Brief Overview of the Root Server System

ICANN Office of the Chief Technology Officer

David Conrad
OCTO-010
6 May 2020

TABLE OF CONTENTS

1 THE DOMAIN NAME SYSTEM RESOLUTION PROCESS 3

2 THE ROOT SERVER SYSTEM 4

3 ICANN COMMUNITY INVOLVEMENT 5

4 CAN MY ORGANIZATION REQUEST A ROOT SERVER ANYCAST INSTANCE? 5

5 THE ICANN ORGANIZATION'S ROLE 5

This document is part of the OCTO document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

A root server answers the first questions in the process that translates domain names into Internet Protocol (IP) addresses or other data that are used in the operation of the Internet.

1 The Domain Name System Resolution Process

While people prefer to use names for identification, computers generally use numbers. When you are surfing the web, your browser needs to know the IP addresses, or the globally unique numbers, of the web servers that host the websites you visit. After you type a website's domain name into a browser navigation bar or click on a URL link, the browser kicks off the process of *DNS resolution* to find these IP addresses.

The browser sends a question to a "resolver", which is software that implements the DNS resolution process. Resolvers maintain a local copy of answers to questions they have previously looked up, known as a *cache*, so the resolver might already be able to respond to the browser with no further work. However, if the answer isn't in that cache, a description of what happens when the resolver has no answers in its local cache illustrates the complete DNS resolution process. The first step[1] is to send a question that includes the domain name[2] of the website to one of the 13 root servers[3], asking for the IP address(es) associated with the website. However, the root servers only contain information about top-level domains (TLDs), specifically a list of TLDs and the name servers that hold the contents, i.e., second-level names, within those TLDs. The root server queried responds back with a "referral", which is a list of the name servers for the TLD of the website's name. For example, if you're trying to visit the website at "www.example.com", your resolver will send a query to one of the root servers asking for the IP address for that domain name and the root server will respond with a list of all the name servers for ".com", the TLD in our example.

The next step in the resolution process is to send the same question to one of the TLD's name servers that were received in the referral response. Similarly to the root servers, the TLD name servers generally only contain information about the name servers for the domains they are responsible for, in the TLD name server case, second-level domains within the TLD. As such, and just like the query that went to the root server, the query to the TLD name server will result in a referral to the list of name servers for the second-level domain in the question. Using our

1 To be technically accurate, in most cases there is a step before this. When the resolver starts up, it (typically) reads a pre-configured file (known as the "root hints" file) that has the 26 (13 IPv4 and 13 IPv6) IP addresses of the 13 root servers. Once that file is read, the resolver sends a query to one of those addresses to see if the addresses of the root servers have changed. This step, known as the "priming query", is how resolvers have up-to-date information about the root servers.

2 A recent standard known as "Query Name Minimization" (see RFC 7816) recommends that to improve privacy, resolvers should only send the part of the name that is relevant to the name servers being asked, e.g., only send a query for the name servers of TLDs to the root servers, the 2nd-level (with the TLD) to the TLD name servers, etc. Details about this standard and its impact are outside the scope of this document.

3 A "name server" is software on a machine that responds to DNS queries. In the case of the root name servers, they are often referred to simply as the root servers, even though a "root server" is actually a number of machines (as will be described later). To confuse things a bit, resolvers are also referred to as name servers, particularly in home routers and various configuration files, but in this document, we'll always refer to them as resolvers.

previous example, the resolver will send a query for “www.example.com” to one of the “.com” name servers, asking for the IP address of that domain name, and the “.com” name server will respond with a list of all the name servers for “example.com”.

This resolution process continues until a query is sent to a name server that either has the answer – that is, the IP address of the web server – or the name server can authoritatively state that the name does not exist. In our example, the resolver would send a query for “www.example.com” to one of the “example.com” name servers, which presumably knows the IP address(es) associated with “www.example.com”, and would respond with those address(es).

Obviously, each of these steps take time, but the local cache of answers described above speeds things up: before a question is sent to a name server, the resolver checks in its local cache to see if the same question has been asked recently. If it has, the response received the last time the question was asked is returned. If not, when the answer is returned from the name server, it is saved in a local cache of answers and that cache is looked at by the resolver before a query is sent to a name server. This caching is critical to the scalability of the DNS. To add more complexity, if DNS Security Extensions (DNSSEC) have been enabled, the resolver will check cryptographic signatures on the data it receives to verify the data hasn’t been modified by an attacker.

2 The Root Server System

As can be seen from the above, the role of the root servers — primarily to respond to the first step in the resolution process — is quite limited. However, despite this limited role, the root servers are critical to the operation of the Internet. Without the ability to obtain the initial referral provided by the root servers, it would not be possible to look up any domain names on the Internet[4].

The Root Server System consists of over 1000 individual machines (known as root server “instances”), which hold DNS root data. These instances respond to queries from the Internet’s resolvers with referrals to the name servers for top-level domains as discussed previously.

Twelve organizations, known as the “root server operators”, administer 13 “identities”[5], each of which is named with the letters 'a' to 'm' within the “root-server.net” domain, i.e., “a.root-servers.net” through “m.root-servers.net”. Each of these root server identities, known as root services, has two unique IP addresses associated with it, an IPv4 address and an IPv6 address. These IP addresses are pre-configured in all resolvers on the Internet and allow those resolvers to find the root services to ask questions. And the root services receive a lot of these questions: over 70 billion per day.

The 13 root services respond to the queries they receive either with information found in the root zone as it is managed by the IANA Functions operated by ICANN or, in the case that the TLD

[4] Some network operators make use of techniques such as those documented in RFC 7706 (https://tools.ietf.org/html/rfc7706) or similar to make a local copy of the root such that their resolvers do not need to query the root servers. However, deployment of these techniques remains relatively rare and outside the scope of this document.

[5] For (primarily) historical reasons, one organization administers two identities.

being queried has not been delegated, a message that indicates the name does not exist. This information is protected by DNSSEC: any modification of the data by anyone will cause resolvers that have DNSSEC enabled to ignore the response, thereby preventing modification of the root zone or attacks that try to insert unauthentic information into a response.

Resiliency of the Root Server System is critical because the system must be able to respond to a vast and continuous stream of questions as well as to resist various cyberattacks. The root server operators have met this resiliency requirement by distributing root server instances around the world using a routing technique called *anycast.* Anycast routing allows machines all over the Internet to use the same IP addresses to provide identical responses, thereby allowing root server instances to be located in hundreds of different cities and countries. Today, with the large number of root server instances around the world, the Root Server System is extremely resilient. More information about the distribution of root server instances can be found at https://root-servers.org.

3 ICANN Community Involvement

One of ICANN's Advisory Committees, the Root Server System Advisory Committee (RSSAC), consists of root server operators along with others. RSSAC provides advice to the ICANN Board and the ICANN community on matters relating to the operation, administration, security, and integrity of the Internet's Root Server System. RSSAC also appoints interested industry experts to the RSSAC Caucus, a group which produces RSSAC documents, including reports and advisories. More information about RSSAC and the RSSAC Caucus can be found at https://www.icann.org/groups/rssac and a list of documents produced by RSSAC can be found at https://www.icann.org/groups/rssac/documents.

4 Can My Organization Request a Root Server Anycast Instance?

A number of root server operators have programs by which you can deploy a root server instance locally. You can find a list of the root server operators at https://root-servers.org.

Hosting a root server instance benefits the users of large networks such as Internet Service Providers (ISPs) and large enterprise networks, and helps improve the security, stability, and resiliency of the Internet's DNS infrastructure in the local country and/or region. One of the benefits of hosting a root server instance is that it can reduce DNS query response times for your networks, particularly for names that do not exist, and can reduce bandwidth usage for DNS queries that would otherwise go to root server instances outside your network.

5 The ICANN Organization's Role

Beyond operating the IANA Functions which (among other activities) makes update to the root zone that is distributed to the 13 root services, operationally, ICANN org administers one of the 13 root server identities ("l.root-servers.net"), known as the ICANN Managed Root Server (IMRS), and participates in discussions among the root server operators. In addition, ICANN org

supports the RSSAC in their policy deliberations and other activities and the RSSAC Caucus in the work they do.

To help maintain a secure, stable, and resilient DNS infrastructure, ICANN org encourages organizations that meet certain operational criteria to deploy an ICANN-managed root server instance. For more information on hosting an anycast instance of an ICANN Managed Root Server (IMRS), please see https://www.icann.org/groups/rssac/faq.