 6
 - Inserting the Adapter 6
 - Card Slot Selection 6
 - ▼ To Install the SunTRI/P Adapter 7

- 2. Installing the SunTRI/P 4.0 Software Packages 9**
 - Software Installation and Configuration 9
 - Installation and Configuration Overview 9

Summary of Installation Steps	10
▼ To Remove Previous Versions of the SunTRI/P Software	10
▼ To Mount the CD-ROM	11
Installing the <code>trp</code> Device Driver	11
▼ To Install the Driver	11
3. Software Configuration	13
Configuring the SunTRI/P Software for TCP/IP	13
Editing <code>trp.conf</code> Parameters	14
Setting the Maximum Transmission Unit (MTU)	15
Configuring the Ring Speed	15
Configuring the Locally Administered Address (LAA)	16
Configuring the System To Use The Embedded MAC Address	17
Configuring the SunTRI/P Adapter as the Sole Network Interface	17
Configuring Internetwork Routing	18
4. Troubleshooting and Diagnostics	19
Troubleshooting Tools	19
Ifconfig	19
Netstat	20
Ping	20
Snoop	21
Spray	21
Trpsr	21
Mver	22
Maintenance Hints	22
System Console Messages	22
Error Messages	23

Tables

TABLE 1-1	Platform Compatibility	6
TABLE A-1	OPEN stages	28
TABLE B-1	Seven Layers of Open Systems Interconnection (OSI)	31
TABLE C-1	Overhead Determining Factors	55

Preface

This guide provides the information needed to install and configure the hardware and software for the SunTRI/P Adapter. These instructions are designed for an experienced system administrator with networking knowledge.

Using UNIX Commands

This document may not contain information on basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook™ online documentation for the Solaris software environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Sun Documentation on the Web

The `docs.sun.com` web site enables you to access Sun technical documentation on

the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

`smcc-docs@sun.com`

Please include the part number of your document in the subject line of your email.

SunTRI/P Adapter Introduction

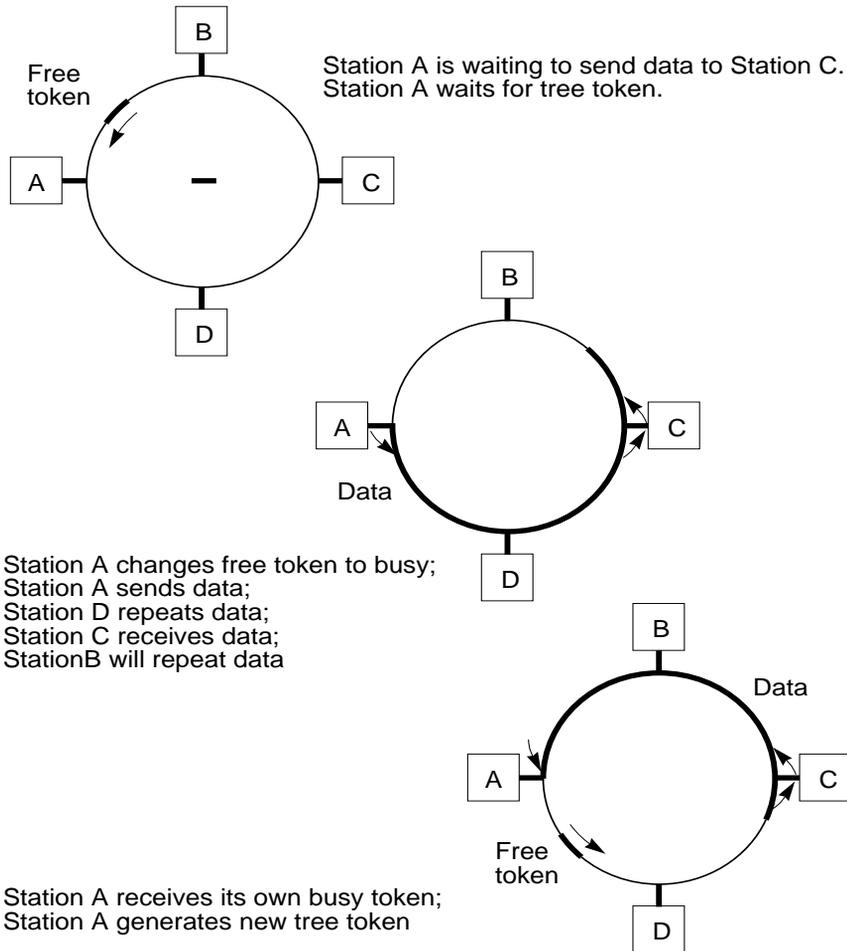
The SunTRI/P adapter supports the Solaris 2.5.1 Hardware: 4/97, 8/97, 11/97, Solaris 2.6, and Solaris 7 operating environments. The function of the SunTRI/P adapter is to provide a means for connecting Sun SPARCstations to a token ring network.

Token Ring Networks

A token ring network is a physically star-wired local area network that interconnects various devices such as personal computers and workstations into a logical ring configuration. The cabling system consists of wiring concentrators, connectors, and end stations.

The Sun token ring protocol conforms to the IEEE 802.5-1988 standard. Token ring refers to the media access control (MAC) portion of the link layer (DLC) as well as the entire physical layer (PHY). Access to the ring is controlled by a bit pattern, called a token, that circulates from station to station around the ring. The circulating token bit pattern means that any station can use the ring. Capturing the token means that a station changes the token bit pattern so that it is no longer that of a token but of a data frame. The sending station then sends its data within the information field of the frame. The frame also includes the destination address of the destination station. The frame is passed from station to station until it arrives at the proper destination. At the destination station, the frame is altered to indicate that the address was recognized and that the data was copied. The frame is then passed back to the original sending station, where the sending station checks to see that the destination station copied the data. If there is no more data to be sent, the sending station alters the frame's bit configuration so that it now functions as a free token available to another station on the ring.

If a station fails, it's physically switched out of the ring, dynamically. The ring is then automatically re-configured. When the station has been repaired, the ring is automatically re-configured to include the added station.



For a detailed description of the token ring protocol, refer to the IEEE 802.5 standard.

SunTRI/P Adapter Overview

Hardware Overview

The ASIC on the adapter is fully compliant with the IEEE 802.5 Token Ring specifications.

The token ring interface has a 9-pin D-connector socket that allows attachment to a token ring network using the IBM STP cable type 1, 2, 6 or 9 cabling system. Both the 4-Mbps and 16-Mbps data rate networks can be attached via these connectors. There is also an RJ-45 Connector for UTP cabling.

Software Overview

SunTRI/P software supports the IEEE 802.5 standards for token ring networks. The IEEE standard specifies the lower two layers of the OSI 7-layer model. The two layers are the Physical layer (layer 1) and the Data Link layer (layer 2). The Data Link layer is further divided into the Logical Link sub-layer (LLC) and the Media Access Control sub-layer (MAC).

The token ring driver is a multi-threaded, loadable, clonable, STREAMS hardware driver that supports the connectionless Data Link Provider Interface, dlpi(7p), over a token ring controller. Multiple token ring controllers installed within the system are supported by the driver.

SunTRI/P software can support different protocol architectures, concurrently, via the SNAP encapsulation technique of RFC1042. From this SNAP encapsulation, high-level applications can communicate through their different protocols over the same SunTRI/P interface. Support also exists for adding different protocol packages (not included with SunTRI/P). These protocol packages include OSI and other protocols available directly from Sun, or through 3rd party vendors. TCP/IP is implicit with the Solaris operating environment.

The software driver also provides source routing which enables the workstation access across multiple ring networks which are connected by source-routing bridges. Locally-administered addressing is also supported and aids in management of certain user/vendor-specific network configurations.

Support for the LAN Manager is provided by MAC-level firmware which complies with the IEEE 802.5 standard.

SunTRI/P Adapter Features

The SunTRI/P Adapter features:

- Automatic Ring Speed detect or manual ring speed setting
- 33 Mhz operating frequency
- 32-bit data width
- 5.0 or 3.3 volt I/O signalling

The SunTRI/P Adapter is compatible with PCI local bus specification 2.1.

Preparing for Installation



Caution – Always turn off the power to your computer before installing or removing any adapter cards.

Requirements

To install the SunTRI/P adapter, you must have a supported Sun SPARCstation with an available 33MHz or 66MHz PCI slot.

Adapter Handling Precautions

Do not remove the adapter from its protective antistatic packaging until you are ready to use it.

Before handling the adapter, discharge the static electricity from your body by touching a protective grounding device or the metal chassis of your computer. As an added precaution, wear a grounding strap.

Always handle the adapter carefully by its edges. Avoid touching any of the components or connections, particularly the gold edge-connectors.

Connecting to the Network

Note – Make sure that you choose the correct cable for the network that you are connecting to.

After installing the adapter into the computer, connect it to a working token-ring network. To do this, use either:

- STP (Shielded Twisted Pair)
- UTP (Unshielded Twisted Pair) cable.

Recommended Types of Cabling

Note – Refer to the information supplied by your hub vendor for maximum supported cable lengths.

The 9 pin 'D' type connector is compatible with IBM STP cable types 1, 2, 6, and 9. The RJ-45 connector is compatible with IBM UTP cable type 3.

The SunTRI/P adapter has the necessary media filters built in to support UTP cabling.

We recommend that you use high-quality cable such as those that meet the requirements of class D links as specified in the ISO/IEC DIS 11801 Generic Cabling for Customer Premises standard.

To be FCC class A compliant you must use the enclosed RJ45 shielded patch cable.

You can also use links that meet the specifications for class C according to the ISO/IEC DIS 11801 standard, but over restricted distances. This class of link also includes most IBM Type 3 cables.

UTP cables are also designated by categories based on the EIA/TIA-TSB-36 standard. There are five categories, with category 5 being the highest performance. The two most common cable categories in use are category 3 and category 5:

- Category 5 cable (4-pair UTP specified up to 100MHz) should meet the requirements for class D links.

Some types of cable are not suitable for token-ring UTP connection. These include:

- Silver satin (also known as flat grey).
- Residential premises wiring (four-conductor parallel cable).
- Older multipaired Public Branch Exchange (PBX) wiring.

Platform Compatibility

The SunTRI/P adapter is compatible with Solaris 2.5.1 Hardware: 4/97, 8/97, 11/97, Solaris 2.6, and Solaris 7 operating environments. TABLE 1-1 describes the platforms and number of supported cards for each adapter.

TABLE 1-1 Platform Compatibility

Platform	Number of Adapters Supported
Ultra 30	4
Ultra 60	4
Ultra E450	4
Ultra 5	2
Ultra 10	2
Ultra Exx00 Series	4

Inserting the Adapter

The following section describes how to install the SunTRI/P adapter.

Note – Before installing the SunTRI/P adapter on an E3000/4000/5000/6000 or an E3500/4500/5500/6500 system, remove the JTAG jumper J1 on the PCI I/O riser card on your system.

Card Slot Selection

When you install the SunTRI/P adapters in your machine for the first time, the instance numbers for the device driver start with 0 (for example trp0) and increase incrementally by one, for each additional adapter. The mapping between the instance number and the adapter follows the PCI slot position. The driver, trp0, will be associated with the first adapter. The driver, trp1, will be associated with the next adapter. Under Solaris 2.5.1 Hardware: 4/97, 8/97, 11/97, Solaris 2.6, or Solaris 7, if a SunTRI/P adapter is moved from its initial slot to a different slot, the system will find a device in a different slot and assume that it is a new instance of the device. The system will then assign a new instance number to the device.

Once the system has seen the device in a particular slot, the instance number remains bound to it. Whether you are installing one or multiple adapters, you can install a SunTRI/P adapter in any available PCI slot.

▼ To Install the SunTRI/P Adapter



Caution – Always turn off the power to your computer before installing or removing any adapter cards.

Read the electronic emissions statement and then follow the installation procedure.

- 1. Power off you system using the standard shut down procedures described in the *Solaris Handbook for Sun Peripherals* or your system service manual.**
The *Solaris Handbook for Sun Peripherals* is shipped with the Solaris operating environment software and is available in the online AnswerBook documentation.
- 2. Open the system unit.**
- 3. Remove the blanking plate from an empty PCI slot.**
- 4. Applying even pressure at both corners of the adapter, push the adapter into the PCI slot until it is firmly seated.**
- 5. Secure the adapter by using the retaining screws from the filler plate that you removed.**
- 6. Connect a data-cable to the appropriate connector on the adapter: the 9-pin D-type connector for STP cabling or the RJ45 connector for UTP cabling. The SunTRI/P adapter automatically detects the type of cabling you connect to the adapter.**
- 7. Connect the other end of the data-cable to a token-ring network.**
- 8. Power-up the machine.**

You can now install the SunTRI/P 4.0 software packages.

Installing the SunTRI/P 4.0 Software Packages

This chapter provides step-by-step instructions for installing the SunTRI/P 4.0 software. The package is installed using `pkgadd`, which installs the software from a mounted CD-ROM.

Software Installation and Configuration

This section describes how to install the SunTRI/P device driver.

Installation and Configuration Overview

Complete installation and configuration of SunTRI/P 4.0 software packages requires these steps:

1. Install the SunTRI/P adapter. See “Inserting the Adapter” on page 6 for installation instructions.
2. Run `pkgrm` to remove old instances of the SunTRI/P software, if necessary.
3. Install the SunTRI/P 4.0 driver software, from a mounted CD-ROM, using `pkgadd`.
4. Configure the software.
5. Reboot the system to bring it up on the token ring network.

Summary of Installation Steps

Following is a summary of the steps required to install SunTRI/P driver software. Complete procedures for these steps are given in the following sections of this chapter.

1. Verify installation requirements as specified in “Requirements” on page 4. If you did not read this section, read it now.
2. Place the CD-ROM in the CD-ROM drive.
3. Mount the CD-ROM.
4. Install the SunTRI/P 4.0 software packages using `pkgadd`.
5. Unmount and eject the CD-ROM.

▼ To Remove Previous Versions of the SunTRI/P Software

Before installing the SunTRI/P 4.0 software on your system, check your system to see if previous versions of the SunTRI/P software are installed. If older SunTRI/P software (earlier than 4.0) exists on your system, you must remove the older software and then install the new SunTRI/P 4.0 software.

1. Use the `pkginfo` command to check the system for an older SunTRI/P software package:

```
# /usr/bin/pkginfo | grep SUNWtrip
system    SUNWtrip SunTRI/P (PCI) Driver and Utilities
```

- If you do not find any SunTRI/P packages, skip to the next section, “To Mount the CD-ROM” on page 11 to continue with the software installation.
- If you do find a SunTRI/P package, you must remove it as described in Step 2.

2. As superuser (root), use the `pkgrm` command to remove the existing SunTRI/P software package.

The following is an example of removing the package `SUNWtrip`:

```
# /usr/sbin/pkgrm SUNWtrip
The following package is currently installed:
    SUNWtrip      SunTRI/P (PCI) Driver and Utilities
                  (sparc) 3.x,REV=1.x
```

▼ To Mount the CD-ROM

- **Insert the SunTRI 4.0 CD-ROM into a CD-ROM drive that is connected to your system.**
 - If your system is running Volume Manager, it should automatically mount the SunTRI 4.0 Software CD-ROM under the `/cdrom/suntri_4_0` directory.
 - If your system is not running Volume Manager, mount the CD-ROM as follows:

```
# mkdir /cdrom
# mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
```

Installing the `trp` Device Driver

Your target machine for installing the SunTRI/P software is a Sun SPARCstation using the Solaris 2.5.1 Hardware: 4/97, 8/97, 11/97, Solaris 2.6, or Solaris 7 operating environment. This section provides instructions for installing with `pkgadd`.

▼ To Install the Driver

1. **As superuser, type the following to install the SunTRI/P 4.0 software packages from the CD:**

```
# pkgadd -d /cdrom/suntri_4_0/Product/pci SUNWtrpd SUNWtrpu
```

2. **If prompted, answer `y` to permit `pkgadd` to launch the post-installation packages.**

```
This package contains scripts which will be executed with super-
user permission during the process of installing this package.
```

```
Do you want to continue with the installation of this package
[y,n,?] y
```

`pkgadd` will continue to install the software and run the post-installation script. When the installation is complete, you will see messages saying that the software has been installed successfully.

3. Use the `pkginfo` command to verify that the software has been installed.

```
# pkginfo SUNWtrpd SUNWtrpu
system SUNWtrpd      SunTRI/P (PCI) Driver
system SUNWtrpu      SunTRI/P (PCI) Man Pages and Utilities
```

4. Unmount and eject the SunTRI 4.0 CD-ROM.

```
# cd /
# umount /cdrom/suntri_4_0
# eject cdrom
```

5. Configure your SunTRI/P interface(s).

You must set the network configuration of your SunTRI/P interface(s) before you can use it. See Chapter 3 for configuration instructions.

Software Configuration

Configuring the SunTRI/P Software for TCP/IP

This section describes how to configure your token ring interface for TCP/IP. You must reboot the system for these changes to become active.

1. **At the command line, use the `grep` command to search the `/etc/path_to_inst` file for `trp` instances.**

In the following examples, the `trp` interfaces have the instance “0” and “1”, respectively.

```
# grep trp /etc/path_to_inst
"/pci@1f,4000/token-ring@2" 0 "trp"
"/pci@1f,2000/token-ring@1" 1 "trp"
```

2. Create an entry in the `/etc/hosts` file for each active `trp` interface.

The following example shows how to edit the `/etc/hosts` file to add the 192.9.200.33 IP address for the `trp0` interface and 193.1.20.1 IP address for the `trp1` interface.

```
# echo "192.9.200.33 cat-sys1" >> /etc/hosts
# echo "193.1.20.1 cat-sys2" >> /etc/hosts
# cat /etc/hosts
#
# Internet host table
#
127.0.0.1      localhost
129.144.131.18 cat-atml4      loghost
192.144.10.18 test-14 loghost
192.9.200.33 cat-sys1
193.1.20.1 cat-sys2
```

3. Set up each `hostname.trp<num>` file to correspond with each `trp` device in the `path_to_inst` file.

The following example shows the `/etc/hostname.trp0` file required for a system which has the `trp0` interface with the hostname `cat-sys1`.

```
# echo "cat-sys1" > /etc/hostname.trp0
# cat /etc/hostname.trp0
cat-sys1
```

The following example shows the `/etc/hostname.trp1` file required for a system which has the `trp1` interface with the hostname `cat-sys2`.

```
# echo "cat-sys2" > /etc/hostname.trp1
# cat /etc/hostname.trp1
cat-sys2
```

Editing `trp.conf` Parameters

This section describes the parameters that can be set in the `/kernel/drv/trp.conf` file. You must reboot the system for these setting to become active.

Setting the Maximum Transmission Unit (MTU)

Sun supports the IEEE 802.5 Token Ring Standard Maximum Transmission Unit (MTU) size of 17800 octets, (LF=0x111). All hosts should use the same MTU size on any particular network. Additionally, if different types of IEEE 802 networks are connected by transparent link layer bridges, all hosts on all of these networks should use the same MTU size.

To modify the MTU, change the MTU parameter in `/kernel/drv/trp.conf` and then reboot the system. The MTU parameter affects all SunTRI/P adapters in the system.

The maximum MTU sizes supported are: 4472 for 4 Mbps operation and 17800 for 16 Mbps operation. The default MTU size is 4472 bytes.

For example, to change the MTU to 2052, edit the MTU parameter value in the `/kernel/drv/trp.conf` file.

```
mtu=2052;
```

Configuring the Ring Speed

The ring speed is the number of mega-bits per second (Mbps) at which the adapter transmits and receives data.

The SunTRI/P software sets the ring speed to auto-detect, by default. When the workstation enters the token ring it will automatically detect the speed at which the ring is running and set itself to that ring speed. If your workstation is the first workstation on the token ring the ring speed is set by the hub. However, if your workstation is the first workstation on the token ring and the token ring has no active hubs you must set the ring speed manually. Additional workstations that join the token ring will set their ring speed automatically.

You can set the Ring Speed using the `trp#_ring_speed` parameter in the `/kernel/drv/trp.conf` file.

The `trp#_ring_speed` parameter can be set on a “per-interface” basis. For example, setting the `trp0_ring_speed` parameter affects the `trp0` adapter.

This parameter can be changed to the following settings.

- 0= auto-detect (default)
- 4= 4-Mbps
- 16= 16-Mbps

trp#_ring_speed is a parameter that affects adapter instance number. For example, the parameter trp0_ring_speed affects the adapter that ifconfig refers to as trp0.

- **To change the value of the ring speed on trp0 to 4 Mbps, and the ring speed on trp1 to 16 Mbps, change the following settings in the /kernel/drv/trp.conf file:**

```
trp0_ring_speed = 4
trp1_ring_speed = 16
```

Note – In most cases, the ring speed should be set to auto-detect. Setting the ring speed to either 4- or 16-Mbps incorrectly will result in the adapter not inserting into the ring.

Configuring the Locally Administered Address (LAA)

The LAA is part of the token ring standard specification. You may need to use an LAA for some protocols, such as DECNET or SNA. To use an LAA, create a file with execute permissions in the /etc/rcS.d directory, such as /etc/rcS.d/S20trpLAA, with the ifconfig trpX ether XX:XX:XX:XX:XX:XX command. The adapter instance is represented by trpX and the LAA for that adapter is used in place of XX:XX:XX:XX:XX:XX. For example:

```
# /sbin/sh
case "$1" in
  'start')
    echo "Configuring Token Ring LAA..."
    /sbin/ifconfig trpX ether XX:XX:XX:XX:XX:XX
    ;;
  'stop')
    echo "Stop of Token Ring LAA is not implemented."
    ;;
  *)
    echo "Usage: $0 { start | stop }"
    ;;
esac
```

For example, to use an LAA of 04:00:ab:cd:11:12 on the trp0 interface, use the following command within the /etc/rcS.d/S20trpLAA file.

```
# /sbin/ifconfig trp0 ether 04:00:ab:cd:11:12
```

Configuring the System To Use The Embedded MAC Address

Each SunTRI/P adapter has a MAC address embedded on its PROM. To use the SunTRI/P adapter's embedded MAC address, instead of the MAC address on the system's IDPROM, set the local-mac-address\? OBP property to true. You must reboot your system for these changes to become active.

- **As superuser, set the local-mac-address\? OBP property to true.**

```
# eeprom local-mac-address\?=true
```

Configuring the SunTRI/P Adapter as the Sole Network Interface

The SunTRI/P interface can be used as your sole network interface. You must reboot your system for these changes to become active.

1. **To do so, delete the existing /etc/hostname.le0 or /etc/hostname.hme0 file on your system.**

```
# rm /etc/hostname.le0
```

2. **Edit the `/etc/nodename` and `/etc/net/*/hosts` file. Change the hostnames in these files to match that of the SunTRI/P adapter.**

In the following example, the SunTRI/P adapter is named `cat-sys1`.

```
# cat /etc/nodename
cat-sys1
#
# cat /etc/net/ticlts/hosts
cat-sys1 cat-sys1
# cat /etc/net/ticots/hosts
cat-sys1 cat-sys1
# cat /etc/net/ticotsord/hosts
cat-sys1 cat-sys1
```

Configuring Internetwork Routing

Any system with two or more network interfaces (for example, Ethernet and Token Ring) may be an internetwork router. By default, the route daemon supports internetwork routing on any such system by broadcasting IP routing information regularly to all directly connected networks and hosts. If your system is not meant to be an internetwork router, you can suppress these broadcasts by creating an `/etc/notrouter` file using `touch /etc/notrouter`.

If you intend to configure the system as a router by installing a Token Ring interface alone with other LAN interfaces, make sure that the MAC addresses between `trp` and other LAN interfaces are different. To change the MAC address or to set the system to use the SunTRI/P PROM's embedded MAC address, see "Configuring the Locally Administered Address (LAA) or "Configuring the System To Use The Embedded MAC Address" on page 17.

Troubleshooting and Diagnostics

This section offers some suggestions to assist you when you are having trouble with your SunTRI/P adapter.

Troubleshooting Tools

There are several Solaris operating environment utilities useful for debugging and monitoring network traffic. These include `ifconfig`, `netstat`, `ping`, `snoop`, and `spray`. The `trpsr` and `mver` utilities are also included with the SunTRI/P packages for troubleshooting.

Ifconfig

The `ifconfig` command can be used to check the status of interfaces used for TCP/IP traffic. `ifconfig` will check whether or not the token ring interfaces were brought up.

To check all interfaces, use the following command:

```
# ifconfig -a
```

To check a particular interface, use a command similar to the following:

```
# ifconfig trp0
trp0: flags=1863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,MULTI-
BCAST> mtu 4472
    inet <IP_address> netmask <netmask> broadcast <address>
    ether <MAC address>
```

The output string:

<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,MULTI-BCAST>
indicates that the TRI/P interface is correctly installed and configured.

Netstat

After installing the `trp` driver software, you can verify that the installation was successful with the `netstat` utility. Use `netstat (8C)` to check for the `trp` interfaces. For example:

```
# netstat -i
```

Ping

`ping` is useful for checking to see if another host is alive, or if communication between two machines is possible. Note that the token ring interface hostname must be used. An example usage is:

```
# ping sake-trp
```

When attempting to get two machines to communicate, first ensure that the hardware and software are configured correctly by using `ifconfig`. Then, attempt to `ping` one system from the other and vice versa. If the `ping` succeeds, the two stations have been configured correctly and are communicating. `ping` might not work with some vendors' IP implementations.

Snoop

The `snoop` utility allows you to monitor network traffic. The `snoop` utility monitors token ring activity by decoding certain token ring fields, such as the source routing information field. For example, to run `snoop` on the `trp0` interface, use the following command:

```
# snoop -d trp0
```

Refer to your system documentation, or the `snoop` man page for a more detailed explanation about the `snoop` utility.

Spray

`spray` is a UCB-developed utility that only works between machines running the `sprayd` process. Do not `spray` a non-Sun system.

Trpsr

The source route table examination utility is used to examine the Source Route Table.

The format for executing this utility is: `trpsr ppa -a | -d`. The `-a` option displays all of the routes in the table. The `-d` option deletes the source route table entry for the specified `mac_addr`.

The source routing table gets updated every 1 minute. This is only effective when IP periodically probes the route table, based on the `ip_ire_flush_interval` (default is 20 minutes). This parameter is configurable in the IP driver through `ndd(1M)`. This is applicable if the network topology is changed.

The following command displays all routes for the `trp0` interface.

```
# trpsr 0 -a
Mac-Addr      : Route
-----      : -----
8:0:20:c:10:ab : 00000000000000000000000000000000
8:0:20:a:8a:17 : 00000000000000000000000000000000
8:0:20:11:e3:7 : 00000000000000000000000000000000
```

The following command deletes a `mac_addr` from the `trp0` interface's source route table.

```
# trpsr 0 -d 08:00:20:90:90:79
Success
```

Mver

This utility allows you to determine the version of the `trp` driver that is running.

- **As superuser, type the following command to display the driver version:**

```
# mver /kernel/drv/trp
```

Maintenance Hints

Occasionally you may find it useful to disconnect a SunTRI/P adapter from the network. To do so, use the `ifconfig trpX down` command, where `X` is the `trp` adapter instance, while logged in to your system as root. In the following example, the user is trying to bring down the `trp0` interface.

```
# ifconfig trp0 down
```

A subsequent `ifconfig trp up` command restarts the firmware.

```
# ifconfig trp0 up
```

System Console Messages

Messages associated with token ring operation are displayed on the host system's console.

Error Messages

trp0: error 12 INSUFFICIENT MEMORY

There is insufficient memory available to allocate transfer buffers. Try reducing the maximum frame size and re-booting your system.

trp0: error 35/21 ADAPTER_INIT_FAILED

The adapter hardware failed to initialize. Please contact your Sun service provider.

trp0: error 35/33 BUD_FAILED

The adapter hardware failed to initialize. Please contact your Sun service provider.

trp0: error 35/22 ADAPTER_TIMEOUT

The adapter took too long to initialize. Ensure that all token-ring lobe cables are properly attached and re-boot your system. If the problem persists please contact your Sun service provider.

trp0: error 35/25 OPEN_FAILED

The adapter failed to open onto the Token Ring. The most likely cause is that the lobe cable is not correctly attached to the adapter or to the hub. Other possible causes are that the driver is set to the wrong ring speed or that the driver is using an alternative address that is already in use on the Token Ring.

WARNING: trp0: lobe cable fault - please check cabling

The lobe cable has been removed from the adapter.

WARNING: trp0: ring is beaconing

The Token Ring has entered beacon state. This is part of an error recovery procedure. If you see this message repeatedly over a short period of time, check the cable and hub connections. If the error continues, then contact your network administrator.

Note – This warning often occurs when the lobe cable is removed.

WARNING: trp0: adapter is single station

Your system is now the only system connected to the Token Ring.

WARNING: trp0: ring has entered recovery mode

The Token Ring has entered recovery. This is part of an error recovery procedure. If you see this message repeatedly over a short period of time, check the cable and hub connections. If the error continues, then contact your network administrator.

Note – This warning often occurs when the lobe cable is removed.

ip_rput: DL_ERROR_ACK for 29, errno 1, unix 0

The token ring driver received an Ethernet encapsulated packet on the streams fast path. This error can be ignored.

```
ip: joining multicasts failed on tr0 - will use link layer
broadcasts for multicast
```

The token ring driver does not support Ethernet-type multicasts. This error can be ignored.

If any other errors are displayed by the trp device, carefully note any error numbers and text and contact your Sun service provider.

Advanced Topics for Programmers

This appendix covers advanced topics that are not needed to install the token ring software. The topics covered here are intended to bridge the gap if you are not familiar with UNIX and/or 802.5. Specific aspects of the Sun implementation of the IEEE 802.5 Token Ring standard over the Solaris operating system are discussed. The areas covered include LLC Packet Encapsulation, MAC Frame Formats, Source Route Support, and the Stages of the Open command.

LLC Packet Encapsulation

This information is a summary of the standard method of encapsulating IP packets on IEEE 802 networks. A complete reference is given in RFC 1042. IP packets are sent on IEEE 802 networks encapsulated within the 802.2 LLC and SNAP data link layer headers. The standard release supports LLC type 1 communication. LLC I support is provided from within the driver and no additional software must be registered above the driver. If an LLC type 2 is registered above the driver, the LLC 1 modifications are omitted and the upper layer LLC II deals with LLC level packet encapsulation issues. Trailer encapsulation is not supported.

If no LLC II is registered above the token ring driver, packets arrive at the device driver in the same format as ethernet packets. See the following example:

DA	SA	Type	Data
-----------	-----------	-------------	-------------

The token ring driver interface modifies the packets to conform to the RFC 1042 standard by inserting the 802.2 LLC and SNAP headers. It also completes most of the MAC header. The format of packets sent to the TMS380 chipset is shown below:

AC	FC	DA	SA	RI*	LCC_SNAP_HDR	Data
-----------	-----------	-----------	-----------	------------	---------------------	-------------

* RI field is optional.

A more detailed inspection of the LLC and SNAP headers is shown below:

DSAP	SSAP	Control	Org Code	
0xaa	0xaa	0x3 (UI)	0x0	

The DSAP and SSAP values are set to 0xaa, the assigned global SAP for SNAP. The control value is 0x3, representing Unnumbered Information Frames. If there is no LLC II registered above the driver and the control field is not equal to Unnumbered Information, the received packet must be either a TEST or XID packet with a DSAP and SSAP of 0x0; Otherwise, the packet is dropped. The EtherType field is the value of the ethernet type field.

MAC Frame Headers

The SunTRI/P driver (`trp`) completes most of the MAC header including the AC, FC, Source and Destination addresses, and the optional Routing information field. The entire frame sent to the chipset is shown below:

AC	FC	DA	SA	RI*	Data
-----------	-----------	-----------	-----------	------------	-------------

* RI field is optional.

The AC, or Access Control field, is set to 0x10 indicating that it is part of a frame. The priority portion is set to zero. Multiple priority levels are not currently supported through a programmatic interface. The FC, or Frame Control field, is set to 0x40 indicating that this is an LLC frame as opposed to a MAC frame.

Sending neither a MAC nor LLC packet type is currently undefined in the 802.5 specification, and the value of 1x for the frame type bits of the FC is reserved for future use. Therefore, using this out-of-hand mechanism may result in fundamental incompatibilities with the 802.5 standard.

Source Route Support

A non-source route host can coexist on a large network with source route stations, and can communicate to the local stations on the ring, but cannot communicate through the source route bridge.

A more detailed inspection of the routing information field is shown below:

RT	LTH	D	LF	r	RD1	RD2	...RDn
-----------	------------	----------	-----------	----------	------------	------------	---------------

The Route Type fields initiated on output are: Specifically Routed (RT=0xx), all Routes Explorer (RT=100), and Spanning Tree Explorer (RT=110). The All-Routes Explorer is used with ARP, TEST, and XID packets to discover a route. Spanning Tree broadcasts are used for datagrams with an IP broadcast address. This is necessary to prevent multiple copies being received by a destination. All route types including Spanning Tree Routed (RT=111) are accepted on input for backwards compatibility.

The Length value (LTH) indicates the length in octets of the Routing field.

The Direction field (D=X) indicates, to a bridge, whether a frame is traveling from the originating station to the target or vice versa.

The Largest Frame (LF) bits specify the largest frame size of the data field. Sun supports the IEEE 802.5 Token Ring Standard Maximum Transmission Unit (MTU) size of 17800 octets, (LF=0x111). All hosts should use the same MTU size on any particular network. Additionally, if different types of IEEE 802 networks are connected by transparent link layer bridges, all hosts on all of these networks should use the same MTU size.

To Modify the MTU:

- Change the MTU value in `/kernel/drv/trp.conf`.
- Reboot.

The maximum mtu sizes supported are: 4472 for 4 Mbps operation and 17800 for 16 Mbps operation.

The next four bits (r) are reserved.

The RDn fields are 16 bit fields that indicate the path or route between nodes on different physical rings.

Initially, an All-Routes Broadcast packet is sent out to determine a route. An entry is created in the Source Route Table indicating that a route is pending. All packets to be transmitted during the period that an entry is pending are sent out immediately without a Source Route Header. If an entry does not complete within three minutes, it is purged. If a packet is received, the entry is completed. The rule is that if the packet is on the local ring, RI length field of 2 (LTH=2), then all further packets sent to that particular station are tagged as non-source route. Otherwise, the routing information is cached into the Source Route Table. Currently, there is no programmatic interface to the Source Route Table, thus the only way to purge the table is to unplumb the interface and then plumb the interface back up again. However, individual entries can be deleted by using the `trpsr` command.

OPEN Stages

On a ring topology many functions must be performed before a station is opened on the ring and is fully operational and ready for data transfer. When you attempt to bring the SunTRI/P adapter up through the `ifconfig` command, there is a short initialization command, and then an `OPEN` command.

There are five stages to the `OPEN` command. TABLE A-1 describes those stages.

TABLE A-1 OPEN stages

Stage	Description
Lobe Media Test	This is a 1-2 second test that sends packets in loopback mode. This tests the receivers and the cabling up to the MSAU.
Physical Insertion	This is the time required to charge up a capacitor and trip a relay on a passive wiring concentrator (MSAU). This can take a few seconds.
Address Verification Participation in Ring Poll Initialization	The last three stages can take up to 18 seconds to complete, when the first station attempts to Request Initialization onto the ring. Active Monitor Present MAC frames are sent every seven seconds (by the active monitor).

If there is any problem opening onto the ring due to an error message indicating that the OPEN command failed, the existence of the error message itself indicates that the software was probably installed correctly but the failure is more likely due to a hardware problem. Recheck all hardware connections including the PCI connection and both ends of the token ring cable. Note that even if everything is installed correctly and there are multiple stations on the ring, the ifconfig command may take approximately 10 seconds to complete.

Token Ring Architecture Overview

This appendix provides a general overview of the token ring architecture.

The International Standards Organization (ISO) has established a seven-layer model that provides a standard reference for intercommunication of computer systems through a network by using common protocols. The seven-layer model has become an international standard.

Each layer of the model defines a specific section of the total architecture. The layers of the Open Systems Interconnection (OSI) are, from top to bottom:

TABLE B-1 Seven Layers of Open Systems Interconnection (OSI)

Layer	Description
Application	Gives the user access to all the lower layer functions and deals with the exchanges between user applications.
Presentation	Represents user data including required translations and code conversions.
Session	Provides the mechanisms for organizing and structuring the interactions between applications and devices.
Transport	Provides transparent end-to-end data transfer relying on lower-level functions for handling any peculiarities of the actual transfer medium.
Network	Contains the definitions that provide the means for establishing connections between networks. The definitions include procedures for the operational control of inter-network communications and the routing of information through multiple networks.
Data Link	Provides the functional and procedural means for transferring data between network entities as well as the means for detecting and correcting errors that occur on the network.
Physical	Provides the mechanical, electrical, and functional and procedural standards for accessing the physical medium.

The Institute of Electrical and Electronic Engineers Computer Society has also drafted standards for local area networks. A reference model similar to that of the OSI model was created. Two layers in the IEEE model correspond to the two lower levels of the OSI model. In the IEEE model, however, the Data Link layer is further divided into two sub-layers: the Logical Link Control (LLC) sub-layer and the Medium Access Control (MAC) sub-layer.

Various media access control techniques have been created and have become standards. Among them is the Token Ring Standard - IEEE 802.5.

The Data Link Control Layer

Standards for local area networks are based on the two lower layers, Data Link Control and Physical Control. As a result, token ring networks can co-exist with SNA systems. In an SNA mode, the Data Link Control layer is replaced by the LAN Data Link Control layer (which is further divided into the Logical Link Control and Medium Access Control sub-layers) and the LAN Physical Control layer.

Each IBM node in an IBM token ring network contains a Data Link Control layer called DLC.LAN. DLC.LAN consists of a manager function called DLC.LAN.MGR, a Logical Link Control sub-layer which includes one or more link stations, an access channel control, and a user datagram service and a Medium Access Control (MAC) sub-layer which includes one medium access control.

A link is a logical connection between two link stations providing data transfer between two nodes. A node is either of the endpoints of the link. A link station is a protocol machine that manages the elements of procedures required for data exchange and schedules data transfer over the link.

DLC.LAN.MGR

DLC.LAN.MGR includes both the LLC and MAC sub-layers. It also provides management functions for both. DLC.LAN.MGR supervises the operation of the DLC.LAN and directs the flow of information through the LLC and MAC sub-layers. It also controls link activation and the attachment of ring stations to the ring and removal of ring stations from the ring.

DLC.LAN.MGR acts as the interface between DLC.LAN and the Physical Unit. It handles the records from the Physical Unit that require action by the DLC.LAN.MGR. It also passes other records from the Physical Unit to the appropriate link stations.

The Logical Link (LLC) Sub-layer

The LLC sub-layer provides sequential, connection-oriented data transfer and non-sequential, connectionless data transfer. The link stations provide sequential connection-oriented data transfer and error recovery for one or more links.

The Path Control layer passes basic transmission units (BTUs) to the link stations which in turn add appropriate control information. On each link, the local and remote stations use the HDLC asynchronous balanced mode of operation to keep the LLC Protocol Data Units (LPDUs) in sequence, and to detect and correct LPDUs that are out of sequence. Link stations pass the BTU and control field to the access channel control which multiplexes message units flowing between link stations and DLC.LAN.MGR and the MAC sub-layer. The access channel control builds LPDUs from the information in its internal table for that link station. It then transmits the LPDUs to the MAC sub-layer. The access channel control also routes the LPDUs it receives from the MAC sub-layer to the appropriate link station or to the DLC.LAN.MGR.

Connectionless data transfer is provided by the user datagram service where data is sent and received without any correlation to previous or subsequent data and without the need for establishment of a data link connection.

A higher-layer protocol passes data to be transferred and the source and destination addressing information to the user datagram service which adds appropriate control information. The user datagram service passes the addressing information, control field, and data to the access channel control. It does not provide acknowledgment of data or flow control or error recovery procedures.

The Medium Access Control (MAC) Sub-layer

The MAC sub-layer controls the routing of information between the Physical layer and the LLC sub-layer. The MAC sub-layer provides:

- Address Recognition - For initiating copying of a frame based on the destination address in the physical header. Each ring station must be able to recognize its own MAC address and an all-stations broadcast address. Additionally, ring stations should be able to recognize one or more group addresses.
- Frame Copying - For copying a frame off the ring.
- Frame Control Recognition - For determining the frame format and the type.
- Delimiting Frames - For delimiting the start and end of a frame.

- Frame Status Generation and Verification - For providing and verifying additional information in each frame to detect transmission errors.
- Priority Management - For gaining access to the transmission medium based on priority.
- Routing - For determining which function in the node should process the frame.
- Timing - For providing timers required by the token and MAC management protocols.
- Token Management - For gaining access to the physical transmission medium including appropriate supervision protocols for handling errors.

The Physical Layer

Each node in a token ring network contains a physical layer. The Physical layer provides attachment to the transmission medium and contains the cable and the circuit switches used to configure the equipment.

The primary function of the Physical layer is to encode, transmit, and recognize bits, code violations, and signal losses.

The Physical layer also provides master timing and voltage control functions.

Basic Operating Concepts

This section presents an overview of the terms and procedures that are basic to the operation of token ring networks. Details are given at the MAC sub-layer level because this level provides the protocols unique to token ring operation.

The Ring

In a token ring network, a ring consists of ring stations and the transmission medium that interconnects them. A ring station consists of functions that allow a device to connect to the ring and use the link access protocols. See Chapter 1, for diagrams illustrate a simple ring configurations.

A ring station transfers data in the form of tokens and frames to the ring where the data travels sequentially from ring station to ring station. The tokens and frames are identified by unique formats that begin with a Starting Delimiter and end with an Ending Delimiter. The actual flow path has a definite direction that is determined by

the physical location of the ring stations on the ring. Each ring station repeats the data it receives, checking it for errors, and copying it if it is the destination (the intended recipient). When the data returns to the originating ring station, the station removes it from the ring.

The protocols provide error detection and recovery functions, such as acknowledgment, beaconing, retransmission, and timeouts.

Each station can serve one or more attached devices, such as terminals or printers, and allow them to communicate with other attached devices on the ring.

A delay of no less than 24 bits is required on the ring to allow the 24-bit token to circle the ring successfully. The total delay on the ring consists of the ring station delays plus the propagation delay caused by the transmission medium.

The token ring network requires addressing functions so that communication between any two rings can be uniquely identified. Addressing is independent of the physical configuration.

The token ring network also requires data checking functions to preserve the integrity of its access control and ring management transmissions.

Multiple Ring Connections

The token ring architecture supports multiple ring connections using bridges, through which data can be transmitted from a ring station on one ring to a ring station on any other ring.

The bridge also acts as a ring station on any ring to which it is attached. The bridge copies frames destined for other rings and transmits frames from other rings that are destined for the local ring, or for rings beyond it. The destination is based on addressing information. The routing information field of the frame determines whether or not the bridge copies the frame.

MAC Sub-layer Operating Modes

This section describes the normal token protocols used to transmit a frame on a token ring network.

- *Normal Repeat Mode*—In this mode, a ring station checks the data in the tokens and frames it receives and sets the Error-Detected, Address-Recognized and Frame-Copied bits, as appropriate, as it repeats the token or frame.
- *Transmit Pending Mode*—This is the mode in which a ring station with data to transmit waits for a free (available) token.

A ring station may receive a token of usable priority and change it to a frame and append data and control information to it. At this point, the station enters the Normal Transmit Mode and transmits the information. If a ring station discovers that it has incorrectly identified a frame as a token, it transmits an Abort sequence and returns to the Transmit Pending Mode. After a station has sent data, it starts its physical trailer (the ending delimiter and frame status field) timer. If the ring station does not receive its physical header by the time it transmits the frame status field, it transmits idles and waits for the physical header to return. If the physical header timer expires and the physical header has not returned, the ring station returns to the Normal Repeat Mode without originating a token.

When the physical header returns, the ring station compares the returned Source Address with the Source Address it transmitted. If the two addresses are identical, the ring station transmits a token followed by idles. Otherwise, the ring station continues to wait for its transmitted physical header. The ring station transmits idles until it receives its transmitted physical trailer and has completely removed its frame from the ring. It then returns to Normal Repeat Mode. If the physical header is never returned, the ring station transmits idles until its physical trailer timer expires, at which time it enters the Normal Repeat Mode.

The Active Monitor

The Active Monitor function is part of the architecture of every ring station. One ring station on each ring is designated as the Active Monitor. It provides token monitoring as well as various other functions. Any operating ring station can be assigned as the Active Monitor. Other ring stations act as standby monitors in case the currently designated Active Monitor fails.

The Active Monitor resolves certain error conditions on the ring such as: lost tokens, frames and priority tokens that circle the ring more than once, more than one Active Monitor on the same ring at the same time, a short ring condition in which the bit delay is too short to hold a token.

Maintaining the Master Clock

The Active Monitor maintains the ring's master clock which controls timing and ensures that all other clocks on the ring are synchronized.

Ensuring Proper Ring Delay

The Active Monitor ensures that a token can be completely transmitted before returning to the originating station. It ensures sufficient delay by inserting a 24-bit delay (the length of a token) into the ring.

Initiating Neighbor Notification

The Active Monitor periodically broadcasts the Active Monitor Present MAC frame to all ring stations on its ring. Each ring station then acquires the address of its Nearest Active Upstream Neighbor (NAUN)

Neighbor Notification begins when the Active Monitor broadcasts an Active Monitor Present MAC frame to all the ring stations on its ring and resets the Neighbor Notification Complete flag. The first ring station that receives the Active Monitor Present MAC frame copies it and sets Address Recognized (A bit) and Frame Copied (C bit). The ring station then saves the source address as its NAUN address and transmits a Standby Monitor Present MAC frame with the A and C bits set, to all other ring stations on the ring.

The next ring station downstream disregards the Active Monitor Present MAC frame because the A and C bits are already set. The ring station then copies its NAUN address from the Standby Monitor Present MAC frame, sets the A and C bits, and transmits its own Standby Monitor Present MAC frame.

In this manner, Neighbor Notification proceeds around the ring with each ring station receiving and transmitting Standby Monitor Present MAC frames until the Active Monitor copies the last Standby Monitor Present MAC frame whose the A and C bits are set. The Active Monitor then sets the Neighbor Notification Complete flag which indicates that the process has been successfully completed.

Monitoring Neighbor Notification

Any time during the Neighbor Notification cycle an exception event could occur and impede the notification process as it makes its way around the ring.

The Neighbor Notification timer has a long time-out. When it expires, the Active Monitor transmits an Active Monitor Present MAC frame. If the previous Neighbor Notification is not complete, the Active Monitor also transmits a Report Neighbor Notification Incomplete MAC frame to the Ring Error Monitor.

The Active Monitor uses its receive notification timer to ensure that during Neighbor Notification each Active Monitor Present MAC frame circles the entire ring. If the timer expires, the Active Monitor Present MAC frame did not circle the ring and the Active Monitor initiates token claiming. Token claiming is a process that allows a ring station to elect an Active Monitor.

If the Active Monitor receives a Standby Monitor Present MAC frame after Neighbor Notification has successfully completed, the frame is ignored.

If the Active Monitor copies an Active Monitor Present MAC frame that contains a source address other than its own, it means that an Active Monitor is already on the ring and that the Active Monitor that copied the frame can now deactivate its Active Monitor functions.

A hard error can cause the suspension of the token protocol and as a result cause the suspension of the Neighbor Notification process. The Active Monitor transmits a Report Neighbor Notification Incomplete MAC frame to the Ring Error Monitor. If the ring is busy transmitting data, especially priority traffic,

Neighbor Notification could be delayed or interrupted. The Active Monitor transmits a Report Neighbor Notification Incomplete MAC frame to the Ring Error Monitor.

Monitoring Token and Frame Transmission

The monitor bit in the access control field is reset in every newly transmitted token or frame. When the Active Monitor repeats a frame or non-zero priority token, it sets the monitor bit. If the bit is already set, the Active Monitor assumes that the token or frame has already circled the ring. The Active Monitor then purges the ring and releases a new token.

Detecting Lost Tokens and Frames

The Any Token Timer has a short time-out but it still exceeds the time for the longest possible frame to circle the ring. The Active Monitor restarts this timer whenever it repeats a Starting Delimiter. If this timer expires, the Active Monitor assumes that the token or frame was lost on the ring. The Active Monitor purges the ring and releases a new token.

Purging the Ring

The Active Monitor broadcasts a Ring Purge MAC frame to all ring stations on the ring before originating a new token. Receipt of the returned frame indicates to the Active Monitor that a frame can circle the ring without incident. The Active Monitor then restarts the token protocol.

The Ring Purge MAC frame resets ring stations to Normal Repeat Mode and cancels or restarts appropriate timers. The Active Monitor ignores the value of the monitor bit in returning Ring Purge MAC frames. The Active Monitor copies the reservation bits of the last returning Ring Purge MAC frame into the new token it originates.

An Active Monitor that is interrupted by a Ring Purge MAC frame that it did not originate, assumes itself to be a duplicate monitor and becomes a standby monitor. (It removes its 24-bit delay and master clock functions.)

Attaching the Ring

This section describes the various phases a ring station must successfully complete to attach and function on a token ring network. The only assumption about the ring station is that it has a pre-assigned address.

Phase 0: Lobe Test and Physical Attachment

Lobe testing is part of the initialization process that is performed before the ring station has attached to the ring. A lobe is the section of cable that attaches a device to an access unit. The access unit allows the devices to access the ring from a central point.

A lobe test consists of sending a series of Lobe Test MAC frames only on the ring station's lobe to ensure that there is no fault on the lobe itself. If the frame passes through the lobe successfully, the ring station attaches to the ring. Otherwise, the ring terminates the attachment process with an error. There is no wait for a token because the Lobe Test MAC frames are only sent on the ring station's lobe.

After the power-on diagnostics and initialization, a device containing a ring station is physically attached to the ring and proceeds to Phase 1.

Phase 1: Monitor Check

The ring station starts an attach timer and the following events occur:

If the ring station receives an Active Monitor Present, Standby Monitor Present, or Ring Purge MAC frame before the timer expires, it assumes that an Active Monitor is present on the ring and proceeds to Phase 2.

If the ring station does not receive any of the above mentioned frames before the timer expires, either it is the first station on the ring or no Active Monitor exists on the ring. In either case, the ring station initiates token claiming.

A ring station does not activate its standby monitor functions until it has attached to the ring.

Phase 2: Duplicate Address Check

In this phase the ring station checks for the presence, on its ring or another ring station, of the same address. The check is performed by sending a Duplicate Address Test MAC frame. If a duplicate address is found, the ring station removes itself from the ring.

Phase 3: Participate in Neighbor Notification

The ring station participates in Neighbor Notification, learns its NAUN, and identifies itself to its nearest downstream neighbor.

If the ring station detects beaconing or the need for beaconing before it has completed Neighbor Notification, it removes itself from the ring and indicates the need for beaconing (described under “Hard Error Detection and Reporting”). After Neighbor Notification is complete, the detection of beaconing, or the need for beaconing, causes the ring station to begin beaconing as though it had completed attaching to the ring. Attachment continues after the ring recovers.

Phase 4: Request Parameters

To maintain consistency across the network, various ring parameters must be used by all ring stations. Changeable parameters include: Physical Location, Soft Error Report Timer Value, and Ring Number. The Ring Parameter Server acts as focus to ensure that ring stations on the ring have the same Soft Error Report Timer and Ring Numbers.

The ring station requests changed operational parameters from a Ring Parameter Server, if one is present. If one is not present, default values are used.

If this information is incorrect or threatens ring integrity, the Ring Parameter Server can notify the LAN manager, which can issue a Remove Ring Station MAC frame, forcing the ring station to remove itself from the ring.

Soft Error Detection and Reporting

Soft errors are usually intermittent faults that temporarily disrupt normal operation of the token ring network. Soft errors are normally tolerated by error recovery procedures. Most soft errors are indicated by architectural inconsistencies, such as cyclic redundancy checks or time-outs, in received or repeated frames or by a ring station's inability to process received frames. If soft errors result in degradation of the ring's performance, the LAN manager can reconfigure the ring to bypass the faulty node.

Each ring station maintains a soft error report timer and a set of counters to measure how often the most critical errors occur. When the soft error report timer expires, the ring station transmits a Report Soft Error MAC frame to the Ring Error Monitor. After reporting the soft error, the ring station resets the error counters.

The Report Soft Error MAC frame reports the number of errors detected since the previous report. The report identifies the transmitting ring station's address and NAUN. It includes all the errors and their values.

Hard Error Detection and Reporting

Hard errors are usually permanent faults in equipment that cause the ring to stop functioning with the normal token ring network protocols. A ring station downstream from the hard fault recognizes a hard error at the receiver side of its attachment. The ring must be reconfigured to bypass the error. The reconfiguration restores the ring to the normal operational state. Repairs might be required to restore the ring to full operation.

When a ring station detects a hard error, it transmits a Beacon MAC frame, with an All-Stations Address to its ring only, until either its input signal is restored or until it removes itself from the ring. This is called beconing. All other ring stations that receive the Beacon MAC frame enter the Beacon Repeat Mode.

The Beacon MAC frame identifies the beaconing ring station's NAUN and the type of error detected. When the beaconing ring station's NAUN has copied eight Beacon MAC frames, the NAUN removes itself from the ring. It then tests itself and its lobe using Duplicate Address Test MAC frames. If the test is successful, the NAUN re-attaches itself to the ring.

If the ring does not recover after the NAUN tests itself, the beaconing ring station removes itself from the ring and follows the same test sequence described above, for its NAUN.

If the ring still does not recover after both the beaconing station and its NAUN have tested, the error requires manual intervention.

If the ring still does not recover after both the beaconing station and its NAUN have tested, the error requires manual intervention.

Addressing and Routing Concepts

Each transmitted frame on a token ring network has Destination Address and a Source Address. Their purposes are to identify the receiving and sending ring stations, respectively.

In a multiple-ring environment in which two or more rings are connected by bridges, the destination addresses for frames originating in one ring station and directed to one or more ring stations on other rings, can be created using a source routing technique. The route is determined at session connection time when an initiating ring starts the process of collecting routing information that will be included in subsequent frames. This method eliminates the need to maintain routing tables at predefined nodes on the ring.

Source Routing

Source routing is the method used within the token ring network architecture to route frames through a multiple-ring local area network. A route is the path taken by a frame as it travels through a network from the originating station to a destination station. Source routing does not require centralized routing tables because each frame carries information about the route it should follow. The routing information is obtained using a search process that originates at the source ring station and fans out through the network, to the destination ring station. The search is initiated by executing an All-Rings Broadcast, (discussed later in this section).

Routing information is acquired using the TEST, or XID, LPDU commands. The originating station sends a TEST, or XID, command LPDU to all rings on the network. The command frame passes through all the interconnected rings by way of bridge stations, searching for the Destination Address. During the search process, the command frame accumulates routing information as it passes through the bridges.

As the LPDU passes through the multiple-ring network, copies are created, and all continue to search for the Destination Address. If more than one route to the Destination Address exists, then more than one LPDU will reach the destination ring station. As the destination ring station receives each LPDU, it returns the acquired routing information to the originating ring station in a TEST, or XID, response LPDU which follows the original route in a reverse manner.

If more than one route to the Destination Address is found, all of them are returned to the originating ring station which then chooses a route. The originating ring station learns the preferred route when it receives the first response from the destination ring station. The originating ring station then uses the preferred route for any subsequent transmissions to the destination ring station.

In addition to the Destination Address and the Source Address carried in each frame sent in a multiple-ring network, an ordered list of the ring and bridge numbers, through which the frame must pass to reach the Destination Address, is sent. Each bridge is responsible for adding this routing information to the routing information filed in All-Rings Broadcast frames. Each bridge must also check the routing information in non-broadcast frames and make forwarding decisions based upon what it finds.

Addresses

The token ring network associates a ring station, or a group of ring stations, with a unique MAC Sub-Layer Address. This enables any ring station to attach to the token ring network.

Individual and Group Addresses

An Individual Address Identifies a particular ring station on the token ring network.

A Group Address identifies a group of destination ring stations on the token ring network.

Universal and Local Address Administration

Universal administration means that all individual addresses are assigned, administered, and guaranteed by the IEEE to be unique across all local area networks by the IEEE. This method eliminates customer involvement in the administration of individual addresses which in turn eliminates the need for site address administrators and address administration programs.

Local administration means that all individual addresses are administered by someone other than the IEEE. Local addresses must still be unique within the token ring network where they occur.

Null Addresses

An address of all zeroes is considered a Null Address. A frame with a Null Address is not addressed to any ring station. It can be sent but not received. When a frame with a Null Address returns to its originating ring station, its data is stripped and a new token is issued.

All-Stations Broadcast Addresses

An All-Stations Broadcast Address means the frame is sent to all stations on a given ring, or interconnected rings. All stations must be able to recognize All-Stations Broadcast Address. Whether or not the frame leaves the source ring, is determined by the routing information in the frame.

Functional Addresses

IBM token ring network architecture provides bit-specific Functional Addresses for various uses. Ring stations use Functional Address masks to identify these functions. The masks consist of the sum of the bits for the functions for a particular ring station.

All Functional Addresses are locally-administered Group Addresses. The architecture allows up to 31 functional addresses to be specified. Refer to IBM bridge documentation for actual address definitions.

Broadcast Terminology

To broadcast a frame is to send it to more than one ring station, or to more than one ring. A frame is also broadcast when a specific Destination Address is known but the location and route to the destination is not.

All-Rings Broadcast is the sending of a frame to all interconnected rings on a local area network. All-Rings Broadcast is indicated by frame with the broadcast bit in the routing information set to "1". The address of particular ring stations on those rings is determined by the Destination Address field.

An All-Rings Broadcast is independent of an All-Stations Broadcast which is the sending of a frame to all stations of a ring.

Within the bridge concept, there is also a Limited Broadcast Frame which refers to the forwarding of specially designed broadcast frames by bridges that have limited broadcast ability enabled.

Network Configuration Planning

Planning your network configuration is an important prerequisite to installing SunTRI/P. Read this appendix thoroughly and plan your configuration before you continue the installation

The performance of your token ring network depends not just on the SunTRI/P adapter itself, but on the overall network configuration and individual system administration. Much of this appendix describes how to configure your network to maximize performance. Topics include:

- Token Ring Network Topologies
- Establishing an all-token ring network
- Incorporating token ring into an existing Ethernet network
- Setting up token ring-to-Ethernet routers
- Improving distributed applications
- Selecting the SunTRI/P adapter's MAC address

Token Ring Network Topologies

Factors that influence the selection of a network topology for a particular installation, include:

- Number of Stations
- Type of Applications
- Affinity Groups
- Backup Paths
- Performance
- Capacity
- Reliability
- Availability
- Expected Growth

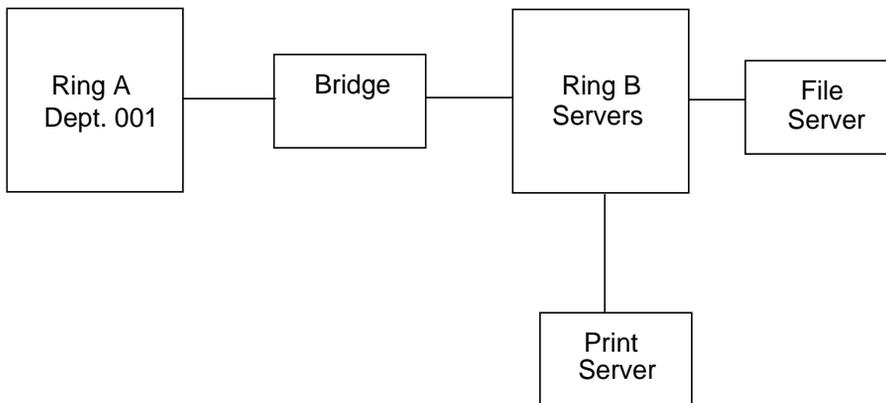
General Guidelines

When the number of workstations is small, less than 40, the distribution of workstations is determined primarily by physical considerations. Selecting a particular physical layout is determined by the locations of various departmental groups, or groups, of users performing related tasks.

The first decision to be made concerns the number of rings to have on the token ring network. The answer depends upon various factors that are examined in the following discussion.

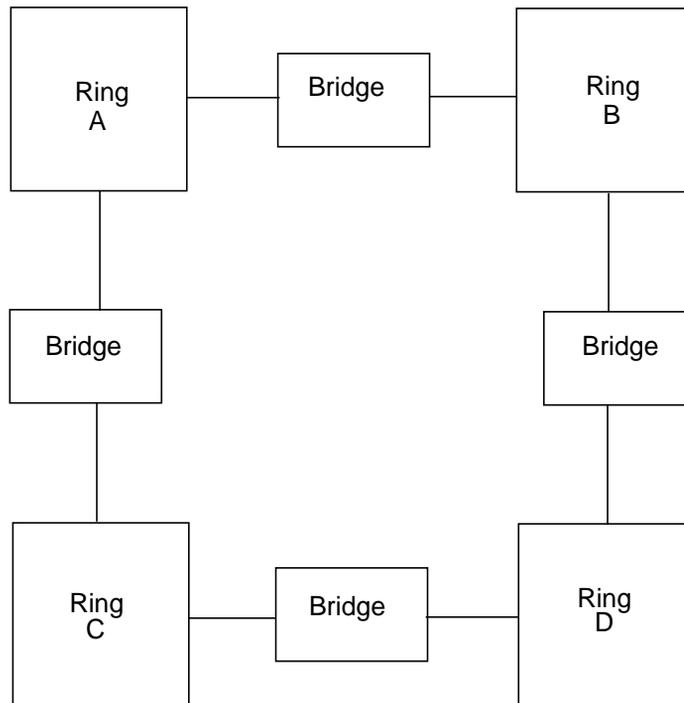
Performance

Depending upon the applications running on each workstation, the performance might be improved by using a multiple-ring network. Performance can often be improved by choosing a two-ring network with the workstation on one ring and the servers on the second ring, as illustrated in the following diagram.



Reliability

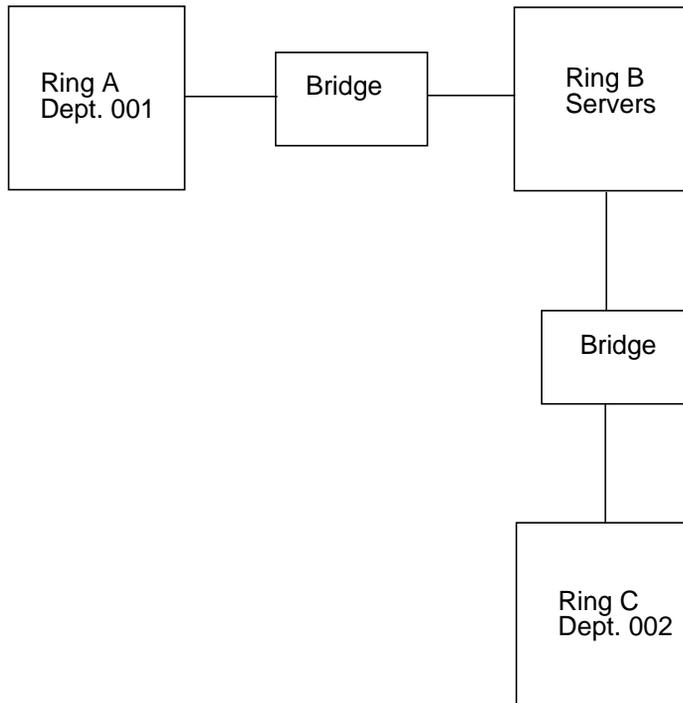
A network configuration of multiple rings might permit easier fault isolation. Within a multiple-ring network, a transient beaconing condition on one ring does not preclude the use of workstations on another ring, though the beaconing ring is inoperable for data transfer. For example, in the following diagram, a four-ring configuration is illustrated. If ring C is beaconing, workstations attached to rings A, B, and D can continue exchanging information with no problem. (Beaconing is sent from a ring station indicating that a serious ring problem exists.)



Affinity Groups

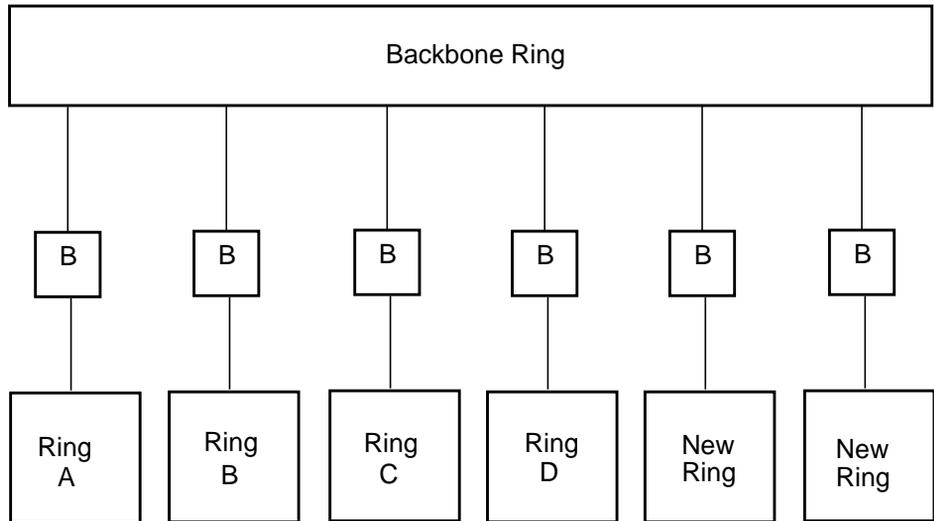
If several end-users perform related tasks on the network, and have little information interchange with groups of other end-users, each of the affinity groups can be attached to different rings within a network. This configuration, as illustrated

below, decreases the network workload and improves overall performance. In this three-ring configuration, Ring A and Ring C are two separate departments with little inter-communication. Ring B contains the servers common to Rings A and B.



Growth

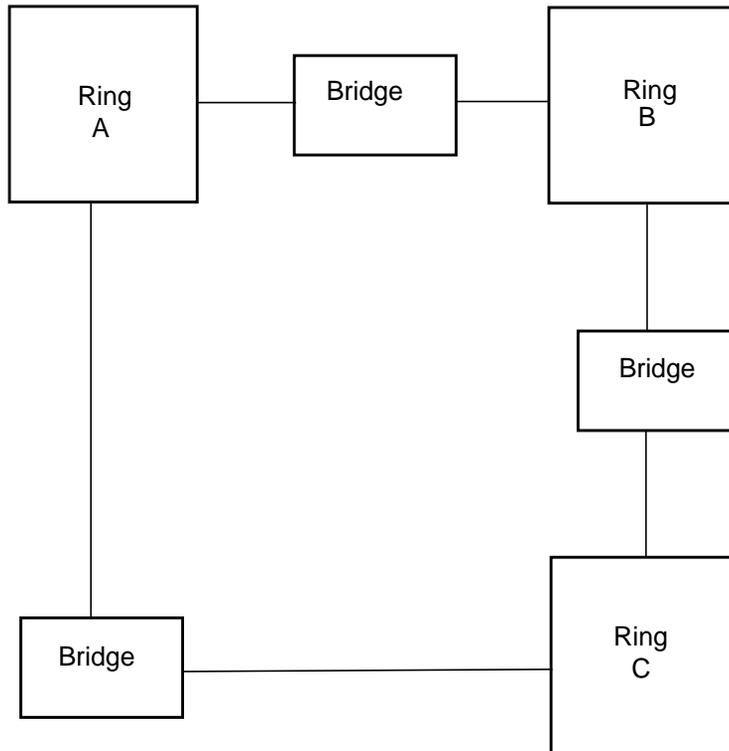
In the course of designing your network, the probability of future expansion should also be considered. Within a multiple-ring network, the addition of new workstations, or new rings, can be performed with a minimum of impact on network activity. The following diagram shows a backbone configuration in which expansion can be added easily.



In the configuration illustrated above, new rings can be added to the Backbone Ring by adding one bridge for each new ring that is added.

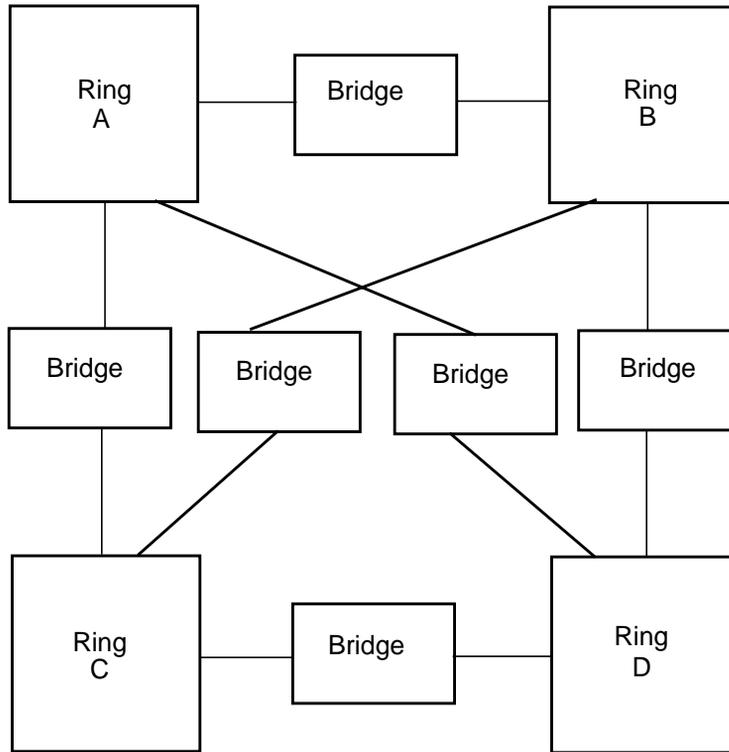
Having decided how many rings to have within the token ring network, the next logical subject deals with the actual inter-connection of multiple rings within a token ring network.

The loop configuration illustrated below, has several unique characteristics. For example, the information exchange between any two workstations on different rings always has only one bridge to cross. This reduces bridge overhead. Further, if an error condition occurs on one bridge and it goes out of operation, an alternative path still exists.



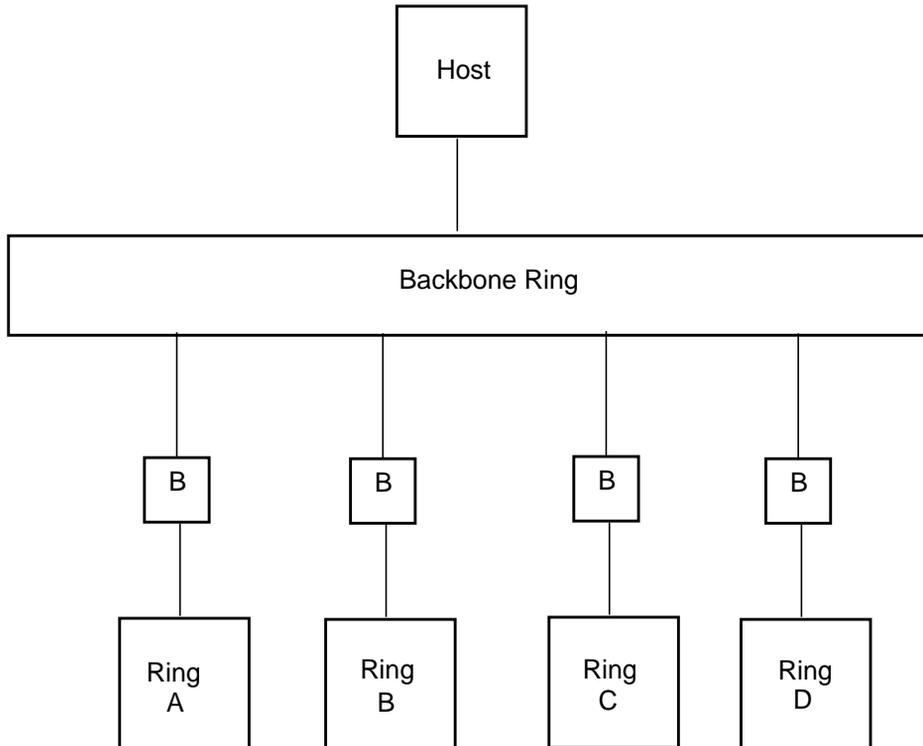
A disadvantage to this loop configuration is that each ring should be bridged to every ring on the network to for maximum performance and connectivity.

Additionally, loop configurations are difficult and expensive to expand. To add one more ring to the loop configuration shown above, would require the addition of three more bridges. See the diagram below.



From this discussion, you can see that if you want to maintain the same loop configuration and add more rings, the number of bridges increases significantly.

For large networks with high growth and host connections, the backbone configuration is more practical than the loop configuration. See the illustration, below for a backbone configuration.



Backbone Configuration Strengths

The method for interconnecting the rings simplifies the physical attachment or relocation of workstations and rings.

Adding a new ring merely requires the addition of one new bridge.

The hosts attached to the backbone ring as are various file, compute, and print servers. This means that rings with low server or host access demands will not load the backbone ring. Additionally, the host and servers are isolated by the bridges from possible beaconing conditions on the other rings.

Backbone Configuration Weaknesses

The exchange of information between lower rings requires crossing two bridges.

The lower-level rings rely on one bridge for connectivity to the host, servers, and to other workstations attached to different rings. So, if one bridge fails, the lower-level ring which it interconnects will lose host and server access.

Lower reliability results because any two workstations have one possible path through which the information can flow. If the path is not available, communication is lost.

Other possible configurations exist and can be evaluated. Actual requirements of your network determine the configuration you design.

Impact on Performance

Bridge processing within multiple-ring token ring networks adds overhead to network throughput. Factors that determine the amount of overhead are listed below.

TABLE C-1 Overhead Determining Factors

Factor	Description
Throughput	An increase in frame size to a certain value, increases bridge throughput.
Hop Delay	This is the elapsed time between the end of the receive stage for a frame entering the bridge and the end of the transmit process for the same frame leaving the bridge. The delay increases as frame size and ring load increase.
Application Delay	Some applications create greater degradation due to bridge processing than others. Applications that involve disk access or host delays will not experience much degradation. Fast applications, such as program loads or memory-to-memory copies, might notice greater impact on performance when the frames cross heavily used bridges.
Frame Flow	Through a bridge passes with higher priority than normal ring station processing because the architecture assigns a higher priority to it.

In general, during normal workload conditions, bridge processing does not cause a major slowdown of activity in a multiple-ring network.

Token Ring Bridge Concept

A bridge is a device that forwards frames from one ring to another within a local area network. It connects multiple rings into one logical network and appears to each ring as a standard ring station. It performs frame checking and participates in token claiming and neighbor identification like any other ring station on the ring.

The bridge's operation on a token ring network is implemented at the Medium Access Control (MAC) sub-layer level. A bridged network is considered to be one large IP, or Network, layer. Thus, all bridged rings would have the same internetwork number. If another ring station on the ring is communicating to the bridge, rather than through it, the bridge is addressed in the same manner as any other ring station - at the MAC sub-layer through the Destination Address field.

Logical link control (LLC) procedures are implemented between pairs of source and destination ring stations. They are not terminated, examined, or affected by intermediate bridges.

Each bridge is assigned a pair of 2-byte ring numbers that correspond to the rings to which it is attached. The Segment Number, also known as the ring address, is used only for routing information and should not be confused with the MAC-level ring station address of the bridge.

During initialization, the bridge requests configuration parameters using the Request Parameter MAC frame. In addition to the MAC configuration parameters, default parameters are used unless their values are modified by an operator.



Caution – Bridges should never be placed in parallel with routers or in a situation where they will interconnect two distinct IP networks already interconnected by a router. You must treat bridged rings as a single subnet for IP addressing purposes.

Source Routing

Each transmitted frame on a token ring network contains a Source Address and a Destination Address. These addresses identify the sending and receiving ring stations. In a multiple-ring environment with two or more rings connected by bridges, the Destination Addresses for frames originating in one ring station and destined to one or more stations in other rings, are built using a unique Source Routing technique. The route is found through bridge stations at session connection time when the initiating station starts collecting routing information that will be included in subsequent frames. This method eliminates the need for maintaining routing tables at predefined nodes on the ring.

Source Routing functionality lies within the Data Link layer of the OSI reference model. IP routing (described on the following pages) lies within the Network layer. The two types of routing are functionally exclusive.

You need not select Source Routing if you are installing your SPARCstation on an isolated token ring network, or on one with interfaces to other Ethernet networks and intend to communicate only with other hosts using TCP/IP. A SPARCstation with multiple network interfaces can route IP packets, but not source-route packets, between the networks.

A more detailed discussion of Source Routing can be found in Appendix B “Token Ring Architecture Overview.”

Setting Up Routers

Whenever you install the SunTRI/P adapter in a machine with an existing Ethernet connection, that machine becomes a router serving both the token ring and Ethernet networks. The router is known to the Ethernet network by its Ethernet IP address and hostname, and to the token ring network by its token ring IP address and hostname.

In general, stations in an Internet can be reached via multiple routes. When the Solaris operating environment is booted, the route daemon (*in.routed*) sends out packets over both the Ethernet and token ring. The route daemon determines the optimum path between stations by relative hop count and not by the speed of the interface. Therefore, in cases where the hop counts are equal, the use of the token ring is not guaranteed.

Users can guarantee the use of token ring for network commands by specifying the target system’s token ring IP name.

Selecting the MAC Address

Each station on the network is identified by its unique MAC address. The SunTRI/P adapter supports IEEE globally-assigned 48-bit MAC addresses. You must select the MAC address for a new station *before* it is brought up on the ring. The following paragraphs describe how to select from the available MAC addresses for diskful stations.

Note – Once a new station starts transactions on the network, the Address Resolution Protocol (ARP) tables on all other stations are updated to include the new station. After this occurs, making changes is more difficult because you would need to delete ARP entries on every active station in the ring.

The choices of MAC address include:

- The MAC address that is resident in the host IDPROM (default). This 48-bit value is administered globally by the IEEE. It is the same IDPROM value used by your SPARCstation's on board ethernet controller.
- A locally-administered MAC address set by ifconfig. This 48-bit value is administered by the network or system administrator.
- The MAC address that is embedded in the SunTRI/P adapter's PROM.

In most cases, you will use the default—the MAC address resident in the host IDPROM. The advantages to using the host IDPROM MAC address are that:

- You do not have to change any settings for the SunTRI/P adapter.
- You conform to Sun's philosophy of using the same host IDPROM value for each LAN attachment that supports 48-bit MAC addresses.
- You avoid conflicts if you use other network protocol stacks that require the host IDPROM MAC address (for example, the XNS networking protocols).
- You need not assure MAC address uniqueness.

Index

SYMBOLS

/etc/hosts file, 14

/etc/path_to_inst file, 13

A

adapter installation instructions, 7

ARP tables, 57

auto-detect, 15

C

cable types, 5

CD-ROM

 mount, 11

compatible platforms, 6

configuration overview, 9

console terminal

 system, 22

D

data rate selection, 15

debugging tools, 19

disconnect from network, 22

driver installation, 11

E

edit

 /etc/path_to_inst file, 13

 hostname.trp file, 14

error messages, 23

Exx00 requirements, 6

F

features, 4

H

hardware overview, 3

hostname.trp file, 14

I

IBM STP cable, 5

ifconfig command, 20

install

 driver, 11

 software, 11

installation overview, 9

installation requirements, 4

instance numbers, 6

J

JTAG jumper J1, 6

M

MAC Address, 57

MAC address, 58

maximum performance, 47

maximum transmission unit (mtu), 15

media filter, 5

messages, 22

mount CD-ROM, 11

mtu

 default, 15

 sizes, 15

N

netstat command, 19, 20

network configuration

 planning, 47

 setting up routers, 57

network traffic

 routing, 57

O

OSI 7 layer model, 3

P

performance factors, 47

ping, 19

pkgadd, 9, 11

pkgadd command, 11

pkgrm command, 10

platform compatibility, 6

R

reconnect to network, 22

removing old instances, 10

ring speed, 15

routing network traffic, 57

S

script, post-installation, 11

set Sole Network Interface, 17

setting up routers, 57

slot assignment, 7

SNAP encapsulation, 3

snoop command, 19

software installation, 11

software overview, 3

source route table, 21

source routing, 3, 21

spray, 19

STP cable, 5

SunTRI/P adapter overview, 3

system messages, 22

T

token ring interface, 3

token ring networks, 1

U

utilities, 19

UTP cable, 5

V

verify install, 20