

Independent Submission  
Request for Comments: 7108  
Category: Informational  
ISSN: 2070-1721

J. Abley  
Dyn, Inc.  
T. Manderson  
ICANN  
January 2014

A Summary of Various Mechanisms Deployed at L-Root for the  
Identification of Anycast Nodes

Abstract

Anycast is a deployment technique commonly employed for authoritative-only servers in the Domain Name System (DNS). L-Root, one of the thirteen root servers, is deployed in this fashion.

Various techniques have been used to map deployed anycast infrastructure externally, i.e., without reference to inside knowledge about where and how such infrastructure has been deployed. Motivations for performing such measurement exercises include operational troubleshooting and infrastructure risk assessment. In the specific case of L-Root, the ability to measure and map anycast infrastructure using the techniques mentioned in this document is provided for reasons of operational transparency.

This document describes all facilities deployed at L-Root to facilitate mapping of its infrastructure and serves as documentation for L-Root as a measurable service.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7108>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

- 1. Introduction . . . . . 2
- 2. Conventions Used in This Document . . . . . 3
- 3. Naming Scheme for L-Root Nodes . . . . . 3
- 4. Identification of L-Root Nodes . . . . . 3
  - 4.1. Use of NSID . . . . . 4
  - 4.2. Use of HOSTNAME.BIND/CH/TXT . . . . . 5
  - 4.3. Use of ID.SERVER/CH/TXT . . . . . 6
  - 4.4. Use of IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT and .../IN/A . 6
  - 4.5. Use of NODES.L.ROOT-SERVERS.ORG/IN/TXT . . . . . 8
- 5. Provisioning of IDENTITY.L.ROOT-SERVERS.ORG . . . . . 9
- 6. Security Considerations . . . . . 9
- 7. Acknowledgements . . . . . 10
- 8. References . . . . . 10
  - 8.1. Normative References . . . . . 10
  - 8.2. Informative References . . . . . 11

1. Introduction

The Domain Name System (DNS) is described in [RFC1034] and [RFC1035]. L-Root, one of the thirteen root servers, is deployed using anycast [RFC4786]; its service addresses, published in the A and AAAA Resource Record (RR) Sets for "L.ROOT-SERVERS.NET", are made available from a substantial number of semi-autonomous servers deployed throughout the Internet. A list of locations served by L-Root can be found at [ROOT-SERVERS].

Various techniques have been used to map deployed anycast infrastructure externally, i.e., without reference to inside knowledge about where and how such infrastructure has been deployed. Motivations for performing such measurement exercises include operational troubleshooting and infrastructure risk assessment. In the specific case of L-Root, the ability to measure and map anycast infrastructure using the techniques mentioned in this document is provided for reasons of operational transparency.

This document describes all facilities currently provided at L-Root to aid node identification.

## 2. Conventions Used in This Document

This document contains several examples of commands typed at a Unix (or Unix-like) command line to illustrate use of the various mechanisms available to identify L-Root nodes. Such examples are presented in this document with lines typed by the user preceded by the "%" prompt character; a bare "%" character indicates the end of the output resulting from the command.

In some cases, the output shown in examples is too long to be represented directly in the text. In those cases, a backslash character ("\") is used to indicate continuation.

## 3. Naming Scheme for L-Root Nodes

Individual L-Root nodes have structured hostnames that are constructed as follows:

```
<IATA Code><NN>.L.ROOT-SERVERS.ORG
```

where

- o <IATA Code> is chosen from the list of three-character airport codes published by the International Air Transport Association (IATA) in the IATA Airline Coding Directory [ACD]; and
- o <NN> is a two-digit numeric code used to distinguish between two different nodes in the vicinity of the same airport.

Where multiple airports exist in the vicinity of a single L-Root node, one is arbitrarily chosen.

More granular location data published for L-Root nodes (e.g., see Section 4.4) is derived from the location of the airport, not the actual location of the node.

## 4. Identification of L-Root Nodes

L-Root service is provided using a single IPv4 address (199.7.83.42) and a single IPv6 address (2001:500:3::42). Note that it is preferable to refer to the service using its DNS name (L.ROOT-SERVERS.NET) rather than literal addresses, since addresses can change from time to time.

At the time of writing, there are 273 separate name server elements ("nodes") deployed in 143 locations: together, these nodes provide L-Root service. A DNS query sent to an L-Root service address will be routed towards exactly one of those nodes for processing, and the corresponding DNS response will be originated from the same node. Queries from different clients may be routed to different nodes. Successive queries from the same client may also be routed to different nodes.

The following sections provide a summary of all mechanisms provided by L-Root to allow a client to identify which L-Root node is being used.

Using `HOSTNAME.BIND/CH/TXT` (Section 4.2), `ID.SERVER/CH/TXT` (Section 4.3), or `IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT` or `IDENTITY.L.ROOT-SERVERS/IN/A` (Section 4.4) to identify a node for the purposes of reporting a problem is frequently reasonable, but it should be acknowledged that there is potential for re-routing between successive queries: an observed problem might relate to one node, whilst a subsequent query using one of those three techniques could be answered by a different node. Use of the Name Server Identifier (NSID) option on the precise queries that yield problematic responses can obviate this possibility (see Section 4.1).

#### 4.1. Use of NSID

L-Root supports the use of the Name Server Identifier (NSID) option [RFC5001] to return the identity of an L-Root node along with the response to a DNS query. The NSID payload of such responses is the fully qualified hostname of the responding L-Root node.

The NSID option allows the identification of a node sending a specific, requested response to the client. This is of particular use if (for example) there is a desire to identify unequivocally what node is responding with a particularly troublesome response; the output of the diagnostic tool "dig" with NSID requested provides the problem response with the node identification, and its output in that case could form the basis of a useful trouble report.

NSID is specified as an EDNS(0) option [RFC6891]. Clients that do not support EDNS(0) signaling (or depend on other systems that do not support EDNS0) may find this mechanism unavailable.

The NSID option can be specified using the widely used diagnostic tool "dig" using the "+nsid" option, as shown below. Note that long lines have been truncated for the purposes of this document ("\" at the end of a line indicates continuation).

```

% dig -4 @L.ROOT-SERVERS.NET . SOA +nsid \
+norec +noall +comments
; <<>> DiG 9.6.-ESV-R3 <<>> -4 @L.ROOT-SERVERS.NET . SOA +nsid \
+norec +noall +comments
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 14913
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 13, ADDITIONAL: 23

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
; NSID: 79 74 7a 30 31 2e 6c 2e 72 6f 6f 74 2d 73 65 72 76 65 72 73 \
2e 6f 72 67 (y) (t) (z) (0) (1) (.) (1) (.) (r) (o) (o) (t) (-) \
(s) (e) (r) (v) (e) (r) (s) (.) (o) (r) (g)
%

% dig -6 @L.ROOT-SERVERS.NET . SOA +nsid \
+norec +noall +comments
; <<>> DiG 9.6.-ESV-R3 <<>> -6 @L.ROOT-SERVERS.NET . SOA +nsid \
+norec +noall +comments
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 33374
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 13, ADDITIONAL: 23

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
; NSID: 79 74 7a 30 31 2e 6c 2e 72 6f 6f 74 2d 73 65 72 76 65 72 73 \
2e 6f 72 67 (y) (t) (z) (0) (1) (.) (1) (.) (r) (o) (o) (t) (-) \
(s) (e) (r) (v) (e) (r) (s) (.) (o) (r) (g)
%

```

#### 4.2. Use of HOSTNAME.BIND/CH/TXT

L-Root supports the use of HOSTNAME.BIND/CH/TXT queries to return the identity of an L-Root node. The TXT RDATA returned is the fully qualified hostname of the responding L-Root node.

The HOSTNAME.BIND/CH/TXT convention is described in [RFC4892].

```
% dig -4 @L.ROOT-SERVERS.NET HOSTNAME.BIND CH TXT +short
"ytz01.l.root-servers.org"
%

% dig -6 @L.ROOT-SERVERS.NET HOSTNAME.BIND CH TXT +short
"ytz01.l.root-servers.org"
%
```

#### 4.3. Use of ID.SERVER/CH/TXT

L-Root supports the use of ID.SERVER/CH/TXT queries to return the identity of an L-Root node. The TXT RDATA returned is the fully qualified hostname of the responding L-Root node.

ID.SERVER/CH/TXT functions identically (apart from the QNAME) to HOSTNAME.BIND/CH/TXT, as discussed in Section 4.2. The discussion there relating to the possibility of re-routing between successive queries also follows for ID.SERVER/CH/TXT.

The ID.SERVER/CH/TXT convention is described in [RFC4892].

```
% dig -4 @L.ROOT-SERVERS.NET ID.SERVER CH TXT +short
"ytz01.l.root-servers.org"
%

% dig -6 @L.ROOT-SERVERS.NET ID.SERVER CH TXT +short
"ytz01.l.root-servers.org"
%
```

#### 4.4. Use of IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT and .../IN/A

The operator of L-Root has distributed a separate DNS service in parallel with L-Root, operating on precisely the same set of nodes but listening on addresses that are different from the L-Root service addresses. Measurements of this separate service should give results that are representative of L-Root. Further discussion of this service can be found in Section 5.

The fully qualified DNS name IDENTITY.L.ROOT-SERVERS.ORG (note the use of ORG, not NET) has associated TXT and A RR Sets that are unique to the responding node. Clients are hence able to issue queries for IDENTITY.L.ROOT-SERVERS.ORG/IN/A and IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT and use the results both to identify individual nodes and to distinguish between responses generated by different nodes.

The TXT record returned in the response to such queries is structured as follows:

1. The fully qualified hostname of the node responding to the query;
2. The city in which the node is located;
3. The region in which the node is located, if applicable;
4. The economy in which the node is located (in most cases, the name of a country); and
5. The Internet Corporation for Assigned Names and Numbers (ICANN) region in which the node is located. A list of ICANN regions at the time of writing can be found at <http://meetings.icann.org/regions>.

The A record returned in the response to such queries is guaranteed to be unique to the responding node. The A RRType was chosen in an effort to make the use of this mechanism as widely available to client environments as possible, and the ability to map a hostname to an IPv4 address seemed more likely to be widespread than the mapping of a hostname to any other value. It should be noted that the availability of this mechanism to any particular client is orthogonal to the local availability of IPv4 or IPv6 transport.

In this case, because identity data is published using IN-class resource records, it is not necessary to send queries directly towards L-Root in order to obtain results. Responses can be obtained through recursive servers, the responses in those cases being the identity of L-Root as observed through the recursive server used rather than the "closest" L-Root node to the client. This facilitates some degree of remote troubleshooting, since a query for IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT or IDENTITY.L.ROOT-SERVERS.ORG/IN/A directed a remote recursive resolver can help illustrate which L-Root node is being used by that server (or was used when the cache was populated).

A related caching effect is that responses to IDENTITY.L.ROOT-SERVERS.ORG/IN/A and IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT may be cached at different times, and may hence persist in a cache for overlapping periods of time. One possible visible effect is that the responses to IDENTITY.L.ROOT-SERVERS.ORG/IN/A and IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT as presented from a cache may appear to be incoherent (i.e., refer to different nodes) despite queries against of the cache happening (near) simultaneously. Caches may also discard the published Times to Live (TTLs) in responses from the authoritative

server and replace them with longer TTLs, as a matter of local policy. Interpretation of responses for these queries from caches should therefore be carried out with these possible effects in mind.

It has been observed that IDENTITY.L.ROOT-SERVERS.ORG/IN/A queries offer a useful mechanism for troubleshooting DNS problems with non-technical users, since such users can often be walked through the process of looking up an A record (e.g., as a side effect of utilities such as ping) far easier than they can be instructed on how to use DNS-specific tools such as dig.

```
% dig IDENTITY.L.ROOT-SERVERS.ORG TXT +short
"ytz01.l.root-servers.org" "Toronto" "Ontario" "Canada" "NorthAmerica"
%
```

```
% dig IDENTITY.L.ROOT-SERVERS.ORG A +short
67.215.199.91
%
```

#### 4.5. Use of NODES.L.ROOT-SERVERS.ORG/IN/TXT

The fully qualified DNS name NODES.L.ROOT-SERVERS.ORG (note again the use of ORG, not NET) provides multiple TXT RRs, one per node, and represents the effective concatenation of all possible responses to the query IDENTITY.L.ROOT-SERVERS.ORG/IN/TXT.

Note that in the example below we have forced dig to send the query over TCP, since we expect the response to be too large for UDP transport to accommodate. Note also that the list shown is truncated for clarity, and can be expected to change from time to time as new L-Root nodes are provisioned and old ones decommissioned.

```
% dig NODES.L.ROOT-SERVERS.ORG TXT +short +tcp | head -10
"abj01.l.root-servers.org" "Abidjan" "" "Cote d'Ivoire" "Africa"
"abj02.l.root-servers.org" "Abidjan" "" "Cote d'Ivoire" "Africa"
"akl01.l.root-servers.org" "Mangere" "" "New Zealand" "AsiaPacific"
"akl41.l.root-servers.org" "Mangere" "" "New Zealand" "AsiaPacific"
"akl42.l.root-servers.org" "Mangere" "" "New Zealand" "AsiaPacific"
"akl43.l.root-servers.org" "Mangere" "" "New Zealand" "AsiaPacific"
"akl44.l.root-servers.org" "Mangere" "" "New Zealand" "AsiaPacific"
"ams01.l.root-servers.org" "Haarlemmermeer" "" "Netherlands" "Europe"
"anc01.l.root-servers.org" "Anchorage" "Alaska" "United States" \
"NorthAmerica"
%
```



## 5. Provisioning of IDENTITY.L.ROOT-SERVERS.ORG

Individual L-Root nodes run a dedicated, separate authority-only DNS server process that serves the IDENTITY.L.ROOT-SERVERS.ORG zone. The contents of that zone are unique to every node; hence, each responding node will generate a node-specific response.

The contents of the IDENTITY.L.ROOT-SERVERS.ORG zone are hence deliberately incoherent, the apparent zone contents depending on the node responding to the corresponding query.

The IDENTITY.L.ROOT-SERVERS.ORG zone is delegated to the single name server BEACON.L.ROOT-SERVERS.ORG, numbered on IPv4 and IPv6 addresses that are covered by the same routing advertisements that cover the L-Root service addresses. Reachability of BEACON.L.ROOT-SERVERS.ORG is hence well-aligned with the reachability of L.ROOT-SERVERS.NET; therefore, measurement of the IDENTITY service ought to give similar results to measurement of the L-Root service.

It is considered best practice always to delegate a DNS zone to more than one name server [RFC2182]; however, as described, the IDENTITY.L.ROOT-SERVERS.ORG zone is delegated to just one server. Ordinarily, this would present a risk of failure if that single server is not available; however, given the purpose of the delegation in this case and that the expected mitigation of a failure in a single node is the routing of a query to a different node, delegation to a single server in this particular use-case is effective.

At the time of writing, the ROOT-SERVERS.ORG zone is not signed with DNSSEC. When DNSSEC is deployed in that zone, the L.ROOT-SERVERS.ORG zone will also be signed. This will facilitate secure responses for queries for BEACON.L.ROOT-SERVERS.ORG and NODES.L.ROOT-SERVERS.ORG.

Secure responses for IDENTITY.L.ROOT-SERVERS.ORG are unlikely to become available even with the deployment of DNSSEC in the parent, since the implementation of the IDENTITY.L.ROOT-SERVERS.ORG service involves widely distributed static zone data. Management of key materials distributed to every L-Root node would be impractical to audit, and signatures returned in secure responses would be correspondingly of low value.

## 6. Security Considerations

Some operators of anycast services choose not to disclose locations (or even numbers) of nodes, citing security concerns. The operator of L-Root considers that none of the published information described in this document is truly secret, since any service element that provides service to the Internet can never truly be obscured from

view. Given that location information can be found regardless of any conscious, deliberate disclosure, and since easy access to this information has diagnostic value, the operator of L-Root has adopted a policy of operational transparency.

The information presented in this document presents no new threat to the Internet.

## 7. Acknowledgements

The aspects of the L-Root service that were deployed to facilitate IN-class mapping were discussed and implemented as part of an informal collaboration with Xun Fan, John Heidemann, and Ramesh Govidan, whose contributions are acknowledged. The motivation to facilitate mapping of L-Root as an anycast service using IN-class queries was inspired by [Fan2013].

Helpful reviews and comments from Gaurab Upadhaya, Hugo Salgado, Brian Dixon, Bob Harold, Paul Hoffman, Jakob Schlyter, Andrew Sullivan, Bruce Campbell, S. Moonesamy, and Stephane Bortzmeyer on earlier versions of this document were very much appreciated.

## 8. References

### 8.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2182] Elz, R., Bush, R., Bradner, S., and M. Patton, "Selection and Operation of Secondary DNS Servers", BCP 16, RFC 2182, July 1997.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, December 2006.
- [RFC4892] Woolf, S. and D. Conrad, "Requirements for a Mechanism Identifying a Name Server Instance", RFC 4892, June 2007.
- [RFC5001] Austein, R., "DNS Name Server Identifier (NSID) Option", RFC 5001, August 2007.
- [RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, April 2013.

## 8.2. Informative References

- [ACD] International Air Transport Association (IATA), "Airline Coding Directory (ACD)", 2013, <<http://www.iata.org/publications/Pages/coding.aspx>>.
- [Fan2013] Fan, X., Heidemann, J., and R. Govidan, "Evaluating Anycast in the Domain Name System", Proceedings of the IEEE Infocom Turin, Italy, April 2013.
- [ROOT-SERVERS] "root-servers.org", <<http://www.root-servers.org>>.

## Authors' Addresses

Joe Abley  
Dyn, Inc.  
470 Moore Street  
London, ON N6C 2C2  
Canada

Phone: +1 519 670 9327  
EMail: [jabley@dyn.com](mailto:jabley@dyn.com)

Terry Manderson  
ICANN  
12025 Waterfront Drive  
Suite 300  
Los Angeles, CA 90094-2536  
USA

EMail: [terry.manderson@icann.org](mailto:terry.manderson@icann.org)