OSA-Express Implementation Guide

Product, planning, and quick start information

Realistic examples and considerations

Hardware and software setup definitions

Bill White
Valirio Braga
Joe Van Dellen

Redbooks

ibm.com/redbooks

International Technical Support Organization

OSA-Express Implementation Guide

January 2003

SG24-5948-02

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

Third Edition (January 2003)

This edition applies to the OSA-Express features installed in the IBM zSeries Servers and IBM 9672 Generation 5 and 6 Servers.

© Copyright International Business Machines Corporation 2000, 2001, 2003. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2003. All rights reserved.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

© Copyright IBM Corp. 2003. All rights reserved.

Trademarks

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

Advanced Peer-to-Peer Networking®
AIX®
CUA®
ESCON®
IBM®
IBM eServer™
Language Environment®
Lotus®
MORE™
MVS™
NetView®
Notes®
OS/2®
OS/390®
Parallel Sysplex®
PC 300®
Perform™
RACF®
Redbooks™
Redbooks(logo)™
RMF™
S/390®
SP™
TCS®
Tivoli®
TME®
VM/ESA®
VSE/ESA™
VTAM®
Word Pro®
z/Architecture™
z/OS™
z/VM™
zSeries™

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM Redbook discusses how to install, tailor, and configure the Open System Adapter-Express features that are available on the 9672 Generation 5 and Generation 6 Parallel Enterprise Servers, and zSeries Servers. It focuses on the hardware installation, as well as the software definitions, needed to provide connectivity to various LAN environments.

This book provides information for planning purposes and system setup. Also included are helpful utilities and commands for monitoring and managing the OSA-Express features.

This document is intended for system engineers, network administrators, and system programmers who will plan and install OSA-Express. A solid background in hardware (9672 Generation 5 and Generation 6 Parallel Enterprise Servers, and zSeries Servers), HCD or IOCP, OSA/SF, SNA/APPN, and TCP/IP, is assumed.

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, Poughkeepsie Center.

Bill White is a Project Leader and Senior Networking Specialist at the International Technical Support Organization, Poughkeepsie Center.

Valirio Braga is a Senior I/T Specialist with IBM in Brazil, and has five years of experience in the networking field. His areas of expertise include VTAM, TCP/IP, CS for z/OS and OSA.

Joe Van Dellen is a Senior I/T Specialist with IBM in Atlanta, Georgia, providing systems and network support. He also develops and teaches courses dealing with large systems networking and OSA.

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

Dave Bennin, Roy Costa
International Technical Support Organization, Poughkeepsie Center

Robert Haimowitz
International Technical Support Organization, Raleigh Center

Connie Beuselinck
IBM zSeries Product Planning

Jim Goethals
IBM zSeries Networking Offering Manager

© Copyright IBM Corp. 2003. All rights reserved.

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYJ Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

Overview

In this chapter, we describe the Open Systems Adapter-Express (OSA-Express) features. These features provide direct connection to server/clients on Ethernet, Fast Ethernet (FENET), Gigabit Ethernet (GbE) and Token Ring local area networks (LANs), and Asynchronous Transfer Mode (ATM) networks.

- OSA-Express modes and topologies:
  - QDIO mode
  - Non-QDIO mode
- OSA-Express elements:
  - OSA-Express Gigabit Ethernet (GbE)
  - OSA-Express Fast Ethernet (FENET)
  - OSA-Express Asynchronous Transfer Mode (ATM)
  - OSA Express Token Ring
  - OSA/SF
- zSeries and 9672 G5/G6 support

© Copyright IBM Corp. 2003. All rights reserved.

1.1 General Information

This section provides information that applies to all OSA-Express features.

1.1.1 Architecture

The OSA-Express features attach directly to the Self-Timed Interface (STI) bus. This design eliminates ESCON bus limitations.

The STI bus, introduced with the 9672 G5/G6 servers, is a 333 MB/second bidirectional bus that carries CHPID data to and from the CEC. The direct-connected STI chaining concept ensures a high level of reliability, availability and serviceability on the 9672 G5/G6 servers.

The STI bus on the zSeries operates at 1GB/second. This is three times that of the 9672 G5 and G6 servers. There are 24 STI buses in the zSeries for 24 GigaBytes of bandwidth between the I/O subsystem and memory. The STI bus is converted into four lower-speed STI buses (called fanout) via an I/O cage, with speeds of 333 MB/second and 500 MB/second. With STI fanout, the OSA-Express feature registers its desired bandwidth of 333 MB/sec to the System Assist Processor.

1.1.2 Characteristics

The OSA-Express feature is shipped with one universal Media Access Control (MAC) address per physical port. The universal MAC address can be replaced by a locally administered address (LAA) by using either the OSA Support Facility (OSA/SF) or OSA Advanced Facilities (OSA/AF) via the Hardware Management Console (HMC) or the Support Element (SE).

As an I/O channel, the OSA-Express feature is identified in the Hardware I/O Configuration Definition (HCD) by its channel path identifier (CHPID). The CHPID equates to the port of the OSA-Express on the 9672 G5/G6 and zSeries servers.

1.1.3 Devices

Even though there are different types of channels, depending on the mode type (QDIO or non-QDIO), the device type for all OSA-Express networking devices is *OSA*.

If you use OSA Support Facility (OSA/SF), you also need one device with a device type of *OSAD*.

In order for a *managing instance* of OSA/SF to recognize an OSA, you must associate one device number with the OSA channel path that is defined in the same logical partition (or system in basic mode) as the system image on which OSA/SF is running.

To grant that, we recommend that the OSA-Express channel and that device always be shared among *all* partitions in the system. Furthermore, this device number must be specified as device type = OSAD with a physical unit address = X’FE’. Because of these requirements, this device number is usually called either the *OSAD device* or the *FE devic*e.

1.1.4 Port sharing

The port on an OSA-Express feature can be shared among logical partitions (LPARs) by using the Enhanced Multiple Image Facility (EMIF). This was previously referred to as the ESCON Multiple Image Facility.

If a system is running logically partitioned, an OSA-Express channel path can be defined as shared among those logical partitions (LPs) to which it is defined in the system hardware I/O configuration (IOCDS). This allows access to a network port on the OSA-Express feature through the system image that is running in the LPAR to which OSA-Express has been defined.

A port can also be shared within the LPAR between TCP/IP stacks, or between TCP/IP and SNA (if it is an OSE CHPID).

For more information on configuring OSA-Express hardware, refer to Chapter 3, “Hardware configuration definitions” on page 37.

1.2 Description

There are two generations of OSA-Express cards. The latest generation was designed to exploit the features and speed of the zSeries servers. The previous generation was designed specifically for the 9672 G5/G6 servers. By installing the OSA-Express cards in the I/O cage of the appropriate server, they become integral components of the zSeries and 9672 G5/G6 servers I/O subsystems.

1.2.1 Operating modes and topologies

The Open Systems Adapter-Express (OSA-Express) Gigabit Ethernet (GbE), Fast Ethernet (FENET), Token Ring, and Asynchronous Transfer Mode (ATM) features are the next generation features beyond OSA-2 (see Figure 1-1 on page 4). OSA-Express features provide significant enhancements over OSA-2 in function, connectivity, bandwidth, data throughput, network availability, reliability, and recovery.

Each OSA-Express feature has one port for 9672 G5/G6 servers, and two ports for zSeries, which can be attached directly to a LAN or ATM network. This integration of channel path with network port makes the OSA-Express a unique channel, recognized by the hardware I/O configuration as one of the following channel types:

- OSD (Queued Direct I/O)
- OSE (non Queued Direct I/O)

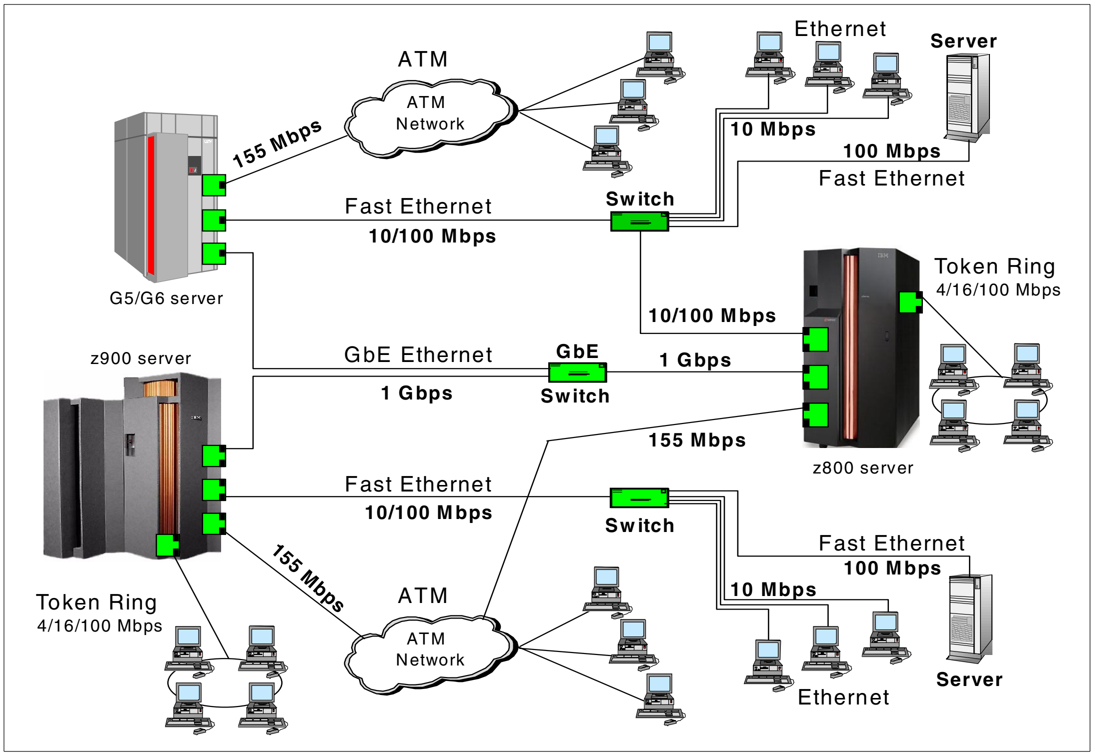

Figure 1-1 OSA-Express connectivity

Depending on the OSA-Express features you are using, two operating modes are supported. Figure 1-2 gives an overview of the Queued Direct I/O (QDIO) and non-QDIO modes of operation based on the supported features.

OSA-Express Feature

MODE	GbE	ATM	FENET	TR
QDIO (IP)	X	X *	X	X
TCP/IP Passthru		X	X	X
SNA		X	X	X
HPDT ATM Native		X		

These modes can run concurrently (TCP/IP Passthru and SNA).

X = Channel Type OSD (QDIO row) X = Channel Type OSE (other rows)

= Supported on z/OS, z/OS.e, OS/390, z/VM, VM/ESA, VSE/ESA, LINUX

= Supported on z/OS, z/OS.e, OS/390, z/VM, VM/ESA, VSE/ESA

= Supported on z/OS, z/OS.e, OS/390, z/VM, VM/ESA

Figure 1-2 OSA -Express operating modes

*QDIO mode is an option for the ATM feature only when emulating an Ethernet LAN.

Figure 1-3 illustrates the much shorter I/O process of the QDIO-enabled feature compared with the non-QDIO (which has the same I/O path as the OSA-2 features). Consequently, I/O interrupts and I/O path-lengths are minimized. The advantages of using QDIO are: 20% improved performance versus non-QDIO; the reduction of System Assist Processor (SAP) utilization; improved response time; CPC cycle reduction.

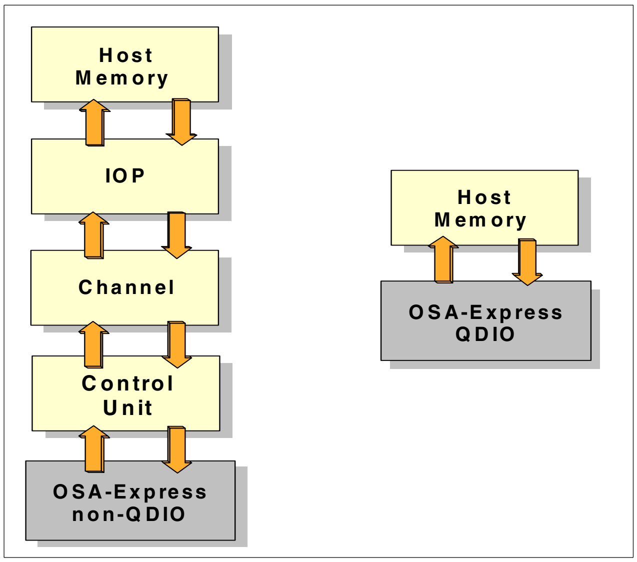

Figure 1-3 QDIO and non QDIO data paths

1.2.2 QDIO mode

Queued Direct Input/Output (QDIO) is a highly efficient data transfer mechanism that satisfies the increasing volume of TCP/IP applications and increasing bandwidth demands. It dramatically reduces system overhead, and improves throughput by using system memory queues and a signaling protocol to directly exchange data between the OSA-Express microprocessor and TCP/IP software. SNA support is provided through the use of Enterprise Extender; see 1.5, “Enterprise Extender” on page 20 for more information.

QDIO is supported with OSA-Express GbE SX and LX, OSA-Express FENET, OSA-Express Token Ring and OSA-Express 155 ATM MM and SM (when configured for Ethernet LAN emulation).

In QDIO mode, the OSA-Express microprocessor communicates directly with the server’s communications program, using data queues in main memory and utilizing Direct Memory Access (DMA).

The components that make up QDIO are Direct Memory Access (DMA), Priority Queuing (z/OS and OS/390 only), dynamic OSA Address Table building, LPAR-to-LPAR communication, and Internet Protocol (IP) Assist functions.

Direct Memory Access (DMA)

OSA-Express and Communications Server share a common storage area for memory-to-memory communication, reducing system overhead and improving performance. Data can move directly from the OSA-Express microprocessor to system memory. There are no read or write channel programs for data exchange. For write processing, no I/O interrupts have to be handled. For read processing, the number of I/O interrupts is minimized.

Priority queuing

Priority queuing is a capability supported by the QDIO architecture and introduced with the Service Policy Server (for z/OS and OS/390 environments only). It sorts outgoing IP message traffic according to the service policy you have set up for the specific priority assigned in the IP header.

This is an alternative to the best effort priority assigned to all traffic in most TCP/IP networks. Priority queuing allows the definition of four different priority levels for TCP/IP traffic through the OSA-Express features defined for QDIO. For example, you can grant interactive communications the highest priority while assigning batch traffic the lowest, with two additional categories in between, perhaps based on particular user groups or projects.

QDIO uses four write (outbound) queues and one read (inbound) queue for each TCP/IP stack sharing the OSA-Express feature.

OSA-Express signals to Communications Server when there is work to do. Communications Server puts outbound datagrams in one of the four queues based on priority settings.

At a certain time, Communications Server signals the OSA-Express feature that there is work to do. The OSA-Express feature searches the four possible outbound queues by priority and sends the datagrams to the network, giving more priority to queues 1 and 2, and less priority to queues 3 and 4.

For example, if there is data on every queue, queue 1 is served first, then portions of queue 2, then fewer portions of queue 3, then even fewer portions of queue 4, and then back to queue 1. This means that if there were four transactions running across the four queues, over time queue 1 would finish first, queue 2 would finish second, and so on.

Dynamic OSA Address Table (OAT) update

With QDIO, this process simplifies installation and configuration setups. The definition of IP addresses is done in one place, the TCP/IP profile, thus removing the requirement to enter the information into the OAT using the OSA Support Facility (OSA/SF).

The OAT entries will be dynamically built when the corresponding IP device in the TCP/IP stack is started.

At device activation, all IP addresses contained in the TCP/IP stack's IP HOME list are downloaded to the OSA-Express feature, and corresponding entries are built in the OAT. Subsequent changes to these IP addresses will cause a corresponding update of the OAT.

LPAR-to-LPAR communication

Access to a port can be shared among the system images that are running in the logical partitions to which the OSA-Express channel path is defined to be shared. Also, access to a port can be shared concurrently among TCP/IP stacks in the same LPAR or in different LPARs.

When port sharing, the OSA-Express features working in QDIO mode have the ability to send and receive IP traffic between logical partitions (LPARs) without sending the IP packets out to the LAN and then back to the destination LPAR.

For outbound packages, OSA-Express uses the next-hop address provided by the TCP/IP stack to determine where to send the packet. If this next-hop address had been registered by another TCP/IP stack sharing this OSA-Express, then OSA will deliver the packet directly to that TCP/IP stack, and not send the packet onto the LAN. This makes possible the routing of IP packets within the same host system.

LPAR-to-LPAR communication also applies to OSA-Express FENET when the mode is non-QDIO.

Internet Protocol (IP) Assist functions

The OSA-Express QDIO microcode assists in IP processing and offloads the TCP/IP stack functions for the following:

- Multicast support

  For sending data to multiple recipients. OSA-Express features support IP multicast destination addresses only in QDIO or IP Passthru mode.

- Broadcast filtering
- Building MAC and LLC headers
- ARP processing

The OSA-Express feature responds to Address Resolution Protocol (ARP) requests for its own IP address and for Virtual IP Addresses (VIPAs) for which the TCP/IP stack has assigned responsibility to send ARP replies. All the downloaded IP addresses are used to forward incoming datagrams to the corresponding TCP/IP stack in case the feature is shared. Also, whenever home IP addresses are dynamically added to or deleted from the stack, TCP/IP downloads these HOME list changes to the feature and updates the OAT.

QDIO is supported with all OSA-Express features, the GbE, the FENET, the Token Ring and the 155 ATM. However, each have different software requirements; see "Software support" on page 25.

ARP statistics

QDIO includes an IP assist (IPA) function, which gathers Address Resolution Protocol (ARP) data during the mapping of IP addresses to media access control (MAC) addresses. CHPIDs defined as OSD maintain ARP cache information in the OSA-Express feature (ARP offload). This is useful in problem determination for the OSA-Express feature.

Note that not all OSA-Express features provide ARP counter statistics and ARP cache information to TCP/IP.

Enhanced IP network availability (IPA)

There are several ways to ensure network availability, should failure occur at either the logical partition or the CHPID/network connection level. Port sharing, redundant paths, and the use of primary and secondary ports all provide some measure of recovery. A combination of these can guarantee network availability regardless of the failing component.

1.2.3 Non-QDIO mode

Like any other channel-attached control unit and device, an OSA-Express feature can execute channel programs (CCW chains) and present I/O interrupts to the issuing applications. For non-QDIO mode, the OSA-Express features are defined as channel type OSE. The non-QDIO mode requires the use of the OSA/SF for setup and customizing of the OSA-Express features.

The OSA-Express FENET, Token Ring and 155 ATM MM and SM features support non-QDIO mode. This mode supports SNA/APPN/HPR and/or TCPIP traffic simultaneously through the OSA-Express port. The non-QDIO mode types are as follows:

TCP/IP Passthru

The OSA-Express FENET, Token Ring, and ATM features can run in this mode. They can run concurrently in TCP/IP Passthru or SNA mode. In TCP/IP Passthru mode, an OSA-Express feature transfers data between a TCP/IP program to which it is defined and clients on the following networks:

- An Ethernet 10/100 Mbps LAN that is attached to the port on an OSA-Express FENET feature and supports one of the following frame protocols:
  - Ethernet II using the DEC Ethernet V 2.0 envelope
  - Ethernet 802.3 using the 802.2 envelope with SNAP
- A Token Ring 4/16/100 Mbps LAN that is attached to the port on an OSA-Express Token Ring feature and supports one of the following frame protocols:
  - IEEE 802.2 LAN MAC
  - IEEE 802.5 MAC (802.5 using the 802.2 envelope)
  - Token Ring 802.5 using the 802.2 envelope with SNAP
- An ATM emulated 155 Mbps LAN on an ATM-based network that is attached to the port of an OSA-Express 155 ATM feature and adheres to one of the following frame protocols:
  - Ethernet II using the DEC Ethernet V 2.0 envelope
  - Ethernet 802.3 using the 802.2 envelope with SNAP
  - Token Ring 802.5 using the 802.2 envelope with SNAP

  The ATM OSA-Express feature port must be attached to a 155 Mbps ATM switch. On each Emulated LAN (ELAN), the ATM OSA-Express feature port provides ATM LAN emulation client (LEC) services by means of one of its two LEC ports.

HPDT MPC (IP)

The OSA-Express FENET feature supports High Performance Data Transfer (HPDT) multipath channel (MPC) mode on 9672 G5/G6 servers when the HSAS (High Speed Access Services) for Communications Server is enabled.

Note: HSAS is not supported on the zSeries. Therefore, the HPDT MPC mode is no longer available for the OSA-Express FENET feature on the zSeries. For IP connectivity, use TCP/IP with the QDIO mode, or LAN Channel Station (LCS) with the non-QDIO mode.

The HPDT MPC mode is available on 9672 G5/G6 servers with Communications Server and HSAS.

HPDT ATM Native

The High Performance Data Transfer (HPDT) ATM Native mode allows you to take full advantage of the facilities of the ATM network to which the ATM OSA-Express feature port is attached. You can specify that the port transfers data across both permanent virtual circuits (PVCs) and switched virtual circuits (SVCs).

An ATM OSA-Express feature can run in the HPDT ATM Native mode to support high-speed networking for classical IP networks (RFC 1577). OSA-Express 155 ATM features running in HPDT ATM Native mode cannot support any other mode at the same time.

SNA/APPN/HPR support

The OSA-Express FENET, Token Ring and ATM features support SNA/APPN/HPR. An OSA-Express 155 ATM feature can support SNA traffic while operating in either ATM Native or LAN emulation.

If an OSA-Express feature is running in the SNA mode, it is viewed by VTAM as an external communications adapter (XCA) that can have either switched or non-switched lines of communication.

In this mode, an OSA-Express feature acts as an SNA passthru agent to clients that:

- Use the SNA protocol on the LAN that is directly attached to the OSA-Express
- Bridge from the ATM network in an emulated LAN (ELAN) configuration with OSA-Express 155 ATM

1.2.4 OSA-Express elements

The OSA-Express features attach directly to the Self-Timed Interface (STI) buses of the zSeries and 9672 G5/G6 servers. This design eliminates ESCON bus limitations.

The OSA-Express feature ports are defined in IOCP and identified by its channel path identifier (CHPID).

OSA-Express link

The transmission medium and cabling requirements for the OSA-Express ports depend on the OSA-Express feature type, OSA-Express port, LAN, and network. Acquiring cables and other connectivity items is the customer's responsibility. Figure 1-1 on page 4 shows connectivity examples of the OSA-Express features.

OSA-Express GbE

The OSA-Express GbE features only support QDIO mode and TCP/IP.

SNA/APPN traffic can be transported through a TCP/IP network by using TN3270 or Enterprise Extender (see “Enterprise Extender” on page 20).

Gigabit Ethernet employs the same Carrier Sense Multiple Access/Collision Detection (CSMA/CD) protocol, supporting the same frame format as the 10 and 100 Mbps Ethernet LANs. There is no need to reinvest in network staff education to support other protocols. The same middleware can be used.

The OSA-Express Gigabit Ethernet features support:

- Queued Direct Input/Output (QDIO) architecture
- Full duplex data transmission, with separate read and write paths
- Multicast traffic (but it will only send Ethernet V2.0 (DIX) formatted frames)
- Two types of Ethernet frames:
  - IEEE 802.3 CSMA/CD (ISO/IEC 802.3)
  - Ethernet V2.0 (DIX)

  Note: Both the 802.3 and the DIX V2 support standard Ethernet frame with MTU of 1492; DIX V2 supports frames larger than 1518 bytes (including Jumbo frame).
- Jumbo frames up to 8992 bytes

When an OSA-Express GbE feature port is configured to use Jumbo frames, fewer frames are generated and processed, achieving increased throughput for data-intensive TCP/IP applications.

OSA-Express FENET

The OSA-Express FENET feature only supports TCP/IP traffic if it is defined in QDIO mode. However, transporting SNA traffic over IP can be accomplished by using TN3270 or Enterprise Extender (see “Enterprise Extender” on page 20).

An OSA-Express feature port will support TCP/IP and SNA/APPN/HPR traffic concurrently when non-QDIO mode is configured. For TCP/IP Passthru mode, the default OAT may be used. In that case, no configuration or setup is required.

The FENET feature provides a direct connection to a 10 Mbps or 100 Mbps Ethernet LAN, running in either half-duplex or full-duplex mode. An OSA-Express FENET supports auto-negotiation with its attached Ethernet switch.

LAN speed and duplex mode default to auto-negotiation, and the OSA-Express FENET feature port and the attached switch automatically negotiate these settings. If the attached switch does not support auto-negotiation, the port enters the LAN at the default speed of 100 Mbps and half-duplex mode.

LAN speed and/or the duplex mode can be set explicitly, using OSA Support Facility (OSA/SF) or the OSA Advanced Facilities (OSA/AF) function of zSeries or 9672 G5/G6 hardware management console (HMC). The explicit settings will override the OSA-Express feature port's ability to auto-negotiate with the Ethernet switch to which it is attached.

Note: The OSA-Express FENET feature port and the attached Ethernet switch must be set to the same speed and duplex mode if auto-negotiation is disabled.

OSA-Express Token Ring

The OSA-Express Token Ring feature supports only TCP/IP traffic if it is defined in QDIO mode. However, transporting SNA traffic over IP can be accomplished by using TN3270 or Enterprise Extender (see “Enterprise Extender” on page 20).

An OSA-Express feature port will support TCP/IP and SNA/APPN/HPR traffic concurrently when non-QDIO mode is configured. For TCP/IP Passthru mode, the default OAT may be used. In that case, no configuration or setup is required.

The Token Ring feature provides a direct connection to a 4 Mbps, 16 Mbps, or 100 Mbps Token Ring LAN. An OSA-Express Token Ring supports autosense and will connect to the Token Ring LAN at the appropriate speed. If first on the ring, it will connect at 16 Mbps.

OSA-Express 155 ATM

The OSA-Express 155 ATM features support TCP/IP Passthru and/or SNA/APPN/HPR traffic concurrently in Ethernet or Token-Ring LAN emulation.

The OSA-Express 155 ATM features support HPDT ATM Native mode in OS/390 and VM environments. This mode requires the exclusive use of the OSA-Express 155 ATM feature. Data transfer is supported via VTAM for both the SNA and TCP/IP functions of the Communications Server.

Note: Two instances of TCP/IP in one LPAR cannot share a single OSA-Express 155 ATM feature running in HPDT ATM Native mode. However, one instance each of TCP/IP and VTAM in the same LPAR can do so.

QDIO support is only available when the OSA-Express 155 ATM feature is running in the ATM Ethernet LAN Emulation mode. QDIO cannot be used in any combination with any other modes of operation. The OSA-Express 155 ATM feature can only be configured for one mode at a time—either for ATM LAN Emulation (TR or Ethernet) or ATM Native.

With the ATM Native Mode, each physical port on the OSA-Express 155 ATM feature supports a virtual connection between the server and an ATM Native network.

With the ATM LAN Emulation Mode, you can define two emulated ports for the OSA-Express 155 ATM feature port to connect to two separate emulated LANs, creating two LAN Emulation Clients (LECs). An *emulated port* is a virtual connection between the server and an existing Ethernet or Token Ring network and provides LEC services for SNA and IP clients. Because the OSA-Express 155 ATM feature port allows two emulated ports to be configured, the CHPID can handle two different kinds of network traffic with no recabling at the physical port.

Tip: An option called *partial activation* enables you to add or change one emulated port without interrupting traffic on the other emulated port. Partial activation applies only to emulated ports.

The ATM LAN Emulation (LANE) can support simultaneous IP and SNA/APPN/HPR traffic over Ethernet LAN or Token Ring (TR) LAN. This means the workstation interface and wiring can remain the same for investment protection. Two Emulated LANs (ELANs) or LAN Emulation Clients (LECs) can be defined per physical port, and each can be connected to an ATM switch that is in turn connected to an Ethernet or Token Ring LAN. A maximum of 4096 PUs for SNA per physical port is supported. Each ELAN will cache up to 2000 simultaneous MAC addresses.

The ATM Native mode can support Classical IP (RFC1577, RFC2225) and APPN, simultaneously.

Note: You must use OSA/SF to set up the ATM features regardless of the mode being customized.

The two physical ports on the OSA-Express 155 ATM features have independent subsystems; each port can be defined for either LAN emulation or native ATM.

1.2.5 OSA Support Facility

The Open Systems Adapter Support Facility (OSA/SF) is a host-based tool to support and manage the Open Systems Adapter 2 (OSA-2) features and the OSA-Express features. One OSA/SF application can communicate with all OSA features in a hardware complex. OSA/SF communicates with an OSA feature through a device predefined on the OSA feature. The device type is OSAD. The OSAD device must have a unit address of x’FE’.

OSA/SF supports a Graphical User Interface (GUI) on a Windows 95, Windows NT, or OS/2 workstation. The communication between the GUI workstation and OSA/SF can be TCP/IP, APPC, or EHLLAPI (3270).

With z/OS and OS/390, a second interface using a set of REXX EXECs through the Time Sharing Options Extensions (TSO/E) can be used to control the OSA-2 features or OSA-Express features defined to the 9672 G5/G6 or zSeries servers on which the TSO/E is running.

OSA/SF is not always required to customize an OSA-Express feature, but is highly recommended to gather operational information and to assist in problem determination. The OSA/SF Query function provides performance information about the OSA-Express CHPIDs.

OSA/SF is required for the configuration of the OSA-Express 155 Mbps ATM features. OSA/SF is not required to configure the GbE and the FENET features in QDIO, and with FENET features in non-shared or default mode.

Refer to 1.10.2, “OSA/SF” on page 28 for information on OSA/SF support in z/OS, OS/390, z/VM, VM/ESA, and VSE environments. Refer to “Setting up and using OSA/SF” on page 47 for details on working with OSA/SF.

1.3 Connectivity

Following are connectivity options in the OSA-Express environment.

1.3.1 OSA-Express support on the zSeries and 9672 G5/G6 servers

The 9672 G5/G6 servers support five OSA-Express feature types, each with one port. zSeries supports six OSA-Express features, each with two ports. All of the OSA-Express features occupy one slot in an I/O cage. Regardless of platform, a maximum of 12 features can be installed; 12 features/ports on 9672 G5/G6 servers and 12 features/24 ports on zSeries servers.

Note: OSA-Express features are not interchangeable between 9672 G5/G6 and zSeries.

1.3.2 zSeries OSA-Express support

The zSeries servers support the following OSA-Express features:

zSeries OSA Express feature type	Feature code
zSeries OSA-Express 155 ATM SM	FC 2362
zSeries OSA-Express 155 ATM MM	FC 2363
zSeries OSA-Express GbE LX	FC 2364
zSeries OSA-Express GbE SX	FC 2365
zSeries OSA-Express FENET	FC 2366
zSeries OSA-Express Token Ring	FC 2367

Note: IBM Statement of Direction - announced April 2002

zSeries (z900 and z800) servers will be the last family of servers to provide OSA-Express 155 ATM features.

zSeries OSA-Express 155 ATM MM feature (FC 2363)

The zSeries OSA-Express 155 ATM MM feature occupies one slot in the zSeries I/O cage. The feature has two independent ports with one CHPID associated with each port.

Each port supports connection to a 155 Mbps ATM network via 62.5 micron or 50 micron multimode fiber optic cable terminated with an SC Duplex connector.

zSeries OSA-Express 155 ATM SM feature (FC 2362)

The zSeries OSA-Express 155 ATM SM feature occupies one slot in the zSeries I/O cage. The feature has two independent ports, with one CHPID associated with each port.

Each port supports connection to a 155 Mbps ATM network via 9 micron single-mode fiber optic cable terminated with an SC Duplex connector.

zSeries OSA-Express GbE LX feature (FC 2364)

The zSeries OSA-Express GbE LX feature occupies one slot in the zSeries I/O cage. The feature has two independent ports, with one CHPID associated with each port. The zSeries OSA-Express GbE LX feature supports the 1000BASE-LX standard transmission scheme.

Each port supports connection to a 1 Gbps Ethernet LAN via 9 micron single-mode fiber optic cable terminated with an SC Duplex connector.

Multimode (62.5 or 50 micron) fiber optic cable may be used with the zSeries OSA-Express GbE LX feature. The use of these multimode cable types requires a mode conditioning patch (MCP) cable to be used at each end of the fiber link. Use of single mode-to-multimode MCP cables reduces the supported optical distance of the link to a maximum end-to-end distance of 550 meters.

The zSeries OSA-Express GbE LX feature must be defined as QDIO mode and supports TCP/IP only. TN3270 or the Enterprise Extender (EE) function of Communications Server for z/OS and OS/390 allows you to run SNA applications and data on IP networks and IP attached clients. See “Enterprise Extender” on page 20 for more information.

zSeries OSA-Express GbE SX feature (FC 2365)

The zSeries OSA-Express GbE SX feature occupies one slot in the zSeries I/O cage. The feature has two independent ports, with one CHPID associated with each port. The zSeries OSA-Express GbE SX feature supports the 1000BASE-SX standard transmission scheme.

Each port supports connection to a 1 Gbps Ethernet LAN via 62.5 micron or 50 micron multimode fiber optic cable terminated with an SC Duplex connector.

The zSeries OSA-Express GbE SX feature must be defined as QDIO mode, and supports TCP/IP only. TN3270 or the Enterprise Extender (EE) function of Communications Server for z/OS and OS/390 allows you to run SNA applications and data on IP networks and IP attached clients. See “Enterprise Extender” on page 20 for more information.

zSeries OSA-Express FENET feature (FC 2366)

The zSeries OSA-Express FENET (Fast Ethernet) feature occupies one I/O slot in the zSeries I/O cage. The feature has two independent ports, with one CHPID associated with each port.

Each port supports connection to either a 100 Mbps or 10 Mbps Ethernet LAN. The LAN must conform either to the IEEE 802.3 (ISO/IEC 8802.3) standard or the Ethernet V2.0 specifications, and the 10BASE-T or 100BASE-TX standard transmission schemes.

Each port has an RJ-45 receptacle for cabling to an Ethernet switch that is appropriate for the LAN speed. The RJ-45 receptacle is required to be attached using EIA/TIA category 5 unshielded twisted pair (UTP) cable with a maximum length of 100 m (328 ft.).

The OSA-Express FENET feature supports auto-negotiation when attached to an Ethernet hub, router, or switch. If you allow the LAN speed to default to auto-negotiation, then the FENET OSA-Express and the attached hub, router, or switch auto-negotiate the LAN speed setting between them. If the attached Ethernet hub, router, or switch does not support auto-negotiation, then the OSA enters the LAN at the default speed of 100 Mbps in half-duplex mode.

If you are not using auto-negotiate, the OSA will attempt to join the LAN at the specified speed/mode; however, the speed/mode settings are only used when the OSA is first in the LAN. If this fails, the OSA will attempt to join the LAN as if auto-negotiate were specified.

You can choose any one of the following settings for the OSA-Express FENET feature:

- Auto-negotiate
- 10 Mbps half-duplex
- 10 Mbps full-duplex
- 100 Mbps half-duplex
- 100 Mbps full-duplex

The HPDT MPC mode is no longer available on the FENET for zSeries; see “HPDT MPC (IP)” on page 8 for more information.

zSeries OSA-Express Token Ring feature (FC 2367)

The zSeries OSA-Express Token Ring feature occupies one I/O slot in the zSeries I/O cage. The feature has two independent ports, with one CHPID associated with each port.

The OSA-Express Token Ring feature supports auto-sensing, as well as any of the following settings: 4 Mbps half- or full-duplex; 16 Mbps half- or full-duplex; 100 Mbps full-duplex.

Note: Regardless of the choice made, the network switch settings must agree with those of the OSA-Express Token Ring feature.

If the LAN speed defaults to auto-sense, the OSA-Express Token Ring feature will sense the speed of the attached switch and insert into the LAN at the appropriate speed. If the Token Ring feature is the first station on the LAN and the user specifies auto-sense, it will default to a speed of 16 Mbps and will attempt to open in full-duplex mode. If unsuccessful, it will default to half-duplex mode. The OSA-Express Token Ring feature conforms to the IEEE 802.5 (ISO/IEC 8802.5) standard.

Each port has an RJ-45 receptacle and a DB-9 D shell receptacle for cabling to a Token Ring MAU or Token Ring switch that is appropriate for the LAN speed. Only one of the port’s two receptacles can be used at any time.

The RJ-45 receptacle is required to be attached using EIA/TIA category 5 unshielded twisted pair (UTP) cable that does not exceed 100 m (328 ft.), or a shielded twisted pair (STP) cable with a DB-9 D Shell connector.

zSeries Channel CHPID assignment facility

OSA-Express is implemented as a channel type (OSD or OSE) on the zSeries servers. It appears to the application software as a channel-attached device.

The zSeries Channel CHPID Assignment facility is a function of the zSeries Support Element microcode which gives the ability to assign channel path identifiers (CHPID) to active channels independent of the installed physical position of the channel feature. This allows customers to fully customize their channel mapping to exploit the reliability, availability, and serviceability characteristics of the system. The Channel CHPID Assignment facility supports all channel feature types installed in the zSeries I/O cage and Compatibility I/O cage.

1.3.3 9672 G5/G6 OSA-Express support

The 9672 G5/G6 servers support the following OSA-Express features:

9672 G5/G6 OSA Express feature type	Feature code
9672 G5/G6 OSA-Express 155 ATM SM	FC 2360

9672 G5/G6 OSA-Express 155 ATM MM	FC 2361
9672 G5/G6 OSA-Express GbE SX	FC 2350
9672 G5/G6 OSA-Express GbE LX	FC 2351
9672 G5/G6 OSA-Express FENET	FC 2340

G5/G6 OSA-Express 155 ATM SM feature (FC2360)

The G5/G6 OSA-Express 155 ATM SM feature occupies one slot in the G5/G6 I/O cage. The feature has one port, with one CHPID associated with that port.

The port supports connection to a 155 Mbps ATM network via 9 micron single-mode fiber optic cable terminated with an SC Duplex connector.

G5/G6 OSA-Express 155 ATM MM feature (FC 2361)

The G5/G6 OSA-Express 155 ATM MM feature occupies one slot in the G5/G6 I/O cage. The feature has one port, with one CHPID associated with that port.

The port supports connection to a 155 Mbps ATM network via 62.5 micron multimode fiber optic cable terminated with an SC Duplex connector.

G5/G6 OSA-Express GbE SX feature (FC2350)

The G5/G6 OSA-Express GbE SX feature occupies one slot in the G5/G6 I/O cage. The feature has one port with one CHPID associated with that port. The G5/G6 OSA-Express GbE SX feature supports the 1000BASE-SX standard transmission scheme.

The port supports connection to a 1 Gbps Ethernet LAN via 62.5 micron or 50 micron multimode fiber optic cable terminated with an SC Duplex connector.

The G5/G6 OSA-Express GbE SX feature must be defined as QDIO mode and supports TCP/IP only. TN3270 or the Enterprise Extender (EE) function of Communications Server for z/OS and OS/390 allows you to run SNA applications and data on IP networks and IP attached clients. See “Enterprise Extender” on page 20 for more information.

G5/G6 OSA-Express GbE LX feature (FC 2351)

The G5/G6 OSA-Express GbE LX feature occupies one slot in the G5/G6 I/O cage. The feature has one port with one CHPID associated with that port. The G5/G6 OSA-Express GbE SX feature supports the 1000BASE-LX standard transmission scheme.

The port supports connection to a 1 Gbps Ethernet LAN via 9 micron single-mode fiber optic cable terminated with an SC Duplex connector.

Multimode (62.5 or 50 micron) fiber cable may be used with the G5/G6 OSA-Express GbE LX feature. The use of these multimode cable types requires a mode conditioner patch (MCP) cable to be used at each end of the fiber link. Use of the single-mode to multimode MCP cables reduces the supported optical distance of the link.

The G5/G6 OSA-Express GbE LX feature must be defined as QDIO mode and supports TCP/IP only. TN3270 or the Enterprise Extender (EE) function of Communications Server for z/OS and OS/390 allows you to run SNA applications and data on IP networks and IP-attached clients. See “Enterprise Extender” on page 20 for more information.

G5/G6 OSA-Express FENET feature (FC 2340)

The G5/G6 OSA-Express FENET feature occupies one I/O slot in the G5/G6 I/O cage. The feature has one port, with one CHPID associated with that port.

The port supports connection to either a 100 Mbps or 10 Mbps Ethernet LAN. The LAN must conform either to the IEEE 802.3 (ISO/IEC 8802.3) standard or the Ethernet V2.0 specifications, and the 10BASE-T or 100BASE-TX standard transmission schemes.

The port has an RJ-45 receptacle for cabling to an Ethernet switch that is appropriate for the LAN speed. The RJ-45 receptacle is required to be attached using EIA/TIA category 5 unshielded twisted pair (UTP) cable with a maximum length of 100 m (328 ft.).

The OSA-Express FENET feature supports auto-negotiation with its attached Ethernet hub, router, or switch. If you allow the LAN speed to default to auto-negotiation, then the FENET OSA-Express and the attached hub, router, or switch auto-negotiates the LAN speed setting between them. If the attached Ethernet hub, router, or switch does not support auto-negotiation, then the OSA enters the LAN at the default speed of 100 Mbps in half-duplex mode.

If you are not using auto-negotiate, the OSA will attempt to join the LAN at the specified speed/mode; however, the speed/mode settings are only used when the OSA is first in the LAN. If this fails, the OSA will attempt to join the LAN as if auto-negotiate were specified.

You can choose any one of the following settings for the OSA-Express FENET feature:

- Auto negotiate
- 10 Mbps half-duplex
- 10 Mbps full-duplex
- 100 Mbps half-duplex
- 100 Mbps full-duplex

1.4 Network considerations

The following sections introduce the rich functions realized from OSA-Express technology that can enhance today's LAN environment. They provide more flexibility, improved security and operations, and increase potential for reduction of cost.

1.4.1 TCP/IP elements

This section describes the TCP/IP elements supported by the OSA-Express features.

Virtual IP address (VIPA)

In the TCP/IP environment, VIPA frees TCP/IP hosts from dependence on a particular network attachment, allowing the establishment of primary and secondary paths through the network. VIPA is supported by all of the OSA-Express features.

An IP address traditionally ties to a physical link at one end of a connection. If the associated physical link goes down, it will be unreachable. The Virtual IP Address, on the other hand, exists only in software and has no association to any physical link. The TCP/IP stack is the destination address instead of the network attachment.

VIPA provides for multiple IP addresses to be defined to a TCP/IP stack, allowing fault tolerant, redundant, backup paths to be established. Applications become insensitive to the condition of the network since the VIPA will always be active, enabling users to route around intermediate points of failure in the network.

VIPA Takeover and Takeback

Since a VIPA is associated with a TCP/IP stack and not a physical network attachment, it can be moved to any TCP/IP stack within its network. If the TCP/IP stack that the VIPA is on fails (due to an outage), the same VIPA can be brought up automatically on another TCP/IP stack (VIPA Takeover) to allow end users to reach the backup server and applications. The original session between the end user and original server is not disrupted. Once the failed TCP/IP stack is restored, the same VIPA can be moved back automatically (VIPA Takeback).

DVIPA and Distributed DVIPA

The VIPA process was improved with Dynamic VIPA (DVIPA) in CS for V2R8. A backup system can be assigned for the Takeover process, which eliminates the need for operator intervention. Then it was improved again with Distributed Dynamic VIPA in CS for V2R10. A Distribution List can be configured to allow multiple stacks to service the many connections within a single stack. This provides for redundancy and workload balancing.

More information on VIPA

For more detailed information on VIPA requirements and implementing VIPA, refer to the following documents:

- *z/OS V1R4.0 Communications Server IP Configuration Guide*, SC31-8775
- *z/OS V1R4.0 Communications Server IP Configuration Reference*, SC31-8776
- *Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 1: Base and TN3270 Configuration,* SG24-5227
- *Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 4: Connectivity and Routing,* SG24-6516

Maximum IP addresses per OAT

The OSA Address Table (OAT) is a component of an OSA feature's configuration. An OAT entry defines the data path between an OSA feature port and an LPAR and device unit address. That is, it manages traffic through the OSA CHPID.

With multiple IP stacks, multiple LPARs, and VIPA, the maximum number of 16 IP addresses per port (with OSA-2) can be easily exceeded.

OSA-Express has increased the maximum number of IP addresses it supports, as follows:

- 512 IP addresses per OSA-Express port (from 16 with OSA-2) for 9672 G5/G6
- 2048 IP addresses per OSA-Express port in the zSeries

ARP Takeover

ARP Takeover provides the capability of switching OSA port operations from one OSA-Express to another OSA-Express running in the same mode.

When TCP/IP is started in QDIO mode, it downloads all the home IP addresses in the stack and stores them in each OSA-Express feature to which it has a connection. This is a service of QDIO architecture and occurs automatically for OSD channels.

For OSA-Express features set up as OSE channels (non-QDIO), you must define multiple IP addresses in the OSA Address Table using OSA/SF. The OSA-Express then responds to ARP requests for its own IP address, as well as for virtual IP addresses (VIPAs). If an OSA-Express feature fails while there is a backup OSA-Express available on the same network or subnetwork, TCP/IP informs the backup OSA which IP addresses (real and VIPA) to take over, and the network connection is maintained.

For more information on this and an example of our test, see Appendix G, "ARP Takeover" on page 223.

HiperSockets Accelerator

HiperSockets Accelerator allows a z/OS TCP/IP router stack to efficiently route IP packets from an OSA-Express (QDIO) interface to a HiperSockets internal QDIO (iQDIO) interface and vice versa. The routing is done by the z/OS Communications Server device drivers at the lowest possible software data link control level. IP packets do not have to be processed at the higher level TCP/IP stack routing function, hence reducing the path length and improving performance.

For more information on this and an example of our test, see Appendix F, “HiperSockets Accelerator” on page 217.

Primary/secondary router function

The primary/secondary router function enables an OSA-Express port to forward datagrams with unknown IP addresses to a TCP/IP stack for routing through another IP network interface, such as HiperSockets (iQDIO) or another OSA-Express feature.

In order for an OSA-Express port to forward IP datagrams to a particular TCP/IP stack for routing to its destination, the PRIRouter must be defined on the DEVICE statement in the TCP/IP profile.

If the TCP/IP stack that has an OSA-Express port defined as PRIRouter becomes unavailable, then a second TCP/IP stack defined as the secondary router (SECRouter on the DEVICE statement in the TCP/IP profile) will receive the datagrams for unknown IP addresses.

For enhanced availability, the definition of one primary router and multiple secondary routers for devices on an OSD-type CHPID is supported; however, only one secondary router is supported for devices on an OSE-type CHPID.

The OSA-Express Gigabit Ethernet and Fast Ethernet features in QDIO mode support multiple secondary routers.

HiperSockets Accelerator requires (at a minimum) that one OSA-Express port has the PRIRouter definition, so IP datagrams can be routed between the OSA-Express IP network and the HiperSockets IP network. You can find an example of this in Appendix F, “HiperSockets Accelerator” on page 217.

IPv6 support for z/OS and Linux

Internet Protocol Version 6 (IPv6) is supported for the OSA-Express GbE and the OSA-Express FENET features when configured in QDIO mode. IPv6 is the protocol designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4 (IPv4). Since there is a growing shortage of IP addresses (which are needed by all new machines added to the Internet), IPv6 was introduced to expand the IP address space from 32 bits to 128 bits, enabling a far greater number of unique IP addresses.

For more information about IPv6, refer to Chapter 12, “IPv6 support” on page 167.

VLAN support for Linux

Virtual Local Area Network (VLAN) is supported for the OSA-Express Ethernet, Fast Ethernet, and Gigabit Ethernet (GbE) features when configured in QDIO mode. This support is applicable to the Linux environment. Null VLAN tagging support is provided.

The IEEE standard 802.1Q describes the operation of Virtual Bridged Local Area Networks. A Virtual Local Area Network (VLAN) is defined to be a subset of the active topology of a Local Area Network. The OSA-Express features provide for the setting of multiple unique VLAN IDs per QDIO data device. They also provide for both tagged and untagged frames to flow from an OSA-Express port. The number of VLANs supported is specific to the operating system.

VLANs facilitate easy administration of logical groups of stations that can communicate as if they were on the same LAN. They also facilitate easier administration of moves, adds, and changes in members of these groups. VLANs are also designed to provide a degree of low-level security by restricting direct contact with a server to only the set of stations that comprise the VLAN.

With zSeries, where multiple stacks may exist potentially sharing one or more OSA-Express features, VLAN support is designed to provide a greater degree of isolation.

For more information concerning VLAN support, refer to Chapter 11, "VLAN support" on page 157.

Direct SNMP query support for z/OS and Linux

Simple Network Management Protocol (SNMP) is supported for all of the OSA-Express features when configured in QDIO mode. The query function is supported using the SNMP `get` command.

Open Systems Adapter Support Facility is no longer required to manage SNMP data for the OSA-Express features. A new SNMP subagent exists on an OSA-Express feature which is part of a direct path between the Linux and z/OS master agents (TCP/IP stacks) and an OSA-Express Management Information Base (MIB).

The OSA-Express features support an SNMP agent by providing data for use by an SNMP management application, such as Tivoli NetView. This data is organized into MIB tables defined in the TCP/IP enterprise-specific MIB, as well as standard RFCs. The data is supported by the SNMP TCP/IP subagent.

More information on SNMP

For more detailed information on SNMP requirements and implementing SNMP support, refer to the following documents:

- *z/OS V1R4.0 Communications Server IP Configuration Guide*, SC31-8775
- *z/OS V1R4.0 Communications Server IP Configuration Reference*, SC31-8776
- *z/OS V1R4.0 Communications Server IP System Administrator's Commands,* SC31-8781
- *Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 6: Policy and Network Management*, SG24-6839
- *Managing OS/390 TCP/IP with SNMP*, SG24-5866

TCP/IP broadcast support for z/OS, z/VM, and Linux

Broadcast support is being enhanced to include support for all of the OSA-Express features when configured in QDIO mode, and supporting the Routing Information Protocol (RIP) Version 1. Broadcast is currently supported for all of the OSA-Express features when carrying TCP/IP traffic and configured in the non-QDIO mode (LAN Channel Station - LCS mode).

A broadcast simultaneously transmits data to more than one destination; messages are transmitted to all stations in a network (for example, a warning message from a system operator). The broadcast frames can be propagated through an OSA-Express feature to all TCP/IP applications that require broadcast support, including applications using RIP V1.

ARP cache management: query and purge ARP

The query and purge ARP enhancements are supported for all OSA-Express features when configured in QDIO mode. The OSA-Express feature maintains a cache of recently acquired IP-to-physical address mappings (or “bindings”). When the binding is not found in the ARP cache, a broadcast (an ARP request “How can I reach you?”) to find an address mapping is sent to all hosts on the same physical network. Because a cache is maintained, ARP does not have to be used repeatedly, and the OSA-Express feature does not have to keep a permanent record of bindings.

Query ARP Table for IPv4 for Linux

Query ARP Table is supported for all of the OSA-Express features when configured in QDIO mode when using Internet Protocol Version 4 (IPv4). The TCP/IP stack already has an awareness of Internet Protocol Version 6 (IPv6) addresses.

Purge ARP entries in cache for IPv4 for z/OS and Linux

Purging of entries in the ARP cache is supported for all of the OSA-Express features when configured in QDIO mode when using IPv4. The TCP/IP stack already has an awareness of IPv6 addresses.

1.4.2 SNA enhanced availability

With the SNA enhanced availability, system availability is increased, network management is simplified, the network is more scalable, and new users can be added non-disruptively. 4096 Physical Units (PUs) can be defined for each OSA-Express feature port.

Load balancing

A single, locally administered MAC address can be defined to multiple physical OSA-Express feature ports attached to different LAN segments. The number of sessions established is monitored, and user session loads distributed; a session is established with the LAN segment that responds the fastest.

Redundancy

A secondary path between a LAN workstation and the host server can be configured as “hot standby”. If the primary path becomes unavailable, the secondary path receives the LAN traffic.

Overflow

User sessions flow through the primary path until the session limit is exceeded. When the port will no longer respond to requests, user sessions will automatically flow to the next OSA port.

Note: SNA Enhanced Availability is only offered for bridged source routing environments. Ethernet does not support source routing, hence only OSA-Express 155 ATM LAN Emulation mode (Token Ring) and OSA-Express Token Ring are supported.

1.5 Enterprise Extender

The Enterprise Extender (EE) function of Communications Server for z/OS and OS/390 allows you to run SNA applications and data on IP networks and IP-attached clients. It can be used with any OSA-Express feature running IP traffic.

EE is a simple set of extensions to the open High Performance Routing technology that integrate HPR frames into User Datagram Protocol/Internet Protocol (UDP/IP) packets, providing:

- SNA application connectivity using an IP backbone support for:
  - SNA-style priority
  - SNA Parallel Sysplex exploitation
- Improved throughput and response times
- Compatible support for TCP and UDP traffic on the IP portion of the application traffic path (SNA/HPR and UDP/IP traffic can coexist on an EE connection)

The EE function is a TCP/IP *encapsulation* technology that carries SNA traffic from an endpoint over an IP network (for example, via the OSA-Express port to Communications Server) to another endpoint where it is *de-encapsulated* and presented to an SNA application.

EE requires APPN/HPR at the endpoints. In order to enable EE, you must configure the TCP/IP stack with a virtual IP address and define an XCA major node. The XCA major node is used to define the PORT, GROUP, and LINE statements for the EE connections. See Chapter 10, “Enterprise Extender (EE)” on page 139, for details.

1.6 TN3270E Server

The TN3270E Server is supported by z/OS and OS/390. It allows desktop users to connect through an IP network directly to SNA applications.

TN3270E provides the following support:

- Secure access using SSL and Client Authentication via RACF
- Over 64,000 sessions per server when using multiple ports
- Over 2,000 transactions/second with sub-second response time
- Reconnect 16,000 sessions in less than a minute using VIPA Takeover support
- IP QoS using Service Policy Server
- Host print support
- Tivoli support provides IP visibility to VTAM
- Manage with your Data Center resources
- More robust than external TN3270 servers

Communications Server also supports a secure, RACF-based single signon logic called Express Logon Facility (ELF). ELF works with IBM TN3270E clients to securely authenticate the client, acquire a passtoken, and then pass it on to the TN3270E Server for replacement/submission to the application.

1.7 OSA-Express performance

The OSA-Express features attach directly to the 9672 G5/G6 server Self-Timed Interface (STI) bus, which has a bandwidth of 333 MBps bidirectional. With the zSeries, the OSA-Express features occupy slots in the new I/O cage running at 333 MBps bidirectional. The I/O cage attaches directly to a zSeries STI bus running at 1 GBps bidirectional.

The OSA-Express GbE feature port supports a link data rate of 1 Gbps and full-duplex transmission. The OSA-Express Fast Ethernet and 155 ATM features are capable of achieving line speeds of 175 Mbps full duplex for Fast Ethernet (the LAN is 100% busy) and 310 Mbps full duplex for 155 ATM.

The OSA-Express features help to ensure that the increasingly high volumes of data traversing the Local Area Network (LAN) do not encounter a bandwidth bottleneck, whether data is exchanged via intranet Web servers, centralized file servers, the Internet, or extranets.

Note: Actual throughput is dependent upon many factors, including the application, the disk subsystems, the traffic direction, the pattern of acknowledgment traffic, the datagram size, and LAN/network switch.

1.7.1 Resource Measurement Facility (RMF) support

RMF (z/OS and OS/390 only) reporting has been enhanced to support OSA-Express. This enables the user to capture performance data of the OSA-Express features:

- Microprocessor utilization (per LPAR image if it applies)
- Physical PCI (Peripheral Component Interconnect) bus utilization
- Bandwidth per port (both read and write directions), per LPAR image

With this enhanced support, a user can, for example, analyze possible bandwidth bottlenecks, and therefore perform root cause analysis.

1.8 Networking protocols

The following protocol standards are applicable for each OSA-Express feature type:

OSA-Express Gigabit Ethernet (GbE)

- Gigabit Ethernet
  - IEEE 802.3z Gigabit Ethernet Standard
  - DIX Version 2 Ethernet (RFC 894)

OSA-Express Fast Ethernet (FENET)

- Ethernet (10BASE-T)
  - IEEE 802.2 LAN MAC protocols
  - IEEE 802.3 CSMA/CD protocols
  - Ethernet V.2 protocols (not supported by SNA/APPN)
  - DIX Version 2 Ethernet (RFC 894)
- Ethernet (100BASE-TX)
  - IEEE 802.3u CSMA/CD protocols

OSA-Express 155 ATM

- ATM Forum UNI (User-to-Network Interface) Specifications Version 3.0 and 3.1 with signaling protocol Q.2931 for Ethernet and Token-Ring emulated LANs

- LAN emulation
  - Ethernet
    - IEEE 802.2 LAN MAC protocols
    - IEEE 802.3 CSMA/CD protocols
    - Ethernet V.2 protocols (not supported by SNA/APPN)
  - Token Ring
    - IEEE 802.2 LAN MAC protocols
    - IEEE 802.5 MAC protocols

OSA-Express Token Ring

- IEEE 802.2 LAN MAC protocols
- IEEE 802.5 MAC protocols

1.9 Cables

Planning the required cable configuration, ordering cables, and making arrangements to have them delivered and installed prior to system implementation is a customer responsibility. Cables are not provided. The Fiber Channel Standard (FCS) connector described in this document may be referred to as the SC Duplex connector. This connector is as defined by the ANSI (X3T9.3) Fiber Channel Physical and Signaling Interface (FC-PH), Rev.4.2.

OSA-Express GbE

The OSA-Express GbE features support the standard transmission scheme, 1000BASE-SX and 1000BASE-LX.

The lowest layer of the Gigabit Ethernet (IEEE 802.3z) protocol stack uses the FC-0 layer from the Fiber Channel specification. FC-0 describes the physical characteristics of the interface and media, including the cables and connectors. Fiber optic cables are required for connecting an OSA-Express GbE feature to the LAN through a hub, switch, or router supporting GbE 1000BASE-SX or GbE 1000BASE-LX interfaces.

If multimode fiber is used with the long wavelength (LX) transceiver, mode conditioning patch (MCP) cables are required, and must be installed in pairs at both the sending and receiving sides of the fiber optic multimode links. The MCP part numbers are:

- 50 micron:
  - PN 21L4172: Terminated with an SC duplex coupler
  - PN 21L4174: Terminated with an ESCON coupler
- 62.5 micron:
  - PN 21L4173: Terminated with an SC duplex coupler
  - PN 21L4175: Terminated with an ESCON coupler

Depending on the type of adapter, fiber bandwidth and use of mode conditioning cables, different maximum cable lengths between the OSA-Express GbE feature and the network device are supported. Table 1-1 on page 24 shows the maximum allowed fiber cable length for the SX and LX features.

Table 1-1 Maximum cable length

Transceiver	Fiber	Distance
SX	62.5 micron multimode	220 m (722 ft)[a] 275 m (902 ft)[b]
SX	50 micron multimode	550 m (1,804 ft)
LX	62.5 micron multimode with a pair of mode-conditioning patch cables	550 m (1,804 ft)
LX	50 micron multimode with a pair of mode-conditioning patch cables	550 m (1,804 ft)
LX	9 micron single mode	5 km (3.1 miles)

a. For 160 MHz-km fiber bandwidth
b. For 200 MHz-km fiber bandwidth

OSA-Express FENET

Connector type: RJ-45

Cable type: Standard transmission scheme, 10BASE-T or 100BASE-TX, which requires two pairs of Category 5 balanced cable, or two pairs of 150 ohm shielded balanced cable as defined by ISO/IEC 11801. One pair of wires is for transmitting and the second pair is for receiving.

The category 5 cable is typically referred to as Cat 5 UTP (Unshielded Twisted Pair) or Cat 5 STP (Shielded Twisted Pair).

OSA Express Token Ring

Connector type: RJ-45

Cable type: Standard Category 5 cable, shielded or unshielded. One pair of wires is for transmitting and the second pair is for receiving.

Connector type: DB-9 D Shell

Cable type: Shielded twisted pair (STP) cable with a DB-9 D Shell connector.

OSA-Express 155 ATM

1. Single mode feature:

   Connector type: SC Duplex

   Cable type: Single mode fiber (9 micron)

2. Multi mode feature:

   Connector type: SC Duplex

   Cable type: Multi mode fiber (62.5 micron)

Trunk cables

Fiber optic trunks terminated in distribution panels are recommended in large single-room data centers, multiple room and multiple floor data centers, and data centers spread among buildings or sites. *Fiber Optic Link Planning*, GA23-0367, should always be consulted when requirements include fiber optic trunk facilities.

Use of trunks requires additional jumper cables and components, which are available from IBM Customer Service or from IBM Authorized Cabling Distributors.

For additional fiber optic components and connectivity solutions, contact your local IBM office Installation Planning Representative or your Area Cabling Specialist.

Assistance in planning, design, and installation of fiber optic cabling systems is available. For more details, contact your IBM representative.

1.10 Software support

Refer to the following tables for the minimum operating system release required to support the various OSA-Express features.

Note the following:

- Additional maintenance may be required for the minimum release.
- It is assumed that the Communications Server (CS), which provides VTAM and TCP/IP support for the z/OS and OS/390 operating systems, is also installed.

1.10.1 Operating systems

The following sections list the supported levels of the operating systems for the OSA-Express features.

OS390, z/OS and z/OS.e

z/OS and z/OS.e

All releases of z/OS and z/OS.e support the OSA-Express features.

9672 G5/G6 servers running OS/390

Use Table 1-2 to identify the release of OS/390 required for the OSA-Express feature to be implemented.

Table 1-2 OS/390 on 9672 G5/G6 servers

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM
QDIO (OSD) TCP/IP only	OS/390 R7 with PTFs	OS/390 R8	OS/390 R8 for Ethernet LANE only
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	OS/390 R6	OS/390 R6
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	OS/390 R6	OS/390 R6
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	OS/390 R6

zSeries running OS/390

Use Table 1-3 on page 26 to identify the release of OS/390 required for the OSA-Express feature to be implemented.

Table 1-3 OS/390 on zSeries

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM	OSA-Express Token Ring
QDIO (OSD) TCP/IP only	OS/390 R7 with PTFs	OS/390 R8	OS/390 R8 for Ethernet LANE only	OS/390 R10
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	OS/390 R6	OS/390 R6	OS/390 R9
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	OS/390 R6	OS/390 R6	OS/390 R9
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	OS/390 R6	

VM/ESA and z/VM

z/VM

All releases of z/VM support the OSA-Express features.

9672 G5/G6 servers running VM/ESA

Use Table 1-4 to identify the release of VM/ESA required for the OSA-Express feature to be implemented.

Table 1-4 VM on 9672 G5/G6 servers

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM
QDIO (OSD) TCP/IP only	VM/ESA V2R4 for QDIO operation on a guest	VM/ESA V2R4 for QDIO operation on a guest	VM/ESA V2R4 for QDIO operation on a guest
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	VM/ESA V2R3	VM/ESA V2R3
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	VM/ESA V2R3	VM/ESA V2R3
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	VM/ESA V2R4 - IP

zSeries running VM/ESA

Use Table 1-5 on page 27 to identify the release of VM/ESA required for the OSA-Express feature to be implemented.

Table 1-5 VM on zSeries

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM	OSA-Express Token Ring
QDIO (OSD) TCP/IP only	VM/ESA V2R4 for QDIO operation on a guest z/VM V3R1 for native QDIO	VM/ESA V2R4 for QDIO operation on a guest z/VM V3R1 for native QDIO	VM/ESA V2R4 for QDIO operation on a guest z/VM V3R1 for native QDIO	VM/ESA V2R4 for QDIO operation on a guest z/VM V3R1 for native QDIO
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	VM/ESA V2R3	VM/ESA V2R3	VM/ESA V2R4
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	VM/ESA V2R3	VM/ESA V2R3	VM/ESA V2R4
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	VM/ESA V2R4 - IP	Not applicable

VSE/ESA

9672 G5/G6 servers running VSE/ESA

Use Table 1-6 to identify the release of VSE/ESA required for the OSA-Express feature to be implemented.

Table 1-6 VSE/ESA on 9672 G5/G6 servers

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM
QDIO (OSD) TCP/IP only	VSE/ESA V2R6	VSE/ESA V2R6	VSE/ESA V2R6
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	VSE/ESA V2R6	VSE/ESA V2R6
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	VSE/ESA V2R6	VSE/ESA V2R6
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	Not supported

zSeries running VSE/ESA

Use Table 1-7 to identify the release of VSE/ESA required for the OSA-Express feature to be implemented.

Table 1-7 VSE/ESA on zSeries

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM	OSA-Express Token Ring
QDIO (OSD) TCP/IP only	VSE/ESA V2R6	VSE/ESA V2R6	VSE/ESA V2R6	VSE/ESA V2R6
Non-QDIO (OSE) TCP/IP Passthru	Not applicable	VSE/ESA V2R6	VSE/ESA V2R6	VSE/ESA V2R6
Non-QDIO (OSE) SNA, APPN, HPR	Not applicable	VSE/ESA V2R6	VSE/ESA V2R6	VSE/ESA V2R6

Mode and (CHPID type)	OSA-Express Gigabit Ethernet	OSA-Express Fast Ethernet	OSA-Express 155 ATM	OSA-Express Token Ring
Non-QDIO (OSE) HPDT ATM Native	Not applicable	Not applicable	Not supported	Not applicable

Transaction Processing Facility (TPF)

TPF 4.1 PUT 13 provides support for the OSA-Express GbE feature; see TPF APAR PJ27333.

Linux

Linux for zSeries and Linux for S/390 can be implemented as a native image, LPAR image, S/390 VIF for Linux image, or as a z/VM or VM/ESA guest with TCP/IP for VM feature. In all these implementations, Linux for zSeries and Linux for S/390 support:

- OSA-Express GbE in QDIO mode (Linux Kernel 2.2.16)
- OSA-Express FENET in QDIO and TCP/IP Passthru non-QDIO (Linux Kernel 2.2.16)
- OSA-Express Token Ring in QDIO mode (Linux Kernel 2.4.0)

1.10.2 OSA/SF

This section provides information concerning the release level of OSA/SF required on the various operating systems.

OSA/SF for z/OS and OS/390

Open Systems Adapter Support Facility Version 2 Release 1 (Program Number 5655–B57) for OS/390 (Program Number 5647-A01), z/OS (5694-A01), and z/OS.e (5655–G52) or later is required.

OSA/SF V2 R1 for z/OS, z/OS.e, and OS/390 requires:

- A host editor to edit data sets (for example ISPF, which is shipped as an element of OS/390)
- Language Environment (LE), shipped as an element of OS/390
- Time Sharing Options Extensions (TSO/E), shipped as an element of OS/390, to access the OSA/SF for OS/390 V2R1 REXX EXEC from the TSO/E command line

OSA/SF for VM/ESA or z/VM

Open Systems Adapter Support Facility for Virtual Machine/Enterprise Systems Architecture (VM/ESA) Version 2 Release 2.0 (Program Number 5654-030) and z/VM Version 3 Release 1 (Program Number 5654-A17) or later is required.

OSA/SF for VSE/ESA or z/VSE

OSA/SF for VSE Version 2 Release 2 (part of VSE Central Functions 6.1.1, 5686-066) in VSE/ESA Version 2 Release 2.6 (5690-VSE) or later is required.

OSA/SF GUI workstation

One of the following is required:

- Microsoft Windows 95 4.00.950B (OSR2) or later
- Microsoft Windows NT 4.0 with Service Pack 3 or later
- OS/2 (R) Warp Version 4.0 fixpack 6 or later

Connection Types (OSA/SF GUI to OSA/SF Server)

The OS/SF GUI supports three different connection types: TCP/IP, APPC, and EHLLAPI. Any one of these can be used.

- TCP/IP only requires that you create an OSA/SF TCP/IP server (IOASRV) and then add the host (server) connection icon on the OSA/SF GUI Hosts window. The GUI communicates to the server using OSA/SF supplied server code running on the server. Although the user does not directly log on to a TSO user ID, access to a TSO user ID is required for this connection. The TCP/IP server uses the user ID to verify the user has security access to the OS/390 OSA/SF server resources.
- APPC is the most complicated connection type to set up. You need to define a transaction program and the LU6.2 connection on the server, and define the LU6.2 connection on the workstation in PCOM. APPC is used to communicate with the OSA/SF server using the supplied transaction program installed with the GUI. As with TCP/IP, access to a TSO user ID is required for security access authorization at the OS/390 OSA/SF server.
- EHLLAPI is the easiest connection to set up. The connection requires only that you add a server connection icon on the OSA/SF GUI Hosts window. EHLLAPI uses a 3270 session and Personal Communications Manager (PCOM 4.2 or later).

  OSA/SF uses EHLLAPI to control a TSO or CMS user ID on the 3270 session. It transfers files between the workstation and the TSO or CMS user ID, and then uses the EHLLAPI interface to enter TSO or CMS commands to communicate with OSA/SF on the server. EHLLAPI requires a TSO or CMS user ID to be logged on to the workstation (window) and to be at the TSO or CMS ready prompt.

Refer to *Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7476, for detailed information concerning OSA/SF.

1.11 Summary

The OSA-Express features provide direct LAN and network connectivity as integrated features of the zSeries and G5/G6 servers. This brings the strengths of z/Architecture and S/390 architecture to the client/server environment.

Following are the key characteristics of the OSA-Express features on zSeries and 9672 G5/G6 servers:

- Support for standard Token Ring, Gigabit Ethernet, and Fast Ethernet LANs, as well as Asynchronous Transfer Mode networks.
- Up to 12 ports on a 9672 G5/G6 server, and 24 ports on a zSeries server, helping to reduce the need for external gateways. Each OSA-Express feature has one port on the 9672 G5/G6 servers, and two ports on the zSeries server.
- OSA-Express features utilize full duplex, direct attachment to the zSeries and S/390 Self-Timed Interface (STI) infrastructure to deliver high bandwidth paths to applications and line speed performance. The STI attachment runs at 333 MBps, far faster than older ESCON devices or OSA-2 that connect to the 17 MBps Channel Request Handler bus.
- OSA-Express features (Token Ring, GbE, FENET, and 155 ATM (Ethernet or Token Ring LANE)) can utilize Queued Direct Input/Output (QDIO). QDIO is a highly efficient data transfer architecture, which dramatically improves data transfer speed and efficiency for TCP/IP traffic.

  QDIO can support SNA/APPN/HPR traffic through the use of Enterprise Extender.

QDIO incorporates a number of features:

– Direct Memory access (DMA) allows data to move directly from the OSA-Express microprocessor to the host memory. This bypasses three layers of processing that are required when using ESCON and OSA-2 features, dramatically improving throughput.

– IP Assist allows compute-intensive functions usually handled by the host system to be performed by the OSA-Express feature; for example, multicast support, broadcast filtering, building MAC and LLC headers, and ARP processing and statistics gathering. This reduces the load on the host system, freeing resources for increased application workloads.

– LPAR-to-LPAR communication via the OSA-Express feature. This allows data being sent from one LPAR to another within the same system to be routed directly through a shared OSA-Express feature, thus completely bypassing the LAN.

– Priority Queuing (for z/OS and OS/390 environments) sorts outgoing IP message traffic according to the priority assigned in the IP header. This priority reflects the user's business priorities assigned to the application, user ID, time of day and other characteristics.

– Dynamic building of the OSA Address Table (OAT). This reduces configuration and setup time, eliminates duplicate data entry, and reduces the chance of data entry errors and incompatible definitions. The TCP/IP profile information only needs to be entered once. It is passed to the OSA-Express feature at device startup, and the OAT is then built and maintained dynamically.

– Enhanced IP network availability (IPA) is a service of QDIO architecture. When TCP/IP is started in QDIO mode, it downloads all the home IP addresses in the stack and stores them in the OSA-Express feature. The OSA-Express feature port then responds to ARP requests for its own IP address, as well as for virtual IP addresses (VIPAs).

  If an OSA-Express feature fails while there is a backup OSA-Express available on the same network or subnetwork, TCP/IP informs the backup OSA-Express feature port which IP addresses (real and VIPA) to take over, and sends a gratuitous ARP which contains the MAC address of the backup OSA-Express. The network connection is maintained.

- The OSA-Express features support Virtual IP Address (VIPA) in the TCP/IP environment. VIPA frees TCP/IP hosts from dependence on a particular network attachment. VIPA provides for multiple IP addresses to be defined to a TCP/IP stack, allowing fault-tolerant, redundant, backup paths to be established. Applications become insensitive to the condition of the network since the VIPA will always be active, enabling users to route around intermediate points of failure in the network.

  VIPA Takeover and Takeback allow for fast connection recovery in the event of a TCP/IP stack services or application failure. Should a stack or configured application fail, the IP addresses are transferred to backup copies and the OSA-Express features using QDIO are updated dynamically. Network access services can continue, uninterrupted.

  This VIPA support also facilitates movement of applications and workload within the sysplex in a manner transparent to end-user connection requests.

- OSA-Express allows network traffic from multiple LPARs to travel through the same port. TCP/IP and SNA/APPN/HPR traffic can be mixed on the same port simultaneously.

2

Quick start

In this chapter, we provide information to help you get a quick start with your OSA-Express installation. The chapter will assist in determining which elements are needed based on your requirements, as well as direct you to the appropriate sections in this book for details.

© Copyright IBM Corp. 2003. All rights reserved.

2.1 OSA-Express definitions

Before you can exploit your OSA-Express feature, it must first be properly defined to hardware and software. Here are items to identify and determine:

1. Which CHPID will the OSA-Express port be assigned to?
2. Will the OSA-Express CHPID be shared or dedicated?
3. Will the OSA-Express CHPID be defined as OSD (QDIO) or OSE (non-QDIO)?
4. Which protocol(s) will be used with the OSA-Express port (TCP/IP, SNA, or both)?

The first three items deal with defining the OSA-Express to hardware. For detailed information regarding this subject, refer to Chapter 3, “Hardware configuration definitions” on page 37.

Table 2-2 on page 34 provides a quick reference relating CHPID type to operation mode.

2.2 OSA/SF requirements

OSA/SF is used to configure the OSA-Express features for non-QDIO mode, with one exception; TCP/IP Passthru mode (non-shared).

You will need OSA/SF for the 155 ATM feature even if it is being configured for QDIO mode (ATM Ethernet LAN Emulation). This is because you must configure the port definitions for the 155 ATM feature.

For a quick check to determine if OSA/SF is required for your installation, refer to Table 2-1.

Table 2-1 OSA/SF requirement

OSA-Express features	OSD (QDIO)	OSE (non-QDIO)	OAT Built	OSA/SF
GbE	Yes		Dynamic	No
FENET	Yes		Dynamic	No
FENET		Yes	Manual	Yes
Token Ring	Yes		Dynamic	No
Token Ring		Yes	Manual	Yes
ATM ENET LANE TR LANE ENET LANE	 Yes 	 Yes Yes	 Dynamic Manual Manual	 Yes[a] Yes Yes
ATM Native		Yes	Manual	Yes

a. use OSA/SF for the port definitions.

For detailed information, see Chapter 4, “Setting up and using OSA/SF” on page 47.

2.3 Policy-based networking

UNIX Policy Agent (PAGENT) is not required for OSA-Express. It is an option that can be implemented to manage and prioritize network traffic. For more information on PAGENT, see “Priority queuing” on page 6.

2.4 TCP/IP quick start

If you have not yet configured your zSeries or 9672 G5/G6 TCP/IP environment, you may want to consider using an IBM wizard or configuration tool.

z/OS TCP/IP Wizard

IBM provides a Web site that has access to a number of wizards that can assist with configuring your z/OS system. Included is a wizard for configuring TCP/IP. By inputting your network definitions, the wizard will provide you with configuration information required to setup your z/OS TCP/IP environment

For information on z/OS wizards and links to specific wizards, go to the following Web site:

`http://www-1.ibm.com/servers/eserver/zseries/zos/wizards/`

z/VM V4R3 TCP/IP Wizard

A TCP/IP configuration wizard is packaged with z/VM V4.3. This new utility automates the connection of a newly installed z/VM system to a TCP/IP-based network. The z/VM TCP/IP configuration wizard requires no knowledge of TCP/IP for z/VM, and is similar to the network configuration utilities used in Linux for zSeries distributions during Linux installation.

This easy-to-use configuration wizard assists the z/VM installer in providing desired IP configuration information such as host and domain name, IP addresses, and subnet mask. From that information, the wizard will generate an initial z/VM TCP/IP configuration and verify that connectivity to the network has been established.

This utility is on the MAINT 193 minidisk. It can be run one time for the initial definition of a single connection. The exec is named IPWIZARD. For more information on IPWIZARD, refer to *z/VM Installation Guide Version 4 Release 3*, GC24-5992.

VSE/ESA

The TCP/IP for VSE/ESA Configuration Support is a PC-based front-end used to create a TCP/IP for VSE/ESA configuration and start up related files. The dialog is available in an OS/2 version and a Windows version. The Windows version runs under Win95/98, Windows NT 4.0, and Windows 2000. Supported languages are English, German, Spanish, and Japanese.

The dialog operates on PC files, that is, it reads and writes from and to a PC hard disk (local or LAN). A batch file is created by the dialog to help upload output files to the host.

Linux

Linux also has a configuration tool that can be used to build your network configuration. The distribution of Linux that you use determines the setup tool used (for example, SuSE Linux uses YAST).

2.5 Quick start tables

In this section, we provide tables that will help you to identify your OSA-Express feature and implementation requirements, and also direct you to installation information and set up examples.

For detailed information about the different types of OSA-Express features, refer to Chapter 1, “Overview” on page 1.

2.5.1 Notes for quick start tables

When reviewing Table 2-2 or Table 2-3, "z/VM quick start table" on page 35, keep the following notes in mind for any additional information related to your OSA-Express implementation:

1. If you are using the *default values* for Fast Ethernet or Token Ring, OSA/SF is not required.

   Default values mean:

   – Non-shared CHPID

   – CHPID type OSE

   – The port is not shared between partitions, TCP/IP stacks, or protocols (TCP/IP and SNA)

   – Using the default unit addresses (starting with 00)

   – Using TCP/IP Passthru only
2. OSA/SF configuration is not required for GbE, FENET, or Token Ring (when defined as OSD).
3. OSA/SF is required to complete the configuration of the ATM feature under all circumstances.
4. If SNA/APPN (LU6.2) is required for an OSA-Express feature that is defined as QDIO (CHPID type OSD), you must use Enterprise Extender. This is described in Chapter 10, "Enterprise Extender (EE)" on page 139.
5. TN3270 can be used in conjunction with SNA (LU2) applications, when the OSA-Express feature is defined as QDIO (CHPID type OSD).

z/OS quick start table

Table 2-2 contains CHPID, TCP/IP profile, and VTAM definitions used when configuring the OSA-Express cards for use in a z/OS environment.

Table 2-2 z/OS quick start table

OSA-Express feature	Operation mode	CHPID type	TCP/IP device type	TCP/IP link type	VTAM definitions	Go to page
GbE	QDIO TCP/IP	OSD	MPCIPA	IPAQENET	TRLE	61
FENET	QDIO TCP/IP	OSD	MPCIPA	IPAQENET	TRLE	61
FENET	Non-QDIO TCP/IP	OSE	LCS	ETHERNet, 802.3, or ETHEROR802.3		95
FENET	Non-QDIO SNA	OSE			XCA, SWNET	103
Token Ring	QDIO TCP/IP	OSD	MPCIPA	IPAQTR	TRLE	61
Token Ring	Non-QDIO TCP/IP	OSE	LCS	IBMTR		95
Token Ring	Non-QDIO SNA	OSE			XCA, SWNET	103

OSA-Express feature	Operation mode	CHPID type	TCP/IP device type	TCP/IP link type	VTAM definitions	Go to page
ATM LANE (ETH)	QDIO TCP/IP	OSD	MPCIPA	IPAQENET	TRLE	61 and 117
ATM LANE (TR or ETH)	Non-QDIO TCP/IP	OSE	LCS	ETHERNet, 802.3, ETHEROR802.3, or IBMTR		117
ATM LANE (TR or ETH)	Non-QDIO SNA	OSE			XCA, SWNET	117
ATM Native IP traffic	Non-QDIO HPDT	OSE	ATM	ATM	TRLE	101
ATM Native SNA traffic	Non-QDIO HPDT	OSE			TRLE, XCA, SWNET	101

z/VM quick start table

Table 2-3 contains CHPID and TCP/IP profile definitions used when configuring the OSA-Express cards for use in a z/VM environment.

Table 2-3 z/VM quick start table

OSA-Express feature	Operation mode	CHPID type	TCP/IP device type	TCP/IP link type	VTAM definitions
GbE	QDIO TCP/IP	OSD	OSD	QDIOETHERNET	
FENET	QDIO TCP/IP	OSD	OSD	QDIOETHERNET	
FENET	Non-QDIO TCP/IP	OSE	LCS	ETHERNET, 802.3, or ETHEROR802.3	
FENET	Non-QDIO SNA	OSE			XCA, SWNET
Token Ring	QDIO TCP/IP	OSD	OSD	QDIOTR	
Token Ring	Non-QDIO TCP/IP	OSE	LCS	IBMTR	
Token Ring	Non-QDIO SNA	OSE			XCA, SWNET
ATM LANE (ETH)	QDIO TCP/IP	OSD	OSD	QDIOATM	

Note: In this redbook, we do not go into detail regarding how to define the OSA-Express features to z/VM. However, by using the information provided in Table 2-3 and the processes provided in this book, as well as the example of a TCP/IP profile for z/VM included in "z/VM TCP/IP profile" on page 216, you should be able to update your z/VM TCP/IP profile with the correct information.

2.6 Linux definitions

There are several files that require updating for OSA-Express implementation in a Linux environment. Refer to “Linux definitions” on page 215 for examples of our implementation.

3

Hardware configuration definitions

As with all channel-attached devices, an OSA-Express feature must be defined by a channel path, a control unit, and I/O devices in the IOCDS.

In this chapter, we discuss the steps needed to define the OSA-Express to the zSeries and/or 9672 G5/G6 environment, using the z/OS Hardware Configuration Definition (HCD) tool. To simplify the configuration process of the environment shown in Figure 3-1 on page 38, we have extracted the portions of the setup that are common to all modes and types of the OSA-Express.

A *CHIPD Report* that specifies where the OSA-Express feature is plugged into your server can be supplied by your IBM representative. The CHPID number will be required for all OSA-Express configuration and setup tasks.

© Copyright IBM Corp. 2003. All rights reserved.

3.1 Configuration chart

In the environment shown in Figure 3-1, we used two zSeries servers, each configured with multiple logical partitions (LPARs). Server SCZP701 has OSA-Express Gigabit Ethernet attached to Channel Path ID (CHPID) F4, OSA-Express Fast Ethernet to CHPID FA, and OSA-Express 155 ATM feature to CHPID F6. These channels are shared between two partitions (SC47 and SC69).

Server SCZP601 has OSA-Express Fast Ethernet attached to CHPID FA and OSA-Express 155 ATM to CHPID F6. These channels are also shared by two partitions (SC53 and SC66).

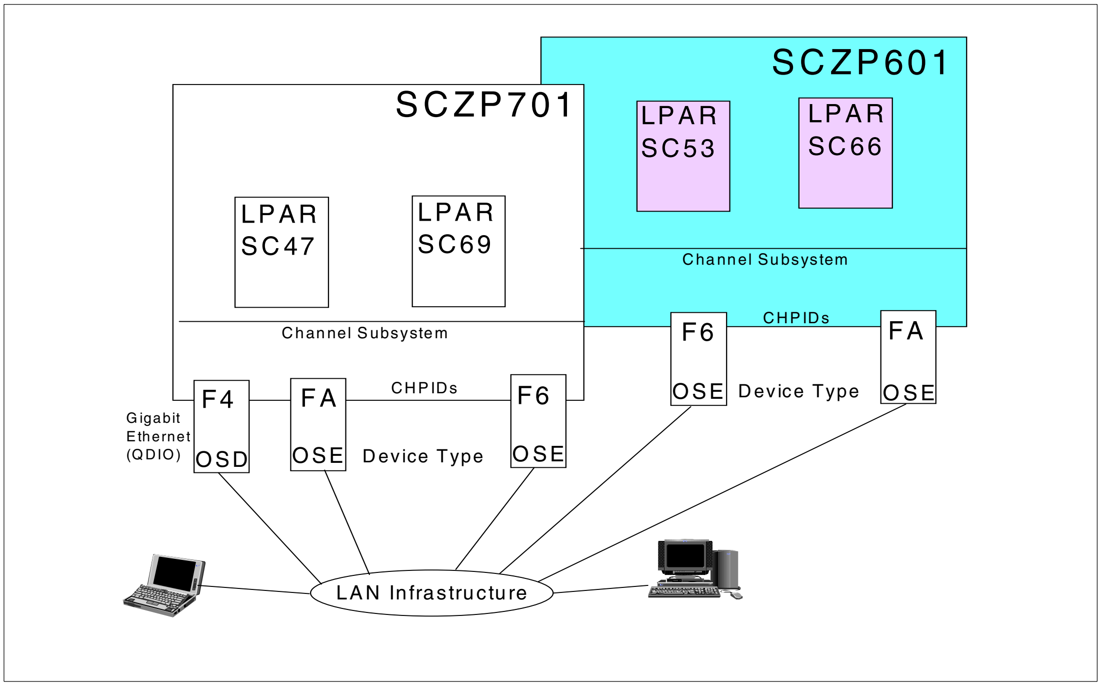

Figure 3-1 HCD definitions for OSA-Express CHPIDs

3.2 HCD definition

HCD is required to define the OSA-Express feature to the I/O hardware configuration. We have included examples of the following definitions:

- Channel path
- Control unit
- Devices

3.2.1 Channel path definition

We did our HCD definitions on a z/OS 1.4 system. Following are the steps we used, starting from the HCD main menu panel shown in Figure 3-2 on page 39.

```
                          OS/390 Release 9 HCD
Command ===> ________________________________________________________

                        Hardware Configuration

Select one of the following.

1  1.  Define, modify, or view configuration data
   2.  Activate or process configuration data
   3.  Print or compare configuration data
   4.  Create or view graphical configuration report
   5.  Migrate configuration data
   6.  Maintain I/O definition files
   7.  Query supported hardware and installed UIMs
   8.  Getting started with this dialog
   9.  What's new in this release

For options 1 to 5, specify the name of the IODF to be used.

I/O definition file . . . 'SYS6.IODF54'                  +
```

Figure 3-2 HCD main menu

1. Select **1**.
2. On the next screen (“Define, modify, or view configuration data”), select Option **3 - Processors**.
3. Select the processor you want to update and press Enter.

We identify the panel selection options by the action code rather than the item number to avoid confusion when a particular HCD menu changes.

4. After you have selected the processor, press Enter.
5. Select “Work with attached channel paths” by using **S** in the option field.
6. On the Channel Path List screen, press F11 to add a channel path.
7. Enter all the required information, as shown in Figure 3-3 on page 40.

   In our scenario, we set:

   – Channel path type to OSD, because we are defining GbE (which uses QDIO).

   – Operation mode to SHR, because the feature will be shared among LPARs.

   We recommend using a meaningful description, which will serve as a reference point in HCD.

```
                     Add Channel Path


Specify or revise the following values.

Processor ID . . . . . : SCZP701
Configuration mode : LPAR

Channel path ID  . . . . . . . . . F4  +
Number of CHPIDs . . . . . . . 1
Channel path type  . . . . . . . OSD  +
Operation mode . . . . . . . . . SHR  +
Managed  . . . . . . . . . . . . . . No   (Yes or No)   I/O Cluster ________  +
Description  . . . . . . . . . . . . Gigabit Ethernet________________

Specify the following values only if connected to a switch:

Dynamic switch ID  . . . . . . __  + (00 - FF)
Entry switch ID  . . . . . . . . . __  +
Entry port  . . . . . . . . . . . . . __  +

F1=Help    F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap
F12=Cancel
```

Figure 3-3   Add Channel Path menu

This example shows how to configure a Gigabit Ethernet OSA-Express feature. The process is identical for Fast Ethernet or ATM with the following exceptions:

– The Channel path type must be OSD for QDIO support, or OSE for non-QDIO support.

– The Operation mode must be DED to dedicate the port to a single LPAR, or SHR to share among LPARs.

8. When the screen definitions are complete, press Enter.
9. Complete the Access List for the partitions sharing the channel, as shown in Figure 3-4 on page 41, and press Enter.

   In this example, two partitions share the OSA-Express channel.

```
                                        Row 1 of 15
Command ===> _________________________________________ Scroll ===> PAGE

Select one or more partitions for inclusion in the access list.

Channel path ID . . : F4     Channel path type  . : OSD
Operation mode . . . : SHR     Number of CHPIDs . . : 1

/ Partition Name   Number Usage Description
_ A10           A     OS
/ A11           B     OS      SC69
/ A12           C     OS      SC47
_ A13           D     OS
_ A2            2     OS
_ A3            3     OS
_ A4            4     OS
_ A5            5     OS
_ A6            6     OS
_ A7            7     OS
_ A8            8     OS
 F1=Help      F2=Split     F3=Exit      F5=Reset      F6=Previous
 F7=Backward   F8=Forward   F9=Swap      F12=Cancel
```

Figure 3-4 Access List Definition

Once you select the partitions for the access list and press Enter, a panel will be displayed for the candidate list.

10. Select the partitions to include in the candidate list and press Enter, or simply press Enter if you do not want any partitions in the candidate list.

3.2.2 Control Unit definition

You should now be back to the Control Unit List panel.

1. Select the CHPID just defined (F4, in our example), and press Enter.
2. Select “Work with attached control units” by using **S** in the option field**.**
3. Press F11 to add a control unit.

   Enter the required information, as shown in Figure 3-5 on page 42.

   In our example, we set:

   – Control unit number to `2380`.

   – Control unit type to `OSA`.

```
                    Add Control Unit

Specify or revise the following values.

Control unit number . . . . . 2380  +
Control unit type  . . . . . . . OSA_________   +

Serial number  . . . . . . . . . __________
Description  . . . . . . . . . . Gigabit Ethernet C.U. (OSD)_____

Connected to switches . . . __ __ __ __ __ __ __ __  +
Ports . . . . . . . . . . . . . . . . __ __ __ __ __ __ __ __  +

If connected to a switch:

Define more than eight ports . . 2   1. Yes
                                     2. No
Propose CHPID/link addresses and
unit addresses . . . . . . . . . . . . . 2   1. Yes
                                     2. No
```

Figure 3-5 Add Control Unit part 1

4. Press Enter when ready to proceed.
5. Select the processor for the control unit, as shown in Figure 3-6.

```
Goto  Filter  Backup  Query  Help
             Select Processor / Control Unit
                                                   Row 1 of 5 More:      >
Command ===> _______________________________________________ Scroll ===> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : 2380     Control unit type . . . : OSA

         Log. Addr. -------Channel Path ID . Link Address + -------
/ Proc. ID Att. (CUADD) +  1---- 2---- 3---- 4---- 5---- 6---- 7---- 8----
_ SCZP601     __        _____ _____ _____ _____ _____ _____ _____ _____
/ SCZP701     __        _____ _____ _____ _____ _____ _____ _____ _____
***************************** Bottom of data ******************************
```

Figure 3-6 Processor selection

6. Select the processor and press Enter.
7. On the “Actions on Selected Processors” screen, select **S** for Select (connect, change), and then press Enter.

   Figure 3-7 on page 43 shows the OSA control unit information. Note that the Unit address must be set to 00 and the number of units must be 255.

```
                    Add Control Unit

Specify or revise the following values.

Control unit number  . : 2380      Type . . . . . . : OSA
Processor ID  . . . . . . : SCZP701

Channel path IDs . . . . F4  __  __  __  __  __  __  __ +
Link address  . . . . . . .__  __  __  __  __  __  __  __ +

Unit address  . . . . . . . 00  __  __  __  __  __  __  __ +
Number of units . . . . . 255  ___ ___ ___ ___ ___ ___ ___

Logical address . . . . . . __  + (same as CUADD)

Protocol  . . . . . . . . . . . __  + (D,S or S4)
I/O concurrency level . . 2  + (1, 2 or 3)
```

Figure 3-7 Add Control Unit part 2

8. Once the values are provided, press Enter.
9. Press Enter again to return to the Control Unit List panel.

3.2.3 Device definition

From the Control Unit List panel:

1. Select the control unit (2380, in our example).
2. Select **S** -- Work with attached devices.
3. Press F11 to add devices.
4. Enter the required information as shown in Figure 3-8 on page 44.

   In our scenario, we set:

   – Device number to `2380`.

   – Number of devices to `15`.

   – Device type to `OSA`.

   How many devices should you define? This depends on things like CHPID type, TCP/IP stacks, and SNA definitions required. For the number of devices, consider the following:

   – Any OSD CHPID requires at least three devices (read, write, and datapath) for each TCP/IP.

   – Any OSE CHPID requires two devices (read and write) for each TCP/IP. SNA use requires one device (with the exception of ATM; for SNA, it requires two devices).

   The defined devices can be shared within a protocol. That is, a device can not be shared between TCP/IP and SNA.

```
                         Add Device


 Specify or revise the following values.

 Device number .........2380   (0000 - FFFF)
 Number of devices ......15___
 Device type ...........OSA__________ +

 Serial number ..........__________
 Description ...........Gigabit Ethernet______________

 Volume serial number ....______ (for DASD)

 Connected to CUs ..2380 ____ ____ ____ ____ ____ +
```

Figure 3-8 Add devices

5. When you press Enter, a panel will be displayed to allow you to edit information for the specific devices. Make any changes you need, and press Enter.
6. You should now have the Device / Processor Definition panel displayed.
7. Select the processor by placing a slash (**/**) to the left of the one you want.

```
                    Define Device / Processor

  Specify or revise the following values.

 Device number .:2380        Number of devices ....:15
 Device type ....:OSA
 Processor ID ...:SCZP701

 Unit address ..........00 + (Only necessary when different
                        from the last 2 digits of device number)
 Time-Out ............No  (Yes or No)
 STADET .............No  (Yes or No)

 Preferred CHPID ........__ +
 Explicit device candidate list .No  (Yes or No)
```

Figure 3-9 Define Device / Processor panel

8. On this panel, you have the option of changing the starting unit address. Verify the value (00 is only required with the default OAT), and press Enter.
9. Press Enter again to display the Define Device to Operating System Configuration panel.
10. Place an **S** to the left of the operating system to which you want to connect the devices and press enter.
11. Accept the default values on the resulting panels by pressing Enter.
12. Repeat the process for each operating system as needed.
13. You should now be at the Device List panel.

14.If you plan to use OSA/SF, an OSAD device must be defined. Press F11 to add a new device and define it, as shown in Figure 3-10.

```
                         Add Device

 Specify or revise the following values.

 Device number . . . . . . . . . 238F   (0000 - FFFF)
 Number of devices . . . . . . 1____
 Device type  . . . . . . . . . . . OSAD_________  +

 Serial number . . . . . . . . . . __________
 Description  . . . . . . . . . . . Gigabit Ethernet FE___________

 Volume serial number . . . . ______  (for DASD)

 Connected to CUs . . 2380  ____ ____ ____ ____ ____  +
```

Figure 3-10 OSAD Definition 1

15.Repeat the process as you did for the other devices, with the following exception:

– Associate the unit address FE with the OSAD device (238F), as shown in Figure 3-11, and then press Enter.

```
                    Define Device / Processor

 Specify or revise the following values.

 Device number . : 238F        Number of devices . . . . : 1
 Device type . . . . : OSAD
 Processor ID . . . : SCZP701

 Unit address . . . . . . . . . . FE  + (Only necessary when different
                           from the last 2 digits of device number)
 Time-Out . . . . . . . . . . . . No  (Yes or No)
 STADET . . . . . . . . . . . . . No  (Yes or No)

 Preferred CHPID . . . . . . . __  +
 Explicit device candidate list . No  (Yes or No)
```

Figure 3-11 OSAD definition 2

16.Define the device with these parameters to the operating system configuration, as shown in Figure 3-12 on page 46, and then press Enter.

```
                 Define Device Parameters / Features

                                                           Row 1 of 3
Command ===> ______________________ Scroll ===> PAGE

Specify or revise the values below.

Configuration ID . : L06RMVS1     Sysplex systems
Device number  . . : 238F         Number of devices  : 1
Device type  . . . : OSAD

Parameter/
Feature           Value  P Req.  Description
OFFLINE            No            Device considered online or offline at IPL
DYNAMIC           Yes            Device has been defined to be dynamic
LOCANY            No             UCB can reside in 31 bit storage
*************************** Bottom of data ****************************
```

Figure 3-12 OS configuration

3.2.4 Generating the input IOCDS from HCD

You can generate input for IOCDS from HCD, which is then used to write the macro definitions to the zSeries server, and can also be used for quick debugging purposes. Figure 3-13 shows an example of IOCDS input, generated by HCD.

```
ID    MSG2='SYS6.IODF64 - 1999-02-15 10:29',SYSTEM=(9672,6),
TOK=('SCZP701',00000001930596721029492000999046F00000000,0000,'99-02-15','10:29:49','SYS6','IODF64')
RESOURCE PARTITION=((A1,1),(A10,A),(A11,B),(A12,C),(A2,2),(A3,3),(A4,4),(A5,5),(A6,6),
(A7,7),(A8,8),(A9,9),(C1,D),(C2,E),(C3,F))
.
CHPID PATH=(F4),SHARED,PARTITION=((A11,A12),(A11,A12)),TYPE=OSD
.
CNTLUNIT CUNUMBR=2380,PATH=(F4),UNIT=OSA
.
IODEVICE ADDRESS=(2380,015),UNITADD=00,CUNUMBR=(2380),UNIT=OSA
IODEVICE ADDRESS=238F,UNITADD=FE,CUNUMBR=(2380),UNIT=OSAD
```

Figure 3-13 IOCDS input

4

Setting up and using OSA/SF

The Open Systems Adapter Support Facility (OSA/SF) is an application to customize and manage the OSA features, such as changing port parameters, and device and LPAR definitions. Also, status and operational information about the HCD-defined OSA features can be obtained to assist in problem determination.

OSA/SF includes a REXX interface and a graphical user interface (GUI) that runs on IBM OS/2 or, beginning with OSA/SF Version 2, either of two Microsoft Windows platforms: Windows 95 or Windows NT.

For more information on using the REXX interface, refer to Appendix D, “Using the OSA/SF REXX interface” on page 201.

From a single OSA/SF GUI (OS/2 or Windows), you can establish connections to all server images (LPARs) that have OSA/SF running. You also do not need to have OSA/SF running on each server image.This, potentially, allows for centralized control of OSA features that are spanned across CPC boundaries, as shown in Figure 4-1 on page 48.

You can have GUI instances within each server image that has OSA/SF, to monitor your OSA features locally.

© Copyright IBM Corp. 2003. All rights reserved.

4.1 Requirements

OSA/SF is required to set up and customize OSA-2 features if the default IP Passthru mode is not used.

OSA/SF is not required for the OSA-Express feature that is configured for the QDIO mode[1], or the default IP Passthru non-QDIO mode.

OSA/SF is a required facility when the OSA-Express feature is being configured for shared non-QDIO mode[2] and where SNA definitions are involved.

OSA/SF is a must for the OSA-Express 155 ATM. You will need OSA/SF when the 155 ATM is being configured for Ethernet or Token Ring LAN Emulations (non-QDIO) where SNA/APPN/HPR definitions are involved; or when it is being configured for 155 ATM Ethernet emulation (QDIO) where OSA/SF is still needed for the definitions of the emulated/logical clients.

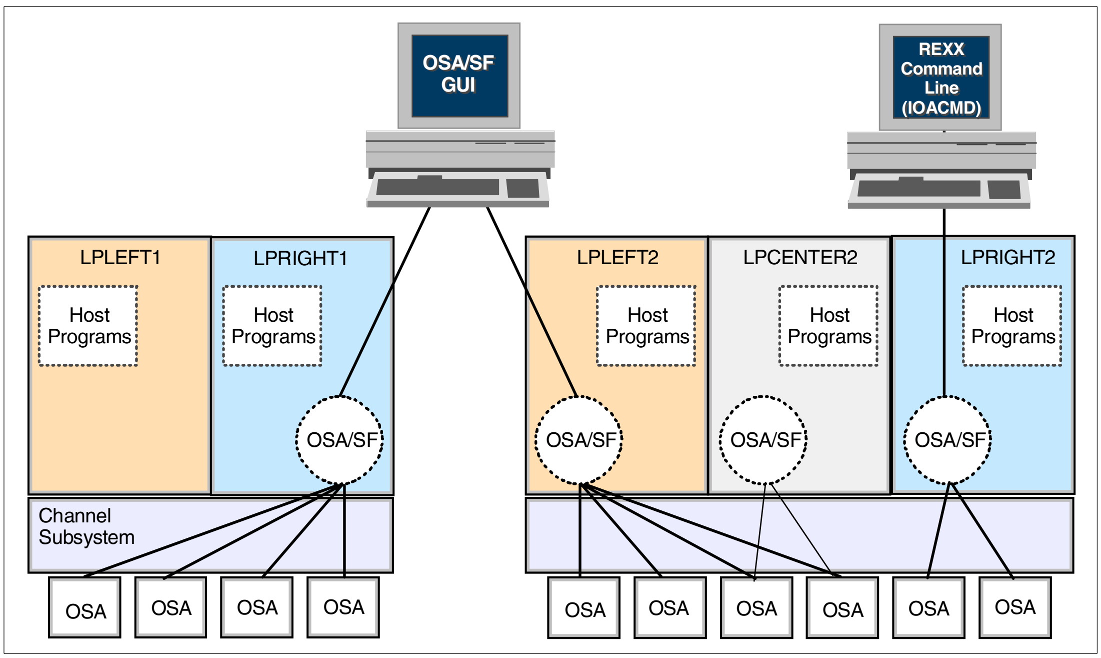

Figure 4-1 Use of OSA/SF and the OSA/SF GUI

OSA/SF V2R1 supports all OSA-2 and OSA-Express features and all OSA modes. Generally, one of these OSA/SFs will be managing the OSAs that are defined to the partitions in which these operating systems are running. However, through the OSA/SF **`Start Managing`** command, you can decide which OSA/SF session controls which OSA channel.

The OSA/SF GUI supports three different connection types:

- TCP/IP
- EHLLAPI
- APPC

In our environment, we used TCP/IP connections.

1 Except for the 155 ATM.
2 Port-sharing on GbE or FENET for QDIO mode does not require OSA/SF.

For information on the other connection types, or on setting up controlling access to OSA/SF (RACF), refer to *Open Systems Adapter-Express Customer's Guide and Reference, SA22-7403,* or *zSeries: Open Systems Adapter-Express Customer's Guide and Reference, SA22-7476.*

4.2 Setting up OSA/SF

Before using OSA/SF to manage your OSA-Express features, you must configure OSA/SF on your system. Additionally, APPC communication is required regardless of the connection type used. Also, the OSAD device has to be defined in HCD.

For more detailed information on installing OSA/SF, refer to the program directory.

4.2.1 Setting up APPC and VTAM

Following are the steps we used to set up our APPC environment:

1. Define the VTAM APPL statement for OSA/SF, edit VTAM Major Node member APPCOSA in the VTAMLST, and add the statements shown in Example 4-1.

Example 4-1 APPCOSA from SYS1.VTAMLST

```
IOASERV APPL   ACBNAME=IOASERV,
APPC=YES,AUTOSES=0,DDRAINL=NALLOW,
DMINWNL=5,DMINWNR=5, DRESPL=NALLOW,
DSESLIM=10,LMDENT=19,MODETAB=MTAPPC,
```

2. Define an APPC local LU for OSA/SF by editing member APPCPMxx in the SYS1.PARMLIB and adding the statements shown in Example 4-2.

Note: There is no dependency on the APPC scheduler for OSA/SF.

Example 4-2 APPCPM00 from SYS1.PARMLIB

```
LUADD
   ACBNAME(IOASERV)
   NOSCHED
   TPDATA(SYS1.APPCTP)
```

Note: SYS1.APPCTP is a VSAM data set. It must be allocated using job ATBTPVSM, included in SYS1.SAMPLIB.

3. Add procedure APPC to the SYS1.PROCLIB, as shown in Example 4-3.

Example 4-3 APPC from SYS1.PROCLIB

```
//APPC PROC APPC=00
//APPC EXEC PGM=ATBINITM,PARM='APPC=&APPC',REGION=0K
```

For automatic startup of the APPC environment, add the following parameters to your COMMNDxx member of SYS1.PARMLIB, as shown in Example 4-4 on page 49.

Example 4-4 COMMND00 from SYS1.PARMLIB

```
COM='S APPC,SUB=MSTR'
```

4.2.2 Setting up OSA/SF

To set up OSA/SF, use the following steps:

1. Create an STC (Started Task):
   a. Copy the sample procedure from IOA.SIOASAMP(IOAOSASF) to the SYS1.PROCLIB.
   b. Edit the procedure and change the name to OSASF.
   c. Verify that the data set names in the STEPLIB and IOALIB DDname statements match your environment.
2. Create a startup profile for OSA/SF:
   a. Allocate a sequential data set.
   b. Copy into this data set the sample provided in IOA.SIOASAMP(IOASPROF).
   c. Edit the profile and change the SYSNAME and CECNAME to suit your installation (verify UNIT and VOLSER).
3. Set up the OSA configuration and master profile:
   a. Allocate data set IOA.&CECNAME.OSAS.CONFIG and copy into the sample provided in IOA.SIOASAMP(IOACFG).
   b. Allocate data set IOA.&CECNAME.MASTER.INDEX and copy into the sample provided in IOA.SIOASAMP(IOAINX).
4. Set up REXX executable for use under TSO. Copy member IOACMD from IOA.SIOASAMP to your local clists or executable data sets.
5. Issue the following commands from the SDSF log:

```
/s appc,sub=mstr
/s osasf
```

You are now able to use OSA/SF from TSO commands, as in Example 4-5.

Example 4-5 Sample commands from TSO

```
IOACMD QUERY HOST "USER.OUTPUT"
IOACMD GET_OAT F8 "USER.OUTPUT"
```

4.2.3 Communicating with OSA/SF using TCP/IP

As mentioned previously, we used TCP/IP as the connection type for communicating with OSA/SF from a workstation. To set up TCP/IP, we used the following steps:

1. Update the TCPIP.TCPPARM(PROFILE) with the following data.
   a. Add the server to the AUTOLOG statement, as shown in Example 4-6.

Example 4-6 IOASRV from TCP/IP Profile

```
AUTOLOG
  .
  IOASRV
  .
```

   b. Add the port statement, as shown in Example 4-7 on page 51.

Example 4-7 Port number from TCP/IP Profile

```
PORT
.
2000 TCP IOASRV              ; OSA/SF Server
.
```

2. Create a procedure in SYS1.PROCLIB(IOASRV), as shown in Example 4-8.

Example 4-8 IOASRV from SYS1.PROCLIB

```
//*
//*   Sample TCP/IP Server Proc
//*
//SERVER   PROC
//SERVER   EXEC PGM=IOAXTSRV,PARM=2000,REGION=4M,TIME=1440
//*STEPLIB  DD  DSN=SYS1.OSASF.R2.SIOALMOD,DISP=SHR
//IOALIB   DD  DSN=SYS1.SIOALMOD,DISP=SHR
//PROFILE  DD  DISP=SHR,DSN=TCPIPMVS.&SYSNAME..TCPPARMS(TCPPROF)
```

3. Restart TCP/IP or use the OBEYFILE TCP/IP sub-command to make these modifications active.

4.3 Installing OSA/SF GUI on a workstation

We did the following to install and set up the OSA/SF GUI on a Windows workstation after we installed and set up OSA/SF on the Host system.

4.3.1 Check the hardware configuration

To use the OSA/SF GUI interface, a workstation with the following hardware features is required:

- A Pentium 200 MHz (or equivalent), 32 MB of RAM, and an SVGA display with a resolution of 1024 by 768 with 16 colors
- A communications adapter that supports an EHLLAPI for 3270 communications protocol, or a TCP/IP communications protocol, or an APPC protocol for a node that supports LU 6.2

In our configuration, we used the following:

- Workstation IBM PC 300PL 256 MBRAM
- Connection by IBM Token-Ring PCI Family Adapter

4.3.2 Check the software configuration

For the software prerequisites, refer to “OSA/SF GUI workstation” on page 28.

In our configuration, we used the following:

- Windows NT 4.0
- Personal Communications/3270 Version 4 Release 3
- TCP/IP

4.3.3 Download the code from the Host system using TCP/IP FTP

To download the OSA/SF GUI to the workstation, do the following:

1. Open a DOS session on your workstation.
2. CD to the directory where you want the exec stored.
3. Issue the FTP command, using the IP address of your Host (for example, `FTP 9.12.14.32`).
4. Provide your userid and password.
5. Issue: `bin` (to set the FTP transfer to Binary)
6. Issue: `CD 'IOA.SIOAWIN'` (the MVS dataset name in single quotes)
7. Issue: `GET IOAWINST IOAWINST.EXE`

We installed the OSA/SF GUI on the workstation by executing IOAWINST under Windows.

4.3.4 Setting up the OSA/SF GUI

To set up the OSA/SF GUI on the workstation, do the following:

1. Start the IBM OSA Support Facility program from Windows Programs to get the primary customization panel, as shown in Figure 4-2.

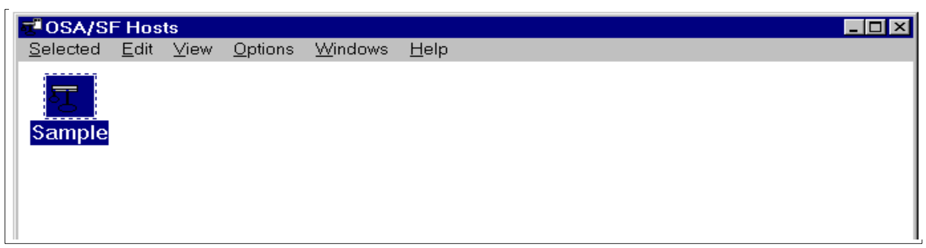

Figure 4-2 Initial OSA/SF display

2. Create an icon for the environment to be monitored, as shown in Figure 4-3.

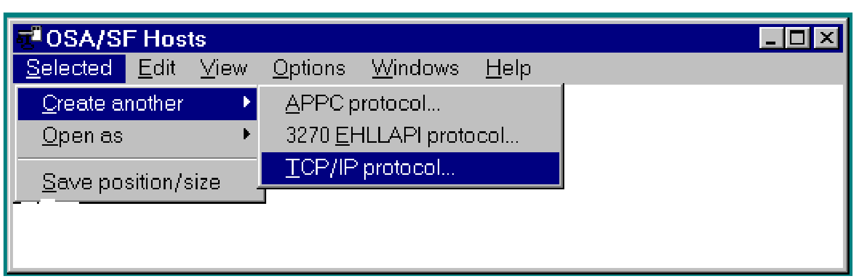

Figure 4-3 TCP/IP connection

3. Define the TCP/IP connection to the Host server, as shown in Figure 4-4 on page 53.

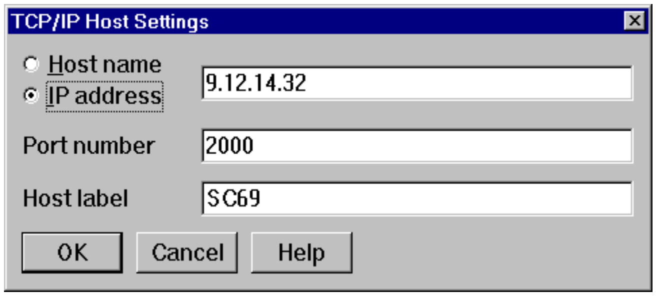

Figure 4-4 Connection to the Host server

4.3.5 Using the OSA/SF GUI

After creating the new Host connection, an icon appears that represents this Host (SC69), as shown in Figure 4-5.

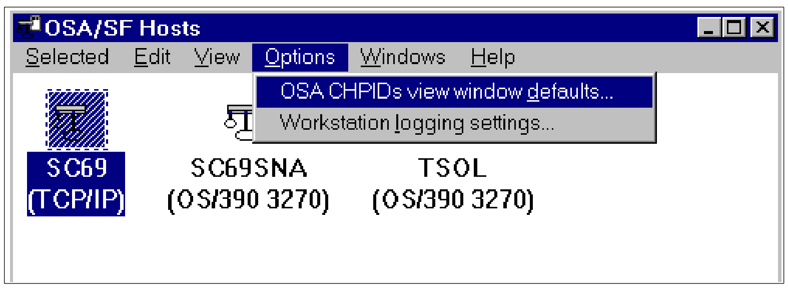

Figure 4-5 OSA/SF Hosts

With the Options tab, you can decide what your initial screen looks like. Figure 4-6 shows an example of the settings for the view screens.

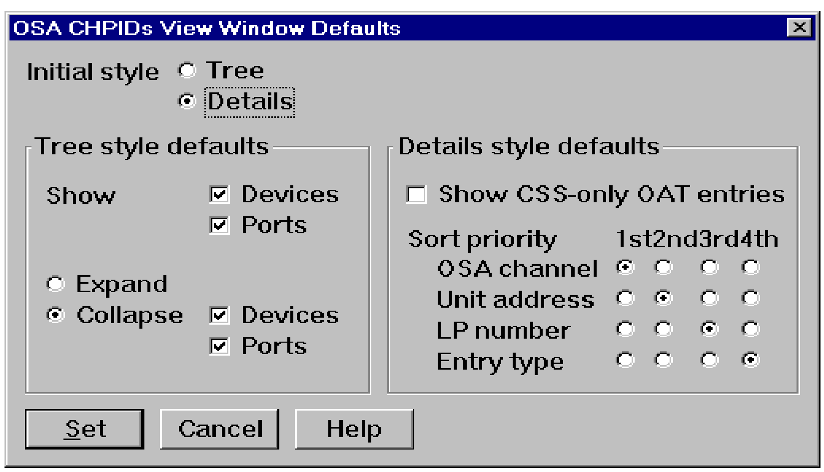

Figure 4-6 OSA/SF View Window Defaults

An example of the Tree View is shown in Figure 4-11 on page 56. If the initial screen is set to Details, you get a screen as shown in Figure 4-8 on page 55.

From the OSA/SF Hosts window, you can log in to your environment and view information such as the following:

1. From the Details View:
   - Defined CHPIDs
   - CHPID type
   - LPAR information
   - Operation mode
   - Operation status
   - HOME IP addresses
2. From the Tree View:
   - Defined CHPIDs
   - Port information
   - Physical data
   - Port statistics
   - Device settings
   - Device statistics

Double-click the icon to access the Host server and you will receive a Host Login panel, as shown in Figure 4-7.

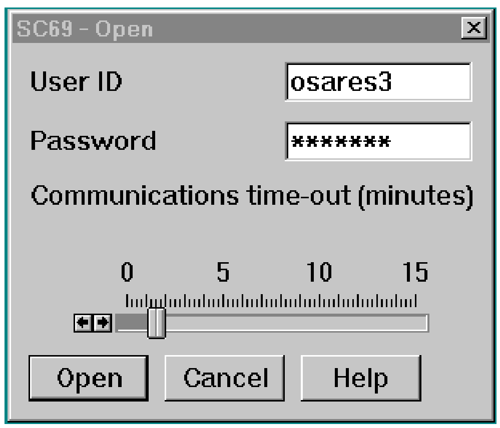

Figure 4-7 Host Login

Once the connection is established and the default view is set to Details, a list of accessible OSA CHPIDs installed is displayed, as shown in Figure 4-8.

Note: To get access to an OSA CHPID, the OSAD must be online.

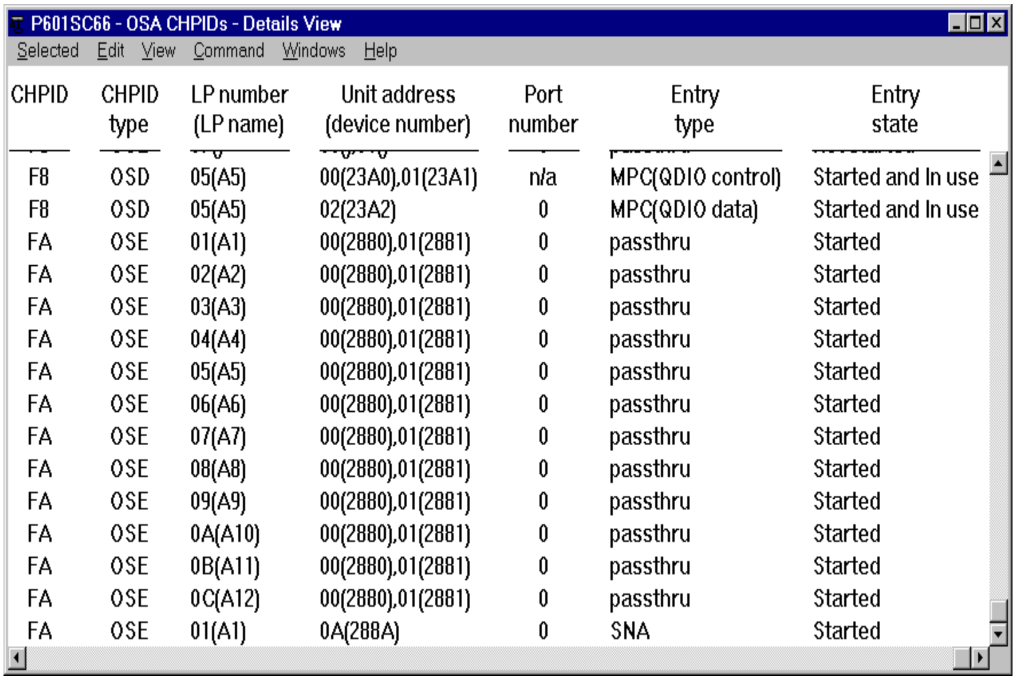

Figure 4-8 OSA CHPIDs - Details View

With a double-click on one row with an entry type used for IP traffic, you get a list of the IP Home addresses that are loaded into the feature. Figure 4-9 shows an example of an IP Home address list of an OSA-Express GbE feature running in QDIO mode.

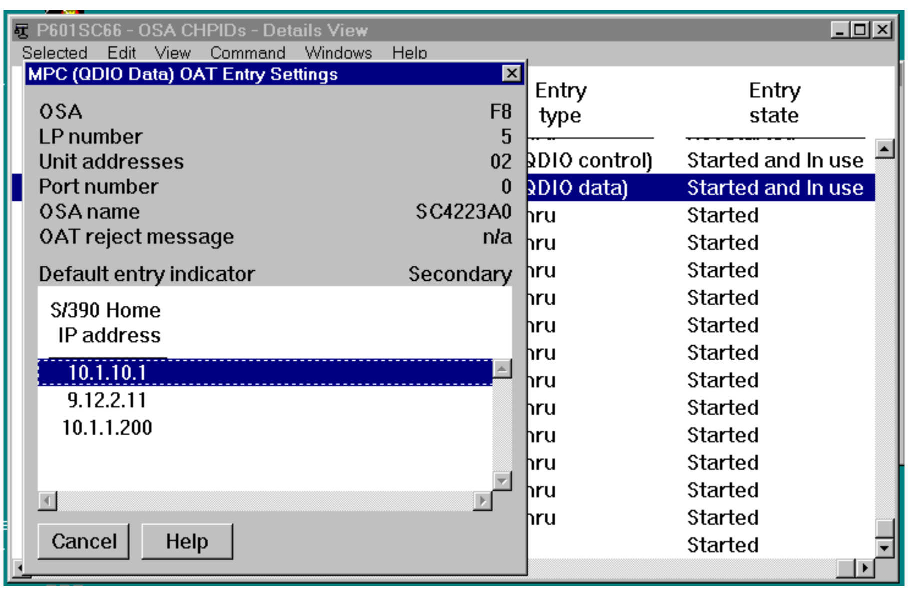

Figure 4-9 OAT entry settings

You can switch to the CHPID tree view by using the options **View -> Style -> Tree view,** as shown in Figure 4-10 on page 56.

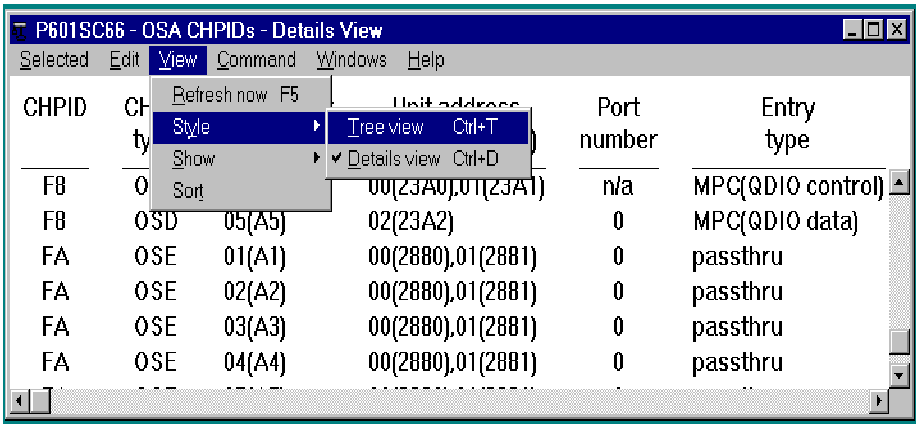

Figure 4-10 Switching between Details View and Tree View

Once the connection is established, a list of OSA CHPIDs installed on the S/390 is displayed, as shown in Figure 4-11.

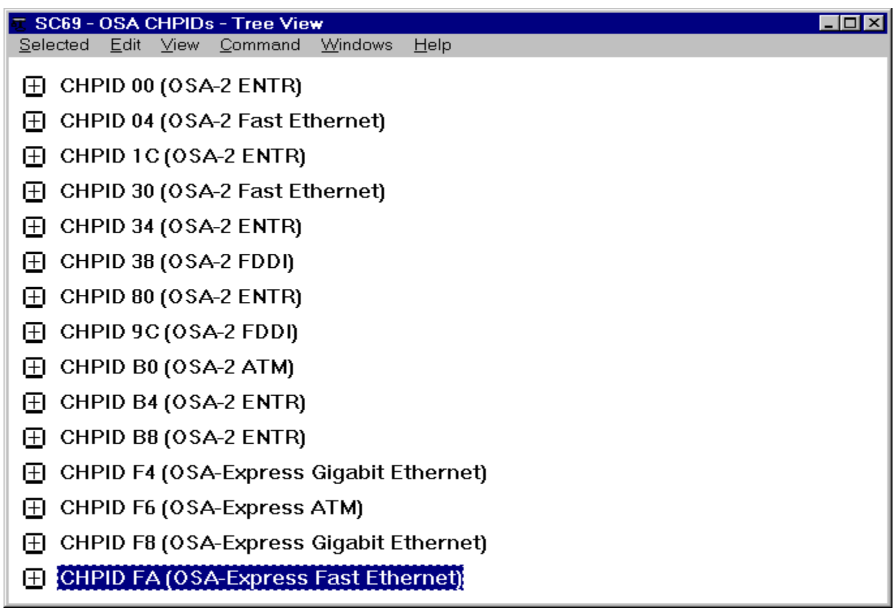

Figure 4-11 OSA CHPIDs - Tree View

For an OSA-Express CHPID view and port information, select the desired CHPID and click the plus sign (+) to expand for port and device information as shown in Figure 4-12 on page 57. If no “plus box” appears next to the CHPID, it means that the OSAD device is offline.

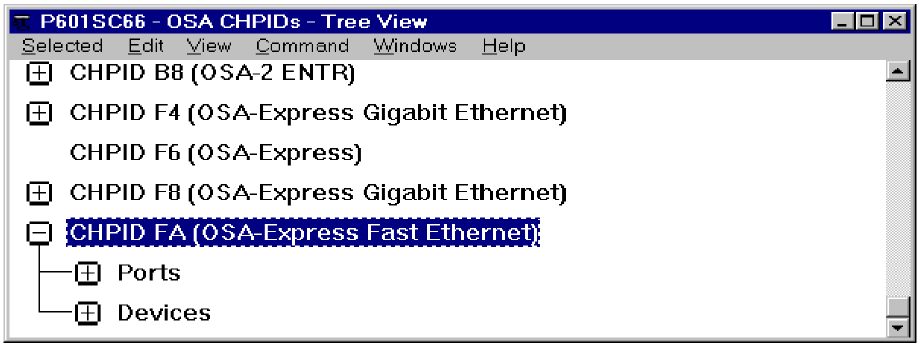

Figure 4-12 Expand and collapse boxes

Double-click **CHPID** to view the OSA-Express settings, as shown in Figure 4-13 and Figure 4-14.

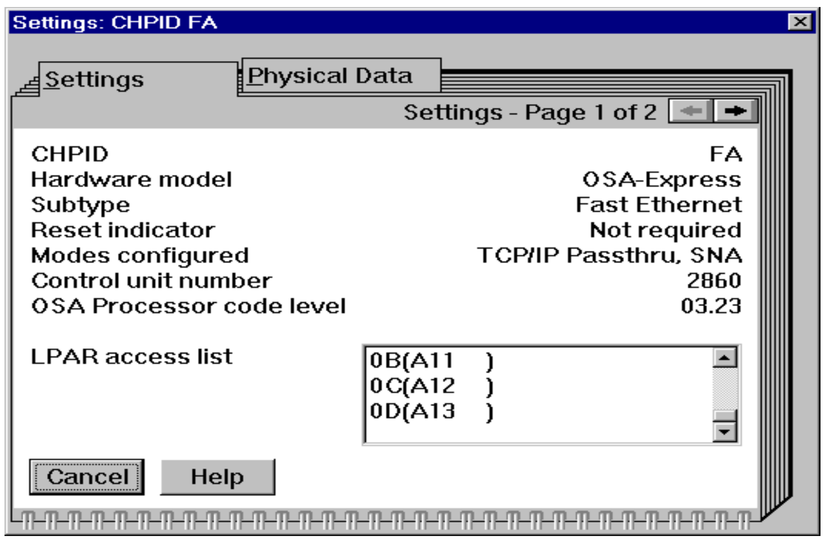

Figure 4-13 Settings Page 1

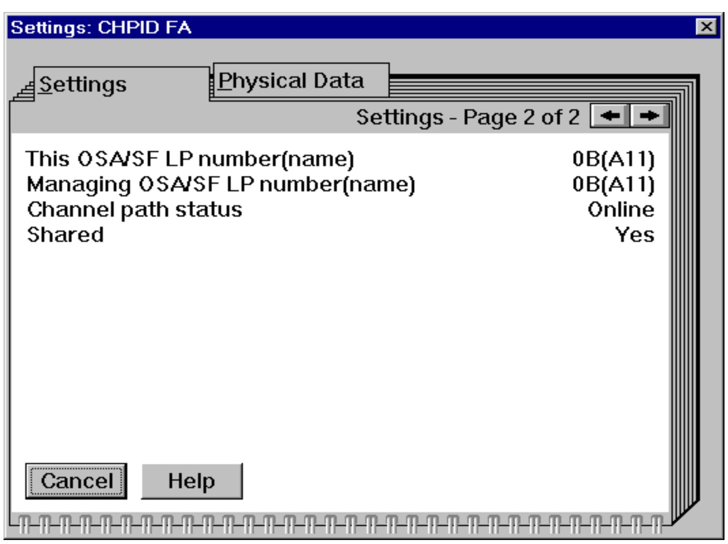

Figure 4-14 Settings Page 2

Click **Physical Data** to display the information shown in Figure 4-15.

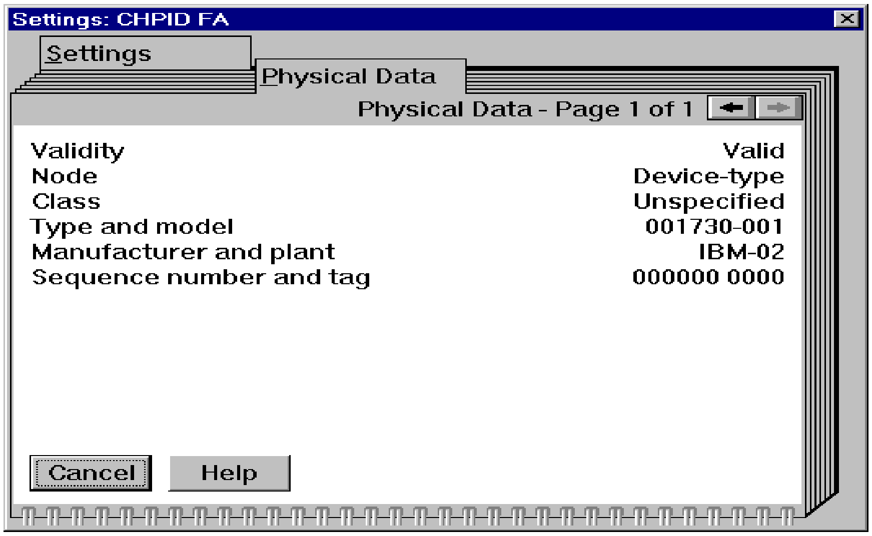

Figure 4-15 Physical Data

Double-click **Port 0** (see Figure 4-16) to view port settings and statistics. They are shown in Figure 4-17 on page 59 and Figure 4-18 on page 59.

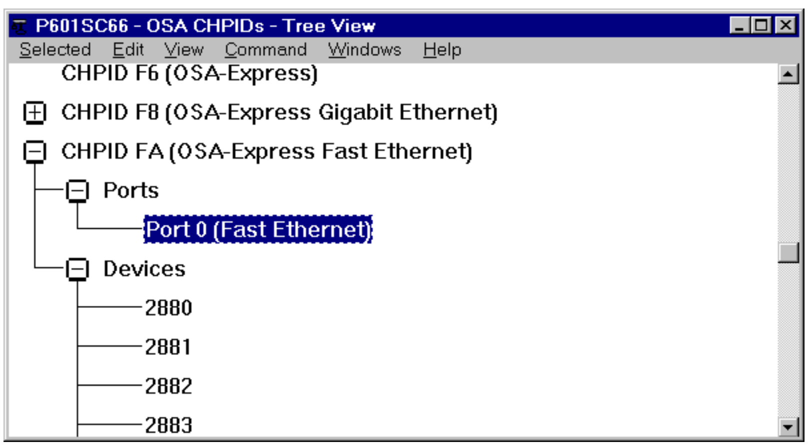

Figure 4-16 Port 0 (Fast Ethernet)

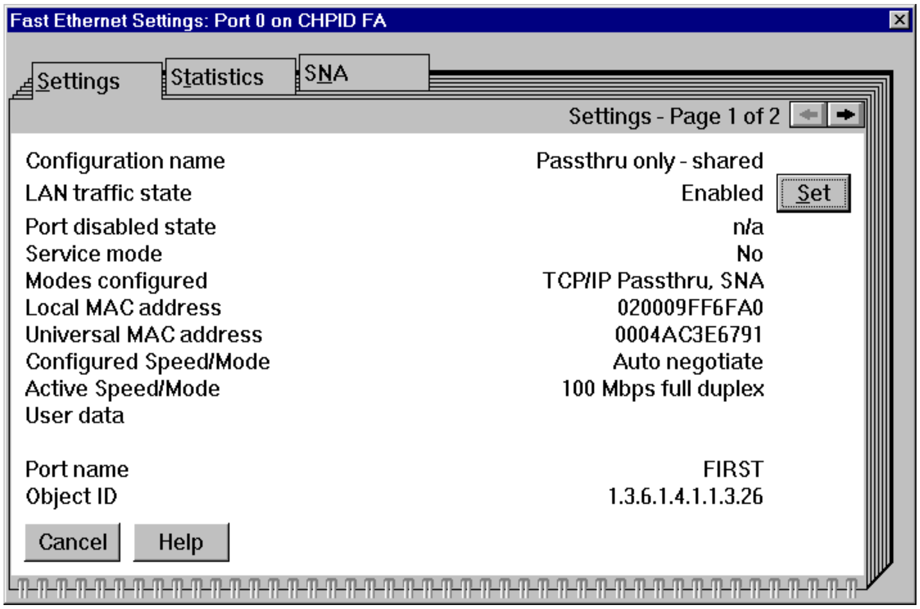

Figure 4-17 Settings Page 1

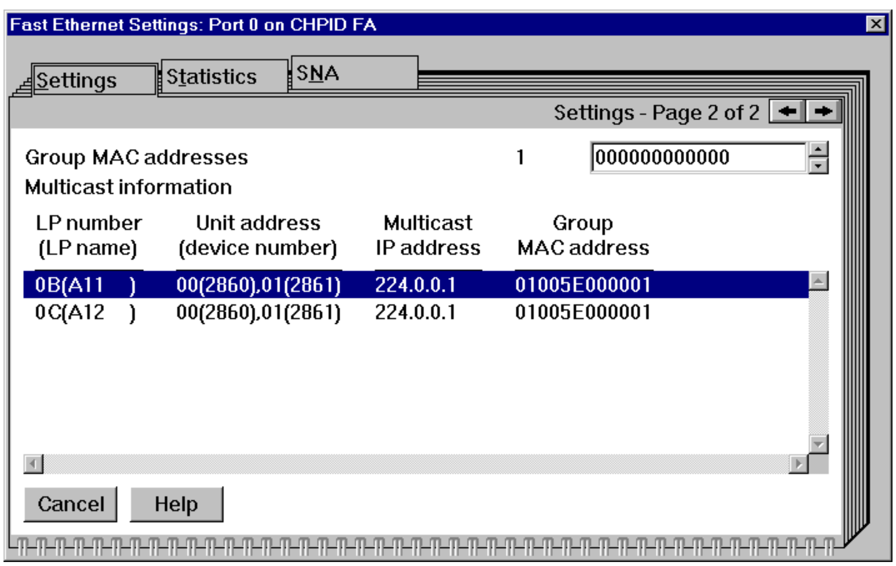

Figure 4-18 Settings Page 2

Click **Statistics** to view the information shown in Figure 4-19 on page 60.

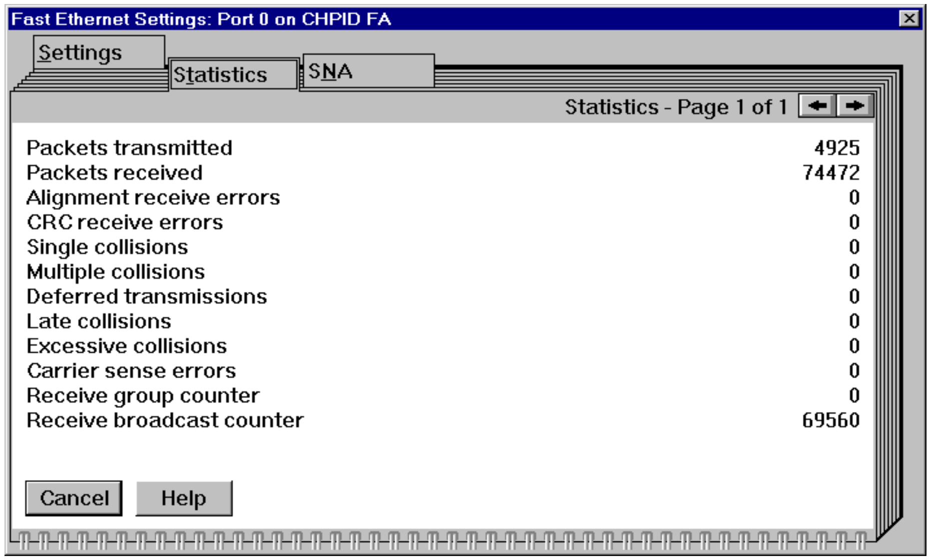

Figure 4-19 Statistics on Port 0

For SNA information as shown in Figure 4-20, select the **SNA** tab.

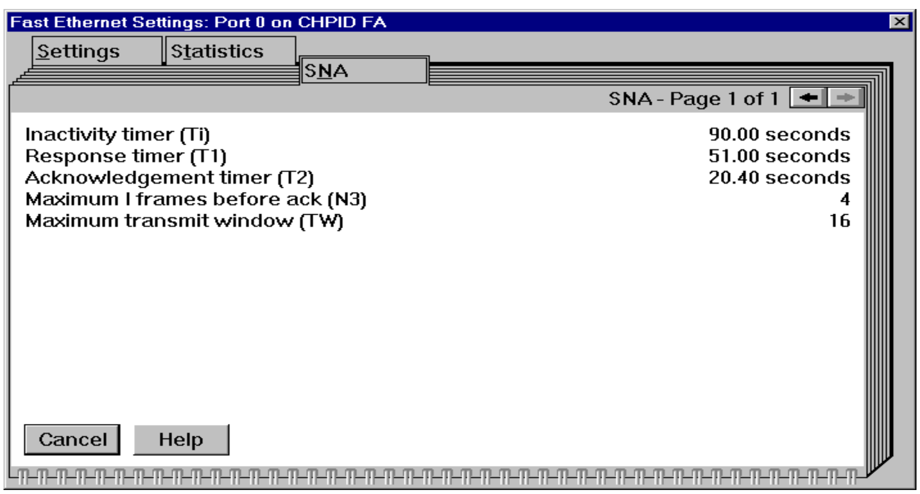

Figure 4-20 SNA Page 1

Go back to the OSA CHPIDs -Tree View panel, as shown in Figure 4-16 on page 58, and double-click any device (for example, 23A0) to see its settings and status information, as shown in Figure 4-21.

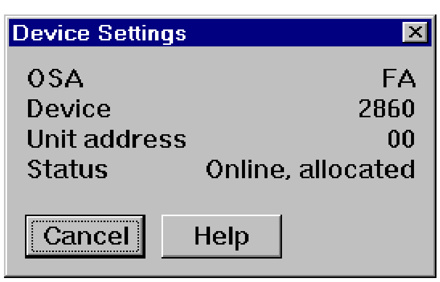

Figure 4-21 Device settings and status

5

QDIO mode

In this chapter, we cover the implementation steps needed to establish network connectivity with an OSA-Express feature using QDIO mode.

We guide you through the steps required to create:

- VTAM definitions to support this configuration
- TCP/IP definitions to support this configuration

OSA/SF is not required because all definitions are set dynamically; however, for monitoring and port controlling, we recommend the use of OSA/SF.

© Copyright IBM Corp. 2003. All rights reserved.

5.1 QDIO environment

Figure 5-1 is a logical representation of the QDIO environment discussed in subsequent sections of this chapter.

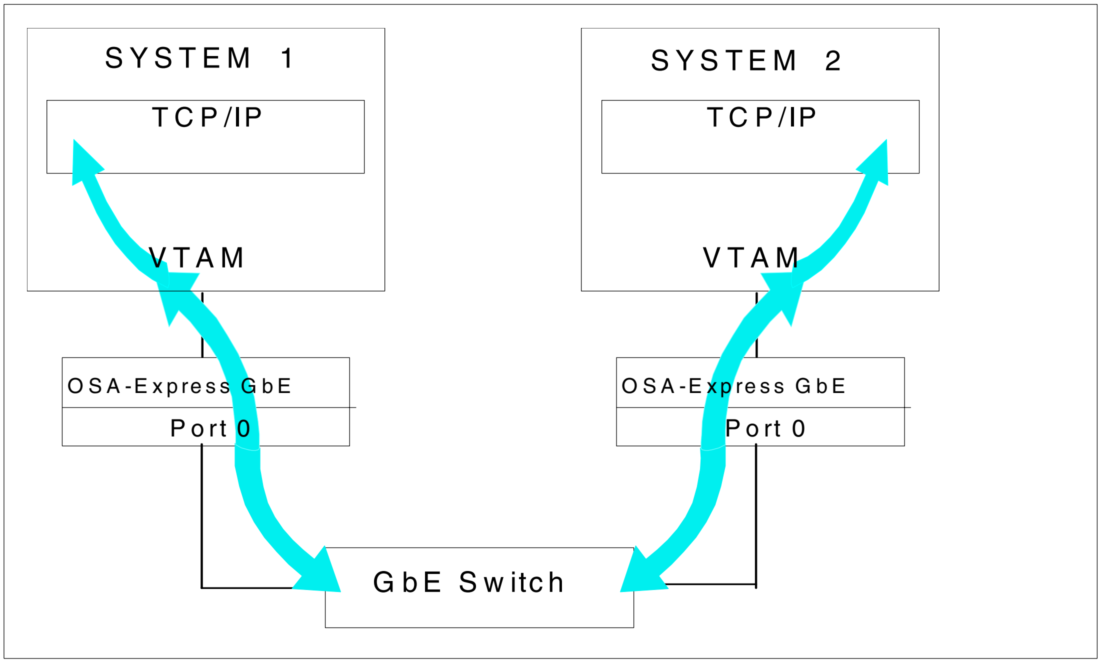

Figure 5-1 QDIO mode example

5.2 HCD definition requirements

The OSA CHPID, the control unit, and the OSA devices must be defined to HCD, and the HCD definitions activated. Refer to Chapter 3, “Hardware configuration definitions” on page 37 for the procedure to create the definitions. The necessary definitions for CHPID F8 on partition SC47 (A11) in the IOCDS used in this chapter are shown in Example 5-1. For future purposes we used 15 OSA devices, although there were only three needed in our configuration.

The OSAD device is only necessary to configure the OSA-Express 155 ATM feature; however, for monitoring and port controlling, we recommend the use of OSA/SF.

Example 5-1 IOCDS input for CHPID F8

```
.
CHPID PATH=(F8),SHARED,PARTITION=((A11,A12),(A11,A12)),TYPE=OSD
.
CNTLUNIT CUNUMBR=2820,PATH=(F8),UNIT=OSA
.
IODEVICE ADDRESS=(2820,15),CUNUMBR=(2820),UNIT=OSA,UNITADD=00
IODEVICE ADDRESS=282F,UNITADD=FE,CUNUMBR=(2820),UNIT=OSAD
```

5.3 Missing Interrupt Handler (MIH) for QDIO

The WRITE devices (as defined in the TRLE) should have an MIH value of at least 15 seconds (or 30 seconds, if running as a guest system on VM).

To determine the current MIH value for the device (2381, in our example), enter the console command:

```
D IOS,MIH,DEV=2381
```

To dynamically change the MIH value, enter the console command:

```
SETIOS MIH,DEV=2381,TIME=00:15.
```

To set these values at IPL time, update the IECIOSxx member in PARMLIB.

5.4 Customizing the z/OS server network subsystem

Note: The discussion here is based on the GbE feature; however, the information presented holds true for the FENET feature, the Token Ring feature, and the 155 ATM feature (in ENET LANE).

Figure 5-2 shows our network configuration: two z/OS systems, with a connection to the switch through OSA-Express Gigabit Ethernet CHPIDs. The switch has a Gigabit Ethernet module that has two ports. Each port is connected to an z/OS system.

The workstation has an IP address of 192.168.104.2, and the GbE switch has an IP address of 192.168.104.1. In this environment, the IP address of the switch is only necessary for configuration purposes.

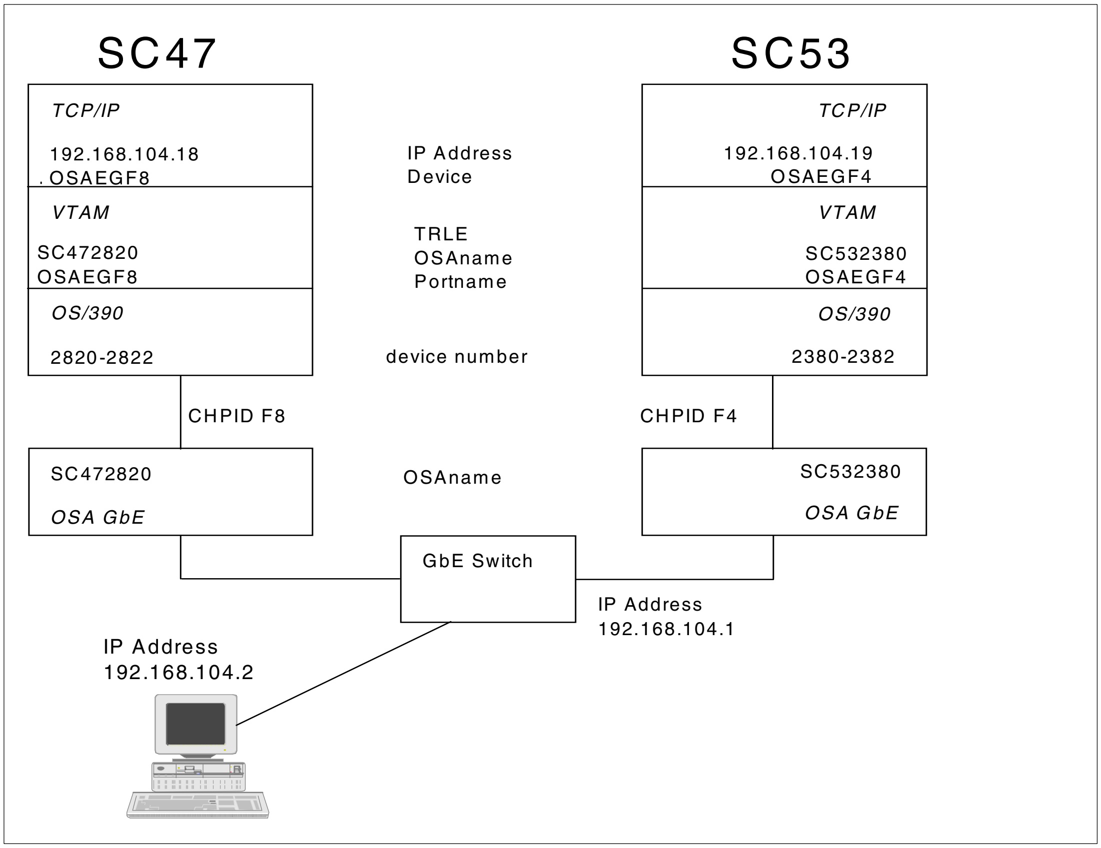

Figure 5-2 Network configuration

5.4.1 VTAM definitions

TCP/IP uses a VTAM interface to run the OSA-Express in QDIO mode. A Transport Resource List (TRL) major node must be defined and activated before TCP/IP starts its QDIO device.

TRLE considerations

The table included in 2.5, “Quick start tables” on page 33 contains various uses of VTAM TRLE definitions. This section lists implementation examples and the associated TRLE definition requirements.

No LPAR, only one TCP/IP stack, only one OSA-Express CHPID

- Only one TRLE (osaname) is needed, containing only one datapath address.
- TCP/IP must refer to the portname in the TRLE.

No LPAR, multiple TCP/IP stacks, only one OSA-Express CHPID

For example, three TCP/IP stacks:

- Only one TRLE (osaname) is needed, containing three datapath addresses.
- Each TCP/IP stack must refer to the portname in the TRLE.
- VTAM assigns one datapath address to each TCP/IP stack.

No LPAR, only one TCP/IP stack, multiple OSA-Express CHPIDs

- One TRLE is needed for each OSA-Express CHPID, containing only one datapath address.
- TCP/IP DEVICE name must match the portname in each of the TRLEs.

No LPAR, multiple TCP/IP stacks sharing multiple OSA-Express CHPIDs

For example, three TCP/IP stacks and two OSA-Express CHPIDs:

One TRLE for each OSA-Express CHPID is needed, each containing three datapath addresses.

- Each TCP/IP stack has two DEVICE and LINK statements. Each DEVICE statement must match the portname in the TRLE.
- VTAM assigns, from both TRLEs, one datapath address to each TCP/IP stack.

Multiple LPARs with only one TCP/IP stack per LPAR sharing only one OSA-Express CHPID

- Only one TRLE is needed in each LPAR, containing only one datapath address. The portname in the TRLE must be the same across the LPARs.
- Each TCP/IP DEVICE name must match the portname in the TRLE.

Multiple LPARs with multiple TCP/IP stacks per LPAR sharing only one OSA-Express CHPID

For example, two TCP/IP stacks per LPAR:

- Only one TRLE is needed in each LPAR, containing two datapath addresses. The portname in the TRLE must be the same across the LPARs.
- Each TCP/IP DEVICE name must match the portname in the associated TRLE.
- VTAM assigns one datapath address to each TCP/IP stack.

Multiple LPARs with only one TCP/IP stack per LPAR sharing multiple OSA-Express CHPIDs

For example, three OSA-Express CHPIDs:

- Three TRLEs are needed in each LPAR, containing only one datapath address. For each single OSA-Express CHPID, the portname in the TRLE must be the same across LPARs, but it must be unique across the OSA-Express features.
- Each TCP/IP stack has three DEVICE and LINK statements. Each DEVICE statement must match the portname of the associated TRLE.

Multiple LPARs with multiple TCP/IP stacks per LPAR sharing multiple OSA-Express CHPIDs

For example, three OSA-Express CHPIDs and two TCP/IP stacks per LPAR:

- Three TRLEs are needed in each LPAR, containing two datapath addresses. For each single OSA-Express CHPID, the portname must be the same across LPARs, but must be unique for each OSA-Express CHPID.
- Each TCP/IP stack has three DEVICE and LINK statements. Each DEVICE statement must match the portname of the associated TRLE.
- VTAM in each LPAR assigns one datapath address to each TCP/IP stack in the same LPAR.

Example 5-2 shows the VTAM TRL major node definition for SC47.

Example 5-2 TRL major node for SC47

```
SC47OSF8 VBUILD TYPE=TRL
SC472820 TRLE  LNCTL=MPC,
               READ=2820,
               WRITE=2821,
               DATAPATH=(2822),
               PORTNAME=OSAEGF8,
               MPCLEVEL=QDIO
```

Example 5-3 shows the VTAM TRL major node definition for SC53.

Example 5-3 TRL major node forSC53

```
SC53OSF4 VBUILD TYPE=TRL
SC532380 TRLE  LNCTL=MPC,
               READ=2380,
               WRITE=2381,
               DATAPATH=(2382),
               PORTNAME=OSAEGF4,
               MPCLEVEL=QDIO
```

Table 5-1 briefly describes the definitions used in TRL major node for SC47; the same applies for SC53.

Table 5-1 TRL major node definition for SC47

Required parameters	Explanation	Remarks
TYPE=TRL	TRL major node	MPC TRL major node known to VTAM.
SC472820	TRLE minor node	The name of the TRLE in VTAM for SC47. This name is downloaded to the OSA-Express feature and used as OSANAME.
READ=2820	Read device	The Read device number must be the even number of the device pair.The Read Write pair of the OSA-Express feature is only used to exchange control information.
WRITE=2821	Write device	The Write device number must be the odd number of the device pair. The Read Write pair of the OSA-Express feature is only used to exchange control information.
DATAPATH=2822	Data device	The device address of DATAPATH of each OSA-Express feature. For QDIO, this device is used for the data transfer in both directions. More than one address can be defined, depending on throughput. In order to add a second TCP/IP stack, an additional DATAPATH device must be added to the TRLE statement and HCD. The DATAPATH address need not be the immediate next after the write address; it can be any address in the range of the defined devices of the OSA-Express in the HCD.

PORTNAME = OSAEGF8	PORTNAME associated with the devices	PORTNAME must match TCP/IP's device name in the TCP/IP profile for this connection. The association between TCP/IP and VTAM is done through the PORTNAME.
MPCLEVEL=QDIO	MPC compatibility level	Indicates the QDIO interface is used for the OSA-Express feature.

5.4.2 TCP/IP definitions

TCP/IP only requires DEVICE and LINK definitions corresponding to the portname in an VTAM TRLE. Table 5-4 shows the TCP/IP profile definitions of the OSA-Express Gigabit Ethernet feature (OSAEGF8) for SC47.

Note: In the TCP/IP profile, link type IPAQENET is used for OSA-Express Fast Ethernet and Gigabit Ethernet features. For the OSA-Express Token Ring feature, the link type must be defined as IPAQTR in the TCP/IP profile.

Example 5-4 TCP/IP profile for SC47

```
DEVICE OSAEGF8     MPCIPA    NONROUTER
LINK   OSAEGF8LNK  IPAQENET OSAEGF8

HOME
192.168.104.18  OSAEGF8LNK

BEGINROUTES
 192.168.104.0/24 =    OSAEGF8LNK   MTU 1500
ENDROUTES

START OSAEGF8
```

Example 5-5 shows the TCP/IP profile definitions for the OSA-Express feature (OSAEGF4) for SC53.

Example 5-5 TCP/IP profile for SC53

```
DEVICE OSAEGF4     MPCIPA    NONROUTER
LINK   OSAEGF4LNK  IPAQENET  OSAEGF4

HOME
192.168.104.19  OSAEGF4LNK

BEGINROUTES
 192.168.104.0/24 =    OSAEGF4LNK   MTU 1500
ENDROUTES

START OSAEGF4
```

The following definitions are used:

OSAEGF8 The device name; it must match the portname in the TRLE on SC47.

OSAEGF4 The device name; it must match the portname in the TRLE on SC53.

OSAEGF8LNK The link name in TCP/IP in system SC47.

OSAEGF4LNK The link name in TCP/IP in system SC53.

192.168.104.18 The IP address that is downloaded to the OSA-Express feature for SC47.

192.168.104.19 The IP address that is downloaded to the OSA-Express feature for SC53.

BEGINROUTES Statement:

If you are migrating your TCP/IP profile from an earlier release, it may use the GATEWAY statement to define static routes instead of the BEGINROUTES - ENDROUTES statements. Gateway will be recognized and used, but you should consider replacing it with BEGINROUTES.

Because it is compatible with UNIX standards, easier to code than GATEWAY, accepts both IPv4 and IPv6 addresses, and has enhanced functionality, BEGINROUTES is the recommended method for defining static routes. Future static route enhancements will only be available with the BEGINROUTES statement.

5.5 Activation

Under the new QDIO architecture interface on OSA-Express, the CHPID is dynamically reset as long as none of the devices are being used in all of the shared LPARs. This rule does not apply for the OSAD device.

Normally the CHPID should be online; if the CHPID is offline, configure it online using the following command:

```
CF CHP(F8),ONLINE
```

Once all the definitions have been added to OSA/SF, VTAM and TCP/IP, you can activate the configuration.

Activation may require several things, such as:

- Verifying the devices are online
- Activating an OSA/SF configuration
- Activating VTAM resources
- Activating TCP/IP

5.5.1 Verify devices are online

The z/OS console display command can verify that the required devices are online; see Figure 5-3 on page 66.

```
D U,,,2820,3
IEE457I 21.29.17 UNIT STATUS 836
UNIT TYPE STATUS        VOLSER     VOLSTATE
2820 OSA  O
2821 OSA  O
2822 OSA  O
```

Figure 5-3 z/OS D U command

If the devices are not online, use the z/OS console vary command:

```
V (2820-2822),ONLINE
```

5.5.2 OSA/SF activation

Since QDIO does not require a configuration to be loaded onto the OSA-Express feature, this task is not needed to manage configurations. However, it is still a useful diagnostic tool.

5.5.3 VTAM activation

Next, activate the TRL using the following VTAM command:

```
V NET,ACT,ID=SC47OSF8
```

TCP/IP requires an active TRL prior to starting its device.

After activation of the TRL, the status of the TRLE is NEVAC or INACT until TCP/IP refers to it. When the TCP/IP device is started, configuration information (for example, HOME IP address) is downloaded to the OSA-Express Gigabit Ethernet feature.

5.5.4 TCP/IP activation

There are two ways to activate the TCP/IP devices: either restart the TCP/IP stack, or use the TCP/IP Obeyfile command. We chose to restart the stack to implement the changes.

5.6 Relevant status displays

Figure 5-4 and Figure 5-5 on page 70 show the status of the OSA-Express Gigabit Ethernet devices on SC47 and SC53, displayed with the TSO command: **NETSTAT DEV**

```
DevName: OSAEGF8            DevType: MPCIPA    DevNum: 0000
  DevStatus: Ready         CfgRouter: Non  ActRouter: Non
  LnkName: OSAEGF8LNK        LnkType: IPAQENET    LnkStatus: Ready
    NetNum: 0   QueSize: 0   Speed: 0000001000
BytesIn: 0000007170      BytesOut: 0000006706
  ArpOffload: Yes  ArpOffloadInfo: Yes
BSD Routing Parameters:
  MTU Size: 00000            Metric: 00
  DestAddr: 0.0.0.0          SubnetMask: 255.255.255.0
Multicast Specific:
  Multicast Capability: Yes
  Group            RefCnt
  -----            ------
  224.0.0.1        0000000001
```

Figure 5-4 NETSTAT DEVLINKS for SC47

Important: Notice that you cannot determine the devices used by TCP/IP from the above field "DevNum: 0000". You must display the VTAM TRLE to see the devices.

```
DevName: OSAEGF4            DevType: MPCIPA    DevNum: 0000
  DevStatus: Ready         CfgRouter: Non  ActRouter: Non
  LnkName: OSAEGF4LNK        LnkType: IPAQENET    LnkStatus: Ready
    NetNum: 0   QueSize: 0   Speed: 0000001000
BytesIn: 0000000000      BytesOut: 0000000000
  ArpOffload: Yes  ArpOffloadInfo: Yes
BSD Routing Parameters:
  MTU Size: 00000            Metric: 00
  DestAddr: 0.0.0.0          SubnetMask: 255.255.255.0
Multicast Specific:
  Multicast Capability: Yes
  Group            RefCnt
  -----            ------
  224.0.0.1        0000000001
```

Figure 5-5 NETSTAT DEVLINKS for SC53

Notice the new addition to the NETSTAT output of *ArpOffload* in Figure 5-4 and Figure 5-5. This indicates that the OSA-Express feature supports ARP offload (ARP caching on the feature) and that the ARP information can be retrieved via NETSTAT ARP or through SNMP.

Figure 5-6 on page 71 shows the TRLE of the active OSA-Express Gigabit Ethernet on system SC47.

```
D NET,TRL,TRLE=SC472820
IST097I DISPLAY ACCEPTED
IST075I NAME = SC472820, TYPE = TRLE 706
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST087I TYPE = LEASED            , CONTROL = MPC , HPDT = YES
IST1715I MPCLEVEL = QDIO       MPCUSAGE = SHARE
IST1716I PORTNAME = OSAEGF8    LINKNUM =   0   OSA CODE LEVEL = 0292
IST1577I HEADER SIZE = 4096 DATA SIZE = 64 STORAGE = ***NA***
IST1221I WRITE DEV = 2821 STATUS = ACTIVE     STATE = ONLINE
IST1577I HEADER SIZE = 4092 DATA SIZE = 64 STORAGE = DATASPACE
IST1221I READ  DEV = 2820 STATUS = ACTIVE     STATE = ONLINE
IST1221I DATA  DEV = 2822 STATUS = ACTIVE     STATE = N/A
IST1724I I/O TRACE = OFF  TRACE LENGTH = *NA*
IST1717I ULPID = TCPIPA
IST314I END
```

Figure 5-6   D NET,TRL,TRLE=SC472820

Figure 5-7 shows the TRLE of the active OSA-Express Gigabit Ethernet on system SC53.

```
 D NET,TRL,TRLE=SC532380
IST097I DISPLAY ACCEPTED
IST075I NAME = SC532380, TYPE = TRLE 913
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST087I TYPE = LEASED            , CONTROL = MPC , HPDT = YES
IST1715I MPCLEVEL = QDIO       MPCUSAGE = SHARE
IST1716I PORTNAME = OSAEGF4    LINKNUM =   0   OSA CODE LEVEL = 0292
IST1577I HEADER SIZE = 4096 DATA SIZE = 64 STORAGE = ***NA***
IST1221I WRITE DEV = 2381 STATUS = ACTIVE     STATE = ONLINE
IST1577I HEADER SIZE = 4092 DATA SIZE = 64 STORAGE = DATASPACE
IST1221I READ  DEV = 2380 STATUS = ACTIVE     STATE = ONLINE
IST1221I DATA  DEV = 2382 STATUS = ACTIVE     STATE = N/A
IST1724I I/O TRACE = OFF  TRACE LENGTH = *NA*
IST1717I ULPID = TCPIPA
IST314I END
```

Figure 5-7   D NET,TRL,TRLE=SC532380

Note that message IST1717I shows the jobname that uses this TRLE (in our case, TCPIPA).

Consideration: If your static TRLE definition is incorrect, you should be aware that an active TRLE entry cannot be deleted. What you can do is vary activate the TRL node with a blank TRLE to cause the deletion of previous entries. Then code the correct TRL with correct TRLE entries and definition and vary activate this corrected TRL/TRLE node.

See “VTAM commands” on page 177 for the vary commands

Figure 5-8 shows a detailed view of datapath device 2382 with IP addresses that have been downloaded from the TCP/IP stack’s HOME statement.

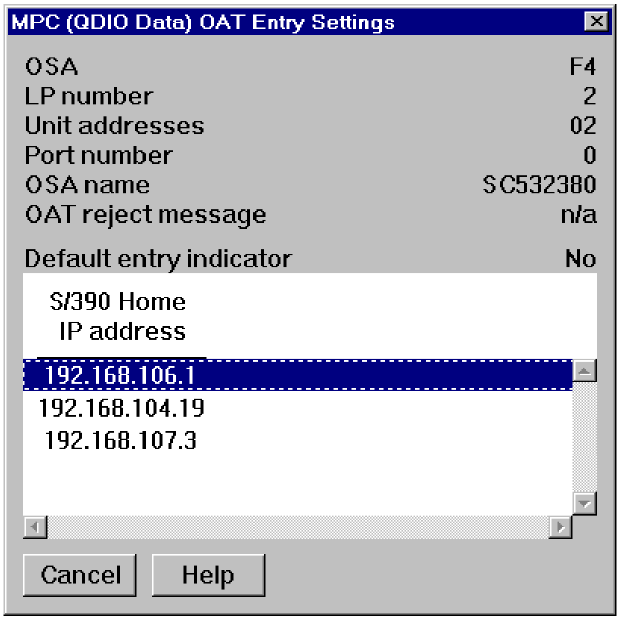

Figure 5-8 Detailed view of device 2382

This view is available through the OSA/SF, OSA CHPIDs - Details View. Simply double-click the CHPID, partition, and device entry in which you are interested.

See 4.3.5, "Using the OSA/SF GUI" on page 53 for more information on displaying OSA information using the OSA/SF GUI.

5.7 SNA support for QDIO mode

In cases where you need to transport SNA traffic over OSA-Express, and you would prefer to leverage the benefits of QDIO, IBM provides two technologies to integrate SNA based traffic via TCP/IP:

1. Enterprise Extender can be used to connect SNA (LU6.2,0,1,2,3) endpoint traffic over TCP/IP directly into the z/Series. See Chapter 10, "Enterprise Extender (EE)" on page 139 for more information.
2. The Telnet 3270 (TN3270) supports TCP/IP host access to SNA(LU2) applications.

6

TCP/IP Passthru (non-QDIO mode)

In this chapter, we describe the process required to use the Fast Ethernet OSA-Express feature in TCP/IP Passthru non-QDIO mode with only one LPAR (non-shared mode).

OSA2 note: The procedures in this chapter deal with configuring an OSA-Express feature. However, these same procedures can also be used to configure the OSA-2 features (with the exception of the HCD definition, CHPID type OSE).

© Copyright IBM Corp. 2003. All rights reserved.

6.1 Default mode

Figure 6-1 shows a functional view of the connectivity we discuss in this chapter.

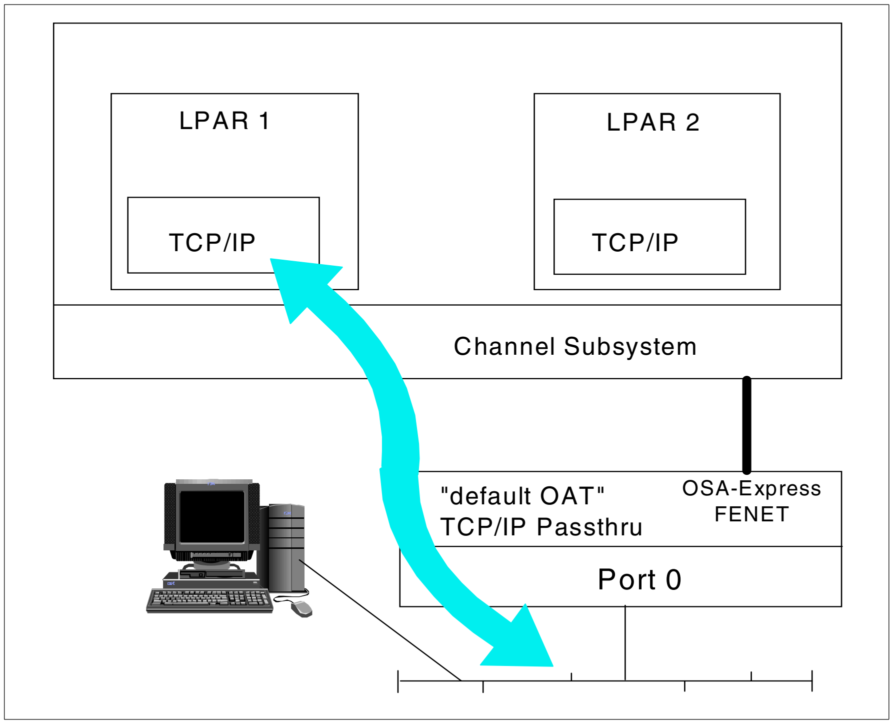

Figure 6-1 TCP/IP in Passthru Mode

For this example we are using the Fast Ethernet OSA-Express, CHPID type OSE, in *default* mode. When an OSA feature is manufactured, a basic configuration is installed that will permit some functionality without the need to build and load an OSA Address Table (OAT).

The default mode permits "passthru" functionality only. One use for this mode is in a new installation, where you can establish that HCD definitions and network connections are correct without the added complication or concern of whether there is a configuration error in the OAT.

The zSeries server sees the Fast Ethernet OSA-Express port as a LAN Channel Station (LCS) device. An LCS device will handle data traffic in either direction for any TCP/IP partition that has an OSA-Express feature defined.

6.2 HCD definition requirements

The OSA CHPID, the control unit, and the OSA devices must be defined to HCD, and the HCD definitions activated. Refer to Chapter 3, "Hardware configuration definitions" on page 37 for the procedure to create the definitions. The necessary definitions for the IOCDS used in this chapter are shown in Example 6-1 on page 75. For future purposes we used 15 OSA-devices, but there are only two needed in our configuration.

The CHPID in our example is dedicated because it is only needed in partition SC69 (A11). However, if you plan to use OSA/SF, we highly recommend that your feature be shared among all partitions where OSA/SF is running. This is why we included the definition for the OSAD (FE) device.

Example 6-1 IOCDS input for CHPID FA

```
.
CHPID PATH=(FA),PARTITION=A11,TYPE=OSE
.
CNTLUNIT CUNUMBR=2860,PATH=(FA),UNIT=OSA
.
IODEVICE ADDRESS=(2860,15),CUNUMBR=(2860),UNIT=OSA,UNITADD=00
IODEVICE ADDRESS=(286F,1),CUNUMBR=(2860),UNITADD=FE,UNIT=OSAD
```

If you are using QDIO (CHPID type=OSD), follow the procedures documented in Chapter 5, “QDIO mode” on page 61.

6.3 Displaying the default OAT

Note that we do not need to use OSA/SF for the activities in this chapter. However, we thought it would be instructive to view the default OAT.

We used the OSA/SF GUI program to show the contents of a default OAT. After starting the OSA/SF GUI, the screen normally comes up with the tree view. If you use the View pull-down and select **Style** -> **Details View,** you will see a display similar to the one in Figure 6-2.

Note: Keep in mind that the use of OSA/SF requires a special device called the OSAD device.

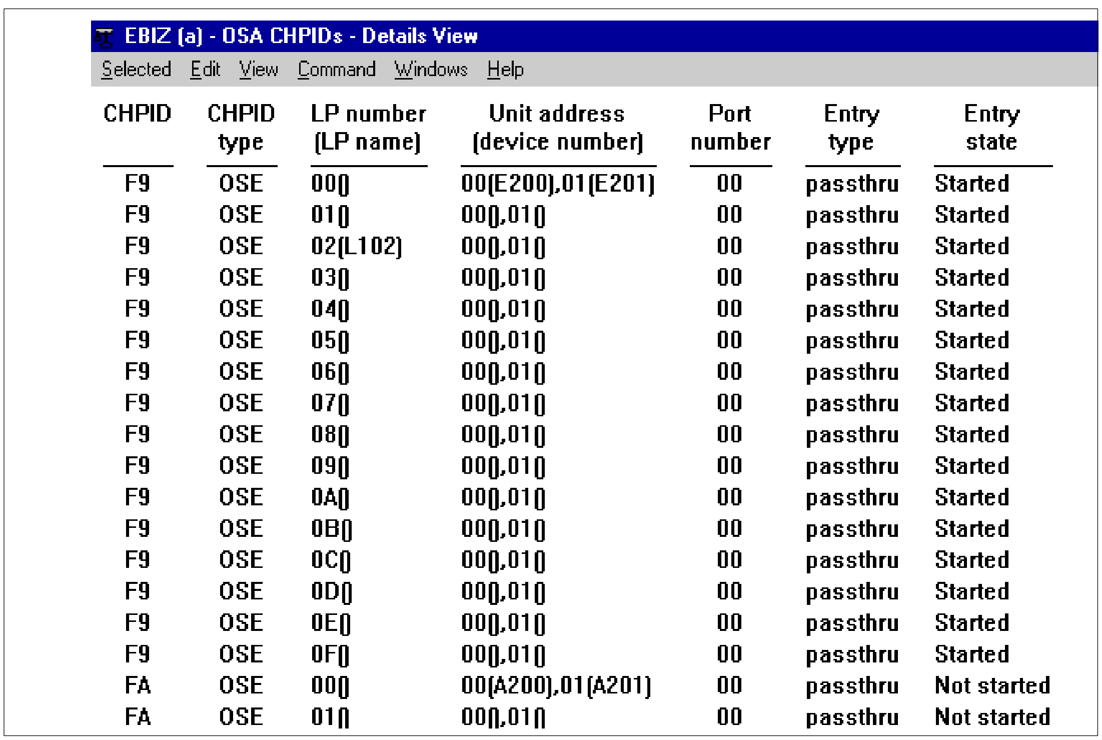

CHPID	CHPID type	LP number (LP name)	Unit address (device number)	Port number	Entry type	Entry state
F9	OSE	00()	00(E200),01(E201)	00	passthru	Started
F9	OSE	01()	00(),01()	00	passthru	Started
F9	OSE	02(L102)	00(),01()	00	passthru	Started
F9	OSE	03()	00(),01()	00	passthru	Started
F9	OSE	04()	00(),01()	00	passthru	Started
F9	OSE	05()	00(),01()	00	passthru	Started
F9	OSE	06()	00(),01()	00	passthru	Started
F9	OSE	07()	00(),01()	00	passthru	Started
F9	OSE	08()	00(),01()	00	passthru	Started
F9	OSE	09()	00(),01()	00	passthru	Started
F9	OSE	0A()	00(),01()	00	passthru	Started
F9	OSE	0B()	00(),01()	00	passthru	Started
F9	OSE	0C()	00(),01()	00	passthru	Started
F9	OSE	0D()	00(),01()	00	passthru	Started
F9	OSE	0E()	00(),01()	00	passthru	Started
F9	OSE	0F()	00(),01()	00	passthru	Started
FA	OSE	00()	00(A200),01(A201)	00	passthru	Not started
FA	OSE	01()	00(),01()	00	passthru	Not started

Figure 6-2 OSA default OAT display

This display is an example of an OSA-Express feature dedicated to one partition. Notice that the devices are identified to LP 0. To see which LPAR the CHPID is dedicated to, search the list until you find the partition for this CHPID that has a name next to it. In this example, you can see that CHPID F9 is dedicated to LP 2 (L102).

LPAR 0 is unique in that this value is used to identify that an OSE CHPID is dedicated. No matter which LPAR the OSE CHPID may be dedicated to, the LPAR number used is always 0. This is especially important if a configuration is being activated on the OSA-Express feature.

6.4 Customizing z/OS TCP/IP

This section describes the definitions required in TCP/IP for Fast Ethernet OSA-Express in TCP/IP Passthru mode.

When an external physical device is defined to the TCP/IP stack, it is defined via a set of DEVICE and LINK statements. For Passthru mode, you can define only one LINK statement for its related DEVICE statement (as compared to other device types, where there may be more than one LINK statement to refer to the same DEVICE statement).

Non-QDIO OSA-Express ports are defined to TCP/IP as LAN Channel Station (LCS) devices. You need to assign an IP address by coding an entry for the LINK name in the HOME statement. Depending on your network design, you will also need to code a route entry. The OSA-Express can be activated via the TCP/IP profile, using a START statement.

Figure 6-3 shows the network and the connections of our configuration example.

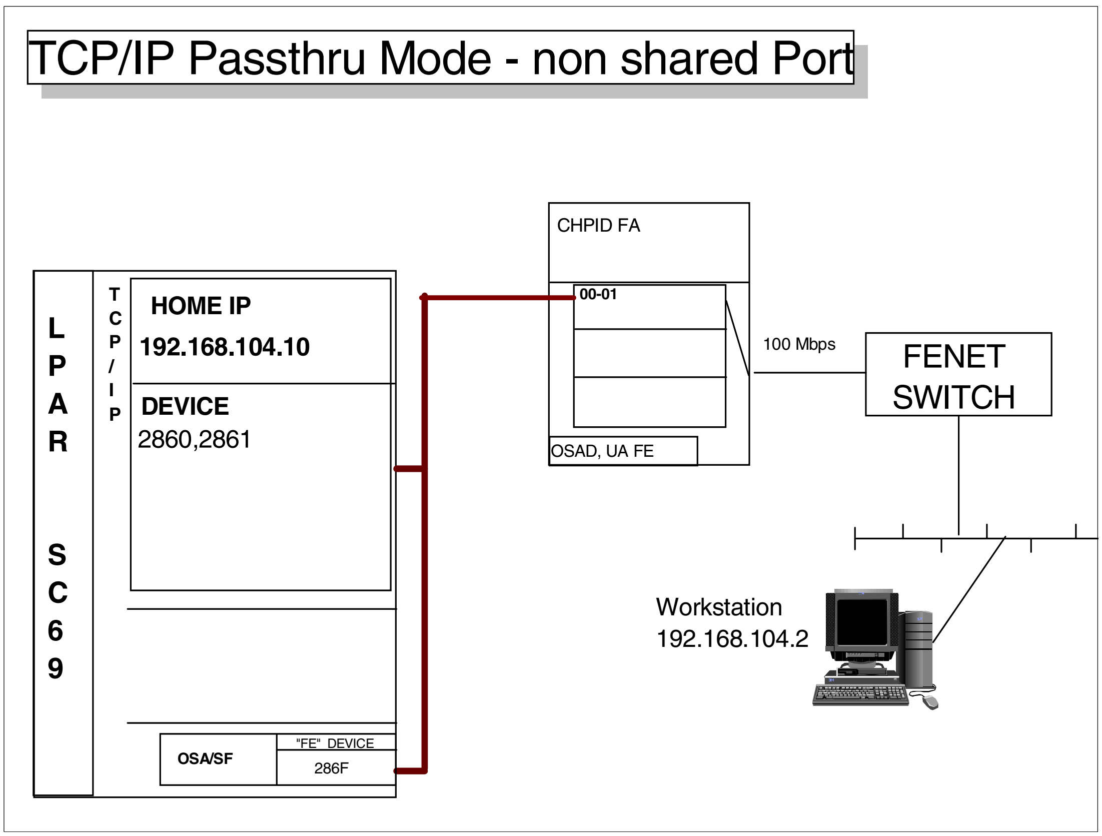

Figure 6-3 TCP/IP Passthru mode with non-shared port

6.4.1 TCP/IP definitions

Our OSA-Express feature is using CHPID FA, with port 0 connected to a Fast Ethernet switch. The z/OS device addresses used for port 0 are 2860 and 2861.

We defined the following:

- One DEVICE statement per OSA port, with the even device number of the two device addresses assigned to the hardware for the port. Note the following considerations:
  - Even though you only define one port, you will use two device numbers per OSA port for TCP/IP Passthru mode. One device is used by TCP/IP for reading, the other device is used for writing.
  - Using the DEVICE statement, you define the DEVICE name, the DEVICE type (LCS) for the OSA port, and the DEVICE number (the read device number, which is the even device number).
- One LINK statement per OSA TCP/IP DEVICE statement.
  - Using the LINK statement, you define the LINK name, the LINK type (in our example, ETHEROR802.3), the PORT number, and the DEVICE name.

    Note: Although the OSA port is known by the device number, the port number in the LINK statement must match the actual OSA port number. With OSA-Express Fast Ethernet, the port will always be 0.

> **Note:** We defined the link type ETHEROR802.3. If the OSA card is Token Ring, you must define the link type IBMTR. Although the OSA port is addressed by the device number, the port number in the LINK statement must match the actual OSA port number.

- A HOME IP address (in our example, we used IP address 192.168.104.10 for LPAR SC69).
  - The HOME statement relates an IP address to the OSA described in the DEVICE and LINK statement pair.
- Routes using static routes through the BEGINROUTES statement.

> **BEGINROUTES statement:**
>
> If you are migrating your TCP/IP profile from an earlier release, it may use the GATEWAY statement to define static routes instead of the BEGINROUTES - ENDROUTES statements. Gateway will be recognized and used, but you should consider replacing it with BEGINROUTES.
>
> Because it is compatible with UNIX standards, easier to code than GATEWAY, accepts both IPv4 and IPv6 addresses, and has enhanced functionality, BEGINROUTES is the recommended method for defining static routes. Future static route enhancements will only be available with the BEGINROUTES statement.

  Dynamic routing can be accomplished using the OMPROUTE daemon, or the BSDROUTINGPARMS statement can be used in conjunction with the RouteD daemon.

> **Note:** Use of the BEGINROUTES statement (static routes) with the OMPROUTE or OROUTED routing daemons is not recommended.

- One START statement per OSA device.
  - The START device statement entry uses the DEVICE statement name.

Example 6-2 on page 78 shows the TCP/IP profile definitions we used to define OSA to TCP/IP for LPAR SC69.

Example 6-2 TCP/IP profile definitions for LPAR SC69

```
DEVICE OSA2860   LCS              2860
 LINK   OSAL2860  ETHEROR802.3   0  OSA2860

HOME
   192.168.104.10   OSAL2860

BEGINROUTES
 192.168.104.0/24 =    OSAL2860 MTU 1500
ENDROUTES

START    OSA2860
```

6.5 Activation

Activation may require several tasks, such as:

- Ensuring the devices are online
- Activating an OSA/SF configuration
- Activating VTAM resources
- Activating TCP/IP

Since the OSA-Express feature is being used in default mode, very little needs to be done here. We only need to ensure that the devices are online, and then activate TCP/IP.

6.5.1 Verify devices are online

The z/OS console display command can verify that the required devices are online; see Figure 6-4.

```
D U,,,2860,2
IEE457I 21.29.17 UNIT STATUS 836
UNIT TYPE STATUS        VOLSER     VOLSTATE
2860 OSA  O
2861 OSA  O
```

Figure 6-4 z/OS D U command

If they are not online, issue the z/OS console vary command:

```
V (2860-2861),ONLINE
```

6.5.2 Activate TCP/IP

There are two ways to make the definitions added to the TCP/IP profile effective. One way is to create an obeyfile using the definition statements we list in this chapter; the other is to restart the TCP/IP stack. There are relative merits for both techniques, but we chose to restart the TCP/IP stack.

We confirmed the status of the TCP/IP devices with a **NETSTAT DEV** command. Figure 6-5 on page 79 shows the results of a **NETSTAT DEV** command; note that both the device and the link are in ready status.

```
DevName: OSA2860           DevType: LCS       DevNum: 2860
   DevStatus: Ready
   LnkName: OSAL2860          LnkType: ETHOR       LnkStatus: Ready
     NetNum: 0   QueSize: 0   ByteIn: 0000659463   ByteOut: 0000632590
     BroadcastCapability: Yes
   BSD Routing Parameters:
     MTU Size: 00000            Metric: 00
     DestAddr: 0.0.0.0          SubnetMask: 255.255.255.0
   Multicast Specific:
     Multicast Capability: Yes
     Group             RefCnt
     -----             ------
     224.0.0.1         0000000001
```

Figure 6-5 TCP/IP Device and Link status

For a summary of related commands, refer to Appendix A, “Commands” on page 175.

7

Mixed TCP/IP and SNA (non-QDIO mode)

In this chapter, we discuss customizing the Fast Ethernet OSA-Express feature for both TCP/IP and SNA (mixed) non-QDIO mode, using OSA/SF. The OSA-Express feature port will be shared between two LPARs, each running a TCP/IP stack and a VTAM region.

OSA-2 note: The procedures in this chapter deal with configuring an OSA-Express feature. However, these same procedures can also be used to configure the OSA-2 version of this feature (with the exception of the HCD definition, CHPID type OSA).

© Copyright IBM Corp. 2003. All rights reserved.

7.1 Configuration Information

Figure 7-1 shows a functional view of the connectivity discussed in this chapter.

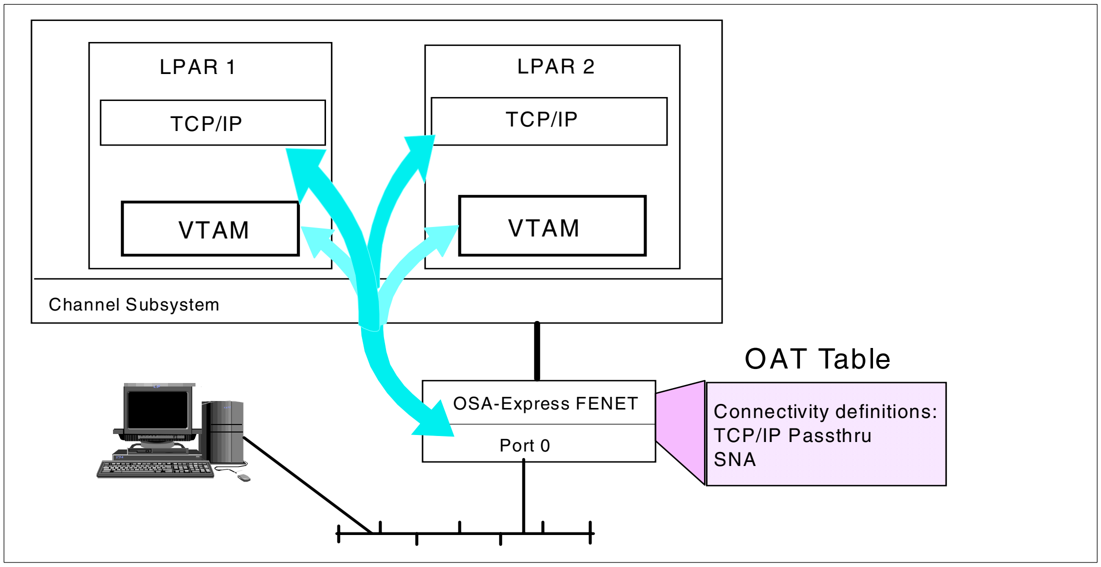

Figure 7-1 TCP/IP mixed mode shared port

The OSA feature configuration file contains the OSA Address Table (OAT). The OAT maps the partition, mode, unit address, and IP address to an OSA port. The OSA port is defined as a LAN Channel Station (LCS) device to TCP/IP. The OAT configuration enables the OSA feature to associate different IP addresses in separate partitions with the same OSA port.

VTAM sees the OSA feature as an External Communications Adapter (XCA). One OSA port is linked to one device for SNA operation.

When the Activate command is issued from OSA/SF, the OSA configuration transfers the SNA configuration and downloads the OAT to the OSA feature by using the OSA feature's FE unit address (OSAD device). After the SNA mode settings and OAT have been downloaded, the OSA feature CHPID is automatically configured offline to all systems, and then configured online. This causes the OSA feature hardware to be reset, and activates the new OSA configuration or OAT.

Following are the definitions needed to support this configuration:

- Hardware definitions
- OAT definitions
- TCP/IP definitions
- VTAM definitions

7.2 HCD definition requirements

The OSA CHPID, control unit, and OSA devices must be defined to HCD, and these HCD definitions must be activated. Refer to Chapter 3, “Hardware configuration definitions” on page 37 for the procedure to create the definitions.

An example of the IOCDS definitions used for CHPID FA is shown in Example 7-1 on page 83. For future considerations, we defined 15 OSA devices. However, only three are needed in our configuration.

Example 7-1 Example of the IOCDS input for CHPID FA

```
.
CHPID PATH=(FA),SHARED,PARTITION=((A11,A12),(A11,A12)),TYPE=OSE
.
CNTLUNIT CUNUMBR=2860,PATH=(FA),UNIT=OSA
.
IODEVICE ADDRESS=(2860,15),CUNUMBR=(2860),UNIT=OSA,UNITADD=00
IODEVICE ADDRESS=286F,UNITADD=FE,CUNUMBR=(2860),UNIT=OSAD
```

If you are using QDIO (CHPID type=OSD), follow the procedures documented in Chapter 5, “QDIO mode” on page 61.

7.3 Create and save the OSA feature configuration with OSA/SF GUI

Creating (adding) and saving an OSA feature configuration is not disruptive. The only time a definition can have any effect on the OAT configuration is when the Activate command is issued. We are implementing the OAT definitions shown in Figure 7-2.

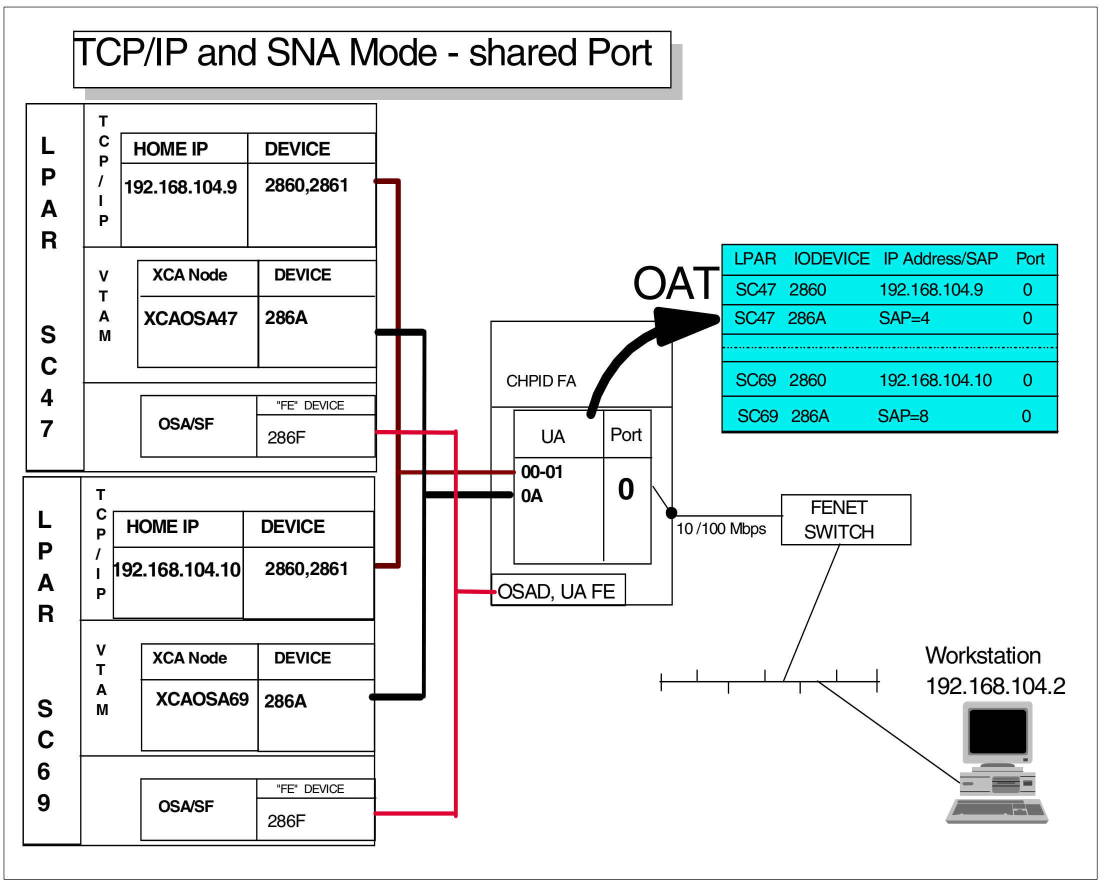

Figure 7-2 Connectivity layout

The OSA/SF GUI must be connected to a z/OS host that has OSA/SF running. The OSA CHPID must also have been defined by HCD, and the HCD definition activated. This step does not require the CHPID to be installed or online; it only creates and saves an OSA feature configuration.

1. Start the OSA/SF GUI program:
   – Double-click the **OSA/SF** folder icon.
   – Double-click the **OSA/SF** program icon.
2. Start the OSA/SF host connection by double-clicking the icon previously created to connect your workstation to the host OSA/SF and logging on.

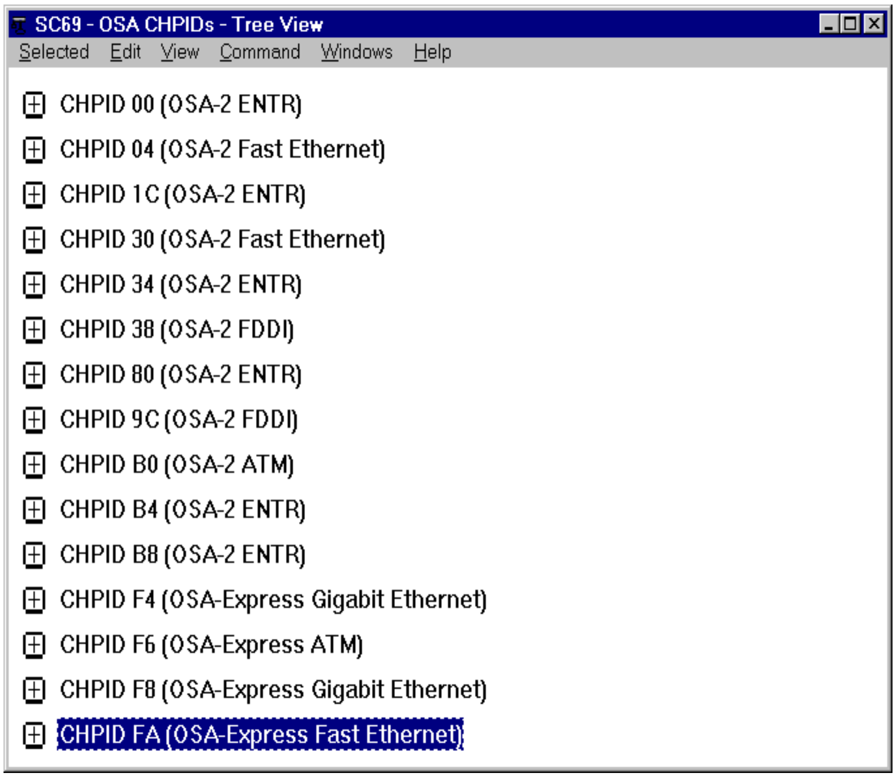

Figure 7-3 List of OSA features

Figure 7-3 shows all the OSA features that were installed on the system used in these examples. CHPID FA is highlighted because we will be working with CHPID FA throughout this chapter.

3. Select the OSA CHPID that represents the OSA port you are going to put into mixed TCP/IP and SNA modes.
4. Choose **Selected** from the OSA/SF menu bar.
5. Select **Configurations** from the Selected pull-down menu.

   Note: If the OSA CHPID has not been defined to the server's channel subsystem (that is, it is not in the list), or if the OSA feature has not been installed, use the Planning option instead of the Configurations option.
6. Select **Configuration lists...** from the cascade menu.
7. On the "Configuration list for CHPID FA" panel, click **Add...** to create a new configuration.

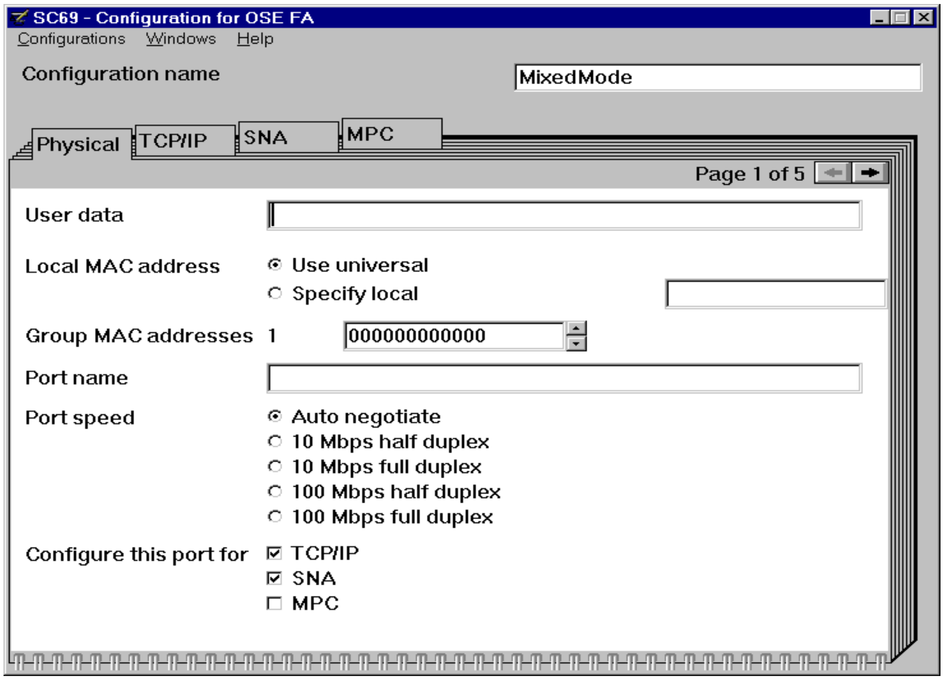

Figure 7-4 Add a new configuration

8. On the "Configuration for OSE FA" panel, shown in Figure 7-4, enter a configuration name in the Configuration name entry field. This is the name that will be displayed in the configuration list panel.

- Optionally, you can enter some comments in the User data field. This field has no effect on the operation of the OSA, but a short description of what this configuration is designated for can be useful later.
- We set the Local MAC address to Universal and the Group MAC address to zeros, which is the default value. For your site, you should check the local networking standards for the correct values.

> **Note:** You may want to use a locally administered address (LAA) in place of the universal address when using SNA. Each workstation connecting to the OSA port through SNA must have the MAC address of the defined OSA port. If the OSA feature is replaced, it is much easier to change the MAC address on the OSA feature than changing all of the workstation profiles to point to the new universal MAC address of the new feature. The MAC address of the OSA feature can also be changed through OSA Advanced Facilities at the HMC.

- A port name is required; it can be up to 8 characters in length. Like the User data field, this is not used by OSA specifically. The value here should not duplicate other port names in your network.
- Port speed was set to Auto negotiate. This will allow the OSA to automatically set its speed to the speed of the LAN to which it is connecting. If you set port speed to a specific value, ensure that the LAN is also capable of this speed.
- Finally, we configured the port for both TCP/IP and SNA traffic by selecting the box for each protocol type.

Note: MPC is not a supported option for the zSeries. If you are using a 9672 G5 or G6 and you want to use MPC, refer to "HPDT MPC (IP)" on page 8 for a discussion of this and a documentation list.

7.3.1 TCP/IP definitions in OSA/SF

1. Select either the **TCP/IP** tab, or the right arrow next to the Page 1 of 5 counter, to proceed to the TCP/IP setup. Figure 7-5 shows the empty TCP/IP OAT entries panel.

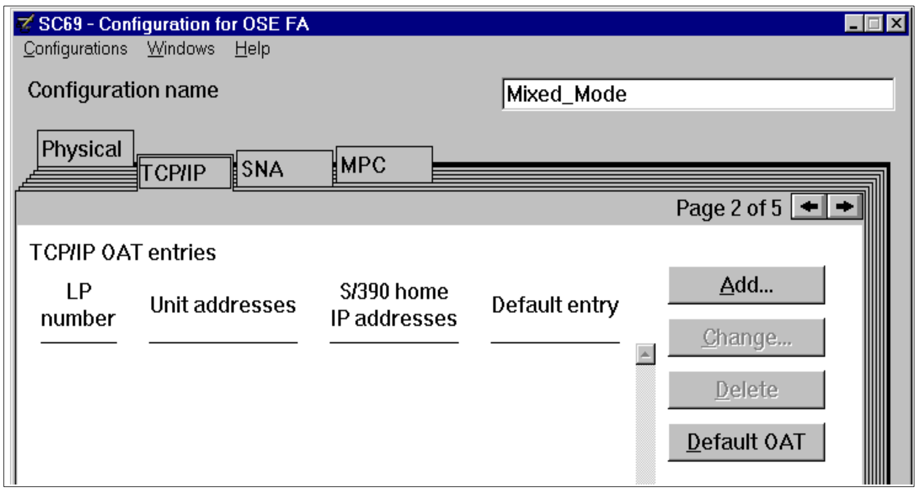

Figure 7-5 TCP/IP Configuration for OSE FA

2. Click **Add...** to move on to the entry panel shown in Figure 7-6.

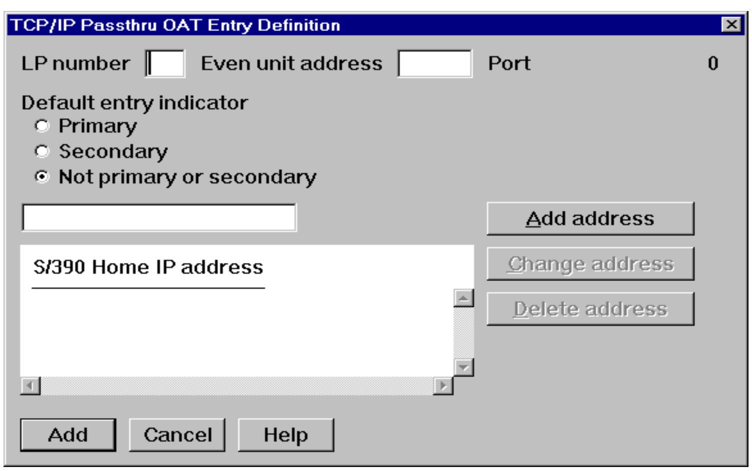

Figure 7-6 Add TCP/IP OAT entry definitions

In Shared Port mode, more than one partition can share an OSA port. The OAT record definition is used to assign a unit address and IP address to each partition/port combination.

3. Enter the LP number (logical partition number) of the partition accessing this port. In our example, partition SC47 is x'B'. To determine the partition number for a partition name, refer to the HCD information in Chapter 3, "Hardware configuration definitions" on page 37.

   Reminder: If you are using a CHPID dedicated to one partition, the LP number must be 0.

4. Enter the Even unit address for this port. The unit address can be any even address, but unit address 00,01 is associated with OSA port 0 in the default OAT table.

   Note: The OSA port unit address was used by HCD in Chapter 3 when defining the OSA devices.

5. You can select **Primary** for one of the partitions using this port.

   The partition designated as the Primary will receive any datagrams that are not specifically addressed to any of the home IP addresses associated with this OSA port.

6. In the field directly below **Not primary or secondary**, enter the home IP address 192.168.104.9 for LPAR SC47 and click on **Add address**.

   Any time an OSA port (in TCP/IP mode) is shared, each partition's TCP/IP home IP address must also be added to the OAT. This allows the OSA feature to forward the received datagrams to the appropriate partition.

7. Click **Add** to add this definition to the OSA/SF OAT.
8. For the next LPAR, enter the LP number of the partition that is to have access to this port. (In this case, partition SC69 is x'C'.)
9. Enter the unit address for this port.
10. The same unit address that was assigned to the other partitions (LP 11) may be used.
11. You can select **Secondary** or **Not primary or secondary**.

    The partition designated as the Secondary will receive any datagrams that are not specifically addressed to any of the home IP addresses associated with this OSA port, when the datagram cannot be sent to the primary partition.

    If Secondary is selected, `secondary` will appear in the Default Entry column.

    If Not primary or secondary is selected, `No` will appear in the Default Entry column.

12. Enter the IP address for the partition, and click **Add address.**
13. Click **Add...** to add this definition to the OSA/SF OAT.
14. Click **Cancel** to leave the OAT record definition dialog panel.

The TCP/IP OAT Settings panel is displayed again; see Figure 7-7.

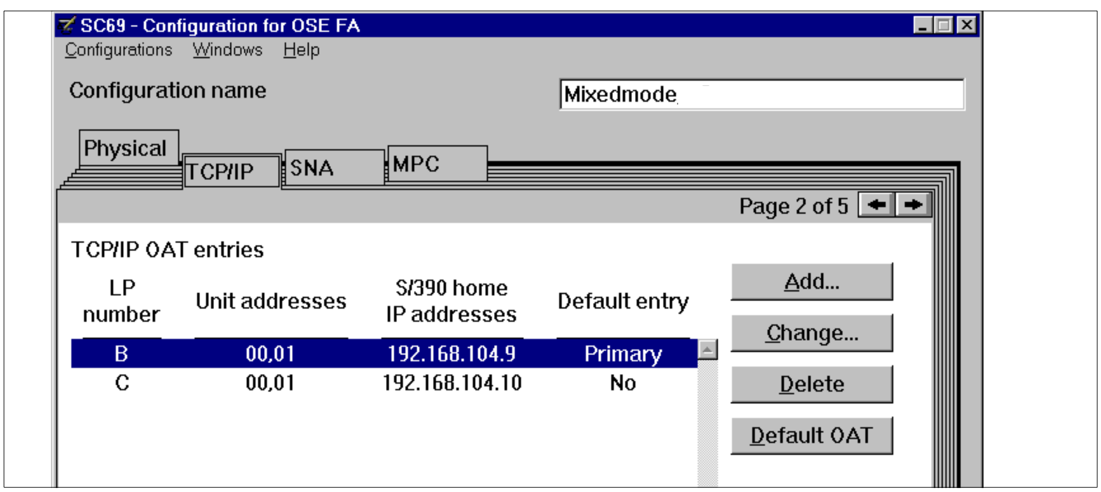

Figure 7-7 Configuration for TCP/IP

Verify the configuration information.

7.3.2 SNA definition in OSA/SF

Because the Open System Adapter Configuration is being defined to operate in both TCP/IP and SNA modes, we now have to define the SNA configuration mode.

1. Click the **SNA** tab, then click the right arrow, which should display page 4 of the configuration panels. Click **Add...** to display the SNA OAT Entry Definition panel; see Figure 7-8.

Figure 7-8 SNA OAT Entry Definition

2. The SNA OAT record definition panel is used to assign a single unit address to each OSA port for every partition that will access the port. In SNA mode, each OSA port that VTAM is sharing must have a unique SAP address. On the SNA OAT record entry panel, do the following:

   a. Enter the LP number of the partition that is to have access to this port. To determine the partition number for a partition, refer to the HCD panel in Chapter 3, “Hardware configuration definitions” on page 37.

      Reminder: If you were using a CHPID dedicated to one partition, the LP number must be 0.

   b. Enter the unit address for this port. We used unit address 0A.

      The unit address can be any address, but unit addresses 00-03 are usually associated with TCP/IP Passthru ports.

   c. Click **Add...** to add this definition to the OSA/SF OAT. The information message `IOAG961I RECORD DEFINITION WAS SUCCESSFULLY ADDED` should be displayed.

   d. To configure OSA/SF to allow SC69 to share the OSA feature ports in SNA mode, do the following on the SNA OAT record definition:

      - Enter the LP number of the partition that is to have access to this port (in this case, C).

      > **Note:** Defining both logical partitions to use the same port is known as Shared Port. In SNA mode, no additional information is required in the OAT definition—unlike TCP/IP Passthru mode, which requires the IP address to be added to the OAT.
      >
      > Shared port is only possible when the VTAMs that are sharing the OSA port each use a unique SAP address. Enter the unit address for this port.

      - Click **Add...** to add this definition to the OSA/SF OAT.

3. On the “Configurations for OSA” panel, do the following:

   a. Select **Configuration** from the menu bar.

   b. Select **Save** from the configuration pull-down menu.

      The Save action takes the configuration information and saves it into a file on the OS/390 host.

The data set name on the host will be *hlq*.PWSOSA*xx*.CFGNUM*y*, where:

- *hlq* is the data set qualifier specified in the OSA/SF startup profile.
- *xx* is the CHPID number.
- *y* is a number from 1 to 99.

4. Message IOAG917A will be displayed when the configuration data has been successfully saved on the OS/390 host.
5. To activate this configuration, use the **Activate** option on the Configuration pull-down menu; or to end, press Alt+F4 (close) until all open panels have been closed.

7.3.3 Activate OSA configuration

All TCP/IP and SNA mode configuration changes are disruptive, so before you activate the TCP/IP and SNA mode configurations, all devices having an active session through the OSA port must be stopped. In contrast with previous versions of OSA/SF, the ports may remain online.

1. Cease all active sessions using OSA ports.
2. Start the OSA/SF GUI program:
   a. Double-click the **OSA/SF folder** icon.
   b. Double-click the **OSA/SF program** icon.
3. Start the OSA/SF host connection by double-clicking the icon previously created to connect your workstation to the host OSA/SF.
4. The “OSA Channels” panel will be displayed; do the following:
   a. Select the OSA CHPID (which represents an OSA port) that you are going to put into TCP/IP and SNA modes.
   b. Select **Selected** from the menu bar.
   c. Select **Configurations** from the Selected pull-down menu.
   d. Select **Configuration List** from the cascade menu.
5. On the “Configuration list for CHPID FA” panel, do the following:
   a. Select the configuration you want to activate.
   b. Click **Change** to select the configuration action panel.
6. The “Configuration for OSA” panel will be displayed; do the following:
   a. Select **Configurations** from the menu bar.
   b. Select **Activate...** from the configurations pull-down menu.

   The activate action performs two functions:

   - The OAT information from the currently selected configuration is reformatted and saved on the OS/390 host in the OSA configuration file. The data set name on the host will be *hlq*.PWSOSA*xx*.OATFILE, where:

     - *hlq* is the data set qualifier specified in the OSA/SF startup profile.

     - *xx* is the CHPID number.

     The OSA/SF config file (also defined in the startup profile) is updated to point to any files that are required to support this configuration and that will be downloaded to the feature during any OSA/SF install action.

   - An OSA/SF install action is done to download any required files (specified in the OSA/SF config), as well as download the OAT table contained in the OATFILE data set.
7. Message IOAG966I will be displayed to indicate that the activation has completed successfully.

8. Verify that the required devices are online (see Figure 7-9).

```
IEE457I 12.00.17 UNIT STATUS 219
 UNIT TYPE STATUS        VOLSER     VOLSTATE
 2860 OSA  O
 2861 OSA  O
 2862 OSA  OFFLINE
 2863 OSA  OFFLINE
 2864 OSA  OFFLINE
 2865 OSA  OFFLINE
 2866 OSA  OFFLINE
 2867 OSA  OFFLINE
 2868 OSA  OFFLINE
 2869 OSA  OFFLINE
 286A OSA  O
 286B OSA  OFFLINE
 286C OSA  OFFLINE
 286D OSA  OFFLINE
 286E OSA  OFFLINE
 286F OSAD O-RAL
```

Figure 7-9 Display of devices 2860-286F

7.3.4 Display the MAC address

In SNA mode, any downstream workstation doing a dial-in operation (like PCOM as a 3270 terminal) needs to know the MAC address of the OSA port. To display this address, you can use the OSA/SF GUI:

1. From the “OSA Channel - Tree View” panel, do the following:

   Click the plus sign (**+**) beside the CHPID of interest, then click the **+** next to the port. A list of ports will be displayed. Double-click the desired port number to display the OSA port settings; see Figure 7-10.

2. Observe the MAC address value at the top of the mode settings notebook to confirm that it matches the MAC address provided in the planning section.

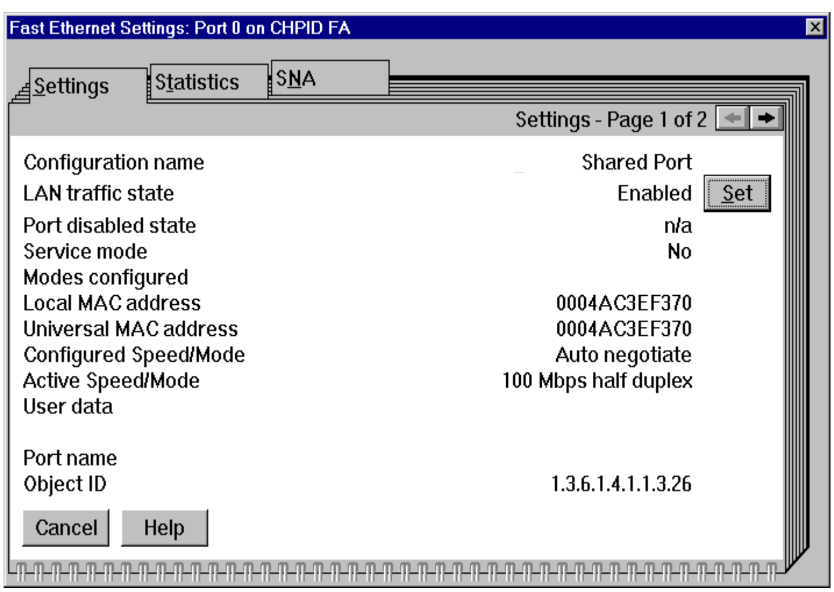

Figure 7-10 Reviewing MAC address values

In our example, we used the Universal Address 0004AC3EF370, which appears in the MAC address field.

> **Reminder:** You may want to use a locally administered address (LAA) in place of the Universal Address for SNA, as discussed earlier.
>
> Each workstation connecting to the OSA port through SNA must have the MAC address of the OSA port defined. If the OSA feature is replaced, it is much easier to change the MAC address on the OSA feature, than it is to change all the workstation profiles to point to the new universal MAC address of the new feature.
>
> The MAC address of the OSA feature can be changed through OSA/SF using the configuration panels, or through OSA Advanced Facilities at the HMC.

Now the Open System Adapter Configuration is ready to operate in the defined mode.

7.4 Customizing the z/OS host network subsystem

In this mixed mode and shared port environment, TCP/IP and VTAM coexist and share the same OSA port without affecting each other. This allows the definitions for TCP/IP and VTAM to be done independently, as if the OSA-Express port was owned by VTAM or TCP/IP exclusively.

Port sharing is not set up in TCP/IP or VTAM, but in OSA/SF.

7.4.1 VTAM definitions

This section describes the definitions required in VTAM to allow SNA applications to access the LAN via the OSA port. To describe the VTAM setup, the network configuration shown in Figure 7-2 on page 83 is used. In this example, both VTAMs in SC47 and SC69 are communicating over a Fast Ethernet switch via port 0 (device number 286A).

You need to define two types of major nodes in VTAM:

- External Communication Adapter major node (XCA)
- Switched major node

XCA major node

Define one XCA major node for each SNA OSA device with the following:

- The node type (VBUILD definition statement)
- The port used by the local area network (Port definition statement)
- The switched peripheral nodes (type 2) attached to an Ethernet LAN through an OSA port (Group, Line and PU definition statements)

The VTAM coding to implement this connection is shown in Example 7-2 on page 92.

Example 7-2 XCA major node definition for SC47

```
************************************************************************
*                                                                      *
* XCA MAJNODE for FAST ETHERNET OSA-Express                            *
*                                                                      *
************************************************************************
XCAOSA47 VBUILD TYPE=XCA
OSASNAM  PORT   MEDIUM=CSMACD,                                          X
                ADAPNO=0,                                               X
                CUADDR=286A,                                            X
                TIMER=60,                                               X
                SAPADDR=04
************************************************************************
OSASNAG  GROUP  DIAL=YES,                                               X
                DYNPU=YES,                                              X
                ANSWER=ON,                                              X
                AUTOGEN=(3,L,P),                                        X
                CALL=INOUT,                                             X
                ISTATUS=ACTIVE
******************************** Bottom of Data************************
```

Table 7-1 lists the Port parameters.

Table 7-1 XCA major node Port definition for XCAOSA47

Required Parameters	Explanation	Remarks
TYPE=XCA	XCA major node	The OSA functions like an XCA to VTAM.
ADAPNO=0	*PORT* statement	Code **ADAPNO=0** for port 0 of OSA-Express FENET.
CUADDR= 286A	Channel unit address	Code the device number defined for this port. In our example, SC47 uses device number 286A for port 0.
MEDIUM= CSMACD	LAN type	Use RING for Token Ring, or CSMACD for Ethernet.
SAPADDR=04	Service access point address	Code a value that is a *multiple of 4*. This address *must* be unique for each VTAM communicating with a port. Use different SAP addresses if a port is shared by multiple VTAMs. See the SAPADDR value of XCAOSA69.

Note: We defined the medium CSMACD; if the OSA card is Token Ring, you must define the medium type RING.

Table 7-2 identifies the significant Group parameters.

Table 7-2 XCA major node Group definition for XCAOSA47

Required Parameters	Explanation	Remarks
DIAL=YES	Switched peripheral node	You *must* code **DIAL=YES** to specify that the switched line control protocol is required.

AUTOGEN=(3,L,P)	Autogeneration of LINE and PU statements	This parameter enables VTAM to automatically generate three sets of LINE and PU statements. The LINE names begin with *L*. The PU names begin with *P*.

Note: The current OSA-Express features support 4096 SNA PU Type 2 connections per port.

The XCA major node (XCAOSA69) definition in VTAM on LPAR SC69 for the Ethernet connection via port 0 of the OSA-Express is given in Example 7-3.

Example 7-3 SNA mode XCA major node definition for XCAOSA69

```
*************************************************************************
*                                                                       *
* XCA MAJNODE for FAST ETHERNET OSA-Express                             *
*                                                                       *
*************************************************************************
XCAOSA69 VBUILD TYPE=XCA
OSASNAM  PORT   MEDIUM=CSMACD,                                           X
                ADAPNO=0,                                                X
                CUADDR=286A,                                             X
                TIMER=60,                                                X
                SAPADDR=08
*************************************************************************
OSASNAG  GROUP  DIAL=YES,                                                X
                DYNPU=YES,                                               X
                ANSWER=ON,                                               X
                AUTOGEN=(3,L,P),                                         X
                CALL=INOUT,                                              X
                ISTATUS=ACTIVE
******************************** Bottom of Data*************************
```

Table 7-3 shows the only parameter that is different between XCAOSA47 and XCAOSA69.

Table 7-3 XCA major node Port definition for XCAOSA69

Parameters	**Explanation**	**Remarks**
SAPADDR=08	Service access point address	Since port 0 is shared with SC47, a different SAP address of 8 is being used.

Switched major node

Define *one* switched major node for each switched connection to the peripheral nodes attached on the LAN. Code the following:

- A remote physical unit (PU definition statement)
- The corresponding logical units (LU definition statements)

Example 7-4 on page 94 shows how 3270 sessions are set up with a switched major node (SWOSA47) for SC47.

Example 7-4 Switched major node definition for SWOSA47

```
* THIS SWITCHED MAJNODE DEFINES THE OSA-Express SNA CONNECTION
* TO VTAM RUNNING IN LPAR SC47
*
         VBUILD TYPE=SWNET
OSA47P   PU    ADDR=02,                                         X
               IDBLK=05D,                                       X
               IDNUM=12863,                                     X
               CPNAME=OSANT,                                    X
               IRETRY=YES,                                      X
               MAXOUT=7,                                        X
               MAXPATH=1,                                       X
               MAXDATA=1024,                                    X
               PACING=0,                                        X
               VPACING=0,                                       X
               PUTYPE=2,                                        X
               DISCNT=(NO),                                     X
               ISTATUS=ACTIVE,                                  X
               MODETAB=NEWMTAB,                                 X
               DLOGMOD=DYNTRN,                                  X
               USSTAB=USSLDYN,                                  X
               SSCPFM=USSSCS
OSA47L0  LU     LOCADDR=0,MODETAB=MTAPPC,DLOGMOD=APPCMODE
OSA47L1  LU     LOCADDR=1                                 3270 SESSIONS
OSA47L2  LU     LOCADDR=2
```

Table 7-4 lists the important parameters in the PU definition for SWOSA47.

Table 7-4 Switched major node PU definition for SWOSA47

Parameters	**Explanation**	**Remarks**
TYPE=SWNET	Switched major node	
ADDR=02	PU's SDLC station address	OSA47P is the PU name. *ADDR* is a required parameter.
PUTYPE=2	PU type	The PUTYPE *must* be 2 for a LAN switched station. Type 2 also denotes a type 2.1 PU.
CPNAME=OSANT	Control point (CP) name of a type 2.1 node	To dial in, a type 2.1 peripheral node on a switched line requires either CPNAME or both IDBLK/IDNUM. However, you can code all three operands.

Table 7-5 presents the important parameters in the LU definition for SWOSA47.

Table 7-5 Switched major node LU definition for SWOSA47

Parameters	**Explanation**	**Remarks**
LOCADDR=0	LU's local address at the PU	LOCADDR=0 denotes an independent LU.
MODETAB=MTAPPC	Logon mode table	Code a separate logon mode table for APPC.
LOCADDR=2		LOCADDR=2 denotes a dependent LU.

7.4.2 TCP/IP definitions

This section describes the definitions required for TCP/IP.

TCP/IP uses the OSA ports as LAN Channel Station (LCS) devices. For OSA, you define only one LINK statement for the associated DEVICE statement. Other device types may define and associate more than one LINK statement to refer to the same device statement. For OSA, each port has its own unique DEVICE and LINK statement defined in the TCP/IP profile.

Figure 7-2 on page 83 shows the network and the connections for our configuration example.

Port 0 is connected to an Ethernet LAN using IP address 192.168.104.9 for LPAR SC47, and IP address 192.168.104.10 for LPAR SC69.

- Define one DEVICE statement per OSA port. Use the even device number of the two device numbers assigned in the hardware to the port.
  - You will have to define two device numbers per OSA port for TCP/IP mode, in HCD, because TCP/IP is running in full-duplex mode. One device is used by TCP/IP for reading, and the other device is used for writing.
  - Using the DEVICE statement, define the DEVICE statement name, the DEVICE type (LCS) for the OSA port, and the DEVICE number (the read number, which is the even number).
- Define one LINK statement per OSA TCP/IP DEVICE statement.

  Using the LINK statement, define the LINK name, the LINK type, the PORT number, and the DEVICE statement name.

Note: We defined the link type `ETHEROR802.3`; if the OSA card is Token Ring, you must define the link type IBMTR. Although the OSA port is addressed by the device number, the port number in the LINK statement must match the actual OSA port number.

- Define the HOME IP address of the OSA port.
- Using the HOME statement, define an IP address referring to a LINK statement name.
- Define static routes through the BEGINROUTES statement.

BEGINROUTES statement

If you are migrating your TCP/IP profile from an earlier release, it may use the GATEWAY statement to define static routes instead of the BEGINROUTES - ENDROUTES statements. Gateway will be recognized and used, but you should consider replacing it with BEGINROUTES.

Because it is compatible with UNIX standards, easier to code than GATEWAY, accepts both IPv4 and IPv6 addresses, and has enhanced functionality, BEGINROUTES is the recommended method for defining static routes. Future static route enhancements will only be available with the BEGINROUTES statement.

Dynamic routing can be accomplished using the OMPROUTE daemon, or the BSDROUTINGPARMS statement can be used in conjunction with the RouteD daemon.

Note: Use of the BEGINROUTES statement (static routes) with the OMPROUTE or OROUTED routing daemons is not recommended.

- Define one START command per DEVICE name.

  After defining an OSA device in the TCP/IP profile, the device still has to be started explicitly. The TCP/IP START device statement entry has one per TCP/IP Read device that is required to be started. It uses the DEVICE statement name.

Example 7-5 shows the TCP/IP PROFILE definitions needed to define OSA to TCP/IP for LPAR SC47.

Example 7-5 OS/390 TCP/IP PROFILE definitions for LPAR SC47

```
; TCP/IP SC47 PROFILE.TCPIP --- Hardware definitions
; OSA2860 is an OSA-Express feature using logical port 0.  Devices are
; 2860-2861.CHPID FA.

DEVICE OSA2860    LCS                 2860

LINK      OSAL2860  ETHERORB02.3 0  OSA2860

HOME
192.168.104.9 OSAL2860

BEGINROUTES
 192.168.104.0/24 =    OSAL2860 MTU 1500
ENDROUTES

;Start all the defined devices.
START     OSA2860
```

The configuration for LPAR SC69 is exactly the same with one exception: the HOME statement is coded to reflect the different IP address:

```
HOME
192.168.104.10 OSAL2860
```

7.5 Activation

Once all the definitions have been added to OSA/SF, VTAM, and TCP/IP, we can activate the configuration.

Activation may require several tasks, such as:

- Verifying the devices are online
- Activating an OSA/SF configuration
- Activating VTAM resources
- Activating TCP/IP

7.5.1 Verify devices are online

The z/OS console display command can verify that the required devices are online.

```
D U,,,2860,2
IEE457I 21.29.17 UNIT STATUS 836
UNIT TYPE STATUS        VOLSER     VOLSTATE
2860 OSA  O
2861 OSA  O

D U,,,286A,1
IEE457I 21.32.17 UNIT STATUS 936
UNIT TYPE STATUS        VOLSER     VOLSTATE
286A OSA  O
```

Figure 7-11 z/OS D U command

If they are not online, issue the OS/390 console vary command:

```
V (2860,2861,286A),ONLINE
```

7.5.2 OSA/SF activation

If the configuration is not activated yet, follow the next procedure. Otherwise, you can skip the OSA-Express activation.

Highlight the CHPID you wish to activate from the Tree View OSA/SF display. Choose **Selected -> Configuration -> Configuration List** (see Figure 7-12).

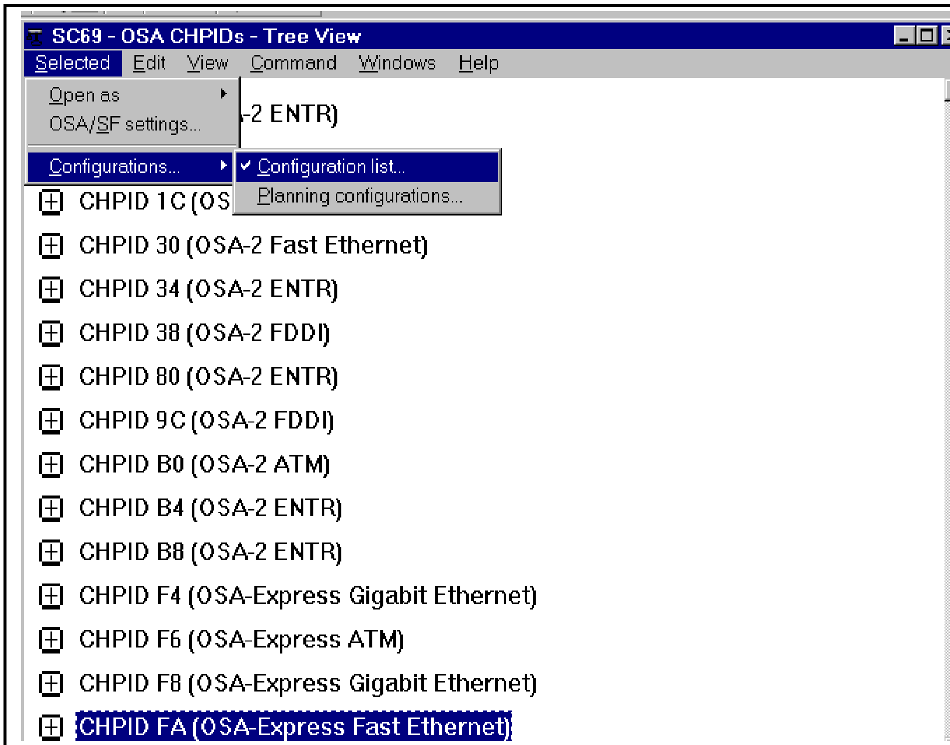

Figure 7-12 Selecting a configuration

On the list of configurations, highlight Mixedmode and click **Change**. This will retrieve the OSA configuration. You can now choose to either **Activate** or **Activate(no install)** from the Configurations pull-down; see Figure 7-13 on page 98.

The difference between these two choices is the timing for the activation. If there are no active sessions using this OSA-Express, the Activate command will load the new configuration down to the OSA-Express feature and recycle the hardware connections. This process can take up to three minutes.

If you select Activate(no install), the new configuration will be loaded to the OSA/SF in preparation for loading the new configuration at a later time. When you use the no install option, you need to return to the Tree View window and select **Install** from the commands pull-down list at a later time.

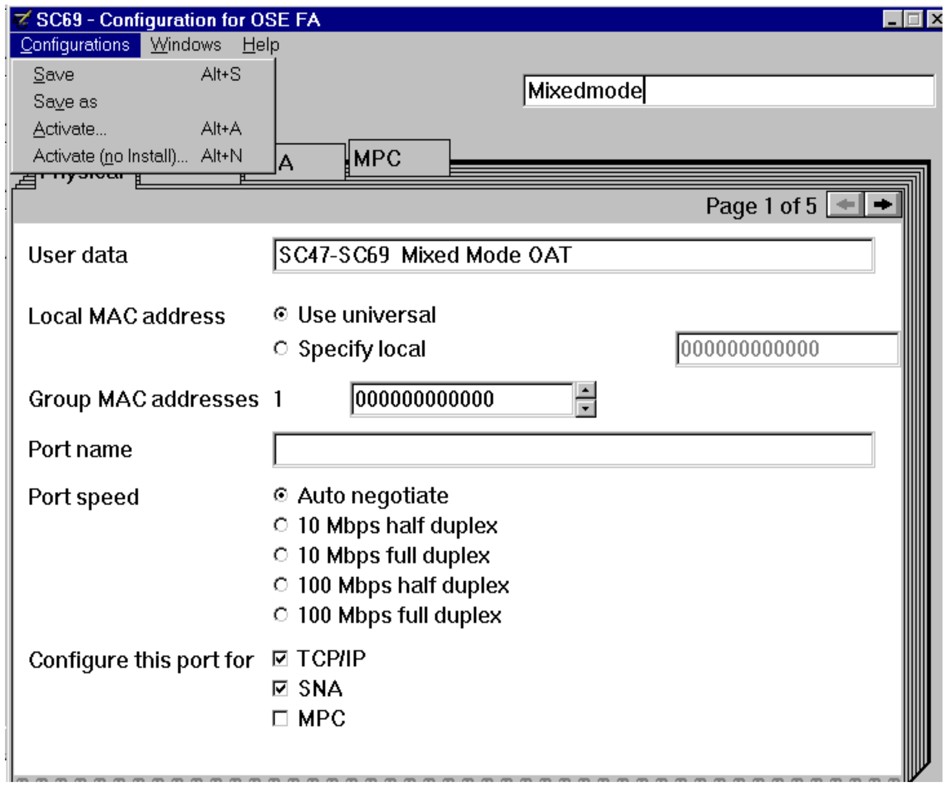

Figure 7-13 Activating the OSA configuration

> **Note:** The configuration defines the access from both LPARs. The activation command makes the OSA configuration active to both LPARs.

7.5.3 VTAM activation

The VTAM activation is no different from any other VTAM resource. Use the VTAM **VARY** command. On each LPAR, activate the XCA major node and the Switched major node. Typically, the commands will be of the following format (for partition SC69):

```
V NET,ID=XCAOSA69,ACT
V NET,ID=SWOSA69,ACT
```

7.5.4 TCP/IP activation

There are two ways to activate the TCP/IP devices: either restart the TCP/IP stack, or use the TCP/IP **Obeyfile** command. We chose to restart the stack to implement the changes.

7.6 Relevant status displays

We monitored the status of the VTAM resources with the VTAM DISPLAY command. Figure 7-14 displays the XCA major node for the Fast Ethernet connection.

```
D NET,ID=XCAOSA69,E
IST097I DISPLAY ACCEPTED
IST075I NAME = XCAOSA69, TYPE = XCA MAJOR NODE 715
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1021I MEDIUM=CSMA/CD,ADAPNO= 0,CUA=286A,SNA SAP= 8
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST232I L286A000 ACTIV
IST232I L286A001 ACTIV
IST232I L286A002 ACTIV
IST314I END
```

Figure 7-14 Display of XCA major node

Figure 7-15 shows the results from the Switched major node.

```
D NET,ID=SWOSA69,E
IST097I DISPLAY ACCEPTED
IST075I NAME = SWOSA69, TYPE = SW SNA MAJ NODE 718
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST084I NETWORK RESOURCES:
IST089I OSASNAP  TYPE = PU_T2           , CONCT
IST089I OSASNAL1 TYPE = LOGICAL UNIT    , CONCT
IST089I OSASNAL3 TYPE = LOGICAL UNIT    , CONCT
IST314I END
```

Figure 7-15 Display of Switched major node

The NETSTAT DEV command displays the TCP/IP devices. Figure 7-16 shows both the device and link in the READY state.

```
DevName: OSA2860           DevType: LCS       DevNum: 2860
   DevStatus: Ready
   LnkName: OSAL2860          LnkType: ETHOR       LnkStatus: Ready
     NetNum: 0   QueSize: 0   ByteIn: 0000659463   ByteOut: 0000632590
     BroadcastCapability: Yes
   BSD Routing Parameters:
     MTU Size: 00000            Metric: 00
     DestAddr: 0.0.0.0          SubnetMask: 255.255.255.0
   Multicast Specific:
     Multicast Capability: Yes
     Group            RefCnt
     -----            ------
     224.0.0.1        0000000001
```

Figure 7-16 Display of TCP/IP device and link

To verify that a connection exists, we did a ping from one TCP/IP stack to the other.

Refer to Appendix A, “Commands” on page 175 for a list of other useful commands.

8

ATM HPDT native (non-QDIO mode)

In this chapter, we cover the implementation steps to set up an OSA-Express 155 ATM feature in High Performance Data Transfer (HPDT) mode. An OSA-Express 155 ATM (native) can only run in HPDT mode (QDIO is not supported). The ATM HPDT Native mode supports communication for TCP/IP as well as for SNA.

OSA-2 note: The procedures in this chapter deal with configuring an OSA-Express feature. However, these same procedures can also be used to configure the OSA-2 version of this feature (with the exception of the HCD definition, CHPID type OSA).

© Copyright IBM Corp. 2003. All rights reserved.

8.1 Configuration information

In our environment, we created a configuration to run in native mode with two ATM features on two different systems. Figure 8-1 shows the logical connectivity. We guide you through the steps required to:

- Create a new OAT table using the GUI interface
- Create VTAM definitions to support this configuration
- Create TCP/IP definitions to support this configuration

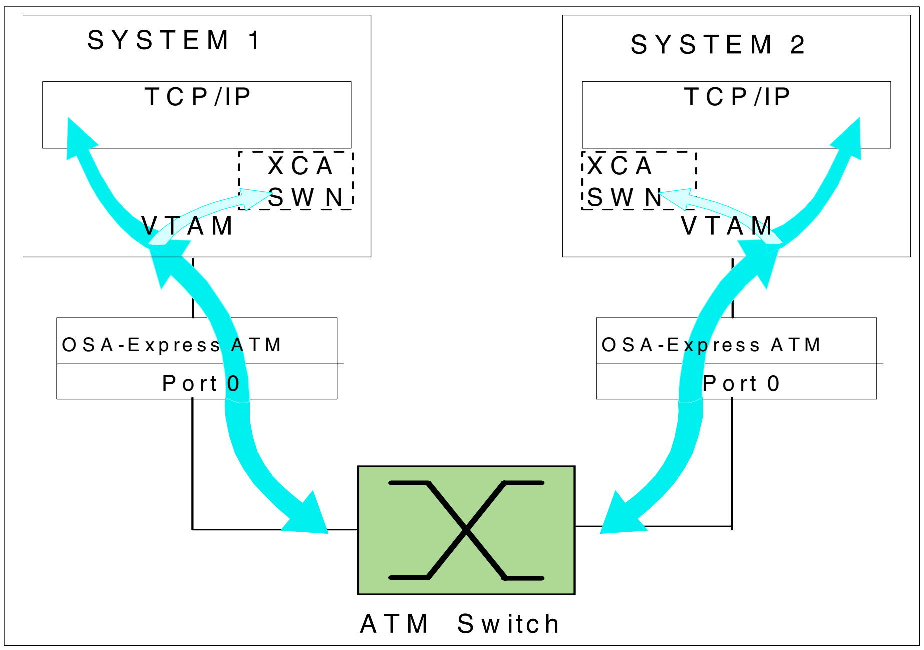

Figure 8-1 HPDT MPC ATM native mode

8.2 HCD definition requirements

The OSA CHPID, the control unit, and the OSA devices must be defined to HCD, and the HCD definitions activated. Refer to Chapter 3, "Hardware configuration definitions" on page 37 for the procedure to create the definitions. Example definitions for the IOCDS used in this chapter are shown in Example 8-1.

Example 8-1 Example of the IOCDS input for CHPID F6

```
.
CHPID PATH=(F6),SHARED,PARTITION=((A1,A2),(A1,A2)),TYPE=OSE
.
CNTLUNIT CUNUMBR=0E90,PATH=(F6),UNIT=OSA
.
IODEVICE ADDRESS=(0E90,15),CUNUMBR=(0E90),UNIT=OSA,UNITADD=90
IODEVICE ADDRESS=0E9F,UNITADD=FE,CUNUMBR=(0E90),UNIT=OSAD
```

8.3 Create and save the OSA configuration with OSA/SF GUI

OSA/SF is required to configure the OSA-Express 155 ATM feature.

Figure 8-2 shows the configuration for the OSA-Express 155 ATM. It shows interdependencies such as—the portname in the TRLE must match the device name in TCP/IP profile, as well as in the XCA major node (if SNA is used). Furthermore, the TRLE name must be the same as the OSA name defined in the OSA feature.

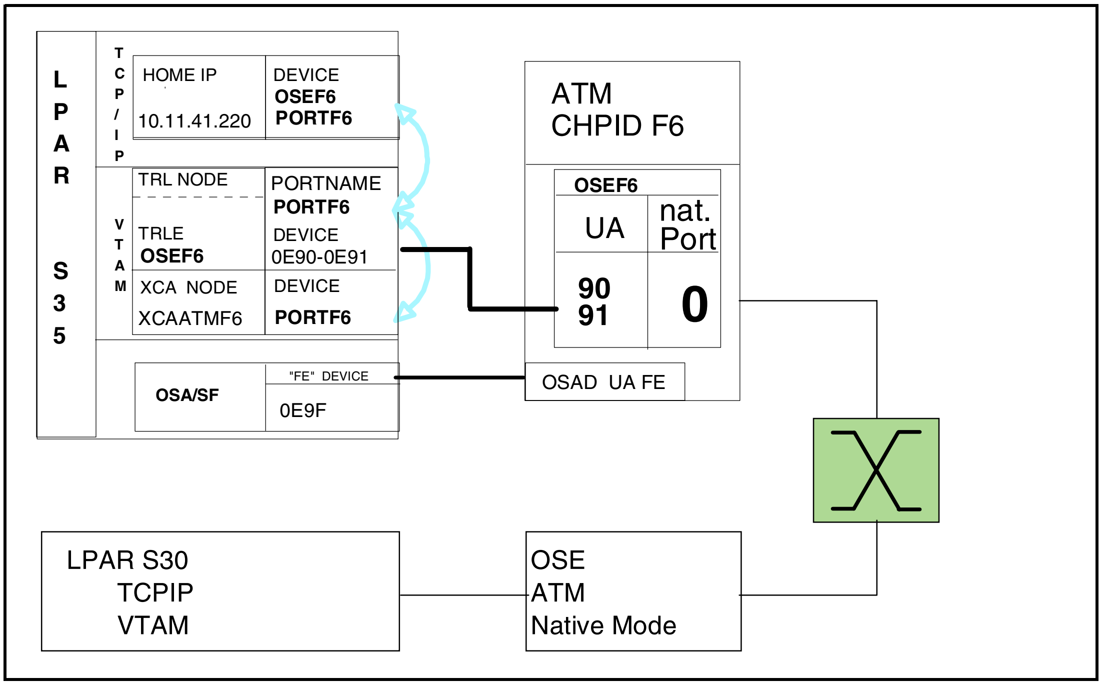

Figure 8-2 Devices in HPDT MPC ATM native mode

Creating (adding) and saving an OSA feature configuration is not disruptive. The OSA/SF GUI code must be connected to a z/OS host which has OSA/SF running. This step does not require the CHPID to be installed or online; we will be creating and saving an OSA feature configuration.

1. Start the OSA/SF GUI program:
   - Double-click the **OSA/SF** folder icon.
   - Double-click the **OSA/SF** program icon.
2. Start the OSA/SF host connection by double-clicking the icon previously created to connect your workstation to the host OSA/SF, and logging on.

The OSA CHPIDs panel will be displayed, as shown in Figure 8-3 on page 104.

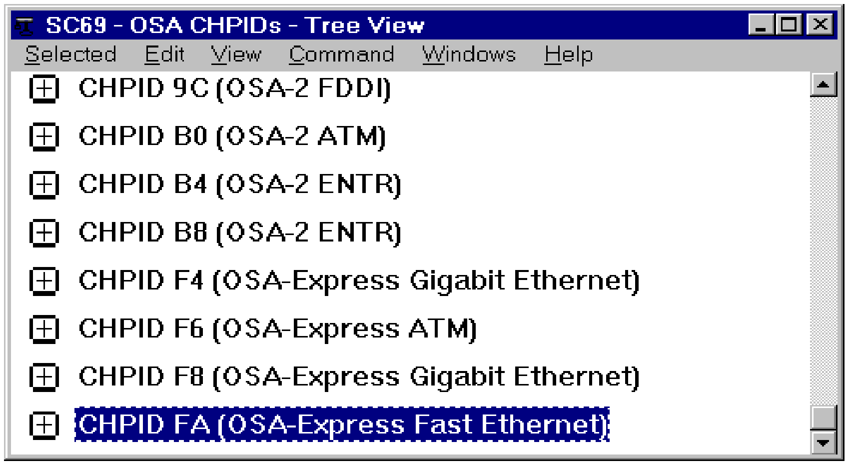

Figure 8-3 OSA CHPIDs Tree View

3. Select the OSA CHPID that represents an OSA port you are going to be configuring (CHPID F6, in our case).
4. Select **Selected** from the menu bar.
5. Select **Configurations....** from the Selected pull-down menu.
6. Select **Configurations list...** from the cascade menu.

Note: If the OSA CHPID has not been defined to the server's channel subsystem (that is, not in the list), or if the OSA feature has not been installed, use the Planning option instead of the Configurations option.

7. Press **Add...** to create a new configuration.

   The Configuration panel is displayed, as shown in Figure 8-4 on page 105. Note the following:

   a. Enter a configuration name in the Configuration name entry field. This is the name that will be displayed in the configuration list panel.
   b. Click **Yes,** to make sure that the physical port changes will be used.

      If you are configuring for partial activation on an OSA-Express 155 ATM feature with ports previously configured, choose **No**.
   c. Optionally, you can enter information in the Port description field. This field has no effect on the operation of the OSA, but a short description of what this physical port is designed to achieve can be useful later.
   d. The value in the Port name field must match the portname value of the TRLE statement and the portname in the device statement of the TCP/IP profile. If SNA will be used, it must also match the portname value in the XCA. It can be up to 8 non-blank characters in length.
   e. The Local end system ID (ESI) field is used to override the default of the universal system ID that was IBM-supplied for this ATM port, with a locally administered system ID.
   f. For the other fields of the physical port configuration panel, check the documentation of the ATM switch to which this ATM OSA-Express is attached, to ensure the correct settings are being used.

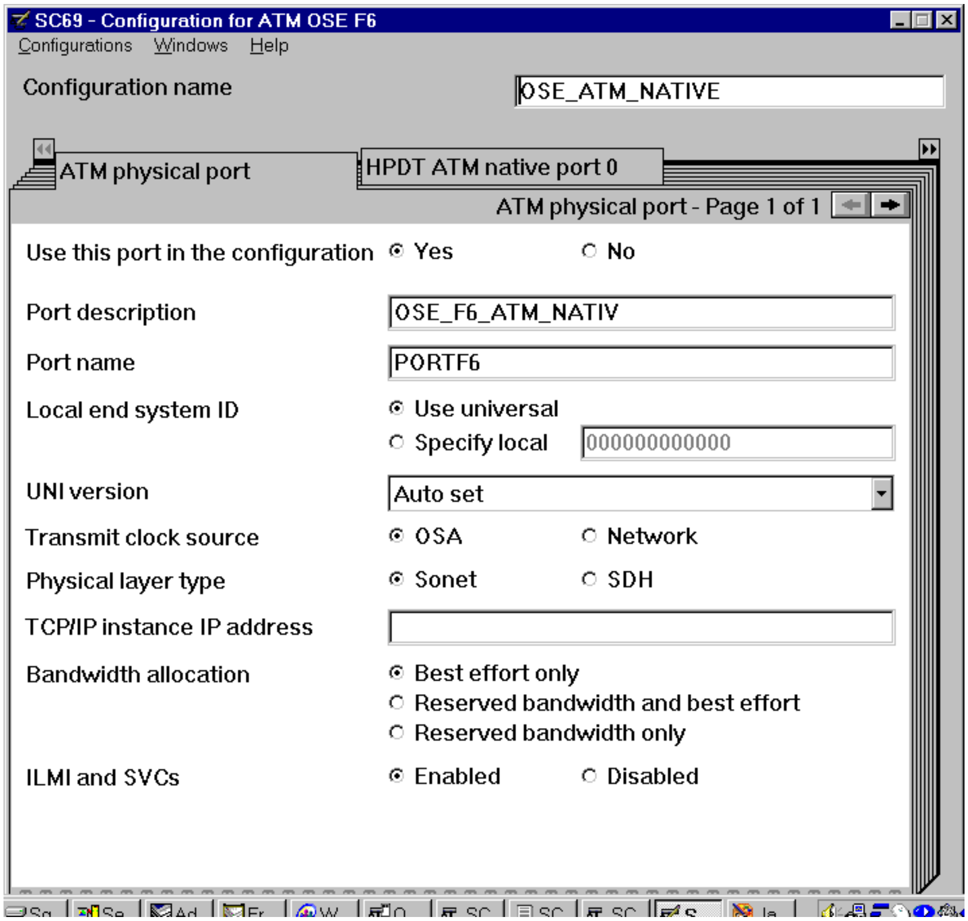

Figure 8-4 Configuration panel (physical port)

8. Select the **HPDT ATM native port 0** tab or click the right arrow in the upper right corner of the notebook until HPDT MPC native port 0 - Page 1 of 2 entries is displayed.
   – Select **Yes** at the Include in this configuration and Enable LAN traffic entries and click **Add**.
9. Specify an MPC OAT entry for this mode as shown in Figure 8-5, and then Press **Add...** to add an OAT port definition.

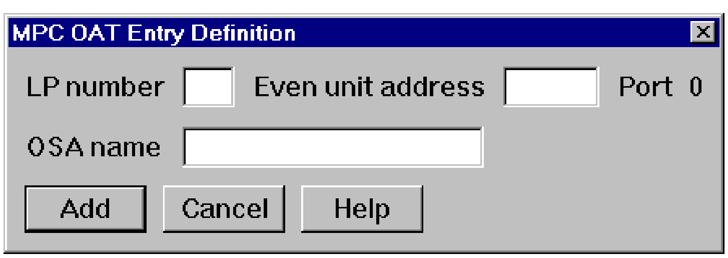

Figure 8-5 MPC OAT entries

Reminder: If you are using a CHPID dedicated to one partition, the LP number must be 0.

The OSA name specified must match the name of the TRLE in the TRL macro. It must also match the DEVICE and START statements in the TCP/IP profile. The OSA name can be the same for different partitions, but must be unique if you use more than one device pair address in the same logical partition.

Now the configuration for MPC mode is complete, and additional entries can be added by pressing the add key. Figure 8-6 shows the completed OAT entries for MPC mode.

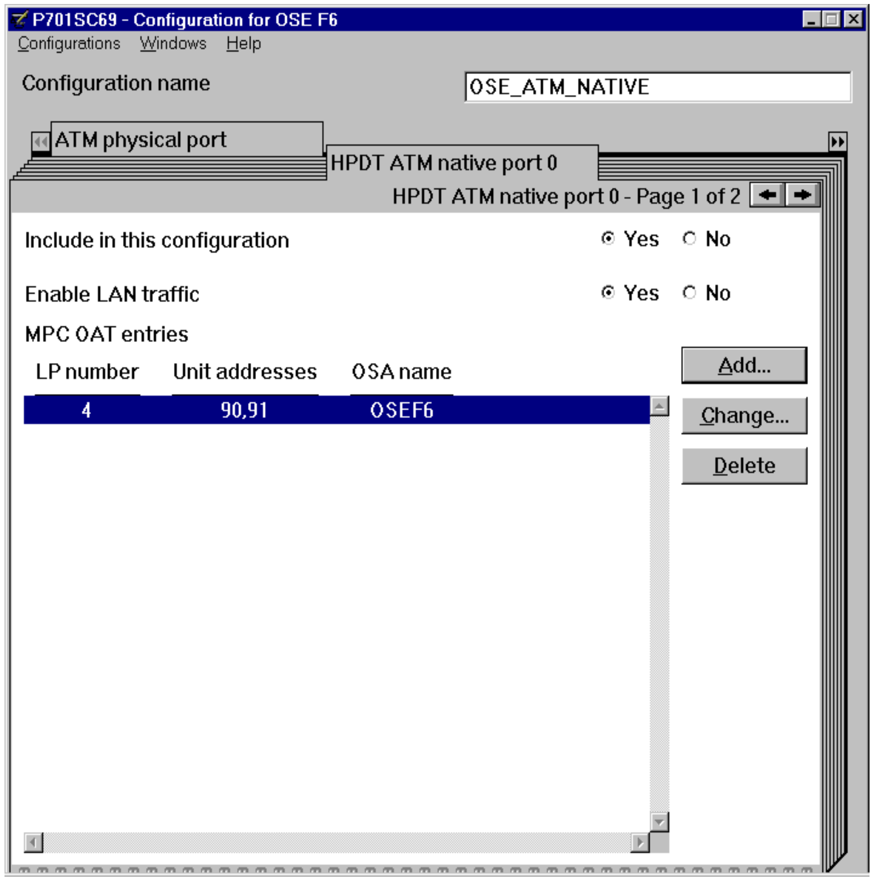

Figure 8-6 MPC ATM native port 0 - Page 1 of 2

10. Click the right arrow in the upper right corner of the notebook until HPDT MPC native port 0 - Page 2 of 2 entry is displayed and click **Add** to configure your permanent virtual circuits (PVCs) for your environment.
    - We configured two PVCs, one for TCP/IP and one for SNA traffic. Note the following:
        - PVCs are only required if you are establishing a permanent virtual circuit between the OSA port and the ATM switch.
        - Switched virtual circuits (SVCs) are not defined. These are built dynamically.
11. Figure 8-7 on page 107 shows our settings for the PVC that is used for the TCP/IP traffic.

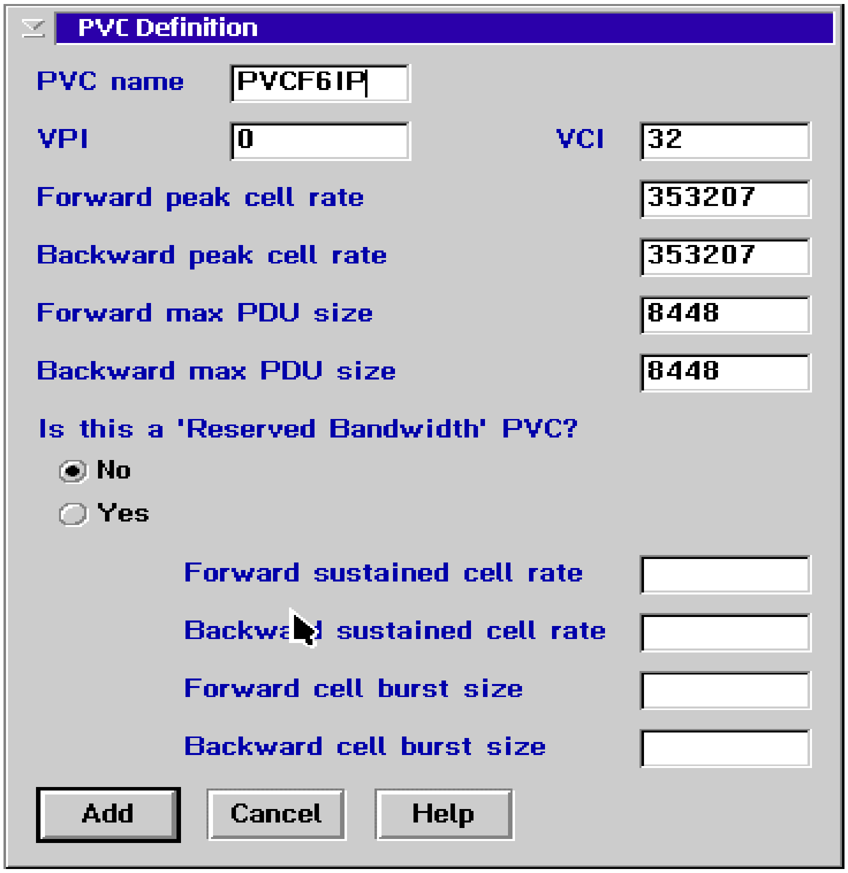

Figure 8-7 PVC Definition panel

a. Fill in the PVC name.

The PVC name must match the name of the ATMPVC statement in the TCP/IP profile.

b. Identify each PVC with a unique virtual path indicator and virtual channel indicator (VPI and VCI) value.

These values must agree with the VPI and VCI values assigned in the ATM switch.

c. The VPI value can be 0 through 15.

d. For each PVC, specify a VCI value from 32 through 8191.

e. Check your ATM switch documentation for the settings of cell rates and PDU sizes. Note the following hints:

- Specify the maximum protocol data unit (PDU) size for each direction in a PVC. The maximum PDU size is equal to the size of one frame that can be processed in that direction for the ATM AAL5 SDU layer.
- For a Best Effort virtual circuit, we recommend that you specify the highest peak cell rate that is acceptable to both endpoints.
- For SNA data transfer, correlate the maximum PDU size with the maximum RU size that you specify in the VTAM logmode table. For more information, see *z/OS V1R4.0 Communications Server SNA Resource Definition Reference*, SC31-8778.

f. Because we are using Best Effort for bandwidth allocation, we do not need to specify any of the sustained cell rates and cell burst size.

12. Figure 8-8 on page 108 shows the completed entries for the PVC records.

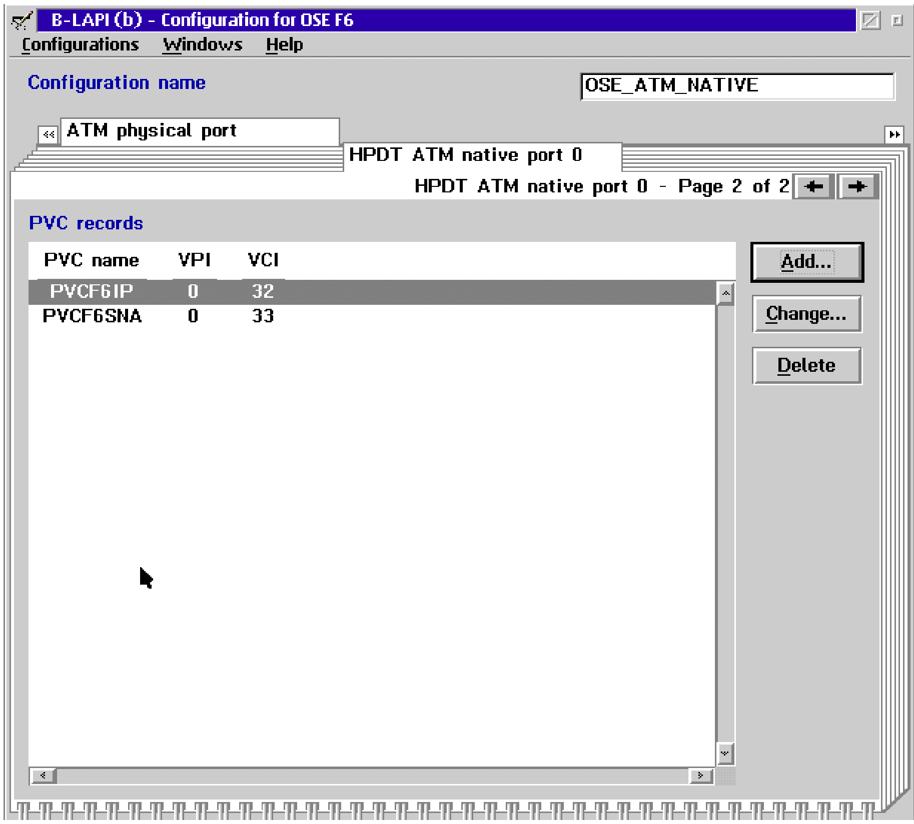

Figure 8-8 HPDT ATM native port - PVC records

8.3.1 Activate OSA configuration

1. From the pull-down menu, select **Configuration -> Save**. This will ensure that the configuration is saved. Once saved, you should see the message displayed in Figure 8-9.

Figure 8-9 Message to indicate configuration saved successfully

2. From the pull-down menu, select **Configurations** followed by **Activate(no install)**.

   Activate(no install) prevents disrupting an OSA that is already running with a different configuration. You can defer the install to a more appropriate time.

3. To complete the install do the following:

   a. From the OSA/SF GUI window OSA CHPIDs - Tree View, select the OSA CHPID that represents an OSA port you are going to install.

   b. Select **Command** from the menu bar.

   c. Select **Install**.

   d. Select **Force**, then **OK**.

8.4 Customizing the z/OS host network subsystem

This section describes the definitions needed in VTAM and TCP/IP, using Figure 8-2 on page 103 as a reference. TCP/IP uses VTAM TRL interfaces to run in MPC mode. The TRL is also shared for SNA traffic, which requires an XCA major node. A TRL major node must be defined and activated before TCP/IP starts its ATM native device, and the XCA major node can be activated to allow SNA traffic.

8.4.1 VTAM TRL definitions

Example 8-2 shows the VTAM TRL major node definition for LPAR S35.

Example 8-2 TRL major node for S35

```
TRLATMF6  VBUILD TYPE=TRL
OSEF6    TRLE LNCTL=MPC,
              READ=E90,
              WRITE=E91,
              PORTNAME=PORTF6,
              STORAGE=ECSA,
              MPCLEVEL=HPDT,
              MAXREADS=8,
              MAXBFRU=16
```

Table 8-1 lists and briefly describes the TRL parameters.

Table 8-1 TRL major node Port definition for OSEF6

Required parameters	Explanation	Remarks
TYPE=TRL	TRL major node	MPC TRL major node known to VTAM.
OSEF6	TRLE minor node	This name must match the OSA name defined in the OSA-Express ATM feature's MPC OAT entry and the DEVICE name in the TCP/IP profile.
READ=E90	Read device number (even number of the device pair)	Code the read device number defined for this port (port 0) and the related OSA name definition. In this example, S35 uses device number 0E90 for port 0 and OSA name OSEF6.
WRITE=E91	Write device number (odd number of the device pair)	Code the write device number defined for this port (port 0) and the related OAS name definition. In this example, S35 uses device number 0E91 for port 0 and OSA name OSEF6.
PORTNAME= PORTF6	Portname coded in this TRLE	The PORTNAME used in the TRLE must match the PORTNAME definition used in the TCP/IP profile, the port name in the ATM physical port panel of OSA/SF, as well as the portname for the XCA major node if using SNA.

8.4.2 TCP/IP definitions

TCP/IP requires device and link definitions corresponding to the TRLE name and the portname in the VTAM TRLE.

The PVC and SVC definitions for the ATM native mode IP traffic must also correspond to the definition made in the native port configuration in the OAS-Express ATM feature. Example 8-3 shows the TCP/IP profile definitions for S35.

Example 8-3 TCP/IP profile for S35

```
; TCPIP.PROFILE.TCPIP
; ===================

  ATMLIS LIS2 10.11.0.0 255.255.0.0 ; describes CIP LIS
  DEVICE OSEF6 ATM PORTNAME PORTF6 ENABLEIN
;
;for the SVC:
  LINK ATMSVCF6 ATM OSEF6 LIS LIS2
  ATMARPSV ARPVS1 LIS2 SVC 10.11.1.40
  NSAP 39999999999999000099990C030004ACAD54C100
;
;for the PVC:
  LINK ATMPVCF6 ATM OSEF6
  ATMPVC PVCF6IP ATMPVCF6

HOME
   10.11.41.240    ATMPVCF6       ; FOR THE PVC.
   10.11.41.241    ATMSVCF6       ; FOR THE SVC.

BEGINROUTES
 10.0.0.0/8        =    ATMPVCF6 MTU 4096 ; For the PVC
 10.11.41.231 HOST =    ATMSVCF6 MTU 9180 ; For the SVC
ENDROUTES

START OSEF6
```

Table 8-2 shows the interdependencies between the OSA-Express 155 ATM feature configuration, the TCP/IP profile, and the TRLE Statements in VTAM listed.

Table 8-2 TCP/IP statements in correlation to TRL and OAT

Required parameters	Explanation	Remarks
OSEF6	TCP/IP device name = OSA name	This name must match the OSA name defined in OSA-Express ATM feature's MPC OAT entry and the TRLE name.
PORTF6	Portname	The PORTNAME in the TCP/IP profile must match the PORTNAME value used in the TRLE and the port name in the OSA configuration.
PVCF6IP	PVC name	The PVC name specified in the OSA native port configuration for TCP/IP and the ATM switch.

Table 8-3 on page 111 describes the required statements for the TCP/IP profile used for ATM native mode IP traffic.

Table 8-3 TCP/IP statements description

Required parameters	Explanation	Remarks
ATMLIS	ATM logical IP subnet (LIS)	This statement is needed to describe the characteristics of the Logical IP Subnet. The LIS is used with ATM Classic IP (CIP).
DEVICE	TCP/IP device name	Device name must match the TRLE name and the OSA name in the OSA configuration.
LINK	Link name	Links must be specified for PVCs and SVCs separately. The SVC Link requires an ATMLIS statement defining the subnet and mask.
ATMARPSV	ATMARP server for CIP	Defines the IP address and ATM address of the ATMARP server using information provided by the ATMLIS statement.
HOME	IP home address	Specifies the home addresses for the defined SVC and PVC links.
BEGINROUTES	static route definition	We defined the home IP address of the corresponding SVC in the other system, because we are not using dynamic routing.
START		Starts the related TCP/IP devices.

The BEGINROUTES statement

If you are migrating your TCP/IP profile from an earlier release, it may use the GATEWAY statement to define static routes instead of the BEGINROUTES - ENDROUTES statements. Gateway will be recognized and used, but you should consider replacing it with BEGINROUTES.

Because it is compatible with UNIX standards, easier to code than GATEWAY, accepts both IPv4 and IPv6 addresses, and has enhanced functionality, BEGINROUTES is the recommended method for defining static routes. Future static route enhancements will only be available with the BEGINROUTES statement.

8.4.3 VTAM XCA and Switched Major Node definitions

This section discusses the XCA major node and switched major node definitions we used in our environment.

Example 8-4 is an example of how we coded our XCA major node.

Example 8-4 VTAM XCA major node for native ATM

```
XCAATMF6 VBUILD  TYPE=XCA
F6APORT  PORT  MEDIUM=ATM,
               PORTNAME=PORTF6
F6AGRP1  GROUP DIAL=NO
F6ALIN1  LINE  PVCNAME=PVCF6SNA
F6APU1   PU    CPNAME=VTAM30,CPCP=YES,HPR=YES,CONNTYPE=APPN,
               PUTYPE=2,TGP=ATMGBD,CAPACITY=49M,NETID=NETA
F6AGRP2  GROUP DIAL=YES,ANSWER=ON,ISTATUS=ACTIVE,CALL=INOUT,DYNPU=YES
F6ALIN2  LINE
F6APU2   PU
```

Table 8-4 lists the relevant definitions used in XCA major node.

Table 8-4 XCA major node description

Required parameters	Explanation	Remarks
TYPE=XCA	XCA major node	XCA major node known to VTAM.
MEDIUM=ATM	ATM Interface	VTAM SNA interface for ATM connection.
PORTNAME= PORTF6	Name of the port connecting to	Required for native ATM support; has to match PORTNAME in TRLE definition and the port name in the OSA configuration.
PVCNAME= PVCF6SNA	Name of PVC registered in the ATM switch	PVCNAME has to match the name defined in the OAT table.
HPR=YES	Enable HPR for this connection	HPR is required for native ATM support.

Next we defined the switched major node, as shown in Example 8-5.

Example 8-5 VTAM switched major node

```
SWATMF6  VBUILD TYPE=SWNET,MAXNO=256,MAXGRP=256
SWF6PU1  PU    MAXPATH=5,MAXDATA=256,ADDR=03,MODETAB=WALKTAB,
               CPNAME=VTAM30,CPCP=YES,HPR=YES,CONNTYPE=APPN,
               PUTYPE=2,TGP=ATMGBA,CAPACITY=49M,NETID=NETA
SWF6PTH1 PATH  DLCADDR=(1,C,ATMSVC,EXCLUSIVE),
               DLCADDR=(7,BCD,00,01,00353207,00353207),
               CALL=INOUT,
               DLCADDR=(8,X,00,00,00),
               DLCADDR=(21,X,0002,39999999999999000099990C01,
               0204357A09B5,60),
               GRPNM=F6AGRP2
```

Table 8-5 lists the relevant definitions used in switched major node.

Table 8-5 Description of switched major node

Required parameters	Explanation	Remarks
TYPE=SWNET	Switched Major	Used to established a switched connection with the PVC and SVC.
DLCADDR=	ATM partner Interface	The last three groups of the numbers represent the ATM that you connect to. The last 2 digits represents the selector. When you connect to a CMOS system, the first of these two digits is the LPAR number minus 1. The 12 digits just prior to the selector is the Local System Identifier (LSI) (or MAC address).

8.5 Activation

Once all the definitions have been added to OSA/SF, VTAM and TCP/IP, we can activate the configuration.

Activation may require several things, such as:

- Verifying the devices are online
- Activating an OSA/SF configuration
- Activating VTAM resources
- Activating TCP/IP

8.5.1 Verify devices are online

The z/OS console display command can verify that the required devices are online:

```
D U,,,E90e,2
IEE457I 21.29.17 UNIT STATUS 836
UNIT TYPE STATUS        VOLSER     VOLSTATE
0E90 OSA  O
0E91 OSA  O
```

If they are not online, issue the z/OS console vary command:

```
V (E90,E91),ONLINE
```

8.5.2 OSA/SF activation

Activation of the configuration loaded on the OSA-Express feature was already accomplished in 8.3.1, “Activate OSA configuration” on page 108.

8.5.3 VTAM activation

Next, activate the TRL using the following VTAM command:

```
V NET,ACT,ID=TRLATMF6
```

TCP/IP requires an active TRL prior to starting its device.

After activation of the TRL, the status of the TRLE is NEVAC or INACT until TCP/IP refers to it or the XCA major node is activated.

8.5.4 TCP/IP activation

There are two ways to activate the TCP/IP devices: either restart the TCP/IP stack, or use the TCP/IP VARY command. We chose to restart the stack to implement the changes. Refer to Appendix A, “Commands” on page 175 for command syntax.

8.6 Relevant status displays

The following displays were taken after we established a TCP/IP connection in OSA-Express native ATM mode. We used the following command:

```
D TCPIP,TCPIP,NETSTAT,DEV
```

Note that we split the PVC and SVC displays into two parts; TCP/IP displays it as one display.

Figure 8-10 shows the status of the PVC after activation. The PVC Status and LnkStatus must be READY.

```
LNKNAME: ATMPVCF6          LNKTYPE: ATM        LNKSTATUS: READY
NETNUM: 0   QUESIZE: 0   BYTEIN: 0000000108   BYTEOUT: 0000000540
BSD ROUTING PARAMETERS:
MTU SIZE: 00000            METRIC: 00
DESTADDR: 0.0.0.0          SUBNETMASK: 255.0.0.0
ATM SPECIFIC:
ATM PORTNAME:  PORTF6
ATM PVC NAME:  PVCF6IP          PVC STATUS: READY
```

Figure 8-10 NETSTAT DEVLINKS for native ATM PVC link

Figure 8-11 shows the status of the SVC.

```
DEVNAME: OSEF6             DEVTYPE: ATM       DEVNUM: 0000
DEVSTATUS: READY
LNKNAME: ATMSVCF6          LNKTYPE: ATM        LNKSTATUS: READY
NETNUM: 0   QUESIZE: 0   BYTEIN: 0000000000   BYTEOUT: 0000000000
BSD ROUTING PARAMETERS:
MTU SIZE: 00000            METRIC: 00
DESTADDR: 0.0.0.0          SUBNETMASK: 255.0.0.0
ATM SPECIFIC:
ATM PORTNAME:  PORTF6
ATM LIS NAME:  LIS2
SUBNETVALUE:   10.11.0.0        SUBNETMASK:    255.255.0.0
DEFAULTMTU:    0000009180       INACTVTIMEOUT: 0000000300
MINHOLDTIME:   0000000060       MAXCALLS:      0000001000
CACHENTRYAGE:  0000000900       ATMARPRETRY:   0000000002
ATMARPTIMEOUT: 0000000003       PEAKCELLRATE:  0000000000
NUMOFSVCS:     0000000000
```

Figure 8-11 NETSTAT DEVLINKS for native ATM SVC link

Next we displayed the status of VTAM TRLE node, as shown in Figure 8-12, by using the following command:

```
D NET,TRL,TRLE=OSEF6
```

```
NAME = OSEF6, TYPE = TRLE 194
STATUS= ACTIV, DESIRED STATE= ACTIV
TYPE = LEASED             , CONTROL = MPC , HPDT = YES
MPCLEVEL = HPDT       MPCUSAGE = SHARE
PORTNAME = PORTF6     LINKNUM = N/A   OSA CODE LEVEL =
HEADER SIZE = 4096 DATA SIZE = 56 STORAGE = ***NA***
WRITE DEV = 0E91 STATUS = ACTIVE     STATE = ONLINE
HEADER SIZE = 4092 DATA SIZE = 60 STORAGE = ECSA
READ  DEV = 0E90 STATUS = ACTIVE     STATE = ONLINE
END
```

Figure 8-12 TRLE display of native ATM connection

Note: If your static TRLE definition is incorrect, you should be aware that an active TRLE entry cannot be deleted. What you can do is vary activate the TRL node with a blank TRLE to cause the deletion of previous entries. Then code the correct TRL with correct TRLE entries and definition and vary activate this corrected TRL/TRLE node. See “VTAM commands” on page 177 for the vary commands.

We activated the SNA resources for ATM with the following commands:

```
V NET,ACT,ID=XCAATMF6
V NET,ACT,ID=SWATMF6
```

We displayed the XCA major node to establish SNA connectivity, as shown in Figure 8-13, by using the following command:

```
D NET,ID=XCAATMF6,E
```

```
NAME = XCAATMF6, TYPE = XCA MAJOR NODE 742
STATUS= ACTIV, DESIRED STATE= ACTIV
MEDIUM = ATM, PORT NAME = PORTF6
ATM ADDRESS                                 TYPE      FORMAT
399999999999999000099990C010020357CD01930    LOCAL     NSAP
I/O TRACE = OFF, BUFFER TRACE = OFF
VTAMTOPO = REPORT, NODE REPORTED - YES
LINES:
F6ALIN1  ACTIV
F6ALIN2  ACTIV
```

Figure 8-13   Display of XCA major node for SNA on ATM native mode

We displayed the status of the line, as shown in Figure 8-14, by using the following command:

```
D NET,ID=F6ALIN1,E
```

```
NAME = F6ALIN1, TYPE = LINE 784
STATUS= ACTIV, DESIRED STATE= ACTIV
TYPE = LEASED             , CONTROL = SDLC, HPDT = *NA*
PVCNAME = PVCF6SNA
GROUP = F6AGRP1, MAJOR NODE = XCAATMF6
TATE TRACE = OFF
VTAMTOPO = REPORT, NODE REPORTED - YES
F6APU1   TYPE = PU_T2.1          , ACTIV
END
```

Figure 8-14   Display of active line for native ATM connection

We displayed the status of PU, as shown in Figure 8-15 on page 115, by using the following command:

```
D NET,ID=F6APU1
```

```
NAME = F6APU1, TYPE = PU_T2.1 053
STATUS= ACTIV, DESIRED STATE= ACTIV
CP NAME = VTAM30, CP NETID = NETA, DYNAMIC LU = YES
XNETALS = YES
RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
F6APU1   AC/R    21 YES   98910000000000000000017100238080
HPR = RTP - OVERRIDE = N/A - CONNECTION = YES
LLERP = NOTPREF - RECEIVED = NOTALLOW
VPCI/VCI = 000021
LINE NAME = F6ALIN1, LINE GROUP = F6AGRP1, MAJNOD = XCAATMF6
I/O TRACE = OFF, BUFFER TRACE = OFF
STATE TRACE = OFF
VTAMTOPO = REPORT, NODE REPORTED - YES
MAJOR NODE VTAMTOPO = REPORT
```

Figure 8-15   Display of active PU for native ATM connection

Then we established a dial connection, as displayed in Figure 8-16, by using the following command:

```
V NET,DIAL,ID=SWF6PU1
```

```
CONNECTOUT ESTABLISHED FOR PU SWF6PU1  ON LINE F6ALIN2
APPN CONNECTION FOR NETA.VTAM30 IS ACTIVE - TGN = 22
VARY DIAL COMMAND COMPLETE FOR SWF6PU1
```

Figure 8-16   Display of messages issued during dial connection

Finally, we displayed the status of the switched node that just got connected, as shown in Figure 8-17, by using the following command:

```
D NET,ID=SWF6PU1,E
```

```
NAME = SWF6PU1, TYPE = PU_T2.1 230
STATUS= ACTIV, DESIRED STATE= ACTIV
CP NAME = VTAM30, CP NETID = NETA, DYNAMIC LU = YES
XNETALS = YES
RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
SWF6PU1  AC/R    22 YES   989100000000000000000017100058080
HPR = RTP - OVERRIDE = N/A - CONNECTION = YES
LLERP = NOTPREF - RECEIVED = NOTALLOW
ATM ADDRESS                                 TYPE      FORMAT
3999999999999999000099990C010020357CD01930    LOCAL     NSAP
3999999999999999000099990C010204357A09B560    REMOTE    NSAP
VPCI/VCI = 000027
SWITCHED SNA MAJOR NODE = SWATMF6
LINE NAME = F6ALIN2, LINE GROUP = F6AGRP2, MAJNOD = XCAATMF6
TRACE = OFF, BUFFER TRACE = OFF
STATE TRACE = OFF
```

Figure 8-17   Display of active switched PU for native ATM connection

9

ATM LAN Emulation (non-QDIO mode)

In this chapter, we cover the implementation steps needed to set up an OSA-Express 155 ATM feature in LANE non-QDIO mode. An OSA-Express 155 ATM can run in TCP/IP Passthru, or in SNA mode, or in both simultaneously.

OSA-2 note: The procedures in this chapter deal with configuring an OSA-Express feature. However, these same procedures can also be used to configure the OSA-2 version of this feature (with the exception of the HCD definition, CHPID type OSA).

© Copyright IBM Corp. 2003. All rights reserved.

9.1 Configuration Information

In our environment, we created a configuration to run in LANE mode with two partitions, both using TCP/IP and SNA. Figure 9-1 shows the logical connectivity. We guide you through the steps required to:

- Create a new OAT table using the GUI interface
- Create TCP/IP definitions to support this configuration
- Create VTAM SNA definitions to support this configuration

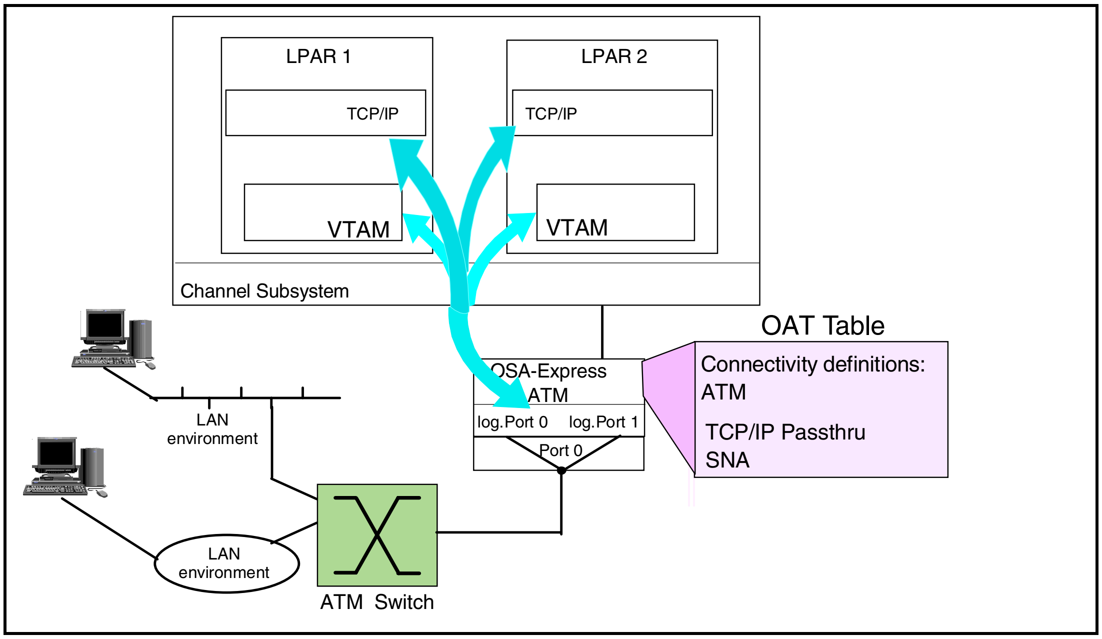

Figure 9-1 ATM LANE mixed mode (shared port)

9.2 HCD definition requirements

The OSA CHPID, the control unit, and the OSA devices must be defined to HCD, and the HCD definitions must be activated. Refer to Chapter 3, “Hardware configuration definitions” on page 37 for the procedure to create the definitions. Example definitions for the IOCDS used in this chapter are shown in Example 9-1.

Example 9-1 IOCDS input for CHPID F6

```
.
CHPID PATH=(F6),SHARED,PARTITION=((A1,A2),(A1,A2)),TYPE=OSE
.
CNTLUNIT CUNUMBR=0E90,PATH=(F6),UNIT=OSA
.
IODEVICE ADDRESS=(09E0,15),CUNUMBR=(09E0),UNIT=OSA,UNITADD=00
IODEVICE ADDRESS=09EF,UNITADD=FE,CUNUMBR=(09E0),UNIT=OSAD
```

Note: If you are implementing Ethernet LAN emulation (ENET LANE) only, the ATM CHPID can be defined as OSD. If you want to use SNA and/or Token Ring LAN emulation (TR LANE), you must define the ATM CHPID as OSE.

9.3 Create and save the OSA configuration with OSA/SF GUI

OSA/SF is required to configure the OSA-Express 155 ATM feature.

Figure 9-2 shows the configuration we used for our ATM LANE environment. We have two logical partitions, S30 and S35, each sharing the OSA-Express 155 ATM feature. The OSA is configured to allow each LPAR to run in TCP/IP and SNA mode.

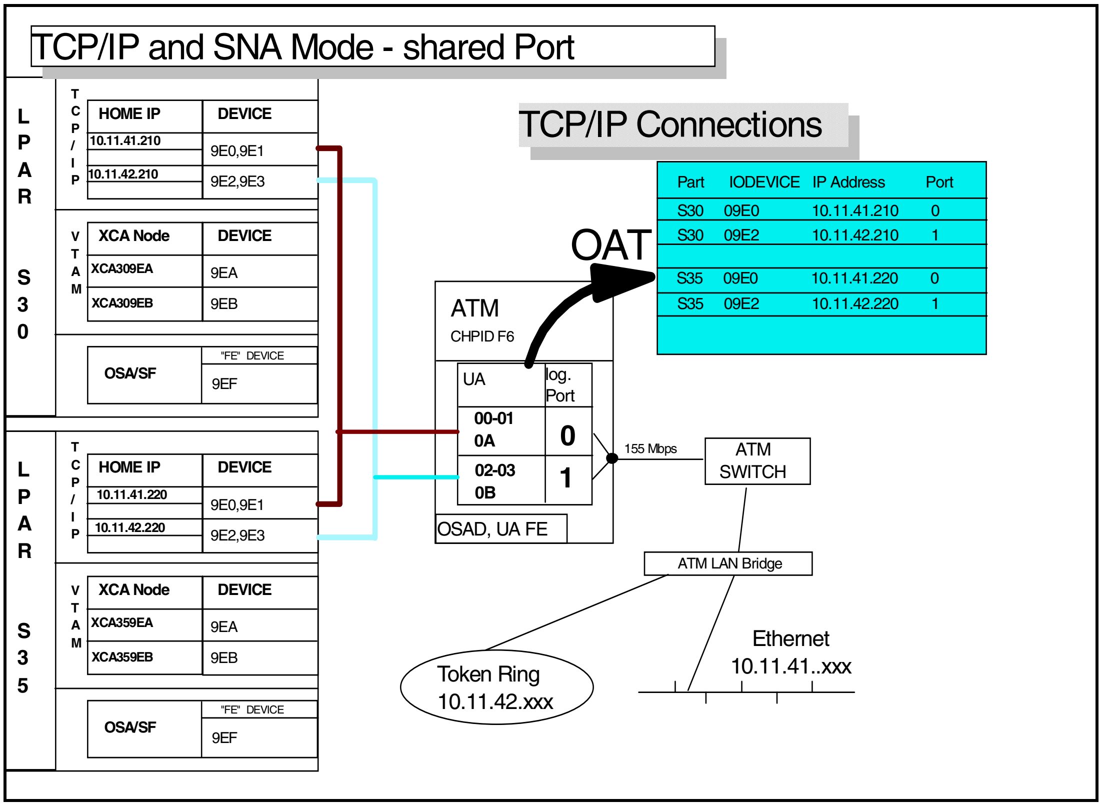

Figure 9-2 Devices in ATM LANE mixed mode

Creating (adding) and saving an OSA feature configuration is not disruptive. The OSA/SF GUI code must be connected to an z/OS host which has OSA/SF running. This step does *not* require the CHPID to be installed or online to create and save an OSA feature configuration.

1. Start the OSA/SF GUI program:
   – Double-click the **OSA/SF** folder icon.
   – Double-click the **OSA/SF** program icon.
2. Start the OSA/SF host connection by double-clicking the icon previously created to connect your workstation to the host OSA/SF and logging on.

   The OSA CHPID panel will be displayed, as shown in Figure 9-3 on page 120.

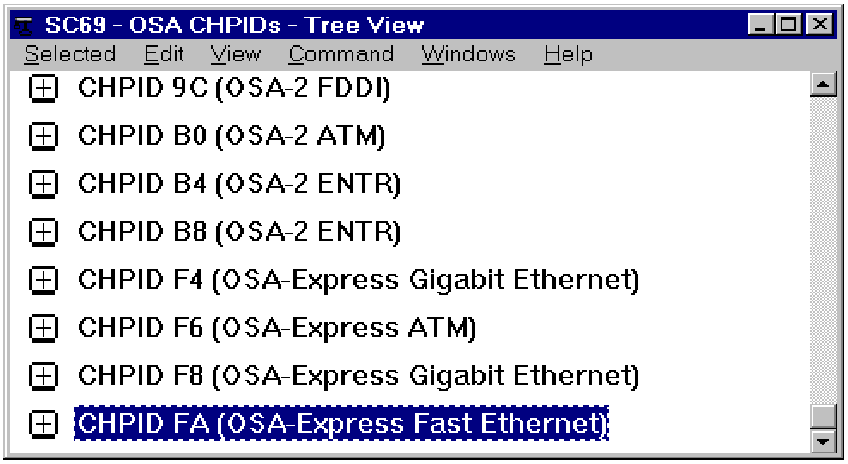

Figure 9-3 OSA CHPIDs Tree View

3. Select the OSA CHPID that represents an OSA port you are going to be configuring (CHPID F6 in our case).
4. Click **Selected** from the menu bar.
5. Select **Configurations....** from the Selected pull-down menu.
6. Select **Configurations list...** from the cascade menu.

> **Note:** If the OSA CHPID has not been defined to the server's channel subsystem (that is, not in the list), or if the OSA feature has not been installed, use the Planning option instead of the Configurations option.

7. The Configuration list for CHPID F6 is displayed.
8. Press **Add...** to create a new configuration.

   The Configuration panel is then displayed as shown in Figure 9-4 on page 121; do the following:

   a. Enter a configuration name in the Configuration name entry field. This is the name that will be displayed in the configuration list panel.
   b. Click **Yes** to make sure that the physical port changes will be used. (If you are configuring for partial activation on an OSA-Express 155 ATM feature with ports previously configured, you can choose **No**.)
   c. Optionally, you can enter configuration information in the Port description field. This field has no effect on the operation of the OSA, but a short description of what this physical port is designed to achieve can be useful later.
   d. A port name is required and can be up to 8 characters in length. Like the User data field, this is not used by OSA specifically. The value here should not duplicate other port names in your network.
   e. The local end system ID (ESI) field is used to override the default of the universal system ID, which was IBM-supplied for this ATM port, with a locally administered system ID.
   f. For the other fields of the physical port configuration panel, check the documentation of the ATM switch to which this ATM OSA-Express is attached to ensure the correct settings are being used.

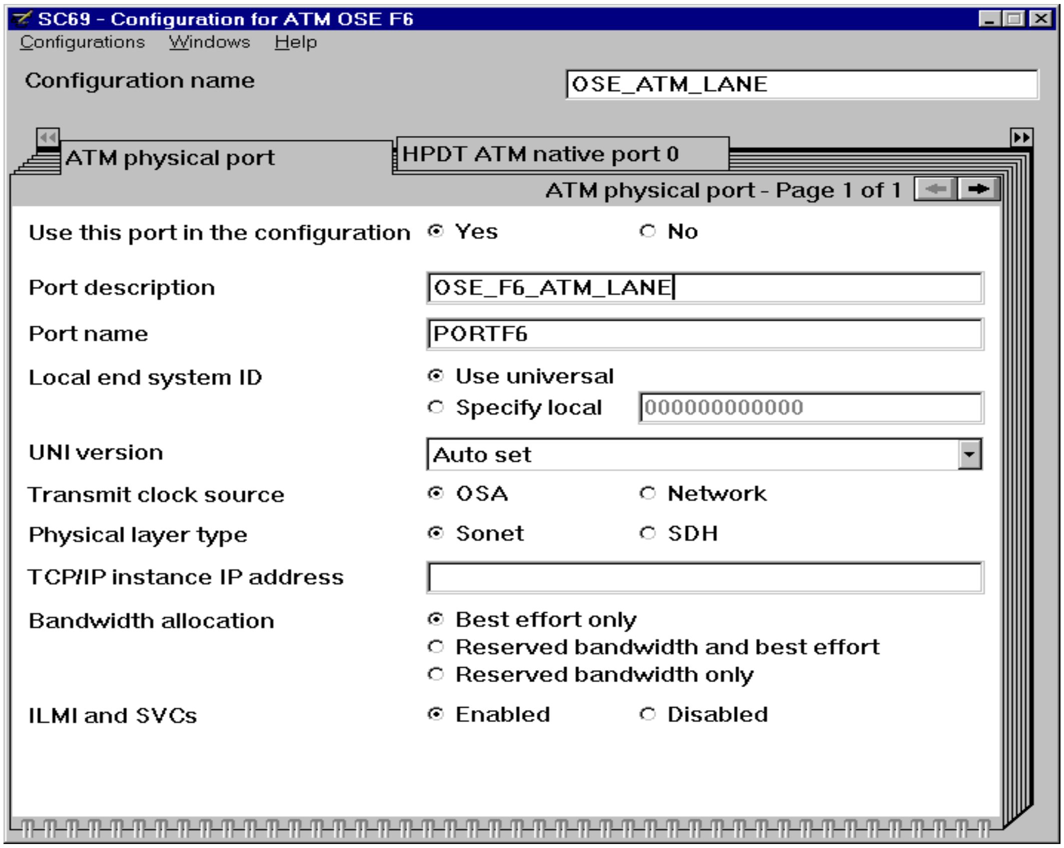

Figure 9-4 Configuration panel (physical port)

9. Click the right double arrow in the upper right corner of the notebook until ATM LEC port 0 appears on top of the workbook. Then select that tab. (The HPDT ATM native port 0 tab is skipped because it is not required for this configuration).
10. The panel ATM LEC port 0 (Implementation values) - page 1 of 5 is displayed, as shown in Figure 9-5 on page 122.

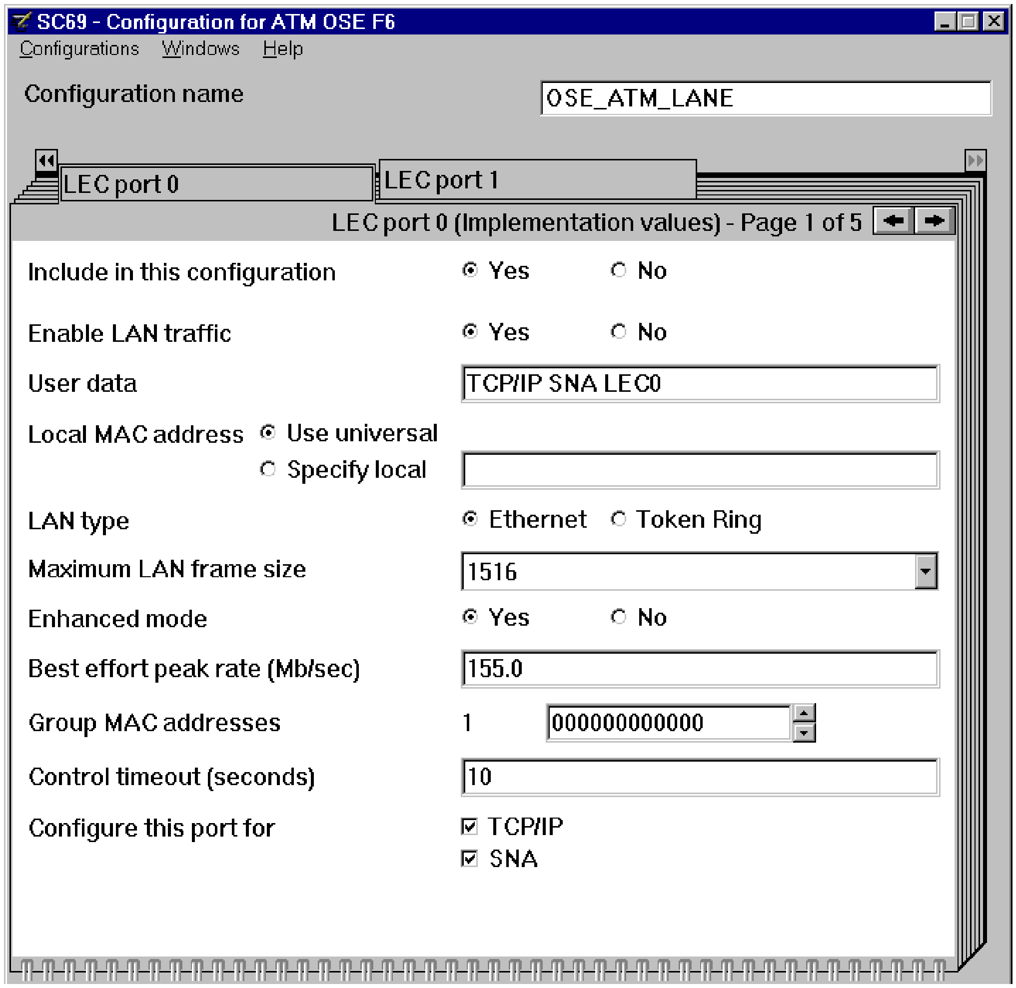

Figure 9-5 LEC port 0 (implementation values)-Page 1 of 5

a. Select **Yes** at the Include in this configuration option and at the Enable LAN traffic option.

b. Optionally, you can enter information in the User data field. This field has no effect on the operation of the OSA, but a short description of what this logical port is designed to achieve can be useful later.

c. The Local MAC address field is used to override the default of the universal MAC address, which was IBM-supplied for this logical port, with a locally administered MAC address.

d. Select the type of LAN to which this port will connect.

e. Specify *one* of the following values for any of the emulated LAN types:

- 1516 bytes, which is the default for an Ethernet LAN.
- 4544 bytes, which is the default for a 4Mbps token-ring LAN.
- 9234 bytes, which is a standard in LAN Emulation Over ATM, which is published by the ATM Forum.
- 8190 bytes, which is the default for a 16 Mbps token-ring LAN.

- For the Enhanced mode choice, select **No** if you want the LEC port to drop its data direct connections to other LAN emulation clients (LECs) if and when it loses its connection with its LAN emulation server (LES).

  Otherwise, select **Yes** if you want this LEC port to keep its data direct connections.

f. Check the documentation of the ATM switch to which this ATM OSA-Express is attached, to ensure the correct Best Effort peak rate is being used. The IBM-supplied default value is 155.0 Mbps, which is the maximum speed supported by an ATM OSA-Express. Make sure that the peak cell rate is acceptable to the ATM switch and to the clients on the LAN to which the ATM OSA-Express is attached.

g. Enter the MAC address of any destination group for which you want this port to receive frames. If you don't want to use group MAC addresses, leave all values set to zeros.

h. Select the emulation mode for this port as TCP/IP, SNA, or both.

i. After you have filled in all desired entries, click the right arrow in the upper right corner of the notebook to change to page 2, as shown in Figure 9-6.

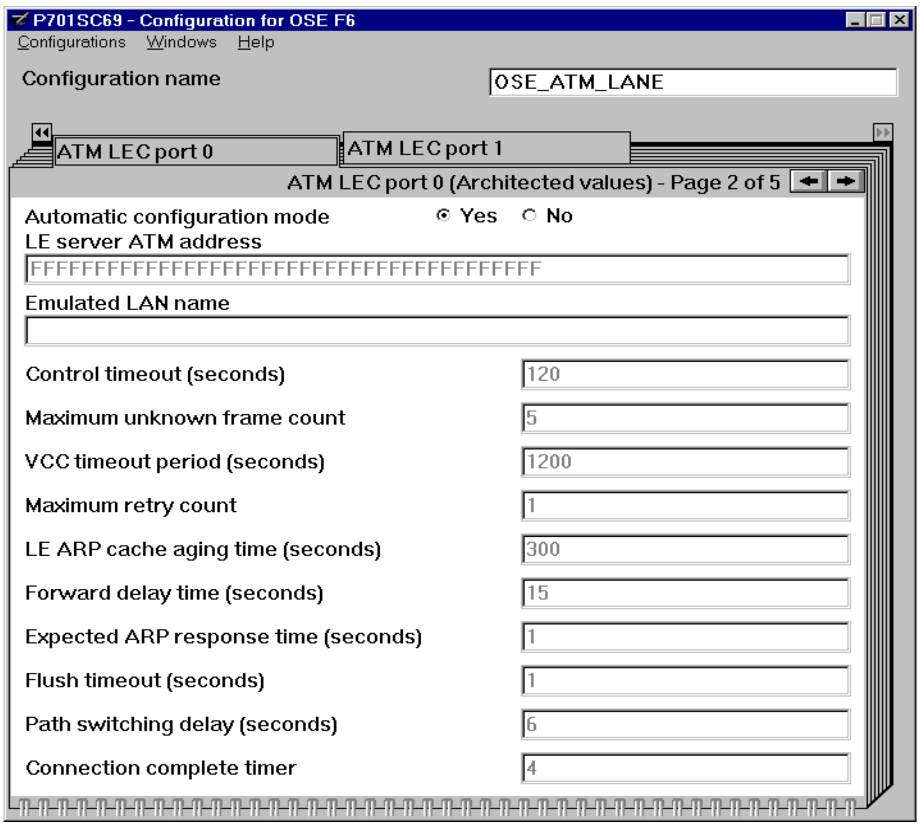

Figure 9-6 ATM LEC port 0 architected values - page 2 of 5

11. Check the documentation of the ATM switch to which this ATM OSA-Express is attached, to ensure the correct settings are being used.

Note: We used the automatic configuration of the LEC port by the LAN emulation configuration server (LECS).

If you do not have a LECS, you must provide the 20 byte ATM address of the LAN Emulation Server (LES).

12. After you have filled in all desired entries, click the right arrow in the upper right corner of the notebook to change to page 3 and click the **Add** button to define TCP/IP Passthru devices. If your ATM feature works only in SNA mode, you can skip this configuration panel.

13. The TCP/IP Passthru OAT entry panel, as shown in Figure 9-7, lets you add the IP addresses related to unit addresses and/or logical partitions.

> **Note:** You can fill in multiple IP addresses, as well as specify the entry used as primary or secondary. For further information about this IP availability function, see "Enhanced IP network availability (IPA)" on page 7.

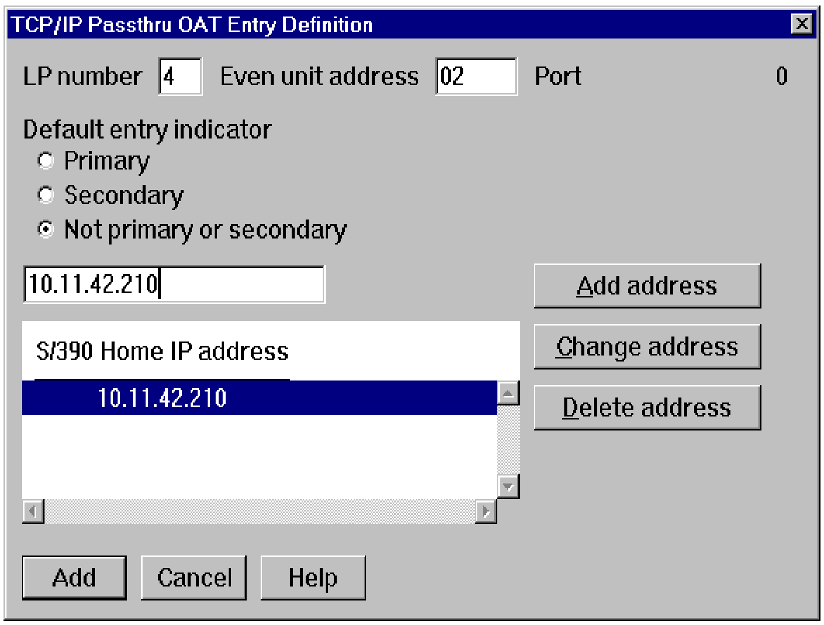

Figure 9-7 TCP/IP passthru OAT entry definition

Reminder: If you are using a CHPID dedicated to one partition, the LP number must be 0.

a. After you have finished the Home IP address list entries for the LP number and unit address combination, click **Add**.This stores the record information to the OAT.

b. After you have stored all your necessary IP addresses, click **Cancel.** This brings you back to the TCP/IP OAT entries, page 3, where you can recheck all TCP/IP OAT entries. For an example, refer to our TCP/IP OAT configuration for the logical port 0 in Figure 9-8 on page 125.

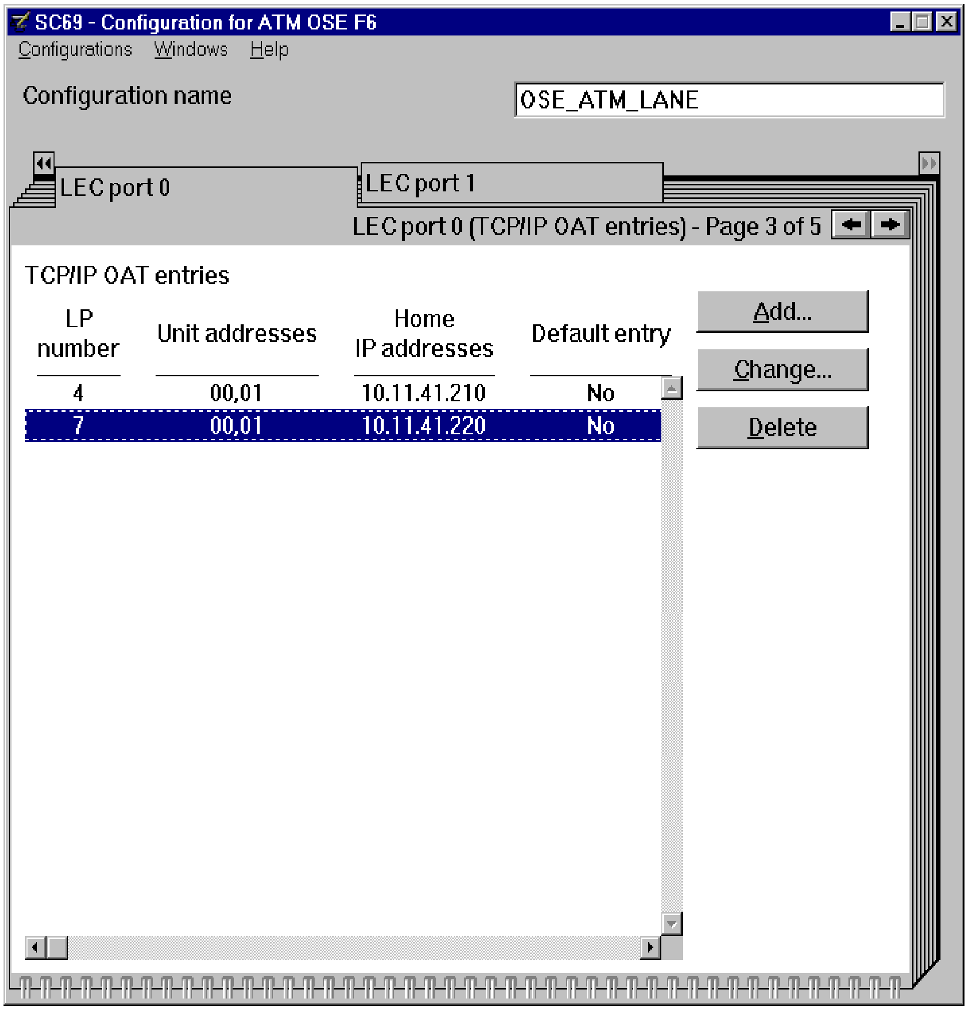

Figure 9-8 ATM LEC port 0 (TCP/IP OAT entries) - page 3 of 5

14. After checking the TCP/IP entries, click the right arrow in the upper right corner of the notebook to change to page 4 (SNA values), which shows the SNA values for Ethernet or token ring set on page 2. Here you are able to define the SNA timer and counter setting.

    Figure 9-9 on page 126 shows the timer and counter slider and the additional enhanced SNA availability feature. The enhanced SNA availability feature is only available for token ring networks. For more information about the SNA enhanced availability, refer to 1.4.2, “SNA enhanced availability” on page 20.

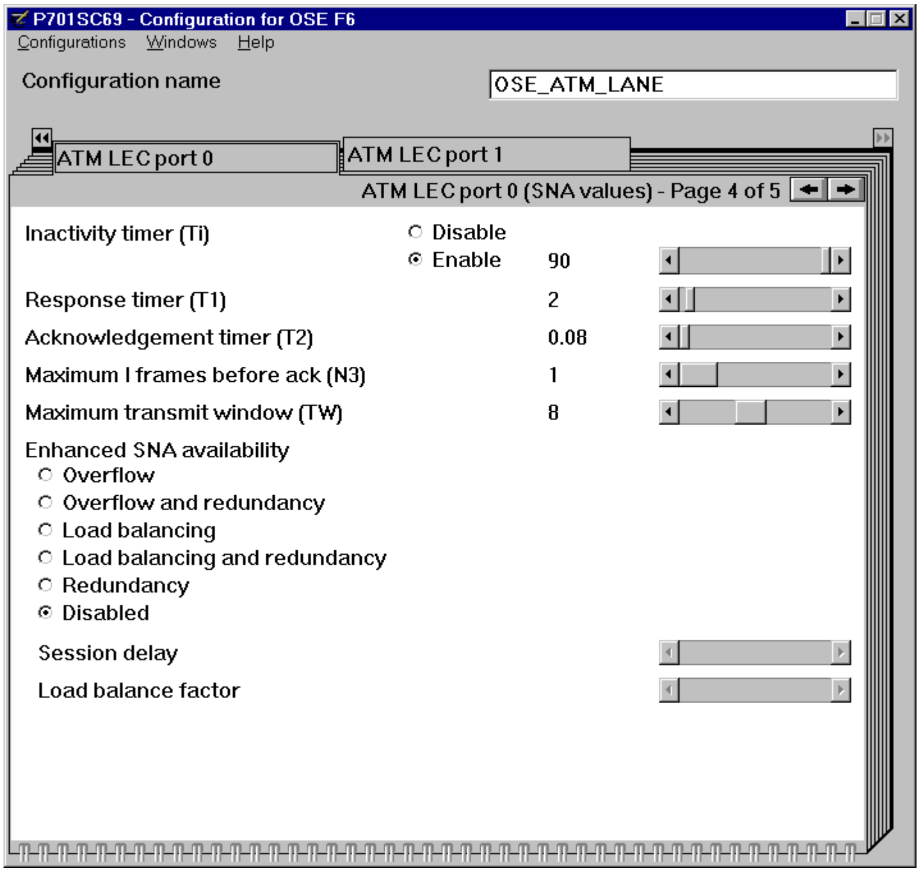

Figure 9-9 ATM LEC port 0 SNA values - page 4 of 5 (set for token ring)

15. After you have filled in all desired entries, click the right arrow in the upper right corner of the notebook to go to page 5. (Figure 9-10 shows an example of the filled-in SNA OAT for logical port 0.)

    Clicking **Add** in page 5 brings up an SNA OAT entry definition panel, where you can fill in your logical partition and unit address combination for that port.

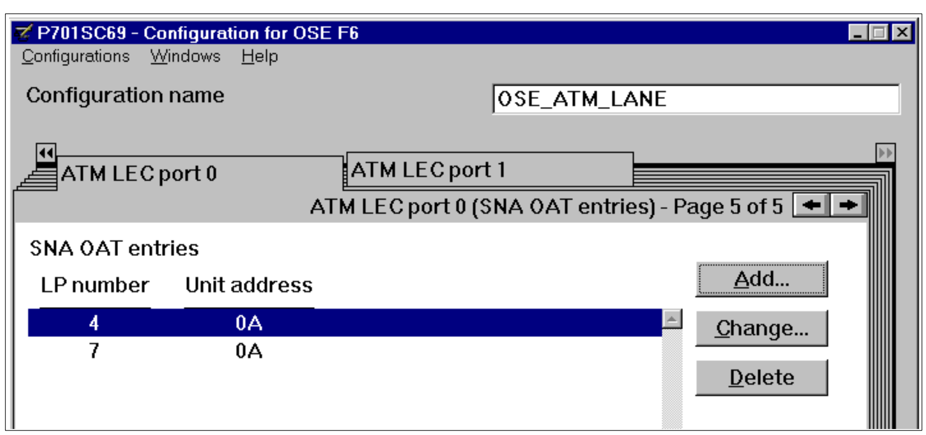

Figure 9-10 ATM LEC port 0 SNA OAT entries - page 5 of 5

16. After adding all your required entries for the SNA mode, you have finished all required definitions for logical port 0.

The arrow button brings you to the starting panel for the definitions for logical port 1. At this point, if you want to configure both logical ports, you have to repeat all steps starting from Step 9 on page 121.

We will not go through all the steps again. In Figure 9-11, and Figure 9-12, you can see the differences between logical port 0 and 1 that are in the OATs.

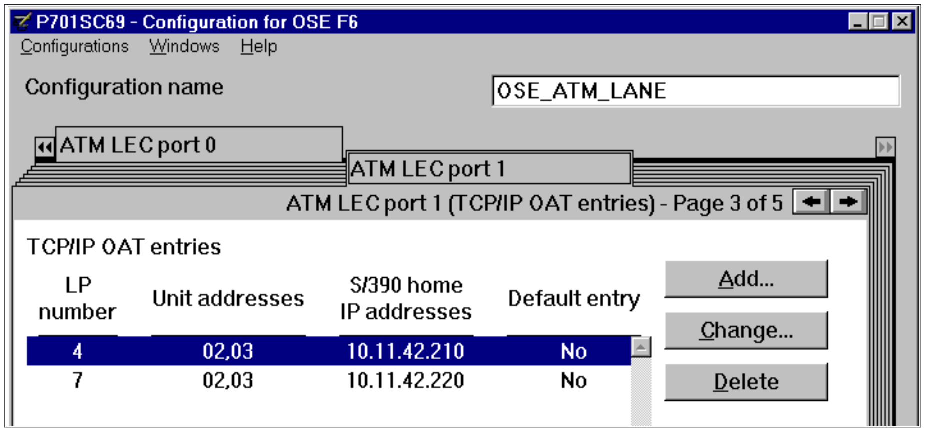

Figure 9-11 ATM LEC port 1 TCP/IP OAT entries - page 3 of 5

Keep in mind that the logical port 1 is now configured for token ring connectivity.

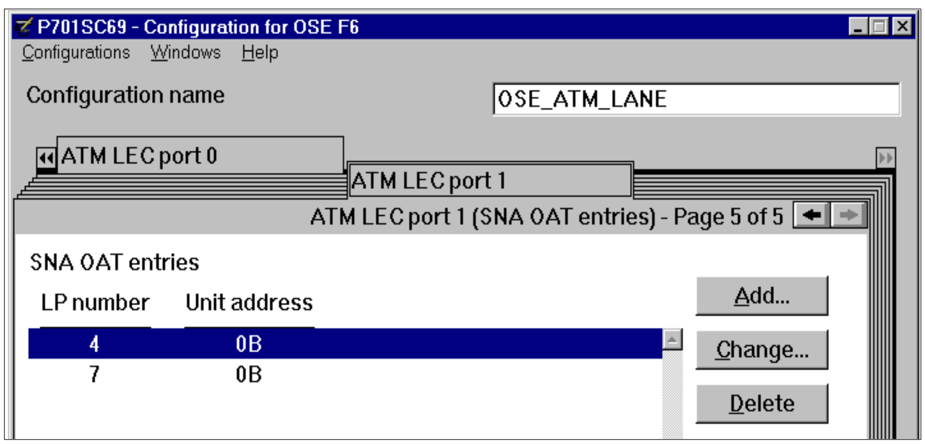

Figure 9-12 ATM LEC port 1 SNA OAT entries - page 5 of 5

9.3.1 Activate the OSA configuration

1. From the pull-down menu, select **Configuration -> Save**. This will ensure that the configuration is saved. Once saved, you should see the message shown in Figure 9-13.

Figure 9-13 Message to indicate configuration saved successfully

2. From the pull-down menu, select **Configurations,** followed by **Activate** or **Activate(no install)**.

   Activate will download the configuration settings and make it usable immediately. (Keep in mind that this is a *disruptive* function.) To be sure that the function is successful, you have to deallocate all devices at the CHPID level from applications in all partitions (except the OSAD device).

   Activate(no install) prevents disrupting an OSA that is already running with a different configuration. You can defer the install to a more appropriate time.

3. To complete the install, do the following:

   a. From the OSA/SF GUI window OSA CHPIDs - Tree View, select the OSA CHPID (which represents an OSA port) that you are going to install.

   b. Select **Command** from the menu bar.

   c. Select **Install.**

   d. Select **Force,** then **OK.**

9.4 Customizing the z/OS host network subsystem

In this mixed mode and shared port environment, TCP/IP and VTAM coexist and share the same logical OSA ports without affecting each other. This allows the definitions for TCP/IP and VTAM to be done independently, as if the OSA-Express 155 ATM logical ports were owned by VTAM or TCP/IP exclusively.

Port sharing is not set up in TCP/IP or VTAM, but rather in OSA/SF.

9.4.1 Customizing VTAM

This section describes the definitions required in VTAM to allow SNA applications to access the LAN connected to the host via an OSA-Express feature. To describe the VTAM setup, the network configuration shown in Figure 9-2 on page 119 is used.

In this example, both VTAMs in S30 and S35 are communicating over an Ethernet LAN via logical port 0 (device number 9EA) and over a token ring connection via logical port 1 (device number 9EB).

You need to define two types of major nodes in VTAM, as follows:

- External Communication Adapter major node (XCA)
- Switched major node

XCA major node

Define one XCA major node for each SNA OSA device for the following characteristics:

- The node type (VBUILD definition statement)
- The port used by the LAN (port definition statement)
- The switched peripheral nodes (type 2) attached to a token ring or an Ethernet LAN through an OSA port (GROUP, LINE, and PU definition statements)

Note: The current OSA-Express features support 4096 SNA PU Type 2 connections per port.

For the configuration example in Figure 9-2 on page 119, the XCA major node (XCAES30) definition in the VTAM for system S30 is for an Ethernet connection via logical port 0 (device number 9EA), and XCATS30 represents the token-ring connection via logical port 1 of the OSA-Express. The actual VTAM coding to implement these connections is shown in Example 9-2 and Example 9-3.

Example 9-2 XCA definition port 0 for system S30

```
**********************************************************************
*                                                                    *
* XCA MAJNODE for ATM LANE OSA EXPRESS Log.PORT 0 on System S30      *
*                                                                    *
**********************************************************************
XCAES30  VBUILD TYPE=XCA
OSE9EAPE PORT   MEDIUM=CSMACD,                                        X
                ADAPNO=0,                                             X
                CUADDR=9EA,                                           X
                TIMER=60,                                             X
                SAPADDR=04
**********************************************************************
*                                                                    *
**********************************************************************
OSE9EAGE GROUP  DIAL=YES,                                             X
                DYNPU=YES,                                            X
                ANSWER=ON,                                            X
                AUTOGEN=(3,L,P),                                      X
                CALL=INOUT,                                           X
                ISTATUS=ACTIVE
******************************** Bottom of Data **********************
```

Example 9-3 XCA definition port 1 for system S30

```
**********************************************************************
*                                                                    *
* XCA MAJNODE for ATM LANE OSA EXPRESS Log.PORT 1 on System S30      *
*                                                                    *
**********************************************************************
XCAES30  VBUILD TYPE=XCA
OSE9EAPE PORT   MEDIUM=RING,                                          X
                ADAPNO=1,                                             X
                CUADDR=9EB,                                           X
                TIMER=60,                                             X
                SAPADDR=04
**********************************************************************
*                                                                    *
**********************************************************************
OSE9EAGE GROUP  DIAL=YES,                                             X
                DYNPU=YES,                                            X
                ANSWER=ON,                                            X
                AUTOGEN=(3,L,P),                                      X
                CALL=INOUT,                                           X
                ISTATUS=ACTIVE
******************************** Bottom of Data **********************
```

Table 9-1 on page 130 provides a brief overview of the Port parameters.

Table 9-1 XCA major node Port definition for XCAES30

Required parameters	Explanation	Remarks
TYPE=XCA	XCA major node	Each OSA port functions like an XCA to VTAM.
ADAPNO=0	*PORT* statement	OSE9APE is the name of port 0 of the OSA feature. Code `ADAPNO=0 for` logical port 0 of OSA-Express 155 ATM.
CUADDR=E9A	Channel unit address	Code the device number defined for this port (port 0). In this example, S30 uses device number E9A for port 0.
MEDIUM= CSMACD	LAN type	Use `RING` for Token Ring, Use `CSMACD` for Ethernet
SAPADDR=4	Service access point address	Code a value that is a *multiple of 4*. This address must be *unique* for each VTAM communicating with a port. Use different SAP addresses if a port is shared by multiple VTAMs. See the SAPADDR value of XCAES35 in Example 9-4 on page 131.

Table 9-2 provides a brief overview of the significant Group parameters.

Table 9-2 XCA major node Group definition for XCAES30

Required parameters	Explanation	Remarks
DIAL=YES	Switched peripheral node	OSE9EAGE is the minor node name of the line group. You *must* code `DIAL=YES` to specify that the switched line control protocol is required.
AUTOGEN= (3,L,P)	Autogeneration of LINE and PU statements	This parameter enables VTAM to generate automatically three sets of LINE and PU statements. The LINE names will begin with L. The PU names will begin with P.

The definitions for XCA major nodes XCAES35 and XCATS35 in VTAM for system S35 via both logical ports of the OSA-Express are given in Example 9-4 and Example 9-5 on page 131.

Example 9-4 XCA definition port 0 for system S35

```
*********************************************************************
*                                                                   *
* XCA MAJNODE for ATM LANE OSA EXPRESS Log.PORT 0 on System S35     *
*                                                                   *
*********************************************************************
XCAES35  VBUILD TYPE=XCA
OSE9EAPE PORT   MEDIUM=CSMACD,                                       X
                ADAPNO=0,                                            X
                CUADDR=9EA,                                          X
                TIMER=60,                                            X
                SAPADDR=08
*********************************************************************
*                                                                   *
*********************************************************************
OSE9EAGE GROUP  DIAL=YES,                                            X
                DYNPU=YES,                                           X
                ANSWER=ON,                                           X
                AUTOGEN=(3,L,P),                                     X
                CALL=INOUT,                                          X
                ISTATUS=INACTIVE
******************************** Bottom of Data **********************
```

Example 9-5 XCA definition port 1 for system S35

```
*********************************************************************
*                                                                   *
* XCA MAJNODE for ATM LANE OSA EXPRESS Log.PORT 1 on System S35     *
*                                                                   *
*********************************************************************
XCAES35  VBUILD TYPE=XCA
OSE9EAPE PORT   MEDIUM=RING,                                         X
                ADAPNO=1,                                            X
                CUADDR=9EB,                                          X
                TIMER=60,                                            X
                SAPADDR=08
*********************************************************************
*                                                                   *
*********************************************************************
OSE9EAGE GROUP  DIAL=YES,                                            X
                DYNPU=YES,                                           X
                ANSWER=ON,                                           X
                AUTOGEN=(3,L,P),                                     X
                CALL=INOUT,                                          X
                ISTATUS=INACTIVE
******************************** Bottom of Data **********************
```

Table 9-3 shows the only parameter that is different between system S30 and S35.

Table 9-3 XCA major node Port definition for XCAES35

Parameters	Explanation	Remarks
SAPADDR=8	Service access point address	A different SAP address of 8 is used because the two VTAMs are sharing the same port 0. See the SAPADDR value of XCAES30 in Example 9-2 on page 129

Switched major node

Define *one* switched major node for the switched connections to the peripheral nodes attached to each LAN. You need to code the following:

- A remote physical unit (*PU* definition statement)
- The corresponding logical units (*LU* definition statement)

Example 9-6 shows how 3270 sessions could be set up with a switched major node (SWNS301) for S30.

Example 9-6 Switched major node for S30

```
* THIS SWITCHED MAJNODE DEFINES THE OSA-EXPRESS SNA CONNECTION
* TO VTAM RUNNING IN LPAR S30
*
SWNS301  VBUILD TYPE=SWNET
OSAS301  PU    ADDR=02,                                          X
               IDBLK=05D,                                        X
               IDNUM=12863,                                      X
               CPNAME=OSANT1                                     X
               IRETRY=YES,                                       X
               MAXOUT=7,                                         X
               MAXPATH=1,                                        X
               MAXDATA=1024,                                     X
               PACING=0,                                         X
               VPACING=0,                                        X
               PUTYPE=2,                                         X
               DISCNT=(NO),                                      X
               ISTATUS=ACTIVE,                                   X
               MODETAB=NEWMTAB,                                  X
               DLOGMOD=DYNTRN,                                   X
               USSTAB=USSLDYN,                                   X
               SSCPFM=USSSCS
OSA30L0  LU     LOCADDR=0,MODETAB=MTAPPC,DLOGMOD=APPCMODE
OSA30L1  LU     LOCADDR=1                                 3270 SESSION
OSA30L2  LU     LOCADDR=2                                 3270 SESSION
```

Table 9-4 presents the important parameters in the PU definition for SWNS301.

Table 9-4 Switched major node PU definition

Parameters	Explanation	Remarks
TYPE=SWNET	Switched major node	
ADDR=02	PU's SDLC station address	OSAS301 is the PU name. ADDR is a required parameter.
PUTYPE=2	PU type	The PUTYPE must be 2 for a LAN switched station. Type 2 also denotes a type 2.1 PU.
CPNAME=OSANT1	Control point (CP) name of a type 2.1 node	To dial in, a type 2.1 peripheral node on a switched line requires either CPNAME or both IDBLK and IDNUM. However, you can code all three operands.

Table 9-5 on page 133 presents the important parameters in the LU definition for SWNS301.

Table 9-5 Switched major node LU definition

Parameters	Explanation	Remarks
LOCADDR=0	LU's local address at the PU	**LOCADDR=0 denotes an independent LU.**
MODETAB=MTAPPC	Logon mode table	Code a separate logon mode table for APPC.
LOCADDR=2		LOCADDR=2 denotes a dependent LU.

For connection from S35 to the OSANT1 workstation, the parameters for the switched major node are exactly the same as the parameters in the switched major node SWNS301. Of course, you should change the names for PUs and LUs to make them more meaningful.

For connections to other PUs, you just have to create additional entries reflecting their CPNAMEs or IDBLOCK/IDNUM combinations.

9.4.2 TCP/IP definitions

TCP/IP uses the OSA ports as LAN Channel Station (LCS) devices. For OSA, you define only one LINK statement for the associated DEVICE statement. Other products may define and associate more than one LINK statement to refer to the same device statement. For OSA, each port has its own unique DEVICE and LINK statement defined in the TCP/IP profile.

Figure 9-2 on page 119 shows the network and the connections for our configuration example.

The following definitions are based on one OSA-Express 155 ATM feature, using CHPID F6.

It has logical port 0 connected to an Ethernet LAN and will be used by LPAR S30 as TCP/IP address 10.11.41.210, and by LPAR S35 as TCP/IP address 10.11.41.220.

Logical port 1 is connected to a token-ring LAN and will be used by LPAR S30 as TCP/IP address 10.11.42.210, and by LPAR S35 as TCP/IP address 10.11.42.220.

In each partition, you have to do the following:

1. Define one DEVICE statement per OSA port. Use the even device number of the two device numbers assigned in the hardware to the port.
   - You will have to define two device numbers per OSA port for TCP/IP mode in HCD since TCP/IP runs in full duplex mode. One device is used by TCP/IP for reading, the other for writing.
   - Using the DEVICE statement, you define the DEVICE statement name, the DEVICE type (LCS) for the OSA port, and the DEVICE number (the read device number, which is the even device number).
2. Define one LINK statement per OSA TCP/IP DEVICE statement.

   Using the LINK statement, you define the LINK name, the LINK type, the PORT number, and the DEVICE statement name.

   Note: Although the OSA port is mainly addressed by the device number, the port number in the LINK statement must match the actual OSA port number.
3. Define the HOME IP address of the OSA port.

   Using the HOME statement, define an IP address referring to a LINK statement name that itself refers to a DEVICE statement name.

4. Define static routes through the BEGINROUTES statement.

The BEGINROUTES statement

If you are migrating your TCP/IP profile from an earlier release, it may use the GATEWAY statement to define static routes instead of the BEGINROUTES - ENDROUTES statements. Gateway will be recognized and used, but you should consider replacing it with BEGINROUTES.

Because it is compatible with UNIX standards, easier to code than GATEWAY, accepts both IPv4 and IPv6 addresses, and has enhanced functionality, BEGINROUTES is the recommended method for defining static routes. Future static route enhancements will only be available with the BEGINROUTES statement.

Dynamic routing can be accomplished using the OMPROUTE daemon or the BSDROUTINGPARMS statement can be used in conjunction with RouteD daemon.

Note: Use of the BEGINROUTES statement (static routes) with the OMPROUTE or OROUTED routing daemons is not recommended.

5. Define one START command per OSA DEVICE.

   After defining an OSA device in the TCP/IP profile, the device still has to be started explicitly. It uses the DEVICE statement name.

Example 9-7 shows the TCP/IP PROFILE definitions needed to define OSA to TCP/IP for LPAR S30.

Example 9-7 TCP/IP profile definitions for LPAR S30

```
; TCP/IP S30 PROFILE.TCPIP --- Hardware definitions
; OSA9E0 is an OSA-Express feature using logical port 0.  Devices are
; 9E0-9E1.
; OSA9E2 is an OSA-Express feature using logical port 1.  Devices are
; 9E0-9E1.
DEVICE OSA9E0      LCS           9E0
DEVICE OSA9E2      LCS           9E2

LINK      OSAL9E0  ETHEROR802.3     0    OSA9E0
LINK      OSAL9E2   IBMTR           1    OSA9E2

HOME
10.11.41.210     OSAL9E0
10.11.42.210     OSAL9E2

BEGINROUTES
 10.11.41.0/24       =    OSAL9E0  MTU 1500
 10.11.42.0/24       =    OSAL9E2  MTU 1500
ENDROUTES

;Start all the defined devices.
START    OSA9E0
START    OSA9E2
```

The configuration for LPAR S35 is exactly the same with one exception: the TCP/IP profile for LPAR S35 has the HOME statements coded to reflect the different IP address, as shown in Example 9-8 on page 135.

Example 9-8 TCP/IP Profile HOME Statement for S35

```
HOME
10.11.41.220   OSAL9E0
10.11.42.220   OSAL9E0
```

9.5 Activation

Once all the definitions have been added to OSA/SF, VTAM, and TCP/IP, we can activate the configuration.

Activation may require several tasks, such as:

- Verifying the devices are online
- Activating an OSA/SF configuration
- Activating VTAM resources
- Activating TCP/IP

9.5.1 Verify devices are online

The z/OS console display command can verify that the required devices are online.

```
D U,,,9E0,4
IEE457I 21.29.17 UNIT STATUS 836
UNIT TYPE STATUS        VOLSER     VOLSTATE
09E0 OSA  O
09E1 OSA  O
09E2 OSA  O
09E3 OSA  O

D U,,,9EA,2
IEE457I 21.32.17 UNIT STATUS 936
UNIT TYPE STATUS        VOLSER     VOLSTATE
09EA OSA  O
09EB OSA  O
```

Figure 9-14 Results of z/OS console display command

If they are not online, use the z/OS console vary command.

```
V (9E0-9E3,9EA,9EB),ONLINE
```

9.5.2 OSA/SF activation

If the configuration is not activated yet, follow the next procedure. Otherwise, you can skip the OSA-Express activation.

Choose **Selected -> Configuration -> Configuration List** (see Figure 9-15 on page 136).

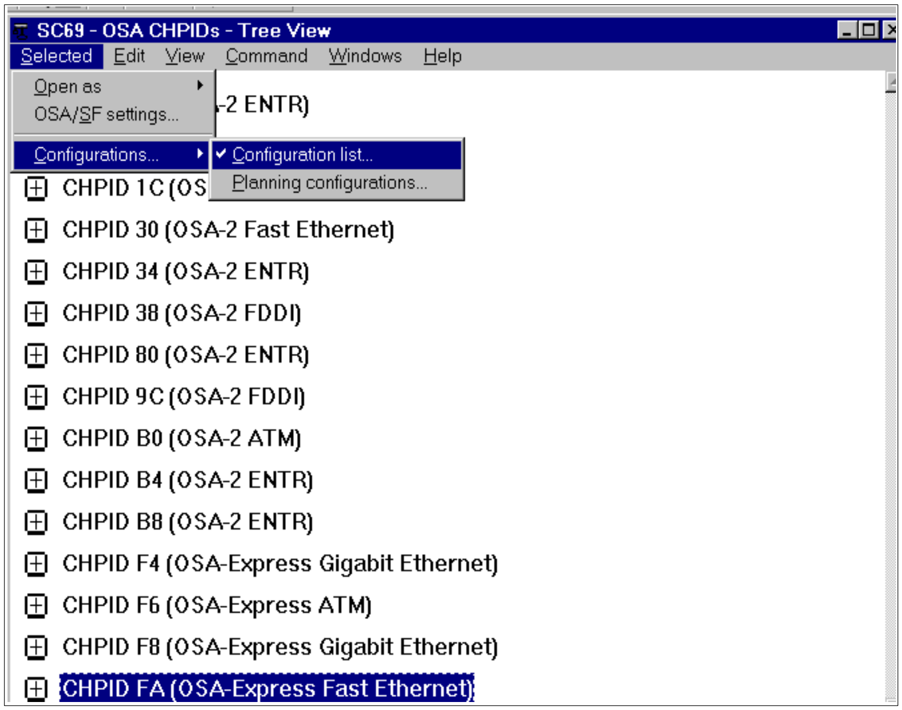

Figure 9-15 Selecting a configuration

From the list of configurations, highlight your saved configuration and press **Change**. This will retrieve the OSA configuration. You can now choose to either Activate or Activate(no install) from the Configurations pull-down.

The difference between these two choices is the timing for the activation. If there are no active sessions using this OSA-Express, the Activate command will load the new configuration down to the OSA-Express feature and recycle the hardware connections. This process can take up to three minutes.

If you select Activate(no install), the new configuration will be loaded to the OSA-Express feature, but will not be operational. When you use the no install option, you will need to return to the tree view window and select **Install** from the commands pull-down.

Note: The configuration defined the access from both LPARs, and the effect of the activation command is to make the OSA configuration active to all LPARs. You do not need to activate OSA-Express from more than one LPAR.

9.5.3 VTAM activation

The VTAM configuration is no different from any other VTAM device. Use the VTAM VARY command. On each LPAR, activate the XCA major node and the switched major node. Typically, the commands will be of the following format (for LPAR S30):

```
V NET,ID=XCAES30,ACT
V NET,ID=XCATS30,ACT
V NET,ID=SWNS301,ACT
```

9.5.4 TCP/IP activation

There are two ways to activate the TCP/IP devices: either restart the TCP/IP stack, or use the TCP/IP **VARY** command. We chose to restart the stack to implement the changes. Refer to Appendix A, “Commands” on page 175 for command syntax.

9.6 Relevant status displays

We monitored the status of the VTAM resources with the VTAM **Display** command. Figure 9-16 displays the XCA major node for the Ethernet connection (port 0).

VTAM information

```
D NET,ID=XCAES30,E
IST097I DISPLAY ACCEPTED
IST075I NAME = XCAES30, TYPE = XCA MAJOR NODE 715
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1021I MEDIUM=CSMA/CD, ADAPNO= 0,CUA=9EA, SNA SAP= 4
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST232I L9EA000 ACTIV
IST232I L9EA001 ACTIV
IST232I L9EA002 ACTIV
IST314I END
```

Figure 9-16 XCA major node logical port 0 display

Figure 9-17 shows the results from the XCA major node for logical port 1 used for the token-ring connection.

```
D NET,ID=XCATS30,E
IST097I DISPLAY ACCEPTED
IST075I NAME = XCATS30, TYPE = XCA MAJOR NODE 715
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1021I MEDIUM=RING, ADAPNO= 1,CUA=9EB,SNA SAP= 4
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST232I L9EB000 ACTIV
IST232I L9EB001 ACTIV
IST232I L9EB002 ACTIV
IST314I END
```

Figure 9-17 XCA major node logical port 1 display

Figure 9-18 on page 138 shows the Switched major node display.

```
D NET,ID=SWNS301,E
IST097I DISPLAY ACCEPTED
IST075I NAME = SWNS301, TYPE = SW SNA MAJ NODE 718
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST084I NETWORK RESOURCES:
IST089I OSAS301 TYPE = PU_T2           , CONCT
IST089I OSA30L1 TYPE = LOGICAL UNIT     , CONCT
IST089I OSA30L2 TYPE = LOGICAL UNIT     , CONCT
IST314I END
```

Figure 9-18 Switched major node display

The displays of the major nodes on system S35 gave us equivalent results.

TCP/IP displays

The TSO **NETSTAT DEV** command will display the TCP/IP devices. Figure 9-19 shows both the device and link in the READY state.

```
DEVNAME: OSA9E0              DEVTYPE: LCS       DEVNUM: 9E0
DEVSTATUS: READY
  LNKNAME: OSAL9E0          LNKTYPE: ETHOR       LNKSTATUS: READY
    NETNUM: 0   QUESIZE: 0
    BYTESIN: 0000206398      BYTESOUT: 0000001596
    BROADCASTCAPABILITY: YES
  BSD ROUTING PARAMETERS:
    MTU SIZE: 00000            METRIC: 00
    DESTADDR: 0.0.0.0          SUBNETMASK: 255.255.255.0
  MULTICAST SPECIFIC:
    MULTICAST CAPABILITY: YES
    GROUP            REFCNT
    -----            ------
    224.0.0.1        0000000001
DEVNAME: OSA9E2             DEVTYPE: LCS       DEVNUM: 9E2
DEVSTATUS: READY
  LNKNAME: OSAL9E2           LNKTYPE: TR          LNKSTATUS: READY
    NETNUM: 0   QUESIZE: 0
    BYTESIN: 0789352831      BYTESOUT: 0324162550
    ARPMACADDRESS: NON-CANONICAL   SRBRIDGINGCAPABILITY: YES
    BROADCASTCAPABILITY: YES       BROADCASTTYPE: ALL RINGS
  BSD ROUTING PARAMETERS:
    MTU SIZE: 00000            METRIC: 00
    DESTADDR: 0.0.0.0          SUBNETMASK: 255.255.255.0
  MULTICAST SPECIFIC:
    MULTICAST CAPABILITY: YES
    GROUP            REFCNT
    -----            ------
    224.0.0.1        0000000001
```

Figure 9-19 Display of the TCP/IP device and link

To verify that a connection existed, we did a ping from both TCP/IP stacks to the workstation.

For a summary of useful commands, refer to Appendix A, “Commands” on page 175.

10

Enterprise Extender (EE)

In this chapter, we introduce and discuss the importance of Enterprise Extender technology. We explain the concept and provide a sample implementation of the OSA-Express in QDIO mode using EE.

Enterprise Extender allows for the use of SNA transport protocols (namely APPN and HPR) over an Internet Protocol (IP) network. It enables the leveraging of IP-based infrastructure network components for use in delivering SNA traffic.

The Enterprise Extender (EE) function of CS for z/OS or OS/390 allows you to run SNA applications and data on IP networks and IP-attached clients. It can be used with any OSA-Express feature running IP traffic. EE is a simple set of extensions to the open High Performance Routing technology that integrates HPR frames into User Datagram Protocol/Internet Protocol (UDP/IP) packets.

© Copyright IBM Corp. 2003. All rights reserved.

10.1 The IP backbone as an APPN connection network

An important aspect of Enterprise Extender is the ability to view the IP network as an Advanced Peer-to-Peer Networking (APPN) connection network. In this case, the benefit comes from the ability to dynamically establish a single one-hop HPR link to any host to which IP connectivity is enabled, provided that the host implements Enterprise Extender. In general, this allows the routing function to be handled entirely within IP. IP routers serve as the only routing nodes (hosts) in the network.

Figure 10-1 illustrates the components of one APPN Network with many host Network Notes (NNs) and End Nodes (ENs) involved.

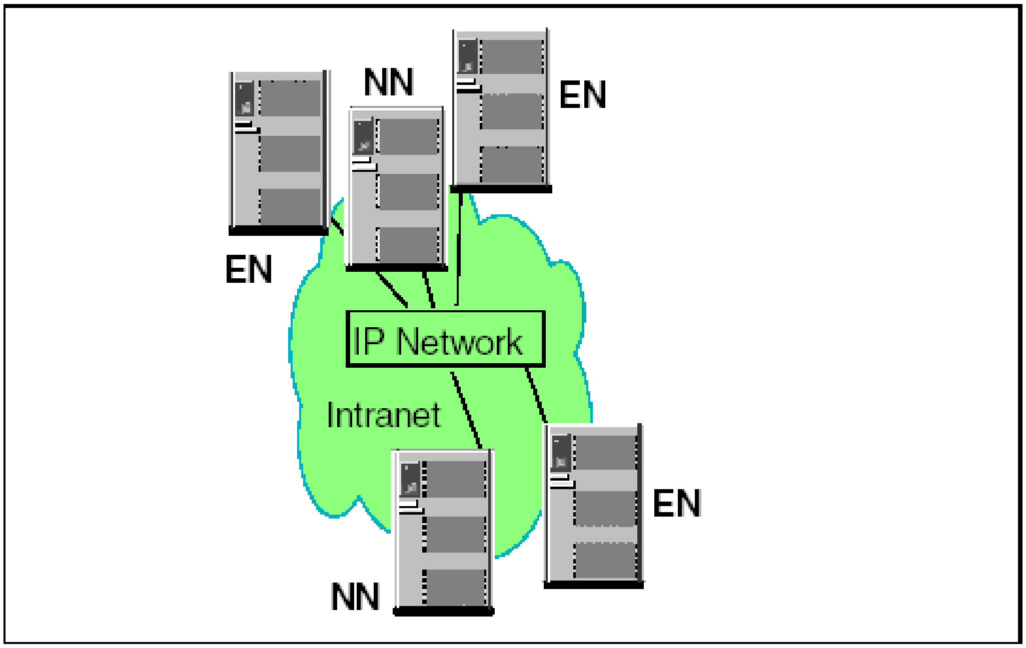

Figure 10-1 Enterprise Extender example

10.2 What is Enterprise Extender (EE)

The Enterprise Extender architecture carries SNA (HPR) traffic of any LU type over an IP infrastructure, without requiring changes to that infrastructure. It essentially treats the IP network as a particular type of SNA logical connection, in much the same way as an ATM or frame relay network is treated. In this manner, these SNA protocols act as transport protocols on top of IP, as does any other transport protocol such as Transmission Control Protocol (TCP).

10.2.1 Why is EE strategic

Enterprise Extender combines features of SNA and IP, to offer the best of both worlds when running SNA traffic over an IP backbone. Because of its design, Enterprise Extender is extremely flexible. It can be used in all networks, from the smallest to the largest, and provides you with a choice of where the SNA/IP boundary is placed. In this section we summarize some of the benefits of Enterprise Extender that make this solution strategic.

10.2.2 Cost effectiveness and resource convergence

The use of Enterprise Extender allows you to avoid costly application rewrites by providing a means of IP-enabling SNA applications that cannot be migrated to IP using technologies such as TN3270. That is, EE allows the continued use of native SNA applications over a different network: the IP network. This allows you to eliminate the SNA infrastructure altogether. The ultimate result is the convergence to a single network infrastructure that carries both IP and SNA application data.

Enterprise Extender has been designed to run over existing IP networks without requiring any change to applications or to IP routers. SNA applications see the same SNA network interfaces as before, whereas IP routers continue to see the same IP (UDP) packets.

> **Note:** Only the nodes at the *edges* of the IP network (potentially, just the host systems) need to be aware of Enterprise Extender.

10.2.3 Flow and congestion control

TCP/IP and HPR both provide their own unique, network-specific mechanisms for flow and congestion control. TCP uses a windowed technique, whereas HPR uses a technique based on data rate. Enterprise Extender introduced a new variant of the HPR flow control method known as responsive mode adaptive rate-based (ARB) flow control. Responsive mode ARB, like basic mode ARB, is designed to prevent network congestion; however, unlike basic mode ARB, it can also ensure a fair division of network capacity between the four SNA priorities and native IP traffic.

10.2.4 Class of service

One of the biggest issues facing those who wish to transport SNA over an IP network is the question of maintaining SNA's class of service. In SNA, the class of service specified for a particular session is used to determine both the route taken by the session and the transmission priority allotted to it.

With an IP backbone, the route is essentially unpredictable because of IP's connectionless property. However, IP provides for a transmission priority using the precedence bits in the IP header. Many routers now support the use of these bits, but in the past they have tended to use the TCP or UDP port number as a means of assigning priorities to packets.

Enterprise Extender supports the use of both the precedence bits and the port numbers to inform the IP network of the transmission priority. Use of the precedence bits is recommended because the UDP or TCP port numbers are carried inside the IP datagram, whereas the precedence bits are in the IP header. Thus, encrypted packets have unreadable port numbers and fragmented packets have no port numbers, after the first fragment. For such encrypted or fragmented packets, intermediate routers cannot determine the appropriate priority.

10.2.5 Internet connectivity exploitation

Enterprise Extender enables remote branches or workstations to be connected to the SNA backbone using the Internet, with no application changes required, while maintaining SNA connectivity from end to end. Dependent LU sessions can be carried on an Enterprise Extender connection as easily as any others, and by utilizing the dependent LU requester function (available on all current IBM workstation and router platforms), they can take advantage of the Enterprise Extender technology all the way into the most remote locations.

10.2.6 Session availability

TCP/IP has always had the ability to reroute packets around failing components, without disrupting the connection, by means of the connectionless property of IP. More recently SNA has implemented the same function, albeit in a rather different fashion. The high-performance routing (HPR) extension to SNA is connection-oriented as SNA has always been, but when it detects a failure, it will move an existing connection around a failing component. The use of HPR transport over an IP network provides non-disruptive rerouting around failed network components using either IP or HPR methods, depending on the location of the failure.

HPR consists of the following layers:

- Automatic Network Routing (ANR) - a connectionless layer that enables rerouting around failures.
- Rapid Transport Protocol (RTP) - a reliable connection-oriented layer that provides the end-to-end functionality.

10.3 What you need to understand about EE

This section describes the mechanism used to carry SNA traffic over an IP network. It includes a discussion of some of the interesting topics that arise when doing so, including the maintenance of Class of Service and TCP-friendly congestion and flow control.

10.3.1 Transporting SNA over IP

The designers of Enterprise Extender had the task of architecting the way in which SNA and IP-based protocols would be layered in order to transport SNA data over the IP network. They essentially had three choices of encapsulation of SNA data units: raw IP datagrams, UDP packets, or a TCP connection. In the following section, we describe each choice in more detail:

- Raw IP datagrams

  Datagrams are completely compatible with the HPR principles, as they flow through the network with minimal overhead and provide no error recovery of any sort. However, raw IP provides no means of multiplexing, particularly with no IETF-designated protocol value for HPR. Using a non-designated protocol value could lead to inconsistencies with security measures that filter IP packets based on this value. Additionally, although raw IP allows priority and type of service to be specified, in practice not all networks or routers are (or can be) configured to support this.

- UDP packets

  These packets provide the multiplexing required because they contain UDP port numbers. This allows Enterprise Extender packets to be distinguished from other IP packets. It also permits a priority scheme to be implemented independent of the type of service bits, since many routers are able to prioritize traffic based on the received port number. UDP, in addition, has low overhead since it does not concern itself with error recovery or flow control.

- TCP connection

  A TCP connection also provides multiplexing by means of port numbers, but it incurs a significantly higher overhead than raw IP or UDP. A TCP connection handles error recovery, retransmission and flow control; none of these is required for an HPR connection because the RTP endpoints are responsible for all of them.

UDP was therefore the method chosen for Enterprise Extender, as Figure 10-2 illustrates. It shows how SNA data is transported over the IP cloud in UDP frames.

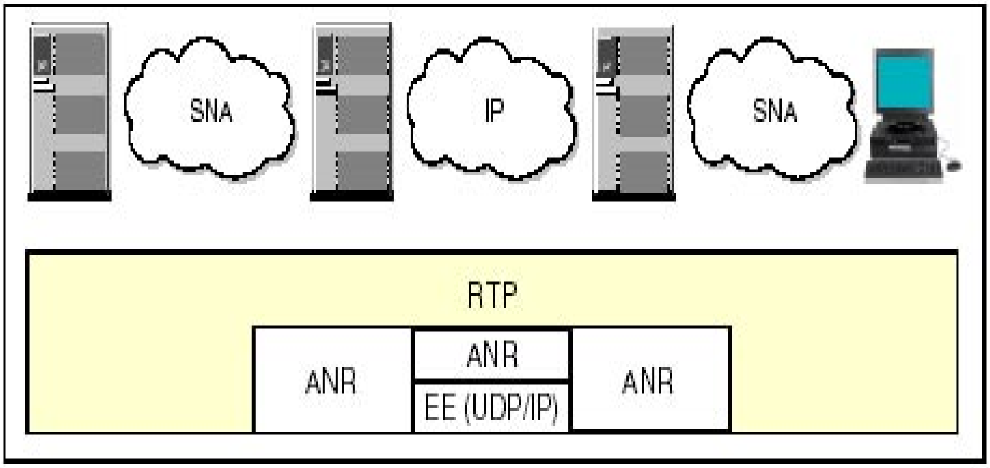

Figure 10-2 SNA transport over IP

10.3.2 Flow and congestion control - enhanced

The original ARB algorithm introduced with HPR works very well with SNA traffic alone, but is less efficient in the Enterprise Extender environment when SNA and IP traffic must coexist. Upon detection of a lost packet (a common occurrence in IP networks), the original ARB would immediately reduce its sending rate by a significant amount, thus impacting performance.

Therefore, an enhanced algorithm known as Responsive Mode ARB was introduced with Enterprise Extender, and is now an option for RTP nodes whether or not their HPR connection includes an Enterprise Extender link. Nodes that support Responsive Mode ARB can negotiate their level of ARB support during route setup exchange, and fall back to the original ARB if their partner does not support Responsive Mode. Responsive Mode ARB provides the following features:

It competes fairly with TCP congestion control.

- It can tolerate a certain level of lost data without significant degradation.
- It gives priority to short transmissions.
- It allocates a fair bandwidth to sustained transmissions, independent of propagation delays.
- It can ramp up its transmission rate faster at startup.

10.4 Enterprise Extender implementation

The following products implement Enterprise Extender:

- IBM Communications Server for OS/390
- IBM Communications Server for z/OS
- IBM Communications Server for Windows
- IBM Communications Server for Linux

- IBM Communications Server for AIX
- Cisco routers with SNA Switching Services (SNASw)
- IBM Personal Communications for Windows.

Restriction: VTAM for z/VM does not support Enterprise Extender.

IBM Communications Server allows multiple TCP/IP stacks to run concurrently with a single VTAM, but the Enterprise Extender function can utilize only one of these at a time. VTAM can change its allegiance from one TCP/IP stack to another, but only when all Enterprise Extender connections have been terminated. A VTAM start option, TCPNAME, allows you to specify which of several stacks VTAM is to use. If this option is not specified, VTAM will choose a suitable stack that supports Enterprise Extender.

If VTAM is to act as an Enterprise Extender node, it must have an IP address associated with it. This IP address is specified using the IPADDR start option, but the default IP address of the chosen stack will be used if the start option is not coded. The IP address needs to be predictable (because partner Enterprise Extender nodes may need to connect to it), but it is not desirable that a single IP port always be used for such connections. Therefore, this address must be a virtual IP address (VIPA). Moreover, it must be defined as a source VIPA so that datagrams sent out by VTAM have this address on them as the source.

The Enterprise Extender connections themselves are defined to VTAM as switched connections. The practical examples shown illustrate how these definitions are coded. The TCP/IP stack also needs a definition for the port represented by the VTAM application, and this is done by using a same-host (IUTSAMEH) statement as shown in the examples. This definition must be active before VTAM can establish any Enterprise Extender connections.

Enterprise Extender can be implemented under all modes of OSA-Express. A router that supports EE could be implemented in a remote site as an endpoint, eliminating the need for EE support at each workstation

10.5 EE configuration examples

Figure 10-3 on page 145 shows the network configuration we used to define EE via OSA-Express. We enabled EE between two Network Nodes (NN) SC64 and SC65, and also defined EE in an End Node (EN) workstation using the IBM Personal Communications for Windows.

Using the OSA-Express for EE

We configured two OSA-Express FENET operating in QDIO mode, to exploit the advantages provided for EE.

In Chapter 5, “QDIO mode” on page 61, we describe the configuration of the OSA-Express to support QDIO mode.

Note: Remember that an active TCP/IP connection, such as an OSA-Express port, is required for EE implementation.

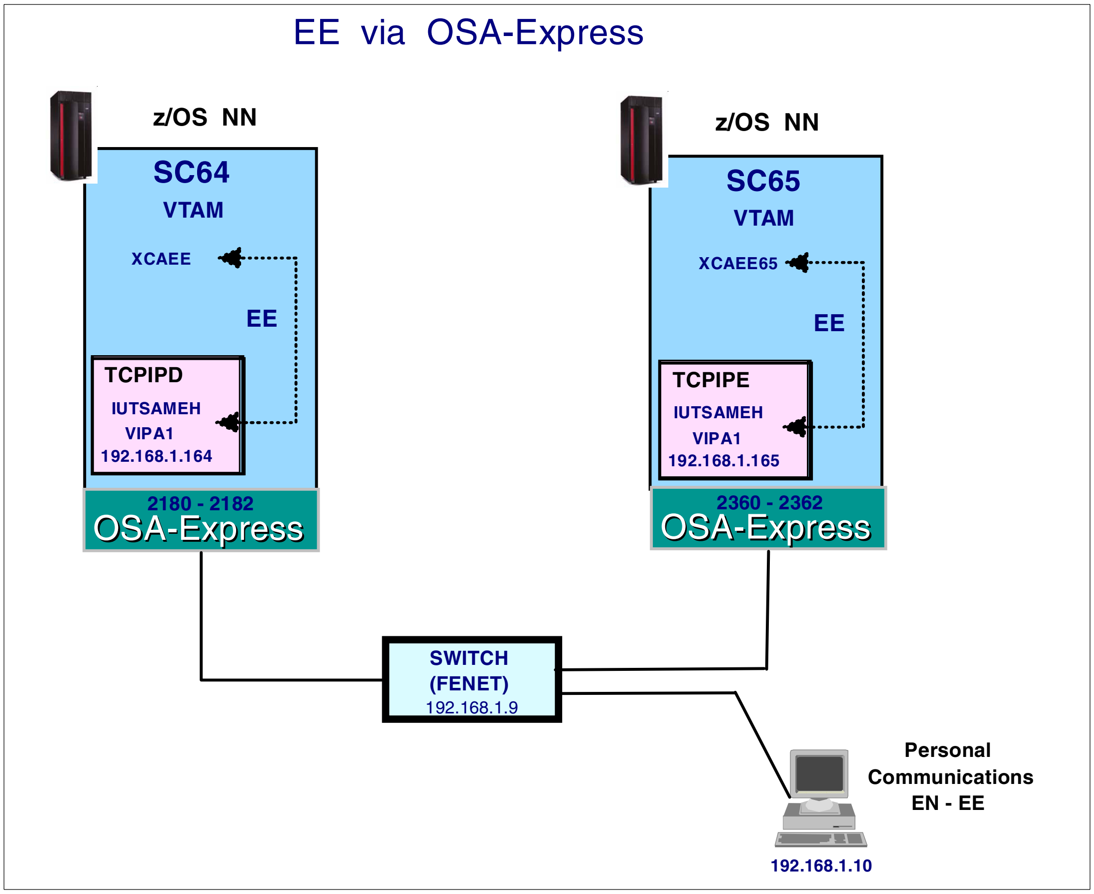

Figure 10-3 Network configuration

10.6 Defining EE to z/OS

This section describes the definitions required for TCP/IP and VTAM to allow access to SNA applications through an OSA-Express port.

10.6.1 TCP/IP definitions

Three types of LINK/DEVICE statements are required in the TCP/IP profile:

- VIRTUAL for the Virtual IP Address (VIPA)
- MPCIPA/IPAQENET for the OSA-Express port operating in QDIO mode
- MPCPTP for IUTSAMEH

> **Note:** The reserved TRLE named IUTSAMEH (same host) in VTAM provides the EE connection between VTAM and TCP/IP. VTAM automatically activates the IUTSAMEH when an MPCPTP device (in TCP/IP) is started.

In the profile for TCPIPD in the SC64 partition, we defined one virtual IP address (VIPA1), one OSA-Express port (OSA2180), and one IUTSAMEH. The same devices were defined in SC65 (TCPIPE), using different IP addresses.

Example 10-1 shows the TCP/IP profile definitions we used for our EE configuration in SC64.

Example 10-1 TCP/IP profile definitions for SC64

```
;TCP/IP SC64 PROFILE. TCPIP
 IPCONFIG
 DATAGRAMFWD
 SOURCEVIPA

;OSA 2180 FENET in QDIO to EE
 DEVICE OSA2180 MPCIPA PRIRouter
 LINK   SC642180LINK IPAQENET OSA2180

;VIPA Definitions
 DEVICE VIPA1 VIRTUAL 0
 LINK VLINK1 VIRTUAL 0 VIPA1

;For Enterprise Extender
 DEVICE IUTSAMEH MPCPTP
 LINK EELINK MPCPTP IUTSAMEH

   HOME
 192.168.1.64 SC642180LINK
 192.168.1.164 VLINK1

BEGINROUTES
ROUTE 192.168.1.0/24 = SC642180LINK MTU 1492
ENDROUTES

START IUTSAMEH
START OSA2180
```

SOURCEVIPA

This parameter is mandatory for Enterprise Extender connections, as it tells TCP/IP to return VIPA addresses on all outbound datagrams.

VIPA1

The VIPA address needs a virtual device and link defined; in our case, we named the device VIPA1 and the link VLINK1. The VIPA address looks like an interface to a virtual network that is concealed behind this TCP/IP stack.

10.6.2 VTAM definitions

In VTAM, two types of major nodes and modifications to the start options are required:

- External Communication Adapter (XCA) major node
- Switched major node
- ATCSRTxx member

External Communication Adapter (XCA)

We defined one XCA major node in each VTAM (SC64 and SC65) for the EE connections.

The XCA major node defines the IP port with the following statements:

PORT This identifies the name of a port through which an HPR connection is made via the IP network. The PORT statement must be set as MEDIUM=HPRIP for EE.

GROUP This must be set as DIAL=YES to define a TG. The group name (EEGROUP) is defined in the PATH statement of the switched major node for the NN-to-NN EE connection.

> **Note:** You should define at least as many lines as the maximum concurrent connections expected.

Example 10-2 shows the XCA definitions for SC64.

Example 10-2 XCA definitions of XCAEE

```
XCAEE        VBUILD TYPE= XCA
EEPORT       PORT MEDIUM=HPRIP,
                     LIVTIME=10,
                     IPTOS=(20,40,80,C0),
                     SAPADDR=4,
                     SRQRETRY=3,
                     SRQTIME=15
EEGROUP   GROUP DIAL=YES,
                     DYNPUPFX=BR,
                     AUTOGEN=(10,L,P),
                     CALL=INOUT
```

We used the same definitions on SC65 (XCAEE65).

Switched major nodes for the NN-to-NN connection

A switched major node is used to define the PU and PATH statements to establish the EE connection:

PU This defines the remote node (physical unit)

PATH This defines the remote IP address or host name, and the group name specified in the XCA major node.

Example 10-3 shows the switched major node we used for SC64 (acting as an NN) that was activated on SC65.

Example 10-3 Switched major node for SC64 (NN)

```
SWT064   VBUILD TYPE=SWNET,
                MAXGRP=4,
                MAXNO=512
 PUT064   PU    MAXDATA=256,ADDR=01,
                CPNAME=SC64M,
                CPCP=YES,HPR=YES,
                DWACT=YES,PUTYPE=2
 PATH64   PATH SAPADDR=08,IPADDR=192.168.1.164,
                GRPNM=EEGROUP
```

Example 10-4 shows the switched major node we used for SC65 (acting as an NN) that was activated on SC64.

Example 10-4 Switched major node for SC65 (NN)

```
SWT065   VBUILD TYPE=SWNET,
                MAXGRP=4,
                MAXNO=512
 PUT065   PU    MAXDATA=256,ADDR=01,
                CPNAME=SC65M,
                CPCP=YES,HPR=YES,
                DWACT=YES,PUTYPE=2
 PATH65   PATH SAPADDR=8,IPADDR=192.168.1.165,
                GRPNM=EEGROUP
```

Switched major node for the EN-to-NN connection

We also connected an EN workstation to SC64 (NN), using Dependent LU Request (DLUR), through EE. In the EN, using IBM Personal Communications (PCOM), we defined two LUs to emulate 3270 sessions.

Example 10-5 shows the switched major node we used for our EN workstation.

Example 10-5 Switched major node for PUDLUR1 (EN)

```
SWNDLUR1 VBUILD TYPE=SWNET,
                      MAXGRP=4,
                      MAXNO=512
PUDLUR1    PU    ADDR=01,
                     IDBLK=05D,
                     IDNUM=11111,
                     ISTATUS=ACTIVE,
                     DISCNT=NO,
                     MAXDATA=265,
                     MAXOUT=7,
                     PACING=0,
                     PUTYPE=2,MODETAB=NEWMTAB,
                     VPACING=0
 LUEE0001   LU    LOCADDR=2,DLOGMOD=D4A32782,
                     LOGAPPL=SC64TS
 LUEE0002   LU    LOCADDR=3,DLOGMOD=D4A32782,
                     LOGAPPL=SC64TS
```

We specified the LOGAPPL=SC64TS to automatically connect the EN LUs to TSO on SC64.

Note: In Example 10-5, the highlighted parameters must match with the configuration in IBM Personal Communications.

VTAM start options

We altered the VTAM start options in the ATCSTRxx member to support EE and NN. We modified the NODETYPE, TCPNAME, and IPADDR parameters.

The VTAM options coded to implement this connection are shown in Example 10-6 on page 149.

Example 10-6 VTAM options for ATCSTR64

```
CONFIG=64,
SSCPID=64,
NOPROMPT,
SSCPNAME=SC64M,
NETID=USIBMSC,
XCFINIT=YES,                    XCF LINK
IQDCHPID=ED,                  XCF LINK CHPID FOR HIPERSOCKETS
HOSTSA=64,                      INTERCHANGE NODE
MAXSUBA=255,
NODETYPE=NN,                NETWORK NODE NN FOR EE
SORDER=SUBAREA,
APPNCOS=#INTER,           DEFAULT APPN COS
CONNTYPE=APPN,
CPCP=YES,
CDSERVR=YES,                 CENTRAL DIRECTORY SERVER
IPADDR=192.168.1.164,      VIPA ADDRESS FOR EE ENTERPRISE EXT
IOPURGE=180,
TCPNAME=TCPIPD,           TCPIP JOB NAME FOR EE ENT.EXT
SUPP=NOSUP,
HOSTPU=SC64MPU,
PPOLOG=YES,
DYNLU=YES,
```

The relevant definitions are listed as follows:

NODETYPE=NN The NODETYPE option must be coded as NN if this VTAM is to act as a Network Node Server (NNS) or a Dependent LU Server (DLUS) for other nodes. An Enterprise Extender connection with dependent LU sessions will work just as well to a VTAM End Node that does not have to act as an NNS or DLUS.

IPADDR The IPADDR option is used to define the IP address of the local VIPA (in the TCP/IP profile).

TCPNAME The TCPNAME option defines the jobname of TCP/IP stack to be utilized.

> **Note:** In the definitions for SC65 (ATCSTR65), we altered the IPADDR to 192.168.1.165 and TCPNAME to TCPIPE. If TCPNAME is not defined, it can be dynamically updated with the MODIFY console command. Example: F VTAM44,VTAMOPTS,TCPNAME=TCPIPE VTAM44 is the jobname of our started task (STC) for VTAM.

The parameters can be verified with the display of the VTAM options. To display the current VTAM options, enter the following command:

```
D NET,VTAMOPTS
```

10.6.3 Activation and verification (NN-to-NN)

Once all the definitions were added to VTAM and TCPIP, we activated the configuration.

Activation requires several tasks, in this order:

1. Verifying that the OSA-Express devices are online.
2. Activating the VTAM resources; TRLE(for OSA-Express), ATCSTR, XCA, and switched major nodes.
3. Starting TCP/IP devices.

Verify EE is running on z/OS

After activation, we verified that EE was running by issuing the following commands:

```
D TCPIP,TCPIPx,NETSTAT,DEV

D NET,ID=XCAEE

D NET,ID=SWTO65 (to verify the NN to NN connection)

D NET,ID=PUTO65,E
```

Example 10-7 shows the status of VIPA1 and IUTSAMEH DEVICEs used for Enterprise Extender.

Example 10-7 D TCPIP,TCPIPD,NETSAT,DEV

```
D TCPIP,TCPIPD,NETSTAT,DEV
DEVNAME: VIPA1            DEVTYPE: VIPA
   DEVSTATUS: READY
   LNKNAME: VLINK1           LNKTYPE: VIPA        LNKSTATUS: READY
     NETNUM: 0   QUESIZE: 0

 DEVNAME: IUTSAMEH         DEVTYPE: MPC
   DEVSTATUS: READY
   LNKNAME: EELINK           LNKTYPE: MPC         LNKSTATUS: READY
     NETNUM: 0   QUESIZE: 0
```

We looked for IUTSAMEH and VIPA1; the link status must be READY.

Example 10-8 that shows a display of XCAEE (XCA majnode) used for Enterprise Extender.

Example 10-8 D NET,ID=XCAEE,E

```
D NET,ID=XCAEE,E
IST097I DISPLAY ACCEPTED
IST075I NAME = XCAEE, TYPE = XCA MAJOR NODE 222
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1679I MEDIUM = HPRIP
IST1685I TCP/IP JOB NAME = TCPIPD
IST1680I LOCAL IP ADDRESS 192.168.1.164   (VIPA1)
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST232I L0000000 ACTIV
IST232I L0000001 ACTIV
IST232I L0000002 ACTIV
IST232I L0000003 ACTIV
IST314I END
```

We looked for an active state and the following:

HPRIP Medium=HPRIP (as opposed to CSMACD, ATM, or RING).
TCPIPD This is the TCP/IP job name.
192.168.1.164 This is the local VIPA.

Example 10-9 on page 151 shows the results of DNET,ID=SWTO65 and PUTO65, to verify the NN-to-NN connection between SC64 to SC65 via Enterprise Extender.

Example 10-9 D NET,ID=SWTO65 and PUTO65

```
D NET,ID=SWTO65,E
IST097I DISPLAY ACCEPTED
IST075I NAME = SWTO65, TYPE = SW SNA MAJ NODE 633
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST084I NETWORK RESOURCES:
IST089I PUTO65   TYPE = PU_T2.1,ACTIV
IST1500I STATE TRACE = OFF
IST314I END
********************************************************************
D NET,ID=PUTO65,E
IST097I DISPLAY ACCEPTED
IST075I NAME = PUTO65, TYPE = PU_T2.1 643
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1043I CP NAME = SC65M, CP NETID = USIBMSC, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1105I RESOURCE STATUS TGN CP-CP TG CHARACTERISTICS
IST1106I PUTO65   AC/R    22 YES   98750000000000000000017100808080
IST1482I HPR = RTP - OVERRIDE = N/A - CONNECTION = YES
IST1510I LLERP = NOTPREF - RECEIVED = NOTALLOW
IST1680I LOCAL IP ADDRESS 192.168.1.164
IST1680I REMOTE IP ADDRESS 192.168.1.165
IST136I SWITCHED SNA MAJOR NODE = SWTO65
IST081I LINE NAME = LOOOOOOO, LINE GROUP = EEGROUP, MAJNOD = XCA
```

We looked for an active state and the following:

PU_T2.1 The link station is a PU type 2.1, APPN connection between the NNs.

TGN CP-CP The link station is in the APPN topology database CP-CP HPR=RTP. The level of HPR supported is RTP and the status of the connection.

LOCAL IP ADDRESS 192.168.1.164 is the VIPA address of the TCP/IP stack being used by this VTAM.

REMOTE IP ADDRESS 192.168.1.165 is the VIPA address of the remote NN.

10.6.4 Defining EE-to-IBM Personal Communications (EN-to-NN)

After installation and verification of EE on z/OS (SC64), the VTAM switched major node SWDLUR (Example 10-5 on page 148) was activated.

Using IBM Personal Communications V5.0 (with EE support) on the workstation, we quickly configured one EE connection to VTAM. In this example, the workstation is the EN and SC64 is the NN.

In the start menu from Windows, we did the following:

Clicked **Start > Program > IBM Personal Communications > Star or Configure Session. (select “Configure Session”).**

We then followed these steps:

1. We selected **S/390** on Type of Host; **IBM-EEDLC** on Interface; **LU(0,1,2,3) via DLUR** for the Attachment (APPN can also be used). Then we clicked **Link Parameter** (see Figure 10-4 on page 152).

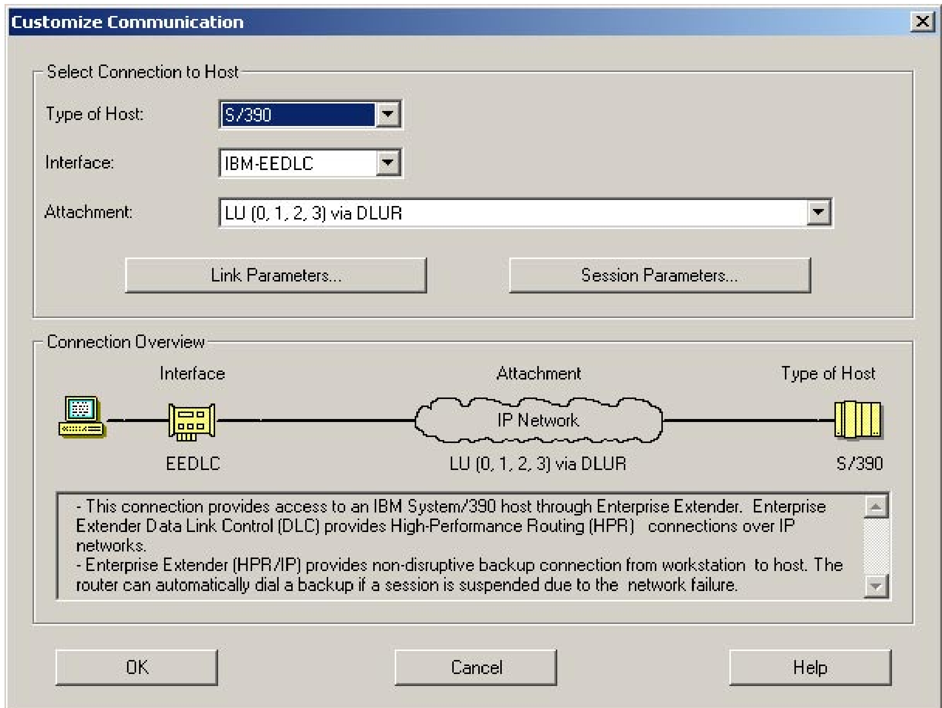

Figure 10-4 Customize Communication step 1

2. We entered our Net ID and added a CP Name, Block ID and Physical ID and clicked **Next** (see Figure 10-5).

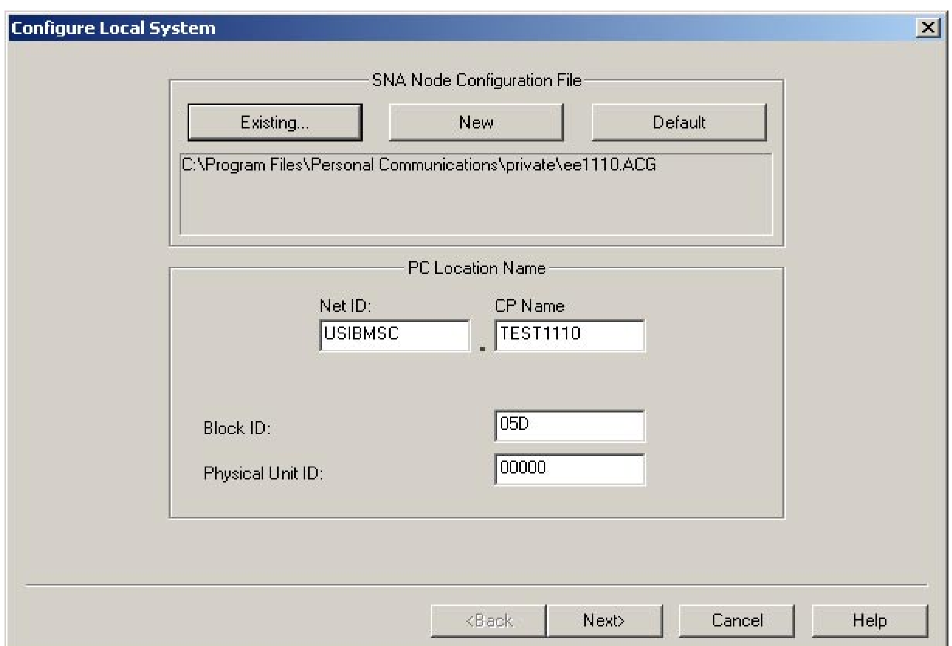

Figure 10-5 Configure Local System step 2

3. We entered our PU name, Block ID, Physical Unit ID, and the DLUS name (Net ID and CPNAME of VTAM), and clicked **Next** (see Figure 10-6 on page 153).

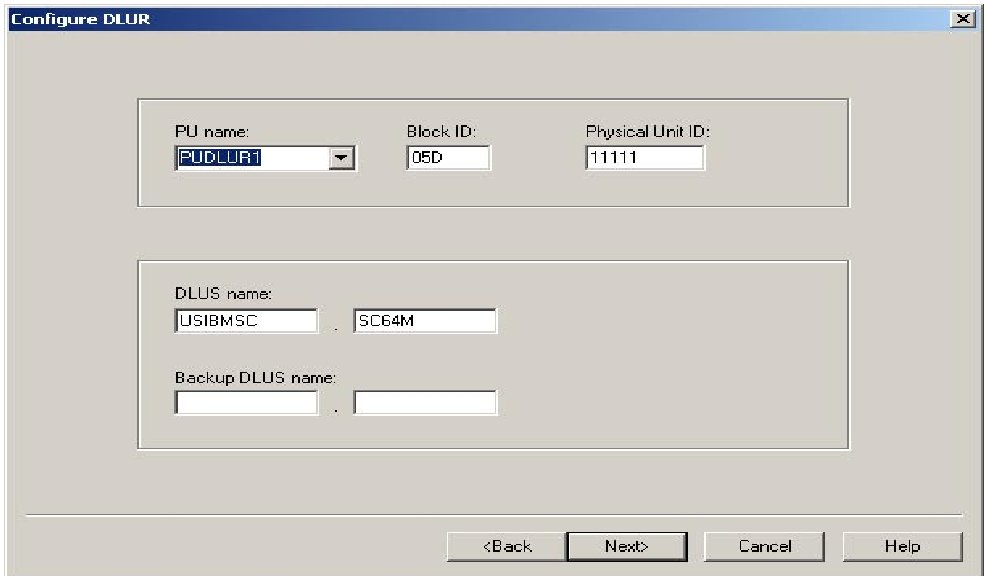

Figure 10-6 Configure DLUR step 3

Important: The PU name, the Block ID, and the Physical Unit ID required in the Configure DLUR panel, *must* match the definitions defined to the associated switched major node in VTAM.

4. We chose Existing device and Existing connection, as shown in Figure 10-7 and clicked **Next**.

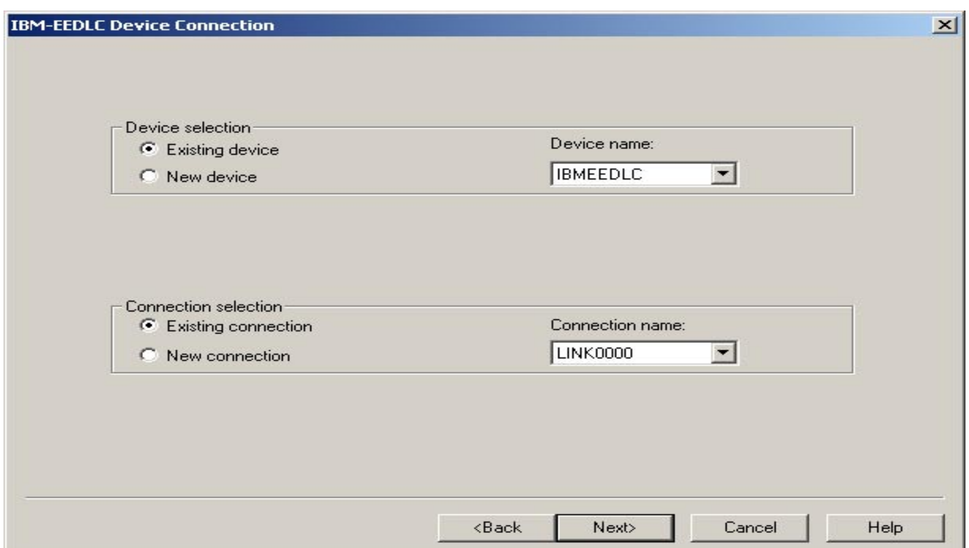

Figure 10-7 IBM EEDLC Device Connection step 4

5. We entered the IP address of SC64 and used the default SAP of 04, as shown in Figure 10-8 on page 154, and clicked **Next**. We then clicked **OK** and saved the new configuration for EE.

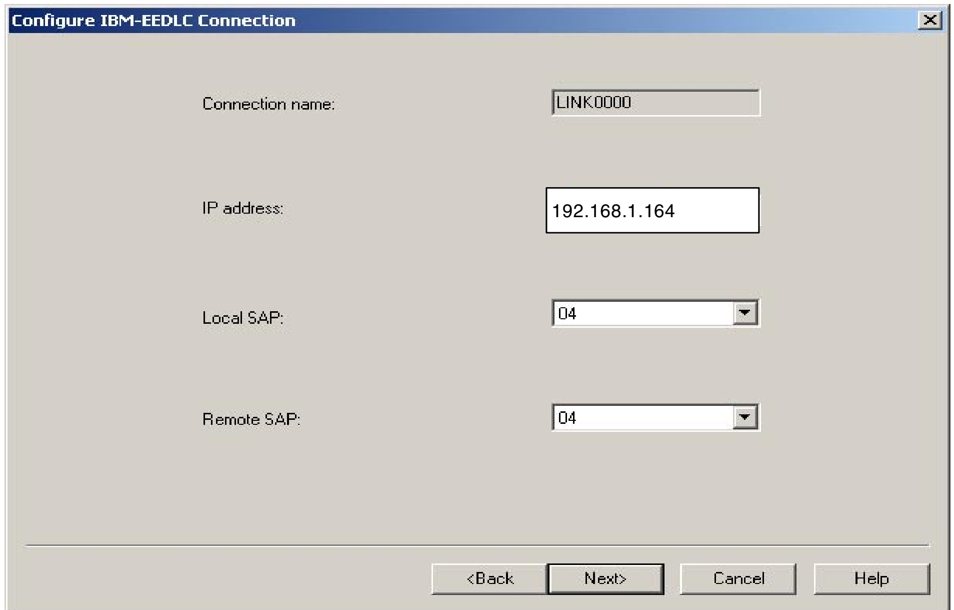

Figure 10-8 IBM-EEDLC connection

Important: The IP address in the Configure IBM-EEDLC Connection panel *must* be the IP address of the VIPA (192.168.1.164).

10.6.5 Activation and verification (EN-to-NN)

Once all the definitions were added to IBM Personal Communications, we activated the configuration.

Activation required these tasks:

- Verifying IP connectivity between the workstation and the VIPA.
- Verifying that the PU DLUR and LUs were active by using the NN VTAM display command on SC64: **`D NET,ID=PUDLUR1,E`**
- Confirming that the correct DLOGMOD,USSTAB and LOGAPPL were in accordance with our environment.

Verifying the EE connection for the workstation

Since we used LOGAPPL=SC64TS (TSO on SC64), we received the TSO prompt (IKJ56700A SC64 ENTER USERID =) after activation.

Example 10-10 shows the VTAM display command we used to verify that the defined LUs were in session.

Example 10-10 Display of LU sessions using EE

```
D NET,ID=PLDLUR1,E
IST453I ID PARAMETER VALUE USIBMSC.PLDLUR1 NOT VALID
D NET,ID=PUDLUR1,E
IST097I DISPLAY ACCEPTED
IST075I NAME = PUDLUR1, TYPE = PU_T2 668
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1043I CP NAME = ***NA***, CP NETID = USIBMSC, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST1354I DLUR NAME = TEST1110           MAJNODE = SWDLUR
IST136I SWITCHED SNA MAJOR NODE = SWDLUR
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST1657I MAJOR NODE VTAMTOPO = REPORT
IST355I LOGICAL UNITS:
IST080I LUEE0001 ACT/S      LUEE0002 ACT/S
IST314I END
```

11

VLAN support

In this chapter, we discuss the IEEE 802.1Q standard, known as Virtual Local Area Network (VLAN). We also explain the VLAN concept and provide a sample implementation of the OSA-Express in conjunction with multiple VLANs.

VLAN technology is supported by the OSA-Express GbE and FENET features running in QDIO mode, in a Linux kernel 2.4.19 (or later) environment.

© Copyright IBM Corp. 2003. All rights reserved.

11.1 VLAN overview

A VLAN is a group of workstations with a common set of requirements, independent of physical location. VLANs have the same attributes as a physical LAN, even though they are not located physically on the same LAN segment. VLANs can be used to increase bandwidth and reduce overhead by allowing networks to be organized for optimum traffic flow.

Figure 11-1 shows an example of VLANs segmented into logically defined networks.

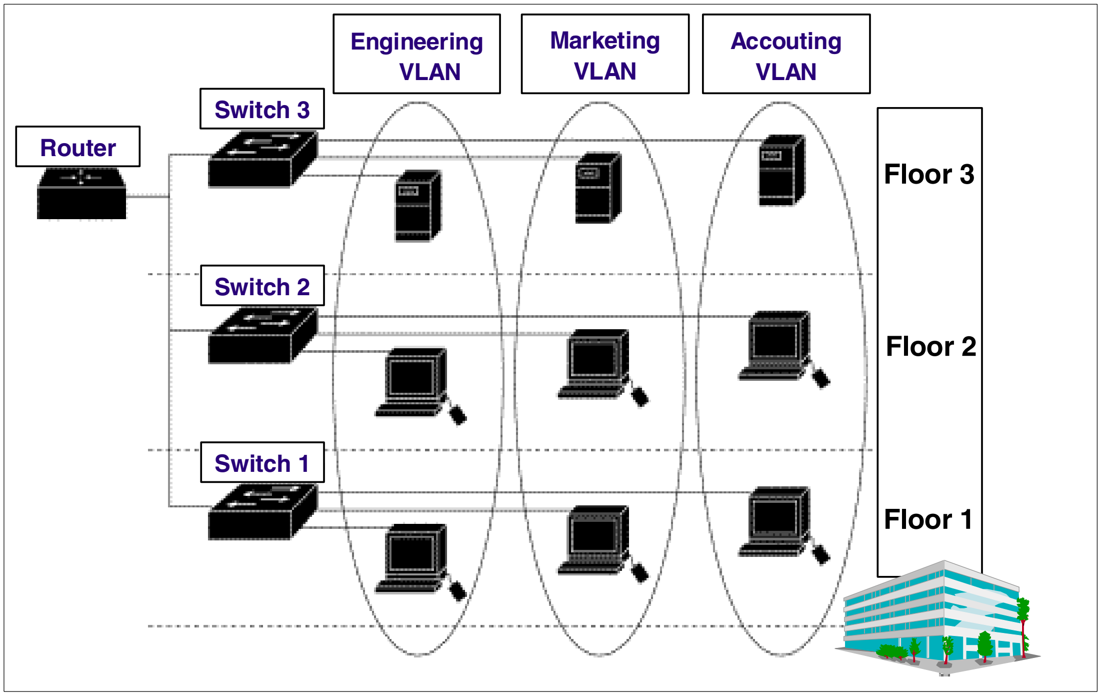

Figure 11-1 Logically defined networks using VLANs

11.1.1 Types of connections

IEEE 802.1Q VLANs operate by defining switch ports as members of virtual LANs. Devices on a VLAN can be connected in three ways, based on whether the connected devices are VLAN-aware or non-VLAN-aware. The VLAN-aware device is one which understands VLAN memberships (which users belong to a particular VLAN) and VLAN formats.

Ports used to attach non-VLAN-aware equipment are called *access* ports, while ports used to connect to other switches are known as *trunk* ports. Network frames generated by VLAN-aware equipment are marked with a *tag,* which identifies the frame to the VLAN.

Trunk port

All the devices connected to a trunk port, including workstations, must be VLAN-aware. All frames on a trunk port must have a special header attached. These special frames are called *tagged* frames. Trunk ports can be defined to multiple VIDs, since the VLAN-aware equipment knows how to handle frames with different VIDs.

Access port

Access ports are defined to belong to a single VLAN. All frames received from devices attached to access ports are “tagged” with the same VLAN ID (VID).

Hybrid port

This is a combination of the previous two ports. This is a port where both VLAN-aware and VLAN-unaware devices are attached. A hybrid port can have both tagged and untagged frames.

Figure 11-2 shows a logical diagram of a VLAN environment with two switches and a number of VLANs.

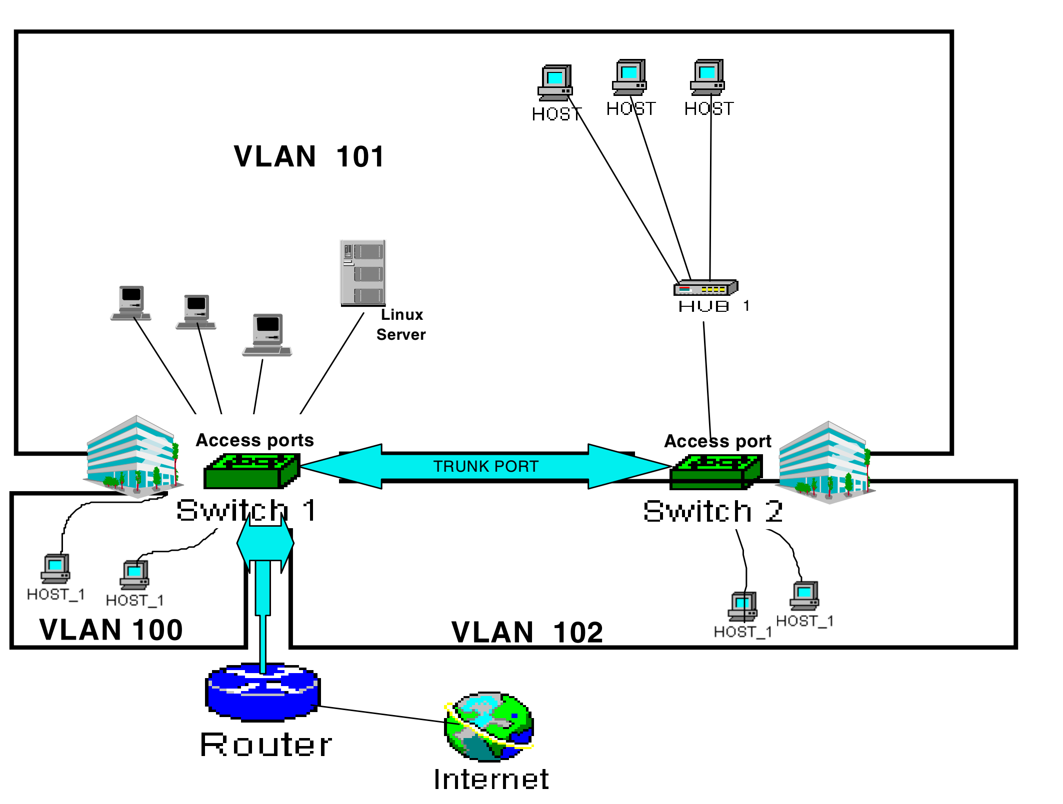

Figure 11-2 VLAN logical diagram

In this example network, VLAN 100 only exists in Switch 1, because the trunk port to Switch 2 is not a member of VLAN 100. VLANs 101 and 102 span the two switches, because the trunk ports in each switch are members of both VLANs.

There is a single VLAN-aware router connected to Switch 1, which provides access to an external network to users in VLANs 100 and 102. The trunk port to which the router is attached is defined to VLANs 100 and 102, not to VLAN 101. Therefore, no workstations in VLAN 101 can reach the router. Hub 1 is a non-VLAN-aware device, so the port in Switch 2 that it connects to is an access port.

This example illustrates two reasons why VLANs are generally used:

1. Staff in different physical locations retain common access to resources.

   The Linux server is used by staff in both buildings. By defining these workstations to the same VLAN, no additional configuration or equipment is required for either location to access the Linux server, while at the same time ensuring that other staff do not get access.

2. Consolidation of resource access.

   The external network has to be accessed by different staff in both buildings. Extending VLAN 102 across to the router port in Switch 1 (and configuring the router correctly) saves having to provide an additional link to the external network or the router from the other building.

Broadcast in VLANs

All ports that are members of the same VLAN, including trunk ports, will operate as if they were part of the same physical network. When a multicast or broadcast frame is received from a device on a particular VLAN, the switch transmits the frame to all ports (both trunk and access ports) belonging to the same VLAN.

The only difference between the trunk and access port in this case is that the frame transmitted onto the trunk port will have the VLAN tag intact (so that the VLAN-aware equipment at the other end of the link will know how to handle it).

VLAN isolation

An important point about VLANs in general is that they provide isolation. VLANs behave like separate physical networks, even though they may be contained within the same switch.

In order for devices in different VLANs to communicate, TCP/IP routing must occur. In the network shown in Figure 11-2 on page 159, workstations in VLAN 100 and VLAN 101 cannot communicate, because there is no routing path between the two VLANs. Workstations in VLANs 100 and 102 can communicate, as long as the router they are attached to is configured appropriately.

11.2 VLAN support for OSA-Express on Linux

To support Virtual Local Area Networks (VLANs), the OSA-Express Gigabit Ethernet and Fast Ethernet features must be at DRV3G with MCL (J11204.022) or later, and running in QDIO mode.

> **Important:** VLAN support was added to the Linux kernel at version 2.4.19. If your Linux guests are not running this level or higher, you will require an updated kernel.
>
> At the time of writing, SuSE SLES8 is the only Linux distribution that supports VLAN.

An example of Linux servers under z/VM, using VLANs is shown in Figure 11-3 on page 161.

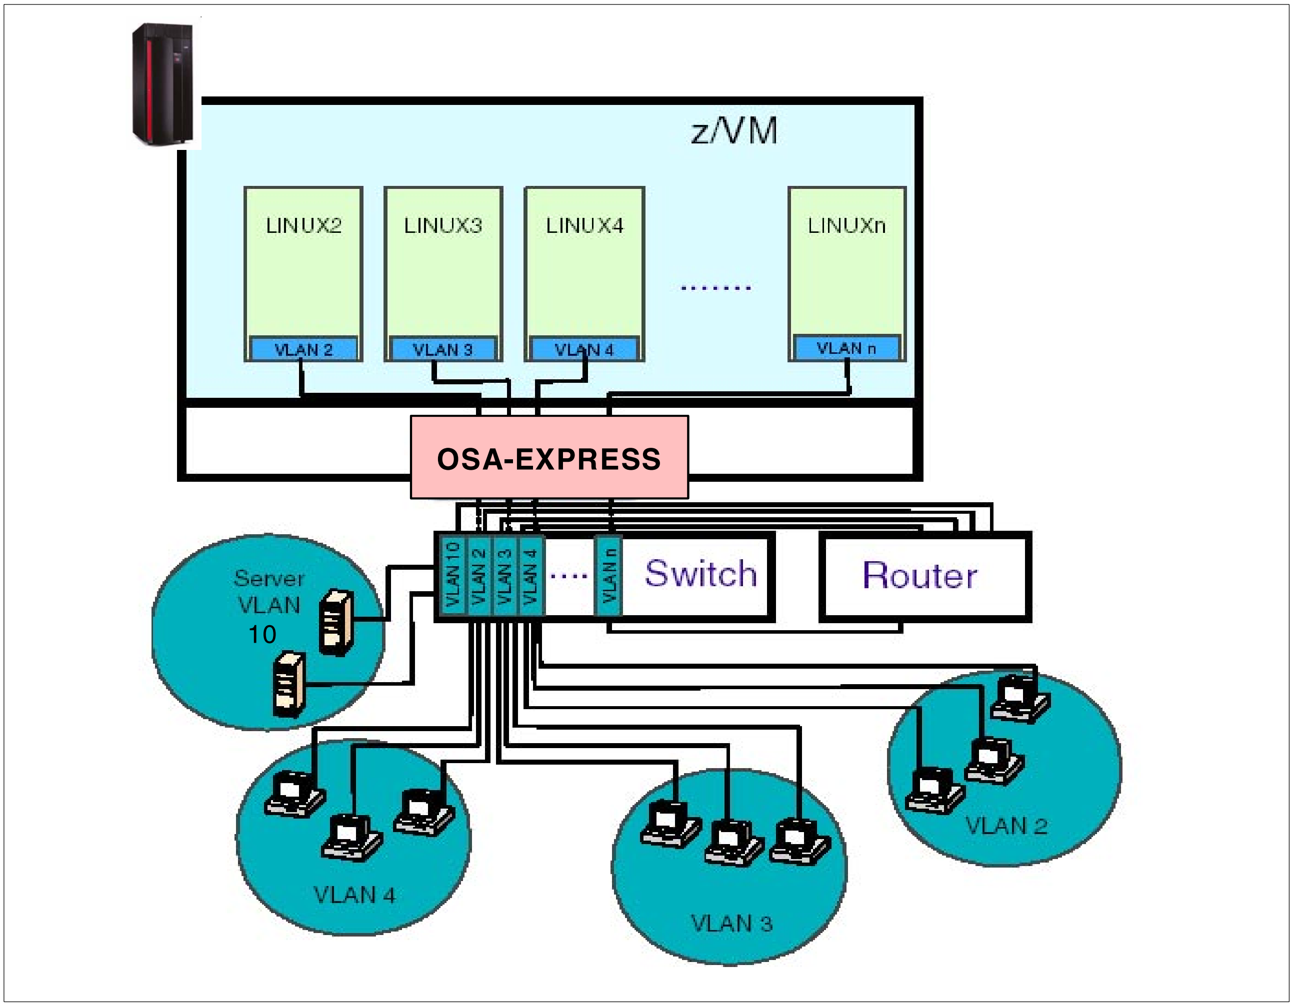

Figure 11-3 Linux using VLAN via OSA

In this diagram, the Linux guests have a logical association with certain groups of workstations in the network. With the use of VLANs, those Linux servers are connected to their client networks without the overhead of routing, and with the advantage of behaving as if they were on the same physical network.

In "VLAN isolation" on page 160, we stated that VLANs provide isolation between networks. Figure 11-3 shows a router with connections to the different VLANs to provide access between the client workstations and the servers in VLAN 10. This can also be done with a router that is VLAN-aware, which means only one physical connection to the switch (from the router) is needed.

Note: Most switch vendors provide routing functions in their switches to allow traffic between VLANs.

11.3 Sharing an OSA-Express when using VLANs

At the time of writing, Linux is the only zSeries operating system to support VLAN tagging using the OSA-Express interface. It is possible, however, to share an OSA-Express between Linux and other zSeries operating systems even though they do not provide VLAN support.

Figure 11-4 shows a network with z/VM and z/OS systems sharing an OSA-Express with Linux guests using VLANs.

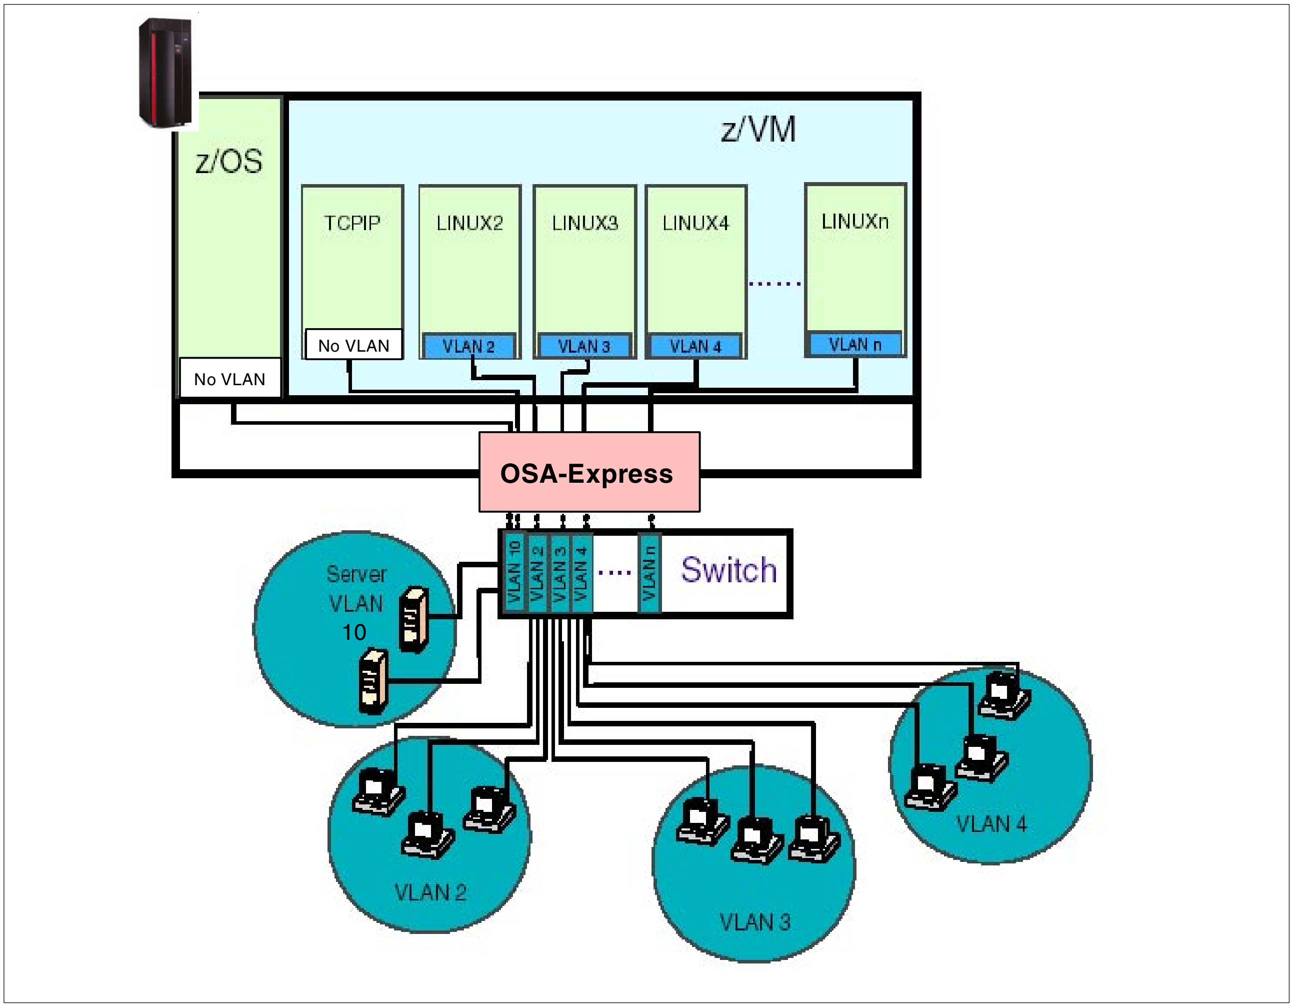

Figure 11-4 VLANs sharing the OSA-Express

Usually, all frames sent on a trunk link are VLAN-tagged. In Figure 11-5, however, we have a z/VM TCP/IP stack and a z/OS Communications Server stack sharing an OSA-Express, with the Linux guests connecting to separate VLANs. Since neither z/VM nor z/OS support IEEE 802.1Q (VLANs), the traffic generated by these stacks will not have VLAN tagging.

According to the 802.1Q standard, this is a valid configuration. Rather than a trunk port, the switch port to which the OSA-Express connects is a *hybrid* port, since it must behave like both an access port and a trunk port. A hybrid port acts like a trunk port for frames that carry a VLAN tag, but the switch tags frames without a VLAN tag with a port VID (like it would with an access port).

11.3.1 Configuring VLANs to Linux

The **vconfig** command is used to add or remove a VLAN configuration to a pre-defined adapter.

Examples

VLANs are allocated to an existing interface representing a physical Ethernet LAN. The following examples show this:

- To add VLAN 100 to an OSA-Express interface (eth1)

  ```
  vconfig add eth1 100
  ```

  This creates a new VLAN interface called eth1.100, which you can now configure as active by using the **ifconfig** command:

  ```
  ifconfig eth1.100 192.168.100.1 netmask 255.255.255.0 up
  ```

- To add two VLANs 2 and 4 to an OSA-Express interface (eth1):

  ```
  vconfig add eth1 2
  vconfig add eth1 4
  ```

  The vconfig commands have added interfaces eth1.2. and eth1.4, which you can then configure:

  ```
  ifconfig eth1.2 10.100.2.3 netmask 255.255.255.0 up
  ifconfig eth1.4 11.111.2.3 netmask 255.255.0.0 up
  ```

  The traffic that flows out of eth1.2 will be in the VLAN with ID=2 (and will not be received by other stacks that listen to VLANs with ID=3).The internal routing table will ensure that every packet for 10.100.2.x goes out via eth1.2 and everything to 11.111.x.x goes via eth1.4.

- To remove a VLAN interface:

  ```
  ifconfig eth1.3 down
  vconfig rem eth1.3
  ```

Startup configuration

At this time, configuring VLANs is a manual process that must be scripted to take place when the Linux guests' are booted up. As VLAN usage grows, you can expect to see that distributors will include VLAN boot-time configuration in their network scripts.

11.4 VLAN implementation

In our environment (see Figure 11-5 on page 164), we configured two VLANs, one for each Linux system. We enabled a connection between the two VLANs through a Cisco switch (6509), and defined a trunk port for the OSA-Express interface. We also installed two workstations, one connected to each VLAN.

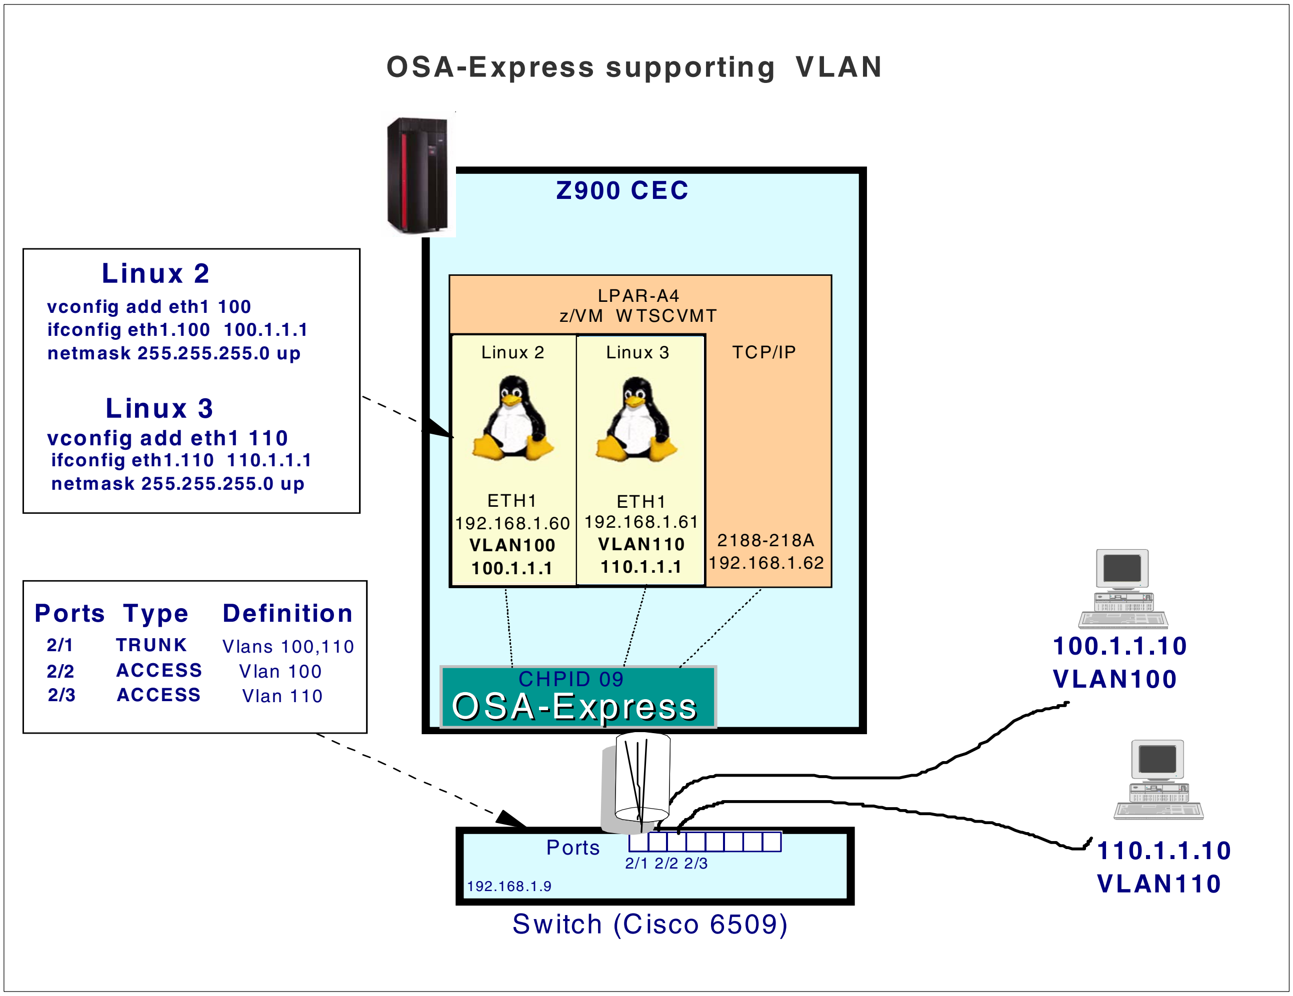

Figure 11-5 OSA-Express VLAN configuration

LINUX commands

On Linux 2 and Linux 3, we defined the VLAN configurations using the following commands:

- For Linux 2

```
vconfig add eth1 100
ifconfig eth1.100 100.1.1.1 netmask 255.255.255.0 up
```

- For Linux 3

```
vconfig add eth1 110
ifconfig eth1.110 110.1.1.1 netmask 255.255.255.0 up
```

After issuing the **`vconfig add`** commands on the respective systems, we received the following messages:

```
Added VLAN with VID == 100 to IF -:eth1:-
```

```
Added VLAN with VID == 110 to IF -:eth1:-
```

Cisco 6509 configuration

We altered the configuration of the Cisco 6509 to support the new VLANs and the trunk. We added the VLAN 100 to port 2/2, the VLAN 110 to port 2/3, and defined the trunk to port 2/1.

Using the command **sh config** in the switch (Cisco 6509), we saw the commands that were used to alter the configuration (see Figure 11-6).

```
begin
!
# ***** NON-DEFAULT CONFIGURATION *****
#time: Tue Nov 5 2002, 13:48:35
!
#version 6.2(2)
!
set password $2$7/k5$gVQbApTTDMWnqrO1fciqL0
set enablepass $2$E0zE$AjJ70DD7nmGAWsLSLpYz01
!
#vtp
set vtp domain ibmtest
set vlan 1 name default type ethernet mtu 1500  state active
set vlan 100,110 type ethernet mtu 1500  state active
set vlan 1003 name token-ring-default type trcrf mtu 1500
!
#ip
set interface sc0 1 192.168.1.9/255.255.255.0 192.168.1.255
set ip route 0.0.0.0/0.0.0.0         192.168.1.11
#module 2 : 48-port 10/100BaseTX Ethernet
set trunk 2/1  on dot1q 1-1005,1025-4094
set vlan 100  2/2
set vlan 110  2/3
!
end
```

Figure 11-6 Cisco 6509 configuration

Note: The 802.1Q trunk protocol is defined in the Cisco 6509 using the parameter dot1q.

Verification

After activation, we verified that the VLANs in our environment were defined correctly, using the **sh trunk** and **sh vlan** commands in the Cisco 6509.

In Figure 11-7 and Figure 11-8 on page 166, we show the result of these commands.

```
Console> (enable) sh trunk
Port     Mode         Encapsulation  Status       Native vlan
-------- -----------  -------------  ------------ -----------
 2/1     on           dot1q          trunking          1
15/1     nonegotiate  isl            trunking          1

Port     Vlans allowed on trunk
-------- ---------------------------------------------------------------------
 2/1     1-1005,1025-4094
15/1     1-99,101-109,111-1005,1025-4094

Port     Vlans allowed and active in management domain
-------- ---------------------------------------------------------------------
 2/1     1,100,110
```

Figure 11-7 Show trunk

```
Console> (enable) sh vlan
VLAN Name                              Status    IfIndex   Mod/Ports
---- --------------------------------  --------- -------   --------------
1    default                           active    5         1/1-2 2/4-48
100  VLAN0100                          active    107       2/2
110  VLAN0110                          active    108       2/3
1002 fddi-default                      active    6
```

Figure 11-8 Show VLAN

We verified the VLAN definitions by issuing the **ifconfig** command. The results of Linux 3 are shown in Example 11-1.

Example 11-1 Results of ifconfig command

```
eth1.110  Link encap:Ethernet  HWaddr 00:06:29:6C:A5:BC
          inet addr:110.1.1.1  Mask:255.255.255.0
          inet6 addr: fe80::6:2900:56c:a5bc/10 Scope:Link
          UP RUNNING MULTICAST  MTU:1492  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:2 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 b)  TX bytes:148 (148.0 b)
```

We then tested the connections using the following **ping** commands:

```
ping 100.1.1.[1 - 10] // Unicast-PING

ping -I eth1.100 224.0.0.1 // Multicast-PING

ping -b 100.1.1.255 // Broadcast-
```

12

IPv6 support

OSA-Express GbE and FENET features running in QDIO mode can use the IPv6 protocol in a z/OS V1R4 or Linux (version 2.4.19) environment. The OSA-Express features must be at DRV3G with MCL (J11204.022) or later.

This chapter describes the implementation of IPv6 for z/OS V1R4, using the OSA-Express features. It describes the basic functions of IPv6 and provides an example configuration.

For details on IPv6 support for Linux, refer to *Device Drivers and Installation Commands (October 24,2002)*, LNUX-1303-02, at the following URL:

`http://www10.software.ibm.com/developerworks/opensource/linux390/docu/lnuxdd02.pdf`

© Copyright IBM Corp. 2003. All rights reserved.

12.1 IPv6

The current IP address space is unable to satisfy the huge increase in the number of users or the geographical needs of the Internet expansion, let alone the requirements of emerging applications such as Internet-enabled personal digital assistants (PDAs), home area networks (HANs), Internet-connected automobiles, integrated telephony services, and distributed gaming.

IPv6 quadruples the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides more than enough globally unique IP addresses for every network device on the planet. The use of globally unique IPv6 addresses simplifies the mechanisms used for reachability and end-to-end security for network devices—functionality that is crucial to the applications and services that are driving the demand for the IPv6 addresses.

12.2 Advantages offered by IPv6

IPv6 provides improved traffic management in the following areas:

- 128–bit addressing

  Eliminates all practical limitations on global address ability. This means that private address space—and the network address translators (NATs) used between private intranet and public internet—are no longer needed.

- Simplified header formats

  Allow for more efficient packet handling and reduced bandwidth cost.

- Hierarchical addressing and routing

  Keep routing tables small and backbone routing efficient by using address prefixes rather than address classes.

- Improved support for options

  Changes the way IP header options are encoded, allowing more efficient forwarding and greater flexibility.

- Address auto-configuration

  Allows stateless IP address configuration without a configuration server. In addition, IPv6 brings greater authentication and privacy capabilities through the definition of new extensions, and integrated Quality of Service (QoS) through a new traffic class byte in the header.

12.3 IPv6 addressing

An IPv6 address is a 128-bit number written in colon hexadecimal notation. This scheme is hexadecimal and consists of eight 16-bit pieces of the address; for example, xxxxxxxx represents a single host on a single network.

Alternate notations described in RFC 2373 are acceptable; for example:

```
FEDC:BA98:7654:3210:FEDC:BA98:7654:321
```

There are three conventional forms for representing IPv6 addresses as text strings:

1. The preferred form is xxxxxxxx, where the x's are the hexadecimal value of the eight 16-bit pieces of the address.

   For example:

   ```
   FEDC:BA98:7654:3210:FEDC:BA98:7654:3210
   1080:0:0:0:8:800:200C:417A
   ```

> **Note:** It is not necessary to write the leading zeros in an individual field, but there must be at least one numeral in every field (except for the case described in item 2).

2. Due to some methods of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier, a special syntax is available to compress the zeros. The use of "**::**" indicates multiple groups of 16 bits of zeros. The "**::**" can only appear once in an address. The "**::**" can also be used to compress the leading and/or trailing zeros in an address.

   For example, the following addresses:

   ```
   1080:0:0:0:8:800:200C:417A (unicast address)
   FF01:0:0:0:0:0:0:101 (multicast address)
   0:0:0:0:0:0:0:1 (loopback address)
   0:0:0:0:0:0:0:0 (unspecified addresses)
   ```

   can be represented as:

   ```
   1080::8:800:200C:417A (unicast address)
   FF01::101 (multicast address)
   ::1 (loopback address)
   :: (unspecified addresses)
   ```

3. An alternative form that is sometimes more convenient when dealing with a mixed environment of IPv4 and IPv6 nodes is x:x:x:x:x:x:d.d.d.d, where the x's are the hexadecimal values of the six high-order 16-bit pieces of the address, and the d's are the decimal values of the four low-order 8-bit pieces of the address (standard IPv4 representation).

   For example:

   ```
   0:0:0:0:0:0:13.1.68.3
   0:0:0:0:0:FFFF:129.144.52.38
   ```

   or in compressed form:

   ```
   ::13.1.68.3
   ::FFFF:129.144.52.38
   ```

12.3.1 IPv6 TCP/IP Network part, also known as prefix

Designers have defined some address types and have left room for future definitions, as unknown requirements may arise. RFC 2373 (July 1998)/IP Version 6 Addressing Architecture defines the current addressing scheme. It defines different types of prefixes (and therefore, address types) as follows:

- Link local address type

  These are special addresses which will only be valid on a link of an interface. Using this address as destination, the packet would never pass through a router. It is used for link communications such as:

  – Anyone else here on this link?
  – Anyone here with a special address (for example, looking for a router)?

For example:

```
fe80:  (currently the only one in use)
fe90:
```

An address with this prefix is found on each IPv6-enabled interface after stateless auto-configuration (which is normally used).

Example: An address of FE80 : : interface ID was added to the OAT of our OSA-Express feature at start of the device. The interface ID is derived from the MAC address of the OSA-Express feature.

- Site local address type

  These are addresses similar to the RFC 1918/Address Allocation for Private Internets in IPv4 today, with the added advantage that everyone who uses this address type has the ability to use the given 16 bits for a maximum number of 65536 subnetworks; this is comparable to the 10.0.0.0/8 address in IPv4 today.

  Another advantage is that, since it is possible to assign more than one address to an interface with IPv6, you can also assign a site local address in addition to a global one.

  It begins with the following prefixes:

  ```
  fec0:  (most commonly used)
  fed0:
  fee0:
  ```

- 6bone test addresses

  These were the first global addresses defined and used for testing purposes; they all begin with the following prefix:

  ```
  3ffe:
  ```

- 6to4 addresses

  These addresses, designed for a special tunneling mechanism (RFC 3056/Connection of IPv6 Domains via IPv4 Clouds and RFC 2893/Transition Mechanisms for IPv6 Hosts and Routers), encode a given IPv4 address and a possible subnet. They begin with the following prefix:

  ```
  2002:
  ```

  For example, representing 192.168.1.1/5:

  ```
  2002:c0a8:0101:5::1
  ```

- Assigned by provider for hierarchical routing

  These addresses are delegated to Internet service providers (ISP) and begin with the following prefix:

  ```
  2001:
  ```

- Multicast addresses

  Multicast addresses are used for related services and always begin with the following prefix:

  ```
  ffxx:
  ```

  (xx is the scope value)

- Anycast addresses

  Anycast addresses are special addresses used to cover things like nearest DNS server, nearest DHCP server, or similar dynamic groups. Addresses are taken out of the unicast address space (aggregatable global or site-local at the moment). The anycast mechanism (client view) will be handled by dynamic routing protocols.

Note: Anycast addresses cannot be used as source addresses; they are only used as destination addresses.

– Subnet-router anycast address

A simple example of an anycast address is the subnet-router anycast address. Assuming that a node has the following global assigned IPv6 address:

```
3ffe:ffff:100:f101:210:a4ff:fee3:9566/64  <- Node's address
```

The subnet-router anycast address will be created blanking the suffix (least significant 64 bits) completely:

```
3ffe:ffff:100:f101::/64  <- subnet-router anycast address
```

12.4 Migrating to IPv6 on z/OS

For a summary of z/OS TCP/IP stack-related functions and the level of support provided in an IPv6 network, consult Table H-1 on page 232. This table will assist you in determining whether a given function is applicable to IPv6.

For more information on the related configuration statements for a particular function, refer to *z/OS Communications Server: IP Configuration Reference,* SC31-8775.

For more details on how IPv4 and IPv6 can co-exist in z/OS environment, refer to *z/OS V1R4.0 CS: IPv6 Network and Application Design Guide*, SC31-8885.

12.5 IPv6 implementation in z/OS

IPv6 support is built in z/OS 1.4, but is not enabled, by default. You will need to specify a NETWORK statement with AF_INET6 in your BPXPRMxx member.

To support IPv4 and IPv6, we added the NETWORK statement shown in Example 12-1 to our BPXPRMxx.

Example 12-1 BPXPRMxx NETWORK statement

```
NETWORK DOMAINNAME (AF_INET6)
DOMAINNUMBER(19)
MAXSOCKETS(2000)
TYPE(CINET)
```

The TYPE option in our case is CINET, because we are using multiple TCP/IP stacks.

Note: The BPXPRMxx member may be updated dynamically via the z/OS command **SETOMVS RESET=(xx)**. After this reset, you will receive the following message:

```
BPXF203I DOMAIN AF_INET6 WAS SUCCESSFULLY ACTIVATED.
```

You must *recycle* the TCP/IP stack to pick up the change.

For more details on the definitions required in BPXPRMxx to provide a dual IPv4 and IPv6 stack, refer to *z/OS V1R4 CS: IP Configuration Guide*, SC31-8775, section “Defining TCP/IP as a UNIX System Services physical file system (PFS)”.

Important: IPv6 is only supported by OSA-Express Ethernet (GbE and FENET) in QDIO mode, with the zSeries LIC level 3 or later. You can verify the OSA-Express code level using a VTAM command, see Figure 12-1 on page 173.

VTAM definitions

Since our OSA-Express FENET will be operating in QDIO mode, a TRLE is required (see Example 12-2).

Example 12-2 TRLE for OSA-Express FENET

```
OSAFE09  VBUILD TYPE=TRL
TRLE2180 TRLE  LNCTL=MPC,
               READ=2180,
               WRITE=2181,
               DATAPATH=2182,
               PORTNAME=OSA2180,
               MPCLEVEL=QDIO
```

TCP/IP definitions

We added one INTERFACE statement for the OSA-Express FENET feature to support IPv6. This statement merges the DEVICE, LINK, and HOME definitions into a single statement. There are a number of different parameters associated with the INTERFACE statement. To determine which of them best fits your requirements, refer to *z/OS Communications Server IP Configuration Reference Version 1 Release 4,* SC31-8776.

We used the following syntax:

INTERFace *interfname* **DEFINE** *linktype* **PORTNAME** *portname* **IPADDR** *ipaddr*

interfname	Specify a name for the interface no more than 16 characters in length.
linktype	Must be IPAQENET6. This is the only DLC that currently supports IPv6.
portname	Specified in the VTAM TRLE definition for the QDIO interface.
ipaddr	Optional for link type IPAQENET6. If not specified, TCP/IP will enable auto-configuration for the interface. If used, one or more prefixes or full IPv6 addresses can be specified.

Note: To configure a single physical device for both IPv4 and IPv6 traffic, you must use DEVICE/LINK/HOME for the IPv4 definition and INTERFACE for the IPv6 definition, such that the PORTNAME value on the INTERFACE statement matches the devicename on the DEVICE statement.

In the TCP/IP profile for SC64 (see Example 12-3 on page 173) we added the INTERFACE statement to implement IPv6 on OSA2180 (portname).

Example 12-3 TCP/IP profile for SC64 to support IPv4 and IPv6

```
DEVICE OSA2180 MPCIPA PRIRouter
LINK   SC642180LINK  IPAQENET  OSA2180
INTERFACE L2180V6 DEFINE IPAQENET6 PORTNAME OSA2180
IPADDR FEC0:0:0:1::3302
       FEC0:0:0:1001::3302

BEGINROUTES
ROUTE FEC0::0/10 = L2180V6 MTU 1492
ROUTE 192.168.1.0/24   =     SC642180LINK     MTU 1492
START OSA2180
START L2180V6
START OSA21A0
START OSA2360
```

12.5.1 Verification

This section shows the procedures we used to verify our environment.

Since the TRLE must be active before the interface is started, first we ensured that the TRLE was in an active state by using the following command:

```
D NET,ID,TRL,TRLE=osatrlename
```

The results are shown in Figure 12-1. You can also verify the OSA-Express code level with this command.

```
D NET,TRL,TRLE=TRLE2180
 DISPLAY ACCEPTED
 NAME = TRLE2180, TYPE = TRLE 631
 STATUS= ACTIV, DESIRED STATE= ACTIV
 TYPE = LEASED             , CONTROL = MPC , HPDT = YES
 MPCLEVEL = QDIO       MPCUSAGE = SHARE
 PORTNAME = OSA2180    LINKNUM =   0   OSA CODE LEVEL = 0320
 HEADER SIZE = 4096 DATA SIZE = 0 STORAGE = ***NA***
 WRITE DEV = 2181 STATUS = ACTIVE     STATE = ONLINE
 SIZE = 4092 DATA SIZE = 0 STORAGE = ***NA***
 READ  DEV = 2180 STATUS = ACTIVE     STATE = ONLINE
 DATA  DEV = 2182 STATUS = ACTIVE     STATE = N/A
```

Figure 12-1 OSA-Express status and code level

The following message (for the IPv6 interface) was written to the z/OS console as the TCP/IP stack was initializing: `EZZ4340I INITIALIZATION COMPLETE FOR INTERFACE L2180V6.`

Figure 12-2 on page 174 shows the output of NETSTAT DEV with the IPv6 interface shown as ready.

```
INTFNAME: L2180V6         INTFTYPE: IPAQENET6 INTFSTATUS: READY
      NETNUM:  0   QUESIZE: 0   SPEED: 0000000100
      BYTESIN: 0                     BYTESOUT: 1304
      MACADDRESS: 0006296CA5BC
      DUPADDRDET: 1
      CFGROUTER: NON                 ACTROUTER: NON
      CFGMTU:  NONE                  ACTMTU:   1492
    MULTICAST SPECIFIC:
      MULTICAST CAPABILITY: YES
```

Figure 12-2 NETSTAT DEV

If the device does not have a LnkStatus or IntfStatus of Ready, this must be resolved before continuing. There are several things that might cause the LnkStatus or IntfStatus to not be ready. For example, the device might not be defined to z/OS correctly, the device might not be defined in the TCP/IP profile correctly, and so on.

Figure 12-3 depicts the results of the OAT, built after starting devices OSA2180 and L2180V6.

```
LP 09 (A9)
00(2180) MPC       n/a     TRLE2180   (QDIO control)       SIU    ALL
01(2181) MPC       n/a     TRLE2180   (QDIO control)       SIU    ALL
02(2182) MPC       00  PRI4 No6  TRLE2180   (QDIO data)    SIU    ALL
                           9.12.2.36
                           192.168.1.164
                           FEC0::1:0:0:0:3302
                           FEC0::1001:0:0:0:3302
```

Figure 12-3 Query CHPID via OSA/SF

We used the **ping** command from various hosts within the network to verify connectivity for IPv4 and IPv6.

```
ping 192.168.1.164
   CS V1R4: Pinging host 192.168.1.164
   Ping #1 response took 0.048 seconds.
   READY

ping FEC0:0:0:1::3302
   CS V1R4: Pinging host FEC0:0:0:1::3302
   Ping #1 response took 0.001 seconds
   READY
```

A

Commands

In this appendix, we list the commands we used when setting up and verifying the various LAN environments described in this redbook.

© Copyright IBM Corp. 2003. All rights reserved.

z/OS commands

For a complete list and description of TCP/IP Console and TSO commands, refer to *IP Systems Administration Commands*, SC31-8781.

z/OS TCP/IP operator commands

Command	Description
D U,,,dddd[a],nnn	Status of the device(s).
D U,,ALLOC,dddd,nnn[b]	Show who the device(s) is allocated to.
D M=DEV(dddd)	Status of paths defined to a device.
D M=CHP	Status and type of all CHPIDs defined to the OS/390.
D M=CHP(cc[c])	Status of path to defined devices.
D A,L	List the jobs running in the system.
D IOS,MIH	Display MIH values for all devices.
SETIOS MIH,DEV=dddd,TIME=mm:ss	Set MIH time for specified device.
V dddd-dddd,ONLINE	Vary a device online.
V dddd-dddd,OFFLINE	Vary a device offline.
CF CHP(cc),ONLINE	Configure a CHPID online.
CF CHP(cc),OFFLINE	Configure a CHPID offline.
D TCPIP	List TCP/IP stacks that have been started since the last IPL and stack status.
D TCPIP,*tcpipstack*,NETSTAT,ARP	Display contents of ARP cache for the TCP/IP stack.
D TCPIP,*tcpipstack*,NETSTAT,DEV	Status of a device(s) or interface(s) defined in TCP/IP stack profile.
D TCPIP,*tcpipstack*,NETSTAT,HOME	Display the home IP address(es) defined in TCP/IP stack profile.
D TCPIP,*tcpipstack*,NETSTAT,ND	Display the contents of the IPv6 neighbor cache.
D TCPIP,*tcpipstack*,NETSTAT,ROUTE	Display the routing information for the TCP/IP stack.
V TCPIP,*tcpipstack*,PURGECACHE,*linkname* (introduced with z/OS 1.4)	Purge ARP cache for the specified adapter (linkname or intfname (IPv6) from NETSTAT,DEV).
V TCPIP,*tcpipstack*,START,*tcpipdev*	Start a device or interface (IPv6) defined in TCP/IP stack.
V TCPIP,*tcpipstack*,STOP,*tcpipdev*	Stop a device or interface (IPv6) defined in TCP/IP stack.

a. dddd...device number
b. nnn.....number of devices to be displayed
c. cc.......CHPID number

TCP/IP TSO commands

Command	Description
NETSTAT ?	Display Netstat options.
NETSTAT ARP ALL	Display ARP cache.
NETSTAT DEV	Display the TCP/IP devices and links
NETSTAT HOME	Display the TCP/IP Home IP addresses.
NETSTAT GATE	Display the TCP/IP Gateway addresses.
PING ipaddress	Perform one PING to specified address.
TRACERTE ipaddress	Trace router hops to specified address.
OBEYFILE	Execute selected TCP/IP profile statements.

VTAM commands

Command	Description
D NET,VTAMOPTS	Display current VTAM start options.
F vtamname,VTAMOPTS, optionname=value	Modify current VTAM options. (vtamname is STC name, optionname is from VTAMOPTS)
D NET,MAJNODES	Display the VTAM major nodes.
D NET,id=*xxxxxxx*,E	Display the information about a specified ID (for example, a Line, PU, or LU).
D NET,TRL	Display the list of TRLEs.
D NET,TRL,TRLE=*trlename*	Display the status of a TRLE.
V NET,ID=ISTTRL,ACT,UPDATE=ALL	Delete inactive TRLEs from the TRL list.
V NET,ID=*mnodename*,ACT	Activate a major node.
V NET,ID=*mnodename*,INACT	Deactivate a major node.

Note: If your static TRLE definition is incorrect, you should be aware that an active TRLE entry cannot be deleted. What you can do is:

1. Vary activate the TRL node with a blank TRLE to cause the deletion of previous entries.
2. Code the correct TRL with correct TRLE entries and definition.
3. Vary activate this corrected TRL/TRLE node.

z/VM Commands

Refer to *z/VM V4R3.0 CP Command and Utility Reference*, SC24-6008 for detailed information.

z/VM TCP/IP operations commands

Command	Description
Q MITIME	Display MIH times for devices.
Q OSA ACTIVE|ALL	Display status of OSA devices
Q rdev|rdev-rdev	Display status of real device(s).
Q PATHS rdev|rdev-rdev	Display path status to real device(s) (PIM, PAM, LPM).
Q CHPID cc	Display real CHPID status.
SET MITIME rdev|rdev-rdev mm:ss	Set MIH time for device(s).
VARY OFF|ON rdev|rdev-rdev	Vary device(s) off or online.
VARY OFF|ON PATH cc FROM|TO rdev|rdev-rdev	Change the status of a path to device(s).
VARY OFF|ON CHPID cc	Configure a CHPID off or on to both hardware and software.

z/VM TCP/IP commands

Refer to *z/VM TCP/IP Level 430 User's Guide*, SSC24-6020, for detailed information.

Command	Description
NETSTAT ?	Display Netstat options.
NETSTAT ARP	Display ARP cache.
NETSTAT DEV	Display the TCP/IP devices and links
NETSTAT HOME	Display the TCP/IP Home IP addresses.
NETSTAT GATE	Display the TCP/IP Gateway addresses.
NETSTAT OBEY START|STOP DEV	Start or stop the device name identified in NETSTAT DEV output.
IFCONFIG (z/VM4.3)	Display the TCP/IP devices and links (Like NETSTAT DEV, but also has other uses; see note).
PING ipaddress	Perform one PING to specified address.
TRACERTE ipaddress	Trace router hops to specified address.
OBEYFILE	Execute selected TCP/IP profile statements.

Note: IFCONFIG can be used to temporarily modify network interfaces in the current TCP/IP stack. Refer to *z/VM TCP/IP Level 430 Planning and Customization*, SC24-6019, for detailed uses. For more detailed information on the other z/VM TCP/IP commands, refer to *z/VM TCP/IP Level 430 User's Guide*, SC24-6020.

Linux for zSeries TCP/IP commands

The following Linux commands should be entered in lower case, as shown.

Command	Description
arp	Display ARP cache. Use -? for options.
dmesg imore	Display device assignments (and more) at kernel initialization.
netstat -i	Display interface table.
netstat -r	Display routes.
ifconfig	Display network interfaces. (LO, eth0, tr0, and so on).
ifconfig interface upidown	Start or stop network interface.
ping ipaddress	Perform one PING to specified address.
route	Display routes.
traceroute ipaddress	Trace router hops to specified address.

B

HMC and SE tasks for OSA-Express

In this appendix, we discuss the tools that can be used from a Hardware Management Console (HMC) or the Support Element (SE) of a 9672 G5/G6 or zSeries server system environment.

In the first part, we describe advanced facilities for the OSA-Express channels. The advanced facilities are now available directly on the HMC as a task under CPC Operational Customization.

The second section can be used for CHPID Off/On guidance.

© Copyright IBM Corp. 2003. All rights reserved.

HMC advanced facilities for OSA-Express

The advanced facilities (OSA/AF) functions available on the HMC are as follows.

Note: These functions are used for troubleshooting under the guidance of IBM Product Engineering.

- Set Trace Buffer
- Read Trace Buffer
- Export Trace/Dump file to diskette
- Card-specific advanced facilities, with the following subdivision:
  - Enable or disable ports
  - Query port status
  - Run port diagnostics
  - View port parameter
  - Display or alter MAC address
  - Set ethernet mode (only for FENET)

Note: Online help is available for each function on the HMC or the SE. Figure B-6 on page 186 shows an example of the online help screens. The online help can be activated in two ways:

- By clicking **Help** on the active window
- By pressing F1 on the keyboard

To get access to the advanced facilities, you must be logged on at the HMC in the system programmer mode. Then follow the listed steps to start the individual functions.

1. Open the CPC group work area.
2. Switch to CPC Operational Customization in the Task Area.
3. Select the appropriate CPC and drag and drop it to the OSA Advanced Facilities Task.
4. A pull-down menu with all OSA-adapters (OSA-2 and OSA-Express) appears at the HMC console, as shown in Figure B-1 on page 183.
5. Select the OSA CHPID you want to work with and click **OK**.

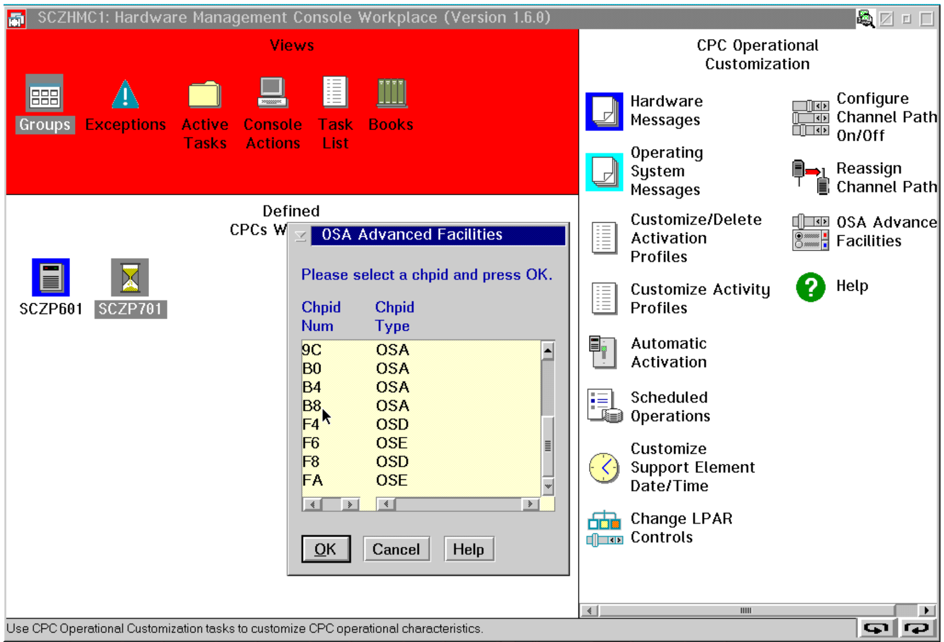

Figure B-1 HMC OSA advanced facilities panel

The Standard Channel Advanced Facilities panel is displayed, as shown in Figure B-2.

Figure B-2 Standard Channel Advanced Facilities panel

Trace functions for OSA-Express

These trace functions are used for troubleshooting under the guidance of IBM Product Engineering.

Set Trace Mask on the OSA-Express Card

On the Standard Channel Advanced Facilities panel, select **Set Trace Mask.**

Figure B-3 shows the Set Trace Mask panel.

Note: The trace mask should be changed only under the guidance of IBM Product Engineering.

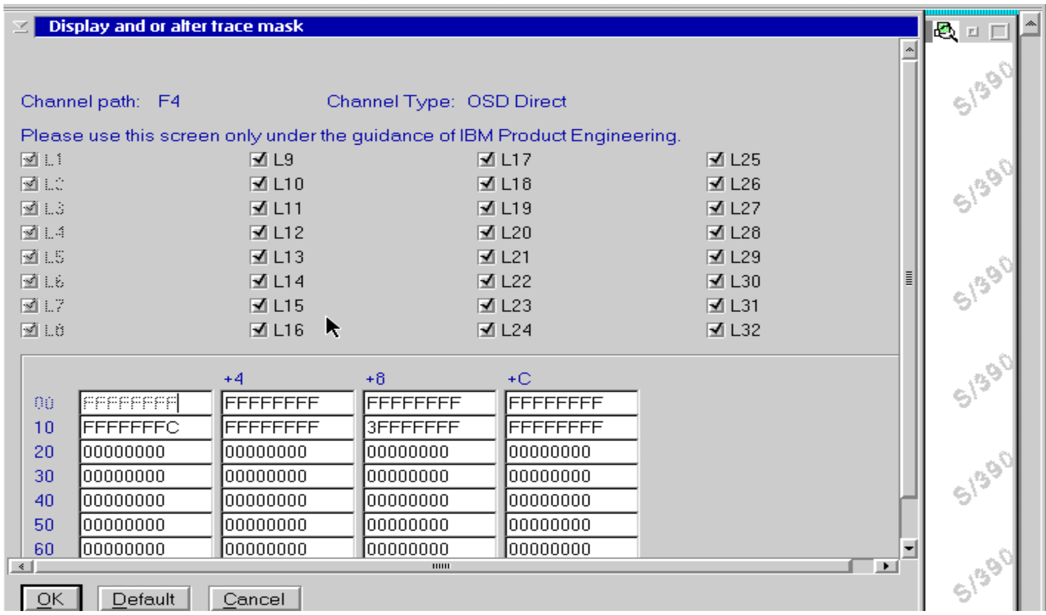

Figure B-3 Display and/or alter trace mask panel

Read Trace Buffer on the OSA-Express Card

On the Standard Channel Advanced Facilities panel, select **Read Trace Buffer.**

The Read Trace Buffer panel is displayed, as shown in Figure B-4.

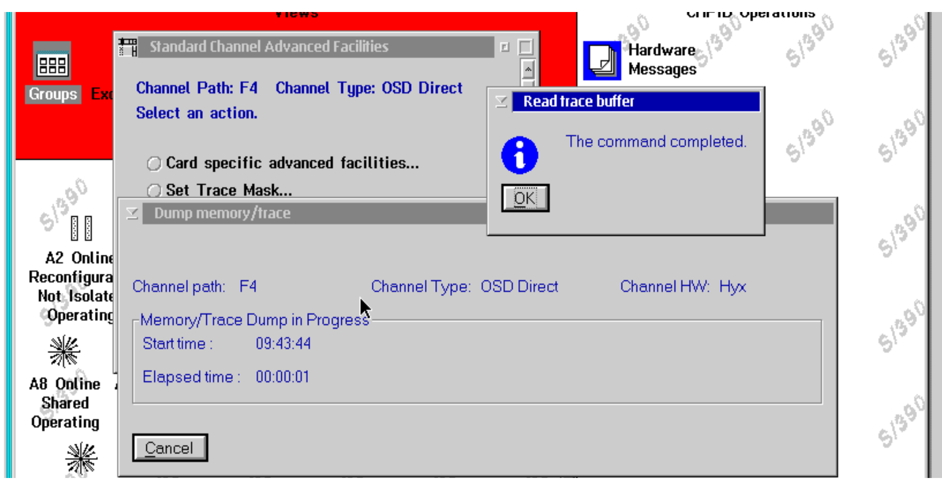

Figure B-4 Read trace buffer confirmation panel

The Read Trace Buffer function collects all data necessary for troubleshooting.

Export Trace/Dump file to diskette

On the Standard Channel Advanced Facilities panel, select **Export Trace/Dump file to diskette**.

After filling in the required information and clicking **OK**, the data is sent to the file you just defined.

Hardware functions for OSA-Express

On the Standard Channel Advanced Facilities panel, select **Card specific advanced facilities...**.

This opens a window with all advanced facilities, as shown in Figure B-5.

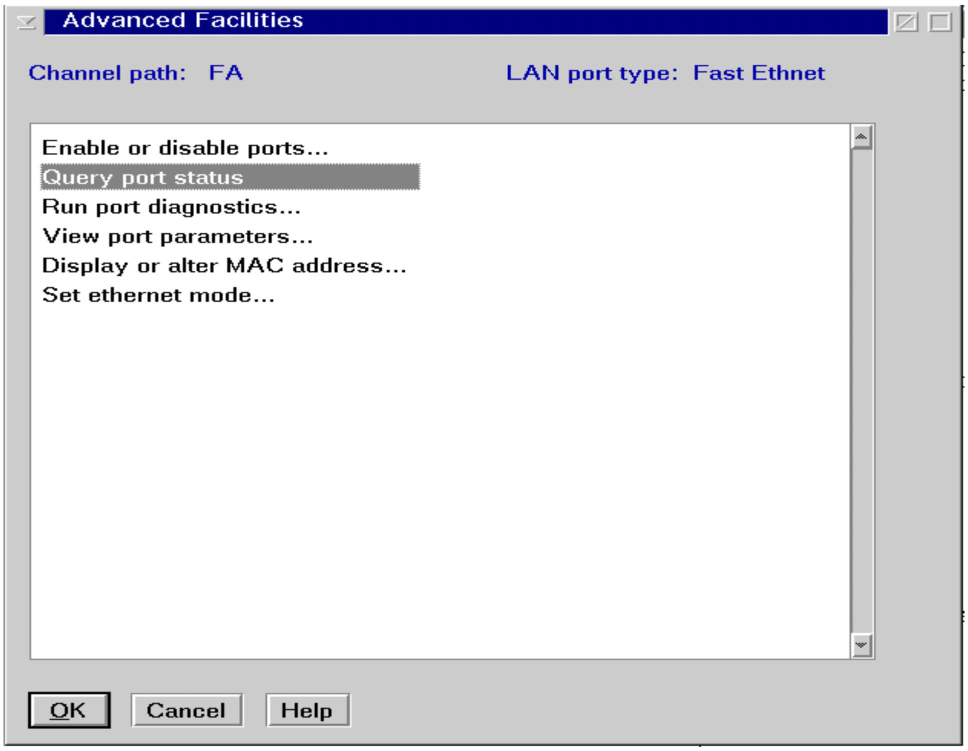

Figure B-5 Card specific advanced facilities

Enable or disable ports on OSA-Express Cards

Select **Enable or disable ports** and click **OK.** The Enable or disable ports panel is displayed, as shown in Figure B-6 on page 186.

Note: Enabling or disabling the port is also possible with an OSA/SF function.

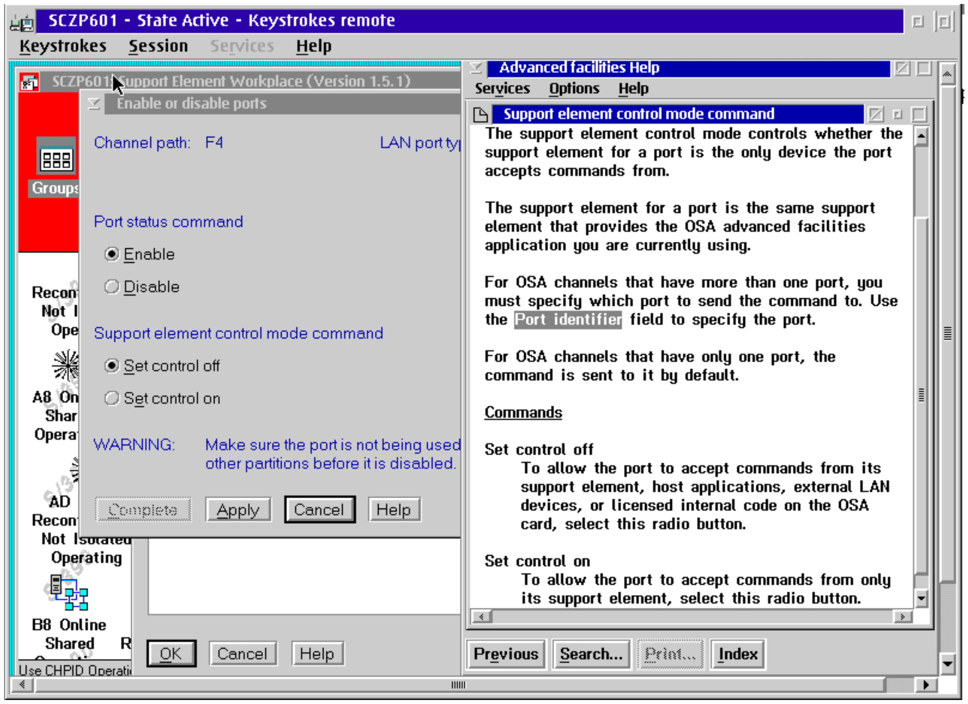

Figure B-6 Enable or disable ports panel

The left part of Figure B-6 shows the Enable or disable ports panel. The right part is an example of the online help for the support element, which can be activated in two ways:

- By clicking **Help** on the active window
- By pressing F1 on the keyboard

There are two tasks related to the enable or disable port function. The first task enables or disables the port. The second task sets the control of enabling and disabling the port either to its Support Element, or to both the Support Element and OSA/SF.

Query the port status on an OSA-Express Card

Select **Query port status** from the Card specific advanced facilities panel and click **OK.** Figure B-7 on page 187 shows the Query port status panel.

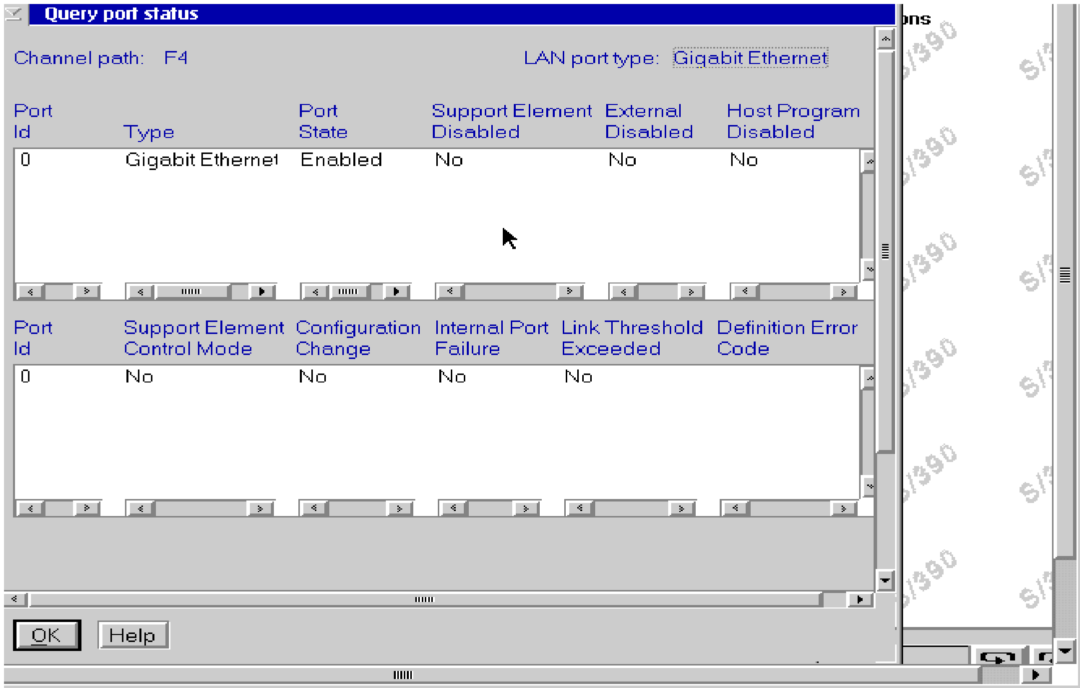

Figure B-7 Query port status panel

You can exit the panel by clicking **OK**.

Run diagnostics on an OSA-Express Card

Select **Run port diagnostics** from the Card specific advanced facilities panel and click **OK.**

Note: The port diagnostics can only be used if the port is disabled.

View port parameters on an OSA-Express Card

Select **View port parameters** from the Card specific advanced facilities panel and click **OK.**

The View port parameters panel is displayed, as shown in Figure B-8 on page 188.

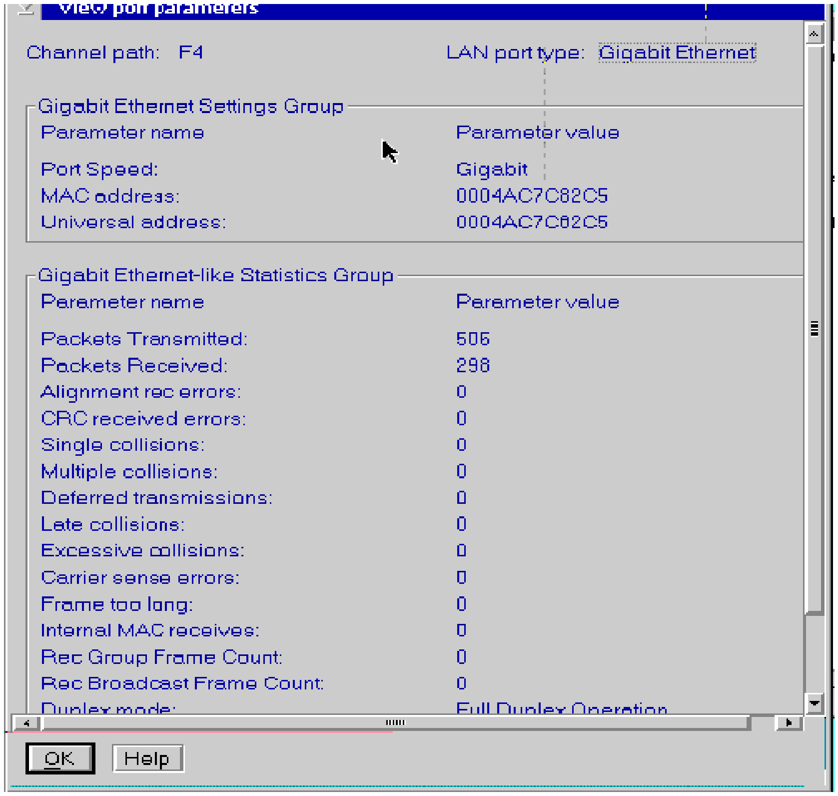

Figure B-8 View port parameters panel

You can use the scroll buttons of the window to view all port parameters. To exit the panel, click **OK**.

Display or alter MAC address on an OSA-Express Card

Select **Display or alter MAC address** from the Card specific advanced facilities panel and click **OK.**

Figure B-9 on page 189 shows a Display or alter MAC address panel.

When you change the displayed MAC address, click **Apply** to activate the new address. To exit the panel without any changes, click **Cancel**.

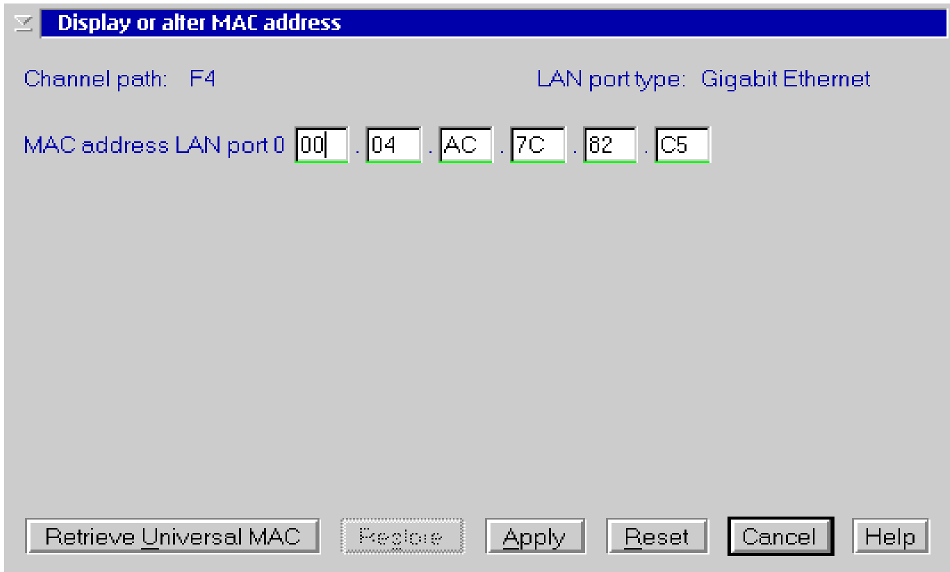

Figure B-9 Display or alter MAC address panel

Settings for speed and mode on an OSA-Express FENET card:

Select **Set ethernet mode** from the Card specific advanced facilities panel and click **OK.**

Figure B-10 shows the screen with the different settings that can be selected.

When you change the current settings, click **Apply** to make the new settings active. To exit the panel without any changes, click **Cancel**.

Note: The settings made here override the capability of the OSA-Express FENET auto negotiation facility.

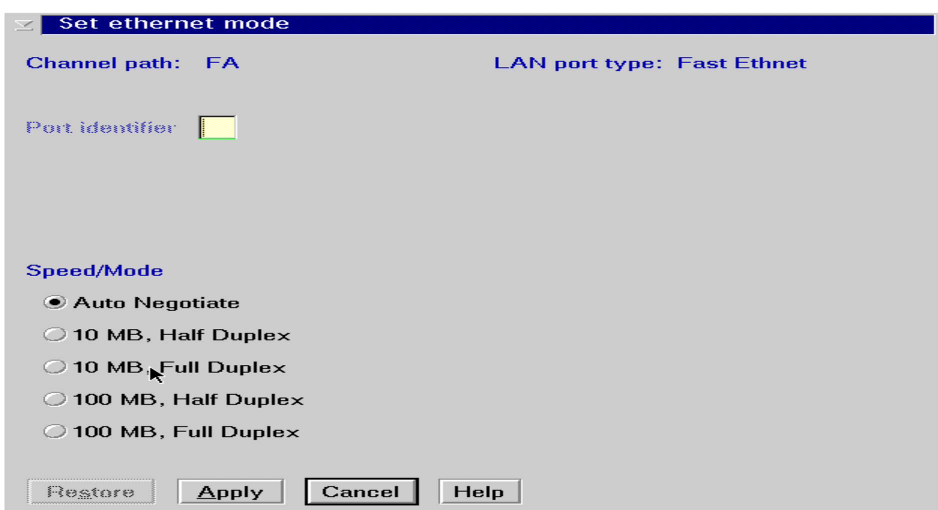

Figure B-10 Ethernet Speed/Mode settings panel

OSA reset to defaults

When you directly log on to the Support Element of an S/390 system, an additional function is displayed as shown in Figure B-11 on page 190.

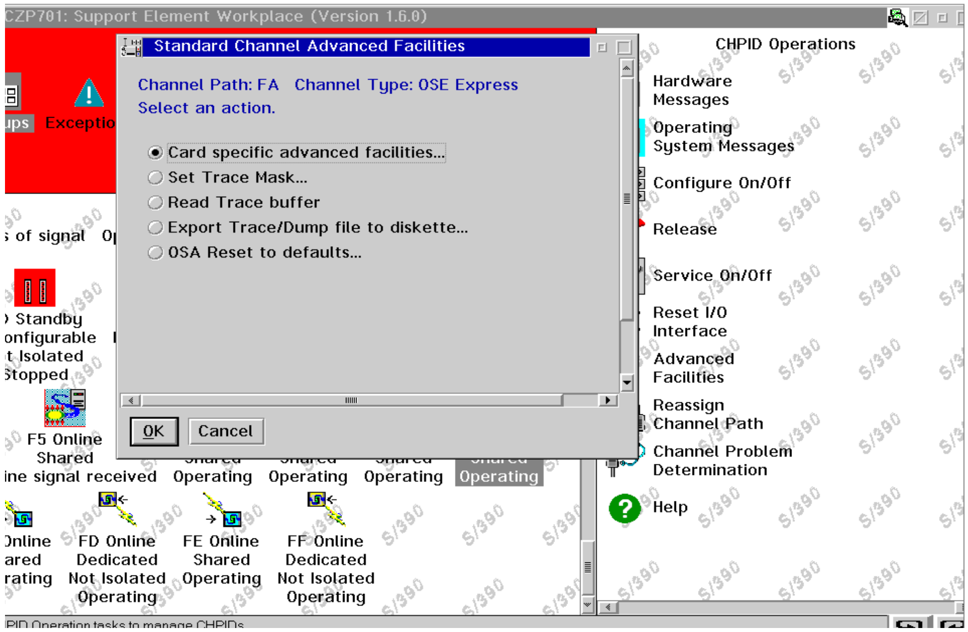

Figure B-11 Advanced Facilities panel at the Support Element

With the “OSA Reset to defaults” function, you can reset all your customized entries in the OSA-Express adapter to the default settings, including the OAT.

The sequence to get access to the CHPID functions is described in the following section. The Advanced Facilities icon from the Support Element is placed under the CHPID Operation in the Tasks Area. To start the advanced facilities, you have to drag the desired CHPID icon and drop it over the Advanced Facilities icon.

Configuring OSA-Express Channels on/off

With the new design of the OSA-Express adapters, there is no longer a need to configure the OSA CHPID offline and online to activate an OAT after changing it. However, for recovery reasons it still may be necessary. Normally, you would use z/OS commands on the operator console to configure a channel off or on.

Since you are not able to configure channels offline to the whole system with one command from an operator console when you are running in LPAR mode, we describe the procedure that is used from the Hardware Management Console (HMC) to configure a CHPID off or on for all LPARs.

Logging on to the Support Element

To get access to all channel functions, a session must be started from the Hardware Management Console (HMC) to the Support Element (SE), as follows:

1. From the Defined CPCs Work Area on the HMC:
   - Drag and drop the appropriate **CPC Object** onto the **Single Object Operations** icon in the CPC Recovery task list.
   - When the confirmation panel appears, click **OK**.

   After a short time, the HMC displays the Support Element Workplace, with the Group Work Area, the Views Area, and the Daily Task Area. Figure B-12 on page 191 shows an example of that display. (The *S/390* in the background of the workplace is the indicator for working on the SE.)

2. In the Groups Work Area:

   Double-click **CPC**.

   The CPC Work Area is displayed, as shown in Figure B-12.

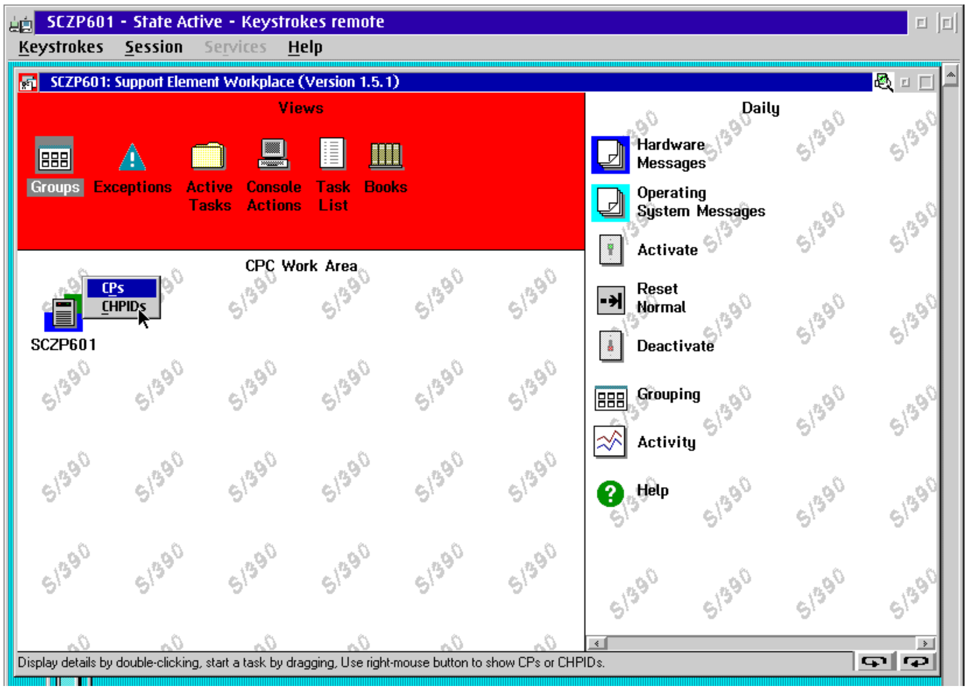

Figure B-12 Support Element CPC Work Area

CHPID Configure on/off

The OSA-Express CHPID is configured on or off by doing the following:

1. In the CPC Work Area:

   Click the right button on the **CPC object** and select **CHPIDs** from the pull-down menu.

   The CPC CHPIDs Work Area is displayed (see Figure B-13 on page 192).

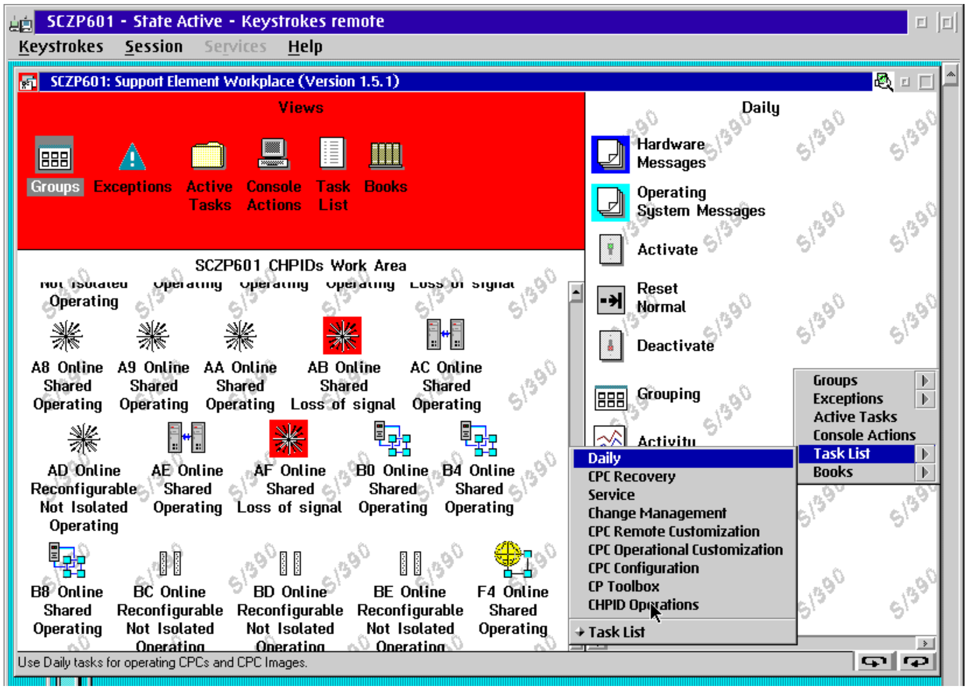

Figure B-13 Support Element CHPIDs Work Area

2. In the CPC CHPIDs Work Area:

- ► Use the scroll bar to display the required CHPID object.
- ► Open the CHPID Operations Task Area in one of two ways:
  - – Use the rotate arrow push buttons below the task list to rotate through other task lists until the CHPID Operations task list is displayed.
  - – Click anywhere in the workplace with the right mouse button and click the arrow next to **Task List** from the pull-down menu. From the next menu, select **CHPID Operations.** (The pull-down menus are shown in Figure B-13.)

After the selection of the CHPID Operations, a display appears, as shown in Figure B-14 on page 193.

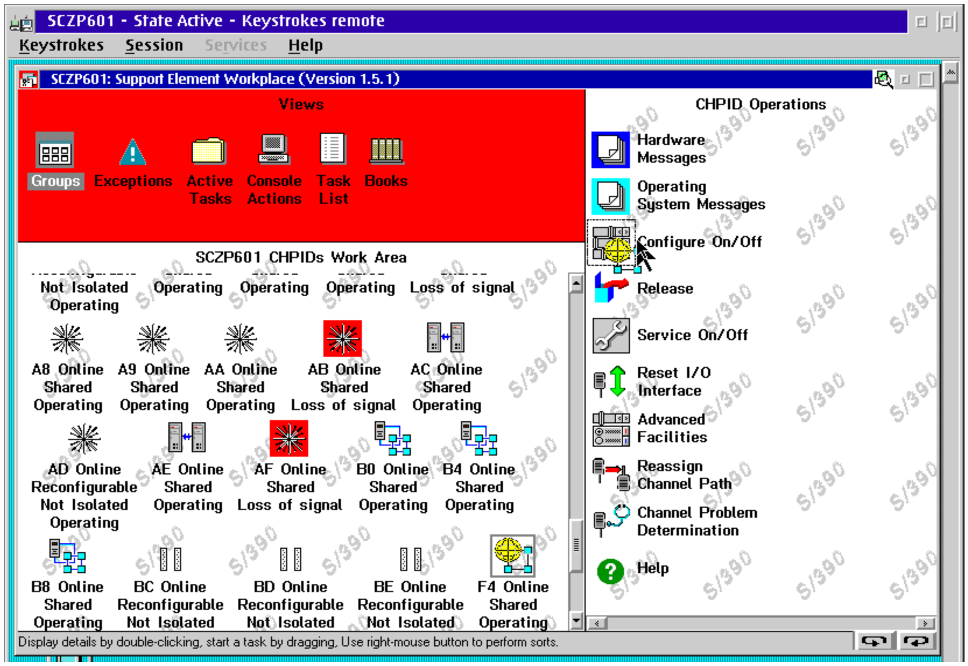

Figure B-14 CHPID Operations

3. In the CPC CHPIDs Work Area:

   Drag the appropriate OSA-Express **CHPID object** and drop it onto the **Configure On/Off** icon. (An alternative is to highlight the CHPID and then double-click the **Configure On/Off** icon.)

   The Configure On/Off panel is displayed (see Figure B-15 on page 194).

   The Current state field shows the status of the channel for each logical partition.

Note: The state Standby is the indicator of an offline CHPID to the LPAR.

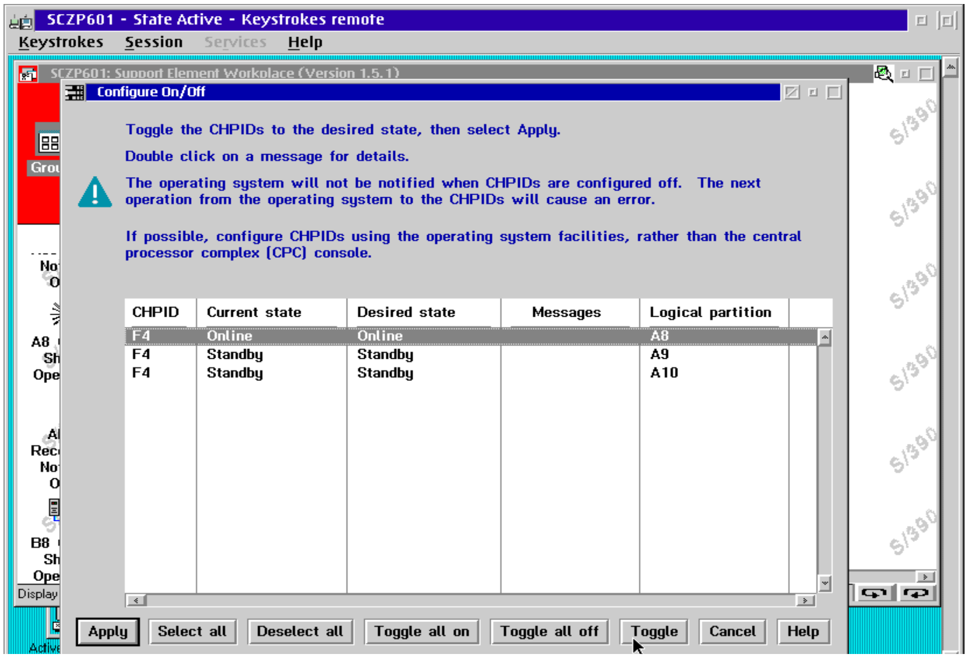

Figure B-15 Configure On/Off panel

4. On the Configure On/Off panel:
   – Read the warning information regarding configuring channels from the Support Element function rather than from an available operating system.

     Reminder: You should only toggle the CHPID if you cannot configure it offline and online from the operating system in each of the LPARs to which the CHPID is assigned.

   – If the channel should be set off to all partitions, **Toggle all off** can be used.
   – If the channel should be set on to all partitions, **Toggle all on** can be used.
   – If the channel should *not* be changed in all partitions, then the affected partitions must be selected. Clicking Toggle will change the desired state of the channel.

     Attention: Before you click Apply in the next step, ensure that only the CHPID(s) you want toggled are highlighted.

   – In all cases, **Apply** must be pressed to immediately change the desired channel state.

The Configure On/Off Progress panel is briefly displayed, with the message: `In progress`. When the message changes to: `Complete`, then click **OK**.

Logging off from the Support Element

When the work with the channels is finished, the Support Element session must be closed, as follows:

- Double-click **Console Action** in the Views Area to open the Console Action Work Area. A display similar to Figure B-16 on page 195 appears.

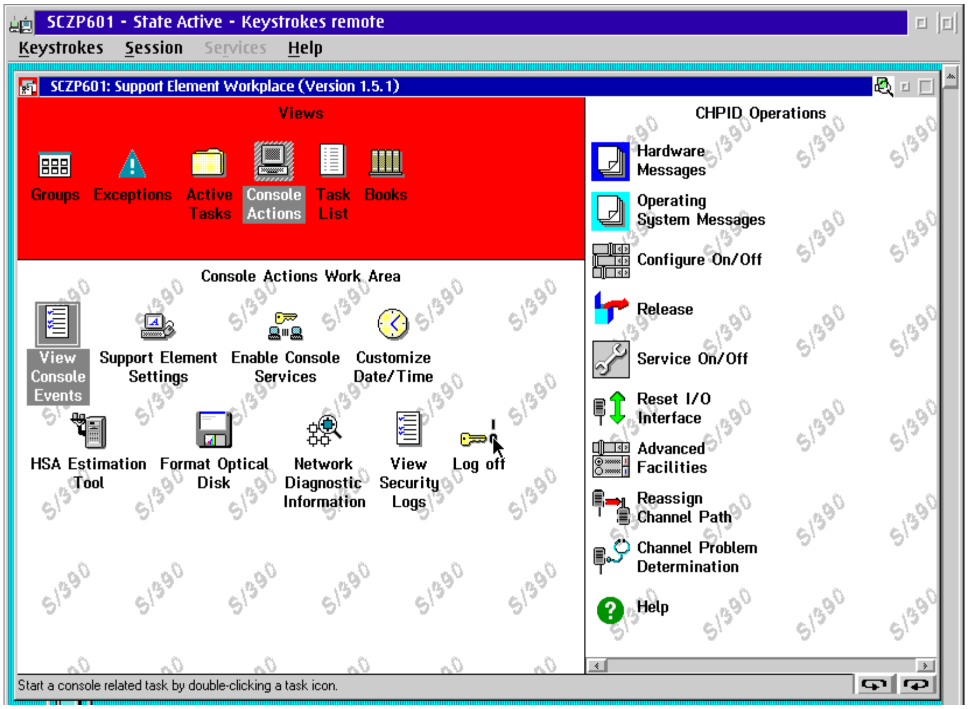

Figure B-16 Log off the Support Element

- Double-click **Log off** in the Console Action Work Area to end the Support Element session.

C

RMF in an OSA-Express environment

Resource Measurement Facility (RMF) measures and reports on the performance and availability of such system resources as processors, channel paths, devices, and storage.

RMF has extended the supported channel types to include the OSA-Express, and to provide statistics about the bus utilization and the transfer rate for both the Read and Write operations in the Channel Activity Report.

© Copyright IBM Corp. 2003. All rights reserved.

RMF for OSA-Express

RMF reports that are associated with OSA-Express are:

- Monitor II Channel Path Activity Report
- Monitor I/Postprocessor Channel Path Activity Report
- Postprocessor Overview Reporting/Recording

Measurements are contained in SMF Record Type 79(13) and are available through the ERBSMFI interface.

SMF has been implemented with the field R79CACR, which contains the channel path acronym.

For OSA-Express CHPIDs (device types OSD and OSE), IBM has expanded the support for performance information by supporting the Extended CPMF, which allows the user to better understand what is occurring within the three main components of OSA-Express:

- Processor utilization
- Physical PCI bus utilization
- The bandwidth per port (both read and write directions)

The Channel Path Measurement Facility (CPMF) provides information about CHPIDs on a per-image basis.

The Extended CPMF offers the enhancements that supply more data about the new channel types.

RMF Monitor II output

The new support for OSA-Express CHPIDs (device types OSD and OSE) can be found in an RMF Monitor report called the Channel Activity Report.

Note that RMF must be defined with the option DEVICE(COMM) in member ERBRMFxx from SYS1.PARMLIB.

For more information, refer to *OS/390 Resource Measurement Facility User's Guide*, SC28-1949.

The Channel Activity Report

To view the OSA-Express channel activity, do the following:

1. In the ISPF Primary Option Menu, enter `6` (Command).
2. In the ISPF Command Shell, enter `RMFMON 2`, then press Enter.
3. In the RMF Display Menu, press F4.

Figure 12-4 on page 199 shows an example of the Channel Activity Report.

```
16:42:43  CHANNEL UTILIZATION(%)   READ(B/S)  WRITE(B/S)  MSG  MSG SEND RECV
ID NO  G  TYPE  S PART  TOT  BUS   PART  TOT  PART  TOT  RATE SIZE FAIL FAIL
08        OSD   Y  0.0 12.1 16.0      0  21K     0 118K
09        OSD   Y  6.8 11.9 15.6      0    0     0    0
0A        OSD   Y  0.0  3.6 15.8      0   3K     0  68K
0B        OSD   Y  7.9 11.9 13.0      0    0     0    0
0C        OSD   Y  0.0  0.0  8.6      0    0     0    0
0D        OSD   Y  0.0  0.0  8.6      0    0     0    0
0E        OSE   Y  0.0  1.4 12.4      0    0     0    0
10        OSD   Y  0.8  3.5 15.6      0  717     0    0
11        OSD   Y  0.0 12.4 16.3      0 127K     0  24K
12        OSD   Y  0.0  3.5 15.6      0    0     0    0
13        OSE   Y  0.0  0.4 10.8      0    0     0    0
14        OSD   Y  0.0  0.0  8.6      0    0     0    0
15        OSD   Y  0.0  0.0  8.6      0    0     0    0
16        OSD   Y  0.0  3.5 15.6      0    0     0    0
17        OSE   Y  0.0  0.4 10.8      0    0     0    0
18        OSE   Y  0.0  0.4 10.8      0    0     0    0
19        OSD   Y  0.0  3.1 15.6      0    0     0    0
```

Figure 12-4   Channel Activity Report

4. To exit the report, enter `QUIT`.

For more information, see *z/OS V1R4.0 Resource Measurement Facility Report Analysis,* SC33-7991.

D

Using the OSA/SF REXX interface

This appendix describes the steps required to configure an OSA-Express feature using the TSO REXX interface for the non-QDIO modes.

© Copyright IBM Corp. 2003. All rights reserved.

Creating the Open System Adapter (OSA) configuration

The OSA configuration in the TSO environment is built with two z/OS sequential data sets containing statements describing the OSA-Express configuration and the OSA Address Table.

The steps required to build an OSA-Express feature configuration areas follows:

1. Create two skeleton definitions: one for the configuration file, and the other with the OAT entries.

   In order to ensure the correct format of the files, first use the `GET Configuration File` and the `GET OSA Address Table` commands to retrieve the current files into sequential data sets, or copy the skeletons from IOA.SIOASAMP.

   The sample ATM configuration file is named IOAATME. The sample FENET configuration file is named IOAFENET. The sample Token Ring configuration files is named IOATR.

   There are also two sample OAT datasets. They are IOAOSHRS (for shared SNA) and IOAOSHRT (for shared TCP/IP).

2. Copy the skeleton definitions to work data sets.
3. Update, delete, or create the configuration and OAT entries as required.
4. Use the `Configure OSA CHPID` command to activate your changes.

Creating the OSA-Express configuration file

1. Create a skeleton configuration file. (An example of a FENET configuration file is shown in Example D-4 on page 207. An example of a Token Ring configuration file is shown in Example D-5 on page 207. An example of an ATM configuration is shown in Example D-6 on page 208.)

   A detailed description of all parameters is included in the example configuration files.

2. Update the skeleton with your configuration requirements.

Creating the OSA-Express OAT file

1. Create a skeleton OAT by using the OSA/SF `GET_OAT` command or by using the sample OAT file from the IOA.SIOASAMP dataset.

   The names of the sample OAT files are:

   IOAOSHRA TCP/IP SNA MPC shared port

   IOAOSHRT TCP/IP shared port (see Example D-1 on page 203)

   IOAOSHRS SNA shared port (see Example D-2 on page 203)

   IOAENTR Default OAT with two ports

   If you retrieve the OAT from the OSA card, the OAT obtained may contain several unneeded entries, especially if it is the default OAT.

Example: D-1 Sample TCP/IP OAT file

```
*****************************************************
*UA  Mode     Port   Default      IP Address
*****************************************************
                       LP 5
00  passthru   00       PRI      105.000.005.005
                                 105.000.005.015
02  passthru   01       SEC      105.001.006.006
                                 105.001.006.016
                                 105.001.006.026

                       LP 7
00  passthru   00       no       107.000.075.075
                                 107.000.075.085
02  passthru   01       PRI      107.001.076.076
                                 107.001.086.086
```

Example: D-2 Sample SNA OAT file

```
*********************************************
*UA  Mode  Port     Entry specific information
*********************************************
                LP 5
0A   sna    00

                LP 7
0A   sna    00
```

2. Copy the required parts of the OAT table entry obtained into a new data set.

   Note: Only the *even* unit address entries are required for TCP/IP Passthru entries.

3. Update the OAT records for your logical partitions with the unit address, port mode and port number, default entry indicator, and IP address for each required device.

   Partition Enter the partition number that is used for that entry. If you are running in basic mode or if the CHPID is dedicated, the partition number is 0.

   UA The unit address can be any even address for TCP/IP passthru, but unit address 00,01 is associated with OSA port 0 in the default OAT table.

   Unit address 0A is usually associated with OSA port 0 in SNA mode.

   Note: Port 1 can be accessed and configured only for the OSA-Express 155 ATM feature running in LAN Emulation mode. Even though ATM has only one physical port, it can be configured as two logical ports.

   The OSA port unit address was used by HCD when defining the OSA devices.

   Mode For TCP/IP Passthru, the port mode is passthru. For SNA mode, the port mode is sna.

   Port Enter the port number you want associated with this unit address.

   Note: Port 1 can be accessed and configured only for the OSA-Express 155 ATM feature running in LAN Emulation mode. Even though this feature is just one physical port, it can be configured as two logical ports.

Default Enter `PRI` or `SEC` to make this the primary or secondary entry for this port, or enter `no` if it is neither the primary or secondary entry.

The entry designated as the primary will receive any datagrams that are not specifically addressed to any of the home IP addresses associated with this OSA port. The secondary entry will overtake that function, if the primary entry is not running.

IP Address Enter the home IP address for the port and unit address. Any time an OSA port (in TCP/IP Passthru Mode) is shared, each partition's TCP/IP home IP address must also be added to the OAT. This allows the OSA feature to forward the received datagrams to the appropriate partition.

4. Update the OAT records for your other logical partitions with the unit address, port mode, port number, partition number, default entry indicator, and IP address for all devices in a similar way as described for the first partition, but substitute the appropriate addressing for your partitions.

Example D-3 shows an example of an OAT file running in two partitions: 9 and A. One TCP/IP stack is running in each partition. UNITADD is defined in both partitions with a value of 00.

Shared SNA mode is defined for both partitions using UNITADD 02.

Example: D-3 Sample OAT, shared port

```
*******************************************************
*UA  Mode     Port   Default      IP Address
*******************************************************
                       LP 9
00  passthru   00      no         9.12.2.36
02  SNA        00
***************************************************
                       LP A
00  passthru   00      no         9.12.2.37
02  SNA        00
***************************************************
```

Activating the OSA Configuration

Note: OSA Configuration changes are disruptive, so all applications running via OSA devices *must not* have active sessions. Also, the OSAD device must be online to the host on which OSA/SF is running.

1. Vary all OSA devices offline (except the OSAD device), or at least those devices that have active sessions for this OSA feature, to all partitions sharing this OSA feature.
2. Log on to TSO from the system on which OSA/SF is running.

   Note: The TSO user ID must be set up to use the OSA/SF TSO interface.
3. Issue the IOACMD command, and you will get the choices shown in Figure D-1 on page 205.

```
IOACMD: Enter the command to be issued

 IOACMD: 0 - End IOACMD
 IOACMD: 1 - Clear Debug
 IOACMD: 2 - Configure OSA CHPID
 IOACMD: 3 - Convert OAT
 IOACMD: 4 - Get Configuration File
 IOACMD: 5 - Get Debug
 IOACMD: 6 - Get OSA Address Table
 IOACMD: 7 - Install
 IOACMD: 8 - Put OSA Address Table (OSA-2 only)
 IOACMD: 9 - Query
 IOACMD:10 - Set Parameter
 IOACMD:11 - Shutdown (VM only)
 IOACMD:12 - Start Managing
 IOACMD:13 - Stop Managing
 IOACMD:14 - Synchronize (OSA-2 only)
```

Figure D-1 IOACMD menu

4. To load a configuration, enter `2`.
5. The list shown in Figure D-2 will be displayed.

```
IOACMD: Enter 'quit' to end IOACMD
 IOACMD: Enter 0 for help
 IOACMD: Enter 1 to configure an OSA-2 ATM CHPID
 IOACMD: Enter 2 to configure an OSA-2 FDDI, ENTR, fast Ethernet CHPID
 IOACMD: Enter 3 to configure an OSA-Express gigabit Ethernet CHPID
 IOACMD: Enter 4 to configure an OSA-Express ATM CHPID
 IOACMD: Enter 5 to configure an OSA-Express fast Ethernet CHPID
 IOACMD: Enter 6 to configure an OSA-Express token ring CHPID
 IOACMD: Anything else to get a list of valid OSA CHPIDs
```

Figure D-2 IOACMD configure list

6. You will be prompted through the configuration steps.
   – Enter the OSA-Express feature type.
   – Enter the OSA CHPID number that the CONFIG and OAT is to be downloaded to.
   – Respond to the prompt: **"**`Is CHPID nn of type OSD (QDIO)? (Y./N)`**"**.
   – Enter the fully qualified data set name containing the CONFIG definitions.
   – Enter the fully qualified data set name containing the OAT definitions.
   – Enter the activation option, as shown in Figure D-3 on page 206.

```
IOACMD: 0 - Quit

IOACMD: 1 - Activate
IOACMD:     Sets up all the files and transfers the data to the CHPID

IOACMD: 2 - Activate, no Install
IOACMD:     Only sets up the files, but does not transfer them to the CHPID
IOACMD:     You must issue the Install command at a later time
IOACMD:     to complete the activation
```

Figure D-3 IOACMD Activation options

- Note the following:
  - `Activate` creates the OAT table, updates the index data set, and downloads the OAT table.
  - `Activate, no Install` creates the OAT table, updates the index data set, but does not download the OAT table. IOACMD INSTALL must be done at a later time.
- Enter `Y` to confirm the download, because it is disruptive to the OSA-Express feature.

7. IOACMD CONFIG_OSA will perform the following functions:
   - The OAT information from the specified data set is reformatted and saved on the z/OS host, in the OSA configuration file. The OSA/SF config file (also defined in the startup profile) is updated to point to any code files that are required to support this configuration, and that will be downloaded to the feature during any OSA/SF install action.
   - An OSA/SF install action is done to download the OAT table contained in the OATFILE data set.

Refer to *Open Systems Adapter-Express Customer's Guide and Reference,* SA22-7476, for details on using OSA/SF.

As previously mentioned, following are sample configuration files.

Example: D-4 Sample Fast Ethernet configuration file

```
enet.0.1 = config file name      /* Configuration name (32-char max)
fenet.0.2 = user data             /* User data (32-char max)
fenet.0.3 = portname              /* Port name (8-char max)
                                  /*  Data ignored for OSD CHPIDs
fenet.0.4 = 000000000000          /* Local MAC address (12 hex digits)
fenet.0.5 = auto                  /* Speed/mode
                                  /*  Auto - auto negotiate
                                  /*  10H  - 10  Mb, half duplex
                                  /*  10F  - 10  Mb, full duplex
                                  /*  100H - 100 Mb, half duplex
                                  /*  100F - 100 Mb, full duplex
fenet.0.6.1 = 000000000000        /* Group address 1 (12 hex digits)
fenet.0.6.5 = 000000000000        /* Group address 5
sna.0.1  = Configuration name     /* Configuration name (32-char max)
sna.0.2  = 90.00                  /* Inactivity timer (ti)
                                  /*   .24-90 in increments of .12
                                  /*   0 disables the inactivity timer
sna.0.3  = 2.00                   /* Response timer (t1)
                                  /*   .20-51 in increments of .20
sna.0.4  = 0.08                   /* Acknowledgement timer (t2)
                                  /*   .08-20.4 in increments of .08
sna.0.5  = 1                      /* N3 (1-4)
sna.0.6  = 8                      /* TW (1-16)
```

Example: D-5 Sample Token Ring configuration file

```
tr.0.1 = config file name      /* Configuration name (32-char max)
tr.0.2 = user data             /* User data (32-char max)
tr.0.3 = portname              /* Port name (8-char max)
                               /*  Data ignored for OSD CHPIDs
tr.0.4 = 000000000000          /* Local MAC address (12 hex digits)
tr.0.5 = 00000000              /* Functional address (8 hex digits)
tr.0.6 = Auto                  /* Speed/mode
                               /*  Auto - Auto sense from the ring
                               /*    4H -   4 Mbs Half duplex
                               /*    4F -   4 Mbs Full duplex
                               /*   16H -  16 Mbs Half duplex
                               /*   16F -  16 Mbs Full duplex
                               /*   100 - 100 Mbs Full duplex
tr.0.7.1 = 000000000000        /* Group address 1 (12 hex digits)
tr.0.7.5 = 000000000000        /* Group address 5
sna.0.1  = Configuration name  /* Configuration name (32-char max)
sna.0.2  = 90.00               /* Inactivity timer (ti)
                               /*   .24-90 in increments of .12
                               /*   0 disables the inactivity timer
sna.0.3  = 2.00                /* Response timer (t1)
                               /*   .20-51 in increments of .20
sna.0.4  = 0.08                /* Acknowledgement timer (t2)
                               /*   .08-20.4 in increments of .08
sna.0.5  = 1                   /* N3 (1-4)
sna.0.6  = 8                   /* TW (1-16)
sna.0.7  = 6                   /* Enhanced availability type
sna.0.8  = 0.00                /* Load balance factor (0-1)
sna.0.9  = 0.00                /* Session delay (0-15)
```

Example: D-6 Sample ATM configuration file

```
phy.0.1  = config file name       /* Configuration name (32-char max)
phy.0.2  = port description       /* Port description (16-char max)
phy.0.3  = Port name              /* Port name (8-char max)
phy.0.4  = 000000000000           /* Local End System ID (12 hex digits)
phy.0.5  = auto                   /* Port UNI version (AUTO, 30 or 31)
phy.0.6  = 0                      /* Control plane use
                                  /*  0 - ILMI & SVC enabled
                                  /*  3 - ILMI & SVC disabled
phy.0.7  = 0                      /* Transmit clock source
                                  /*  0 - OSA generated
                                  /*  1 - Network generated
phy.0.8  = 0                      /* Physical layer type
                                  /*  0 - Sonet
                                  /*  1 - SDH
phy.0.9  = 0.0.0.0                /* TCP/IP instance IP address
phy.0.10 = 1                      /* Bandwidth allocation
                                  /*  1 - Best effort only
                                  /*  2 - Reserve bandwidth
                                  /*      & best effort
                                  /*  3 - Reserved bandwidth
nat.0.1  = configuration name     /* Configuration name (32-char max)
nat.0.2  = Yes                    /* Enable LAN traffic (Yes, No)
pvc.0.1.1 = PVC name              /* PVC name (8-char max)
pvc.0.1.2 = 353207                /* Forward peak cell rate (0-353207)
pvc.0.1.3 = 353207                /* Backward peak cell rate(0-353207)
pvc.0.1.4 = 0                     /* VPI for this PVC entry (0-255)
pvc.0.1.5 = 35                    /* VCI for this PVC entry (32-65535)
pvc.0.1.6 = 8448                  /* Forward Max PDU size (64-9188)
pvc.0.1.7 = 8448                  /* Backward Max PDU size(64-9188)
pvc.0.1.8 = 0                     /* Reserved bandwidth
                                  /*  0 - Use defaults
                                  /*  1 - Specify parameters 9-12
pvc.0.1.9 = 353207                /* Forward sustain cell rate (0-353207)
pvc.0.1.10= 353207                /* Backward sustain cell rate(0-353207)
pvc.0.1.11= 353207                /* Forward cell burst rate (0-353207)
pvc.0.1.12= 353207                /* Backward cell burst rate(0-353207)
emul.0.1  = configuration name    /* Configuration name (32-char max)
emul.0.2  = Yes                   /* Enable LAN traffic (Yes, No)
emul.0.3  = 1                     /* Emulated port type
                                  /*  1 - Ethernet
                                  /*  2 - Token ring
emul.0.4  = user data             /* User data (32-char max)
emul.0.5  = ELAN name             /* ELAN name (32-char max)
emul.0.6  = 000000000000          /* Local MAC address (12 hex digits)
emul.0.7  = 155.0                 /* Best effort peak rate (1-155)
                                  /*  in 0.1 increments
emul.0.8  = 0                     /* IBM Enhanced mode
                                  /*  0 - drop direct connect
                                  /*  Not 0 - keep connections
.
.                                  /* Several other parameters ...
.
sna.0.8  = 0.00                   /* Load balance factor (0-1)
sna.0.9  = 0.00                   /* Session delay (0-15)
```

E

Sample definitions

This appendix lists sample of definitions for TCP/IP in various operating systems environments. It also includes VTAM definitions used in z/OS.

© Copyright IBM Corp. 2003. All rights reserved.

Sample environment

Figure E-1 is a logical representation of the environment used for the definition examples.

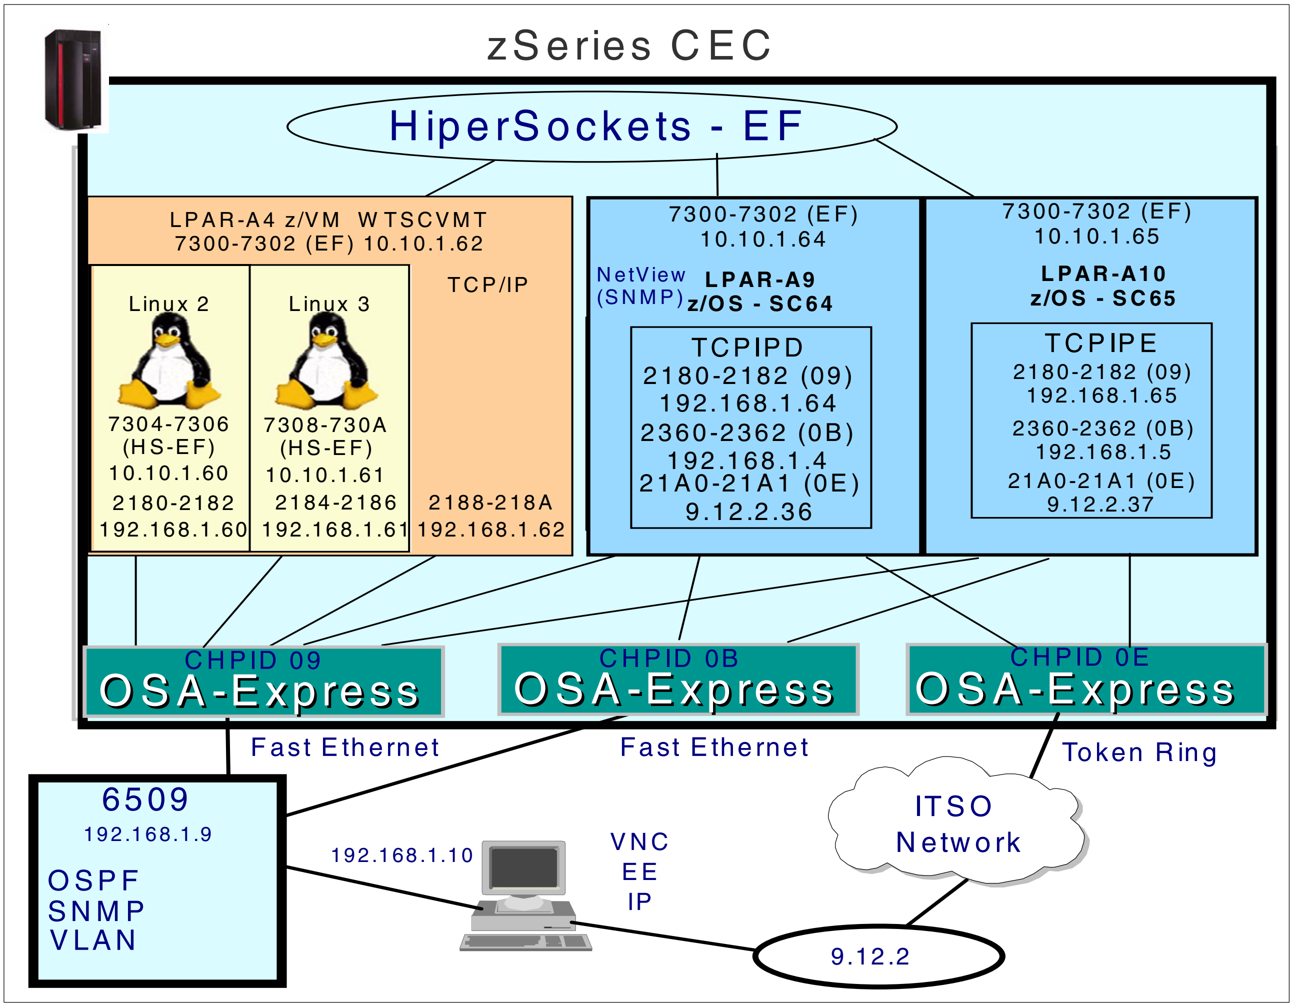

Figure E-1 ITSO test environment

z/OS definitions

This section includes examples of the definitions we used in our z/OS LPARs.

TCP/IP profiles

This TCP/IP profile for SC64 only shows the lines we added or changed.

Example: E-1 SC64 TCP/IP profile

```
IPCONFIG MULTIPATH IQDIOR DATAGRAMFWD

DEVICE OSA2180 MPCIPA PRIRouter          ; OSD Fast Ethernet Devices on CHPID 09
 LINK   SC642180LINK  IPAQENET  OSA2180

INTERFACE L2180V6 DEFINE IPAQENET6 PORTNAME OSA2180 ; IPv6 definition for CHPID09
IPADDR FEC0:0:0:1::3302
       FEC0:0:0:1001::3302

 DEVICE VIPA1 VIRTUAL 0                  ; VIPA for Enterprise Extender
 LINK VLINK1  VIRTUAL 0 VIPA1

 DEVICE IUTSAMEH MPCPTP                  ; For Enterprise Extender
 LINK EELINK MPCPTP IUTSAMEH

 DEVICE OSA2360 MPCIPA SECRouter         ; OSD Fast Ethernet Devices on CHPID 0B
 LINK   SC642360LINK  IPAQENET OSA2360

DEVICE IUTIQDEF MPCIPA                   ; HiperSockets CHPID EF
 LINK   HIPERLEF IPAQIDIO IUTIQDEF

DEVICE OSA21A0 LCS   21A0                ; OSE Token Ring Devices on CHPID 0E
LINK   SC6421A0LINK IBMTR 0 OSA21A0

HOME
   192.168.1.64  SC642180LINK
   192.168.1.4   SC642360LINK
   192.168.1.164 VLINK1
   10.10.1.64    HIPERLEF
   9.12.2.36     SC6421A0LINK

BEGINROUTES
ROUTE FEC0::0/10      =     L2180v6          MTU 1500
ROUTE 192.168.1.0/24  =     SC642180LINK     MTU 1492
ROUTE 192.168.1.0/24  =     SC642360LINK     MTU 1492
ROUTE 9.12.2.0/24     =     SC6421A0LINK     MTU 2000
ROUTE 10.10.1.0/24    =        HIPERLEF      MTU 16384
ROUTE DEFAULT    192.168.1.64 SC642180LINK MTU 1492
ENDROUTES

START IUTIQDEF
START IUTSAMEH
START OSA2180
START L2180V6
START OSA21A0
START OSA2360
```

The following TCP/IP profile for SC65 shows only the lines we added or changed.

Example: E-2 SC65 TCP/IP profile

```
IPCONFIG MULTIPATH IQDIOR DATAGRAMFWD

DEVICE OSA2360  MPCIPA SECRouter          ; OSD Fast Ethernet Devices on CHPID 0B
 LINK   SC652360LINK  IPAQENET      OSA2360

 DEVICE OSA2180  MPCIPA                   ; OSD Fast Ethernet Devices on CHPID 09
 LINK   SC652180LINK  IPAQENET      OSA2180

 DEVICE VIPA1 VIRTUAL 0                   ; VIPA for Enterprise Extender
 LINK VLINK1  VIRTUAL 0 VIPA1

 DEVICE IUTSAMEH MPCPTP                   ; For Enterprise Extender
 LINK EELINK MPCPTP IUTSAMEH

 INTERFACE L2180V6 DEFINE IPAQENET6 PORTNAME OSA2180  ; IPv6 definition for CHPID 09
 IPADDR FEC0:0:0:1::3365

  DEVICE IUTIQDEF MPCIPA                  ; Hipersockets CHPID EF
  LINK   HIPERLEF IPAQIDIO IUTIQDEF

DEVICE OSA21A0 LCS  21A0                  ; OSE Token Ring Devices on CHPID 0E
LINK   SC6521A0LINK IBMTR 0 OSA21A0

HOME
  192.168.1.5    SC652360LINK
  192.168.1.65   SC652180LINK
  192.168.1.165    VLINK1
  10.10.1.65     HIPERLEF
  9.12.2.37     SC6521A0LINK

BEGINROUTES
ROUTE FEC0::0/10       =     L2180v6          MTU 1492
ROUTE 192.168.1.0/24   =     SC642180LINK     MTU 1492
ROUTE 192.168.1.0/24   =     SC642360LINK     MTU 1492
ROUTE  9.12.2.0/24     =     SC6521A0LINK     MTU 2000
ROUTE 10.10.1.0/24     =       HIPERLEF       MTU 16384
ROUTE DEFAULT    192.168.1.5 SC652360LINK     MTU 1492
ENDROUTES

START OSA2360
START OSA2180
START L2180V6
START OSA21A0
START IUTIQDEF
```

VTAM definitions

These are examples of VTAM definitions used for OSD and Enterprise Extender.

VTAM TRL definitions

We used the same TRLs in both SC64 and SC65.

Note: TRLs are built dynamically for Hipersockets.

Example: E-3 TRL for CHPID 09

```
OSAFE09  VBUILD TYPE=TRL
TRLE2180 TRLE  LNCTL=MPC,                                          X
               READ=2180,                                          X
               WRITE=2181,                                         X
               DATAPATH=2182,                                      X
               PORTNAME=OSA2180,                                   X
               MPCLEVEL=QDIO
```

Example: E-4 TRL for CHPID 0B

```
OSAFE0B  VBUILD TYPE=TRL
TRLE2360 TRLE  LNCTL=MPC,                                          X
               READ=2360,                                          X
               WRITE=2361,                                         x
               DATAPATH=2362,                                      X
               PORTNAME=OSA2360,                                   X
               MPCLEVEL=QDIO
```

VTAM definitions for Enterprise Extender

Switched Dynamic LU definitions for Enterprise Extender.

Example: E-5 SWNDLUR1

```
SWNDLUR1 VBUILD TYPE=SWNET,            REQUIRED FOR TR              X
               MAXGRP=4,                                           X
               MAXNO=512
*
PUDLUR1  PU    ADDR=01,                ADDR NOT USED                X
               IDBLK=05D,                                          X
               IDNUM=11111,                                        X
               ISTATUS=ACTIVE,                                     X
               DISCNT=NO,                                          X
               MAXDATA=265,                                        X
               MAXOUT=7,                                           X
               PACING=0,                                           X
               PUTYPE=2,MODETAB=NEWMTAB,                           X
               VPACING=0
LUEE0001 LU    LOCADDR=1,DLOGMOD=D4A32782,LOGAPPL=SC64TS
LUEE0002 LU    LOCADDR=2,DLOGMOD=D4A32782,LOGAPPL=SC64TS
```

Example: E-6 SWTO64

```
SWTO64   VBUILD TYPE=SWNET,                                        X
               MAXGRP=4,                                           X
               MAXNO=512
PUTO64   PU    MAXDATA=256,ADDR=01,                                X
               CPNAME=SC64M,                                       X
               CPCP=YES,HPR=YES,                                   X
               DWACT=YES,PUTYPE=2
PATH64   PATH  SAPADDR=08,IPADDR=192.168.1.164,                    X
               GRPNM=EEGROUP
```

Example: E-7 SWTO65

```
SWTO65   VBUILD TYPE=SWNET,                                       X
               MAXGRP=4,                                          X
               MAXNO=512
PUTO65   PU    MAXDATA=256,ADDR=01,                               X
               CPNAME=SC65M,                                      X
               CPCP=YES,HPR=YES,                                  X
               DWACT=YES,PUTYPE=2
PATH65   PATH  SAPADDR=8,IPADDR=192.168.1.165,                    X
               GRPNM=EEGROUP
```

Example: E-8 XCAEE for SC64

```
XCAEE    VBUILD TYPE=XCA
EEPORT   PORT  MEDIUM=HPRIP,                                      X
               LIVTIME=10,                                        X
               IPTOS=(20,40,80,C0),                               X
               SAPADDR=4,                                         X
               SRQRETRY=3,                                        X
               SRQTIME=15
EEGROUP  GROUP DIAL=YES,ANSWER=ON,                                X
               DYNPU=YES,                                         X
               DYNPUPFX=BR,                                       X
               AUTOGEN=(10,L,P),                                  X
               CALL=INOUT
```

Example: E-9 XCAEE65 for SC65

```
XCAEE65  VBUILD TYPE=XCA
EEPORT   PORT  MEDIUM=HPRIP,                                      X
               LIVTIME=10,                                        X
               IPTOS=(20,40,80,C0),                               X
               SAPADDR=4,                                         X
               SRQRETRY=3,                                        X
               SRQTIME=15
EEGROUP  GROUP DIAL=YES,ANSWER=ON,                                X
               DYNPU=YES,                                         X
               DYNPUPFX=BR,                                       X
               AUTOGEN=(10,L,P),                                  X
               CALL=INOUT
```

Example E-10 and Example E-11 only contain the lines changed or added to support Enterprise Extender.

Example: E-10 ATCSTR64 for SC64

```
NODETYPE=NN,                NETWORK NODE NN FOR  EE               X
IPADDR=192.168.1.164,       VIPA ADDRESS FOR EE ENTERPRISE EXT    X
TCPNAME=TCPIPD,             TCPIP JOB NAME FOR EE ENT.EXT         X
```

Example: E-11 ATCSTR65 for SC65

```
NODETYPE=NN,                NETWORK NODE NN FOR  EE               X
IPADDR=192.168.1.165,       VIPA ADDRESS FOR EE ENTERPRISE EXT    X
TCPNAME=TCPIPE,             TCPIP JOB NAME FOR EE ENT.EXT         X
```

Linux definitions

This section includes examples of the definitions we used in our Linux guests

Example E-12 has changed or added lines in modules.conf (the same in both Linux guests).

Example: E-12 modules.conf

```
alias hsi0 qeth
```

Example E-13 and Example E-14 only show the lines we added or changed in the chandev.conf files.

Example: E-13 LINUX2 chandev.conf

```
noauto;qeth0,0x7106,0x7107,0x7108
qeth1,0x2180,0x2181,0x2182
add_parms,0x10,0x2180,0x2182,portname:OSA2180
qeth2,0x7304,0x7305,0x7306
```

Example: E-14 LINUX3 chandev.conf

```
noauto;qeth0,0x7109,0x710a,0x710b
qeth1,0x2184,0x2185,0x2186
add_parms,0x10,0x2184,0x2186,portname:OSA2180
qeth2,0x7308,0x7309,0x730a
```

The next two examples only show the lines we added or changed in the rc.config file.

Example: E-15 LINUX2 rc.config

```
# Number of network cards: "_0" for one, "_0 _1 _2 _3" for four cards
#
NETCONFIG="_0 _1 _2"
# IP Adresses
#
IPADDR_0="9.12.9.3"
IPADDR_1="192.168.1.60"
IPADDR_2="10.10.1.60"
# Network device names (e.g. "eth0")
#
NETDEV_0="hsi0"
NETDEV_1="eth1"
NETDEV_2="hsi2"
#
IFCONFIG_0="9.12.9.3 broadcast 9.12.9.31 netmask 255.255.255.224 mtu 1492 up"
IFCONFIG_1="192.168.1.60 broadcast 192.168.1.255 netmask 255.255.255.0 mtu 1492 up"
IFCONFIG_2="10.10.1.60 broadcast 10.10.1.255 netmask 255.255.255.0 mtu 1492 up"
```

Example: E-16 LINUX3 rc.config

```
# Number of network cards: "_0" for one, "_0 _1 _2 _3" for four cards
#
NETCONFIG="_0 _1 _2"
# IP Adresses
#
IPADDR_0="9.12.9.4"
IPADDR_1="192.168.1.61"
IPADDR_2="10.10.1.61"
# Network device names (e.g. "eth0")
```

```
#
NETDEV_0="hsi0"
NETDEV_1="eth1"
NETDEV_2="hsi2"
#
IFCONFIG_0="9.12.9.4 broadcast 9.12.9.31 netmask 255.255.255.224 mtu 1492 up"
IFCONFIG_1="192.168.1.61 broadcast 192.168.1.255 netmask 255.255.255.0 mtu 1492 up"
IFCONFIG_2="10.10.1.61 broadcast 10.10.1.255 netmask 255.255.255.0 mtu 1492 up"
```

Example E-17 - route.conf is the same in both Linux guests.

Example: E-17   route.conf

```
9.12.9.0            0.0.0.0             255.255.255.224     hsi0
192.168.1.0         0.0.0.0             255.255.255.0       eth0
default             9.12.9.1            0.0.0.0             hsi0
```

z/VM TCP/IP profile

The TCP/IP profile in Example E-18 only shows the statements related to OSA-Express and HiperSockets.

Example: E-18   Sample z/VM TCP/IPprofile

```
DEVICE OSA2188  OSD 2188 PORTNAME OSA2180          ; OSD Fast Ethernet devices on CHPID 09
LINK   OSA2188L    QDIOETHERNET  OSA2188
;
DEVICE HIPERDEF HIPERS 7300 PORTNAME HIPERPEF      ; HiperSockets CHPID EF
LINK   HIPERLEF QDIOIP HIPERDEF
;
HOME

  192.168.1.62   OSA2188L
  10.10.1.62     HIPERLEF
  ;
GATEWAY
; (IP) Network  First        Link      Max. Packet  Subnet      Subnet
; Address       Hop          Name      Size (MTU)   Mask        Value
; -----------   ------------ -------   -----------  ----------- --------
;
; THESE ARE THE CURRENT GATEWAYS THAT ARE BEING USED TODAY IN PRODUCTION
192.168.1         =        OSA2188L     1500         0
10                =        HIPERLEF     8192     0.255.255.0  0.10.1.0
;
START HIPERDEF
START OSA2188
```

F

HiperSockets Accelerator

This appendix provides a description of HiperSockets Accelerator, definition requirements, and verification procedures.

© Copyright IBM Corp. 2003. All rights reserved.

HiperSockets Accelerator description

HiperSockets Accelerator allows a z/OS TCP/IP router stack to efficiently route IP packets from an OSA-Express (QDIO) interface to a HiperSockets (iQDIO) interface and vice versa. The routing is done by the z/OS Communications Server device drivers at the lowest possible software data link control level. IP packets do not have to be processed at the higher level TCP/IP stack routing function, hence reducing the path-length and improving performance.

HiperSockets Accelerator is activated by configuring the IQDIORouting option in the TCP/IP profile using the IPCONFIG statement.

The TCP/IP stack automatically detects IP packet prerouting across a HiperSockets Accelerator eligible route. Eligible routes are from OSA-Express (QDIO) to HiperSockets (iQDIO), and from HiperSockets (iQDIO) to OSA-Express (QDIO).

The TCP/IP routing stack creates IQDIORouting route entries for packets for which it is not the destination stack. These entries are added to the IQDIORouting table. The destination stack must be reachable through HiperSockets.

All subsequent packets for the same destination will take the optimized device driver path, and will not traverse the routing function of the TCP/IP routing stack. No change is required for target stacks. There is a timer built in to the HiperSockets Accelerator function.

Figure F-1 shows our network configuration and the HiperSockets Accelerator flow.

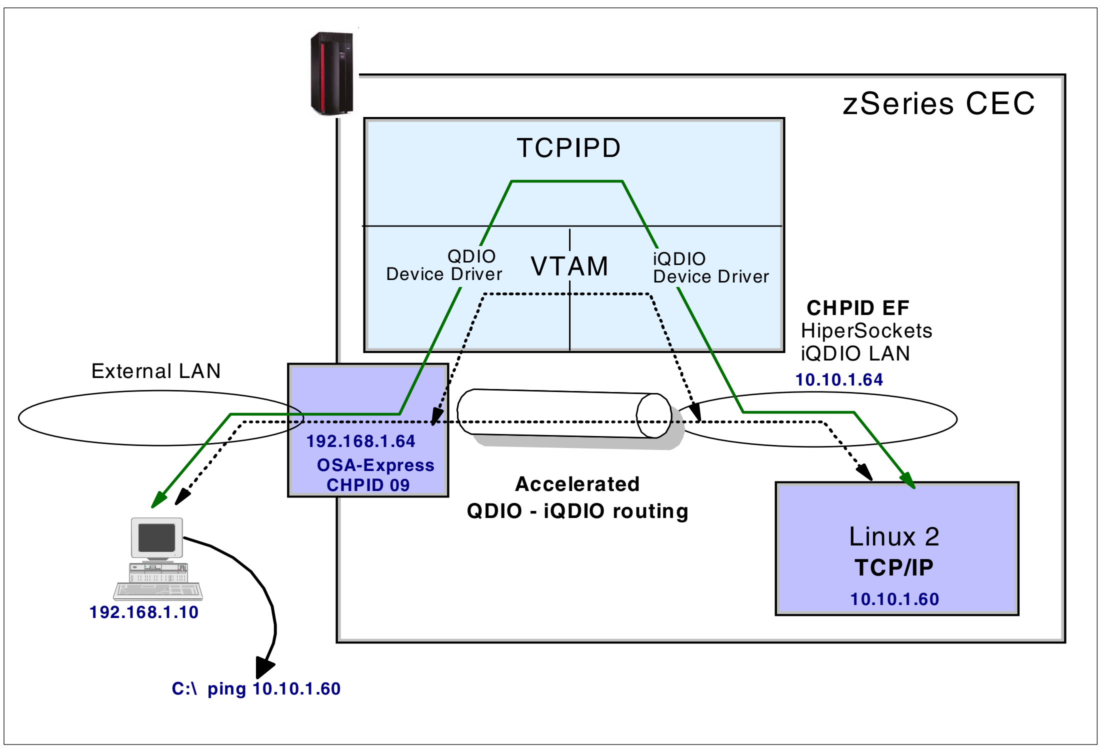

Figure F-1 HiperSockets Accelerator flow

HiperSockets definitions

Implementation of HiperSockets Accelerator is quite simple. Take a look at our sample TCP/IP profile in Example F-1.

Example: F-1 TCP/IP Profile for HiperSockets Accelerator

```
IPCONFIG IQDIOR DATAGRAMFWD

DEVICE OSA2180 MPCIPA PRIRouter          ; OSD Fast Ethernet Devices on CHPID 09
 LINK   SC642180LINK  IPAQENET  OSA2180

DEVICE IUTIQDEF MPCIPA                   ; HiperSockets CHPID EF
 LINK   HIPERLEF IPAQIDIO IUTIQDEF

HOME
   192.168.1.64  SC642180LINK
   10.10.1.64    HIPERLEF

BEGINROUTES
ROUTE 192.168.1.0/24   =     SC642180LINK     MTU 1492
ROUTE 10.10.1.0/24     =        HIPERLEF      MTU 16384
ROUTE DEFAULT    192.168.1.64 SC642180LINK MTU 1492
ENDROUTES

START IUTIQDEF
START OSA2180
```

The profile in Example F-1 will allow us to use the HiperSockets Accelerator to connect between the OSA-Express card on CHPID 09 and HiperSockets on CHPID EF.

Note: HiperSockets must be defined in the IOCDS as CHPID type IQD.

Note the following:

- IQDIOR activates HiperSockets Accelerator.
- DATAGRAMFWD is required to forward the datagrams between the two networks.
- PRIRouter is required to route IP addresses between the networks.

Verifying HiperSockets Accelerator

We verified that HiperSockets Accelerator was working by issuing z/OS console commands and reviewing the resulting messages.

When TCP/IP was started, we saw the message shown in Figure F-2 confirming that HiperSockets Accelerator was enabled.

```
EZZ0688I IQDIO ROUTING IS ENABLED
```

Figure F-2 Message when starting TCP/IP

Next, we looked at the active TCP/IP configuration as shown in Figure F-3.

```
D TCPIP,TCPIPD,N,CONFIG
EZD0101I NETSTAT CS V1R4 TCPIPD 841
TCP CONFIGURATION TABLE:
DEFAULTRCVBUFSIZE:  00016384  DEFAULTSNDBUFSIZE: 00016384
.
.
.
IQDIOROUTE: YES  QDIOPRIORITY: 1
```

Figure F-3 Results of D TCPIP,,N,CONFIG command

Since IQDIOROUTE is YES, we know that HiperSockets Accelerator is enabled.

We can verify this in the dynamically built VTAM TRLE as well, as shown in Figure F-4.

```
D NET,TRL,TRLE=IUTIQDEF
IST097I DISPLAY ACCEPTED
IST075I NAME = IUTIQDEF, TYPE = TRLE 881
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST087I TYPE = LEASED            , CONTROL = MPC , HPDT = YES
IST1715I MPCLEVEL = QDIO       MPCUSAGE = SHARE
IST1716I PORTNAME =            LINKNUM =   0   OSA CODE LEVEL = D3GF
IST1577I HEADER SIZE = 4096 DATA SIZE = 65536 STORAGE = ***NA***
IST1221I WRITE DEV = 7301 STATUS = ACTIVE     STATE = ONLINE
IST1577I HEADER SIZE = 4092 DATA SIZE = 0 STORAGE = ***NA***
IST1221I READ  DEV = 7300 STATUS = ACTIVE     STATE = ONLINE
IST1221I DATA  DEV = 7302 STATUS = ACTIVE     STATE = N/A
IST1724I I/O TRACE = OFF  TRACE LENGTH = *NA*
IST1717I ULPID = TCPIPD
IST1814I IQDIO ROUTING ENABLED
```

Figure F-4 HiperSockets TRLE

You can see from message IST1814I that VTAM recognizes that HiperSockets Accelerator is enabled.

To verify it works, we must communicate between the IP network on the OSA-Express card and the IP network on the HiperSockets CHPID. Figure F-5 on page 221 shows the HiperSockets routing table before any communication has taken place.

```
D TCPIP,TCPIPD,N,ROUTE,IQDIO
EZD0101I NETSTAT CS V1R4 TCPIPD 874
IPV4 DESTINATIONS
DESTINATION         GATEWAY          INTERFACE
0 OF 0 RECORDS DISPLAYED
```

Figure F-5 HiperSockets Route (before communication)

Since our workstation has multiple network adapters connected to different networks, we had to build an indirect route from the 192.168.1 network to the 10 network. We accomplished this with the command shown in Figure F-6.

```
ROUTE ADD 10.10.1.0 MASK 255.255.255.0 192.168.1.64 -p
```

Figure F-6 Add indirect route to workstation

The -p in Figure F-6 makes this a persistent route across boots.

From the workstation we did a ping through the 192.168.1 network to our two Linux guests on the HiperSockets 10 network to dynamically build the HiperSockets routing table. Figure F-7 shows the results.

```
D TCPIP,TCPIPD,N,ROUTE,IQDIO
EZD0101I NETSTAT CS V1R4 TCPIPD 902
IPV4 DESTINATIONS
DESTINATION         GATEWAY          INTERFACE
10.10.1.60/0        10.10.1.60       HIPERLEF
10.10.1.61/0        10.10.1.61       HIPERLEF
2 OF 2 RECORDS DISPLAYED
```

Figure F-7 HiperSockets Route (after communication)

The HiperSockets routing table is very short-lived. After about 90 seconds of no use, the display would once again look like Figure F-5.

For more implementation information on HiperSockets and HiperSockets Accelerator, refer to *zSeries HiperSockets*, SG24-6816.

G

ARP Takeover

This appendix provides a description of ARP Takeover, definition requirements, and verification procedures.

© Copyright IBM Corp. 2003. All rights reserved.

ARP Takeover description

ARP Takeover is a mechanism that allows traffic to be redirected from a failing OSA-Express connection to another OSA-Express connection.

When TCP/IP is started, all of the IP addresses for devices defined as connecting to an OSA-Express feature in QDIO mode (device type MPCIPA in the TCP/IP profile) are dynamically downloaded to the OSA-Express card. If the OSA-Express card is running in non-QDIO mode (device type LCS in the TCP/IP profile), then OSA/SF must be used to build and activate a configuration that identifies the multiple IP addresses that are used with this feature.

The OSA-Express card maintains the ARP table for all of the IP addresses to which it is connected. Maintaining the ARP table within the OSA-Express card improves performance, since fewer I/O operations are required to direct IP traffic to the desired destination.

To take advantage of ARP Takeover, the following conditions must be met:

- If this is an OSD-type CHPID, IP addresses will be dynamically downloaded to the feature.
- Or—if this is an OSE-type CHPID, a configuration must be activated that includes all of the IP addresses that are used by this feature.
- TCP/IP must have connections to at least two similar OSA-Express cards.
- The TCP/IP profile must be defined properly.

Figure G-1 depicts our test environment for ARP Takeover. We defined CHPID 09 and 0B as OSD-type CHPIDs.

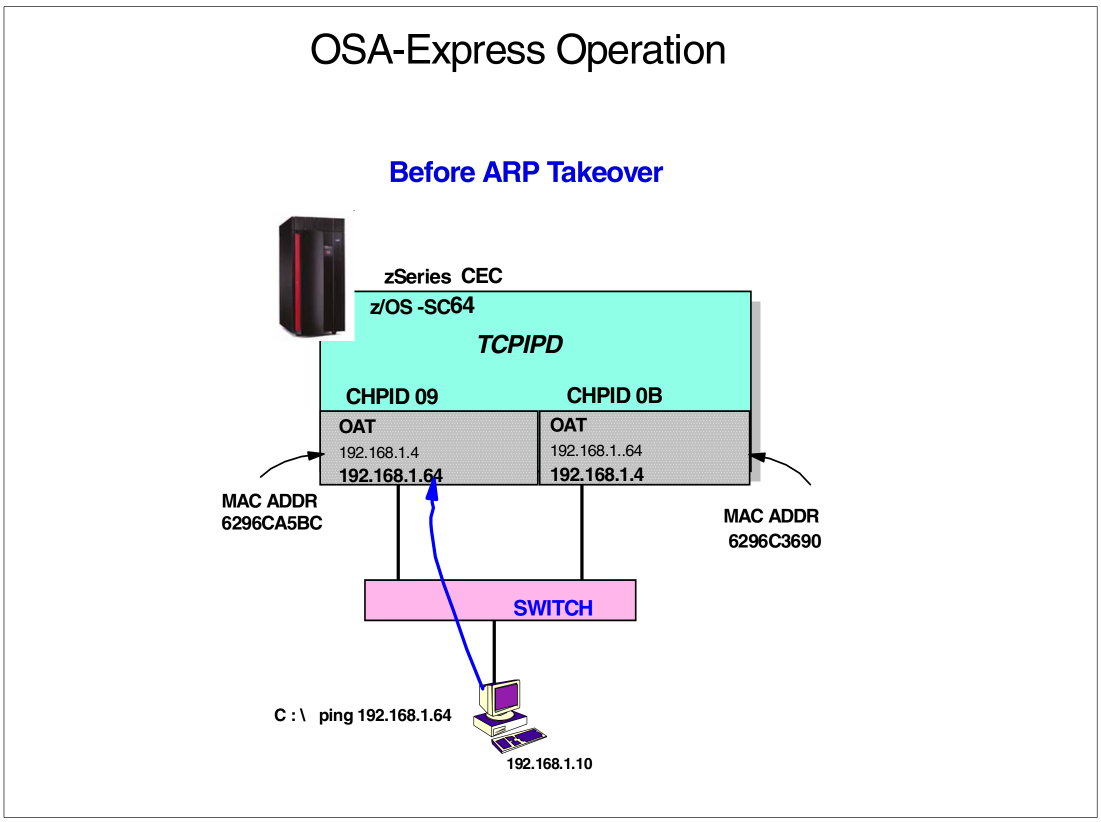

Figure G-1 Test environment before ARP Takeover

ARP Takeover definitions

We had to update the TCP/IP profile to support ARP Takeover. We also had to modify the Cisco switch (6509) configuration to support ARP Takeover.

TCP/IP definitions

Example G-1 shows the items we changed or added in our TCP/IP profile for ARP Takeover support.

Example: G-1 SC64 TCP/IP profile

```
IPCONFIG MULTIPATH IQDIOR DATAGRAMFWD

DEVICE OSA2180 MPCIPA PRIRouter          ; OSD Fast Ethernet Devices on CHPID 09
 LINK   SC642180LINK  IPAQENET  OSA2180

 DEVICE OSA2360 MPCIPA                   ; OSD Fast Ethernet Devices on CHPID 0B
 LINK   SC642360LINK  IPAQENET OSA2360
HOME
   192.168.1.64  SC642180LINK
   192.168.1.4   SC642360LINK

BEGINROUTES
ROUTE 192.168.1.0/24   =     SC642180LINK     MTU 1492
ROUTE 192.168.1.0/24   =     SC642360LINK     MTU 1492
ROUTE DEFAULT    192.168.1.64 SC642180LINK MTU 1492
ENDROUTES
```

On the IPCONFIG statement, we added MULTIPATH to allow multiple path definitions to the same network (or subnetwork).

A DEVICE and LINK statement was needed for each of the two OSA-Express features that we will switch between for the test.

Note the ROUTE statements that define two paths to get to the same network.

Cisco switch definitions

We used the Cisco command **set port host to** do this, because it accomplishes all the tasks required to define a port as an end-station port.

The **set port host** command sets the following:

- **channel mode** off
- **trunk mode** off
- **port fast start** to enabled

Example G-2 shows an example of executing the command.

Example: G-2 Set port host

```
6500-top> (enable) set port host 1/1
2000 Jan 15 04:01:32 %SYS-6-CFG_CHG:Module 1 block changed by Console//
Port(s) 1/1 channel mode set to off.

Warning: Spantree port fast start should only be enabled on ports connected to
```

```
a single host. Connecting hubs, concentrators, switches, bridges, etc. to a
fast start port can cause temporary spanning tree loops. Use with caution.

Spantree port 1/1 fast start enabled.
Port(s) 1/1 trunk mode set to off.
```

Restriction: If you intend to use an OSA-Express CHPID as a trunk for a VLAN, you can not connect it to a host port. This means that ARP Takeover will not work through a trunk port of a VLAN.

Verifying ARP Takeover

Note: The results documented here are for our test with a specific switch. You should consult the documentation provided by the manufacturer of your switch to determine the requirements to support ARP Takeover.

We verified that ARP Takeover worked by performing two different tasks.

1. Pulling the CAT5 cable from the OSA-Express feature
2. Stopping the device in the TCP/IP stack

Pulling the CAT5 cable

Figure G-2 on page 227 shows our environment after pulling the CAT5 cable, and the route that the **ping** will take as a result of this.

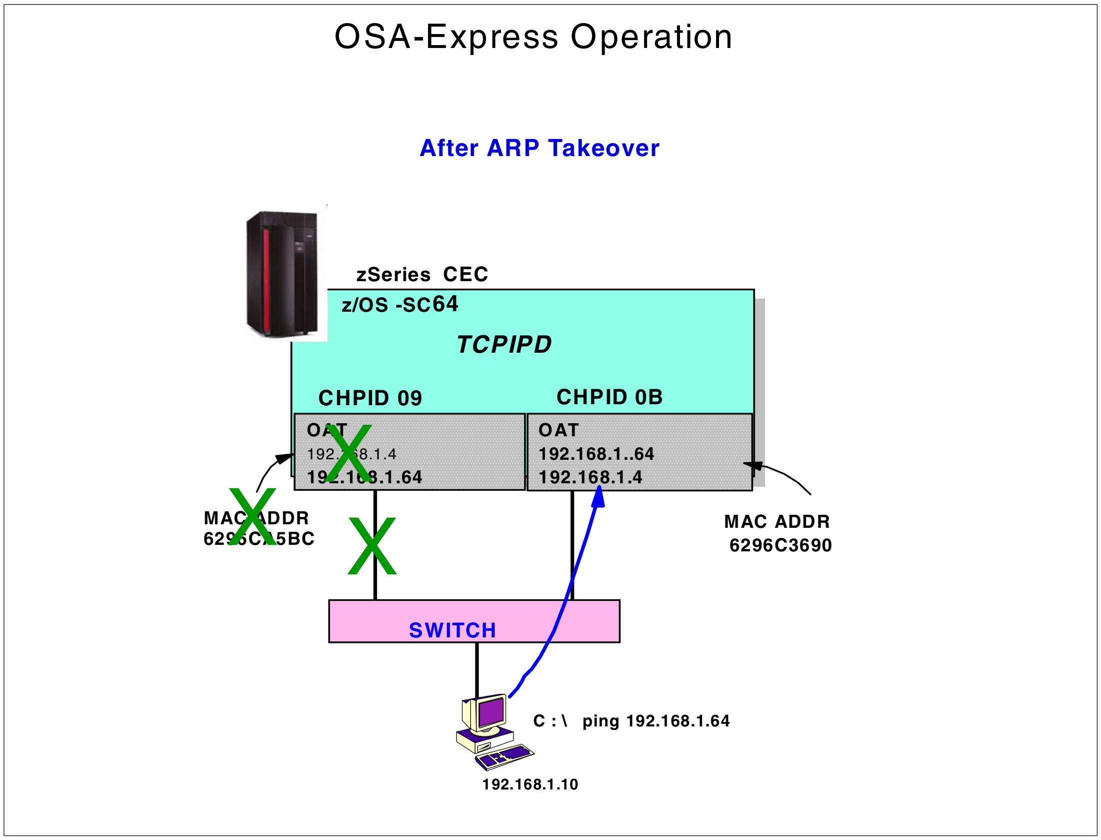

Figure G-2 Test environment after pulling the CAT5 cable

Example G-3 shows the contents of ARP cache on a Windows workstation after pinging each of the IP addresses defined in TCPIPD before any error condition was introduced.

Example: G-3 Windows workstation APR cache (normal)

```
C:\>arp -a

Interface: 192.168.1.10 on Interface 0x3000004
  Internet Address      Physical Address      Type
  192.168.1.4           00-06-29-6c-36-90     dynamic
  192.168.1.64          00-06-29-6c-a5-bc     dynamic
```

Notice that 192.168.1.4 is currently assigned to the MAC address associated with CHPID 0B and 192.168.1.64 is currently assigned to the MAC address associated with CHPID 09.

Example G-4 on page 228 shows the contents of the ARP table from the OSA-Express cards before any error condition was introduced.

Example: G-4 ARP Table from SC64 TCPIPD (normal)

```
MVS TCP/IP NETSTAT CS V1R4       TCPIP Name: TCPIPD         15:34:44
Querying ARP cache for address 192.168.1.4
Link: SC642360LINK      ETHERNET: 0006296C3690

Querying ARP cache for address 192.168.1.64
Link: SC642180LINK      ETHERNET: 0006296CA5BC

Querying ARP cache for address 192.168.1.10
Link: SC642180LINK      ETHERNET: 0004AC1D18E9
```

We pulled the CAT5 cable connecting the Fast Ethernet port of CHPID 09 from the switch.

Example G-5 shows the resulting messages from the z/OS console.

Example: G-5 z/OS console messages after pulling CHPID 09 CAT5 cable

```
EZZ4329I LINK SC642360LINK HAS TAKEN OVER ARP RESPONSIBILITY FOR
INACTIVE LINK SC642180LINK
EZZ4311I LINK SC642180LINK HAS FAILED ON DEVICE OSA2180
```

Example G-6 shows the contents of the ARP cache of the workstation after pulling the cable for CHPID 09.

Example: G-6 Workstation ARP cache after pulling CHPID 09 CAT5 cable

```
C:\>arp -a

Interface: 192.168.1.10 on Interface 0x3000004
  Internet Address      Physical Address      Type
  192.168.1.4           00-06-29-6c-36-90     dynamic
  192.168.1.64          00-06-29-6c-36-90     dynamic
```

Notice that both IP addresses now point to the same MAC address which is associated with CHPID 0B.

Nothing had to be done at the workstation to update ARP cache. The TCP/IP running on z/OS initiated a gratuitous ARP to all hosts on the LAN when it was notified that the connection on CHPID 09 was lost.

Example G-7 shows the contents of the ARP table on z/OS immediately after pulling the CAT5 cable for CHPID 09.

Example: G-7 z/OS TCP/IP ARP table after CAT5 cable for CHPID 09 pulled

```
MVS TCP/IP NETSTAT CS V1R4       TCPIP Name: TCPIPD         15:37:16
Querying ARP cache for address 192.168.1.64
Link: SC642180LINK      ETHERNET: 0006296CA5BC

Querying ARP cache for address 192.168.1.10
Link: SC642180LINK      ETHERNET: 0004AC1D18E9

Querying ARP cache for address 192.168.1.4
Link: SC642360LINK      ETHERNET: 0006296C3690
```

Notice that there is no change yet in the MAC addresses in this ARP table.

After a few minutes, the ARP table contained entries associating the same IP address with both MAC addresses, as shown in Example G-8 on page 229.

Example: G-8 z/OS TCP/IP ARP table minutes after CAT5 cable for CHPID 09 pulled

```
MVS TCP/IP NETSTAT CS V1R4       TCPIP Name: TCPIPD         15:39:23
Querying ARP cache for address 192.168.1.64
Link: SC642180LINK      ETHERNET: 0006296CA5BC

Querying ARP cache for address 192.168.1.10
Link: SC642180LINK      ETHERNET: 0004AC1D18E9

Querying ARP cache for address 192.168.1.4
Link: SC642360LINK      ETHERNET: 0006296C3690

Querying ARP cache for address 192.168.1.64
Link: SC642180LINK      ETHERNET: 0006296C3690
```

We then cleaned up the ARP table by using the **purgecache** command available with z/OS 1.4. Example G-9 shows an example of the command.

Example: G-9 z/OS PURGECACHE command

```
V TCPIP,TCPIPD,PURGECACHE,SC642180LINK
EZZ0060I PROCESSING COMMAND: VARY TCPIP,TCPIPD,PURGECACHE,SC642180LINK
EZZ9786I PURGECACHE PROCESSED FOR LINK SC642180LINK
EZZ0053I COMMAND PURGECACHE COMPLETED SUCCESSFULLY
```

When the CAT5 cable was plugged back in to the switch, we received the messages in Example G-10 at the z/OS operator's console.

Example: G-10 CAT5 cable for CHPID 09 plugged into switch

```
EZZ4313I INITIALIZATION COMPLETE FOR DEVICE OSA2180
EZZ4313I INITIALIZATION COMPLETE FOR DEVICE OSA2180
```

Another gratuitous ARP was issued by TCPIPD to the hosts on the LAN that updated the ARP cache with the correct MAC addresses.

On our Windows workstation, the gratuitous ARP only updated existing entries in ARP cache; it did not create entries for IP addresses that were not currently in ARP cache.

Stopping the device in the TCP/IP stack

In this test, we created an error condition by stopping the device in the TCP/IP stack. Example G-11 shows the command that was issued and the resulting messages.

Example: G-11 Induced error by stopping the device in TCPIPD

```
V TCPIP,TCPIPD,STOP,OSA2180
EZZ0060I PROCESSING COMMAND: VARY TCPIP,TCPIPD,STOP,OSA2180
EZZ0053I COMMAND VARY STOP COMPLETED SUCCESSFULLY
EZZ4315I DEACTIVATION COMPLETE FOR DEVICE OSA2180
EZZ4329I LINK SC642360LINK HAS TAKEN OVER ARP RESPONSIBILITY FOR
INACTIVE LINK SC642180LINK
```

A Gratuitous ARP was sent and the changes to the ARP tables were identical to those shown in Example G-6 on page 228 and Example G-7 on page 228.

We started the device again, as shown in Example G-12 on page 230.

Example: G-12   Starting the device in TCPIPD

```
V TCPIP,TCPIPD,START,OSA2180
EZZ0060I PROCESSING COMMAND: VARY TCPIP,TCPIPD,START,OSA2180
EZZ0053I COMMAND VARY START COMPLETED SUCCESSFULLY
EZZ4313I INITIALIZATION COMPLETE FOR DEVICE OSA2180
```

Once again, the gratuitous ARP was sent to update existing ARP cache entries to the correct MAC addresses.

H

IPv6 function list

This appendix summarizes the z/OS TCP/IP stack-related functions and the level of support provided in an IPv6 network. It is anticipated that many more of these functions will be enabled for IPv6 support in subsequent releases of z/OS Communications Server.

© Copyright IBM Corp. 2003. All rights reserved.

IPv6 Table

Table H-1 can be used to determine whether a given function is applicable to IPv6.

Table H-1 z/OS TCP/IP stack function support

z/OS TCP/IP stack function	IPv4 support	IPv6 support	Comments
Link-layer device support	Y	Y	IPv4 devices are defined with the DEVICE and LINK configuration statements. In IPv6, interfaces are defined with the INTERFACE statement.
Ethernet LAN connectivity using OSA-Express in QDIO mode	Y	Y	To define an MPCIPA device for IPv4, use the DEVICE statement with the MPCIPA parameter and the LINK statement with the IPAQENET parameter. For IPv6 traffic, OSA-Express QDIO fast ethernet and gigabit ethernet support is configured using an INTERFACE statement of type IPAQENET6
Virtual IP addressing support			
Virtual device/interface configuration	Y	Y	With IPv4, a static virtual device is configured using DEVICE and LINK statements with the VIRTUAL parameter. An IPv6 virtual interface is configured with an INTERFACE statement of type VIRTUAL6.
Sysplex support	Y	N	Related configuration statements: VIPADYNAMIC IPCONFIG
IP routing functions			
Dynamic routing - OSPF and RIP	Y	N	
Dynamic routing - Auto Configuration	N	Y	
Static route configuration using BEGINROUTES statement	Y	Y	Related configuration statements: BEGINROUTES
Static route configuration using GATEWAY statement	Y	N	Related configuration statements: GATEWAY
Multipath routing groups	Y	Y	Related configuration statements: IPCONFIG IPCONFIG6

z/OS TCP/IP stack function	IPv4 support	IPv6 support	Comments
Miscellaneous stack functions			
Path MTU discovery	Y	Y	Path MTU discovery is mandatory in IPv6. Related configuration statements: IPCONFIG IPCONFIG6
Link-layer address resolution	Y	Y	In IPv4, performed using Address Resolution Protocol (ARP). In IPv6, performed using neighbor discovery protocol. Related configuration statements: DEVICE and LINK (LAN Channel Station and OSA devices) INTERFACE (IPAQENET6 interfaces)
ARP/Neighbor cache PURGE capability	Y	Y	Use the V TCPIP,PURGECACHE command. For information, see *z/OS Communications Server: IP System Administrator's Commands*
Datagram forwarding enable/disable	Y	Y	Related configuration statements: IPCONFIG IPCONFIG6
Transport-layer functions			
Fast response cache accelerator	Y	N	
Enterprise extender	Y	N	
Server-BIND control	Y	Y	Related configuration statements: PORT
UDP Checksum disablement option	Y	N	UDP checksum is required when operating over IPv6. Related configuration statements UDPCONFIG
Network management and accounting functions			
SNMP	Y	N	
Policy-based networking	Y	N	

z/OS TCP/IP stack function	IPv4 support	IPv6 support	Comments
SMF SMFCONFIG	Y	Y	Type 118 records do not support IPv6 addresses. IPv6 support in type 119 records is being phased in. Currently, only the following records provide IPv6 support: TCP connection initiation TCP connection termination UDP socket close FTP client transfer completion FTP server transfer completion FTP server logon failure Related configuration statements: SMFCONFIG
IPSec	Y	N	Related configuration statements: IPCONFIG
IP filtering	Y	N	
Network access control	Y	N	Related configuration statements: NETACCESS
Stack and port access control	Y	Y	Related configuration statements: PORT DELETE
Intrusion detection services	Y	N	

Related publications

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

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 236.

- *zSeries 900 Open Systems Adapter-Express Implementation Guide,* SG24-5948
- *Migrating Subarea Networks to an IP Infrastructure Using Enterprise Extender,* SG24-5957
- *Large Scale Linux Deployment*, SG24-6824
- *zSeries HiperSockets*, SG24-6816
- *IBM zSeries Connectivity Handbook*, SG24-5444
- *Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 6: Policy and Network Management*, SG24-6839
- *Managing OS/390 TCP/IP with SNMP*, SG24-5866

Other resources

These publications are also relevant as further information sources:

- *Open Systems Adapter-Express Customer’s Guide and Reference for zSeries,* SA22-7476
- *Open Systems Adapter-Express Customer’s Guide and Reference for S/390,* SA22-7403
- *z/OS Open Systems Adapter Support Facility User’s Guide,* SC28-1855
- *VM/ESA Open Systems Adapter Support Facility User’s Guide,* SC28-1992
- *VSE/ESA Open Systems Adapter Support Facility User’s Guide,* SC28-1946
- *OS/390 eNetwork Communications Server: SNA Network Implementation Guide,* SC31-8563
- *z/OS V1R4.0 Resource Measurement Facility (RMF) Report Analysis*, SC33-7991
- *Planning for the Open Systems Adapter-2 Feature for zSeries,* GA22-7477
- *Network and e-business Products Reference booklet,* GX28-8002
- *z/OS V1R4.0 Communications Server IP Configuration Guide*, SC31-8775
- *z/OS V1R4.0 Communications Server IP Configuration Reference*, SC31-8776
- *z/OS V1R4.0 Communications Server IP System Administrator's Commands,* SC31-8781
- *z/OS V1R4.0 Communications Server SNA Resource Definition Reference*, SC31-8778
- *z/OS V1R4.0 Resource Measurement Facility Report Analysis,* SC33-7991
- *z/OS V1R4.0 Communications Server IP User's Guide and Commands,* SC31-8780
- *z/OS V1R4.0 Communications Server IP Programmer's Reference,* SC31-8787
- *z/VM Installation Guide Version 4 Release 3.0*, GC24-5992

© Copyright IBM Corp. 2003. All rights reserved.

- *z/VM TCP/IP Level 430 Planning and Customization*, SC24-6019
- *z/VM TCP/IP Level 430 User's Guide*, SC24-6020
- *z/VM V4R3.0 CP Command and Utility Reference*, SC24-6008

Referenced Web sites

These Web sites are also relevant as further information sources:

- zSeries Networking

  `http://www.ibm.com/servers/eserver/zseries/networking/`
- zSeries Networking White Papers

  `http://www.ibm.com/servers/eserver/zseries/networking/wpapers.html`
- IBM Documentation and tools

  `http://www.ibm.com/servers/resourcelink/`
- z/OS information wizardry

  `http://www.ibm.com/servers/eserver/zseries/zos/wizards/`
- Linux IBM

  `http://www.ibm.com/servers/eserver/zseries/os/linux/`
  `http://www10.software.ibm.com/developerworks/opensource/linux390/index.shtml`
  `http://www10.software.ibm.com/developerworks/opensource/linux390/overview24.shtml`
- *OSA-Express for zSeries 900 and S/390 Specification Sheet,* G221-9110

  `http://www.ibm.com/eserver/zseries`

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

© Copyright IBM Corp. 2003. All rights reserved.

IBM Redbooks

OSA-Express Implementation Guide

(0.2"spine)
0.17"<->0.473"
90<->249 pages

OSA-Express Implementation Guide

Product, planning, and quick start information

Realistic examples and considerations

Hardware and software setup definitions

This IBM Redbook discusses how to install, tailor, and configure the Open System Adapter-Express features that are available on the 9672 Generation 5 and Generation 6 Parallel Enterprise Servers, and zSeries Servers. It focuses on the hardware installation as well as the software definitions needed to provide connectivity to various LAN environments.

This book provides information for planning purposes and system setup. Also included are helpful utilities and commands for monitoring and managing the OSA-Express features.

This document is intended for system engineers, network administrators, and system programmers who will plan and install OSA-Express. A solid background in hardware (9672 Generation 5 and Generation 6 Parallel Enterprise Servers, and zSeries Servers), HCD or IOCP, OSA/SF, SNA/APPN, and TCP/IP, is assumed.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-5948-02 ISBN 0738427934