Public Summary of Reports Provided Under Cooperative Research and Development Agreement CN-1634 Between the Internet Corporation for Assigned Names and Numbers and the United States Department of Commerce

In 1999, the Internet Corporation for Assigned Names and Numbers (ICANN) and the United States Department of Commerce, as represented by the National Institute of Standards and Technology (NIST) and the National Telecommunications and Information Administration (NTIA), entered into Cooperative Research and Development Agreement (CRADA) CN-1634, which establishes a joint project entitled "Improvements to Management of the Internet Root Server System". Under the CRADA, the parties have been collaborating on a study and process for making the management of the Internet (DNS) root server system more robust and secure.

In late 2002, two reports under ICANN's CRADA were submitted to the Department of Commerce. A 30 November 2002 report provided a description of the current status of the root server system. A 31 December 2002 report concerned a proposal for enhanced architecture for root server security, a procedural plan for the transition to that enhanced architecture, and a schedule for the transition.

Both reports contain security-sensitive and proprietary information, as well as information that may be disclosed publicly without compromising security of the root nameserver system or proprietary information. This public summary has been prepared to provide the Internet community with information from the reports that may appropriately be disclosed publicly.

Internet Corporation for Assigned Names and Numbers
14 March 2003

Introduction

In 1999, the Internet Corporation for Assigned Names and Numbers (ICANN), the National Institute of Standards and Technology, and the National Telecommunications and Information Administration entered into Cooperative Research and Development Agreement (CRADA) CN-1634, under which the parties agreed to collaborate on a study and process for making the management of the Internet (DNS) root server system more robust and secure. Since that time, the parties have been collaborating on that study, with the additional assistance of the operators of the thirteen root servers, who are members of ICANN's Root Server System Advisory Committee (RSSAC).

The CRADA's Statement of Work (SoW) framed the topics that the study is intended to address:

> "Operational requirements of root name servers, including host hardware capacities, operating system and name server software versions, network connectivity, and physical environment.
>
> "Examination of the security aspects of the root name server system and review of the number, location, and distribution of root name servers considering the total system performance, robustness, and reliability.
>
> "Development of operational procedures for the root system, including formalization of contractual relationships under which root servers throughout the world are operated.
>
> "The study will address the technical management of the entire Internet (DNS) root server system, including all (currently thirteen) root servers located throughout the world and the techniques and equipment for generating, maintaining, and distributing authoritative root zone files. The study will include formulation of the operational procedures, requirements, and protocols referenced above through engineering analysis and will be accomplished with appropriate consultation with affected parties, including existing operators of the Internet root server system, through use of collaborator's Root Server System Advisory Committee and otherwise. After operational procedures, requirements, and protocols are formulated, they will be evaluated by implementing them in a controlled manner on the Internet (DNS) root server system. Any change(s) in the designated authoritative source for root zone files must be approved by the Department of Commerce in a separate document and nothing in this SoW or this Agreement is intended to direct such a change."

The study is being conducted in stages, resulting in several reports. The first report, dated 30 November 2002, provided a description of the current status of

the root-server system. A second report, dated 31 December 2002, discussed these three topics:

> a. A written description of the enhanced architecture incorporating a dedicated primary root server and standards for physical protection;
>
> b. A procedural plan for transition to the enhanced architecture;
>
> c. An implementation schedule for transition to the enhanced architecture;

These two topics will be the subject of further work:

> d. The documentation of IANA procedures for root zone editing, root zone generation, and root zone WHOIS service; and
>
> e. An agreement between ICANN and root-server operators that formalizes stable, secure, and professional operation of the root-servers in accordance with the enhanced architecture.

(Topics a-e are specified in paragraph II(C)(5) of Amendment 5 to the ICANN-Commerce Department Memorandum of Understanding.)

With regard to the first three topics, some key aspects of the enhanced architecture contemplated by the CRADA have already been implemented by the root-nameserver operators. In view of this circumstance, this summary is presented in five parts:

A. A background description of the current root-nameserver system.
B. A description of the historical development of the architecture of the root-nameserver system from its deployment in the mid-1980s until the initiation of the CRADA study in 1999-2000.
C. An overall description of the target architecture developed in the CRADA study in the framework of the RSSAC and the root-nameserver operators. This part describes the benefits of and motivations for the new architecture.
D. A description of the currently deployed architecture (which includes some key aspects of the target architecture), including a discussion regarding the process for its implementation.
E. A discussion of the steps remaining to complete implementation of the proposed architecture, including the specifications and requirements for the new facilities required to complete the architecture.

A. Background Description of the Current Root Nameserver System

Currently, the domain-name system has thirteen nameservers that provide nameservice for the root zone. The provisioning of the DNS with thirteen root

nameservers reflects technical limitations of the DNS specification. The basic technical specifications for the DNS are set forth in RFC 1034, "Domain Names - Concepts And Facilities", and RFC 1035, "Domain Names - Implementation And Specification", both written by Paul Mockapetris and published in November 1987. (Other RFCs also discuss aspects of the DNS, but RFC 1034 and RFC 1035 describe its basic features and have been designated Internet Standard STD 13.)

Queries and responses can be transported between DNS clients and nameservers using either the UDP or the TCP protocol. While TCP can be used for any DNS activity, UDP is the recommended method for DNS transactions (other than zone transfer) due to its lower overhead and better performance. [RFC 1035, page 32] As stated in RFC 1035, however, DNS messages transported by the better-performing UDP protocol are limited to 512 bytes in length.

The current limitation to thirteen root nameservers arises from the practical need to have most DNS responses fit within the maximum 512 bytes that UDP can accommodate. If a DNS query results in a response that cannot fit within a 512-byte DNS message, the response is truncated and returned to the client with an indicator (the "TC bit") set to indicate that truncation has occurred. Truncated responses can result in the client sending the query again using TCP [RFC 2181, page 11].

TCP sessions require a much higher connectivity and processing overhead than do UDP queries and responses. Accordingly, having a significant proportion of a nameserver's DNS traffic transported by TCP is a very undesirable feature, and it is important that TCP traffic be kept at a relatively low level.

The choice of having thirteen root nameservers results in a highly robust root-nameserver system while keeping TCP traffic to manageable levels. Having more than thirteen root nameservers would increase the size of certain responses to DNS queries received by the root nameservers so that they could not fit within 512 bytes, which would result in a greater incidence of truncation and a consequent increase in the load on the root nameservers.

In view of the technical considerations described above, it is not practical under current conditions to have more than thirteen root nameservers. It should be noted, however, that this number limitation applies to the number of nameservers listed as authoritative for the root zone, not to the total number of servers that operate within the root nameserver system. In the case of most of the thirteen root nameservers, the nameserver’s query load is actually distributed among multiple servers. The determination of how many cooperating servers to configure to perform a particular listed root nameserver's work is made by the operator responsible for operating that root nameserver, taking into account the server's query load and performance, as well as the need to provision for peak

query conditions.

With the exception of some recent cases in which all or a portion of the activities of an operator's organization were transferred to another organization, which assumed the responsibility to operate root nameserver as part of that transfer, the current operators of the root nameservers were all selected in 1997 or earlier, before ICANN was formed. Shortly after ICANN was formed in late 1998, the Root Server System Advisory Committee (RSSAC) was established, with the responsibility for advising the ICANN Board about the operation of the DNS root nameservers. RSSAC's membership includes the root nameserver operators. RSSAC's charter includes reviewing "the number, location, and distribution of root nameservers considering the total system performance, robustness, and reliability," and RSSAC has been engaged in gathering and analyzing data to assess whether a change in root nameserver locations would, as a technical matter, improve the service levels provided by the root nameserver system to the overall Internet.

Since ICANN's creation no operators of new root nameservers have been selected, and except in the cases involving transferred operations mentioned above there has not been any change in the organizations designated as responsible for a root nameserver's operation. Should the need arise to select a new or successor operator, it is anticipated that the RSSAC would provide advice concerning the technical qualifications and characteristics that operators of new root nameservers should possess.

B. Historical Development of the Architecture of the Root-Nameserver System

From the time of its deployment in the mid-1980s until 2000, the authoritative root-nameserver system consisted of several nameservers organized in a constellation with:

- a single primary nameserver (initially named ns.internic.net and later a.root-servers.net) and
- multiple secondary nameservers (initially with various names and later named b.root-servers.net through m.root-servers.net).

Since 1993 Network Solutions, Inc. (NSI) (acquired by VeriSign in 2000) has had the responsibility under its cooperative agreement with the United States Government to perform edits to and generate the root-zone file. This has been done once or twice daily.

In the initial configuration, revised root-zone files were introduced into the root-nameserver system by loading them on the primary root nameserver; the secondary nameservers retrieved each new version of the root-zone file from the

primary root nameserver. By 2000, this configuration had been changed slightly, so that the various secondary root nameservers obtained their root zones by various means from either the primary root nameserver or a shadow host operated by NSI.

The configuration present in early 2000 may be depicted as follows:

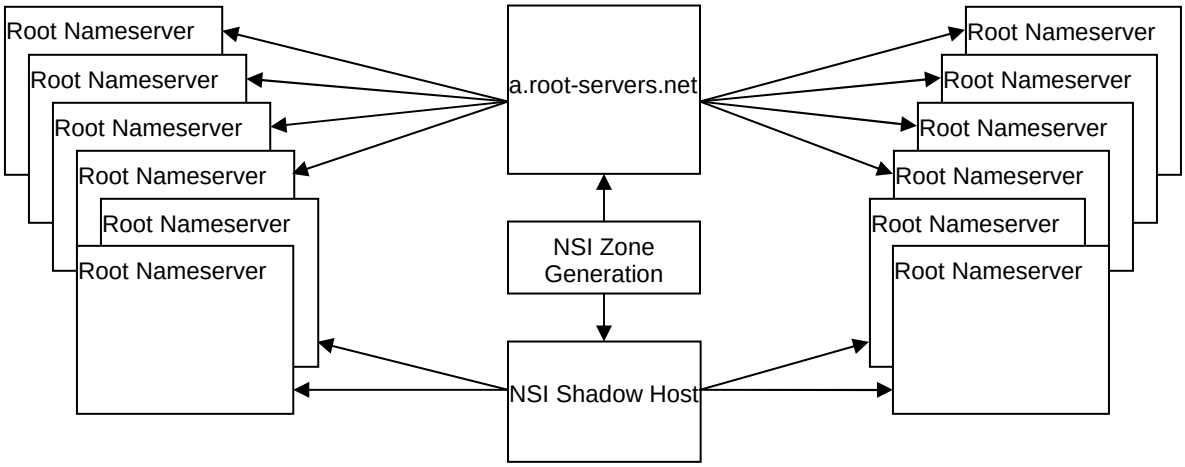

Figure 1 – 1999 Configuration with Public Primary Nameserver

C. Description of the Target Enhanced Architecture for the Root-Nameserver System

Since 1999, work has been ongoing within the framework of the root-nameserver operators and RSSAC on the design of an enhanced root-nameserver-system architecture. This work crystallized in a specific design by mid-2000, and was reported by the RSSAC at ICANN's Yokohama meeting in July 2000. The most prominent enhancements proposed for incorporation into this architecture were:

1. Implementation of a hidden (non-public) primary nameserver for distribution of the root-zone file to all of the root nameservers.
2. Implementation of secure mechanisms for ensuring the authenticity of the root-zone file that is distributed to all root nameservers.
3. Shifting of the responsibility for root-zone-file editing and generating from VeriSign/NSI to ICANN.

These enhancements are described in turn below.

1. Implementation of a hidden primary root nameserver

The key advantage of the first enhancement is greater and stricter protection from possible security breaches relating to the authoritative root-zone files. As shown in the following diagram, the hidden primary nameserver need only be accessible to the root-zone generation system and the public root nameservers:

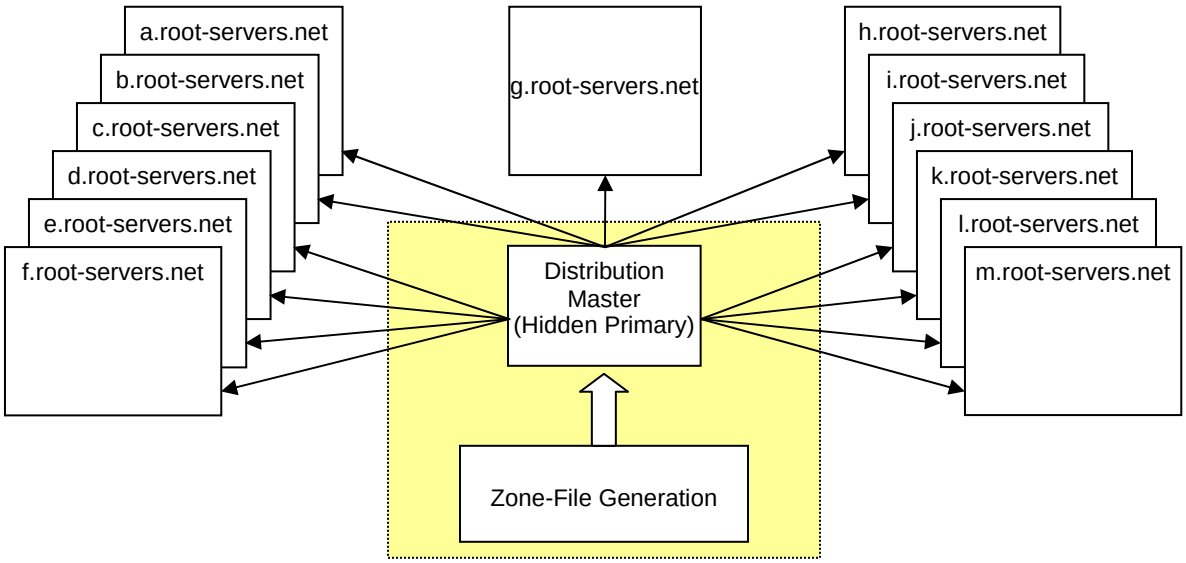

Figure 2 – Proposed Configuration with Hidden Primary Nameserver

Because access to the servers in the shaded box need only be given to the thirteen root nameservers, those servers can be tightly secured through implementation of access-control lists and otherwise. In contrast, the configuration shown in Figure 1 includes a distribution nameserver (a.root-servers.net) that also must be capable of acting as a public nameserver, and therefore must be visible to all users of the Internet. This precludes the strict limitations on access to the distribution source that are possible with a hidden primary nameserver. The elimination of general access to this nameserver has no negative effect on the function of the domain-name system (DNS) or the Internet.

In addition, the separation of the two functions formerly served by a.root-servers.net—distribution of root-zone files to the other root nameservers and provision of public nameservice—is also beneficial because it eliminates an unnecessary (though relatively small) burden on that machine, thereby slightly increasing its capacity to provide public nameservice.

2. Implementation of authenticated distribution mechanisms

The second enhancement is the implementation of mechanisms for ensuring that the root-zone file obtained by each of the root nameservers comes from the genuine source. This feature employs appropriate features of the protocol under development by the Internet Engineering Task Force known as DNSSEC. The two most useful parts of DNSSEC in relationship to the DNS root are:

> (a) the ability to authenticate the transactions between the root distribution source and each of the root nameservers through the shared-secret Transaction Signature (TSIG) mechanism described in RFC 2845; and

(b) the ability to digitally sign the root zone through public-private key technology, so that the authenticity of answers to queries for root-zone information may be verified by security-aware systems using the DNS.

By allowing root nameservers to verify securely and with confidence that the root-zone files they receive come from the expected source, the TSIG mechanism provides protection against potential security breaches such as man-in-the-middle attacks that could result in one or more of the designated root nameservers loading a false root zone and thereby providing responses to DNS queries based on false data. The TSIG mechanism is a private mechanism between the servers involved in the zone transfers, and is invisible to public users transmitting DNS queries (unknowingly as a rule) to the root nameservers. TSIG is described in RFC 2845 and provides very significant protection against these types of exploits.

Use of public-private key technology allows other users to authenticate the content of the root zone using a publicly available key. The signature will give the guarantee that the zone they are seeing is signed by the correct authority. This capability, however, involves complexities (including the deployment of security-aware DNS resolver and other software) precluding its immediate implementation. Therefore, it is not part of the enhanced architecture discussed in this report.

3. Shifting of root-zone editing/generation responsibility from VeriSign/NSI to ICANN

The third enhancement is to shift responsibility for editing and generating zone file from VeriSign/NSI to ICANN. In the current implementation, root-zone change requests from top-level domain (TLD) operators are received by ICANN, which is responsible for reviewing the appropriateness of these requests as part of its performance of the IANA function. Once their appropriateness is verified, ICANN sends these requests to the United States Department of Commerce for approval; these approvals are then transmitted to VeriSign, which makes the changes as requested by ICANN and approved by the Commerce Department. Although the communications among ICANN, the Commerce Department, and VeriSign are cryptographically authenticated (thereby providing protection against spoofing attacks), this complex sequence involves extensive manual handling, which has on a few occasions resulted in incorrect data being introduced into the root-zone file due to clerical error. Many in the Internet community have also expressed concern that conferring root-zone editing responsibility on VeriSign or, indeed, any commercial TLD operator, may not be appropriate in the long term. Operators of TLDs have a significant self-interest in the content of the root zone that could conflict with the responsibility to faithfully make root-zone edits to carry out changes initiated by the IANA and approved by the Commerce Department.

To minimize the potential for errors arising from the current manual mechanism for communicating root-zone changes, as well as to alleviate concerns regarding placing a commercial TLD operator in a situation of conflicting interests, the enhanced architecture proposes to shift responsibility for making root-zone edits from VeriSign to ICANN. As provided in paragraph II)(C)(6) of Amendment 5 to the Memorandum of Understanding between ICANN and the Commerce Department,

> 6. Following Departmental review and approval of [specified] documentation [ICANN would] test and implement the enhanced root-server system architecture, including ICANN's operation of the authoritative root, under appropriate terms and conditions.

It is contemplated that the terms and conditions that would be in place at the time of the shifting of root-zone-editing responsibility to ICANN would involve a mechanism for approval by the Commerce Department substantively similar to that in place presently for changes performed by VeriSign. Thus, root-zone changes would be initiated by ICANN as part of its IANA responsibilities, but would only be executed according to the approval protocols specified by the Commerce Department. By avoiding the need for several steps of manual handling, including VeriSign's entry of the requests into its system, having changes made by ICANN would reduce the potential for clerical errors. As noted above, it would also avoid placing a commercial TLD operator in the potentially conflicted position of also being responsible to faithfully make root-zone edits.

It should be noted that a portion of the specified documentation required before the shifting of responsibility to ICANN can occur is "documentation of IANA procedures for root zone editing, root zone generation, and root zone WHOIS service". That documentation is beyond the scope of the reports already submitted, and is a subject for future work.

D. Current State of Deployment of Enhanced Architecture

As noted above, discussions among the root nameserver operators since the commencement of the CRADA project have led to partial implementation of the enhanced architecture developed in 1999-2000 and described in part C of this summary. Specifically, at the beginning of 2002 the role of primary root nameserver was migrated away from a.root-servers.net to another nameserver, which operates as a distribution master for all thirteen public root nameservers. Also, TSIG has been implemented to authenticate transfers of the root-zone file from the distribution master to each of the thirteen public root nameservers.

The revised configuration was initially implemented in two stages in 2Q and 4Q 2002. Initially, a first group of root nameservers were converted to receive TSIG-signed updates from the distribution master; after this conversion was successfully completed the remaining root nameservers were completed. Zone

transfers from a.root-servers.net were disabled as of November 2002. There were no major engineering issues encountered in the transition from the old architecture to the present distribution-master architecture, nor were there any service interruptions experienced by Internet users.

E. Remaining Steps to Complete the Enhanced Architecture

The successful implementation of a dedicated distribution master, as described in part D of this summary, has demonstrated that it is technically feasible to move the source from which root-zone files are distributed to the public root nameservers without adverse effects on Internet operation. It also considerably simplifies completing the implementation of the proposed enhanced architecture.

To complete the proposed enhanced architecture, the following additional technical steps are required:

1. Final testing of the zone-generation module for the IANA registry system.
2. Deployment and testing of a primary distribution master operated by ICANN.
3. Deployment and testing of an alternate distribution master.
4. Shifting of root nameservers to receive new root zones from the ICANN-operated distribution master rather than the current distribution master.

In addition, the following non-technical steps will be involved:

5. Establishment of protocols for Commerce Department approval of root-zone changes. (These are intended to be logically equivalent to the approval protocols followed in the current IANA-Commerce Department-VeriSign arrangement, but will be clerically simplified due to the fact that changes as approved by the Commerce Department will not require manual entry into VeriSign systems.)
6. Revision of existing IANA procedures for maintenance of the root-zone registry and Whois services to incorporate generation and verification of the root zone for loading on the distribution masters.
7. Establishment of suitable agreements or similar documents formalizing the stable, secure, and professional operation of the root nameservers.
8. Commerce Department approval of steps 5 and 6, as contemplated by paragraph II(C)(6) of Amendment 5 to the ICANN-Commerce Department Memorandum of Understanding.)

This summary covers only steps 2-4 above. Steps 1 and 3-8 will be covered in future work.

Step 2: Deployment and testing of a primary distribution master operated by ICANN.

Implementing a primary distribution master requires arranging a suitably secure physical environment for the server, deploying the server hardware, and arranging for appropriate connectivity for the server.

Facility. A suitable facility has been identified for deployment of the primary distribution master.

Deploying the server hardware. The requirements for the server hardware for the distribution master are relatively undemanding. The distribution master normally must handle only two incoming and 26 outgoing data transfers per day (each update is distributed to each of thirteen root nameservers). Conventional server equipment fully meets the requirements.

The distribution master will also be provisioned with appropriate routing equipment, but again conventional equipment clearly meets the requirements.

Connectivity. The distribution master will be provisioned with a separate connection to the Internet obtained under contract from a major carrier.

Step 3: Deployment and testing of an alternate distribution master.

Arrangements have not yet been made for a facility to house an alternate distribution master. However, the security environment of the alternate facility will be similar to those for the primary facility.

One option for an alternate facility would be having an organization other than ICANN operate it, thereby providing organizational diversity for the operation of the zone-distribution function. This diversity would ensure that the function of a distribution master would be available not only in the event of a technical failure of the primary systems, but also in the event of an organizational failure of ICANN itself. Although a distribution master operated by another organization would not achieve the goal of minimizing the potential for clerical errors to the same extent as an ICANN-operated alternate facility, this trade-off for the organizational diversity is likely worthwhile in view of the fact that the distribution master would only be employed in the event of failure of the ICANN-operated primary distribution master.

Step 4: Shifting of root nameservers to receive new root zones from the ICANN-operated distribution master rather than the current distribution master.

Once the new distribution masters are deployed and tested, shifting from the current distribution masters to the ICANN-operated distribution masters will be a relatively straightforward matter. A phased process will be employed. The new distribution master will initially load root-zone updates from the existing distribution master as they become available. These transfers will be authenticated with the already-deployed TSIG process. After proper operation of

this process is fully verified, groups of root-nameservers will begin to receive root zones from the new distribution master instead of the existing one. (During the timeframe of this process, the existing distribution master will serve as a back-up source for the root nameservers to obtain root zones.) All transfers from the new distribution master will be authenticated using TSIG. After the proper functioning of transfers to each group of nameservers is verified, a new group of root nameservers will be transitioned.

After all root nameservers are transitioned to receive root-zone updates from the new distribution master, the old distribution master will be retained as an alternate source for direct loads to the root nameservers for a specified period, during which operation of the transfers from the new distribution master will be monitored.