OVERSIGHT AND MULTIPLE ROOT SERVER SYSTEMS

Vittorio Bertola

The WGIG report addresses different options for the oversight of public policy issues in Internet governance, including the most problematic issue – the authorization of changes in the root zone file of the Domain Name System (DNS). All these options were developed under the assumption that it is possible to reach global agreement on a model for the management of a single root server system for the DNS.

While this is definitely the most desired outcome, it is important to provide an analysis of other options, in view of the possible scenario in which the only agreement that can be reached is that different stakeholder groups desire to set up different processes to manage the root server system, and, consequently, different root server systems.

In that scenario, the results for the stability and the functioning of the Internet can be very different according to modalities and relationships between the resulting root systems. There are ways of splitting the single root server system without damaging the Internet, which should be embraced in such a situation.

Reasons for a Single Root Server System

The discussion on whether the existence of multiple root server systems should be allowed or even encouraged has been going on for over ten years.

The existence of a single root server system is traditionally justified in terms of the two following objectives:

a. *Preventing confusion.* It is important that, whenever a user enters one URL or domain name into any Internet application, he is pointed at the same IP address. If different users were pointed at different places in different moments or by different service providers, global communications would be threatened.

b. *Ensuring stability.* Once a domain name is established, users start to build content, services, activities, companies that rely on its functioning. The sudden disruption of an entire Top Level Domain (TLD), or problems with the resolution of its names, could cause significant economic, social and even military disruption.

These two objectives have a lot of merit; certainly, a single root, if coupled with appropriate policies, ensures that they can be met. However, *it has never been demonstrated that a single root server system is the only way to meet these objectives.*

During the WSIS discussions, a third feature of the single root became evident: the centralization of control. Whichever entity controls a single root has the power to prevent confusion and ensure stability, but also to create confusion and disrupt stability. While this never happened in the past, and is unlikely to happen in the present arrangements, the increased importance of the Internet has brought many countries and stakeholders to demand a redefinition of control processes over the root zone files.

However, another option to prevent centralized control from being used badly – that is, against those who do not exert such control, but are affected by it – is to distribute control over the root zone.

Different Ways of Splitting the Root

The root server system is, in the end, a set of pointers, much like a central turnpike of the Internet. It is a sort of central place to which all users can go and ask, "Where do I have to go to find this name?" The pointers point the user at another place in which the name can be found, or where further pointers can be accessed. The different root servers of a single root server system correspond to multiple places in which the pointers on the turnpike are ensured to be the same.

As a first note, there is no practical effect caused by having multiple root server systems, managed by different entities, if all pointers continue to point at the same places. These multiple systems could be defined as *cooperative*: they share the burden of the service, and are in fact beneficial to the Internet. However, distribution of control in this scenario would be only partial: there would still be the need for global agreements, as changes in the pointers must happen similarly and at the same time in all the roots.

Confusion arises only when the uniqueness of the pointer breaks; that is, when different root server systems point the user at different places when asking for the same name. Two root server systems that do so can be defined as *confrontational*; this is a situation in which also the stability of the Internet would be severely disrupted, especially if the lack of agreement happens on major TLDs[1].

However, the case in which, while varying the root server system, the user is either told that the name does not exist, or pointed at the *same* place, does not create confusion. The uniqueness of the pointer would still be ensured, and the only consequence would be inaccessibility from users of those root server systems in which the TLD does not exist.

[1] Possibly, the only practical result would be getting someone to create "gateways" or "translators" between names in one root and in the other, and thus to create a "super-root" encompassing all roots in a new, single naming system.

This happens when someone establishes a new root server system by taking the content of the existing root and removing some TLDs, possibly with the purpose of making them inaccessible. The result would be a *decremental* root server system; it would not generate confusion, but it could disrupt the functioning of the TLDs that were deleted, or, as a minimum, prevent users of the decremental root server system from accessing content in those TLDs. This situation is generally undesirable, except that those who control a root server system might want to decrement it as a form of attack or retaliation against users and operators of the deleted TLDs.

However, this also happens when someone takes the content of the existing root zone file and only adds further TLDs of its liking, using strings that are not used anywhere else in any other root server system. This can be defined as an *incremental* root server system, and would not cause the same concerns as in the decremental case.

The four different situations, plus the single root, are assessed in the following table:

Table 1: Evaluation of different types of multiple root server systems

	Prevents confusion?	Preserves stability?	Distributes control?
Single root	Yes	Yes	No
Cooperative roots	Yes	Yes	Partly
Confrontational roots	No	No	Yes
Decremental roots	Yes	No	Yes
Incremental roots	Yes	Yes	Yes

Source: self

Lessons from the Past

The possibility of multiple root server systems is inherent to the decentralized nature of the Internet; as with any other application and protocol, any user is technically and practically free to pick which servers to use, and even to create a new root server system.

In the last ten years, while control on the "mainstream" root server system was being moved from Jon Postel to the Internet Corporation for Assigned Names and Numbers (ICANN), there have been a number of attempts to establish other root server systems. While none of them could have an appreciable impact, still some of them reached considerable user bases. Commercial companies such as New.net and UNIDT have established alternate root server systems[2] and reached agreements with major ISPs, software makers and even governmental

[2] This, however, was often coupled with the adoption of browser plug-ins to be installed by those who still used the "mainstream" root server system. The plug-in would intercept names in "non-

entities that started to use them, so that new TLDs, introduced by them in their own roots without any agreement with ICANN, are now visible to hundreds of millions of Internet users. Similarly, some countries are known to have established separate root server systems for the experimentation and deployment of Internationalized Domain Name (IDN) TLDs, given the somewhat slow pace at which this kind of domain names is being introduced in the ICANN root.

This shows that *the "single root" is in fact a myth: a significant number of different root server systems have been in use for years and already coexist without problems.*

The reason why these events were not subject to extensive public notice and review is that… they work. They meet the need of their users, without creating confusion or threatening the stability of the TLDs in the mainstream root. This, in turn, is due to the fact that all these root server systems, none excluded, were created in an *incremental* way: supporting ICANN for the mainstream Generic Top Level Domains (gTLDs) and Country Code Top Level Domains (ccTLDs), and only adding new TLDs, managed by the alternate root server system operator, to them.

In comparison, the only cases in which there has been extensive public concern over the DNS are related to decisions from ICANN, and especially creations and redelegations of gTLDs. The main cause for this, rather than mismanagement, is simply the major visibility and impact that the affected gTLDs have.

However, rather than defining the present situation as one in which there is one global root server system, it is more correct to say that in the present situation, for historical and practical reasons, one root server system has a de facto monopoly on the "market" for name resolution services at the root level. This monopoly situation, in turn, has made power struggles over its management so important and heated.

Having multiple root server systems, as long as users are freely able to pick which one to use, would constitute an intrinsic "check and balance" over how each root server system is managed, and would possibly make strict and direct oversight over the DNS less necessary.

The freedom of choice of the end user is a key element for this mechanism to work. If the user encounters a name that cannot be resolved and used through his own root server system, he will simply try to gain access to another system that works. This is why all alternate root server systems were born in an incremental way: if users can choose, they will pick the system that

mainstream" TLDs and convert them into a conventional subdomain of a domain name in the mainstream TLD which was owned and used to that purpose by the alternate root server system operator. In this way, also users of the "mainstream" root could access the new TLDs.

ensures the broadest accessibility of TLDs and the less concerns on its policies and management.

Dynamics of the Root Server Market

Users of the Internet pick which root server system to use when configuring the DNS servers that they will then use for the resolution of domain names into Internet Protocol (IP) addresses. This operation is typically done by ISPs and by corporate system administrators, rather than by individual end-users; however, smarter users – especially those who use the GNU/Linux operating system, and other Unix flavours – typically run their own DNS server themselves.

There is no simple way to force any user to pick one root server system against another. This would imply mass filtering of traffic through firewalls, as is done by some countries, and capillary control of each and every DNS server. However, it is fairly easy to bring the majority of users to use one root server system: most DNS server software – first of all BIND, the free DNS software which has a quasi-monopoly over this market – already comes with the choice of a root server system preconfigured into it, and, de facto, DNS software makers have the highest share of control in determining which root system is used. Another important role is that of ISPs: forcing all ISPs to use a given root system is a manageable task for a country and would possibly bring 90% of the users under control.

However, it must also be said that even now, in the present situation, nothing can prevent a country from passing legislation or adopting technical measures to force users onto a separate root server system, nationally controlled. In fact, some countries are said to be already doing it.

Apart from force, however, the only way to bring significant numbers of users and ISPs to adopt a different root server system than the current mainstream one is to give them a compelling reason to spend time in changing the configuration of their DNS servers. In the past, attempts have been focused on providing additional content and services through new TLDs. Whether an argument based on control and theoretical risks would be sufficient to push people to action, it is still to be determined.

Root Servers and Freedom of Expression

An additional argument often made in the press against multiple roots is that some countries would like to deploy their own root server system to increase opportunities for censorship.

Political censorship by governments is usually accomplished by altering results from search engines and filtering traffic at the application level – for example, matching all web pages that a user tries to access against a list of blacklisted URLs. (By the way, this usually happens thanks to the well-paid technical efforts of leading technology companies from the United States and other Western countries.) The DNS does not have much to do with that.

Commercial censorship – that is, preventing anyone from talking negatively about a company and its products – is, on the other hand, often exerted by corporations through extra-judiciary actions, including challenging the registration of the domain name through alternate dispute resolution policies. To this regard, many experts from civil society consider ICANN's Uniform Dispute Resolution Policy (UDRP)[3] and the deriving policies to be imbalanced in favour of corporations, and to have been exploited for commercial censorship in a number of cases. In the end, the possibilities to exert commercial censorship through the policies for management of a root server system strongly depend on how much these policies are designed to protect freedom of expression; in that regard, the current mainstream root has a bad track of results.

In conclusion, freedom of expression is not endangered or fostered by whether there are a single or multiple root server systems, but by the policies enacted by each of them. The existence of multiple roots, however, would make it more likely to find at least one root server system whose policies are more friendly to freedom of expression or other human rights.

An Example of a Non-Damaging Split

Due to the architecture of the Internet, the deployment of alternate root server systems is likely to happen if no agreement can be found to change the management of the current mainstream root, or if some stakeholders are significantly unhappy with the agreement found.

This event, however, can be managed in a way so that it does not threaten the Internet, and it still meets the traditional objectives of preventing confusion and ensuring stability, while removing control from that party that exerted it traditionally, and distributing it.

The main condition, as evident from Table 1 above, is that the resulting set of root server systems is of the incremental type. This would mean, in practice, that:

a) All root server systems should agree to "carry" each other's TLDs, and let their users access all TLDs from all different root systems;

[3] The UDRP is a set of rules, established by ICANN after discussions with WIPO, that introduces extra-judicial procedures to arbitrate conflicts over the assignment of domain names, and fight against *cybersquatting* – the practice of registering domain names corresponding to known names so to later sell them to the companies or people who have a direct interest in such names. Parties who think to have special rights to a domain name may hire a UDRP provider (WIPO is the most frequently used one) which in turn will appoint arbitrators who will decide over the request to reassign the name. In practice, the most frequent use of the UDRP is by corporations who want to recover possession of domain names corresponding to, or containing, their name or the names of their products.

b) Different root server systems, with different regulating entities or processes, would be in charge for different sets of TLDs;

c) All root server systems should commit not to interfere with the management of TLDs that belong to other systems;

d) All root server systems would be free to introduce new TLDs, as long as the related strings are not already in use by any other root server system.

While all root server systems could manage all types of TLDs, there could be some merit and simplification in agreeing on a subdivision along specific categories. For example, there could be one root server system in charge for the management of ccTLDs, and one or more root server systems that manage existing gTLDs and introduce new ones on a "first come, first served" basis.

From the practical point of view, a *mechanical* registry would be necessary to keep formal account of which TLDs already exist, and which root server system is in charge for each of them. By "mechanical" we mean that no power of decision should be given to this registry – it should just record new additions as soon as they are requested, for example like Internet Assigned Numbers Authority (IANA)[4] does with port and protocol numbers.

Speaking from a user's point of view, in such a system names would exist and point at the same place independently from which root server system is used. The change would thus have no effect for users of the Internet.

However, as an additional note, one could imagine the case in which the managers of one root server system have an issue with a TLD created or managed by another root system. In that case, by removing rule number 1, each root server system would be free not to make other TLDs accessible if they deem this necessary. This, however, would fall into the "decremental" case, and potentially constitute a form of attack and "information warfare" between managers of different root server systems. Careful evaluation needs to be applied to determine whether advantages of this specific option are more than disadvantages.

Conclusion

An agreement between all parties, that addresses and resolves the concerns of all stakeholders through an evolution of the processes for the management of the existing "mainstream" root server system, is the most desirable outcome of the WSIS process for what concerns the DNS, provided that such agreement is compatible with the natural freedom of choice given to all stakeholders by the architecture of the Internet. Any agreement causing significant

[4] The IANA, currently run by ICANN, is the registry of protocol and port numbers, and of TLD delegations. It does not make policy decisions, but simply records decisions taken by others (e.g. by ICANN itself, for TLDs) or assigns resources (e.g. numbers) through a fixed and objective algorithm.

dissatisfaction in a stakeholder group would likely bring to the creation of alternate root server systems.

If no satisfying agreement can be found, however, a plan for the creation and deployment of multiple root server systems could be devised without hampering the functioning of the Internet, as long as the split is done in an incremental fashion and with some clear commitments by all. A decremental structure could be considered if the possibility for root system managers to "opt out" from the creation of new TLDs needs to be preserved, at the risk of hampering the stability of the DNS.