

Objective

The assignment has been designed to develop the understanding of IPv6 and UDP.

Instructions:

Please read the following instructions carefully before solving & submitting assignment:
Assignment should be in your own wordings not copied from internet, handouts or books.
It should be clear that your assignment will not get any credit (zero marks) if:

- The assignment is submitted after due date.
- The submitted assignment does not open or file is corrupt.
- The assignment not submitted in MS Word file format.
- The assignment is copied (from other student or copied from handouts or internet).
- Student ID is not mentioned in the assignment file or name of file is other than student ID.

For any query about the assignment, contact at CS610@vu.edu.pk

GOOD LUCK

Q.1. Why does IPv6 use separate extension headers instead of fields in a single, fixed header? [5 Marks]

Solution:

There are two reasons:

- Economy
- Extensibility

Economy is easiest to understand: partitioning the datagram functionality into separate headers is economical because it saves space. Although IPv6 includes many features, designers expect a given datagram to use only a small subset. Having separate headers makes it possible to define a large set of features without requiring each datagram header to have at least one field for each feature. For example, although many IPv4 datagrams are not fragmented, the IPv4 header has fields used to hold fragmentation information. In contrast, IPv6 does not waste space on fields for fragmentation unless the datagram is fragmented. Because most datagrams only need a few headers, avoiding unnecessary header fields can save considerable space. Smaller datagrams also take less time to transmit. Thus, reducing datagram size also reduces the bandwidth consumed.

To understand extensibility, consider adding a new feature to a protocol. A protocol like IPv4 that uses a fixed header format requires a complete change — the header must be redesigned to accommodate fields needed to support the new feature. In IPv6, however, existing protocol headers can remain unchanged. A new *NEXT HEADER* type is defined as well as a new header format.

Q.2. What endpoint values must be specified by an application that engages in 1-to-1 communication? In 1-to-many? In many-to-1? [5 Marks]

Solution:

An application that engages in,

1-to-1 communication specifies, local port number, remote IP address, remote protocol port number

An application that engages in,

Many-to-1 communication specifies local port number and informs UDP that the remote endpoint can be any system.

An application that engages in,

1-to-many communication specifies local port number and informs UDP that the remote endpoint can be any system Multicast address.

Q.3. Calculate the size of the largest possible UDP message. (Hint: the entire UDP message must fit in an IP datagram.) [5 Marks]

Solution:

If a UDP message is larger than the network MTU, IP will fragment the resulting datagram, which reduces efficiency. As a consequence, a message size that produces data grams that fit in a standard MTU is used to avoid fragmentation. In particular, because most parts of the Internet now support an MTU of 1500 octets, often a message size of **1400 or 1450 octets** are chosen to leave plenty of space for IP and UDP headers.