Caucus feedback:

1. In all the initial tables 1-2 the software was included, later in 3 there isn't any mention. I wonder if this is because there isn't information or if it wasn't considered relevant anymore.

Response: The first two tables showed software diversity in the root server system, and how it moved from mainframe to Unix and from JEEVEs to BIND. Table 3 forward mainly emphasize the root server organizations, not so much about software diversity so it is not necessary to list software again.

We will add a footnote to Table 3 noting this emphasis.

2. Before moving to Section 3 it would be good to have a simple table summarizing the 13 root servers (you have the table at the Appendix but I think it would be interesting to have it at the end of the section as a nice conclusion.

Response: This information is repeated in section 5.1. A sentence is added to refer to the appendix.

3. I know that you provide the statements verbatim from the operators, but there are some that do not read as well as the rest of the document. Could you ask some of the operators to rewrite a little some of the statements or provide them with some editorial suggestions?

   The ones that I really like and could be used as an example are B, F, I and M. The ones that I suggest some rewrite are A (J it is much better and less commercial. A it is like a PR statement ...) For others I will suggest to make it more about the node history, why it was important to deploy and not so much about the people (which are important but not central to the document).

Response: Root server operators will consider this feedback. In some cases it would be difficult to change due to company policies.

4. In section 2. "History of Root Servers" I think we should be more

descriptive, tell more history, give more context, what the problem was, problem distributing of the hosts.txt file, limitations, scalability, things like that. Saying this, IMHO I believe the history in the document can be better organized. It's a kind of shocking that it starts talking about the RFCs and just in the second paragraph it jumps into implementation of the first root server. I believe our target is technical people but the document is about the History of DNS. I wonder if we could include things like how much time took the first developments, how many people was involve, language program used for, etc. I personally enjoy these details.

Response: Add the following to the text:

Prior to the development of the Domain Name System(DNS), hosts in the ARPA research and Defense Data Network (DDN) operational communities were assigned names in a flat or global name space of character strings (e.g. USC-ISIF). The name to address translation was done by looking up the information in a table of all hosts. The maintenance of this table was centralized at the Network Information Center (NIC) at SRI and each host is expected to obtain a current copy of the table on a timely basis from SRI-NIC.

As the size of the network grew, so did the number of hosts. The size of this table, and especially the frequency of updates to the table are near the limit of manageability. What was needed is a distributed database that performs the same function, and hence avoids the problems caused by a centralized database. To address this bottleneck, In 1983, Jon Postel and Paul Mockapetris published a series of RFCs that laid out the design of the Domain Name System (DNS) [foonote 1] and the transition plan to DNS for the ARPAnet.

[footnote 1] The proposed domain name system has three major components: 1) The Domain Name Space, which is a specification for a tree structured name space. Each node and leaf on the tree has an associated label and corresponds to a resource set. The domain name of a node or leaf is the path from the root of the tree to the node or leaf. When domain names are printed, labels in a path are separated by dots (“.”); 2) Name servers – server programs which hold information about the domain tree’s structure and set information; and 3) Resolvers – programs that extract information from name servers in response to user requests.

5. In few parts of the documents it's mentioned that (for example A-Root) is one of the 13 logical Internet Root Servers....., also in 2.7 under the "Adding Root Letters", this information is terrific!, I would love to expand in why the "limitation" is 13?. It's mention Maybe this is not history but this is something that probably is not widely well explained. I don't think this point is particularly relevant, I just wanted to mention it just in case.

Response: Add the following in the footnote.

By moving to root-servers.net, operators were able to take advantage of DNS label compression, leaving room for four additional root servers to fit within a 512 byte DNS response [footnote 1].

[footnote 1]: The limitation is specified in RFC 1035 because at the time, there were networks that could not handle DNS packets larger than 512 bytes without fragmenting, as well as known firewall rules to drop DNS packets more than 512 bytes in size.

6. What do you think about dedicating a small paragraph about what DNS label compression is.., this concept is mentioned twice in the document. I might be wrong but my guess is that many people don't know about this.

Response: Add the following footnote:

Domain name compression was introduced in RFC1035 as an optional protocol feature and later mandated by RFC1123. In this scheme, an entire domain name or a list of labels at the end of a domain name is replaced with a pointer to a prior occurrence of the same name in the same message, thus eliminating the repetition of domain names in a message and reducing the size of the message. In the case of responses to root server priming queries, the domain root-servers.net appears only once in the response, instead of 13 times, one for each root server.

7. And finally, I wonder if at the end, along with the appendix we could add something like an infographic image

Response: This is a good idea. We can engage ICANN communication to produce infographic and that can be a supplement to the document, but not in the main body of the text.

8. IIRC the concern was less about limited diversity in name server software than about limited diversity in open source name server software, i.e. the root server operators have always had a preference for open source because of the auditability it provides: it's always possible to see what your software is doing. This seemed important for the root, both for debugging (serving the root is sometimes a corner case) and for general transparency.

   Both BIND and NSD are open source to this day. Other quality open source name server code bases have also appeared since those days, but I won't attempt to list them in case I leave someone out!

Response: Text revised as the following:

“Most root server operators have always had a preference for open source name server software because of the auditability it provides. This is important for the root, both for debugging and for general transparency. In early 2000s, there were increasing concerns about the lack of diversity in open source name server software. The RIPE NCC partnered with NLnet Labs to design and develop an authoritative name server (NSD) from scratch. The RIPE NCC contributed requirements, input to the design and lab testing to the initial development of NSD. NSD was deployed on K-Root in 2003."