

September 2007 , Version 2

By Daniel Karrenberg

Dear non-Experts,

This is for you who always wanted to know what the DNS root name servers actually do and how they are operated. These servers are a small but essential part of the Internet Domain Name System (DNS), a part that is often talked about these days but little understood and sometimes even mysticised. After reading this briefing DNS root name servers will be no mystery to you any more and you will also understand what these servers are *not* involved in.

Dear DNS Experts,

This is not for you. For the sake of explaining the principles we will not mention technical or engineering details. You can find these elsewhere. We will also not mention any details that could be used by an adversary to disrupt DNS root name server operations.

What They Serve

Essentially the DNS root name servers reliably publish the contents of one small file to the Internet. The file is called the root zone file. On 12-Dec-2004 there were 5335 lines of text in the file which had a size of 119KB; you can find this particular version here: www.isoc.org/briefings/020/zonefile.shtml

This file is smaller than the ISOC home page <http://www.isoc.org/>. The root zone file also does not change very frequently; it changed only 7 times between 11-Nov-2004 and 12-Dec-2004 and only 90 times in the year preceding 12-Dec-2004. The [individual changes](#) are all localized and relatively small.

Compared to many other services on the Internet this service is rather unspectacular. The function and content of the root zone file are what make it special and cause it to be at the focus of increased attention. The root zone file is at the apex of a hierarchical distributed database called the Domain Name System (DNS). This database is used by almost all Internet applications to translate worldwide unique names like www.isoc.org into other identifiers; the web, e-mail and many other services make use of the DNS.

The root zone file lists the names and numeric IP addresses of the authoritative DNS servers for all top-level domains (TLDs) such as ORG, COM, NL and AU. On 12-Dec-2004 there were 258 TLDs and 773 different authoritative servers for those TLDs listed. Other name servers forward queries for which they do not have any information about authoritative servers to a root name server. The root name server answers with a referral to the authoritative servers for the appropriate TLD or with an indication that no such TLD exists. For a more detailed description of how the DNS works, see the companion ISOC Member Briefing #16 (The Internet Domain Name System Explained for Non-Experts) which can be found here: www.isoc.org/briefings/016/index.shtml.

The Operators

Root servers are operated by twelve organisations often referred to as the "root server operators". They are

- A - VeriSign Global Registry Services
- B - Information Sciences Institute
- C - Cogent Communications
- D - University of Maryland
- E - NASA Ames Research Center
- F - Internet Systems Consortium, Inc.
- G - U.S. DOD Network Information Center
- H - U.S. Army Research Lab
- I - Autonomica/NORDUnet
- J - VeriSign Global Registry Services

This article is also available in [PDF](#) and [ASCII](#)

Version 1 is [available here](#).

Expanded Coverage from ISOC

In-depth articles, papers, links and other resources on a variety of topics are available from the ISOC site at: www.isoc.org/internet/issues

About the Author

Daniel Karrenberg currently serves the [RIPE NCC](#) as Chief Scientist. His interests include Internet measurements, the development of the DNS and the evolution of what others often call "Internet Governance".



Daniel is one of the founders of [RIPE](#). In the 1990s Daniel led the establishment of the RIPE NCC, the first of the Regional Internet Registries. He has helped to shape [Internet address space distribution policy](#), transferring both policy development and implementation to the [region](#).

Daniel helped to design [NSD](#), designed and implemented [dnsmon](#) and deployed the initial [K-root](#) server.

In the 1980s Daniel helped to build EUnet and led the effort to transition it to Internet protocols, making EUnet the first pan-European ISP and bringing Internet connections to many places in and around Europe.

Acknowledgments

The ISOC Member Briefing series is made possible through the generous assistance of ISOC's Platinum Program Sponsors: Afiliias, APNIC, ARIN, Microsoft, and the RIPE NCC, Sida. [More information on the Platinum Sponsorship Program...](#)

About the Background Paper Series

Published by:
The Internet Society
1775 Wiehle Avenue, Suite 102
Reston, Virginia 20190 USA
Tel: +1 703 326 9880
Fax: +1 703 326 9881

4, rue des Falaises
CH-1205 Geneva
Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445

Email: info@isoc.org
Web: www.isoc.org

Series Editor: Martin Kupres

K - RIPE NCC

L - ICANN

M - WIDE Project

The letters A-M represent the 13 numeric IPv4 addresses at which the service is provided. Each operator is responsible for providing reliable DNS service to the Internet at large from their address. Some operators still provide the service from one location with one or more physical machines. Other operators provide the service from multiple locations using a method called "anycast" which is explained in the FAQ referenced below. Actual root name servers machines are located at more than 130 locations within 53 countries (ISO3166 definition of country) worldwide (September 2007). The list of ISO3166 country codes with root name server locations in September 2007 is:

AE AR AU BD BE BG BR CA CH CN CZ DE EC EG ES FI FJ FR GR HK
HU ID IE IL IN IS IT JP KE KR LT MX MY NL NO NZ PA PH PK PL PT QA
RO RU SE SG TH TR TW UK US VE ZA.

For comparison: in December 2004 there 34 countries in this list. There is a map showing the approximate positions of all root name servers at <http://stupid.domain.name/node/407> and more such maps exist.

More information about most operators can be found via www.root-servers.org, or specifically via <http://X.root-servers.org> where X stands for one of the letters listed above.

Some Common Misconceptions

The root name server operators do not determine the content of the root zone file. The file is edited by the IANA according to a process described on the IANA web site. The root name server operators publish the file as received from the IANA. See: www.iana.org/root-management.htm

No Internet traffic passes through the root name servers at all. They have nothing to do with routing, note the difference in spelling. Name servers just answer queries from other parts of the DNS.

The root name servers do not store all the information in the DNS. Storing all the information in one place would be totally infeasible today. This is exactly why the DNS was developed as a distributed database. So if you register thatnewdomain.org the root zone file will not change and the root name servers will not give different answers. The ORG zone file will be changed.

The root name servers are not queried every time you browse the web or send mail. Information is cached in the DNS. Your computer will query a caching DNS server to resolve domain names. A well behaved DNS server needs to query the root name servers only once every 48 hours for each particular TLD.

In the meantime it can resolve names for that TLD without involvement of the root name servers. Because of this caching almost all DNS queries are answered without involvement of the root name servers.

Diversity

While the root zone file represents the apex of a hierarchical naming system, the root name servers that publish this zone file are organised in a distributed and diverse fashion. No single entity has authority or control over the operation of these servers. This diversity and the distributed authority has been a key element of the reliability of the root name service. Therefore this diversity should be maintained in the face of increasing pressure for more hierarchical "Internet Governance".

For More Information

For those seeking more detailed information about the root name servers I have compiled an extensive FAQ on the subject. It tries to answer questions I have received from outside the technical community over the last few years. While it repeats some material from this briefing the FAQ is much more extensive than the briefing; it is also a living document that will be updated as new questions arise and better answers become available. You can find this FAQ here: <http://www.isoc.org/briefings/020/>

Copyright © Internet Society 2005.
All rights reserved.