TCP/IP

Connectivity Services

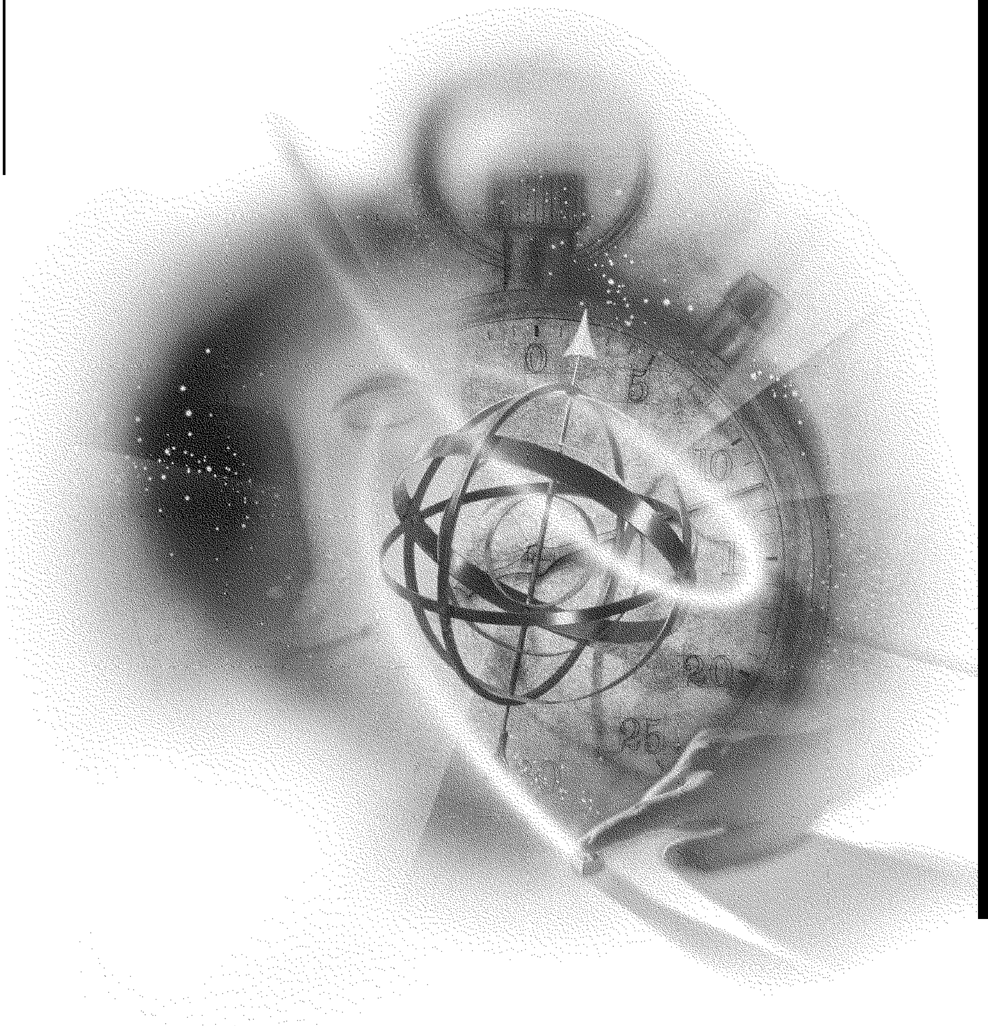

Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

This product may require export authorization from the U.S. Department of Commerce prior to exporting from the U.S. or Canada.

U.S. Patent Nos. 4,555,775; 5,157,663; 5,349,642; 5,455,932; 5,553,139; 5,553,143; 5,594,863; 5,608,903; 5,633,931; 5,652,854; 5,671,414; 5,677,851; 5,692,129; 5,758,069; 5,758,344; 5,761,499; 5,781,724; 5,781,733; 5,784,560; 5,787,439; 5,818,936; 5,828,882; 5,832,275; 5,832,483; 5,832,487; 5,859,978; 5,870,739; 5,873,079; 5,878,415; 5,884,304; 5,893,118; 5,903,650; 5,905,860; 5,913,025; 5,915,253; 5,925,108; 5,933,503; 5,933,826; 5,946,467; 5,956,718; 5,974,474. U.S. and Foreign Patents Pending.

Novell, Inc.
122 East 1700 South
Provo, UT 84606
U.S.A.

www.novell.com

TCP/IP
January 2000
104-001261-001

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

Third-Party Trademarks

Contents

About This Guide

This guide provides the information you need to configure and manage the Novell Internet Access Server 4.1 TCP/IP routing software. In addition to planning information, this guide provides troubleshooting tips, techniques, and tools, as well as the symptoms of and solutions to commonly occurring problems for the TCP/ IP components of Novell Internet Access Server 4.1.

1 Understanding

This section introduces TCP/IP and lists the files used by the Novell® TCP/IP implementation. It also provides an overview of the TCP/IP suite of protocols.

In addition, this section discusses using TCP/IP to connect dissimilar computers, describes assigning network addresses and subnet addresses, and examines TCP/IP routing protocols and different topologies.

This section discusses the four network database files used to convert internal data into identifiable and workable names. It describes TCP/IP network management using the SNMP and SNMPLOG NetWare® Loadable Module™ (NLM™) files.

The TCP/IP Suite of Protocols

The protocols in the TCP/IP suite roughly correspond to a network communications model defined by the International Organization for Standardization (ISO). This model is called the Open Systems Interconnection (OSI) reference model. The OSI model describes an ideal computer network system in which communication on the network occurs between processes at discrete and identifiable layers. Each layer on a given host provides services to the layers above it and receives services from the layers below it. Figure 1 illustrates the seven layers of the OSI reference model, as defined by ISO, and the roughly corresponding layers of the TCP/IP protocol suite.

Figure 1 OSI Reference Model and Corresponding TCP/IP Layers

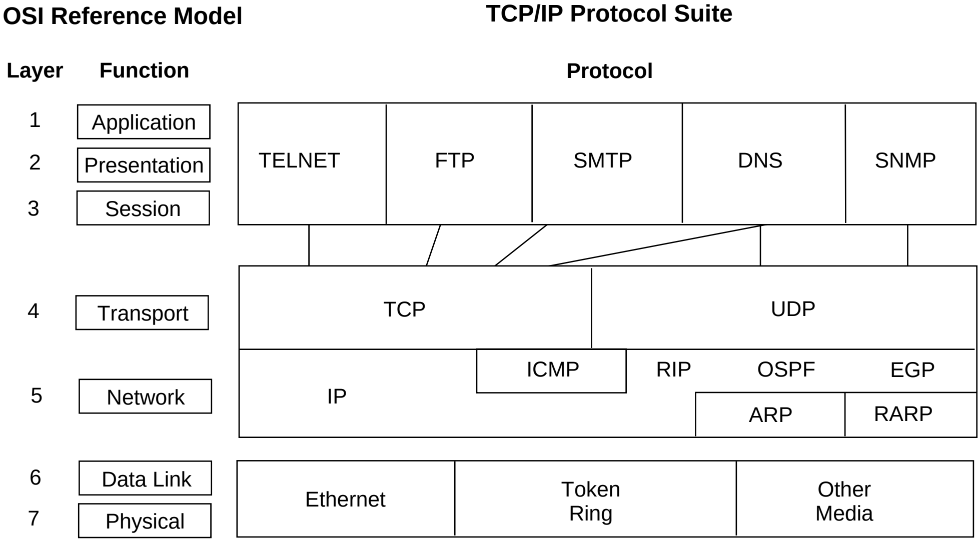

The layering system lets developers concentrate their efforts on the functions in a given layer. It is not necessary for designers to create all the mechanisms to send information across the network. They have to know only what services the software needs to provide to the layer above it, what services the layers below it can provide to the software, and which protocols in the suite provide those services.

Table 1 lists some of the more common protocols in the TCP/IP suite and the services they provide.

Table 1 TCP/IP Protocols

Protocol	Service
Internet Protocol (IP)	Provides packet delivery services (routing) between nodes.
Internet Control Message Protocol (ICMP)	Provides transmission of error and control messages between hosts and routers.
Address Resolution Protocol (ARP)	Maps IP addresses to physical addresses.
Transmission Control Protocol (TCP)	Provides reliable data-stream delivery service between end nodes.

Protocol	Service
User Datagram Protocol (UDP)	Provides unreliable datagram delivery service between end nodes.
File Transfer Protocol (FTP)	Provides application-level services for file transfer.
TELNET	Provides terminal emulation.
Routing Information Protocol (RIP)	Enables the exchange of distance vector routing information between routers.
Open Shortest Path First (OSPF)	Enables the exchange of link state routing information between routers.
Exterior Gateway Protocol (EGP)	Enables the exchange of routing information between exterior routers.

Overview of TCP/IP Protocol Usage

Applications developed for TCP/IP generally use several of the protocols in the suite. The sum of the layers of the protocol suite is also known as the *protocol stack*. User applications communicate with the top layer of the protocol suite. The top-level protocol layer on the source computer passes information to the lower layers of the stack, which in turn pass it to the physical network. The physical network transfers the information to the destination computer. The lower layers of the protocol stack on the destination computer pass the information to higher layers, which in turn pass it to the destination application.

Each protocol layer within the TCP/IP suite has various functions; these functions are independent of the other layers. Each layer, however, expects to receive specific services from the layer beneath it, and each layer provides specific services to the layer above it.

Figure 2 shows the TCP/IP protocol layers. The layers at the same level on the source and destination computers are *peers*. For example, the application on the source computer and the application on the destination computer are peers. Each layer of the protocol stack on the source computer communicates with its peer layer on the destination computer. From the perspective of the software developer or user, the transfer takes place as if the peer layers sent their packets directly to one another.

Figure 2 TCP/IP Protocol Layers

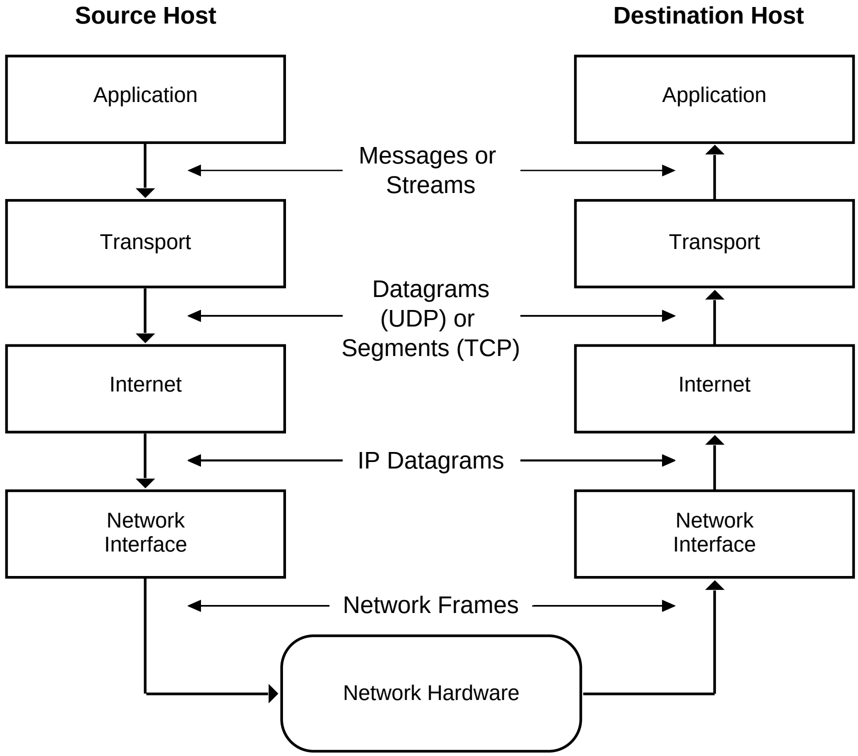

An application for transferring files with TCP, for instance, performs the following operations to send the file contents:

1. The Application layer passes a stream of bytes to the Transport layer on the source computer.
2. The Transport layer divides the stream into TCP segments, adds a header with a sequence number for that segment, and passes the segment to the Internet (IP) layer. A checksum is computed over the TCP header and data.
3. The IP layer creates a packet with a data portion containing the TCP segment. The IP layer adds a packet header containing source and destination IP addresses.
4. The IP layer also determines the physical address of the destination computer or intermediate computer on the way to the destination host. It

passes the packet and the physical address to the Data-Link layer. A checksum is computed on the IP header.

5. The Data-Link layer transmits the IP packet in the data portion of a data-link frame to the destination computer or an intermediate computer. If the packet is sent to an intermediate computer, steps 4 through 7 are repeated until the destination computer is reached.
6. At the destination computer, the Data-Link layer discards the data-link header and passes the IP packet to the IP layer.
7. The IP layer checks the IP packet header. If the checksum contained in the header does not match the checksum computed by the IP layer, it discards the packet.
8. If the checksums match, the IP layer passes the TCP segment to the TCP layer.
9. The TCP layer computes a checksum for the TCP header and data. If the computed checksum does not match the checksum transmitted in the header, the TCP layer discards the segment. If the checksum is correct and the segment is in the correct sequence, the TCP layer sends an acknowledgment to the source computer and passes the data to the application.
10. The application on the destination computer receives a stream of bytes, just as if it were directly connected to the application on the source computer.

Internet Protocol

In the TCP/IP protocol suite, all packets are delivered by the IP *datagram delivery* service. Packet delivery is not guaranteed by this service. A packet can be misdirected, duplicated, or lost on the way to its destination. The service is *connections* because all packets are transmitted independently of any other packets. This is in contrast to a telephone network, for instance, where a circuit is established and maintained.

To keep track of the delivery status, TCP/IP applications using the IP datagram delivery service expect to receive replies from the destination node.

IP defines the form that packets must take and the ways that packets are handled when they are transmitted or received. The form the packet takes is called an I*P datagram*. It is the basic unit of information that is passed across a TCP/IP network. The IP datagram consists of a header and a data section.

The header section contains the sender’s (source) IP address and the receiver’s (destination) IP address and other information. Figure 3 shows the general form of an IP datagram.

Figure 3 IP Datagram Structure

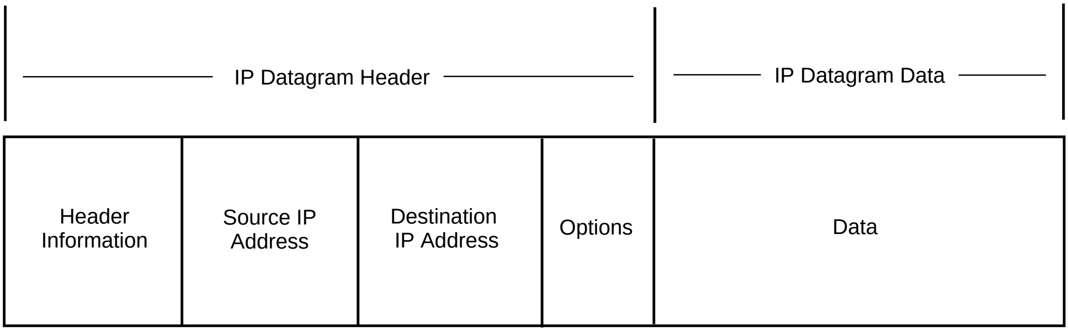

The Data-Link layer transmits IP packets in the data section of its physical frame. Because IP supports a 64-KB packet length, an IP datagram might not fit in a data-link frame. Also, in traveling to its destination, a datagram can traverse many different media with different physical frame lengths. An IP router might have to forward a packet across media in which the inbound and outbound frame lengths differ.

To handle these potential problems with packet transmission, IP specifies a method for breaking datagrams into *fragments*. The fragments are *reassembled* when they arrive at the final destination. Reassembling fragments reconstructs the entire IP datagram.

Path Maximum Transfer Unit

The maximum transfer unit (MTU) is the largest amount of data that can be transferred across a given physical network. For local area networks, such as Ethernet, the MTU is determined by the network hardware. For wide area networks that use serial lines to interconnect packet switches, the MTU is determined by software.

The Path MTU is the smallest MTU of all MTUs for the hops along a path from the source host to the destination host. The Path MTU governs the size of the largest IP packet that can be sent across the path without fragmentation. This feature conforms to RFC 1191.

There are two advantages to this feature. The Path MTU avoids fragmentation anywhere along the path. In addition, it reduces the protocol overhead.

No configuration is required for the Path MTU discovery. This feature is enabled when you enable TCP/IP.

Path MTU Discovery Process

The Path MTU Discovery process prevents fragmentation between two routers. Figure 4 illustrates the following description of the Path MTU Discovery process.

Figure 4 Path MTU Discovery

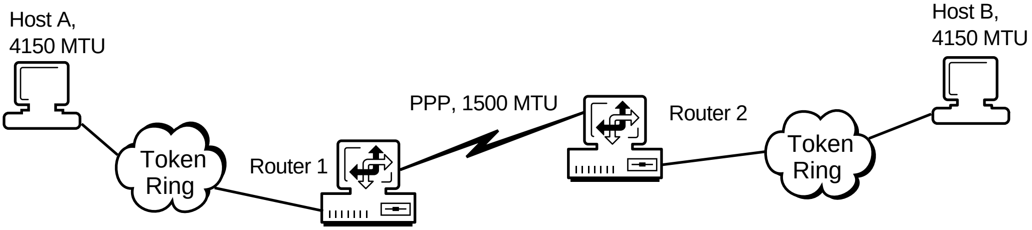

The following describes the steps involved in the Path MTU Discovery process:

1. Host A opens a File Transfer Protocol (FTP) connection to Host B.
2. Host A and Host B negotiate the maximum segment size (MSS) during their connection. This is the largest TCP segment that a host can send across a network. The MSS in Figure 4 is 4,110 bytes, which is 4,150 bytes minus 40 bytes for the IP and TCP headers.
3. Host A sends a 4,150-byte packet to Host B. This packet consists of 4,110 bytes of data and 40 bytes of header information. The don’t fragment (DF) flag in the IP header is set to yes in Host A.
4. Router 1 receives the packet from Host A. Then Router 1 determines that the packet is larger than 1,500 bytes, which is the maximum packet size that can be sent over a PPP network.
5. Router 1 sends Host A an ICMP destination unreachable error message. This message indicates that Router 1 must fragment packets larger than 1,500 bytes.

6. Host A receives the error message from Router 1. In response, it adjusts the maximum segment size to 1,460 bytes.
7. Host A resends the data from Step 3. Each packet consists of 1,460 bytes of data and 40 bytes of header information.
8. Router 1 accepts the packets and forwards them to Router 2 and then to Host B.

Routing

The term *routing* refers to the transmission of a datagram from one node to another on the same or a different network. The route refers to the path that is chosen to transmit an IP datagram from its origin to its destination, based on the IP addresses contained in the datagram.

When a datagram is sent to a node on another network, the network portions of the source and the destination IP addresses are different. When the packet is received by a router that connects the source to the destination network, the router forwards the packet on the correct interface to reach the destination, as shown in Figure 5. Two networks are connected if at least one router is attached to both networks.

Figure 5 Network Router

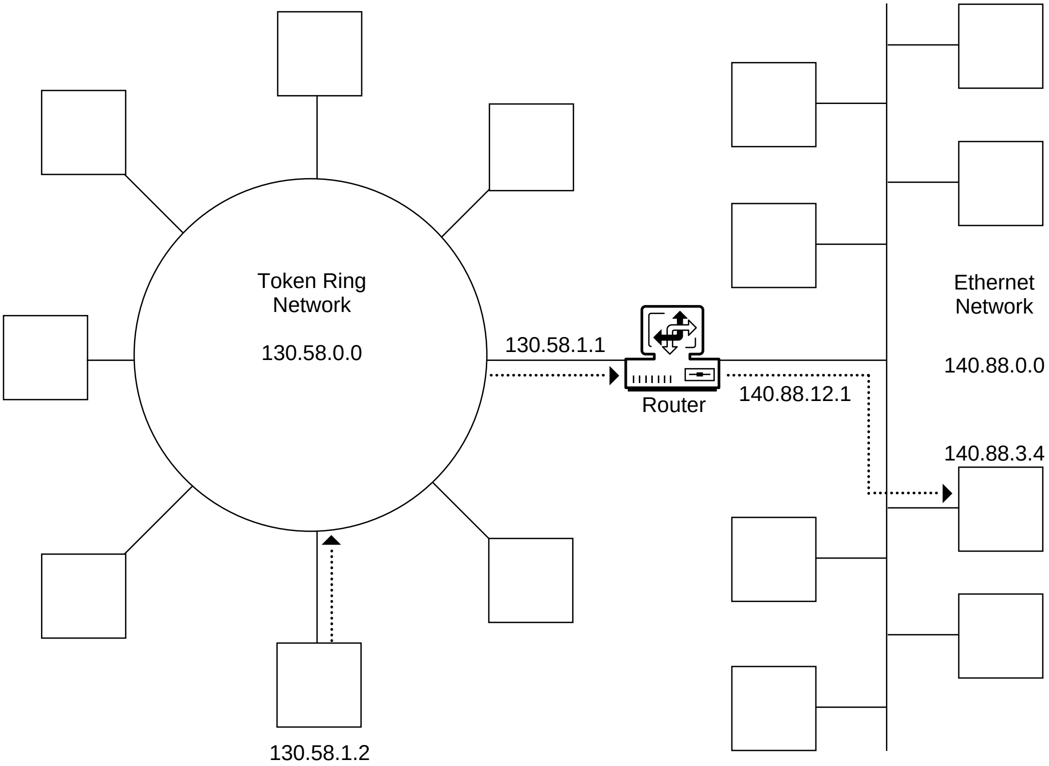

Each host has a default router or a list of routers in other networks. When IP sends a datagram, it performs the following steps:

1. IP searches the routing table of the sending node for a default route or a path to the destination IP address.
2. IP extracts the address of the default router or next-hop router from the route entry.
3. IP requires ARP to map the next-hop address to its hardware address.
4. IP transmits the packet to the next hop.
5. IP repeats Steps 1 through 4 until the final destination is reached.

Error and Control Messages

Another protocol in the TCP/IP suite is the Internet Control Message Protocol (ICMP). ICMP packets contain information about failures on the network: inoperative nodes and gateways, packet congestion at a gateway, and so on. The IP software, rather than the application, interprets an ICMP message. The IP software then takes the appropriate action with respect to the ICMP message, independently of the application. Because an ICMP message might need to travel across several networks to reach its destination, it is encapsulated in the data portion of an IP datagram.

ICMP is also used to test connectivity between two nodes. The originating node uses PING to send an ICMP echo request and waits for an ICMP echo response from the destination.

Transport Layer Protocols

The Transport layer of the TCP/IP protocol suite consists of two protocols, UDP and TCP. UDP provides an unreliable connectionless delivery service to send and receive messages. TCP adds reliable byte stream-delivery services on top of the IP datagram delivery service.

UDP

UDP identifies applications through *ports*. The protocol defines two types of protocol ports: well-known port assignments and dynamically bound ports. For well-known port assignments, certain UDP port numbers are reserved for particular applications. The ports numbered between 1 and 1,023 are well-known port numbers. For dynamically bound ports, an application requests that UDP assign a port to identify which port the process uses. The port must be in the range of 1,024 to 65,535. Then the application can direct UDP datagrams to that port.

UDP enables multiple clients to use the same port number and different IP addresses. The arriving UDP datagrams are delivered to the client that matches both the destination port number and address. (A socket consists of an IP address and the port number.) If there is no matching client, an ICMP Destination Unreachable, Port Unreachable message is sent and the packet is dropped.

The UDP datagram is encapsulated in an IP datagram that, in turn, is encapsulated in physical frames. Figure 6 shows a UDP datagram encapsulated in an IP datagram, which, in turn, is encapsulated in an Ethernet frame. Figure 6 also illustrates how the concept of *layering* , discussed at the beginning of this section, affects the construction of packets sent across the network.

Figure 6 UDP Datagram Encapsulation

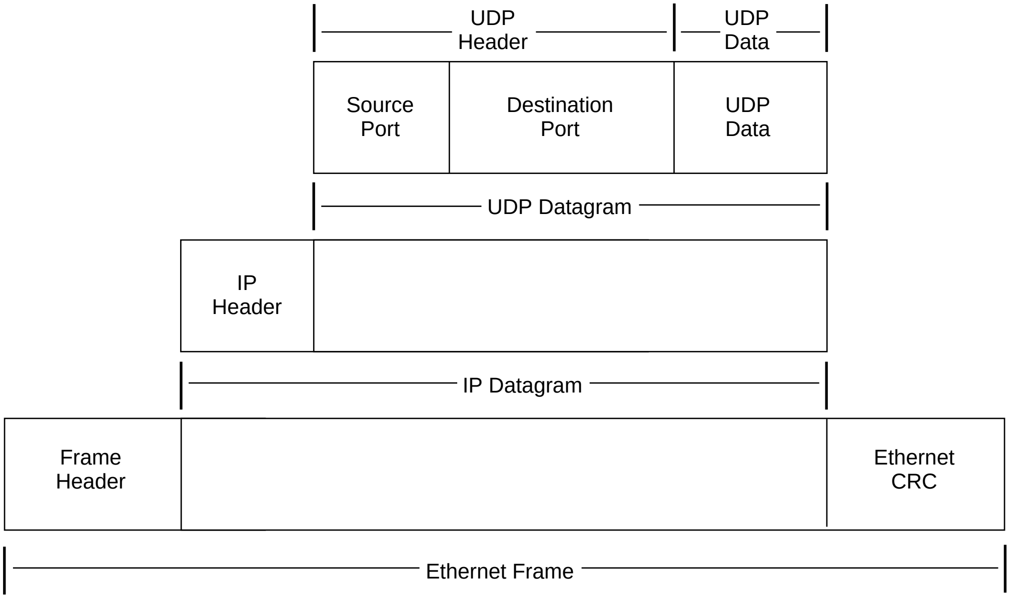

In this example, the IP address transmits the IP datagram to the node. At that destination, the IP software extracts the UDP datagram and delivers it to the UDP-layer software. The UDP-layer software delivers the UDP data through the destination port to the receiving application. The process at that port uses the data in the UDP datagram. The UDP datagram also contains a source port to ensure that the destination process can reply correctly.

TCP

For applications that must send or receive large volumes of data, unreliable datagram delivery can become burdensome. Application programmers might have to develop extensive error handling and status information modules to track the progress and state of data transfer for every application. The TCP/IP suite of protocols avoids this problem by using TCP, a *reliable byte-stream delivery protocol*. TCP establishes a connection between two applications and sends a stream of bytes to a destination in exactly the same order that they left the source. Before transmission begins, the applications at both ends of transmission obtain a TCP *port* from their respective operating systems. These are analogous to the ports used by UDP. The application initiating the transfer, known as the client side, generally obtains a port dynamically. The application responding to the transfer request, known as the server side, generally uses a well-known TCP port. The client side is typically the active side and initiates the connection to the passive server side.

Like the UDP datagrams, TCP *segments* are encapsulated in an IP datagram. TCP *buffers* the stream by waiting for enough data to fill a large datagram before sending the datagram. The stream is *unstructured* , which means that before transmission of data, both the sending and receiving applications must agree on the meaning of the contents of the stream. The TCP protocol uses *full-duplex* transmission. Full duplex means that two data streams can flow in opposite directions simultaneously. Thus, the receiving application can send data or control information back to the sending application while the sending application continues to send data.

The TCP protocol gives each segment a sequence number. At the receiving end of the connection, TCP checks successive sequence numbers to ensure that all the segments are received and processed in the order of the sequence numbers. The receiving end sends an acknowledgment to the sender for the segments received. TCP enables the sender to have several outstanding segments before the receiver must return an acknowledgment. If the sending node does not receive an acknowledgment for a segment within a certain time, it retransmits that segment. This scheme, called *positive acknowledgment with retransmission* , ensures that the stream delivery is reliable.

Physical and IP Addresses

Each node has a *physical address* for the specific hardware device that connects it to a network. For instance, a physical address on an Ethernet network is a 6-byte numeric value, such as 08-00-14-57-69-69. It is assigned by the manufacturer of the Ethernet interface hardware. X.25 networks, which conform to the specification of the ITU-T (International Telecommunications Union, Telecommunications sector), previously CCITT, use the X.121 standard for physical addresses, which consist of 14-digit numbers.

NOTE: Physical addresses are also called media access control (MAC) addresses. Throughout the rest of this section, all references to MAC or physical addresses assume physical addresses on Ethernet, token ring, or FDDI networks.

The IP address for a node is a logical address and is independent of any particular hardware or network topology. It has the same form, regardless of the media type. The IP address is a 4-byte (32-bit) numeric value that identifies both a network and a local host or node (computer or other device) on that network. The 4-byte IP address is usually represented in dotted decimal notation. Each byte is represented by a decimal number, and periods separate the bytes, for example, 129.47.6.17.

A conflict arises because IP uses a 32-bit address and Ethernet uses a 48-bit Ethernet address. To associate the IP address to a physical address on an Ethernet network, a mapping must occur between the two types. The address resolution protocol (ARP) provides a mapping between the two different forms of addresses. As a result of the mapping performed by ARP, the IP address is mapped to a physical address. ARP mapping is limited to networks that support hardware broadcast.

IP Address to Physical Address Translation

Each physical medium has its own *physical* address for nodes on that medium. The physical addresses are also called MAC addresses. Ethernet and token ring networks use 6-byte MAC addresses. ARCnet uses a 1-byte MAC address.

IP addresses are independent of the hardware. When an IP packet is transmitted on the network, it is first *encapsulated* within the physical frame used by that network. Figure 7 shows an IP packet encapsulated in an Ethernet frame. The IP packet contains an Internet address for a node, but the Ethernet frame must have a physical address for it to be delivered on the data-link network. Therefore, the sending node must be able to map an IP address to a physical hardware address.

Figure 7 IP Datagram Encapsulation

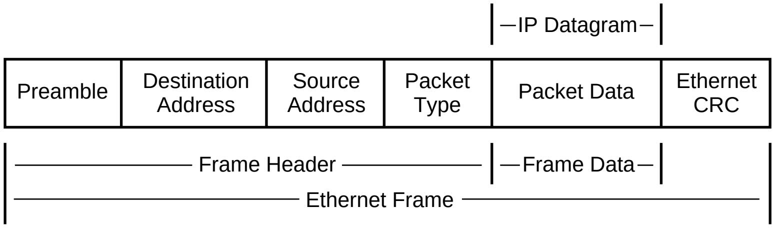

Mapping Internet Addresses to Physical Addresses

When an IP address is mapped to a physical, or MAC, address, ARP is used on broadcast networks such as Ethernet, token ring, and ARCnet. When a node uses IP to send a packet, it must determine which physical address on the network corresponds to the destination IP address. To find the physical address, the node broadcasts an ARP packet containing the destination IP address. The node with the specified destination IP address sends its physical address back to the requesting node.

Address Resolution Cache

To speed packet transmissions and reduce the number of broadcast requests that must be examined by every node on the network, each node keeps an *address resolution cache* , or ARP table. Each time the node broadcasts an ARP request and receives a response, it creates an entry in its address resolution cache. The entry maps the IP address to the physical address.

When the node sends an IP packet, it looks up the IP address in its cache and uses the physical address, if found. The node broadcasts an ARP request only if the IP address is not in its cache.

Assigning IP Network Addresses

IP network addresses should be assigned by one person at your company. We recommend that a network administrator assign IP network addresses. Therefore, to obtain a new address, see your network administrator. If you are a network administrator, use this section to help you assign IP network addresses.

For a node using the TCP/IP protocol suite to communicate with other nodes, including nodes on other private networks and on the Internet, an IP network address is required. Your IP network address could be determined in one of the following ways:

- If you are accessing the Internet through an Internet Service Provider (ISP), you can be assigned an IP address by your ISP.
- If you are connected directly to the Internet community or if you cannot connect to the Internet using the registered IP address range you were assigned by your ISP, contact the following organization:

  Network Solutions, Inc.
  Attn: InterNIC Registration Services

505 Huntmar Park Dr.
Herndon, VA, USA 20170

E-mail: hostmaster@internic.net

Web address: http://nic.ddn.mil or http://192.112.36.5

- If your network is not attached to the public Internet community, you can select an arbitrary IP network number. However, if you plan to attach your network to the Internet later, you should use the guidelines in RFC 1918.

The addresses for all the nodes on the network must meet the following criteria:

- All addresses within a network must use the same prefix. For example, any node on network 129.47 must have an address in the form 129.47.x.x.
- Each node must have a unique IP address.

For information about selecting and assigning IP classes and addresses, refer to:

- Historic IP Address Classes
- Class A Addresses
- Class B Addresses
- Class C Addresses
- Identifying Network Classes
- Selecting an Appropriate Address Class
- Reserved IP Addresses
- Dynamically Assigning IP Addresses

Historic IP Address Classes

Each 4-byte IP address is divided into two parts:

- A *network* portion, which identifies the network
- A *host* portion, which identifies the node

IP addresses are differentiated into three classes, based on the two most significant bits of the first byte. This is done so that routers can efficiently extract the network portion of the address.

This division can occur at any one of three locations within the 32-bit address. These divisions correspond to the three IP address classes: Class A, Class B, and Class C. Regardless of address class, all nodes on any single network share the same network portion; each node has a unique host portion.

Class A Addresses

A Class A IP address consists of a 1-byte network portion followed by a 3-byte host portion, as shown in Figure 8. The highest-order bit of the network portion is always set to 0. Thus, within an internetwork, there can be a total of 126 Class A networks (1 through 126), with more than 16 million nodes in each (networks 0 and 127 are reserved).

The format of a Class A address is as follows:

```
0nnnnnnn.hhhhhhhh.hhhhhhhh.hhhhhhhh
```

where *n* represents the network address and *h* represents the host address.

Class A addresses contain 7 bits of network address and 24 bits of host address.

Figure 8 Class A Address Example

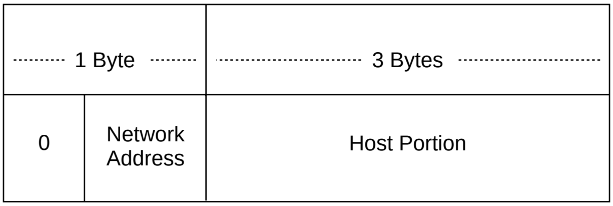

Class B Addresses

A Class B IP address consists of a 2-byte network portion followed by a 2-byte host portion, as shown in Figure 9. The two highest-order bits of the network portion are always set to 10. Thus, within a single internetwork there can be approximately 16,000 Class B networks (128.0 through 191.255), with more than 65,000 nodes in each.

The format of a Class B address is as follows:

```
10nnnnnn.nnnnnnnn.hhhhhhhh.hhhhhhhh
```

where *n* represents the network address and *h* represents the host address.

Class B addresses contain 14 bits of network address and 16 bits of host address.

Figure 9 Class B Address Example

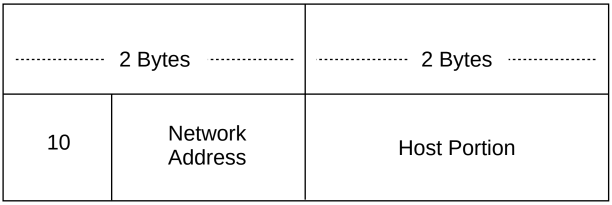

Class C Addresses

A Class C IP address consists of a 3-byte network portion followed by a 1-byte host portion, as shown in Figure 10. The three highest-order bits of the network portion are always set to 110. Within a single internetwork, there can be approximately 2 million Class C networks (192.0.0 through 223.255.255), with up to 254 nodes in each.

The format of a Class C address is as follows:

```
110nnnnn.nnnnnnnn.nnnnnnnn.hhhhhhhh
```

where *n* represents the network address and *h* represents the host address.

Class C addresses contain 21 bits of network address and 8 bits of host address.

Figure 10 Class C Address Example

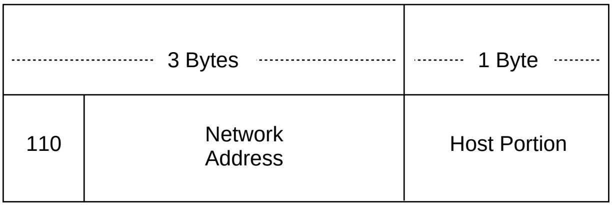

Identifying Network Classes

The first byte of an IP address identifies which of the three network classes that address belongs to. The ranges for that first byte are as follows:

- Class A: 1 to 126 (1.h.h.h to 126.h.h.h)
- Class B: 128 to 191 (128.n.h.h to 191.n.h.h)
- Class C: 192 to 223 (192.n.n.h to 223.n.n.h)

An IP address beginning with 154 is a Class B address. The first two bytes of the address are represented by *n* for the network portion of the address, and the last two bytes are represented by *h* for the host portion. For example, an IP address of 154.1.0.3 means the IP network portion is 154.1.0.0 and the host portion on that network is #.#.0.3.

The network portion of an IP address should be the same for all nodes on that network. Each node connected to the network must have a unique IP host address assigned to it.

HINT: The key to selecting a number for the host portion of the IP address is to ensure that the number selected is unique, that is, that no other host on the network has the same IP address.

Selecting an Appropriate Address Class

When selecting an IP address class, you must decide on both network numbers and host address portions. Because the first 1, 2, or 3 bits of the IP address determine how the entire address is to be interpreted and where the division between the network address and host address portion is to occur, you should know the consequences of your choice. When deciding on a network class, you should consider the number of IP nodes to be supported on your network and the number of networks you plan to configure.

For example, if you use Class C addresses (the first 3 bits of the IP address are 110 binary), then you are restricted to 254 nodes. However, the number of nodes available can be altered by using subnets. Before selecting an IP address class, refer to “Creating Subnets.”

Reserved IP Addresses

The IP addressing rules reserve the following types of IP addresses for special purposes:

- **Network addresses** —IP addresses in which the host portion is set to all zeros. For example,129.47.0.0 is the network address (or network number) for a Class B network. Network addresses identify networks rather than nodes on a network. By convention, no node is ever assigned a host portion consisting of all zeros.
- **Broadcast addresses** —Addresses in which the host portion is set to all ones. A packet with a broadcast address is destined for every node on the network. By convention, no node is ever assigned a host portion consisting of all ones.
- **Loopback addresses** —Addresses that cause the protocol software to return data without sending traffic across a network. Network address 127.0.0.0 and all host addresses on that network (for example, 127.0.0.1) are reserved.
- **Multicast addresses** —Addresses that are used to send packets to a group of hosts or routers. They range from 224.0.0.1 to 239.255.255.255.
- **Reserved addresses** —Addresses in which the network portion consists of all zeros or all ones.

Dynamically Assigning IP Addresses

Routing can be configured to use IPCP in one of two modes: dynamically assigning IP addresses or automatically being assigned an IP address. In the first mode, your router is configured with a range of IP addresses that it uses to dynamically assign addresses to dial-up routers. In the second mode, your router is configured to automatically accept an IP address assigned by your ISP.

IPCP can be used only on PPP interfaces. For more information about configuring IPCP, refer to "Setting Up."

Creating Subnets

One IP network can be divided into smaller networks, called subnets. The following are reasons to divide your network:

- **Use multiple media** —It can be impossible, inconvenient, or too expensive to connect all nodes to a single network medium when these nodes are too far apart or already connected to different media.

- **Reduce congestion** —Traffic between nodes on a single network uses network bandwidth. As a result, more bandwidth is required when you have more nodes. Splitting nodes into separate networks reduces the number of nodes on a data-link network. Fewer nodes generate less traffic and, as a consequence, less congestion.
- **Reduce CPU use** —Reducing CPU use on connected nodes is similar to reducing congestion. More nodes on a network cause more broadcasts on that network. Even if a broadcast is not sent to a particular node, each node on a network must react to every broadcast before deciding to accept it or discard it.
- **Isolate a network** —By splitting a large network into small networks, you limit the impact of one network's problems on another. Such problems can include network hardware failures, such as an open Ethernet tap, or software failures, such as a broadcast storm.
- **Improve security** —On a broadcast network medium such as Ethernet, each node on a network has access to all packets sent on that network. By enabling sensitive network traffic on only one network, other network monitors can be prevented from accessing this sensitive traffic.
- **Make efficient use of IP address space** —If you are using a Class A or B network number and have multiple small physical networks, you can divide the IP address space into multiple IP subnets and assign them to individual physical networks. Another option is to obtain several Class C network numbers, although this is less desirable.

  For more information about creating subnets, refer to:

  - Subnet Addresses and Masks
  - Subnet Zero
  - Variable Size Subnets
  - Assigning Subnet Addresses
  - Broadcast Addresses
  - Multicast Addresses

Subnet Addresses and Masks

Communication between a node on a local subnet and a node on a different subnet is similar to communication between nodes on two different networks. To a user, routing between subnets is transparent. Internally, the IP software recognizes any IP addresses that are destined for a remote subnet and sends those packets to the router on that subnet.

As in network-to-network communication, the routing information for communication between subnets is maintained in the routing table (by IP).

When a network is divided into subnets, the host address portion of the IP address is divided into two parts, just as the IP address itself is divided into two parts. The host address portion specifies both the subnet of the IP network and the node on that subnet.

The 4-byte IP address consists of a network address and a host portion, as shown in Figure 11.

Figure 11 Subnet IP Address

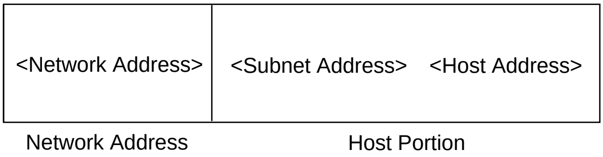

For instance, if a network has the Class B IP network address portion 129.47, the remainder of the IP address can be divided into subnet addresses and host addresses. Controlled by the local network administrator, this division allows the most flexibility for network development at the local site. For example, the subnet address could comprise 4 bits of the remaining 2 bytes. This allows 15 subnets, each with 4,094 nodes. Or, in another case, the subnet address could comprise 8 bits, allowing 255 subnets (a subnet address of all ones is not valid), each with 254 nodes.

NOTE: NetWare routing software supports the use of all zeros in the subnet field (subnet zero). However, a subnet field with all ones denotes all subnets of a particular network; therefore, a subnet field with all ones cannot be used as a local IP address.

Figure 12 shows a single IP network divided into two subnets. The router shown has physical attachments and IP addresses on both subnets (129.47.128.1 and 129.47.192.1). It might also have physical devices and IP addresses (*nn.nn.nn.nn*) connecting it to other networks.

Figure 12 Network with Two Subnets

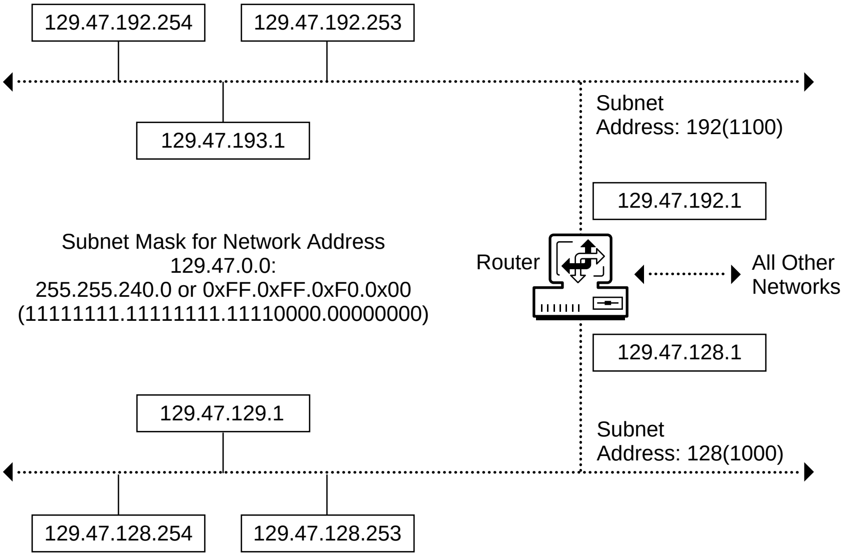

A *subnet mask* indicates how the host portion of the IP address is divided into a subnet address and a local host portion. The network mask is a 32-bit number with all ones for all network and subnet address portions, and all zeros for the host field. With a Class B network portion of 129.47 and a 4-bit subnet address, for instance, the subnet mask consists of 20 ones and 12 zeros. In essence, a subnet mask locally extends the network address portion of an IP address and reduces the host portion.

Table 2 shows an example of a Class C subnet with an IP address of 200.2.1.209. To create a subnet address, bits are taken from the local host portion. As the size of the subnet mask increases, the number of hosts decreases and the number of subnets increases.

Table 2 Subnet Masks with Class C Addresses

Class C IP Address 200.2.1.209	**Network Number**	**Subnet Number**	**Host Number**	**Available Networks, Subnets, and Hosts**
FF.FF.FF.0	200.2.1.0	None	0.0.0.209	1 network, 0 subnets, and 254 hosts
FF.FF.FF.E0	200.2.1.0	200.2.1.192	0.0.0.17	7 subnets and 30 hosts per subnet

Class C IP Address 200.2.1.209	Network Number	Subnet Number	Host Number	Available Networks, Subnets, and Hosts
FF.FF.FF.F0	200.2.1.0	200.2.1.208	0.0.0.1	15 subnets and 14 hosts per subnet

Figure 13 shows examples of IP network addresses, their relationship to the subnet mask, and the corresponding subnets.

Figure 13 Subnet Mask and IP Addresses

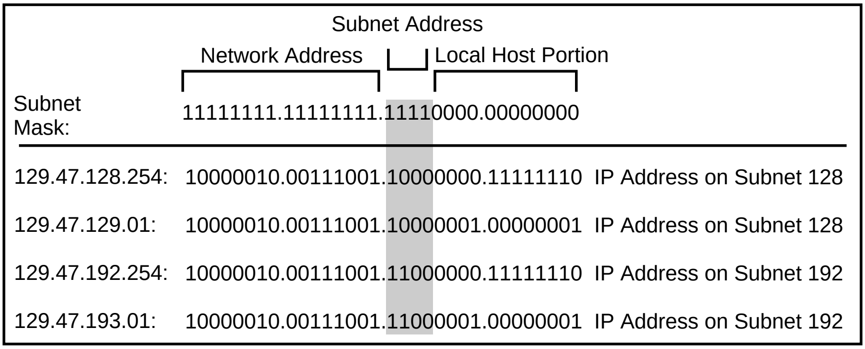

```
Subnet
Mask:            11111111.11111111.11110000.00000000

129.47.128.254:  10000010.00111001.10000000.11111110  IP Address on Subnet 128
129.47.129.01:   10000010.00111001.10000001.00000001  IP Address on Subnet 128
129.47.192.254:  10000010.00111001.11000000.11111110  IP Address on Subnet 192
129.47.193.01:   10000010.00111001.11000001.00000001  IP Address on Subnet 192
```

Subnet Zero

Subnet zero is a subnet with all the bits in the subnet field of the IP address set to 0. For example, subnet 130.57.0.0, with a mask of 255.255.240.0, is a subnet zero of network 130.57, as shown in Figure 14.

Figure 14 Subnet Zero

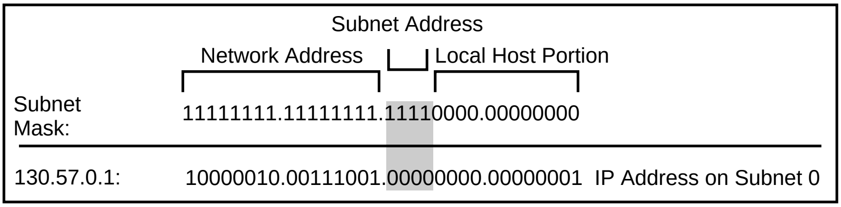

```
Subnet
Mask:            11111111.11111111.11110000.00000000

130.57.0.1:      10000010.00111001.00000000.00000001  IP Address on Subnet 0
```

The official IP specification reserves the subnet addresses with all zeros and all ones and does not allow them to be used as subnet addresses. However, this policy wastes one subnet in the IP address space. To counteract this limitation, Novell's TCP/IP implementation enables the use of subnet zero.

Variable Size Subnets

The subnets of a network can have different length subnet masks, called *variable length subnet masks*. These subnets are called variable because the size, or length, of the subnet varies from subnet to subnet.

A subnet mask defines the number of bits that can be used to define the subnet and the number of bits to define the host. As the subnet mask increases, the number of hosts on a subnet decreases. As the subnet mask decreases, the number of hosts that can be defined increases.

Some network configurations have individual subnets with a large number of hosts and other subnets with a small number of hosts. Using the same subnet masks on all subnets can mean either of the following:

- The mask is too small and you do not have enough subnet numbers for all your subnets.
- The mask is too big and you do not have enough host IDs for all your hosts on a subnet.

If the mask is too small or too big, use a variable size subnet. By varying the size of the subnet mask used on a network, you can match the number and size of subnets to your configuration.

For example, subnet 16 of network 130.57.0.0 with mask 255.255.240.0, 130.57.16.0, can be further divided into 16 sub-subnets with 256 hosts each. (Actually, this division creates 15 sub-subnets with 254 hosts each because sub-subnet 130.57.31.0, host 0, and host 255 are not used.)

NOTE: OSPF and RIP II recognize subnet masks and support variable size subnets. RIP I does not work when the network is partitioned into variable length subnets because RIP I assumes that all subnets belonging to the same network use the same subnet mask.

Assigning Subnet Addresses

NOTE: Because RIP I packets do not carry subnet mask information, the RIP I routing protocol imposes several restrictions on the use of subnets. If you are using RIP I, use the same subnet mask for all subnets belonging to the same network. Using RIP II lifts this restriction.

If you are installing the routing software on a network with subnets, use the subnet mask already established for the network.

Subnet addresses and host addresses are typically assigned in numeric order, where both the subnet and host addresses are assigned from the right edge of their field. By this method, the border between the subnet address and the host address becomes fixed when the first subnet (subnet address = 1) is assigned. If the number of hosts on a subnet or the number of subnets required exceeds the limits of the subnet mask, using this method makes it difficult to adjust the subnet mask because each host must be renumbered.

To prepare for changes in the size of the subnet mask, RFC 1219 suggests that subnets be assigned from the *left* of the subnet address field, and that hosts be assigned, in numeric order, from the *right* of the host address field. In this way, the subnet bits become a *mirror image* of the host bits. (You must still select an initial subnet mask and use it for all subnets in the network.) For example, to apply this method to a Class B IP network with a subnet mask of 255.255.255.0, you assign subnet addresses as follows:

1000 0000 (Decimal 128)
0100 0000 (Decimal 64)
1100 0000 (Decimal 192)
0010 0000 (Decimal 32)
...

Then, you assign host addresses on each subnet as follows:

0000 0001 (Decimal 1)
0000 0010 (Decimal 2)
0000 0011 (Decimal 3)
0000 0100 (Decimal 4)

. . .

Using this method leaves a buffer zone between the subnet and host addresses, which enables future network growth.

The method of assigning subnet addresses described in this section summarizes the method suggested in RFC 1219, *On the Assignment of Subnetwork Numbers*. For a complete description of this method, refer to RFC 1219.

Broadcast Addresses

There are four types of broadcast addresses: directed broadcasts, subnet directed broadcasts, all-subnets directed broadcasts, and limited broadcasts. A directed broadcast has a destination IP address with the network portion of the IP address set to Class A, B, or C network, and the host field set to all ones. Directed broadcasts are sent to all hosts on the specified network.

If the network is divided into subnets, each subnet has a subnet directed broadcast. A subnet directed broadcast has an IP address with the network field set to the network identifier, the subnet field set to the subnet identifier, and the host field set to all ones.

An IP address with both the subnet and host field set to all ones is interpreted as a broadcast directed to all the subnets on the network. That is, the first router on the specified network broadcasts the IP address to one of its subnets. If broadcast forwarding is enabled, the receiving routers in that network forward the broadcast to other subnets.

An IP address with all ones, 255.255.255.255, is called a limited address. It is directed to all hosts on the subnet from which the broadcast originated.

Multicast Addresses

A multicast address is used to send packets to a group of hosts or routers. A packet with a multicast address is received by all hosts and routers belonging to that multicast group. Class D addresses are reserved for multicast addresses. They range from 224.0.0.1 to 239.255.255.255.

Novell’s TCP/IP implementation uses five multicast addresses. Two are used by OSPF to multicast packets to OSPF routers. These addresses are 224.0.0.5 and 224.0.0.6. Two are used by Router Discovery messages to multicast router advertisements and solicitation messages. These addresses are 224.0.0.1 and 224.0.0.2. RIP II uses multicast address 224.0.0.9.

Router Discovery Protocol

The Router Discovery Protocol, an Internet Control Message Protocol (ICMP) extension, allows hosts to discover routers on their networks and determine which router to use as the default router. When a host needs to send a packet to another network, it first sends the packet to a router that forwards the packet toward the destination. To accomplish this, the host needs to know where the routers are on its network and which one to send packets to.

When you configure the router discovery mechanism, the router advertises itself with periodic ICMP router advertisement messages. Then the host listens to this message and decides whether to use a router as the default router.

You can configure the host to solicit the router advertisement on attached networks. All participating routers then reply to the inquiry. By collecting those replies, the host discovers the routers on the network and determines which router to use.

A host might not select the best router (the router with the optimal path) to forward packets to a specific destination. When a router receives a packet from a host that is better forwarded to another router on the network, the router uses an ICMP Redirect message to notify the host of the optimal path.

NetWare routing software provides both host and router implementations of the Router Discovery Protocol. The mode of operation of the Router Discovery Protocol is determined by whether the IP Packet Forwarding parameter is enabled. If IP Packet Forwarding is enabled, the Router Discovery Protocol will send Router Advertisement messages. If IP Packet Forwarding is disabled, the Router Discovery Protocol will send Router Solicitation Messages. these messages are explained in the next section.

Router Discovery Messages

The two message types that are used by the Router Discovery Protocol to communicate between hosts and routers are discussed in the following sections.

ICMP Router Advertisement Message

The ICMP Router Advertisement Message is ICMP message type 9. This message is used by routers to advertise their presence on the network and is broadcast or multicast to all hosts on the network.

This message type carries the IP address of the router and its preference level. Hosts use the preference level to determine which router to use for forwarding. The router with the highest preference becomes the default router. A value of 0x80000000 indicates the router is not to be used. Routers with this value are used *only* when other routers send ICMP Redirect messages to the host.

ICMP Router Solicitation Message

The ICMP Router Solicitation Message is ICMP message type 10. Hosts use this message to solicit router advertisements from all participating routers on the network.

Router Discovery Multicast Address

Router Discovery uses two IP multicast addresses. The IP address 224.0.0.1 is reserved to multicast the Router Advertisement Message to the hosts. The IP address 224.0.0.2 is reserved to multicast the Router Solicitation Message to the routers. If the network does *not* support multicast, then broadcast address 255.255.255.255 is used for both the Router Advertisement and Router Solicitation messages.

Multihoming

Multihoming enables an interface to assume multiple IP addresses on the same network. Multihoming can be used for all IP networks bound to a router, whether the networks are bound to on the same interface or different interfaces. The most common use of multiple addresses on the same network is to enable a Web server to operate as though it is several Web servers. One application is to use each secondary IP address to point to a different Web page on the same Web server, depending on the Domain Name System (DNS) name that is used to reach the server.

Multihoming is also commonly used with Network Address Translation (NAT), the proxy server, and the Virtual Private Network (VPN). In all cases, the secondary IP address can be configured on the same interface that has the primary IP address, or the secondary address can be configured on a different interface. When there are multiple existing interfaces, the secondary address is associated with the interface that is bound to the network that uses the same address. If the secondary address is not valid on any of the networks bound to existing interfaces, the address is rejected and an error message is produced.

When multihoming is used with the proxy server, VPN, or NAT, the secondary addresses must be configured manually as described in "Setting Up."

Configuring Database Files

TCP/IP uses four database files to convert internal data, such as IP addresses, into more identifiable and workable names. The user interface for TCPCON and other NLM files uses these database files. To inform TCP/IP of names and addresses of local nodes and networks, you must add that information to these files. The files—HOSTS, NETWORKS, PROTOCOL, and SERVICES—are cached in memory so that disk access is avoided during lookup. Because of this, TCP/IP takes up more memory. If this is not desirable, keep the size of the database small or simply delete the files.

TCP/IP finds the following four database files in the SYS:ETC\ directory:

- HOSTS File—Maps hostnames to IP addresses.
- NETWORKS File—Maps network names to network addresses.
- PROTOCOL File—Maps protocol names to IP protocol numbers.
- SERVICES File—Maps service names to TCP and UDP ports.

These files are described in the sections following this discussion.

If you are configuring TCP/IP for the first time, we recommend that you start by copying the sample database files from SYS:ETC\SAMPLES to SYS:ETC. This provides you with some examples to refer to as you add your own entries, and also provides TCP/IP with the PROTOCOL and SERVICES files.

You can modify these files with a standard text editor from any NetWare client, or you can use EDIT.NLM from the NetWare system's console. The following sections describe the formats of the files, which are compatible with the same files on standard 4.3BSD UNIX systems. The examples in the sample files can also help you create your own entries.

The files have the same names and format as the files on UnixWare™ systems and other UNIX systems. You can use FTP to transfer the files from a UNIX host.

Each database file describes a table. Each line of the file describes a separate table entry. Blank lines and comments are ignored. Comments begin with a number sign (#) anywhere in a line and include the number sign and any characters following it on the same line.

IMPORTANT: Do not use the sample addresses provided in the database files if you are connected to the Internet; these addresses are for example only.

For more information about configuring database files, refer to:

- HOSTS File
- NETWORKS File

- PROTOCOL File
- SERVICES File

HOSTS File

The SYS:ETC\HOSTS file contains information about the known hosts on the IP internetwork. Typically, it is centrally administered and distributed to all local hosts. Its format, as shown in Figure 15, is identical to */etc/hosts* on UNIX systems. Each entry provides information about a single host. An entry cannot extend beyond one line.

Figure 15 Sample HOSTS File

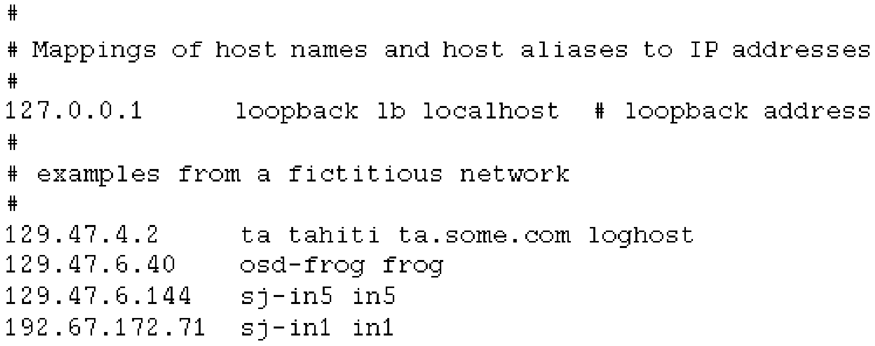

```
#
# Mappings of host names and host aliases to IP addresses
#
127.0.0.1      loopback lb localhost  # loopback address
#
# examples from a fictitious network
#
129.47.4.2     ta tahiti ta.some.com loghost
129.47.6.40    osd-frog frog
129.47.6.144   sj-in5 in5
192.67.172.71  sj-in1 in1
```

The HOSTS file entry has the following format:

```
IP_address host_name  [alias  [...]]
```

The *IP_address* is a 4-byte address in standard dotted decimal notation. Each byte is a decimal, hexadecimal, or octal value and is separated by a period. Hexadecimal numbers must start with the character pair 0x or 0X; octal numbers must start with 0.

The *host_name* is the name of the system associated with this IP address. The name cannot contain a space, tab, number sign (#), or end-of-line character. Each hostname must be unique.

The *alias* is another name for the same system. Typically, this is a shorter name. A single host can have from 1 to 10 aliases. For example, the host sales could have the following address and aliases:

```
129.0.9.5 sales sa saleshost
```

The sample file SYS:ETC\SAMPLES\HOSTS is included with the TCP/IP software. When you are configuring TCP/IP for the first time, copy the sample HOSTS file from SYS:ETC\SAMPLES to SYS:ETC. You then edit the SYS:ETC\HOSTS file. You can change your configuration at any time by editing your existing SYS:ETC\HOSTS file.

NETWORKS File

The SYS:ETC\NETWORKS file contains information about the networks in your internetwork. Each entry provides information about one network. An entry cannot extend beyond one line. Figure 16 shows a sample NETWORKS file.

Figure 16 Sample NETWORKS File

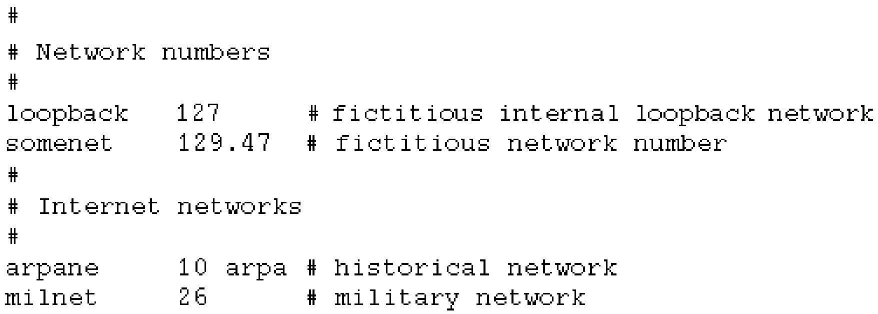

```
#
# Network numbers
#
loopback   127      # fictitious internal loopback network
somenet    129.47   # fictitious network number
#
# Internet networks
#
arpane     10 arpa # historical network
milnet     26      # military network
```

The NETWORKS file entry has the following format:

network_name network_number [/network_mask] [*alias* [...]]

The *network_name* is the name of the network associated with this network number. The name cannot contain a space, tab, number sign (#), or end-of-line character. The network name must be unique.

The *network_number* is the number of the network. Hexadecimal numbers must start with the character pair 0x or 0X. The *network_number* can be specified with or without trailing zeros. For example, the addresses 130.57 and 130.57.0.0 denote the same IP network.

The *network_mask* is the subnet mask of the network. Like IP addresses, it can be specified in octal, decimal, or hexadecimal notation. This field is optional. If not specified, the subnet mask is deduced from existing routing table entries.

The *alias* is another name for the same network; you can specify up to 10 aliases for a network.

The sample file SYS:ETC\SAMPLES\NETWORKS is included with the TCP/IP software. When you are configuring TCP/IP for the first time, copy the sample NETWORKS file from SYS:ETC\SAMPLES to SYS:ETC. Then edit the SYS:ETC\NETWORKS file. You can change your configuration at any time by editing your existing SYS:ETC\NETWORKS file.

PROTOCOL File

The SYS:ETC\PROTOCOL file, shown in Figure 17, contains information about the known protocols used on the internetwork. Each line provides information about one protocol. An entry cannot extend beyond one line.

NOTE: The PROTOCOL file is called PROTOCOLS on UNIX systems. The name is shortened to PROTOCOL because of the DOS eight-character limit.

Figure 17 Sample PROTOCOL File

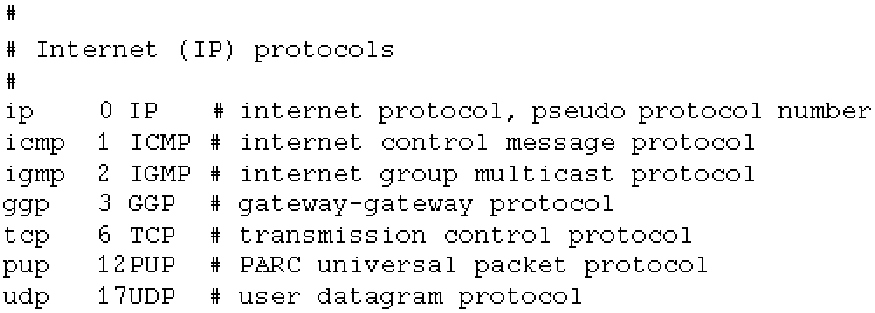

```
#
# Internet (IP) protocols
#
ip     0 IP    # internet protocol, pseudo protocol number
icmp   1 ICMP  # internet control message protocol
igmp   2 IGMP  # internet group multicast protocol
ggp    3 GGP   # gateway-gateway protocol
tcp    6 TCP   # transmission control protocol
pup    12PUP   # PARC universal packet protocol
udp    17UDP   # user datagram protocol
```

The PROTOCOL file entry has the following format:

```
protocol_name  protocol_number  [alias  [...]]
```

The *protocol_name* is the name of the Internet protocol associated with this protocol number. The name cannot contain a space, tab, number sign (#), or end-of-line character.

The *protocol_number* is the number of the Internet protocol.

The *alias* is an alternate name for the protocol.

The sample file SYS:ETC\SAMPLES\PROTOCOL is included with the TCP/IP software. When you are configuring TCP/IP for the first time, copy the sample PROTOCOL file from SYS:ETC\SAMPLES to SYS:ETC. You can then edit the SYS:ETC\PROTOCOL file. You can change your configuration at any time by editing your existing SYS:ETC\PROTOCOL file.

SERVICES File

The SYS:ETC\SERVICES file, shown in Figure 18, contains information about the known services used on the IP internetwork. Each entry provides information about one service. An entry cannot extend beyond one line.

Figure 18 Sample SERVICES File

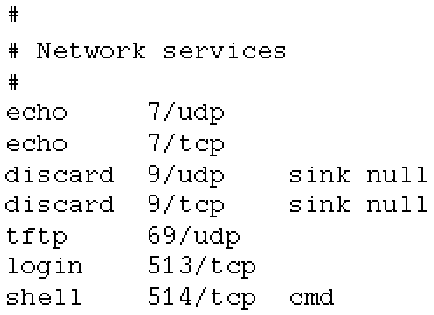

```
#
# Network services
#
echo     7/udp
echo     7/tcp
discard  9/udp    sink null
discard  9/tcp    sink null
tftp     69/udp
login    513/tcp
shell    514/tcp  cmd
```

The SERVICES file entry has the following format:

```
service_name  port_number /protocol_name  [alias  [...]]
```

The *service_name* is the name of the service associated with this port number and protocol name. The name cannot contain a space, tab, number sign (#), or end-of-line character. These are generally Application-layer, Presentation-layer, or Session-layer services, such as TFTP, FTP, SMTP, and TELNET.

The *port_number* is the number of the Internet port used by the service.

The *protocol_name* is the protocol with which the service is associated. This is generally a Transport- or Network-layer protocol, such as TCP or UDP. You must put a slash between the port number and the protocol name (for example, SMTP 25/TCP MAIL).

The *alias* is an alternate name for the service.

The sample file SYS:ETC\SAMPLES\SERVICES is included with the TCP/IP software. When you are configuring TCP/IP for the first time, you should copy the sample SERVICES file from SYS:ETC\SAMPLES to SYS:ETC. You can then edit the SYS:ETC\SERVICES file. You can change your configuration at any time by editing your existing SYS:ETC\SERVICES file.

2 Planning

This section explains what decisions must be made before you can configure TCP/IP beyond its most basic configuration.

Configuration Decisions

How you configure TCP/IP beyond the most basic configuration depends on the following decisions:

- **Whether to use the computer as a router or an end node (that is, a host)**

  The IP Packet Forwarding parameter, which controls IP packet routing, is enabled by default. This parameter permits your computer to operate as an IP router. When you want your computer to operate as an end node only, disable this parameter.

- **If you are configuring a WAN connection, whether to configure any of the following:**
  - Permanent or on-demand calls
  - WAN network mode, which can be unnumbered point-to-point, numbered single point-to-point, or multiaccess
  - Individual WAN calls
  - Static routes
  - TCP/IP header compression
  - Binding IP to an interface group
  - Dynamic address assignments

How you configure a WAN connection depends on how you want to use it and whether you use ATM (Asynchronous Transfer Mode), PPP (Point-to-Point Protocol), PPP/ISDN (Integrated Services Digital Network), X.25, or frame relay.

To configure a WAN connection, refer to "Configuring IP for a WAN Connection."

- **Whether to use Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or a mixed RIP-OSPF environment**

  RIP and OSPF are IP routing protocols. If you already have IP routers in your network environment, use the same routing protocol they use. If your network currently has no other routers, use OSPF.

  To configure your router as a RIP router, refer to "Configuring RIP."

  To configure your router as an OSPF router, refer to "Configuring OSPF."

  To configure a mixed RIP-OSPF environment, refer to both of the preceding procedures.

- **Whether to use static routes on a router**

  Static routes are useful for reducing routing traffic, providing security, accessing isolated networks, and operating as backup routes on routers. Static routes are required for on-demand connections.

  To configure static routes on a router, refer to "Configuring Static Routes for WAN Connections" on page 51 and to "Configuring Static Routes for LANs."

- **Whether to filter routes or various TCP/IP packets**

  Enable filters when you want to do either of the following:

  - Control access to any services, such as File Transfer Protocol (FTP), on your network
  - Reduce the bandwidth consumed by routing traffic

  To configure TCP/IP filters, you must enable the Filtering Support parameter in NIASCFG and then load the Filter Configuration utility (FILTCFG). For more information, refer to Setting Up in the *Filters* documentation.

- **Whether to configure router discovery**

  Router discovery enables end nodes to find an IP router on their network. If your computer is operating as a router, it can advertise itself periodically as a router. If your computer is operating as an end node, it can send queries to locate a router.

To configure router discovery, refer to "Configuring Router Discovery."

- **Whether to disable Address Resolution Protocol (ARP) or enable Proxy ARP**

  ARP is a LAN protocol that maps Internet addresses to physical addresses. IP routers and end nodes use ARP to determine a destination node’s physical address.

  An IP router using *Proxy ARP* replies to ARP requests it receives through an interface on behalf of an end node on a network attached to another interface.

  To change the default settings of the ARP or Proxy ARP features, refer to "Configuring ARP and Proxy ARP."

- **Whether to enable the router to forward directed broadcasts**

  A *directed broadcast* is a broadcast intended only for a specific group of nodes rather than all nodes on the network.

  To enable directed broadcast forwarding, refer to "Configuring Directed Broadcast Forwarding."

- **Whether to configure the router or end node as a BOOTP forwarder**

  The BOOTP protocol enables end nodes to obtain an IP address at startup time. If there is a BOOTP or Dynamic Host Configuration Protocol (DHCP) server on your internetwork, any IP routers that are configured to act as a *BOOTP forwarder* accept and forward BOOTP or DHCP requests to the server. The BOOTP or DHCP server then assigns an IP address to the end station.

  To configure BOOTP forwarding, refer to "Configuring BOOTP Forwarding."

- **Whether to configure multiple logical interfaces on a single board**

  Using multiple logical interfaces enables you to bind more than one IP network to a LAN or WAN board. Each binding operates as a separate logical interface.

  To configure multiple logical interfaces on a board, refer to "Configuring Multiple Logical Interfaces."

- **Whether to use Multihoming**

Multihoming enables an interface to be bound to multiple IP addresses on the same network. Multihoming can be used for all IP networks bound to a router, whether the networks are bound to on the same interface or different interfaces. The most common use of multiple addresses on the same network is to enable a Web server to operate as though it is several Web servers. In this application, each secondary IP address is used by a different virtual host on the same Web server. The Domain Name System (DNS) can be used to access these virtual hosts using unique host names.

Multihoming is also commonly used with Network Address Translation (NAT), the proxy server, and the Virtual Private Network (VPN).

To configure multihoming, refer to "Multihoming."

- **Whether to use the Novell IP Gateway**

  The Novell IP Gateway is used to enable IPX and IP clients on your private network to access the Internet (or other TCP/IP services) without being required to assign globally unique IP addresses to all your private systems.

3 Setting Up

The Novell® Internet Access Server 4.1 routing software provides a set of configurable parameters with which you can modify operational characteristics of the Internet Protocol (IP). You can select its routing protocol and configure it to run over a LAN or WAN connection.

To configure IP for Novell Internet Access Server 4.1 routing software, you enable the protocol, set its parameters, and bind it to a network interface. You configure all IP parameters from the Novell Internet Access Server Configuration utility (NIASCFG).

NOTE: The configuration you specify with NIASCFG does not take effect automatically. To activate the configuration, save your changes and press Esc until you see the Internetworking Configuration menu. You can then select Reinitialize System and Yes to activate your changes.

Configuring IP for a WAN Connection

This section explains the advanced features available for running IP over a WAN connection. To configure an individual WAN call, use the procedures provided under the following topics:

- Configuring IP for permanent and on-demand calls
- Configuring the WAN network mode
- Configuring individual WAN calls
- Configuring static routes for WAN routers
- Enabling TCP/IP header compression

This section also provides procedures that apply to WAN calls in general. To configure additional advanced WAN features, use the procedures provided under the following topics:

- Binding IP to an interface group
- Assigning OSPF neighbors
- Configuring dynamic address assignments

NOTE: Before you can configure IP for a WAN connection, you must configure the following information: the WAN board, the network interface, and the WAN Call Directory.

This topic contains the following sections:

- Configuring IP for Permanent and On-Demand Calls
- Configuring the WAN Network Mode
- Configuring Individual WAN Calls
- Configuring Static Routes for WAN Connections
- Enabling TCP/IP Header Compression
- Binding IP to an Interface Group
- Assigning OSPF Neighbors
- Configuring Dynamic Address Assignments

Configuring IP for Permanent and On-Demand Calls

When you installed Novell Internet Access Server 4.1 routing software, you probably accepted the default configuration for the WAN interfaces on your router. This default configuration specified the numbered single point-to-point mode; however, it did *not* specify the following:

- Permanent call
- On-demand call
- Static routes

You can continue using this default configuration, or you can change it on one or more interfaces according to the requirements of the connection. The next two sections discuss permanent and on-demand calls and explain the configuration options available for each call type.

Configuring Permanent Calls

A *permanent call* is always active between the local router and the remote peer router associated with a WAN call destination. When IP is the only protocol active on this call, the call remains active until IP is unbound from the board. Or, the call remains active until the call is disconnected manually from the Call Manager utility (CALLMGR).

A routing protocol, such as RIP or OSPF, is commonly configured to send routing traffic across a permanent WAN connection.

There are two types of permanent calls, automatic and manual. Automatic calls are brought up when a router comes up. If the connection fails when an automatic call is in process, the router immediately tries to reestablish the call. Manual calls must be brought up through CALLMGR. If the connection fails when a manual call is in process, the router does not reestablish the call.

The permanent call configuration for IP is presented in "Configuring Individual WAN Calls."

HINT: Some network modes are not suitable for multiple permanent calls over some WAN media. To decide which network mode is appropriate, refer to "Configuring the WAN Network Mode."

You can also configure the following features for permanent calls:

- **Static routes** —Static routes are optional for permanent calls. They are most often used when you do not want routing traffic on the connection. In this case, you would disable the routing protocol on the interface over which the call operates.

  The static routes for a permanent call are put in the routing table only when the call is up.

  To configure static routes for a permanent call, refer to "Configuring Static Routes for WAN Connections."

- **Call type** —Permanent WAN connections can be configured to come up automatically when the router is restarted or to require manual input to come up.

  To configure the call type, refer to "Configuring Individual WAN Calls."

Configuring On-Demand Calls

An *on-demand call* is a WAN connection between two routers that becomes active only when one router has data to send to the other. On-demand calls are well-suited for use with connections that use expensive telecommunications carriers who charge based on the amount of time the link is up.

NOTE: On-demand calls are activated by OSPF and Exterior Gateway Protocol (EGP) packets, but not by RIP packets. To avoid keeping the connection up unnecessarily, disable OSPF and EGP on the WAN interface.

The on-demand call configuration for IP is presented in "Configuring Individual WAN Calls." To use an on-demand call instead of a permanent call, change the call type of the WAN call destination to on-demand. In addition, you can configure the following features:

- **Static routes** —Static routes are required for *on-demand calls* regardless of the call type (described in the next bullet). Static routes for on-demand calls are added to the routing table whether or not the call is connected. As a result, when the router receives a packet that is destined for a network that is defined in a static route, the on-demand call is activated.
- **Call type** —On-demand WAN connections can be configured to use a dynamic routing protocol to exchange routes or to use a static routing table.

  To configure the call type, refer to "Configuring Individual WAN Calls."

For on-demand WAN calls, IP considers only data to be traffic. Maintenance data (for example, RIP updates and ICMP messages) is not considered to be data. Because maintenance data is not considered to be traffic, it does not keep the on-demand WAN link active. When only maintenance data is sent over the link, the link is brought down after the idle-timer value expires.

For each on-demand connection, you configure this time period in the Idle Connection Timeout parameter in the WAN call destination configuration, which has a default of 10 minutes.

To configure the Idle Connection Timeout parameter for an on-demand call, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > WAN Call Directory > an on-demand WAN call destination

2. Select Idle Connection Timeout.

   Select a time, between 0 and 18 hours.

3. Select Outbound Authentication.

   Specify PAP (Password Authentication Protocol), CHAP (Challenge Handshake Authentication Protocol), or PAP or CHAP. This field determines the type of authentication protocol used with the outbound connection. You cannot select NONE for on-demand calls because it prevents an authentication protocol from being used.

4. Select Password.

Enter 1 to 47 ASCII characters to specify the authentication password. This value is offered to the remote system during outbound authentication.

5 Select Local System ID.

Enter 1 to 47 ASCII characters to specify the symbolic name of this system when you place an outbound call. The default is the local server name.

6 Select Remote System ID.

Enter 1 to 47 ASCII characters to specify the symbolic name of the remote system when you place an outbound call. Usually, this value is the name of the remote server.

7 Press Esc , then select Yes to save your changes.

8 Press Esc to return to the Internetworking Configuration menu.

9 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring the WAN Network Mode

The *WAN network mode* governs how IP operates over a WAN connection. Depending on which WAN medium you use—PPP (including ISDN over PPP), X.25, frame relay, or ATM—and how you want to use the connection, you can use any of the following network modes:

- Unnumbered point-to-point

  Use this mode if you do not want the connection to occupy an entire IP network or subnet address.

- Numbered single point-to-point

  Use this mode when you have an IP network or subnet address available and you want a single, dedicated connection to a peer router. This mode is usually used only when the unnumbered mode is not supported by the remote router.

- Multiaccess

  Use this mode when you want to use multiple connections to several peer routers through a single interface or an interface group.

Table 3 indicates the modes that can operate over each WAN media.

Table 3 WAN Media and Compatible Network Modes

	WAN Network Mode		
WAN Medium	**Unnumbered Point-to-Point**	**Numbered Single Point-to-Point**	**Multiaccess**
PPP (including ISDN over PPP)	Yes	Yes	Can be used only with interface groups.
X.25	Yes	Limited to one connection	Yes
Frame relay	Yes	Limited to one connection	Yes
ATM	Yes	Limited to one connection	Yes

Use this table as a guide when you are choosing a network mode for your WAN connections.

IMPORTANT: If you are configuring interfaces that are part of a PPP interface group, you must bind to the interface group to configure the WAN network mode.

The rest of this section provides additional information about each network mode and explains how to configure them for permanent and on-demand calls.

IMPORTANT: If you are configuring your router to act as an ISP router or to connect to an ISP router, refer to "Configuring Dynamic Address Assignments" on page 68 before configuring the WAN network mode.

Understanding Unnumbered Point-to-Point Mode

Unnumbered point-to-point mode is so named because the router's WAN interfaces do not use IP addresses. This mode is useful when you do not want the connection to occupy an entire IP network or subnet address.

You can use unnumbered point-to-point mode with any WAN medium that supports multiple connections to remote peer routers, such as X.25, frame relay, or ATM. If you are using unnumbered point-to-point mode over X.25 or ATM, you can select several WAN call destinations for simultaneously active permanent calls.

With PPP as a single interface, you can configure either one permanent call or multiple on-demand calls. A permanent call is always active when IP is bound to the interface. Because PPP supports only a single call on an interface, an on-demand call cannot be made when a permanent call is configured. You can configure multiple on-demand calls on a PPP interface; however, only one call can be active at a time.

With PPP as an interface group, you can configure multiple permanent and on-demand calls as long as there are enough interfaces for each of the calls.

NOTE: If any of the remote peers is a third-party router, make sure it supports unnumbered point-to-point mode for IP. Some third-party routers do not.

How to Configure Unnumbered Point-to-Point Mode

To configure unnumbered point-to-point mode, complete the following steps:

NOTE: All interface-specific configurations, such as routing protocol and header compression, apply to *all* connections through the same unnumbered interface.

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > Press Ins > TCP/IP from the list of configured protocols

2 Select Network Interface or Each Interface in a Group.

Select Network Interface to bind to a specific interface. Select Each Interface in a Group to bind to an interface group.

Either option can be selected for unnumbered point-to-point mode.

3 Select a configured network interface or an interface group.

The Binding TCP/IP to a WAN Interface menu is displayed.

4 Select the WAN Network Mode field.

The default, Numbered Point-to-Point , is displayed.

5 Press Enter , then select Unnumbered Point-to-Point.

Because unnumbered point-to-point mode does not use IP addresses, you cannot select the Local IP Address and Subnetwork Mask of Connected Network fields.

IMPORTANT: Each router must have an IP address configured on at least one LAN or WAN interface.

6 Press Esc until you are prompted to save your changes, then select Yes.

7 Press Esc to return to the Internetworking Configuration menu.

8 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

NOTE: When configuring an on-demand or permanent ATM or X.25 call, if either end of the connection can initiate the call, enable inbound authentication on the network interface as described in Setting Up in the *NetWare Link/ATM* documentation or in Setting Up in the *NetWare Link/X25* documentation. This automatically creates an authentication entry for an inbound call from the other end of the connection.

9 Configure a WAN call destination as described in "Configuring Individual WAN Calls."

Understanding Numbered Single Point-to-Point Mode

Numbered single point-to-point mode is typically used with PPP, which supports either a single permanent or single on-demand dedicated connection to a remote peer router. You can also use this mode with WAN media that support multiple connections, such as X.25, frame relay, or ATM, but you are limited to having only one dedicated connection. Numbered single point-to-point mode is well-suited for a connection that has just one destination, such as a link between a local branch office and the main office.

Numbered single point-to-point mode uses a single IP address for the connection; therefore, you can bind IP only once to the interface.

How to Configure Numbered Single Point-to-Point Mode

To configure numbered single point-to-point mode, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > Press Ins > TCP/IP from the list of configured protocols

2 Select Network Interface or Each Interface in a Group.

Select Network Interface to bind to a specific interface. Select Each Interface in a Group to bind to an interface group.

Either option can be selected for Numbered Single Point-to-Point mode.

3 Select a configured network interface or an interface group.

The Binding TCP/IP to a WAN Interface menu is displayed. The WAN Network Mode field is displayed with a default value of Numbered Point-to-Point.

NOTE: The Remote Router Will Dynamically Assign the IP Address parameter should be left at the default value of No.

4. Configure the following parameters:
   - Local IP Address —Enter the IP address of the local interface.
   - Subnetwork Mask of Connected Network —Enter the subnet mask of the IP network to which the interface is connected. This mask should match the mask on the remote router.
5. Press Esc until you are prompted to save your changes, then select Yes.
6. Press Esc to return to the Internetworking Configuration menu.
7. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.
8. Configure a WAN call destination as described in "Configuring Individual WAN Calls."

Understanding Multiaccess Mode

Multiaccess mode is typically used with X.25, frame relay, and ATM, which support multiple, simultaneous connections to remote peer routers. You can also use multiaccess mode with PPP interface groups to accept calls for a group of interfaces.

Multiaccess mode supports multiple logical networks; that is, you can bind IP to the local WAN interface for each IP network represented on the connection. Figure 19 shows an example of two logical networks, 1.0.0.0 and 2.0.0.0, each supporting two remote peer routers, which are attached to a local router interface through an X.25 WAN.

Figure 19 Two Logical IP Networks on an X.25 WAN

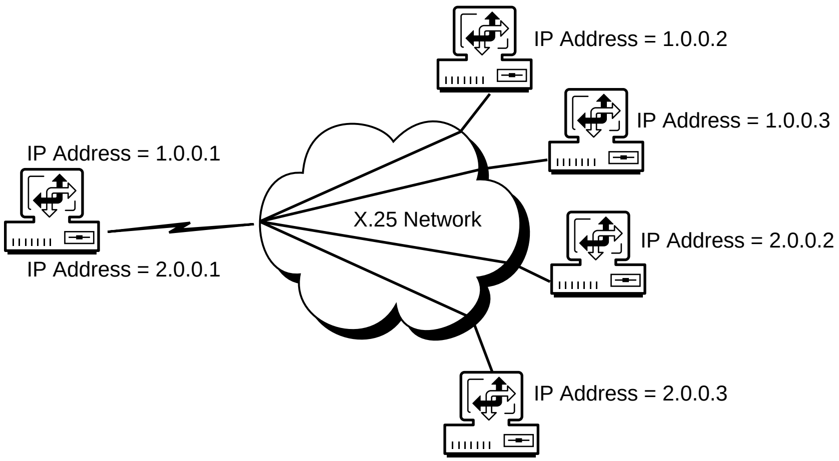

When you bind IP to the local address 1.0.0.1, you add the remote peer routers, 1.0.0.2 and 1.0.0.3, to the WAN Call Destination List. In the same way, when you bind IP to the local address 2.0.0.1, you add the remote peer routers, 2.0.0.2 and 2.0.0.3, to the WAN Call Destination List. For this configuration, you have two bindings—one for each logical network—and two remote peers per network, each mapped to a WAN call destination.

With X.25 and ATM, you can configure multiple, simultaneous permanent calls. Additionally, you can configure multiple on-demand calls for X.25 and ATM. You do not have to use the same call type for these multiaccess connections; you can use any combination of permanent and on-demand calls.

With a PPP single interface, you can configure either one permanent call or multiple on-demand calls. A permanent call is always active when IP is bound to the interface. Because PPP supports only a single call on an interface, an on-demand call cannot be made when a permanent call is active or established. You can configure multiple on-demand calls on a PPP interface; however, only one call can be active at a time.

With frame relay, because all calls are incoming calls, you do not need to configure any WAN call destinations unless you are using static routes over the WAN.

Broadcasts are not supported on multiaccess interfaces. Therefore, routing information must be sent directly to each peer router on the interface. If you want to run a routing protocol over one of the connections, you must do the following:

- Enable RIP or OSPF at the local interface.
- Provide the IP address of the remote peer OSPF router in the OSPF neighbor list.
- Enable RIP under the WAN call destination to the remote peer.

On an incoming connection, the local router must have a way to discover the IP address of the remote peer router. Frame relay uses the Inverse Address Resolution Protocol (Inverse ARP) for this purpose. PPP uses the Internet Protocol Control Protocol (IPCP). If a remote peer router does not support Inverse ARP or IPCP for an incoming call, or if you are using X.25 or ATM for an incoming or outgoing call, you must map the WAN call destination associated with the remote peer router to its IP address using the procedure described in "Configuring Individual WAN Calls." Although frame relay does not use WAN call destinations, Novell Internet Access Server 4.1 routing software enables you to configure them if the remote peer router does not support Inverse ARP. You are not required to provide this mapping for frame relay if the remote peer router is running Novell Internet Access Server 4.1 routing software, unless you are using static routes over the WAN.

IMPORTANT: Multiaccess mode operates best when all routers are connected in a mesh topology. If the routers are not connected in a mesh topology, use unnumbered point-to-point mode for each connection between the local router and a remote peer router. For configuration instructions, refer to "How to Configure Unnumbered Point-to-Point Mode."

How to Configure Multiaccess Mode

To configure multiaccess mode, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > Press Ins > TCP/IP from the list of configured protocols

2 Select Network Interface or Each Interface in a Group.

Select Network Interface to bind to a specific interface. Select Each Interface in a Group to bind to an interface group.

Either option can be selected for multiaccess mode.

3 Select a configured network interface or an interface group.

The Binding TCP/IP to a WAN Interface menu is displayed.

4 Select the WAN Network Mode field. Press Enter , then select Multi-Access.

5 Configure the following parameters:

- Local IP Address —Enter the IP address of the local interface.
- Subnetwork Mask of Connected Network —Enter the subnet mask of the IP network to which the interface is connected.

6 Press Esc until you are prompted to save your changes, then select Yes.

7 Press Esc to return to the Internetworking Configuration menu.

8 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Individual WAN Calls

Use this procedure to configure individual WAN calls. The parameters in this procedure apply only to one WAN call. Depending on how you have configured the WAN network mode, you might or might not see all the parameters that appear in this procedure. With frame relay, because all calls are incoming calls, you do not need to configure any WAN call destinations unless you are using static routes over the WAN.

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > an existing binding with the WAN network mode configured

2 Select WAN Call Destinations , then press Ins.

The parameters in this menu apply only to this WAN call. You can select one or more WAN call destinations, depending on how you configured the WAN Network Mode parameter as described in “Configuring the WAN Network Mode.”

Configure the following parameters:

- WAN Call Destination —Name of the WAN call destination that you want to configure.
- Type —Select Automatic or Manual for permanent calls. Select Routed On Demand or Static On Demand for on-demand calls. The Automatic value enables the WAN call to be brought up with the router. Also, if the WAN call connection is broken, the software automatically tries to reestablish the call. The Manual value requires that the WAN call is brought up manually from CALLMGR. When you select Routed On Demand , RIP is automatically enabled. When you select Static On Demand , RIP is automatically disabled.

IMPORTANT: If multiaccess mode is used for X.25 or ATM, you must assign a value to the Remote IP Address field.

- Remote IP Address —Enter an IP address that is associated with the WAN call destination. When TCP/IP sends an IP packet to this address, it uses this mapping to determine the WAN call destination for the packet.
- Verify Remote Address —Select Yes to verify that the remote IP address, specified previously, is announced by the remote router during IPCP negotiations. This option is only for PPP connections.
- Header Compression —Select Enabled to compress the Transmission Control Protocol (TCP) and Internet Protocol (IP) headers on serial point-to-point connections. This parameter applies only to PPP interfaces.
- Static Routing Table —Select this option to configure static routes for this WAN call. Refer to "Configuring Static Routes for WAN Connections."

3 If you want to customize RIP parameters for this WAN call, complete the following steps:

3a Select RIP Bind Options. You must configure these parameters for each WAN call.

The RIP parameters in this menu apply to the WAN call destination. The RIP parameters configured here override the RIP parameters configured under the Binding TCP/IP to a WAN Interface menu, except Status. For example, if you enable RIP under WAN Call Destination, you also must set Status to Enabled under the Binding TCP/IP to a WAN Interface menu. Configure the following parameters:

- Status —The default is Enabled. If this system is configured as a router, this parameter allows RIP to exchange routing information with other routers. If this system is configured as a host, this parameter allows RIP to discover routers on the assigned WAN call destination. As a host, it listens to RIP messages, but it does not send them. If you do not want to run RIP over this connection, select Disabled.
- RIP Version —Select the version of RIP that is used on this WAN call destination. RIP I is the standard RIP used by many end nodes and routers. If there are nodes on your network that support only RIP I, select either RIPI or RIPI & RIPII. RIP II is

an enhanced version of RIP I that includes the subnet mask in the routing information. If your network consists of subnets of varying sizes, RIP II improves reachability.

- RIP Mode —Select the RIP mode that is used on this WAN call destination. Select Normal to send and receive RIP packets (RIP I, RIP II, or both). Select Send Only to send RIP packets. Select Receive Only to receive RIP packets.
- RIP II Options —Select this option to view or modify RIP II options.

  Authentication —Enable authentication when there are routers that you do not want this router to exchange RIP II routing information with.

  Authentication Password —Enter a password to allow access to your router. Authentication works only when this password matches the password on another router. The default is the null string.
- Cost of Interface —Specifies the cost that RIP associates with this network. It is used when advertising a path to other routers. RIP allows a maximum cost of 15. Usually, you do not need to change the default unless you want to discourage other routers from using this path.
- Originate Default Route —Select Enable to cause RIP packets sent on this interface to contain only the default route.
- Poison Reverse —Select Enable to allow RIP to use poison reverse in RIP updates. If you disable this field, RIP traffic is reduced a small amount at a small cost in stability.
- Split Horizon —Select Enable to reduce loops between two routers. Split horizon prohibits a router from propagating a route over the same port that supplied the route.
- Update Time —Enter the number of seconds that the router sends RIP update messages. The default value is to broadcast an update message every 30 seconds. If a router does not receive an update within six times the value of this parameter, the route is invalidated.
- Expire Time —Enter the time after which the route is invalidated.

- Garbage Time —Enter the time an invalidated route is saved. After the value of the Garbage Time parameter expires, the route is discarded.

3b If you have enabled RIP on this WAN call destination, make sure RIP is enabled at the remote interface and uses the same RIP version.

NOTE: You cannot configure OSPF for individual WAN calls. If you do not want to run OSPF over this WAN connection, disable OSPF by selecting OSPF Bind Options on the Binding TCP/IP to a WAN Interface menu.

4 Press Esc until you are prompted to save your changes, then select Yes.

5 Press Esc to return to the Internetworking Configuration menu.

6 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Static Routes for WAN Connections

Static routes can be configured for both on-demand and permanent WAN calls, but they are most commonly used with on-demand links. When you do not want routing traffic to cross a WAN link, you can manually configure remote routes on the local router as *static routes*. With the static route configured, an on-demand link can remain inactive until data needs to cross it. IP hosts that need to reach remote destinations send packets to their local IP router that has the static routes configured, assuming the packets can reach their destination. The local router stores the packets and tries to establish a connection to the remote router. After the local router completes the call, it forwards the stored packets to the remote router, which then forwards them to their destination. Static routes for on-demand calls are always present in the routing table.

You can also configure static routes for permanent connections to provide access to isolated networks, reduce routing traffic, provide security, and operate as backup routes. In addition, using static routes and disabling dynamic routing protocols, such as RIP, over slow links improves performance. Static routes for permanent calls are in the routing table only when the permanent calls are established.

NOTE: Use this procedure to specify static routes for any WAN connection.

How to Configure Static Routes

To configure one or more static routes for an on-demand or permanent call, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > an existing binding with the WAN network mode configured

2 Select WAN Call Destinations , then press Ins or select an existing call destination.

The parameters in this menu apply only to this WAN call. Configure WAN call destination parameters if you want to override the WAN interface parameters.

Configure the following parameters:

- WAN Call Destination —Select the name of the WAN call destination that you want to configure.
- Type —Select Automatic or Manual for permanent calls. Select Routed On Demand or Static On Demand for on-demand calls. When you select Routed On Demand , RIP is automatically enabled. When you select Static On Demand , RIP is automatically disabled.
- Remote IP Address —Enter an IP address that is associated with the WAN call destination. When TCP/IP sends an IP packet to this address, it uses this mapping to determine the WAN call destination for the packet.

3 Select Static Routing Table , then press Ins.

4 Configure the following static route parameters:

- Route to Network or Host —Enter the destination that can be reached through WAN connection, which can be a default route, a single IP host, or an IP network (that is, a group of hosts).
- IP Address of Network/Host —Enter the address of the destination network or host. To select from a list of symbolic network or hostnames and addresses, press Ins. The list of symbolic network names and addresses comes from the SYS:\ETC\NETWORKS file. The list of symbolic host names and addresses comes from the SYS:\ETC\HOSTS file.
- Subnetwork Mask —Enter the subnet mask of the destination if the destination is an IP network. If you do not specify a value, the natural mask is used.
- Metric for this route —Enter the number of hops to the destination. This metric is directly proportional to the cost of the route. Given two

routes to the same destination, the router chooses the lower-cost route.

If you want to use the static route as a *backup route* to a dynamic route, select a value that is higher than the cost associated with the dynamic route so that the dynamic route remains the preferred route under typical conditions.

Do not set this metric value to 16 unless you want to disable the route.

- Type of route —Specify whether the static route is *Active* or *Passive*. This parameter specifies whether the next hop router for this route actively advertises the route to this network.

  Usually, static routes are not advertised and are categorized as passive routes. When a route is marked as active, TCP/IP expects the next hop router to advertise the route regularly. If a router stops advertising an active static route, TCP/IP assumes the route is no longer available and deletes it from the routing table.

  If the static route is active and the router discovers a lower-cost dynamic route to the same destination, it uses the lower-cost route instead of the active static route. If the lower-cost route becomes unavailable, the router returns to using the active static route.

  If you want to use the static route as a backup route, select Active.

  A passive static route is always used, regardless of whether the router discovers a lower-cost route to the same destination.

5. Press Esc until you are prompted to save your changes, then select Yes.
6. Press Esc to return to the Internetworking Configuration menu.
7. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Enabling TCP/IP Header Compression

When used on a PPP interface, TCP/IP header compression increases the effective throughput of TCP/IP packets. Header compression reduces the size of the combined TCP/IP packet headers to just a few bytes. UDP/IP packet headers are not compressed.

NOTE: TCP/IP header compression can be used only on PPP interfaces.

Enabling TCP/IP Header Compression at the Interface Level

To enable TCP/IP header compression on this interface, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > an existing binding > Expert TCP/IP Bind Options

2. Select the Header Compression field, then select Enabled.
3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.
5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Enabling TCP/IP Header Compression on Individual WAN Calls

To enable TCP/IP header compression for a WAN call, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > an existing WAN binding > WAN Call Destinations > an existing WAN call destination

2. Select the Header Compression field, then select Enabled.

   This value overrides the value configured for header compression under the Binding TCP/IP to a WAN Interface menu.

3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.
5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Binding IP to an Interface Group

An interface group is a grouping of several PPP or X.25 interfaces with similar characteristics. Interface groups are defined during configuration. Interfaces that belong to a group can be used interchangeably by a WAN call. To configure an interface group, load NIASCFG and select Configure NIAS > Protocols and Routing > Network Interfaces > Group and enter the same group name for each interface that you want to belong to the group.

Defining an interface group lets you make an on-demand call on any of several network interfaces without creating an individual WAN call destination for each interface. By specifying an interface group name in place of the interface name in the WAN call destination, an available interface is selected automatically from the group when a call is made. Interface groups are most commonly used for asynchronous on-demand connections.

To create a WAN connection, you can bind to an interface group as a whole without binding to an individual interface. For an interface group, only the unnumbered and multiaccess modes are practical options for a WAN network.

Assigning OSPF Neighbors

Use this procedure to run OSPF on a multiaccess frame relay, X.25, or ATM WAN connection. It allows you to assign remote IP addresses to OSPF neighbors when you have configured the network mode as multiaccess. Before assigning OSPF neighbors, enable OSPF. Refer to "Configuring OSPF" on page 74, then complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > a multiaccess WAN binding > OSPF Bind Options > Neighbor List > Press Ins

2 Enter the IP address of the OSPF router at the other end of the connection, then press Enter.

For X.25 and ATM, this is the same IP address as the one you mapped to the WAN call destination associated with this connection.

2a Make sure OSPF is enabled at the local interface.

Press Esc until you return to the OSPF Bind Options menu. Make sure the Status field is set to Enabled. If it is not, select the field, then select Enabled.

2b Make sure OSPF is enabled at the remote interface.

3 Press Esc until you are prompted to save your changes, then select Yes.

4 Press Esc to return to the Internetworking Configuration menu.

5 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Dynamic Address Assignments

Use this procedure to configure your router so that it can dynamically obtain an IP address from your Internet Service Provider (ISP), or to configure your router with a range of IP addresses to dynamically assign to dial-up routers through IPCP. This procedure is valid only on a PPP connection.

How to Configure Your Router to Connect to a Remote Router or ISP Router

To configure your router to connect to a remote router or ISP router, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > press Ins > TCP/IP > a WAN network interface or interface group

2. Select Remote Router will Dynamically Assign the IP Address.

   If you need an IP address for your WAN link for use with Network Address Translation (NAT) or the IP gateways, you have the following two options:

   - Select Yes for this parameter to configure your local router to obtain an IP address from the ISP router. In this case, the WAN network mode is automatically set to dynamic numbered point-to-point. Note that this WAN network mode is not suitable for use with the proxy server, the Virtual Private Network (VPN), or any other feature that requires static addresses.

   - Select No for this parameter. After you complete this procedure, set the WAN network mode to numbered point-to-point as described in "How to Configure Numbered Single Point-to-Point Mode."

   If you do not need an IP address on the WAN link and the ISP allocates a block of IP addresses to you for your hosts, select No for this option. After you complete this procedure, set the WAN network mode to unnumbered point-to-point as described in "How to Configure Unnumbered Point-to-Point Mode." The block of IP addresses is then used by the hosts on your LAN segment to access the Internet.

   If you need an IP address on the WAN link and want to use the rest of the block of IP addresses the ISP allocated to you for your hosts, select No for this option. After you complete this procedure, set the WAN network mode to numbered point-to-point as described in or "How to Configure Numbered Single Point-to-Point Mode." You must subnet the addresses as described in RFC 1918, as determined by the size of the block of IP addresses given to you by the ISP. For a brief description of RFC 1918, refer to the description of subnetting in "Understanding."

3 Select WAN Call Destinations and press Ins.

4 Configure the WAN call destination as described in "Configuring Individual WAN Calls."

5 Select Static Routing Table and press Ins.

6 If you have only one WAN link to the ISP, set Route to Network or Host to Default Route. Otherwise, configure any needed network or host routes as described in "Configuring Static Routes for WAN Connections."

We strongly recommend that you use static routes instead of a dynamic routing protocol. Because ISPs tend to assign addresses that belong to a subnet or network that is different from its WAN address, the local and remote routers do not accept RIP packets from the other side of the WAN connection to update their routing tables. Therefore, you should configure static routes to reach hosts on the Internet.

7 Press Esc until you are prompted to save your changes, then select Yes.

8 Press Esc to return to the WAN Call Destination to IP Address Mapping Configuration menu.

9 Select RIP Bind Options and set Status to Disabled.

We recommend that you disable RIP for a WAN call to the ISP for the following reasons:

- To avoid maintaining a large routing table
- To avoid RIP updates every 30 seconds over the WAN

10 Press Esc until you are prompted to save your changes, then select Yes.

11 Press Esc to return to the Internetworking Configuration menu.

12 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

13 After successfully connecting to the ISP, you can use TCPCON, PPPCON, or the CONFIG command to determine the IP address that is bound to your WAN interface.

How to Configure Your Router to Assign IP Addresses

To configure your router to dynamically assign IP addresses using IPCP, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > press Ins > TCP/IP > a WAN network interface or interface group

2. Select Remote Router will Dynamically Assign the IP Address.

   To configure your router to act as an ISP router and dynamically assign IP addresses to dial-up routers, select No. If you have only one incoming interface for this router, leave the WAN network mode at the default value of numbered single point-to-point. If you are using interface groups to allow multiple connections to the router, set the WAN network mode to multiaccess as described in "How to Configure Multiaccess Mode."

3. Set Local IP Address to the address that will be used for the WAN link.

4. Set Subnetwork Mask of Connected Network to the appropriate value the IP address used for the WAN link.

5. Select WAN Call Destinations and press Ins.

6. Configure the WAN call destination with Type set to Manual as described in "Configuring Individual WAN Calls."

7. If you are not using a dynamic routing protocol on both the local and remote routers, select Static Routing Table and press Ins.

   Configure static network or host routes on your router for the dial-up router's networks or hosts. To configure a static network or host routes, refer to "Configuring Static Routes for WAN Connections."

8. Press Esc until you are prompted to save your changes, then select Yes.

9. Select RIP Bind Options , set Status to Enabled , and set Originate Default Route to Enabled.

10. Press Esc until you are prompted to save your changes, then select Yes.

11. Select Expert TCP/IP Bind Options.

12. Select IPCP Address Assignment Range.

13. Select Range Start.

    The IP addresses you assign to Range Start must be within the local network address and network mask for the interface. The value you enter here must be less than the value in the Range End field.

    You can include the local address in the range; however, it will not be used for address assignment.

14 Select Range End.

The IP addresses you assign to Range End must be within the local network address and network mask for the interface. The value you enter here must be greater than the value in the Range Start field.

15 Press Esc until you are prompted to save your changes, then select Yes.

16 Press Esc to return to the Internetworking Configuration menu.

17 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring RIP

RIP is probably the most common IP routing protocol in use. It is widely available and presents few obstacles to interoperability with other IP internetworks, most notably the Internet.

RIP performs sufficiently well in small IP internetworks that have simple architectures and few routers. However, RIP reveals its limitations in the large, complex internetworks that have become common in government and private-sector organizations throughout the world. Its most apparent limitations are the following:

- All subnets must be contiguous.
- The entire network must use the same subnet mask.
- RIP routes are limited to 15 hops.

To overcome or ease some of these limitations, the internetworking community developed various enhancements to RIP. *RIP II* , for example, is an enhanced version of RIP that supports variable-length subnet masks. It carries a field that contains the subnet mask of the destination network. RIP II also supports the use of subnet zero, whose addresses were reserved under the original IP specification. When configuring RIP on your router, you can run RIP, RIP II, or both on a single interface.

NOTE: Not all third-party routers support RIP II.

You can also enable *poison reverse* on an interface. This is a mechanism that causes RIP to advertise a route back through the same path from which it learned the route, but with a hop count of 16—that is, unreachable. Although poison reverse prevents routing loops, the unreachable routes carried in each RIP packet increase the bandwidth consumed by RIP traffic. This increase becomes significant in large internetworks.

RIP enables you to assign a *cost* value between 1 and 15 to each network interface you configure. This enables you to establish a preferred route according to the type of network media connected to the interface. For example, you might want to increase the cost of an interface that uses a slow link so that, given the choice, RIP uses the interface to a faster, less costly link. The default cost for each interface is 1. Do not increase this value on an interface unless you want to discourage its use as an eligible routing path.

RIP can run over most WAN connections, depending on which call type you use. On-demand calls, for example, typically use static routes instead of an active routing protocol. While using RIP over on-demand calls, RIP updates will not activate the call. Permanent calls on an IP network typically use a routing protocol, such as RIP, to communicate routing information. However, they can also use static routes to conserve bandwidth. RIP can also run over a nonbroadcast multiaccess network, such as X.25. For more information about using RIP over WAN connections, refer to “Configuring IP for a WAN Connection.”

When choosing an IP routing protocol, consider the following guidelines:

- If the IP internetwork is small and uses no routing protocol besides RIP, continue using RIP.

  To configure RIP on the router, refer to “How to Configure RIP.”

  However, if the network will continue to grow and perhaps become part of a larger IP internetwork, you should consider *migrating* the network from RIP to OSPF.

- If the internetwork uses variable-length subnets or has third-party routers that support RIP II, use RIP II or OSPF.

  To configure RIP II, refer to “How to Configure RIP.” To configure OSPF, refer to “How to Configure OSPF.”

- If the internetwork has some third-party routers that support RIP II and others that do not, use RIP I *and* RIP II.

  For instructions on enabling RIP I and RIP II simultaneously on a network interface, refer to “How to Configure RIP.”

- If you are currently building a large IP internetwork, use OSPF.

  You can also run RIP and OSPF concurrently; for more information, refer to “How to Configure OSPF.”

How to Configure RIP

To enable RIP routing on the router and to configure RIP on a network interface, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

2. Make sure RIP routing is enabled globally. Set the RIP field to Enabled.

   This is the default setting.

   If you want to disable RIP routing on a single interface, set the Status parameter in the RIP bind options to Disabled. This action is described in Step 3.

3. Press Esc twice to return to the Internetworking Configuration menu, then select the following parameter path:

   Select Bindings > an existing binding > RIP Bind Options

   Configure the following parameters:

   - Status —Status of RIP routing on this interface. RIP routing is enabled by default; to disable RIP routing only on this interface, select this parameter, then select Disabled.

   - RIP Version —Version of RIP to use on this interface. Select one of the following options:

     RIPI —Standard version of RIP used by most IP routers and end nodes. This is the default option.

     RIPI & RIPII —Both versions of RIP. Select this option if your internetwork has nodes that support both RIP I and RIP II.

     RIPII —Enhanced version of RIP that supports variable-length subnet masks.

   - RIP Mode —Mode of the RIP version you selected in RIP Version.

     Normal —Causes the router to send and accept RIP packets, RIP I, RIP II, or both.

     Receive Only —Causes the router to only receive RIP packets.

     Send Only —Causes the router to broadcast, in RIP packets, only the entries in its own routing table.

     Some end nodes learn routes only by listening to RIP, even if portions of the internetwork run OSPF. Select Send Only if you want the router to broadcast the OSPF routes in its RIP I packets so that every end node can learn *all* available routes.

The RIP Bind Options menu also includes the following parameters:

- Cost of Interface
- Originate Default Route
- Poison Reverse
- Split Horizon
- Update Time
- Expire Time
- Garbage Time
- RIP II Options

IMPORTANT: Because the default settings for these parameters are suitable for most IP networks, you should change them only for a specific purpose. Incorrectly configuring these parameters can increase routing traffic or cause loss of connectivity on your network.

For a WAN interface, you can configure the parameters for each WAN call. Refer to Step 3 of "Configuring Individual WAN Calls."

4 Press Esc until you are prompted to save your changes, then select Yes.

5 Press Esc to return to the Internetworking Configuration menu.

6 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring OSPF

OSPF was developed to satisfy the need for a scalable, open-standards routing protocol for large IP internetworks. It is a *link state* protocol that provides highly efficient routing and fast convergence.

OSPF makes large internetworks more manageable by enabling you to partition them into administrative domains called *areas*. Areas impose a hierarchy to the internetwork. All OSPF areas are connected to a central *backbone* area by an *Area Border Router* (ABR). The ABR shares OSPF routing information between the area and the backbone.

When configuring an OSPF area, you assign to it a 4-byte decimal number called the *Area ID*. You also indicate which of the router's network interfaces belong to the area and whether the area is a *stub area*.

Novell Internet Access Server 4.1 routing software supports the use of *virtual links* between OSPF routers. A virtual link patches together a partitioned backbone. It creates a direct point-to-point link between the ABRs that connect the partitioned backbone areas through the *transit area*.

WARNING: Because of the complexity and high probability of misconfiguring virtual links, we recommend against using them on your network.

Most IP internetworks in use today are not pure OSPF networks; that is, portions of these internetworks still employ other routing protocols, such as RIP. OSPF uses an *Autonomous System Boundary Router* (ASBR) to import and propagate routing information from these protocols. ASBRs are always located on the border of an OSPF domain. When configuring OSPF, you can enable your router to operate as an ASBR. For an ASBR to import RIP routes learned through an interface, RIP must be enabled on that interface.

Each OSPF router has its own *Router ID* , a 4-byte number that uniquely identifies the router and enables it to participate in informational exchanges with neighboring routers. The default Router ID is the IP address of the first interface bound to IP on the router. Although NIASCFG enables you to change the Router ID, you should use the default unless you need a simpler numbering scheme for administrating several hundred routers on an internetwork.

WARNING: If you are using an unnumbered point-to-point interface, we recommend that you configure a unique router ID.

Optionally, OSPF can be configured to *authenticate* its packets by providing an *authentication key* —an 8-byte, alphanumeric password—in each OSPF packet header. OSPF authentication gives you administrative control over which routers participate in link state exchanges on the internetwork. A router without proper authentication is excluded from these exchanges and, essentially, from performing any OSPF routing whatsoever. Novell Internet Access Server 4.1 routing software enables you to provide authentication for an area and to provide an authentication key for each network to which the router is connected. By default, authentication is turned off.

OSPF enables you to assign a *cost* value to each network interface you configure. This enables you to establish a preferred route according to the type of network media connected to the interface. For example, you might want to increase the cost of an interface that uses a slow link so that, given the choice, OSPF uses the interface to a faster, less costly link.

Like RIP, OSPF can run over most WAN connections, depending on which call type you use. On-demand calls, for example, typically use static routes instead of an active routing protocol.

WARNING: An active routing protocol, such as OSPF, should not be used on an on-demand link because it will periodically bring up the link and will cause the link to continue to stay up.

Permanent calls on an IP network typically use a routing protocol, such as OSPF or RIP, to communicate routing information. However, they can also use static routes to conserve bandwidth. OSPF can also run over a nonbroadcast multiaccess network, such as X.25 or frame relay, but you must provide the IP address of the peer OSPF router at the other end of each connection. For more information about configuring OSPF for use over WAN connections, refer to "Assigning OSPF Neighbors."

WARNING: Novell Internet Access Server 4.1 routing software enables you to run OSPF and RIP on the same router, but under normal circumstances, you should run them separately on different interfaces. Although an ASBR must run both protocols so that it can import RIP routes and propagate them to other OSPF routers, you should not run both on too many other routers in your OSPF domain. Doing so consumes additional network bandwidth and router memory, and might even create routing loops.

How to Configure OSPF

The extent to which you must configure OSPF depends on the characteristics of your network, such as its size and topology, and whether it uses other IP routing protocols besides OSPF. To help you configure only what is necessary, this section provides the following procedures:

- Basic OSPF configuration
- Advanced OSPF configuration

Basic OSPF Configuration

To enable OSPF routing on the router and to configure OSPF on a network interface, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

2. Select the OSPF field, then select Enabled.

This action enables OSPF routing globally on the router. If you want to disable OSPF routing on a single interface, set the Status parameter to Disabled as described in Step 3.

3 Press Esc repeatedly to return to the Internetworking Configuration menu, then select the following parameter path:

Select Bindings > an existing binding > OSPF Bind Options

The Status field indicates whether OSPF routing is active on this interface. OSPF routing is enabled by default; to disable OSPF routing only on this interface, select Status , then select Disabled.

The OSPF Bind Options menu also includes the following parameters:

- Cost of Interface
- Area ID
- Priority
- Authentication Password
- Hello Interval
- Router Dead Interval
- Neighbor List

IMPORTANT: Because the default settings for these parameters are suitable for most IP networks, you should change them only for a specific purpose. Misconfiguring these parameters can increase routing traffic or cause loss of connectivity on your network.

The Neighbor List parameter is used when you want to run OSPF over a WAN connection that uses multiaccess mode. Configuring this parameter is explained in "Assigning OSPF Neighbors."

4 Press Esc until you return to the Internetworking Configuration menu. Select Yes if you are prompted to save your changes.

5 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Advanced OSPF Configuration

To configure advanced OSPF features, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

2 Select OSPF Configuration.

The OSPF Configuration menu is displayed. This menu includes the following parameters:

- Router ID
- Virtual Link Configuration
- IP Load Sharing

IMPORTANT: Most network configurations do not require you to change these parameters.

3 To configure an ASBR, select Autonomous System Boundary Router , then select Enabled.

Enabling this parameter enables the router to operate as an ASBR. In this capacity, the router advertises non-OSPF routes, such as those generated by RIP and EGP. In addition, static routes and direct routes to the OSPF domain are advertised. This is necessary to preserve connectivity throughout an internetwork that uses routing protocols other than OSPF. This parameter should be configured only on routers that connect an OSPF area to an area that uses a different routing protocol.

NOTE: Do not enable this parameter on an internetwork that uses only OSPF. Doing so causes unwanted traffic on the route.

4 To configure an OSPF area, select Area Configuration. Otherwise, go to Step 11.

The OSPF Areas menu is displayed.

This menu lists the IDs of all areas to which the router belongs. If you have not configured an OSPF area on this router, the only area listed is 0.0.0.0, the *backbone area*.

5 Select an existing area or press Ins to create a new area.

The OSPF Area Configuration menu is displayed.

6 Configure the following area parameters:

- Area ID —Four-byte decimal number that identifies the area. For example, a valid Area ID is 85.8.0.11. However, the Area ID does not have to be an IP address. You can enter any number, but it must be in the format of an IP address. If you enter a hexadecimal number, NIASCFG converts it to decimal.

For the router to belong to an area, the Area ID that identifies that area must be assigned to at least one of the router's interfaces. You assign an Area ID to an interface in Step 8.

- Authentication —Switch that enables or disables authentication for the area.

  If you enable authentication on this router, you must enable authentication on all other routers in the area. Also, all interfaces belonging to that area must have an *authentication key*. You provide the authentication key in Step 8.

- Route Aggregation —Network number of a group of networks that is aggregated into one network number. Press Ins to assign the Network and Mask values of this network number. Because supernetting is not supported, the aggregated network must be the same length as the natural mask of the network class.

- Area Type —Type of OSPF area, which can be Normal or Stub. All routers in the same area must agree on the area type.

  NOTE: The backbone area (0.0.0.0) cannot be a stub area.

- Stub Cost —Cost of the default route advertised to the stub area. This parameter is used only if the Area Type is set to Stub.

7 Press Esc until you are prompted to save your changes, then select Yes.

8 Press Esc until you return to the Internetworking Configuration menu, then select the following parameter path:

Select Bindings > an existing binding > OSPF Bind Options

9 If you are configuring an OSPF area, configure the following area parameters:

- Area ID —ID of the area to which this interface belongs. Press Enter to determine the list of available areas. Use the Up-arrow and Down-arrow keys to select an area, then press Enter to select it.

- Authentication Password —Eight-byte password that authenticates the router's OSPF packets to the area to which this interface belongs. Valid characters are 0 to 9, A to Z, a to z, underscore, and dash.

  This parameter is required only if you enabled the Authentication parameter for the area you select, as described in Step 6 on page 78.

  IMPORTANT: Not all interfaces within the same area are required to have the same authentication key; however, all interfaces *connected to the same network* must have the same authentication key.

10 Press Esc until you are prompted to save your changes, then select Yes.

11 Press Esc to return to the Internetworking Configuration menu.

12 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Load Sharing over Equal-Cost OSPF Routes

IP maintains multiple equal-cost OSPF routes. Load sharing enables a router to divide traffic over equal-cost routes. The router can have several next hops available toward any destination. With this configuration, the router can divide the traffic among the various equal-cost routes to the destination. As a result, load sharing increases the effective bandwidth of an end-to-end path. In addition, it can improve the traffic distribution on an internetwork.

NOTE: Load sharing is performed only on equal-cost routes learned from OSPF.

You enable load sharing within OSPF. IP maintains a maximum of four equal-cost routes to each destination network. The OSPF equal-cost routes are maintained internally and are not displayed in TCPCON.

IMPORTANT: Because OSPF networks tend to be large and complex, we recommend that you do not manually adjust the cost of the interface to create equal-cost routes. It is best to let OSPF automatically determine the equal-cost routes to the destination network.

How to Configure Load Sharing

To configure load sharing on the router, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

2 Select OSPF.

Select Enabled to enable OSPF.

3 Select OSPF Configuration.

4 Select IP Load Sharing , then select Enabled.

This action activates the load-sharing feature.

5 Press Esc until you are prompted to save your changes, then select Yes.

6 Press Esc to return to the Internetworking Configuration menu.

7 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Static Routes for LANs

Static routes are useful if you want to do any of the following on your network:

- Eliminate routing traffic, which increases the bandwidth available for data.
- Limit user access to one portion of the network. For example, if a static route for a network is configured on a router, any packets that are received by the router are forwarded only to the destination network specified by that static route.
- Gain access to isolated areas of the network, which is useful if dealing with legacy network topologies.
- Gain access to a network more than 15 hops away.
- Use a static route as a backup route to dynamic routes.

IMPORTANT: Use this procedure to configure static routes when the next hop router is on the same LAN as the router you are configuring. When the next hop router is across a WAN connection, refer to "Configuring Static Routes for WAN Connections."

How to Configure a LAN Static Route

To configure a static route for a LAN, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

2 Configure the following static route parameters:

- LAN Static Routing —Enables LAN static routing on the router.

  Select this field, then select Enabled.

- LAN Static Routing Table —Entry point to the LAN static route configuration parameters.

  Press Ins and configure the following parameters:

  Route to Network or Host —Destination at the other end of the static route, which can be a single IP host or an IP network (that is, a group of hosts). Or, you can select Default Route. If the router must forward a packet for which it can find no destination in its routing table, it sends the packet to the address specified by the next hop for the default route. This type of blind forwarding keeps a packet on the network until a router can forward it to its final destination.

IP Address of Network/Host —Enter the address of the destination network or host. To select from a list of symbolic network or host names and addresses, press Ins. The list of symbolic network names and addresses comes from the SYS:\ETC\NETWORKS file. The list of symbolic host names and addresses comes from the SYS:\ETC\HOSTS file.

Subnetwork Mask —If the destination is an IP network, the subnet mask of that network.

Next Hop Router on Route —Explicit destination of the next hop.

Enter the IP address of the next-hop router. To select from a list of symbolic hostnames and addresses, press Ins.

Metric for this route —Number of hops to the destination. This metric is directly proportional to the cost of the route. Given two routes to the same destination, the router chooses the lower-cost route.

If you want to use the static route as a *backup route* to a dynamic route, select a value that is higher than the cost associated with the dynamic route. This selection ensures that the dynamic route remains the preferred route under typical conditions.

Do not set this metric value to 16 unless you want to disable the route.

Type of route —Specify whether the static route is *Active* or *Passive*. This parameter specifies whether the next hop router for this route actively advertises the route to this network.

Usually, static routes are not advertised and are categorized as passive routes. When a route is marked as active, TCP/IP expects the next hop router to advertise the route regularly. If a router stops advertising an active route, TCP/IP assumes the route is no longer available and deletes it from the routing table.

If the static route is active and the router discovers a lower-cost dynamic route to the same destination, it uses the lower-cost route instead of the active static route. If the lower-cost route becomes unavailable, the router returns to using the active static route.

If you want to use the static route as a backup route, select Active.

A passive static route is always used, regardless of whether the router discovers a lower-cost route to the same destination.

3 Press Esc twice, then select Yes to save your changes.

4 If you want to disable the routing protocol on this interface to reduce routing traffic, complete the following steps:

4a Select the following:

Select Bindings > an existing binding

4b Select RIP Bind Options.

Select Status > Disabled

4c Press Esc, then select OSPF Bind Options.

Select Status > Disabled

5 If your router has multiple interfaces and you want to disable them, repeat Step 4.

6 Press Esc until you are prompted to save your changes, then select Yes.

7 Press Esc to return to the Internetworking Configuration menu.

8 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Router Discovery

Both IP routers and end nodes can use the ICMP Router Discovery Protocol. Routers use it to advertise themselves as an IP router and to answer queries from end nodes. End nodes use it to locate an IP router on their network. Your system acts as a router when Packet Forwarding is enabled for IP, and acts as an end node when Packet Forwarding is disabled for IP.

NOTE: For an end node to locate an IP router by this method, it must also support the ICMP Router Discovery Protocol.

How to Configure Router Discovery

To configure router discovery on an interface, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Bindings > an existing binding > Expert TCP/IP Bind Options > Router Discovery Options

2 Select the Status field, then select Enabled.

3 Select Destination Address.

This is the method by which the IP router or end node sends router discovery packets. Select one of the following options:

- Broadcast —Sends the packets to all nodes on the network.
- Router Discovery Multicast —Sends the packets to an IP multicast address used specifically for router discovery exchanges. The packets are received only by nodes that understand this multicast address.

4 Press Esc until you are prompted to save your changes, then select Yes.

5 Press Esc to return to the Internetworking Configuration menu.

6 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring ARP and Proxy ARP

IP routers and end nodes use ARP to determine the physical address of a node to which they want to send a packet. ARP is enabled by default. For one node to send a packet to another, the sending node must know the physical address of the destination node. The sending node, knowing only the destination IP address, first checks its *ARP table* for an entry that maps the destination IP address to the destination physical address. If the sending node finds the entry, it inserts the physical address into the packet and sends it. If the sending node does not find the entry in its ARP table, it broadcasts an ARP address request to the network. The destination node replies to the request with its own physical address, which the sending node uses to send the packet and adds to its ARP table for future use.

An IP router uses *Proxy ARP* when devices attached to one of its interfaces do not support IP subnetting and are unaware that they must go through the router to reach devices on other subnets of the same IP network. A router using Proxy ARP replies to ARP requests intended for devices on other subnets, but does so only if the device is reachable through the router. To determine whether the device is reachable, the router examines its own routing table.

Proxy ARP is required on the parent network of a stub subnet. The parent network has an IP address range that includes the IP address range of the stub subnet. The router responds to ARP requests sent on the parent network on behalf of devices on the stub subnet.

When both the parent and stub subnet are bound to IP interfaces, the router can detect the parent/stub subnet and automatically enable Proxy ARP for the appropriate interfaces. Even if Proxy ARP is not required, and not automatically enabled, you can still force it to be enabled with the Force Proxy ARP parameter.

You must enable Force Proxy ARP on each LAN interface on which the router must reply to ARP requests for destinations it can reach.

Force Proxy ARP is disabled on each interface by default.

This topic contains the following sections:

- How to Disable ARP
- How to Enable Proxy ARP

How to Disable ARP

To disable ARP on a LAN network interface, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > an existing binding > Expert TCP/IP Bind Options

2. Select Use of ARP , then select Disabled.
3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.
5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

How to Enable Proxy ARP

To enable Proxy ARP on a network interface, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Bindings > an existing binding > Expert TCP/IP Bind Options

2. Select Force Proxy ARP , then select Enabled.
3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.

5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Directed Broadcast Forwarding

A *directed broadcast* is a broadcast intended for all nodes on a non-local network. For example, the broadcast address 255.255.255.255 reaches all nodes on a network; the directed broadcast address 128.1.255.255 is intended for all nodes whose network address is 128.1.0.0. A router not directly attached to 128.1.0.0 simply forwards the directed broadcast packet to the next hop. A router on network 128.1.0.0—if it has directed broadcast forwarding enabled—accepts and forwards the packet to all nodes whose network address is 128.1.0.0. Routers connecting subnets of 128.1.0.0 also accept and forward the packet to the nodes on their respective subnets.

IMPORTANT: For all nodes on network 128.1.0.0 to receive the directed broadcast, each router attached to network 128.1.0.0 must have Directed Broadcast Forwarding enabled.

How to Enable Directed Broadcast Forwarding

To enable the router to forward directed broadcasts for its network, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > Expert Configuration Options

2. Select Directed Broadcast Forwarding , then select Enabled.
3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.
5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Source Route Packet Forwarding

Using source route packets enables you to determine the route a packet takes to reach its destination. This feature is disabled by default.

How to Enable Forwarding Source Route Packets

To permit forwarding source route packets, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > Expert Configuration Options

2. Select Forward Source Route Packets.

   Select Enabled to permit forwarding IP source route packets.

3. Press Esc until you are prompted to save your changes, then select Yes.
4. Press Esc to return to the Internetworking Configuration menu.
5. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring BOOTP Forwarding

BOOTP is a protocol that enables end nodes to receive their IP addresses from a *BOOTP server* at startup time. If your internetwork has a BOOTP or DHCP server, you can configure your IP router to accept and forward BOOTP or DHCP requests to that server.

How to Configure the Router as a BOOTP Forwarder

To configure the router as a BOOTP forwarder, complete the following steps:

1. Load NIASCFG, then select the following parameter path:

   Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > Expert Configuration Options > BOOTP Forwarding Configuration

2. Select BOOTP Server List , then press Ins.
3. Enter the IP address of the BOOTP or DHCP server at the prompt, or press Ins to display a list of symbolic hostnames and addresses from the SYS:\ETC\HOSTS file.

   The server address appears in the BOOTP Servers screen.

4. Press Esc.
5. Select BOOTP Packet Forwarding , then select Enabled.
6. If you want to record the activity of the BOOTP forwarder, select Log Operation , then select one of the following options:

- Log to BOOTP Screen —Logs BOOTP activity to the BOOTP screen. This is a separate screen that you can select and monitor from the NetWare console. The information logged to this screen is not saved to a file.
- Log to File —Logs BOOTP activity to the SYS:\ETC\BOOTP.LOG file by default.

  To use a different file, type its full path name in the Log File field.

7 If you do not want to record the activity of the BOOTP forwarder, select Do Not Log.

8 Press Esc until you are prompted to save your changes, then select Yes.

9 Press Esc to return to the Internetworking Configuration menu.

10 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring EGP

The Exterior Gateway Protocol (EGP) is an exterior routing protocol that is supported by the TCP/IP software. Exterior routing protocols exchange information between different Autonomous Systems (ASs). The local EGP gets the information about its own AS from the local Interior Gateway Protocols (IGPs). Usually, exterior routing protocols are used only when different companies or commercial services are being connected.

The information EGP receives from the IGP must be explicitly configured. The exterior routing protocol shares only the information specified in the outgoing route filters. This is desirable because you generally want to limit the information exchanged between different ASs.

To enable the EGP, complete the following steps:

1 Load NIASCFG, then select the following parameter path:

Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > Expert Configuration Options

2 Select EGP , then select Enabled.

3 Select EGP Configuration. Configure the following parameters:

- Autonomous System —Enter the autonomous system number. It identifies the autonomous system to which the router belongs. The router establishes an EGP neighbor relationship with routers in other autonomous systems.

- Maximum Neighbors to Acquire —Enter the maximum number of concurrent EGP neighbors with which this router can exchange EGP network reachability information.
- Neighbor List —Select this field to add, modify, or delete EGP neighbors. This router attempts to establish a relationship with the configured EGP neighbors to exchange network reachability information. Press Ins. Configure the following parameters:

  Neighbor's Address —Press Ins to display a list of symbolic hostnames from the SYS:\ETC\HOSTS file. Select a host here or enter the address.

  Neighbor's Autonomous System —Enter the number of the autonomous system to which this EGP neighbor belongs. The router is able to be a neighbor with the EGP peer only when the router and the EGP peer are in different autonomous systems.

4. Press Esc until you are prompted to save your changes, then select Yes.
5. Press Esc to return to the Internetworking Configuration menu.
6. If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Configuring Multiple Logical Interfaces

Novell Internet Access Server 4.1 routing software enables you to bind more than one IP network to a LAN board—or a WAN board with the WAN network mode set to multiaccess. The networks can operate as separate logical interfaces. The ability to configure multiple logical interfaces simplifies the task of managing a growing network in the following ways:

- You can merge network when a there is a router failure.

  For a description, refer to “Merging Two Networks When the Connecting Router Fails.”

- You can move hosts from one IP network to another without losing connectivity.

  For a description, refer to “Reassigning IP Addresses.”

- You can add new nodes to a nearly full subnet.

For a description, refer to “Adding New Nodes to a Full Subnet.”

To attach more than one IP network to a LAN or WAN board, bind IP to the board as many times as necessary; then supply a different IP address for each network.

IMPORTANT: To attach more than one IP network to a WAN board, the WAN network mode must be set to Multi-Access.

Configuring multiple logical interfaces is different from multihoming, which enables you to bind multiple addresses belonging to the same IP network to the same interface or different interfaces. To configure multihoming, refer to “Multihoming.”

Merging Two Networks When the Connecting Router Fails

Suppose a router that connects IP networks 130.81.0.0 and 167.10.0.0 fails. For simplicity, assume that the physical medium is Ethernet. If the router cannot be repaired quickly, you can temporarily fix the problem by completing the following steps:

1 Join the two networks into a single network segment using a barrel connector, a repeater, or other appropriate means.

2 Find an operating Novell Internet Access Server 4.1 system connected to the joined network.

3 Load NIASCFG and select the following path:

Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP

4 Set IP Packet Forwarding to Enabled (Router).

5 Press Esc until you are prompted to save your changes, then select Yes.

6 Press Esc to return to the Internetworking Configuration menu.

7 Select Bindings , then bind IP to the joined network twice.

7a Select an existing binding to an interface connected to the joined network.

7b Set Local IP Address to an available host address on the first network.In this example, enter an available host address on the 130.81.0.0 network.

7c Press Esc , then save your change when prompted.

7d Press Ins to create a new binding and select the same interface connected to the joined network.

7e Set Local IP Address to an available host address on the second network.In this example, enter an available host address on the 167.10.0.0 network.

8 Press Esc until you are prompted to save your changes, then select Yes.

9 Press Esc to return to the Internetworking Configuration menu.

10 If you want these changes to take effect immediately, select Reinitialize System and select Yes to activate your changes.

Reassigning IP Addresses

Suppose you must change network number 89 to 130.57. If the system does not have multiple logical interfaces, you must change all IP addresses on network 89 at the same time or lose connectivity to any host that did not have its address changed. With multiple logical interfaces, you can assign the new IP addresses gradually. Networks 89 and 130.57 can coexist on the same network segment until the transition is complete. The router interfaces, attached to *both* logical networks, forward packets for each network and route packets between the two.

Adding New Nodes to a Full Subnet

Suppose you want to add several new nodes to a subnet that has no more available IP addresses. Assume that the network has enough free connectors available to physically attach the nodes.

First, you assign a new subnet number to the cable so that both subnets share the cable. Then to add new nodes, you bind their IP address to the new logical subnet. The router whose interface is bound to both subnet addresses provides connectivity between the two subnets and to the rest of the internetwork.

Multihoming

Multihoming enables a system to assume multiple IP addresses on the same network. A secondary IP address can be configured on the same interface that has the primary IP address, or a secondary address can be configured on a different interface. When multiple interfaces exist, the secondary address is associated with the interface that is bound to an address that is on the same network. If the secondary address is not valid on any of the networks bound to existing interfaces, the address is rejected and an error message is produced.

When multihoming is used with the proxy server, Virtual Private Network (VPN), or Network Address Translation (NAT), the secondary addresses must be configured manually as described in this section.

To configure secondary IP addresses, complete the following steps:

1. Load NIASCFG and select the following path:

   Select Configure NIAS > Protocols and Routing

2. If you have not done so previously, configure TCP/IP under Protocols and assign one IP address to an interface under Bindings.

3. Press Esc until you are prompted to save your changes, then select Yes.

4. Select the following parameter path:

   Select Manage Configuration > Edit AUTOEXEC.NCF

5. Add a secondary IP address by entering the following command at the end of the file:

   ```
   add secondary IPAddress x.x.x.x
   ```

6. To delete or display secondary IP addresses, press Alt + Esc to display the server console prompt.

   You can delete the secondary IP address by entering the following command:

   ```
   del secondary IPAddress x.x.x.x
   ```

   You can display the secondary IP addresses by entering the following command:

   ```
   display secondary IPAddress
   ```

4 Managing

This topic describes the diagnostic utilities used to mange the Novell® Internet Access Server 4.1 TCP/IP software. These utilities enable you to manage, optimize, and troubleshoot the product and its connections.

Using the TCPCON Utility

TCPCON is an NLM utility that provides access to statistics and information about the status of various components of the TCP/IP protocol suite. TCPCON uses SNMP to access this information from any local or remote system on the network. TCPCON operates over TCP/IP and IPX networks.

To launch TCPCON, enter **`LOAD TCPCON`** at the system console prompt or load NIASCFG and follow this path:

Select View Status for NIAS > Protocols and Routing > TCP/IP Protocol Stack

To monitor a remote system, select SNMP Access Configuration, change the Transport Protocol option to TCP/IP, and set the Host option to the IP address of the remote host you want to monitor. Press Esc to exit and save the options. If details from that remote host are displayed, there is a bidirectional route available.

You can use TCPCON to perform the following tasks:

- Monitor activity in the TCP/IP network segments of your internetwork
- Display configuration information and statistics about the following TCP/IP protocols—IP, ICMP, UDP, TCP, OSPF, and EGP
- Display the IP routes currently known to a TCP/IP node
- Display the network interfaces supported by a TCP/IP node

- Access the trap log maintained by SNMPLOG (for the local system only)
- Access TCP/IP information in any remote protocol stack supporting the TCP/IP Management Information Base (MIB)

NOTE: TCPCON requires SNMP to be loaded on the remote host; otherwise, you receive an error message that the host is unavailable. Another cause of the Host unavailable message might be a routing error. To check for errors in the routing table, accept the default value of 127.0.0.1 in the Host option under SNMP Access Configuration. Select Routing Table to view the routing information table that the routing software has received from routing protocols (RIP and OSPF) or static routes. Compare this to the address topology of the network.

Viewing TCP/IP Configuration Information

To see how TCP/IP is configured, load TCPCON and select the following options:

- SNMP Access Configuration to view and change SNMP access configuration
- Protocol Information to view and change the run-time configuration of TCP/IP protocols
- IP Routing to view, change, and create IP routes
- Statistics to view detailed TCP/IP statistics
- Interfaces to view information about network interfaces
- Display Local Traps to view the local system SNMP trap log

Determining Whether a Remote TCP/IP Node Is Reachable

To determine whether a remote node is reachable, run an Echo test. To run an Echo test, load ping and perform the following steps:

1. Specify the remote node address in the Host Name field.
2. Specify the number of seconds between each transmission in the Seconds to pause between pings field.
3. Specify the packet size to be transmitted in the IP packet size to send in bytes field.

4. Press Esc to begin transmitting.

If you receive an echo reply packet, the remote node is reachable.

Monitoring Error Counters

Error counters are monitored to make sure they are not increasing rapidly, because a rapid increase indicates a problem. For information about troubleshooting these problems, refer to "Troubleshooting.". You can monitor error counters for TCP/IP interfaces in the following ways:

- By using MONITOR to view counters such as Checksum Errors, Send and Receive Packet Errors, and interface-specific errors. To view these counters, load MONITOR and follow this path:

  Select LAN/WAN Information > interface you want to view

- By using PPPCON to view the following PPP counters:

  - Bad Address Fields
  - Bad Control Fields
  - Bad FCS Values
  - Packets Too Long

  To view these counters, load PPPCON and follow this path:

  Select PPP Interfaces > interface you want to view > PPP Error Statistics

- By using TCPCON to view the following TCP/IP counters:

  - IP Errors
  - IP Address Errors
  - Unknown Protocol Errors
  - Local Errors
  - Reassembly Failures Detected
  - Fragmentation Failures Detected

  To view these counters, load TCPCON and follow this path:

  Select Statistics > IP > More IP Statistics

Monitoring TCP/IP Information

Monitoring TCP/IP information can give you a clear view of the status of your TCP/IP network and whether the router is configured properly to run efficiently in the network. This information can also be helpful in troubleshooting and optimizing of the network. This topic contains the following sections:

- Checking the TCP/IP Routing Table
- Monitoring the Configured TCP/IP Protocols

Checking the TCP/IP Routing Table

To check the TCP/IP routing table and information associated with each route, load TCPCON and follow this path:

Select IP Routing Table > Proceed > entry you want to view

The IP Routing Table window shows you all known TCP/IP destination networks. The list has the following information about each item:

- IP address of the destination
- IP address of the next hop router
- Type of the route (direct, remote)
- Primary cost for the route
- Interface used to reach a route

The IP Route Information window expands on this by showing information about the mask used, the routing protocol through which the destination was learned, and the age of the route.

Monitoring the Configured TCP/IP Protocols

You can view, and sometimes change, the configuration of TCP/IP protocols configured for use in your router. You can reach this information by loading TCPCON and selecting Protocol Information. You can configure and view statistics and other information for the following protocols:

- EGP
- ICMP
- IP
- OSPF
- TCP
- UDP

For additional information about each protocol, press F1 to access on-line help.

5 Troubleshooting

This section contains IP troubleshooting information that is divided into three categories:

- Troubleshooting tools
- Troubleshooting checkpoints
- Common problems

If a problem that is general in nature occurs, the procedure described in "Troubleshooting Checkpoints" on page 100 will help you isolate and resolve the problem. If a problem with a specific symptom occurs, refer to "Common Problems."

Troubleshooting Tools

TCPCON is an NLM utility that provides access to statistics and information about the status of various components of the TCP/IP protocol suite. It uses SNMP to access this information from any local or remote system on the network. TCPCON operates over TCP/IP and IPX networks. Use TCPCON to monitor and test a remote system (parameter path: Select SNMP Access Configuration > Transport Protocol > TCP/IP). Set the Host option to the IP address of the remote host you want to test. If details from that remote host are displayed, the remote host is functioning.

You can use TCPCON to perform the following tasks:

- Monitor activity in the TCP/IP network segments of your internetwork
- Display configuration information and statistics about the following TCP/IP protocols: IP, ICMP, UDP, TCP, OSPF, and EGP
- Display the IP routes currently known to a TCP/IP node

- Display the network interfaces supported by a TCP/IP node
- Access the trap log maintained by SNMPLOG (for local systems only)
- Access TCP/IP information in any remote protocol stack supporting the TCP/IP Management Information Base (MIB)

NOTE: TCPCON requires SNMP to be loaded on the remote host; otherwise, you receive an error message that the host is unavailable. Another cause of the error message might be a routing error. To check for errors in the routing table on a Novell Internet Access Server 4.1 machine, accept the default value of 127.0.0.1 in the Host option under SNMP Access Configuration. Select Routing Table to view the routing information table that the routing software has received from routing protocols (RIP and OSPF) or static routes. Compare this to the address topology of the network.

Troubleshooting Checkpoints

To isolate and resolve TCP/IP problems, complete the following steps:

1. Load TCPCON to verify that IP is bound to the desired interfaces with the correct addresses and masks for your internetwork (parameter path: Select Protocol Information > IP > IP Addresses).

   Use NIASCFG to make any required corrections.

2. Load TCPCON to check the routing table for routes to the required network (parameter path: Select IP Routing Table > Proceed and press Enter).

   If routes are missing, verify that the required routing protocols have been enabled and bound to the correct interfaces in NIASCFG. Also verify that the routing protocol in use on an interface is correctly configured on other routers that are accessible through that interface. Finally, use FILTCFG to check for filters that would interfere with the propagation of routes.

3. Load NIASCFG to verify that static routing is configured if other third-party routers that do not use RIP or OSPF are connected on the network (parameter path: Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > Static Routing).

4. Load NIASCFG to verify that the IP Packet Forwarding statistic is set to Enabled (parameter path: Select Configure NIAS > Protocols and Routing > Protocols > TCP/IP > IP Packet Forwarding).

   Use NIASCFG to make any required corrections and then reinitialize the system.

5. Use PING or TPING to test connectivity.

Perform Step 1 through Step 4 on any routers that cannot be reached. Start with the router that is closest to the local node.

6 Use CALLMGR to determine the state of WAN connections (Retry, Out-Queued, or Disconnected).

If the WAN connection is in the wrong state, perform the corrective action described in "WAN Connectivity Problems."

7 Verify that all client software has the Default Router parameter configured to match the IP address of the network board inside the router that is connected to the local segment.

If you are using the Novell LAN WorkPlace® for DOS product, the IP_ROUTER option in the NET.CFG file sets this parameter.

8 Load TCPCON for the following IP statistics:

- Local errors (memory error)

  If a problem exists, see refer to Resolving Memory Problems in the *Overview* documentation.

- IP errors (unexpected protocol errors)

  If a problem exists, check the configuration of other IP nodes on the network. Reduce IP traffic or use a network analyzer to identify the source of invalid packets.

- IP address errors (misdirected packets)

  If a problem exists, check the Address Translation tables on other IP nodes to determine the source of the errors.

- Unknown protocol errors (unsupported IP clients)

  If a problem exists, load the required applications.

- No route found (router failure)

  If a problem exists, check the configuration of the routing protocols.

9 Load TCPCON for the following ICMP statistics:

- Destination unreachable (network failure)

  If a problem exists, use a network analyzer to determine the unreachable destination. Check that the routers on the path to the destination advertise the route.

- Time exceeded (network failure)

If a problem exists, reduce the excessive delays by reducing the size of the internetwork or increasing the speed of WAN links.

- Redirects (router failure)

  If a problem exists, check that all routers on the network are properly configured and advertising routes. Verify that the correct Default Router is configured on the clients.

10. Load NIASCFG to verify that all configuration options are set correctly (parameter path: Select Configure NIAS > Protocols and Routing > View Configuration).

Common Problems

This topic discusses the following common problems and their potential solutions:

- A Dial-In Router is Losing Connectivity
- Router Cannot Ping a Remote Router or the Internet
- WAN Link Comes Up, but Ping Receives No Reply
- LAN Connectivity Problems
- WAN Connectivity Problems
- Routing Table Maintenance Problems

A Dial-In Router is Losing Connectivity

If multiple remote routers are dialing into a PPP group interface with WAN Network Mode set to Multiaccess, each router must have a unique IP address. If a duplicate address exists, one of the dial-in routers will lose connectivity, but no error message is issued.

Router Cannot Ping a Remote Router or the Internet

Load TCPCON and verify that there is a destination that is specified as the default route (parameter path: Select IP Routing Table). If there is no default route you must configure it. Load NIASCFG to permanently configure the default route.

To create a default route, complete the following steps:

1. Load NIASCFG and select Configure NIAS > Protocols and Routing > Bindings > TCP/IP.

2 Select the Interface Group for your WAN card.

3 Select WAN Call Destination.

The Configured WAN Call Destinations menu appears.

4 Press Ins on the WAN Call Destination option and select the WAN card defined earlier.

5 Press Ins on the Static Routing Table option.

The Static Routing entry sets up the default route that points to the Internet Service Provider (ISP).

6 Select Route to Network or Host and select Default Route.

7 Press Esc to save your changes and exit the menus.

WAN Link Comes Up, but Ping Receives No Reply

Load CALLMGR.NLM and initiate the call. If the link comes up and PPP and IP are negotiated, CALLMGR will show the calls as being Out-Connected. However, unless the routing tables on both sides are properly configured, IP packets will not be routed correctly.

Load PING.NLM on the server and ping the ISP's IP address. You should receive a reply. If you do not receive a reply, ensure that the default route appears in the routing table by loading TCPCON.NLM at the server console. Select IP Routing Table and press Enter. You should see an entry with the Default selected for Destination, Unspecified selected for Next Hop, and Remote selected for Type.

NOTE: For a permanent link, the line must be up for the default route to be visible through TCPCON. If the default entry does not show up, ensure that it is properly defined in the Binding for the WAN card.

For an unnumbered link, make sure that your network information card (NIC) is bound to the IP address assigned to you by the ISP.

If you can ping the ISP, but not beyond, check your IP addresses. For an unnumbered link, you must bind the local IP address the ISP gave you to the NIC in your system. Also, check with your ISP to make sure that it has a static route that points back to your network. The ISP must have a routing table entry that specifies that your destination network/host is reachable through the local IP address you have been assigned. This entry ensures that packets are properly routed back to your network.

LAN Connectivity Problems

- The router does not forward IP packets.

  Verify that the IP Packet Forwarding statistic is set to Router (Router indicates that it is enabled) in TCPCON (parameter path: Protocol Information > IP). If routing is not enabled, enable IP Packet Forwarding under Protocols in NIASCFG, then issue the REINITIALIZE SYSTEM command.

- A TCP/IP host cannot reach the router on the local network.

  - Verify that the router and host use the same frame type.

  - Verify that the network portion of the IP address and the subnet mask are the same on the router and the host.

  - Use PING from the router to verify connectivity to the TCP/IP host and verify that the IP Address Translation table has an entry for the host.

    If there is no entry, use MONITOR to check the status of the LAN driver.

  - Use PING from the router to verify connectivity to the TCP/IP host and check for Echo Requests in TCPCON (parameter path: Select Statistics > ICMP).

    If the value of the Echo Requests statistic is not incrementing, check the IP statistics for errors and perform Step 8 in "Troubleshooting Checkpoints."

  - Use PING from the router to verify connectivity to the TCP/IP host and check for Echo Replies in TCPCON (parameter path: Select Statistics > ICMP).

    If the value of the Echo Replies statistic is not incrementing, verify that IP is bound to the host's interface with the correct address and mask. Also, verify that the interface driver is loaded with the correct frame type. If required, check the IP statistics for errors and perform Step 8 in "Troubleshooting Checkpoints."

- A TCP/IP host cannot reach a remote host.

  - Verify that the local TCP/IP host has the local router listed as the default router.

  - Using PING, verify that the local TCP/IP host can reach each router on the path to the remote host.

    - Verify that each router has a routing protocol enabled and that it has not been disabled on the interface.
    - Starting at the local router, verify that each router has a route to the remote host’s network.
    - Verify that there are no filters capable of blocking IP traffic configured on any routers along the path.
    - Verify that the remote host has a route to the local host’s network.
    - Using PING, verify that the remote host can reach each router on the path to the local TCP/IP host.
    - Starting at the router closest to the remote host, verify that each router has a route to the local TCP/IP host’s network.
- The router cannot initiate IP traffic to a remote router through a LAN interface.
    - Verify that IP is bound to the right interface with the correct address and mask.
    - Check whether the interface driver is loaded with the correct frame type.
    - Check whether a route exists to the network on which the destination router resides. This can be done through the IP Routing Table window of TCPCON. If the destination router is accessible, also verify that it has a route to the source router’s network.

WAN Connectivity Problems

- A permanent WAN call destination is defined in the WAN Call Directory window of NIASCFG, but the call is not initiated through the interface when IP is bound to that interface.

  The permanent WAN call destination defined must be specified when configuring the binding (parameter path: Select Configure NIAS > Protocols and Routing > Bindings > a specific binding > Permanent WAN Call Destination).
- An on-demand WAN call destination to a remote router that is on the same LAN is defined in the WAN Call Directory window of NIASCFG. When an attempt is made to initiate IP traffic to this remote router, the WAN connection does not come up.

Although the remote router is on the same network as the local router, a static route must be configured to the remote router address with the next hop as the on-demand WAN call destination.

- An X.25 WAN connection does not come up between two routers, and IP traffic cannot be initiated, even when one of the routers has a permanent WAN call destination configured to access the other router.

  To bring up a WAN connection through IP on an X.25 link in multiaccess mode, both routers must have WAN Call Destination/Remote IP Address Map configured for the link to the other router (parameter path: Select Bindings).

- The router cannot reach a remote WAN call destination.

  - Check the Boards, Network Interfaces, and WAN Call Directory configurations in NIASCFG.
  - Verify that the Remote System ID option and other authentication options are properly configured.
  - Verify that IP is bound to the interface.
  - For on-demand WAN connections, numbered point-to-point and unnumbered point-to-point links must have a static route with at least one IP address. Multiaccess links must have a static route or WAN mapping configured.
  - For permanent WAN connections, a permanent WAN call destination must be configured. X.25 multiaccess links must have WAN mapping configured in addition to a permanent WAN call destination.

- A remote WAN call destination cannot reach the router.

  - Check the Boards, Network Interfaces, and WAN Call Directory configurations in NIASCFG.
  - If a static route has been configured, verify that Remote System ID is properly configured.
  - For on-demand WAN connections (incoming), numbered point-to-point and unnumbered point-to-point links must have a static route configured, and the routing protocol must be enabled. When X.25 is being used, multiaccess links must have WAN mapping configured or the remote node will not support Internet Protocol Control Protocol (IPCP) or Address Resolution Protocol (ARP).

- An on-demand WAN connection is connected continually.

    - The routing protocol or router discovery is enabled on the interface.

    - TCPCON or PING to the remote host was left running.

    - Applications using IP or other protocols tunneled through IP might be generating IP traffic.

- An on-demand WAN connection connects and disconnects continually.

  Traffic is being generated by TCPCON, PING, router discovery, a routing protocol, or an application using IP or other protocols tunneled through IP, and the Idle Connection Timeout option is set to a value that is too low (10 minutes or less).

- Router software can ping the Internet, but clients cannot.

  Ensure that the client has a default route that points to the IP address of the NIC in the routing software. Also, verify that the IP address on the workstations are correct and belong to the same network to which the NIC is bound. On the server, verify that IP Packet Forwarding is set to Enabled.

  If you are using a numbered WAN link, the IP address bound to the NIC must be assigned to you by the ISP, and the local IP address bound to the NIC must be on a network different from the one that the remote IP address of the ISP is on.

  If you are using an unnumbered link, ensure that the IP address bound to the NIC is the correct address assigned by the ISP.

Routing Table Maintenance Problems

- Route updates are not sent to other routers on the frame relay cloud, even though IP and the routing protocol are enabled over the multiaccess WAN interface.

  Configure the Neighbor List option by specifying the IP address of routers to which route updates must be sent (parameter path: Select Configure NIAS > Protocols and Routing > Bindings > a specific binding > RIP (or OSPF) Bind Options).

- Routes are not exchanged on a LAN or WAN.

    - Use NIASCFG to verify that the IP Packet Forwarding option is enabled.

    - Use NIASCFG to verify that a routing protocol has been enabled.

- Use NIASCFG to verify that the routing protocol has not been disabled on an interface.
- Use TCPCON to examine the routing table and determine which routes are missing.
- Check TCPCON for IP errors.
- Check TCPCON for ICMP errors.
- If you are using RIP, then in NIASCFG under Bindings, verify that the RIP Mode option is not set to Send Only or Receive Only.
- If the RIP Version option is set to RIPII, verify that the other routers also support RIP II.
- Verify that no route filters are configured that would block route information packets for that interface.
- If you are using OSPF, verify that the following conditions have been met:

  Routers in the area have the same Authentication Type configured.

  All routers on the same network have the same Authentication Password configured for the interface to the network.

  All routers on the same network have the same Hello Intervals configured for the interface to the network.

  The state of each neighbor is either two-way or full in TCPCON (parameter path: Select Protocol Information > OSPF > Neighbors). If it is not, one of the two conditions described next will occur. Refer to the next two paragraphs for an explanation of the corrective actions required.

  In TCPCON, there is a router link state advertisement for each router in your area (parameter path: Select Protocol Information > OSPF > Link State Advertisements). If these advertisements are not present, verify that the missing router is active and the correct area ID is configured for the network interface.

  In TCPCON, the number of link state advertisements, Area Boundary Routers, and Autonomous System Boundary Routers are the same for each router in your area (parameter path: Select Protocol Information > OSPF > Areas). Verify that the problem routers are active. Bring down any router whose routing database is not synchronized with the databases of its routing neighbors. If the problem persists, reduce the size of your network or add more memory to the router.

- Routes are not exchanged on a LAN.

  Verify that the broadcast address is correct.

- Routes are not exchanged across a WAN connection.

  - Verify that the routing protocol is enabled on the WAN connection.

  - If you are using OSPF and the WAN Mode option is set to Multi-Access, verify that the correct routing neighbors are defined and that the address mappings are defined.

- RIP routes are not accessible to hosts on OSPF networks.

  - Check the status of the Autonomous System Boundary Router statistic in TCPCON (parameter path: Select Protocol Information > OSPF).

  - Verify that no filters are configured that would block access to the network.

A Novell Trademarks

Access Manager is a registered trademark of Novell, Inc. in the United States and other countries.

Advanced NetWare is a trademark of Novell, Inc.

AlarmPro is a registered trademark of Novell, Inc. in the United States and other countries.

AppNotes is a registered service mark of Novell, Inc. in the United States and other countries.

AppNotes is a registered service mark of Novell, Inc. in the United States and other countries.

AppTester is a registered service mark of Novell, Inc. in the United States and other countries.

BrainShare is a registered service mark of Novell, Inc. in the United States and other countries.

C-Worthy is a trademark of Novell, Inc.

C3PO is a trademark of Novell, Inc.

CBASIC is a registered trademark of Novell, Inc. in the United States and other countries.

Certified NetWare Administrator in Japanese and CNA-J are service marks of Novell, Inc.

Certified NetWare Engineer in Japanese and CNE-J are service marks of Novell, Inc.

Certified NetWare Instructor in Japanese and CNI-J are service marks of Novell, Inc.

Certified Novell Administrator and CNA are service marks of Novell, Inc.

Certified Novell Engineer is a trademark and CNE is a registered service mark of Novell, Inc. in the United States and other countries.

Certified Novell Salesperson is a trademark of Novell, Inc.

Client 32 is a trademark of Novell, Inc.

ConnectView is a registered trademark of Novell, Inc. in the United States and other countries.

Connectware is a registered trademark of Novell, Inc. in the United States and other countries.

Corsair is a registered trademark of Novell, Inc. in the United States and other countries.

CP/Net is a registered trademark of Novell, Inc. in the United States and other countries.

Custom 3rd-Party Object and C3PO are trademarks of Novell, Inc.

DeveloperNet is a registered trademark of Novell, Inc. in the United States and other countries.

Documenter's Workbench is a registered trademark of Novell, Inc. in the United States and other countries.

ElectroText is a trademark of Novell, Inc.

Enterprise Certified Novell Engineer and ECNE are service marks of Novell, Inc.

Envoy is a registered trademark of Novell, Inc. in the United States and other countries.

EtherPort is a registered trademark of Novell, Inc. in the United States and other countries.

EXOS is a trademark of Novell, Inc.

Global MHS is a trademark of Novell, Inc.

Global Network Operations Center and GNOC are service marks of Novell, Inc.

Graphics Environment Manager and GEM are registered trademarks of Novell, Inc. in the United States and other countries.

GroupWise is a registered trademark of Novell, Inc. in the United States and other countries.

GroupWise XTD is a trademark of Novell, Inc.

Hardware Specific Module is a trademark of Novell, Inc.

Hot Fix is a trademark of Novell, Inc.

InForms is a trademark of Novell, Inc.

Instructional Workbench is a registered trademark of Novell, Inc. in the United States and other countries.

Internetwork Packet Exchange and IPX are trademarks of Novell, Inc.

IPX/SPX is a trademark of Novell, Inc.

IPXODI is a trademark of Novell, Inc.

IPXWAN is a trademark of Novell, Inc.

LAN WorkGroup is a trademark of Novell, Inc.

LAN WorkPlace is a registered trademark of Novell, Inc. in the United States and other countries.

LAN WorkShop is a trademark of Novell, Inc.

LANalyzer is a registered trademark of Novell, Inc. in the United States and other countries.

LANalyzer Agent is a trademark of Novell, Inc.

Link Support Layer and LSL are trademarks of Novell, Inc.

MacIPX is a registered trademark of Novell, Inc. in the United States and other countries.

ManageWise is a registered trademark of Novell, Inc. in the United States and other countries.

Media Support Module and MSM are trademarks of Novell, Inc.

Mirrored Server Link and MSL are trademarks of Novell, Inc.

Mobile IPX is a trademark of Novell, Inc.

Multiple Link Interface and MLI are trademarks of Novell, Inc.

Multiple Link Interface Driver and MLID are trademarks of Novell, Inc.

My World is a registered trademark of Novell, Inc. in the United States and other countries.

N-Design is a registered trademark of Novell, Inc. in the United States and other countries.

Natural Language Interface for Help is a trademark of Novell, Inc.

NDS Manager is a trademark of Novell, Inc.

NE/2 is a trademark of Novell, Inc.

NE/2-32 is a trademark of Novell, Inc.

NE/2T is a trademark of Novell, Inc.

NE1000 is a trademark of Novell, Inc.

NE1500T is a trademark of Novell, Inc.

NE2000 is a trademark of Novell, Inc.

NE2000T is a trademark of Novell, Inc.

NE2100 is a trademark of Novell, Inc.

NE3200 is a trademark of Novell, Inc.

NE32HUB is a trademark of Novell, Inc.

NEST Autoroute is a trademark of Novell, Inc.

NetExplorer is a trademark of Novell, Inc.

NetNotes is a registered trademark of Novell, Inc. in the United States and other countries.

NetSync is a trademark of Novell, Inc.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare 3270 CUT Workstation is a trademark of Novell, Inc.

NetWare 3270 LAN Workstation is a trademark of Novell, Inc.

NetWare 386 is a trademark of Novell, Inc.

NetWare Access Server is a trademark of Novell, Inc.

NetWare Access Services is a trademark of Novell, Inc.

NetWare Application Manager is a trademark of Novell, Inc.

NetWare Application Notes is a trademark of Novell, Inc.

NetWare Asynchronous Communication Services and NACS are trademarks of Novell, Inc.

NetWare Asynchronous Services Interface and NASI are trademarks of Novell, Inc.

NetWare Aware is a trademark of Novell, Inc.

NetWare Basic MHS is a trademark of Novell, Inc.

NetWare BranchLink Router is a trademark of Novell, Inc.

NetWare Care is a trademark of Novell, Inc.

NetWare Communication Services Manager is a trademark of Novell, Inc.

NetWare Connect is a registered trademark of Novell, Inc. in the United States.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

NetWare Distributed Management Services is a trademark of Novell, Inc.

NetWare Document Management Services is a trademark of Novell, Inc.

NetWare DOS Requester and NDR are trademarks of Novell, Inc.

NetWare Enterprise Router is a trademark of Novell, Inc.

NetWare Express is a registered service mark of Novell, Inc. in the United States and other countries.

NetWare Global Messaging and NGM are trademarks of Novell, Inc.

NetWare Global MHS is a trademark of Novell, Inc.

NetWare HostPrint is a registered trademark of Novell, Inc. in the United States.

NetWare IPX Router is a trademark of Novell, Inc.

NetWare LANalyzer Agent is a trademark of Novell, Inc.

NetWare Link Services Protocol and NLSP are trademarks of Novell, Inc.

NetWare Link/ATM is a trademark of Novell, Inc.

NetWare Link/Frame Relay is a trademark of Novell, Inc.

NetWare Link/PPP is a trademark of Novell, Inc.

NetWare Link/X.25 is a trademark of Novell, Inc.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

NetWare LU6.2 is trademark of Novell, Inc.

NetWare Management Agent is a trademark of Novell, Inc.

NetWare Management System and NMS are trademarks of Novell, Inc.

NetWare Message Handling Service and NetWare MHS are trademarks of Novell, Inc.

NetWare MHS Mailslots is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare Mirrored Server Link and NMSL are trademarks of Novell, Inc.

NetWare Mobile is a trademark of Novell, Inc.

NetWare Mobile IPX is a trademark of Novell, Inc.

NetWare MultiProtocol Router and NetWare MPR are trademarks of Novell, Inc.

NetWare MultiProtocol Router Plus is a trademark of Novell, Inc.

NetWare Name Service is trademark of Novell, Inc.

NetWare Navigator is a trademark of Novell, Inc.

NetWare Peripheral Architecture is a trademark of Novell, Inc.

NetWare Print Server is a trademark of Novell, Inc.

NetWare Ready is a trademark of Novell, Inc.

NetWare Requester is a trademark of Novell, Inc.

NetWare Runtime is a trademark of Novell, Inc.

NetWare RX-Net is a trademark of Novell, Inc.

NetWare SFT is a trademark of Novell, Inc.

NetWare SFT III is a trademark of Novell, Inc.

NetWare SNA Gateway is a trademark of Novell, Inc.

NetWare SNA Links is a trademark of Novell, Inc.

NetWare SQL is a trademark of Novell, Inc.

NetWare Storage Management Services and NetWare SMS are trademarks of Novell, Inc.

NetWare Telephony Services is a trademark of Novell, Inc.

NetWare Tools is a trademark of Novell, Inc.

NetWare UAM is a trademark of Novell, Inc.

NetWare WAN Links is a trademark of Novell, Inc.

NetWare/IP is a trademark of Novell, Inc.

NetWire is a registered service mark of Novell, Inc. in the United States and other countries.

Network Navigator is a registered trademark of Novell, Inc. in the United States.

Network Navigator - AutoPilot is a registered trademark of Novell, Inc. in the United States and other countries.

Network Navigator - Dispatcher is a registered trademark of Novell, Inc. in the United States and other countries.

Network Support Encyclopedia and NSE are trademarks of Novell, Inc.

Network Support Encyclopedia Professional Volume and NSEPro are trademarks of Novell, Inc.

NetWorld is a registered service mark of Novell, Inc. in the United States and other countries.

Novell is a service mark and a registered trademark of Novell, Inc. in the United States and other countries.

Novell Alliance Partners Program is a collective mark of Novell, Inc.

Novell Application Launcher is a trademark of Novell, Inc.

Novell Authorized CNE is a trademark and service mark of Novell, Inc.

Novell Authorized Education Center and NAEC are service marks of Novell, Inc.

Novell Authorized Partner is a service mark of Novell, Inc.

Novell Authorized Reseller is a service mark of Novell, Inc.

Novell Authorized Service Center and NASC are service marks of Novell, Inc.

Novell BorderManager is a trademark of Novell, Inc.

Novell BorderManager FastCache is a trademark of Novell, Inc.

Novell Client is a trademark of Novell, Inc.

Novell Corporate Symbol is a trademark of Novell, Inc.

Novell Customer Connections is a registered trademark of Novell, Inc. in the United States.

Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.

Novell Distributed Print Services is a trademark and NDPS is a registered trademark of Novell, Inc. in the United States and other countries.

Novell ElectroText is a trademark of Novell, Inc.

Novell Embedded Systems Technology is a registered trademark and NEST is a trademark of Novell, Inc. in the United States and other countries.

Novell Gold Authorized Reseller is a service mark of Novell, Inc.

Novell Gold Partner is a service mark of Novell, Inc.

Novell Labs is a trademark of Novell, Inc.

Novell N-Design is a registered trademark of Novell, Inc. in the United States and other countries.

Novell NE/2 is a trademark of Novell, Inc.

Novell NE/2-32 is a trademark of Novell, Inc.

Novell NE3200 is a trademark of Novell, Inc.

Novell Network Registry is a service mark of Novell, Inc.

Novell Platinum Partner is a service mark of Novell, Inc.

Novell Press is a trademark of Novell, Inc.

Novell Press Logo (teeth logo) is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Replication Services is a trademark of Novell, Inc.

Novell Research Reports is a trademark of Novell, Inc.

Novell RX-Net/2 is a trademark of Novell, Inc.

Novell Service Partner is a trademark of Novell, Inc.

Novell Storage Services is a trademark of Novell, Inc.

Novell Support Connection is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Technical Services and NTS are service marks of Novell, Inc.

Novell Technology Institute and NTI are registered service marks of Novell, Inc. in the United States and other countries.

Novell Virtual Terminal and NVT are trademarks of Novell, Inc.

Novell Web Server is a trademark of Novell, Inc.

Novell World Wide is a trademark of Novell, Inc.

NSE Online is a service mark of Novell, Inc.

NTR2000 is a trademark of Novell, Inc.

Nutcracker is a registered trademark of Novell, Inc. in the United States and other countries.

OnLAN/LAP is a registered trademark of Novell, Inc. in the United States and other countries.

OnLAN/PC is a registered trademark of Novell, Inc. in the United States and other countries.

Open Data-Link Interface and ODI are trademarks of Novell, Inc.

Open Look is a registered trademark of Novell, Inc. in the United States and other countries.

Open Networking Platform is a registered trademark of Novell, Inc. in the United States and other countries.

Open Socket is a registered trademark of Novell, Inc. in the United States.

Packet Burst is a trademark of Novell, Inc.

PartnerNet is a registered service mark of Novell, Inc. in the United States and other countries.

PC Navigator is a trademark of Novell, Inc.

PCOX is a registered trademark of Novell, Inc. in the United States and other countries.

Perform3 is a trademark of Novell, Inc.

Personal NetWare is a trademark of Novell, Inc.

Pervasive Computing from Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Portable NetWare is a trademark of Novell, Inc.

Presentation Master is a registered trademark of Novell, Inc. in the United States and other countries.

Print Managing Agent is a trademark of Novell, Inc.

Printer Agent is a trademark of Novell, Inc.

QuickFinder is a trademark of Novell, Inc.

Red Box is a trademark of Novell, Inc.

Reference Software is a registered trademark of Novell, Inc. in the United States and other countries.

Remote Console is a trademark of Novell, Inc.

Remote MHS is a trademark of Novell, Inc.

RX-Net is a trademark of Novell, Inc.

RX-Net/2 is a trademark of Novell, Inc.

ScanXpress is a registered trademark of Novell, Inc. in the United States and other countries.

Script Director is a registered trademark of Novell, Inc. in the United States and other countries.

Sequenced Packet Exchange and SPX are trademarks of Novell, Inc.

Service Response System is a trademark of Novell, Inc.

Serving FTP is a trademark of Novell, Inc.

SFT is a trademark of Novell, Inc.

SFT III is a trademark of Novell, Inc.

SoftSolutions is a registered trademark of SoftSolutions Technology Corporation, a wholly owned subsidiary of Novell, Inc.

Software Transformation, Inc. is a registered trademark of Software Transformation, Inc., a wholly owned subsidiary of Novell, Inc.

SPX/IPX is a trademark of Novell, Inc.

StarLink is a registered trademark of Novell, Inc. in the United States and other countries.

Storage Management Services and SMS are trademarks of Novell, Inc.

Technical Support Alliance and TSA are collective marks of Novell, Inc.

The Fastest Way to Find the Right Word is a registered trademark of Novell, Inc. in the United States and other countries.

The Novell Network Symbol is a trademark of Novell, Inc.

Topology Specific Module and TSM are trademarks of Novell, Inc.

Transaction Tracking System and TTS are trademarks of Novell, Inc.

Universal Component System is a registered trademark of Novell, Inc. in the United States and other countries.

Virtual Loadable Module and VLM are trademarks of Novell, Inc.

Writer’s Workbench is a registered trademark of Novell, Inc. in the United States and other countries.

Yes, It Runs with NetWare (logo) is a trademark of Novell, Inc.

Yes, NetWare Tested and Approved (logo) is a trademark of Novell, Inc.

ZENworks is a trademark of Novell, Inc.