Briefing Paper on Alternate DNS Roots

Prepared for Names Council
of the Domain Name Supporting Organization
June 2, 2001

Contents:

Delving Into Multiple DNS Roots

Karl Auerbach, *Cisco Systems, ICANN At-large Board member for North America, founder of several technology start-ups and expert in SNMP and network management.*

Note: This is a small part of what will eventually become a much more detailed paper

I'd like to use this e-mail to discuss some of the technical aspects I've uncovered with respect to multiple root systems for the domain name system.

As some of you may know, on some of the nets I operate I've been running a long term (more than three year) experiment to see, first hand, what things might happen if one uses other root systems or runs their own root. These are admittedly small tests of limited scope but I believe that they are representative of the experience others would have should they subscribe to any of the various tiny root systems that have appeared on the net. Since I tend to interact with my systems from both within and without I've had the ability to perceive how a non-legacy root environment behaves when observed by those who use the dominant root.

Some Background

We are going to have multiple roots whether we want them or not – any kid with a Linux box can set up a DNS root. So it is important that DNS be robust and not disintegrate and collapse should some group of folks set up their own DNS playground. Fortunately, DNS is pretty solid and largely immune.

But *largely* immune is not *totally* immune. DNS can suffer data contamination. Because DNS implementations try to optimize their behavior by caching ancillary data there have been situations in which intentionally false information was introduced. DNS implementations have been hardened against these kinds of problems. It is worthwhile to note, however, that this problem was one that happens even without there being multiple DNS root.

But, as we will see later in this document, it is this data caching that is the source of one of the issues that can arise when there are multiple DNS roots.

There are several forces that are driving people to consider establishing DNS roots.

Three Cases

There appear to be three distinct cases to describe the events that transpire when users of the Internet use different DNS roots.

In order to simplify things let me adopt some simple terminology: "Root-D" stands for the dominant (NTIA controlled) DNS root – this is the one that serves the vast majority of Internet users. "Root-X" stands for any of the other root systems.

The three cases seem to be:

A. Root-D and Root-X have identical contents.

B. Root-X has more top-level domains than does Root-D but for those TLDs in common, the contents are identical.

C. Root-X and Root-D contain at least one top-level domain with the same name but with different contents.

Most people consider Case-A to be essentially a mirroring situation and reasonably benign.

Case B represents the situation that obtains today between the NTIA controlled root and the other root systems. There are clearly some technical issues that arise however I believe that these may be limited to the adoption of some simple procedures. I will get to these and moment. Case B contains the most often mentioned issue with regard to multiple root systems: the fact that the name on the Internet may or may not work as expected depending on which root system the user trying to use the name subscribes to. This will be discussed in more detail later.

Case C represents a situation that may readily occur but which most people consider pathological. The real issue here is not a technical one. Rather the question is what are the most effective means to either prevent the situation from arising at all or to create pressures that work so that these situations tend to occur only in somewhat private backwater areas of the net.

There is a related situation that is created by the existence of devices that violate end-to-end principles, most particularly web caches and “content management” devices. This situation is distinct from multiple DNS roots but it does contain one of the most oft complained-of aspects of multiple roots – the fact that a name may not have the same meaning everywhere. These devices often intercept user accesses to the World Wide Web. This is often to improve efficiency or improve response time. However it is also done to create tailored responses – sometimes based on geography, sometimes on based on personal characteristics of the person making the request.

Some Thoughts On The Notion of Universality of Names

There is no denying that the net would be made less convenient if DNS names were not uniquely and identically meaningful no matter where or by whom they are uttered.

Even without the presence of multiple roots, DNS names are already potentially ambiguous. Among the sources of such ambiguity are the following:

- Content management mechanisms, web caches, redirectors
- MX vs A records
- Personalized services

We need to distinguish between two cases:

- A given name may have meaning to some set of users and no meaning (i.e. it is not resolvable via DNS) to some other set of users.
- A given name has different meanings to different sets of users.

Competition in the Root Zone: A Basic Economic Analysis

A selection from Ruling the Root (forthcoming, MIT Press), Chapter 3.
Dr. Milton Mueller

What is the DNS Root?

The term "DNS root" refers to two distinct things: the *root zone file* and the *root name servers*. The root zone file is a list of top-level domain name assignments, with pointers to primary and secondary name servers for each top-level domain. The root server system, on the other hand, is the operational means of *distributing* the information contained in the root zone file in response to resolution queries from other name servers on the Internet. Currently, the root server system consists of 13 name servers placed in various parts of the world. The server where the root zone file is first loaded is considered authoritative; the others merely copy its contents. The additional servers make the root zone file available more rapidly to spatially distributed users, and provide redundancy in case some root servers lose connectivity or crash.

Technically, the most important thing about the DNS root is that it provides a single, and therefore globally consistent, starting point for the resolution of domain names. As long as all the world's name servers reference the same data about the contents of the root zone, the picture of the name resolution hierarchy in one part of the world will continue to match closely the picture in any other part of the world.

Aside from its technical significance, administration of the DNS root is important politically and economically. Defining the root zone file determines who is delegated authority over top-level domains. If top-level domain assignments are economically valuable, and many people believe that they are, then the decision about who gets one and who doesn't can be contentious. Monopoly control of top-level domain name assignments can also provide the leverage needed to regulate registry policies, second-level domain name assignments and other aspects of Internet use.

Competing Roots: A Definition

Recently there has been much discussion of alternate, competing, or multiple DNS roots. Some parties believe that competition is the solution to many of the policy problems posed by ICANN. Others contend that such competition is "impossible" or undesirable. That debate can be clarified by starting with a more precise definition of competition in this arena, and by applying known concepts from economics.

Competition at the root level means competition for the *right to define the contents of the root zone file.* More precisely, it means that organizations compete for the right to have *their* definition of the content of the root zone recognized and accepted by the rest of the Internet.

Competing roots are a form of standards competition. Economic theory has a lot of interesting things to say about how that kind of competition works. In standards competition, user choices are affected by the value of compatibility with other users, not just by the technical and economic features of the product or service itself. A simple example would be the rivalry between the IBM and Apple computer platforms in the mid-1980s. During that time period the two computer systems were almost completely incompatible. Thus, a decision to buy a personal computer had to be based not only on the intrinsic features of the computer itself, but also on what platform other people were using. If all of a consumer's co-workers and friends were using Macs, for example, a buyer's choice of an IBM compatible PC would lead to difficulties in exchanging files or communicating over a network.

There are many other historical examples of competition based on compatibility. Studies of competing railroad gauges (Friedlander, 1995), alternate electric power grid standards (Bunn and David, 1988), separate telegraph systems (Brock, 1981), non-interconnected telephone networks (Mueller 1997), and alternate broadcast standards (Farrell and Shapiro, 1992; Besen, 1992) all have shown that the need for compatibility among multiple users led to convergence on a single standard or network, or in interconnection arrangements among formerly separate systems.

This feature of demand is called the *network externality* in economic jargon. It means that the value of a system or service to its users tends to increase as other users adopt the same system or service. A more precise definition characterizes them as demand-side economies of scope that arise from the creation of complementary relationships among the components of a system. (Economides, 1996) A rich economic literature on the network externality has developed in the past 25 years.

One of the distinctive features of standards competition is the need to develop *critical mass.* A product with network externalities must pass a minimum threshold of adoption to survive in the market. Another key concept is known as *tipping* or *the bandwagon effect*. This means that once a product or service with network externalities achieves critical mass, what Shapiro and Varian (1998) call "positive feedback" can set in. Users flock to one of the competing standards in order to realize the value of universal compatibility and eventually most users to converge on a single system, or on interconnected systems. However, network externalities can also be realized by the development of gateway technologies that interconnect or make compatible technologies that formerly were separate and distinct.

Features of Competing Roots

What does all this have to do with DNS? The need for unique name assignments and universal resolution of names creates strong network externalities in the selection of a DNS root. If all ISPs and users rely on the same public name space – the same delegation hierarchy – it is likely that all name assignments will be unique, and one can be confident that one's domain name can be resolved by any name server in the world. Thus, a public name space is vastly more valuable as a tool for internetworking if all other users also rely on it, or coordinate with it. Network administrators thus have a strong tendency to converge on a single DNS root.

Alternate roots face a serious chicken and egg problem when trying to achieve critical mass. The domain name registrations they sell have little value to an individual user unless many other users utilize the same root zone file information to resolve names. But no one has much of an incentive to point at an alternate root zone when they have so few users. As long as other people don't use the same root zone file, the names from an alternate root will be incompatible with other users' implementation of DNS. That is, other users will be unable to resolve the name.

Network externalities are really the *only* barrier to all-out competition over the right to define the root zone file. A root server system is just a name server at the top of the DNS hierarchy. There are hundreds of thousands of name servers being operated by various organizations on the Internet. In principle, any one of them could set up a public name space, assign top-level domain names to users, and either resolve the names or point to other name servers that resolve them at lower levels of the hierarchy. The catch, however, is that names in an alternate space are not worth much unless a critical mass of name servers on the Internet recognize the alternate root and point their name servers at it.

There already are, in fact, several alternative root server systems. Most were set up to create new top-level domain names. However, until recently only an estimated 0.3% of the world's name servers pointed to them.[1] That changed when New.net, a company with venture capital financing, created 20 new top-level domains in the Spring of 2001 and formed alliances with mid-sized Internet service providers to support the new domains.[2]

Until now, alternate roots have been promoted only by small entrepreneurs unable to establish critical mass. If an alternate root was supported by major Internet industry players, on the other hand, the story could be very different. An America Online, a Microsoft, a major ISP such as MCI WorldCom, all possess the economic and technical clout to establish an alternate DNS root should they choose to do so. If the producers of Internet browsers, for example, pre-configured their resolvers to point to a new root with an alternate root zone file that included or was compatible with the legacy root zone, millions of users could be switched to an alternate root. It is also possible that a national government with a large population that communicated predominantly with itself could establish an alternate root zone file and require, either through persuasion or regulation, national ISPs to point at it. Indeed, the Peoples Republic of China is offering new top-level domains based on Chinese characters on an experimental basis. The whole transition to multilingual names is creating numerous opportunities for establishing alternate roots.

Recall, however, that the value of universal connectivity and compatibility on the Internet is immense. Thus, those who attempt to establish alternate roots have powerful incentives

[1] Joe Baptista, 2000, root server estimates.

[2] Karen Kaplan, "Start-up offers alternative system for Net addresses," Los Angeles Times, March 6, 2001.

to a) retain compatibility with the existing DNS root, and b) offer something of considerable value to move industry actors away from the established root.

Why would anyone want to start an alternate root? Here again, it is useful to refer to the literature on standards competition. New standards create new capabilities and features that may not be possible under the old standard. In ICANN's case, there are two immediate drivers of change: new TLDs and internationalized domain names. In both cases, the existing DNS administration is not responding properly to clear market signals.

ICANN's restrictive policies toward new top-level domains has created a market for alternate root zones. By tightly restricting the number of new top-level domains, by arbitrarily rejecting many valid applications, and by imposing extremely costly regulations on the small number of businesses lucky(?) enough to receive delegations, ICANN encouraged companies such as New.net to risk capital on achieving the critical mass needed to create an alternate name space.

ICANN itself is not directly responsible for the standardization impasse over internationalized domain names, but in general the inability of the IETF and the rest of the Internet community to agree on a common standard has brought us to the brink of an all-out standards competition.

Standards wars are an inevitable part of technological change. Standards changes always bring risks and costs associated with incompatibility. People who invested in digital audio recorders a few years ago probably wasted their money. The transition to DVDs or to High Definition Television will impose additional investment costs on consumers and broadcasters, respectively, and make some equipment obsolete. Technological innovation almost inevitably leads to standards competition in some form or another. Furthermore, monopolies have a tendency to become unaccountable, overly expensive, or unresponsive. Competition has a very good record of making monopolies more responsive to technical and business developments that they would otherwise ignore. So even when a competitor fails to displace the dominant standard or network, they may succeed in substantially improving it.

Need for a DNSO Policy

If alternate roots do exist, and if ISPs and end users choose to use them, ICANN cannot avoid adopting a policy about how it relates to them. The policy decisions can be made consciously or by default, but they must be made. The following aspects of the relationship between ICANN and alternate roots must be confronted:

- Whether to exclude the contents of alternate roots zones from the ICANN root; i.e., whether to compete with alternate roots or seek compatibility with them. By including them it will achieve wider compatibility and foreclose many opportunities for losing control of the root altogether.
- Whether, when it adopts new TLDs, ICANN will choose names that collide with TLD strings already in use by alternate roots, or avoid collisions. ICANN's current approach to this problem has been inconsistent. The assignment of the

.BIZ top-level domain knowingly collided with a string in use by an alternate root; however, the assignment of the .WEB top-level domain to Afilias was explicitly avoided because of its prior use by Image Online Design, an operator of an alternate registry.

- Whether it will respond to the challenge of New.net by becoming more liberal in its policies toward creating and regulating new TLDs.

Network Working Group   Internet Architecture Board
Request for Comments: 2826   May 2000
Category: Informational

IAB Technical Comment on the Unique DNS Root

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Summary

To remain a global network, the Internet requires the existence of a globally unique public name space. The DNS name space is a hierarchical name space derived from a single, globally unique root. This is a technical constraint inherent in the design of the DNS. Therefore it is not technically feasible for there to be more than one root in the public DNS. That one root must be supported by a set of coordinated root servers administered by a unique naming authority.

Put simply, deploying multiple public DNS roots would raise a very strong possibility that users of different ISPs who click on the ame link on a web page could end up at different destinations, against the will of the web page designers.

This does not preclude private networks from operating their own private name spaces, but if they wish to make use of names uniquely defined for the global Internet, they have to fetch that information from the global DNS naming hierarchy, and in particular from the coordinated root servers of the global DNS naming hierarchy.

1. Detailed Explanation

There are several distinct reasons why the DNS requires a single root in order to operate properly.

1.1. Maintenance of a Common Symbol Set

Effective communications between two parties requires two essential preconditions:

- The existence of a common symbol set, and
- The existence of a common semantic interpretation of these symbols.

Failure to meet the first condition implies a failure to communicate at all, while failure to meet the second implies that the meaning of the communication is lost.

In the case of a public communications system this condition of a common symbol set with a common semantic interpretation must be further strengthened to that of a unique symbol set with a unique semantic interpretation. This condition of uniqueness allows any party to initiate a communication that can be received and understood by any other party. Such a condition rules out the ability to define a symbol within some bounded context. In such a case, once the communication moves out of the context of interpretation in which it was defined, the meaning of the symbol becomes lost.

Within public digital communications networks such as the Internet this requirement for a uniquely defined symbol set with a uniquely defined meaning exists at many levels, commencing with the binary encoding scheme, extending to packet headers and payload formats and the protocol that an application uses to interact. In each case a variation of the symbol set or a difference of interpretation of the symbols being used within the interaction causes a protocol failure, and the communication fails. The property of uniqueness allows a symbol to be used unambiguously in any context, allowing the symbol to be passed on, referred to, and reused, while still preserving the meaning of the original use.

The DNS fulfills an essential role within the Internet protocol environment, allowing network locations to be referred to using a label other than a protocol address. As with any other such symbol set, DNS names are designed to be globally unique, that is, for any one DNS name at any one time there must be a single set of DNS records uniquely describing protocol addresses, network resources and services associated with that DNS name. All of the applications deployed on the Internet which use the DNS assume this, and Internet users expect such behavior from DNS names. Names are then constant symbols, whose interpretation does not specifically require knowledge of the context of any individual party. A DNS name can be passed from one party to another without altering the semantic intent of the name.

Since the DNS is hierarchically structured into domains, the uniqueness requirement for DNS names in their entirety implies that each of the names (sub-domains) defined within a domain has a unique meaning (i.e., set of DNS records) within that domain. This is as true for the root domain as for any other DNS domain. The requirement for uniqueness within a domain further implies that there be some mechanism to prevent name conflicts within a domain. In DNS this is accomplished by assigning a single owner or maintainer to every domain, including the root domain, who is responsible for ensuring that each sub-domain of that domain has the proper records associated with it. This is a technical requirement, not a policy choice.

1.2. Coordination of Updates

Both the design and implementations of the DNS protocol are heavily based on the assumption that there is a single owner or maintainer for every domain, and that any set of resources records associated with a domain is modified in a single-copy serializable fashion. That is, even assuming that a single domain could somehow be "shared" by uncooperating parties, there is no means within the DNS protocol by which a user or client could discover, and choose between, conflicting definitions of a DNS name made by different parties. The client will simply return the first set of resource records that it finds that matches the requested domain, and assume that these are valid. This protocol is embedded in the operating software of hundreds of millions of computer systems, and is not easily updated to support a shared domain scenario.

Moreover, even supposing that some other means of resolving conflicting definitions could be provided in the future, it would have to be based on objective rules established in advance. For example, zone A.B could declare that naming authority Y had been delegated all subdomains of A.B with an odd number of characters, and that naming authority Z had been delegated authority to define subdomains of A.B with an even number of characters. Thus, a single set of rules would have to be agreed to prevent Y and Z from making conflicting assignments, and with this train of actions a single unique space has been created in any case. Even this would not allow multiple non-cooperating authorities to assign arbitrary sub-domains within a single domain.

It seems that a degree of cooperation and agreed technical rules are required in order to guarantee the uniqueness of names. In the DNS, these rules are established independently for each part of the naming hierarchy, and the root domain is no exception. Thus, there must be a generally agreed single set of rules for the root.

1.3. Difficulty of Relocating the Root Zone

There is one specific technical respect in which the root zone differs from all other DNS zones: the addresses of the name servers for the root zone come primarily from out-of-band information. This out-of-band information is often poorly maintained and, unlike all other data in the DNS, the out-of-band information has no automatic timeout mechanism. It is not uncommon for this information to be years out of date at many sites.

Like any other zone, the root zone contains a set of "name server" resource records listing its servers, but a resolver with no valid addresses for the current set of root servers will never be able to obtain these records. More insidiously, a resolver that has a mixed set of partially valid and partially stale out-of-band configuration information will not be able to tell which are the "real" root servers if it gets back conflicting answers; thus, it is very difficult to revoke the status of a malicious root server, or even to route around a buggy root server.

In effect, every full-service resolver in the world "delegates" the root of the public tree to the public root server(s) of its choice.

As a direct consequence, any change to the list of IP addresses that specify the public root zone is significantly more difficult than changing any other aspect of the DNS delegation chain. Thus, stability of the system calls for extremely conservative and cautious management of the public root zone: the frequency of updates to the root zone must be kept low, and the servers for the root zone must be closely coordinated.

These problems can be ameliorated to some extent by the DNS Security Extensions [DNSSEC], but a similar out-of-band configuration problem exists for the cryptographic signature key to the root zone, so the root zone still requires tight coupling and coordinated management even in the presence of DNSSEC.

2. Conclusion

The DNS type of unique naming and name-mapping system may not be ideal for a number of purposes for which it was never designed, such a locating information when the user doesn't precisely know the correct names. As the Internet continues to expand, we would expect directory systems to evolve which can assist the user in dealing with vague or ambiguous references. To preserve the many important features of the DNS and its multiple record types -- including the Internet's equivalent of telephone number portability -- we would expect the result of directory lookups and identification of the correct names for a particular purpose to be unique DNS names that are then resolved normally, rather than having directory systems "replace" the DNS.

There is no getting away from the unique root of the public DNS.

3. Security Considerations

This memo does not introduce any new security issues, but it does attempt to identify some of the problems inherent in a family of recurring technically naive proposals.

4. IANA Considerations

This memo is not intended to create any new issues for IANA.

5. References

[DNS-CONCEPTS] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, November 1987.

[DNS-IMPLEMENTATION] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, November 1987.

[DNSSEC] Eastlake, D., "Domain Name System Security Extensions", RFC 2535, March 1999.

6. Author's Address

Internet Architecture Board

EMail: iab@iab.org

7. Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Network Working Group
Internet Draft
Category: Informational
Document: <draft-higgs-root-defs-01.txt>

S. Higgs
Higgs Communications

May 2001

Root Zone Definitions

Status of this Memo

1. Abstract

The purpose of this memo is to provide guidelines to prevent a root zone fragmentation. This memo is provided as a supplement to Request For Comments 2826 (RFC2826)[1]. RFC2826 states that there is a single unique root of the public DNS. This memo attempts to resolve outstanding issues pertaining to a unique root while maintain the unicity of the DNS across any variation of the actual data contained in a root zone. In other words, the total sum of DNS data from all variations of root zone data is a single unique root. This root zone is defined in this memo as the "Virtual Inclusive Root".

This memo also attempts to further refine the concepts of RFC2826 by defining the relationship between the U.S. Government Root Zone and the Private and Inclusive Root Zones.

This memo does not provide guidelines for the introduction of new Top Level Domains, nor does it address the various issues that have delayed the introduction of new TLDs since the first requests were submitted to IANA in 1995[2].

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119[3].

For the purposes of this document, the term "non-U.S. Government" will be referred to as "Inclusive".

3. Unresolved issues pertaining to a unique root

Domain Name Service (DNS) is a hierarchical distributed database architecture[4]. Because it is hierarchical, the assumption is made in RFC2826 that there can be only one unique root zone.

RFC2826 mentions the use of private networks creating private name spaces but does not define the relationship between the private name space and the U.S. Government-published name space.

RFC2606 (also known as Best Current Practice 32 / BCP32)[5] also mentions four reserved top level domains (TLDs) which are used for configuration and testing purposes. These are deliberately left out of the U.S. Government-published name space, and their use immediately creates an "non-U.S. Government" or "Inclusive" root zone.

RFC2826 does not mention enhancements to the U.S. Government-published name space that are provided by non-U.S. Government Root Servers. These are also known as "non-ICANN Root Servers", "Alternative Root Servers", "Enhanced Root Servers", and "Inclusive Root Servers".

This document does not refute the technical findings of RFC2826. In all the variations of root servers examined, there is only one root zone being published for each root server cluster.

The reality is that for various reasons that are beyond the scope of this document, multiple root servers exist within the publicly visible segments of the Internet. It is a simple matter for any DNS Server operator or end user to change their DNS configuration settings to see any of these non-U.S. Government root servers.

It is also possible for DNS information to be altered, at any level within the DNS hierarchy, on any DNS Server, at any time. This is entirely at the discretion of each DNS Server operator. Consequently the DNS Server operator MUST, at all times, act in a responsible manner consistent with the stable operation of the Internet.

Most modern operating systems provide a mechanism (such as the resolv.conf file or a "Network" control panel) that pre-defines the local trusted DNS Servers that will be initially queried. Each computer therefore has the ability to query a unique combination of DNS Servers.

Consequently the end user MAY change their DNS settings and bypass their local ISPs DNS Servers. This allows Inclusive Root Zones to be viewed in the public Internet space.

4. Stability of the root zone and criminal consequences

It should be recognized that in the United States, altering DNS records to the detriment of a pre-existing organization is covered under federal computer fraud statute, 18 United States Code, Section 1030[6]. As a result, criminal convictions have resulted from the alteration of DNS information[7]. Most countries now have similar laws.

5. U.S. Government Root Zone

U.S. Government root servers are identified by the ROOT-SERVERS.NET domain name. Historically, these servers resolve the default root zone that is shipped with DNS server software. The zone file for the U.S. Government root servers can be found here:

ftp://rs.internic.net/domain/named.ca

The authoritative host for the U.S. Government-published TLDs is A.ROOT-SERVERS.NET.

U.S. Government authorized root servers publish the root zone described in RFC2826. This document uses this zone as the baseline to determine the relationships to other published DNS root zones.

Use of the U.S. Government root zone is RECOMMENDED. It is used as the baseline for the Inclusive Root zones.

6. Private Root Zone

Private root zones do not reflect the publicly viewable Internet name space. They MAY carry a sub-set (or none at all) or the U.S. Government-published baseline TLDs.

They are NOT required to carry the complete U.S. Government-published Root Zone.

They MUST NOT be directly accessible from the public Internet. The only exception is when they are accessed through a secure and authenticated gateway (such as a Virtual Private Network (VPN)) in order to identify hosts which are only accessible within an organization's intranet infrastructure.

Use of a Private Root Zone is OPTIONAL. In certain circumstance use may be required to meet the specific operational needs of a particular organization.

7. Inclusive Root Zones

Inclusive Root Zones utilize the U.S. Government root zone as a baseline and add additional TLDs to enhance the name space.

Inclusive Zoot zones SHOULD include the complete U.S. Government-published zone.

Inclusive Root Servers SHOULD peer the name space extended beyond the U.S. Government-published baseline. This can be achieved by reciprocal agreements of non-U.S. Government published TLDs between Inclusive Root Zone operators.

Use of an Inclusive Root Zone is OPTIONAL.

8. Virtual Inclusive Root

The "Virtual Inclusive Root" is the sum of all variations of all publicly-accessable root zone data. It is the gross manifestation of the unicity in the global DNS.

Each root zone MUST pay the same respect to all other root zones.

Each root zone MUST NOT create top level domain conflicts with other root zones.

Pre-existing top level domains MUST be recognized by other root zones as part of the Virtual Inclusive Root zone.

Peering of top level domains amongst root zones is highly RECOMMENDED.

9. Security Considerations

There is an inherent trust relationship created between a DNS Server and DNS Client. By convention, all DNS Servers are expected to return correct information to the DNS Client.

Both Private and Inclusive Root Zone servers become authoritative for subservient DNS Servers and Clients. They will produce results different from the U.S. Government Root Zone servers for non-U.S. Government-published TLDs.

Private or Inclusive Root Zone servers MAY be employed in order to enhance network security of a particular organization. Several well known companies use additional TLDs within their local area networks. These _hidden_ TLDs are used to protect the identity of network assets and do not resolve outside of the company's intranet.

Other Security Considerations for root servers are described in detail in RFC2870[8]. This document RECOMMENDS full compliance with RFC2870.

9. References

1 Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, May 2000

2 Postel, J., "The IANA's File of iTLD Requests", http://www.gtld-mou.org/gtld-discuss/mail-archive/00990.html

3 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

4 Mockapetris, P., RFC1034, "Domain Names - Concepts and Facilities", November 1983

5 D. Eastlake, A. Panitz, "Reserved Top Level DNS Names", BCP32,

RFC 2606, June 1999
6 United States Code, Title 18, Chapter 47, Sec. 1030. "Fraud and related activity in connection with computers" http://www.usdoj.gov/criminal/cybercrime/1030_new.html
7 U.S. vs. Kashpureff (NY) http://www.usdoj.gov/criminal/cybercrime/kashpurepr.htm
8 Bush, R., Karrenberg, D., Kosters, M., Plzak, R., "Root Name Server Operational Requirements", RFC2870, June 2000

10. Acknowledgments

The author would like to thank Karl Auerbach, Scott Bradner, Milton Mueller, Brian Reid, Richard Sexton, and Einar Stefferud for their constructive comments.

11. Author's Address

Higgs Communications
P.O. Box 4519
Sunland, CA 91041-4519

Phone: 818-352-3208
Fax: 818-352-0030
Email: simon@higgs.net

12. Full Copyright Statement

CONTRACT, TORT, WARRANTY, OR OTHERWISE, ARISING IN ANY MANNER RELATING TO THIS DOCUMENT, WHETHER OR NOT SUCH PARTY HAD ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.

13. Expires: November 2001

###

2001.5.21
Kilnam Chon

Issues for Next Generation Root Servers

1. Introduction

There are 13 ICANN/IANA root servers in operation now. They are mostly in the USA now with two in Europe and one in Asia. ICANN is responsible for the operation of the root servers. DNS Root Server System Advisory Committee (RSSAC) of ICANN gives technical and operational advice to the ICANN Board. ICANN and RSSAC look to the IETF to provide engineering standards.[2]

When we look into future of the Internet, we need to consider several issues on the root servers such as

a. How many root servers do we need in future? Are 16 servers, which are the current maximum number of servers, sufficient?

b. Current arrangement on the root servers is centralized with the original copy called Copy A in USA. Can we have a distributed root server architecture?

c. Is the current root server arrangement appropriate for internationalized top level domains?

2. Number of root servers

Although there are 13 root servers, the current specification allows up to 16 root servers. With proliferation of the Internet, which is becoming social infrastructure, it may make sense to expand the number of the root servers to more than 16. This would require changes in the protocol specification, which would take some time, and we may start preparation for potential changes on this matter.

3. Centralized vs distributed

The current root server arrangement is as follows;

- Copy A in USA
- 12 other copies in USA(9), Europe(2), and Asia(1)

Is it technically necessary to have the original copy in one place? Or is the distributed architecture by having the original content in two or more places. We also need to consider social requirement on this matter as the Internet is becoming global social infrastructure.

4. Internationalized domain names

Internationalized domain names are being incorporated starting from the internationalized domain name access. It is matter of time before we have to consider the internationalized top level domains for both ccTLD and gTLD. Is the current arrangement of the ICANN root servers appropriate to facilitate the internationalized top level domains? We need to consider appropriate arrangment for the internationalized top level domains.

5. Remarks

In order to answer the above and other issues for the future technical and social requirements on the root servers, series of meetings and workshops may be necessary.

There are non-ICANN root servers for various reasons; to offer additional TLDs, to offer an inclusive root server, to offer internationalized TLDs, and so on. They may be taken into consideration with respect to the future requirement of the root servers.

References

[1] IAB, IAB Technical Comment o the Unique DNS Root, RFC 2826, IETF, May 2000.

[2] R. Bush, et al, Root Name Server Operational Requirements, RFC 2870, IETF, June 2000.

[3] S. Higgs, Root Server Definitions, Internet Draft, IETF, February 2001.

[4] ICANN, DNS Root Server System Advisory Committee, www.icann.org, March 2001.