Routing Traffic

Q12.1 What is a bit in an IP address, and how does that relate to the subnet?

Typically, we refer to the length of an IP address as the number of bits that we use to represent it. If you remember from Part I, a bit is simply a 0 or a 1. Every decimal number can be represented using a sequence of bits, forming a **binary number**, and vice versa. Each of the four numbers making up an IPv4 address consists of eight bits, for a total of 32. An IPv6 address, by contrast, has 128 bits.

In the book, we used 24-bit IP address prefixes, *e.g.*, 127.12.5.0/24. The prefix doesn't have to be 24 bits long. The number after the slash will always tell us how many bits make up the subnet. For instance, with 127.12.0.0/16, the subnet is 127.12, and for 127.0.0.0/8, it is 127. Note that when we are specifying the subnet, we fill digits that would make up the host ID with 0's.

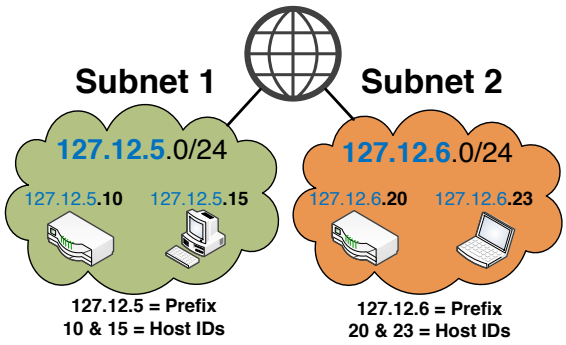

Illustration 35: An IP address can be divided into two parts: the prefix, which indicates the subnet, and the host identifier, indicating the specific machine within the subnet.

Q12.2 Can we go through another example of subnets?

In Illustration 35, we have two subnets, with prefixes 127.12.5 and 127.12.6, respectively. Within the first subnet, we show two machines with host IDs 10 and 15, and in the second, we show two with IDs 20

and 23.

Q12.3 In the Bellman-Ford Algorithm, why is the number of source nodes the number of steps needed to guarantee that the shortest paths have been discovered?

If the number of source nodes is n, then any path traversing $n + 1$ or more hops will have to go around a **cycle** in the graph that starts and stops at the same node before arriving at the destination. A cycle can only add to the total cost, and would never be part of a shortest path.