Network Working Group
Internet-Draft
Obsoletes: 6304 (if approved)
Intended status: Informational
Expires: December 28, 2014

J. Abley
Dyn, Inc.
W. Maton
OttIX

June 26, 2014

AS112 Nameserver Operations
draft-ietf-dnsop-rfc6304bis-03

Abstract

Many sites connected to the Internet make use of IPv4 addresses that are not globally-unique. Examples are the addresses designated in RFC 1918 for private use within individual sites.

Devices in such environments may occasionally originate Domain Name System (DNS) queries (so-called "reverse lookups") corresponding to those private-use addresses. Since the addresses concerned have only local significance, it is good practice for site administrators to ensure that such queries are answered locally. However, it is not uncommon for such queries to follow the normal delegation path in the public DNS instead of being answered within the site.

It is not possible for public DNS servers to give useful answers to such queries. In addition, due to the wide deployment of private-use addresses and the continuing growth of the Internet, the volume of such queries is large and growing. The AS112 project aims to provide a distributed sink for such queries in order to reduce the load on the corresponding authoritative servers. The AS112 project is named after the Autonomous System Number (ASN) that was assigned to it.

RFC6304 described the steps required to install a new AS112 node, and offered advice relating to such a node's operation. This document updates that advice to facilitate the addition and removal of zones for which query traffic will be sunk at AS112 nodes, using DNAME, whilst still supporting direct delegations to AS112 name servers.

This document obsoletes RFC6304.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-

Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 28, 2014.

Table of Contents

1. NOTE TO RFC EDITOR AND REVIEWERS

This document uses the phrases "TBA-prefix-v4" and "TBA-prefix-v6" in anticipation of the address assignments requested in the IANA Considerations section of [I-D.ietf-dnsop-as112-dname].

"TBA-address-v4" refers to the address within TBA-prefix-v4 assigned to the name server BLACKHOLE.AS112.ARPA. Similarly, "TBA-address-v6" refers to the corresponding address within TBA-prefix-v6.

All of "TBA-prefix-v4", "TBA-address-v4", "TBA-prefix-v6" and "TBA-address-v6" in this document should be replaced with their assigned values prior to publication.

The delegation of the AS112.ARPA zone is specified (and requested) in [I-D.ietf-dnsop-as112-dname].

See Appendix B for an abridged revision history, and discussion of an appropriate venue for discussion.

This section should be removed prior to publication.

2. Introduction

Many sites connected to the Internet make use of IPv4 addresses that are not globally unique. Examples are the addresses designated in [RFC1918] for private use within individual sites.

Devices in such environments may occasionally originate Domain Name System (DNS) [RFC1034] queries (so-called "reverse lookups") corresponding to those private-use addresses. Since the addresses concerned have only local significance, it is good practice for site administrators to ensure that such queries are answered locally [RFC6303]. However, it is not uncommon for such queries to follow the normal delegation path in the public DNS instead of being answered within the site.

It is not possible for public DNS servers to give useful answers to such queries. In addition, due to the wide deployment of private-use addresses and the continuing growth of the Internet, the volume of such queries is large and growing. The AS112 project aims to provide a distributed sink for such queries in order to reduce the load on the IN-ADDR.ARPA authoritative servers [RFC5855].

The AS112 project encompasses a loosely coordinated collection of independently operated name servers. Each name server functions as a single node in an AS112 anycast cloud [RFC4786], and is configured to answer authoritatively for a particular set of nominated zones.

The AS112 project is named after the Autonomous System Number (ASN) that was assigned to it (see Appendix A).

3. AS112 DNS Service

3.1. Approach

3.1.1. Direct Delegation

[RFC6304] describes an approach whereby zones whose traffic should be directed towards an AS112 sink should be directly delegated to AS112 name servers. Correspondingly, each AS112 node is manually configured to answer appropriately for those zones.

The guidance in this document preserves this capability for the zones that were originally delegated in this fashion. AS112 nodes that were implemented in accordance with the guidance in [RFC6304] will continue to provide service for those zones.

3.1.2. DNAME Redirection

[I-D.ietf-dnsop-as112-dname] describes a different approach whereby queries towards specific zones are redirected to an empty zone also hosted on AS112 servers, using DNAME [RFC6672].

The guidance in this document introduces this capability, allowing any zone administrator to sink query traffic in AS112 infrastructure without requiring changes to any AS112 node.

3.2. Zones

To support Direct Delegation AS112 service, AS112 name servers answer authoritatively for the following zones, corresponding to [RFC1918] private-use netblocks:

- 10.IN-ADDR.ARPA
- 16.172.IN-ADDR.ARPA, 17.172.IN-ADDR.ARPA, ..., 31.172.IN-ADDR.ARPA
- 168.192.IN-ADDR.ARPA

and the following zone, corresponding to the "link local" netblock 169.254.0.0/16 described in [RFC6890]:

- 254.169.IN-ADDR.ARPA

To support DNAME Redirection AS112 service, AS112 name servers answer authoritatively for the following zone, as specified in [I-D.ietf-dnsop-as112-dname]:

- EMPTY.AS112.ARPA

To aid identification of AS112 anycast nodes, each node also answers authoritatively for the following zones:

- HOSTNAME.AS112.NET

- HOSTNAME.AS112.ARPA

See Section 4.5 for the recommended contents of all these zones.

3.3. Nameservers

To support Direct Delegation AS112 service, the relevant zones listed in Section 3.2 are delegated to the two name servers BLACKHOLE-1.IANA.ORG (192.175.48.6, 2620:4f:8000::6) and BLACKHOLE-2.IANA.ORG (192.175.48.42, 2620:4f:8000::42).

Additionally, the server PRISONER.IANA.ORG (192.175.48.1, 2620:4f:8000::1) is listed in the MNAME field of the SOA records of the IN-ADDR.ARPA zones served by AS112 name servers. PRISONER.IANA.ORG receives mainly dynamic update queries.

The addresses of all these name servers are covered by the single IPv4 prefix 192.175.48.0/24 and the IPv6 prefix 2620:4f:8000::/48. To date, IPv6 transport for these nameservers has only been available for pre-production testing. Direction to the IANA to add AAAA RRSets for the owner names of these name servers can be found in Section 8.

To support DNAME Redirection AS112 service, the single zone EMPTY.AS112.ARPA is delegated to the single name server BLACKHOLE.AS112.ARPA (TBA-address-v4, TBA-address-v6). The addresses of that name server are covered by the single IPv4 prefix TBA-prefix-v4, and the single IPv6 prefix TBA-prefix-v6.

4. Installation of a New Node

4.1. Useful Background Knowledge

Installation of an AS112 node is relatively straightforward. However, experience in the following general areas may prove useful:

- inter-domain routing with BGP [RFC4271];
- DNS authoritative server operations;
- anycast [RFC4786] distribution of DNS services.

4.2. Topological Location

AS112 nodes may be located anywhere on the Internet. For nodes that are intended to provide a public service to the Internet community (as opposed to private use), it may well be advantageous to choose a location that is easily (and cheaply) reachable by multiple providers, such as an Internet exchange point.

AS112 nodes may advertise their service prefix to BGP peers for local use (analogous to a conventional peering relationship between two providers) or for global use (analogous to a customer relationship with one or more providers).

It is good operational practice to notify the community of users that may fall within the reach of a new AS112 node before it is installed. At an Internet Exchange, local mailing lists usually exist to facilitate such announcements. For nodes that are intended to be globally reachable, coordination with other AS112 operators is highly recommended. See also Section 6.

4.3. Operating System and Host Considerations

Examples in this document are based on UNIX and UNIX-like operating systems, but other operating systems exist which are suitable for use in construction of an AS112 node.

The chosen platform should include support for either cloned loopback interfaces, or the capability to bind multiple addresses to a single loopback interface. The addresses of the name servers listed in Section 3.3 will be configured on these interfaces in order that the DNS software can respond to queries properly.

A host that is configured to act as an AS112 anycast node should be dedicated to that purpose, and should not be used to simultaneously provide other services. This guidance is provided due to the

unpredictable (and occasionally high) traffic levels that AS112 nodes have been seen to attract.

System startup scripts should be arranged such that the various AS112-related components start automatically following a system reboot. The order in which interfaces are configured and software components started should be arranged such that routing software startup follows DNS software startup, and DNS software startup follows loopback interface configuration.

Wrapper scripts or other arrangements should be employed to ensure that the anycast service prefix for AS112 is not advertised while either the anycast addresses are not configured, or while the DNS software is not running.

4.4. Routing Software

AS112 nodes signal the availability of AS112 name servers to the Internet using BGP [RFC4271]: each AS112 node is a BGP speaker, and announces the prefixes 192.175.48.0/24 and 2620:4f:8000::/48 to the Internet with origin AS 112 (see also Section 3.3).

The examples in this document are based on the Quagga Routing Suite [1] running on Linux, but other software packages exist which also provide suitable BGP support for AS112 nodes.

The "bgpd.conf" file is used by Quagga's bgpd daemon, which provides BGP protocol support. The router id in this example is 203.0.113.1; the AS112 node peers with external peers 192.0.2.1, 192.0.2.2, 2001:db8::1 and 2001:db8::2. Note the local AS number 112, and the origination of the prefixes 192.175.48.0/24 and 2620:4f:8000::/48 to support Direct Delegation AS112 service; the IPv4 prefix TBA-prefix-v4 and the IPv6 prefix TBA-prefix-v6 support DNAME Redirection.

For clarity, an IPv4-only AS112 node need not configure any of the IPv6 elements that follow; similarly, an IPv6-only AS112 node need not configure any of the IPv4 elements. Such single-stack hosts can still contribute usefully to IPv4 and IPv6 AS112 services, however, and single-stack operation is not discouraged.

```
! bgpd.conf
!
hostname as112-bgpd
password <something>
enable password <supersomething>
!
! Note that all AS112 nodes use the local Autonomous System Number
```

```
! 112, and originate the IPv4 prefixes 192.175.48.0/24 and
! TBA-prefix-v4 and the IPv6 prefixes 2620:4f:8000::/48 and
! TBA-prefix-v6.
!
! All other addresses shown below are illustrative, and
! actual numbers will depend on local circumstances.
!
! IPv4-only or IPv6-only AS112 nodes should omit advertisements
! for address families they do not support.
!
router bgp 112
 bgp router-id 203.0.113.1
 neighbor 192.0.2.1 remote-as 64496
 neighbor 192.0.2.1 next-hop-self
 neighbor 192.0.2.1 prefix-list AS112-v4 out
 neighbor 192.0.2.1 filter-list 1 out
!
 neighbor 192.0.2.2 remote-as 64497
 neighbor 192.0.2.2 next-hop-self
 neighbor 192.0.2.2 prefix-list AS112-v4 out
 neighbor 192.0.2.2 filter-list 1 out
!
 neighbor 2001:db8::1 remote-as 64498
 neighbor 2001:db8::1 next-hop-self
 neighbor 2001:db8::1 prefix-list AS112-v6 out
 neighbor 2001:db8::1 filter-list 1 out
!
 neighbor 2001:db8::2 remote-as 64499
 neighbor 2001:db8::2 next-hop-self
 neighbor 2001:db8::2 prefix-list AS112-v6 out
 neighbor 2001:db8::2 filter-list 1 out
!
 network 192.175.48.0/24
 network TBA-prefix-v4
!
 address-family ipv6 unicast
  network 2620:4f:8000::/48
  network TBA-prefix-v6
 exit-address-family
!
ip prefix-list AS112-v4 permit 192.175.48.0/24
ip prefix-list AS112-v4 permit TBA-prefix-v4
!
ipv6 prefix-list AS112-v6 permit 2620:4f:8000::/48
ipv6 prefix-list AS112-v6 permit TBA-prefix-v6
!
ip as-path access-list 1 permit ^$
```

The configuration above includes two restrictions on what the AS112 should advertise to its BGP neighbours: a prefix filter that permits only the service prefixes, and an AS_PATH filter that matches only locally-originated routes. Together, these measures prevent the node from becoming a transit point for its adjacent ASes.

The "zebra.conf" file is required to provide integration between protocol daemons (bgpd, in this case) and the kernel.

```
! zebra.conf
!
hostname as112
password <something>
enable password <supersomething>
!
interface lo
!
interface eth0
!
```

4.5. DNS Software

Although the queries received by AS112 nodes are definitively misdirected, it is important that they be answered in a manner that is accurate and consistent. For this reason AS112 nodes operate as fully-functional and standards-compliant DNS authoritative servers [RFC1034], and hence require appropriate DNS software.

Examples in this document are based on ISC BIND9 [2], but other DNS software exists which is suitable for use in construction of an AS112 node.

The following is a sample BIND9 "named.conf" file for a dedicated AS112 server. Note that the name server is configured to act as an authoritative-only server (i.e. recursion is disabled). The name server is also configured to listen on the various AS112 anycast name server addresses, as well as its local addresses.

```
// named.conf

// global options

options {
  listen-on {
    127.0.0.1;         // localhost

  // the following address is node-dependent, and should be set to
  // something appropriate for the new AS112 node
```

```
      203.0.113.1;       // local address (globally-unique, unicast)

    // the following addresses are used to support Direct Delegation
    // AS112 service, and are the same for all AS112 nodes

      192.175.48.1;      // PRISONER.IANA.ORG (anycast)
      192.175.48.6;      // BLACKHOLE-1.IANA.ORG (anycast)
      192.175.48.42;     // BLACKHOLE-2.IANA.ORG (anycast)

    // the following address is used to support DNAME Redirection
    // AS112 service, and is the same for all AS112 nodes

      TBA-address-v4;    // BLACKHOLE.AS112.ARPA (anycast)
    };

    listen-on-v6 {
      ::1;               // localhost

    // the following addresses are used to support Direct Delegation
    // AS112 service, and are the same for all AS112 nodes

      2620:4f:8000::1;   // PRISONER.IANA.ORG (anycast)
      2620:4f:8000::6;   // BLACKHOLE-1.IANA.ORG (anycast)
      2620:4f:8000::42;  // BLACKHOLE-2.IANA.ORG (anycast)

    // the following address is used to support DNAME Redirection
    // AS112 service, and is the same for all AS112 nodes

      TBA-address-v6;    // BLACKHOLE.AS112.ARPA (anycast)
    };

    directory "/var/named";
    recursion no;        // authoritative-only server
  };

  // log queries,  so that when people call us about unexpected
  // answers to queries they didn't realise they had sent, we
  // have something to talk about.  Note that activating this
  // naively has the potential to create high CPU load and consume
  // enormous amounts of disk space.

  logging {
    channel "querylog" {
      file "/var/log/query.log" versions 2 size 500m;
      print-time yes;
    };
    category queries { querylog; };
  };
```

```
// Direct Delegation AS112 Service

// RFC 1918

zone "10.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "16.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "17.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "18.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "19.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "20.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "21.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "22.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "23.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "24.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "25.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "26.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "27.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "28.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "29.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "30.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "31.172.in-addr.arpa" { type master; file "db.dd-empty"; };
zone "168.192.in-addr.arpa" { type master; file "db.dd-empty"; };

// RFC 6890

zone "254.169.in-addr.arpa" { type master; file "db.dd-empty"; };

// DNAME Redirection AS112 Service

zone "empty.as112.arpa" { type master; file "db.dr-empty"; };

// also answer authoritatively for the HOSTNAME.AS112.NET and
// HOSTNAME.AS112.ARPA zones, which contain data of operational
// relevance

zone "hostname.as112.net" {
  type master;
  file "db.hostname.as112.net";
};

zone "hostname.as112.arpa" {
  type master;
  file "db.hostname.as112.arpa";
};
```

The "db.dd-empty" file follows, below. This is the source data used to populate all the IN-ADDR.ARPA zones listed in Section 3.2 that support Direct Delegation AS112 service. Note that the RNAME

specified in the SOA record corresponds to hostmaster@root-servers.org, a suitable e-mail address for technical queries about these zones.

```
; db.dd-empty
;
; Empty zone for Direct Delegation AS112 service.
;
$TTL    1W
@  IN  SOA  PRISONER.IANA.ORG. HOSTMASTER.ROOT-SERVERS.ORG. (
                                1       ; serial number
                                1W      ; refresh
                                1M      ; retry
                                1W      ; expire
                                1W )    ; negative caching TTL
;
       NS     BLACKHOLE-1.IANA.ORG.
       NS     BLACKHOLE-2.IANA.ORG.
;
; There should be no other resource records included in this zone.
;
; Records that relate to RFC 1918-numbered resources within the
; site hosting this AS112 node should not be hosted on this
; name server.
```

The "db.dr-empty" file follows, below. This is the source data used to populate the EMPTY.AS112.ARPA zone that supports DNAME Redirection AS112 service. Note that the RNAME specified in the SOA record corresponds to noc@dns.icann.org, a suitable e-mail address for technical queries about this zone.

```
; db.dr-empty
;
; Empty zone for Direct Delegation AS112 service.
;
$TTL    1W
@  IN  SOA  BLACKHOLE.AS112.ARPA. NOC.DNS.ICANN.ORG. (
                                1       ; serial number
                                1W      ; refresh
                                1M      ; retry
                                1W      ; expire
                                1W )    ; negative caching TTL
;
       NS     BLACKHOLE.AS112.ARPA.
;
; There should be no other resource records included in this zone.
;
; Records that relate to RFC 1918-numbered resources within the
; site hosting this AS112 node should not be hosted on this
; name server.
```

The "db.hostname.as112.net" and "db.hostname.as112.arpa" files follow, below. These zones contain various resource records that provide operational data to users for troubleshooting or measurement purposes, and should be edited to suit local circumstances. Note that the responses to the queries "HOSTNAME.AS112.NET IN TXT" and "HOSTNAME.AS112.ARPA IN TXT" should fit within a 512 octet DNS/UDP datagram: i.e. it should be available over UDP transport without requiring EDNS0 support by the client.

The optional LOC record [RFC1876] included in each zone apex provides information about the geospatial location of the node.

```
; db.hostname.as112.net
;
$TTL    1W
@       SOA     SERVER.EXAMPLE.NET. ADMIN.EXAMPLE.NET. (
                        1               ; serial number
                        1W              ; refresh
                        1M              ; retry
                        1W              ; expire
                        1W )            ; negative caching TTL
;
        NS      BLACKHOLE-1.IANA.ORG.
        NS      BLACKHOLE-2.IANA.ORG.
;
        TXT     "Name of Facility or similar" "City, Country"
        TXT     "See http://www.as112.net/ for more information."
;
```

```
        LOC     45 25 0.000 N 75 42 0.000 W 80.00m 1m 10000m 10m


; db.hostname.as112.arpa
;
$TTL    1W
@       SOA     SERVER.EXAMPLE.NET. ADMIN.EXAMPLE.NET. (
                        1               ; serial number
                        1W              ; refresh
                        1M              ; retry
                        1W              ; expire
                        1W )            ; negative caching TTL
;
        NS      BLACKHOLE.AS112.ARPA.
;
        TXT     "Name of Facility or similar" "City, Country"
        TXT     "See http://www.as112.net/ for more information."
;
        LOC     45 25 0.000 N 75 42 0.000 W 80.00m 1m 10000m 10m
```

4.6. Testing a Newly-Installed Node

The BIND9 tool "dig" can be used to retrieve the TXT resource records associated with the names "HOSTNAME.AS112.NET" and "HOSTNAME.AS112.ARPA", directed at one of the AS112 anycast name server addresses. Continuing the example from above, the response received should indicate the identity of the AS112 node that responded to the query. See Section 4.5 for more details about the resource records associated with "HOSTNAME.AS112.NET".

```
% dig @PRISONER.IANA.ORG HOSTNAME.AS112.NET txt +short +norec
"Name of Facility or similar" "City, Country"
"See http://www.as112.net/ for more information."
%
```

If the response received indicates a different node is being used, then there is probably a routing problem to solve. If there is no response received at all, there might be host or name server problem. Judicious use of tools such as traceroute, and consultation of BGP looking glasses might be useful in troubleshooting.

Note that an appropriate set of tests for a new server will include queries sent from many different places within the expected service area of the node, using both UDP and TCP transport, and exercising all three AS112 anycast name server addresses.

5. Operations

5.1. Monitoring

AS112 nodes should be monitored to ensure they are functioning correctly, just as with any other production service. An AS112 node that stops answering queries correctly can cause failures and timeouts in unexpected places and can lead to failures in dependent systems that can be difficult to troubleshoot.

5.2. Downtime

An AS112 node that needs to go off-line (e.g. for planned maintenance or as part of the diagnosis of some problem) should stop advertising the AS112 service prefixes to its BGP peers. This can be done by shutting down the routing software on the node altogether or by causing the routing system to withdraw the route.

Withdrawal of the service prefixes is important in order to avoid blackholing query traffic in the event that the DNS software on the node is not functioning normally.

5.3. Statistics and Measurement

Use of the AS112 node should be measured in order to track long-term trends, identify anomalous conditions, and to ensure that the configuration of the AS112 node is sufficient to handle the query load.

Examples of free monitoring tools that might be useful to operators of AS112 nodes include:

- bindgraph [3]
- dnstop [4]
- DSC [5]

6. Communications

It is good operational practice to notify the community of users that may fall within the reach of a new AS112 node before it is installed. At Internet Exchanges, local mailing lists usually exist to facilitate such announcements.

For nodes that are intended to be globally reachable, coordination with other AS112 operators is especially recommended. The mailing list <mailto:as112-ops@lists.dns-oarc.net> is operated for this purpose.

Information pertinent to AS112 operations is maintained at <http://www.as112.net/>.

Information about an AS112 node should also be published within the DNS, within the "HOSTNAME.AS112.NET" and "HOSTNAME.AS112.ARPA" zones. See Section 4.5 for more details.

7. On the Future of AS112 Nodes

It is recommended practice for the operators of recursive name servers to answer queries for zones served by AS112 nodes locally, such that queries never have an opportunity to reach AS112 servers [RFC6303]. Operational experience with AS112 nodes does not currently indicate an observable trend towards compliance with those recommendations, however.

It is expected that some DNS software vendors will include default configuration that will implement measures such as those described in [RFC6303]. If such software is widely deployed, it is reasonable to assume that the query load received by AS112 nodes will decrease; however, it is safe to assume that the query load will not decrease to zero, and consequently that AS112 nodes will continue to provide a useful service for the foreseeable future.

The use of DNAME Redirection to provide AS112 service is new, and hence is informed by minimal operational experience. The use of DNAME means that queries for many source zones could be redirected to AS112 infrastructure with no real opportunity for coordination.

If the DNAME Redirection approach is successful, and in the absence of any operational concerns, the community might well recommend the retirement of the original Direct Delegation AS112 service. This document makes no such recommendation, however.

8. IANA Considerations

8.1. General

The name servers associated with Direct Delegation AS112 service are all named under the domain IANA.ORG (see Section 3.3). However, the anycast infrastructure itself is operated by a loosely-coordinated, diverse mix of organisations across the Internet, and is not an IANA function.

The autonomous system number 112, the IPv4 prefix 192.175.48.0/24 and the IPv6 prefix 2620:4f:8000::/48 were assigned by ARIN.

The IPv4 prefix TBA-prefix-v4 and the IPv6 prefix TBA-prefix-v6, used for DNAME Redirection AS112 service, were assigned by the IANA [I-D.ietf-dnsop-as112-dname].

The three nameservers BLACKHOLE-1.IANA.ORG, BLACKHOLE-2.IANA.ORG and PRISONER.IANA.ORG are also reachable over IPv6, as described in Section 3.3. Following a substantial period of pre-production testing by AS112 operators, the IANA is directed to add AAAA RRSets to those owner names in Section 8.2.1, to allow the servers to receive queries and generate responses over IPv6 transport.

8.2. IANA Actions

8.2.1. IPv6 Transport for Direct Delegation AS112 Servers

The IANA is directed to add the following AAAA resource records for the three Direct Delegation AS112 name servers named under IANA.ORG:

Owner Name	AAAA RDATA
PRISONER.IANA.ORG	2620:4f:8000::1
BLACKHOLE-1.IANA.ORG	2620:4f:8000::6
BLACKHOLE-2.IANA.ORG	2620:4f:8000::42

8.2.2. Registration in the Special-Purpose AS Numbers Registry

The IANA is directed to add AS112 to the "Special-Purpose AS Numbers" registry specified in [RFC7249] as follows:

AS Numbers	Reason for Reservation
112	Used by the AS112 project; see [THIS DOCUMENT]

9. Security Considerations

Hosts should never normally send queries to AS112 servers; queries relating to private-use addresses should be answered locally within a site. Hosts that send queries to AS112 servers may well leak information relating to private infrastructure to the public network, and this could present a security risk. This risk is orthogonal to the presence or absence of authoritative servers for these zones in the public DNS infrastructure, however.

Queries that are answered by AS112 servers are usually unintentional; it follows that the responses from AS112 servers are usually unexpected. Unexpected inbound traffic can trigger intrusion detection systems or alerts by firewalls. Operators of AS112 servers should be prepared to be contacted by operators of remote infrastructure who believe their security has been violated. Advice to those who mistakenly believe that responses from AS112 nodes constitutes an attack on their infrastructure can be found in [RFC6305].

The deployment of AS112 nodes is very loosely coordinated compared to other services distributed using anycast. The malicious compromise of an AS112 node and subversion of the data served by the node is hence more difficult to detect due to the lack of central management. Since it is conceivable that changing the responses to queries received by AS112 nodes might influence the behaviour of the hosts sending the queries, such a compromise might be used as an attack vector against private infrastructure.

Operators of AS112 should take appropriate measures to ensure that AS112 nodes are appropriately protected from compromise, such as would normally be employed for production name server or network infrastructure. The guidance provided for root name servers in [RFC2870] may be instructive.

The zones hosted by AS112 servers are not signed with DNSSEC [RFC4033]. Given the distributed and loosely-coordinated structure of the AS112 service, the zones concerned could only be signed if the private key material used was effectively public, obviating any security benefit resulting from the use of those keys.

10. Acknowledgements

The authors wish to acknowledge the assistance of Bill Manning, John Brown, Marco D'Itri, Daniele Arena, Stephane Bortzmeyer, Frank Habicht, Chris Thompson, Peter Losher, Peter Koch, Alfred Hoenes, S. Moonesamy and Mehmet Akcin in the preparation of [RFC6304], which this document supercedes.

11. References

11.1. Normative References

[I-D.ietf-dnsop-as112-dname]
Abley, J., Dickson, B., Kumari, W., and G. Michaelson, "AS112 Redirection using DNAME", draft-ietf-dnsop-as112-dname-03 (work in progress), March 2014.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2870] Bush, R., Karrenberg, D., Kosters, M., and R. Plzak, "Root Name Server Operational Requirements", BCP 40, RFC 2870, June 2000.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, December 2006.

11.2. Informative References

[RFC1876] Davis, C., Vixie, P., Goodwin, T., and I. Dickinson, "A Means for Expressing Location Information in the Domain Name System", RFC 1876, January 1996.

[RFC5855] Abley, J. and T. Manderson, "Nameservers for IPv4 and IPv6 Reverse Zones", BCP 155, RFC 5855, May 2010.

[RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC 6303, July 2011.

[RFC6304] Abley, J. and W. Maton, "AS112 Nameserver Operations", RFC 6304, July 2011.

[RFC6305] Abley, J. and W. Maton, "I'm Being Attacked by PRISONER.IANA.ORG!", RFC 6305, July 2011.

[RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", RFC 6672, June 2012.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, April 2013.

[RFC7249] Housley, R., "Internet Numbers Registries", RFC 7249, May 2014.

URIs

[1] <http://www.quagga.net/>

[2] <http://www.isc.org/software/BIND/>

[3] <http://www.linux.it/~md/software/>

[4] <http://dns.measurement-factory.com/tools/dnstop/>

[5] <http://dns.measurement-factory.com/tools/dsc/>

Appendix A. History

Widespread use of the private address blocks listed in [RFC1918] followed that document's publication in 1996. At that time the IN-ADDR.ARPA zone was served by root servers.

The idea of off-loading IN-ADDR.ARPA queries relating to [RFC1918] addresses from the root name servers was first proposed by Bill Manning and John Brown.

The use of anycast for distributing authoritative DNS service for [RFC1918] IN-ADDR.ARPA zones was subsequently proposed at a private meeting of root server operators.

ARIN provided an IPv4 prefix for the anycast service and also the autonomous system number 112 for use in originating that prefix. This assignment gave the project its name.

In 2002, the first AS112 anycast nodes were deployed.

In 2011, the IN-ADDR.ARPA zone was redelegated from the root servers to a new set of servers operated independently by AfriNIC, APNIC, ARIN, ICANN, LACNIC, and the RIPE NCC and named according to [RFC5855].

[RFC6304], the precursor to this document, was published in July 2011.

The use of anycast name servers in the AS112 project contributed to the operational experience of anycast DNS services, and it can be seen as a precursor to the anycast distribution of other authoritative DNS servers in subsequent years (e.g., various root servers).

Appendix B. Revision History and Venue

A suitable venue for discussion of this document is the dnsop working group. Private comments may also be directed at the authors.

This section (and sub-sections) should be removed prior to publication.

B.1. draft-jabley-dnsop-rfc6304bis-00

Initial revision of [RFC6304] intended to provide guidance consistent with [I-D.ietf-dnsop-as112-dname].

B.2. draft-ietf-dnsop-rfc6304bis-00

Change of filename following working group adoption.

B.3. draft-ietf-dnsop-rfc6304bis-01

Correct "Obsoletes" header in document metadata from "RFC6304" to "6304" as requested by Tim Wicinski.

B.4. draft-ietf-dnsop-rfc6304bis-02

Add IPv6 details for Direct Delegation AS112 service, including IANA considerations to add AAAA RRs to PRISONER.IANA.ORG and friends.

Add IANA considerations that will add AS112 to the "Special-Purpose AS Numbers" registry, as created in [RFC7249].

Merge in some AUTH48 changes from RFC6304 that had been overlooked.

Authors' Addresses

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

Phone: +1 519 670 9327
Email: jabley@dyn.com

William F. Maton Sotomayor
Ottawa Internet Exchange
Constitution Square
1400-340 Albert Street
Ottawa, ON K1R 0A5
Canada

Email: wmaton@ottix.net