5.4 DOMAIN NAME SYSTEM (DNS)

The DNS is a distributed database that resides on multiple machines on the internet. It provide e-mail routing information.The DNS protocol runs over UDP and uses port 53.

Figure 5.4.1 shows how TCP/IP uses a DNS client and a DNS server to map a name to an address. A user wants to use a file transfer client to access the corresponding file transfer server running on a remote host. The user knows only the file transfer server name, such as a filesource.com. The TCP/IP suite needs the IP address of the file transfer server to make the connection.

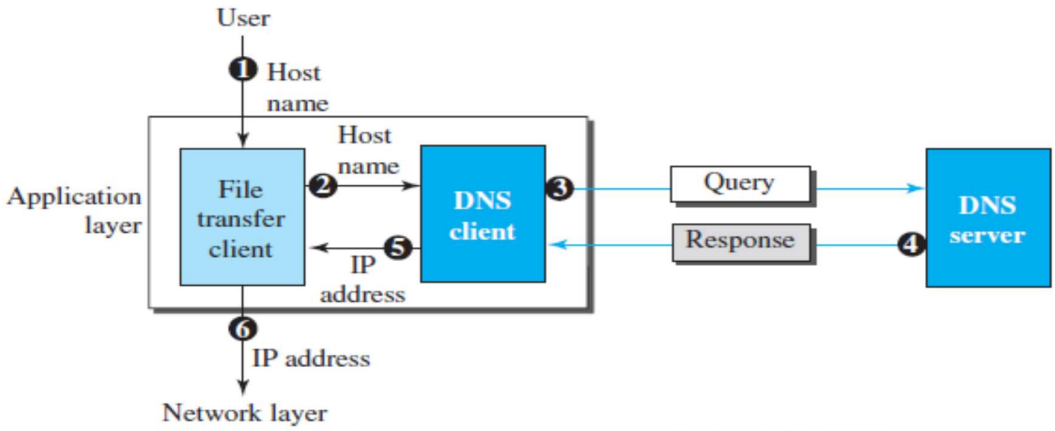

Fig5.4.1: Purpose of DNS

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-910]

The following six steps map the host name to an IP address:

The user passes the host name to the file transfer client.

The file transfer client passes the host name to the DNS client.

Each computer, after being booted, knows the address of one DNS server. The DNS client sends a message to a DNS server with a query that gives the file transfer server name using the known IP address of the DNS server.The DNS server responds with the IP address of the desired file transfer server. The DNS server passes the IP address to the file transfer client. The file transfer client now uses the received IP address to access the file transfer server.

Domain Name Space

To have a hierarchical name space, a domain name space was designed. In this design the names are defined in an inverted-tree structure with the root at the top. The tree has 128 levels: level 0 (root) to level 127 (see Figure 5.4.2).

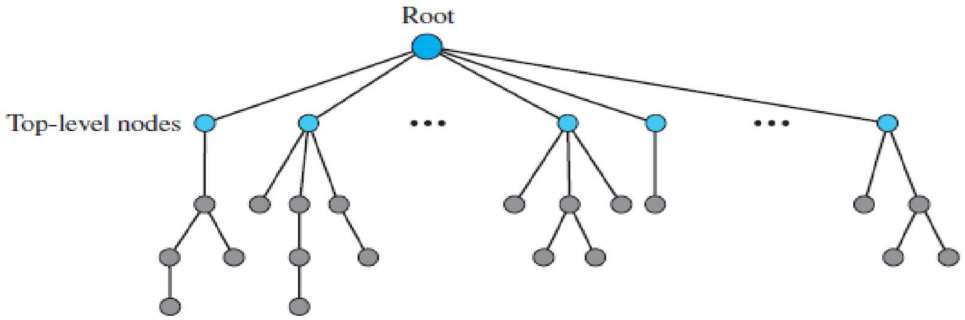

Fig5.4.2: DNS tree.

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-912]

Label

Each node in the tree has a label, which is a string with a maximum of 63 characters. The root label is a null string (empty string).

Domain Name

Each node in the tree has a domain name. A full domain name is a sequence of labels separated by dots (.). The domain names are always read from the node up to the root as in figure 5.4.3. The last label is the label of the root (null). This means that a full domain name always ends in a null label, which means the last character is a dot because the null string is nothing.

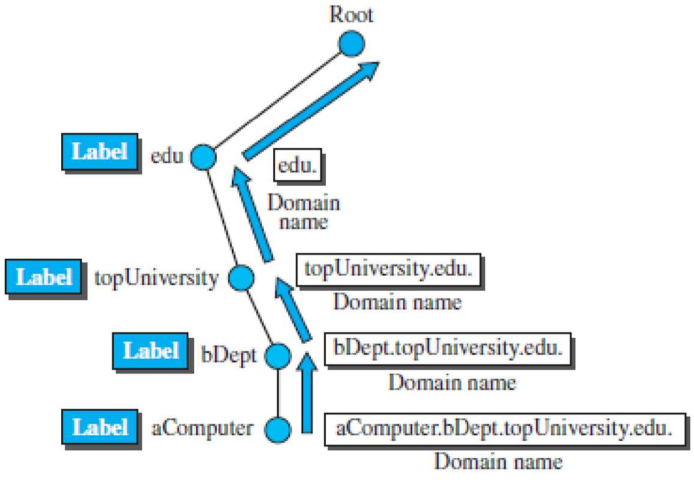

Fig5.4.3: Domain names.

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-913]

If a label is terminated by a null string, it is called a fully qualified domain name (FQDN). The name must end with a null label, but because null means nothing, the label ends with a dot.

If a label is not terminated by a null string, it is called a partially qualified domain name (PQDN). A PQDN starts from a node, but it does not reach the root. It is used when the name to be resolved belongs to the same site as the client. Here the resolver can supply the missing part, called the suffix, to create an FQDN.

A domain is a sub tree of the domain name space. The name of the domain is the name of the node at the top of the sub tree. Figure (below) shows some domains. A domain is divided into domains.

Distribution of Name Space

The information contained in the domain name space must be stored. It is inefficient and not reliable to have just one computer store such a huge amount of information. It is inefficient because responding to requests from all over the world places a heavy load on the system.

Hierarchy of Name Servers

To distribute the information among many computers called DNS servers as shown in figure 5.4.4. One way to do this is to divide the whole space into many domains based on the first level. DNS allows domains to be divided further into smaller domains (subdomains). Each server can be responsible (authoritative) for either a large or small domain.

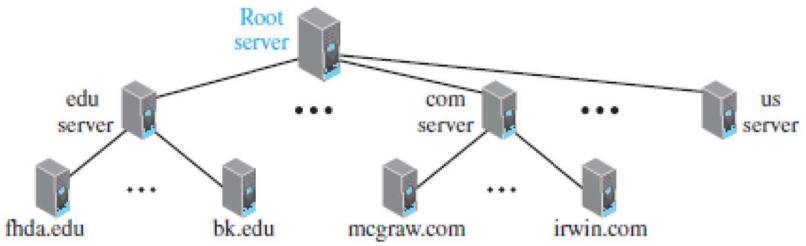

Fig5.4.4: Domain name hierarchy.

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-914]

Zone

If a server accepts responsibility for a domain and does not divide the domain into smaller domains, the "domain" and the "zone" refer to the same thing. The server makes a data base called a zone file and keeps all the information for every node under that domain as shown in figure 5.4.5.The information about the nodes in the subdomains is stored in the servers at the lower levels, with the original server keeping some sort of reference to these lower-level servers.

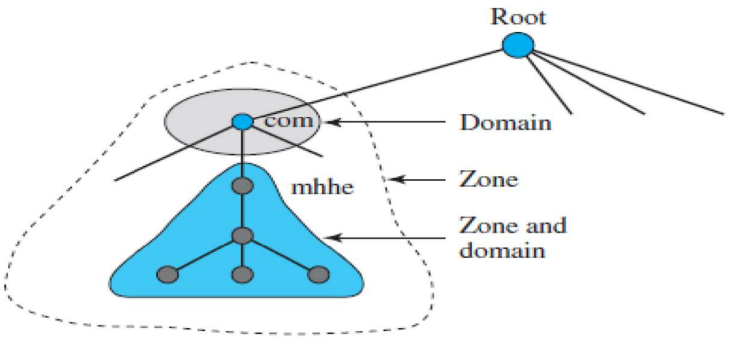

Fig5.4.5: Zone

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-914]

Root Server

If zone consists of the full tree then that zone server is called root server. A root server does not store any information about domains. A primary server is a server that stores a file about the zone for which it is an authority. It is responsible for creating, maintaining, and updating the zone file.

A secondary server is a server that loads all information from the primary server. Secondary server cannot perform any operation on zone file.

DNS in the Internet

In the Internet, the domain name space (tree) was divided into three sections: generic domains, country domains, and the inverse domains.

Generic Domains

The generic domains define registered hosts according to their generic behavior. Each node in the tree defines a domain, which is an index to the domain name space database as shown in figure 5.4.6. Looking at the tree, the first level in the generic domains section allows 14 possible labels.

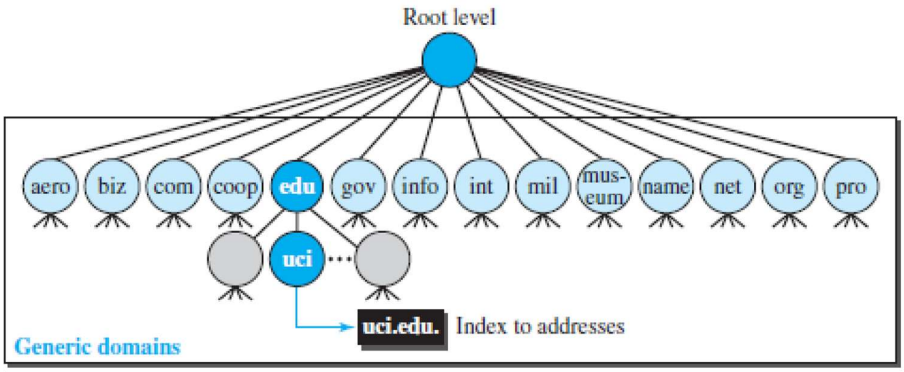

Fig5.4.6: Generic domain.

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-915]

Country Domains

The country domains section uses two-character country abbreviations (e.g., us for United States) as shown in figure 5.4.7. Second labels can be organizational, or they can be more specific national designations.

Example for the country domains section. The address uci.ca.us. can be translated to University of California, Irvine, in the state of California in the United States.

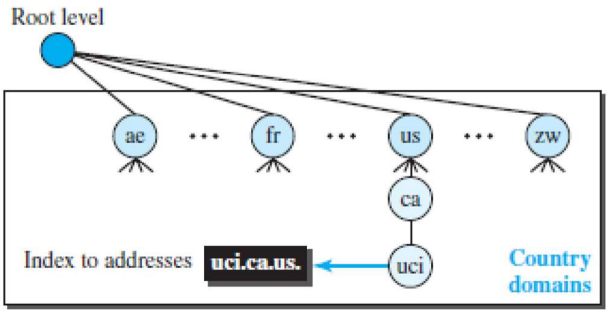

Fig5.4.7: Country domain

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-916]

Inverse domain

Inverse domain is used to find the name of a host when given the IP address.

Resolution: Mapping a name to an address is called name-address resolution.

DNS is designed as a client-server application. A host that needs to map an address to a name or a name to an address calls a DNS client called a resolver.The resolver accesses the closest DNS server with a mapping request. If the server has the information, it satisfies the resolver; otherwise, it either refers the resolver to other servers or asks other servers to provide the information.

Recursive Resolution

A client request complete translation. If the server is authority for the domain name, it checks its database and responds. If the server is not authority, it sends the request to another server and waits for the response.When the query is resolved ,the response travel back until it finally reaches the requesting client. This is called recursive resolution.

DNS Messages

To get information about hosts, DNS uses two types of messages: query and response. Both messages have the same format as shown in Figure 5.4.8.

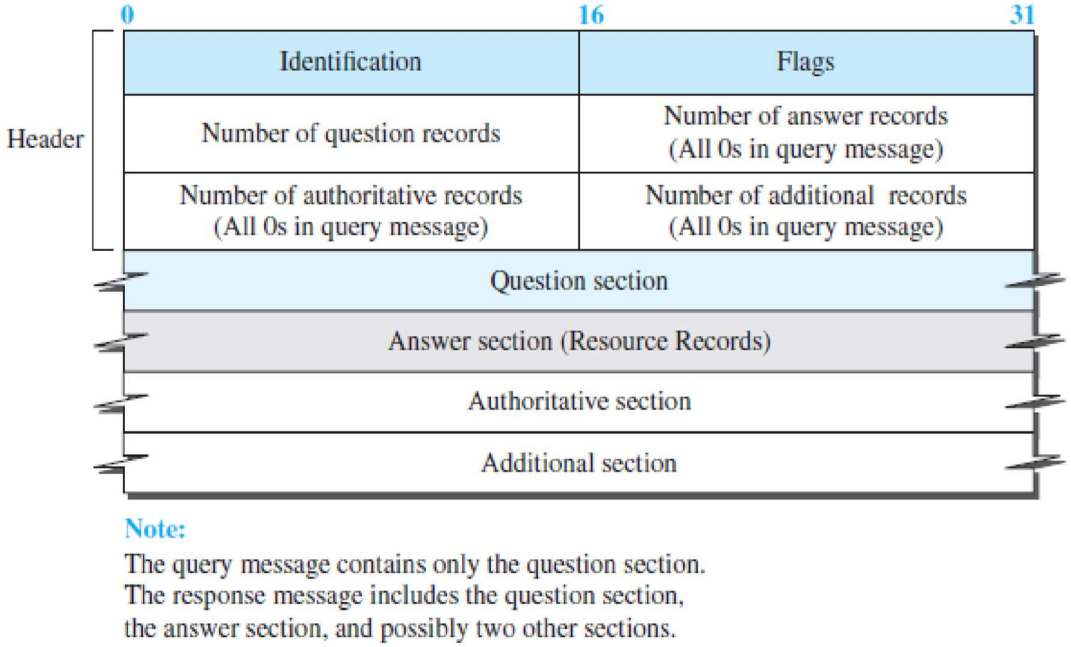

Fig5.4.8: DNS messages.

[Source :"Data Communications and Networking" by Behrouz A. Forouzan,Page-919]

The identification field is used by the client to match the response with the query. The flag field defines whether the message is a query or response. It also includes status of error. The next four fields in the header define the number of each record type in the message. The question section consists of one or more question records. It is present in both query and response messages. The answer section consists of one or more resource records. It is present only in response messages. The authoritative section gives information (domain name) about one or more authoritative servers for the query. The additional information section provides additional information that may help the resolver.

Encapsulation

DNS can use either UDP or TCP. The port used by the server is port 53. UDP is used when the size of the response message is less than 512 bytes because most UDP packages have a 512-byte packet size limit. If the size of the response message is more than 512 bytes, a TCP connection is used.

Security of DNS

Applications such as Web access or e-mail are dependent on the proper operation of DNS. DNS can be attacked in several ways .

1.The attacker may read the response of a DNS server to find the nature or names of sites the user mostly accesses. This type of information can be used to find the user's profile. To prevent this attack, DNS messages need to be confidential .

2. The attacker may intercept the response of a DNS server and change it or create a totally new bogus response to direct the user to the site or domain the attacker wishes the user to access. This type of attack can be prevented using message origin authentication and message integrity.

3**.** The attacker may flood the DNS server to overwhelm it or eventually crash it. This type of attack can be prevented using the provision against denial-of-service attack.

To protect DNS, IETF has devised a technology named DNS Security (DNSSEC) that provides message origin authentication and message integrity using a security service called digital signature.
