

Planning for Network Deployment in Oracle® Solaris 11.2

ORACLE®

Part No: E37474
July 2014

Copyright © 2011, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2011, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Using This Documentation	5
1 Planning For Network Deployment	7
Determining the Network Hardware	7
Network Topology Overview	8
Using Subnets on Your Network	10
IPv4 Autonomous System Topology	11
Planning for Routers on Your Network	12
How Routers Transfer Packets	13
Deciding on an IP Addressing Format for Your Network	14
IPv4 Addresses	15
Private Addresses	16
DHCP Addresses	16
IPv6 Addresses	17
Documentation Prefixes	17
Obtaining Your Network's IP Number	17
Using Naming Entities on Your Network	18
Domain Names	18
Selecting a Naming Service and Directory Service	18
Administering Host Names	19
2 Planning for Using IPv6 Addresses	21
IPv6 Planning Tasks	21
IPv6 Network Topology Overview	22
Ensuring Hardware Support for IPv6	24
Preparing an IPv6 Addressing Plan	24
Obtaining a Site Prefix	25
Creating the IPv6 Numbering Scheme	25
Configuring Network Services to Support IPv6	26
▼ How to Prepare Network Services for IPv6 Support	27

▼ How to Prepare DNS for IPv6 Support	28
Planning for Tunnel Use in the Network	28
Security Considerations for the IPv6 Implementation	29
Index	31

Using This Documentation

- **Overview** – Includes basic topics and tasks to assist you in planning for deploying IPv4 and IPv6 networks.
- **Audience** – System administrators.
- **Required knowledge** – Basic understanding of network administration concepts and practices.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E36784>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Planning For Network Deployment

This chapter describes the different considerations when planning for deployment of a TCP/IP network. The planning tasks that are described can assist you in deploying your network in an organized and cost-effective manner. Note that the details of planning the network are outside the scope of this book. Only general directions are provided. This book also assumes that you are familiar with basic networking concepts and terminology.

For a description of how TCP/IP is implemented in Oracle Solaris and an overview of network administration in this release, see [Chapter 1, “About Network Administration in Oracle Solaris,”](#) in “Configuring and Administering Network Components in Oracle Solaris 11.2”.

For further assistance with planning your site's overall networking scheme, review the networking strategies that are described in [Chapter 1, “Summary of Oracle Solaris Network Administration,”](#) in “Strategies for Network Administration in Oracle Solaris 11.2”.

This chapter contains the following topics:

- [“Determining the Network Hardware”](#) on page 7
- [“Network Topology Overview”](#) on page 8
- [“Using Subnets on Your Network”](#) on page 10
- [“IPv4 Autonomous System Topology”](#) on page 11
- [“Planning for Routers on Your Network”](#) on page 12
- [“Deciding on an IP Addressing Format for Your Network”](#) on page 14
- [“Obtaining Your Network's IP Number”](#) on page 17
- [“Using Naming Entities on Your Network”](#) on page 18

Determining the Network Hardware

The number of systems that you expect to support affects how you configure your network. Your organization might require a small network of several dozen standalone systems that are located on one floor of a single building. Or, you might need to set up a network with more than 1,000 systems in several buildings. This setup can require you to further divide your network into subdivisions that are called *subnets*.

Some of the planning decisions that you must make about hardware include the following:

- Network topology, the layout, and connections of the network hardware
- Type and number of host systems your network can support, including the virtual systems that might be required on your server
- Network devices to be installed in these systems
- Type of network media to use, such as Ethernet, and so on
- Use of bridges, routers, and firewalls to extend the network media or connect the local network to external networks

For information about how bridges work, see [“Overview of Bridged Networks” in “Managing Network Datalinks in Oracle Solaris 11.2”](#).

For a description of how routers function, see [“Planning for Routers on Your Network” on page 12](#).

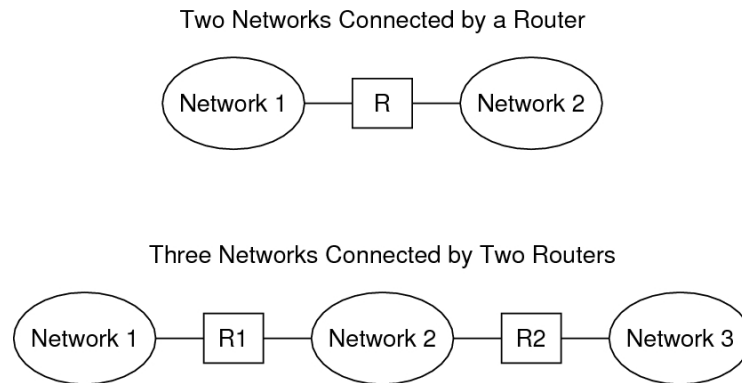
For information about firewalls, see [Chapter 4, “About IP Filter in Oracle Solaris,” in “Securing the Network in Oracle Solaris 11.2”](#).

Network Topology Overview

Network topology describes how networks fit together. Routers are the entities that connect networks to each other. A router is any machine that has two or more network interfaces and implements IP forwarding. However, the system cannot function as a router until properly configured, as described in [Chapter 2, “Configuring a System as a Router,” in “Configuring an Oracle Solaris 11.2 System as a Router or a Load Balancer”](#).

Routers connect two or more networks to form larger internetworks. You must configure the routers to pass packets between two adjacent networks. The routers also should be able to pass packets to networks that lie beyond the adjacent networks by forwarding them to other routers that are connected to adjacent networks.


The following figure shows the basic parts of a network topology. The top portion of the illustration shows a simple configuration of two networks that are connected by a single router. The bottom portion of the illustration shows a configuration of three networks, interconnected by two routers. In the first example, Router R joins Network 1 and Network 2 into a larger internetwork. In the second example, Router R1 connects Networks 1 and 2. Router R2 connects Networks 2 and 3. The connections form a network that includes Networks 1, 2, and 3.

FIGURE 1-1 Basic Network Topology

In addition to joining networks into internetworks, routers route packets between networks that are based on the addresses of the destination network. As internetworks grow more complex, each router must make more and more decisions about the packet destinations.

The following figure shows a more complex case. Router R3 directly connects networks 1 and 3. The redundancy improves reliability. If network 2 goes down, router R3 still provides a route between networks 1 and 3. You can interconnect many networks. However, the networks must use the same network protocols.

FIGURE 1-2 A Network Topology That Provides an Additional Path Between Networks



Routers are discussed in more detail in [“Planning for Routers on Your Network”](#) on page 12.

Using Subnets on Your Network

The use of subnets is connected with the need for administrative subdivisions to address issues of size and control. The more hosts and servers that you have on a network, the more complex your management task. By creating administrative divisions and using subnets, managing complex networks becomes easier.

The decision about setting up administrative subdivisions for your network is determined by the following factors:

- **Size of the network**

Subnets are also useful even in a relatively small network whose subdivisions are located across an extensive geographical area.

- **Common needs shared by groups of users**

For example, you might have a network that is confined to a single building and supports a relatively small number of machines. These machines are divided among a number of subnetworks. Each subnetwork supports groups of users with different needs. In this example, you might use an administrative subdivision for each subnet.

- **Security**

You might want to segregate your mission critical servers, desktop systems, and Internet facing web servers into separate subnets where you can establish firewalls between them.

IPv4 Autonomous System Topology

Sites with multiple routers and networks typically administer their network topology as a single routing domain or an *autonomous system* (AS). [Figure 1-3](#) shows an AS that is divided into three local networks: `10.0.5.0`, `172.20.1.0`, and `192.168.5.0`.

The network is comprised of the following types of systems:

- **Routers**

Routers use routing protocols to manage how network packets are directed or routed from their source to their destinations within the local network or to external networks. For information about the routing protocols that are supported in Oracle Solaris and instructions on configuring a system as a router, see [“Routing Protocols”](#) in [“Configuring an Oracle Solaris 11.2 System as a Router or a Load Balancer”](#).


Types of routers include the following:

- **Border routers** – Connect the local network, such as `10.0.5.0`, externally to a service provider.
- **Default routers** – Manage packet routing in the local network, which itself can include several local networks. For example, in [Figure 1-3](#), Router 1 serves as the default router for `192.168.5`. Contemporaneously, Router 1 is also connected to the `10.0.5.0` internal network. Router 2's interfaces connect to the `10.0.5.0` and `172.20.1.0` internal networks.
- **Packet-forwarding routers** – Forward packets between internal networks but do not run routing protocols. In [Figure 1-3](#), Router 3 is a packet-forwarding router with connections to the `172.20.1` and `192.168.5` networks.
- **Client systems**
 - **Multihomed systems** or systems that have multiple NICs. In Oracle Solaris, these systems by default can forward packets to other systems in the same network segment.
 - **Single-interfaced systems** rely on the local routers for both packet forwarding and receiving configuration information.

For task-related information, see [Chapter 3, “Configuring and Administering IP Interfaces and Addresses in Oracle Solaris,”](#) in [“Configuring and Administering Network Components in Oracle Solaris 11.2”](#).

Use the following figure as a reference when configuring additional network components.

FIGURE 1-3 Autonomous System With Multiple IPv4 Routers



Planning for Routers on Your Network

Recall that in TCP/IP, two types of entities exist on a network: hosts and routers. All networks must have hosts, while not all networks require routers. The physical topology of the network determines if you need routers. This section introduces the concepts of network topology and routing. These concepts are important when you decide to add another network to your existing network environment.

Note - For complete details and tasks for router configuration on IPv4 and IPv6 networks, see [Chapter 2, “Configuring a System as a Router,”](#) in [“Configuring an Oracle Solaris 11.2 System as a Router or a Load Balancer ”](#).

How Routers Transfer Packets

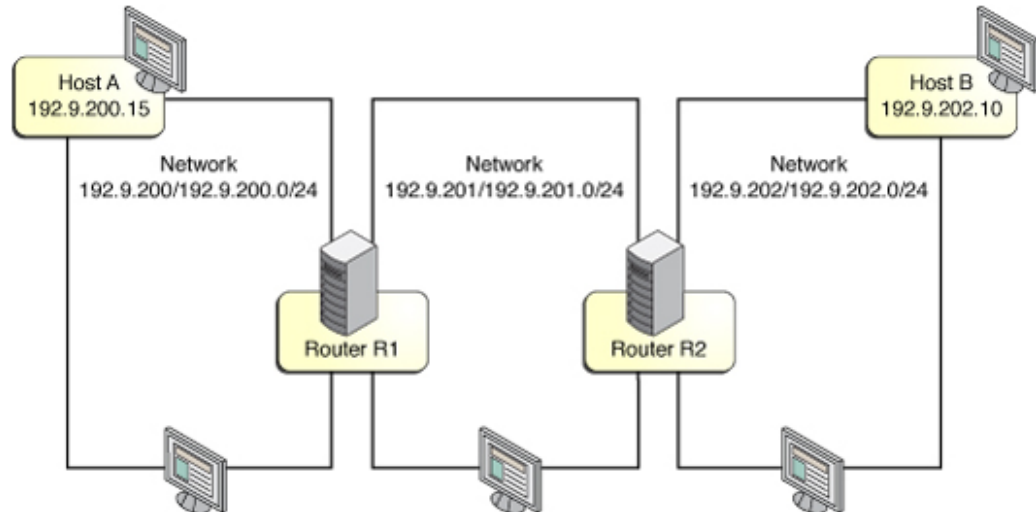
Routers transfer packets in the following manner:

- All nodes on an IP network maintain routing information in routing tables. These tables contain information of how to reach systems that are attached to both local and remote networks. The routing tables are generated from local configuration information and from routing protocol messages that is exchanged with neighboring systems.
- When a host system initially sends a packet, it looks up the packet's destination address in its routing table to determine if the destination is on the local network. If yes, the packet goes directly to the host with that IP address. If not, the packet goes to a router on the local network.
- When a router receives a packet, the router checks its routing table to determine if the destination address is for a system on one of its attached networks or if the message must be forwarded through another router. It then sends the message to the next system in the path to the destination.
- This process is repeated on each router that receives the message until the message reaches the destination system.

Refer to [Chapter 2, “Configuring a System as a Router,”](#) in [“Configuring an Oracle Solaris 11.2 System as a Router or a Load Balancer ”](#).

The following figure shows a network topology with three networks that are connected by two routers.

FIGURE 1-4 A Network Topology With Three Interconnected Networks



Router R1 connects networks 192.9.200.0/24 and 192.9.201.0/24. Router R2 connects networks 192.9.201.0/24 and 192.9.202.0/24.

If Host A on network 192.9.200.0/24 sends a message to Host B on network 192.9.202, the following events occur:

1. Host A examines its routing tables for the path to 192.9.202.10. The local network address range does not cover this address, but there is a previously learned default route through router R1 that covers the address. Therefore, Host A sends the packet to Router R1.
2. Router R1 examines its routing tables. No local network's address range covers the destination address, but there is a known route to network 192.9.202.0/24 through Router R2 that covers the address, Router R1 sends the packet to Router R2.
3. Router R2 is connected directly to network 192.9.202.0/24. The routing table lookup reveals that 192.9.202.10 is on the attached network. Router R2 sends the packet directly to Host B.

Deciding on an IP Addressing Format for Your Network

When you plan your network addressing scheme, consider the following factors:

- Type of IP address that you want to use: IPv4 or IPv6

- Number of potential systems on your network
- Number of systems that are multihomed, or routers, which require multiple network interface cards (NICs), each with individual IP addresses
- Whether to use private addresses on your network
- Whether to have a DHCP server that manages pools of IP addresses

IPv4 Addresses

This is the original address format that is used on TCP/IP networks. IPv4 addresses are 32 bits in length. IPv4 addresses were originally allocated to various organizations in contiguous blocks of 16777216 (Class A), 65536 (Class B), or 256 addresses (Class C) addresses. Each organization that requested an address block received a fixed address prefix and an implied prefix mask, both specified in dotted decimal notation. For example, the Internet Assigned Numbers Authority (IANA) allocated the Class A address block 156.0.0.0 netmask 255.0.0.0 to the American Registry for Internet Numbers (ARIN). All addresses whose first byte equals 156 are within this address block. ARIN sub-allocated the Class B address block 156.151.0.0 netmask 255.255.0.0 from its Class A block to Sun Microsystems (now Oracle).

Later, The Internet Engineering Task Force (IETF) developed *Classless Inter-Domain Routing* (CIDR) addresses as an interim remedy for the shortage of IPv4 addresses and the limited capacity of the global Internet routing tables. CIDR address allocations are subdivided on whatever bit boundary best meets an organization's requirements. Address blocks are specified as a dotted decimal IPv4 address followed by a slash and the address prefix length in bits.

For more information, refer to the following resources:

- [Internet Protocol DARPA Internet Program Protocol Specification \(http://tools.ietf.org/html/rfc791\)](http://tools.ietf.org/html/rfc791)
- [Classless Inter-domain Routing \(CIDR\): The Internet Address Assignment and Aggregation Plan \(http://tools.ietf.org/html/rfc4632\)](http://tools.ietf.org/html/rfc4632)

The following table provides example subnet length specifications in both CIDR notation and dotted decimal format, as well as the total number of hosts that are possible on a network with that prefix length.

TABLE 1-1 CIDR Prefixes and Their Decimal Equivalents

CIDR Network Prefix Length	Corresponding Dotted Decimal Subnet Mask	Available IP Addresses
/19	255.255.224.0	8,192
/20	255.255.240.0	4,096
/21	255.255.248.0	2,048
/22	255.255.252.0	1,024

CIDR Network Prefix Length	Corresponding Dotted Decimal Subnet Mask	Available IP Addresses
/23	255.255.254.0	512
/24	255.255.255.0	256
/25	255.255.255.128	128
/26	255.255.255.192	64
/27	255.255.255.224	32

Private Addresses

The IANA has reserved a block of IPv4 addresses. These private addresses are used for network traffic within a private network. Organizations that request a block of IPv4 addresses from their Internet Service Provider are unlikely to receive a sufficient allocation to have unique address for each of their systems. Organizations typically assign private addresses to the systems on their internal networks. Their systems can effectively share the limited Internet Service Provider (ISP) provided addresses by using Network Address Translators (NAT) and application proxy servers to communicate with other sites on the Internet.

The following table lists the private IPv4 address ranges and their corresponding netmask.

Network Prefix/Length	IPv4 Address Range
10.0.0.0/8	10.0.0.0 - 10.255.255.255
172.16.0.0/125	172.16.0.0 - 172.31.255.25
192.168.0.0/16	192.168.0.0 - 192.168.255.255

DHCP Addresses

The Dynamic Host Configuration Protocol (DHCP) protocol enables a system to receive configuration information from a DHCP server, including an IP address, as part of the boot process. DHCP servers maintain pools of IP addresses from which to assign addresses to DHCP clients. A site that uses DHCP can use a smaller pool of IP addresses than the number of IP addresses that would be required if every client was assigned a permanent IP address, provided that the clients are not continuously connected. In this case, you can share addresses between clients, and thus reduce the total required number of IP addresses. However, if you do not have a sufficient coming and going of clients, you ultimately will require the same number of IP addresses. The more general advantage of using DHCP addresses is that you do not need to perform as much configuration of individual hosts because you set up a DHCP server with the configuration details. In this way, hosts require very minimal or even no manual

configuration. You can set up the DHCP service to manage your site's IP addresses or a portion of the addresses. For more information, refer to “[Working With DHCP in Oracle Solaris 11.2](#)”.

IPv6 Addresses

These 128-bit IPv6 addresses provide greater address space than is available with IPv4. IPv6 addresses are presented as eight groups of four hexadecimal digits, each group separated by a colon. Leading zeros in each group can be suppressed. One or more consecutive groups of all zeros can be replaced by a double colon, as shown in the following example:

```
2001:db8:2f32:27:214:4fff:fe4a:9926
```

As with IPv4 addresses in CIDR format, IPv6 addresses are classless and use prefixes to designate the portion of the address that defines the site's network, as shown in the following example:

```
2001:db8:2f32::/48
```

For details about IPv6 addressing, see [IP Version 6 Addressing Architecture \(http://tools.ietf.org/html/rfc4291\)](http://tools.ietf.org/html/rfc4291).

Documentation Prefixes

For IPv6 addresses, the prefix `2001:db8::/32` is a special IPv6 prefix that is used specifically for documentation examples. The examples in this book use private IPv4 addresses and the reserved IPv6 documentation prefix.

Obtaining Your Network's IP Number

An IPv4 network is defined by a combination of an IPv4 network number plus a network mask, or *netmask*. An IPv6 network is defined by its *site prefix*, and its *subnet prefix*, if subnets are used.

To enable the private network to communicate to external networks in the Internet, you must obtain a registered IP number for your network from the appropriate organization. This address becomes the network number for your IPv4 addressing scheme or the site prefix for your IPv6 addressing scheme.

ISPs provide IP addresses for networks with pricing that is based on different levels of service. Investigate with various ISPs to determine which provides the best service for your network. ISPs typically offer dynamically allocated addresses or static IP addresses to businesses. Some ISPs offer both IPv4 and IPv6 addresses.

If your site is an ISP, you obtain IP address blocks for your customers from the Internet Registry (IR) for your locale. IANA is ultimately responsible for delegating registered IP addresses to IRs around the world. Each IR has registration information and templates for the locale that the IR services. For information about the IANA and its IRs, refer to the [IANA's IP Address Service page \(http://www.iana.org/ipaddress/ip-addresses.htm\)](http://www.iana.org/ipaddress/ip-addresses.htm).

Using Naming Entities on Your Network

The TCP/IP protocols locate a system on a network by using its IP address. However, a host name enables you to identify systems more easily than IP addresses.

From a TCP/IP perspective, a network is a set of named entities. A host is an entity with a name. A router is an entity with a name. The network is an entity with a name. A group or department in which the network is installed can also be given a name, as can a division, a region, or a company. In theory, the hierarchy of names that can be used to identify a network has virtually no limit.

Domain Names

Many networks organize their hosts and routers into a hierarchy of administrative domains. If you are using the Network Information Service (NIS) or the Domain Name System (DNS) naming service, you must select a domain name for your organization that is unique worldwide. To ensure that your domain name is unique, you should register the domain name with *InterNIC*. A unique domain name is required if you plan to allow other sites on the Internet to locate your systems through DNS.

A domain name that is located below another domain is often referred to as a sub-domain. The domain name structure is hierarchical. A new domain typically is located below an existing, related domain. For example, the domain name for a subsidiary company can be located below the domain of the parent company. If the domain name has no other relationship, an organization can place its domain name directly under one of the existing top-level domains such as .com, .org, .edu, .gov, and so forth.

Selecting a Naming Service and Directory Service

In Oracle Solaris you can select from three types of naming services: local files, NIS, and DNS. Naming services maintain critical information about the machines on a network, such as the host names, IP addresses, and so forth. You can also use the LDAP directory service in addition to or instead of a naming service. LDAP is a secure network protocol that is used to access directory servers for distributed naming and other directory services. This standard based protocol supports a hierarchal database structure. The same protocol can be used to provide

naming services in both UNIX and multi-platform environments. For an introduction to naming services in Oracle Solaris, refer to [Chapter 1, “About Naming and Directory Services,”](#) in [“Working With Oracle Solaris 11.2 Directory and Naming Services: DNS and NIS ”](#).

The configuration of the network databases is critical. Therefore, you need to decide which naming or directory service to use as part of the network planning process. Moreover, the decision to use naming services also affects whether you organize your network into an administrative domain.

For a naming or directory service, you can select from the following:

- NIS or DNS – The NIS and DNS naming services maintain network databases on several servers on the network. [“Working With Oracle Solaris 11.2 Directory and Naming Services: DNS and NIS ”](#) describes these naming services and explains how to configure the databases. In addition, the guide explains the *namespace* and *administrative domain* concepts in more detail.
- LDAP – You can also use the LDAP directory service in addition to or instead of a naming service. LDAP is a secure network protocol that is used to access directory servers for distributed naming and other directory services.
- Local files – If you do not implement NIS, DNS, or LDAP, the network uses *local files* to provide the naming service. The term “local files” refers to the series of files in the `/etc` directory that the network databases use. The procedures in this book assume you are using local files for your naming service, unless otherwise indicated.

Note - If you decide to use local files as the naming service for your network, you can set up another naming service at a later date.

Administering Host Names

Plan a naming scheme for the systems that will comprise the network. Each machine on the network should have a TCP/IP host name that corresponds to the IP address on its primary network interface. The host name must be unique within the system's sub-domain. Just like physical machines, virtual systems should also have a unique IP address and host name.

A system can have the following:

- Multiple host names that map to the system's IP address. For example, `systema.mycompany.com` can also be known as `www.mycompany.com`.
- The same host name for both IPv4 and an IPv6 addresses.
- A new IP address and an old deprecated IP address that are configured with the same host name for a period of time to support network renumbering.
- Multiple network interfaces on different subnets, each with a unique IP address and host name.

When planning your network, make a list of IP addresses and their associated host names for easy access during the setup process. The list can help you verify that all host names are unique.

Note - The primary interface's TCP/IP host name is a distinct entity from the *system host name* that you set with the `hostname` command. Although not required by Oracle Solaris, the same name is normally used for both. Some network applications depend on this convention. See the [hostname\(1\)](#) man page for more information.

Planning for Using IPv6 Addresses

This chapter supplements [Chapter 1, “Planning For Network Deployment”](#) by describing additional considerations if you decide to use IPv6 addresses on your network.

If you do plan to use IPv6 addresses in addition to IPv4 addresses, ensure that your current ISP supports both address types

For an introduction to IPv6 concepts, refer to [Internet Protocol, Version 6 \(IPv6\) Specification \(http://www.ietf.org/rfc/rfc2460.txt\)](http://www.ietf.org/rfc/rfc2460.txt).

For IPv6 configuration tasks, see [“Configuring IPv6 Interfaces”](#) in [“Configuring and Administering Network Components in Oracle Solaris 11.2 ”](#).

For information about troubleshooting IPv6 networks, see [“Troubleshooting Issues With IPv6 Deployment”](#) in [“Troubleshooting Network Administration Issues in Oracle Solaris 11.2 ”](#).

This chapter contains the following topics:

- [“IPv6 Planning Tasks”](#) on page 21
- [“IPv6 Network Topology Overview”](#) on page 22
- [“Ensuring Hardware Support for IPv6”](#) on page 24
- [“Preparing an IPv6 Addressing Plan”](#) on page 24
- [“Configuring Network Services to Support IPv6”](#) on page 26
- [“Planning for Tunnel Use in the Network”](#) on page 28
- [“Security Considerations for the IPv6 Implementation”](#) on page 29

IPv6 Planning Tasks

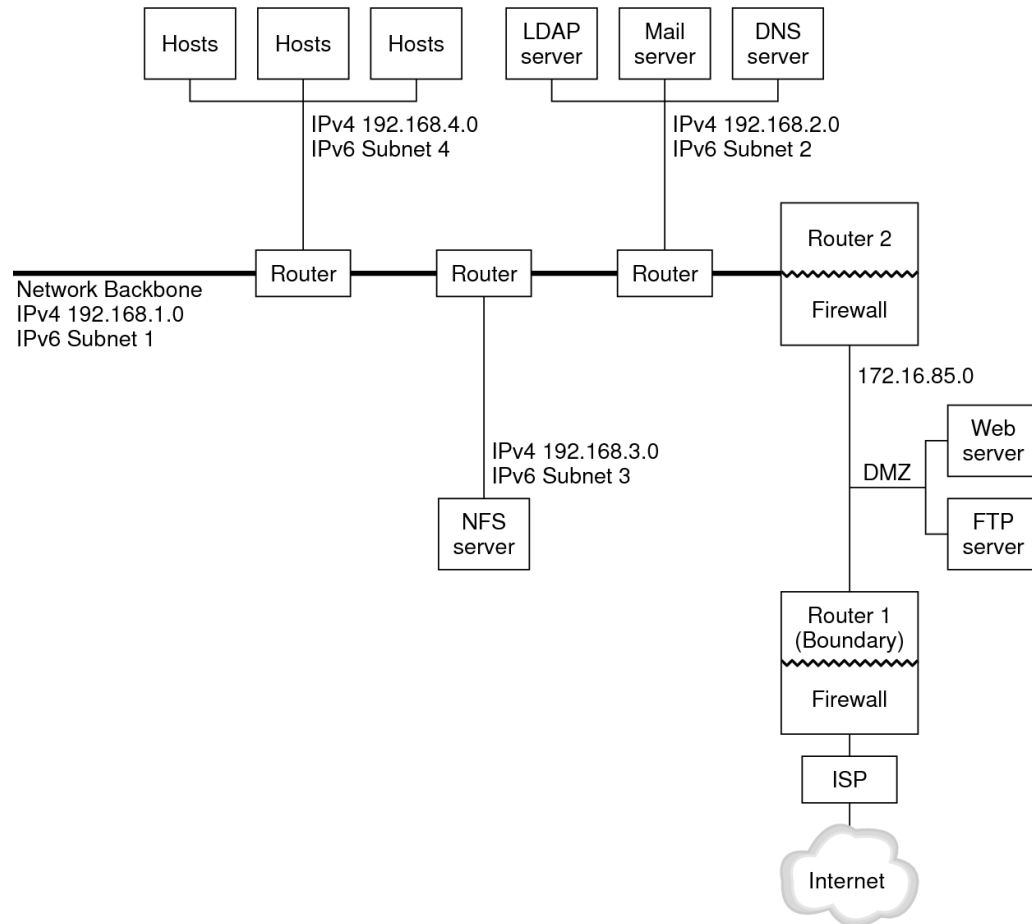
The following table describes different considerations when planning to implement IPv6 on your network. If you are migrating from an existing IPv4 network to an IPv6 network, see [“Migrating From an IPv4 Network to an IPv6 Network”](#) in [“Configuring and Administering Network Components in Oracle Solaris 11.2 ”](#) for additional instructions.

Task	Description	For Instructions
Prepare your hardware to support IPv6.	Ensure that your hardware can be upgraded to IPv6.	“Ensuring Hardware Support for IPv6” on page 24
Ensure that your applications are IPv6 ready.	Verify that your applications can run in an IPv6 environment.	“Configuring Network Services to Support IPv6” on page 26
Design a plan for tunnel usage.	Determine which routers should run tunnels to other subnets or external networks.	“Planning for Tunnel Use in the Network” on page 28
Plan how to secure your networks and develop an IPv6 security policy.	For security purposes, you need an addressing plan for the DMZ and its entities before you configure IPv6. Decide how you would implement security, such as using IP Filter, IPsec security architecture (IPsec), Internet Key Exchange (IKE), and other security features of this release.	“Security Considerations for the IPv6 Implementation” on page 29 “Securing the Network in Oracle Solaris 11.2 ”
Create an addressing plan for systems on the network.	Your plan for addressing servers, routers, and hosts should be in place before IPv6 configuration. This step includes obtaining a site prefix for your network as well as planning IPv6 subnets, if needed.	“Creating an IPv6 Addressing Plan for Nodes” on page 25

IPv6 Network Topology Overview

Typically, IPv6 is used in a mixed network topology that also uses IPv4, such as shown in the following figure. The following figure is used as reference in the description of IPv6 configuration tasks that are described in this chapter.

FIGURE 2-1 IPv6 Network Topology Scenario



The enterprise network scenario depicted in the figure consists of five subnets with existing IPv4 addresses. The links of the network correspond directly to the administrative subnets. The four internal networks are shown with RFC 1918-style private IPv4 addresses, which is a common solution for the lack of IPv4 addresses.

These internal networks use the following address scheme:

- Subnet 1 is the internal network backbone 192.168.1.
- Subnet 2 is the internal network 192.168.2, with LDAP, sendmail, and DNS servers.
- Subnet 3 is the internal network 192.168.3, with the enterprise's NFS servers.

- Subnet 4 is the internal network 192.168.4, which contains hosts for the enterprise's employees.

The external, public network 172.16.85 functions as the corporation's DMZ. This network contains web servers, anonymous FTP servers, and other resources that the enterprise offers to the outside world. Router 2 runs a firewall and separates public network 172.16.85 from the internal backbone. On the other end of the Demilitarized Zone (DMZ), Router 1 runs a firewall and serves as the enterprise's boundary server.

In [Figure 2-1](#), the public DMZ has the RFC 1918 private address 172.16.85. In the real world, the public DMZ must have a registered IPv4 address. Most IPv4 sites use a combination of public addresses and RFC 1918 private addresses. However, when you introduce IPv6, the concept of public addresses and private addresses changes. Because IPv6 has a much larger address space, you use public IPv6 addresses on both private networks and public networks.

The Oracle Solaris dual protocol stack supports concurrent IPv4 and IPv6 operations. You can successfully run IPv4-related operations during and after deploying IPv6 on your network. When you deploy IPv6 on an operating network that is already using IPv4, ensure that you do not disrupt ongoing operations.

Ensuring Hardware Support for IPv6

Check the manufacturers' documentation for IPv6 readiness regarding the following classes of hardware:

- Routers
- Firewalls
- Servers
- Switches

Note - All of the procedures in the this book assume that your equipment, particularly routers, can be upgraded to IPv6. However, some router models cannot be upgraded to IPv6. For more information and a workaround, refer to [“Cannot Upgrade IPv4 Router to IPv6”](#) in [“Troubleshooting Network Administration Issues in Oracle Solaris 11.2”](#).

Preparing an IPv6 Addressing Plan

A major part transitioning from IPv4 to IPv6 includes developing an addressing plan, which involves the following preparations:

- “Obtaining a Site Prefix” on page 25
- “Creating the IPv6 Numbering Scheme” on page 25

For actual migration tasks, see “Migrating From an IPv4 Network to an IPv6 Network” in “Configuring and Administering Network Components in Oracle Solaris 11.2”.

Obtaining a Site Prefix

Before you configure IPv6, you must obtain a site prefix. The site prefix is used to derive IPv6 addresses for all the nodes in your IPv6 implementation.

Any ISP that supports IPv6 can provide your organization with a 48-bit IPv6 site prefix. If your current ISP only supports IPv4, you can use another ISP for IPv6 support while retaining your current ISP for IPv4 support. In such an instance, you can use one of several workarounds. For more information, see “Current ISP Does Not Support IPv6” in “Troubleshooting Network Administration Issues in Oracle Solaris 11.2”.

If your organization is an ISP, then you obtain site prefixes for your customers from the appropriate Internet registry. For more information, see the [Internet Assigned Numbers Authority \(IANA\) \(http://www.iana.org\)](http://www.iana.org).

Creating the IPv6 Numbering Scheme

Unless your proposed IPv6 network is entirely new, use your existing IPv4 topology as the basis for the IPv6 numbering scheme.

Creating an IPv6 Addressing Plan for Nodes

For most hosts, stateless autoconfiguration of IPv6 addresses for their interfaces is an appropriate, time saving strategy. When the host receives the site prefix from the nearest router, Neighbor Discovery automatically generates IPv6 addresses for each interface on the host.

Servers need to have stable IPv6 addresses. If you do not manually configure a server's IPv6 addresses, a new IPv6 address is autoconfigured whenever a NIC card is replaced on the server.

Keep the following tips in mind when you create addresses for servers:

- Give servers meaningful and stable interface IDs. One strategy is to use a sequential numbering scheme for interface IDs. For example, the internal interface of the LDAP server in [Figure 2-1](#) might become 2001:db8:3c4d:2::2.

- Alternatively, if you do not regularly renumber your IPv4 network, consider using the existing IPv4 addresses of the routers and servers as their interface IDs. In [Figure 2-1](#), suppose Router 1's interface to the DMZ has the IPv4 address 192.168/16. You can convert the IPv4 address to hexadecimal and use the result as the interface ID. The new interface ID would be ::7bc8:156F.

Only use this approach if you own the registered IPv4 address, rather than having obtained the address from an ISP. If you use an IPv4 address that was provided to you by an ISP, you create a dependency that would create problems if you change ISPs.

Due to the limited number of IPv4 addresses that are available, in the past, a network designer had to consider where to use global, registered addresses and private, RFC 1918 addresses. However, the notion of global and private IPv4 addresses does not apply to IPv6 addresses. You can use *global unicast addresses*, which include the site prefix, on all links of the network, including the public DMZ.

Creating a Numbering Scheme for Subnets

Begin your numbering scheme by mapping your existing IPv4 subnets into equivalent IPv6 subnets. For example, consider the subnets that are shown in [Figure 2-1](#). Subnets 1–4 use the RFC 1918 IPv4 private address designation for the first 16 bits of their addresses, in addition to the digits 1–4 to indicate the subnet. For illustrative purposes, assume that the IPv6 prefix 2001:db8:3c4d/48 has been assigned to the site.

The following table shows how the private IPv4 prefixes map into IPv6 prefixes.

IPv4 Subnet Prefix	Equivalent IPv6 Subnet Prefix
192.168.1.0/24	2001:db8:3c4d:1::/64
192.168.2.0/24	2001:db8:3c4d:2::/64
192.168.3.0/24	2001:db8:3c4d:3::/64
192.168.4.0/24	2001:db8:3c4d:4::/64

Configuring Network Services to Support IPv6

The following typical IPv4 network services are also IPv6 ready:

- DNS
- HTTP (Apache 2 releases or Orion)
- LDAP

- NFS
- sendmail

The IMAP mail service is for IPv4 only.

Nodes that are configured for IPv6 can run IPv4 services. When you turn on IPv6, not all services accept IPv6 connections. Services that have been ported to IPv6 will accept a connection. Services that have not been ported to IPv6 continue to work with the IPv4 portion of the protocol stack.

Some issues can arise after you upgrade services to IPv6. For details, see [“Problems Encountered When Upgrading Services to Support IPv6”](#) in [“Troubleshooting Network Administration Issues in Oracle Solaris 11.2”](#).

▼ How to Prepare Network Services for IPv6 Support

1. Update the following network services to support IPv6:

- Mail servers
- NIS servers
- NFS

Note - LDAP supports IPv6 without requiring IPv6-specific configuration tasks.

2. Verify that your firewall hardware is IPv6 ready.

Refer to the appropriate firewall-related documentation for instructions.

3. Verify that other services on your network have been ported to IPv6.

For more information, refer to marketing collateral and associated documentation for the software.

4. If your site deploys the following services, make sure that you have taken the appropriate measures for these services:

- **Firewalls** – Consider strengthening the policies that are in place for IPv4 to support IPv6. For more security considerations, see [“Security Considerations for the IPv6 Implementation”](#) on page 29.
- **Mail**– In the mail exchanger record (MX record) for DNS, consider adding the IPv6 address of your mail server.
- **DNS**– For DNS-specific considerations, see [“How to Prepare DNS for IPv6 Support”](#) on page 28.

- **IPQoS** – Use the same *Diffserv* policies on a host that were used for IPv4.
5. **Audit any network services that are offered by a node prior to converting that node to IPv6.**

▼ How to Prepare DNS for IPv6 Support

Oracle Solaris supports DNS resolution on both the client side and the server side. Use the following procedure to prepare DNS services for IPv6.

For more information that is related to DNS support for IPv6, refer to [“Working With Oracle Solaris 11.2 Directory and Naming Services: DNS and NIS”](#).

1. **Ensure that the DNS server that performs recursive name resolution is dual-stacked (IPv4 and IPv6) or for IPv4 only.**
2. **On the DNS server, populate the DNS database with relevant IPv6 database AAAA records in the forward zone.**

Note - Servers that run multiple critical services require special attention. Ensure that the network is working properly. Also ensure that all critical services are ported to IPv6. Then, add the server's IPv6 address to the DNS database.

3. **Add the associated PTR records for the AAAA records into the reverse zone.**
4. **Add either IPv4 only data, or both IPv6 and IPv4 data into the NS record that describes zones.**

Planning for Tunnel Use in the Network

The IPv6 implementation supports a number of tunnel configurations to serve as transition mechanisms as your network migrates to a mix of IPv4 and IPv6. Tunnels enable isolated IPv6 networks to communicate. Because most of the Internet runs IPv4, IPv6 packets from your site need to travel across the Internet through tunnels to destination IPv6 networks.

Here are some major scenarios for using tunnels in the IPv6 network topology:

- The ISP from which you purchase IPv6 service allows you to create a tunnel from your site's boundary router to the ISP network. [Figure 2-1](#) shows such a tunnel. In such a case, you would run a manual, IPv6 over IPv4 tunnel.

- You manage a large, distributed network with IPv4 connectivity. To connect the distributed sites that use IPv6, you can run an automatic 6to4 tunnel from the edge router of each subnet.
- Sometimes, a router in your infrastructure cannot be upgraded to IPv6. In this case, you can create a manual tunnel over the IPv4 router, with two IPv6 routers as endpoints.

For procedures for configuring tunnels, refer to [Chapter 5, “Administering IP Tunnels,”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2 ”](#). For conceptual information regarding tunnels, refer to [“IP Tunnel Feature Summary”](#) in [“Administering TCP/IP Networks, IPMP, and IP Tunnels in Oracle Solaris 11.2 ”](#).

Security Considerations for the IPv6 Implementation

When you introduce IPv6 into an existing network, you must take care not to compromise the security of the site.

Be aware of the following security issues as you phase in your IPv6 implementation:

- The same amount of filtering is required for both IPv6 packets and IPv4 packets.
- IPv6 packets are often tunneled through a firewall.
Therefore, you should implement either of the following scenarios:
 - Have the firewall do content inspection inside the tunnel.
 - Put an IPv6 firewall with similar rules at the opposite tunnel endpoint.
- Some transition mechanisms exist that use IPv6 over User Datagram Protocol (UDP) over IPv4 tunnels. These mechanisms might prove problematic by short-circuiting the firewall.
- IPv6 nodes are globally reachable from outside the enterprise network. If your security policy prohibits public access, you must establish stricter rules for the firewall. For example, consider configuring a *stateful firewall*.

This book includes information about following security features, which can be used within an IPv6 implementation:

- The IP security architecture (IPsec) feature enables you to provide cryptographic protection for IPv6 packets. For more information, refer to [Chapter 6, “About IP Security Architecture,”](#) in [“Securing the Network in Oracle Solaris 11.2 ”](#).
- The Internet Key Exchange (IKE) feature automates keys management for IPsec. For more information, refer to [Chapter 8, “About Internet Key Exchange,”](#) in [“Securing the Network in Oracle Solaris 11.2 ”](#).

Index

A

autonomous system (AS) *See* network topology

B

border router, 11

C

CIDR notation, 15

class A, B, and C network numbers, 15

D

default router

definition, 11

designing the network

domain name selection, 18

IP addressing scheme, 14

naming hosts, 19

domain name system (DNS)

preparing, for IPv6 support, 28

selecting as naming service, 19

domain names

selecting, 18

H

hosts

host name

administering, 19

I

internetworks

definition, 8

packet transfer by routers, 13

redundancy and reliability, 9

topology, 8, 9

IP addresses

CIDR notation, 15

designing an address scheme, 14

network classes

network number administration, 15

IPQoS

policies for IPv6-enabled networks, 28

IPv6

addressing plan, 25

DNS support preparation, 28

security considerations, 29

L

local files

selecting as naming service, 19

M

multihomed systems

definition, 11

N

naming services

selecting, 18

network administration

host names, 19

network planning

adding routers, 12

IP addressing scheme, 14

- registering your network, 17
- network topology, 8, 9
 - autonomous system, 12
- NIS
 - selecting as naming service, 19

P

- packet forwarding router, 11
- packets
 - transfer
 - router, 13

R

- registering
 - networks, 17
- routers
 - adding, 12
 - network topology, 8, 9
 - packet forwarding router, 11
 - packet transfer, 13

S

- security considerations
 - IPv6-enabled networks, 29
- site prefix, IPv6
 - how to obtain, 25
- subnets, 10
 - IPv6
 - suggestions for numbering, 26

T

- task maps
 - IPv6
 - planning, 21
- topology, 8, 9
- tunnels
 - planning, for IPv6, 28