Domain Name System

Uniform Resource Locator(URL)

URLs reflect the address of the webpage. Each of their parts has a meaning. The first one is the protocol, most commonly HTTP or HTTPS. The next part is the fully qualified domain name(www.sample.com), which in turn has three parts: the hostname(www), the domain(sample), and the TLD or top-level domain(com). The hostname is the name for the server hosting the webpage, the domain stands for the name of the organization that owns the webpage, and the TLD tells about the type of services provided by that webpage or the company it belongs to, though many companies use TLDs that do not match their company's services. Then comes the directory, which is the name of the folder in the server; there might be more than one directory. Then comes the file name, which indicates the name of the file that contains the webpage. The last part of the URL is the query parameters, which indicate what the user can see on that page. But to connect to any server, we need its IP address. This is where we need the DNS.

DNS

When we want to access a webpage, we type its domain name. A software built into our computers checks whether it already knows the IP address of that webpage in its cache. If it doesn't, it asks the DNS server for it. In most cases, it happens through the ISP. The ISP server the client connects to sends a request to its DNS servers. These are special types of DNS servers called Recursive resolvers. It does not have a database of IPs but maintains the cache of previously resolved domain names. If it has the IP for that request, it passes it back to the client. If the resolver doesn't have it, it contacts one of 13 root DNS servers, which are distributed across many servers worldwide. Though there are 750 servers, you can only access them via 13 addresses, so logically there are only 13 of them. Root name servers store information about 1000+ TLDs, which are looked up by ICANN, whom we should contact to register a new hostname. The root server looks at the top-level domain part of the request and sends the ISP the IP address of the top-level domain server. Then the ISP requests the TLD server for authoritative name servers' IPs. This contains information about thousands of IP addresses matched with domain names for a particular top-level domain. The ISP server chooses an Authoritative Name Server at random and requests the IP address for a specific domain name. After receiving the requested IP address, the ISP's Resolver passes the IP to the ISP server and saves the cache for that request. Then the ISP server passes the IP to the client. The use of DNS has a mass of benefits for the end users. The first one is

that they do not need to remember the IPs of webpages; instead, they can only remember their domain names. If the IPs change, the URL can stay the same. Also, DNS can be used to get the domain name if we have the address. This can be helpful for tracking attacks.

There are a few types of DNS. The first one is the Dynamic DNS(DDNS). DDNS services automatically update their caches when IP addresses change to ensure that the IP addresses returned are always valid. Next comes the Static DNS; most of the DNS will assign the user a static IP address, which never changes, and which uses a static name server. Users have to update the IPs manually.

Moreover, there are three types of DNS queries. In the Recursive query, the DNS client requires a response, either the requested record or a message signaling that the record can’t be found. In an Iterative query, the DNS server is allowed to return the best answer it can. If it does not have a queried DNS that has a match for the request, it will return a referral to the authoritative DNS for a lower level of the domain namespace. Then the client would refer to the referral address. Until an error or a timeout occurs. In the Non-recursive DNS, the resolver client has access to the record it requests.

Moreover, there are three types of DNS queries. In the Recursive query, the DNS client requires a response, either the requested record or a message signaling that the record can’t be found. In an iterative query, the DNS server is allowed to return the best answer it can. If it does not have a queried DNS that has a match for the request, it will return a referral to the authoritative DNS for a lower level of the domain namespace. Then the client would refer to the referral address. Until an error or a timeout occurs. In the Non-recursive DNS, the resolver client has access to the record it requests.