


Transport Layer

- Purpose of transport layer services:
 - multiplexing/demultiplexing
 - reliable data transfer
 - flow control
 - congestion control
- Connection-less transport: UDP
- Connection-oriented transport: TCP
 - reliable transfer
 - flow control
 - connection management

Transport services and protocols

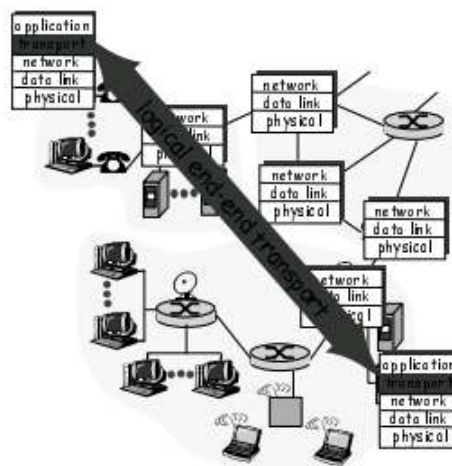
- provide logical communication between application processes running on different hosts
- transport protocols run in end systems via software
- transport vs network layer services:
 - network layer: data transfer between end systems
 - transport layer: data transfer between processes
 - relies on, enhances, network layer services



Transport-layer protocols

Internet transport services:

- reliable, in-order unicast delivery (TCP)
 - congestion
 - flow control
 - connection setup
- unreliable (“best-effort”), unordered unicast or multicast delivery: UDP
- services not available:
 - real-time
 - bandwidth guarantees
 - reliable multicast



UDP: User Datagram Protocol [RFC 768]

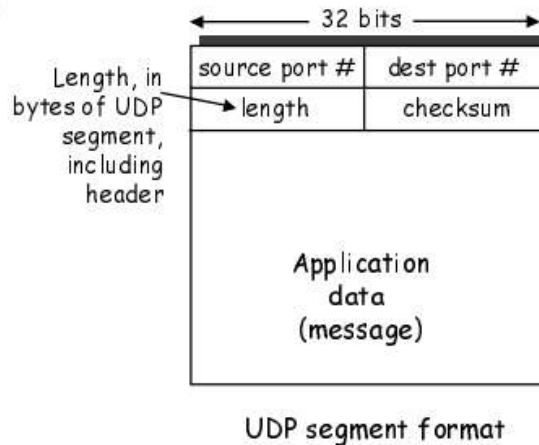
- “no frills,” “bare bones” Internet transport protocol
- “best effort” service, UDP segments may be:
 - lost
 - delivered out of order to app
- *connectionless*:
 - no handshaking between UDP sender, receiver
 - each UDP segment handled independently of others

Why is there a UDP?

- no connection establishment (which can add delay)
- simple: no connection state at sender, receiver
- small segment header
- no congestion control: UDP can blast away as fast as desired

UDP: more

- often used for streaming multimedia apps
 - Controversial: no congestion control
- other UDP uses (why?):
 - DNS
 - SNMP
- reliable transfer over UDP: add reliability at application layer
 - application-specific error recover!



UDP checksum

Goal: detect “errors” (e.g., flipped bits) in transmitted segment

Sender:

- treat segment contents as sequence of 16-bit integers
- checksum: addition (1’s complement sum) of segment contents
- sender puts checksum value into UDP checksum field

Receiver:

- compute checksum of received segment
- check if computed checksum equals checksum field value:
 - NO - error detected
 - Toss data OR
 - Pass to app with warning
 - YES - no error detected.

Connection Oriented Transport Protocol Mechanisms

- Properties of connection-oriented Transport Protocols:
 - Logical connection
 - Establishment
 - Maintenance termination
 - Reliable
 - e.g. TCP

Connection-Oriented Transport via Reliable Network Layer

- Transport Layer Services like TCP are complicated – to start, let's first assume we are working with a reliable network layer service
 - e.g. reliable packet switched network using X.25
 - e.g. frame relay using LAPF control protocol
 - e.g. IEEE 802.3 using connection oriented LLC service
 - NOT IP! IP is unreliable
- Assume arbitrary length message
- Transport service is end to end protocol between two systems on same network

3.0 Standard TCP Congestion Control Algorithms

The standard fare in TCP implementations today can be found in RFC 2581 [2]. This reference document specifies four standard congestion control algorithms that are now in common use. Each of the algorithms noted within that document was actually designed long before the standard was published [9], [11]. Their usefulness has passed the test of time.

The four algorithms, Slow Start, Congestion Avoidance, Fast Retransmit and Fast Recovery are described below.

3.1 Slow Start

Slow Start, a requirement for TCP software implementations is a mechanism used by the sender to control the transmission rate, otherwise known as sender-based flow control. This is accomplished through the return rate of acknowledgements from the receiver. In

other words, the rate of acknowledgements returned by the receiver determine the rate at which the sender can transmit data.

When a TCP connection first begins, the Slow Start algorithm initializes a *congestion window* to one segment, which is the maximum segment size (MSS) initialized by the receiver during the connection establishment phase. When acknowledgements are returned by the receiver, the congestion window increases by one segment for each acknowledgement returned. Thus, the sender can transmit the minimum of the congestion window and the advertised window of the receiver, which is simply called the *transmission window*.

Slow Start is actually not very slow when the network is not congested and network response time is good. For example, the first successful transmission and acknowledgement of a TCP segment increases the window to two segments. After successful transmission of these two segments and acknowledgements completes, the window is increased to four segments. Then eight segments, then sixteen segments and so on, doubling from there on out up to the maximum window size advertised by the receiver or until congestion finally does occur.

At some point the congestion window may become too large for the network or network conditions may change such that packets may be dropped. Packets lost will trigger a timeout at the sender. When this happens, the sender goes into congestion avoidance mode as described in the next section.

3.2 Congestion Avoidance

During the initial data transfer phase of a TCP connection the Slow Start algorithm is used. However, there may be a point during Slow Start that the network is forced to drop one or more packets due to overload or congestion. If this happens, Congestion Avoidance is used to slow the transmission rate. However, Slow Start is used in conjunction with Congestion Avoidance as the means to get the data transfer going again so it doesn't slow down and stay slow.

In the Congestion Avoidance algorithm a retransmission timer expiring or the reception of duplicate ACKs can implicitly signal the sender that a network congestion situation is occurring. The sender immediately sets its transmission window to one half of the current window size (the minimum of the congestion window and the receiver's advertised window size), but to at least two segments. If congestion was indicated by a timeout, the congestion window is reset to one segment, which automatically puts the sender into Slow Start mode. If congestion was indicated by duplicate ACKs, the Fast Retransmit and Fast Recovery algorithms are invoked (see below).

3.3 Fast Retransmit

When a duplicate ACK is received, the sender does not know if it is because a TCP segment was lost or simply that a segment was delayed and received out of order at the receiver. If the receiver can re-order segments, it should not be long before the receiver sends the latest expected acknowledgement. Typically no more than one or two duplicate ACKs should be received when simple out of order conditions exist. If however more than two duplicate ACKs are received by the sender, it is a strong indication that at least one segment has been lost. The TCP sender will assume enough time has lapsed for all segments to be properly re-ordered by the fact that the receiver had enough time to send three duplicate ACKs.

When three or more duplicate ACKs are received, the sender does not even wait for a retransmission timer to expire before retransmitting the segment (as indicated by the position of the duplicate ACK in the byte stream). This process is called the Fast Retransmit algorithm and was first defined in [11]. Immediately following Fast Retransmit is the Fast Recovery algorithm.

3.4 Fast Recovery

Since the Fast Retransmit algorithm is used when duplicate ACKs are being received, the TCP sender has implicit knowledge that there is data still flowing to the receiver. Why? The reason is because duplicate ACKs can only be generated when a segment is received. This is a strong indication that serious network congestion may not exist and that the lost segment was a rare event. So instead of reducing the flow of data abruptly by going all the way into Slow Start, the sender only enters Congestion Avoidance mode.

Rather than start at a window of one segment as in Slow Start mode, the sender resumes transmission with a larger window, incrementing as if in Congestion Avoidance mode. This allows for higher throughput under the condition of only moderate congestion [23].

To summarize this section of the paper, figure 2 below depicts what a typical TCP data transfer phase using TCP congestion control might look like. Notice the periods of exponential window size increase, linear increase and drop-off. Each of these scenarios depicts the sender's response to implicit or explicit signals it receives about network conditions.

Dynamic Host Configuration

Protocol(DHCP) is an application layer protocol which is used to provide:

1. Subnet Mask (Option 1 – e.g., 255.255.255.0)
2. Router Address (Option 3 – e.g., 192.168.1.1)
3. DNS Address (Option 6 – e.g., 8.8.8.8)
4. Vendor Class Identifier (Option 43 – e.g., ‘unifi’ = 192.168.1.9 ##where unifi = controller)

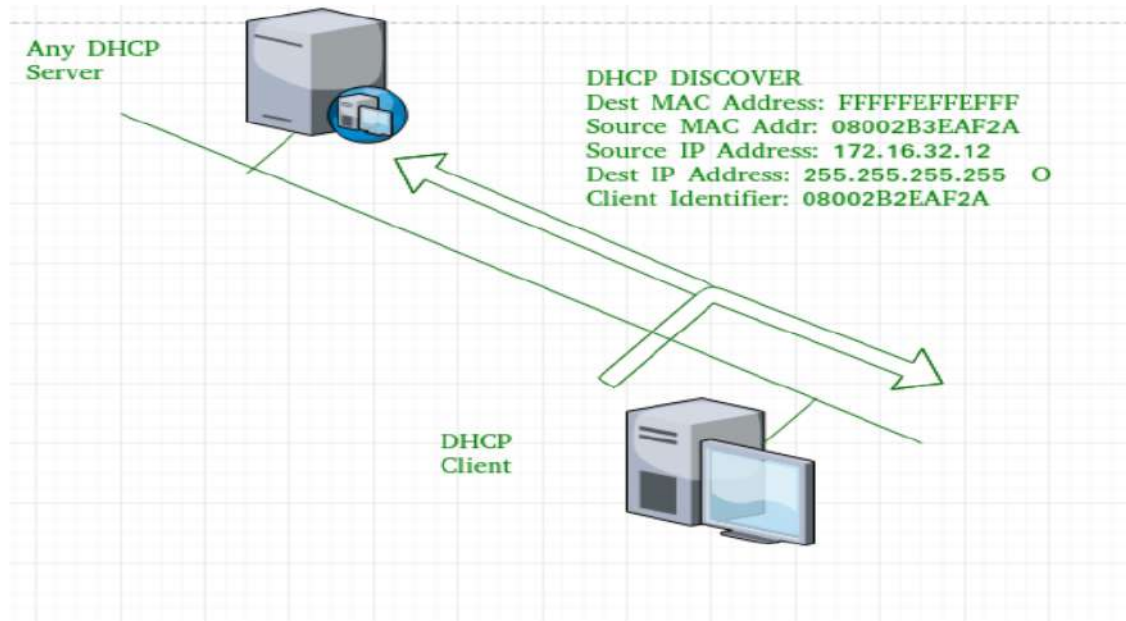
DHCP is based on a client-server model and based on discovery, offer, request, and ACK.

DHCP port number for server is 67 and for the client is 68. It is a Client server protocol which uses UDP services. IP address is assigned from a pool of addresses. In DHCP, the client and the server exchange mainly 4 DHCP messages in order to make a connection, also called DORA process, but there are 8 DHCP messages in the process.

These messages are given as below:

1. DHCP discover message –

This is a first message generated in the communication process between server and client. This message is generated by Client host in order to discover if there is any DHCP server/servers are present in a network or not. This message is broadcasted to all devices present in a network to find the DHCP server. This message is 342 or 576 bytes long



As shown in the figure, source MAC address (client PC) is 08002B2EAF2A, destination MAC address(server) is FFFFFFFF00000000, source IP address is 172.16.32.12 (because PC has no IP address till now) and destination IP address is 255.255.255.255 (IP address used for broadcasting). As the discover message is broadcast to find out the DHCP server or servers in the network therefore broadcast IP address and MAC address is used.

2. DHCP offer message –

The server will respond to host in this message specifying the unleased IP address and other TCP configuration information. This message is broadcasted by server. Size of message is 342 bytes. If there are more than one DHCP servers present in the network then client host will accept the first DHCP OFFER message it receives. Also a server ID is specified in the packet in order to identify the server.

Now, for the offer message, source IP address is 172.16.32.12 (server's IP address in the example), destination IP address is 255.255.255.255 (broadcast IP address) ,source MAC address is 00AA00123456, destination MAC address is FFFFFFFFFFFFFFFF. Here, the offer message is broadcast by the DHCP server therefore destination IP address is broadcast IP address and destination MAC address is FFFFFFFFFFFFFFFF and the

source IP address is server IP address and MAC address is server MAC address.

Also the server has provided the offered IP address 192.16.32.51 and lease time of 72 hours(after this time the entry of host will be erased from the server automatically) . Also the client identifier is PC MAC address (08002B2EAF2A) for all the messages.

3. DHCP request message –

When a client receives a offer message, it responds by broadcasting a DHCP request message. The client will produce a gratuitous ARP in order to find if there is any other host present in the network with same IP address. If there is no reply by other host, then there is no host with same TCP configuration in the network and the message is broadcasted to server showing the acceptance of IP address

.A Client ID is also added in this message.


Now, the request message is broadcast by the client PC therefore source IP address is 0.0.0.0(as the client has no IP right now) and destination IP address is 255.255.255.255 (broadcast IP address) and source MAC address is 08002B2EAF2A (PC MAC address) and destination MAC address is FFFFFFFFFFFFFFFF.

Note – This message is broadcast after the ARP request broadcast by the PC to find out whether any other host is not using that offered IP. If there is no reply, then the client host broadcast the DHCP request message for the server showing the acceptance of IP address and Other TCP/IP Configuration.

SNMP

- SNMP stands for **Simple Network Management Protocol**.
- SNMP is a framework used for managing devices on the internet.
- It provides a set of operations for monitoring and managing the internet.

SNMP Concept



- SNMP has two components Manager and agent.
- The manager is a host that controls and monitors a set of agents such as routers.
- It is an application layer protocol in which a few manager stations can handle a set of agents.
- The protocol designed at the application level can monitor the devices made by different manufacturers and installed on different physical networks.
- It is used in a heterogeneous network made of different LANs and WANs connected by routers or gateways.

Managers & Agents

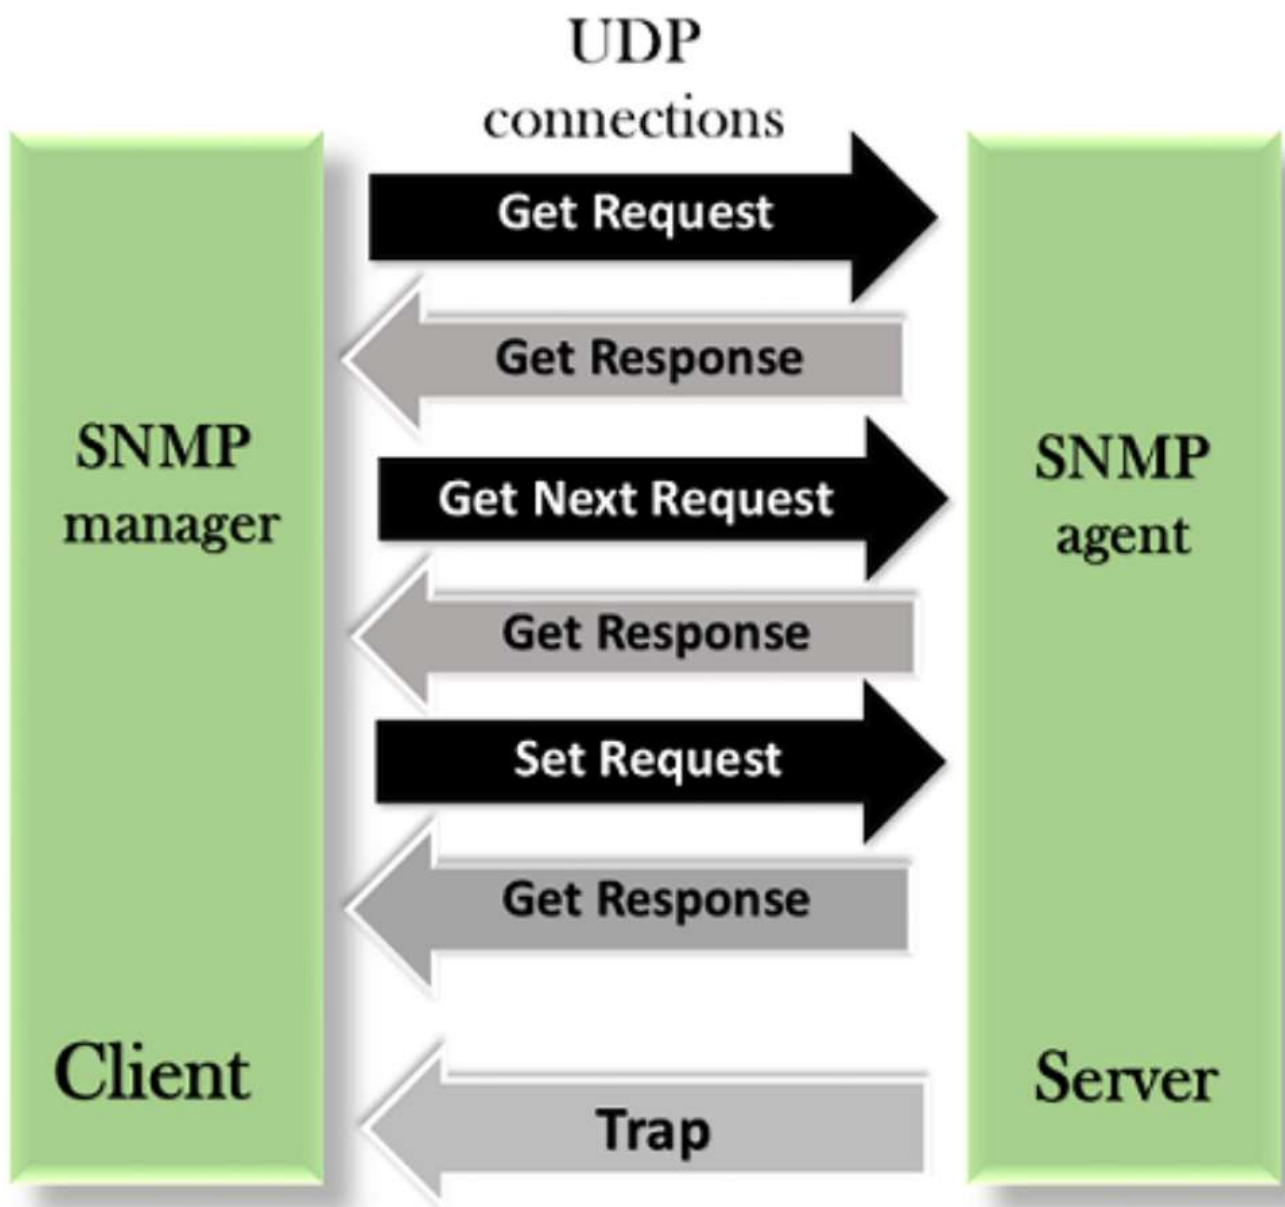
- A manager is a host that runs the SNMP client program while the agent is a router that runs the SNMP server program.
- Management of the internet is achieved through simple interaction between a manager and agent.
- The agent is used to keep the information in a database while the manager is used to access the values in the database. For example, a router can store the appropriate variables such as a number of packets received and forwarded while the manager can compare these variables to determine whether the router is congested or not.

Management Components

- Management is not achieved only through the SNMP protocol but also the use of other protocols that can cooperate with the SNMP protocol. Management is achieved through the use of the other two protocols: SMI (Structure of management information) and MIB(management information base).
- Management is a combination of SMI, MIB, and SNMP. All these three protocols such as abstract syntax notation 1 (ASN.1) and basic encoding rules (BER).

SNMP

SNMP defines five types of messages:
GetRequest, GetNextRequest, SetRequest,
GetResponse, and Trap.



GetRequest: The GetRequest message is sent from a manager (client) to the agent (server) to retrieve the value of a variable.

GetNextRequest: The GetNextRequest message is sent from the manager to agent to retrieve the value of a variable. This type of message is used to retrieve the values of the entries in a table. If the manager does not know the indexes of the entries, then it will not be able to retrieve the values. In such situations, GetNextRequest message is used to define an object.

GetResponse: The GetResponse message is sent from an agent to the manager in response to the GetRequest and GetNextRequest message. This message contains the value of a variable requested by the manager.

SetRequest: The SetRequest message is sent from a manager to the agent to set a value in a variable.


Trap: The Trap message is sent from an agent to the manager to report an event. For example, if the agent is rebooted, then it informs the manager as well as sends the time of rebooting.

SMTP


- SMTP stands for Simple Mail Transfer Protocol.
- SMTP is a set of communication guidelines that allow software to transmit an electronic mail over the internet is called **Simple Mail Transfer Protocol**.
- It is a program used for sending messages to other computer users based on e-mail addresses.
- It provides a mail exchange between users on the same or different computers and it also supports:
 - It can send a single message to one or more recipients.

- The main purpose of SMTP is used to set up communication rules between servers. The servers have a way of identifying themselves and announcing what kind of communication they are trying to perform. They also have a way of handling the errors such as incorrect email address. For example, if the recipient address is wrong, then receiving server reply with an error message of some kind.

Components of SMTP



- First, we will break the SMTP client and SMTP server into two components such as user agent (UA) and mail transfer agent (MTA). The user agent (UA) prepares the message, creates the envelope and then puts the message in the envelope. The mail transfer agent (MTA) transfers this mail across the internet.



- SMTP allows a more complex system by adding a relaying system. Instead of just having one MTA at sending side and one at receiving side, more MTAs can be added, acting either as a client or server to relay the email.

1. **Composition of Mail:** A user sends an e-mail by composing an electronic mail message using a Mail User Agent (MUA). Mail User Agent is a program which is used to send and receive mail. The message contains two parts: body and header. The body is the main part of the message while the header includes information such as the sender and recipient address. The header also includes descriptive information such as the subject of the message. In this case, the message body is like a letter and header is like an envelope that contains the recipient's address.

2. **Submission of Mail:** After composing an email, the mail client then submits the completed e-mail to the SMTP server by using SMTP on TCP port 25.
3. **Delivery of Mail:** E-mail addresses contain two parts: username of the recipient and domain name. For example, vivek@gmail.com, where "vivek" is the username of the recipient and "gmail.com" is the domain name.

If the domain name of the recipient's email address is different from the sender's domain name, then MSA will send the mail to the Mail Transfer Agent (MTA). To relay the email, the MTA will find the target domain. It checks the MX record from Domain Name System to obtain the target domain. The MX record contains the domain name and IP address of the recipient's domain. Once the record

is located, MTA connects to the exchange server to relay the message.

4. **Receipt and Processing of Mail:** Once the incoming message is received, the exchange server delivers it to the incoming server (Mail Delivery Agent) which stores the e-mail where it waits for the user to retrieve it.
5. **Access and Retrieval of Mail:** The stored email in MDA can be retrieved by using MUA (Mail User Agent). MUA can be accessed by using login and password.

Hyper Text Transfer Protocol-


- HTTP is short for **Hyper Text Transfer Protocol**.
- It is an application layer protocol.

Purpose-

- It is mainly used for the retrieval of data from websites throughout the internet.
- It works on the top of TCP/IP suite of protocols.

Working-

- Web browser is the client.
- Client communicates with the web server hosting the website.



Whenever a client requests some information (say clicks on a hyperlink) to the website server.

The browser sends a request message to the HTTP server for the requested objects.


Then-

- HTTP opens a connection between the client and server through **TCP**.
- HTTP sends a request to the server which collects the requested data.
- HTTP sends the response with the objects back to the client.
- HTTP closes the connection.

HTTP Connections-

HTTP connections can be of two types-

1. Non-persistent HTTP connection
2. Persistent HTTP connection



Non-persistent HTTP connection	Persistent HTTP connection
<p>Non-persistent HTTP connection is one that is used for serving exactly one request and sending one response.</p>	<p>Persistent HTTP connection is one that can be used for serving multiple requests.</p>
<p>HTTP server closes the TCP connection automatically after sending a HTTP response.</p>	<p>HTTP server closes the TCP connection only when it is not used for a certain configurable amount of time.</p>
<p>A new separate TCP connection is used for each object.</p>	<p>A single TCP connection is used for sending multiple objects one after the other.</p>

<p>HTTP 1.0 supports non-persistent connections by default.</p>	<p>HTTP 1.1 supports persistent connections by default.</p>
<p><u>Example-</u></p> <p>Suppose a request has been made for a HTML page that contains 10 images (called objects).</p> <p>Then,</p> <p>With non-persistent connection, all the 11 objects (1 page + 10 images) will be sent one by one.</p> <p>For getting each object, a new separate connection will be opened and used.</p>	<p><u>Example-</u></p> <p>Suppose a request has been made for a HTML page that contains 10 images (called objects).</p> <p>Then,</p> <p>With persistent connection, all the 11 objects (1 page + 10 images) will be sent one after the other using a single TCP connection.</p>