


RS/6000 and Asynchronous Transfer Mode

December 1996



**International Technical Support Organization
Austin Center**



International Technical Support Organization

SG24-4796-00

RS/6000 and Asynchronous Transfer Mode

December 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, "Special Notices" on page 57.

First Edition (December 1996)

This edition applies to Version 4.1.4, Version 4.1.5 and Version 4.2 of the AIX Operating System.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. JN9B Building 045 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1996. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|---|------|
| Figures | v |
| Preface | vii |
| How This Redbook Is Organized | vii |
| The Team That Wrote This Redbook | viii |
| Comments Welcome | viii |
| Chapter 1. Asynchronous Transfer Mode (ATM) and RS/6000 | 1 |
| 1.1 Introduction to ATM | 1 |
| 1.1.1 What is ATM? | 1 |
| 1.1.2 ATM Network | 4 |
| 1.1.3 Classical IP and ARP over ATM - RFC 1577 | 7 |
| 1.2 ATM and AIX | 8 |
| 1.2.2 AIX ATM Capabilities | 8 |
| Chapter 2. Hardware Installation | 11 |
| 2.1 Prerequisites | 11 |
| 2.1.1 Installing the Adapter | 11 |
| 2.2 Connecting Cables to the Adapter | 12 |
| Chapter 3. Software installation | 13 |
| 3.1 Prerequisites | 13 |
| 3.2 Installing the ATM Device Drivers | 13 |
| Chapter 4. Connecting Two RISC System/6000s Back-to-Back | 15 |
| 4.1 Prerequisites | 15 |
| 4.1.1 SONET Clocking | 15 |
| 4.2 Adding an ATM Network Interface | 15 |
| 4.3 Configuring a PVC for IP over ATM | 17 |
| 4.3.1 Testing the TCP/IP Connection over ATM | 18 |
| Chapter 5. Connecting RISC System/6000s with an ATM Switch | 21 |
| 5.1 Prerequisites | 21 |
| 5.2 Configuring the ARP Servers Network Interface | 22 |
| 5.3 Getting the ARP Server Address | 23 |
| 5.4 Configuring the ARP Client Network Interface | 24 |
| 5.4.1 Testing the TCP/IP Connection over ATM | 25 |
| Chapter 6. Using ATM as a Backbone with RISC System/6000s | 27 |
| Chapter 7. Multiple Interfaces on the Same ATM Adapter | 31 |
| 7.1 More Than Eight Interfaces | 31 |
| 7.2 Configure Multiple Network Interfaces on the Same ATM Adapter | 32 |
| Chapter 8. An Approach to a Backup ARP Server | 35 |
| 8.1 Address Resolution | 35 |
| 8.2 ARP Server | 35 |
| 8.3 Backup ARP Server | 35 |
| 8.4 ARP Clients | 36 |
| 8.5 Takeover Situation | 36 |

| | |
|---|----|
| Chapter 9. Troubleshooting | 37 |
| 9.1 Helpful Commands | 37 |
| 9.1.1 errpt | 37 |
| 9.1.2 iptrace and tcpdump | 39 |
| 9.1.3 atmstat | 39 |
| 9.2 Troubleshooting | 44 |
| 9.2.1 Back-to-Back Configuration | 44 |
| 9.2.2 Switched Configuration | 45 |
| | |
| Chapter 10. Performance and Tuning | 49 |
| 10.1 Overview | 49 |
| 10.2 Tuning TCP/IP | 49 |
| 10.3 Tuning the ATM Adapter Parameters | 50 |
| | |
| Appendix A. End Node Impact on Throughput | 55 |
| | |
| Appendix B. Special Notices | 57 |
| | |
| Appendix C. Related Publications | 59 |
| C.1 International Technical Support Organization Publications | 59 |
| C.2 Redbooks on CD-ROMs | 59 |
| C.3 Other Publications | 59 |
| | |
| How To Get ITSO Redbooks | 61 |
| How IBM Employees Can Get ITSO Redbooks | 61 |
| How Customers Can Get ITSO Redbooks | 62 |
| IBM Redbook Order Form | 63 |
| | |
| List of Abbreviations | 65 |
| | |
| Index | 67 |

Figures

| | |
|--|----|
| 1. ATM Cell Distribution | 3 |
| 2. ATM Layers | 3 |
| 3. An ATM Network | 5 |
| 4. VCs and VPs on a Physical Link | 6 |
| 5. Connector Types | 12 |
| 6. RISC System/6000s Connected Back-to-Back | 15 |
| 7. RISC System/6000s Connected to the ATM Switch | 21 |
| 8. RISC System/6000s as an ATM Backbone | 27 |
| 9. Routing TCP/IP on One Adapter | 31 |
| 10. Processor Impact on Throughput | 55 |
| 11. Disk I/O Impact on ATM Throughput | 56 |

Preface

This redbook addresses Asynchronous Transfer Mode (ATM) technology for the RS/6000. ATM is an emerging high-speed communications technology that is consistent in both wide area and local area networks.

This redbook was written for everyone who needs to implement ATM and Classical TCP/IP on the AIX Operating System.

In this book we document our experiences in connecting RS/6000s back-to-back and to ATM networks using ATM technology. TCP/IP is configured for both PVCs and SVCs and examples are given of routing between two Logical IP Subnets and TCP/IP on other media.

Some knowledge of AIX, ATM, and TCP/IP is assumed.

How This Redbook Is Organized

This redbook contains 68 pages. It is organized as follows:

- Chapter 1, "Asynchronous Transfer Mode (ATM) and RS/6000"
This provides a brief overview of ATM for AIX.
- Chapter 2, "Hardware Installation"
This describes the installation of the Turboways 100 and 155 adapters.
- Chapter 3, "Software installation"
This chapter describes the installation of ATM software.
- Chapter 4, "Connecting Two RISC System/6000s Back-to-Back"
This gives an example of how to connect two RS/6000s without an ATM switch. PVCs are also configured in this chapter.
- Chapter 5, "Connecting RISC System/6000s with an ATM Switch"
This chapter describes connection to ATM switches and the configuration of ATMARP server and clients.
- Chapter 6, "Using ATM as a Backbone with RISC System/6000s"
This provides an explanation of routing TCP/IP among ATM, Ethernet, Token-Ring, and X.25.
- Chapter 7, "Multiple Interfaces on the Same ATM Adapter"
In this chapter multiple interfaces are configured on one adapter.
- Chapter 8, "An Approach to a Backup ARP Server"
This is one approach to providing a backup ATMARP server.
- Chapter 9, "Troubleshooting"
This chapter covers problem determination.
- Chapter 10, "Performance and Tuning"
This chapter gives some suggestions to maximize the ATM throughput.
- Appendix A, "End Node Impact on Throughput"

This chapter describes some of the things that can affect throughput.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization Austin Center.

Al Mitchell is a Senior SE Support Representative at the International Technical Support Organization, Austin Center. He writes extensively and teaches IBM classes worldwide on all areas of AIX communications. Before joining the ITSO two years ago, he worked in USM&S as a Senior AIX Marketing Support Specialist.

Hartmut Sprandl is an AIX Systems Support Specialist in Germany.

Sue Foulds is an AIX Systems Support Specialist in the UK. She has five years of experience in the AIX Networking field. Her areas of expertise include X.25 and SNA. She has written on these subjects.

Praben Primasaputra is an AIX Country Support Systems Engineer in Jakarta, Indonesia. He has been with IBM for six years, and has five years of experience in AIX, with the last two years specializing in Networking. He holds a degree in Telecommunication Engineering from Sepuluh November Institute of Technology Surabaya, Indonesia. His areas of expertise include, HACMP/6000, SNA Server, X.25, AIX Performance Tuning, and TME10 (Tivoli). He has written extensively on AIX System Management for IBM Indonesia operation.

Thanks to the following people for their invaluable contributions to this project:

Daisy Chang
RISC System Division Austin

Dave Cosby
Networking System Division Raleigh

Tom Adams
RISC Systems Division Austin

Cindy Young
AIX Systems Center Dallas

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Chapter 1. Asynchronous Transfer Mode (ATM) and RS/6000

In this section we will present an overview of ATM, Classical IP and how these relate to AIX offerings. This manual will address ATM for AIX Version 4.1.4, 4.1.5 and 4.2. This information should be valid for later releases.

1.1 Introduction to ATM

In this manual, we will not attempt to provide a detailed explanation of ATM. Please refer to *Asynchronous Transfer Mode (ATM) Technical Overview*, SG24-4625 for a more complete description of the technology. There is an abundance of other books on this technology.

1.1.1 What is ATM?

Asynchronous Transfer Mode is relatively new communications technology. It is technology that was selected to deliver the Broad Band Integrated Services Digital Network, B-ISDN. This concept was developed starting in 1988.

ATM is a set of rules and standards that define a communications method to satisfy the needs of current and evolving applications. ATM defines switched, cell-based, connection-oriented networks. The major components are end nodes, in our case RS/6000s, and network nodes, the switches. ATM defines the interfaces between these nodes. ATM specifies the communications layers, services and application mapping to the ATM services on the end nodes.

In a switched network the full bandwidth of the physical connection is available between nodes. With shared bandwidth LANs, such as token ring and ethernet, the nodes connected to a segment share the bandwidth of that segment.

Although most network technologies are packet-based, ATM differs in that a fixed length packet 53 byte cell is specified. Of those bytes, 48 bytes are payload and 5 bytes are header. The packet size is a compromise to facilitate different application requirements. Fixed length packets in a switched network will be received at the end node at regular intervals. This allows applications that are timing-dependent.

In connection-oriented networks, the path between the end nodes is established at the time the connection is made. During a session, the connection appears to be point-to-point. In a switched connection-oriented network, cells are guaranteed to be delivered in the same sequence they are sent. Cells can be lost in the network but a cell can never pass the cells sent before it.

Since ATM networks use highly reliable switches and media, minimal error detection and correction is done within the network. Error correction is usually performed by the end node application, such as TCP. This has a positive performance impact on networks with very low error rates.

1.1.1.1 LAN, MAN and WAN

ATM networks can be LAN, MAN, WAN or any combination. ATM networks are not limited by geography.

ATM defines the interfaces between end nodes and switches for private ATM networks, LANs, networks within one common carrier, MANs, and among common carriers, WANs. This means a connection can be established between end nodes anywhere in the ATM network without using routers or gateways. To the end user it all appears to be local. In Figure 3 on page 5 all the RS/6000s can be defined in the same subnet. Connections can be established between any 2 nodes without using routers. The end nodes in the figure can be dispersed throughout the world.

1.1.1.2 Information Types and Requirements

The evolving world of communications has produced the desire to transmit several types of information across a single communications technology. The most common types of information are:

- Voice - Real Time and/or Compressed
- Video - Real Time and/or Compressed
- Data - Large and Small Transmissions
- Multimedia - Any/All of the Above
- Emulated Leased Line - Looks Like Today's Lines

The requirements of the different information types are radically or slightly different among themselves. The only constants are reliability and cost. Even these vary since video can tolerate an occasional lost packet, while data integrity is the holy grail of the computer industry, and cost is relative in each situation. ATM attempts to provide a network solution for all these information types. An ATM network can integrate all these data types in a single network.

Real time voice and video and emulated leased line require Constant Bit Rate (CBR) traffic, information sent and received at regular intervals. In ATM terminology, this is service class A.

Compressed voice and video and multimedia applications can use Variable Bit Rate (VBR) traffic, an amount of information during an interval. In ATM terminology, this is service class B.

Data, the stuff we send around now with TCP/IP and other protocols, just has to get there reliably. In ATM terminology, this is service class C or D. Class C is connection-oriented while class D allows for connectionless protocols.

How fast information is sent and received is determined by the application and how much money you are willing to spend.

To oversimplify how ATM works, let's assume that a Class A application that requires every third cell is sharing a physical link with a Class B application that requires two cells of every 10 cells and two Class C applications, C1 and C2. The cells on the link might look like Figure 1 on page 3.




Figure 1. ATM Cell Distribution

1.1.1.3 ATM Layers

In Figure 2 you can see the layers of ATM and how the ATM adaptation layers (AALs) relate to the classes of service.

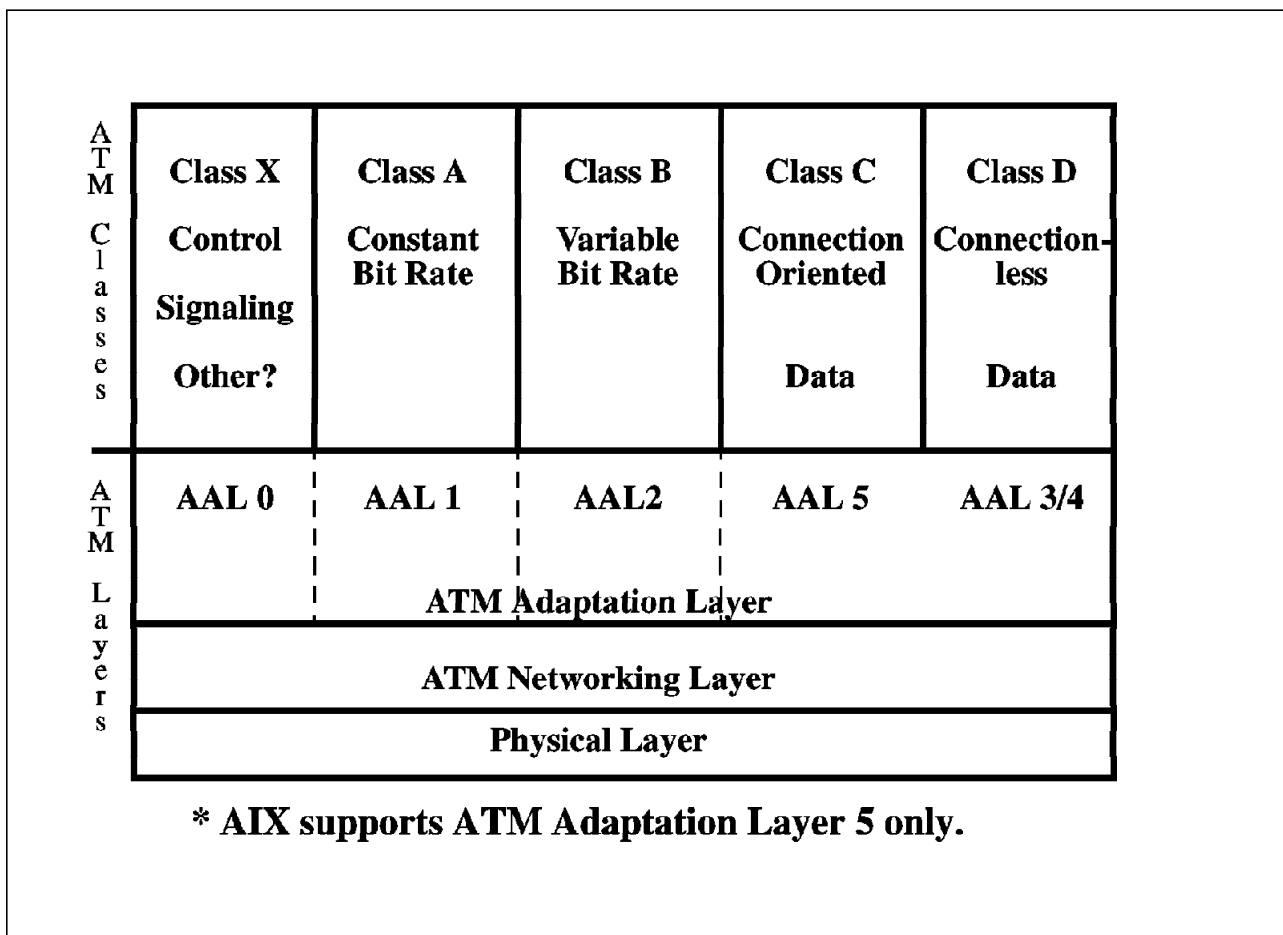


Figure 2. ATM Layers

1.1.1.4 AAL 5

The purpose of the ATM Adaptation Layer, AAL, is to adapt the internal network characteristics to the traffic types.

You can also see from Figure 2 that AIX supports AAL 5 only. This is the most widely implemented AAL among ATM end nodes. This is the least complicated and most efficient AAL. AAL 5 is considered sufficient for Classical TCP/IP and LAN Emulation, the two most common uses of ATM today.

Note: AIX supports only Classical TCP/IP as of this writing.

Once a connection is established between 2 end nodes, the major functions of AAL 5 for the sending end node are:

- Interface with the higher level application, TCP/IP in our case, at the service interface.
- Add fill bytes so that the user data plus the trailer to be added is a multiple of 48 bytes.
- Add the trailer. This is now a PDU.
- Segment the PDU to 48 byte segments.
- Indicate the final segment.
- Pass the segments to the ATM layer.

The ATM layer adds the 5 byte ATM header and passes the cell to the physical layer for transmission.

The receiving end node reverses the above process.

For more information see Asynchronous Transfer Mode (ATM) Technical Overview, SG24-4625.

1.1.1.5 ATM Addresses

ATM addresses are 20 bytes in length. The address is made up of 13 bytes that are set at the switch, 6 bytes that is the MAC address of the adapter and 1 byte that is used for the interfaces configured on that adapter. This byte is called the selector byte. AIX allows up to 256 ATM interfaces for one RS/6000.

The ATM Forum has defined 3 address formats for the first 13 bytes of the address.

- ITU-T (E.164)
- DCC (Data Country Code)
- ICD Format

These address formats are in the ATM Forum UNI 3.0 Specification. AIX supports all three formats.

Once the session is established, ATM addresses are no longer used. Each end node will be assigned a Virtual Path:Virtual Channel number. This will be the interface for all cells for the duration of the session.

1.1.2 ATM Network

Figure 3 on page 5 shows an ATM network with end nodes, switches and physical links. An RS/6000 can be attached to a local switch or a switch belonging to an ATM network provider. RS/6000s can also be connected back-to-back as seen in Chapter 4, "Connecting Two RISC System/6000s Back-to-Back" on page 15.

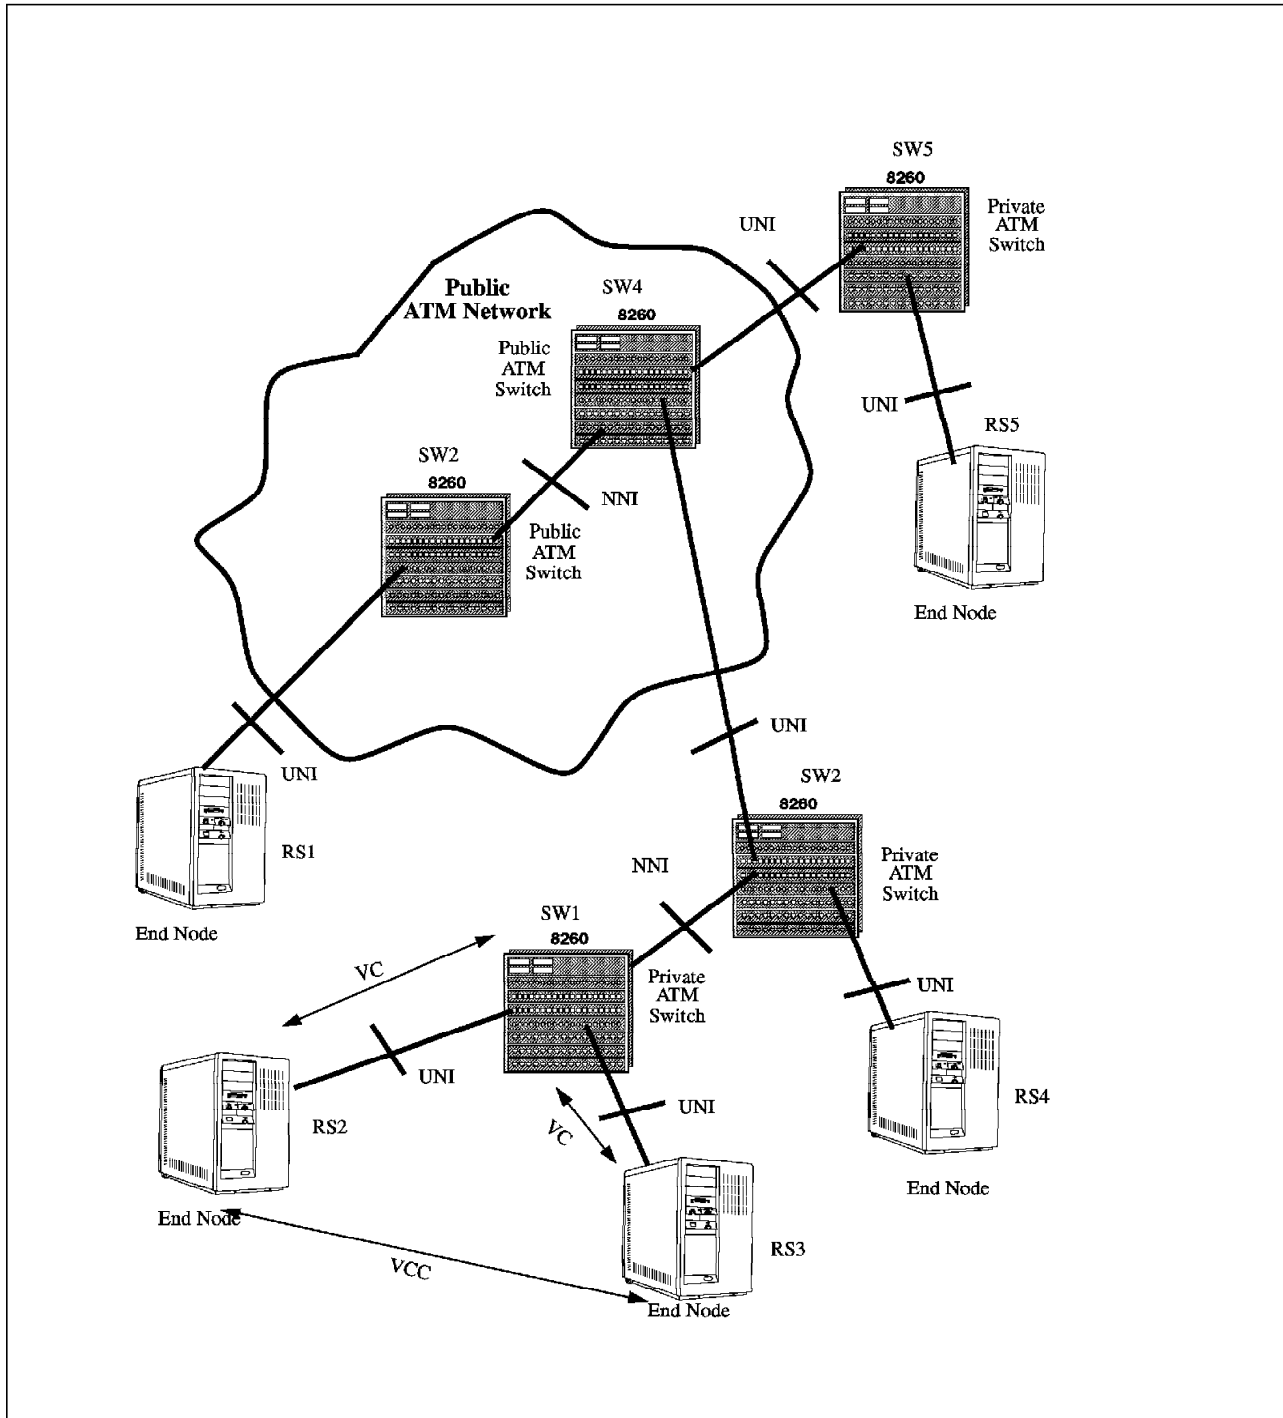


Figure 3. An ATM Network

1.1.2.1 Interfaces

The ATM Forum defines the interfaces between all nodes in the network. This assures that the devices can work together. The two basic interface types are User Network Interface, UNI, and Network Node Interface, NNI. There are several types of UNIs and NNIs, all beyond the scope of this manual. See *Asynchronous Transfer Mode (ATM) Technical Overview* for more information on interfaces.

AIX supports ATM Forum UNI Specification V3.0. This allows RS/6000s to be end nodes connected to private or public ATM switches that conform to the same standard UNI.

1.1.2.2 Virtual Channels and Virtual Paths

AIX versions 4.1.4, 4.1.5 and 4.2 support up to 1024 virtual channels per adapter. The switch or service provider that your RS/6000 is connected to may allow a different number.

Information flows on virtual channels. Virtual channels can be grouped together in virtual paths. A physical link can have many virtual paths each with many virtual channels.

When a connection is established between 2 end nodes, VP:VCs are connected throughout the network to establish a Virtual Channel Connection, VCC. The VP:VC pairs of the two ends are now connected.

In Figure 3 on page 5 there are VCs between end node RS2 and switch SW1 and SW1 and RS3. When a connection is made, the VCC exists between RS2 and RS3.

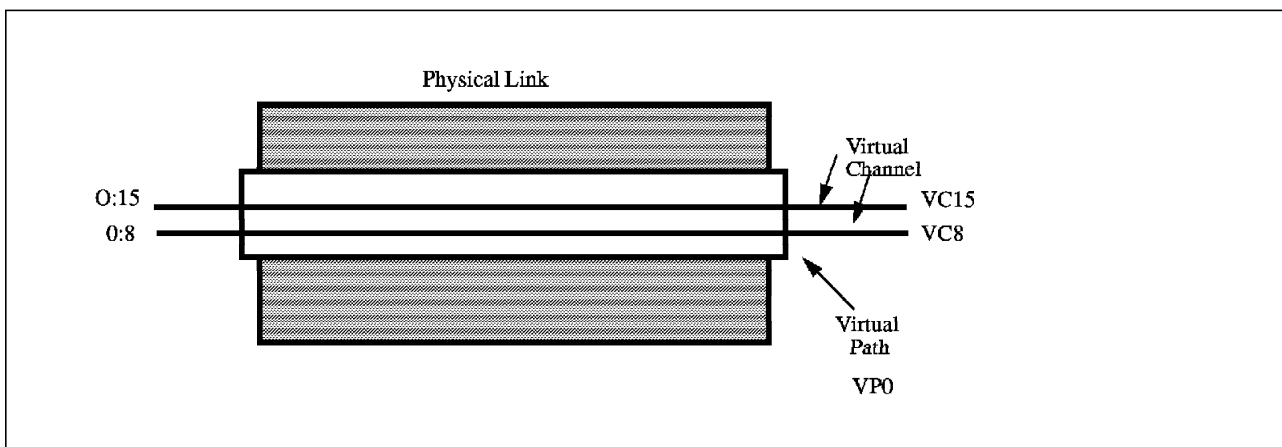


Figure 4. VCs and VPs on a Physical Link

1.1.2.3 PVC and SVC

Virtual channels may be switched, SVCs, or semipermanent, PVCs. SVCs are dynamically created when an end node requests a connection to another end node. A VCC is created through the ATM network and a VP:VC pair is assigned at each end node. It is assumed that this connection will be terminated. From a TCP/IP viewpoint, applications such as telnet, ftp, etc. would request that SVCs be allocated and terminated. Classical IP uses SVCs. TCP/IP sessions can be established using PVCs also.

PVCs are predefined VP:VC pairs. The VCC is configured in the switch nodes and is always assigned to the end nodes. The PVCs are available as long as the network is intact.

1.1.3 Classical IP and ARP over ATM - RFC 1577

Broadcasts and arps are not sent over connection-oriented networks such as ATM, PPP and X.25. Clients that broadcast to discover the addresses of servers will not work with ATM.

1.1.3.1 RFC 1577

Since TCP/IP addresses, x.x.x.x, must be resolved to 20-byte ATM addresses RFC 1577 defines IP operation and ATM address resolution protocol, ATMARP, over ATM AAL5. This RFC defines the operation ARP clients, ARP servers and Logical IP Subnetworks, LIS.

1.1.3.2 Logical IP Subnet - LIS

RFC 1577 defines a Logical IP Subnetwork, LIS, for ATM. The LIS consists of members in an ATM network with a common subnet address and a single ARP server. Since an ATM network can span many switches, an LIS may be geographically dispersed. An ATM network is not limited to one LIS. Moreover, an end node can be a member of several LISs. This member can function as a router between the subnets. Members using TCP/IP can only establish VCCs with other members of the same LIS.

Please see Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21 for information on creating an RS/6000 LIS. Also see Chapter 7, "Multiple Interfaces on the Same ATM Adapter" on page 31 for information on an RS/6000 belonging to more than one LIS.

Some of the requirements for LIS members are:

- same subnet address and subnet mask
- same MTU
- direct connection to the ATM network
- support ATMARP and InATMARP

1.1.3.3 ARP Server and Client

Every LIS must have an ATMARP server to resolve IP addresses to ATM addresses. This mechanism is required for SVCs. This ARP Server can be a member of the LIS that is capable of this function.

Note:

Some ATM switches now have the ARP server function as an option.

An ARP client uses the destination 20-byte address obtained from the ARP server to request an SVC be created to the destination end node. When the call is accepted by the destination end node, a VCC is established and VP:VC pairs are assigned at each end.

See Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21 and Chapter 8, "An Approach to a Backup ARP Server" on page 35 for information on configuring RS/6000s as ATM ARP servers and clients.

1.1.3.4 TCP/IP Access Outside a LIS

TCP/IP packets may be routed to hosts that are not members of the LIS. One of the LIS members can serve as a router to another LIS or to a TCP/IP subnet using another media. Routing methodologies that use broadcast to resolve addresses will not work in ATM networks. See Chapter 6, "Using ATM as a Backbone with RISC System/6000s" on page 27 for examples of routing to other types of networks.

1.2 ATM and AIX

As you can see from the discussion so far, not all of the ATM technology is currently implemented on AIX V4.1.4, V4.1.5 and V4.2.

1.2.1.1 ATM Adapters

AIX V4.1.4, V4.1.5 and V4.2 currently support the following adapters:

- IBM Turboways 100 ATM Adapter - FC 2984
- IBM Turboways 155 ATM Adapter - FC 2989

Both adapters require 62.5 micron multi-mode fiber cable with an SC-type connector.

Note: There is a 155Mbps ATM adapter available for AIX V3.2.5 that supports connections via twisted pair. Use a CAT 5 10 base T cable only.

1.2.1.2 AIX and RS/6000 ATM Standards

These are the standards currently adhered to by the AIX V4.1.4, 4.1.5 and V4.2 ATM subsystems and adapters:

- CCITT 1.361
- ATM Forum UNI Specification V3.0
- ANSI T1S1.5/92-002R3
- TCP/IP RFC 1577 "Classical IP and ARP over ATM"

1.2.2 AIX ATM Capabilities

At this writing, AIX 4.1.4, 4.1.5 and 4.2 support ATM as follows:

- 100Mbps and 155Mbps - Fiber Only
- AAL 5 Only
- UNI 3.0
- Classical IP Only - TCP/IP - RFC 1577
- TCP/IP over SVCs
- Sockets API (AIX 4.2)
- SNMP MIB

Be certain that your AIX system is at the current PTF level for ATM components to ensure correct operation.

ATM LAN Emulation will be supported in a future release of AIX.

This soft copy for use by IBM employees only.

Although this list appears to be restrictive, most AIX applications requiring high speed connections will not be impacted. This is a relatively new technology that will evolve.

Chapter 2. Hardware Installation

In this chapter you will find a step-by-step description of the installation process for the ATM adapters. It covers the Installation of the ATM 100 and ATM 155 Adapter.

2.1 Prerequisites

For the currently available ATM Adapters for the RISC System/6000 you have to have at least one available Micro Channel slot in your machine. Due to the Adapter's height, it cannot be installed in all RISC System/6000 models.

2.1.1 Installing the Adapter

The ATM Adapter may be installed in any available Micro Channel slot. If you are not immediately connecting the adapter to a switch or another machine, make sure that the wrap connector is attached to the card.

Note: It is good practice to execute the `diag -a` command before installing any adapters.

Once the card is installed, use SMIT to check and configure the adapter.

Once installed, check the adapter with:

```
lsparent -C -k atm -H
```

If you get an output like:

```
name status    location description
atm0 Available 00-05    100 Mbps ATM Adapter
```

or:

```
name status    location description
atm0 Available 00-05    155 Mbps ATM Adapter
```

then the adapter is installed correctly and the microcode has been loaded onto the card.

If you get an output like:

```
name status    location description
atm0 Defined   00-05    100 Mbps ATM Adapter
```

or:

```
name status    location description
atm0 Defined   00-05    155 Mbps ATM Adapter
```

then either the microcode has not been installed on the machine before the adapter was installed or the adapter needs to be reseated. To install the microcode see Chapter 3, "Software installation" on page 13.

2.2 Connecting Cables to the Adapter

For the RISC System RS/6000 and AIX V4, only adapters for fiber optic cables are supported. Connections between the RISC System RS/6000 ATM adapter and the switch, hub device or, in a back-to-back configuration, a second RISC System RS/6000, are made with a 62.5 micron multi-mode duplex fiber optic cable that is terminated with SC-type connectors. The other end of the fiber cable will have to have the appropriate connector for the device at that end. The RS/6000 ATM adapters support only SC-type connectors.

The RISC System RS/6000 ATM adapter has two fiber sockets, one for receiving and one for transmitting. To connect two ATM devices in the correct way, you have to connect the transmitting socket of one device to the receiving socket of the other device. In other words you have to cross the cables.

Note: It is very important to connect the cable in the correct way, otherwise the connection will not work. Most switches have indicators that show that the cables are connected correctly. The RS/6000 error report will show wire fault for the atm(x) device when cables are incorrectly connected. Use SMIT or the errpt command to view the error log.




Figure 5. Connector Types

Fiber cables are available with different connectors on each end. This is sometimes required to connect an RS/6000 to a switch.

Chapter 3. Software installation

This section will describe the step-by-step process for installing the microcode for the ATM adapters under AIX Versions 4.1.4 and AIX 4.2. This information should be correct for later releases.

3.1 Prerequisites

AIX V4.1.4 or AIX V4.2 have to be installed. It is good practice to have the current PTFs applied.

3.2 Installing the ATM Device Drivers

The following filesets on the base operating system install media provides the ATM support at AIX 4.1.4 and 4.2:

Turboways 100 ATM software:

```
devices.mca.8f7f.com
devices.mca.8f7f.rte
devices.mca.8f7f.unicode
```

Turboways 155 ATM software:

```
devices.mca.8f67.com
devices.mca.8f67.rte
devices.mca.8f67.unicode
devices.mca.8f7f.com
```

On AIX 4.2 there is some additional software which also needs to be installed:

```
devices.common.IBM.atm.rte
```

Before installing the ATM device driver software check that it is not already installed by issuing the following commands:

```
lslpp -l devices.mca.8f7f      for the Turboways 100 ATM Adapter
lslpp -l devices.mca.8f67     for the Turboways 150 ATM Adapter
lslpp -l devices.common.IBM.atm.rte only if you have AIX 4.2 installed
```

If the output does not show the filesets use the following SMIT screen:

```
smit install_latest
```

```

                                Install Software Products at Latest Level

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                [Entry Fields]
* INPUT device / directory for software      /dev/rmt0.1
* SOFTWARE to install                        [Common ATM >]
PREVIEW only? (install operation will NOT occur)  no
COMMIT software updates?                      yes
SAVE replaced files?                          no
ALTERNATE save directory                      []
AUTOMATICALLY install requisite software?       yes
EXTEND file systems if space needed?            yes
OVERWRITE same or newer versions?              no
VERIFY install and check file sizes?           no
Include corresponding LANGUAGE filesets?       yes
DETAILED output?                              no
```

Scroll down to the end of the SMIT output, to check if the software was successfully installed.

If the ATM adapter is already installed in the machine then the `cfgmgr` command will need to be run or the machine rebooted in order to enable the microcode to be loaded onto the adapter.

Check if the adapter is now available:

```
lsparent -C -k atm -H
```

Chapter 4. Connecting Two RISC System/6000s Back-to-Back

This chapter describes how to set up a TCP/IP connection over ATM between two RISC System/6000s in a back-to-back configuration. For this configuration no ATM switch is needed. The TCP/IP connection will be established over a PVC connection. Only PVC connections server are supported for a back-to-back configuration. An ARP server is not needed. See Chapter 1, "Asynchronous Transfer Mode (ATM) and RS/6000" on page 1 for a description of a PVC connection.




Figure 6. RISC System/6000s Connected Back-to-Back

4.1 Prerequisites

Both systems must have ATM adapters installed with the same speed. Adapters with different speeds will not work in this configuration. Both adapters must be available and the fiber optic cable must be connected to both machines. See Chapter 2, "Hardware Installation" on page 11 and Chapter 3, "Software installation" on page 13 for installing the ATM adapter and the microcode and for connecting the fiber optic cable.

4.1.1 SONET Clocking

The Turboways 155 adapters use SONET clocking. This function is normally provided by the switch node. Since there is no switch in the back-to-back configuration, one of the adapters will provide the clocking. This is done by setting the Provide SONET clock field to 1 on the Change / Show Characteristics of an ATM Adapter SMIT screen. The SMIT fastpath is `smit chg_atm`. This should be set on only one of the adapters.

Turboways 100 adapters do not use SONET clocking.

4.2 Adding an ATM Network Interface

TCP/IP uses a network interface to talk to the ATM adapter hardware using the TCP/IP protocols. On both systems you have to add a network interface. To add this network interface do the following on system1:

```
smit mkinetat
```

or

smit

- >Communications Applications and Services
- >TCP/IP
- >Further Configuration
- >Network Interfaces
- >Network Interface Selection
- >Add a Network Interface
- >Add an ATM Network Interface

You will be presented with the following screen which should be filled in as shown:

```

                                Add an ATM Network Interface
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* INTERNET ADDRESS (dotted decimal)           [Entry Fields]
Network MASK (hexadecimal or dotted decimal) [9.4.5.58]
Network Interface                             [at0]
Connection Type                               pvc
ATM Server Address                            []
Alternate Device                              [atm0]
Idle Timer                                    []
Best Effort Bit Rate (UBR) in Kbits/sec      []
* ACTIVATE the Interface after Creating it?   yes
```

Description of the input fields:

INTERNET ADDRESS

As for any TCP/IP interface the Internet Address must be unique to this adapter and should be in dotted decimal format.

Network MASK

The Network Mask is an optional field but is usually the subnet mask assigned relative to this adapter.

Network Interface

The Network Interface in this example is at0. AIX V4 supports 256 interfaces.

Connection Type

For a back-to-back configuration, the connection type should always be pvc.

ATM Server Address

Leave this field blank for back-to-back configurations. This field contains the 20-byte ARP address of the ARP server. SVC connections do not use ATM ARP servers.

Alternate Device

The Alternate Device field is the device name associated with this interface (atm0).

Idle Timer

Leave this field blank. This parameter is not used with PVC connections. Specify a decimal number representing the minutes that a SVC connection will stay up without any data traffic. The default is 60 minutes.

Best Effort Bit Rate (UBR) in Kbits/sec

Leave this field blank. This field is not available in AIX 4.1.4. It specifies the bit rate for the Best Effort connections that are initiated from the configuring interface. This network interface supports the Classical IP over ATM connections. All of the outgoing SVC connections will be set as Best Effort connections with the specified bit rate as the forward and backward peak rate. When this attribute is 0 (by default), the network interface uses the peak rate supported by the associated device as the Best Effort bit rate. Specify a decimal integer in K bits per second. For example, the value 155000 represents 155M bits per second. This field does not apply to PVC connections.

ACTIVATE the Interface after Creating it?

If yes, the added interface will be activated.

Add the network interface on system2 as you have on system1. Only the TCP/IP ADDRESS field will be different.

4.3 Configuring a PVC for IP over ATM

In a back-to-back configuration, only PVCs are used. Refer to Chapter 1, "Asynchronous Transfer Mode (ATM) and RS/6000" on page 1 for information on PVCs. You have to configure both systems with PVCs. To configure a PVC for IP over ATM use the following smit path on system1:

```
smit mkatmpvc
```

or

```
smit
```

```
->Communications Applications and Services  
->TCP/IP  
->Further Configuration  
->Network Interfaces  
->Network Interface Selection  
->PVCs for IP over an ATM Network  
->Add a PVC for IP over an ATM Network
```

You will be presented with the following screen which needs to be completed as follows:

| Add a PVC for IP over an ATM Network | |
|---|----------------|
| Type or select values in entry fields. Press Enter AFTER making all desired changes. | [Entry Fields] |
| PVC Description (Optional) | [] |
| * VPI:VCI | [0:94] |
| * Network Interface | [at0] |
| Destination IP Address | [] |
| Automatically Discover Destination IP Address | yes |
| LLC Encapsulation | yes |

Description of the input fields:

PVC Description

This field is optional. You can enter any text you wish.

VPI:VCI

The most important field on this SMIT screen is the VPI:VCI parameter. This entry has the format x:y, where 'x' refers to the Virtual Path Identifier and 'y' refers to the Virtual Channel Identifier. The 'x' value will ALWAYS be '0' and the 'y' value will be the number of the PVC being added. The VCI number is limited by the maximum number of virtual connections supported by the device. VCIs 0 through 31 are reserved. Do not use these reserved VPI:VCI numbers for the PVCs you define.

Network Interface

Specify the network interface associated with this PVC (for example, at0).

Destination IP Address

Leave this field at the default. Specify the IP address of the destination host for this PVC. The address format is the standard IP address format, 4-bytes, each byte separated by a . (period). This field is optional because the ATM network interface has the capability of obtaining the remote IP address by using the INverseARP protocol. If you want the network interface to obtain the remote IP address for the PVC connection, do not enter a value in this field. If you do not enter a value in this field, make sure you define the Network Interface, and make sure that both sides of the PVC connection support the ATM ARP protocol and LLC/SNAP encapsulation. If you want to enter an IP address, you have to specify the IP address of system2.

Automatically Discover Destination IP Address

Leave this field at the default. Specify whether the ATM ARP protocol will be supported on this PVC. The default value is yes. If the remote host does not support the ATM ARP protocol on this PVC, you can specify no to ensure compatibility. If you specify no, you must define the Destination IP address for this PVC.

LLC Encapsulation

Leave this field at the default. Specify whether LLC/SNAP encapsulation is supported on this PVC. The default value is yes. If the remote host does not support LLC/SNAP encapsulation on this PVC, you can specify no to ensure compatibility. Without LLC/SNAP encapsulation support, the ATM ARP protocol cannot be supported, and the Destination IP address of this PVC must be defined.

Configure a PVC on system2 with the same values as on system1. Make sure that you enter the same VPI:VCI value as you entered on system1.

The systems should now be able to communicate with each other via TCP/IP over ATM.

4.3.1 Testing the TCP/IP Connection over ATM

The output from the ping command on a successful TCP/IP over ATM configuration is as follows:

```
PING 9.4.5.151: (9.4.5.151): 56 data bytes
64 bytes from 9.4.5.151: icmp_seq=0 ttl=255 time=1 ms
64 bytes from 9.4.5.151: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 9.4.5.151: icmp_seq=2 ttl=255 time=1 ms
```

```
----9.4.5.151 PING Statistics----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
```

If the ping fails then follow the normal TCP/IP debug procedures to find out where the failure is occurring. The AIX error report will show wire fault if the fiber cables are incorrectly connected.

The arp table output for a correctly configured TCP/IP over ATM connection is as follows when the arp -t atm -a command is issued.

For AIX V4.1.4 the output is:

```
PVC
      VPI:VCI      Handle IP Address      llc      arp
      0:94         3      NS1(9.4.5.151)      True     True
```

The ifconfig output for the at0 interface is as follows for

```
at0: flags=8000861<UP,NOTRAILERS,RUNNING,SIMPLEX>
      inet 9.4.5.151 netmask 0xff000000
```

The netstat -i output for the at0 interface should look as follows for AIX V4.x:

| Name | Mtu | Network | Address | Ipkts | Ierrs | Opkts | Oerrs | Coll |
|------|-------|---------|-----------------|-------|-------|-------|-------|------|
| lo0 | 16896 | <Link> | | 290 | 0 | 290 | 0 | 0 |
| lo0 | 16896 | 127 | loopback | 290 | 0 | 290 | 0 | 0 |
| at0 | 9180 | <Link> | 8.0.5a.99.10.a9 | 4 | 0 | 4 | 0 | 0 |
| at0 | 9180 | 9 | system1 | 4 | 0 | 4 | 0 | 0 |

Chapter 5. Connecting RISC System/6000s with an ATM Switch

This chapter describes how to set up a TCP/IP connection over ATM between RISC System/6000s connected to an ATM switch. The configuration consists of 3 RISC System/6000s which are connected to an IBM 8260 Multiprotocol Intelligent Switching Hub. The RISC System/6000s were equipped either with a Turboway ATM 100 adapter or with a Turboway 155 adapter. The ATM switch had an ATM100 and an ATM155 blade installed. Because any available ATM switch can be used, as long as the switch supports the adapter speed of the connected RISC System/6000s, refer to the corresponding documentation for installation and configuration of the ATM switch.

An IP subnet over an ATM Network is called a Logical IP Subnet LIS. An LIS is analogous to a traditional LAN segment and consists of one ARP Server and several ARP clients. So you can define one RISC System/6000 as an ARP Server while the others are ARP clients. The RS/6000 defined as the server has all the capabilities of the client RS/6000s in the LIS.




Figure 7. RISC System/6000s Connected to the ATM Switch

5.1 Prerequisites

The IBM 8260 Multiprotocol Intelligent Switching Hub must be installed and configured. Because we are using Turboway ATM 100 and Turboway ATM 155 adapters, an ATM100 blade and an ATM155 blade must be installed in the IBM 8260 Multiprotocol Intelligent Switching Hub. For information on how to install and set up 8260 Hub, please refer to *SA33-0251-2 Multiprotocol Intelligent Switching Hub Installation and Operation Guide*, and also *Nways 8260 ATM 155 Mbps Flexible Concentration Module Installation and User's Guide*. All RISC System/6000s must have ATM adapters installed. All ATM adapters must be available and the fiber optic cables must be connected to the correct blade in the ATM switch. See Chapter 2, "Hardware Installation" on page 11 and Chapter 3, "Software installation" on page 13 for installing the ATM adapter and the microcode and for connecting the fiber optic cables. Make sure that the RISC System/6000s Turboway ATM 100 adapters are connected to the ATM100 blade

and the RISC System/6000s Turboway ATM 155 adapters are connected to the ATM155 blade.

5.2 Configuring the ARP Servers Network Interface

The ARP Server is the central point of an LIS which stores the ATM addresses and IP addresses of all ARP clients of the LIS in its ARP cache. At system boot each client sends its 20-byte ATM address to the responsible ARP server. The first 13 bytes of the address were determined by the switch. Therefore, if a client system is moved and rebooted, its IP address will now resolve to the new ATM address. In our configuration, system1 will be the ARP server.

Note: If you intend to define more than eight ATM interfaces on a RS/6000, see 7.1, "More Than Eight Interfaces" on page 31.

TCP/IP uses a network interface to talk to the ATM adapter hardware using the TCP/IP protocols. To add this network interface do the following on system1:

```
smit mkinetat
```

or

```
smit
->Communications Applications and Services
->TCP/IP
->Further Configuration
->Network Interfaces
->Network Interface Selection
->Add a Network Interface
->Add an ATM Network Interface
```

You will be presented with the following screen which should be filled in as below:

| Add an ATM Network Interface | |
|---|---------------------------------|
| Type or select values in entry fields. Press Enter AFTER making all desired changes. | |
| * INTERNET ADDRESS (dotted decimal) | [Entry Fields] [192.4.5.151] |
| Network MASK (hexadecimal or dotted decimal) | [255.255.255.0] |
| Network Interface | [at0] |
| Connection Type | svc_s |
| ATM Server Address | [] |
| Alternate Device | [atm0] |
| Idle Timer | [] |
| Best Effort Bit Rate (UBR) in Kbits/sec | [] |
| * ACTIVATE the Interface after Creating it? | yes |

Description of the input fields:

INTERNET ADDRESS

As for any TCP/IP interface, the Internet Address must be unique to this adapter and should be in dotted decimal format.

Network MASK

The Network Mask is an optional field but is usually the subnet mask assigned relative to this adapter.

Network Interface

The network Interface in this example is at0 as the adapter is atm0.

Connection Type

Specify the Virtual Connection type to be configured for this network interface. Following options are valid:

- pvc** Only Permanent Virtual Connections are supported on this interface.
- svc_s** The ARP server on the Switched Virtual Connection type.
- svc_c** The ARP client on the Switched Virtual Connection type. You must specify the ATM 20-byte address of the client's designated ARP server in the ATM Server Address field.

ATM Server Address

This field contains the 20-byte ATM address of the ARP Server. If this is an ARP Server or PVC, leave this field blank. If this is an ARP Client, you must enter the ARP Server ATM address.

Alternate Device

The Alternate Device field is the device name of the ATM adapter, (atm0), associated with this interface.

Idle Timer

Specify a decimal number representing the minutes that an SVC connection will stay up without any data traffic. The default is 60 minutes.

Best Effort Bit Rate (UBR) in Kbits/sec

This field is not available in AIX 4.1.4. It specifies the bit rate for the Best Effort connections that are initiated from the configuring interface. This network interface supports the Classical IP over ATM connections. All of the outgoing SVC connections will be set as Best Effort connections with the specified bit rate as the forward and backward peak rate. When this attribute is 0 (by default), the network interface uses the peak rate supported by the associated device as the Best Effort bit rate. Specify a decimal integer in K bits per second. For example, the value 155000 represents 155M bits per second. This field does not apply to PVC connections.

ACTIVATE the Interface after Creating it?

If yes, the added interface will be activated.

5.3 Getting the ARP Server Address

Before you configure the ARP clients, you have to get the ATM address of the ARP server. The ATM address is a 20-byte address assembled at the ATM switch. The first 13 bytes are the ATM address prefix type the switch implements. The next 6 bytes is the unique burned-in hardware MAC address of the ATM adapter at the RISC System/6000. The final byte, the selector byte, is ignored by the switch, but the AIX clients and Servers use it to identify the network interface.

When the network interface of the ARP server is being activated, the server sends its 6-byte MAC address to the switch and gets back the assembled 20-byte ATM address, which is loaded in the ARP cache of the server.

This address can be retrieved by issuing the following command on the ARP server if your ARP server is an RS/6000:

```
arp -t atm -a
```

For AIX V4.1.4 the output is:

```
PVC
      VPI:VCI      Handle IP Address      llc      arp
SVC
      IP Addr      ATM Address
NS1(192.4.5.151)  39.11.ff.22.99.99.99.0.0.0.0.1.49.8.0.5a.99.10.a9.0
```

For AIX V4.2.4 the output is:

```
PVC
===
      VPI:VCI Handle      IP Address      llc      arp

SVC - at0 on device atm0
=====
at0(192.4.5.151)      39.11.ff.22.99.99.99.0.0.0.0.1.49.8.0.5a.99.10.a9.0
  IP Addr      VPI:VCI Handle ATM Address
```

5.4 Configuring the ARP Client Network Interface

On each RISC System/6000 that wants to be part of the LIS, an ATM network interface must be configured. Each RISC System/6000 in an LIS must have a unique IP address and all clients in a LIS must have the same ATM server address configured.

Note: If you intend to define more than 8 ATM interfaces on an RS/6000, see 7.1, "More Than Eight Interfaces" on page 31.

To do so, do the following on each client, in our case system2 and system3:

```
smit mkinetat
```

or

```
smit
->Communications Applications and Services
->TCP/IP
->Further Configuration
->Network Interfaces
->Network Interface Selection
->Add a Network Interface
->Add an ATM Network Interface
```

You will be presented with the following screen which should be filled in as below:

```

                                Add an ATM Network Interface
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* INTERNET ADDRESS (dotted decimal)           [Entry Fields]
Network MASK (hexadecimal or dotted decimal) [192.4.5.58]
Network Interface                             [255.255.255.0]
Connection Type                               [at0]
ATM Server Address                            svc_c
Alternate Device                              [39.11.ff.22.99.99.0>]
Idle Timer                                    [atm0]
Best Effort Bit Rate (UBR) in Kbits/sec      []
* ACTIVATE the Interface after Creating it?   []
                                                yes
```

For a description of the input fields see 'Configuring the ARP-Servers Network Interface' in this chapter.

Note: The INTERNET ADDRESS for each system is different, so when you configure the network interface on system3 you have to enter 9.4.5.187 in this field.

The systems should now be able to communicate with each other via TCP/IP over ATM.

5.4.1 Testing the TCP/IP Connection over ATM

The standard way to test any TCP/IP configuration is to use ping. TCP/IP over ATM is no different. The output from the ping command on a successful TCP/IP over ATM configuration is as follows:

```

PING 9.4.5.151: (9.4.5.151): 56 data bytes
64 bytes from 9.4.5.151: icmp_seq=0 ttl=255 time=1 ms
64 bytes from 9.4.5.151: icmp_seq=1 ttl=255 time=1 ms
64 bytes from 9.4.5.151: icmp_seq=2 ttl=255 time=1 ms

----9.4.5.151 PING Statistics----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 1/1/1 ms
```

If the ping fails, refer to Chapter 9, "Troubleshooting" on page 37.

Chapter 6. Using ATM as a Backbone with RISC System/6000s

This chapter discusses the items you have to consider when you are connecting existing LANs over an ATM backbone.




Figure 8. RISC System/6000s as an ATM Backbone

As shown in the figure, the ATM backbone consists of an ATM switch and the RISC System/6000s named ATMSYS1, ATMSYS2 and ATMSYS3. ATMSYS1 is chosen to act as the ARP Server, while ATMSYS2 and ATMSYS3 are ARP clients. The ATM configuration is exactly the same as described in Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21. The ATM

backbone must have its own IP subnet, i.e. 192.4.5. In order to transfer IP packets from one LAN to another LAN, the 3 RISC System/6000s in the ATM backbone must also act as routers. So ATMSYS1, connected to the X.25 network acts as the router for the X.25 subnet, ATMSYS2 acts as the router for the Token Ring subnet and ATMSYS3 acts as the router for the ethernet subnet. To activate the router functionality on ATMSYS1, ATMSYS2 and ATMSYS3 you have to enter the following command on all 3 systems:

```
no -o ipforwarding=1
```

To activate the router automatically on each reboot you have to add this command to the end of the script /etc/rc.net. To accommodate communications between the different networks, you have to define static routes on each system for each not directly attached network. Routes to directly attached networks will be added automatically, when you configure the network interfaces. In the example for all network addresses the netmask 255.255.255.0 is used. See the TCP/IP documentation for details.

Note: Since Classical IP over ATM as defined in RFC1577 does not support protocols that resolve addresses using broadcast, routed and gated will not work across the ATM subnet. Static network routes must be issued.

To verify the routes, enter the following command on each system:

```
netstat -rn
```

The routing tables should look like the following:

on ATMSYS1

| | | | | | |
|---------|-------------|----|---|-----|-----|
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 192.4.5 | 192.4.5.151 | U | 1 | 42 | at0 |
| 9.4.5 | 9.4.5.12 | U | 1 | 42 | xs0 |
| 2.4.5 | 192.4.5.58 | UG | 1 | 42 | at0 |
| 19.4.5 | 192.4.5.187 | UG | 1 | 42 | at0 |

on ATMSYS2

| | | | | | |
|---------|-------------|----|---|-----|-----|
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 192.4.5 | 192.4.5.58 | U | 1 | 42 | at0 |
| 9.4.5 | 192.4.5.151 | UG | 1 | 42 | at0 |
| 2.4.5 | 2.4.5.15 | U | 1 | 42 | tr0 |
| 19.4.5 | 192.4.5.187 | UG | 1 | 42 | at0 |

on ATMSYS3

| | | | | | |
|---------|-------------|----|---|-----|-----|
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 192.4.5 | 192.4.5.187 | U | 1 | 42 | at0 |
| 9.4.5 | 192.4.5.151 | UG | 1 | 42 | at0 |
| 2.4.5 | 192.4.5.58 | UG | 1 | 42 | at0 |
| 19.4.5 | 19.4.5.5 | U | 1 | 42 | en0 |

on X25SYS1

| | | | | | |
|---------|-----------|----|---|-----|-----|
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 9.4.5 | 9.4.5.13 | U | 1 | 42 | xs0 |
| default | 9.4.5.12 | UG | 1 | 42 | xs0 |

on TRSYS1

| | | | | | |
|-----------|-----------|----|---|-----|-----|
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 2.4.5 | 2.4.5.1 | U | 1 | 42 | tr0 |
| default | 2.4.5.15 | UG | 1 | 42 | tr0 |
| on TRSYS2 | | | | | |
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 2.4.5 | 2.4.5.2 | U | 1 | 42 | tr0 |
| default | 2.4.5.15 | UG | 1 | 42 | tr0 |
| on TRSYS3 | | | | | |
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 2.4.5 | 2.4.5.3 | U | 1 | 42 | tr0 |
| default | 2.4.5.15 | UG | 1 | 42 | tr0 |
| on ENSYS1 | | | | | |
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 19.4.5 | 19.4.5.10 | U | 1 | 42 | en0 |
| on ENSYS2 | | | | | |
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 19.4.5 | 19.4.5.11 | U | 1 | 42 | en0 |
| default | 19.4.5.5 | UG | 1 | 42 | en0 |
| on ENSYS3 | | | | | |
| 127 | 127.0.0.1 | U | 3 | 285 | lo0 |
| 19.4.5 | 19.4.5.12 | U | 1 | 42 | en0 |
| default | 19.4.5.5 | UG | 1 | 42 | en0 |

Now you should be able to reach every system from every system via TCP/IP; ping away.

Chapter 7. Multiple Interfaces on the Same ATM Adapter

With ATM it is possible to configure more than one TCP/IP network interface on the same adapter. This is particularly useful, if you have multiple LISs.

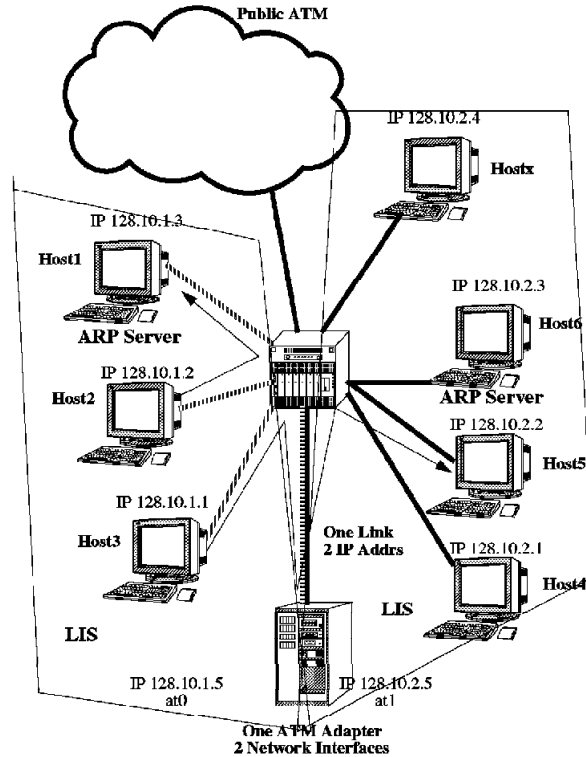


Figure 9. Routing TCP/IP on One Adapter

As shown in the figure, there are two IP networks and two ARP servers. One RS/6000 acts also as a router. For installing the ATM adapter and the software, see Chapter 2, "Hardware Installation" on page 11 and Chapter 3, "Software installation" on page 13.

7.1 More Than Eight Interfaces

The maximum number of interfaces allowed under AIX for any network type is defined by the `ifsize` option of the `no` command. The default is `ifsize=8`. This means 8 ethernets, 8 token rings, 8 fddi, 8 ATM, etc... When you change the `ifsize` you change the maximum number of interfaces allowed for all network types. With the default, valid ATM interfaces are `at0 - at7`. AIX supports 256 ATM interfaces on one RS/6000. If you intend to have more than 8 interfaces on an RS/6000, you must change the `ifsize`. You change the `ifsize` by:

```
no -o ifsize=16
```

In this example, the maximum number of ATM interfaces is set to 16. The number must be equal or greater than the number of interfaces you intend to configure on an ATM adapter. Valid ATM interfaces are now `at0 - at15`.

The `no` options are only valid until the next time the system is booted. After which, any `ifsize` is returned to the default. To avoid this, add the `no -o`

ifsize=xxx in the /etc/rc.net file before the section of the file that starts the interfaces. Do not add it at the end where the other no commands are shown.

7.2 Configure Multiple Network Interfaces on the Same ATM Adapter

To configure the first interface on the router, do the following:

```
smit mkinetat
```

or

```
smit
->Communications Applications and Services
->TCP/IP
->Further Configuration
->Network Interfaces
->Network Interface Selection
->Add a Network Interface
->Add an ATM Network Interface
```

You will be presented with the following screen which should be filled in as below:

| Add an ATM Network Interface | |
|---|--------------------------------|
| Type or select values in entry fields. Press Enter AFTER making all desired changes. | |
| * INTERNET ADDRESS (dotted decimal) | [Entry Fields] [128.10.1.5] |
| Network MASK (hexadecimal or dotted decimal) | [255.255.255.0] |
| Network Interface | [at0] |
| Connection Type | svc_c |
| ATM Server Address | [39.11.ff.33.9.a4.>.] |
| Alternate Device | [atm0] |
| Idle Timer | [] |
| Best Effort Bit Rate (UBR) in Kbits/sec | [] |
| * ACTIVATE the Interface after Creating it? | yes |

Description of the input fields:

INTERNET ADDRESS

As for any TCP/IP interface the Internet Address must be unique to this adapter and should be in dotted decimal format.

Network MASK

The Network Mask is an optional field but is usually the subnet mask assigned relative to this adapter.

Network Interface

Specify the name of the network interface that you are adding. The interface name for the ATM Classical IP network interface is atN, where N is a decimal number from 0 to 255 (for example, at1). See 7.1, "More Than Eight Interfaces" on page 31 to define more than 8 interfaces.

Connection Type

Specify the Virtual Connection type to be configured for this network interface. Following options are valid:

- pvc** Only Permanent Virtual Connections are supported on this interface.
- svc_s** The ARP server on the Switched Virtual Connection type.
- svc_c** The ARP client on the Switched Virtual Connection type. You must specify the ATM 20-byte address of the client's designated ARP server in the ATM Server Address field.

ATM Server Address

This field contains the 20-byte ARP address of the ARP server.

Alternate Device

The Alternate Device field is the device name associated with this interface (atm0).

Idle Timer

Specify a decimal number representing the minutes that an SVC connection will stay up without any data traffic. Specify a decimal number representing the minutes that an SVC connection will stay up without any data traffic. The default is 60 minutes.

Best Effort Bit Rate (UBR) in Kbits/sec

This field is not available in AIX 4.1.4. It specifies the bit rate for the Best Effort connections that are initiated from the configuring interface. This network interface supports the Classical IP over ATM connections. All of the outgoing SVC connections will be set as Best Effort connections with the specified bit rate as the forward and backward peak rate. When this attribute is 0 (by default), the network interface uses the peak rate supported by the associated device as the Best Effort bit rate. Specify a decimal integer in K bits per second. For example, the value 155000 represents 155M bits per second. This field does not apply to PVC connections.

To configure the second interface on the router, do the following:

```
smit mkinetat
```

or

```
smit
```

```
->Communications Applications and Services
```

```
->TCP/IP
```

```
->Further Configuration
```

```
->Network Interfaces
```

```
->Network Interface Selection
```

```
->Add a Network Interface
```

```
->Add an ATM Network Interface
```

You will be presented with the following screen which should be filled in as below:

```

                                Add an ATM Network Interface
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

* INTERNET ADDRESS (dotted decimal)           [Entry Fields]
Network MASK (hexadecimal or dotted decimal) [128.10.2.5]
Network Interface                             [255.255.255.0]
Connection Type                               [at1]
ATM Server Address                            svc_c
Alternate Device                              [38.2a.7f.31.9.a4.>.]
Idle Timer                                    [atm0]
Best Effort Bit Rate (UBR) in Kbits/sec      []
* ACTIVATE the Interface after Creating it?   [yes]
```

The value for the Network Interface field can be at0 to at7. Because you can't specify a name twice, and at0 is already specified, you may enter at1. For the Alternate Device field enter the device name of the adapter you want to use. In this example atm0.

See Chapter 6, "Using ATM as a Backbone with RISC System/6000s" on page 27 to implement routing.

Chapter 8. An Approach to a Backup ARP Server

To support IP over ATM Switched Virtual Connections, a single ARP server must be located within the LIS. This ARP server keeps all ATM addresses of all available clients in this LIS. Each client who wants to connect to another client has to contact the ARP server for address resolution of the other client. Since only one ARP server is allowed, it becomes a possible single point of failure along with the switch. Since an HACMP solution is beyond the scope of this redbook, this chapter only discusses a possible solution. It's up to the HACMP specialists to implement this solution.

8.1 Address Resolution

Address resolution within an ATM logical IP subnet makes use of the ATM Address Resolution Protocol (ATMARP) and the Inverse ATM Address Resolution Protocol (InATMARP). Each ARP client in an LIS connects to the ARP server using a VCC. Once the connection between the ARP client and the ARP server is established, the server transmits an InARP_REQUEST to determine the IP address of the ARP client. The InARP_REPLY from the client contains the information necessary for the ARP server to build its ARP table cache. This information is used by the server to generate replies to the ARP requests it receives from the clients. After registration with the ARP server, the ARP client may send ARP requests to the ARP server in order to find out the ATM address of other ARP clients when it needs to communicate with them. The ARP client adds the ATM addresses received by the ARP server to its ARP table cache. The next time it wants to communicate with the other ARP server, it checks its ARP table cache for an entry and, if present, connects to the destination client without asking the ARP server. The ARP cache entries of the ARP server are aging and deleted after 20 minutes. To avoid this, the ARP server generates an InARP_REQUEST on any open VCC associated with an entry, prior to aging and deleting the entry.

The ARP server mechanism requires that each ARP client be administratively configured with the ATM address of the ARP server. This is done when the network interface is configured.

8.2 ARP Server

The ARP server may be configured as usual. It should have a separate connection for the heartbeat to the backup ARP server.

8.3 Backup ARP Server

The backup ARP server must be first configured as an ARP client, so it is reachable by all ARP clients in the LIS. All available ARP clients must have an ARP cache entry. The backup ARP server also needs the following scripts or programs:

- an RPC program that calls on each available client a corresponding RPC program.
- a script to change the network interface to act as an ARP server.

8.4 ARP Clients

The ARP clients may be configured as usual.

Since the backup ARP server needs an ARP cache entry for every potential client, the clients have to ping the backup ARP server often enough to prevent the entries from being deleted from the backup servers ARP cache. The default time for the ARP cache entry aging is 20 minutes. During the takeover phase, when the ARP server is no longer available and the backup ARP server is not yet available, a connection request to clients not present in the ARP cache results in invalid entries i.e. (192.4.5.23) 0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0. The clients also need the following scripts or programs:

- a script that pings the backup ARP server every 10 minutes, so the ARP table entry on the backup ARP server will never be deleted.
- a script that deletes invalid ARP cache entries
- a script to change the ATM server address
- an RPC program that can be called by the backup server and that calls the script.

8.5 Takeover Situation

In case the ARP server fails, the backup ARP server will be informed via the heartbeat line.

On the backup ARP server:

1. the IP addresses of all available ARP clients are extracted from the ARP cache.
2. All available ARP clients are notified via RPC program or another procedure.
3. The Network Interface will be changed to act as an ARP server.

On the ARP clients:

1. The ARP client gets the notification from the backup ARP server via RPC or some other method.
2. The ARP server address is changed to the backup ARP servers address.
3. All invalid ARP cache entries are deleted.

Chapter 9. Troubleshooting

In this section we will describe our experiences determining the causes of problems.

9.1 Helpful Commands

Several commands are available with AIX to get information for troubleshooting IP connections over ATM. These commands are:

- `errpt` to view error report
- `iptrace` to trace IP over ATM
- `tcpdump` to view iptrace output
- `atmstat` to view ATM statistics
- `arp` to view the ARP table
- `ps` to see which ATM processes are running
- `lsdev` to list ATM devices and interfaces
- `lsattr` to list ATM configuration information

Note: ATM support for `iptrace` and `tcpdump` is not available under AIX 4.1.4.

9.1.1 `errpt`

`errpt` generates a report from the errors detected by operating system modules and collected in the system errorlog. For ATM several new Error Labels were added. These Error Labels are:

| | |
|------------------------------|--|
| <code>ATM_DMAFAIL</code> | DMA Error |
| <code>ATM_FAIL</code> | Microcode Program abnormally terminated |
| <code>ATM_LIGHT_LOST</code> | Wire Fault |
| <code>ATM_LOAD_UCODE</code> | Failed loading microcode |
| <code>ATM_NOBUFS_PERM</code> | Error allocating memory |
| <code>ATM_UCODE</code> | Adapter Error |
| <code>ATM_PIOFAIL</code> | This Error is logged if a programming input or output exception has occurred, indicating possible adapter failure. |

To get a list of all errors concerning ATM enter the following command:

```
errpt | grep atm
```

You can also view `errpt` output using SMIT. If there were already some errors, the list could look like this:

| IDENTIFIER | TIMESTAMP | T | C | RESOURCE_NAME | DESCRIPTION |
|------------|------------|---|---|---------------|-------------|
| 2FB174CD | 1014182796 | T | H | atm0 | DMA ERROR |
| 26492630 | 1010153996 | T | H | atm0 | WIRE FAULT |
| 26492630 | 1010084096 | T | H | atm0 | WIRE FAULT |
| 26492630 | 1009171796 | T | H | atm0 | WIRE FAULT |
| 26492630 | 1009155796 | T | H | atm0 | WIRE FAULT |
| 26492630 | 1009151996 | T | H | atm0 | WIRE FAULT |

To take a closer look at a particular error i.e the DMA ERROR enter the following:

```
errprt -a -j 2FB174CD
```

which results in:

```
LABEL:          ATM_DMAFAIL
IDENTIFIER:     2FB174CD
```

```
Date/Time:      Mon Oct 14 18:27:51
Sequence Number: 347
Machine Id:     000141023000
Node Id:       praben1
Class:         H
Type:          TEMP
Resource Name:  atm0
Resource Class: adapter
Resource Type: atm1
Location:      00-04
VPD:
```

```
Part Number.....10H4538
EC Level.....D50653
Serial Number.....008355R6
FRU Number.....73G9821
Manufacturer.....IBM982
Diagnostic Level.....01
ROS Level and ID.....(base)014210
ROS Level and ID.....(daughter) 104
Displayable Message.....ATM 100Mbps Fiber Adapter
Device Driver Level.....01
Loadable Microcode Level....01
Network Address.....08005A9910A9
```

Description
DMA ERROR

Probable Causes
CONTROL PROGRAM
LOCAL CSMA/CD ADAPTER

Failure Causes
SYSTEM I/O BUS
COMMUNICATIONS SUBSYSTEM CONTROLLER
COMMUNICATIONS SUBSYSTEM CONTROLLER INTERFACE

Recommended Actions
PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data
FILE NAME
line: 2174 file: atmE-87-!P
SENSE DATA
0000 0001 0000 0000
DEVICE ADDRESS
0000 0000
SENSE DATA

For a further description of errprt use the InfoExplorer facility of AIX.

9.1.2 iptrace and tcpdump

The iptrace and tcpdump commands provide interface level packet tracing for Internet Protocols.

For a further description of errpt use the InfoExplorer facility of AIX. Note, that iptrace and tcpdump do not support ATM on AIX 4.1.4.

9.1.3 atmstat

The atmstat command provides information about an ATM adapter. The command can be entered in 3 ways:

- | | |
|------------------------|--|
| atmstat atm0 | Lists the Transmit Statistics, the Receive Statistics and the General Statistics of adapter atm0. |
| atmstat -d atm0 | Lists the Transmit Statistics, the Receive Statistics, the General Statistics and the Turboways ATM Adapter Specific Statistics of adapter atm0. |
| atmstat -r atm0 | Lists then resets the Transmit Statistics, the Receive Statistics and the General Statistics of adapter atm0. |

Following is an example output from the command atmstat -d atm0:

```
-----  
ATM STATISTICS (atm0) :  
Device Type: Turboways 100 MCA ATM Adapter  
Hardware Address: 08:00:5a:99:10:a9  
Elapsed Time: 0 days 0 hours 22 minutes 51 seconds  
  
Transmit Statistics:                Receive Statistics:  
-----  
Packets: 452                        Packets: 452  
Bytes: 19160                        Bytes: 18682  
Interrupts: 0                       Interrupts: 527  
Transmit Errors: 0                  Receive Errors: 0  
Packets Dropped: 0                 Packets Dropped: 0  
Max Packets on S/W Transmit Queue: 0  Bad Packets: 0  
S/W Transmit Queue Overflow: 0  
Current S/W+H/W Transmit Queue Length: 0  
  
Cells Transmitted: 721              Cells Received: 724  
Out of Xmit Buffers: 0              Out of Rcv Buffers: 0  
Current HW Transmit Queue Length: 0  CRC Errors: 0  
Current SW Transmit Queue Length: 0  Packets Too Long: 0  
                                      Incomplete Packets: 0  
                                      Cells Dropped: 0  
  
General Statistics:  
-----  
No mbuf Errors: 0  
Adapter Loss of Signals: 0  
Adapter Reset Count: 0  
Driver Flags: Up Running Simplex  
Virtual Connections in use: 5  
Max Virtual Connections in use: 5  
Virtual Connections Overflow: 0  
  
Turboways ATM Adapter Specific Statistics:  
-----
```

Packets Dropped - No small DMA buffer: 0
Packets Dropped - No medium DMA buffer: 0
Packets Dropped - No large DMA buffer: 0
Receive Aborted - No Adapter Receive Buffer: 0
Transmit Aborted - No small DMA buffer: 0
Transmit Aborted - No medium DMA buffer: 0
Transmit Aborted - No large DMA buffer: 0
Transmit Aborted - No MTB DMA buffer: 0
Transmit Aborted - No Adapter Transmit Buffer: 0
Max Hardware transmit queue length: 4
Small Mbuf in Use: 0
Medium Mbuf in Use: 0
Large Mbuf in Use: 64
Huge Mbuf in Use: 0
MTB Mbuf in Use: 0
Max Small Mbuf in Use: 0
Max Medium Mbuf in Use: 0
Max Large Mbuf in Use: 64
Max Huge Mbuf in Use: 0
MTB Mbuf in Use: 0
Small Mbuf overflow: 0
Medium Mbuf overflow: 0
Large Mbuf overflow: 0
Huge Mbuf overflow: 0
MTB Mbuf overflow: 0

The following is a description of each field of the output of the atmstat command.

Note: Some fields are only present on AIX 4.2 and are marked as such.

Transmit Statistics:

Packets

This field contains the number of packets (or PDUs) transmitted.

Bytes

This field contains the count of bytes transmitted. These are the user bytes. The ATM overhead, for example, ATM cell header, AAL5 PDU trailer, etc., are excluded.

Interrupts

This field is not used.

Transmit Errors

This field contains the number of Transmit Errors for this device.

Packets Dropped

This field contains the number of Transmit Packets that were dropped, for instance, due to out of buffer condition.

Max Packets on S/W Transmit Queue

This field is not used.

S/W Transmit Queue Overflow

This field is not used.

Current S/W + H/W Transmit Queue Length

This field contains the current transmit queue length.

Cells Transmitted

This field contains the number of cells transmitted.

Out of Xmit Buffers

This field contains the number of packets dropped because of out-of-xmit buffers condition.

Current HW Transmit Queue Length

This field contains the number of transmit packets on the hardware queue.

Current SW Transmit Queue Length

This field is not used.

Receive Statistics:

Packets

This field contains the number of packets (or PDUs) received.

Bytes

This field contains the count of bytes received. These are the user bytes. The ATM overhead, for example, ATM cell header, AAL5 PDU trailer, etc., are excluded.

Interrupts

This field contains the number of Interrupts taken by the system for the Adapter-to-System indications. Some of the events that cause these interrupts are packets received, transmit done indication, and so on.

Receive Errors

This field contains the number of Receive Errors for this device.

Packets Dropped

This field contains the number of Received Packets dropped, for instance, due to out-of-buffer condition.

Bad Packets

This field is not used.

Cells Received

This field contains the number of cells received by this device.

Out of Rcv Buffers

This field contains the number of packets dropped because of out-of-receive buffers condition.

CRC Errors

This field contains the number of received packets that encountered CRC errors.

Packets Too Long

This field contains the number of received packets that exceeded the maximum PDU size.

Incomplete Packets

This field contains the number of incomplete received packets.

Cells Dropped

This field contains the number of dropped cells. Cells could be dropped for a number of reasons, such as bad HEC, out of buffer condition, and so on.

General Statistics:

No mbuf Errors

This field contains the number of mbuf requests that were denied.

Adapter Loss of Signals

This field contains the number of times the adapter encountered Loss of Signal.

Adapter Reset Count

This field contains the number of times the adapter has been reset.

Driver Flags

This field contains the NDD (Network Device Driver) flags. i.e. Up Running Simplex

Virtual Connections in use

This field contains the number of VCs that are currently allocated or in use.

Max Virtual Connections in use

This field contains the maximum number of VCs allocated since the last reset of the statistics.

Virtual Connections Overflow

This field contains the number of allocate VC requests that have been denied.

Additional ATM Statistics:

Packets Dropped - No small DMA buffer

This field contains the number of received packets dropped because the adapter did not have small system buffers for DMA.

Packets Dropped - No medium DMA buffer

This field contains the number of received packets dropped because the adapter did not have medium system buffers for DMA.

Packets Dropped - No large DMA buffer

This field contains the number of received packets dropped because the adapter did not have large system buffers for DMA.

Receive Aborted - No Adapter Receive buffer

This field contains the number of received packets aborted because the adapter did not have on-card receive buffers.

Transmit Aborted - No small DMA buffer

This field contains the number of transmit packets aborted because of the lack of small system buffers for DMA.

Transmit Aborted - No medium DMA buffer

This field contains the number of transmit packets aborted because of the lack of medium system buffers for DMA.

Transmit Aborted - No large DMA buffer

This field contains the number of transmit packets aborted because of the lack of large system buffers for DMA.

Transmit Aborted - No Adapter Transmit buffer

This field contains the number of transmit packets aborted because of the lack of adapter on-card transmit buffers.

Max Hardware Transmit Queue Length

This field contains the maximum number of transmit packets queued in the hardware queue.

Small Mbufs in Use

This field contains the number of small mbufs currently in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Medium Mbufs in Use

This field contains the number of medium mbufs currently in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Large Mbufs in Use

This field contains the number of large mbufs currently in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Huge Mbufs in Use (AIX 4.2)

This field contains the number of huge mbufs currently in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Max Small Mbufs in Use

This field contains the maximum number of small mbufs that have been in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Max Medium Mbufs in Use

This field contains the maximum number of medium mbufs that have been in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Max Large Mbufs in Use

This field contains the maximum number of large mbufs that have been in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Max Huge Mbufs in Use (AIX 4.2)

This field contains the maximum number of huge mbufs that have been in use. The adapter device driver allocates these buffers according to the configuration information provided by system administrators. This information can be used to tune the configuration information.

Small Mbufs overflow

This field contains the number of times that a small mbuf could not be allocated.

Medium Mbufs overflow

This field contains the number of times that a medium mbuf could not be allocated.

Large Mbufs overflow

This field contains the number of times that a large mbuf could not be allocated.

Huge Mbufs overflow (AIX 4.2)

This field contains the number of times that a huge mbuf could not be allocated.

9.2 Troubleshooting

Here is a guideline to solve problems concerning ATM. In most cases these procedures should provide a solution. Choose the problem you have and follow the instructions in sequence.

9.2.1 Back-to-Back Configuration

Problem: No IP connection between the two systems.

1. Log in on both systems and check if the ATM adapter is available.

```
lsparent -C -k atm -H
```
2. You should get an output like this:

```
name status    location description
atm0 Available 00-04    100 Mbps ATM Fiber Adapter (8f7f)
```
3. If the status is define, check that the software is installed correctly. See Chapter 3, "Software installation" on page 13.
4. If there is no adapter listed, check that the adapter is installed correctly. See Chapter 2, "Hardware Installation" on page 11.
5. If the adapter has the status of available, verify that the network connection type is set to pvc.

```
lsdev -Cc if
```
6. You should get an output like this:

```
lo0 Available Loopback Network Interface
tr0 Available Token Ring Network Interface
at0 Available ATM Network Interface
```
7. If there is no ATM Network Interface listed, see Chapter 4, "Connecting Two RISC System/6000s Back-to-Back" on page 15 to add one.
8. If there is an ATM Network Interface listed, but the status is stopped, enter the command:

```
chdev -l at0 -a state=up
```
9. If there is an ATM Network Interface listed and the status is available, check that the Connection Type is set to PVC, the INTERNET ADDRESS is correct and, if specified, that the Network MASK is correct. If you have multiple ATM adapter installed and multiple network interfaces to the adapter configured, check also, that the correct Alternate Device is selected.

```
smit chinet
```

or

```
smit
->Communications Applications and Services
->TCP/IP
->Further Configuration
->Network Interfaces
->Network Interface Selection
->Change/Show Characteristics of a Network Interface
```

10. Check if there is a PVC configured.

```
mkatmpvc -S
```

SMIT can also be used to show a PVC.

11. You should get an output like this:

| #Interface | VPI:VCI | IP Address | Description |
|------------|---------|------------|-------------|
| at1 | 0:33 | 123.6.7.9 | Test PVC |

12. If there is no PVC listed, see Chapter 4, "Connecting Two RISC System/6000s Back-to-Back" on page 15 and add a PVC.

13. Check if there is an entry for the PVC in the ARP cache.

```
arp -t atm -a
```

14. You should get an output like this:

```
PVC
      VPI:VCI      Handle IP Address      llc   arp
      0:33         3      (123.6.7.9)      True  True
```

Note the IP address is the address of the other machine. If there is no connection, the IP address may be 0.0.0.0.

15. If there are no entries in the ARP cache, check the cabling between the system and the switch and check the system error log for ATM errors with the `errpt` command mentioned at the beginning of this chapter or reverse the connection at one end.

16. If entries are listed as shown above, the system should be correctly installed and configured.

9.2.2 Switched Configuration

Problem: The ATM network doesn't work at all. All systems in an LIS can't connect via IP to each other.

1. Check if the switch is powered on and operational. See the appropriate manual of the switch for how to check this. Normally there are indicators like green LEDs.
 - a. Check if all systems are properly connected to the switch.
 - b. Check if you have an ARP server defined for this LIS.
 - c. Check if the ARP server is correctly installed and configured.
2. Log in to the ARP server and check if the ATM adapter is available.

```
lsparent -C -k atm -H
```

You should get an output like this:

```
name status    location description
```

```
atm0 Available 00-04    100 Mbps ATM Fiber Adapter (8f7f)
```

3. If the status is defined, check that the software is installed correctly. See Chapter 3, "Software installation" on page 13.
4. If there is no adapter listed, check that the adapter is installed correctly. See Chapter 2, "Hardware Installation" on page 11.
5. If the adapter has the status available, check if a network interface for an ARP server is configured.

```
lsdev -Cc if
```

You should get an output like this:

```
lo0 Available Loopback Network Interface
tr0 Available Token Ring Network Interface
at0 Available ATM Network Interface
```

If there is no ATM Network Interface listed, see Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21 to add one.

6. If there is an ATM Network Interface listed, but the status is stopped, enter the command:

```
chdev -l at0 -a state=up
```
7. If there is an ATM Network Interface listed and the status is available, check that the Connection Type is set to `svc_s`, the INTERNET ADDRESS is correct and, if specified, that the Network MASK is correct. If you have multiple ATM adapters installed and multiple network interfaces to the adapter configured, check also, that the correct Alternate Device is selected.

```
smit chinet
```

or

```
smit
```

```
->Communications Applications and Services
->TCP/IP
->Further Configuration
->Network Interfaces
->Network Interface Selection
->Change/Show Characteristics of a Network Interface
```

Check if the ARP server has an entry for this network interface in its ARP cache.

```
arp -t atm -a
```

You should get an output like this:

```
SVC - at0 on device atm0 -
=====
at0(192.4.5.151)      39.11.ff.22.99.99.99.0.0.0.0.1.49.8.0.5a.99.10.a9.0
```

Note that the output is slightly different between AIX 4.1.4 and AIX 4.2, but both should show an entry with the local IP address and a valid ATM address.

8. If the output does not show an entry for the local IP address, check the cabling between the system and the switch and check the system error log for ATM errors with the `errpt` command mentioned at the beginning of this chapter.
9. If an entry is listed as shown above, the ARP server should be correctly installed and configured.
10. Check that the clients are correctly installed and configured.

11. Log in each ARP client and check if the ATM adapter is available.

```
lsparent -C -k atm -H
```

You should get an output like this:

```
name status    location description
```

```
atm0 Available 00-04    100 Mbps ATM Fiber Adapter (8f7f)
```

12. If the status is defined, check that the software is installed correctly. See Chapter 3, "Software installation" on page 13.

13. If there is no adapter listed, check that the adapter is installed correctly. See Chapter 2, "Hardware Installation" on page 11.

14. If the adapter has the status available, check if a network interface for an ARP client is configured.

```
lsdev -Cc if
```

You should get an output like this:

```
lo0  Available          Loopback Network Interface
tr0  Available          Token Ring Network Interface
at0  Available          ATM Network Interface
```

15. If there is no ATM Network Interface listed, see Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21 to add one.

16. If there is an ATM Network Interface listed, but the status is stopped, enter the command:

```
chdev -l at0 -a state=up
```

17. If there is an ATM Network Interface listed and the status is available, check that the Connection Type is set to svc_c, the INTERNET ADDRESS is in the same LIS as the ARP server, if specified, that the Network MASK is correct and that you entered the correct ATM Server Address. For the ATM Server Address see Chapter 5, "Connecting RISC System/6000s with an ATM Switch" on page 21. If you have multiple ATM adapters installed and multiple network interfaces to the adapter configured, check also, that the correct Alternate Device is selected.

```
smit chinet
```

or

```
smit
```

```
->Communications Applications and Services
```

```
->TCP/IP
```

```
->Further Configuration
```

```
->Network Interfaces
```

```
->Network Interface Selection
```

```
->Change/Show Characteristics of a Network Interface
```

18. Check if the ARP client has an entry for this network interface and an entry with the ARP servers IP address in its ARP cache.

```
arp -t atm -a
```

You should get an output like this:

```
SVC - at0 on device atm0 -
```

```
=====
```

```
at0(192.4.5.187)          39.11.ff.22.99.99.99.0.0.0.0.1.49.8.0.5a.99.c.ca.0
  IP Addr          VPI:VCI Handle ATM Address
?(192.4.5.151)          0:548 3 39.11.ff.22. --- .8.0.5a.99.10.a9.0
```

Note that the output is slightly different between AIX 4.1.4 and AIX 4.2, but both should show an entry with the local IP address and a valid ATM address. The local IP address in this example is 192.4.5.187 and the ARP server's IP address is 192.4.5.151.

19. If there are no entries in the ARP cache, check the cabling between the system and the switch and check the system error log for ATM errors with the `errpt` command mentioned at the beginning of this chapter.

If entries are listed as shown above, the ARP client should be correctly installed and configured.

Problem: One system can't make IP connections to other systems while the other systems can connect to each other.

For problem determination, see the preceding problem.

Chapter 10. Performance and Tuning

In this chapter we will look at methods available to increase the throughput of ATM networks.

10.1 Overview

The ATM adapters can perform quite close to their designed limits, provided that they are used in a properly configured environment with resources capable of sustaining the desired throughputs. While the media speed may be 100 or 155 Mbps, there are a number of factors that will determine the final achieved maximum throughput observed by a user of an ATM adapter. One is the overhead inherent to ATM, such as the 5-byte ATM header that accompanies every 48 bytes of data sent. Others are inherent in the protocol used in the communication layers above ATM, like TCP/IP. Still others are dependent upon the processor speed and operating system used by the adapter host system.

The end-to-end throughput observed by a user at the application level will typically be that of the device or process (hardware or software) with the lowest throughput capacity in the path, also referred to as the bottleneck. With the ATM media speeds of 100 and 155Mbps the bottleneck will generally not be the network but some component of the workstation. The processor speed of the workstation in which the adapter is installed can have a dramatic effect on throughput. Less powerful machines are generally incapable of delivering frames at media speeds. Effects of processor speed on TURBOWAYS 100 adapter throughput are shown in Figure 10 on page 55.

Most RS/6000s can maintain maximum ATM throughput. A model 59H has sufficient processor power to handle a full complement of ATM adapters.

Since Micro Channel has a maximum throughput of 40 to 160 MBps it will not become a bottleneck, as long as multiple ATM adapters are installed.

For the end-to-end throughput, disk drives are likely to be a serious bottleneck when transferring data residing on a disk. The transfer rates of some disk drives are far lower than the media speed of the TURBOWAYS 155 adapter. Among the most important factors impacting final throughput are upper layer protocols such as TCP/IP. See Figure 11 on page 56. To investigate and eliminate bottlenecks may be quite complex in nature and is beyond the scope of this book. However this redbook will provide some tuning guidelines concerning TCP/IP and the ATM adapter.

10.2 Tuning TCP/IP

See *AIX Performance Tuning Guide*, SR28-5930 for information on tuning AIX systems and network interfaces.

Several TCP/IP parameters can be changed to increase throughput. These changes should be entered in the `/etc/rc.net` file if you want the changes to be permanent. The changes affect **all** IP interfaces, not just ATM.

Increase the TCP sendspace to 56k or larger:

```
no -o tcp_sendspace=57344
```

Increase the TCP receivespace to 56k or larger:

```
no -o tcp_receivespace=57344
```

Increase the socket buffer space:

```
no -o sb_max=2*57344
```

If you set `tcp_receivespace` and `tcp_receivespace` to more than 64k, you have to enable TCP Extensions for High Performance RFC1323:

```
no -o rfc1323=1
```

Note: The `sendspace` and `receivespace` is taken from your main memory. Choose the values depending on the amount of main memory you have installed and on the RS/6000 model. For example, if you have an RS/6000 Mod 520 with 64MB main memory, use the values above. For an RS/6000 Mod 590 with 512MB main memory you can set `tcp_sendspace` and `tcp_receivespace` to 128k and `sb_max` to 256k.

10.3 Tuning the ATM Adapter Parameters

In most cases the default values of the adapter are well defined and shouldn't be changed. However from time to time you should observe if your adapter parameters still meet the needs of your traffic. To do this, use the `atmstat -d` command, described in detail in Chapter 9, "Troubleshooting" on page 37.

Following values of the Turboways ATM Adapter Specific Statistics are interesting:

- Packets Dropped - No small DMA buffer
- Packets Dropped - No medium DMA buffer
- Packets Dropped - No large DMA buffer
- Receive Aborted - No Adapter Receive Buffer
- Transmit Aborted - No small DMA buffer
- Transmit Aborted - No medium DMA buffer
- Transmit Aborted - No large DMA buffer
- Transmit Aborted - No MTB DMA buffer
- Transmit Aborted - No Adapter Transmit Buffer
- Small Mbuf overflow
- Medium Mbuf overflow
- Large Mbuf overflow
- Huge Mbuf overflow
- MTB Mbuf overflow

If one or more of these values are nonzero, the corresponding parameter should be increased. For example if the value of Medium Mbuf overflow is nonzero the adapter parameter 'Maximum Medium ATM mbufs' should be increased. If you increase the maximum number of mbufs, the 'DMA bus memory width' parameter must be recomputed and increased accordingly. An increase is not normally required. An error will be detected if this parameter is too small.

Note: The mbufs are located in the main memory, increasing the maximum number of mbufs will decrease the amount of main memory available for user processes. Only increase the values if there is really a need for it.

To change adapter parameters, all interfaces using this adapter must be stopped. Before you stop a network interface, make sure that no user process uses this interface.

Note: It may be that the system has to be rebooted in order to get the network interfaces active again.

To stop network interface at0 use the following command:

```
ifconfig at0 detach
```

Repeat this command for all network interfaces on this adapter.

The adapter parameters are changed using SMIT:

```
smit chg_atm
```

Select the ATM adapter you wish to change and the following SMIT screen is presented:

```
Change / Show Characteristics of an ATM Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

ATM Adapter          [Entry Fields]
Description          atm0
Status              155 Mbps ATM Fiber Ada>
Location            Defined
User Best Effort Peak Rate (kbits/sec) 00-03
Enable ALTERNATE ATM MAC address       [1500]
ALTERNATE ATM MAC address (12 hex digits) no
ATM Adapter Maximum PDU size (bytes)   [9180]      +#
DMA bus memory width                    [0x1000000]  +X
Maximum Small ATM mbufs                  [50]        +#
Maximum Medium ATM mbufs                 [100]       +#
Maximum Large ATM mbufs                  [300]       +#
Maximum Huge ATM mbufs                   [50]        +#
Maximum ATM MTB mbufs                    [4]         +#
Max Transmit Block (MTB) size (kbytes)   [64]        +#
Minimum Small ATM mbufs                   [20]        +#
Minimum Medium ATM mbufs                  [30]        +#
Minimum Large ATM mbufs                   [70]        +#
Minimum Huge ATM mbufs                    [5]         +#
Minimum ATM MTB mbufs                     [4]         +#
Minimum ATM interface buffer size        [2048]      +#
Minimum Guaranteed VCs Supported         [16]        +#
Maximum Number of VCs Needed             [1000]     +#
Sonet or SDH interface                    [1]         +#
Provide SONET clock                       [1]         +#
```

Following is a description of the fields. Note that some fields are only available for the ATM 155Mbps adapter.

ATM Adapter

Identifies the logical name of an adapter.

Description

Provides a short text description of the adapter.

Status

Indicates the current status of the adapter. Possible values are available, indicating that the adapter is configured in the system and

ready to use, and defined, indicating that the adapter is defined to the system but not configured.

Location

The location code for an adapter consists of two pairs of digits with the format AA-BB where AA identifies the location code of the drawer containing the adapter card and BB identifies both the I/O bus and slot containing the card.

User Best Effort Peak Rate (kbits/sec)

Best Effort is for traffic that does not require quality of service and bandwidth reservation. The ATM network does not reject a Best Effort connection. There are 4 best effort rate queues, one of which is user-definable. Use this option to set the user-definable rate queue for Best Effort connections. The 3 predefined values for the ATM 100 adapter accommodate 25Mbps, 45Mbps, and 100Mbps. For the ATM 155 adapter, the values are 25Mbps, 100Mbps and 155Mbps. The adapter will automatically make connections using the fastest of the 4 queues that does not exceed the requested rate.

Enable ALTERNATE ATM MAC address

If yes, will use the user-defined MAC address rather than the address found in the read-only memory of the adapter. This is most useful when changing the card on which an ARP server was running.

ALTERNATE ATM MAC address (12 hex digits)

The alternate MAC address for the adapter to be used in place of the address found in the read-only memory of the adapter:

0x0123456789ab

ATM Adapter Maximum PDU size (bytes)

This is the maximum AAL5 PDU (protocol data unit) size. This is used to select the adapter buffer size. Setting this value larger than what is needed by the network applications will result in wasted adapter buffer space. Most installations will perform well with the default of 9180.

DMA bus memory width

The DMA bus width specifies the amount of bus space that can be used when preparing memory for DMA access from the ATM adapter. This value can be determined by adding the products of the size and maximum number of each ATM buffer size as configured. The DMA bus width for the device must be greater than or equal to the sum. Use 4 kB as the minimum buffer size when calculating this value.

Maximum Small ATM mbufs

The maximum number of small ATM mbufs this adapter should allocate. The small buffers are 256 bytes.

Maximum Medium ATM mbufs

The maximum number of medium ATM mbufs this adapter should allocate. The medium buffers are 4kB.

Maximum Large ATM mbufs

The maximum number of large ATM mbufs this adapter should allocate. The large buffers are 8kB.

Maximum Huge ATM mbufs

The maximum number of huge ATM mbufs this adapter should allocate. The huge buffers are 16kB.

Maximum Special ATM mbufs

The maximum number of special ATM mbufs this adapter should allocate. The special buffers are 64kB. Special mbufs are not supported on AIX V4.1.4.

Maximum ATM MTB mbufs

The maximum number of ATM MTB mbufs this adapter should allocate. This parameter also determines the maximum transmit block size.

Max Transmit Block (MTB) size (kbytes)

The maximum transmit block mechanism allows the transfer of data blocks that are larger than the PDU size from the system to adapter. These blocks are then segmented by the adapter. This method reduces the host processor utilization.

Minimum Small ATM mbufs

The minimum number of small ATM mbufs; the adapter will allocate this many buffers at configuration and may free buffers if more than the minimum number of buffers is unused. Small ATM mbufs are 256 bytes.

Minimum Medium ATM mbufs

The minimum number of medium ATM mbufs; the adapter will allocate this many buffers at configuration and may free buffers if more than the minimum number of buffers is unused. Medium ATM mbufs are 4kB.

Minimum Large ATM mbufs

The minimum number of large ATM mbufs; the adapter will allocate this many buffers at configuration and may free buffers if more than the minimum number of buffers is unused. Large ATM mbufs are 8kB.

Minimum Huge ATM mbufs

The minimum number of huge ATM mbufs; the adapter will allocate this many buffers at configuration and may free buffers if more than the minimum number of buffers is unused. Huge ATM mbufs are 16kB

Minimum Special ATM mbufs

The minimum number of special ATM mbufs this adapter should allocate. The Special buffers are 64kB. Special mbufs are not supported on AIX V4.1.4.

Minimum ATM MTB mbufs

The minimum number of ATM MTB mbufs; the adapter will allocate this many buffers at configuration and may free buffers if more than the minimum number of buffers is unused. The ATM MTB buffer size is user-defined in KB.

Minimum ATM interface buffer size

This ATM adapter supports the use of interface buffers. The socket layer will attempt to obtain an interface buffer and copy data directly into a buffer that has already been prepared for DMA by the adapter. Interface buffers are not used for transmits with lengths less than this value. The value is set to 2kB.

Minimum Guaranteed VCs Supported

The minimum number of VCs that the system is guaranteed to accommodate at any time.

Maximum Number of VCs Needed

The maximum number of VCs that the system is expected to accommodate at any given time.

SONET or SDH interface (only 155 Mbps adapter)

The 155 Mbps adapter can generate SONET or SDH framing. 0 is the default and provides SONET framing. A value of 1 provides SDH framing.

Provide SONET clock (only 155 Mbps adapter)

The 155 Mbps adapter can provide the SONET clock or use the Network's clocking. Clocking is provided to the network if set to 1, if 0 the network clock is used.

After you have changed the values, the network interfaces must be restarted. To start network interface at0 use the following command:

```
/usr/lib/methods/cfgif at0
```

Note: If the network interfaces can't be reactivated, the system must be rebooted.

Appendix A. End Node Impact on Throughput

As you can see from the following chart, the higher throughput available from ATM networks can be lost due to system performance. This chart shows the impact of processor speed. The I/O performance of the end nodes has a greater impact. A very slow disk throughput on either end will significantly degrade performance. Archiving to a network share tape via ATM is not efficient use of the network.

Processor, memory, bus, disk, application, etc. can all affect throughput. Some ATM adapters cannot reach the same speed as others. The maximum throughput is limited by the slowest component.

Figure 10 shows the throughput of the Turboways 100 adapter using different RS/6000s as end nodes.

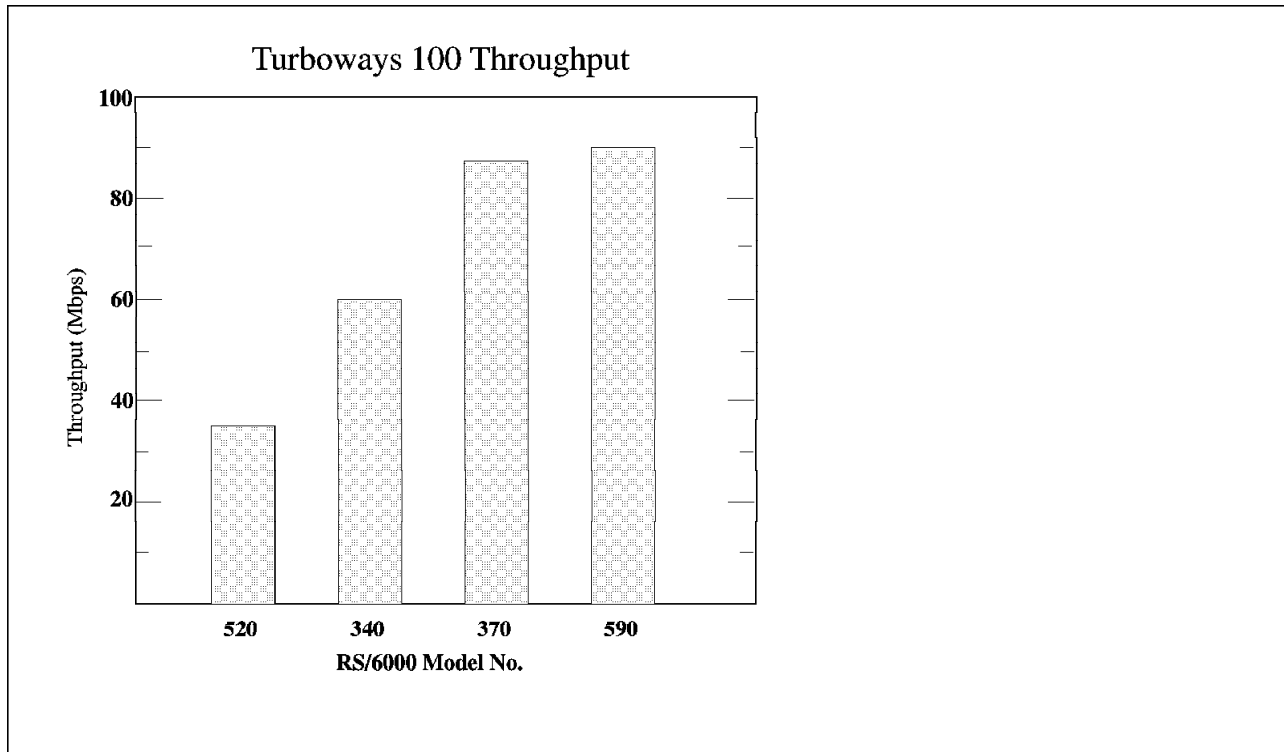


Figure 10. Processor Impact on Throughput

Figure 11 on page 56 shows the impact of disk I/O on high speed communications throughput.

| | Standard SCSI Low Speed Disk | Fast/Wide SCSI Hi-Perf Disk | Fast/Wide SCSI H-Perf Disks Striped (2 Disks) |
|---------|---------------------------------|--------------------------------|---|
| ATM 100 | 3.8MBs | 6.99MBs | 9.8MBs |
| ATM 155 | 3.8MBs | 6.99MBs | 13.1MBs |
| FDDI | 3.7 MBs | 6.8MBs | 9 MBs |

FTP was used on two RS/6000 Model 59Hs.

RFC 1323 was ON.

There was not other activity on the systems.

The networks were point to point

AIX Version 4.

Files were not memory cached.

87% CPU utilization on receiving system to attain 13.1MBs.

Figure 11. Disk I/O Impact on ATM Throughput

The throughput for 100MBs ATM and FDDI with striped disks is near the maximum for the ATM network. Striping to more or faster disks should improve the throughput of the 155MBs ATM network.

These figures represent experience with ATM, not controlled performance testing. Your results may not be consistent with ours.

Just remember that something is limiting the throughput. It may not be the network. When you remove that limiter, you will find another.

Appendix B. Special Notices

This publication is intended to help anyone intending to configure or maintain AIX end nodes in ATM network. The information in this publication is not intended as the specification of any programming interfaces that are provided by AIX Operating System. See the PUBLICATIONS section of the IBM Programming Announcement for AIX Version 4.1.4 or AIX Version 4.2 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific

information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

| | |
|------------------|--------------|
| AIX | HACMP/6000 |
| IBM | InfoExplorer |
| Micro Channel | Nways |
| RISC System/6000 | RS/6000 |
| TURBOWAYS | |

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Java and HotJava are trademarks of Sun Microsystems, Inc.

| | |
|--------|--------------------------------------|
| ATM | Adobe Systems, Incorporated |
| Tivoli | Tivoli Systems, Inc., an IBM Company |

Other trademarks are trademarks of their respective companies.

Appendix C. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How To Get ITSO Redbooks" on page 61.

- *Asynchronous Transfer Mode (ATM) Technical Overview*, SG24-4625

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

| CD-ROM Title | Subscription Number | Collection Kit Number |
|---|---------------------|-----------------------|
| System/390 Redbooks Collection | SBOF-7201 | SK2T-2177 |
| Networking and Systems Management Redbooks Collection | SBOF-7370 | SK2T-6022 |
| Transaction Processing and Data Management Redbook | SBOF-7240 | SK2T-8038 |
| AS/400 Redbooks Collection | SBOF-7270 | SK2T-2849 |
| RISC System/6000 Redbooks Collection (HTML, BkMgr) | SBOF-7230 | SK2T-8040 |
| RISC System/6000 Redbooks Collection (PostScript) | SBOF-7205 | SK2T-8041 |
| Application Development Redbooks Collection | SBOF-7290 | SK2T-8037 |
| Personal Systems Redbooks Collection | SBOF-7250 | SK2T-8042 |

C.3 Other Publications

These publications are also relevant as further information sources:

- */AIXtra*, March/April 1996, URL <http://pscc.dfw.ibm.com/aixtra/>
- *The Best of AIXtra Volume II*, SR28-5911
- *System Management Guide: Communications and Networks*, SC23-2526
- *AIX Performance Tuning Guide*, SR28-5930
- *Multiprotocol Intelligent Switching Hub Installation and Operation Guide*, SA33-0251
- *Nways 8260 ATM 155 Mbps Flexible Concentration Module Installation and User's Guide*. SA33-0358-01

How To Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL <http://www.redbooks.ibm.com>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get lists of redbooks:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

To register for information on workshops, residencies, and redbooks:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
```

For a list of product area specialists in the ITSO:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Home Page on the World Wide Web**

<http://w3.itso.ibm.com/redbooks>

- **IBM Direct Publications Catalog on the World Wide Web**

<http://www.elink.ibm.link.ibm.com/pb1/pb1>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibm.link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

| | IBMAIL | Internet |
|------------------------|--------------------|----------------------|
| In United States: | usib6fpl at ibmail | usib6fpl@ibmail.com |
| In Canada: | caibmbkz at ibmail | lmannix@vnet.ibm.com |
| Outside North America: | dkibmbsh at ibmail | bookshop@dk.ibm.com |

- **Telephone orders**

| | |
|---------------------------|-------------------------------|
| United States (toll free) | 1-800-879-2755 |
| Canada (toll free) | 1-800-IBM-4YOU |
| Outside North America | (long distance charges apply) |
| (+45) 4810-1320 - Danish | (+45) 4810-1020 - German |
| (+45) 4810-1420 - Dutch | (+45) 4810-1620 - Italian |
| (+45) 4810-1540 - English | (+45) 4810-1270 - Norwegian |
| (+45) 4810-1670 - Finnish | (+45) 4810-1120 - Spanish |
| (+45) 4810-1220 - French | (+45) 4810-1170 - Swedish |

- **Mail Orders** — send orders to:

| | | |
|--|--|--|
| IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA | IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada | IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark |
|--|--|--|

- **Fax** — send orders to:

| | |
|---------------------------|---|
| United States (toll free) | 1-800-445-9269 |
| Canada | 1-403-267-4455 |
| Outside North America | (+45) 48 14 2207 (long distance charge) |

- **1-800-IBM-4FAX (United States)** or **(+1) 415 855 43 29 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

| | |
|---------------------------------|---|
| Redbooks Home Page | http://www.redbooks.ibm.com |
| IBM Direct Publications Catalog | http://www.elink.ibm.com/pbl/pbl |

- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibm.com with the keyword `subscribe` in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

| Title | Order Number | Quantity |
|-------|--------------|----------|
|-------|--------------|----------|

- Please put me on the mailing list for updated versions of the IBM Redbook Catalog.
-

| | |
|------------|-----------|
| First name | Last name |
|------------|-----------|

Company

Address

| | | |
|------|-------------|---------|
| City | Postal code | Country |
|------|-------------|---------|

| | | |
|------------------|----------------|------------|
| Telephone number | Telefax number | VAT number |
|------------------|----------------|------------|

- Invoice to customer number _____

- Credit card number _____
-

| | | |
|-----------------------------|----------------|-----------|
| Credit card expiration date | Card issued to | Signature |
|-----------------------------|----------------|-----------|

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

List of Abbreviations

| | | | |
|---------------|---|---------------|---|
| AAL | ATM Adaptation Layer | MIB | Management Information Block |
| API | Application Program Interface | MTU | Maximum Transmission Unit |
| ARP | Address Resolution Protocol | Mbps | Million bits per second |
| ATM | Asynchronous Transfer Mode | NNI | Network Node Interface |
| B-ISDN | Broadband Integrated-Services Digital Network | PDU | Protocol Data Unit |
| CBR | Constant Bit Rate | PVC | Permanent Virtual Channel |
| DCC | Data Country Code | RFC | Request for Comment |
| DMA | Direct Memory Access | RPC | Remote Procedure Call |
| FC | Feature Code | SDH | Synchronous Data Hierarchy |
| HACMP | High Availability Cluster Multi-Processing | SMIT | System Management Interface Tool |
| IBM | International Business Machines Corporation | SNMP | Simple Network Management Protocol |
| ISDN | Integrated-Services Digital Network | SONET | Synchronous Optical Network |
| ITSO | International Technical Support Organization | SVC | Switched Virtual Channel |
| ITU-T | International Telecommunication Union - Telecommunication | TCP/IP | Transmission Control Protocol/Internet Protocol |
| LAN | Local Area Network | UNI | User Network Interface |
| LIS | Logical IP Subnet | VBR | Variable Bit Rate |
| LLC | Logical Link Control | VC | Virtual Channel |
| MAC | Medium Access Control | VCC | Virtual Channel Connection |
| MAN | Metropolitan Area Network | VCI | Virtual Channel Identifier |
| | | VP | Virtual Path |
| | | VPC | Virtual Path Connection |
| | | VPI | Virtual Path Identifier |

Index

Special Characters

/etc/rc.net 31

A

AAL 5 3
AAL5 7
abbreviations 65
acronyms 65
Address formats 4
arp 7, 19
ARP cache 22, 45, 48
ARP Client 7, 47
ARP Server 7, 23, 31, 45
arp table 19, 35
ARP-client 21
ARP-Server 21
ATM 1
ATM 100 11
ATM 155 11
ATM adaptation layer 3
ATM Address 4
ATM layer 4
ATM Network Interface 15
ATM switch 21
ATM100 21
ATM155 21
ATMARP 7, 35
ATMARP client 7
ATMARP server. 7
atmstat 37

B

B-ISDN 1
backbone 27
bibliography 59
blade 21
bottleneck 49
broadcast 7

C

cables 12
CCITT 8
cfgmgr 14
chinet 44
Classical IP 4, 6
command 39

D

DCC 4

devices 13
disk 49
DMA 50

E

errpt 37
Ethernet 27

F

filesets 13

H

HACMP 35
Hardware 11

I

IBM 8260 21
IDC 4
ifconfig 19, 51
ifsize 31
InARP_REPLY 35
InARP_REQUEST 35
InATMARP 7
InfoExplorer 39
Interface 5, 31, 33
IP subnet 31
IP-subnet 21
ipforwarding 28
iptrace 37, 39
ISDN 1

L

LAN 2
LIS 7, 21
Logical IP Subnetwork 7
lsdev command 11
lsparent 44

M

MAN 2
mbuf 44, 50
media 49
 Micro Channel 49
microcode 14
 back-to-back 15
mkatmpvc 45

N

netstat 19
Network Interface 15, 22, 46
NNI 5
no 31

P

Prerequisites 21
PVC 6, 15, 45

R

receivespace 49
RFC 1577 7, 8
rfc1323 49
route 28
router 8, 31
Routing 31

S

SC-type 12
SDH 51
sendspace 49
socket buffer 49
Software 13
SONET 51
Specification V3.0 6
ST-type 12
standards 8
static route 28
status 47
subnet address 7
subnet mask 7
SVC 6
switch 21

T

TCP/IP 4, 15, 22
tcpdump 37, 39
throughput 49
TokenRing 27
Turboway 21
Turboways 39
Turboways 100 8
Turboways 155 8

U

UNI 5
UNI V3.0 8

V

Variable Bit Rate 2
VC 6

VCC 6, 35
Virtual Channel 4
Virtual Channel Connection 6
Virtual Path 4
VP 6

W

WAN 2
wire fault 12



This soft copy for use by IBM employees only.

Printed in U.S.A.

SG24-4796-00

