

*SunLink[®] OSI 8.1
Communication Platform Administrator's
Guide*



SunSoft

A Sun Microsystems, Inc. Business

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

Part No.: 801-4975-13
Revision A, March 1995

© 1995 Sun Microsystems, Inc.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX[®] and Berkeley 4.3 BSD systems, licensed from UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and the University of California, respectively. Third-party font software in this product is protected by copyright and licensed from Sun's font suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 (c)(1)(ii) and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, the Sun logo, Sun Microsystems, Solaris and SunLink are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and certain other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. OPEN LOOK is a registered trademark of Novell, Inc. PostScript and Display PostScript are trademarks of Adobe Systems, Inc. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCstorage, SPARCware, SPARCcenter, SPARCclassic, SPARCcluster, SPARCdesign, SPARC811, SPARCprinter, UltraSPARC, microSPARC, SPARCworks, and SPARCcompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

X Window System is a product of the Massachusetts Institute of Technology.

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

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Contents

1. Software Architecture Overview	1
OSI Reference Model	1
Lower Layer Infrastructure	3
Physical Layer	3
Data Link Layer	3
Network Layer	3
Transport Layer	4
Upper Layer Infrastructure	4
Session Layer	4
Presentation Layer	4
Application Layer	4
Introducing SunLink OSI 8.1	4
Components of SunLink OSI 8.1	6
Stack Boot File	6
Startup Daemon	6

Configuration File	6
STREAMS Modules	6
OSI Administration Tool	7
Chapter Summary	7
2. Using the OSI Administration Tool.....	9
Before You Start	9
RPC-Based Utilities	10
Starting the Stack.....	10
Starting <code>ositool</code>	11
Introducing the OSI Administration Tool (<code>ositool</code>).....	12
<code>ositool</code> Main Window	12
Save	13
Command Menu	14
Stack Manager	15
ES-IS Configuration.....	17
Network Layer Addresses	18
Route Manager.....	19
Configuration Steps.....	20
Chapter Summary	20
3. Using Stack Manager	21
Stack Parameters Summary	22
Stack Manager Main Window	23
Presentation & ACSE Entity.....	24
Session Entity	26

Additional Session Options	29
Transport & CLNS Entity	30
Additional Transport & CLNS Options.....	33
Transport over CONS Entity	36
Additional Transport over CONS Options	40
CONS Entity.....	43
Configuration Menu	45
Application Selectors Window.....	47
SAP List.....	49
Device Configuration	51
LAN Device.....	52
X.25 Subnetwork Device	53
RFC1006 Subnetwork Device	55
Resource Configuration	56
Chapter Summary	59
4. End System to Intermediate System Configuration	61
Subnet List	63
Subnet ID	64
Route Table Size Options	65
CLNP Options	66
ES-IS Options	68
Chapter Summary	71
5. Network Layer Addresses	73
NSAP Family Window	74

nbs Address Structure	75
osinet Address Structure.	77
us-gossip-v1 Address Structure.	79
us-gossip-v2 Address Structure.	81
user-defined Address Structure	83
hex-pub Address Structure.	85
free-form Address Structure.	86
Chapter Summary	87
6. Route Manager	89
Route Manager Main Window	89
Host Route Configuration.	92
Prefix Routes.	94
Remote X.25 Features	98
Chapter Summary	101
7. Addressing.	103
Overview of OSI Addressing	103
Subnetwork Point of Attachment (SNPA)	105
Network Service Access Point (NSAP)	106
Initial Domain Part (IDP)	107
Domain Specific Part (DSP).	110
NSAP Address Field Lengths and Syntax.	110
Types of NSAP Addresses.	112
Selectors	115
Network Entity Title (NET)	116

Chapter Summary	116
8. Dynamic Routing	117
ES-IS Protocol	117
How ES-IS Works	118
End System Responsibilities	120
Intermediate System Responsibilities	121
Routing Sequence	122
Using Token Ring Networks	122
CLNP over X.25	123
Chapter Summary	123
9. Troubleshooting and Diagnostics	125
Troubleshooting Overview	125
Consistency	125
Installation	126
Terminating Applications to Restart the Daemon	126
Overview of Diagnostics	127
The <code>osi_ping</code> Function	127
The <code>osi_trace</code> Function	129
<code>osi_trace</code> Options	130
<code>osi_trace</code> Filter Expressions	131
Trace Examples	132
The <code>osi_decode</code> Program	135
Problems Interworking with SunPro Products	138
Chapter Summary	138

A. Configuration File	139
Command Components	140
Tokens	140
Commands	142
Verbose Trace Command	142
Opening a Device Driver	142
Set Command	146
Link Command	147
Operator Statements	148
Close Command	148

Figures

Figure 1-1	OSI Reference Model.	2
Figure 1-2	SunLink OSI 8.1 Entities	5
Figure 2-1	OSI Administration Tool Window	13
Figure 2-2	The Command Menu	14
Figure 2-3	Console Window	15
Figure 2-4	Stack Manager	16
Figure 2-5	ES-IS Configuration.	17
Figure 2-6	Network Layer Addresses	18
Figure 2-7	Route Manager.	19
Figure 3-1	Stack Manager Configuration Parameter Summary	22
Figure 3-2	Stack Manager Window	23
Figure 3-3	Presentation & ACSE Configuration	24
Figure 3-4	Session Configuration.	26
Figure 3-5	Additional Session Options	29
Figure 3-6	Transport & CLNS Configuration.	30
Figure 3-7	Additional Transport & CLNS Options	33

Figure 3-8	Transport over CONS Configuration	36
Figure 3-9	Additional Transport over CONS Options	40
Figure 3-10	CONS Configuration.	43
Figure 3-11	Configuration Menu	45
Figure 3-12	Application Selectors Window	47
Figure 3-13	SAP List Window	49
Figure 3-14	Device Configuration Window	51
Figure 3-15	LAN Subnetwork Device	52
Figure 3-16	X.25 Subnetwork Device.	54
Figure 3-17	RFC1006 Subnetwork Device.	55
Figure 3-18	Resource Configuration Window	56
Figure 4-1	ES-IS Configuration.	62
Figure 4-2	Subnet List	63
Figure 4-3	Subnet Identifiers.	64
Figure 4-4	Route Table Size Options	65
Figure 4-5	CLNP Protocol Options	66
Figure 4-6	ES-IS Options	68
Figure 5-1	NSAP Family Window	74
Figure 5-2	nbs NSAP Format	75
Figure 5-3	osinet NSAP Format	77
Figure 5-4	us-gossip-v1 NSAP Format	79
Figure 5-5	us-gossip-v2 NSAP Format	81
Figure 5-6	user-defined NSAP Format	83
Figure 5-7	hex-pub NSAP Format	85
Figure 5-8	free-form NSAP Format	86

Figure 6-1	Route Manager Window	90
Figure 6-2	Host Routes Window	92
Figure 6-3	Prefix Routes Window	95
Figure 6-4	X.25 Features Menu.....	98
Figure 7-1	Network Addressing.....	105
Figure 7-2	NSAP Components	107
Figure 7-3	Network Connection.....	115

Tables

Table P-1	Typographic Conventions	xxii
Table 4-1	Default Values for CLNP/LLC1 Subnetworks	70
Table 4-2	Default Values for CLNP/X.25 Subnetworks	70
Table 7-1	IDI Descriptions	108
Table 7-2	IDI Leading Significance of Zeroes	109
Table 7-3	NSAP Address Field Lengths	111
Table A-1	Keyword Parameters	140

Preface

This book describes the procedure for the configuration and administration of the SunLink OSI 8.1 stack and its associated modules. It describes:

- How to use the graphical user interface to configure your system for network applications. The OSI Administration Tool, `ositool`, provides a menu system with which you can change the parameters and addresses associated with the stack, end systems and intermediate systems, network layer, and routes.
- Fundamental rules and concepts about how to deal with addressing and routing in your network, and explains some procedures for avoiding problems with your network.

Who Should Use This Book

This book is written for the system administrator who is configuring SunLink OSI 8.1 for its applications. You need to know your system and the network devices and setup that you have, for example, how each network device is connected and what applications you need to configure. You will also find that a good understanding of OSI and other networking principles and terminology are necessary, such as network addressing and routing techniques.

You should have installed the SunLink OSI 8.1 package according to *Installing and Licensing SunLink 8.1*.

How This Book Is Organized

The *SunLink OSI 8.1 Communication Platform Administrator's Guide* is organized as follows:

Chapter 1, “Software Architecture Overview” is an overview of the OSI architecture and how it relates to the SunLink OSI stack architecture. It introduces some networking concepts and the structure of the software.

Chapter 2, “Using the OSI Administration Tool” describes the OSI Administration Tool, its components, and how it is used.

Chapter 3, “Using Stack Manager” describes how to set parameters for the network entities, specify access points for applications, configure network devices, and determine connection resources.

Chapter 4, “End System to Intermediate System Configuration” describes how to configure end systems, intermediate systems, and subnetworks.

Chapter 5, “Network Layer Addresses” describes how to set network addresses.

Chapter 6, “Route Manager” describes how to update the routes used for CONS and CLNP subnetworks.

Chapter 7, “Addressing” explains some of the main addressing components of SunLink OSI 8.1.

Chapter 8, “Dynamic Routing” explains how the ES-IS routing protocol, based on the ISO 9542 standards, is implemented in SunLink OSI 8.1.

Chapter 9, “Troubleshooting and Diagnostics” discusses some problems that you might encounter with the configuration of SunLink OSI 8.1 and provides some suggestions for avoiding error conditions.

Appendix A, “Configuration File” provides examples of the configuration file that is updated when you change the configuration parameters.

Related Documentation

The other documents for SunLink OSI 8.1 are:

- *Installing and Licensing SunLink 8.1* describes the procedure for installing and licensing SunLink OSI 8.1. It also describes the components and modules that make up the SunLink OSI 8.1 package. Use this book to ensure that the software is installed correctly before proceeding to the configuration procedure explained in the *SunLink OSI 8.1 Communication Platform Administrator's Guide*.
- *SunLink OSI 8.1 TLI Programmer's Reference* describes the Transport Level Interface (TLI), allowing you to develop applications that access the transport layer services directly. It describes its functions and explains how to compile and link application programs.
- *SunLink OSI 8.1 APLI Programmer's Reference* describes the ACSE Presentation Library Interface (APLI), so you can develop applications that access the presentation layer services directly. It describes its functions and how to compile and link application programs.

Standards Reference

SunLink OSI 8.1 is based on the following specifications and standards:

General

- ISO/IEC-7498 *Basic Reference Model for OSI* (CCITT X.200)
- ISO/IEC-7498 (ADD1) *Basic Reference Model Connectionless – Mode Transmission* (CCITT X.200)
- ISO/IEC-7498-2 (Part 2) *Security Architecture* (CCITT X.200)
- ISO/IEC-7498-3 (Part 3) *Naming and Addressing* (CCITT X.200)
- ISO/IEC-7498-4 (Part 4) *Management Framework* (CCITT X.200)

ACSE

- ISO/IEC-8649 *Service Definition for the Association Control Service Element (ACSE)* (CCITT X.217)
- ISO/IEC-8650 *Protocol Specification for the Association Control Service Element (ACSE)* (CCITT X.227)

ROSE

- ISO/IEC-9072-1 (Part 1) *Remote Operations – Model, Notation, and Service Definition* (CCITT X.219)
- ISO/IEC-9072-2 (Part 2) *Protocol Specifications* (CCITT X.229)

Presentation Layer

- ISO/IEC-8822 *Connection-Oriented Presentation Services Definition* (CCITT X.216)
- ISO/IEC-8823 *Connection-Oriented Presentation Protocol Specification* (CCITT X.226)
- ISO/IEC-8824 *Specification of Abstract Syntax Notation 1 (ASN.1)* (CCITT X.208-88)
- ISO/IEC-8824 (ADD 1) *Extensions to ASN.1* (CCITT X.208-88)
- ISO/IEC-8825 *Specification of Basic Encoding Rules for Abstract Syntax Notation 1 (ASN.1)* (CCITT X.209-88)
- ISO/IEC-8825 (ADD 1) *Extensions to ASN.1 Basic Encoding Rules* (CCITT X.209-88)

Session Layer

- ISO/IEC-8326 *Basic Connection-Oriented Session Service Definition* (CCITT X.215)
- ISO/IEC-8326 (ADD 2) *Incorporation of Unlimited User Data* (CCITT X.215)
- ISO/IEC-8327 *Basic Connection-Oriented Session Protocol Specification* (CCITT X.225)
- ISO/IEC-8327 (ADD 2) *Incorporation of Unlimited User Data* (CCITT X.225)

Transport Layer

- ISO/IEC-8072 *Transport Service Definition* (CCITT X.214)
- ISO/IEC-8073 *Connection-Oriented Transport Protocol Specification* (CCITT X.224)
- ISO/IEC-8073 (ADD 2) *Class 4 Operation Over Connectionless Network Service* (CCITT X.224)
- ISO/IEC-8602 *Connectionless Transport Protocol (CLTP)*

Network Layer

- ISO/IEC-8348 *Network Service Definition*
(CCITT X.213)
- ISO/IEC-8348 (ADD 1) *Connectionless-Mode Transmission*
(CCITT X.213)
- ISO/IEC-8348 (ADD 2) *Network Layer Addressing*
(CCITT X.213)
- ISO/IEC-8473 *Protocol for Providing the Connectionless-Mode Network Service*
Internet Protocol (IP)
- ISO/IEC-8648 *Internal Organization of the Network Layer*
- ISO/IEC-8878 *Use of X.25 to Provide the OSI Connection-Oriented Network Service*
(CCITT X.223)
- ISO/IEC-9542 *End System to Intermediate System Routing Exchange Protocol*

Data Link Layer

- ISO/IEC-8802-2 *Logical Link Control (LLC)*

Further Reading

The following books are useful for more detailed information about OSI networking concepts and implementation:

- Tanenbaum, Andrew S., “*Computer Networks*” Second Edition. Prentice-Hall International Editions, 1988.
- Black, Uyless, “*OSI A Model for Computer Communications Standards*” Prentice-Hall, 1991.
- Black, Uyless, “*Network Management Standards (The OSI, SNMP and CMOL Protocols)*” McGraw-Hill on Computer Communications, 1992.
- Cypser, R.J., “*Communications for Cooperating Systems OSI, SNA, and TCP/IP*” Addison-Wesley Publishing, 1992.
- Rose, Marshall, “*The Open Book—A Practical Perspective on OSI*” Prentice-Hall, 1990.

Typographic Conventions

The following table describes the fonts and symbols used in this book.

Table P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
Typewriter	The names of commands, files, and directories; computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. system% You have mail.
boldface	User input; what you type	<div>system% su Password:</div>
<i>italic</i>	Command-line placeholder: replace with an actual name or value	To delete a file, type <code>rm filename</code> .
	Book titles, new words or terms requiring emphasis	See Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.
Code samples are placed in boxes and may display the following output:		
%	UNIX C shell prompt	system%
\$	UNIX Bourne and Korn shell prompt	system\$
#	Superuser prompt, all shells	system#

The mouse buttons used in this manual are described as:

- SELECT, to select a window or options from it, or to work a window control. This is usually the left-hand button on your mouse.
- MENU, to open a pull-down menu and select one of its choices. This is usually the right-hand button on your mouse.

Software Architecture Overview

1 

<i>OSI Reference Model</i>	<i>page 1</i>
<i>Introducing SunLink OSI 8.1</i>	<i>page 4</i>
<i>Chapter Summary</i>	<i>page 7</i>

This chapter provides an overview of the OSI reference model and how it relates to the architecture of the SunLink OSI 8.1 software. To learn more about the OSI reference model, refer to the books listed in “Further Reading” on page xxi. For more detailed information about the components in the SunLink OSI 8.1 software, refer to the *Installing and Licensing SunLink 8.1* document.

OSI Reference Model

The OSI model for network communications defines seven layers, each of which performs specific communications operations independent of the other layers. By dividing the function of communicating between applications into specific simplified tasks, each layer deals only with operations related to its set of tasks, providing the next layer up with a transparent service.

For example, the transport layer takes information from the network layer, performs its defined transport operation on the information, and then passes it to the session layer. The session layer sees only the service provided by the transport layer, it does not see the operations that occurred in the lower layers.

The lower, network-dependent, layers provide services related to the physical connections, types of links, and routing functions. The upper application-oriented layers perform services related to session management, data abstraction, and applications.

Figure 1-1 illustrates the seven layers and highlights the services that they perform.

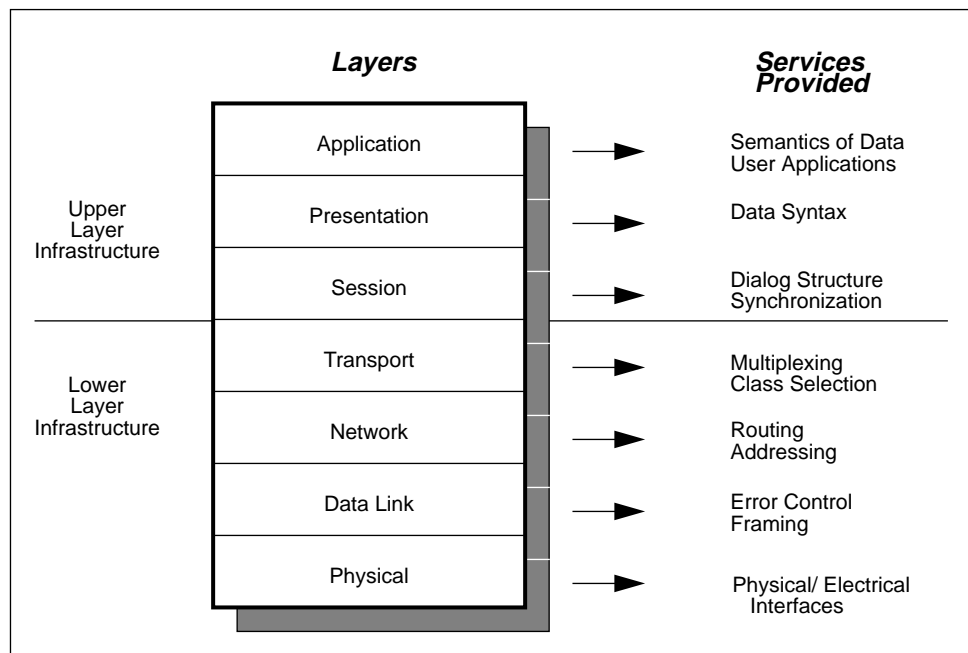


Figure 1-1 OSI Reference Model

Since an application uses each of these seven layers to communicate with another application, each layer has a corresponding peer protocol in another application. Peer protocols perform equivalent functions in the sequence of operations as information is passed through the seven layers. Applications communicating with each other across peer protocols obtain a transparent communications interface. Each layer sees only the service provided by the interface of the entity below it.

The implementation of the interface between each of these layers is defined by the set of standards for services and protocols which make up the *OSI Reference Model*. By using this as a framework, the SunLink OSI 8.1 software allows communication between diverse network configurations and applications.

To illustrate the SunLink OSI 8.1 implementation, it is useful to have an overview of the functions and services provided by the OSI standard layers.

Lower Layer Infrastructure

The lower layers provide transparent connections over diverse network configurations and provide a consistent interface to the upper layers.

Physical Layer

This layer defines the physical and electrical interface connections between communicating systems. For example, for an X.25 network, it defines the physical connection between the Data Terminating Equipment and the Data Circuit Terminating Equipment.

Data Link Layer

This layer takes the information provided by the physical layer and adds error detection and retransmission functions. At this stage data is treated as units of data.

Network Layer

This layer provides addressing and routing functions, and may also include flow control between networks. It provides two types of service:

- Connectionless Network Service (CLNS)—treats each data frame separately for data transfer. For example, IP uses a connectionless service.
- Connection Oriented Network Service (CONS)—establishes a connection over the network and then transfers the data. For example, X.25 uses a connection-oriented service.

Transport Layer

This layer provides an interface between the upper layers and the lower layers, concealing the detailed functional operation of the physical network connections to provide a network-independent service to the application-oriented upper layers. It allows a choice of transport class, according to the type of network, and dependent on the Quality of Service (QOS) options required.

The transport layer handles the multiplexing of connections, that is, it allows a choice of connections through the layered protocol. The Connectionless Transport Protocol (CLTP) is supported.

Upper Layer Infrastructure

The upper layers provide services that handle the applications, and the structuring, and encoding of data.

Session Layer

This layer provides organizing functions for synchronizing dialog and session recovery from lower layer problems.

Presentation Layer

This layer negotiates a common syntax used to encode data for data transfer. It allows data to be transferred, independent of hardware considerations.

Application Layer

This layer provides services to the user and applications, such as job control, file transfer facilities, electronic mail, virtual terminal and directory services. It concerns the semantics of the data, that is, the meaning or use of the data, rather than the techniques involved in transferring it.

Introducing SunLink OSI 8.1

SunLink OSI 8.1 implements the OSI model to allow communications between different applications over diverse networks.

You configure the protocol characteristics for your network by modifying the *entities*. These are shown in Figure 1-2 on page 5 in relation to the OSI layers. The entities can be considered as approximately equivalent to the protocols, apart from the Transport and CLNS entity. Using these entities, you can define how applications at the top layers communicate through peer layers with another application. The total of these entities is known as the *stack*.

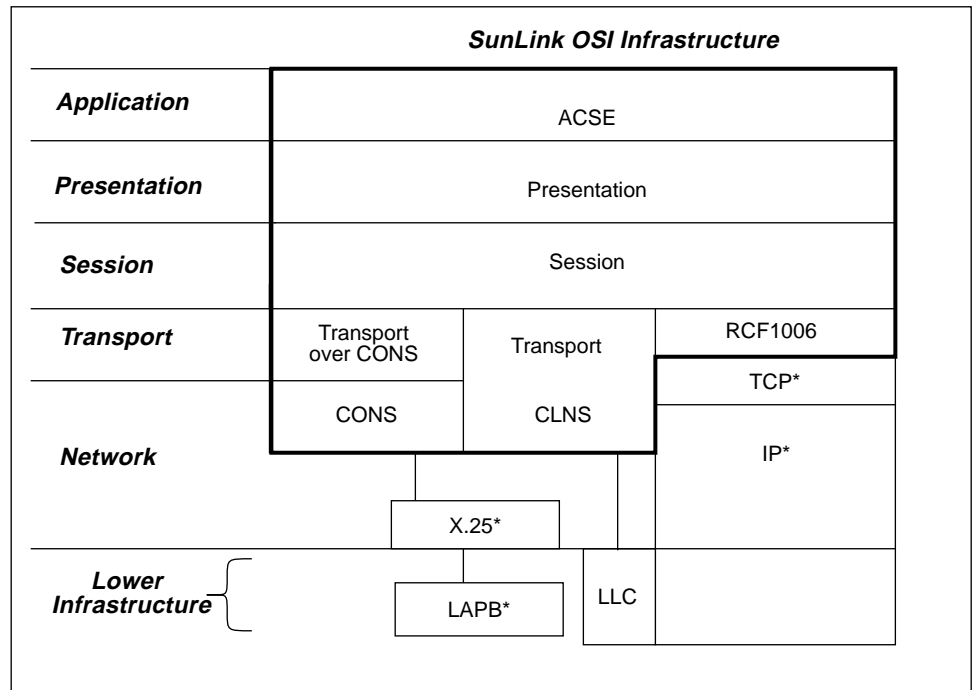


Figure 1-2 SunLink OSI 8.1 Entities

Note – Items marked with * are not included in the SunLink OSI software:
 TCP/IP is included in the Solaris operating system
 X.25 software is included in the SunLink X.25 package

Components of SunLink OSI 8.1

The SunLink OSI 8.1 software is divided into packages, each of which performs a specific task. These are described in more detail in the *Installing and Licensing SunLink 8.1*. It is useful though, to recognize which files and programs are relevant to the configuration process.

Stack Boot File

The `S90osinet` module boots the stack, that is, the software which comprises the entities shown Figure 1-2. It performs various checks on the stack and then calls the startup daemon.

Startup Daemon

The `osinetd` daemon loads the software and its configuration, linking the relevant STREAMS modules. It is invoked by the stack boot file, `S90osinet` and uses the information about your setup contained in the configuration file.

Configuration File

The `osinetd.conf` file contains the parameters and options that customize the software according to your network setup. You can modify this file with the OSI administration tool (`ositool`), where necessary. The default values have been chosen to be valid for most networks, so that many parameters do not require modification.

STREAMS Modules

The STREAMS modules are used to improve the manageability and performance of the interface with other STREAMS modules. ACSE, presentation, session, and transport entities are contained within one STREAMS module, the LLC is contained in a separate module.

OSI Administration Tool

The `osinetd.conf` file is updated using the OSI administration tool, `ositool`. It is a graphical user interface that provides a series of windows and menus with which you can customize the configuration of the SunLink OSI 8.1 for your network.


Chapter Summary

This chapter introduced the main concepts of OSI layered protocols and how they relate to the way in which SunLink OSI 8.1 is implemented.

It introduced briefly the main components of SunLink OSI 8.1 that are relevant for its configuration. More details of the components can be found in *Installing and Licensing SunLink 8.1*.

Chapter 2 describes the OSI administration tool in more detail, the pre-requisites for running it, and the configuration process.

Using the OSI Administration Tool



<i>Before You Start</i>	<i>page 9</i>
<i>Introducing the OSI Administration Tool (ositool)</i>	<i>page 12</i>
<i>Configuration Steps</i>	<i>page 20</i>
<i>Chapter Summary</i>	<i>page 20</i>

This chapter describes the tool provided with SunLink OSI 8.1, which is used to configure your software. The OSI administration tool (`ositool`) is a graphical user interface that provides easy access to the options and parameters that you can set for your system. This chapter lists the prerequisites for the configuration process and how to start `ositool`. A summary of each `ositool` configuration component is provided. Detailed information for each configuration component is given in the relevant chapter.

Before You Start

Ensure that the software has been correctly installed as described in *Installing and Licensing SunLink 8.1*.

You can only use `ositool` to configure the system where you installed the SunLink OSI 8.1 software; it cannot be used to configure a remote system. You can however, *display* `ositool` and modify the configuration for a remote system where the software is installed. Note that only *one* version of `ositool` can be running at a time.

RPC-Based Utilities

If you want to run RPC-based utilities (such as `spray(1)`) over OSI, you should have installed the `SUNWcorpc` package, as explained in *Installing and Licensing SunLink 8.1*. To enable TI-RPC, you have to edit the configuration file. Use a text editor to delete the comment mark (#) from the following line in `/etc/netconfig`:

```
#oclt      tpi_clts      v      osi      datagram_v /dev/oclt      osiaddr.so
```

Now save the configuration file and reboot the machine.

You can add new hosts by editing the `/etc/netconfig/oclt/hosts` file. This should already contain the local host, as defined by the `ositool` configuration. Add other hosts by adding the NSAP and the hostname to the list, as explained in the hosts file.

You can add services by editing the `/etc/netconfig/oclt/services` file. Add the service name and number to the list, as described in the file.

More detailed information about this can be found in the *SunOS 5.3 Network Interfaces Programmer's Guide*.

Starting the Stack

To use `ositool`, you need to adapt your environment variables to ensure that the software can locate your files.

Set the following environment variables to include the specified pathnames, if you have used the defaults paths:

```
PATH          /opt/SUNWconn/bin
MANPATH       /opt/SUNWconn/man
HELPPATH      /opt/SUNWconn/osinet/lib/locale/C
```

You can do this by editing the appropriate file (usually `.cshrc`) or by setting them temporarily using the `setenv` command.

Check if the stack is running by entering:

```
prompt% ps -ef | grep osinetd
```


If the active process `/usr/sbin/osinetd` owned by root is displayed, then the stack is running. Otherwise, perform the following steps to start the stack boot file:

```
prompt% su
prompt# (enter your root password)
prompt# /etc/rc2.d/S90osinet start
```

If it does not start, check “Troubleshooting and Diagnostics” on page 125 for possible causes.

Starting ositool

To start `ositool` you must have root access on the host and your `PATH` environment variable must include the location of the software. If you selected the default base directory during the installation procedure (that is, `/opt/SUNWconn/bin`), then you can copy the command exactly as shown.

```
prompt% su
prompt# (enter your root password)
prompt# cd /opt/SUNWconn/bin
prompt# ./ositool
```

If you did not install the software in `/opt/SUNWconn/bin`, then you need to change your `PATH` environment variable to include this location, or change to that directory and type `./ositool`.

When you start `ositool` the main OSI administration tool window is displayed. If `ositool` does not start correctly, an error message is displayed. Check “Troubleshooting and Diagnostics” on page 125 for possible causes. The example windows and menus shown in this book may vary slightly from the ones displayed on your screen, depending on how you have set up your own environment.

Note – Only one `ositool` can run on a host at a time.

Introducing the OSI Administration Tool (`ositool`)



The `ositool` graphical user interface provides four configuration components:

- Stack Manager—used to set the configurable parameters for entities, access points for applications, subnetwork device types, and channel resources.
- ES-IS Configuration—used to set the configurable parameters for the end system to intermediate system protocol for CLNP LANs and WANs.
- Network Addressing—used to set the access points to the network.
- Route Manager—used to configure routes to remote systems (NSAPs and SNPs).

These components are described in more detail in the following sections. The default values provided with the configuration file should be adequate for most network configurations. Where modification of the parameters and options is required, it is important that you understand the principles of the OSI reference model before you start the configuration process. If you need more information, there is a list of useful books in “Further Reading” on page xxi.

All windows have a context-related help panel. Just press the help key to display the help information relating to your cursor position.

`ositool` Main Window

The main window that is displayed when you start `ositool` displays the four configuration components, a Save button, and a Command pull-down menu. This is illustrated in Figure 2-1.

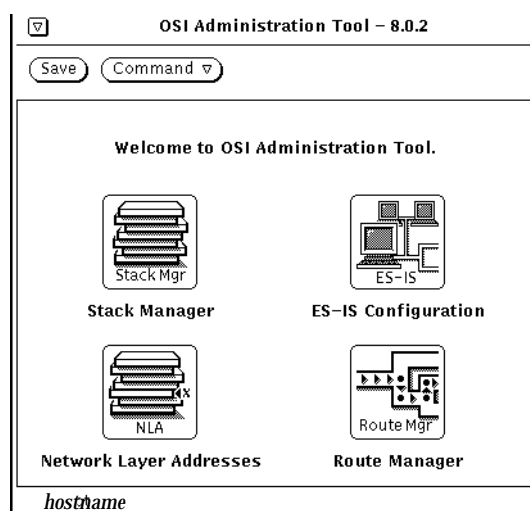


Figure 2-1 OSI Administration Tool Window

Save

The Save button saves the current configuration settings to the `osinetd.conf` file. Some changes that you have made in the configuration will have already been implemented, while other modifications will only take place when you save the configuration and restart the `osinetd` daemon.

A *dynamic* change is implemented as soon as you modify configuration parameters and press the Apply button. Any modifications on that screen are implemented immediately. Dynamic changes take place, for example, when you switch the status of an entity off and press Apply, the stack is updated immediately and the entity is no longer authorized to be used. Dynamic changes are lost when you restart the stack, if you do not save them from the OSI administration tool main window.

A *static* change is only implemented when you save the configuration, (that is, updated the `osinetd.conf` file), and restart the stack. For example, when you add a new subnetwork device to the configuration, the modification takes place when you press Apply, return to the main menu and save the configuration, and then restart the `osinetd` daemon.

Note – If you are migrating to SunLink OSI 8.1 from a previous release, you should save the configuration before you start modifying any configuration parameters. This converts the configuration file to the format required by this release of the software.

Command Menu

The Command menu illustrated in Figure 2-2, allows you to display the system console and to restart the `osinetd` daemon. You need to restart the daemon when you want to implement static changes, since the restart process reads the current version of the `osinetd.conf` file.

Press MENU on the Command button to display the pull-down menu shown in Figure 2-2.




Figure 2-2 The Command Menu

It is useful to display the console before restarting the `osinetd` daemon to monitor any messages that are produced as a result of the restart.

You can only restart the `osinetd` daemon if all other OSI applications have been terminated.



Caution – All applications that access the stack **must** be terminated before the Restart `osinetd` command is selected. Check the console window for error messages. Refer to “Troubleshooting and Diagnostics” on page 125 if you have problems starting `ositool`.

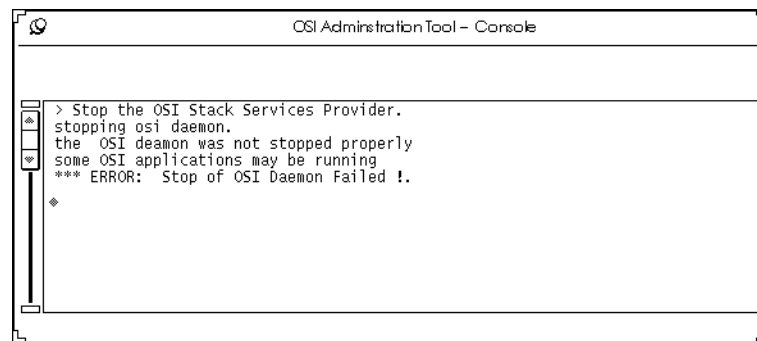


Figure 2-3 Console Window

Figure 2-3 shows an example of the Console window after the `Restart osinetd` command is issued. It indicates that another application was still running when the restart command was issued.

Stack Manager

The Stack Manager allows you to manage and control parameter flags for the OSI layers, Service Access Points (SAP), stack resources, and network access methods. These need to be configured for access between applications (such as messaging and file transfer) and to application programming interfaces (API and TLI). The specific options available from this menu are illustrated in Figure 2-4.

You need to change these parameters according to your subnetwork connections and the types of applications that you want to install.

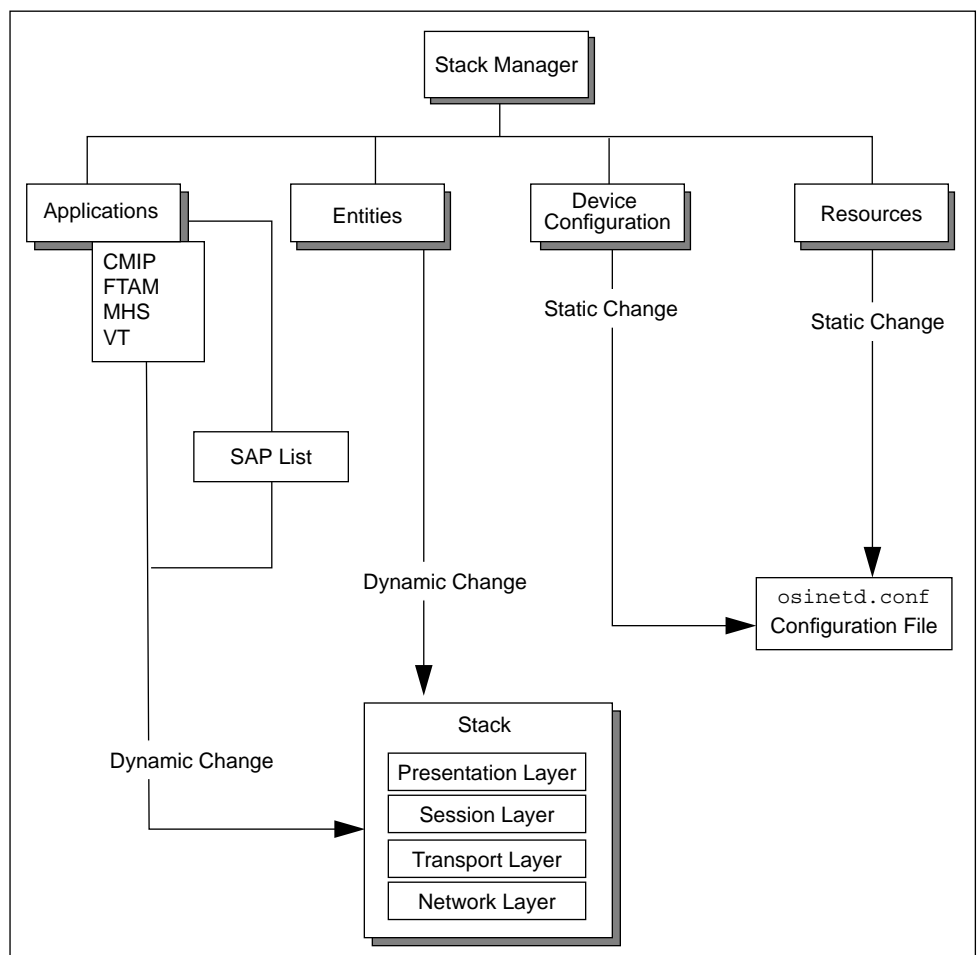


Figure 2-4 Stack Manager

ES-IS Configuration

The End System to Intermediate System (ES-IS) Configuration Manager is used to configure:

- ES-IS options for end system and intermediate system hello and timer settings.
- Buffer options for setting the size of the routing table.
- CLNP options for specifying the type of protocol.

Figure 2-5 illustrates the specific options that you can modify. Note that all changes are static, and are only implemented when you save the configuration and restart the stack.

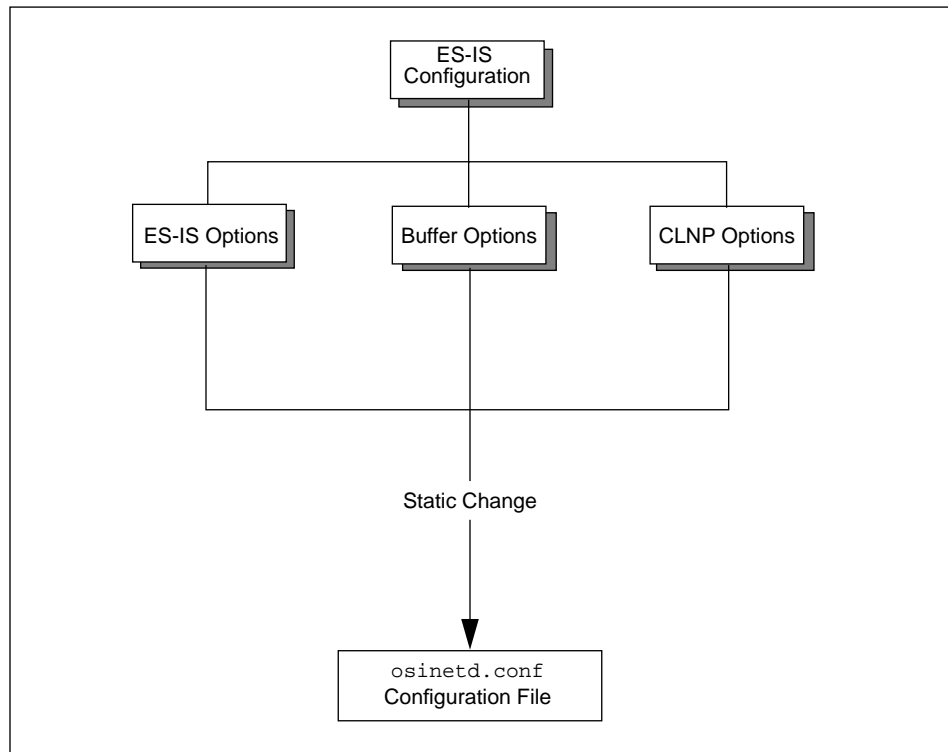


Figure 2-5 ES-IS Configuration

Network Layer Addresses

The Network Layer Address Manager is used to configure:

- Connectionless Network Service Protocol (CLNP) Network Service Access Points (NSAP).
- Connection Oriented Network Service (CONS) NSAPs.
- Network Entity Titles (NET).

Figure 2-6 illustrates the specific options available from this window.

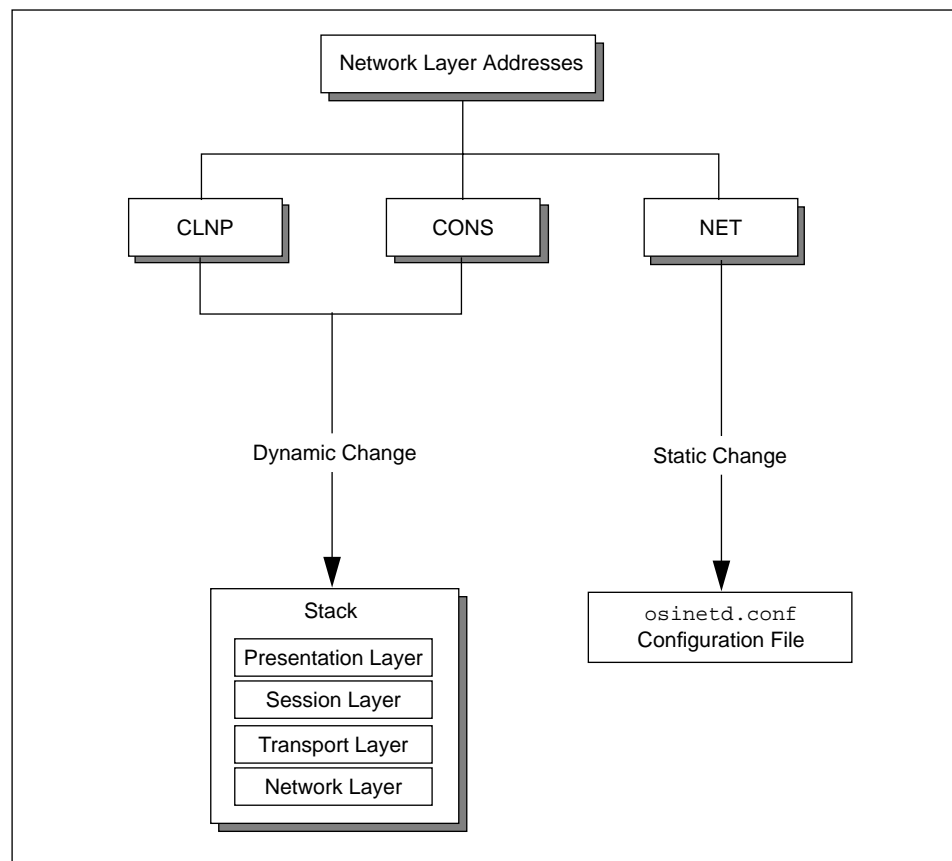


Figure 2-6 Network Layer Addresses

Route Manager

The Route Manager is used to specify:

- Host routes that define a full address for one host.
- Prefix routes that use a masked address, which can be used for a group of hosts.

Figure 2-7 illustrates the specific options that can be modified from this window. All these changes are implemented dynamically.

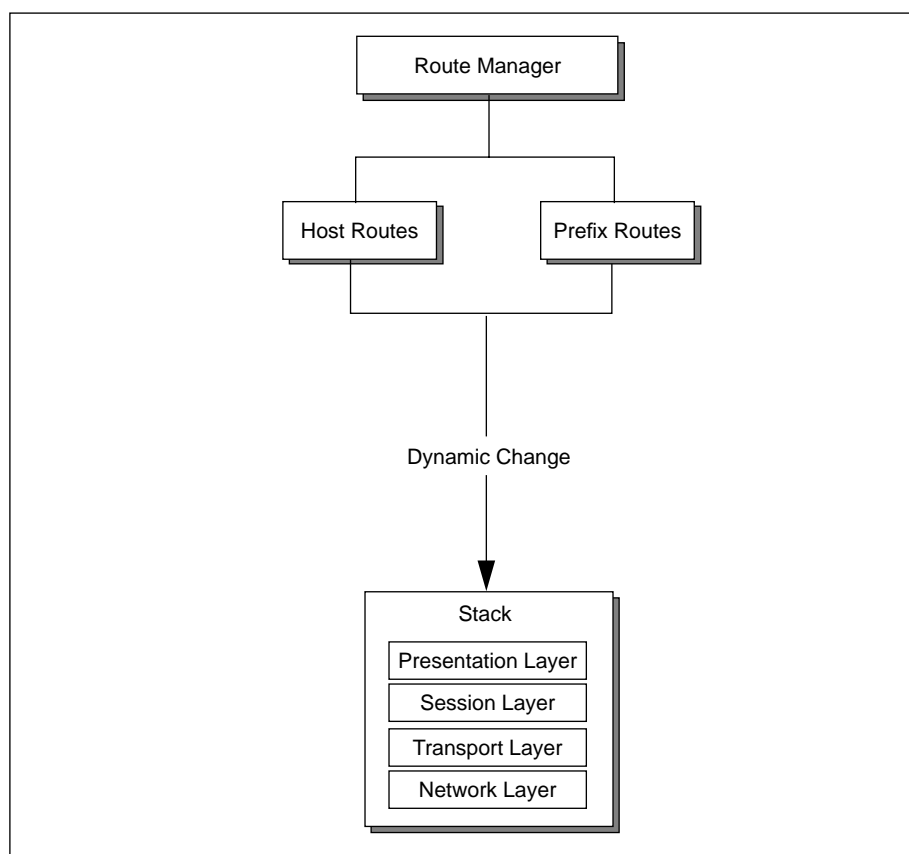


Figure 2-7 Route Manager

Configuration Steps

Each configuration component is described in separate chapters of this book. Some configuration parameters can affect other parameters that are defined in other configuration components. It is therefore a good idea to follow the sequence of this book, that is, configure the Stack Manager first, continuing through to ES-IS Configuration, Network Layer Addressing and Routing management.

Chapter Summary

This chapter introduced the tool with which you configure the SunLink OSI 8.1 software, where necessary. It described the main configuration components of `ositool`, and explained how the dynamic and static changes are implemented.

The procedure for starting up `ositool` and a summary of the steps involved in the configuration process were explained. If you have any problems, check “Troubleshooting and Diagnostics” on page 125 for possible causes, or *Installing and Licensing SunLink 8.1*, in case the software was not installed correctly.

When the `ositool` main window is displayed, you can now go on to modify the appropriate parameters. Chapters 3 through 6 describe each component available in `ositool`.

Using Stack Manager

3



<i>Stack Parameters Summary</i>	<i>page 22</i>
<i>Stack Manager Main Window</i>	<i>page 23</i>
<i>Chapter Summary</i>	<i>page 59</i>

This chapter explains how to use the Stack Manager to configure access points to applications, subnetwork device connections, channel resources, and layer entities. Use the Stack Manager to determine the characteristics of protocol entities and resources in the stack.

Click SELECT on the Stack Manager icon to access the configuration parameters.

Stack Parameters Summary

A summary of the parameters and options for stack configuration is given in Figure 3-1. Examples of the configurable parameters are illustrated in boldface.

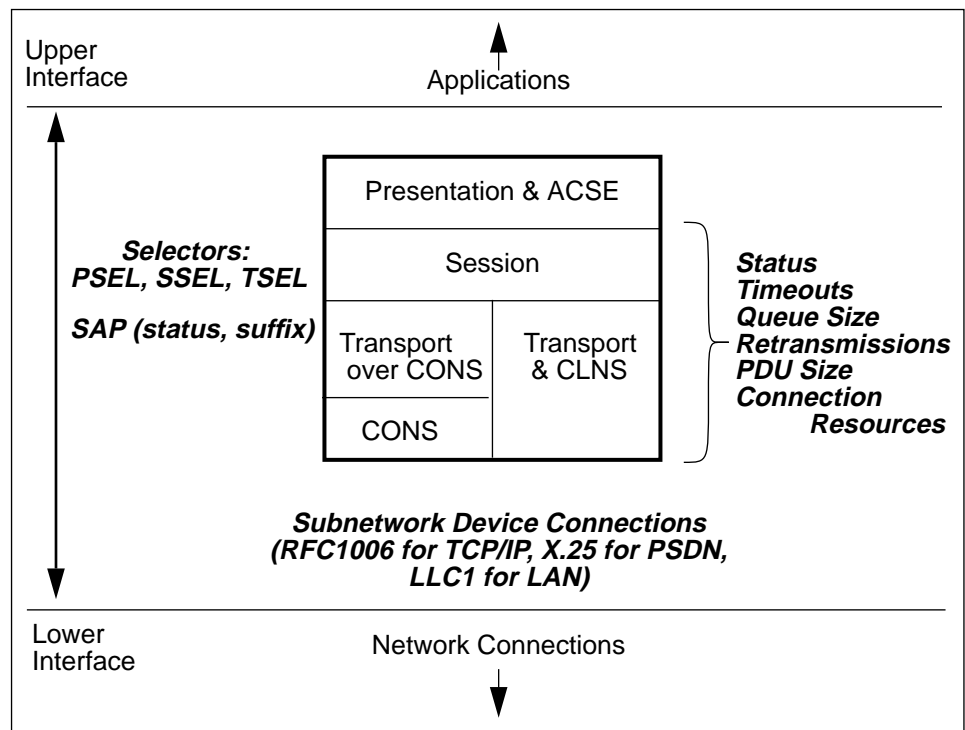


Figure 3-1 Stack Manager Configuration Parameter Summary

Most of the entity characteristics should not require modification, since the default values should be adequate. However, the options that might require customization for your network are:

- Transport class (for transport over CONS)
- Entity status
- Adding new subnetwork devices and their associated options

Each of these areas of stack configuration are explained in this chapter.

Stack Manager Main Window

When you select the Stack Manager from the `ositool` main window, a window similar to that shown in Figure 3-2 is displayed. From this window, you can set the stack entity characteristics for each entity, for example, timeouts, queue size, and packet size, depending on your network.

These parameters describe the fundamental setup of the stack. Default values are provided with the SunLink OSI 8.1 package that are designed to be valid for a typical network configuration. However, you still need to add specific information about your particular network during the configuration process.

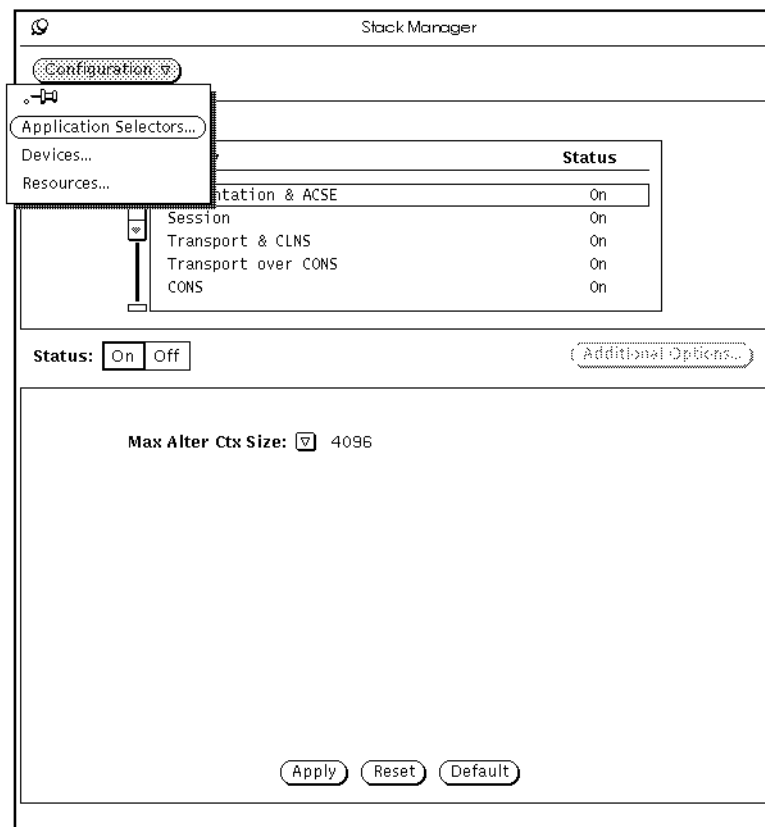


Figure 3-2 Stack Manager Window

To display the configurable parameters for Application Selectors, Devices and Resources, press MENU on the Configuration button. These configuration parameters are described later in this chapter.

Presentation & ACSE Entity

The Presentation and ACSE parameter configuration defines how this entity deals with application and user interfaces. Select Presentation & ACSE from the Entity list to display the configurable parameters, as shown in Figure 3-3.

The screenshot shows the 'Stack Manager' window with the 'Configuration' tab selected. A list of entities is displayed, with 'Presentation & ACSE' selected. Below the list, the 'Status' is set to 'On'. A 'Max Alter Ctx Size' field is set to '4096'. At the bottom, there are 'Apply', 'Reset', and 'Default' buttons.

Entity	Status
Presentation & ACSE	On
Session	On
Transport & CLNS	On
Transport over CONS	On
CONS	On

Status: ☒ On ☐ Off (Additional Options...)

Max Alter Ctx Size:

Figure 3-3 Presentation & ACSE Configuration

The parameters that you can configure are described below:

Parameter	Description
Status	Turns the selected Presentation & ACSE entity on or off. If you turn the entity status off, the Presentation & ACSE entity is isolated, and data no longer passes through this entity. A warning message is displayed for confirmation. You might want to do this for security or debugging purposes.
Max Alter Ctx Size	Sets the maximum size for PDUs accepted by the Presentation & ACSE entity for an Alter Context (AC) or an Alter Context Acknowledgement (ACA) PDU. This size only applies when context management has been negotiated for the connection. Press MENU to display the available settings and choose one of the following values: 0, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, 65536. The default value is 4096.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Session Entity

The session parameter configuration defines how the session entity deals with the data received from the transport layer. Select Session from the entity list to display the configurable parameters, as shown in Figure 3-4.

The screenshot shows the 'Stack Manager' application window. At the top, there is a 'Configuration' dropdown menu. Below it, a table lists the entities and their status. The 'Session' entity is selected and highlighted. Below the table, there is a 'Status' section with 'On' and 'Off' radio buttons, and an 'Additional Options...' button. The main configuration area contains several settings: 'Timeout (*10s)' set to 6, 'Max TSDU Queue' set to 4, 'Initiator Reuse Timeout (*10s)' set to 3, and 'Acceptor Reuse Timeout (*10s)' set to 3. To the right, under 'Protocol Options', there are four checkboxes: 'Accept reuse of TC', 'Propose reuse of TC', 'Include SSAP-ID in AC SPDU', and 'Use TRS Expedited Data'. At the bottom, there are three buttons: 'Apply', 'Reset', and 'Default'.

Entity	Status
Presentation & ACSE	On
Session	On
Transport & CLNS	On
Transport over CONS	On
CONS	On

Status: ☒ On ☐ Off Additional Options...

Timeout (*10s): 6 ▲▼
Max TSDU Queue: 4 ▲▼
Initiator Reuse Timeout (*10s): 3 ▲▼
Acceptor Reuse Timeout (*10s): 3 ▲▼

Protocol Options:
☐ Accept reuse of TC
☐ Propose reuse of TC
☐ Include SSAP-ID in AC SPDU
☐ Try to reuse TC
☐ Use TRS Expedited Data

Apply Reset Default

Figure 3-4 Session Configuration

The parameters that you can configure are described below:

Parameter	Description
Status	Turns the selected session entity on or off. If you turn the entity status off, the session entity is isolated, and data no longer passes through the session entity. A warning message is displayed for confirmation. You might want to do this for security or debugging purposes.
Timeout	Defines the time during which a session connection can wait for an accept, after which the connection aborts. Specify a value between 3 and 12 (*10s). The default value is 6 (equivalent to one minute).
Max TSDU Queue	This parameter is not implemented in this release.
Initiator Reuse Timeout	If the transport connection can be reused (according to the outcome of the Propose Reuse of TC protocol negotiation), then when the timeout between the termination of a session connection and the initiation of another session connection expires, the associated transport connection is terminated. This timeout specifies how long to maintain the transport connection open after the associated session connection is closed, if the connection was initiated locally. Specify a value between 3 and 120 (*10s). The default value is 3.
Acceptor Reuse Timeout	If the transport connection can be reused (according to the outcome of the Propose Reuse of TC protocol negotiation), then when the timeout between the termination of a session connection and the initiation of another session connection expires, the associated transport connection is terminated. This timeout specifies how long to maintain the transport connection open after the associated session connection is closed, if the connection was initiated remotely. Specify a value between 3 and 120 (*10s). The default value is 3.

Parameter	Description
Protocol Options	<p>Click SELECT on the checkboxes to switch the following protocol options on or off:</p> <p><i>Accept Reuse of TC</i> accepts the reuse of a transport connection when proposed by the remote system.</p> <p><i>Propose Reuse of TC</i> proposes to the peer session entity that the transport connection should be used again after the specified timeout. If the peer session entity accepts, then the transport connection can be reused. If you choose this option, make sure that the Accept Reuse of TC is also set on.</p> <p><i>Include SSAP-ID in AC SPDU</i> adds the session service access point identifier to the AC SPDU (Session layer Protocol Data Unit Accept PDU).</p> <p><i>Try to reuse TC</i> specifies that new connection requests can attempt to reuse a transport connection to the peer session entity. If you choose this option, make sure that the Propose Reuse of TC is also set on. The peer session entity must also allow reuse of transport connections. Ensure that the Accept Reuse of TC and Propose Reuse of TC are both selected if you choose this option.</p> <p><i>Use TRS Expedited Data</i> specifies that the Transport Expedited Data service should be used if it is available.</p>
Additional Options...	<p>Displays a list of additional parameters that do not normally require updating. Refer to “Additional Session Options” on page 29 for a description of these.</p>

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Additional Session Options

Click SELECT on the Additional Options button to display the extra configuration parameters that are available. The default settings are shown in Figure 3-5.

The screenshot shows a dialog box titled "Additional Options - Session". It is divided into two main sections: "Protocol Options" and "Debug FU Options".

Protocol Options:

- ☒ Version 1
- ☒ Version 2
- ☒ Extended Concatenation

Debug FU Options:

- ☒ Negotiated Release
- ☒ Half-Duplex
- ☒ Duplex
- ☒ Expedited Data
- ☒ Typed Data
- ☒ Capability Data
- ☒ Minor Synchronize
- ☒ Major Synchronize
- ☒ Resynchronize
- ☒ Exceptions
- ☒ Activity Management

At the bottom of the dialog box are three buttons: "Apply", "Reset", and "Default".

Figure 3-5 Additional Session Options

The parameters that you can configure are described below:

Parameter	Description
Protocol Options	Select the version of the session entity that is supported. <i>Version 1</i> specifies certain length restrictions on PDUs. <i>Version 2</i> specifies no length restriction for PDUs. <i>Extended Concatenation</i> allows the extended concatenation of PDUs and TPDUs. You must choose at least version 1 or version 2.
Debug FU Options	The functional unit options are for debugging purposes. The default is that all functional units are supported by the local session entity. These can be overridden by negotiated values. These options do not normally need to be updated.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Transport & CLNS Entity

The Transport & CLNS parameter configuration specifies the transport entity in the connectionless network for a LAN or X.25 network. Select Transport & CLNS from the entity list to display the configurable parameters as shown in Figure 3-6.

The screenshot shows the 'Stack Manager' window with the 'Configuration' tab selected. A table lists the entities and their status. The 'Transport & CLNS' entity is highlighted. Below the table, the 'Status' is set to 'On'. The 'Additional Options...' section contains several parameters: Max PDU Size (512), Credit Window (5), Retransmission Limit (5), Retransmission Timer (*1s) (40), Window Timer (*1s) (50), and Default CLNP Lifetime (50). There are also checkboxes for 'Propose Checksum', 'Use Extended Formats', and 'Disconnect on Protocol Error'. At the bottom are 'Apply', 'Reset', and 'Default' buttons.

Entity	Status
Presentation & ACSE	On
Session	On
Transport & CLNS	On
Transport over CONS	On
CONS	On

Status: ☒ On ☐ Off

Additional Options...

Max PDU Size: ☒ 512

Credit Window: 5

Retransmission Limit: 5

Retransmission Timer (*1s): 40

Window Timer (*1s): 50

Default CLNP Lifetime: 50

Protocol Options:

☐ Propose Checksum

☐ Use Extended Formats

☐ Disconnect on Protocol Error

Figure 3-6 Transport & CLNS Configuration

The maximum number of PDUs that can be queued is calculated from the Credit Window and the Allocation Quantum that you can set with `ositool`. The PDU size is a value that is negotiated between corresponding systems.

The actual value is calculated according to the following equation:

$$\frac{(\text{Credit Window} * \text{Allocation Quantum})}{\text{Negotiated Size of PDUs}} = \text{Credit Buffer for PDUs Received}$$

You can set the Credit Window on this screen, and the Allocation Quantum in the Additional Options for Transport & CLNS. Since the PDU size is a negotiated value, the maximum buffer space allowed for PDUs received is calculated using this equation, from the values you set. The PDU size is negotiated between the local system, where you can set a proposed size on this screen, and the maximum size allowed by the remote system.

Parameter	Description
Status	Turns the selected Transport entity on or off. If you turn the entity status off, the transport entity is isolated, and data no longer passes through this transport layer. A warning message is displayed for confirmation. You might want to do this for security or debugging purposes.
Max PDU Size	Specifies the maximum size for a transport PDU. The actual value is negotiated between the local and the remote system, the lower value being accepted. Press MENU to display the available settings and choose one of the following values: 128, 256, 512, 1024, 2148, 4096, or 8192. The default size is 512.
Credit Window	This value is used in the calculation to define the maximum buffer space allowed for PDUs that can be received before sending an acknowledgement. The equation is shown above, and the Allocation Quantum can be set from the Additional Options window described on page 33. Specify a value between 1 and 15. The default value is 5.
Retransmission Limit	Specifies the maximum number of TPDU retransmission attempts that can be made before aborting the connection. Specify a value between 2 and 20. The default value is 5.
Retransmission Timer	Specifies the initial time between retransmissions in the case of an error. The time between retransmissions doubles each time a retransmission attempt fails. Specify a value between 1 and 100. The default value is 40 seconds.

Parameter	Description
Window Timer	Specifies a timer that can be used to keep a connection open if the connection is inactive. An acknowledgement PDU is sent when this timer expires, if there is no activity. If this timer is less than the inactivity timer on the remote system, then the connection remains open. Specify a value between 1 and 100. The default value is 50 seconds.
Default CLNP Lifetime	Specifies the maximum lifetime that a CLNP PDU can exist in the network, after which it is discarded. Specify a value between 2 and 255. The default is 50 (*10s).
Protocol Options	<p>Click SELECT on the checkboxes to switch the following protocol options on or off:</p> <p><i>Propose Checksum</i> allows checksum to be negotiated if you are initiating a transport class 4 connection. The use of checksum is negotiated and will only be used if both the local and the remote systems select this option.</p> <p><i>Use Extended Formats</i> proposes the use of extended formats for PDUs to the remote system. This is negotiated and will only be used if both the local and the remote systems select this option.</p> <p><i>Disconnect on Protocol Error</i> when set on, sends a disconnection request TPDU when a protocol error has occurred. When set off, it sends an error TPDU or ignores the bad PDU.</p>
Additional Options...	Displays a list of additional parameters which do not normally require updating. Refer to "Additional Transport & CLNS Options" on page 33 for a description of these.

Note – The local inactivity timer is calculated by multiplying:
*Retransmission Limit * Window Timer * 2*

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Additional Transport & CLNS Options

Click SELECT on the Additional Options button to display the extra configuration parameters that are available. The default settings are shown in Figure 3-7.

Additional Options - Transport & CLNS

Allocation Quantum Size: 512

TIDU Size (512*n+1100): 0

Acknowledgment Timer (*1s): 10

Receiving Window: 5

Sending Queue Size: 10

Max Reassemble Size (CLNP): 5

Control Block Quantum (CLNP): 2

Protocol Options:

☐ Ignore Criteria for Ack

☒ Ignore Acknowledgment Timer

☐ Set Checksum on Spurious Response

☐ Ignore Checksum when Receiving

☐ Send Burst of TPDU as NUDTRQ

Apply Reset Default

Figure 3-7 Additional Transport & CLNS Options

These options and parameters are described below.

Parameter	Description
Allocation Quantum Size	This is used to control credit in terms of buffer space rather than in terms of PDUs since the PDU size is negotiated. It is used in conjunction with the Credit Window parameter. The equation used in the calculation is stated in "Transport & CLNS Entity" on page 30. Press MENU to display the available settings and choose one of the following sizes: 128, 256, 512, 1024, 2048, 4096, or 8192. The default value is 512 octets.

Parameter	Description
TIDU Size (512*n+1100)	<p>Since the size of Transport Service Data Units (TSDUs) is unlimited, they might need to be divided into Transport Interface Data Units (TIDUs), so that the session entity can deal with them. However, if these TIDUs are too small, the session control PDUs are also divided. This TIDU size specifies the minimum size allowed for a TIDU, where the minimum size is the specified value * 512 plus 1100. Even if you set a TIDU size of zero, the minimum size will be $0 * 512 + 1100 = 1100$.</p> <p>Note that this parameter is a default value that may be overridden by the user of the transport connection at connection time.</p> <p>Specify a value between 0 and 64. The default value is 0.</p>
Acknowledgement Timer	<p>Specifies the maximum delay allowed before sending a TPDU acknowledgement in response to a data TPDU. This can be used to minimize the number of TPDU acknowledgements sent, by sending one acknowledgement for several TPDU received.</p> <p>It should be used in conjunction with the Ignore Criteria for AK and Ignore AK Timer Protocol Options.</p> <p>Specify a value between 1 and 100 seconds. The default value is 10 seconds.</p>
Receiving Window	<p>Implements flow control for data passed from the transport layer to the user of the transport service. This value defines the maximum number of TSDUs that can be received and queued internally before invoking flow control.</p> <p>Specify a value between 1 and 20. The default value is 5 TSDUs.</p>
Sending Queue Size	<p>Specifies a limit on the number of data TPDU waiting to be transmitted or waiting to be retransmitted if not acknowledged.</p> <p>Specify a value between 1 and 20. The default value is 10 TPDU.</p>

Parameter	Description
Max Reassemble Size (CLNP)	<p>Specifies the maximum number of CLNP PDUs that can be queued for reassembly. Any additional PDUs are discarded.</p> <p>This is only used when the Full CLNP protocol is being used. (See “CLNP Options” on page 66 for information about setting the CLNP protocol options.)</p> <p>Specify a value between 1 and 100 (*10). The default value is 5 (*10).</p>
Control Block Quantum (CLNP)	<p>Specifies the maximum number of segments that can be handled when reassembling a Network Service Data Unit (NSDU).</p> <p>This is only used when the Full CLNP protocol is being used. (See “CLNP Options” on page 66 for information about setting the CLNP protocol options.)</p> <p>Specify a value between 1 and 100 (*10). The default value is 2 (*10).</p>
Protocol Options	<p>Click SELECT on the checkboxes to switch the following protocol options on or off:</p> <p><i>Ignore Criteria for Ack</i> when set on, this delays the sending of a TPDU acknowledgement for the amount of time specified by the Acknowledgement Timer. The default is off.</p> <p><i>Ignore Acknowledgement Timer</i> when set on, this ignores the Acknowledgement Timer and sends PDU acknowledgements immediately. The default is on.</p> <p><i>Set Checksum on Spurious Response</i> when set on, if an invalid TPDU is received (such as a data request sent in response to a connection request), then a checksum is included in the responding TPDU. The default is off.</p> <p><i>Ignore Checksum when Receiving</i> when set on, this ignores the checksum when receiving TPDU's, even if checksum was negotiated during the connection procedure. The default is off.</p> <p><i>Send Burst of TPDU as NUDTRQ</i> when set on, allows concatenation of TPDU's into a single NSDU. The default is off.</p>

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Transport over CONS Entity

The Transport over CONS parameter configuration specifies the transport entity in a connection-oriented network for a LAN or X.25 network. Select Transport over CONS from the entity list to display the configurable parameters as shown in Figure 3-8.

The screenshot shows the 'Stack Manager' window with the 'Configuration' tab selected. A list of entities is displayed, with 'Transport over CONS' highlighted. Below the list, the 'Status' is set to 'On'. The 'Additional Options...' section is expanded, showing various parameters and checkboxes.

Entity	Status
Presentation & ACSE	On
Session	On
Transport & CLNS	On
Transport over CONS	On
CONS	On

Status: ☒ On ☐ Off

Additional Options...

Max PDU Size:

Credit Window:

TS1/TS2 Timer (*10s):

TTR Timeout (*10s):

TWR Increment (*10s):

Retransmission Timer (*1s):

Retransmission Limit:

Class Options:

- ☐ Class 0 Only
- ☐ Propose Class 3
- ☐ Propose Class 4

Protocol Options:

- ☐ Include TSAP_ID in CC
- ☐ NULL PID if OSI

Buttons: Apply, Reset, Default

Figure 3-8 Transport over CONS Configuration

Parameter	Description
Status	<p>Turns the selected Transport entity on or off. If you turn the Transport over CONS status off, the entity is isolated, and data no longer passes through it. A warning message is displayed for confirmation.</p> <p>You might want to do this for security or debugging purposes.</p>
Max PDU Size	<p>Specifies the maximum size a transport PDU is allowed to be on this system. The actual value is negotiated between the local and the remote systems, the lower value being accepted.</p> <p>Press MENU to display the available settings and choose one of the following values: 128, 256, 512, 1024, 2148, 4096, or 8192. The default size is 512.</p>
Credit Window	<p>This value is used in the calculation to define the maximum buffer space allowed for PDUs that can be received before sending an acknowledgement. The equation is stated in "Transport & CLNS Entity" on page 30, and the Allocation Quantum can be set from the Additional Options window described on page 40.</p> <p>Specify a value between 1 and 15. The default value is 5.</p>
TS1/TS2 Timer (*10s)	<p>Specifies the TS1/TS2 timers (as defined in ISO 8073) for disconnection of an inactive network connection.</p> <p>Specify a value between 3 and 12. The default value is 6.</p>
TTR Timeout	<p>Specifies the time between trying to resynchronize or reassign the transport connection after a transmission failure.</p> <p>It is only used when the local system initiated a class 3 protocol connection.</p> <p>Specify a value between 3 and 12 (*10s). The default value is 6 (*10s).</p>

Parameter	Description
TWR Increment	<p>This is used to specify the time to wait for resynchronization or reassignment of the transport connection after a transmission failure.</p> <p>It is only used when the remote system initiated a class 3 connection. The stated value is the difference between the TTR timer (plus the maximum disconnection wait and the maximum transmission time) and the TWR timer.</p> <p>Specify a values between 1 and 11 (*10s). The default value is 2 (*10s).</p>
Retransmission Timer (*1s)	<p>Specifies the initial time between retransmissions in the case of an error for transport class 4 connections. The time between retransmissions doubles each time a retransmission takes place.</p> <p>Specify a value between 30 and 120 seconds. The default value is 30 seconds.</p>
Retransmission Limit	<p>Specifies the maximum number of TPDU retransmission attempts that can be made before aborting the connection.</p> <p>Specify a value between 2 and 20. The default value is 3.</p>
Class Options	<p>Click SELECT on the checkboxes to determine the transport class to be used:</p> <p><i>Class 0 Only</i> specifies that connections always use class 0 protocol which provides basic transport connection support that can detect errors but not recover from them.</p> <p><i>Propose Class 3</i> means that class 3 is proposed as the preferred class and if the remote system accepts class 3, the transport connection is established according to class 3 transport protocol specifications. This includes recovery from network disconnections.</p> <p><i>Propose Class 4</i> means that class 4 is proposed as the preferred class and if the remote system accepts class 4, the transport connection is established according to class 4 transport protocol specifications. This includes flow control and full error control.</p>

Parameter	Description
Class Options (continued)	If you specify both class 3 and 4, then the class chosen depends on the remote system. When a remote system initiates the connection, the local system accepts the higher class proposed, that is, class 4 is preferred.
	If you do not specify any class options, then class 2 is selected.
Protocol Options	Click SELECT on the checkboxes to switch the following protocol options on or off: <i>Include TSAP_ID in CC</i> includes the transport service access point identifier in PDUs for operability with systems that require it. <i>NULL PID if OSI</i> should be set on if you do not want to include the protocol identifier in a connection request PDU.
Additional Options...	Displays a list of additional parameters which do not normally require updating. Refer to “Additional Transport over CONS Options” on page 40 for a description of these.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Additional Transport over CONS Options

Click SELECT on the Additional Options button to display the extra configuration parameters that are available. The default settings are shown in Figure 3-9.

Figure 3-9 Additional Transport over CONS Options

These options do not normally need to be changed:

Parameter	Description
Allocation Quantum Size	This is used to control credit in terms of buffer space rather than in terms of PDUs since the PDU size is negotiated. It is used in conjunction with the Credit Window parameter. The equation used in the calculation is stated in "Transport & CLNS Entity" on page 30. Press MENU to display the available settings and choose one of the following sizes: 128, 256, 512, 1024, 2048, 4096, or 8192. The default value is 512 octets.
Max Multiplexing (initiator)	Specifies the maximum multiplexing for locally initiated transport connections onto a single network connection. Specify a value between 1 and 126. The default value is 5 connections.

Parameter	Description
Max Multiplexing (acceptor)	<p>Specifies the maximum multiplexing for remotely initiated transport connections onto a single network connection. It should be a greater value than the maximum multiplexing (initiator) parameter.</p> <p>Specify a value between 1 and 126. The default value is 10 connections.</p>
TIDU Size (512*n+1100)	<p>Since the size of Transport Service Data Units (TSDUs) is unlimited, they might need to be divided into Transport Interface Data Units (TIDUs), so that the session entity can deal with them. However, if these TIDUs are too small the session control PDUs are also divided. This TIDU size specifies the minimum size allowed for a TIDU, where the minimum size is the specified value * 512 plus 1100. Even if you set a TIDU size of zero, the minimum size will be $0 * 512 + 1100 = 1100$.</p> <p>Note that this parameter is a default value that may be overridden by the user of the transport connection at connection time.</p> <p>Specify a value between 0 and 61. The default value is 0.</p>
Long NC Timeout (*10s)	<p>Specifies an inactivity timer for a network connection when:</p> <ul style="list-style-type: none">—the local system is not the initiator of the network connection, and—multiplexing is being used for the network connection. <p>Specify a value between 3 and 48 (*10s). The default value is 24 (*10s).</p>
Short NC Timeout (*10s)	<p>Specifies an inactivity timer for a network connection when:</p> <ul style="list-style-type: none">—the local system initiated the network connection, and—multiplexing is not being used for the network connection. <p>Specify a value between 3 and 6 (*10s). The default value is 3 (*10s).</p>
QOS Threshold/nfc	<p>Sets a threshold for the use of flow control in the class 2 transport protocol. The default value of zero suppresses flow control. A value of 8 switches flow control on.</p>
QOS Threshold/mpx	<p>Sets a threshold for the use of multiplexing in class 2 transport protocol. A value of zero suppresses multiplexing. The default value is 1.</p>

Parameter	Description
Max Size/nofc	For class 2 transport protocol, specifies a maximum limit for the receiving queue size where no flow control has been negotiated at this level, since flow control can be exercised at the TIDU level. When this limit is exceeded, an error PDU is generated. Specify a value between 1 and 100 (*1024 bytes). The default value is 4.
Protocol Options	Click SELECT on the checkboxes to switch the following protocol options on or off: <i>Disconnect Unused NC</i> when set on, disconnects an inactive network connection to allow it to be reassigned if a new connection is required. The default is off. <i>Support Network Reset</i> set on to configure the local system to send an N-RESET when an error occurs and to take the appropriate action when an N-RESET is received. If this is set on and the DR or DR/ER options are set off, then an N-RESET is sent when an error occurs. If this option and the DR and DR/ER error options are all set off, a disconnect request is sent. <i>Suppress Expedited</i> when set on prevents expedited data from being proposed or accepted.
Error Options	Click SELECT on the checkboxes to switch the following protocol options on or off: <i>Use DR in Case of Error</i> when set on, sends a disconnect request TPDU when an error occurs, according to the type of error. <i>Use DR or ER in Case of Error</i> when set on sends a disconnect request or error TPDU when an error occurs. These error options are not normally required. However, some remote systems may require the use of a disconnect request or error PDU.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

CONS Entity

The CONS parameter configuration specifies the CONS entity in a connection-oriented network for an X.25 network. Select CONS from the entity list to display the configurable parameters as shown in Figure 3-10.

The screenshot shows the 'Stack Manager' window with the 'Configuration' tab selected. A table lists the entities: Presentation & ACSE, Session, Transport & CLNS, Transport over CONS, and CONS. The 'CONS' entity is selected and highlighted. Below the table, the 'Status' is set to 'On'. A button labeled 'Additional Options...' is visible. The main configuration area contains several parameters with up/down arrows for adjustment:

Entity	Status
Presentation & ACSE	On
Session	On
Transport & CLNS	On
Transport over CONS	On
CONS	On

Status: ☒ On ☐ Off Additional Options...

Default X.25 Packet Size: 128

Max NSDU Length (*256): 8

Connection Timer (*10s): 30

Connection Timer (X25/80) (*10s): 25

Deconnection Timer (X25/80) (*10s): 20

Grcb Pool Size (*10): 3

Rccb Pool Size (*10): 3

Recb Pool Size (*10): 3

Apply Reset Default

Figure 3-10 CONS Configuration

The parameters and their possible values are described below.

Parameter	Description
Status	Turns the selected network entity on or off. If you turn the entity status off, the CONS entity is isolated, and data no longer passes through it. A warning message is displayed for confirmation. You might want to do this for debugging or security purposes.
Default X.25 Packet Size	Specifies the default packet size for an X.25 network connection. Press MENU to display the available settings and choose one of the following sizes: 16, 32, 64, 128, 256, 512, 1024, 2048, or 4096. The default value is 128 octets.
Max NSDU Length (*256)	Specifies the maximum length allowed for a Network Service Data Unit (NSDU). This value should be greater than or equal to the maximum PDU size specified for transport. Specify a value between 1 and 40. The default value is 8 (*256 octets).
Connection Timer (*10s)	Specifies the maximum time allowed for establishing a connection. If the connection is not established within the specified time, the connection attempt is aborted. This timer is used for all versions of X.25 networks. Specify a value between 0 and 60 (*10s). The default value is 30 (*10s), equivalent to a length of 5 minutes.
Connection Timer (X25/80) (*10s)	Specifies the maximum time allowed for establishing a connection in an X.25/80 (SNDP) network. It should be less than or equal to the connection timer defined above. If the connection is not established within the specified time, the connection attempt is aborted. Specify a value between 0 and 30 (*10s). The default value is 25 (*10s).
Deconnection Timer (X25/80) (*10s)	Specifies the time for completion for disconnecting a network connection, in an X.25/80 (SNDP) network. If the disconnection is not complete within the specified time, the disconnection is aborted. Specify a value between 1 and 30 (*10s). The default value is 20 (*10s).

Parameter	Description
Grcb Pool Size (*10)	This is used to define the size of the routing table for the network. It specifies the maximum number of entries in the general routing information table. Specify a value between 0 and 255 (*10). The default value is 3 (*10). You should not need to change this value, unless a large number of routes is required.
Rccb Pool Size (*10)	This is used to define the routing table for the network. It specifies the size of the cache to be used for routes that are currently or frequently used. Specify a value between 0 and 255 (*10). The default value is 3 (*10). You should not need to change this value, unless a large number of routes is required.
Recb Pool Size (*10)	This is used to define the routing table for the network. It specifies the maximum number of routes for each entry (described in Grcb Pool Size above) in the general routing information table. Specify a value between 0 and 255 (*10). The default value is 3 (*10). You should not need to change this value, unless a large number of routes is required.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Configuration Menu

Press MENU on the configuration button to display the pull-down menu shown in Figure 3-11.




Figure 3-11 Configuration Menu

The following configuration options are available:

Option	Description
Application Selectors	This window is used to specify the access points for applications that you want to install, for example, FTAM, CMIP, and X.400. The configuration of the application selectors is explained in “Application Selectors Window” on page 47.
Devices	This allows you to specify the types of devices used for connection through the lower layers to LAN, X.25, and TCP/IP networks. The configuration of the device connections is explained in “Device Configuration” on page 51.
Resources	This allows you to define internal connection resources for each entity in the stack. The configuration of connection resources is explained in “Resource Configuration” on page 56.

Application Selectors Window

The Application Selectors window is used to configure the access points for applications used over SunLink OSI 8.1. The screen shown in Figure 3-12 is displayed when you select Application Selectors from the Configuration pull-down menu.

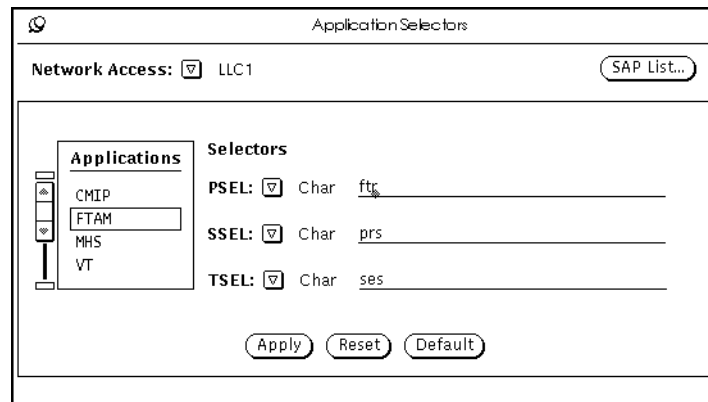


Figure 3-12 Application Selectors Window

Some of the selectors for different applications are used by more than one application. You should be careful that, if you want to configure for more than one application, modifications to the selectors for one are compatible with the other applications. Enter the following information for each application that you want to run over SunLink OSI 8.1:

Parameter	Description
Network Access	Press MENU to display the available settings and choose the type of network device used for the connection: LLC1 to access LANs; X.25 to access a PSDN; or RFC1006 to access TCP/IP network.
Applications	Lists the applications available from SunConnect. Use the scroll bar to move up and down the list, and choose an application.
PSEL	Press MENU to choose a hexadecimal or character format for the access point. Enter the presentation selector for the selected application.

Parameter	Description
SSEL	Press MENU to choose a hexadecimal or character format for the access point. Enter the session selector for the selected application.
TSEL	Press MENU to choose a hexadecimal or character format for the access point. Enter the transport selector for the selected application.
SAP List...	Click SELECT on the SAP List button to display the defined Service Access Points. Further information about the SAP List window is given in “SAP List” below.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

SAP List

The Service Access Point (SAP) is a unique identification of a link between two entities. The characteristics of each SAP is described by its associated parameters. Figure 3-13 shows an example of the SAP List window.

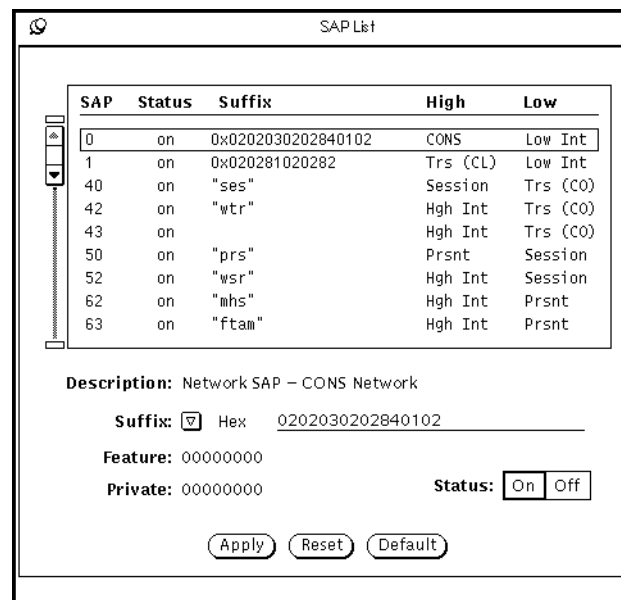


Figure 3-13 SAP List Window

The SAP list indicates the following information:

- The Service Access Point (SAP) for an entity. This is a fixed internal address that cannot be changed.
- Status indicates whether the entity is switched on or off.
- Suffix is the external address reference for the SAP. It identifies a particular SAP address.
- High indicates the entity above the SAP.
- Low indicates the entity below the SAP.

The Description field cannot be changed. It gives a one-line description of the SAP.

The Feature and Private fields are also read-only, and are mainly used for debugging purposes.

Parameter	Description
Suffix	Press MENU to choose a hexadecimal or character format for the SAP suffix. Enter the suffix that references the SAP.
Status	Choose the appropriate box to turn the specified SAP connection on or off. Changes to the SAP here are reflected in the SAP list.

Press the Apply button to save the configuration changes. The changes are implemented immediately. Otherwise, press Reset to return to the last saved settings, or Default to redisplay the default settings.

Device Configuration

This window is used to specify which types of devices are used for connections to a subnetwork. Select Devices from the Configuration pull-down menu to display a screen similar to that shown in Figure 3-14.

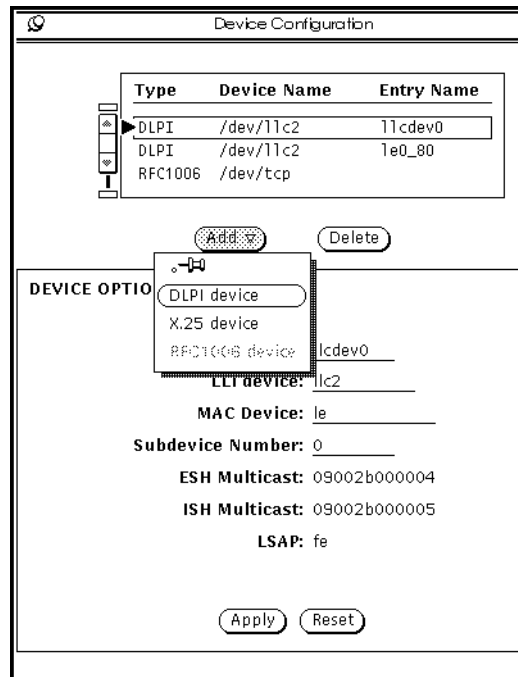


Figure 3-14 Device Configuration Window

The device list displays the subnetwork devices that are configured and you can do one of the following:

- Use the scroll bar to move up and down the list and choose a device. Its current settings are displayed.
- Press MENU on Add and choose the type of device that you want to add to the list from the pull-down menu. You can add up to 64 LAN devices (DLPI), one X.25 device and one TCP/IP device (RFC1006). The parameters that need to be specified are displayed.

- Choose a device and press the Delete button to remove it from the list. The deletion is implemented when you save the configuration and restart the stack.

LAN Device

If you highlight a LAN device in the list, a screen similar to the one shown in Figure 3-15 is displayed.

Type	Device Name	Entry Name
DLPI	/dev/l1c2	11cdev0
DLPI	/dev/l1c2	1e0_80
RFC1006	/dev/tcp	

Add Delete

DEVICE OPTIONS

Name: llcdev0
 LLI device: llc2
 MAC Device: le
 Subdevice Number: 0
 ESH Multicast: 09002b000004
 ISH Multicast: 09002b000005
 LSAP: fe

Apply Reset

Figure 3-15 LAN Subnetwork Device

These parameters are described below:

Parameter	Description
Name	Enter a name of up to 8 characters for each LAN device. This name is used to identify the subnetwork device.

Parameter	Description
LLI Device	Specifies the DPLI device which implements link level 1. For SunConnect products, this value is <code>llc2</code> .
MAC Device	Specifies the Media Access Control (MAC) name. This is the name of the device that implements the type of network. The default is for an Ethernet device and is <code>le</code> . Change this to <code>tr</code> for token ring, <code>bf</code> for FDDI, or <code>ie</code> for the implementation of Internet Ethernet link.
Subdevice Number	This specifies the subnetwork device number which identifies the physical interface. The default is 0.
ESH Multicast	Displays the End-System-Hello (ESH) multicast address defined by the LAN device. You cannot change this field.
ISH Multicast	Displays the Intermediate-System-Hello (ISH) multicast address defined by the LAN device. You cannot change this field.
LSAP	Displays the Link Service Access Point (LSAP). The default value is <code>fe</code> . You might need to change this if your system is not OSI conformant.

X.25 Subnetwork Device

When you select an X.25 subnetwork device, a screen similar to that shown in Figure 3-16 is shown.

These parameters are explained below.

Parameter	Description
Link Number	Specifies the link number. This value should be the same as the link number specified in your SunLink X.25 configuration. Specify a value between 0 and 254 to specify a particular link, or 255 to specify that X.25 selects the link according to the defined NSAP (this is described in the X.25 documentation). The default value is 0.
Connection Pool	Specifies the maximum number of X.25 connections allowed. Specify a value between 1 and 128. The default value is 3.

Parameter	Description
SNPA Address	Specifies the Subnetwork Point of Attachment address. This is the address that uniquely identifies the X.25 device on the subnetwork, and is described in the SunLink X.25 configuration tool. Enter an X.121 address of up to 15 digits.

Device Configuration

Type	Device Name	Entry Name
DLPI	/dev/l1c2	1e0_80
RFC1006	/dev/tcp	
X.25	/dev/x25	

Add
Delete

Link Number: 0
Connection Pool: 3
SNPA Address:

Apply
Reset

Modified

Figure 3-16 X.25 Subnetwork Device

RFC1006 Subnetwork Device

If you select an RFC1006 device from the list, a screen similar to the one shown in Figure 3-17 is displayed.

Device Configuration

Type	Device Name	Entry Name
DLPI	/dev/l1c2	l1cdev0
DLPI	/dev/l1c2	le0_80
RFC1006	/dev/tcp	

Add Delete

Connection Pool: 4

Apply Reset

Figure 3-17 RFC1006 Subnetwork Device

You can enter the following information:

Parameter	Description
Connection Pool	Specifies the maximum number of connections allowed per TCP/IP subnetwork. Specify a value between 1 and 128. The default value is 3.

Press Apply to save the changes or Reset to return to the last saved settings. The changes are implemented when you save the configuration and restart the stack.

Resource Configuration

The resource configuration window is used to adjust resources within the stack for connections. These parameters probably do not need to be modified unless you have applications that require a large number of simultaneous connections. The default values should be adequate for most systems.

Select Resources from the Configuration pull-down menu to display the options shown in Figure 3-18.

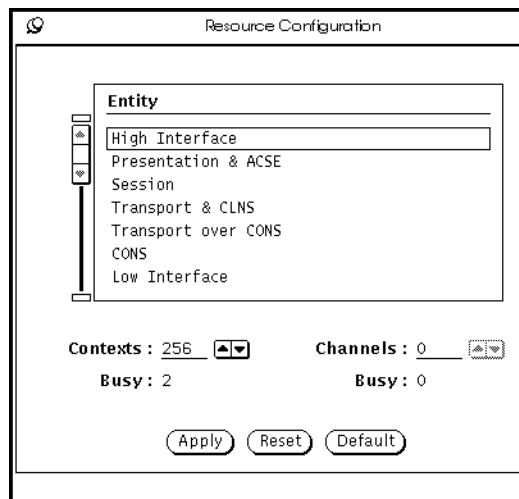


Figure 3-18 Resource Configuration Window

There are two related resource parameters that you can configure for each entity:

- Channels define a pool of connections between entities. You specify the maximum number of channels that might be required between two entities.
- Contexts define control information specifying the channel and the state of the connection. A channel is described by the context in the entity below it.

The total number of contexts and channels is related to the maximum number of simultaneous connections allowed. Some entities require more contexts than channels. For example the possibility for multiplexing in the transport entity requires twice as many contexts as channels. The number of contexts should be equal to or greater than the number of channels.

Context and channel resources are defined per entity and are allocated an appropriate amount of memory. The total number of channels and contexts that you specify here should allow for the maximum number of simultaneous connections required and the resources required for these. Memory is allocated for the total resources that you configure, therefore you should be careful not to overestimate the resources required.

Entity	Description
High Interface	<p>The high entity configuration affects connections to applications. You do not need to set any channels for this, but a context is required for each TLI and each APLI connection. You should add a minimum of two contexts for administrative purposes (such as <code>ositol</code> and console functions).</p> <p>In addition to this, configure the following contexts for applications:</p> <ul style="list-style-type: none"> —30 contexts for SunLink CMIP —3 contexts for SunLink FTAM (one each for LAN, X.25 and TCP/IP subnetworks) —3 contexts for SunLink VT (one each for LAN, X.25 and TCP/IP subnetworks) —3 contexts for SunLink MHS (one each for LAN, X.25 and TCP/IP subnetworks) <p>The default value in total is 256 contexts and 0 channels.</p>
Presentation & ACSE	Configure one channel and one context per connection. The default values for both contexts and channels is 128.
Session	Configure one channel and one context per connection. The default values for both contexts and channels is 128.
Transport & CLNS	<p>Due to the possibility of multiplexing in the transport layer, you require up to twice as many contexts as channels.</p> <p>The default values are 256 contexts and 128 channels.</p>

Entity	Description
Transport over CONS	Due to the possibility of multiplexing in the transport layer, you require up to twice as many contexts as channels. The default values are 256 contexts and 128 channels.
CONS	Configure one channel and one context per connection. The default value for both contexts and channels is 128.
Low Interface	This figure is dependant on the total simultaneous connections required for the pool. Each connection requires one channel and one context.

These entities are defined with the following parameters:

Parameter	Description
Context	Specify a value equal to or greater than the number of channels required. The default value depends on the selected entity.
Channel	Specify a value equal to or less than the number of contexts required. The default value depends on the selected entity.
Busy	Displays the channels and contexts that were in use when you entered this screen.

Note – When you are adding a device, the number of contexts are updated by `ositool` automatically. If the device is removed later, the value is not updated.

Press Apply to save, Reset to return to the last saved settings or Default to display the default settings. The changes saved are not implemented until you saved the whole configuration and restart the `osinetd` daemon.

Chapter Summary

This chapter described how to update the configuration parameters in the Stack Manager. You need to update some of the parameters; for example, to change some protocol elements such as timers for LAN, X.25, and TCP/IP network connections.

Now turn to Chapter 4, “End System to Intermediate System Configuration” if you need to update the configuration parameters defining end systems and intermediate systems and routing across subnetworks.

End System to Intermediate System Configuration

4



<i>Subnet List</i>	<i>page 63</i>
<i>Chapter Summary</i>	<i>page 71</i>

This chapter explains how to configure the End System to Intermediate System protocol (ES-IS), set the size of the routing table, and specify the CLNP protocol type.

For more information about how dynamic routing is defined by the ES-IS protocol, refer to “Dynamic Routing” on page 117.

Select the ES-IS Configuration Manager from the `ositool` main menu to display the screen shown in Figure 4-1.

Note – The multicast address used for SunLink OSI 8.1 ES-IS is not valid for token ring networks. For this reason, you cannot use the ES-IS protocol to route dynamically over token ring networks. You must either configure all token ring routes manually, using route manager, or set the default subnetwork for the token ring connection. Refer to “Using Token Ring Networks” on page 122 for a complete explanation.

ES-IS Configuration

SUBNETS

No	Type	Entry Name	Subset
1	X25	x25	Full Protocol
2	LLC	llcdev0	Null Protocol
3	LLC	llcdev0	Full Protocol

Add
Delete

Default: 2
Status: On Off

SUBNET ID

Type: X25
Entry Name: x25

ROUTE TABLE SIZE OPTIONS

Static Entries: 50
ESH Entries: 50
SII Entries: 50
ISH Entries: 50

CLNP OPTIONS

Protocol Subset: Full Protocol

☐ Use Checksum
☐ Use Error Reporting

ES-IS OPTIONS

Send ESH Timer: 30
Send ISH Timer: 30
Holding Timer: 30

☐ Send Redirect
☐ Send to all ES
☐ Send ES Hello
☐ Send IS Hello
☒ Send to default IS

☐ Record IS Hello
☐ Record ES Hello
☒ Process Redirect
☐ Refresh Cache Entry

Apply
Reset

Figure 4-1 ES-IS Configuration

Subnet List

The scrollable list displayed in Figure 4-2 summarizes the subnetwork devices that you configured in Device Configuration (see “Device Configuration” on page 51).

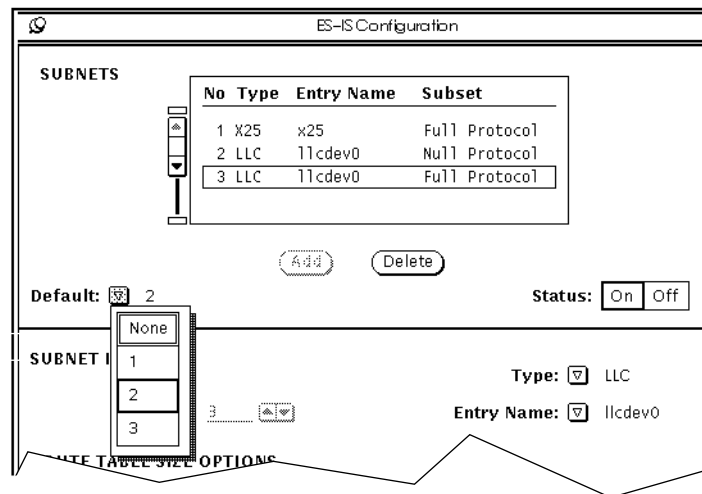


Figure 4-2 Subnet List

It displays the following information:

- Subnet Number—shows the subnetwork number that is associated with each subnetwork device. This is generated automatically by the stack.
- Type—shows the type of subnetwork connection configured for this device.
- Entry Name—shows the name associated with the subnetwork device. This is specified in “Device Configuration” on page 51.
- Subset—shows whether the full, non-segmenting, or null (inactive) protocol is being used. This is defined in this window.

Press MENU on Default to select a subnetwork which will be chosen as the default if no route is specified for an NSAP. This is also known as a “black hole.” Since null CLNP performs no routing function, you may need to set the default to a direct null CLNP route. See “Using Null CLNP” on page 114 for more details.

From this window you can also switch the status of a subnetwork connection on or off. Remember that changes made to the ES-IS configuration are static, so you have to save the configuration and restart the daemon to implement the changes. Therefore, if you have added a subnetwork device, you cannot switch its status on or off until you have saved the whole configuration and restarted the daemon.

Do one of the following:

- Select one of the displayed subnetwork devices to display its current settings. You can edit these settings if necessary, and press Apply to implement them. These options are described in the following sections.
- Click SELECT on Add to add another subnetwork entry to the list. You can have up to two subnetworks per device entry specified; one using full or non-segmenting protocol, and one using null protocol. Then add the relevant information regarding the ES-IS protocol and press Apply to implement the new settings. The options available are described in the following sections.
- Highlight one of the subnetwork devices in the list and then click SELECT on Delete to remove a subnetwork device from the list.

The information that you can set to specify the characteristics of an ES-IS route are described in the following sections.

Subnet ID

The Subnet ID information identifies the characteristics relating to a particular device subnetwork connection. Figure 4-3 shows the options that you can change.

The image shows a configuration window titled "SUBNET ID". Inside the window, there are three main fields: "Number" with the value "2", "Type" with a dropdown menu showing "LLC", and "Entry Name" with a dropdown menu showing "llcdev0". Below these fields is a section labeled "ROUTE TABLE SIZE OPTIONS". The window has a decorative border with a wavy top and bottom edge.

Figure 4-3 Subnet Identifiers

These are described below:

Parameter	Description
Number	Displays the automatically generated number that is associated with every subnetwork entry. You cannot change this value.
Type	Press MENU to display the types of subnetwork connections that can be configured for this subnetwork entry. Choose LLC or X25.
Entry Name	Press MENU to display the names of subnetwork devices (defined in "Device Configuration" on page 51). If the type of subnetwork is X.25, then the entry name is also X.25. If the type of subnetwork is LLC, then the pull-down menu displays all those subnetwork devices with less than two connections configured. A subnetwork device can have up to two connections, one using full or non-segmenting protocol, and the other using null protocol.

Route Table Size Options

The Route Table Size information specifies the size of the routing table that is used for the ES-IS dynamic and static routing functions. It is illustrated in Figure 4-4.

```

ROUTE TABLE SIZE OPTIONS

Static Entries: 50 [▲▼]   ESH Entries: 50 [▲▼]
SII Entries: 50 [▲▼]    ISH Entries: 50 [▲▼]

CLNP OPTIONS
  
```

Figure 4-4 Route Table Size Options

Note – Modifications to buffer space for routing tables are only implemented when you save the configuration and restart the `osinetd` daemon.

Enter the following information, if necessary..

Parameter	Description
Static entries	Specifies the size of the routing table for static routes for the selected subnetwork. Specify a value to set how many static entries are required
SII entries	Specifies the size of the routing table for static IS entries. You might need to change this if you have many intermediate systems.
ESH entries	Specifies the size of routing table for recording ESHs on the selected subnetwork. Specify a value to set how many ESH entries are required.
ISH entries	Specifies the size of the routing table for recording ISHs on the selected subnetwork. Specify a value to set how many ISH entries are required.

The number of default entries for the routing tables depend on the type of route you are configuring.

CLNP Options

This section explains how to specify particular CLNP protocol characteristics, as shown in Figure 4-5.

CLNP OPTIONS

Protocol Subset: ☒ Full Protocol

☐ Use Checksum

☐ Use Error Reporting

Figure 4-5 CLNP Protocol Options

You can specify the following options.

Parameter	Description
Protocol Subset	<p>Press MENU and choose one of the following options from the pull-down menu:</p> <p><i>Full Protocol</i> to specify that your local system should run the full CLNP protocol according to ISO 8473.</p> <p><i>Null Protocol</i> can be used if the source and destination end systems are connected by a single subnetwork. The protocol options are not used. If you want to use null CLNP over X.25, you must specify the X.121 address, or over a LAN you must use the MAC address, to identify the remote system over a specific subnetwork. See “Using Null CLNP” on page 114.</p> <p><i>Non-Segmenting Protocol</i> is a subset of the protocol that permits the simplification of data transmission and reception by suppressing segmentation. NSDUs must fit into a single data PDU.</p>
Protocol Options	<p>Click SELECT on the checkboxes to choose the following protocol options:</p> <p><i>Use Checksum</i> causes the checksum parameter to be used for error control in CLNP PDUs.</p> <p><i>Use Error Reporting</i> causes error reporting to be used for CLNP PDUs.</p>

Note – Using CLNP Null Protocol

If you are using the null protocol for CLNP, the NSAP address is not known by the remote system. You therefore need to specify the X.121 address for X.25 or the MAC address for a LAN, instead of its NSAP. Otherwise, during connection, the connect confirm acknowledgement is not received by the remote system. Use a default direct route to specify the subnet for a null CLNP connection, or define a new direct route using the Route Manager.

ES-IS Options

This section explains how to configure the timers for the ES-IS protocol and the specific protocol options, as shown in Figure 4-6.

Figure 4-6 ES-IS Options

These options are described below.

Option	Description
Send ESH Timer	Determines how often the local ES sends an End System Hello (ESH). Specify a value between 1 and 1000 (seconds). The default value is 30 for both LAN and X.25 subnetworks.
Send ISH Timer	Determines how often the local IS sends an Intermediate System Hello (ISH). Specify a value between 1 and 1000 (seconds). The default value is 30 for a LAN and 50 for an X.25 subnetwork.

Option	Description
Holding Timer	Determines the maximum time for the receiving network entity to retain the routing information in your local ES-IS PDUs. Specify a value between 1 and 1000 (seconds). The default value is 30 for both LAN and X.25 subnetworks.

Click SELECT on the checkboxes to set the following protocol options.

Option	Description
Send Redirect	If your local system does not recognize the destination of the PDU, the packet is forwarded with a redirect PDU to the originating system.
Send to all ES	If there are no ISs on a subnetwork and your local system does not know an immediate destination for an outgoing packet, a query configuration PDU is broadcast to all ES multicast addresses.
Send ES Hello	Tells your local system to send periodical End System Hellos (ESH) on the local subnetwork.
Send IS Hello	Tells your local system to send periodical Intermediate System Hellos (ISH) on the local subnetwork.
Send to Default IS	Specifies that if your local system does not know the SNPA address to the corresponding NSAP address of the destination host, it sends the PDU to the first available IS on the local network.
Record IS Hello	Tells your local system to listen for ISHs on the all-ES multicast addresses.
Record ES Hello	Tells your local system to listen for ESHs on the all-IS multicast addresses.
Process Redirect	Tells your local system to listen for redirect PDUs and update the routing table in response to a redirect PDU.
Refresh Cache Entry	Sets the local system to update the route in the cache routing table for all routing decisions and incoming and outgoing PDUs. The performance of the network is affected by this change.

The default values for each type of subnetwork are shown in Table 4-1 and Table 4-2.

Table 4-1 Default Values for CLNP/LLC1 Subnetworks

Type	ES	IS	ES+IS
Send to Default IS	✓		
Send to All ES	✓		
Record ES Hello		✓	✓
Record IS Hello	✓		✓
Send ES Hello	✓		✓
Send IS Hello		✓	✓
Process Received Redirect			
Send Redirect		✓	✓

Table 4-2 Default Values for CLNP/X.25 Subnetworks

Type	ES	IS	ES+IS
Send to Default IS	✓		✓
Send to All ES			
Record ES Hello			
Record IS Hello			
Send ES Hello			
Send IS Hello			
Process Received Redirect	✓		
Send Redirect		✓	✓

Press Apply to save, or Reset to return to the last saved settings. The changes are not implemented until you save the whole configuration and restart the `osinetd` daemon.

Chapter Summary

This chapter explained how to set the configurable parameters for the ES-IS protocol. You can set timers for ESHs and ISHs, configure buffer space for the static routing tables and specify which protocol subsets should be used. These options need to be set for both LAN and X.25 connectionless subnetworks. For more information about how the ES-IS protocol performs dynamic routing, refer to “Dynamic Routing” on page 117.

Now go to Chapter 5, “Network Layer Addresses” to specify details of network addresses.



<i>NSAP Family Window</i>	<i>page 74</i>
<i>Chapter Summary</i>	<i>page 87</i>

This chapter describes how to configure network addresses (NSAPs and NETs), which define the point where a network service is made available to the transport layer.

You can configure a network address in seven different formats, depending on your network. Further information about the format of network addresses is given in “Addressing” on page 103.

Changes to the Network Service Access Points (NSAPs) are dynamic changes that are implemented as soon as you press **Apply**. Modifications to the Network Entity Title (NET) are static changes and are only implemented when you save the configuration and restart the `osinetd` daemon.

Note – Using CLNP Null Protocol

If you are using the null protocol for CLNP, the NSAP address is not known by the remote system. See “Using Null CLNP” on page 114 for details.

NSAP Family Window

Choose Network Layer Addresses from the `ositol` main menu and press MENU on the NSAP Family menu button. The screen shown in Figure 5-1 is displayed. Choose the type of format that you want to use to specify the NSAP or NET. Each address format and its options are described later in this chapter.

The screenshot shows a window titled "Network Layer Addresses". Inside, there's a section for "NSAP Family" with a pull-down menu currently showing "nbs". To the right of this is a "Type" section with three buttons: "CLNP", "CONS", and "NET". Below the "NSAP Family" pull-down menu is a list of options: "nbs", "osinet", "us-gossip-v1", "us-gossip-v2", "user-defined", "hex-pub", and "free-form". To the right of this list are three labels: "Identifier (AFI): 49", "Identifier (IDI): NULL", and "imum 5 octets): _____". Below these are two more labels: "xactly 7 octets): _____" and "Exactly 1 octet): _____". At the bottom of the window, there's a large text area labeled "NSAP:" and two buttons: "Apply" and "Reset". At the very bottom of the window, there's a status bar that says "CLNP NSAP: 497270740701".

Figure 5-1 NSAP Family Window

The following options are available:

- The NSAP Family pull-down menu displays a choice of seven NSAP addressing schema: *nbs*, *osinet*, *us-gossip-v1*, *us-gossip-v2*, *user-defined*, *hex-pub*, and *free-form*. Once a selection is made, its corresponding options are displayed.
- Type specifies which type of network address is being configured: CLNP or CONS network, or with a NET.
- NSAP field displays the NSAP address formed from the specific information that you enter.

Press Apply to save or Reset to return to the last saved settings. For CLNP and CONS NSAPs, the applied settings are implemented immediately. For a NET type, the modifications are implemented when you save the configuration and restart the `osinetd` daemon.

nbs Address Structure

Choose *nbs* from the NSAP Family pull-down menu. The screen shown in Figure 5-2 is displayed. You should use the nbs format for networks that conform to the National Institute of Standards and Technology (NIST) specifications.

Figure 5-2 nbs NSAP Format

Specify the following information to build the NSAP:

Parameter	Description
Authority and Format Identifier (AFI)	For an nbs NSAP format, the AFI of 49 specifies a local binary DSP format. This cannot be changed.
Initial Domain Identifier (IDI)	For an nbs NSAP format, this is null. This cannot be changed.

Parameter	Description
Subnetwork number (Maximum 5 octets)	Specifies a number that identifies the subnetwork, that is: —1 byte long (for example, between 00 and FF) —5 bytes long where the first octet is non-zero (for example, between 01 00 00 00 00 and FF FF FF FF FF)
Station identifier	This is a 7-byte station identifier whose last two digits are FE.
NSEL	A 1-byte network selector that identifies the user of the network layer service, that is, the transport layer.

The nbs NSAP is built from the information that you enter and is displayed on the screen.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

osinet Address Structure

Choose *osinet* from the NSAP Family pull-down menu. The screen shown in Figure 5-3 is displayed. You should use the osinet format if your network conforms to the Open Systems Interconnection Network (OSINET) specifications.

The screenshot shows a window titled "Network Layer Addresses". Inside, there's a section for "NSAP Family" with a dropdown menu set to "osinet". To the right, there's a "Type" section with three buttons: "CLNP", "CONS", and "NET", where "CLNP" is selected. Below this, there are several fields for configuration: "Authority and Format Identifier (AFI): 47", "Initial Domain Identifier (IDI): 0004", "Organization number (Exactly 2 octets):", "Subnetwork number (Exactly 2 octets):", "Station id (Exactly 6 octets):", and "NSEL (Exactly 1 octet):". Each of these last four fields has a corresponding input line. At the bottom left, there's a label "NSAP:". At the bottom center, there are two buttons: "Apply" and "Reset". At the very bottom, a status bar displays "CLNP NSAP: 497270740701".

Figure 5-3 osinet NSAP Format

Specify the following information to build the NSAP:

Parameter	Description
Authority and Format Identifier (AFI)	For the osinet format, the AFI of 47 specifies an IDI format according to OSI 6523-ICD and a binary DSP. You cannot change this value.
Initial Domain Identifier (IDI)	For an osinet format, the IDI of 0004 identifies conformance to OSINET. You cannot change this value.
Organization number	Enter the unique 2-byte number that identifies your organization. It is determined by OSINET.
Subnetwork number	Enter a 2-byte number that identifies the subnetwork.

Parameter	Description
Station id	Enter a 6-byte station identifier: this is often the MAC address.
NSEL	Enter a 1-byte NSEL. This is not specifically used to access the transport layer, since any NSEL is accepted.

The osinet NSAP is built from the information that you enter and is displayed on the screen.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

us-gosp-v1 Address Structure

Choose *us-gosp-v1* from the NSAP Family pull-down menu. The screen shown in Figure 5-4 is displayed. You should use the us-gosp-v1 format when the network conforms to the U.S. Government Open Systems Interconnection Profile (GOSIP) version 1.

The screenshot shows a window titled "Network Layer Addresses". Inside, there is a section for "NSAP Family" with a dropdown menu set to "us-gosp-v1". To the right, there are three radio buttons for "Type": "CLNP" (which is selected), "CONS", and "NET". Below this, there are several fields for configuration: "Authority and Format Identifier (AFI): 47", "Initial Domain Identifier (IDI): 0005", "Organization id (Exactly 2 octets):", "Subnetwork id (Exactly 2 octets):", "End System id (Max 8 octets):", and "NSEL (Exactly 1 octet):". Each of these last four fields has a corresponding input line. At the bottom left, there is a label "NSAP:". At the bottom center, there are two buttons: "Apply" and "Reset". At the bottom of the window, there is a status bar that reads "CLNP NSAP: 497270740701".

Figure 5-4 us-gosp-v1 NSAP Format

Specify the following information:

Parameter	Description
Authority and Format Identifier (AFI)	For the us-gosp-v1 format, the AFI of 47 specifies an IDI format according to OSI 6523-ICD and a binary DSP. You cannot change this value.
Initial Domain Identifier (IDI)	For a us-gosp-v1 format, the IDI of 0005 identifies conformance to U.S. General Services Administration (GSA). You cannot change this value.
Organization id	Enter the 2-octet organization identifier. For us-gosp-v1 format, the organization identifier is determined by GSA.

Parameter	Description
Subnetwork id	Enter the 2-octet subnetwork number, a logical identifier of the subnetwork.
End system id	Enter the end system identifier. This is a number between 4 and 8 octets long which identifies a unique system within an area. It can be the physical or logical address.
NSEL	Enter the 1-byte network selector. This identifies the user of the network services, that is the transport layer.

The us-gossip-v1 NSAP is built up from the information that you enter and is displayed on the screen.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

us-gosp-v2 Address Structure

Choose *us-gosp-v2* from the NSAP Family pull-down menu. The screen shown in Figure 5-5 is displayed. You should use the *us-gosp-v2* format when the network conforms to the U.S. Government Open Systems Interconnection Profile (GOSIP) version 2.

Network Layer Addresses

NSAP Family: us-gosp-v2

Type: CLNP CONS NET

Authority and Format Identifier (AFI): 47

Initial Domain Identifier (IDI): 0005

Data format identifier (Exactly 1 octet):

Admin authority (Exactly 3 octets):

Reserved (Exactly 2 octets): 0000

Routing domain (Exactly 2 octets):

Area id (Exactly 2 octets):

End System id (Exactly 6 octets):

NSEL (Exactly 1 octet):

NSAP:

Apply

Reset

CLNP NSAP: 497270740701

Figure 5-5 us-gosp-v2 NSAP Format

Specify the following information:

Parameter	Description
Authority and Format Identifier (AFI)	For the <i>us-gosp-v1</i> format, the AFI of 47 specifies an IDI format according to OSI 6523-ICD and a binary DSP. You cannot change this value.
Initial Domain Identifier (IDI)	For a <i>us-gosp-v1</i> format, the IDI of 0005 identifies conformance to U.S. General Services Administration (GSA). You cannot change this value.
Data Format Identifier	For <i>us-gosp-v2</i> you should enter 0x80.

Parameter	Description
Administration Authority	Enter the 3-octet administration authority identifier. This is determined by U.S. General Services Administration (GSA).
Reserved	This 2-octet field is reserved. You cannot change it.
Routing Domain	Enter the 2-octet routing domain identifier. It identifies a set of end systems and intermediate systems that use the same procedures and belong to the same administrative domain.
Area Identifier	Enter the 2-octet area identifier. This identifies a specific subdomain within the routing domain.
End System Identifier	Enter the 6-octet end system identifier. This identifies a unique system within an area. It can be the physical or logical address.
NSEL	Enter the 1-octet network selector. This identifies the user of the network service, that is the transport layer.

The us-gosp-v2 NSAP is built from the information that you enter and is displayed on the screen.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

user-defined Address Structure

Choose *user-defined* from the NSAP Family pull-down menu. The screen shown in Figure 5-6 is displayed. Use the user-defined format for network addresses that do not conform to one of the previously described formats. This format verifies the AFI, IDI and DSP individually.

The screenshot shows a window titled "Network Layer Addresses". At the top, there is a "NSAP Family:" dropdown menu set to "user-defined" and a "Type:" section with three buttons: "CLNP", "CONS", and "NET". The "CLNP" button is selected. Below this, there are three input fields labeled "Authority and Format Identifier (AFI):", "Initial Domain Identifier (IDI):", and "Domain Specific Part (DSP):". Below these fields is a large text area labeled "NSAP:". At the bottom of the text area are two buttons: "Apply" and "Reset". At the very bottom of the window, the text "CLNP NSAP: 497270740701" is displayed.

Figure 5-6 user-defined NSAP Format

Specify the following information:

Parameter	Description
Authority and Format Identifier (AFI)	The AFI determines the format of the Initial Domain Identifier (IDI) and the abstract syntax of the Domain Specific Part (DSP). Enter the 1-byte AFI.
Initial Domain Identifier (IDI)	The IDI contains the addresses used in a subnetwork or identifies the authority responsible for allocating the DSP. The length and contents of the IDI are determined by the AFI.
Domain Specific Part (DSP)	The DSP and its syntax are determined by the authority defined by the IDI. It is a unique part of an NSAP address within a domain.

The user-defined NSAP is built from the information that you enter and is displayed on the screen.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

hex-pub Address Structure

Choose *hex-pub* from the NSAP Family pull-down menu. The screen shown in Figure 5-7 is displayed. Use the Hexadecimal Reference Publication format for networks whose addressing conforms to ISO 8348/AD2. This is the same as user-defined formats, but the components in the NSAP are not verified individually. The NSAP is accepted if the AFI is valid and if there is an even number of digits.

Network Layer Addresses

NSAP Family: hex-pub

Type: CLNP CONS NET

NSAP:

Apply Reset

CLNP NSAP: 497270740701

Figure 5-7 hex-pub NSAP Format

Specify the following information:

Parameter	Description
NSAP	Enter a valid 20-byte network service access point.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

free-form Address Structure

Choose *free-form* from the NSAP Family pull-down menu. The screen shown in Figure 5-8 is displayed. You should use the free-form format for networks whose addressing does not conform to any of the standards described. For example, you might want to use the hostname as the NSAP. Since the NSAP is not verified as valid, this is acceptable in free-form.

Note – The free-form NSAP family does not verify that the entered NSAP address conforms to ISO 8348/AD2, “Network service Definition, Addendum 2, Covering Network Layer Addressing.”

The screenshot shows a window titled "Network Layer Addresses". Inside, there is a section for "NSAP Family" with a pull-down menu set to "free-form". To the right, there is a "Type" section with three buttons: "CLNP", "CONS", and "NET". Below this is a large empty rectangular area for input. At the bottom left of this area, there is a label "NSAP:" followed by a pull-down menu set to "Hex" and a long text input field. At the bottom right of the input area are two buttons: "Apply" and "Reset". Below the main window, the text "CLNP NSAP: 497270740701" is displayed.

Figure 5-8 free-form NSAP Format

Specify the following information:

Parameter	Description
NSAP	Press MENU to choose Hex (for a hexadecimal format) or Char (for an ASCII format), and enter an NSAP address up to 20 octets long.

Press Apply to save, or Reset to return to the last saved settings. The changes saved are implemented immediately for a CLNP or CONS network. For a NET, the changes are implemented when you save the configuration and restart the `osinetd` daemon.

Chapter Summary

This chapter explained the different formats that you can use for the NSAP address and how to specify it. Refer to “Addressing” on page 103 for more information about NSAPs.

When you have saved this part of the configuration, go onto Chapter 6, “Route Manager” to configure the network routes.



<i>Route Manager Main Window</i>	<i>page 89</i>
<i>Chapter Summary</i>	<i>page 101</i>

This chapter explains how to configure static routes in Connection Oriented Network Service (CONS), Connectionless Network Protocol CLNP/LLC1 and CLNP/X.25 subnetworks, using the Route Manager.

All the changes you make in Route Manager are dynamic, and are implemented as soon as you press Apply.

Route Manager Main Window

When you choose Route Manager from `ositool`, the screen shown in Figure 6-1 is displayed. You use this screen to add the details of routes to the routing table.

The screenshot shows the 'Route Manager' window. At the top, there's a title bar with a small icon and the text 'Route Manager'. Below the title bar, there are two sections: 'Category:' with a dropdown menu set to 'Host Routes', and 'Network:' with two buttons, 'CONS' and 'CLNP'. The main area contains a table with three columns: 'Type', 'NSAP/NET', and 'SNPA/NET'. Below the table, there are 'Add' and 'Delete' buttons. At the bottom, there's a configuration section with labels and dropdown menus: 'Route:' (set to 'Direct Route'), 'NSAP:' (set to 'Hex'), 'SNPA:' (set to 'Hex'), and 'Subnet:' (set to '0'). There are also 'Apply' and 'Reset' buttons. A button labeled 'X.25 Service...' is located to the right of the 'Subnet:' field.

Figure 6-1 Route Manager Window

From this screen you can configure host or prefix routes, where:

- Host routes define a full address to one host.
- Prefix routes use a masked address, which are used for routing to a group of hosts.

If any routes are already configured, the following information is displayed:

Category	Description
Type	Indicates the type of route, which can be: —DD for direct, dynamic routes —ID for intermediate dynamic routes —DS for direct static routes —IS for intermediate static routes —RS for relayed static routes —XS for extracted static routes
NSAP/NET	Displays the NSAP address for direct or relayed routes or the NET address for intermediate routes. For null CLNP, this is the same as the SNPA.
SNPA/NET	Displays the SNPA address for direct or intermediate routes, or the NET address for relayed routes.
Subnet	Displays the number of the subnetwork that is configured for this address, where: —0 is a CONS/X.25 subnetwork (CLNP is not applicable) —1 is a CLNP/X.25 subnetwork (CONS is not applicable) —2 to 64 are CLNP/LLC1 subnetworks The subnetwork numbers relate to those displayed on the Subnet List in the ES-IS configuration.

Note – Using CLNP Null Protocol

If you are using the null protocol for CLNP, the NSAP address is not known by the remote system. You therefore need to specify the X.121 address for an X.25 connection or the MAC address for a LAN connection, instead of its NSAP and SNPA. Otherwise, during connection, the connect confirm acknowledgement is not received by the remote system. Since the null protocol performs no routing, you must specify a direct route for that system, or specify the connection as the default route

For more information regarding the administration of routes and addresses, see “Addressing” on page 103 and “Dynamic Routing” on page 117.

Host Route Configuration

Press MENU on Category and choose Host Routes from the pull-down menu. The screen shown in Figure 6-2 is displayed. The Host Routes window is used to define routes to a single host.

Route Manager

Category: ☐ Host Routes Network: ☐ CONS ☐ CLNP

Type	NSAP/NET	SNPA/NET	Subnet
D D	6775656e6965767265	08002018f94b	2
D D	4955407c9f00	0800200e035f	2
D D	4955407c9f1e	0800200e035f	2
D D	491e14a8e401	0000c0b7a267	2
D D	64616c6169	08002018f5ca	2
D D	6b6f6469616b	0800201d4f18	2
D D	7361626f756b	0800200ecca8	2

Add Delete Update

Route: ☐ Direct Route

NSAP: ☐ Hex 491e14a8e401

SNPA: ☐ Hex 0000c0b7a267

Subnet: ☐ 2

X.25 Service...

Apply Reset

Figure 6-2 Host Routes Window

Click SELECT on CONS or CLNP to define the type of subnetwork and then perform one of the following tasks:

- Click SELECT on Add to add a new route to the list, then enter the relevant route information.
- Choose a route and edit its relevant information.
- Choose a route and click SELECT on Delete to remove a route from the list.
- Click SELECT on Update for CLNP subnetworks to display up-to-date information about dynamic routes.

Enter the following information about the host route:

Parameter	Description
Route	<p>Press MENU on Route to choose one of the following:</p> <p><i>Direct Route</i> describes a route that sends PDUs for an NSAP to a specific SNPA on a CONS or CLNP subnetwork.</p> <p><i>Relayed Route</i> describes a route that sends PDUs for an NSAP to a specific NET on a CLNP subnetwork.</p> <p><i>IS Route</i> describes a route that sends PDUs for a NET to a specific SNPA on a CLNP subnetwork.</p>
NSAP	<p>Press MENU on NSAP and choose Hex (for a hexadecimal format) or Char (for an ASCII format). Then enter the NSAP address.</p> <p>This option is only valid for a CONS direct route, CLNP direct route, or CLNP relayed route. For a null CLNP route, you must specify the MAC or X.121 address of the remote system. This is the same as the SNPA.</p>
SNPA	<p>Press MENU on SNPA and choose Hex (for a hexadecimal format) or Char (for an ASCII format). Then enter the SNPA address.</p> <p>This option is only valid for a CONS direct route, CLNP direct route, or a CLNP intermediate route. For a null CLNP route, you must specify the MAC or X.121 address of the remote system.</p>
NET	<p>Press MENU on NET and choose Hex (for a hexadecimal format) or Char (for an ASCII format). Then enter the NET address.</p> <p>This option is only valid for a CLNP relayed route or a CLNP intermediate route.</p>
Subnet	<p>Press MENU and specify:</p> <ul style="list-style-type: none">- 0 for a CONS/X.25 subnetwork- 1 for a CLNP/X.25 subnetwork- 2 to 64 for CLNP/LLC1 subnetworks <p>The subnetwork numbers relate to those displayed on the Subnet List in the ES-IS configuration.</p>

Parameter	Description
X.25 Service	For CONS/X.25 subnetworks, click SELECT on the X.25 Service button, to set the additional protocol requirements. Displays a Remote X.25 Feature menu window. These options are described in “Remote X.25 Features” on page 98.

Note – If you are configuring a null CLNP, you must specify the X.121 or MAC address in place of both the NSAP and the SNPA. See “Using Null CLNP” on page 114.

Press Apply to save the changes or Reset to return to the last saved settings. Changes are implemented immediately.

Prefix Routes

Press MENU on Category and choose Prefix Routes from the pull-down menu. The screen shown in Figure 6-3 is displayed.

The Prefix Routes window is used to define routes to a group of hosts, by only using the part of the NSAP that is common to all hosts in the group.

Route Manager

Category: Network:

Type	NSAP Prefix	SNPA/NET	Subnet
X S	360	<Off: 3 Len: 48>	0
X S	36	<Off: 2 Len: 48>	0
X S	521	<Off: 3 Len: 48>	0
X S	52	<Off: 2 Len: 48>	0

Route: SNPA Offset:

NSAP Prefix: SNPA Length:

SNPA:

Subnet:

Figure 6-3 Prefix Routes Window

Click SELECT on CONS or CLNP to define the type of subnetwork and then perform one of the following tasks:

- Click SELECT on Add to add a new route to the list, then enter the relevant route information.
- Choose a route and edit its relevant information.
- Choose a route and click SELECT on Delete to remove a route from the list.

Enter the following information:

Parameter	Description
Route	<p>Press MENU on Route to choose one of the following:</p> <p><i>Direct Route</i> describes a route that sends PDUs for an NSAP to a specific SNPA on a CONS or CLNP subnetwork.</p> <p><i>Relayed Route</i> describes a route that sends PDUs for an NSAP to a specific NET on a CLNP subnetwork.</p> <p><i>Extract Route</i> describes a route that sends PDUs for NSAPs with leading digits that match the NSAP Prefix to the extracted SNPA. An SNPA is extracted from the NSAP address by specifying an offset and a length of the SNPA.</p>
NSAP Prefix	<p>Press MENU on NSAP and choose Hex (for a hexadecimal format) or Char (for an ASCII format). Then enter the full NSAP address. For a null CLNP route, you must specify the MAC or X.121 address of the remote system. This is the same as the SNPA.</p>
SNPA	<p>Press MENU on SNPA and choose Hex (for a hexadecimal format) or Char (for a character format). Then enter the SNPA address. For a null CLNP route, you must specify the MAC or X.121 address of the remote system.</p> <p>Note: This option is only valid for CONS direct prefix routes or CLNP direct prefix routes.</p>
NET	<p>Press MENU on NET and choose Hex (for a hexadecimal format) or Char (for an ASCII format). Enter the NET address.</p> <p>This option is only valid for CLNP relayed prefix routes.</p>
Subnet	<p>Press MENU and specify:</p> <ul style="list-style-type: none"> - 0 for a CONS/X.25 subnetwork - 1 for a CLNP/X.25 subnetwork - 2 to 64 for CLNP/LLC1 subnetworks <p>The Subnetwork numbers relate to those displayed on the Subnet List in the ES-IS configuration.</p>

Parameter	Description
SNPA Offset	Specify a value to indicate the offset of the SNPA from the beginning of the NSAP address. The extracted address begins at the byte following the number specified. For a CONS extracted route, specify the offset in digits. For a CLNP extracted route, specify the offset in bytes.
SNPA Length	Specify a value to indicate the length of the SNPA address for an extracted route. For a CONS network, specify the length in digits. An SNPA length of 48 extracts the SNPA address to the end of the NSAP. For a CLNP network, the length of the SNPA address is not relevant.
X.25 Service	For CONS/X.25 subnetworks or CLNP direct static routes, click SELECT on the X.25 Service button to set the additional protocol requirements. These options are described in “Remote X.25 Features” on page 98.

Note – If you are configuring a null CLNP, you must specify the X.121 or MAC address in place of both the NSAP and the SNPA. See “Using Null CLNP” on page 114 for more details.

Press Apply to save the changes or Reset to return to the last saved settings. Changes are implemented immediately.

Remote X.25 Features

When you choose X.25 Service from the host or prefix routes windows, the screen shown in Figure 6-4 is displayed. You need to specify the type of X.25 service that is used for the network and any protocol options that are required.

Remote X.25 Features

Protocol Options:

☐ Use Reverse Charging ☐ Use Throughput Class Negotiation

☐ Use Flow Control Parameter Negotiation ☐ Use D-bit

☐ Use Fast Select Facility ☐ Use Closed User Group

CUG Format: **X.25 Link Type:** ☒ 1988

CUG Value: **Addressing Mode:** ☒ CONS-84/88

Local Packet Size: Remote Packet Size:

Local Window Size: Remote Window Size:

Local Throughput Class: Remote Throughput Class:

Figure 6-4 X.25 Features Menu

The Link Type and Addressing Mode options can be modified independent of the protocol options. Other settings that you can change depend on which protocol options are selected. For example, the Closed User Group Format and Value can only be changed if the Use Closed User Group checkbox is selected.

The settings that you choose here should agree with the way in which you have set up your X.25 software package.

You can specify the following information, independent of the protocol options that you select:

Parameter	Description
Link Type	<p>Press MENU to choose one of the following X.25 versions:</p> <ul style="list-style-type: none">—1980 supports the 1980 CCITT text of recommendation X.25. This can only be selected if the Addressing Mode is CONS-80.—1984 supports the 1984 CCITT text of recommendation X.25.—1988 supports the 1988 CCITT text of recommendation X.25. <p>Note that these options are dependent on the addressing Mode selected and should agree with the settings for the X.25 software package.</p>
Addressing	<p>Press MENU to choose one of the following X.25 addressing options:</p> <p><i>CONS-84/88</i> if the X.25 network supports CCITT Address Extension Facilities (AEF) and the remote host supports full ISO 8878 CONS. That is, it uses CCITT-specified Calling Address Extension and Called Address Extension facilities to convey full NSAP addresses in the network connection establishment phase. You can only select this if you have chosen a 1984 or 1988 link type.</p> <p><i>CONS-80</i> if the X.25 network does not support CCITT AEF. If you select this parameter, you must use the special X.25-type address. This uses the X.121 address and the Protocol Identifier (PID) to form the NSAP. The formation of the NSAP is explained in more detail in “Addressing” on page 103. You can select this option with any link type.</p>

The protocol options-dependent items that you can modify are:

Protocol Option	Description
Use Reverse Charging	Requests Reverse Charging in outgoing call request PDUs as defined in the CCITT X.25 recommendations.
Use Flow Control Parameter Negotiation	Supports an X.25 subnetwork with the Flow Control Parameter Negotiation as defined in the CCITT X.25 recommendations. These values should agree with those set in your X.25 software package. When you select this option, you can modify the following values: — <i>Local Packet Size</i> - press MENU to select the maximum packet size allowed for outgoing packets. The actual size is negotiated. — <i>Local Window Size</i> - specify a value between 2 and 127 for buffering packets. — <i>Remote Packet Size</i> - press MENU to select the maximum packet size allowed for incoming packets. The actual size is negotiated. — <i>Remote Window Size</i> - specify a value between 2 and 127 for buffering packets.
Use Fast Select	Requests the Fast Select in outgoing call request PDUs as defined in the CCITT X.25 recommendations.
Use Throughput Class Negotiation	Supports the X.25 subnetwork for Throughput Class Negotiation as defined in the CCITT X.25 recommendations. These values should agree with those set in your X.25 software package. When you select this protocol option, you can modify the following values: — <i>Local Throughput Class</i> - press MENU to select the maximum line speed for the outgoing connection. — <i>Remote Throughput Class</i> - press MENU to select the maximum line speed for the incoming connection. The line speeds are defined in your X.25 configuration as a number between 3 for speeds of 75 bps, up to 13 for 64000 bps. The default value is 19200 bps (equivalent to 11 in your X.25 configuration).
Use D-bit	Indicates that the DTE wants to receive an end-to-end acknowledgement of delivery for data it was transmitted using the Delivery Confirmation Bit (D-bit).

Protocol Option	Description
Use Closed User Group	Enables the closed user group facility. When you select this option, you can modify the following values: <ul style="list-style-type: none">— <i>CUG Format</i> - press MENU to specify basic, extended or bilateral use of CUG.— <i>CUG Value</i> - select a value between 0 and 99 to identify the closed user group to which this system belongs. This value is used in the call request to specify the CUG for a virtual call.

Note – The values and options that you set for X.25 should agree with the configuration settings for the X.25 software package.

Chapter Summary

This chapter described the parameters used to configure the host and prefix routes over the CONS or CLNP networks. If you are using X.25 over CONS or CLNP, you might also need to modify the X.25 Service options.

If you have completed the steps in chapters 3 to 6, then you have completed your configuration. Return to the main `ositool` window, save your changes and restart the `osinetd` daemon to implement any static changes that you have made (see “Using the OSI Administration Tool” on page 9 for advice on how to do this).

Chapters 7 and 8 provide some additional information about addressing and routing, to help you with network administration.

<i>Overview of OSI Addressing</i>	<i>page 103</i>
<i>Chapter Summary</i>	<i>page 116</i>

This chapter describes the formats and conventions used for addressing in SunLink OSI 8.1. The general rules for OSI addressing are defined and then each component of the OSI address is discussed in further detail. The following topics are described:

- General overview of OSI addressing
- Subnetwork Point of Attachment (SNPA) addresses
- Network Service Access Point (NSAP) addresses
- Selectors (TSEL, SSEL, and PSEL)
- Network Entity Title (NET)

Overview of OSI Addressing

Addresses in the OSI network are used to uniquely identify originators and destinations for applications used over the network. The OSI addressing scheme uses a mixture of identifiers for the type of network, access points into the network and the stack, from the subnetwork.

OSI addresses can be divided into the following parts:

- Network Services Access Point (NSAP) is a unique address in an OSI network that identifies the network service.
- Selectors (TSEL, SSEL, and PSEL) identify services of peer entities. For example, the TSEL identifies the service provided by the transport layer.

SunLink OSI 8.1 also requires a Subnetwork Point of Attachment (SNPA) to identify a unique point at which the subnetwork is connected to the NSAP.

These are described in more detail in the following sections. Figure 7-1 on page 105 shows a general view of the parts of addresses and how they relate to the OSI Reference Model.

The OSI address format uses the NSAP address plus the selectors (TSEL, SSEL, and PSEL), to connect applications through peer entities in the OSI layers. The SNPA, although not part of the OSI address format, identifies the mapping between the NSAP and the subnetwork. These address components are described in more detail in the following sections.

The NSAP is independent of the subnetwork to which it is connected, while the format and content of the SNPA depends on the subnetwork.

Note – Using CLNP Null Protocol

If you are using the null protocol for CLNP over X.25, the NSAP address is not known by the remote system. You therefore need to specify the X.121 address of the remote site, instead of its NSAP and SNPA. Otherwise, during connection, the connect confirm acknowledgement is not received by the remote system.

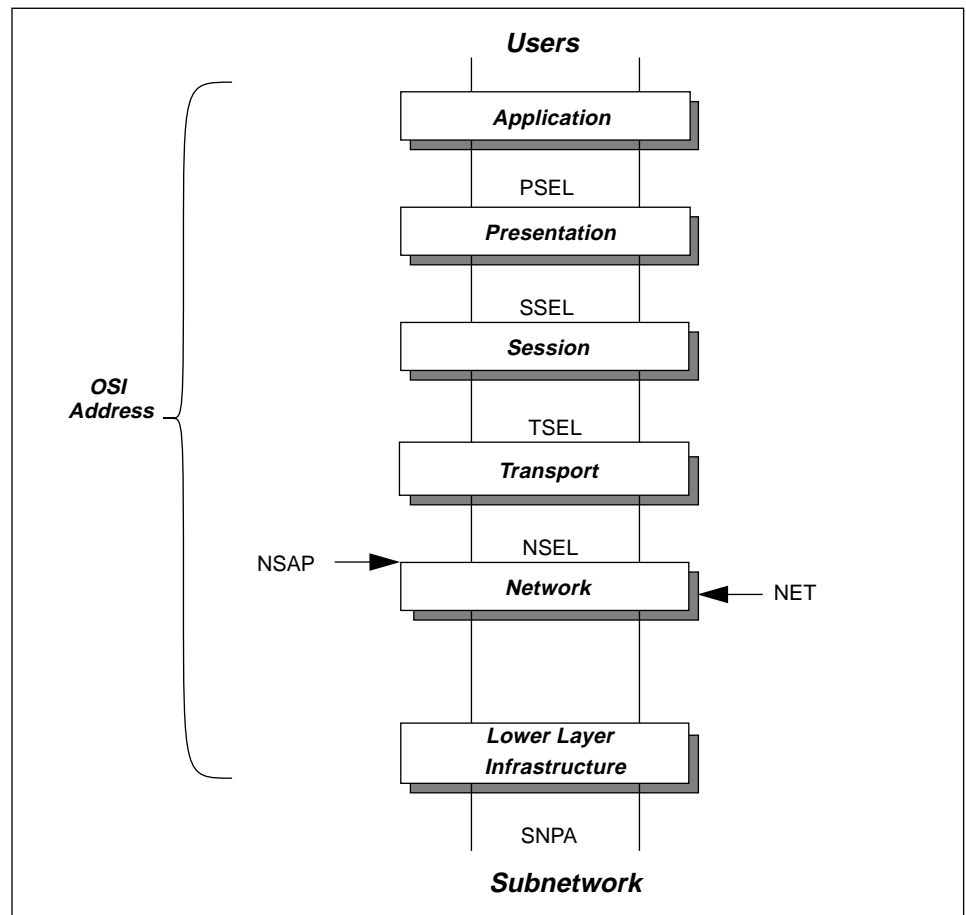


Figure 7-1 Network Addressing

Subnetwork Point of Attachment (SNPA)

The SNPA address uniquely identifies a method of access in an OSI subnetwork. It maps the subnetwork to the NSAP. For example, the SNPA can be the physical address. The format of the SNPA depends on the subnetwork:

- For 802.x, 802.3 and FDDI LANs, the SNPA is the six-byte Ethernet-style address, described by the Medium Access Control layer address (MAC) and the Link Selector (LSEL) 0xfe.

- For a PSDN, the SNPA is the X.121 address. A 14-digit X.121 address, plus 1-digit for the X.121 prefix is supported by SunLink OSI 8.1 for the X.25 over CONS. The X.121 address is derived from the NSAP address for CONS 1980 or the AEF for CONS 1984/1988. This is described in more detail later in this chapter.

Note – The SNPA is not a defined part of the OSI scheme. It is required in all networks as the mapping between the NSAP address and the subnetwork.

Network Service Access Point (NSAP)

The NSAP address is used to identify a network layer service provider. It is the entry point of a network layer entity for a network service user, and is unique in the OSI network domain. The global OSI domain is divided into network addressing domains for which an addressing authority ensures unique address identifiers within that domain.

The NSAP address can be a maximum of 20 bytes in length.

The addressing used in SunLink OSI 8.1 conforms to ISO 8348/AD2, “Network Services Definition, Addendum 2, Covering Network Layer Addressing” and Annex A of the X.213 Recommendation in relation to handling NSAP addresses.

The general format of the NSAP is composed of several components, that indicate the addressing authority, the format of the address, and the access point for network services. The addressing authority determines the actual format of the NSAP. Figure 7-2 illustrates how these components fit together.

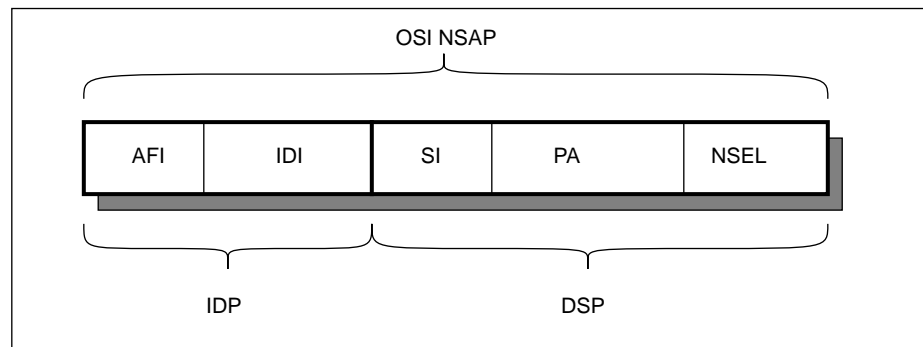


Figure 7-2 NSAP Components

where:

- AFI - Authority and Format Identifier (2 bytes)
- IDI - Initial Domain Identifier
- SI - Subnetwork Identifier
- PA - Point of Attachment
- NSEL - Network Selector (1 byte)
- IDP - Initial Domain Part
- DSP - Domain Specific Part
- NSAP - Network Service Access Point (maximum 20 bytes)

The NSAP components are explained in the following sections:

Initial Domain Part (IDP)

The IDP identifies the addressing authority for the overall NSAP address. Its syntax is in BCD-encoded decimal digits. It comprises the AFI and the IDI:

Authority and Format Identifier (AFI)

The AFI is a 2-digit number in the range 36 to 59 which:

- Identifies the authority that allocates and identifies the format for this NSAP address.

- Determines if a leading zero in an IDI is significant. For example, both AFI 36 and 52 specify an IDI that is an X.121 address and specify decimal syntax for the DSP. However, the AFI 36 indicates that a leading zero in the IDI is not significant, while AFI 52 indicates that a leading zero in the IDI is significant.
- Specifies the syntax of the DSP. For example, an AFI of 44 specifies a syntax of decimal digits for the DSP. An AFI of 50 specifies a syntax of national characters for the DSP

The authority and format identified by the valid AFIs are indicated in Table 7-3 on page 111.

Initial Domain Identifier (IDI)

The IDI format and syntax is determined by the AFI. It can determine the type of addresses used in a subnetwork (such as a telephone number), or it can identify the authority that allocated it. For example, an AFI of 38 specifies an IDI that is a three-digit country code whereas an AFI of 47 specifies an IDI determined by OSI 6523-ICD. The authority and format identified by valid AFIs are indicated in Table 7-1.

Table 7-1 IDI Descriptions

IDI Format	Description
X.121	Up to 14-digit address used across Packet-Switched Data Networks plus 1 digit for the prefix digit. Allocated according to CCITT Recommendation X.121.
ISO DCC	3-digit Data Country Code.
F.69	Up to 8-digit telex number allocated according to CCITT Recommendation F.69.
E.163	Up to 12-digit Public Switched Telephone Number (PSTN) allocated according to CCITT Recommendation E.163.
E.164	Up to 15-digit ISDN number allocated according to CCITT Recommendation E.164.
ISO 6523-ICD	4-digit International Code Designator.
Local	Null IDI.

For the currently valid AFIs, ISO 8348/AD2 specifies the maximum length and syntax for the IDI and DSP fields.

For IDIs with variable-length formats, ISO requires that IDIs are padded to the maximum length to identify the end of the IDP and the beginning of the DSP. For each variable-length IDI (binary and decimal syntax), ISO allocates two AFI values, where:

- Leading zeroes in the IDI are significant. Therefore, the network layer pads with ones.
- Leading zeroes in the IDI are not significant. Therefore, the network layer pads with zeroes.

For example, AFIs of 36 and 52 both specify a decimal-syntax IDI that is a variable-length X.121 address. However, an AFI of 36 indicates that leading zeroes are not significant, and zeroes are used to pad the IDI to its maximum size. The AFI of 52 indicates that leading zeroes are significant, so ones are used to pad IDIs.

Table 7-2 lists the AFIs that specify variable-length IDIs and whether leading zeroes are significant.

Table 7-2 IDI Leading Significance of Zeroes

AFI Value	IDI Format	DSP Syntax	Leading Zero Significant
36	X.121	Decimal	No
37		Binary	No
52		Decimal	Yes
53		Binary	Yes
40	F.69	Decimal	No
41		Binary	No
54		Decimal	Yes
55		Binary	Yes
42	E.163	Decimal	No
43		Binary	No
56		Decimal	Yes
57		Binary	Yes
44	E.164	Decimal	No
45		Binary	No
58		Decimal	Yes
59		Binary	Yes

Domain Specific Part (DSP)

The Domain Specific Part (DSP) is unique within a given addressing authority's domain. Its semantics are determined by the authority defined in the IDI. In some cases, the DSP is null and the IDI number is a public number, similar to a telephone number. There are three network address formats:

- CCITT
- ISO
- Local

Depending on the value of the AFI, a DSP can have a syntax of decimal digits, binary octets (hexadecimal numbers), ISO 646 characters, or characters from a national character set.

While authorities beneath the authority identified by the AFI can control the format and meaning of different parts of the DSP, ISO determines the syntax for the entire DSP. The DSP is comprised of the following components:

Subnetwork Identifier (SI)

This is a global identifier for the attached subnetwork.

Point of Attachment (PA)

This describes the physical attachment address in relation to the network.

Network Selector (NSEL)

The last two digits of the NSAP address describe the Network Selector. This identifies the user of the network layer service, that is, it identifies the transport layer. For all real NSAP address families, SunLink OSI 8.1 assumes that the last byte (two hexadecimal digits) of the NSAP address is the NSEL. While the NSEL is part of the NSAP address, its handling is distinct from the way the rest of an NSAP address is handled.

NSAP Address Field Lengths and Syntax

Table 7-3 is an extension of the *Maximum NSAP Address Lengths* table in ISO 8348/AD2. This information is provided for general reference to ensure that you use the correct number of digits for each component in the NSAP address.

The IDP length can be obtained by adding the length of the AFI (2 bytes) to the IDI length. The total length of the NSAP is equal to the sum of the IDP and the DSP.

Table 7-3 NSAP Address Field Lengths

AFI Value	IDI Format	DSP Format	NSAP (Binary DSP)	NSAP (Dec)	IDI Length	IDP Length	DSP Length
36	X.121	Decimal	20	40	14	16	24
37		Binary	17	39	14	16	9
38	ISO DCC	Decimal	20	40	3	5	35
39		Binary	17	40	3	5	14
40	F.69	Decimal	20	40	8	10	30
41		Binary	17	40	8	10	12
42	E.163	Decimal	20	40	12	14	26
43		Binary	17	39	12	14	10
44	E.164	Decimal	20	40	15	17	23
45		Binary	18	40	15	17	9
46	ISO 6523-ICD	Decimal	20	40	4	6	34
47		Binary	16	39	4	6	13
48	Local	Decimal	20	40	Null	2	38
49		Binary	16	40	Null	2	15
50		Character	20	40	Null	2	19
51		National Char	15	37	Null	2	7
52	X.121	Decimal	20	40	14	16	24
53		Binary	17	39	14	16	9
54	F.69	Decimal	20	40	8	10	30
55		Binary	17	40	8	10	12
56	E.163	Decimal	20	40	12	14	26
57		Binary	17	39	12	14	10
58	E.164	Decimal	20	40	15	17	23
59		Binary	18	40	15	17	9

Types of NSAP Addresses

The NSAP that you use to configure your system depends on the type of network and the standards to which it conforms. The types of NSAPs are known as families and the following can be used in SunLink OSI 8.1:

- *nbs* for network addresses conforming to the National Institute of Standards and Technology (NIST).
- *osinet* for network address conforming to the Open Systems Interconnection Network (OSINET).
- *us-gosip-v1* for network addresses conforming to U.S. Government Open Systems Interconnection Profile, Version 1 (U.S. GOSIP, v1).
- *us-gosip-v2* for network addresses conforming to U.S. Government Open Systems Interconnection Profile, Version 2 (U.S. GOSIP, version 2).
- *user-defined* to specify NSAPs that are not covered by the above specifications, in decimal.
- *hex-pub* to specify NSAPs that not covered by the above specifications, in hexadecimal format.
- *free-form* to specify NSAPs that do not conform to other specifications. The contents of the NSAP is not verified.

The configuration of these is explained in more detail in “Network Layer Addresses” on page 73. All types of NSAPs are treated in exactly the same way by SunLink OSI 8.1. Some combinations of families are interchangeable. Any valid NSAP address can be in the *user-defined* or *hex-pub* families, whereas the *nbs*, *osinet*, *us-gosip-v1*, and *us-gosip-v2* families require specific values for the IDP and parts of the DSP. All NSAP addresses in all families are handled identically in SunLink OSI 8.1.

The characteristics and requirements of each family as used in SunLink OSI 8.1 are described in the following examples:

nbs NSAP Address

- AFI of 49 and a null IDI
- A subnetwork number that is 1-byte long (00 to FF) or 5 bytes long where the first octet is non-zero (for example, 01 00 00 00 to FF FF FF FF)
- A seven-byte station id
- A one-byte NSEL

osinet NSAP Address

- An AFI of 47 and an IDI of 4
- A two-byte organization number
- A two-byte subnetwork number
- A six-byte station id (often a MAC address)
- A one-byte NSEL

us-gosip-v1 NSAP Address

- An AFI of 47 and an IDI of 5
- A two-byte organization id
- A two-byte subnetwork id
- An End System id of between four and eight bytes
- A one-byte NSEL

us-gosip-v2 NSAP Address

- An AFI of 47 and IDI of 5
- A data format identifier of 0x80
- A two-byte reserved field for 0
- A three-byte administration authority id
- A two-byte routing domain id
- A two-byte area id
- A six-byte End System id
- A one-byte NSEL

user-defined NSAP Address

All valid BCD NSAP addresses of up to 20-bytes in length. That is, all NSAP addresses that meet the requirements of ISO 8348/AD2 can be in this family. Each component (AFI, IDI and DSP) of the NSAP address is checked separately for validity.

hex-pub NSAP Address

All valid hexadecimal NSAP address of up to 20-bytes in length. hex-pub is the same as the user-defined family because all valid NSAP addresses can be in this family. A check is made that the first two digits are a valid AFI and that there is an even number of digits. Other checks are also made based on the value of the AFI to ensure that a valid NSAP has been specified.

free-form NSAP Address

The *free-form* address can be any 40-digit hexadecimal or 20-character string. SunLink OSI 8.1 does not verify that the entered NSAP address conforms to ISO 8348-AD2.

X.25 over CONS NSAP

For networks with 1984 X.25 supporting Address Extension Facilities (AEF), you can use a real NSAP address. For networks with 1980 X.25 over CONS, you need to use a special X.25 address that is based on the SNPA (that is, the X.121 address) plus the following:

- A prefix of 52, if the X.121 address begins with a zero. For example, if the X.121 address is 04325223, then the NSAP is 5204325223. If the address is specified in binary, use a prefix of 53.
- A prefix of 36, if the X.121 address begins with a non-zero number. For example, if the X.121 address is 31311432, then the NSAP is 3631311432. If the address is specified in binary, use a prefix of 37.
- Padding digits where the resulting NSAP address does not have an even number of digits. For example, an X.121 address that starts with a zero and has an odd number of digits requires the addition of a 1 after the prefix of 52, thus 024094322 becomes 52102409432. An X.121 address that starts with a non-zero and has an odd number of digits requires the addition of a 0 after the prefix of 36; thus, 24037121005 becomes 36024037121005.

Using Null CLNP

For CLNP over X.25 networks which use the inactive or null protocol, you need to specify the X.121 address of the remote system in place of the NSAP and SNPA. For systems using null CLNP over a LAN, you must specify the MAC address of the remote system, instead of the NSAP and SNPA.

Since the null protocol does not perform routing functions, you must either specify a direct route for the remote system (using Route Manager), or specify that the default route is a direct route to that system (using ES-IS Configuration).

Selectors

The selector identifies the user of a service, that is, the entity in the next layer above. For example, the Transport Selector (TSEL) identifies the session entity.

Figure 7-3 indicates the relationships between the selectors and the access points that describes them.

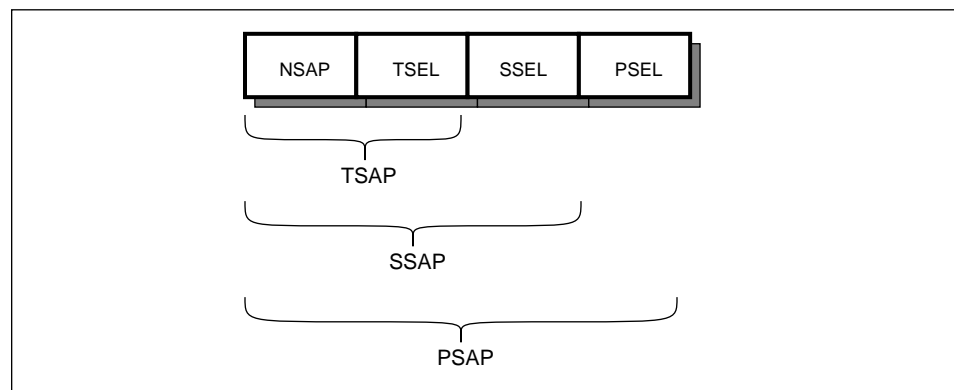


Figure 7-3 Network Connection

where:

- NSAP - Network Service Access Point
- TSEL - Transport Selector
- SSEL - Session Selector
- PSEL - Presentation Selector
- TSAP - Transport Service Access Point
- SSAP - Session Service Access Point
- PSAP - Presentation Service Access Point

The following selectors are used in OSI addressing:

- Transport Selector (TSEL) is used by the transport service to identify a session layer entity. The TSEL plus the NSAP address is known as the TSAP address.
- Session Selector (SSEL) is used by the session layer to identify a presentation layer entity. The SSEL plus the TSAP address is known as the SSAP address.

- Presentation Selector (PSEL) is used by the presentation service to identify an application layer entity. The PSEL plus the SSAP is known as the PSAP address.
- TSAP, SSAP, and PSAP addresses communicate with an entity in the named layer or service and offers its services to the next higher layer.

The maximum length for each selector that can be configured with `ositool` is:

- 4 octets for PSEL
- 16 octets for SSEL
- 32 octets for TSEL

Network Entity Title (NET)

A Network Entity Title (NET) is the single network layer address by which an Intermediate System (IS) is known on all subnetwork connections. It is often defined as the NSAP minus the NSEL in many networks. However, the NET that you specify is not verified for this format by the software. You can define your own format for the NET.

Chapter Summary

This chapter described some of the components that are used to form network addresses. SunLink OSI 8.1 uses OSI addressing for connecting to LAN and X.25 subnetworks, and TCP/IP addressing for connecting to TCP/IP subnetworks.

Each component used in an OSI address (SNPA, NSAP and selectors) was defined.

<i>ES-IS Protocol</i>	<i>page 117</i>
<i>Routing Sequence</i>	<i>page 122</i>
<i>Using Token Ring Networks</i>	<i>page 122</i>
<i>CLNP over X.25</i>	<i>page 123</i>
<i>Chapter Summary</i>	<i>page 123</i>

This chapter describes the way that the End System to Intermediate System (ES-IS) protocol works to support dynamic routing functions over a LAN or PSDN subnetwork. End Systems and Intermediate Systems (which perform routing functions) use the ES-IS protocol to exchange information about their own addresses, and about routes to reach other systems.

Refer to “End System to Intermediate System Configuration” on page 61 for information about configuring the ES-IS timers and parameters, and “Route Manager” on page 89 for configuring host and prefix routes.

ES-IS Protocol

The ES-IS protocol is an International Standard defined in ISO 9542. It provides a means of dynamically and automatically updating the routing information maintained by hosts located on an OSI subnetwork. It does this by a process of

broadcasting messages and exchanging network addresses that can be added, updated, and deleted automatically in the routing tables maintained by each system.

This release of the software does not support the IS-IS protocol, that is, it does not allow dynamic routing between intermediate systems.

Multicast addresses are used to broadcast information about a system's address and routes. The stack uses the all-ES and all-IS multicast addresses recommended in the NIST December 1988 agreements (Phase II). All systems on a given subnetwork must use the same ES-IS multicast addresses.

How ES-IS Works

The ES-IS protocol uses three methods to distribute routing information:

- End Systems that operate ES-IS periodically exchange their SNPA addresses with each other to inform systems on the network of their presence.
- Intermediate Systems (ISs) act as routers on a subnetwork, to optimize the routes used between systems. If an intermediate system receives PDUs destined for another system, and it knows a better route, it informs the originator of the better route. That is, the intermediate system sends the originator the SNPA address of the ultimate receiving system.
- If there are no ISs on a subnetwork and an ES does not know the SNPA address of a destination machine, the ES can broadcast a query packet to all ESs on the subnetwork, asking "Whose address is this?"

If the address belongs to a machine, that machine accepts the packet and sends its address back to the sending machine.

ESs periodically send End System Hellos (ESHs) to all ISs on the local subnetwork. This is a broadcast announcing its existence and its address to others on the networks. An ESH contains:

- A list of NSAP addresses for the sending ES
- An implicit SNPA address for the sending ES
- A holding time for this information

The ESHs are sent to the all-IS multicast address.

ISs periodically send an Intermediate System Hello (ISHs) to all ESs on the local subnetwork. This is a broadcast announcing its existence and its address to other systems on the networks. An ISH contains:

- A Network Entity Title (NET) for the sending IS (equivalent to an NSAP address for an ES)
- An implicit SNPA address for the sending IS
- A holding time for this information

The ISHs are sent to the all-ES multicast address. ESs listen for ISHs on the all-ES multicast address. ISs listen for ESHs on the all-IS multicast address.

When an ES receives routing information from an ISH, it updates its routing table to include the SNPA and NET of the sending IS. Similarly, when an IS receives an ESH, it updates its own routing database to include the SNPA and NSAP of the sending ES.

In addition to the periodic exchange of Hellos, the ES-IS protocol provides the following services:

- *Redirection*

If an ES does not know the SNPA address that corresponds to the NSAP address of the destination machine, it sends the packet to an IS on the local subnetwork.

If the IS knows the next-hop SNPA address (the next SNPA on the route), and the address is on the same subnetwork as the originating ES, it forwards the packet to the next-hop machine and sends a Redirect PDU back to the originating ES. The Redirect PDU contains the SNPA and logical network addresses of the next-hop host (ES or IS).

The Request Redirect function handles the IS sending of a Redirect PDU and the Record Redirect function handles the ES listening to a Redirect PDU. Together, these request and record functions make up the redirect subset.

- *Query Configuration*

If there are no ISs on a subnetwork and an ES does not know an immediate destination for an outgoing packet, the ES sends a Query Configuration PDU to the all-ES multicast address.

A query configuration PDU is identical to a CLNP data PDU except that the immediate destination field (MAC address) is set to the all-ES multicast address. If an ES receives a query configuration PDU with an NSAP address that belongs to it, it accepts the packet and sends an ESH to the originating ES.

- *Adding and Removing Systems from the Subnetwork*

When an ES or IS receives new routing information (presumably from a system that has just come up), it immediately responds with a Hello that is directed to the new system, instead of waiting to send a Hello at the expiration of its configuration timer. This feature allows newcomers to the network to quickly attain full connectivity on the subnetwork.

Whenever an ES or IS is terminating, it sends a Hello with the holding timer set to zero. This causes the receiving ESs or ISs to remove the routing entries that are associated with the sending system. This notification is sent to another when an ES or IS on the subnetwork is down.

End System Responsibilities

End systems in the subnetwork perform the following functions:

- Sends ESHs (one for each source address) to the all-IS multicast address, at the expiration of its send ES timer.
- Listens for ISHs on the all-ES multicast address and updates the IS list in response to an ISH. Once the ISH is received, the IS list is updated with the SNPA address of the sending IS. It also records the Network Entity Title (NET) of the IS.
- Sends query configurations PDUs to the all-ES multicast address if the network layer cannot find a route for an outgoing packet and if there is no IS available.
- Listens for query configuration PDUs on the all-ES multicast address. If a query configuration PDU is intended for a host, the network layer accepts the packet and sends the originator an ESH containing the SNPA address of the specified host.
- Listens for Redirect PDUs and updates the host routing table in response to a Redirect PDU. If the IS that sends the Redirect PDU used a prefix routing entry to route the original packet toward its destination, the IS includes it in

the Redirect. In addition to host-route information, the prefix route information is sent in the form of a bit mask. The ES that receives such a Redirect updates both its host and prefix routing tables.

- Removes entries in the host routing table and the IS list when the holding timer is expired.
- Sends an ESH a newly active IS on the reception of a new routing entry (from an ISH), indicating that an IS was added.
- Sends an ESH with the holding timer set to zero to the all-IS multicast address when an ES leaves the network.

Intermediate System Responsibilities

Intermediate systems in the subnetwork perform the following functions:

- Sends periodically ISHs to the all-ES multicast address.
- Listens for ESHs on the all-IS multicast address and updates the host routing table in response to an ISH. When the ESH is received, the host routing table is updated with the NSAP and SNPA addresses of the sending IS.
- Forwards the packet and sends a Redirect PDU back to the originating ES if it knows a more efficient route. The Redirect PDU contains the more efficient route. If the IS that sends the Redirect PDU used a prefix routing entry to route the original packet toward its destination, the IS will also include a bit mask from the Redirect host-route and prefix-route information.
- Sends a directed ISH to the newly active ES when it receives a new routing entry (from an ESH), indicating that an ES was added.
- Sends an ISH with the holding timer set to zero to the all-ES multicast address when an IS leaves the network.

In the ES-IS protocol document, the function by which an ES or IS sends out Hellos is called the Report Configuration function. When an IS sends a Redirect PDU it is called the Request Redirect function. IS can perform either of these functions on a per-interface basis.

Routing Sequence

A destination NSAP address in an outgoing packet requires the network layer to route the packet through the following steps:

1. It matches the NSAP address in the host routing table defined during Route Manager configuration, or defined dynamically with directed ESHs and Route Redirect PDUs.
2. If no match is found in the host routing table, a prefix match is searched for in the prefix routing table defined in the Route Manager configuration. For an IS, a packet is undeliverable if the prefix match fails.
3. If no match is found for an ES in the prefix table, it looks for an entry in the IS list. The IS list is built from information received in ISHs.
4. If there are no matching IS list entries, the ES invokes the query configuration function and sends a query configuration PDU.

Use of the ES-IS protocol over networks without multicast capabilities, such as PSDNs, is possible in subnetworks with or without routers or ISs. When a query configuration PDU is sent, it uses the same steps for selecting a source address as it does for any other packet.

Using Token Ring Networks

The multicast address used for SunLink OSI 8.1 ES-IS is not valid for token ring networks. SunLink OSI 8.1 uses the standard ES multicast address of 09002b000004 and IS multicast address of 09002b000005.

These are not accepted by token ring networks, which require multicast addresses of 030000000200 and 030000000100. SunLink OSI 8.1 continues to work over token ring networks, but you cannot use the ES-IS protocol to route dynamically.

Configure all token ring routes manually, using route manager, or set the default subnetwork for the token ring connection.

CLNP over X.25

If you have multiple X.25 links or are a member of a Closed User Group (CUG), the use of the redirect subset of ES-IS is not recommended. This implementation of ES-IS treats Connectionless Network Protocol (CLNP) over X.25 as a single subnetwork, regardless of the number of links. Therefore, an ES connected to one PSDN might receive (by redirection), a CLNP route to an ES on a different PSDN where the first ES has no X.25 connectivity. Similarly, an ES might receive a route to an ES outside the first ES or CUG domain, where the ES is prevented from establishing a virtual circuit.

Chapter Summary

This chapter described some of the characteristics of the ES-IS protocol and how it works to ensure up-to-date addressing and routing information throughout the subnetwork.

The ES-IS protocol uses the configured routing tables (defined when you use the Route Manager configuration tool) as a basis, and adds new information that is broadcast over the subnetwork.

Troubleshooting and Diagnostics

9 

<i>Troubleshooting Overview</i>	<i>page 125</i>
<i>Overview of Diagnostics</i>	<i>page 127</i>
<i>Problems Interworking with SunPro Products</i>	<i>page 138</i>
<i>Chapter Summary</i>	<i>page 138</i>

This chapter explains some of the ways in which you can avoid problems in your network configuration and helps to solve some problems that can occur. Network problems can often be solved and efficiency can be improved by adjusting the configuration parameters described in this book.

Troubleshooting Overview

This section deals with the prevention and solution of minor problems that you might encounter while setting up and configuring the network.

Consistency

To ensure that your system works properly, you should plan your network thoroughly to allow consistent and logical conventions throughout your administrative domain. In particular:

- Use consistent naming conventions and addressing techniques.

- Maintain a logical order to your administrative domain in the network, as well as noting the conventions used in other parts of the network.
- Use hexadecimal or ASCII notations, and quantity notations (bytes or octets, and digits) carefully.

Installation

Ensure that your system is installed correctly and that the license requirements are specified correctly. Refer to *Installing and Licensing SunLink 8.1* for full details.

Terminating Applications to Restart the Daemon

When you have finished modifying the configuration with `ositool`, you need to save the changes and restart the `osinetd` daemon to implement static changes. If any associated applications are already running, you need to terminate them before you can restart the daemon.

Use the following command to search for associated applications that may be running:

```
hostname% ps -ef | egrep "osi|vt"
```

Other applications may be running that also use the stack, such as those written for the APLI or TLI, could also be running. This command only checks for SunLink applications.

If any OSI applications are shown (other than the `grep` command), then you should terminate them using:

```
hostname% /etc/rc2/<startup script name> stop
```

If any associated applications are still running, the daemon will not restart and an error message is displayed on the console. Check if there are any other applications running using:

```
hostname% modinfo | grep -i osi
```

Terminate the listed applications with the following command. Note that you need to become superuser or root to use this command:

```
hostname# modunload -i <application identifier>
```

If this does not work, reboot your system.

Repeat this command for each application that is still running. Now you can stop ositool:

```
hostname% /etc/rc2.d/S90osinet stop
```

Restart the osinetd daemon using the start option:

```
hostname% /etc/rc2.d/S90osinet start
```

Any static modifications that were saved are now implemented.

Overview of Diagnostics

Problems in a network are often related to routing. There are two main tools that you can use for diagnosing problems:

- The `osi_ping` function is used to determine if a route exists to a specified host on a CLNP network, by sending packets to the host, which are then echoed back to the originator.
- The `osi_trace` function handles the 802.x-related problems by displaying packets at selected layers, starting at the link layer and proceeding upward.

Full details of these and their associated options can be obtained with the online manual pages provided.

The `osi_ping` Function

The `osi_ping` program is used to diagnose CLNP routing problems. It checks to see whether there is a route between a local machine and a remote host. You can use `osi_ping` to ping any host on your OSI network that supports CLNP

echo requests according to RFC1139. Such hosts include any machine that runs the current release of SunLink OSI 8.1. There are two implementations of the `osi_ping` function as defined by RFC1139:

- Long-term implementation defines two new ISO 8473 PDU types: an echo request (ERQ) and an echo reply (ERP). Long-term implementation is the default version, if you do not specify any options for `osi_ping`.
- The short-term version uses an existing OSI 8473 data PDU as the ERQ and ERP, but uses an NSEL of 0x1e to identify an ERQ and 0x1f to identify an ERP.

Note – SunLink 7.0 systems only support the short-term version of `osi_ping`.

When you use the `osi_ping` function, you specify a destination host and an optional timeout parameter. You can either specify a hostname as an alias for the NSAP or the NSAP address itself. The general format of the command is:

```
hostname% osi_ping [-s] [-n <NSAP address>] [-N <NSAP address>] [-d
<data length>] [host <hostname>]
```

where the options are:

`-s`

Specifies that the short-term version is used.

`-n <NSAP address>`

Specifies the remote NSAP address in hexadecimal.

`-N <NSAP address>`

Specifies the remote NSAP address in ASCII.

`-d <data length>`

Specifies the number of bytes to be sent as in the PDU to test this route. The value can be in the range of 4 to 8192. The default is 4 bytes.

`host <hostname>`

Specifies that `osi_ping` should look for the NSAP associated with the `hostname` in the alias file. When used with the `-n` or `-N` options, it defines a host as an alias for a given NSAP address.

The alias is stored in a temporary file for future use. If you run `osi_ping` as root or superuser, then this file is not temporary, and new host aliases are not lost when the daemon is restarted.

Note – You cannot use `osi_ping` over the inactive subset (null) of CLNP.

Five echo requests are sent to the remote host. If the host responds, `osi_ping` will report “*host is alive.*” If the host does not respond, `osi_ping` will report one of the following messages:

no route to `<hostname>`

The CLNP entity does not know how to route to `<hostname>`. Use `ositool` to verify that the route to the host is correctly defined in the ES-IS configuration and routing tables.

unknown host `<hostname>`

There is no entry for host name `<hostname>` in the alias file. Use `-n` or `-N` plus the host option to specify an alias for the host.

no answer from `<hostname>`

`<hostname>` did not respond, the configuration is correct, but the host is not active.

The `osi_trace` Function

The `osi_trace` program is used to monitor OSI and X.25 activity throughout the stack. You can use `osi_trace` to trace the incoming and outgoing Protocol Data Units (PDUs) for a specified host. In this function, `osi_trace` is identical to `x25trace`. You can use the `osi_decode` program to decode the hexadecimal output of `osi_trace` and `x25_trace`.

Note – You cannot use the trace function for TCP/IP networks since there is no NIT (Network Interface Tap) in the operating system. Tracing can only be done at the LLC level, or at the LAPB or X.25 levels, if X.25 is being used. Tracing TCP/IP packets when using RFC1006 is not possible.

You must log in as root or become superuser to run the program. The general form of an `osi_trace` command is:

```
hostname# osi_trace [options] <filter_expression>
```

The options and filter expressions are described below.

`osi_trace` *Options*

The following options can be used for all types of trace:

`-a`

This displays the number of user data bytes transferred, not the actual user data values. By default it displays this number in hexadecimal for the highest protocol specified. For example, if you enter:

```
prompt# osi_trace -a -i /dev/llc2 tp ses
```

`osi_trace` displays only the number of user data bytes transferred by the session layer.

`-i interface`

This specifies one of the following interfaces:

`/dev/llc2` — the LLC interface for CLNP over LLC1. This is the default interface.

`/dev/lapb` — the LAPB interface for CONS or CLNP over X.25.

`/dev/x25` — the X.25 packet layer interface for CONS or CLNP over X.25.

`-t`

This turns off the default operation that displays a time stamp at the beginning of each line.

`-T`

This applies a time stamp to the delta times between packets; it is measured in microseconds.

-u

This outputs the buffer file line-by-line instead of using the default, which is buffered packet-by-packet.

-x

This decodes the packet and displays a hexadecimal value. Error packets are identified with an error message starting with two asterisks (**).

-l length

This specifies the number of packets to print with the -x option.

-c count

This limits the total number of packets displayed to the value specified by count.

`osi_trace` *Filter Expressions*

You can specify the highest layer to be tested in the `osi_trace` command using one of the following options:

`acse`

Decodes all ACSE Layer headers contained in the Session user data.

`ber`

Decodes session user data according to the Basic Encoding Rules (BER).

`ber_debug`

Specifies that the BER tag and length bytes should be printed before the BER symbolic trace.

`clnp`

Traces all CLNP PDUs.

`cmip`

Decodes all CMIP PDU headers contained in the Session user data.

`cons`

Traces CONS PDUs.

`esis`

Traces all ES-IS PDUs.

`hdlc`

Traces the LABP PDUs.

`inactive`

Traces the inactive (null) protocol.

`llc`

Traces LLC1 and LLC2 PDUs.

`pres`

Specifies Basic Encoding Rules (BER) for decoding the session user data.

`pres+`

Specifies the decoding of all headers contained in the Session user data. Specifying this option is the equivalent of specifying `pres`, `acse` and `cmip`.

`ses`

Specifies session layer decoding of the transport user data.

`tp`

Specifies transport layer decoding of the network user data.

`x25`

Traces X.25 PLP PDUs.

If you specify a layer higher than the one with which an application interfaces, then an error occurs. The `osi_trace` function tries to decode the packets at the higher level, but cannot. For example, if you specify `ses` when an application interfaces directly with the transport layer, then `osi_trace` displays many error messages.

Trace Examples

This section provides some examples of using `osi_trace` for tracing incoming and outgoing PDUs.

To trace HDLC and X.25 using `/dev/lapb`, the following example of trace command packets are sent or received by X.25. Replace HDLC with X.25, if required.

```
hostname# osi_trace -i /dev/lapb hdlc
```

To trace using a single X.25 connection (logical channel number):

```
hostname# osi_trace -i /dev/lapb x25lcn <logical channel number> x25
```

To trace only the next X.25 connection that is set up:

```
hostname# osi_trace -i /dev/lapb x25lcn + x25
```

To trace using CLNP packets using `/dev/l1c2`:

```
hostname# osi_trace -i /dev/l1c2 clnp
```

or for CLNP packets using `/dev/x25`:

```
hostname# osi_trace -i /dev/x25 clnp
```

To trace using ES-IS packets:

```
hostname# osi_trace -i /dev/l1c2 esis
```

To trace using CLNP packets to or from a specific NSAP address:

```
hostname# osi_trace -i /dev/llc2 clnp nsap 47000400210000108200001020300
hostname# osi_trace -i /dev/llc2 clnp srcnsap 47000400210000108200001020300
hostname# osi_trace -i /dev/llc2 clnp dstnsap 47000400210000108200001020300
hostname# osi_trace -i /dev/llc2 clnp betweennsap 47000400210000108200001020300\
4700040021000108200011223300
```

The example above traces PDUs to and from an NSAP; PDUs from the specified source NSAP; PDUs to the specified destination NSAP; and PDUs between two particular NSAPs, respectively.

To trace using CLNP, inactive, or ES-IS to or from specific MAC addresses on llc2:

```
hostname# osi_trace -i /dev/llc2 srcmac 8:20:0:1:2:3 clnp
hostname# osi_trace -i /dev/llc2 dstmac 8:20:0:1:2:3 inactive
hostname# osi_trace -i /dev/llc2 betweenmac 8:20:0:1:2:3 9:0:2b:18:21:5 esis
```

The example above traces packets from a particular source MAC address, packets going to a particular destination MAC address, and packets going between two specified MAC addresses, respectively.

To trace using ES-IS packets that are not multicast (such as redirect packets or directed ESHs):

```
hostname# osi_trace -i /dev/llc2 esis not multicast
```

To trace using all transport protocol PDUs:

```
hostname# osi_trace -i /dev/llc2 pdu tp
```

To trace a specific TP/CONS connection on /dev/x25:

```
hostname# osi_trace -i /dev/x25 cons tp x251cn 0x3fe
```

The above command specifies the TP/CONS connection with x251cn.

To trace the next TP/CONS connection that is being set up on /dev/x25:

```
hostname# osi_trace -i /dev/x25 cons tp x25lcn +
```

To trace session PDUs:

```
hostname# osi_trace -i /dev/l1c2 ses
```

With the `pres` option, `osi_trace` does not decode and display presentation headers. Instead, it decodes and displays the ASN.1 Basic Encoding Rules (BER) without any knowledge of the presentation (or any application layer) protocol.

The osi_decode Program

The `osi_decode` program decodes the hexadecimal output of `osi_trace` and `x25_trace`.

`osi_decode` has the following syntax:

```
hostname% ?????/osi_decode options filter
```

The options are:

-a

Displays the number of bytes of user data, rather than the data itself. By default `osi_decode` displays the user data in hexadecimal for the highest protocol specified.

-l *line_length*

By default, `osi_decode` truncates lines following column 51. Use `-l` to specify a different line length.

-i *level*

Specifies the protocol level at which hex user data is to be interpreted. The available levels and expected formats are:

protocol layer	data interpreted as
clnp	CLNP PDUs
cons	CONS PDUs
esis	CLNP PDUs
hdlc	LAPB PDUs
inactive	CLNP PDUs
llc	LLC1 and LLC2 PDUs
net	CLNP PDUs
tp	Transport PDU
ses	Session PDU
pres	Presentation PDU

The default level is `net`.

`-u`

Causes display output to be buffered line-by-line.

The valid filters are shown below. You can set more than one filter.

`acse`

Decodes all ACSE Layer headers contained in the Session user data.

`ber`

Decodes session user data according to the Basic Encoding Rules (BER).

`ber_debug`

Specifies that the BER tag and length bytes should be printed before the BER symbolic trace.

`clnp`

Traces all CLNP PDUs.

`cmip`

Decodes all CMIP PDU headers contained in the Session user data.

`cons`

Traces CONS PDUs.

`esis`

Traces all ES-IS PDUs.

`hdlc`

Traces the LABP PDUs.

`inactive`

Traces the inactive (null) protocol.

`llc`

Traces LLC1 and LLC2 PDUs.

`pres`

Specifies Basic Encoding Rules (BER) for decoding the session user data.

`pres+`

Specifies the decoding of all headers contained in the Session user data. Specifying this option is the equivalent of specifying `pres`, `acse` and `cmip`.

`ses`

Specifies session layer decoding of the transport user data.

`tp`

Specifies transport layer decoding of the network user data.

`x25`

Traces X.25 PLP PDUs.

Problems Interworking with SunPro Products

If you install SunLink OSI 8.1 on the same system as one or more SunPro products, you may not be able to obtain a license for one or both of the products.

To solve the problem, look in the `/etc/opt/licenses/licenses_combined` file and see whether there is a line that begins `DAEMON lic.SUNW`. If the line is missing, re-enter it, as shown below:

```
DAEMON lic.SUNW /etc/opt/licenses/lic.SUNW /etc/opt/licenses/daemon_option
```

Chapter Summary

This chapter provided some suggestions for solving problems that you might have during configuration or may have been caused by specific configuration options. It described the `osi_ping` and `osi_trace` commands that you can use for solving routing problems.

Configuration File



<i>Command Components</i>	<i>page 140</i>
<i>Commands</i>	<i>page 142</i>

The `osinetd.conf` file contains all the configuration information about your network regarding routing, addressing, and stack parameters that you set with the OSI administration tool (`ositool`).

This chapter describes the components of the configuration language used in the `osinetd.config` file, and provides some examples of parts of the configuration file. It is not a full description of all the commands used in the `osinetd.config` file.

Warning – This information is provided only for your reference—if you need to update the configuration, use `ositool`. You should not edit the configuration file directly.

The `osinetd` daemon opens, loads, and initializes relevant STREAMS modules. This is required to implement the OSI STREAMS multiplexor during the start-up of the stack. After start-up, it continues running to maintain the links between the STREAMS modules that provide the OSI service to user applications, such as X.400 and FTAM. The configuration file contains information required by the `osinetd` daemon for creating these links and maintaining the STREAMS multiplexor.

Command Components

The `osinetd` daemon configuration commands are specified in a simple language that consists of tokens, keywords, identifiers, numeric values, and strings. These are described in the following sections.

Tokens

The language is divided into a number of *tokens*. Tokens can be considered as any strings that are separated by a white space, horizontal-tab, carriage return, new line, vertical-tab, form feed, end-of-file, or start of comment character (`#`). For example, keywords, identifiers, numeric values and character strings are considered as tokens.

Keyword

A *keyword* describes how a command works. These are always specified in lowercase characters. Table A-1 shows the strings that are reserved as keywords.

Table A-1 Keyword Parameters

address	close	conind_number	default
dlpi	initoper	link	lsap
macdev	multicast	npi	npiloop
nsap	off	on	oper
osiamx	ositcp	sendoper	set
timer	trace	type	under

Identifier

An *identifier* can be an alphabetic or numeric string of up to 64 characters in length, where the initial character must be an alphabetic character. Although you can mix uppercase and lowercase characters, the use of them in identifiers is case sensitive. You must use them consistently.

Some examples of acceptable identifiers are:

- llcdev0
- LLCdev0
- This_Is_An_Identifier
- other_character!!*

Some examples of illegal identifiers are:

- 99ident
- Embedded#comment
- embedded space

Numeric Value

The *numeric value* is defined as a positive integer or a hexadecimal number.

An integer can be a numeric value between the range of 0 to ($2^{32} - 1$).

A hexadecimal number has a maximum length of 62 digits with a character range of 0-9 and A-F (or a-f), and is identified with a prefix of 0x, or 0X. For example:

- 0X01
- 0x012abCDA

String

A *string* is defined as a string of characters between double quotation marks (for example, "string"). You cannot place another set of quotation marks within the quotation marks. You can break continuous long quotes in several lines with a backslash (\). The maximum number of characters is 509.

Commands

Tokens are grouped together to form commands. Commands are separated by new lines. You can break a command into several lines using a backslash (\). The following commands are supported:

- `trace` - enables or disables verbose trace output
- `opdv` - opens a device driver
- `set` - sets a multicast address for the LLC1 driver
- `link` - links a STREAMS driver
- `sendoper/initoper` statements - operator statements that send commands to the stack
- `close` - closes a previously opened device

Verbose Trace Command

The verbose `trace` command is used to enable or disable a trace output. The trace output is issued on `stdout`. The default is off and the format is:

```
trace { on | off }
```

Opening a Device Driver

To build the OSI multiplexor, the daemon must open a number of STREAMS device drivers. An example of a device open command is:

```
opdv = "/dev/oopi0" type oper
```

You must specify the type of device that is to be opened and provide any initialization information to the daemon. Do this by adding qualifiers to the basic device open command.

LLC1 Device

The general format of the command is:

```
<identifier> = <string1> type dlpi lsap <lsap> macdev <string2>
```

where:

<identifier> is a string identifying the device.

<string1> is the pathname of the LAN device module.

<lsap> is a hexadecimal number that identifies the LSAP to the device.

<string2> specifies that the MAC device is to be linked under the LLC1 device driver.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the LLC over le Device.  
#  
llcdev0 = "/dev/llc2" type dlpi lsap 0xfe macdev "le"
```

X.25 Device

The general format of the command to open an X.25 device driver is:

```
<identifier> = "/dev/x25" type npci address <address>  
[conind_number <conind_number>]
```

where:

<identifier> is a string identifying the device.

<address> is a hexadecimal number that represents a device-specific Network Provider Interface (NPI) address that is passed to the driver during the bind operation.

<conind_number> is a decimal number that specifies the maximum number of outstanding connections permitted for listener streams. If this is not specified, then the opened stream to the device is used by the listener to service outgoing/incoming connections.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the NPI X25 Device.
#
# Listener Streams queue, handling 3 concurrent connections
#
npidevC0 = "/dev/x25" type npid address
0x000000000000b12345654321000000000000000000000000000000000000000000
000000000000 conind_number 3
```

RFC1006 Device

The general format of a command to open an RFC1006 device is:

```
<identifier> = "/dev/tcp/" type ositcp
[ conind_number <conind_number> ]
```

where:

<identifier> is a string that identifies the device.

<conind_number> is a decimal number specifying the maximum number of outstanding connections permitted for listener streams. If this is not specified, then the opened stream to the device is used by the listener to service outgoing/incoming connections.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the TCP/IP STREAMS Driver.
#
# Listener Streams queue, handling 4 concurrent connections
#
tcpdev = "/dev/tcp" type ositcp conind_number 4
```

Timer Device

The general format of a command to open a timer device driver is:

```
<identifier> = "/dev/otmr" type timer
```

where:

<identifier> is a string that identifies the device.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the TIMER STREAMS Driver.  
#  
timerdev = "/dev/otmr" type timer
```

OSI STREAMS Device Driver

The general format of the command to open the OSI STREAMS device driver is:

```
<identifier> = "/dev/ooip1" type osiamx
```

where:

<identifier> is a string that identifies the device.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the OSIAM STREAMS Multiplexor.  
#  
stack = "/dev/ooip1" type osiamx
```

Device for Operator Commands

The general format of the command to open the device for operator commands is:

```
<identifier> = "/dev/oopi0" type oper
```

where:

<identifier> is a string that identifies the device.

This type *oper* is used for `sendoper` and `initoper` commands. (see “Operator Statements” on page 148).

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Opens the operator Streams queue.  
#  
opdv = "/dev/oopi0" type oper
```

Set Command

Currently, the only `set` command supported is the setting of a multicast address for an LLC1 driver. The format for this command is:

```
set multicast <number> <identifier>
```

where:

<number> is a driver-specific hexadecimal number that is given to the driver as a multicast address.

<identifier> is an identifier associated with an already opened stream to an LLC1 device.

Example of this command in the `osinetd.conf` file made by `ositool` for ESHs and ISHs are:

```
### Multicast : Accepts ESH (Hello ES) frames.  
#  
set multicast 0x09002b000004 llcdev0  
  
### Multicast : Accepts ISH (Hello IS) frames.  
#  
set multicast 0x09002b000005 llcdev0
```

The OSI multicast address is specific to the ES-IS protocol and must *not* be changed.

Link Command

This command is used to link an opened STREAMS under the OSI multiplexing driver. The format of the command is:

```
link <identifier1> under <identifier2>
```

where:

<identifier1> is an identifier associated with an already opened stream to a device. Upon successful completion of the link command, this identifier is no longer accessible.

<identifier2> is an identifier associated with an already opened stream to the OSI STREAMS device. This device must be of the *osiamx* type.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Links the TCP/IP STREAMS Driver.  
#  
link tcpdev under stack
```

Operator Statements

The daemon uses the operator statements to send commands to the OSI stack. No syntax checking is performed on the command line and the string is sent to the OSI multiplexor as it is received. If the command must be sent to the OSI stack only at initialization time, that is, “e2i res”, (reserving memory for internal tables), then the statement must be introduced by the `initoper` keyword; otherwise, use the `sendoper` statement.

The format of the two commands is:

```
initoper <string> on <identifier>
sendoper <string> on <identifier>
```

where:

`<string>` is a character string recognized by the OSI stack as a command.

`<identifier>` is a string that identifies the device. This device must be of the *oper* type.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Configures the CLNP address.
#
sendoper "def suf 130 H490155421f4e00" on opdv

### Starts the OSIAM Stack.
#
initoper "start" on opdv
```

Close Command

This command is used to close a previously opened device. The general format is:

```
close <identifier>
```

where:

<identifier> is an identifier associated with an already opened stream to a device.

An example of this command in the `osinetd.conf` file made by `ositool` is:

```
### Closes the operator Streams queue
#
close opdv
```

≡ A

Index

A

- accept reuse of TC, 28
- acceptor reuse timeout, 27
- acknowledgement timer, 34, 35
- ACSE, 24
- add
 - device, 58
 - route, 92, 95
 - subnetwork route, 64
 - system, 120
- additional options
 - session, 29
 - transport & CLNS, 33
 - transport over CONS, 40
- address
 - authority, 106
 - Ethernet, 105
 - IP, 106
 - network, 18
 - NSAP, 106
 - variable length, 109
 - X.121, 106, 108
- Address Extension Facility - See AEF
- addressing, 103
- administration authority, 82
- AEF, 99, 106, 114
- AFI, 75, 77, 79, 81, 84, 106

- all-ES multicast address, 118
- all-IS multicast address, 118
- allocation quantum, 30, 33, 40
- API connections, 57
- application layer, 4
- Application Presentation Library Interface
 - See API
- application selectors, 45
 - applications, 47
 - network access, 47
 - selectors, 47
- application-dependent layers, 2
- architecture, 1
- area identifier, 82
- Authority and Format Identifier - See AFI

B

- basic CUG format, 100, 101
- bilateral CUG format, 101
- block diagram
 - ES-IS Configuration, 17
 - Network Layer Addresses, 18
 - Route Manager, 19
 - Stack Manager, 16
- boot file, 6
- buffer space, 30, 33, 37, 40, 65

busy, 58

C

cache entry, 69

CCITT, 110

change

- dynamic, 13
- static, 13

channel configuration, 56, 58

checksum, 32, 35, 67

class

- negotiation, 100
- options, 39

CLNP, 18, 89, 123

- ES-IS defaults, 70
- lifetime, 32
- LLC1 subnetwork route, 93
- null, 63, 91, 93, 96, 114
- options, 66
- route, 64
- routing problems, 127
- X.25 subnetwork route, 93

CLNS, 3

close command, 148

Closed User Group - See CUG

CLTP, 4

CMIP resource configuration, 57

command

- configuration, 142
- pull-down menu, 14

configuration

- commands, 142
- file, 6, 139
- language components, 140
- menu, 45
- process, 20

connection

- per device, 65
- pool, 53, 55
- TCP/IP, 55
- timeout, 27
- timer, 44

Connection Oriented Network Service -
See CONS

connection status, 49

- CONS, 44
- presentation & ACSE, 25
- session, 27
- transport, 31
- transport over CONS, 37

Connectionless Network Protocol - See
CLNP

Connectionless Network Service - See
CLNS

Connectionless Transport Protocol - See
CLTP

CONS, 3, 18, 89

- connection timer, 44
- deconnection timer (X25/80), 44
- default X.25 packet size, 44
- entity, 43
- grcb pool size, 45
- max NSDU length, 44
- rccb pool size, 45
- recb pool size, 45
- resource configuration, 58
- status, 44
- transport, 36
- X.25 route configuration, 98
- X.25 subnetwork route, 93

consistency, 125

console, 14, 15

contexts, 56, 58

control block quantum (CLNP), 35

conventions used, xxii

credit window, 30, 31, 33, 37, 40

CUG, 101, 123

D

daemon, 6, 14

data

- format identifier, 81

Data Country Code - See DCC

data link

- layer, 3

d-bit, 100
DCC, 108
debug FU options, 29
deconnection timer, 44
default
 installation, 11
 subnetwork, 63, 91
 values for ES-IS, 70
delete a route, 92, 95
delivery confirmation bit, 100
device
 add, 58
 driver, 142
 name, 63
 open command, 142
 operator commands, 146
device configuration, 45, 51
 ESH multicast, 53
 ISH multicast, 53
 LSAP, 53
 MAC device, 53
diagnosing problems, 125
direct route, 93, 96
disconnect
 on error, 32, 42
 unused NC, 42
DLPI device, 52
document set, xvii
Domain Specific Part - See DSP
DSP, 106, 109, 110
dynamic
 change, 13
 route information, 92
 routing, 61, 65, 117

E

E.163, 108, 111
E.164, 108, 111
edit route information, 92, 95
End System Hello - See ESH
End System-Intermediate System - See
 ES-IS

entity, 5
entry name, 63, 65
environment variable
 HELPPATH, 10
 MANPATH, 10
 PATH, 10
error
 options, 42
 reporting, 67
 TPDU, 42
ES, 117, 120
 default, 69
 identifier, 80, 82
ESH, 118
 multicast address, 53
 record, 69
 route entries, 66
 send, 69
ES-IS
 configuration, 12, 17, 61
 default values, 70
 protocol, 64, 117, 119
 protocol options, 69
 timers, 68
Ethernet address, 105
expedited data, 28, 42
extended
 concatenation, 29
 CUG format, 101
 format, 32
extracted route, 96

F

F.69, 108, 111
fast select, 100
FDDI, 105
flow control, 34, 41, 100
free-form network address, 86, 114
FTAM resource configuration, 57
FU, 29
full protocol, 63, 67
Functional Units - See FU, 29

G

grcb pool size, 45

H

help, 12

HELPPATH variable, 10

hex-pub network address, 85

high interface, 57

holding timer, 69, 119

host route, 90

- configuration, 92, 93

- NET, 93

- NSAP, 93

- SNPA, 93

- subnet, 93

- X.25 service, 94

I

identifier, 140

IDI, 75, 77, 79, 81, 84, 106, 108

- format, 111

IDP, 106, 111

ignore

- acknowledgement timer, 35

- checksum, 35

- criteria for ACK, 35

inactivity timer, 32

include SSAP-ID in AC SPDU, 28

Initial Domain Identifier - See IDI

Initial Domain Part - See IDP

initiator reuse timeout, 27

initoper keyword, 148

Intermediate System - See IS

Intermediate System Hello - See ISH

international code designator, 108

Internet Protocol address, 106

IS, 116, 117, 121

- route, 93

- static entries, 66

ISDN number, 108

ISH, 118

- multicast address, 53

- record, 69

- route entries, 66

- send, 69

ISO, 110

- 6523-ICD, 111

- DCC, 108, 111

- international code designator, 108

K

keyword, 140

L

LAN connection, 52, 53

language components, 140

leading zero significance, 109

license requirements, 9

line speed, 100

link

- number, 53

- type, 99

link command, 147

Link Selector - See LSEL

Link Service Access Point - See LSAP

LLI device, 53

local IDI format, 111

local packet size, 100

long NC timeout, 41

low interface

- layers, 3

- resource configuration, 58

LSAP, 53

LSEL, 105

M

MAC, 53, 105

manager

- route, 19

- stack, 15

MANPATH variable, 10

- maximum
 - alter ctx size, 25
 - multiplexing, 40
 - PDU size, 37
 - PDUs, 30
 - queue size, 42
 - reassemble size, 35
 - TSDU queue, 27

Medium Access Control - See MAC

MENU button, xxii

MHS resource configuration, 57

migrating the configuration file, 14

multicast address, 118

multiplexing, 40, 41, 58

N

name, 52

naming conventions, 125

national character set, 110

National Institute of Standards and Technology, 112

nbs address structure, 75, 112

NC timeout, 41

NET, 18, 73, 91, 116, 119

- host route, 93

- prefix route, 96

network

- access, 47

- addressing, 12, 105

- connection, 41, 42, 115

- format, 110

- layer, 3

- problems, 125

- reset, 42

Network Entity Title - See NET

network layer addresses, 73

- administration authority, 82

- AFI, 75, 77, 79, 81, 84

- area id, 82

- data format identifier, 81

- end system identifier, 80, 82

- free-form, 86

- hex-pub, 85

- IDI, 75, 77, 79, 81, 84

- manager, 18

- nbs, 75

- NSAP, 74

- NSEL, 78, 80, 82

- organization id, 77, 79

- osinet, 77

- routing domain, 82

- station identifier, 76, 78

- subnetwork number, 76, 77, 80

- type, 74

- user-defined, 83

- us-gossip-v1, 79

- us-gossip-v2, 81

Network Selector - See NSEL

Network Service Access Point - See NSAP

Network Service Data Unit - See NSDU

network-dependent layers, 2

next-hop, 119

NIST, 112

non-segmenting protocol, 67

NSAP

- address, 103, 106, 119

- family, 73, 74

- format, 106, 112, 115

- host route, 93

- length, 110, 111

- manager, 18

- prefix route, 96

- route, 91

NSDU, 35

- length, 44

- non-segmenting protocol, 67

NSEL, 78, 80, 82, 106, 110

null

- CLNP, 63, 91, 93, 96, 114

- PID, 39

- protocol, 63, 67

number of subnetwork, 64

numeric value, 141

O

online help, 12

- opdv, 142
- Open Systems Interconnection Network, 112
- organization identifier, 77, 79
- OSI
 - addressing, 103
 - NSAP address structure, 106
 - STREAMS device driver, 145
- OSI Administration Tool - See `ositool`
- OSI Reference Model, 1, 2
- `osi_ping`, 127
- `osi_trace`, 127, 129
 - filter expressions, 131
 - options, 130
- `osinet` network address, 77, 113
- `osinetd` daemon, 6, 140
 - restarting, 14
- `osinetd.conf`, 6
 - identifier, 140
 - keyword, 140
 - numeric value, 141
 - string, 141
 - tokens, 140
- `ositool`, xv, 7, 9, 12, 125
- outgoing packet, 122
- overview of configuration process, 20

P

- PA, 104, 106, 110
- packet size, 44, 100
- padding characters, 109
- PATH variable, 10
- PDU
 - extended format, 32
 - length restrictions, 29
 - maximum size, 31, 37
 - queue size, 30
 - size, 25
 - tracing, 132
- physical interface identifier, 53
- physical layer, 3
- PID, 39

- ping host, 128
- Point of Attachment - See PA
- prefix for X.121, 114
- prefix route, 90, 94, 96
 - NET, 96
 - NSAP, 96
 - SNPA, 96
 - SNPA offset, 97
 - subnetwork type, 96
- presentation & ACSE
 - entity, 24
 - resource configuration, 57
 - status, 25
- presentation layer, 4
- Presentation Selector - See PSEL
- Presentation Service Access Point - See PSAP
- problems, 125
- process redirect, 69
- propose
 - checksum, 32
 - reuse of TC, 28
- Protocol Identifier - See PID
- protocol options
 - CLNP, 66
 - session, 28
 - subset, 63
 - transport & CLNS, 32, 35
 - transport over CONS, 42
 - X.25, 98
- PSAP, 115
- PSDN, 53, 106
- PSEL, 47, 115
- PSTN number, 108
- Public Service Data Network - See PSDN

Q

- QOS threshold, 41
- Quality of Service - See QOS
- query PDU, 119
- queue size, 34, 35, 42

R

- rcb pool size, 45
- reassemble size -CLNP, 35
- reassign transport connection, 37
- recb pool size, 45
- receiving
 - system, 118
 - window, 34
- record
 - ESH, 69
 - ISH, 69
- redirect
 - PDU, 69, 119
 - process, 69
- refresh cache, 69
- relayed route, 93, 96
- remote
 - packet size, 100
 - window size, 100
 - X.25, 98
- remove
 - route, 92, 95
 - system, 120
- request redirect, 119
- requirements, 9
- resource configuration, 45, 56
 - busy, 58
 - channels, 58
 - CMIP, 57
 - CONS, 58
 - contexts, 58
 - FTAM, 57
 - high interface, 57
 - low interface, 58
 - MHS, 57
 - presentation & ACSE, 57
 - session, 57
 - transport & CLNS, 57
 - transport over CONS, 58
- restart daemon, 14, 126
- retransmission
 - limit, 31, 38
 - timer, 31, 38

- reuse connection, 27
- reverse charging, 100
- route
 - add, 92, 95
 - direct, 93, 96
 - dynamic, 92, 117
 - edit, 92, 95
 - extracted, 96
 - host, 90, 93
 - IS, 93
 - manager, 19
 - relayed, 93, 96
 - remove, 92, 95
 - type, 91
- route manager, 89
 - configuration, 12
 - window syntax, 89
- router, 118, 119
- routing
 - domain, 82
 - problems, 127
 - sequence, 122
 - table size, 45, 65, 66
- RPC utilities, 10

S

- S90osinet module, 6
- SAP, 49
 - list, 49
 - suffix, 49
- Save button, 13
- segments, 35
- SELECT button, xxii
- selectors, 47, 104, 115
- send
 - burst of TPDU, 35
 - ESH, 69
 - ESH timer, 68
 - ISH, 69
 - ISH timer, 68
 - redirect, 69
 - to all ES, 69
 - to default ES, 69

sending
 queue size, 34
 system, 118
sendoper, 148
Service Access Point - See SAP
session
 entity, 26
 functional units, 29
 layer, 4
 protocol options, 28
 resource configuration, 57
 status, 27
 version, 29
Session Selector - See SSEL
Session Service Access Point - See SSAP
set checksum, 35
set command, 146
setenv command, 10
short NC timeout, 41
SI, 64, 80, 106, 110
significance, 109
SII entries, 66
SNPA, 96, 104, 105, 119
 address, 54
 host route, 93
 prefix route, 97
specifications, xviii
SSAP, 115
SSEL, 48, 104, 115
stack, 5
 boot file, 6
 manager, 12, 15, 21
 starting, 11
standards, xviii
start
 ositool, 11
 stack, 10
startup daemon, 6
static
 change, 13
 route entries, 66
 routing, 65
station identifier, 76, 78

status, 31, 37, 44, 49
STREAMS modules, 6
string, 141
subdevice number, 53
subnetwork
 access, 105
 connection, 55, 63, 65
 device, 47
 list, 63
 number, 63, 76, 77
 type, 91, 93, 96
Subnetwork Identifier - See SI
Subnetwork Point of Attachment - See SNPA
suffix, 49
support network reset, 42
system console, 14

T

TCP/IP
 subnetwork connection, 55
terminating applications, 126
throughput class negotiation, 100
TIDU size, 34, 41
timeout
 session connection, 27
 TTR, 37
timer
 device driver, 145
 ES-IS, 68
 retransmission, 38
 TS1/TS2, 37
TI-RPC, 10
TLI connections, 57
tokens, 140
TPDU, 35
 timer, 34
trace
 command, 142
 PDUs, 129
 using ES-IS packets, 134
Transmission Control Protocol - See TCP

-
- transport
 - class, 38, 39
 - layer, 4
 - reuse connection, 28
 - transport & CLNS
 - additional options, 33
 - credit window, 31
 - default CLNP lifetime, 32
 - max PDU size, 31
 - protocol options, 32
 - resource configuration, 57
 - retransmission limit, 31
 - retransmission timer, 31
 - status, 31
 - window timer, 32
 - Transport Interface Data Units - See TIDU
 - Transport Layer Interface - See TLI
 - transport over CONS, 36
 - allocation quantum size, 40
 - class, 39
 - credit window, 37
 - error options, 42
 - long NC timeout, 41
 - max multiplexing, 40
 - max PDU size, 37
 - max size/nofc, 42
 - protocol options, 39, 42
 - QOS threshold/mpx, 41
 - QOS threshold/nfc, 41
 - resource configuration, 58
 - retransmission limit, 38
 - retransmission timer, 38
 - short NC timeout, 41
 - status, 37
 - TIDU size, 41
 - TS1/TS2 timer, 37
 - TTR timeout, 37
 - TWR increment, 38
 - Transport PDU - See TPDU
 - Transport Selector - See TSEL
 - Transport Service Access Point - See TSAP
 - Transport Service Data Units - See TSDU
 - troubleshooting, 125
 - TRS expedited data, 28
 - TS1/TS2 timers, 37
 - TSAP, 115
 - TSAP ID, 39
 - TSDU, 34, 41
 - TSEL, 48, 104, 115
 - TTR timeout, 37
 - TWR increment, 38
 - type
 - network, 74
 - route, 91
 - subnetwork, 63, 65
 - subnetwork connection, 47
- ## U
- U.S. Government Open Systems
 - Interconnection Profile - See us-gossip
 - upper layers, 4
 - user-defined network address, 83
 - us-gossip-v1 network address, 79, 113
 - us-gossip-v2 network address, 81, 113
- ## V
- variable-length format, 109
 - verbose trace command, 142
- ## W
- window
 - Additional Session Options, 29
 - Additional Transport & CLNS Options, 33
 - Additional Transport over CONS Options, 40
 - Application Selector, 47
 - Command Menu, 14
 - Configuration Menu, 45
 - CONS, 43
 - Console, 15
 - Device Configuration, 51
 - ES-IS Configuration, 62
 - free-form network address, 86

- hex-pub network address, 85
- Host Routes, 92
- nbs network address, 75
- NSAP family, 74
- osinet network address, 77
- Prefix Routes, 95
- Presentation & ACSE, 24
- Resource Configuration, 56
- SAP List, 49
- Session, 26
- size, 100
- Stack Manager, 23
- timer, 32
- Transport & CLNS, 30
- Transport over CONS, 36
- user-defined network address, 83
- us-gossip-v1 network address, 79
- us-gossip-v2 network address, 81
- X.25 Features Menu, 98

window example

- OSI Administration Tool, 13

X

- X.121 address, 54, 106, 108, 114
- X.25
 - addressing, 99
 - CLNP, 123
 - connection pool, 53
 - ES-IS defaults, 70
 - link number, 53
 - link type, 99
 - over CONS, 114
 - packet size, 44
 - route configuration, 98
 - service, 94
 - set CUG, 101
 - SNPA, 54
 - subnetwork device, 53
 - version, 99