

Tru64 UNIX Best Practice

Deploying IPv6 in Your Network

June 2001

Product Version: **Tru64 UNIX Version 5.1, 5.1A**

This Best Practice describes how to deploy IPv6 for the Tru64 UNIX operating system in your network.

Contents

Deploying IPv6 in Your Network

Is This Best Practice Right for You?	1
Before You Begin	1
Understand IPv6	2
Understand Tunnels and How They Work	2
Develop an Implementation Plan	3
Intranet Scenario	3
Intranet-to-Internet Scenario	8
Intranet-to-Internet-to-Intranet Scenario	10
Include IPv6 Support in the Kernel	11
Applying the Best Practice	12
Port Existing IPv4 Applications	12
Obtain IPv6 Addresses	12
Install IPv6-Capable Routers	13
Configure Domain Name Service (DNS) Servers	13
Configure Tru64 UNIX IPv6 Routers	14
Configure Tru64 UNIX IPv6 Hosts	15
Verifying Success	16
Troubleshooting	16
Comments and Questions	16
Legal Notice	16

Deploying IPv6 in Your Network

Internet Protocol Version 6 (IPv6), as defined in RFC 2460, is the replacement network layer protocol for the Internet, designed to replace the current Internet Protocol Version 4 (IPv4). IPv6 also changes the structure of the Internet architecture. That does not mean that you have to deploy IPv6 all at once across your network; rather it is something that you can do in stages because IPv6 and IPv4 were designed to interoperate. This Best Practice provides you with guidelines for deployment, presents deployment scenarios, and describes the steps to follow whether for a single system or your entire network.

See the Tru64 UNIX Best Practices Web page for more information about Best Practices documentation.

Is This Best Practice Right for You?

Not all Best Practices apply to all configurations, so you must be sure that it is appropriate for your system and circumstances. To use this Best Practice, you must meet the requirements described in the following table:

Requirement	Description
Operating System	Tru64 UNIX Version 5.1
System Configuration	IPV6 and IPTUNNEL (if needed) options need to be configured in the kernel.
Impact on Availability	System will be unavailable while the new kernel is booted.

Before You Begin

Before you apply the Best Practice for deploying IPv6 in your network, you must understand some background information and perform some preliminary tasks.

Understand IPv6

You should be familiar with IPv4 and IPv6. The following is a summary of IPv6 features:

- Address

The IPv6 address is 128 bits in length (compared to the 32-bit IPv4 address) and uses a new text representation format. In addition, there are three types of IPv6 addresses: unicast, anycast, and multicast. The unicast address consists of an address prefix and a 64-bit interface identifier. See the *Network Administration* manual and RFC 2373 for information about IPv6 addresses.

- Neighbor Discovery

A mechanism by which IPv6 nodes on the same link discover each other's presence, determine each other's link-local addresses, find routers, and maintain reachability information about paths to active neighbors and remote destinations.

- Stateless address autoconfiguration

The process by which IPv6 nodes listen for Router Advertisement packets from routers and learn IPv6 address prefixes. The node creates IPv6 unicast addresses by combining the prefix with a datalink-specific interface identifier that is typically derived from the datalink address of the interface. The Tru64 UNIX operating system performs this process automatically.

Understand Tunnels and How They Work

Tunneling IPv6 packets in IPv4 is a mechanism that enables the gradual deployment of IPv6 in your network by allowing IPv6 nodes to interoperate with IPv4 hosts and routers. Tru64 UNIX systems support both IPv4 and IPv6, and therefore can have both an IPv4 address and an IPv6 address. An end system with both addresses is considered a v4/v6 host; a router with both addresses is considered a v4/v6 router. A v4/v6 host can use IPv6 to communicate with other v4/v6 hosts on the same communications link. However, when these hosts need to communicate over an IPv4 network, the hosts need to tunnel the IPv6 packets in IPv4 packets in order for the IPv4 routing infrastructure to route the packets to the destination host.

The Tru64 UNIX implementation uses bidirectional configured tunnels to carry IPv6 packets through an IPv4 routing infrastructure; unidirectional tunnels are not supported. This means that a configured tunnel must be created on the nodes at both ends of the tunnel. A bidirectional configured tunnel behaves as a virtual point-to-point link. For the remainder of

this Best Practice, the term **configured tunnel** refers to a bidirectional configured tunnel.

A configured tunnel has a source IPv4 address and a destination IPv4 address. The following configured tunnels are possible:

- Router-to-router tunnel — In this case, the v4/v6 routers are connected by an IPv4 infrastructure. For end-to-end communications, this represents only one segment of the total path. (The Intranet-to-Internet-to-Intranet Scenario section contains a picture that shows this type of tunnel.)
- Host-to-router tunnel — In this case, the v4/v6 host and v4/v6 router are connected by an IPv4 infrastructure. For end-to-end communications, this represents the first segment of the total path. (The Intranet Scenario section contains a picture that shows this type of tunnel.)
- Host-to-host tunnel — In this case, the v4/v6 hosts are connected by an IPv4 infrastructure. For end-to-end communications, this represents the total path; the tunnel spans the total path.
- Router-to-host tunnel — In this case, the v4/v6 router and v4/v6 host are connected by an IPv4 infrastructure. For end-to-end communications, this represents the final segment of the total path. (The Intranet-to-Internet Scenario section contains two pictures that show this type of tunnel.)

See RFC 2893 for more information on tunnels.

Develop an Implementation Plan

The following three scenarios, in order of increasing complexity, serve as models for deploying IPv6 in your network:

- Intranet
- Intranet-to-Internet
- Intranet-to-Internet-to-Intranet

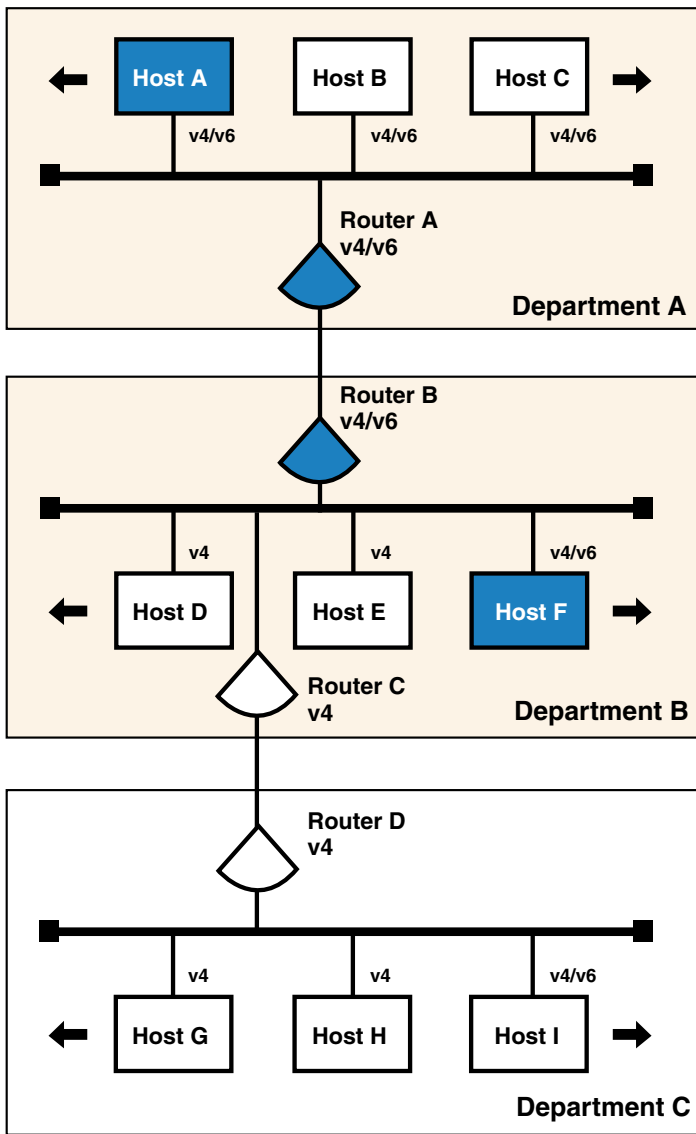
Intranet Scenario

In this scenario, you deploy IPv6 hosts on a small subnet in your network. These hosts communicate with each other using link-local addresses. If you add an IPv6 router to the subnet and advertise an address prefix, each IPv6 host autoconfigures a global IPv6 address and uses that to communicate with other IPv6 hosts.

As you become more experienced with using IPv6, for the next phase you can add an IPv6 host or hosts on other subnets in your network. Communications between IPv6 hosts on different subnets would be accomplished using configured router-to-host tunnels and host-to-router tunnels. The existing IPv4 routing infrastructure is used to get the packets end to end.

The following figures illustrate an intranet scenario in which a corporation has three departments in a local geographic area. Department A has deployed v4/v6 hosts and a v4/v6 router. Departments B and C have deployed only one v4/v6 host each, with a majority of v4 hosts.

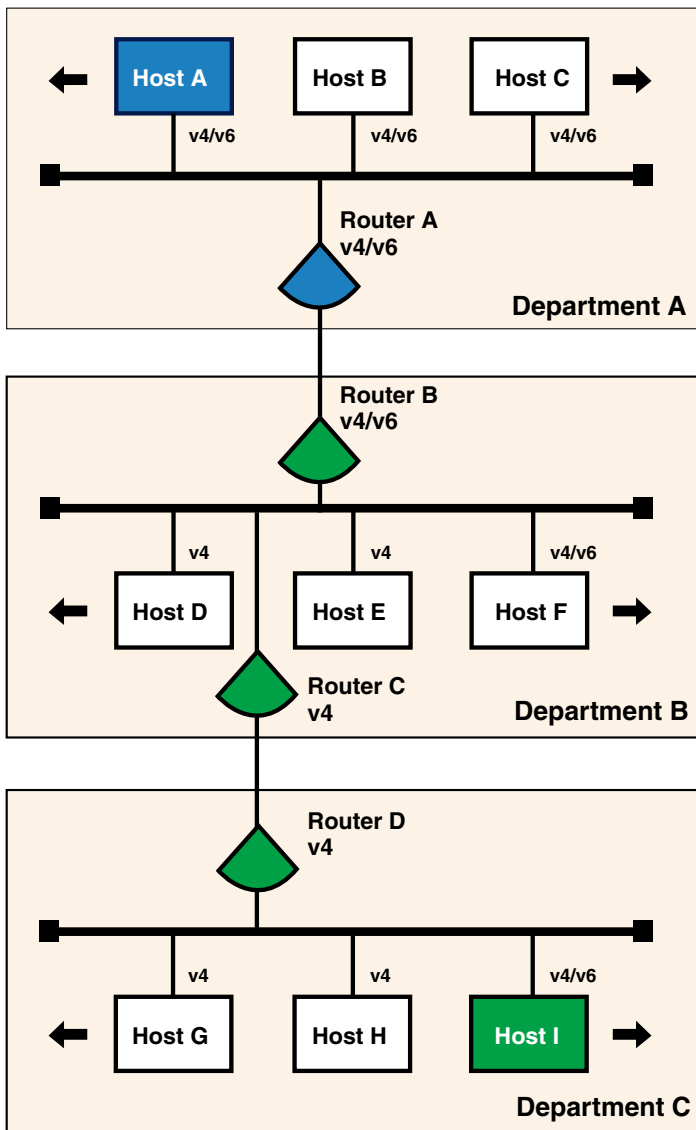
To communicate with Host F, native IPv6 traffic is routed from Host A to Host F.



ipv6_flow1

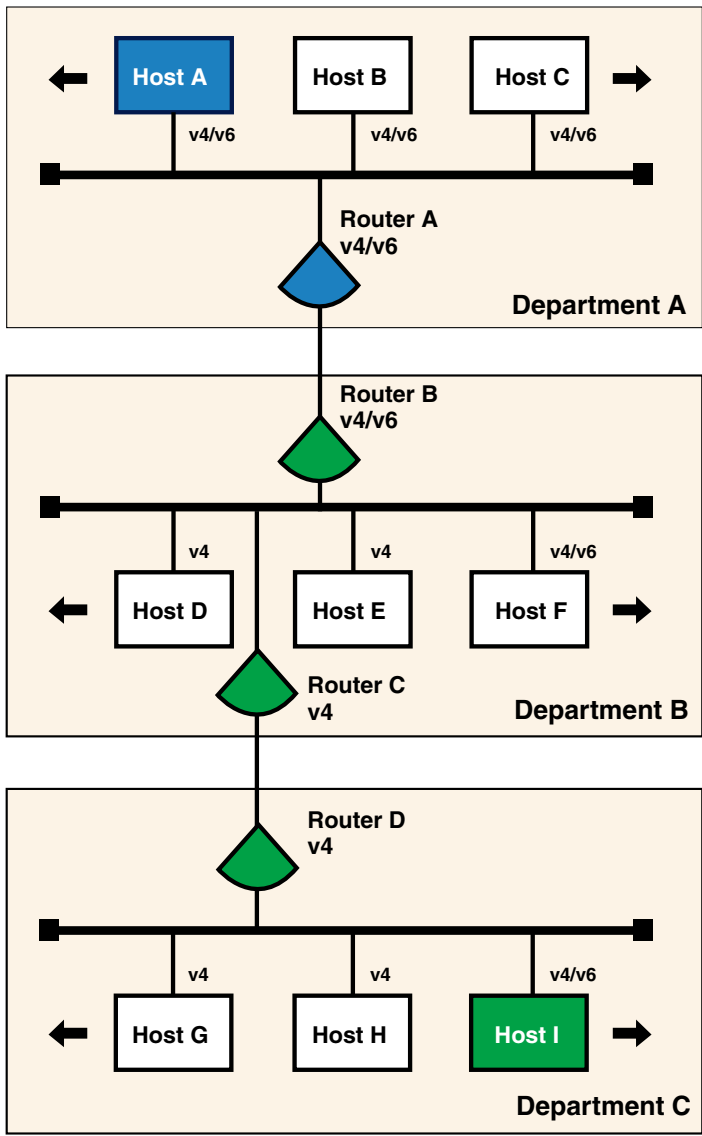
To communicate with Host I, Host A sends an IPv6 packet to Router A. Router A forwards the IPv6 packet to Router B. Router B encapsulates the IPv6 packet and sends the IPv4 packet over a router-to-host tunnel to Host

I, which decapsulates the IPv4 packet. The IPv4 infrastructure routes the packet to Host I. For hosts, the host-to-router tunnel is more efficient because it saves the Host A, Host B, and Host C administrators from having to create individual host-to-host tunnels for each destination host.



ipv6_flow2

To communicate with Host A, Host I encapsulates the IPv6 packet and sends the IPv4 packet over a host-to-router tunnel to Router B. From there, Router B decapsulates the IPv4 packet and routes the IPv6 packet to Host A. For hosts, the host-to-router tunnel is more efficient because it saves the Host I administrator from having to create individual host-to-host tunnels for each destination host.



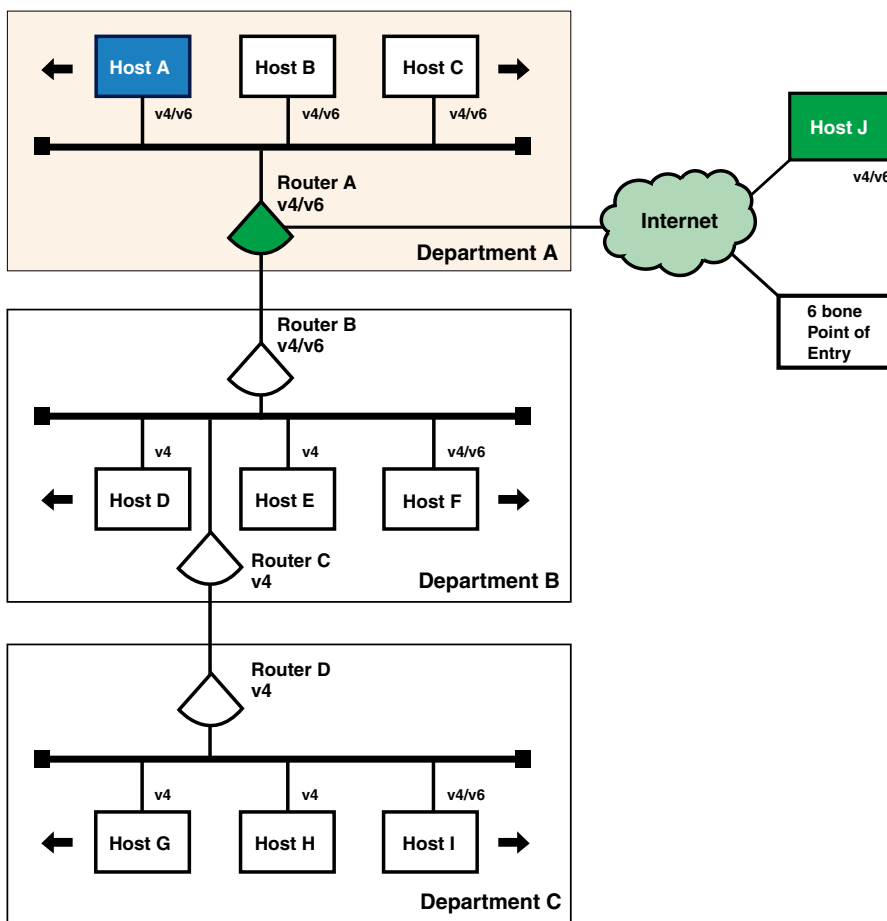
ipv6_flow3

Intranet-to-Internet Scenario

In this scenario, you add a v4/v6 router to your network and use it to communicate with the global Internet. The IPv6 hosts communicate with the v4/v6 router using IPv6. For IPv6 traffic to v4/v6 hosts on the 6bone

or the Internet, you configure router-to-host tunnels. The next figure illustrates a scenario in which the corporation described in the Intranet Scenario section adds a connection from Router A to the Internet. Potential destination nodes are in turn connected to the Internet.

To communicate with Host J, Host A sends the IPv6 packet to Router A. Router A encapsulates the IPv6 packet and sends the IPv4 packet over a router-to-host tunnel to Host J, which decapsulates the IPv4 packet.



ipv6_flow4

To communicate with the 6bone, Host A sends the IPv6 packet to Router A. Router A encapsulates the IPv6 packet and sends the IPv4 packet over a

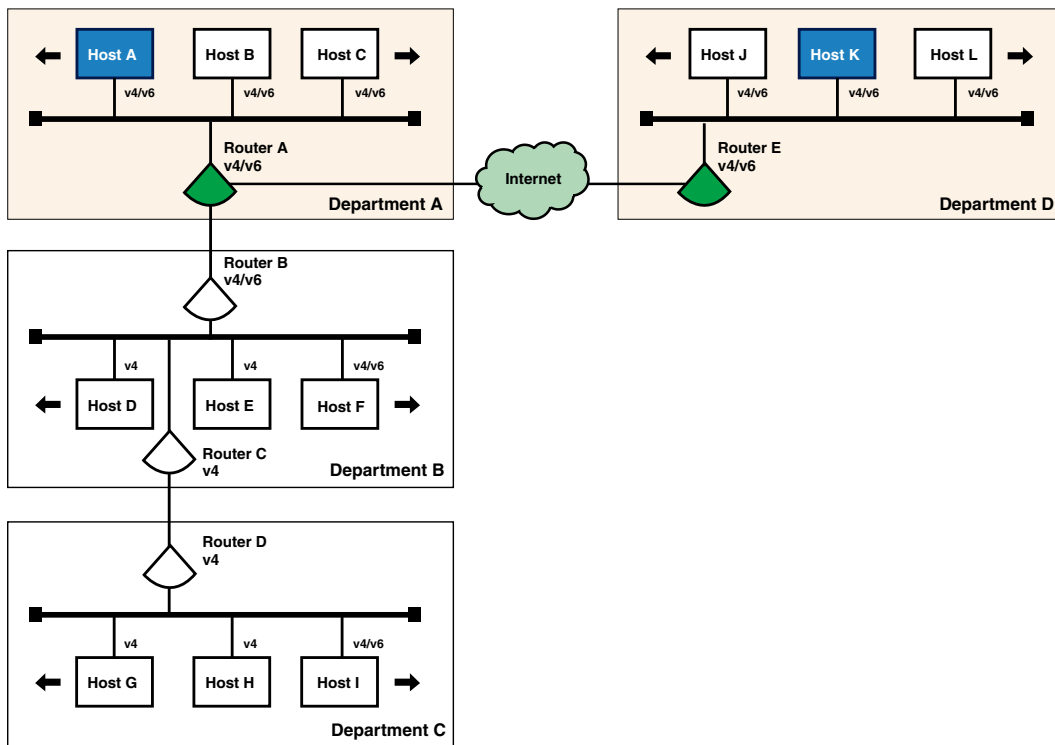
router-to-host tunnel to the 6bone point of entry. The point of entry router decapsulates the IPv4 packet and routes the IPv6 packet to its destination.

Intranet-to-Internet-to-Intranet Scenario

In this scenario, you add additional v4/v6 routers on remote subnets and connect the two of them through the Internet to create a virtual private network (VPN). An example of this might be a global corporation with manufacturing in one country and a design center in another country. The IPv6 hosts communicate with the v4/v6 routers using IPv6. For IPv6 traffic between the v4/v6 routers on each subnet, you configure router-to-router tunnels.

The next figure illustrates a scenario in which the corporation described in the previous sections wants to connect its corporate network with one of its geographically remote departments to create a VPN.

To communicate with Host K, Host A sends the IPv6 packet to Router A. Router A encapsulates the IPv6 packet and sends the IPv4 packet over a router-to-router tunnel to Router E, which decapsulates the IPv4 packet and routes the IPv6 packet to Host K. For routers, the router-to-router tunnel is more efficient because it saves the Router A administrator from having to create individual router-to-host tunnels for each destination host.



ipv6_flow5

Include IPv6 Support in the Kernel

To include the IP Version 6 (IPV6) and IP-in-IP Tunneling (IPTUNNEL) options in the kernel, do the following:

1. Build a new kernel by using the following command:

```
# doconfig -c SYSTEM_NAME
```

Choose the IPV6 and IPTUNNEL options in addition to any others you want.

2. Save the original kernel, then copy the new kernel to the root directory:

```
# cp /vmunix /vmunix.save
# cp /sys/SYSTEM_NAME/vmunix /vmunix
```

3. Reboot the system. Make sure that there are no other users on the system. Use a command similar to the following:

```
# shutdown -r +5 "Adding IPv6 and IPTUNNEL kernel options ..."
```

Applying the Best Practice

Before you deploy IPv6 in your network, be sure to follow the recommendations in *Before You Begin*.

Port Existing IPv4 Applications

The Tru64 UNIX operating system provides the basic applications programming interfaces (APIs) as defined in RFC 2553. You can use the APIs and the AF_INET6 sockets in your existing applications (or in new applications) to communicate with IPv4 nodes today. Your ported applications will continue to communicate with IPv4 nodes and be ready to communicate with IPv6 nodes. See the *Network Programmer's Guide* for guidelines for developing applications that use AF_INET6 sockets and client/server code examples.

Obtain IPv6 Addresses

IPv6 addresses are now being deployed by the regional registries. To obtain an IPv6 address or block of addresses, contact your Internet Service Provider (ISP).

If you are an Internet Service Provider, contact your upstream registry or one of the registries at the following locations:

```
APNIC (Asia-Pacific Network Information Center)
ARIN (American Registry for Internet Numbers)
RIPE NCC (Réseau IP Européens)
```

Because of the need to test various implementation of the IPv6 RFCs, the Internet Engineering Task Force (IETF) has defined a temporary IPv6 address allocation scheme. You can assign the addresses in this scheme to hosts and routers for testing IPv6 on the 6bone. See the 6bone home page at the following location for more information on 6bone address allocation and assignment:

```
http://www.6bone.net
```

After you contract with your ISP for a block of addresses, your deployment of IPv6 in your network begins the process of renumbering of your network. In IPv4, network renumbering was a difficult and time-consuming

process. In IPv6, network renumbering is more dynamic. This enables you renumber your network for any of the following reasons:

- Your enterprise is growing and needs more address space.
- Your network needs are changing.
- Your enterprise wants a global presence.
- You are outgrowing your ISP.

Whatever the reason, when your current ISP contract expires, your right to use the block of IPv6 addresses also expires. Although network renumbering is simplified in IPv6, the following guidelines will help ease the process:

- Have your routers advertise new network prefixes and deprecate the old prefixes by setting a lifetime.
- Change DNS servers to advertise node names and the new addresses.
- Do not hard code addresses in configuration files, as this makes the process more complex and labor intensive.
- Clear all server caches as appropriate.

Install IPv6–Capable Routers

This process depends on the hardware vendor you have chosen. You will need to define what address prefixes the router will advertise and the interfaces over which to advertise them.

Configure Domain Name Service (DNS) Servers

The Tru64 UNIX operating system supports AAAA lookups over IPv4 (AF_INET) connections only. The resolver and server have not been ported to IPv6, but IPv6 applications can make `getaddrinfo` and `getnameinfo` calls to retrieve the AAAA records.

To configure a DNS/BIND server to operate in an IPv6 environment, review the following guidelines:

- Select a node to function as an IPv6 name server.
- Dedicate a zone to IPv6 addresses or add IPv6 addresses to your enterprise's current zone. If you want global IPv6 name services, you must delegate a domain under the `ip6.int` domain for the reverse lookup of IPv6 addresses.

Note

Do not point different zone names to the same zone database file.

- Send mail to the following address to request a domain for reverse lookups:

`bmanning@isi.edu`

See RFC 1886 for more information.

- If the system is configured as a DNS/BIND server, change the `/etc/resolv.conf` file to point to the local node for name lookups, as follows:

```
nameserver 127.0.0.1
```

To enable dynamic updates for a DNS/BIND server, do the following:

1. Edit the `named.conf` file and add the `allow-update` substatement to the `zone` statements for those zones you want to dynamically update and for the reverse lookup zone. See Chapter 8 of the *Network Administration* manual and the `named.conf(4)` reference page for more information.
2. Start the `named` daemon. See the SysMan online help for more information.

Configure Tru64 UNIX IPv6 Routers

To configure a Tru64 UNIX IPv6 router, run the `ip6_setup` script and complete the following steps:

1. Identify the interface or interfaces over which to run IPv6.
2. Decide if you want to enable IPv6 routing over Point-to-Point Protocol (PPP) interfaces.
3. Decide if you need a configured IPv4 tunnel for communications with other IPv6 nodes or networks. You will need the remote node's IPv4 address (the remote end of the tunnel) and your node's IPv4 address (this end of the tunnel).
4. Decide if you want to configure static routes. You might want to configure static routes if one of the following conditions is true:
 - You want a configured tunnel and you are not advertising an address prefix on the tunnel link.

- You want a configured tunnel and the router on the other end of the tunnel is not running the RIPng protocol.
 - Your system is not running the RIPng protocol.
5. Identify the interface (LAN, PPP, or configured tunnel) on which you want to run the RIPng protocol or advertise an address prefix. If the latter, you must decide on the address prefix to advertise.

See the *Network Administration* manual for more information.

Configure Tru64 UNIX IPv6 Hosts

To configure a Tru64 UNIX IPv6 host, run the `ip6_setup` script and complete the following steps:

1. Identify the interface or interfaces over which to run IPv6.
2. Decide if you want this system to update the DNS/BIND database automatically. If you want to do this, you must also provide a fully qualified domain name for the IPv6 host.

Note

The DNS/BIND server must also be configured to allow the dynamic updates.

3. Decide if you need a configured IPv4 tunnel for communications with other IPv6 nodes or networks. You will need the remote node's IPv4 address (the remote end of the tunnel) and your node's IPv4 address (this end of the tunnel).
4. Decide if you want to configure static routes. You might want to configure static routes if you want a configured tunnel to a router and the router is not advertising itself as a default router on the tunnel link.

See the *Network Administration* manual for more information.

Verifying Success

After you apply the Best Practice for deploying IPv6 for the Tru64 UNIX operating system in your network, you can verify whether it was successful by completing the following steps:

1. Issue a `ping` command to an on-link node.
2. Issue a `ping` command to an off-link node.
3. Verify that your application, if using IPv6, can connect to a remote node.

Troubleshooting

If you determine that the Best Practice was not successful, as described in *Verifying Success*, use the extensive problem solving information contained in the *Network Administration* manual.

Comments and Questions

We value your comments and questions on the information in this document. Please mail your comments to us at this address:

`best_practices@zk3.dec.com`

Legal Notice

Compaq and the Compaq logo Registered in U.S. Patent and Trademark Office. Tru64 is a trademark of Compaq Information Technologies Group, L.P. in the United States and other countries. UNIX is a trademark of The Open Group in the United States and other countries. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq 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.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.