The ISP Column
A column on various things Internet

August 2025
Geoff Huston

The Governance of the Root of the DNS

Background

The arrangements regarding the composition and organisation of the provision and operation of authoritative root servers are one of the more long-lasting aspects of the public Internet. In the late 1980's, Jon Postel, as the IANA, worked with a small set of interested organisations to provide this service. It was informally arranged, without contracts and without payment of any form. It appeared to me that the selection parameters used by Jon were a combination of interest in undertaking the task, expertise and capability in performing the task, and geographic diversity in the set of operating root nameservers covering the major areas of Internet deployment at the time.

It was determined that there was an upper bound on the number of nameservers whose names and IPv4 service addresses could be packed into a 512-byte DNS response to a priming query, and this upper bound is 13 (512 bytes is the maximal DNS over UDP payload size that every DNS implementation is assured of receiving (sec 4.2.1 of RFC1035). Thus, the number 13 was inscribed into DNS mythology as the fixed upper limit of the number of uniquely named authoritative nameservers for the root zone of the DNS.

As the Internet expanded, the pressure placed on this seemingly arbitrary limit increased. A maximum of 13 discrete DNS service platforms was clearly not viable as a long-term plan, and the introduction of IPv6 as a supported protocol also implied that the size of the response to a priming query would inevitably exceed 512 bytes.

The technical response was to introduce the use of *anycast* into the DNS root server set (see RFC7094), where each individual instance of a root zone server could be replicated across multiple locations using the same IP service address. It was left to the Internet's routing system to direct a DNS client to the "closest" root server for each uniquely named service platform.

Operational experience with the root server system to date indicates that these arrangements have proved to be adequate in terms of meeting the demands being placed upon them so far. However, the Internet continues to expand both in the population of connected devices and the volume of DNS queries this device population generates, and the result has been that the pressure to continue to expand the capabilities of this framework of root zone authoritative servers has been relatively constant. But this is not a pressure that is visible to users of the Internet, either practically or financially. To date the system operates on the basis that each root service operator meets its own costs in providing this service, so meeting the costs of further scaling this service is a task that is left to the Root Service Operators (RSOs).

With the incorporation of ICANN there were a number of changes in these arrangements. The IANA remained the entity that was responsible for the contents of the root zone, but in relation to the labels in the root zone that did not refer to individual countries, then the IANA was directed by various ICANN processes under the generic classification of generic Top Level Domains (gTLDs). The designation of two-letter top level codes that refer to countries continued to be informed by the ISO 3166 standard.

The production and maintenance of the root zone itself is provided by the *Root Zone Maintainer* (RZM), a function performed by Verisign (who also operate the a.root-servers.net and j.root-servers.net root server systems) in an agreement with ICANN.

The RSOs are a set of independent autonomous entities who undertake to faithfully publish the root zone and answer all DNS queries that are directed to root zone servers to the best of their capability. Their interactions with ICANN are organised largely within the framework of the Root Server System Advisory Committee, which advises the ICANN community and the Board of ICANN on matters relating to the operation, administration, security, and integrity of the Root Server System. The RSSAC consists of representatives from the root server operator organizations and liaisons from the partner organizations involved in the technical and operational management of the root zone (see https://www.icann.org/en/rssac).

In 2018 the RSSAC published a Proposed Governance Model for the DNS Root Server System (see RSSAC037), and presented it to the ICANN Board of Directors and the larger community of ICANN and Internet users. The proposal noted that the Root Server System has scaled and adapted to the growth of the network and continues to provide resilient service so far, but it felt that the time had come for the RSS to adopt new governance structures and business models to meet the more rigorous requirements of governance, accountability, and transparency.

It's useful at this stage to remind ourselves that the role of determining what is contained in the Root Zone, namely the labels that are in effect top level domains, are determined by the ICANN community through the operation of various ICANN policy fora.

It's also useful to remember that the resource records that are contained in the root zone are also not determined by the RSOs. Such records are specified through the IETF-administered Internet Standards Process and adopted again through consultation within the ICANN community.

The processes used to alter resources records in the root zone are operated by the PTI and the RZM.

So while the term "Governance Model for the DNS Root Server System" may appear to encompass the entirety of the functions that play a role in generating and distributing DNS Root Zone, this exercise is limited to the distribution function, an in particular to the deployment and operation of the collection of authoritative nameservers that respond to DNS queries to the root zone.

In response to this RSSAC037 report, the ICANN Board chartered the Root Server System Governance Working Group (GWG) as a special purpose working group, chartered to make recommendations about the ongoing evolution of Root Server System (RSS) governance in line with recommendations originally adopted and published by RSSAC in RSSAC038.

The GWG brought together participants nominated by identified stakeholder communities with a special interest in the RSS. Participants in the GWG included representatives of each RSO);

individuals nominated by the generic top level domains (gTLD) registry community, the country code top-level domains (ccTLD) registry community, the Internet Engineering Task Force (IETF), and the ICANN Security and Stability Advisory Committee (SSAC); and liaisons from the ICANN Board of Directors, the Internet Assigned Numbers Authority (IANA), and the Root Zone Maintainer (RZM).

The GWG Report

The draft GWG report has now reached the public comment stage. The background is explained reasonably well in the document's introduction, but to summarize, the "Functional Model" is the outcome of lengthy deliberations among the Root Server Operators and representatives from other community groups (including SSAC, as well as the IAB and gTLD and ccTLD registry operators) as to what structure would best serve the governance needs of the RSOs as a group and the Root Server System overall.

The document provides an initial structure, a functional model for how it ought to develop (consistent with the previously published "Governance Principles for the Root Server System" document, and some guidance on how to evaluate the outcomes of each phase.

This work matters, in that there have been major gaps in the governance of the RSOs and the RSS for about 30 years now, ever since the DNS became prominent as critical Internet infrastructure in the late 1990s. The Internet has been living on borrowed time with respect to two key issues, namely the arrangements for funding of root servers (which so far has been entirely provided by the RSOs themselves by various means), and the undefined procedures for adding, modifying or deleting specific RSOs in the root zone. There is also the consideration of the technical evolution of the service model of the root zone in the context of the continuing evolution of the DNS itself. The GWG's Principles and Functional Model documents do not directly resolve these questions but intended to provide a basis for transparent, legitimate, multi-stakeholder decisions in these areas to be made in the future.

The current GWG functional model document may disappoint people who expected more concrete outcomes, but long deliberations and difficult consensus-building went into it, and it's by no means clear that more substantive measures were possible to propose; it has been evident to many participants in this process that we need legitimate structures and processes that don't exist yet to make decisions regarding matters such as evaluation criteria for root server performance, and the issues of selection and removal of root service operators.

The specific weaknesses in this report as I see it are:

1. The report lacks implementation detail: for instance, where and how the proposed GWG Secretariat is incorporated, what its budget would be and how that would be paid for, and who is expected to provide this funding. The mode of approach of the GWG group was to avoid discussing such details of implementation and spend some years on an extended discussion of principles (see the Governance Princples Document). The Functional Model was derived from these principles and in many areas is extremely vague on organisational detail. The funding description is a good example of this, and the treatment of finances (Pages 30 and 45 of this report) is somewhat cursory in nature.

2. The relationship with ICANN, which one could assume would be one of the more significant sections in this report, was addressed using a somewhat dismissive treatment, addressed in a single sentence on page 19 of the Functional Model report.

3. The GWG has focused on preserving a form of status quo for the Root Server System, instead of exploring evolutionary models for serving the DNS root zone and its governance. On the other hand, it is entirely unclear what the bounds on scope of this working group were. While the initial incarnation of the GWG canvassed a broader scope of organisational models, the request from the Root Server Operators (RSOs) for all RSOs to directly participate in the RSS GWG as individual entities without indirect representation skewed the group's composition so that any course other than a relatively anodyne restatement of the status quo was practically infeasible for the group.

4. The more challenging questions, such as who and how to make decisions as to RSO designation and removal are pushed down the road to this future body. The essential characterisation of the collection of RSOs as autonomous entities who operate in a manner that lacks visible accountability to the broader set of stakeholders and the global community of users who are indirectly dependant on the service that these entities operate remains a feature of this proposed structure.

5. Much effort has been spent on introspection on the current arrangements relating to the provision of Root Service via the existing RSOs, but little in the way of substantive suggestion for changes. It is not that the current system is perfect - far far from it, and the reliance on the altruism of a select cadre of entities without associated accountabilities to operate the core of the DNS is concerning. But small-scale incremental changes to the current framework are not going have much of an impact, and there is no viable commitment to undertake anything further than largely cosmetic and insubstantive changes to the current framework.

A reader of this report might reach the conclusion that the objective in this report is little more than proposing the ongoing sustenance of the current arrangement for serving the root zone of the DNS through this small collection of operators who operate with a degree of independence and autonomy. But the underlying background to this work is that questions as to the legitimacy of the undertaking this critical role by a small set of historically nominated entities have been raised and will continue to be raised. The proposed governance structures do not appear to offer any significant level of assurance that the concerns that apparently have motivated such questions are adequately addressed in the measures proposed in this report.

Is the objective here the addition of more organisational structure to provide a wrapping around the current arrangements of the provision of the root service, largely preserving the status quo as the overriding consideration? Or was this exercise an opportunity to explore mechanisms that could evolve the provision of the root service with the necessary attributes of fidelity, scalability and resilience without exclusively relying on the altruistic efforts of a small cadre of service entities? If so, then this report is obviously disappointing in its failure to look beyond the current arrangements, and it's likely that in the face of further scaling pressure on the DNS and the root services, then the internet community will be returning to the same conversation about the framework for the provision of root services in the DNS in the near term future.

I'm not sure that there is much of enduring value in this report's proposals. Adding further layers of administrative process is not going to engender confidence in the stable and resilient nature of the operation of the service if one already had concerns in this space about the somewhat random legacy of decisions made some decades ago over which entities were chosen to operate this service and the relatively informal nature of the relationship between each operator and the IANA. The Internet has matured as it has scaled over this period, and we observe many of the previously informal relationships and arrangements have been formalised into enumerations of mutual commitments in the form of contracts and similar formal instruments.

Are there other options for the RSS?

The Internet was constructed during the wave of deregulation of the previous telecommunication model of national monopolies. It had been argued that this model out outlived its utility, and the burden of monopoly rentals was exacting too high a price on consumers and enterprises, and the related resistance to technological evolution was hindering any moves to improve the efficiency and utility of the telecommunications endeavour. It was believed that the scale of ongoing investment in telecommunications infrastructure was well within the scope of private sector investment, and there was no need to sustain a public sector investment model. A deregulated private sector-led activity would also be attuned to the evolving needs of the consumers of this service. It was this thinking that permeated the telecommunications sector in the last two decades of the twentieth century.

These days the Internet can be regarded as a collection of quite conventional markets, where consumers are able to express their preferences in the selection of providers, and the role of regulation io intended to ensure that consumer interests are protected through the operation of open and stable markets that are not distorted or undermined by the aberrant actions of a few self-interested parties. Is there any residual role for altruism in this space? Or does altruism undermine the conventional operation of such markets? Market dynamics would suggest that escalating demand for a service would be associated with increased service revenue, which would fund further expansion of the service platform and motivate further suppliers to enter the market. In the model of root service provision further scaling of the capacity and reach of the root server system is made by making further calls on the donation of service and support by this select group of root service operators. In other contexts, the provision of goods and services into a market at a price well below the cost of service provision is termed "dumping" and has a negative connotation of a market actor attempting to bring down the price to a point that drives out competitors.

We operate the DNS root service in its current framework because it represents a set of compromises that have been functionally adequate so far. That is to say the predominate query-based approach to root zone distribution with a select group of authoritative service operators hasn't visibly collapsed in a screaming heap of broken DNS yet! And it will probably continue to operate in a robust manner for many years to come.

But we don't have to continue relying on this query-based approach just because it hasn't broken so far. Our need to further scale this function is an ongoing need, and it makes a whole lot of sense to take a broader view of available options and investigate alternatives to the just-in-time delivery process used by DNS's incremental query name resolution algorithm.

We have some options as to how the root service can evolve and scale.

We can wait for the DNS system to fracture and then try and salvage the DNS from the broken mess, or we could explore some alternatives now, and look at how we can break out of a query-based incremental root content promulgation model and view the root zone as just another content "blob" in the larger ecosystem of content distribution in the Internet. If we can efficiently load every recursive resolver with a current copy of the root zone, and these days that's not even a remotely challenging target, then perhaps we can put aside the issues of how to scale the root server system to serve ever greater volumes of queries to ever more demanding clients, and perhaps also provide an alternate answer to the continual questions about the politics and finances relating to root servers and their operation.

The reason why content distribution networks have revolutionised the Internet in recent years is that pre-provisioning at the edge makes for a faster, cheaper and more scalable network in the current context of abundant computing and storage capabilities. If we are prepared to allow this same thinking to intrude into the way we provision the DNS, then I suspect there are similar benefits that could be realised for the DNS as well.

The technique is described in RFC8806, which describes a method for the operator of a recursive resolver to have a complete root zone locally and not to make specific queries to the authoritative root servers. The basic approach is to create an up-to-date root zone service on the same host as the recursive resolver and use that service when the recursive resolver looks up root information. The recursive resolver validates all responses from the root service on the same host, just as it would validate all responses from a remote root server. An alternative implementation with a similar outcome is to load the validated contents of the root zone into the resolver's local cache.

If these were to be adopted as a default mode of operation by all recursive resolvers in the public Internet, then the query load on the authoritative root servers would be significantly reduced. Furthermore, if the root zone itself were to be served as a web object over HTTPS, then the object could be served through the existing CDN services, posing a relatively minor addition load to existing CDN platforms. A recursive resolver need only use the ZONEMD message digest record to assure itself of the authenticity of the retrieved zone object, and the SOA record in the zone can be used to ensure the currency of the data. The intrinsic property of the zone's message digest is that it no longer matters how or where a client obtains a copy of the root zone. As long as the message digest can be validated and the date field in the zone's SOA record is current, then it does not matter how the zone file was retrieved.

This is by no means a new or novel approach, and many recursive resolvers, from the scale of the very largest of systems (operated by Google as their Public DNS Service) through to single instances in end sites already use this.

The essential point here is that we are not in a position where the continuation of the existing arrangements with the Root Service Operators is the only option available to us. Nor is it even clear that further investment in organisational structure surrounding these legacy arrangements represents the best possible use of our time and resources, particularly if the aim is one of improvement in the resilience and performance of the root system of the DNS. It strikes me that we can leverage the signed root zone to move beyond the current trust model of "I can trust the answer I receive if I query one of the 13 IP addresses that I learned from a priming query" to a trust model of "I can load my cache with a complete current and authentic root zone because the ZONEMD record validates the authenticity of this zone file I've obtained via a locally present CDN." If we can focus our collective attention in this direction, then perhaps we can move beyond perpetuation of a largely historical arrangement in a direction that adds a few million authoritative sources of the root zone in the guise of the entire set of recursive resolvers.

ICANN's Call for Public Comment on the Functional Model for Root Server System Governance can be found at: https://www.icann.org/en/public-comment/proceeding/functional-model-for-root-server-system-governance-11-08-2025, and a link to submit comments can be found at that page. This comment period closes on 22 September 2025.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

The author was one of two liaisons from the IETF to the RSS GWG. The views expressed here are his personal views and are not endorsed by anyone else!

Author

Geoff Huston AM, M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net