



# Technology Backgrounder

## IP Environment

1.	Introduction to IP .....	1
2.	IP Networks, IP Hosts and IP Ports .....	1
3.	IP Packet Structure .....	2
4.	IP Address Structure.....	2
	Network Portion .....	2
	Host Portion .....	3
	Global vs. Private IP Addresses .....	3
	Subnetting .....	4
	Subnet Masks .....	5
5.	Dynamic Allocation of IP Addresses .....	6
	DHCP Services .....	6
	NAT/NAPT Services .....	6
	PAT Services .....	7
6.	IP Routing Principles.....	7
	IP Communication between Hosts within the Same Network .....	7
	IP Communication between Hosts Located on Different Networks .....	8
	Tools for Checking IP Connectivity .....	9

## 1. Introduction to IP

The information presented in this section refers to Version 4 of the IP protocol (IP4), currently the most widely used protocol version.

IP means “Internet Protocol”. The term **IP protocol** is often used to indicate a standardized set of rules and procedures that enable data exchange through a packet-switched network.

Accordingly, the term **Internet** indicates the set of networks that use the IP protocol and are interconnected in a way that, at least in principle, permits any entity on one network to communicate with any entity on another network.

---

**Note** *The term “suite of IP protocols” is also often used, in recognition of the fact that the operation of the Internet is actually defined by many related protocols.*

---

---

## 2. IP Networks, IP Hosts and IP Ports

Any entity that can communicate using the IP protocol is called an **IP host**.

The connection point between an IP host and an IP network is called **IP port**.

An **IP network** forms when a number of IP ports can communicate directly (peer to peer) using the IP protocol, without any intermediaries.

An IP host can have any number of IP ports. Moreover, the ports may be located on different IP networks.

To enable IP communication between two IP hosts, it is necessary to find a route between their IP ports. For this purpose, each IP port is assigned an IP address.

An IP address is a number selected in accordance with the IP protocol. The only purpose of an IP address is to permit unambiguous identification of an IP port. Therefore, each IP port must be assigned a distinct and unique IP address.

The IP protocol does not require the IP port to be related in an unambiguous way to a physical (communication) port. This has two main implications:

- Since the IP port is actually a connection to an IP network, any number of IP ports can share a given physical port.
- An IP port may be reached through several physical ports.

---

**Note** *By convention, the scope of IP addresses has been extended in two ways:*

- *To permit identification of IP networks*
- *To permit simultaneous addressing of all the ports connected to an IP network (this operation is called broadcasting).*

---

### 3. IP Packet Structure

The information exchanged through IP networks is organized in packets. The structure of an IP packet, as specified by IP protocol Version 4, is as follows (the numbers are byte numbers):

0	4	8	12	16	20	24	28	31
IP Version (4)	IP Header Length	IP Type of Service (IP TOS)		Total IP Packet Length (total number of octets in header + payload)				
Fragment Identification (16 bits)				Flags (3 bits)	Fragment Offset (13 bits)			
(These fields are used for IP packet fragmentation)								
Time to Live (Range: 0 to 255; when 0, packet is discarded)		Number of Upper-Layer Protocol Carried in Payload (IGMP = 2) (UDP = 17)		IP Header Checksum				
Source IP Address								
Destination IP Address								
Options (when used)						Padding (as required)		
Payload (maximum bytes: 65535 – “header length”)								
.								
.								
.								

### 4. IP Address Structure

An IP address is a 32-bit number, represented as four 8-bit bytes. Each byte represents a decimal number in the range of 0 through 255.

The address is written in decimal format, with the bytes separated by decimal points, e.g., 164.90.70.47. This format is called ***dotted quad notation***.

An IP address is logically divided into two main portions:

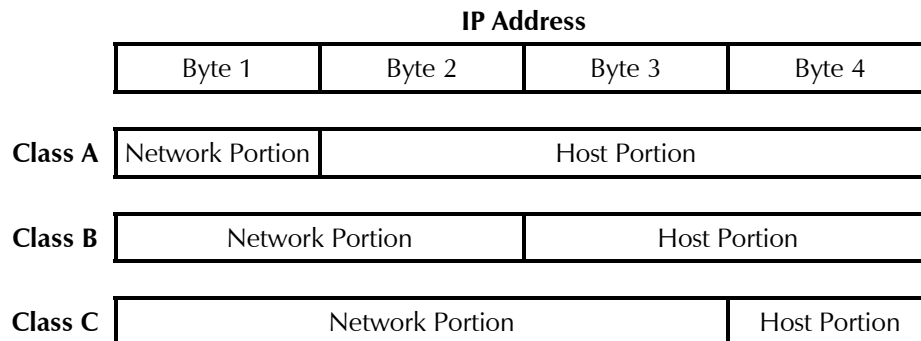
- Network portion
- Host portion.

#### Network Portion

In general, the network portion is assigned by the Internet Assigned Numbers Authority (IANA), and its main purpose is to identify a specific IP network. For exceptions, see the [Global vs. Private IP Addresses](#) section below.

There are five IP address classes: A, B, C, D, and E. However, only the A, B and C classes are used for IP addressing. Consult your network manager with respect to the class of IP addresses used on your network.

The network portion of an IP address can be one, two, or three bytes long, in accordance with the IP address class. This arrangement is illustrated below:



The class of each IP address can be determined from its leftmost byte, in accordance with the following chart:

Address Class	First Byte	Address Range
Class A	0 through 127	0.H.H.H through 127.H.H.H
Class B	128 through 191	128.N.H.H through 191.N.H.H
Class C	192 through 223	192.N.N.H through 223.N.N.H

where:

- N - indicates bytes that are part of the network portion
- H - indicates bytes that are part of the host portion.

## Host Portion

In general, the host portion is used to identify an individual host connected to an IP network. For exceptions, see [Subnetting](#) section below.

After obtaining an IP network address, the using organization is free to assign host identifiers in accordance with its specific needs.

**Note** *The following host identifiers have special meanings, and must not be assigned to an actual host:*

- The “all-zeros” host identifier is interpreted as a network identifier.
- The “all-ones” host identifier is interpreted as a broadcast address. Therefore, a message with an “all-ones” host identifier is accepted by all the hosts in the network.

## Global vs. Private IP Addresses

Given the current number of users already having access to the Internet, and the rapid increase in this number, the 32-bit IP space address available in Version 4 of the IP protocol is rather limited.

On the other hand, an IP address must permit unambiguous identification of any host in the Internet.

That is the reason the allocation of IP addresses to networks is globally controlled by a universally-accepted IP registry organization (IANA).

Although any address used on the Internet must be unique, there are many IP networks, called private networks, which are not connected to the Internet. A private network is also created when the access of hosts to the Internet is controlled by protocols and procedures that do not permit an outsider to find and use directly the actual address of the hosts connected to that network. A typical example of a private network is the internal IP network of an enterprise (such networks are often called *intranets*).

In recognition of this fact, IANA permits using two types of addresses:

- **Global addresses**, i.e., addresses that are unique in the whole Internet
- **Private addresses**, i.e., addresses allocated for internal use only and therefore cannot be used on the Internet.

Although no restrictions need to be imposed on private addresses except for conforming to the structure specified in the *Network Portion* section above, the following address spaces have been specifically put aside by IANA for use as private addresses:

- The Class A addresses in the range of 10.0.0.0 to 10.255.255.255 (this group of addresses is formally referred to as **10/8**). This address space is actually one Class A network number.
- The Class B addresses in the range of 172.16.0.0 to 172.31.255.255 (this group of addresses is formally referred to as **172.16/12**). This address space defines 16 contiguous Class B network numbers.
- The Class C addresses in the range of 192.168.0.0 to 192.168.255.255 (this group of addresses is formally referred to as **192.168/16**). This address space defines 256 contiguous Class C network numbers.

---

**Note** *Three techniques are available to improve the utilization of scarce global address space:*

- *Dynamic Host Configuration Protocol (DHCP).*
- *Network Address Translation (NAT) and Network Address/Port Translation (NAPT).*
- *Port Address Translation (PAT).*

*See the [Dynamic Allocation of IP Addresses](#) section below.*

---

## Subnetting

Given the scarcity of IP network addresses, for organizations operating several relatively small, physically separated, IP networks, e.g., several departmental networks, it is advantageous to enable several physical networks to share a common IP network address. **Small** in this context means that the number of IP ports connected to each of these networks is small relative to the host address space for the corresponding IP address class.

The approach taken to enable the sharing of an IP network address by two or more networks is called **subnetting**, which means **use of subnets**. The subnetting is relevant only within the using organization, and therefore can be freely selected to meet its specific needs.

To enable subnetting, the meaning of the bits in the host portion of the IP address is further sub-divided into two portions:

- **Subnet number.** For example, subnet numbers can be used to identify departmental subnets. The subnet number follows the network identifier.
- **Host number** - the last bits of the IP address.

This subdivision is illustrated below:



For example, when the subnet includes 16 IP hosts, only the last four bits need to be reserved for the host number. For an organization which obtained one global Class C network address, this means that four bits are available to identify subnets. Therefore, this organization can implement 16 IP subnets, each comprising up to 16 hosts (except for two subnets that are limited to 15 hosts).

## Subnet Masks

Subnet masks are used to indicate the division of the IP address bits between the net and subnet portion and the host portion.

The mask is a 32-bit word that includes “**ones**” in the positions used for net and subnet identification, followed by “**zeros**” up to the end of the IP address.

For example, the default subnet mask for any Class C address (i.e., all the eight bits in the host address space are used for hosts in the same net) is 255.255.255.000.

However, if the same address is used in a subnet comprising up to 16 hosts and for which the host numbers range is 00 to 15, the subnet mask changes as follows:

<b>IP Address (Dotted-Quad)</b>	192	70	55	13
<b>IP Address (Binary)</b>	1011 1111	0100 0110	0011 0111	0000 0111
<b>Subnet Mask (Binary)</b>	1111 1111	1111 1111	1111 1111	1111 0000
<b>Subnet Mask (Dotted-Quad)</b>	255	255	255	240

In most applications, the binary subnet mask is built as a contiguous string of “**ones**”, followed by a number of “**zeros**” (the number of “**zeros**” is selected as needed, to complete the number of subnet mask bits to 32). Therefore, when this conventional approach is used, the subnet mask can also be specified simply by stating the number of “**ones**” in the mask. For example, the subnet mask shown above is specified by stating that it comprises 28 bits.

## 5. Dynamic Allocation of IP Addresses

To improve the utilization of the IP address space, several protocols have been developed:

- Use of the Dynamic Host Configuration Protocol (DHCP), used to provide configuration parameters to IP hosts, including assignment of an IP address and subnet mask.
- Network Address Translation (NAT) and Network Address/Port Translation (NAT/NAPT), enables converting a large number of private IP addresses to a smaller number of global IP addresses.
- Use of Port Address Translation (PAT). PAT complements the NAT by allowing outside access to certain users on the network using private addresses.

### DHCP Services

Two types of services DHCP can be provided:

- DHCP relay services: in this mode, the IP router relays DHCP requests to a predefined DHCP server. The user can specify the maximum number of hops that a DHCP request can traverse before being discarded.
- DHCP server services: in this mode, the IP router itself serves as the DHCP server, which provides in response to DHCP requests an IP address, an IP subnet mask, a default gateway, and the IP addresses of two DNS servers (primary and secondary). The user can define different DHCP address pools. For each pool, the user specifies the IP address range, the default gateway, the primary and secondary DNS servers, and the lease time.

### NAT/NAPT Services

The routers can also provide network address translation (NAT) and network address/port translations (NAPT). The translations can apply to either the LAN or the WAN port:

- When the address translation is defined on a LAN port, the real IP addresses are located on the LAN side, and the virtual addresses are located on the WAN side.
- When the address translation is defined on a WAN port, the real IP addresses are located on the WAN side, and the virtual addresses are located on the LAN side.

Four types of translations can be defined:

- Dynamic (temporary) translation of a group of virtual IP addresses to a smaller group of real IP addresses, in accordance with the usage requirements received from the hosts using virtual IP addresses. This type is similar to the basic traditional NAT, as described in RFC2663 and RFC3022.

- Static (permanent) translation of a specific virtual IP address to a specific real IP address. This is a form of bidirectional NAT, as described in RFC2663.
- Transparent translation (no translation at all): the real and virtual IP addresses are identical. This is a form of bidirectional NAT, as described in RFC2663.
- Dynamic utilization of a single IP address in accordance with usage requirements received from the virtual IP addresses, using transport identifiers (port numbers) for multiplexing. In this case, This type is similar to NAPT, a form of traditional NAT, as described in RFC2663 and RFC3022.

All the translations, except for the transparent translation, hide the virtual addresses from the outside world (“outside” is determined by the type of interface, as explained above).

## PAT Services

PAT is a static translation that specifies a unique mapping between a [real IP address; port; protocol] and a [virtual IP address; port]. Its purpose is to enable access from the real IP side to a host using a virtual IP address that is included in an existing dynamic NAPT definition.

Therefore, a PAT translation can be defined only if a matching dynamic NAPT translation exists (*matching* means that a SINGLE PAT entry has the same real IP and virtual entries).

---

---

## 6. IP Routing Principles

### IP Communication between Hosts within the Same Network

The exchange of information between IP hosts is made in packets using the structure specified by the IP protocol. As explained in the *IP Packet Structure* section above, IP frames carry, within their header, the IP addresses of the destination and source hosts.

In accordance with the IP protocol, an IP host checks the addresses of all the received frames, and accepts only frames carrying its own IP address as the destination. The source address is then used to enable the destination to respond to the source.

An IP host will also respond to broadcasts (frames whose destination host identifier is “*all-ones*”).

**Note** *IP hosts support additional protocols within the IP suite, e.g., protocols used for connectivity checking, maintenance, etc. Therefore, IP hosts will accept additional types of messages, which are beyond the scope of this description.*

When checking the destination address of an IP frame, an IP host starts by checking the network identifier. If the network identifier is different, the host will immediately reject the frame. Therefore, IP hosts can communicate only if they have the same network identifier.

## IP Communication between Hosts Located on Different Networks

When the destination host is on a different IP network, before IP communication can be established it is necessary to find a route between the local IP network to the destination network.

When the user knows the route within the IP network that connects the two hosts (that is, all the intermediate hosts through which the IP packets exchanged between the two hosts must pass), it is possible to specify this route directly. This method is called **static routing**, and its main advantage is security.

However, in most cases it is not possible to specify a static route. Therefore, special procedures are needed to find a suitable route. These procedures are called routing, and the function that implements them is called **IP router**, or just **router**.

The most common type of routing protocol in use is called **Routing Information Protocol** (RIP), and the current version is referred to as RIP2. RIP1 is also still used, and is supported by some RAD's products.

The router can be a software program running on any IP host, but most often it is implemented on dedicated hardware, e.g., a special purpose PC.

Any router has at least two ports: a local LAN port and a WAN port. Accordingly, a router keeps two basic types of tables:

- Local LAN table. This table includes the IP addresses of the hosts connected to the LAN port of the router. These addresses are learned from the source addresses of the IP frames appearing on the local LAN. The LAN table is dynamically updated: this means that new hosts are added as they appear on the LAN, and hosts which are inactive for a specified (long) interval are removed.

In addition to the dynamic entries in the table, many routers permit the user to define static addresses in the LAN table. These static addresses are not removed from the table even if they are inactive for long periods.

- Routing table. This table includes information on the IP routes to other routers. Routers periodically exchange special messages that enable to build the routing information appearing in this table. The table is dynamically updated by these exchanges, therefore new routes are added whereas old, or disconnected, routes are removed.

In addition to the dynamic entries in the routing table, most routers permit the user to define static entries to the desired destinations in the routing table.

These entries are permanent, and in addition they have priority over dynamic (learned) entries: this means that the user-entered route will always be used (provided it is active), even if a dynamic route is also available.

A router receiving an IP frame with a destination address not located on the local LAN must find a route through the WAN or Internet to the destination. In general, this involves "asking" the other routers whether the destination is on their local LAN.

After the appropriate remote router is found, the process ends. At this stage, the local router can send the IP frame to the next router on the route to the

destination; the next router then sends the frame to the third router on the path to the destination, and so on. When the remote router gets the frame, it places it on the local LAN, where the addressed host can read it. Responses from the host automatically follow the same route through the Internet, but in reverse.

### Default Gateways

The general IP routing process described above has two shortcomings:

- It requires significant time to complete, because it involves a rather lengthy exchange of information
- The destination cannot influence the selection of a route, which is automatically performed.

To partially overcome these shortcomings, the IP protocol permits the destination host to specify one of the following routing options:

- The IP address (and the subnet mask) of the router to be used as a default to establish IP communication. The router selected for this purpose is called **default gateway**.
- The next hop address: this is the address to which any packets not addressed to hosts in the same net (or subnet) will be sent. This is usually the address of a specific IP router, which is “known” to be able to provide the route to the desired destinations.

For example, a default gateway must always be specified when the RADview station is not on the same IP network with the managed RAD’s unit. This is often the case when the RAD’s unit is managed inband.

## Tools for Checking IP Connectivity

The IP protocol is referred to as an unreliable connectionless packet delivery protocol, because each packet transmitted by an IP host travels on its own through the network until it eventually reaches its destination. To ensure reliable delivery, higher layer protocols are used, for example, the widely-used TCP protocol.

However, using higher layer protocols cannot provide an answer to the need to check that it is indeed possible to reach the desired IP destination address, because configuration errors or a network fault, for example, temporary congestion or failure of critical communication links, may still prevent the establishment of an IP connection. Therefore, it is often necessary to check for IP connectivity.

The IP protocol suite includes a special protocol, the Internet Control Message Protocol (ICMP), that enables IP hosts connected to the Internet to report a wide range of errors and provide information about the conditions that caused the errors. Support for this protocol is mandatory on every IP host.

ICMP includes a dedicated connectivity test procedure, implemented by means of two types of ICMP messages: **echo request** and **echo reply**. This procedure is often referred to as **pinging**: the host wanting to check IP connectivity to a destination sends one or more **ping** (echo request) messages, and the destination returns an echo reply message for each request. By comparing the number of

*pings* sent to the number received and the time needed for each reply to reach the *ping* source, the source host can obtain useful information regarding the transmission conditions.