Hyperlocal Root Zone Technical Analysis

ICANN Office of the Chief Technology Officer

Roy Arends and Nicolas Antoniello
OCTO-027
25 August 2021

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.

This document supports ICANN's strategic goal to improve the shared responsibility for upholding the security and stability of the DNS by strengthening DNS coordination in partnership with relevant stakeholders. It is part of ICANN's strategic objective to strengthen the security of the domain name system and the DNS root server system.

1 Executive Summary

This document provides a technical analysis of hyperlocal root service. Hyperlocal root service is an approach to make root zone content more available to resolvers by making its content local, one approach of which is described in “Running a Root Server Local to a Resolver,” RFC 8806.[1]

This technical analysis is laid out in several sections that signify the key differences between the current root server system (RSS) and a hyperlocal root zone deployment. Each section contains a discussion on the benefits and drawbacks of those deployments.

The primary benefits of a hyperlocal root zone are that access to the root zone is more reliable and private, and responses are faster (it has lower latency).

The drawbacks of this approach are that it requires additional configuration of an already complex resolver system, which can impact root zone scalability, and reduce telemetry for the purpose of operations, analysis, and research by root server operators (RSOs) and others.

2 Introduction

The Domain Name System (DNS) can be seen as comprising two independently operable parts: a publication and a lookup part. The publication part is the hierarchy of names from the root down, in which the registrant ultimately decides what information is published, e.g., the IP address or mail exchanger associated with a domain name. The lookup part, also known as the resolution, is typically performed by software known as a recursive resolver. These resolvers are usually operated by network operators for their users, but some organizations, such as Google, Cloudflare, Cisco, UncensoredDNS, and others provide public resolution services unassociated with the users’ network operations. In contrast to the publication side of the DNS, in which ICANN delegates authority to administer portions of the DNS name space to registry operators (top-level domains), ICANN generally does not have a relationship with resolver operators. However, resolver operators that have enabled the DNS Security Extensions (DNSSEC) depend on ICANN for the root zone “trust anchor” (the public portion of the root zone’s key signing key), and the root hints (the addresses of the root servers), both of which are shipped with all major resolver implementations today. Importantly, since the root zone contains top-level domain referral information, DNS resolution depends on the secure and stable operation of the root of the DNS.

Information fetched by resolvers from the root zone is critical to the operation of the DNS: without information about the name servers of top-level domains obtained from the root zone, resolution of any domain names is impossible. Almost twenty years ago, some resolver operators observed that since the root zone is relatively small and critically important, it could be copied locally, thereby reducing the amount of time root lookups take and ensuring the ability to

[1] See https://www.rfc-editor.org/info/rfc8806

do lookups even if the root servers are unavailable.[2,3] This approach to ensure root zone content availability for resolvers is known as hyperlocal root service.

The Oxford English Dictionary defines hyperlocal as "relating to or focusing on matters concerning a small community or geographical area."[4] While originally used in the context of local news and weather forecast provisioning, this term has been applied to provisioning data pertaining to locally used applications. Thus, the term "hyperlocal root service" is intended to convey the concept of making the root zone available locally by a recursive resolver.

OCTO-016, "ICANN's Root Name Service Strategy and Implementation"[5] contains a section on "Supporting Root Service Decentralization with Hyperlocal." The inclusion of hyperlocal root service is aimed at supporting ICANN's mission of "ensuring the security and stability" of the DNS and, in particular, facilitating "openness, interoperability, resilience, security and/or stability of the DNS."

In RSSAC045, "RSSAC Statement on Threat Mitigation for the Root Server System,"[6] the ICANN Root Server System Advisory Committee (RSSAC) and the RSOs have acknowledged that threats to the RSS include denial of service (DoS) attacks on network bandwidth, CPU, and memory consumption. In "RSSAC Statement on Threat Mitigation for the Root Server System,"[7] RSOs stated that "Availability and data integrity of the root zone are currently the primary concerns of the Root Server System," and this position was formally endorsed by RSSAC. These concerns underscore the need for solutions to address the risk of DoS attacks at the root, of which the broader root zone distribution system beyond the widely distributed anycast root servers currently deployed**,** should be considered.

One motivation for the technology discussed in this paper is its use to defend against the threats of reduced availability of the root zone by making the root zone's content available locally to the millions of resolvers on the Internet. This technology is not new, [8,9] and some large public resolver operators such as Cisco's OpenDNS, UncensoredDNS, and hundreds of others already make use of this approach, despite it not having been documented by the Internet

---

[2] "As for making your local resolver a slave for the root zone, that suggestion has some merit", Doug Barton on the FreeBSD-stable list, January 2003, https://groups.google.com/g/fa.freebsd.stable/c/wPmILQUpz4E/m/LDWgwhHxipIJ

[3] "FreeBSD default configuration is now to slave the root zone from the root name servers," Doug Barton on his motivation to change the default FreeBSD configuration, August 2007, https://lists.dns-oarc.net/pipermail/dns-operations/2007-August/001858.html

[4] Definition of the term hyperlocal, Lexico, accessed July 2021, https://www.lexico.com/definition/hyperlocal

[5] See https://www.icann.org/octo-016-en.pdf

[6] See https://www.icann.org/en/system/files/files/rssac-045-03dec19-en.pdf

[7] Root Server Operators, "Threat Mitigation for the Root Server System." 19 August 2019. https://root-servers.org/media/news/Threat_Mitigation_For_the_Root_Server_System.pdf

[8] "Paul Mockapetris, chief scientist at Nominum Inc. and author of the original DNS protocol specifications, recently suggested that DNS operators keep a current copy of root zones in order to isolate themselves from future root-server attacks.", Computer World, February 2003, https://www.computerworld.com/article/2579981/how-to-use-a-personal-dns-for-root-server-attack-isolation.html

[9] Malone, David, "The root of the matter: hints or slaves," Communications Network Research Institute Dublin Institute of Technology, 25–27,October 2004, http://conferences.sigcomm.org/imc/2004/papers/p15-malone.pdf

Engineering Task Force (IETF) before “Decreasing Access Time to Root Servers by Running One on Loopback,” RFC 7706,[10] which was the precursor to RFC 8806.

Hyperlocal root service can be seen to be indirectly supported by “A Proposed Governance Model for the DNS Root Server System,“ RSSAC037,[11] which states that “Architectural changes should result from technical evolution and demonstrate technical need. RSOs should embrace emerging technologies affecting the RSS, as long as the Internet’s globally unique public namespace is preserved.”

Hyperlocal root service moves root service to the party that has the most incentive to operate that service: the resolver operator. The resolver operator, in order to keep their customers happy, has sole authority to ensure their service is upgraded as demands require.

3 Technical Analysis

This section describes the benefits and drawbacks of the current system of root servers and compares that system with the benefits and drawbacks of a hyperlocal root zone on the areas of privacy, availability, latency, integrity, telemetry, and timeliness.

The existing root server system has operated without noticeable disruption for end users since the creation of the DNS. The description of drawbacks of the current system is merely an observation of the effects of deploying a decades old protocol and are independent of and unrelated to RSOs.

There is no intent within this document to suggest that existing RSOs are flawed in any way.

3.1 Privacy

As “Privacy Considerations for Internet Protocols,” RFC 6973[12] explains, DNS servers are enablers (a protocol entity that facilitates communication between an initiator and a recipient), which become observers (an entity that is able to observe and collect information from communications, potentially posing privacy threats) when they start collecting data. Many programs exist to collect and analyze DNS data.[13] This data is often kept for a long time and distributed to third parties for research purposes.[14] DNS data is also collected passively at observation points (passive DNS) by commercial enterprises such as Farsight Security and RiskIQ for security research and services, as described in “DNS Privacy Considerations,” RFC 7626.[15] A number of high-level analysis techniques can be used to quickly extract interesting information from DNS traffic. These techniques expose data about underlying networks and

[10] See https://www.rfc-editor.org/info/rfc7706
[11] See https://www.icann.org/en/system/files/files/rssac-037-15jun18-en.pdf
[12] See https://www.rfc-editor.org/info/rfc6973
[13] "NSA's MORECOWBELL: Knell for DNS," (Christian Grothoff et.al., GNUnet Git Repositories, July 2021), https://git.gnunet.org/bibliography.git/plain/docs/mcb-en.pdf
[14] Castro, S., Wessels, D., Fomenkov, M., & Claffy, K. (September 30, 2008). “A day at the root of the internet.” ACM Sigcomm Computer Communication Review, 38, 5, 41-46. http://www.sigcomm.org/sites/default/files/ccr/papers/2008/October/1452335-1452341.pdf
[15] See https://www.rfc-editor.org/info/rfc7626

their users and are highly useful for reconnaissance to determine what attacks to launch.[16] There is little transparency or consistency surrounding how DNS query data is collected, stored, processed, analyzed, used, shared, and sold.[17]

Architecturally, the RSS as a whole is in a position to see a significant fraction of all queries on the public DNS. Historically, those queries included the full domain being resolved and always included the source IP address of the resolver issuing the query. As the DNS is increasingly implemented, there are various techniques to limit exposure to authoritative servers, such as those described in "Aggressive Use of DNSSEC-Validated Cache," RFC 8198,[18] and "DNS Query Name Minimisation to Improve Privacy," RFC 7816,[19] implementation of these techniques is not ubiquitous. As of May 2021, about 35% of resolvers observed at the root use query name minimization (QNAME minimization).[20]

For hyperlocal root service, a query not sent is a query that cannot be collected, stored, processed, analyzed, used, shared, or sold, and as a result, "Recommendations for DNS Privacy Service Operators," RFC 8932,[21] recommends "Run a local copy of the root zone [RFC8806] to avoid making queries to the root servers that might leak information."

3.2 Availability

Root zone content is crucial for the operation of resolvers. When root zone content becomes unavailable and cached entries expire, resolvers cannot function.

In the current system, availability of root service is extremely high, due to the heavy use of anycast. If one anycast instance fails to function, it has little-to-no impact on the rest of the anycast instance for that specific root server. As a result, DoS attacks on the root server's IP addresses are contained to the anycast catchment of the instances that use those IP addresses; a catchment is the set of IP prefixes that are routed to a particular anycast site. Resolvers in that catchment may experience a small delay, as now fewer of the root servers are available to those resolvers and queries may time out before the resolvers select a new root server to query. Even in the very unlikely event that all root servers become unavailable, some resolvers have special strategies to temporarily cope with potentially unavailable information, such as "serve-stale" as described in "Serving Stale Data to Improve DNS Resiliency," RFC 8767,[22] to use previously cached root-zone content after time-to-live (TTL) expiration. These are not used by default, and these strategies are not without risk; RFC 8767 states, "The most obvious security issue is the increased likelihood of DNSSEC validation failures when using stale data because signatures could be returned outside their validity period. Additionally, bad actors have been known to use DNS caches to keep records alive even after their authorities have gone away."

---

[16] W.Hardaker, "Analyzing and Mitigating Privacy with the DNS Root Service." 2018, https://www.isi.edu/%7ehardaker/papers/2018-02-ndss-analyzing-root-privacy.pdf

[17] Bradshaw, S., & DeNardis, L. (March 01, 2019). "Privacy by Infrastructure: The Unresolved Case of the Domain Name System. Policy and Internet," 11, 1, 16-36, https://onlinelibrary.wiley.com/doi/abs/10.1002/poi3.195

[18] See https://www.rfc-editor.org/info/rfc8198

[19] See https://www.rfc-editor.org/info/rfc7816

[20] See metric M3.6 from https://ithi.research.icann.org/graph-m3.html

[21] See https://www.rfc-editor.org/info/rfc8932

[22] See https://www.rfc-editor.org/info/rfc8767

A natural reaction to the increasing risk of DoS attacks would be to increase the availability of root servers by each RSO scaling their anycast constellation. However, this obviously comes at a real cost in hardware, bandwidth and management. Worse, since the cost to attackers is typically minimal, increasing expenditures to scale anycast deployment risks becoming an unwinnable arms race by the RSOs against state-of-the-art DOS scaling.

The root zone changes regularly, and it is essential that a copy is kept up to date. Serving stale data leads to failures in resolvers.[23] This is just as essential for hyperlocal root service deployments as it is with the existing RSS however, hyperlocal root service deployments will have significantly less users than the existing root servers and, thus the impact of a stale hyperlocal root zone deployment is limited to its users. The onus of monitoring hyperlocal root service is on the operator of the resolver who may not be as well versed in obtaining and serving root zones as RSOs in these functions.

With hyperlocal root service, the root zone must be consistently and continuously available for retrieval. The Internet Assigned Numbers Authority's (IANA's) naming function is responsible for management of the DNS root zone.[24] The complete root zone is available for download from the Internic website and the FTP site.

The Root Zone Maintainer[25] makes the root zone available to the RSOs. Also, some RSOs allow for zone transfers from their root servers. At the time of writing, these are:

- University of Southern California, Information Sciences Institute (USC-ISI) (b.root-servers.net)
- Cogent Communications (c.root-servers.net)
- University of Maryland (UMD) (d.root-servers.net)
- National Aeronautics and Space Administration (NASA) Ames (e.root-servers.net)
- Internet Systems Consortium (ISC) (f.root-servers.net)
- Defense Information Systems Agency, U.S. Government (DISA) (g.root-servers.net)
- Réseaux IP Européens Network Coordination Centre (RIPE NCC) (k.root-servers.net)

Additionally, some RSOs, such as ISI (at https://localroot.isi.edu) and ICANN (at https://www.dns.icann.org/services/axfr/) provide alternative locations from which the root zone can be retrieved.

In theory, hyperlocal root zone provisioning can make use of existing content delivery networks (CDNs), as the provisioning method is not limited to DNS zone transfers, but can also be delivered over HTTP(S) if the resolver supports it. While some CDNs can provide DNS zone transfer via anycast instances, all CDNs can provide HTTP(S). However, not all resolver implementations can currently make use of HTTP(S) as a way to get the root zone. A hyperlocal root zone provisioning system can make use of existing generic CDNs, but only when HTTP(S) is supported in resolvers. Such a system would likely be more cost effective than building a dedicated root zone provisioning system.

---

[23] "Root zone on c.root-servers.net is obsoleted and this causes a lot of DNSSec-related errors." lists.dns-oarc.net, August 2019, https://lists.dns-oarc.net/pipermail/dns-operations/2019-August/019119.html

[24] Root Zone Management, IANA, July 2021, https://www.iana.org/domains/root

[25] Root Zone Maintainer Agreement (RZMA), ICANN org, September 2016, https://www.icann.org/iana_imp_docs/129-root-zone-maintainer-service-agreement-v-28sep16

3.3 Latency

When a query is sent to the root, it is often the very first query in the resolution process when resolving a domain name. This query blocks subsequent steps in the resolution process until it is answered. Different resolver configurations might send a large volume of queries to the RSS for various reasons; if those resolvers instead were using a hyperlocal root zone, they would respond to their stub resolver clients more quickly.

At the time of this writing, there are over a thousand anycast instances of root servers deployed around the world. The latency between resolvers and root servers is generally under 100 milliseconds.[26] Aggressive use of a DNSSEC-validated cache reduces the amount of time it takes for a recursive resolver to respond to queries from stub resolvers for non-existent domains.

Deploying hyperlocal root service will avoid sending these queries to root servers and will optimize the overall speed of resolving names. Research shows that the deployment of hyperlocal root service saves time and improves throughput.[27] Research on the use of root servers in Latin America concluded that hyperlocal root service deployment will improve the user experience in Internet browsing.[28]

3.4 Integrity

In the current system, responses from root servers contain DNS records that are not signed, such as delegation point nameserver (NS) records and glue address records. Delegation point NS records and glue records are not authoritative in a zone. NS records are signed in the child zone. Glue records should be signed in the zone they reside as authoritative.

Because the transport between resolvers and root servers is not secure, this unsigned data presents an attack vector to poison a resolver's caches with false data. The resolver has no way to check the integrity of unsigned records in the root zone. While this attack has not been identified to date, as other parts of the Internet infrastructure are hardened, it is possible that more obscure attacks such as this may be attempted.

The deployment of hyperlocal root service involves retrieving the root zone file regularly from a set of provisioning servers, applying potentially out-of-band integrity checks or transport security between the provisioning servers and the resolver, with the potential ability to fall back to regular root servers if any element in this process fails.

Transport Layer Security (TLS) provides integrity checks during transport. However, trust in those integrity checks is ultimately provided by the certificate issuer (for TLS), not by the root

[26] J.H. Kuipers. "Latency across root servers is pretty good," RIPE Labs, May 2017, https://labs.ripe.net/Members/jh_kuipers/anycast-latency-how-many-sites-are-enough

[27] B. Rajendran, et. al. "Reducing RTT of DNS Query Resolution using RFC 8806," Centre of Excellence for DNS Security, October 2020, https://coednssecurity.in/pdf/ReducingRTTofDNSQueryResolution-V1.pdf

[28] H. Salgado. "Use of DNS Root Servers in Latin America," Latin American and Caribbean Internet Addresses Registry (LACNIC), October 2019, https://www.lacnic.net/innovaportal/file/1031/1/informe-rootservers-lac-en.pdf

zone manager. Additionally, the integrity provided by TLS, based on the channel of communication, is transient. There are no guarantees that the root zone file is intact before or after transport, when the root zone is "in situ" waiting to be used locally or transferred elsewhere.

Similarly, Transaction Signatures (TSIG) such as those provided by ISI's LocalRoot project requires a shared secret between the authoritative server and the resolver. In the worst case, the management of the secret keys at the authoritative servers given to the root zone would need to scale to the number of resolvers that would be fetching the zone, which may be infeasible should that number grow significantly.

The only cryptographic method currently available to verify the integrity of the root zone is PGP signatures. This method is currently provided by the Root Zone Maintainer. However, as currently deployed, this integrity check has limitations. At this time, the root zone integrity proof is a SHA1 digest over the root zone, signed by a 1024-bit DSA key. However, the "NIST Policy on Hash Functions"[29] recommends against using SHA-1, while NIST's "Digital Signature Standard (DSS)"[30] says that DSA is no longer approved for signing. Trust in the PGP key of the Root Zone Maintainer needs to be established manually.

The IETF has recently extended the DNS with "Message Digest for DNS Zones," RFC 8976.[31] The protocol commonly referred to as "ZONEMD" generates a cryptographic digest over the contents of a zone file and signs this digest using the zone signing key. This signature is part of a DNSSEC signed zone and makes integrity provision an inherent part of the zone file. Running code exists for ZONEMD, and there is a commitment by the larger resolver developers to implement this new standard. It is expected that the ZONEMD record will be deployed in the root zone in the near future.

Assuming implementation of ZONEMD in the root zone, hyperlocal root service deployment would be able to check the integrity of the unsigned records in the root zone. Since the hyperlocal root service is deployed at or close to the resolver, the attack potential on unsigned root zone content would virtually disappear.

3.5 Telemetry

The root servers are a vantage point for useful telemetry including telemetry that does not infringe on privacy, such as the deployment uptake of modern features known as "DNS Cookies," RFC 7873,[32] the ratio between TCP and UDP queries, the ratio between DNS queries over IPv4 and IPv6, etc. This data can be helpful in understanding the inner workings of the DNS at a global scale.

Hyperlocal root service deployment reduces telemetry at the root. In particular, when resolvers using hyperlocal root service stop sending queries to root servers, there is a reduction in telemetry used to inform the Key Signing Key (KSK) rollover performed in 2018. This telemetry

[29] National Institute of Standards and Technology, August 2015, https://csrc.nist.gov/Projects/Hash-Functions/NIST-Policy-on-Hash-Functions
[30] National Institute of Standards and Technology "FIPS 186-5 (Draft)," October 2019, https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-5-draft.pdf
[31] See https://www.rfc-editor.org/info/rfc8976
[32] See https://www.rfc-editor.org/info/rfc7873

includes “Signaling Trust Anchor Knowledge in DNSSEC,” RFC 8145,[33] and “Signaling Cryptographic Algorithm Understanding in DNSSEC,” RFC 6975.[34] However, in practice, the latter is not deployed at all, and while trust anchor signals provided data for the 2018 KSK rollover, a major challenge was the sparse and distorted telemetry from resolvers. Other telemetry will be reduced as well, such as the aforementioned deployment uptake of modern features known as DNS cookies, the ratio between TCP and UDP queries, or the ratio between IPv4 and IPv6 DNS queries.

The implication of the reduction of telemetry is that root server traffic becomes less useful as an indicator of features deployed by resolvers. In many cases there are alternative sources of telemetry that can be used as indicators, such as resolver traffic,[35] or traffic to top level domains; however, one likely consequence of significant hyperlocal root service deployment will be a general decrease in knowledge about how the global DNS operates.

3.6 Timeliness

The root zone maintainer publishes two or more versions of the root zone every day and has on occasion published as many as five versions on a single day.

Due to the “loose coherency” of the DNS and the way DNS zone replication is handled in the current system of root servers, the root zone data is mostly up to date, though regularly, instances may be minutes or hours behind the latest version, as shown by data collected from the RSOs in fulfilment of “RSSAC Advisory on Measurements of the Root Server System,” RSSAC002.[36] In one documented instance, a root server served a stale root zone for a few days.[37]

Depending on implementation, hyperlocal root service deployments follow the refresh timing from the SOA record, retry, and expire mechanism described in “Domain Names - Implementation and Specification,” RFC 1035,[38] or will prime their cache regularly with a complete root zone. Assuming the root zone provisioning system is able to provide the root zone soon after the root zone maintainer publishes it and the network operator becomes aware of the updated root zone, the fundamental difference between them, depending on root servers and a hyperlocal root service deployment is that the network operator has more control over timely deployment of the root zone in the resolver(s) under the operator’s control. The standard mechanism in the DNS for prompt notification of zone changes is the DNS NOTIFY protocol, described in “A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY),” RFC 1996.[39] A future root zone provisioning system should support signup for DNS NOTIFY service as has been done with ISI’s Local Root project.

[33] See https://www.rfc-editor.org/info/rfc8145
[34] See https://www.rfc-editor.org/info/rfc6975
[35] “DNS Recursive Server Analysis” ICANN org, July 2021, https://ithi.research.icann.org/graph-m4.html
[36] See https://www.icann.org/en/system/files/files/rssac-002-measurements-root-06jun16-en.pdf
[37] P. Vixie, “Problem with c.root-servers.net,” https://lists.dns-oarc.net/pipermail/dns-operations/2019-August/019128.html
[38] See https://www.rfc-editor.org/info/rfc1035
[39] See https://www.rfc-editor.org/info/rfc1996

3.7 Root Zone Scalability

The limited number of players in the existing RSS and the close coordination between those players suggests that increasing the scale of the infrastructure would be feasible. In comparison, given that hyperlocal root service is independently and unilaterally deployed by resolver operators, coordinating an increase in infrastructure scale would likely be intractable.

For example, as ICANN org proceeds to deploy additional new top level domains, the size of the root zone will continue to grow. At some point in the future, root zone size could result in resolvers being unable to load the root zone and thereby unable to resolve names. As it is difficult to establish an upper bound at which point a significant number of resolvers would have difficulty, or even what form that difficulty may take, growth of the root zone would likely be increasingly conservative.

4 Conclusion

Hyperlocal root service is a method that involves running a root zone locally to the resolver as an alternative to using root servers.

Hyperlocal root service does not send queries for the root zone to root servers. This provides improved privacy as these queries cannot be collected. However, modern resolvers already have the ability to reduce the amount of information in a single query such as DNS QNAME minimization, and reduce the number of queries such as Aggressive Use of DNSSEC-Validated Cache.

Queries for the root zone are answered locally in the hyperlocal root service. This provides reduced latency for responses to these queries. However, modern resolvers have optimized the frequency of queries for the root zone by cleverly making use of the aforementioned Aggressive Use of DNSSEC-Validated Cache. Hence, the provided lower latency would likely have only a small additional benefit and it is unlikely that the end user will experience a noticeable difference.

Currently, delegation point NS records and glue address records do not have DNSSEC signatures. As a result, responses containing these records can be spoofed. In a hyperlocal root service deployment, there are opportunities for an integrity check on the entire zone. Consequently, the window of opportunity to spoof these responses is limited to the path between the validating resolver and the hyperlocal root service deployment, which is very small. However, this root-zone integrity check is currently not fully implemented in resolvers. This will improve once the ZONEMD protocol is implemented and deployed.

In hyperlocal root service, control over obtaining a fresh copy of the root zone is the responsibility of the operator of the resolver. This additional control requires additional management and monitoring by the operator.

Running a hyperlocal root zone is not a trivial exercise. While there are definitely privacy benefits and more control for the resolver operator, it comes with additional responsibilities for the resolver operator.