Domain Name System

Complicated technology explained in simple terms

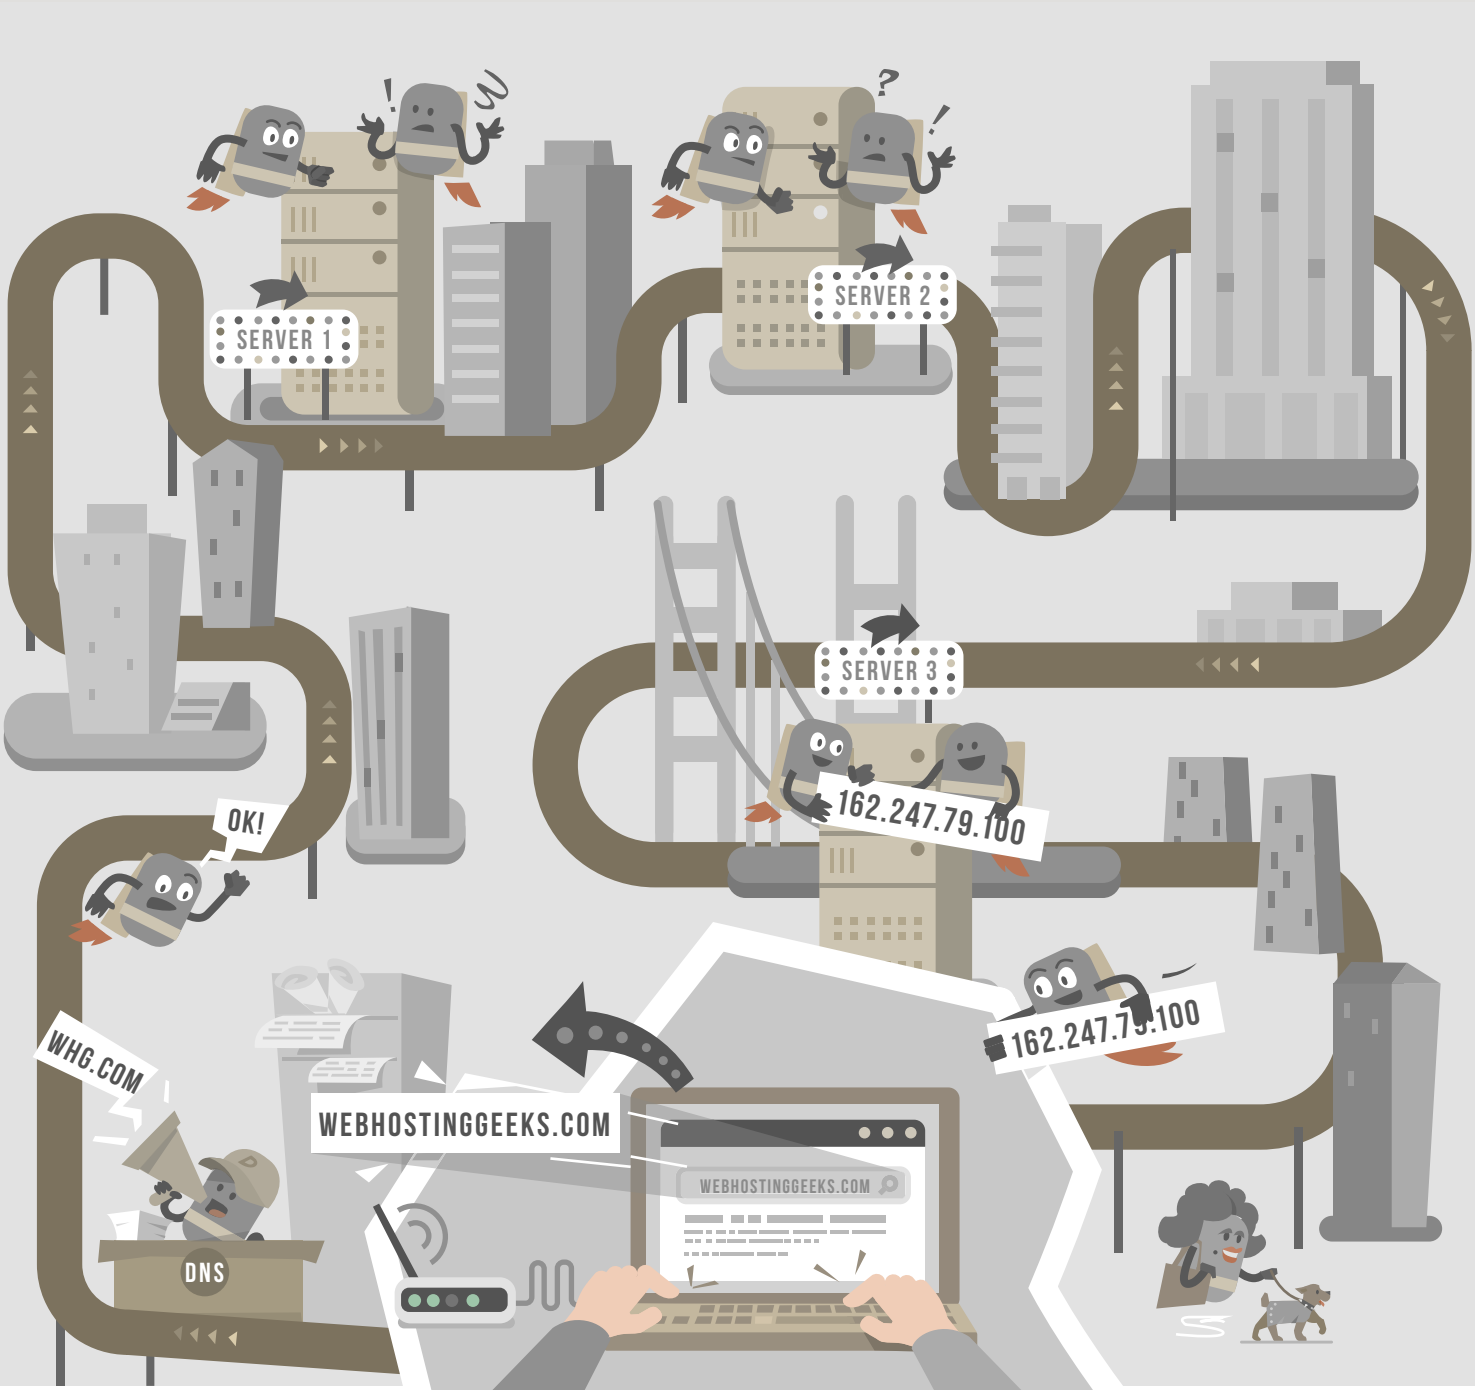

By Web Hosting Geeks

2016

Table of Content

Introduction

Have you ever wondered how the Internet really works? Many people do, from simple web surfing to sharing pictures on social media. In fact, the Internet heavily relies on something called a DNS: a database of network names and IP addresses. These three little letters hold huge weight. Without DNS, the Internet as we know it would simply not exist, and we would be left dealing in ones and zeroes. Without DNS, everyday activities such as shopping, web browsing, research, communications, or downloading would not be possible. That is why experts usually refer to DNS as the Phonebook of the Internet.

So, what is DNS and why is it important? In brief, DNS is a comprehensive translation system used to search the Internet. You might wonder, naturally, what it translates. Well, in the simplest definition, DNS is the term used to describe a system that assigns user-friendly names to unique IP addresses. It translates unfathomable amounts of data into words and phrases in order to provide clear and accurate search results.

While computers communicate using strings of numbers, humans, obviously, do not. DNS translates such number strings into human-friendly phrases. You see, each IP address must be distinct in a network, which allows users to reach a particular website. An IP address could be a set of any four numbers, from 0 to 255, like 162.247.79.100. When you type a domain name into your browser, the DNS system bursts into action, translating the browser name into the IP address associated with the website. Once the website IP address is found, your computer connects with the web host and the requested page is displayed on your computer. While the concept might seem basic, DNS is a cornerstone in how the Internet functions.

Jump to Section:

Download Infographic

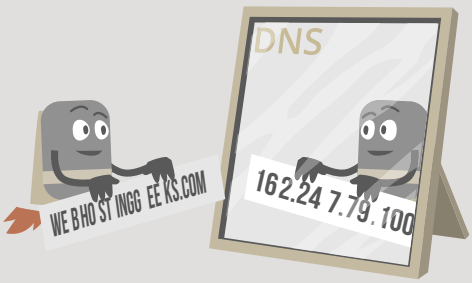

WebHostingGeeks.com = 162.247.79.100

History of DNS

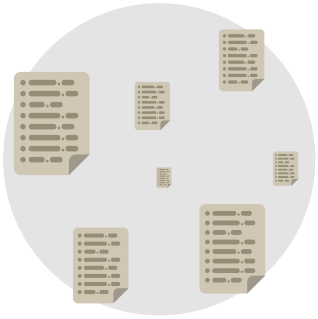

Before 1973

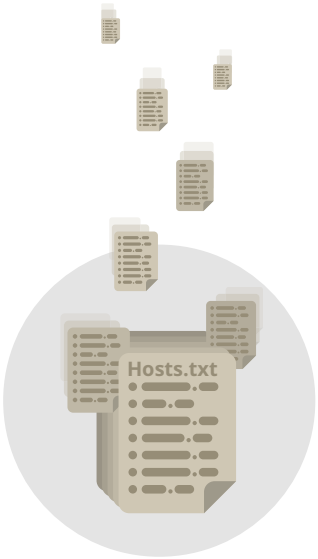

1973

It is imperative for today's Internet users to be aware of the evolution and history of DNS. This system was initially conceptualized to support the growth of communication via email on the ARPANET. Now, it supports the Internet on a global scale, yet effectively understanding its early history and development can be challenging, to say the least. However, due to its pivotal function in how the Interest operates, it is essential to understand DNS' characteristics and components in their entirety.

Initially, working with a few sets of numbers leads to assigning alphabetic hosts to ARPANET. Afterwards, the use of alphabetic names is enhanced since they are easier to remember. The development of host names is useful for the growth of computer programs, and being aware of how they network is important. Since the body of each host name was built by numbers, each site was awarded a host name to provide a guide of network addresses in simple text records.

On the other hand, as early data types began to communicate, Internet mail was re-establishing its attempts to make mail systems benefit from the use of DNS. These attempts included adding application features; however, these proved unsuccessful as it was not yet achievable to hook other applications to DNS roots. In fact, it took nearly a decade to create the first major update to the DNS protocol.

What was the update?

Well, it was the inclusion of a more flexible and dynamic method through the use of Incremental Zone Transfer (IXFR) and NOTIFY, which were both important mechanisms at the time.

However, users soon realized that keeping multiple copies of hosts is inefficient and becomes vulnerable to human error. Therefore, in 1973, a central system was allocated to be the official source of host master files. This system worked well for a decade, but by the 1980s, the disadvantages of a centralized management were becoming obvious, and the need to incentivize interest in the domain concept was growing.

A group of programmers held a meeting in 1982 to come up with a solution to relaying emails. Initially, emails were sent site-to-site and would have to go through several different links. Consequently, sending emails became a tedious task. In a bid to solve this matter, domain names were constructed to give individuals the same address, regardless of the destination of the email.

Hence, there was a need to construct a registered administrative domain, which could be maintained better. After a series of communica-

tions, the concept was developed in November 1983. It was published under the name Domain Names Plan.

First Generation DNS

The most effective way to enhance first generation DNS was by ensuring continuity when multiple servers answered numerous queries simultaneously. This renamed a server as “master”, denoting the other servers as “slave” servers. Practically, each slave followed instructions to keep updated with the master, determining changes in data periodically.

1982

Second Generation DNS

The game changer in the second generation DNS was NOTIFY. This prevented the master from waiting on slaves for feedback. Moreover, delaying problems were solved as well, as previously the master was unable to send notification messages to its respective slaves to prompt them to acquire fresh data. Meanwhile, IXFR highlighted the way data was to be communicated through records, notifying hundreds of changes instead of just the primary. It changed the system of sending central messages, making it so that with each specific change, changes could be sent rather than multiple messages at a time.

1983

Third Generation DNS

The third generation was a turning point for the dynamic changes later adopted, mentioned as RFC 2136. Comparatively, in the first generation, an administrator accessed the master server, did file editing, and then waited till the master reloaded the file before slaves finished with their updates. With this iteration, administrators were no longer required to log into the master, as they could carry out their updates across the network.

Although this sounds like a minor accomplishment, its effect was significant in the long run. Updates now reused messages with their original format for other purposes. Meanwhile, other efforts to define extensions were added, and this modernized the system overall. Additionally, the structural integrity of the protocol increased with the codes being added, and this led to DNS security, which would become the main focus for future modification.

The IETF

PARTICIPATE IN THE INTERNET ENGINEERING TASK FORCE

HTTP://WWW.IETF.ORG

THE INTERNET ENGINEERING TASK FORCE (IETF) IS THE INTERNET'S PREMIER TECHNICAL STANDARDS BODY. IT GATHERS A LARGE OPEN INTERNATIONAL COMMUNITY OF NETWORK DESIGNERS, OPERATORS, VENDORS, AND RESEARCHERS CONCERNED WITH THE EVOLUTION OF THE INTERNET ARCHITECTURE AND THE SMOOTH OPERATION OF THE INTERNET.

The IETF seeks broad participation. The work of the IETF takes place online, largely through email lists, reducing barriers to participation and maximizing contributions from around the world. IETF Working Groups (WGs) are organized by topic into several areas (e.g., routing, transport, security, etc.).

MISSION AND PRINCIPLES

The mission of the IETF is make the Internet work better by producing high quality, relevant technical documents that influence the way people design, use, and manage the Internet. The IETF pursues this mission in adherence to the following principles:

Open Process. Any interested person can participate in the work, know what is being decided, and make his or her voice heard on an issue. Part of this principle is our commitment to making documents, working group mailing lists, attendance lists, and meeting minutes publicly available on the Internet.

Technical Competence. The issues on which the IETF produces its documents are issues where the IETF has the competence needed to speak to them. The IETF is willing to listen to technically competent input from any source. Technical competence also means that IETF output follows sound network engineering principles—this is also often referred to as "engineering quality."

Volunteer Core. IETF participants and leadership are people who come to the IETF because they want to do work that furthers the IETF's mission of "making the Internet work better."

Rough Consensus and Running Code. The IETF makes standards based on the combined engineering judgment of participants and real-world experience in implementing and deploying IETF specifications.

Protocol Ownership. When the IETF takes ownership of a protocol or function, it accepts the responsibility for all aspects of the protocol, even though some aspects may rarely or never be seen on the Internet.

IETF PARTICIPATION AND ORGANIZATION

The IETF is really about the individuals to participants. There is no membership in the IETF—the IETF is made up of volunteers. Anyone may register for and attend any meeting. The closest thing to being an IETF member is being on the IETF or Working Group mailing lists.

The IETF Working Groups are grouped into areas, and managed by Area Directors, or ADs. The ADs are members of the Internet Engineering Steering Group (IESG). Providing architectural oversight is the Internet Architecture Board, (IAB). The IAB also adjudicates appeals when someone complains that the IESG has failed. The IAB and IESG are chartered by the Internet Society (ISOC) for these purposes. The General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB.

While the actual technical work of working groups is accomplished largely through email lists, IETF meetings are held three times a year with the primary goal of supporting IETF Working Groups in getting their tasks done. A secondary goal is to promote a fair amount of mixing between the WGs and the areas. However, the results of Working Group sessions at a face-to-face meeting must gain consensus on the WG mailing list for a decision to be made.

FOR MORE INFORMATION

Much more information is available online at http://www.ietf.org and, especially for those new to the IETF, at http://www.ietf.org/tao.html

THE IETF AND THE INTERNET SOCIETY THE INTERNET SOCIETY IS THE ORGANIZATIONAL HOME OF THE IETF AND SUPPORTS IT WITH FUNDING AND THROUGH PROGRAMMES SUCH AS THE ISOC FELLOWSHIP TO THE IETF.

Internet Society

The Internet Engineering Task Force (IETF) is the name given to a global Internet community that consists of network designers, operators and researchers. It is concerned with developments in the field of Internet. The membership of this community is open to anyone who might be interested. The organization holds meetings three times a year and much of the work is distributed via emails.

Additionally, the technical work is carried out by working groups that are divided into further specific areas, and which come under the command of area directors. Therefore, they are members of the Internet Engineering Steering Group. An area director's job is to provide an overview of all the tasks carried out by their group. They are also responsible for any failure the group might encounter, which the board would have to investigate for an appeal.

The other organization that is involved in the regulation of this system is the Internet Assigned Numbers Authority (IANA). It is the key coordinator for the guidelines of specific Internet projects and their respective standards. The body is governed by the Internet society and acts as the regulator to allocate and coordinate the innumerable Internet protocols. These guidelines are presented in the IETF Standards Process.

For the most part, creating an Internet standard is very basic. It requires a specification, and careful analysis of the information by the Internet community. This is adopted to uphold the standard. However, in reality, the process is much more complicated, since it demands creating high-tech specifications, consulting all the stakeholders, and the need of an established community to evaluate.

RFCs

A Request for Comments (RFC) is a term used to describe an official request from the IETF, which occurs after the committee has constructed rules. Usually, it is done when the stakeholders present a review. Each RFC is of a different nature. While some are informational, others are intended to construct Internet standards. Once the RFC has been finalized, no further comments can be made to alter it. If a change is required, it can be done by suppressing other RFCs.

Interestingly, RFCs were first constructed in 1969 and are currently a part of the official functions of the IETF. They often comprise large portions of the global Internet research community. The first RFC was drafted and its copies distributed among leading IT experts, with the earlier versions of RFCs aimed at encouraging discussion. Conversely, its form of writing did not indicate authority, and the less formal style has become a common form of writing of RFCs.

The University of California was responsible for some of the earlier RFCs, as it became the face of the interface message processors. It also became home to the Augmentation Research Center (ARC) and was one of the first sources of early transmitted RFCs as well as other network information. After the original contract with the United States Government had expired, the Internet society, acting on behalf of the IETF, assumed an editorship role and took the responsibilities of working on the RFC. The IETF working groups, under the IEFT director, handles the publication of RFC documents. In 2008, a new model was proposed to split the task into several different stages. This also included a new role for the RFC series advisory group and, subsequently, it was revised again 2009 with new standards. Up until late 2011, the system has been additionally revised, when Heather Flanagan was appointed as the permanent RFC editor.

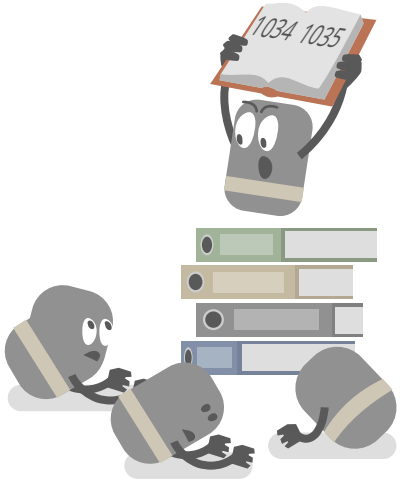

DNS related RFCs

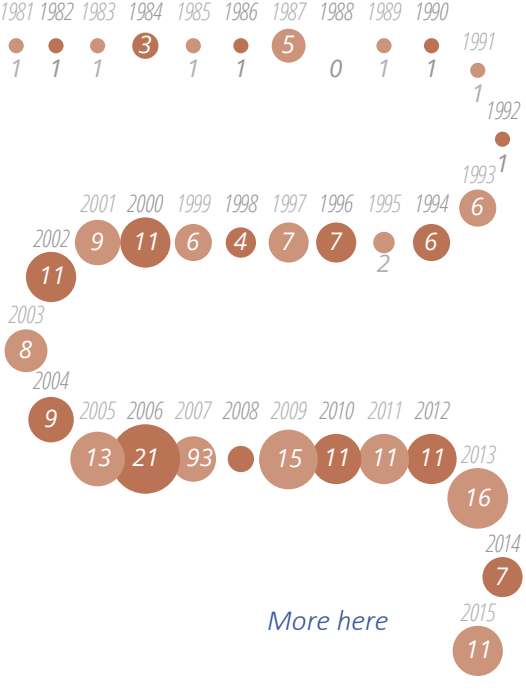

Function

In its simplest form, the DNS is a database that maintains the names of websites, such as webhostinggeeks.com, and links them to particular IP addresses that consist of a number pattern. However, this can be understood as its simplest task. Linking addresses to names is the basic function of DNS, as is it used for a variety of services, apart from host-to-address mapping.

Some of the major functions of DNS include locating IP addresses to specific site names, and then storing this data. This process is also known as “maintaining records”. A second function is to distribute the DNS over a vast network of connections, and a DNS can also store a vast library of records. For many experts, DNS is the term used to define a database and, most importantly, a database that can be easily shared. This is because each server holds only a minor portion of the host name to IP address mapping details.

DNS servers are configured with a special record that informs where the DNS server is located. Due to this process, each DNS server holds a small part of the host to IP mapping address. This collection of host to IP address mapping is also called the namespace. When looking up a name in the DNS system, the user must first check the high-level database, which tells the client how to check the DNS server host. As a next step in the process, it specifies queries the client can address through the hostname given by the DNS server. The process continues until the user finds the correct server that hosts the DNS required.

Additionally, finding the correct DNS and identifying the correct mapping of records stored by the database permits the DNS to maintain records. These record types are useful for several other purposes and may help other applications. For example, the record of the Mail Exchanger provides mail servers with the data needed to pass on sender-to-recipient emails. Another important record used by Microsoft Active Directory is to locate network services accurately.

Although it may seem as if DNS is complicated, its importance lies in the fact that other processes solely rely on it to function.

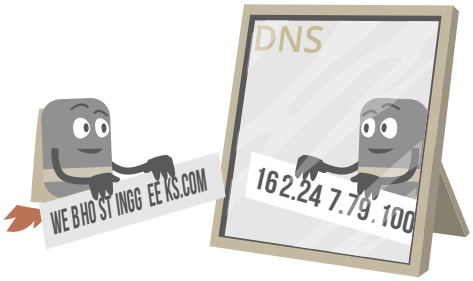

DNS provider market share

The market share of DNS providers is calculated based on the number of domain names that use their service.

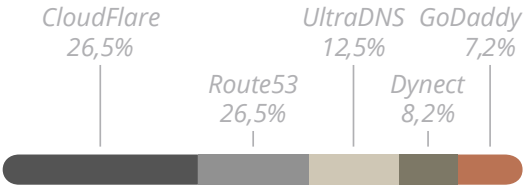

Domain Name Registrars Market Share

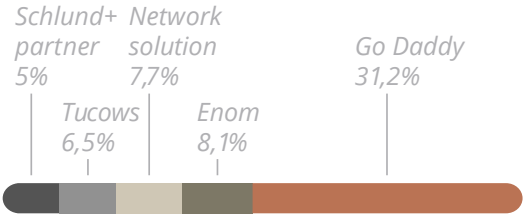

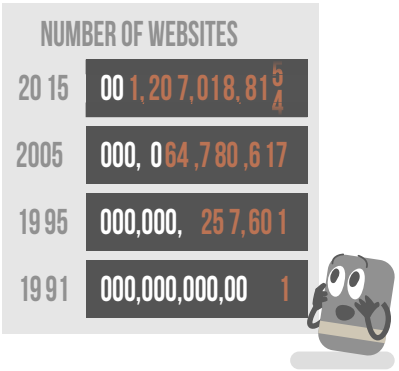

World Wide Web

The WWW relies on DNS for human-friendly navigation. Users can easily access a website by entering the IP address of a particular site or web browser. However, remembering several numbers is not the best way to approach the site. Therefore, it is much easier to remember the DNS name for a website that will present user-friendly names, such as webhostinggeeks.com.

E-Mails

E-mail is the main reason the DNS was developed and is one of the most popular functions of the DNS. Through the web, DNS links the names to IP addresses for various sites, although email servers need a more advanced record than what is required of basic host names. For instance, when an email is sent by a user through Outlook or Gmail, it can either be sent to the recipient at their domain or to another email server that is providing a similar service. If the email specifies an outgoing mail server which is not the target domain, then the user is using a reliable process.

An email address contains two portions: a host and a recipient. For instance, in the address postmaster@domain.tld, postmaster is the recipient and the mail transfer agent is responsible for ensuring that the message reaches the recipient. In actuality, any application that requires the Internet connects two or more hosts, which then shares information or communicates using DNS services.

Other uses of DNS servers include the more recent upgrade in 2008 that supports a zone type called the Stub Zone. This is a zone that contains features and records of resources that are used to identify contained DNS servers. The zone operates in such a way that lets the parent zone be aware of a forceful DNS server for its child zone. Another key feature of the DNS is that it provides integration with other Microsoft networking services. These features include connection with services, such as Windows Internet Name Service and Dynamic Host Configuration Protocol. With its improved ease of administration, DNS now allows a graphical user interface to manage DNS server services, in addition to other applications.

Structure

The DNS architecture is defined by a hierarchical distributed database and a set of protocols. It is a mechanism for updating, replicating information and a schema of the database. DNS was conceptualized in the Internet's early days when it was just a minor network established by the United States Department of Defense. The various host names in DNS were administered by a single host that was located in the central server, and anyone that required the host name downloaded this file. On the other hand, as the Internet grew, the size of this file expanded with the traffic it generated. The need for a new host soon rose, which further featured support for various data types.

Therefore, the Domain Name System was introduced in 1984, becoming the new system depended upon for the Internet. For the DNS, the host name is stored in a database that can be distributed among multiple servers. This will then decrease the pressure on any one server

and will also allow access to the database without any location constraints. DNS is said to support hierarchical names and allows the use of various data, in addition to mapping. Since the data is shared and the size of the host is unlimited, the performance of the DNS does not degrade when more servers are added.

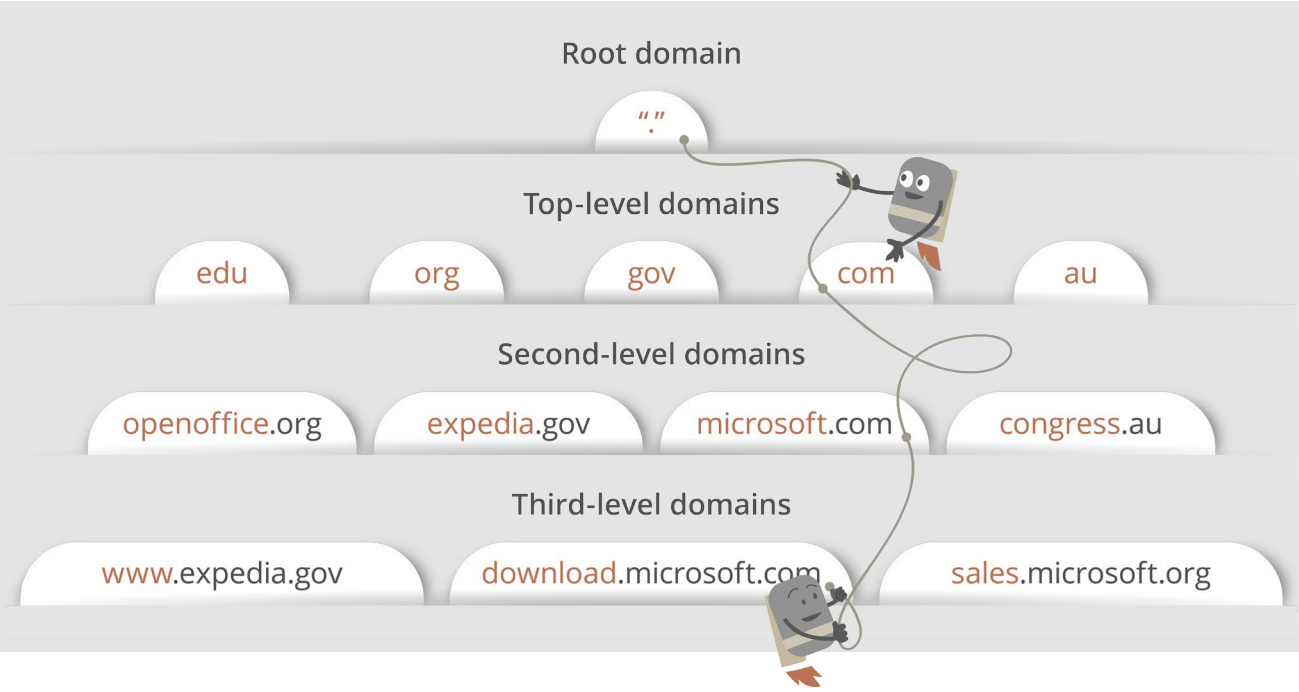

The names in the DNS form a hierarchical tree structure; this is called the domain namespace. The domain name lies at the top of the hierarchy. These names are of individual labels, which are subsequently divided through dots. A fully qualified domain name is unique enough to be easily identified by the host's position in the DNS's structure. This can be done through the hierarchical tree or by specifying the dots that state the path from the host to the root. The namespace is dependent on the concept of a tree that consists of named domains. Each level, branch or leaf can represent a different stage of the hierarchy. Adding on a branch is a stage in which more than one name is used to identify the collection of named resources. A leaf represents a single name that is used only once to mention a specific resource.

Any name that is used in the tree is technically a domain. However, experts have found that there are five main levels for domains. For example, a DNS domain name assigned to Microsoft is a second-level domain. This occurs due to the name having two parts that indicate whether they are located near the root or the top of the tree. Several DNS names have two or more labels, each of which indicate an additional stage in the tree.

Internet domain names are managed by a name registration authority on the Internet, which is responsible for maintaining the profile of

There are five standard categories used to describe the domain names and their functions:

- *Root domain, which can be found at the highest level of the tree, representing an unnamed level. It can sometimes be shown as two empty quotation marks. When used in a DNS, it shows that the name is located at the highest level of the tree, and it is often stated by a trailing period. It points to an exact location on the tree.*
- *Root domains are followed by the top-level domain, a term used to describe a country, a region, or a type of an organization.*
- *Second-level domains are variable-length names that are assigned to appropriate top-level domains, depending on the organization or the geographical location.*

top-level domains (TLDs) that are allocated by countries and regions. These follow international compliant standards and often exist in abbreviations reserved for organizations, as well as for countries.

A DNS database can be divided into several zones, and each zone carries a portion of the DNS database. These contain the resource records of the owner names that are part of the namespace. Zone files are part of the DNS servers, and these can be configured to host zero or multiple zones. Characteristically, each zone is then part of a particular domain name, which is referred to as its root. This zone contains all the information about the names and ends in the zone's domain root name. A name within the zone can also be associated with different zones, which are hosted by a different DNS server. This delegation is a process of giving the responsibility of the DNS namespace to a DNS server owned by a separate entity. This can be another organization or working group.

Zones and root name servers

The root zone is a global list of top domain levels. The information that root zones contain can vary from generic top-level domains to top country code level domains. These include two letter codes, which represent each country, e.g. .se to symbolize Sweden. In addition to this, internationalized top-level domains are incorporated, which indicates that countries are coded and grouped together. Individually, each of these top-level domains contains its own root zone in the numeric addresses of name servers. These aid with the top-level domain's subjects, and the root servers respond to reports when requested about a top-level domain.

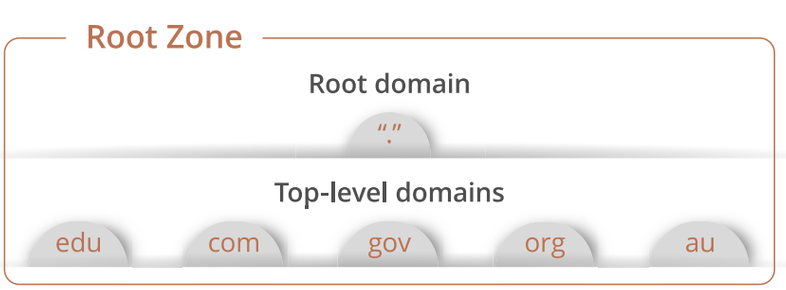

Collectively, each of those top-level domains, contains its own root zone in the numeric addresses of name servers. These aid with the top level domain's subjects, and the root servers respond to reports when requested about a top level domain.

Some organizations that operate these root servers are US Army Research Lab, Internet Systems Consortium, NASA AMES Research Center, US Department of Defense, University of Maryland, Cogent, University of Southern California, Net nod, RIP, Verisign, ICANN, and WIDE. Currently, these are the top 12 organizations using the root servers, and some of these firms have been using the root servers since the invention of the Domain Name System.

In other words, there are over three hundred root servers that have been distributed globally and onto the six most populated organizations. Moreover, each one can be reached through thirteen different IP addresses. Each organization can have one or two IP addresses, such as Verisign, which has two. In addition, any DNS query sent

- *Sub-domains are used to describe any organization that is created and derived from the second-level domain names. These include names that are added to enhance the tree of names in terms of organization.*
- *Hosts, or resources, are the names that are represented in the leaf of the DNS tree of names, and are from a particular resource. Typically, the left-most label of the DNA tree identifies a particular computer on a network.*

13 root name servers worldwide

- *Verisign*
- *USC-ISI*
- *Cogent Communications*
- *University of Maryland*
- *NASA*
- *Internet Systems Consortium*
- *Defense Information Systems Agency*
- *U.S. Army Research Lab*
- *Netnod*
- *Verisign*
- *RIPE NCC*
- *ICANN*
- *WIDE Project*

through these addresses will get a fast response. The number of root servers has increased significantly since the start of the last decade, when there were only 13 worldwide.

The root name server is a name server for the root zone of the DNS. It is known to answer requests directly through the root zone, as well as to record other requests through several authoritative name servers by assigning proper top-level domains, also referred to as TLD. These root servers are essential, as they are used primarily to solve or interpret human decipherable host terms into IP addresses. This is key for communicating between different Internet hosts. The translation is done through a resolver, which answers users' queries directly. Likewise, it tries to identify each and every command word by word.

UDP (User Datagram Protocol) is the combination of several protocols and certain limits in the DNS. The practical size of non-fragmented UDP led to the conclusion that the number of root servers can be limited to thirteen server addresses. However, it should be noted that if any cast is used, then the root server number tends to be higher than predicted.

gTLD and ccTLD

TLD (top-level domain) can be seen whenever one writes the domain name, the web address, or the URL. To be exact, wherever your email address ends up is where the top-level domain resides. TLD is commonly known as the last part of the name of any website, domain, or email address. Some examples of TLD include .com, .biz, .org, .net, and so on.

These TLDs can be categorized into two basic forms, mainly the gTLDs and the ccTLDs. TLDs are taken care of by the Internet Assigned Number Authority, popularly known as IANA. This is the administration that is responsible for the root of the Domain Name System or DNS. The IANA is being operated by the ICANN, which stands for the Internet Corporation for Assigned Names and Numbers. It should be considered that the second part of the TLD is the dot, which helps us separate the TLDs. This is known as the second level domain and is supposed to be registered with a registrar.

The generic top-level domains, or gTLDs, as the name suggests, are generic. Hence, they are not for any specific country. These can be used by anyone who is surfing the Internet. Some of the top-level domains include .com, .org, .net, .gov, and .mil. These are generic top-level domains that can be expanded to 22 gTLDs. Therefore, gTLDs tend to be more restricted, dictating that only a specific group can register and access them, after which they will be eligible. However, they are never bound to a specific country.

On the other hand, ccTLDs denote country code top-level domains. These are more commonly known as the two-letter TLDs, which means they are allotted to countries established customarily on the ISOC 3166 list of country codes.

Largest TLDs by zone size

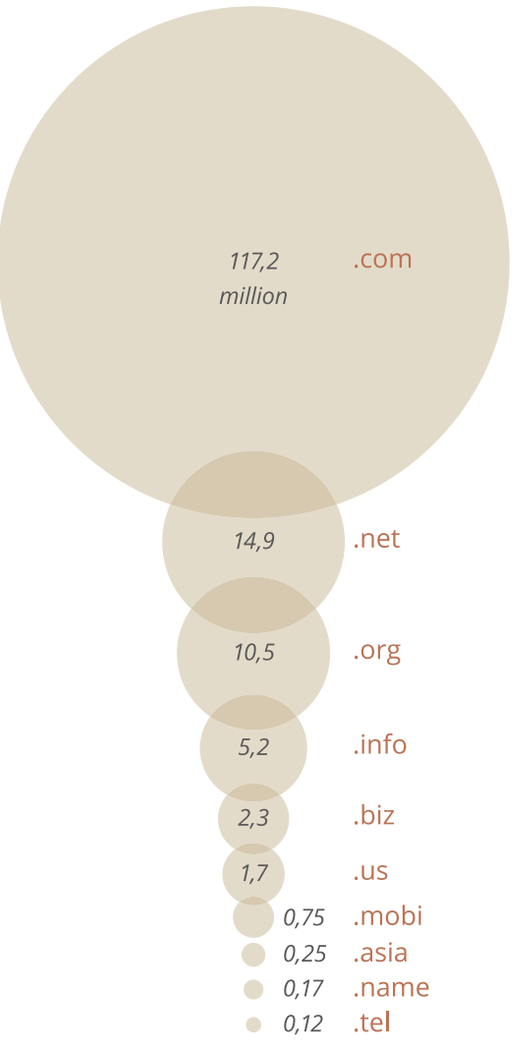

The Current 22 gTLDs

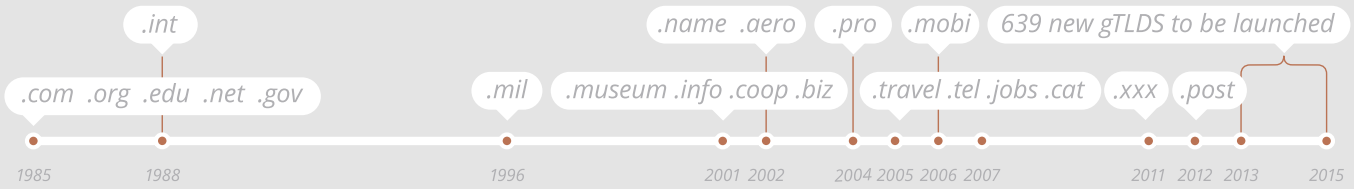

Some countries have opted to function their ccTLD solely for domains that will be used inside their country or within its geographic territory. It should be considered that some countries do not permit individuals to record the second-level domains under the TLD. However, as an alternative, they permit individuals to register third-level domains under one of the wide range of different second-level domains available. Some countries, such as the United Kingdom, are required to register their domains of .uk, such as .co.uk or .org.uk. This will basically change the generic top-level domain to a country code top-level domain.

Total country-code TLD (ccTLD)

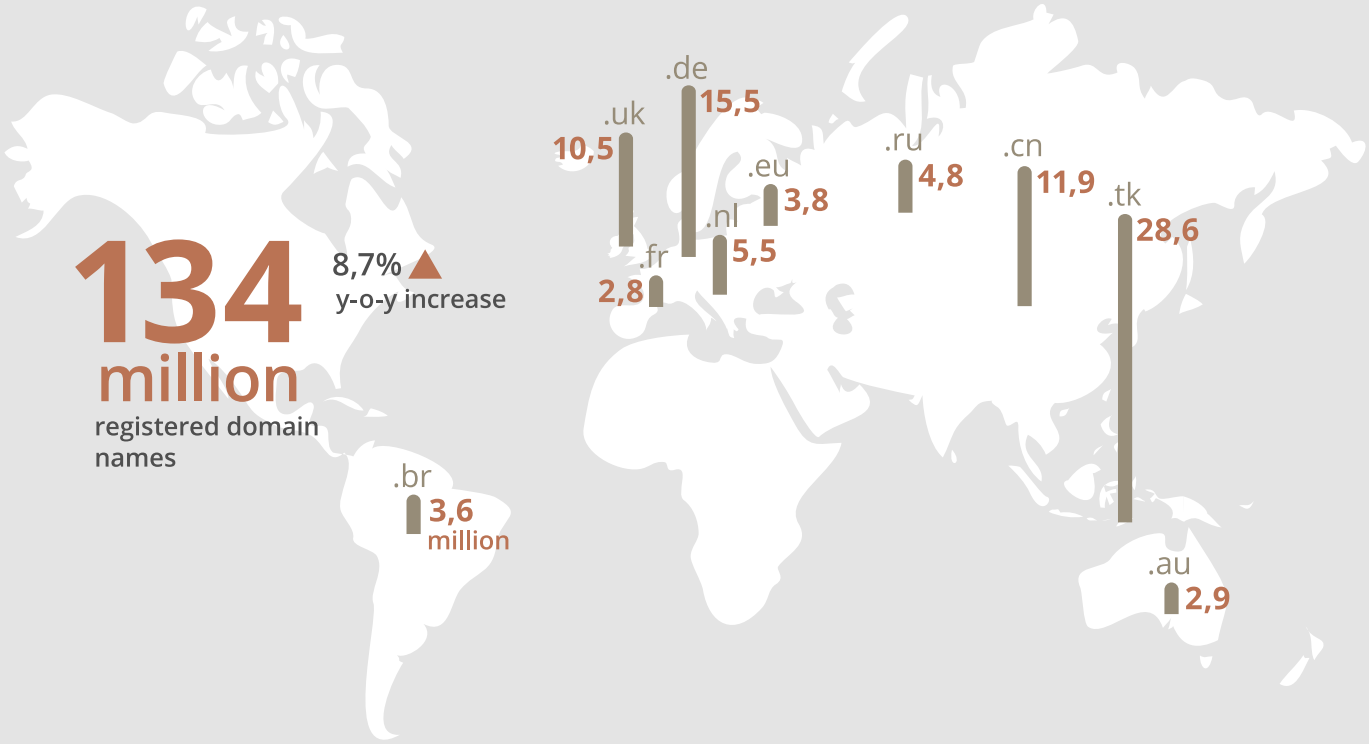

The country code top-level domain is specific to certain countries. Hence, each domain is based on the country extension. Although some have restrictions on who can register, most do not have this formality. For example, .tv, .me, .cc, and .ws are some of the extensions

that are said to be open for registration by the common public. Some of these extensions have also been repurposed for general usage.

New Developments

Ever since the Internet became a phenomenon, ICANN has been constantly asked to approve the support for character sets in the top level of the DNS, other than the 26 letters of the basic Latin alphabet. With the approval of Internationalized Domain Names, TLDs can now include characters other than the traditional ASCII characters (A through Z).

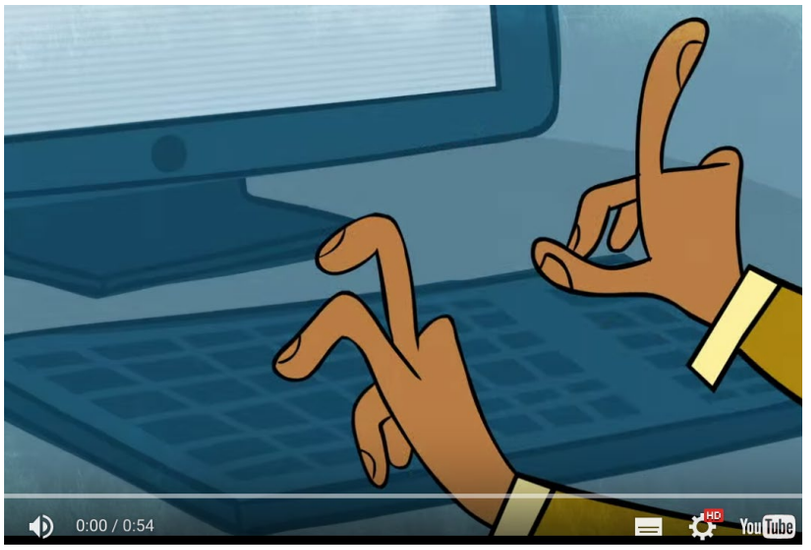

Currently, the Internet is undergoing its largest expansion with more than 1,300 new gTLDs set to be online by the end of 2016, which would represent a major milestone in the development of the Internet namespace. As of January 2016, almost 900 new TLDs are already online to create opportunities for both businesses and consumers.

The New gTLD Program

The new gTLD Program is aimed at adding an unlimited number of new gTLDs to the root zone, the Internet's authoritative database. The first round of applications began on the 12th of January, 2012, and it ended on the 20th of April, 2012. Applicants applied through TLD Application System (TAS) to run the registry for the TLD of their choice. Although the application window should have closed on the 12th of April, a glitch in the TAS system caused a shutdown for a while before it was reopened for another week to allow applicants to complete their applications.

On the "Reveal Day" (June, 13th), there were 1,930 applications: This means it is possible that the first round of the new gTLD program will create 1,409 new TLDs, including:

- Brand TLDs (.samsung, .google, .toyota, .alibaba, .bentley).
- Community TLDs (.catholic, .islam, .art, .music).
- Geographical TLDs (.berlin, .nyc, .tokyo, .amsterdam).

876 New gTLDs and Counting

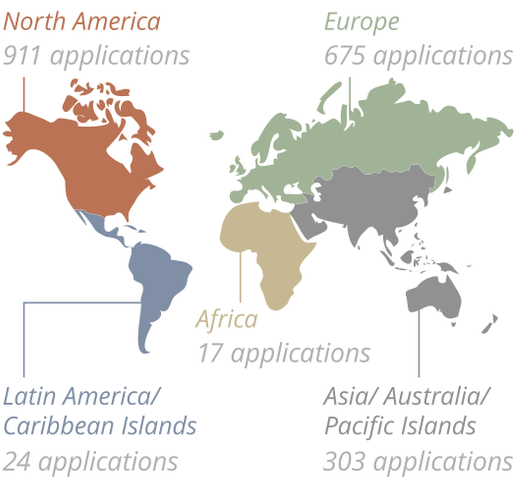

Top 10 new gTLD domains, thousand

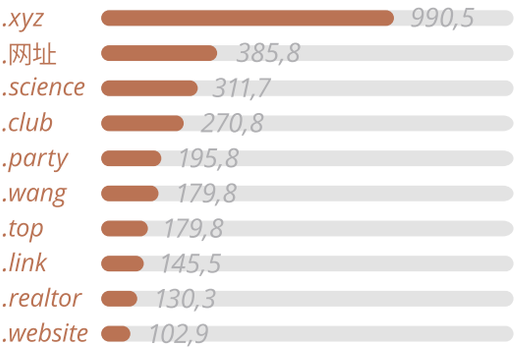

The Internet Protocol (IP)

The Internet Protocol (IP) is an addressing scheme through which computers communicate through a given network. Some of the networks, to get a better connection, combine these IPs with a Transmission Control Protocol (TCP) which is a higher level protocol. This helps create a virtual connection between the endpoint and the source. To understand this concept better, the IP can be compared to a postal system where a labelled package is dropped into the system, which helps you connect the sender to the receiver. In other words, IP is just the connection that is formed between the two hosts.

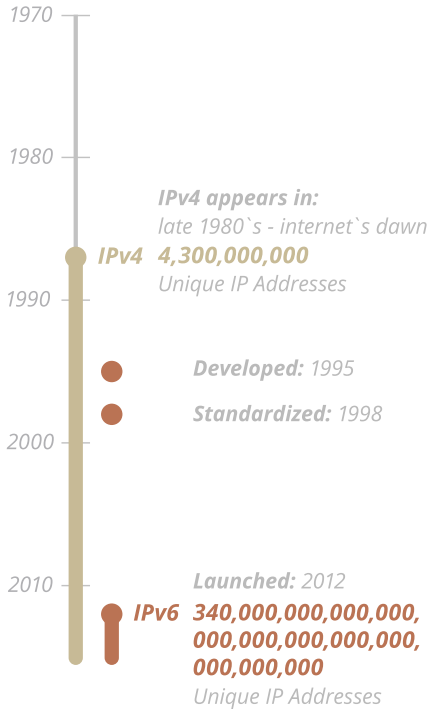

IPv4 vs IPv6

IPv4 is simply the fourth version of the Internet Protocol. The main purpose of this is to recognize the devices in the addressing system that are passing through the network. In essence, it has been designed to perform as a link would in an interconnected system.

IPv4 is one of the most common versions of IP used today to connect devices over the Internet. This version uses a 32-bit address scheme and allows over four billion addresses. However, because of the growth of the Internet and the requirement of having an address on every device, remaining IPv4 addresses will eventually run out.

The newest version of IP is IPv6. Also known as IPng, which stands for Internet Protocol next generation, it has effectively replaced IPv4. This successor is designed in such a way that the Internet and IPv6 will, eventually, go hand-in-hand, in terms of the total amount of data that is being transferred and the amount of hosts that are being connected. However, it should be noted that IPv4 and IPv6 will coexist together for at least a few years.

The next generation of Internet Protocol, IPv6, has been in the development phase since 1990. The main reason behind its creation was the concern regarding the gap between the demand and the supply of IP addresses. However, many people fear that the transition from IPv4 to IPv6 will not be easy, partly due to the uncertainty surrounding the new technology.

The main difference between the IPv4 and the IPv6 is that the IP addresses are different. The IP address is a set of binary numbers which are different for both versions. The IPv4 is written in four numbers, which are separated by periods in the 32-bit address, and each of the numbers can be anything starting from zero to 255. On the other hand, IPv6 is a 128-bit IP address, which means it is written in hexadecimal and separated by colons, rather than dots. This makes the entire procedure easier to use and implement.

An IPv4 address

(dotted-decimal notation)

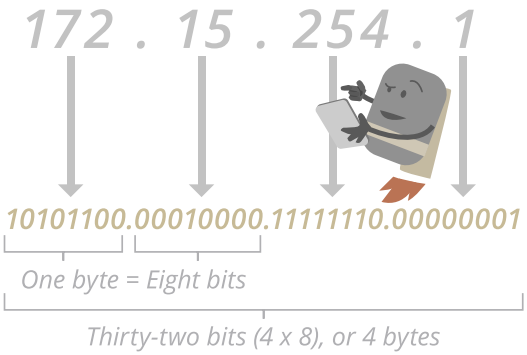

IPv4 was basically used to transfer data from one device to the other. As mentioned earlier, every device, such as a PC, Mac, or even smartphone, will have its own address and is assigned a unique numerical IP address. These are vital; without an IP a device would not be able to communicate or transfer any data.

Although the function of IPv6 is largely the same as that of the IPv4, there are still significant differences. IPv4 utilizes 32-bits while IPv6 is 128-bit. The former means that IPv4 can support up to 2^32 IP addresses, which totals up to 4.29 billion. While this may seem like a huge number, the number of devices requiring an IP has grown exponentially over the past 20 years. In short, we are running out of available addresses. This is where IPv6 steps in. The IPv6 can support over 2^128 addresses, which equates to a significantly higher amount. This will keep the Internet operational for centuries to come. The problem, however, lies with the switch between the two. Although the progress started over a decade ago, only a small fraction of devices have switched over to IPv6. In conclusion, both the IPv4 and IPv6 run parallel to each other due to exchanging data requiring special gateways, but this in turn is slowing down the process.

An IPv6 address (in hexadecimal)

Name Servers

The Domain Name System (DNS) is simply a server-based software designed to match and connect easy-to-read web addresses to officially registered numerical IP addresses. DNS uses a network of servers to carry out these matchups. Of course, you can simply enter the IP address of a webpage into the browser's address bar. However, DNS was created for a user-friendly Internet, wherein websites could be identified by recognizable names.

Managing the entire directory of the Internet can get slightly complicated, due to the billions of daily requests. This is made simpler by the use of specific Internet protocols mentioned in the last section, for instance, the IPv4 and IPv6.

It is important to understand that a DNS adds an additional server process, increasing webpage load times. Thankfully, this does not happen every single time you visit a website. Instead, computers cache DNS results. Once a computer learns that a certain domain name is translated into a specific IP address, it saves that information for a certain period of time.

Resolving a hostname IP address query will provide a more in-depth understanding of the minute steps involved in processing domain name resolutions, involved various types of servers. We can start by identifying important terms and then consider the mechanisms behind a domain name resolution.

As you may recall from prior reading, the management of the DNS is broken down into regions known as the DNS zones. The DNS root zone includes 13 clusters of root servers that are the authoritative, or the go-to servers for queries of TLDs.

Recursive vs Authoritative

Recursive DNS name servers are responsible for providing the proper IP address of the intended domain name to the requesting host. Think of it as a search engine which searches other pages; it is one that responds to each query by asking other name servers for the answer. When you type a website name into your browser, such as webhostinggeeks.com, your computer will then make a request to a recursive DNS server to find the correct IP address associated with the requested website. From there, the recursive server will check to see if it has any DNS records cached for the domain you are trying to reach. If not, the recursive server then queries the root DNS server for the TLD of the domain.

The purpose of authoritative DNS servers is to respond to recursive DNS servers, providing answers with the IP "mapping" of the requested website. Their responses contain all the essential DNS information for each website, such as corresponding IP addresses, a list of mail servers, and other necessary DNS records.

This entire resolver operation is carried out in a span of nanoseconds. DNS caches store DNS resolutions for a fixed period of time, known as time-to-live (TTL). Such DNS caches are usually maintained by an ISP. However, home routers have similar built-in DNS caches that improve overall speed and network efficiency.

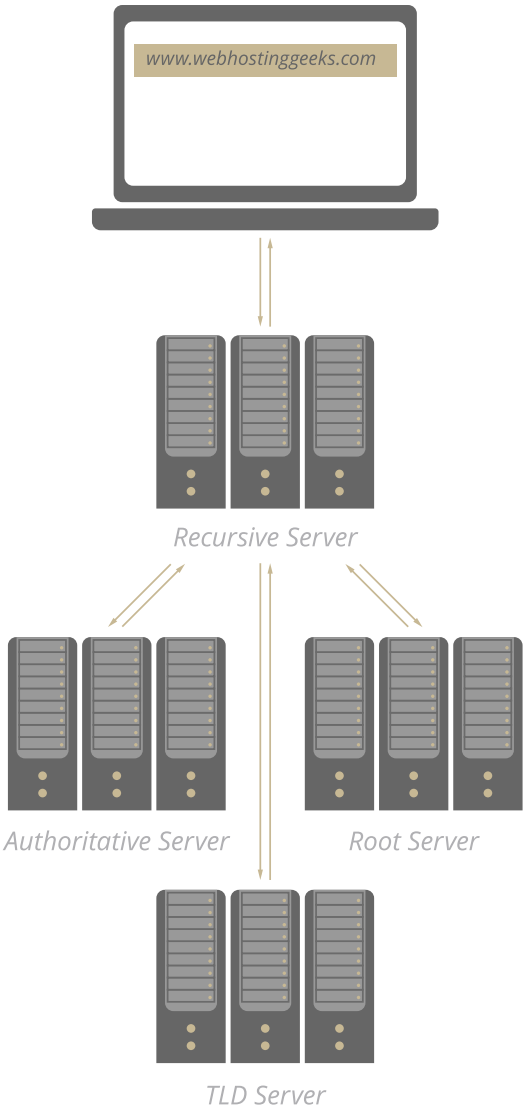

Primary vs Secondary

Every domain name has at least two name servers, provided by the hosting provider in order to get a website online.

A primary DNS server is in charge of perusing information related to the domain zone from a record that is stored on the web server of a hosting account. The primary server is additionally in charge of corresponding with the secondary DNS server, which is known as a zone exchange or zone transfer. Every domain name is given its DNS records for redundancy, and to make the recovery procedure of server administration easier. There's a possibility that a primary server already has the zone data for a particular domain. In such a case, the data would not have to be replicated as both primary and secondary server share zone data without any interruption. In simple terms, when a request is issued to a domain name it goes through the primary DNS server first to reach the website's server.

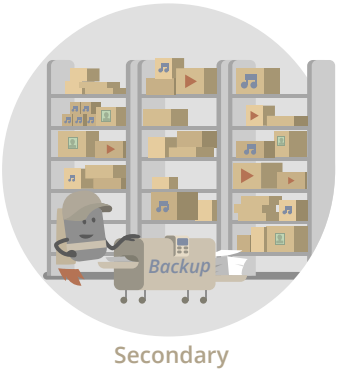

Secondary DNS servers act as a backup when primary servers fails to direct a user to the web hosting server. A secondary DNS server, also known as a slave server, is in charge of acquiring zone data from the primary DNS server quickly. Every time a secondary DNS server performs a function, it gets data from the primary DNS server. It ought to be noted that a secondary DNS server does not always have to get data from a primary DNS server, as secondary servers can also be made master servers. In general, secondary servers are equally as essential as primary DNS servers since they offer redundancy, as well as alleviate the collective resource load put on the primary DNS server.

Relationship between primary and secondary DNS:

- Primary DNS servers hold the master copy of the zone record while secondary DNS servers usually obtain data from the primary DNS.
- Secondary DNS provides redundancy to primary DNS servers, improving the security level.

Record Types

Record types are part of a bigger structure known as DNS zones. DNS zones are configurations implemented on Domain Name Servers. Continuing our last example, when a specific authoritative DNS server directs a recursive server for a specific TLD, it directs it to a certain TLD zone, a form of a hierarchical layout of several domains and/or subdomains.

If we were to view DNS zones as a building, then DNS record types should be considered as the individual rooms. A DNS record is a single data point which provides directions to DNS zones on how to process incoming queries. For example, the DNS zone for google.com can have multiple DNS records, such as www.google.com, mail.google.com, or maps.google.com.

A DNS record has three details attached: a record name, a record data or record type, and time-to-live (TTL). Record data/value is basically the instructor of various operations, while TTL is import, as it specifies how quickly a record is refreshed.

Specifically, TTL is a fundamental part of DNS records that sets the time delay before a DNS record is refreshed by defining the cache timeframe of DNS records in seconds. The TTL process starts out with a name server inquiry for a DNS record, after which the name server confirms to see if it has provided a cached DNS record within the TTL. If it has, it will do so again for the new query. If it has not, it will request the DNS zone for the record again and cache that for the period of the record TTL.

Another aspect of TTL to consider is that any record value changes will only start to take effect once the TTL expires. Until that time, the record will remain stale with older data or value.

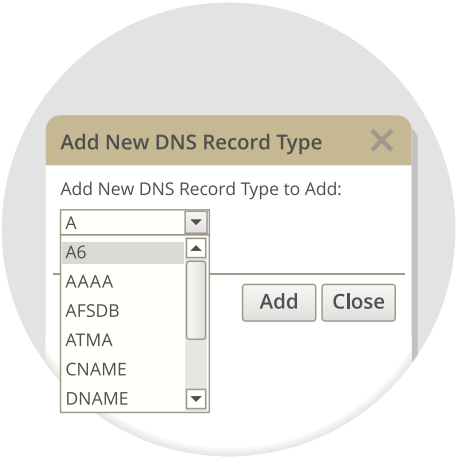

The most common record types are the following:

- *A Records - represents IP. Guides on where the Record name is pointing and how to get there.*
- *CNAMES - represents domain. Guides on where the Record name is pointing and how to get there MX Records - specifies email servers to route emails.*
- *TXT Records - stores text-based data and accesses it when needed.*
- *SPF Records - permits specific servers to use its domain name for email.*

Default NS Records

1and1	ns-us.1and1-dns.com	ns-us.1and1-dns.us	reference
Arvixe	ns1.arvixe.com	ns2.arvixe.com	reference
Bluehost	ns1.blueHost.com	ns2.blueHost.com	reference
Dreamhost	ns1.dreamhost.com	ns3.dreamhost.com	reference
Fatcow	ns1.fatcow.com	ns2.fatcow.com	reference
GoDaddy	ns13.domaincontrol.com	ns26.domaincontrol.com	reference
GreenGeeks	ns1.greengeeks.com	ns2.greengeeks.com	reference
HostGator	ns1.hostgator.com	ns2.hostgator.com	reference
HostMetro	ns1.hostmetro.com	ns2.hostmetro.com	reference
Hostmonster	ns1.hostmonster.com	ns2.hostmonster.com	reference
HostPapa	ns1.hostpapa.com	ns2.hostpapa.com	reference
Inmotion Hosting	ns.inmotionhosting.com	ns2.inmotionhosting.com	reference
iPage	ns1.ipage.com	ns2.ipage.com	reference
IX Webhosting	ns19.ixwebhosting.com	ns20.ixwebhosting.com	reference
JustHost	ns1.Justhost.com	ns2.Justhost.com	reference
LiquidWeb	ns.liquidweb.com	ns1.liquidweb.com	reference
Lunarpages	ns1.lunarpages.com	ns9.lunarmania.com	reference
MediaTemple	ns1.mediatemple.net	ns2.mediatemple.net	reference
MyHosting	ns.myhosting.com	ns2.myhosting.com	reference
Netfirms	ns1.Netfirmsonline.com	ns2.Netfirmsonline.com	reference
SiteGround	ns1.m21.siteground.biz	ns2.m21.siteground.biz	reference
UK2	ultra103.uk2.net	ultra104.uk2.net	reference
WebHostingBuzz	ns1.whbdns.com	ns101.whbdns.com	reference
WebHostingHub	ns1.webhostinghub.com	ns2.webhostinghub.com	reference
WebHostingPad	ns1.webhostingpad.com	ns2.webhostingpad.com	reference

Common Errors

These are common errors faced while managing networks with DNS.

TCP/IP Configuration Points to Public DNS Servers

Many people face this particular error. TCP/IP settings are part of a network's interface, which includes a list of DNS servers used by it. If the settings for a specific computer are of an IP address belonging to a public DNS server, as an ISP, then the TCP/IP resolver will not be able to view Service Locator (SRV) records. These advertise domain controller services: Global Catalog, LDAP, and Kerberos. Without these, authentication problems will arise that complicate the operations of DNS.

Fixing this problem is a case of entering the correct DNS entries in TCP/IP settings at the DC, and then populate the zone with SRV records by stopping and starting the Net Logon service. Additional changes to the DHCP scope option should also be made, as well as manually correcting DNS entries for any statically mapped servers and desktops.

Erroneous DNS Suffix Handling

DNS servers require each query to specify a target domain in order to select the proper zone file. Some DNS resolvers accept the regular domain name from the user, and then append a suffix to form Fully Qualified Domain Names (FQDN). This can then be sent to the DNS server. Usually, it is done by the resolver, as it can obtain the DNS suffix from the active domain name, among others.

The domain to which a certain desktop or server belongs has a DNS name, as well as a simple host name. This can be found in the Properties of the local system, also known as the primary suffix, as per the TCP/IP Settings window. If that query fails and the "Append Parent Suffixes" option has been checked, the resolver strips the leftmost element from the primary suffix before trying again. As an example, for www.google.com, the resolver first appends www.google.com then google.com.

No Registered Records

This is a common error that usually originates from an incomplete website setup. To prevent this error, clients must adopt a holistic approach with a suite of offerings for database management, centralized domain, easy integration options, and a full range of diagnostics and auditing for verification and data integrity.

SNAME Error

SNAME errors are other common DNS errors. They occur due to domain names not having a valid IP address, and arise to advise users to validate all IP addresses before settings are finalized.

DNS Hijacking

Typically, this aspect of malware programs hijacks traffic and redirects it to another malicious site. DNS hijacking is achieved through programs containing viruses. These end up changing the designated DNS server to a malicious DNS server, often occurring when the user visits websites run by scammers.

Such issues can be rectified by running regular antivirus software checks and upgrades. Users should be on the lookout for error messages pertaining to websites with encryptions certificates (HTTPS), such as bank websites. If a user is visiting a bank's website but is seeing "invalid certificate" messages for the website, the user is mostly likely a victim of DNS hijacking where culprits have successfully misguided the user to a fake website, masquerading as the user's bank website to gain login credentials.

Security Issues and Hijacking

A very common problem with a name server and the DNS resolver operations is that it can be susceptible to security issues. The most common type of security issue is DNS hijacking.

In practice, a user attempting to visit www.google.com enters the URL into his webpage address bar and a DNS resolution inquiry is launched. The ISP name servers respond back with the correct IP address. However, in cases of DNS hijacking, a pre-installed malicious software goes into action and directs the user to a malicious DNS server operated by the scammers, prompting the malicious DNS server to reply back with their own, alternative IP address.

Although, the user would see google.com in their web browser's address bar, they may actually be on an entirely different site. It may even look like the original, but the facsimile is intended solely to gain the user's credentials. This also opens up access to other sources of personal information of users and their devices.

To rectify these issues, users are advised to run regular antivirus software checks, as well as antivirus software upgrades. Users should also be on the lookout for error messages pertaining to websites with encryptions certificates (HTTPS), such as bank websites. If a user believes he is visiting a bank's website but is seeing "invalid certificate" messages for the website, the user is mostly likely a victim of DNS hijacking, whereby culprits have successfully guided the user to a fake website masquerading as the user's bank website to gain their login information.

Another alternative precaution is to use third party DNS servers. As you learned earlier, users by default use their ISP's DNS servers for domain name resolutions. However, they can also use third party DNS servers, most popular of which is OpenDNS. Such servers are excellent for providing extra layers of protection and improved speed via use of filter.

Improved speed is accomplished as more servers are utilized by the third party. This creates a higher probability of accessing DNS servers in closer proximity to the user, thereby reducing hops and latency of domain name resolution. Obviously, such gains in efficiency will depend on how far the third party servers are to the user, relative to his or her current ISP DNS servers.

Additionally, the use of filtering by third party DNS server providers has other advantages. For example, third party servers have parental controls, allowing the filtering of pornographic material. After which, the third party DNS server will return a "blocked" message for websites containing pornographic material.

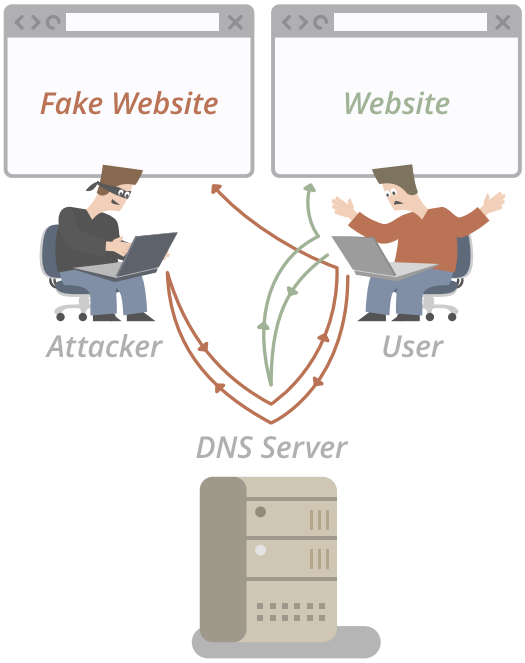

Top 10 DNS attacks

TCP/ UDP/ ICMP floods: *Flood victim's network with large amounts of traffic*

DNS cache poisoning: *Corruption of a DNS cache database with a rogue address*

DNS tunneling: *Tunneling of another protocol through DNS for data ex-filtration*

DNS based exploits: *Exploit vulnerabilities in DNS software*

DNS reflection/DrDos: *Use third party DNS servers to propagate DDoS attack*

DNS amplification: *Use amplification in DNS reply to flood victim*

Protocol anomalies: *Malformed DNS packets causing server to crash*

DNS hijacking: *Subverting resolution of DNS queries to point to rogue DNS server*

Reconnaissance: *Probe to get information on network environment before launching attack*

Fragmentation: *Traffic with lost of small out of order fragments*

Common Terms and Meanings

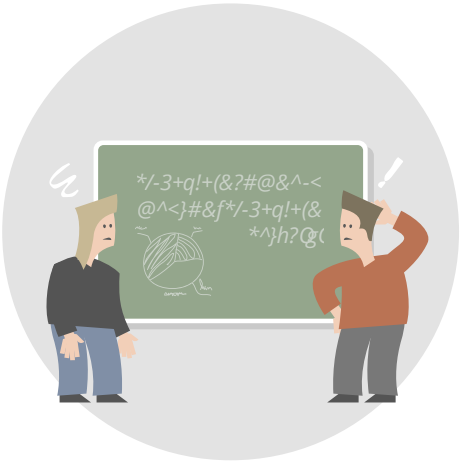

Below is a guideline of the most common terms and meanings associated with DNS that you can use as a refresher:

A Record - A single data point based on a certain type that directs DNS zones on how to process incoming queries. For example, the DNS zone can have multiple DNS records, such as www.google.com, mail.google.com, or maps.google.com.

Authoritative - The purpose of the Authoritative DNS server is to provide an answer for the recursive resolver, also known as the recursive server. An authoritative DNS server has the mapping of the IP addresses of requested websites.

CNAME - This is a record that can be used as an alias for a hostname. For example, maps.google.com is a CNAME for the host name google.com.

Delegation - This is the process of assigning responsibility of handling certain domains and sub-domains to a name server.

DNS Query - An inquiry from a user to translate, or resolve, a domain name for an IP address.

DNS Zone - A specified section of the DNS namespace that has been broken up into sections, or zones; for the better management of DNS queries in the DNS zone. Each DNS Zone has specific DNS records that include information mapped to that zone about a domain.

IP Address - This is a specific identifier for a computer system or device that enables computers on the Internet to locate and communicate together. It is exactly as the name implies: an address.

MX Server - The MX Server is the server responsible for handling emails for a specific domain. MX stands for Mail Exchange.

Name Server - The name server forms part of the domain name system that has been set up to answer queries regarding domains. This is a DNS server designated to handle DNS queries and/or provide additional information about the domain.

Recursive Query - This identifies requests from a user for information pertaining to domain names to identify IP addresses.

Resolver - The recursive resolver, or recursive server, sends out requests for information and refers back to itself until the requested information is acquired.

Root - These are name servers that are known among all name servers. They forward the ISP's recursive DNS server to an authoritative DNS server, which is responsible for handling that specific domain.

This ultimately provides the corresponding IP address the website being sought.

Start of Authority Record (SoA) - The start of authority record provides details of the basic properties of a zone, and is the first resource record in the system for that zone. Some of the details it includes are the host name, email of the person responsible for the domain, the zone serial number, and TTL.

Top-level domain (TLD) - TLD is the highest level of the DNS hierarchy, examples of which include .com, .org, .net, etc.

Time-to-Live (TTL) - TTL is a fundamental part of DNS records, as it sets the time lag before a DNS record is refreshed. It does this by defining the cache timeframe in seconds.

Useful Tools

Name	URL	Description
DNSCAP	A DNS traffic capture utility that provides DNS-specific functionality beyond that of tcpdump.	www.dns-oarc.net
DSC	DNS Stats Collector, a tool that creates statistical information for DNS traffic.	www.dns-oarc.net
fpdns	A tool used to fingerprint DNS resolvers.	www.dns-oarc.net
dnstop	A tool that builds statistics based on DNS traffic seen on the network.	dns.measurement-factory.com
dnsstat	A DNS-specific tool that builds statistics based on DNS traffic seen on the network.	www.caida.org
dig	A powerful command line utility for debugging and troubleshooting DNS.	www.isc.org
host	A DNS lookup command line utility.	www.isc.org
nslookup	A command line DNS lookup utility.	www.webhostinggeeks.com
dnsdump	A tool that will monitor and display DNS messages seen on the network.	dns.measurement-factory.com
dnsmap	A tool that collects all available information for a sub-domain.	code.google.com
dnshealth	DNS configuration analysis.	webhostinggeeks.com
TXDNS	A multithreaded Win32 tool used primarily to send many DNS queries at a time for testing DNS servers.	www.txdns.net
Open Resolver	A web-based tool that will check DNS servers to determine if they support recursion from the Internet.	dns.measurement-factory.com
dnsenum	A tool that attempts to collect all possible information available for a domain.	code.google.com

Sources

1. Kahn, R., “Communications Principles for Operating Systems”, Internal BBN Memorandum, 1972.

2. Dunlap, K. J., Bloom, J. M., “Experiences Implementing BIND, A Distributed Name Server for the DARPA Internet”, Proceedings USENIX Summer Conference, Atlanta, Georgia, 1986.

3. Dunlap, K. J., “Name Server Operations Guide for BIND”, Unix System Manager’s Manual, SMM-11. 4.3 Berkeley Software Distribution, Virtual VAX-11 Version. University of California, 1986.

4. Quarterman, S.J., Hoskins,C.J., “Notable computer networks, Communications of the ACM”, v.29 n.10, pp.932-971, 1986.

5. Murray, A.D., “Internet Domain Names: The Trade Mark Challenge”, International Journal of Law and Information Technology 6(3): 285-312, 1988.

6. Middleton, G., “Australia: Intellectual Property-Domain Names”, Computer and Telecommunications Law Review 11, 2005.

7. Middleton, G., “Electronic Commerce-Domain Names”, Computer and Telecommunications Law Review 9, 2003.

8. Zhou, Tao, “Web Server Load Balancers”. Windows & .NET Magazine, 2000.

9. Rose, S. and Nakassis, A. “Minimizing Information Leakage in the DNS” IEEE Network Magazine vol. 22 no. 2, 2008.

10. Rastegari, S., Saripan M. I., and M. Rasid, F. A., “Detection of Denial of Service Attacks against Domain Name System Using Neural Networks”, IJCSI, Volume 6, Issue 1, 2009.

11. Liu, C., Albitz, P., “DNS and BIND”, 5th Edition, 2006.

RFCs

Domain Name System Operations Working Group
http://tools.ietf.org/wg/dnsop

DNS Extensions Working Group
http://tools.ietf.org/wg/dnsext/

Measures for making DNS more resilient against forged answers
http://tools.ietf.org/html/draft-ietf-dnsext-forgery-resilience

Use of Bit 0x20 in DNS labels to Improve Transaction Identity
http://tools.ietf.org/html/draft-vixie-dnsext-dns0x20

Domain names - Concepts and Facilities
http://tools.ietf.org/html/rfc882

Domain names - Implementation and Specification
http://tools.ietf.org/html/rfc883

Domain System Changes and Observations
http://tools.ietf.org/html/rfc973

Domain Administrators Guide
http://tools.ietf.org/html/rfc1032

Domain Administrators Operations Guide
http://tools.ietf.org/html/rfc1033

Domain names - Concepts and Facilities
http://tools.ietf.org/html/rfc1034

Domain names - Implementation and Specification
http://tools.ietf.org/html/rfc1035

Domain Name System Structure and Delegation
http://tools.ietf.org/html/rfc1591

Domain Name System Protocol Security Extensions
http://tools.ietf.org/html/rfc2065

Secure Domain Name System Dynamic Update
http://tools.ietf.org/html/rfc2137

Negative Caching of DNS Queries (DNS NCACHE)
http://tools.ietf.org/html/rfc2308

IAB Technical Comment on the Unique DNS Root
http://tools.ietf.org/html/rfc2826

Secret Key Transaction Signatures for DNS (TSIG)
http://tools.ietf.org/html/rfc2845

Domain Name System (DNS) IANA Considerations
http://tools.ietf.org/html/rfc2929

Role of the Domain Name System (DNS)
http://tools.ietf.org/html/rfc3467

Threat Analysis of the Domain Name System (DNS)
http://tools.ietf.org/html/rfc3833

DNS Security Introduction and Requirements
http://tools.ietf.org/html/rfc4033

What's in a Name: False Assumptions about DNS Names
http://tools.ietf.org/html/rfc4367

Observed DNS Resolution Misbehavior
http://tools.ietf.org/html/rfc4697