HP-UX IPv6 Transport Administrator's Guide for TOUR 1.0

HP-UX 11i

Manufacturing Part Number : 5187-4159
E0703

United States

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

A copy of the specific warranty terms applicable to your Hewlett- Packard product and replacement parts can be obtained from your local Sales and Service Office.

U.S. Government License

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

Contents

Contents

Contents

About This Document

This document describes how to install, configure, and troubleshoot HP-UX 11i IPv6 transport bundled as part of TOUR 1.0.

The document printing date and part number indicate the document's current edition. The printing date will change when a new edition is printed. Minor changes may be made at reprint without changing the printing date. The document part number will change when extensive changes are made.

Document updates may be issued between editions to correct errors or document product changes. To ensure that you receive the updated or new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

The latest version of this document can be found on line at:
`docs.hp.com/hpux/netcom/index.html#IPv6`.

Intended Audience

This document is intended for system and network administrators responsible for installing, configuring, and managing IPv6 transport. Administrators are expected to have knowledge of Transmission Control Protocol/Internet Protocol (TCP/IP) networking concepts and network configuration. As well it is helpful to have knowledge of operating system concepts, commands, and configuration.

This document is not a tutorial.

New and Changed Documentation in This Edition

Overall, the document has been updated to reflect that IPv6 transport is included as part of the TOUR 1.0 bundle. The previous version of this guide was for a web download “pull” version of (only) the HP-UX 11i IPv6 product software.

Publishing History

Table 1 Publishing History Details

Document Manufacturing Part Number	Title	OS Supported/ Release	Publication Date
5187-4159	HP-UX IPv6 Transport Administrator's Guide for TOUR 1.0	HP-UX 11i : TOUR 1.0 (Web Download)	July 2003
T1306-90001	Installing & Administering HP-UX 11i IPv6 Software	HP-UX 11i: IPv6 Standalone Product (IPv6NCF11i) (Web Download)	Sep 2001

What Is in This Document

This manual provides information for administering HP-UX 11i IPv6 transport software bundled as part of TOUR 1.0. HP-UX 11i IPv6 transport software uses the next generation Internet Protocol (IPv6) to connect HP-UX Servers and Workstations with other systems running IPv4 or IPv6 over IEEE 802.3, Ethernet or FDDI Local Area Networks. An IPv6 for HP-UX 11i network can extend over routers into a Wide Area Network.

This manual is organized as follows:

Chapter 1 **Features Overview:** provides a features summary and overview of HP-UX 11i IPv6 transport bundled as part of TOUR 1.0.

Chapter 2 **Installation:** provides a brief summary of the procedures necessary to install HP-UX 11i IPv6 bundled as part of TOUR 1.0 on your system.

Chapter 3 **Configuration:** describes how to automatically or manually configure HP-UX 11i IPv6 transport.

Chapter 4 **Troubleshooting:** provides flowcharts to help diagnose HP-UX 11i IPv6 software problems.

Chapter 5 **Utilities:** describes useful tools for installing, configuring, and maintaining IPv6 for HP-UX 11i software.

Chapter 6 **IPv6 Addressing and Concepts:** defines networking terms and explains network interface names, network addresses, names and prefixes.

Chapter 7 **IPv6 Software and Interface Technology:** defines terms used by the I/O system to identify HP-UX 11i IPv6 bundled as part of TOUR 1.0.

Appendix A **IPv6 ndd Tunable Parameters:** provides a list of new ndd Ipv6 tunable paramters that allow for advanced performance tuning.

If you are unfamiliar with IPv6 networking concepts, refer to Chapter 6, “IPv6 Addressing and Concepts,” on page 51 and Chapter 7, “IPv6 Software and Interface Technology,” on page 65, before configuring IPv6 interfaces.

Related Documents

HP Documentation

Additional information about HP-UX IPv6 can be found on *docs.hp.com* in the *networking and communications* collection under *IPv6* at:

`www.docs.hp.com/hpux/netcom/index.html#IPv6`

Other documents in this collection (besides this guide) include:

HP-UX IPv6 Porting Guide

Related RFCs

The following table lists RFCs supported by HP-UX 11i IPv6 bundled as part of TOUR 1.0. Note that the IETF (Internet Engineering Task Force) RFCs listed below can be located at: `http://www.ietf.org/rfc.html`.

Table 2 IPv6 RFCs Supported

RFCs	Description
RFC 1933	Transition Mechanisms for IPv6 Hosts and Routers
RFC 1981	Path MTU Discovery for IPv6
*RFC 2292	Advanced Sockets API for IPv6
RFC 2373	IPv6 Addressing Architecture
RFC 2374	IPv6 Aggregatable Global Unicast Address Format
RFC 2375	IPv6 Multicast Address Assignments
RFC 2452	IPv6 MIB for TCP
RFC 2454	IPv6 MIB for UDP

Table 2 IPv6 RFCs Supported (Continued)

RFCs	Description
RFC 2460	IPv6 Specification
RFC 2461	Neighbor Discovery for IPv6
RFC 2462	IPv6 Stateless Address Autoconfiguration
RFC 2463	ICMPv6 for IPv6 Specification
RFC 2464	Transmission of IPv6 Packets over Ethernet Networks
RFC 2465	MIB for IPv6: Textual Conventions and General Group
RFC 2466	MIB for IPv6: ICMPv6 Group
RFC 2553	Basic Socket Interface Extensions for IPv6
RFC 2710	Multicast Listener Discovery (MLD) for IPv6 (Host Part)
RFC 3056	Connection of IPv6 Domains via IPv4 Clouds (6to4)

- * Advanced Socket API features from RFC 2292bis, such as Routing Header, Hop-by-Hop, and Destination Option processing are also supported. The `inet6_rth_*()` and `inet6_opt_*()` functions provide these features. The `inet6_rth_*()` and `inet6_opt_*()` functions are in `/usr/lib/libipv6.1`.

HP Welcomes Your Comments

HP welcomes your comments concerning this document. HP is committed to providing documentation that meets your needs.

Please send comments to: netinfo_feedback@cup.hp.com

Please include document title, manufacturing part number, and any comment, error found, or suggestion for improvement you have concerning this document. Also, please include what you like, so we can incorporate it into other documents.

1 Features Overview

This chapter provides an overview of TOUR 1.0 and a features summary of HP-UX 11i IPv6 transport for TOUR 1.0.

TOUR 1.0 Overview

Typically, transport functionality is delivered as part of the HP-UX OS. TOUR (Transport Optional Upgrade Release) is an additional, fully supported product bundle, available from the web, to deliver transport functionality and other networking functionality independently from the major HP-UX releases. All features of a TOUR release will be rolled into a later HP-UX release. New versions of TOUR will completely supersede previous TOUR versions.

TOUR 1.0 is a web release that enables HP-UX 11i customers to obtain additional Transport related functionality on HP-UX 11i. It is available from `www.software.hp.com`. TOUR 1.0 is only supported on HP-UX 11i (also referred to as HP-UX 11i v1).

TOUR 1.0 uses an overlay model to install new features. It overlays all the delivered files, libraries, commands, and scripts in both user space and kernel space. TOUR 1.0 includes the following components:

- IPv4 Enhancements and IPv6 Transport
- Point-to-Point Protocol (PPP) enhancements:PPPoE and PPPv6
- Router Advertisement Daemon

NOTE IPv4 transport enhancements delivered by TOUR 1.0 supersede base HP-UX 11i OS IPv4 transport functionality; IPv6 delivered by TOUR 1.0 supersedes the intitial HP-UX 11i IPv6 product offering, the stand-alone product bundle, IPv6NCF11i, if it had been installed.

IPv6 Transport (bundled as part of TOUR 1.0)

IPv6 Overview: IPv6 is the next generation Internet Protocol. The IPv6 protocol is also referred to as "IPng" (IP next generation). It provides the infrastructure for the next wave of Internet devices, such as PDAs, mobile phones and appliances. It also provides greater connectivity for existing devices such as laptop computers.

IPv6 was designed by the Internet Engineering Task Force (IETF) to improve upon the scalability, security, ease of configuration, and network management capabilities of IPv4. HP-UX 11i IPv6 network transport software provides host support for IPv6.

IPv6 bundled as part of TOUR 1.0: The following sections highlight IPv6 features available as part of TOUR 1.0. Some of these IPv6 features are new for TOUR 1.0 and will be highlighted as such, the rest have been offered since HP first made available a web release HP-UX 11i IPv6 product.

NOTE IPv6 functionality is not included in the default (base) version of HP-UX 11i. It is only available as a separate add-on product, from the web, as part of TOUR 1.0.

HP-UX 11i IPv6 Transport Features New with TOUR 1.0

Starting with TOUR 1.0, the following IPv6 transport features are available:

- Host MLD Support: The host part of Multicast Listener Discovery (MLD) protocol for IPv6 based on RFC 2710 “Multicast Listener Discovery (MLD) for IPv6”, is supported. MLD is automatically enabled when an IPv6 interface is initialized. The Management Information Base for MLD, based on RFC 3019, is also supported.

  RFC 2710 specifies the protocol used by an IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically, which multicast addresses are of interest to those

neighboring nodes. This protocol is referred to as Multicast Listener Discovery or MLD. MLD is derived from version 2 of IPv4's InternetGroup Management Protocol, IGMPv2. One important difference to note is that MLD uses ICMPv6 (IP Protocol 58) message types, rather than IGMP (IP Protocol 2) message types.

For more MLD information refer to RFC 2710, "Multicast Listener Discovery (MLD) for IPv6".

- Router Advertisement: Router Functionality as specified in RFC 2461 "Neighbor Discovery for IP Version 6 (IPv6)", is implemented with a new daemon, `rtradvd`, and an accompanying configuration file, `/etc/rtradvd.conf`. The `rtradvd` daemon listens to router solicitation and sends router advertisement messages on demand or periodically (as described in RFC 2461). These advertisements allow any listening host to configure their addresses and some other parameters automatically without manual intervention. They can also choose a default router based on these advertisements.

  Router advertisement is configured on a per interface basis. Refer to the `rtradvd.conf(4)` man page for more information.

- FDDI Link Support.
- MC/SG Enablement.

Additional HP-UX 11i IPv6 Transport Features

The following IPv6 transport features are also available as part of TOUR 1.0 (and have been since the first HP-UX 11i IPv6 product was available):

- IPv4/IPv6 Dual Stack: HP-UX 11i IPv6 supports both IPv4 and IPv6 applications. Programmers can write IPv6 applications that communicate with both IPv6 and IPv4 peers. Existing IPv4 applications do not need to be modified.
- Tunneling: IPv6 tunneling enables IPv6/IPv4 hosts and routers to connect with other IPv6/IPv4 hosts and routers over the existing IPv4 network. IPv6 tunneling encapsulates IPv6 datagrams within IPv4 packets. The encapsulated packets travel across an IPv4 infrastructure until they reach their destination host or router. The IPv6-aware host or router decapsulates the IPv6 datagrams,

forwarding them as needed. IPv6 tunneling eases IPv6 deployment by maintaining compatibility with the large existing base of IPv4 hosts and routers.

- Full Ethernet Link Support.
- IPv6 Stateless Address Autoconfiguration.
- IPv6 Neighbor Discovery (which includes Router Discovery and Duplicate Address Detection).
- TCP/UDP over IPv6, PMTUv6, ICMPv6, IPv6 MIBs and Sockets APIs.
- Network Configuration and Troubleshooting Utilities for both IPv4 and IPv6: `ifconfig`, `netstat`, `ping`, `route`, `ndd`, `ndp` (neighbor-discovery command for IPv6 only) and `traceroute`. There have also been enhancements to `nettl` and `netfmt` for IPv6 tracing and formatting.
- New `netconf-ipv6` file stores IPv6 settings. The `/etc/rc.config.d/netconf-ipv6` configuration file stores IPv6 configuration information similar to IPv4's `/etc/rc.config.d/netconf` file.
- The `/etc/hosts` file now supports IPv4 and IPv6 addresses. The `/etc/hosts` file contains IP addresses and corresponding host names. The file can contain IPv4 and IPv6 addresses for the same host. Lookup policies are identical to IPv4.
- Name Service Switch: `/etc/nsswitch.conf` is a configuration file for the name service switch. A new entity, `ipnodes`, specifies which name services resolve IPv6 addresses and host names. Refer to the `nsswitch.conf(4)` man page for more information.

IMPORTANT

If you have an `/etc/nsswitch.conf` file, and if you have set a `hosts` entry in that file, then you must set a similar `ipnodes` entry. You must do this even if you do not use IPv6 addresses. Also, be aware that you cannot use `NIS` or `NIS+` on the `ipnodes` entry.

- The following IPv6-capable Internet Services are included with HP-UX 11i IPv6 bundled as part of TOUR 1.0:
  - `inetd`, Internet Daemon

— `telnet`

— r* commands

— name and address resolution resolver routine

— `inetd.sec` over IPv6

The following Internet Services have been IPv6-enhanced (but are not included as part of TOUR 1.0):

— WU-FTPD 2.6.1

— Sendmail 8.11.1

— BIND 9.2.0

These IPv6-enhanced products are available separately at *www.software.hp.com*. Download these web releases if you want to use FTP, Sendmail, or DNS server to handle IPv6 addresses.

IPv6 Bundled as Part of TOUR 1.0 Replaces “HP-UX 11i IPv6” Product Bundle (IPv6NCF11i)

Prior to the release of TOUR 1.0, HP made IPv6 available as “HP-UX 11i IPv6”, a product bundle downloadable from the web. HP-UX 11i IPv6 bundled as part of TOUR 1.0 supersedes this initial HP-UX 11i IPv6 product offering (IPv6NCF11i).

To obtain IPv6 functionality on your HP-UX 11i system, you need not install this earlier version (IPv6NCF11i), just install TOUR 1.0.

If you have already installed this earlier IPv6 product, you can install TOUR 1.0 on your system and the IPv6 bundled as part of TOUR 1.0 will supersede it. (In this case a `swlist` output will no longer display `IPv6NCF11i` since this product has been superseded by TOUR 1.0.)

2 Installation

This chapter briefly describes how to install HP-UX 11i IPv6 bundled as part of TOUR 1.0 on your system. This chapter discusses the following topics:

- Installation of TOUR 1.0
- Verification of the TOUR 1.0 installation
- Verification of IPv6 Loopback

NOTE The entire TOUR 1.0 bundle must be installed to obtain HP-UX 11i IPv6 as part of TOUR 1.0. In addition to HP-UX 11i IPv6, the TOUR 1.0 bundle also contains: IPv4 transport enhancements and PPP enhancements (PPPoE and PPPv6). For more information on TOUR 1.0 components, features and installation, refer to the *TOUR 1.0 Release Notes*.

If you have already completed installation of TOUR 1.0, you may skip this chapter and proceed to Chapter 3, “Configuration,” on page 13.

Installing HP-UX 11i IPv6 (bundled as part of TOUR 1.0)

HP-UX 11i IPv6 is a component of TOUR 1.0. The complete TOUR 1.0 product bundle must be installed. You can download the TOUR 1.0 depot from *www.software.hp.com* (Keyword “TOUR”).

You can follow the Installation instructions provided on the web. These steps are also summarized below. For additional information, refer to the *TOUR 1.0 Release Notes*.

System Requirements

TOUR 1.0 must be installed on an HP-UX 11i system. You can install it on a 32-bit or 64-bit environment. It requires approximately 8-9 Mbytes of free disk space. (~7 Mbytes is used in `/usr`, ~1 Mbyte is used in `/var`.) The approximate file size of the TOUR 1.0 depot file is less than 4.5 Mbytes. (Note that this disk space information does not include disk space requirements for prerequisite patches.)

Installation Instructions

To install TOUR 1.0 on HP-UX 11i, complete the following steps:

1. Install the following patches:

NOTE A few patches are kernel patches. Therefore, these patches require a system reboot.

- Install the latest (March 2003 - B.11.11.0212.4 or later) HP-UX 11i Quality Pack (GOLDQPK11i) patch bundle.
- In addition, you must install the following patches (or later):
  - PHCO_24287 - “patch for the `syslogd(1M)` command”
  - PHNE_28476 - “cumulative STREAMS patch”
  - PHNE_27796 - “libnss_dns DNS backend patch”

— PHNE_28895 - “cumulative ARPA Transport patch”

You can retrieve individual patches from *www.itrc.com*.

NOTE: Future versions of Quality Pack (GOLDQPK11i) may include these four patches. In this case, you will not have to retrieve the patches individually.

2. Complete the registration form at *www.software.hp.com* with appropriate information.
3. Download the TOUR 1.0 software depot to a directory, for example, `/tmp/TOUR.depot`.
4. Use swinstall to install TOUR 1.0. Execute the following at the command prompt:

   ```
   swinstall -x autoreboot=true -s /tmp/TOUR.depot TOUR
   ```

   The system reboots automatically if the `swinstall` execution completes without errors.

NOTE After installing TOUR 1.0, you will no longer be able to remove or install additional Transport patches. If you need to add or remove a Transport patch, you will first need to remove TOUR 1.0. Separate patches for TOUR 1.0 (distinct from Transport patches) will be available, as needed.

NOTE TOUR 1.0 is not compatible with any version of IPSec. To use both IPv6 and IPSec functionality, you must install the initial HP-UX 11i IPv6 product offering, the stand-alone product bundle IPv6NCF11i. This is available on the Software Pack (Optional HP-UX 11i Core Enhancements) CD or from your HP representative.

Verifying Your TOUR 1.0 Installation

To verify your TOUR 1.0 installation, execute the following at the command prompt:

```
swlist -l bundle TOUR
```

You should see the following:

```
TOUR A.01.00 Transport Optional Upgrade Release for B.11.11
```

To verify that there were no errors during the TOUR1.0 installation, execute the following at the command prompt:

```
swverify TOUR
```

If TOUR 1.0 installed successfully, the following message is displayed:

```
Verification succeeded
```

Verifying IPv6 Loopback

After installing TOUR 1.0, to test IPv6 internal loopback capability execute the following at the command prompt:

```
ping -f inet6 ::1 -n 2
```

You should see output similar to the following:

```
# ping -f inet6 ::1 -n 2
PING ::1: 64 byte packets
64 bytes from ::1: icmp_seq=0. time=2. ms
64 bytes from ::1: icmp_seq=1. time=0. ms
----::1 PING Statistics----
2 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/2
```

Refer to later chapters for more detailed configuration information.

3 Configuration

This chapter summarizes the steps for configuring LAN interfaces, assigning IPv6 addresses, optionally enabling IPv6 tunneling through IPv4 networks, and assigning host names to IPv6 addresses.

The first interface configured on a physical LAN interface is called the **primary interface**. Additional interfaces configured on the same physical device are called **secondary interfaces**. You must configure an IPv6 primary interface to use IPv6 over that interface.

Configuring IPv6 Interfaces and Addresses

This section describes IPv6 interface and address configuration tasks that involve editing the `/etc/rc.config.d/netconf-ipv6` file.

Before configuring IPv6 interfaces, remember:

- To edit the `netconf-ipv6` file and to activate the configuration, you must be superuser.
- The `netconf-ipv6` file and the script that is executed are shell programs; therefore, shell programming rules apply.
- To activate the `netconf-ipv6` configuration, you must either reboot the system or use `ifconfig`, `route` commands with appropriate equivalent values. (Note: `ifconfig` and `route` configuration changes are ephemeral and do not permeate across reboots.) Refer to the “Activating netconf-ipv6 file Configuration” on page 20 for more information.
- To configure HP-UX IPv6, bundled as part of TOUR 1.0, you may use SAM. You will need to download the latest HP-UX 11i SAM patch (`PHCO_24118` or later.) Open SAM, proceed to the “Networking and Communications” area and navigate (proceed) as needed for your configuration tasks (Network Interface Cards, Name Service Switch, Hosts).

NOTE `setparms` has not been enhanced to support IPv6 configuration.

Configure IPv6 interfaces and routing using one of the following methods:

- Stateless autoconfiguration
- Manual configuration

These methods are described in more detail below.

Stateless Autoconfiguration

Addresses on IPv6 interfaces, unlike IPv4 interfaces, can be configured without manual intervention. With stateless address autoconfiguration, the primary interface (`lanX:0`) is automatically assigned a link-local IPv6 address by the system when the interface is configured (marked “up”). This link-local IPv6 address is generated by prepending a fixed local address prefix (`fe80::`) to a token derived from the MAC address. (The address is verified to be unique.) This allows each IPv6 interface to have at least one source address that can be used by Neighbor Discovery.

If an IPv6 router on the network advertises network prefixes in router advertisements, IPv6 derives secondary IPv6 addresses based on the network interface identifier of the primary interface and on the network prefixes advertised. IPv6 assigns this address to a secondary interface for the network interface.

Refer to “Stateless Address Autoconfiguration” on page 56 in Chapter 6 of this guide, and the `ifconfig(1M)` man page for more information.

Configuring a Primary Interface

To configure a primary interface, edit the `IPV6_INTERFACE[0]` statement in the `/etc/rc.config.d/netconf-ipv6` file to specify the interface name, such as lan0. The interface name must be the name of the physical interface card, as reported by `lanscan`.

A sample `netconf-ipv6` file entry is as follows:

```
IPV6_INTERFACE[0]="lan0"
IPV6_INTERFACE_STATE[0]="up"
```

Again, in the above example, the address is automatically assigned. Note that autoconfiguration is not mandatory, manual specification of the address is also allowed and is described below.

Configuring Secondary Interfaces

If an IPv6 router that advertises network prefixes resides on the LAN, a secondary interface is automatically configured after the primary interface comes up. IPv6 builds additional secondary interfaces for each network prefix advertised.

If you manually configure a link-local address for the primary interface, then autoconfigured secondary addresses are derived from the interface identifier part of the manually configured address for the primary interface.

For example, if an IPv6 router on the LAN advertises two prefixes (such as `fec0::/64` and `2000::/64`), HP-UX 11i IPv6, bundled as part of TOUR 1.0, configures two secondary interfaces.

Configuring Route Information

HP-UX 11i IPv6, bundled as part of TOUR 1.0, automatically configures network routes based on the prefix information received from an IPv6 router. HP-UX 11i IPv6, bundled as part of TOUR 1.0, automatically adds the router to its list of default gateways if the router advertises a non-zero router-lifetime value.

Manual Configuration

The following section describes manual configuration processes for HP-UX 11i IPv6, bundled as part of TOUR 1.0.

Configuring a Primary Interface

To configure an IPv6 link-local address for a primary interface, edit the `IPV6_INTERFACE[0]` statement in the `/etc/rc.config.d/netconf-ipv6` file to specify the interface name (for example, lan0) and the interface state, either `up` or `down`. The interface name must be the name of the physical interface card, as reported by `lanscan`.

To manually specify a link-local address for the primary interface, note that the universal/local "U" bit must be set to 0. That implies, the manually configured address for the primary interface must match the pattern FE80::xMxx:xxxx:xxxx:xxxx where x are hexadecimal digits, and M is either 0, 1, 4, 5, 8, 9, C, or D. (To be more specific, break M down to the bit level and thus, M = yy0y, where y can be 0 or 1.)

A sample `netconf-ipv6` file entry is as follows:

```
IPV6_INTERFACE[0]="lan0"
IPV6_INTERFACE_STATE[0]="up"
IPV6_LINK_LOCAL_ADDRESS[0]= "fe80::1"
```

Note that if you do not specify a link-local address, then as described earlier in the autoconfiguration section, a link-local address is automatically configured for the primary interface based on the interface's 48-bit MAC address.

Configuring Secondary Interfaces

If no IPv6 Router on the LAN advertises network prefixes, you can add secondary interface entries to the `/etc/rc.config.d/netconf-ipv6` file. Editing the `netconf-ipv6` file allows you to identify the network interface name, IPv6 address, and prefix length and also to add entries to the network routing table.

A sample `netconf-ipv6` file entry is as follows:

```
IPV6_SECONDARY_INTERFACE_NAME[1]="lan0:1"
IPV6_ADDRESS[1]="2345::5432"
IPV6_PREFIXLEN[1]="64"
IPV6_SECONDARY_INTERFACE_STATE[1]="up"
DHCPV6_ENABLE[1]=0
```

Always set DHCPV6_ENABLE to 0.

For more information about specifying interface names for multiple interfaces, refer to Chapter 6, “IPv6 Addressing and Concepts,” on page 51.

Configuring a Default IPv6 Route

In the absence of router advertisements, you can add the default IPv6 router information to the /etc/rc.config.d/netconf-ipv6 file. The routing configuration parameters have an index value, [x], that groups the routing parameters together.

A sample netconf-ipv6 file entry is as follows:

```
IPV6_DESTINATION[0]="default"
IPV6_GATEWAY[0]="2008:7:6:5:4:3:2:1"
IPV6_ROUTE_COUNT[0]="1"
IPV6_ROUTE_ARGS[0]=""
```

Enabling Tunneling

To enable tunneling set the IPV6_TUNNEL variable to 1.

A sample netconf-ipv6 file entry is as follows:

```
IPV6_TUNNEL="1"
```

Refer to the section on “Tunneling” on page 68 in Chapter 7 of this guide for more information.

Creating a Configured Tunnel

If you regularly expect to exchange data between isolated IPv6 networks over an IPv4 network, you may want to create a configured tunnel. Do the following:

- Enable tunneling. Refer to “Enabling Tunneling” section above. (Basically, set IPV6_TUNNEL variable to 1.)

- Assign the following parameters: destination IPv6 network, gateway address (for configured tunnels this is an IPv4-compatible IPv6 address), route_count (if set to 1 the gateway is remote, if set to 0 the gateway is local), and any `route` arguments (for configured tunnels, specify the `-t` option). Refer to the commented text in the `netconf-ipv6` file and the `route(1m)` man page for more information.

  A sample `netconf-ipv6` file entry is as follows:

```
IPV6_DESTINATION[0]="C001::/64"
IPV6_GATEWAY[0]="::192.1.1.1"
IPV6_ROUTE_COUNT[0]="1"
IPV6_ROUTE_ARGS[0]="-t"
```

Enabling rtradvd (Router Advertiser Daemon)

Background: The `rtradvd` daemon, is now available on HP-UX 11i IPv6 bundled as part of TOUR 1.0. When configured, it sends router advertisement messages to a local LAN periodically, and, when requested, by a node sending a router solicitation message. For more information refer to the `rtradvd(1M)` man page.

Configuration for `rtradvd` is set, on a per interface basis, by editing the `/etc/rtradvd.conf` file. The `rtradvd.conf` file allows for setting global defaults as well as interface specific settings for both interface options and prefixinfo specific options. Refer to the `rtradvd.conf(4)` man page for more information.

Required Steps: To configure the HP-UX system to run `rtradvd`, and enable the Router Advertisement functionality, the following steps must be taken: the `/etc/rtradvd.conf` file must be edited as needed; the "private" interface flag must be cleared ("-private" for each enabled interface) and the `rtradvd` daemon must be enabled. More specifically:

- Edit the `/etc/rtradvd.conf` file as needed.

  The example below shows the minimum configuration needed to send router advertisement packets containing the prefix `2008:65::/64` on `lan0`.

```
#example begins
defaults {
        AdvSendAdvertisement on ;
 };

  interface lan0 {
        prefixinfo 2008:65::/64 {
        };
 };

#example ends
```

For more examples, refer to the `rtradvd.conf(4)` man page.

- Edit the `/etc/rc.config.d/netconf-ipv6` file to enable `rtradvd` to start up at boot. Also, clear the "private" interface flag (`-private`), on the appropriate interface(s) to enable forwarding. [Note: when the "private" flag is set (`private`), which it is by default, forwarding on the interface is disabled.] For more information, refer to the relevant commented text in the `/etc/rc.config.d/netconf-ipv6` file that is included with HP-UX 11i IPv6 bundled as part of TOUR 1.0.

  A sample `netconf-ipv6` file entry, which clears the private flag and enables `rtradvd`, is as follows:

```
IPV6_INTERFACE[0]="lan0"
IPV6_INTERFACE_STATE[0]="up"
IPV6_INTERFACE_FLAG[0]= "-private"
#
#
RTRADVD=1
```

Activating netconf-ipv6 file Configuration

You can activate the `netconf-ipv6` configuration in one of the following ways:

- By rebooting the system.

  HP recommends rebooting your system to activate any changes you

made to your netconf-ipv6 file. A reboot is the cleanest way to reconfigure an interface because the reboot handles any network initialization dependencies.

- Or alternatively, by executing the ifconfig and route commands, as needed, to make equivalent configuration settings.

  HP recognizes that system reboots are disruptive to end users. To delay or schedule the reboot, but still make your configuration changes active, you may execute the ifconfig and route commands with the appropriate values for your network. These values are ephemeral however, and will not last across reboots. After the reboot, the values in your netconf-ipv6 file will be used. Refer to the examples that follow and the ifconfig(1M), and route(1M) man pages for more information on using these commands.

Example ifconfig and route Commands

HP recommends editing the /etc/rc.config.d/netconf-ipv6 file to preserve IPv6 interface and address configurations across system reboots. For reference, the commands equivalent to the netconf-ipv6 edits described earlier are listed below. Refer to the ifconfig(1M) and route(1M) man pages for more information.

To configure a primary interface, enter:

```
ifconfig lan0 inet6 up
```

To configure a secondary interface, enter:

```
ifconfig lan0:1 inet6 2345::5432 up
```

To add a default IPv6 route, enter:

```
route inet6 add net default 2008:7:6:5:4:3:2:1
```

To enable tunneling, enter:

```
ifconfig tu0 inet6 up
```

To create a configured tunnel, enter:

```
route inet6 -t add c001::/64 ::192.1.1.1l
```

NOTE Remember that configuration using `ifconfig` and `route` is ephemeral, and not maintained after a system reboot.

Host Names and IPv6 Addresses

The following section provides additional information on how addressing works on HP-UX 11i IPv6 bundled as part of TOUR 1.0.

Creating the /etc/hosts File

It is generally recommended to add IPv6 addresses (known as AAAA records) to a DNS Name Server only when the following conditions are true:

- The IPv6 address is assigned to the interface on the node
- The address is configured on the interface
- The interface is on a link which connects to the IPv6 infrastructure

HP recommends beginning with IPv6 addresses and host names in the `/etc/hosts` file on a development network; then adding IPv6 addresses and hosts to a Domain Name Service when moving IPv6 to a production backbone network.

This subsection describes how to edit the `/etc/hosts` file to add an IPv6 address and host name for the network interface you are configuring.

NOTE If using the name service DNS over IPv6, add the IP address and host name to the appropriate databases on the name server system. Refer to BIND v9.2.0 (or later) documentation on `docs.hp.com` for more information on DNS over IPv6.

The `/etc/hosts` file associates IP host addresses with mnemonic host names and alias names. It contains the names of other nodes in the network with which your system can communicate.

Example Host Name Entry

The example below shows how a system with the name, `host3`, might be referenced in the `/etc/hosts` file:

System name in `swinstall` screen: `host3`

`/etc/hosts` file:

```
3ffe:ffff:101::230:6eff:fe04:d9ff  host3  host3.site2.region4
192.1.2.34  hpfcrm loghost
```

NOTE HP-UX 11i IPv6 bundled as part of TOUR 1.0 is a dual stack implementation. A single host name can have entries for both an IPv6 address and an IPv4 address in `/etc/hosts`.

Name and Address Lookup for IPv6

`/etc/nsswitch.conf` (`nsswitch.conf(4)`) is a configuration file for the name service switch. A new entity, `ipnodes`, specifies which name services resolve IPv4 and IPv6 addresses and host names on HP-UX 11i IPv6 bundled as part of TOUR 1.0.

More specifically, the new keyword "`ipnodes`" specifies the resolver policy for the new library functions `getnameinfo(3N)`, `getaddrinfo(3N)`, `getipnodebyname(3N)` and `getipnodebyaddr(3N)` for both IPv4 and IPv6 addresses. The existing keyword "`hosts`" specifies the resolver policy for the library functions `gethostbyname()` and `gethostbyaddr()` for IPv4 addresses.

NOTE Internet Services applications (such as `telnet`, `r*` commands, etc.) use these new library functions to resolve IPv4 and IPv6 addresses.

By default, the `/etc/nsswitch.conf` is not on the system. The default `ipnodes` policy (same as default `hosts` policy) is as follows:

```
dns [NOTFOUND=return] files
```

This policy implies that `dns` is the authoritative resolver and will only try `files` if `dns` is down. If `dns` is available but returns `NOTFOUND`, the search stops.

Thus, if DNS has **not** been set up as the definitive source, and files (`/etc/hosts`) may need to be used for address and host name resolution, HP recommends adding the following entry to `/etc/nsswitch.conf`:

```
ipnodes: dns [NOTFOUND=continue] files
```

Or if /etc/hosts is to be the primary Name Service, the entry should be set as follows:

```
ipnodes: files [NOTFOUND=continue] dns
```

IMPORTANT

If you have an /etc/nsswitch.conf file, and if you have set a hosts entry in that file, then you must set a similar ipnodes entry. You must do this even if you do not use IPv6 addresses. Also, be aware that you cannot use NIS or NIS+ on the ipnodes entry.

Manually editing nsswitch.conf

If the current system has no nsswitch.conf file, use a text editor to create an /etc/nsswitch.conf file containing the following line, or copy the /etc/nsswitch.defaults file and modify as needed.

Refer to the nsswitch.conf(4) man page for hosts syntax.

4 Troubleshooting

This chapter provides guidelines for troubleshooting HP-UX 11i IPv6 bundled as part of TOUR 1.0. It contains a troubleshooting overview and diagnostic flowcharts.

Troubleshooting Overview

Troubleshooting problems on HP-UX 11i IPv6, bundled as part of TOUR 1.0, can involve a variety of hardware and software components. The problem impacting your system might originate in another part of the network.

Because HP-UX 11i IPv6, bundled as part of TOUR 1.0, supports an IPv4/IPv6 Dual Stack, test IPv4 connectivity before testing IPv6 connectivity. Refer to the *HP-UX LAN Administrator's Guide* (available on docs.hp.com) for IPv4 troubleshooting advice.

If you are still unable to identify your problem, proceed to the troubleshooting flowcharts. The troubleshooting flowcharts provide logical steps to follow when troubleshooting HP-UX 11i IPv6, bundled as part of TOUR 1.0. Use the diagnostic flowcharts provided in this chapter to verify your assumptions and to try to identify whether the problem is with HP-UX 11i IPv6, bundled as part of TOUR 1.0, or router configuration.

Diagnostic Flowcharts

-Below is a summary of the types of network tests in the diagnostic flowcharts. To diagnose your problem, first check the connections and configuration on your system (Flowcharts 1 through 5). If this does not solve your problem, use flowcharts 6 through 8 to test and/or verify connectivity with a remote system.

-1	Transport Level Test using Internet Services
2	Network Connectivity Test
3	Name Services Test
4	Interface Test
5	Interface Test (continued)
6	Router Remote Loopback Test

Transport Level Loopback Test using Internet Service: Checks round-trip communication between Transport Layers on the source and target host using telnet.

Network Connectivity Test: Checks round-trip communication between Network Layers on the source and target host using the `ping(1M)` diagnostic.

Name Services Test: Checks host name and IPv6 address resolution.

Interface Test: Verifies the configuration of the network interface on a host using the `lanscan`, and `ifconfig` commands.

Router Remote Loopback Test: Verifies the connection between local and remote nodes through IPv6 routers using the `ping` and `netstat` commands.

Flowchart 1: Transport Level Testing using Internet Services

Figure 4-1 Flowchart 1

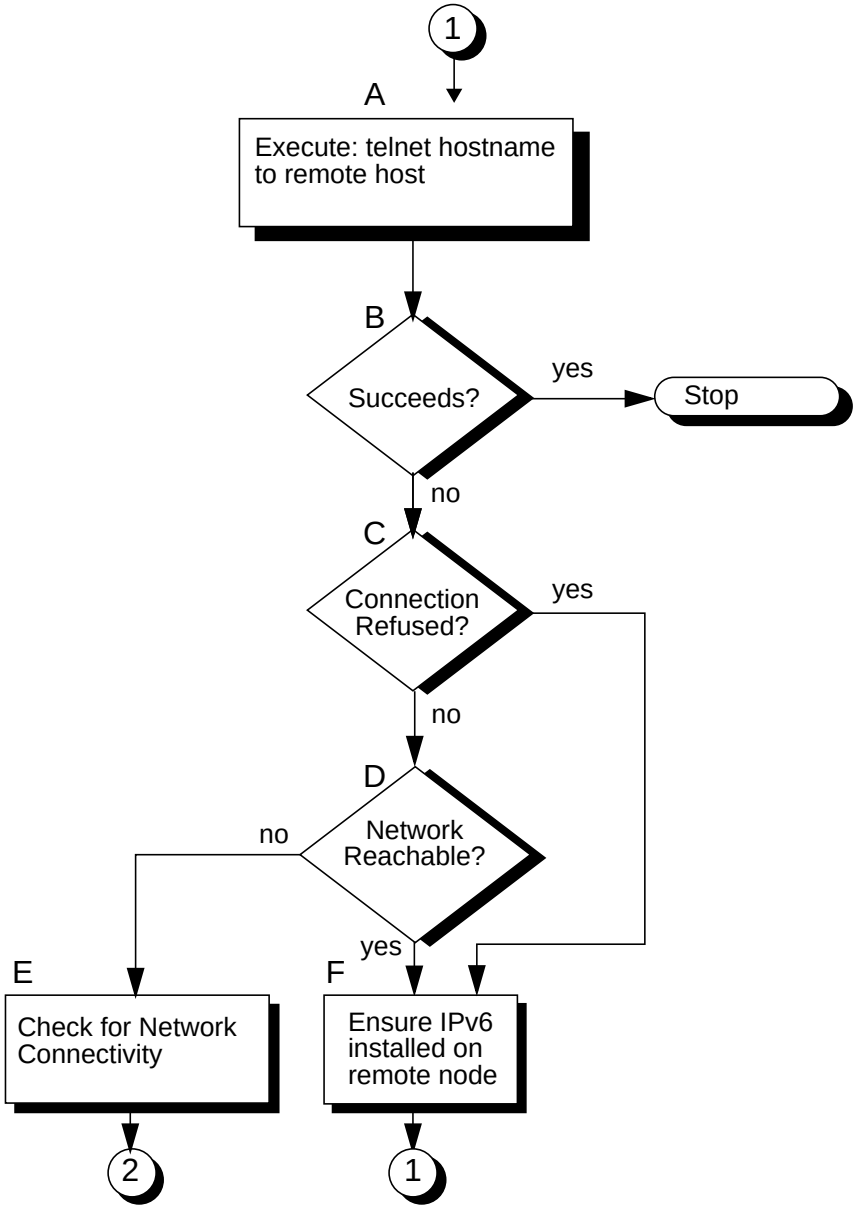

Flowchart 1 Procedures

A. *Execute:* `telnet <hostname> to remote host`. Try to connect using `telnet` to a remote host.

B. *Succeeds?* If telnet succeeds, stop. The system connects using TCP6 over IPv6 through the Transport Layer (OSI Layer 4).

C. *Connection Refused?* Trying to connect to a remote system where HP-UX 11i IPv6 is not installed can cause this message.

D. *Network Reachable?* If it is, go to F, otherwise continue to E.

E. *Check for Network Connectivity.* Ensure network connectivity by following the steps in Flowchart 2.

F. *Ensure IPv6 installed on remote node.* If `telnet` still fails, check the `etc/inetd.conf` file.

Flowchart 2: Network Connectivity Test

Figure 4-2 Flowchart 2

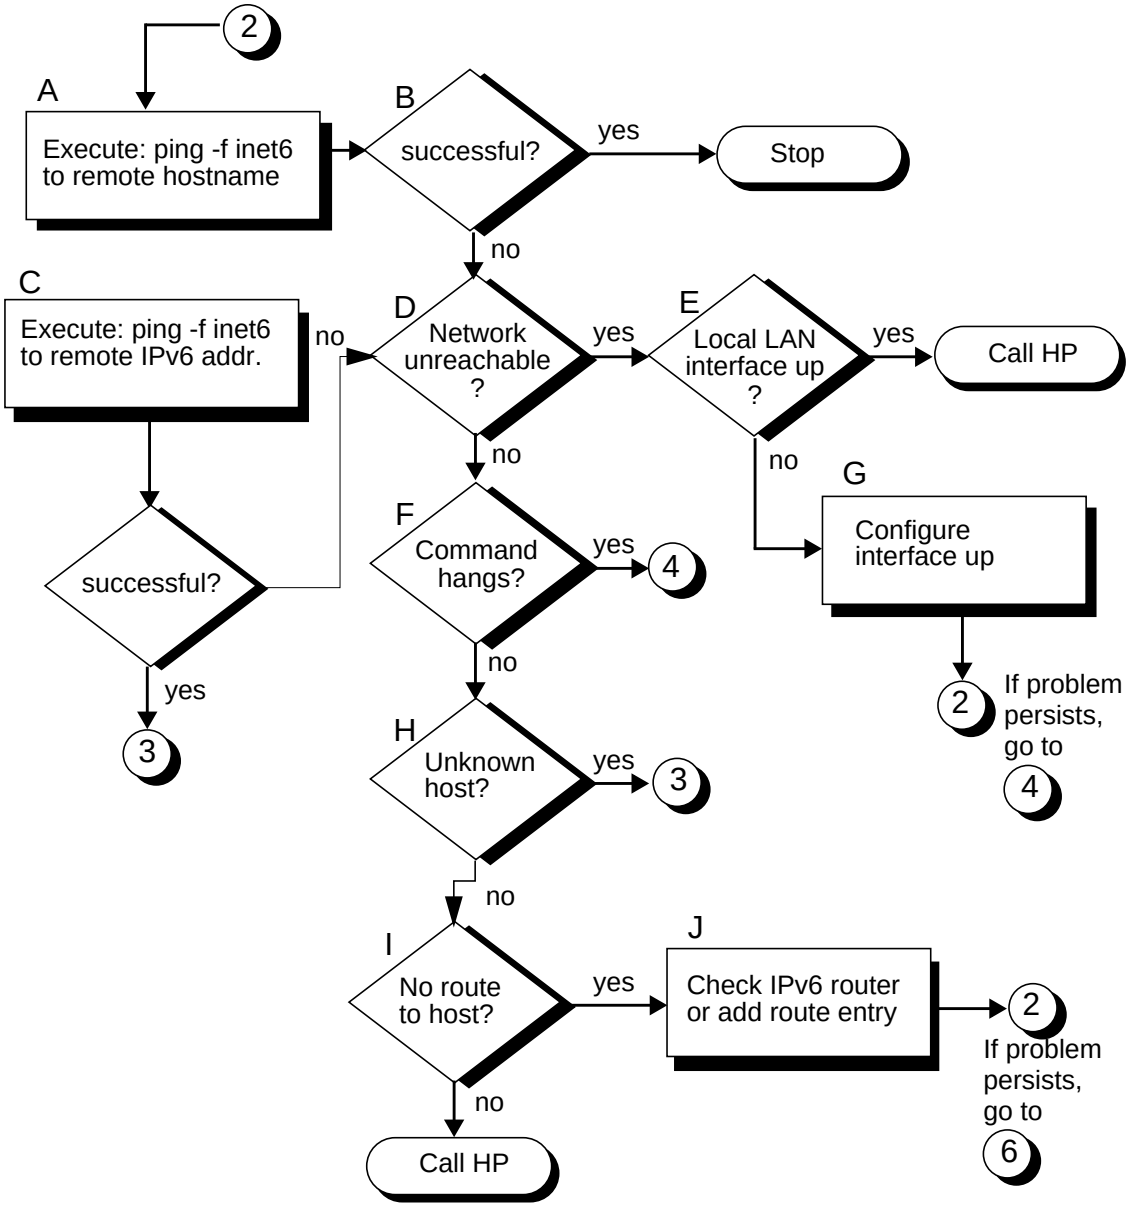

Flowchart 2 Procedures

A. *Execute ping to remote IPv6 "host name".* Using `ping`, send an ICMPv6 message to the remote host with which you are having problems connecting. For example, the remote host name is `hpindon`. Enter:

```
ping -f inet6 hpindon
```

B. *ping successful?* A message is printed on stdout for each ping packet returned by the remote host. If packets are being returned, your system has network level connectivity to the remote host.

C. *Execute ping to remote IPv6 address.* Using `ping`, send a message to the IPv6 address of the remote host. For example,

```
ping -f inet6 8:7:6:5:4:3:2:1
```

D. *Network unreachable?* If so, check the status of the local LAN interface first. If not, proceed to F.

E. *Local LAN interface up?* Execute `ifconfig` on the local interface to be sure it is configured up. If it is not, go to G. If it is up, call your HP representative for help.

F. *Command hangs?* If a message is not returned after executing ping, go to Flowchart 4, otherwise go to H.

G. *Configure interface up.* If you find the local interface is not up, execute `ifconfig` with the appropriate flags set. Begin Flowchart 2 again. If the problem persists, go to Flowchart 4.

H. *Unknown host?* (Error= Unknown host <hostname>?) If so, there is a problem with the IPv6 address configuration for the host <hostname> in the `/etc/hosts` file or on the name server. Go to Flowchart 3. Otherwise, proceed to I.

I. *No route to host?* (Error= Sendto: No route to host?) Use `netstat -rn` to check the routing table. If there is no route to host, go to J. Otherwise, call your HP representative for help.

J. *Check IPv6 Router or add route table entry.* Add a route table entry to that host, or ensure that the IPv6 router advertises correct prefixes. Then try Flowchart 2 again. If the problem persists, go to Flowchart 6.

Flowchart 3: Name Service Test

Figure 4-3 Flowchart 3

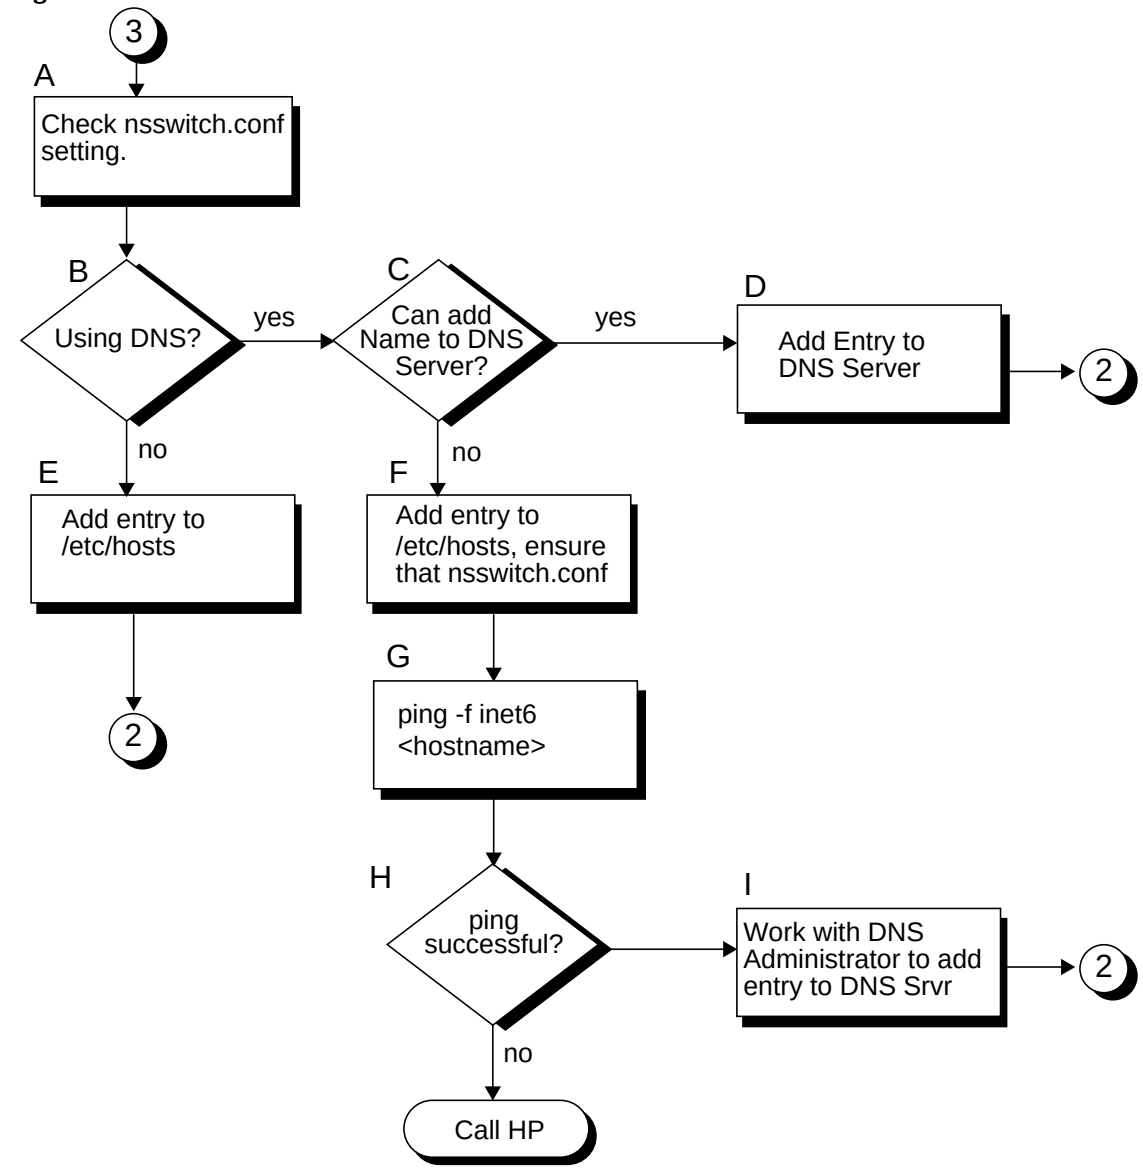

Flowchart 3 Procedures

A. *Check /etc/hosts and /etc/nsswitch.conf files.* If needed, add a missing host name or IPv6 address. If the IPv6 address for the host is in `/etc/hosts`, ensure that you have an `/etc/nsswitch.conf` file entry with an appropriate `ipnodes` policy. For example,
`ipnodes: DNS [NOTFOUND=continue] files`

and start again with Flowchart 3.

B. *Using DNS?* If your name and IPv6 address resolution policy use DNS as the primary resolver, go to C. Otherwise, proceed to E.

C. *Can you add a Host Name to the DNS Server?* Are you a DNS administrator? If so, continue on to D, otherwise proceed to F.

D. *Add Entry to DNS Server.* Refer to the BIND 9.2.0 information in the *HP-UX IP Address and Client Management Administrator's Guide* for details (available at `http://docs.hp.com`). Then retry Flowchart 2.

E. *Add entry to /etc/hosts.* If your name and IPv6 address resolution policy uses `/etc/hosts` as the primary resolver, add a correct IPv6 address and host name to the local `/etc/hosts` file. Then retry Flowchart 2.

F. *Add entry to /etc/hosts and ensure that nsswitch.conf is configured properly.* Add a correct IPv6 address and host name to the local `/etc/hosts` file. Ensure that your IPv6 address resolution policy, specified with the `ipnodes` keyword in `/etc/nsswitch.conf` includes using "files" (`/etc/hosts`) in the policy. Then retry Flowchart 2.

G. *ping -f inet6 hostname.* Test connectivity to the remote host using the `ping` command.

H. *ping successful?* If `ping -f inet6 <hostname>` succeeds using a host name and IPv6 address from `/etc/hosts`, DNS needs updating, proceed to I. If `ping` fails, check the `/etc/hosts`, `/etc/resolv.conf`, and `/etc/nsswitch.conf` files on both the local and remote hosts. If all look correct, call your HP representative for help.

I. *Work with DNS Administrator to add entry to DNS Server.* When entry is added, retry Flowchart 2 to ensure that DNS correctly resolves host names and IPv6 addresses.

Flowchart 4: Interface Test

Figure 4-4 Flowchart 4

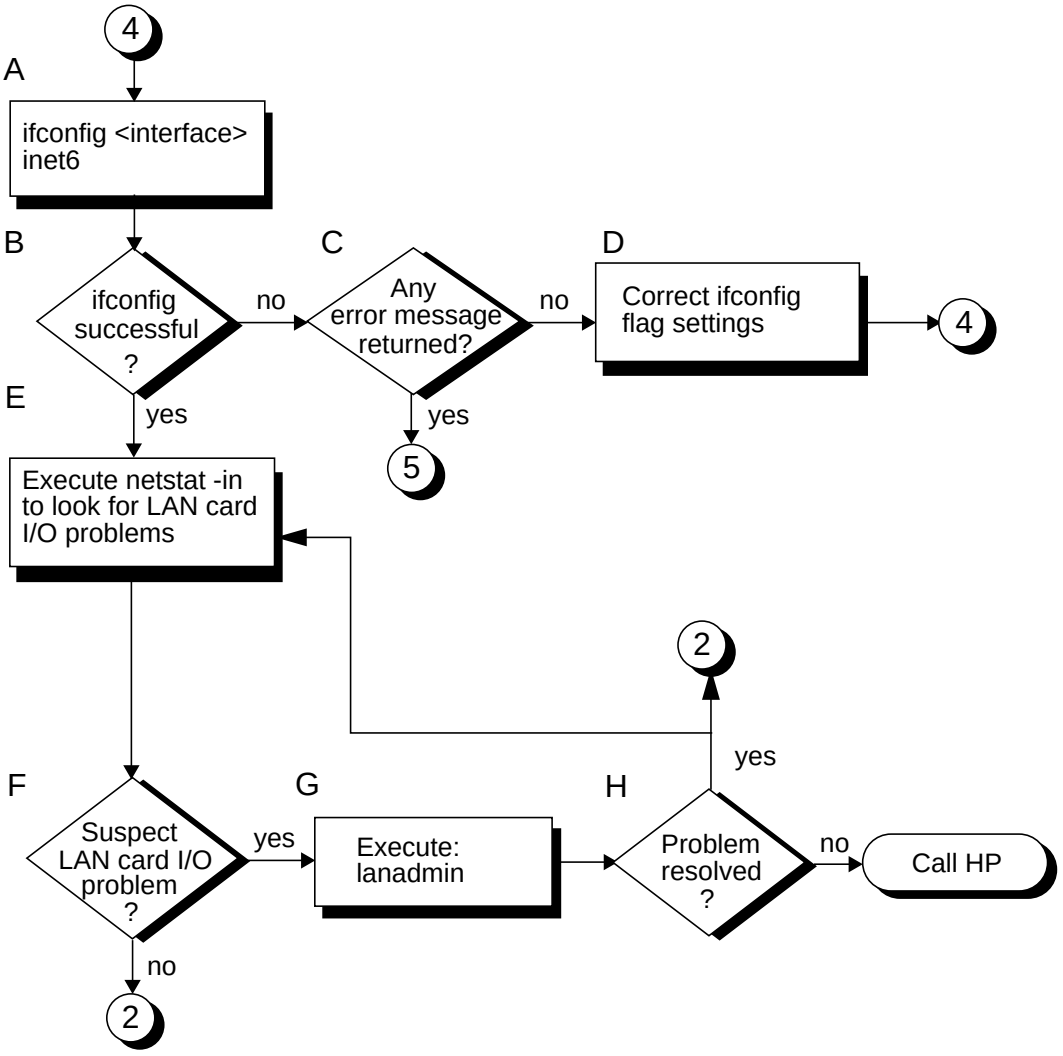

Flowchart 4 Procedures

A. *Execute: ifconfig <interface>inet6.* Execute `ifconfig` on the interface you want to test. For example, to check LAN interface `lan0`, enter:

```
ifconfig lan0 inet6
```

B. *ifconfig successful?* ifconfig succeeds when the output shows an Internet address and the flags: UP, RUNNING, MULTICAST, ONLINK. If successful, go to E, if not continue to C.

C. *Any error message returned?* If `ifconfig` fails and displays an error message, go to Flowchart 5. Flowchart 5 shows what to do based on the error message. Otherwise continue to D.

D. *Correct `ifconfig` with non-default flag settings.* If `ifconfig` returns an unexpected flag setting, re-execute the command with the proper setting. For more information, refer to the `ifconfig(1M)` man page. Start again with Flowchart 4.

E. *Execute: netstat -inf -inet6.* If `ifconfig` succeeds, then the network interface is configured correctly. `netstat -i` displays the number of incoming (`Ipkts`) and outgoing (`Opkts`) packets passed through an interface. No increase in the number of incoming or outgoing packets could indicate LAN card I/0 problems.

F. *Suspect LAN card I/O problems?* If the statistics indicate possible LAN card problems, go to G, otherwise go to Flowchart 2 to test Network Connectivity.

G. *Execute: lanadmin.* Use `lanadmin` to ensure the LAN card is operational. A substantial increase in the number of the `Ierrs` and `Oerrs` during a file transfer attempt might indicate transmission problems.

H. *Problem resolved?* If you found and corrected the LAN card problem, return to step E to verify the correction. If corrected, re-execute `ifconfig` to bring up the interface, then go to Flowchart 2. If the problem persists, call your HP representative for help.

Flowchart 5: Interface Test continued

Figure 4-5 Flowchart 5

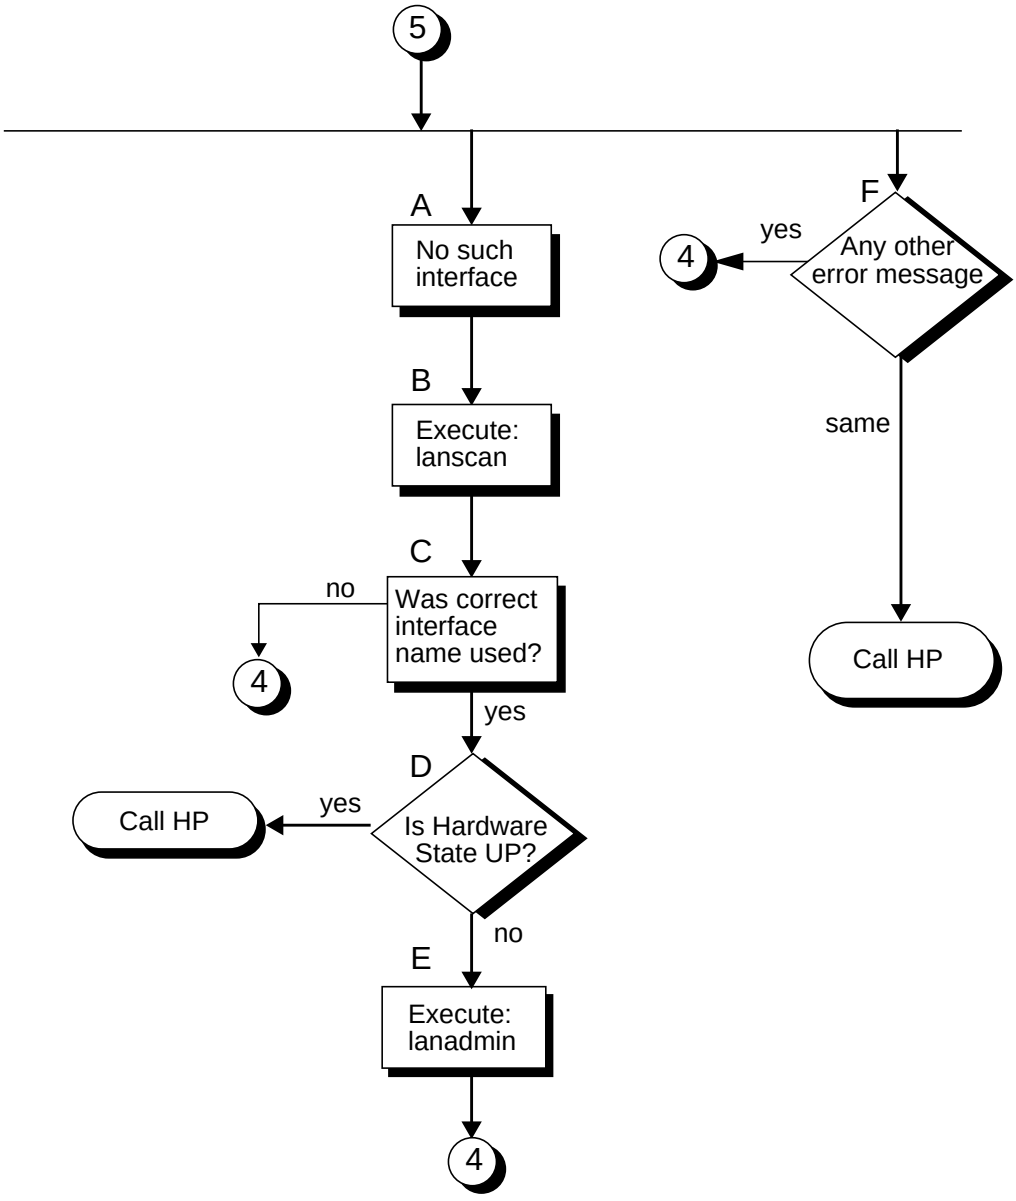

Flowchart 5 Procedures

A. *Is error message "No such interface name"?* If not, go to F. If so, the interface name passed to `ifconfig` does not exist on the system. Using `lanscan`, check the spelling and names of the interfaces on the system.

If the system contains more than one LAN card, make sure the correct number of LAN cards was configured into the kernel and that an `ifconfig` command was executed for each interface.

B. *Execute: lanscan.* Execute `lanscan` to display information about the LAN cards in your system.

C. *Was correct interface name used?* Configure interface using `ifconfig` with the correct interface name. After reconfiguring using the correct interface name, start again with Flowchart 4.

D. *Is Hardware State UP?* Verify the state of the hardware with the output from the `lanscan` command. If the Hardware State is UP call your HP representative for help, otherwise continue to E.

E. *Execute: lanadmin.* Use the `lanadmin` command to reset the LAN card. Go to Flowchart 4.

F. *Any other error message.* Interpret any other error message and take the appropriate action. Then repeat flowchart 4. If you receive the same error message again, call your HP representative for help.

Flowchart 6:Router Remote Loopback Test

Figure 4-6 Flowchart 6

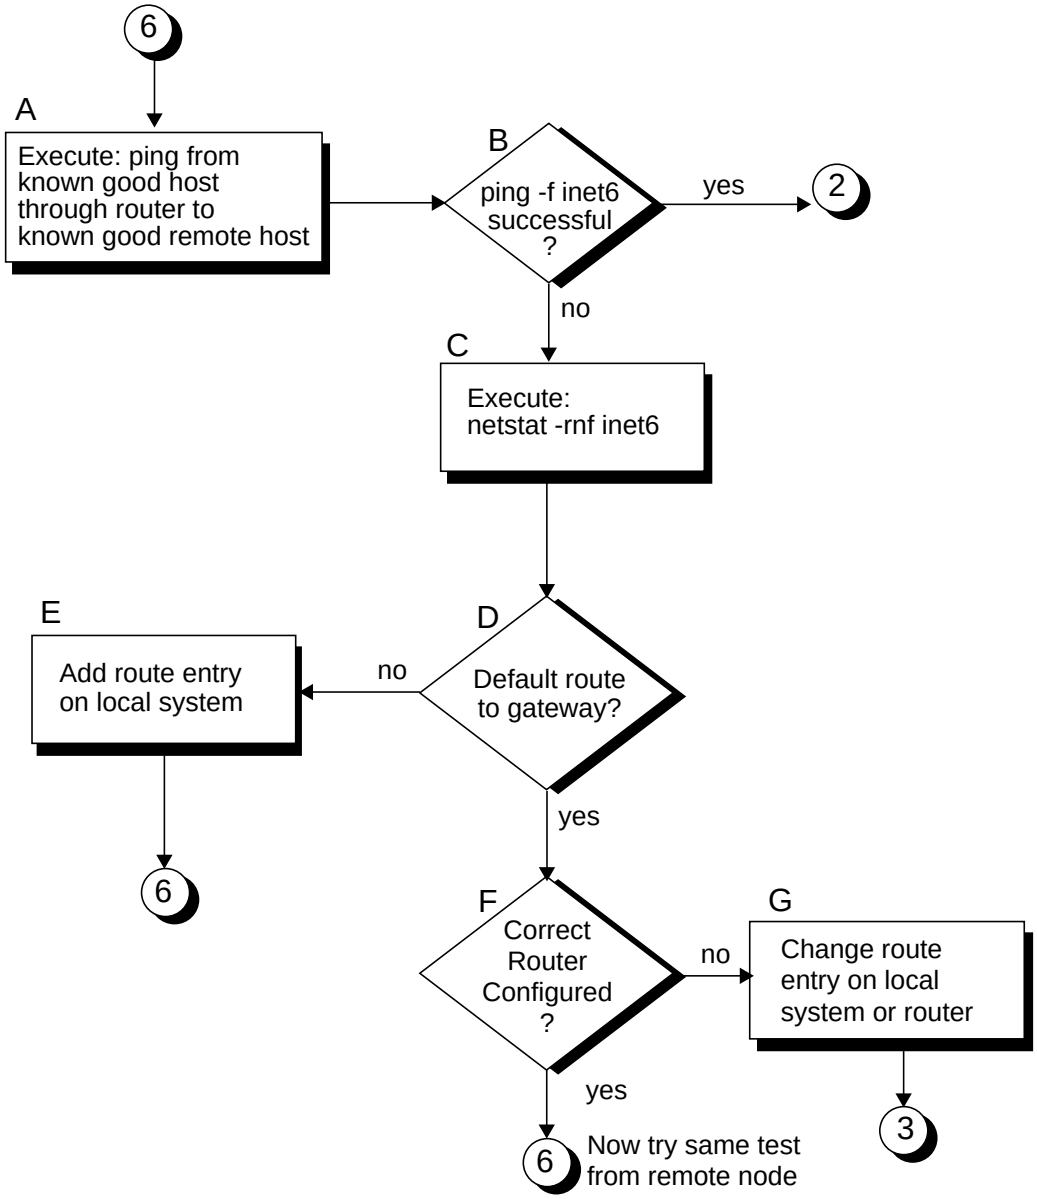

Flowchart 6 Procedures

A. *Execute: ping from known good host through gateway to known good host on remote network*. This tests router connectivity to the remote network. For more information on `ping`, refer to the `ping(1M)` man page.

B. *ping successful?* If `ping -f inet6` succeeded, return to Flowchart 2. If `ping -f inet6` failed, the problem may exist in the routing table for the problem host. Continue to C.

C. *Execute: netstat -rnf inet6*. To display gateway routing information in numerical form, execute: `netstat -rnf inet6`

D. *Direct route to remote or default route to gateway?* If the route exists, go to F. If not, continue to E to add a new route.

E. *Add route entry on local system*. Use the `route` command to add a route entry to the route table on the local system. Refer to `route(1M)` for a complete description of the command. Or if an IPv6 router on the LAN advertises default routes, wait a few minutes to see if a route advertisement is added to the default router list. Start again with Flowchart 6.

F. *Correct router configured?* If your local host has a route to the correct router, then retry Flowchart 6 from the remote node. If the remote node's routing is configured properly, and both the local and remote nodes can connect to their respective routers, then contact your ISP or network administrator to verify network-to-network connectivity.

G. *Change route entry on local system or router.* If the routing information is incorrect, correct it using `route`, or verify that the IPv6 router is advertising proper subnet prefixes. Then retry Flowchart 2 to test network connectivity.

5 Utilities

HP-UX 11i IPv6, bundled as part of TOUR 1.0, for the most part uses IPv6-enhanced IPv4 network utilities. This section summarizes the utilities required for administration of HP-UX 11i IPv6 bundled as part of TOUR 1.0.

Configuration Utilities

This section describes configuration utilities available to configure HP-UX 11i IPv6 bundled as part of TOUR 1.0.

New ifconfig “inet6” address family

Use `ifconfig` to assign an IPv6 address to an interface and configure parameters, such as the network prefix. (InIPv6, prefix replaces netmask.)

The new `ifconfig keyword` **inet6** is required to configure IPv6 interfaces. It is not required to examine IPv6 interfaces. Refer to the `ifconfig(1M)` man page for details.

Neighbor Discovery Protocol Replaces arp in IPv6

The Neighbor Discovery Protocol (`ndp`) replaces `arp` in IPv6. Refer to “Neighbor Discovery” on page 55 in Chapter 6 for details.

New “inet6” route Option

`route(1M)` adds and deletes entries to the network routing table, allowing your system to communicate through a router. In IPv6, routing entries are automatically added when router advertisements are received from an IPv6 router. A configured tunnel route can be added using the `route` utility. Refer to `route(1m)` for more information.

Network Diagnostic Utilities

This section lists network diagnostic utilities available for use as part of the process of administering HP-UX 11i IPv6 bundled as part of TOUR 1.0.

- `lanadmin(1M)` resets or reports status of the LAN card.
- `lanscan(1M)` displays LAN device configuration and status.
- `ndd(1M)` displays and modifies network driver parameters.
- `ndp(1M)` displays and modifies the IPv6 neighbor discovery cache.
- `netstat(1)` provides network statistics and information about network connections.
- `ping(1M)` verifies network connectivity through the Network Layer and reports round-trip time of communication time between hosts.
- `traceroute(1M)` traces the path between hosts at the Network Layer.

IPv6 Additions to Network Tracing and Logging

Use `nettl` to trace traffic through new IPv6 Subsystems, or use `nettladm`. Table 5-1 below lists the new subsystems available for IPv6 packet tracing.

Table 5-1 New Network Trace Subsystems

Description	Subsystem Name
IPv6 Packets	NS_LS_IPV6
ICMPV6 Packets	NS_LS_ICMPV6
IPv6 Loopback packets	NS_LS_LOOPBACK6

Use `netfmt` to format trace records captured by `nettl` from the new IPv6 subsystems. `netfmt` can also filter `nettl` output according to the following IPv6 criteria:

Table 5-2 New IPv6 Network Filter Criteria

Filter Description	Entry in the netfmt configuration file
IPv6 Packets	NS_LS_IPV6
ICMPV6 Packets	NS_LS_ICMPV6
IPv6 Source Address	`filter ip6_saddr ::abcd` where ::abcd is the source address
IPv6 Destination Address	`filter ip6_daddr ::fedc` where ::fedc is the destination address

Table 5-2 New IPv6 Network Filter Criteria (Continued)

Filter Description	Entry in the netfmt configuration file
Connection per port and IPv6 address	filter connection6 <local_IPv6addr> | <port> <remote_IPv6addr | port

Contacting Your HP Representative

If you do not have a service contract with HP, you may follow the procedure described below, but you will be billed accordingly for time and materials.

If you have a service contract with HP, document the problem as a Service Request (SR) and forward it to your HP representative. Include the following information where applicable:

- A characterization of the problem. Describe the events and symptoms leading up to the problem. Attempt to describe the source of the problem.

  Your characterization should include: HP-UX commands; communication subsystem commands; functionality of user programs; result codes and messages; and data that can reproduce the problem.

- Obtain the version, update, and fix information for all software.

  To check the version of your HP-UX Operating System, execute the command:

  ```
  uname -a >> /tmp/filename
  ```

  This allows HP to determine if the problem is already known, and if the correct software is installed at your site.

- Illustrate as clearly as possible the context of any message(s). Record all error messages and numbers that appear at the user terminal and the system console.

- Prepare a listing of the HP-UX I/O configuration you are using for your HP representative to further analyze.

- Try to determine the general area within the software where you think the problem exists. Refer to the appropriate reference manual and follow the guidelines on gathering information for that product.

- Document your interim, or “workaround,” solution. The cause of the problem can sometimes be found by comparing the circumstances in which it occurs with the circumstances in which it does not occur.

- Create copies of any Internet Services or HP-UX 11i IPv6 Software link trace files that were active when the problem occurred, for your HP representative to further analyze.
- In the event of a system failure, obtain a full memory dump. If the directory `/var/adm/crash` exists, the HP-UX utility `/usr/sbin/savecore` automatically executes during reboot to save the memory dump. HP recommends that you create the `/var/adm/crash` directory after successfully installing this product. Send the output of your system failure memory dump to your HP representative.
- Prepare copies of the name service files such as `/etc/hosts`, etc/`nsswitch.conf`, `named.conf` and `resolv.conf`. Prepare a copy of the IPv6 configuration file /`etc/rc.config.d/netconf-ipv6`.
- Verify the software: `/usr/sbin/swverify > /tmp/swv-out`
- Execute the `display` command of the `lanadmin` diagnostic on the LAN interface and record the output.
- Record the troubleshooting flowchart number and step number where you are unable to resolve the problem.
- Save all network log files. Make sure that ERROR and DISASTER log classes are enabled when log files are collected in `/var/adm/nettl.LOG000`.
- Execute the following commands and record the output:

```
uname -a >> /tmp/filename
what /stand/vmunix >> /tmp/filename
lanscan >> /tmp/filename
netstat -sf inet6 >> /tmp/filename
netstat -inf inet6 >> /tmp/filename
netstat -rnf inet6 >> /tmp/filename
ndp -an >> /tmp/filename
ndd -get /dev/tcp6 tcp_status >> /tmp/filename
ndd -get /dev/ip6 ip6_ill_status >> /tmp/filename
ndd -get /dev/ip6 ip6_ipif_status >> /tmp/filename
ndd -get /dev/ip6 ip6_ire_status >> /tmp/filename
ndd -get /dev/ip6 ip6_ill_config_status >> /tmp/filename
```

  Prepare the formatted output and a copy of the log file for your HP representative to further analyze.

6 IPv6 Addressing and Concepts

This chapter introduces network addressing concepts for IPv6. It contains sections on Obtaining IPv6 Addresses, IPv6 Address Formats, Neighbor Discovery, Stateless Address Autoconfiguration and some basic general Networking Terminology.

Where to Get IPv6 Addresses

Contact a local ISP or the Regional Internet Registries below

ARIN - American IPv6 registration services
APNIC- Asia Pacific Network Information Center
RIPE - European Regional Internet Registry

The amount of addresses allocated varies according to your network requirements. Small Internet Service Providers (ISPs) or end nodes acquire IPv6 addresses from their upstream provider. Large ISPs, for example can receive from ARIN a minimum prefix of /48 with a second-level allocation of 16 bits for subnets. The remaining 64 bits are for a network interface

For Development Networks, contact the 6bone (www.6bone.net). The 6bone provides a testbed for deployment of IPv6.

IPv6 Address Formats

IPv6 addresses are 128-bit entities. IPv4 addresses are 32-bit addresses normally written as four decimal numbers (dotted decimal), one for each byte of the address.

Example: `192.1.2.34`

IPv6 Node Addresses are 128-bit records represented as eight fields of up to four hexadecimal digits. A colon separates each field (:). Example: `3ffe:ffff:101::230:6eff:fe04:d9ff`.

To indicate a subnetwork address, IPv6 uses subnet prefixes similar to IPv4 CIDR format. Figure 5-1 shows a 128-bit IPv6 node address with a 64-bit subnet prefix.

Figure 6-1 IPv6 128-bit Addresses; HP-UX Default Prefix 64

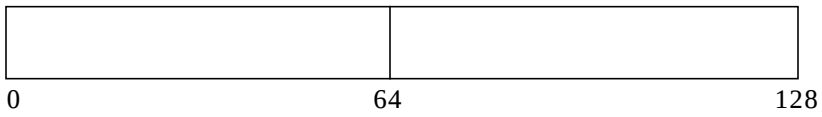

An IPv6 node address and its subnet prefix length can be combined in the following format:

<IPv6-Node-Address>/<Prefix-Length>

Where <*IPv6-Node-Address*> is an IPv6 address and <*Prefix-Length*> is a decimal value specifying how many of the leftmost contiguous bits of the address compose the subnet prefix.

In Figure 5-2, prefix length 48 specifies that the leftmost 48 bits of the IPv6 address compose the subnet prefix.

Figure 6-2 Example Prefix Length 48

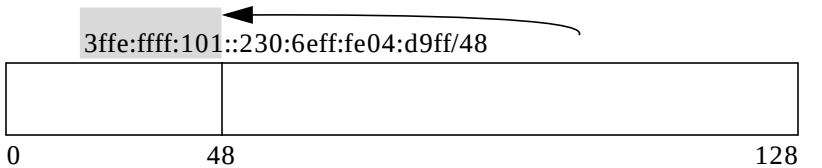

Address Scope

Link-local	An IPv6 address used on a single link.
Site-local	An IPv6 address used inside a single site.
Global	An IPv6 address that uniquely identifies a node on the Internet such that packets can be routed to the node from any other node on the Internet.

Address Type

Unicast	Identifies a single interface. Notable unicast addresses are:	
	Loopback	`::1` Address internal to IPv6 stack
	Unspecified	`::` Not a legally defined address
Anycast	Identifies a group of interfaces, possibly belonging to different nodes. A packet sent to an anycast address is delivered to only one of the interfaces in the group. Anycast addresses are currently not supported by HP-UX 11i IPv6, bundled as part of TOUR 1.0.	
Multicast	Identifies a group of interfaces, possibly belonging to different nodes. A packet sent to a multicast address is delivered to all the interfaces in this group.	

Neighbor Discovery

IPv6 hosts and routers use the IPv6 Neighbor Discovery Protocol to:

- advertise their link-layer address on the local link
- find neighbors' link-layer addresses on the local link
- find neighboring routers able to forward IPv6 packets
- actively track which neighbors are reachable
- search for alternate routers when a path to a router fails

The IPv6 Neighbor Discovery Protocol (ndp) uses ICMPv6. An IPv6-only utility, `ndp` and the Neighbor Discovery Protocol encompass the functionality of the IPv4 Address Resolution Protocol (ARP) and the `arp` utility. ndp also provides some of the address-configuration functionality found in protocols BOOTP and DHCP.

A network device connecting to a network for the first time can learn all parameters necessary to function, solely through Neighbor Discovery information. Both IPv6 hosts and routers advertise their presence using neighbor advertisements and route advertisements, respectively. When an IPv6 host first comes up, it advertises its link-layer address, and solicits neighbor and router information.

For more information, see the `ndp(1m)` and `ndp(7p)` man pages and RFC 2461, "Neighbor Discovery for IP Version 6 (IPv6)."

Stateless Address Autoconfiguration

Stateless address autoconfiguration requires no manual configuration of hosts, minimal configuration of routers, and no additional servers. The primary interface (`lanX:0`) is automatically assigned a link-local address by the system when the interface is configured. This allows each IPv6 interface to have at least one source address that can be used by Neighbor Discovery. Therefore, it is not advisable to assign other addresses to the primary interface besides the link-local address. See RFC 2373 “IP Version 6 Addressing Architecture” for details.

Link-Local Address Assigned Automatically

A link-local address is formed by prepending the well-known link-local prefix `FE80::/10` to the interface identifier which is typically 64 bits long and based on EUI-64 identifiers. Link-local addresses are sufficient for allowing communication among IPv6 hosts attached to the same link.

Figure 5-3 shows the Primary Interface Autoconfiguration steps performed after using the `ifconfig` command, which is as follows:

```
ifconfig lan0 inet6 up
```

Figure 6-3 Primary Interface Address Autoconfiguration

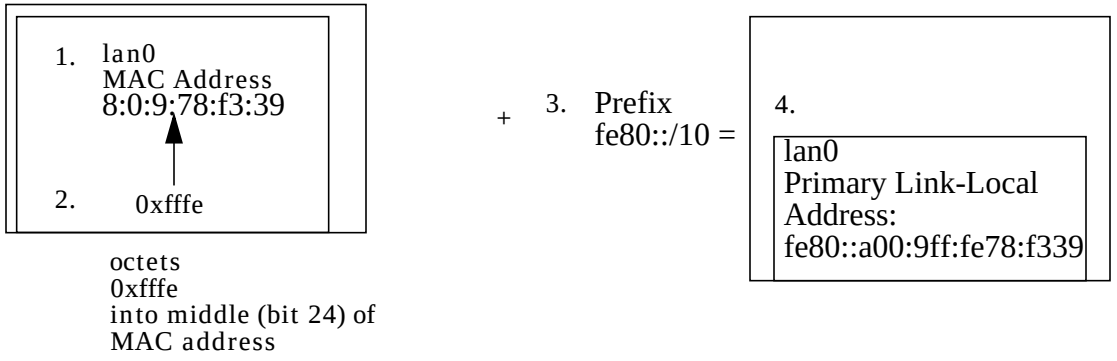

If you mark an interface “up” without assigning a primary address, the system derives a link-local address by performing the following 4 steps:

1. Taking the LAN card’s 48-bit link-level address (“MAC address” `8:0:9:78:f3:39`)

   `0000 1000 0000 0000 0000 1001 0111 1000 1111 0011 0011 1001`

   and putting it into an EUI-64 identifier by:

2. Putting two bytes (`0xffee`) into the middle (bit 24) of the 48-bit link-level address `8:0:9:`**`ff:fe`**`:78:f3:39`;

`0000 10`**`00`**` 0000 0000 0000 1001 `**`1111 1111 1111 1110`**` 0111 1000 1111 0011 0011 1001`

   and flipping the Universal/local bit (as described in RFC 2373) to form a 64-bit EIU-64 **interface identifier** `a:0:9:ff:fe:78:f3:39`

`0000 10`**`1`**`0 0000 0000 0000 1001 1111 1111 1111 1110 0111 1000 1111 0011 0011 1001`

3. Prepending the well-known prefix `fe80::/10`

4. Forming a 128-bit link-local unicast address for the primary interface `fe80::a00:9ff:fe78:f339`

Check the configuration by typing

```
ifconfig lan0 inet6
```

```
lan0: flags=4800841<UP,RUNNING,MULTICAST,ONLINK>
inet6 fe80::a00:9ff:fe78:f339 prefix 10
```

Secondary Interface Autoconfiguration

If an IPv6 router on the network advertises network prefixes in router advertisements, IPv6 derives a second IPv6 address based on the interface identifier. IPv6 assigns this address to a secondary interface for the network interface. The host adds the router as one of its default gateways.

Figure 5-4 shows a general example of Secondary Interface Autoconfiguration.

Figure 6-4 Secondary Interface Autoconfiguration From an IPv6 Router

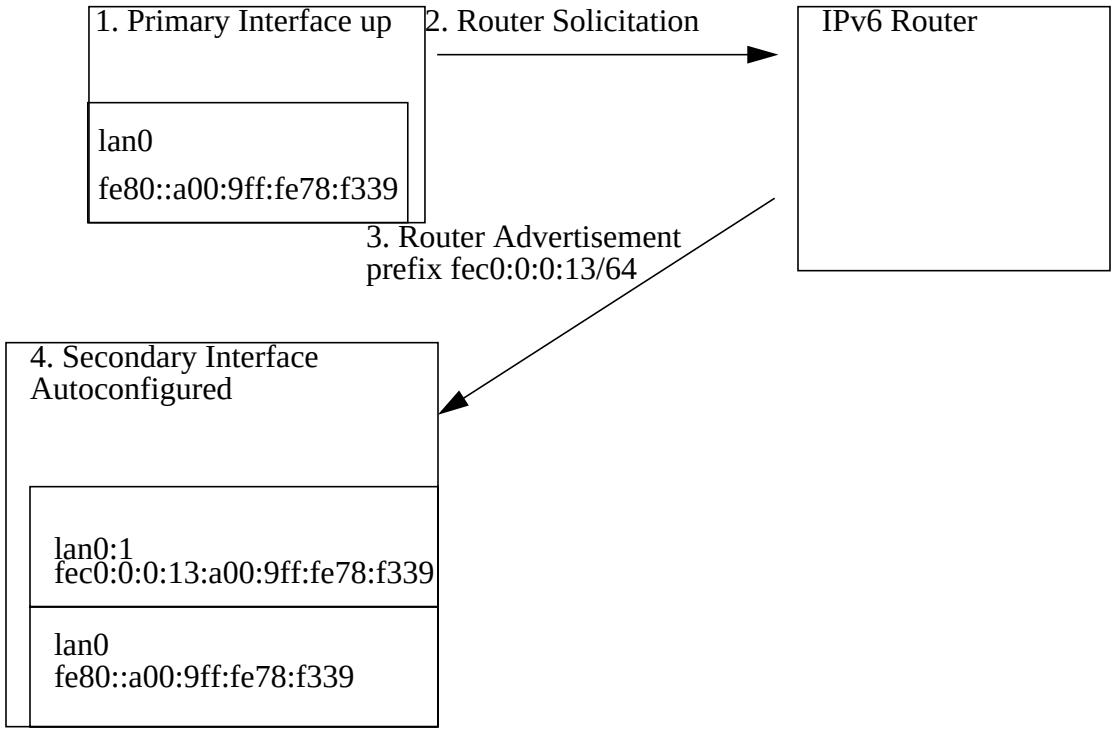

1. Primary interface comes up with the link-local address autoconfigured.
2. Host multicasts Router Solicitation.
3. IPv6 Router sends Router Advertisement to host.
4. Host autoconfigures secondary interface (`lan0:1`) by prepending prefix (`fec0:0:0:13/64`) sent by router to interface identifier (`a00:9ff:fe78:f339`). Refer to RFC 2461 “Neighbor Discovery for IP Version 6 (IPv6)” for details.

Manual Configuration and Router Advertisements

Note that even if a primary interface is manually configured, if the host receives prefixes from router advertisements, then secondary interfaces are autoconfigured. In this case, the addresses on the secondary interfaces are derived from the interface ID portion of the manually specified primary interface address.

Manual Configuration Overwriting Autoconfiguration

Manual configuration can overwrite autoconfiguration. When a secondary interface is configured with a manually assigned address, and if the user chooses an interface index number that has been used for an already autoconfigured secondary interface, the manual configuration overwrites the autoconfiguration. When this happens, network connectivity through the overwritten autoconfigured IP address is temporarily lost. At a later time, when the host receives the next router advertisement, the host will bring up another secondary interface with a different IP index number, but with the same IP address, and network connectivity through that IP address is restored. Normally, a user can avoid this by checking used IP index numbers. However, there is always a possibility that address autoconfiguration due to router advertisement is happening concurrently while the user manually configures secondary interfaces.

Disabling Specific IPv6 Interfaces

To disable communication through a specific IP address on an autoconfigured secondary interface, that secondary interface should be marked down, not removed or overwritten with a different IP address. If that interface is removed or overwritten, the host will reconfigure another secondary interface with the same IP address when it receives the next router advertisement. Alternatively, the router can be configured to stop advertising the prefix that corresponds to the offending IP address.

Removing An Interface

A primary interface cannot be removed from the system until all secondary interfaces are removed. You can remove secondary interfaces from the system using the `ifconfig inet6` command, as in the following example:

```
ifconfig lan1:1 inet6 ::
```

The primary interface (for example, lan1) can then be removed from the system with the `ifconfig` command, as in the following example:

```
ifconfig lan1 inet6 unplumb
```

A loopback interface does not have a hardware device associated with it. The name of the loopback interface is lo0. A loopback interface is automatically created by the system. You cannot delete it.

Networking Terminology

The following are descriptions of some important IPv6 networking terms.

Node

A node is a device that implements IP on the network. A node can be a host or a router.

A local node (or host) is the computer (or host) where you have logged-in. A remote note is a computer on the IP network where you are not logged in. A remote node does not have to be directly attached to your terminal.

Router

A **router** is a node that forwards IP packets not explicitly addressed to itself. It is a device that can forward packets between two or more IP networks. An IPv6 router can advertise prefixes. IPv6 router guidelines are beyond the scope of this manual. Refer to RFC 2461 for IPv6 router guidelines.

Host

A **host** is any node that is not a router.

Network Interface Name

A **network interface** is a communication device through which messages can be sent and received. An IPv6 address is associated with an interface name. Find the interface name(s) for a network interface by running the lanscan command and looking at the “Net-Interface Name PPA” field. For example,

```
lanscan
Hardware Station        Crd Hdw   Net-Interface  NM  MAC       HP-DLPI DLPI
Path     Address        In# State NamePPA        ID  Type      Support Mjr#
2/0/2    0x08000978F339 0   UP    lan0 snap0     1   ETHER     Yes     119
```

The interface name may include a colon (:), followed by an interface index number that denotes the interface number. The interface index number 0 is the first interface number for a card/encapsulation type and is known as the primary interface. The interface name lan0 is equivalent to lan0:0. The syntax is as follows:

nameX[:*interface-index-number*]

In the preceding syntax, *name* is the class of the interface. Valid name is lan (Ethernet LAN). *X* is the Physical Point of Attachment (PPA). *interface-index-number* is the number of the interface.

You must configure the **primary interface** for a LAN card before you can configure subsequent interfaces, known as **secondary interfaces**, for the same card. For example, you must configure lan0 before you configure lan0:1 and lan0:2.

7 IPv6 Software and Interface Technology

The topics discussed in this section concern HP-UX IPv6 deployment and migration.

Name and Address Lookup for IPv6

It is generally recommended to add IPv6 addresses (known as AAAA records) to a DNS Name Server only when the following conditions are true:

- The IPv6 address is assigned to the interface on the node.
- The address is configured on the interface.
- The interface is on a link which connects to the IPv6 infrastructure.

HP recommends beginning with IPv6 addresses and host names in the `/etc/hosts` file on a development network; then adding IPv6 addresses and hosts to a Domain Name Service when moving IPv6 to a production backbone network.

`/etc/nsswitch.conf` (`nsswitch.conf(4)`) is a configuration file for the name service switch. A new entity, `ipnodes`, specifies which name services resolve IPv4 and IPv6 addresses and host names on HP-UX 11i IPv6 bundled as part of TOUR 1.0.

More specifically, on HP-UX 11i IPv6 bundled as part of TOUR 1.0, the new keyword `ipnodes` specifies the resolver policy for the new library functions `getnameinfo(3N)`, `getaddrinfo(3N)`, `getipnodebyname(3N)` and `getipnodebyaddr(3N)` for both IPv4 and IPv6 addresses. The existing keyword `hosts` specifies the resolver policy for the library functions `gethostbyname()` and `gethostbyaddr()` for IPv4 addresses.

NOTE Internet Services applications (such as `telnet`, `r*` commands, etc.) use these new library functions to resolve IPv4 and IPv6 addresses.

By default, the `/etc/nsswitch.conf` is not on a system. The default `ipnodes` policy (same as default `hosts` policy) is as follows:

```
dns [NOTFOUND=return] files
```

This policy implies that `dns` is the authoritative resolver and will only try `files` if `dns` is down. If `dns` is available but returns `NOTFOUND`, the search stops.

Thus, if DNS has **not** been set up as the definitive source, and files (/etc/hosts) may need to be used for address and host name resolution, HP recommends adding the following entry to /etc/nsswitch.conf:

```
ipnodes: dns [NOTFOUND=continue] files
```

Or if /etc/hosts is to be the primary Name Service

```
ipnodes: files [NOTFOUND=continue] dns
```

IMPORTANT

If you have an /etc/nsswitch.conf file, and if you have set a hosts entry in that file, then you must set a similar ipnodes entry. You must do this even if you do not use IPv6 addresses. Also, be aware that you cannot use NIS or NIS+ on the ipnodes entry.

Migrating Name and IPv6 Address Lookup

Most sites test IPv6 on a development subnetwork before deploying it on a larger scale. These sites typically add IPv6 address and host names to the /etc/hosts files on IPv6 hosts, then change their hosts lookup policy to search files.

HP recommends that you maintain at least a minimal /etc/hosts file that includes important addresses like gateways, root servers, and your host's own IP address. HP also recommends that you include the word files in the hosts and ipnodes lines of /etc/nsswitch.conf to help ensure a successful system boot using the /etc/hosts file when DNS is not available.

Migrating from IPv4 to IPv6

To successfully migrate to IPv6, maintain compatibility with the large installed base of IPv4 hosts and routers. Staying compatible with IPv4 when deploying IPv6 eases the task of moving the Internet to IPv6. HP-UX 11i IPv6, bundled as part of TOUR 1.0, supports three of the many transition mechanisms recommended by the IETF:

- Dual-Stack
- Automatic and Configured Tunneling IPv6 traffic through IPv4 networks
- “6to4” Connection of IPv6 Domains via IPv4 Clouds (RFC 3056). Note that only host functionality is supported by HP-UX 11i IPv6 bundled as part of TOUR 1.0.

Dual Stack

Dual Stack nodes support both IPv6 and IPv4 functionality. HP-UX 11i IPv6 bundled as part of TOUR 1.0 supports dual stack. IPv6 applications can coexist with IPv4 applications.

Tunneling

Tunneling encapsulates IPv6 packets within IPv4 packets. IPv6 transmission across the IPv4 Internet is transparent.

Tunneling

IPv6 tunneling enables IPv6 hosts and routers to connect with other IPv6 hosts and routers over the existing IPv4 Internet. IPv6 tunneling encapsulates IPv6 datagrams within IPv4 packets. The encapsulated packets travel across an IPv4 Internet until they reach their destination host or router. The IPv6-aware host or router decapsulates the IPv6 datagrams, forwarding them as needed. IPv6 tunneling eases IPv6 deployment by maintaining compatibility with the large existing base of IPv4 hosts and routers.

IPv6/IPv4 hosts and routers can tunnel IPv6 datagrams over regions of IPv4 routing topology by encapsulating them within IPv4 packets. Tunneling can be used in a variety of ways:

Router-to-Router

IPv6/IPv4 routers interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the end-to-end path that the IPv6 packet takes.

Host-to-Router IPv6/IPv4 hosts can tunnel IPv6 packets to an intermediary IPv6/IPv4 router that is reachable via an IPv4 infrastructure. This type of tunnel spans the first segment of the packet's end-to-end path.

Host-to-Host IPv6/IPv4 hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire end-to-end path that the packet takes.

Router-to-Host IPv6/IPv4 routers can tunnel IPv6 packets to their final destination IPv6/IPv4 host. This tunnel spans only the last segment of the end-to-end path.

Configured Tunneling

Tunneling techniques are classified by the way the encapsulating node determines the address of the node at the end of the tunnel. In router-to-router and host-to-router tunneling methods, the IPv6 packet is tunneled to a router. The tunnel endpoint is an intermediary router. The intermediary router decapsulates the IPv6 packet, then forwards it to its final destination. When tunneling to a router, the tunnel endpoint differs from the tunneled packet's destination. So the addresses in the tunneled IPv6 packet do not provide the IPv4 address of the tunnel endpoint. Instead, the node performing the tunneling provides configuration information that determines the tunnel endpoint address. This type of tunneling is called “configured tunneling.”

Automatic Tunneling

In the host-to-host and router-to-host tunneling methods, the IPv6 packet is tunneled all the way to its final destination. The tunnel endpoint is the node to which the IPv6 packet is addressed. Because the tunnel endpoint is the IPv6 packet's destination, the IPv6 packet's destination address determines the tunnel endpoint: if that address is an IPv4-compatible IPv6 address, then the low-order 32-bits hold the destination node's IPv4 address. That can be used as the tunnel endpoint address. This technique avoids the need to explicitly configure the tunnel endpoint address. Deriving the tunnel endpoint address from the embedded IPv4 address of the packet's IPv6 address is called “automatic tunneling.”

Connecting IPv6 Domains over IPv4 Clouds (6to4)

Isolated IPv6 nodes and networks can communicate over an IPv4 network without explicitly configuring tunnels, by using “6to4” mechanism (RFC 3056). “6to4” effectively uses the IPv4 wide area network as a unicast point-to-point layer. “6to4” requires no end-node reconfiguration and minimal router configuration.

“6to4” Well-Known Prefix

The IANA permanently assigned one IPv6 address prefix for “6to4”. It is `2002::/16`. The “6to4” prefix `2002:/16` can be prepended to a host or router’s globally-unique 32-bit IPv4 address (<IPv4-Addr>) to form a 48-bit “6to4” prefix `2002:<IPv4-Addr>`. The “6to4” prefix provides a network prefix for the local IPv6 host or network. The IPv4 address is the endpoint for all external IPv4 connections.

Figure 7-1 **“6to4” Prefix**

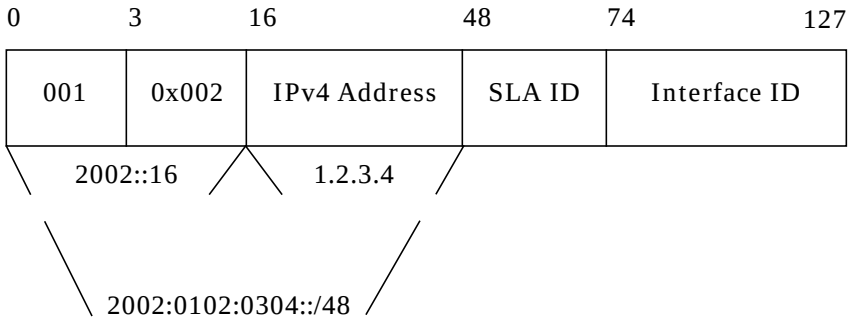

Encapsulating in IPv4

IPv6 packets from a “6to4” site are encapsulated in IPv4 packets when they leave the site over its external IPv4 connection. IPv6 packets are transmitted in IPv4 packets with an IPv4 protocol type of 41, the same protocol type set when IPv6 packets tunnel inside IPv4 frames according to RFC 1933.

Example "6to4" Topology

Figure 6-2 shows two IPv6 subnetworks. The end nodes have their routers globally unique IPv4 addresses embedded in their network prefixes. The routers have "6to4" addresses and corresponding globally unique IPv4 addresses. From the IPv6 end-node view, each host's subnetwork is connected to the other's through a 6to4 router. All IPv4 tunneling is transparent to the IPv6 end nodes.

Figure 7-2 **6to4 IPv6 End Node View**

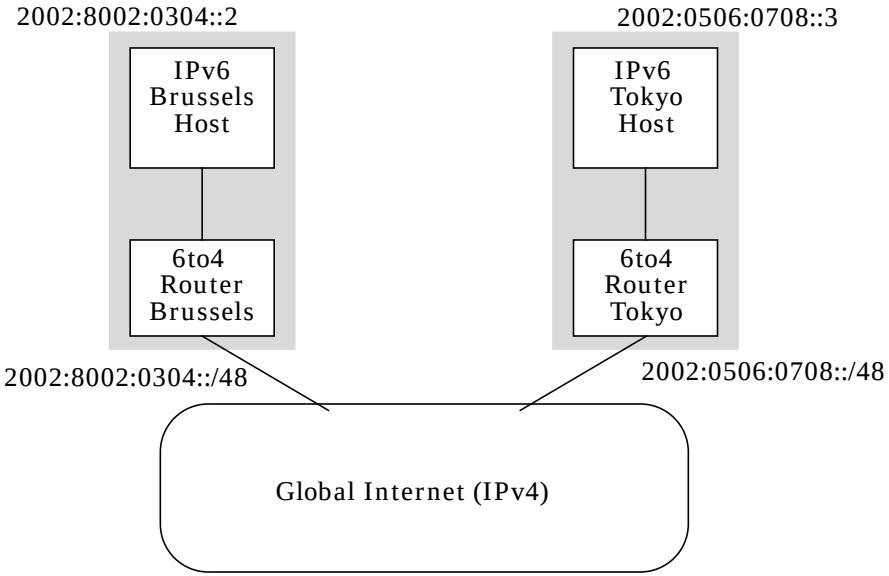

Figure 6-3 shows the same routers from the IPv4 network perspective. Router Brussels and Router Tokyo are IPv4 Routers. All IPv6 traffic is tunneled transparently through the IPv4 network.

Figure 7-3 6to4 IPv4 Internet View

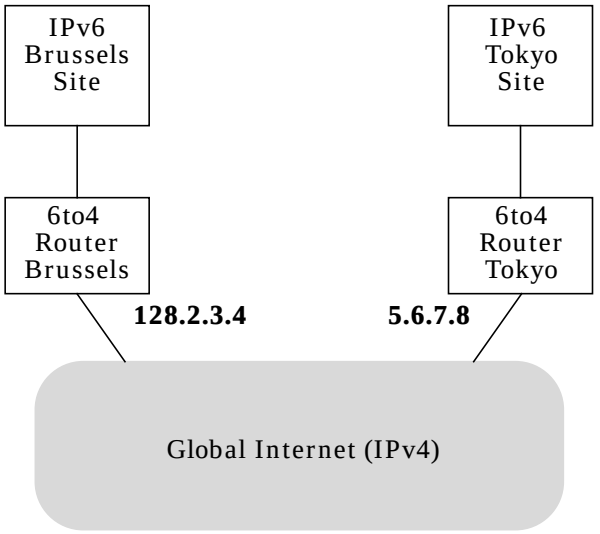

A IPv6 ndd Tunable Parameters

The following new IPv6 tunable parameters allow advanced performance fine-tuning of HP-UX 11i IPv6, bundled as part of TOUR 1.0.

NOTE

To obtain a list of supported IPv6 ndd parameters, enter:

```
ndd /dev/ip6 ?
```

For more information on a specific parameter, enter:

```
ndd -h <parameter>
```

Below is the output received from entering, "ndd /dev/ip6 ?":

```
ip6_ill_status                (read only)
ip6_ipif_status               (read only)
ip6_ifhash_status             (read only)
ip6_ire_status                (read only)
ip6_raw_status                (read only)
ip6_tcp_status                (read only)
ip6_udp_status                (read only)
ip6_ire_hash                  (read only)
ip6_ire_hash_summary          (read only)
ip6_rput_pullups              (read and write)
ip6_fragment_timeout          (read and write)
ip6_ill_config_command        (write only)
ip6_ill_config_status         (read only)
ip6_reass_mem_limit           (read and write)
ip6_forwarding                (read and write)
ip6_send_redirects            (read and write)
ip6_debug                     (read and write)
ip6_ire_reachable_interval    (read and write)
ip6_ire_cleanup_interval      (read and write)
ip6_ire_redirect_interval     (read and write)
ip6_ire_pathmtu_interval      (read and write)
ip6_def_hop_limit             (read and write)
```

```
ip6_wroff_extra               (read and write)
ip6_dl_sap                    (read and write)
ip6_dl_snap_sap               (read and write)
ip6_bogus_sap                 (read and write)
ip6_encap_hop_limit           (read and write)
ip6_nd_dad_solicit_count      (read and write)
ip6_nd_multicast_solicit_count(read and write)
ip6_nd_unicast_solicit_count  (read and write)
ip6_nd_advertise_count        (read and write)
ip6_rd_solicit_count          (read and write)
ip6_nd_transmit_interval      (read and write)
ip6_nd_anycast_delay          (read and write)
ip6_nd_probe_delay            (read and write)
ip6_rd_solicit_delay          (read and write)
ip6_rd_transmit_interval      (read and write)
ip6_min_random_factor         (read and write)
ip6_max_random_factor         (read and write)
ip6_icmp_interval             (read and write)
ip6_ipsec_enable              (read and write)
ip6_ipsec_polmask_action      (read and write)
ip6_ipsec_samask_type         (read and write)
ip6_ipsec_pollist             (read only)
ip6_ipsec_salist              (read only)
```