ICANN's Root Name Service Strategy and Implementation

ICANN Office of the Chief Technology Officer

OCTO-016
26 October 2020

TABLE OF CONTENTS

This document is part of ICANN's Office of the Chief Technical Officer (OCTO) document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

Executive Summary

This document describes, at a high level, the ICANN organization's strategy and implementation plans for the ICANN Managed Root Server (IMRS). The strategy has two goals, which are associated with its implementation plans.

The goals of the IMRS strategy are:

- Supporting the Internet community by placing root server instances in diverse locations
- Protecting the confidentiality, integrity, and availability of the root server system during attack

This document is a public version of the full strategy with certain confidential, financial, and operational details omitted.

1 Introduction

The root server system (RSS) of the global Domain Name System (DNS) faces growing volumes of traffic generated by legitimate users, mostly through the millions of recursive resolvers that are operated by Internet Service Providers (ISPs), network operators, and other organizations. This increase is driven by many factors, such as the growing number of new generic top-level domains (gTLDs), the steady increase in the complexity of web pages with embedded domain names, and the growing number of connected devices that perform DNS queries.

While the RSS has operated successfully since its inception, it is increasingly at risk of being unable to keep pace with the increase of attack traffic launched by malicious entities, misconfiguration of the RSS, misuse, or bugs. Some measurements suggest that attackers' ability to launch larger and more disruptive attacks increases every year, and the cost of implementing those attacks decreases.[1] At the same time, the costs incurred by the operators of the RSS continue to climb to mitigate these attacks using the traditional approach, such as provisioning sufficient instances or ensuring that instances are able to handle much more than typical traffic .

This document gives an overview of ICANN org's strategy aimed at providing improved ICANN Managed Root Server (IMRS) availability, consistency, and resiliency. It describes a multi-pronged strategy that expands and enhances existing approaches It also facilitates the standardization and implementation of technologies, such as "hyperlocal" (described later in this document), which improves the decentralization of the root name service to mitigate risks that the RSS may face over time.

Because this strategy is comprehensive in nature, some of its aspects may impact ICANN org, the ICANN community, and the Internet as a whole. Careful planning, significant resources both from the ICANN org as well as the community and care during implementation will be required to meet the stated goals. However, since ICANN's Bylaws require the org (specifically section

[1] For example, see https://www.helpnetsecurity.com/2020/03/27/ddos-attacks-increase-2020/, https://www.techrepublic.com/article/major-ddos-attacks-increased-967-this-year/, and https://www.infosecurity-magazine.com/news/ddos-attacks-on-the-rise-1-1-1-1/.

1.1(a)(ii)[2]) to ensure the security, stability, and resilience of the DNS, such a strategy is necessary due to the continued evolution and growth of the root system.

Each root server operator independently creates its own instance, placement, and operational strategies within the cooperative root server system. IMRS's overarching strategy is to be a useful and reliable participant in the root server system based on ICANN's strengths and mission as an organization. The other root server operators will have their own independent strategies based on their own strengths.

2 Goals of the Strategy

The IMRS strategy laid out here is based on two high-level goals. The strategy for each goal is described in detail later in this document, as are the plans for implementing the strategies associated with each goal.

2.1 Support the Internet Community by Placing Root Server Instances in Diverse Locations

The primary way to achieve this goal is to make new IMRS single instances available at a low initial cost to organizations that have good connectivity relative to local standards even in places with limited connectivity that can be expected to be good long-term stewards of the instances. ICANN focuses primarily on Internet service provider (ISP) organizations, because they are located closer to recursive resolvers used by end users, and often have good interconnections with other local ISPs. Note that non-ISPs are eligible even though they are not the primary focus of ICANN.

Because of the way the Internet's routing system works, "location" in this strategy typically refers to Internet topology, not geographic location. A resolver in a particular autonomous system can provide a measurable number of hops from various root servers. IMRS instances are located in autonomous systems of many sizes and types. This does not preclude instances that are geographically located. There may be reasons why placing IMRS singles in particular geographic areas makes sense; however, care is taken when selecting these areas to ensure such instances are not overly redundant.

2.2 Protect the Confidentiality, Integrity, and Availability of the Root Server System During Attack

Attacks on the root server system as a whole are relatively rare but can have significant consequences. These attacks are also described in a recent published report from the root server operators (RSOs).[3]

[2] See https://www.icann.org/resources/pages/governance/bylaws-en/#article1

[3] See https://root-servers.org/media/news/Threat_Mitigation_For_the_Root_Server_System.pdf

Since the DNS namespace is hierarchical, successfully attacking the root name service could make name resolution impossible for many users. To reduce the potentially negative impact of future attacks on the RSS, it is necessary to increase the availability of the service and the resiliency of the servers that provide access to that service in a cost-effective manner. Although some potential mitigation strategies exist (e.g., the root zone records could be cached and the lifetime of these records in the cache could be extended by resolvers when they are unable to reach a root server), these mitigations are only temporary solutions. They may not be effective in the face of a sustained attack.

To reduce attacks on the integrity of the data delivered, Domain Name System Security Extensions (DNSSEC) allows the signing of DNS data by the data owners. While the adoption of DNSSEC has been slow, despite strong encouragement and support from ICANN, it remains the standard way to help resolver operators ensure that the data they are receiving from authoritative servers has not been altered by attackers.

Interest in the confidentiality of the DNS has grown over time, particularly in light of the disclosures by Edward Snowden. New technologies to protect DNS traffic from being watched are now being widely deployed. As these technologies mature, more users will be protected from attackers who want to observe their DNS queries.

3 IMRS Placement

All twelve RSOs have, either directly or in partnership with other organizations, increased capacity by deploying multiple servers, referred to as “instances,” which use the same Internet Protocol (IP) addresses to respond to root name service queries. This deployment is done using an operational routing technique known as *anycast*, which enables adding servers anywhere on the Internet where the IP addresses of the root server can be announced into the global routing system. Because each instance is independent of the others, except for operational control and the data being served, the ability of each instance to withstand attacks can be tailored to the needs of that instance.

In the case of ICANN org, IMRS instances have been deployed globally in more than (as of this writing) 165 locations, using two different name server codebases. Operational management of these servers is centralized and performed by ICANN org’s Security and Network Engineering (SaNE) department. The Office of the Chief Technical Office (OCTO) provides the deployment strategy. These instances consist of two types of deployments:

- *IMRS singles* are a single server for each location, hosted by third parties, and with a signed agreement between the host and ICANN. These instances are easy to set up at low costs. Hosting organizations bear these initial costs and deploy the servers in networks of various sizes throughout the world. Having many IMRS singles in diverse locations creates more catchments for the sources of distributed denial of service (DoS) traffic. Although these instances may be overwhelmed by attack traffic, they will tend to keep the attack traffic from the rest of the root server system.
- *IMRS clusters* are large installations consisting of multiple servers and significant networking equipment in a single location; they are deployed at major interconnection points. When under a DoS attack, IMRS clusters are likely to answer queries long before being overwhelmed compared to IMRS singles. Thus they can prevent resolvers from timing out and attempting to move to other root server identifiers.

To fulfill the goal to "support the Internet community by placing root server instances in diverse locations," ICANN org will both expand the range of IMRS singles and choose the placement location of IMRS instances by performing the necessary monitoring and measurements on the root service.

3.1 Increasing Deployment of IMRS singles

ICANN org will continue to promote and deploy IMRS singles throughout the world. To gain the most benefit for increased IMRS single deployment, a major focus will be on areas currently underserved by existing root server instance deployment. ICANN will continue its research in measuring and forecasting the demand for IMRS instances.

In addition to geographic concerns, ICANN org is measuring the number of routers between the currently deployed IMRS singles and the wider Internet. This research will lead to a more data-driven model for proposing new IMRS single locations that will augment the current policy of adding new instances based on requests from potential hosting organizations. The end result will be a more detailed placement strategy for IMRS singles.

3.2 Enhancing Root System Monitoring

Generally, it is difficult to protect a system that is not measured and monitored, because it is challenging to know when the system is under attack, how the attacks are being implemented, and what the attacks are accomplishing. As such, a critical component of the strategy for ICANN org in reducing the effects of attacks on the RSS as a whole is to enhance the monitoring of the RSS.

Wherever possible, a core part of this strategy will be to leverage other work related to the monitoring of the RSS, such as the nascent efforts to implement "RSSAC047: RSSAC Advisory on Metrics for the DNS Root Servers and the Root Server System".[4] In some cases, such as the uptime monitoring of root servers, little work is likely needed by ICANN org. In other cases such as monitoring the selection of root servers by resolvers, ICANN org will need to develop and deploy monitoring systems.

Gaining access to data in order to provide monitoring will be a key challenge. In many cases, data will not be available without in-network active probes on other organizations' networks or by gaining access to the system or resolver logs of other organizations' servers. As such, some monitoring that would be beneficial will not be possible given the practical difficulties of enabling such access.

Additional research is needed to determine what can and cannot be monitored and how monitoring will be performed. This research, conducted in cooperation and collaboration with the RSOs and other parts of the community, is part of the strategy.

[4] See https://www.icann.org/en/system/files/files/rssac-047-12mar20-en.pdf

4 Resilience when Under Attack

The assumption for this goal is that at some point the root server system will be under attack and that this will significantly affect many of the root server identifiers. This assumption is based on a history of relatively rare but significant DoS attacks against the root server system.
For this goal, supporting the resilience of the root server system does not mean that IMRS must never go down during an attack on the root server system: that would be a goal that incurs unbounded costs and yet is also unmeasurable. Instead, the goal is to have measurable resilience against attacks that include attacks on IMRS.

To develop a strategy aimed at reducing the effects of attacks on the RSS, it is necessary to understand the categories of attacks to which the RSS is subject. For this document, the commonly used "confidentiality, integrity, availability" (CIA) model of security is used for categorization, although the categories are listed in this section in reverse order because the availability attacks are the most common by far and, thus, the most significant for the RSS.

The sections below list the types of attacks that ICANN org can help mitigate. There are other types of attacks, such as compromising the DNSSEC key signing key, that are expected to be both exceedingly rare and for which ICANN can provide no effective mitigation. This is why these attacks are not listed in this report.

The diverse strategies listed here all serve the overall goal of protecting the confidentiality, integrity, and availability of the root server system during attack.

4.1 Availability Attacks and their Mitigation

Most attacks against the RSS are some form of DoS attacks. Although the infrastructure of the RSS, and thus its capacity to withstand DoS attacks, is growing because of the independent efforts of the twelve RSOs, there is a significant and increasing risk that the growth in root server capacity could be overtaken by the growth in the capacity of DoS attacks. Given the poor state of device security, particularly end-user systems such as the "Internet of Things" (IoT) devices being deployed in ever greater numbers, it is safe to assume that the increases in attack capacity will continue for the foreseeable future.[5] This suggests a greater risk that the root zone will not be available to all users during sustained DoS attacks. It is also likely that at some point in the future, it will no longer be cost-effective to add more capacity to address the DoS risk; thus, it will be necessary to adopt new strategies.

Separately, there have been cases in which software bugs in the name servers or operating systems could allow for the exploitation of vulnerabilities that would permit DoS attacks by crashing or at least severely slowing down the name server or operating system. Alternatively, a software bug could allow for a "remote code execution" that could permit the compromise of a root server instance itself, where that compromise could at least bring a loss of availability of root service by that instance, or possibly have much worse effects.

[5] See https://www.icann.org/en/system/files/files/sac-105-en.pdf

4.1.1 Investigation into IMRS Cloud

OCTO will continue to investigate adding cloud-based service from one or more content delivery networks (CDNs) or DNS service providers. The IMRS cloud approach could take advantage of third-party cloud vendors' infrastructure to increase the number of IMRS instances as other root server operators have done.

IMRS cloud presents a different set of tradeoffs compared to a IMRS single and a IMRS cluster. IMRS cloud could provide the same, better, or worse protection than IMRS clusters, depending on the type of DoS attacks that they are intended to thwart and the specific infrastructure of the cloud provider. OCTO plans to investigate the various parameters and characteristics, such as money, energy, cooling and other resources needed per query as well as performance characteristics, management overheads, and security, stability, and resiliency (SSR) requirements of IMRS singles, IMRS clusters, and any planned IMRS cloud. These investigations will play a partial role in deciding how to best support the strategy of the root server system when it is under DoS attack.

Before making a decision to deploy an IMRS cloud, ICANN org will investigate the cloud services landscape and potential costs associated with providing root service via cloud services. ICANN org will engage cloud providers to discuss costs and features, and to investigate how the org's stringent requirements, particularly around anycast routing announcements, could be met. In addition, any agreement with a cloud service provider would need to minimize the risk of an over-concentration of root service capacity in any cloud service provider. If multiple root operators were deployed with only a single cloud service provider, an unacceptable single point of failure risk could be created. It is possible that ICANN org could partner with one or more cloud service vendors to obtain the advantages of a significantly wider coverage at lower costs than traditional deployments, but further investigation is required.

4.1.2 Supporting Root Service Decentralization with Hyperlocal

Architecturally, the root of the DNS name space serves as a single point through which the lookup of any name within that name space must pass at least once. This fundamental design aspect poses a risk of a single point of failure for the entire DNS. Historically, this risk was mitigated by adding more root server IP addresses; 26 addresses (13 IPv4 addresses and 13 IPv6 addresses) are in use today.

More recently, anycast technology has been used to greatly expand the number and location of servers for each IP address, known as instances. There are now over 1,000 individual instances. However, as described earlier in this document, the approach of depending on RSOs to voluntarily add more and more capacity is unlikely to be sustainable in the face of the rapidly increasing availability of attack capacity.

The DNS as a system has an architectural requirement of a single namespace: multiple namespaces can result in name collisions, where multiple uncoordinated entities are responsible for the same name. But the architectural requirement of a single namespace does not dictate a particular implementation of that namespace. Indeed, the proliferation of instances funded and operated by the RSOs demonstrates that the implementation of the namespace can

be decentralized across over 1,000 machines. As such, a logical mitigation of root server availability attacks would be to further decentralize the root service.

Beyond deploying instances controlled by the RSOs, one decentralization technique would be for recursive resolvers to obtain and use a copy of the root zone themselves. This approach, termed *hyperlocal*, is already in use in some networks and, at the current scale, does not require any actions by ICANN. Hyperlocal gives recursive resolver operators a way to ensure the availability of the root zone for their local users, even when root service is unavailable to the resolver, perhaps because of a significant attack. A recursive resolver can use the hyperlocal root technique by making software configuration changes, similar to what is outlined in RFC 8806, "Running a Root Server Local to a Resolver".[6]

ICANN org has already made preliminary efforts to facilitate hyperlocal deployments by encouraging and funding developers of recursive resolvers to implement features making hyperlocal configurations easier and less error-prone. Possibilities for further development will be a topic of ongoing research by OCTO. In addition, hyperlocal deployments require the ability for resolvers to obtain the root zone. ICANN's SaNE function is already providing a root zone distribution service to address this requirement, albeit the distribution service, as currently implemented, was not designed for a significant scale.

The existing root zone distribution mechanism offered by the Root Zone Maintainer (RZM), through which the root zone is made available, is designed and scaled to meet the needs of answering normal DNS zone transfer requests by the existing RSOs. If hyperlocal were to see a significant uptake, a new system for root zone distribution would need to be devised to satisfy the reliability and scalability requirements associated with the widespread hyperlocal deployment in recursive resolvers. As part of the strategy to encourage hyperlocal deployment, ICANN org will investigate approaches to the scalable distribution of the root zone for resolver operators who want to use a hyperlocal strategy.

The strategy to encourage hyperlocal deployment would be supported by promoting and publicizing how to configure a recursive resolver to use the hyperlocal technique. In addition, a "sign-up" mechanism through which resolver operators can request to be notified of current and expected changes to the root zone, could be deployed. Such a system could then facilitate wider communications with resolver operators for other root-related activities such as future or emergency KSK rollovers.

4.2 Integrity Attacks and Mitigation

The goal of this form of attack is to corrupt the data associated with a DNS query or response on the root server instance itself, or in the network that provides Internet connectivity to the instance. Although the effects of this attack could be mitigated with the use of DNSSEC, the attack remains viable given the relatively low deployment of DNSSEC to date, both in terms of DNSSEC signing of zones and of enabling DNSSEC validation in resolvers. While there have been no known compromises of root server instances that permitted an integrity attack, given the potential impact of such an attack (an Internet-wide corruption of DNS responses to specific questions facilitating global man-in-the-middle attacks), it would be prudent to prepare for this form of attack.

[6] See https://datatracker.ietf.org/doc/rfc8806/

ICANN org's mitigations for integrity attacks focus on protecting both the IMRS instances and the networks they rely upon, as well as on the deployment of DNSSEC at the root.
ICANN org's SaNE team applies strict security controls on access to the physical systems of the root instances. On instances hosted by other organizations, e.g., IMRS single hosting organizations, the terms of the agreement under which SaNE engages the hosting organizations to deploy IMRS single instances mandate strong security controls and limitations on physical access to the systems. Despite being hosted by external organizations, SaNE maintains the software and administrative control of IMRS singles.

With respect to deploying DNSSEC at the root, the root zone of the DNS has been signed with DNSSEC since July 2010. Key ceremonies in which sets of zone signing keys (ZSKs) created by Verisign acting as the Root Zone Maintainer are signed with ICANN org's key signing key (KSK) every three months, except under extreme circumstances, such as the COVID-19 lockdown. Verisign then includes the KSK-signed ZSKs in the signed root zone. The signed root zone enables anyone who has configured the public key of ICANN's KSK as their DNSSEC-validating resolver's trust anchor to verify that root zone data in a DNS response has not been corrupted.

4.2.1 Encouraging DNSSEC Validation

As described earlier, the RSS is subject to integrity attacks in a variety of forms. Given that DNSSEC was specifically designed to address the integrity of DNS responses, and the DNS root zone was signed in 2010, one of the most effective mitigations of these attacks will be to increase the prevalence of DNSSEC validation by resolvers and DNSSEC-signing of DNS data.

As part of the strategy to reduce the effects of attacks on the RSS, increased efforts will be undertaken to encourage resolver operators to enable DNSSEC validation, to work with resolver software developers and vendors to enable DNSSEC validation by default, and to provide training and capacity building in the area of DNSSEC configurations and operations to the community.

4.3 Confidentiality Attacks and Mitigation

Confidentiality attacks aim to expose sensitive information. The risk to the RSS of these forms of attacks is limited since the DNS does not generally rely upon secrets. Due to increased concerns about privacy, however, the fact that the original DNS protocol suite transmits and receives data without encryption provides for information leakage that can be seen as a breach of confidentiality.

Mitigation of confidentiality attacks typically revolves around ensuring that the stream of DNS queries and responses is encrypted. ICANN org has provided funding to several organizations implementing various technologies aimed at improving the privacy of the DNS that have been or are being standardized in the Internet Engineering task force (IETF). Deployment of these technologies in the context of the root servers can mitigate confidentiality attacks. However, given the need for backwards compatibility, it is likely that deployment of confidentiality in the root server system will take a significant time.

4.3.1 Encouraging Implementation and Deployment of Technologies to Reduce DNS Data Leakage

The DNS protocol was not designed to provide for the confidentiality of queries or responses. The DNS namespace hierarchy is assumed to be public, as well as the data associated with the names in that hierarchy. The recent increased interest in privacy, however, has resulted in concerns about the lack of confidentiality in the DNS.

As an outcome of this increased interest in privacy, efforts focused on increasing the privacy of the DNS have resulted in standards from the IETF to add encryption between users' systems and resolvers. These standards may eventually help reduce the effects of confidentiality attacks on the RSS. Specifically, in the case where the query/response streams to the root servers are subject to eavesdropping, the deployment of privacy-enhancing mechanisms that may be standardized in the future would mitigate the risk.

On a different front, the IETF is working on making the "QNAME minimization" operational guidance into a standard.[7] Query name minimization reduces the amount of personally identifiable information that may appear at the root servers by sending only part of the full DNS query that the root servers know about, typically the name servers for top level domains, to the root server. This feature has recently been made the default operating configuration for some resolver software packages. ICANN org can encourage deployment of this operational practice in resolvers and through measurements of its prevalence at the IMRS.

5 Conclusion

The IMRS strategy described in this document, being comprehensive in nature and impacting the ICANN org, the ICANN community, and the Internet as a whole, will require careful planning, significant resources, particularly of ICANN org, and care during implementation. ICANN's goals of greater availability of root service and of reducing the effects of attacks on the root system will continue to be important, given the importance of the DNS and the role that the RSS plays within it. Since ICANN is committed through its Bylaws to ensure the security, stability, and resilience of the DNS, these strategies are necessary as the value of, and risks to, the root system continue to evolve.

[7] See https://datatracker.ietf.org/doc/draft-ietf-dnsop-rfc7816bis/