

Internet Protocol: an Introduction

DoZ*

©*HackersCenter*

<http://www.hackerscenter.com>

18/12/05

Permission to make digital or hard copies of all or part of this work for personal use is granted without fee provided that copies are not made for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2005 © Obsidis n°1 22/12/2005

* doz@hackerscenter.com

1. Introduction

IP is a protocol developed to allow computers to share resources that support network communications. It was developed by a community of researchers centered on the Arpanet. Certainly the Arpanet is the best-known IP network. The Internet protocol is the world's most popular open-system protocol suite because it can be used to communicate across any set of interconnected networks and it is equally well suited for LAN and WAN communications. There is Dynamic IP which means your IP switches and there is Network IP that usually stays with you for while such as Cable. IP addresses on a network are given so that you can tell the location of the Internet Provider network or subnet where the host lives. IP addresses can be divided into two parts, the network ID and The host ID. The Internet protocols consist of a suite of communication protocols. The two best known Protocols are the Transmission Control Protocol (TCP) and the Internet Protocol (IP). In the latter there are four important layers you should know about: Application layer, Transport Layer, Internet Layer and Network Access Layer. IP is the 3rd layer of a network that contains addressing information and some control information that enables packets to be routed. IP/TCP has two main duties: delivering datagrams trough internet and translating (data fragment and reassembly) to support data links with different maximum-transmission unit sizes. An IP transmission is performed by a software component known as the vendor's implementation of TCP/IP. An IP implementation is a software component performing the functions that enable a computer to participate in a TCP/IP network. The IP standard ensure the compatibility of all IP implementations regardless of version or vendor. There are many TCP/IP Protocols but all of them have IP standards, thus a PC can communicate with any kind of Operating System. Internet Protocol is always being improved and advanced which leads to more changes. One new technology is the Ipv6, let's give a quick look into what's IPv6. The world is running out of addresses, the 32-bit address field of the current IP format (i.e. ipv-4) can

provide over three billion of possible host Ids, but it is important to remember how many of these three billion addresses are actually unusable. A network ID is typically assigned to a company, and that company controls the host IDs associated with its own network. The IP address format in IPv6 calls for 128-bit addresses. The reason for this, larger address space is supposedly to support one billion networks.

NOTE: in this article IP is also referred as IP/TCP.

2. TCP/IP Features

2.1 Logical addressing

A network adapter has a unique and permanent physical address. The physical address is a number that was given to the card at the factory. On a local area network, low-lying hardware-conscious protocols deliver data across the physical network using the adapter's physical address. There are many network types, and each has a different way of delivering data. On a basic Ethernet network, for example, a computer sends messages directly into the transmission medium. The network adapter of each computer listens to every transmission of the local network to determine whether a message is addressed to its own physical address. On large networks, of course, every network adapter can't listen to every message. As the transmission medium becomes more populated with computers, a physical addressing scheme cannot work correctly. Network admins often improve networks using devices such as routers to reduce network traffic. On routed networks, administrators need a way to equally divide network into smaller subnetworks (called sub-networks) and impose a class structure so that a message can travel efficiently to its destination. TCP/IP provides this subnetting capability through logical addressing. A logical address is an address configured through the network software. In TCP/IP, a computer's logical address is called an IP address and it is resolved from the corresponding hardware-specific physical address, using the ARP and RARP protocols.

2.2 Routing

A router is a special Tool that reads logical addressing information and direct data across the network to its destination. At the simplest level, a router divides a local subnet from the larger network. Data addressed to another computer or device on the local subnet does not cross the router and therefore doesn't clutter up the transmission lines of the greater network. If data is addressed to a computer outside the subnet, the router accordingly forwards the data. Very large networks such as Internet include many routers and provide multiple paths from the source to the destination. Network devices such as bridges, switches, and smart hubs also can filter traffic and reduce network traffic. Because these devices work with physical addresses rather than logical addresses, they cannot perform the complex routing functions.

2.3 Error control and flow control

The TCP/IP protocol suite provides features that ensure the reliable delivery of data across the network. These features include checking data for transmission errors and acknowledging successful receipt of a network message. The Transport Layer defines many of these error-checking, flow-controls, and acknowledgment functions through the TCP protocol. Lower-level protocols at TCP/IP's Network Access layer play a part in error control, too.

2.4 Name Resolution

Although the numeric IP address is probably easier than the network adapter's prefabricated physical address, the IP address is still designed for the convenience of the computers rather than the convenience of the user. People might have trouble remembering whether a computer's address is 111.121.131.146 or 111.121.131.156. TCP/IP, therefore, provides for a parallel structure of user-oriented alphanumeric names, called domain names or DNS names. This mapping of domain names to an IP address is called name resolution. Special servers called name servers, store tables showing how to translate these domain names to and from IP addresses. The computer addresses commonly

associated with email or the World Wide Web is expressed as DNS names. For example, www.Yahoo.com is 216.109.118.69 and 216.239.37.99 is www.Google.com. TCP/IP's name service system provides for a hierarchy of name servers that supply domain name/IP address mappings for DNS-registered computers on the network. DNS is the name resolution system for the Internet and is the most common name resolution method. However, some TCP/IP networks also support other methods for resolving alphanumeric names to IP addresses. Another common name resolution scheme is the Windows Internet Name Services (WINS) for resolving Microsoft Windows NetBIOS names to IP addresses. So, basically, instead of remembering IP addresses of DNS you have a Name for the DNS such as Google and Yahoo.

2.4 Application support

Several network applications might be running on the same computer. The protocol software must provide some means for determining which incoming packet belongs to each application. In TCP/IP, this interface from the network to the applications is accomplished through a system of logical channels called ports. Each port is a number that is used to identify the application. You can think of these ports as logical doors within the computer through which data can flow from the application to (and from) the protocol software. TCP/IP is actually entering into a new phase at the time of this writing.

2.5 IP Classes

IP addressing supports five different address classes: A, B, C, D, and E. only classes A, B, and C are available for commercial uses.

Most TCP/IP communication is either sent from one source computer to one destination computer or sent to all computers on the segment or network. Class D addresses, on the other hand, are used for multicasting. A multicast is a single message sent to a subset of the network. The four leftmost bits of a Class D network address always starts with the binary

pattern 1110, which corresponds to decimal numbers 224 through 239.

Class E networks are considered experimental. They are not normally used in any production environment.

2.6 IP Addressing

As with any other network-layer protocol, the IP addressing scheme is important for routing IP datagrams through an internetwork. Each IP address has specific components and follows a basic format. These IP addresses can be subdivided and used to create addresses for subnetworks. Each host on a TCP/IP network is assigned to a unique 32-bit logical address that is divided into two main parts: the network number and the host number. The network number identifies a network and must be assigned by the Internet Network Information Center. An IP address is a 32-bit binary address. This 32-bit address is subdivided into four 8-bit segments called octets. The IP address is almost always expressed in what is called dotted decimal format. In dotted decimal format, each octet is given as an equivalent decimal number. The four decimal values ($4 \times 8 = 32$ bits) are then separated with periods. Eight binary bits can represent any whole number from 0 to 255, so the segments of a dotted decimal address are decimal numbers from 0 to 255. A dotted decimal IP address looks like this: 50.78.5.150.

In Class A addresses, the first 8 bits of the IP address are used for the network ID. The final 24 bits are used for the host ID. In Class B addresses the first 16 bits of the IP address are used for the network ID while final 16 bits are used for the host ID. Finally in Class C addresses, the first 24 bits of the IP address are used for the network ID and the final 8 bits are used for the host ID.

More bits lead to more bit combinations. The Class A format provides a small number of possible network IDs and a huge number of possible host IDs for each network. A Class A network can support approximately 2^{24} , or 16,777,216 hosts. A Class C network, on the other hand, can provide host IDs for only a

small number of hosts (approximately 2^8 , or 256), but many more combinations of network IDs are available in the Class C format. (I would recommend reading more about Binary numbers).

A computer or router knows whether to interpret an IP address as a Class A, Class B, or Class C address. The designers of TCP/IP wrote the address rules such that the class of an address is obvious from the address itself. The first few bits of the binary address specify whether the address should be interpreted as a Class A, Class B, or Class C address. The owner of a network can divide the network into smaller subnetworks called subnets. Subnetting essentially borrows some of the bits of the host ID to create additional networks within the network. Class A and B networks, with their large host ID address spaces, make extensive use of subnetting.

2.7 IP Header Fields

Every IP datagram begins with an IP header. The TCP/IP software on the source computer constructs the IP header. The TCP/IP software at the destination uses the information enclosed in the IP header to process the datagram. The header contains a great deal of information, including the IP addresses of the source and destination computers, the length of the datagram, the IP version number, and special instructions to routers. The smallest size for an IP header is 20 bytes.

Version: This 4-bit field indicates which version of IP is being used. The current version of IP is 4 and the binary pattern for 4 is 0100.

IHL (Internet Header Length): This 4-bit field gives the length of the IP header in 32-bit words. The minimum header length is five 32-bit words. The binary pattern for 5 is 0101.

Type of Service: The source IP can designate special routing information. Some routers ignore the Type of Service field, although this field recently has received more attention with the emergence of Quality of Service (QoS) technologies. The main purpose of this 8-bit field is to provide priority for datagrams that are

waiting to pass through a router.

Total Length: This 16-bit field identifies the length, in octets, of the IP datagram. This length includes the IP header and the data payload.

Identification: This 16-bit field is an incrementing sequence number assigned to messages sent by the source IP. When a message is sent to the IP layer and it is too large to fit in one datagram, IP fragments the message into multiple datagrams, giving all datagrams the same identification number. This number is used from the receiving end to reassemble the original message.

Flags: The Flags field indicates fragmentation possibilities. The first bit is unused and should always have a value of zero. The next bit is called the **DF** (Don't Fragment) flag. The **DF** flag signifies whether fragmentation is allowed (value = 0) or not (value = 1). The next bit is the **MF** (More Fragments) flag, which tells the receiver that more fragments are on their way. When **MF** is set to 0, no more fragments need to be sent or the datagram never was fragmented.

Fragment Offset: This 13-bit field is a numeric value assigned to each successive fragment. IP at the destination uses the fragment offset to reassemble the fragments into the proper order. The offset value found here expresses the offset as a number of 8-byte units.

Time to Live: This bit field indicates the amount of time in seconds or router hops that the datagram can survive before being discarded. Every router examines and decrements this field by at least 1, or by the number of seconds the datagram is delayed inside the router. The datagram is discarded when this field reaches zero.

Protocol: The 8-bit Protocol field indicates the protocol that will receive the data payload. A datagram with the protocol identifier 6 (binary 00000110) is passed up the stack to the TCP module, for example. The following are some common protocol values:

Protocol Identification numbers: ICMP = 1
TCP = 6 UDP = 17

Hop: A hop or a router hop correlates to a

router that a datagram travels through on its way to its destination. If a datagram passes through five routers before arriving at its destination, the destination is said to be five hops, or five router hops away.

Header Checksum: This field holds a 16-bit calculated value to verify the validity of the header only.

Source IP Address: This 32-bit field holds the address of the source of the datagram.

Destination IP Address: This 32-bit field holds the destination address of the datagram and is used by the destination IP to verify correct delivery.

IP Data Payload: This field typically contains data destined for delivery to TCP or UDP (in the Transport layer), ICMP, or IGMP. The amount of data is variable but could include thousands of bytes.

2.8 TCP

Transmission Control Protocol provides reliable transmission of data for an IP environment. TCP corresponds to the transport layer (Layer 4) of the ISO/OSI reference model. Among its services TCP provides stream data transfer, reliability, efficient flow control, full-duplex operation, and multiplexing. With stream data transfer, TCP delivers an unstructured stream of bytes identified by sequence numbers. This service benefits applications because they do not have to cut data into blocks before handing it off to TCP. Instead, TCP groups bytes into segments and passes them to the IP level for delivery. TCP offers reliability by providing connection-oriented, end-to-end reliable packet delivery through an internet network. It does this by sequencing bytes with a forwarding acknowledgment number that indicates to the destination the next acknowledge the source expects to receive. Bytes not acknowledged within a specified time period are retransmitted. The reliability mechanism of TCP allows devices to deal with lost, delayed, duplicated, or misread packets. A time-out mechanism allows devices to detect lost packets and request the retransmission. De facto, we can say that TCP is

the Connection-Oriented Transport Protocol. When sending acknowledgments back to the source, the receiving TCP process indicates the highest sequence number it can receive without overflowing its internal buffers. Everything in TCP happens in the context of a connection. TCP sends and receives data through a connection, which must be requested, opened, and closed according to the rules of TCP. When it is time to close the connection, the closing end-host, let's say Computer A, places a segment in the queue with the `FIN` flag set to one. The application then enters what is called the `fin-wait` state. In the `fin-wait` state, Computer A TCP software continues to receive segments and processes the segments already in the queue, but no additional data is accepted from the application. When Computer B receives the `FIN` segment, it returns an acknowledgment to the `FIN`, sends any remaining segments, and notifies the local application that a `FIN` was received. Computer B sends a `FIN` segment to Computer A, which Computer A acknowledges, and the connection is closed.

3. UDP

UDP is much simpler than TCP, and it doesn't perform any of the functions listed in the **TCP** section. UDP is usually described as having no error-checking capabilities, in fact, it is capable of performing only rudimentary error checking. However, it is better to characterize UDP as having the capability for limited error checking. The UDP datagram includes a checksum value that the receiving machine can use to test the integrity of the data. (Often, this checksum test is optional and can be disabled on the receiving machine to speed up processing of incoming data.) The UDP datagram includes a pseudo-header that encompasses the destination address for the datagram, thus providing a mean of checking for misdirected datagrams. Also, if the receiving UDP module receives a datagram directed to an inactive or undefined UDP port, it returns an ICMP message notifying the source machine that the port is unreachable. UDP's lean, connectionless design makes it the protocol of choice for network broadcast

situations. A broadcast communication is a single message that will be received and processed by all hosts of the subnet. Obviously, if the source host had to simultaneously open a TCP-style connection with every computer on the subnet in order to send a single broadcast, the result could be a significant erosion of network performance. Second, UDP does not offer the resequencing of data provided by TCP. Resequencing is most significant on a large network, such as the Internet, where the segments of data might take different paths and experience significant delays in router buffers. For local networks, the lack of a resequencing feature in UDP typically does not lead to unreliable reception. The User Datagram Protocol (UDP) is a connectionless transport-layer protocol (Layer 4) that belongs to the Internet Protocol family. UDP is basically an interface between IP and upper-layer processes. UDP protocol's ports distinguish multiple applications running on a single device from one another. Unlike the TCP, UDP adds no reliability, flow-control, or error-recovery functions to IP. Because of UDP simplicity, UDP headers contain fewer bytes and consume less network overhead than TCP. UDP is useful in situations where the reliability mechanisms of TCP are not necessary, such as in cases where a higher-layer protocol might provide error and flow control. UDP is the transport protocol for several well-known application-layer protocols, including Network File System (NFS), Simple Network Management Protocol (SNMP), Domain Name System (DNS), and Trivial File Transfer Protocol (TFTP). The UDP packet format contains four fields, these include source and destination ports, length, and checksum fields.

4. ICMP

Data sent to a remote host often travels through one or more routers. These routers can encounter a number of problems sending the message to its ultimate destination. Routers use Internet Control Message Protocol (ICMP) messages to notify the source IP about these problems.

ICMP is also used for other diagnosis and

troubleshooting functions.

Here are the most common ICMP messages. Few other conditions generate ICMP messages but their frequency of occurrence is quite low.

Echo Request and Echo Reply — ICMP is often used during testing. When a technician uses the ping command to check connectivity with another host, he is using ICMP. Ping sends a datagram to an IP address and requests the destination computer to return the data sent in a response datagram. The commands actually being used are the ICMP Echo-Request and Echo-Reply.

Source Quench — if a fast host is sending large amounts of data to a remote host, the traffic volume can overwhelm the router. The router might use ICMP to send a Source Quench message to the source IP to ask it to slow down the rate at which it is shipping data. If necessary, additional source quenches can be sent to the source IP.

Destination Unreachable — if a router receives a datagram that cannot be delivered, ICMP returns a Destination Unreachable message to the source IP.

Time Exceeded — an host sends this message to the source IP if a datagram is discarded because TTL reaches zero. This indicates that the destination is too many router hops away to reach with the current TTL value, or it indicates router table problems that make the datagram to loop through the same routers continuously.

A routing loop occurs when a datagram circulates through the same routers continuously and never reaches its destination. Suppose three routers are located in New York, Moscow, and Rome. The New York router sends datagrams to Moscow, which sends them to Rome, which sends them back to New York again. The datagram becomes trapped and will circulate continuously through these three routers until the TTL reaches zero. A routing loop should not occur, but occasionally it does. A routing loop sometimes occurs when a network administrator places bad static routing entries in a routing

table.

Fragmentation needed — an host sends this message if it receives a datagram with the Don't Fragment bit set and if the router needs to fragment the datagram in order to forward it to the next router or the destination.

5. IP Security

It's easy to intercept and read an unprotected packet of data traveling over a public network. Some times, data might contain user's password or sensitive information you don't want anyone else to read, like credit card numbers or Bank accounts. The fact is that even if the data isn't particularly secret, users are exposed to hackers or any criminals willing to get those information. Encryption is the best way of altering data to make it unreadable to unauthorized users. Data is encrypted by the sender and then travels over the network in encoded, unreadable form. The receiving computer then decrypts the data in order to read it. You might wonder: what if Hackers have a tool to decrypt the packets and read them? That's why you should trust you receiver and know who you're sending packets because there are tools to decrypt data.

An encryption algorithm is essentially a set of mathematical steps used to transform the data into its unreadable form. The secret part of the process is called the key. The result of the encryption process depends on the value of the key. Therefore, as long as the value of the key is kept secret, unauthorized users will not be able to read the data even if they have the necessary decryption tool.

Some ways to make TCP/IP more secure is by using (SSL) Secure Sockets Layer or the (IPSec) IP Security. The purpose of SSL is to provide a layer of security between the sockets at the Transport layer. The purpose of IP Security is an alternative security protocol system used on TCP/IP networks. IPSec operates inside the TCP/IP protocol stack, beneath the Transport layer. Because the security system is implemented beneath the Transport layer, the applications operating above the Transport layer do not need

knowledge of the security system.

6. Running IP Configuration from Command Prompt (on Windows)

Have you ever played with Command prompt? Well, if you did you probably know the basics like FTP or Telnet and maybe IPconfig or Ipconfig /all.

Let's say you want to ping some one, you should type into command prompt: ping 127.0.0.1 (make sure you put the right IP address).

Connectivity Utilities:

1. IPConfig - A Windows utility that displays TCP/IP configuration settings.
1. Ping - An utility that tests for network connectivity.
2. Arp - An utility that lets you view (can modify) the ARP cache of a local or remote computer. The ARP cache contains the physical address to IP address mappings.
3. Traceroute - An utility that traces the path of a datagram through the internet.
4. Route - An utility that lets you view,

add, or edit entries in a routing table.

5. Netstat - An utility that displays IP, UDP, TCP, and ICMP statistics.
6. NBTstat - An utility that displays statistics on NetBIOS and NBT.
7. Hostname - An utility that returns the hostname of the local host.

File Transfer Utilities:

1. Ftp - A basic file transfer client.
1. Tftp - A basic file transfer utility using UDP.
2. Rcp - A simple remote file transfer utility.

Remote Utilities:

1. Telnet - A remote terminal utility.
2. Rexec - An utility which runs commands on a remote computer through the rexecd daemon.
3. Rsh - An utility invoking a shell on the remote computer, in order to execute commands.
4. Finger - An utility which displays user information.