

Network Working Group
Internet-Draft
Updates: [6304](#) (if approved)
Intended status: BCP
Expires: December 1, 2012

W. Kumari
Google
W. Sotomayor
NRC-CNRC
J. Abley
ICANN
May 30, 2012

Omniscient AS112 Servers
draft-wkumari-dnsop-omniscient-as112-00

Abstract

The AS112 Project loosely coordinates Domain Name System (DNS) servers to which DNS zones corresponding to private use addresses are delegated. Queries for names within those zones have no useful responses in a global context. The purpose of the project is to reduce the load of such junk queries on the authoritative name servers that would otherwise receive them, directing the load instead to name servers operated within the AS112 project.

Adding and dropping zones from the AS112 servers is difficult, due to the loosely-coordinated nature of the project. This document proposes a mechanism by which AS112 name servers could answer authoritatively for all possible zones, reducing the add/drop problem to one of delegation within the DNS without operational impact on the servers themselves.

This document updates [RFC 6304](#).

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 1, 2012.

Copyright Notice

Copyright (c) 2012 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Protocol Considerations	4
4.	Operational Considerations	4
5.	Addressing Considerations	5
6.	Updates to RFC 6304	5
6.1.	Changes to Section 3.4 , Routing Software	5
6.2.	Changes to Section 3.5 , DNS Software	7
6.3.	Changes to Section 3.6 , Testing a Newly Installed Node	9
7.	IANA Considerations	9
8.	Security Considerations	10
9.	Acknowledgements	10
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	10
Appendix A.	Document Notes	11
A.1.	Venue	11
A.2.	Textual Substitutions	11
A.3.	Open Questions	11
A.4.	Change History	11
A.4.1.	draft-wkumari-dnsop-omniscient-as112-00	11
A.4.2.	draft-wkumari-omniscient-as112-00	12
	Authors' Addresses	12

1. Introduction

The AS112 Project loosely coordinates Domain Name System (DNS) servers [[RFC1034](#)] to which DNS zones corresponding to private use addresses are delegated. Queries for names within those zones have no useful responses in a global context. The purpose of the project is to reduce the load of such junk queries on the authoritative name servers that would otherwise receive them, directing the load instead to name servers operated within the AS112 project.

To date, AS112 nameservers have been used exclusively for names corresponding to the reverse mapping for private-use IPv4 addresses. A description of current advice for AS112 operators, including motivations and guidance for technical deployment and operations can be found in [[RFC6304](#)].

Other DNS domains have analogously local significance. Examples corresponding to the reverse-mapping of special-use IPv4 and IPv6 addresses can be found in [[RFC6303](#)].

It is to be expected that new domains will be identified from time to time that fit the use pattern for which delegation to AS112 servers might be desirable. There is currently no mechanism by which particular zones can be reliably added to or dropped from AS112 servers, however. This is principally a consequence of the loosely-coordinated nature of the project, coupled with a desire to avoid lame delegations which might have unforeseen operational consequences.

This document proposes a mechanism by which AS112 servers could provide consistent, reliable negative responses for all DNS queries, eliminating the operational requirement to add or drop particular zones from all AS112 servers.

2. Terminology

An "Existing AS112 Server" is a DNS name server configured according to the guidance provided in [[RFC6304](#)] and listening on the IPv4 addresses 192.175.48.1 (PRISONER.IANA.ORG), 192.175.48.6 (BLACKHOLE-1.IANA.ORG) and 192.175.48.42 (BLACKHOLE-2.IANA.ORG).

An "Omniscient AS112 Server" is a DNS nameserver configured according to the guidance provided in [[RFC6304](#)], as extended by this document. Such servers listen on the same addresses as Existing AS112 Servers, but also additional addresses as described in [Section 5](#).

Where discussions apply equally to Existing AS112 Servers and Omniscient AS112 Servers, the unqualified phrase "AS112 Server" is

used.

An "AS112 Zone" is a DNS zone which has been delegated to an AS112 Server.

An "Existing AS112 Zone" is an AS112 Zone which has been delegated to an existing AS112 Server.

3. Protocol Considerations

An AS112 Server responds with an authoritative (AA=1) name error (NXDOMAIN, RCODE=3) for any query request whose (QNAME, QCLASS) falls within an AS112 Zone [[RFC1035](#)].

AS112 Servers do not respond to zone transfer requests (QTYPE=252).

The name error (NXDOMAIN) response from an Omniscient AS112 Server differs from that sent by an Existing AS112 Server in that the closest enclosing SOA returned has a different owner name. Existing AS112 Servers return an authority-section SOA with an owner name corresponding to the apex of the AS112 Zone concerned; Omniscient AS112 Servers return an SOA with an owner name of ".". This difference has not been shown to cause any practical change in behaviour in commonly-deployed DNS resolver software.

4. Operational Considerations

Existing AS112 Servers address the protocol considerations described in [Section 3](#) by serving each Existing AS112 Zone explicitly. In each case the zone contents are identical, containing only required apex SOA and NS records. Adding or dropping a delegation for an Existing AS112 Zone requires coordination amongst all deployed Existing AS112 Server operators in order to add or drop the zone.

There is no practical expectation that AS112 Server operators coordinate the configuration of their infrastructure or even make their existence known in any systematic way. Delegation of new zones to Existing AS112 Servers is hence problematic; there is an expectation that such delegations would be lame for a significant client population. Since the predictable behaviour of AS112 Servers from clients is desirable, and it is possible that significant variation would have operational consequences, no new zones should be delegated to existing AS112 Servers.

Omniscient AS112 Servers serve an unsigned root zone, containing only required apex SOA and NS records. Adding or dropping a delegation

for an AS112 Zone requires imposes no operational requirements on Omniscient AS112 Server operators.

Delegation of new AS112 Zones should only be made to Omniscient AS112 Servers. The desire to delegate new AS112 Zones therefore imposes a requirement on Omniscient AS112 Servers to listen on addresses which are different to those used by Existing AS112 Servers. Addressing is discussed in [Section 5](#).

By ensuring that Omniscient AS112 Servers listen on Existing AS112 Servers' addresses as well as the new addresses specified in [Section 5](#) a smooth migration is possible, allowing Existing AS112 Servers to be reconfigured as Omniscient AS112 Servers. Omniscient AS112 Servers are therefore a superset of AS112 Servers.

5. Addressing Considerations

Omniscient AS112 Servers listen on the following addresses:

- o IPv4-TBA1 (A.AS112.NET)
- o IPv6-TBA1 (A.AS112.NET)
- o IPv4-TBA2 (B.AS112.NET)
- o IPv6-TBA2 (B.AS112.NET)
- o IPv4-TBA3 (C.AS112.NET)
- o IPv6-TBA3 (C.AS112.NET)

IPv4-TBA1, IPv4-TBA2 and IPv4-TBA3 are covered by a single IPv4 prefix, IPv4-PREFIX-TBA. Similarly, IPv6-TBA1, IPv6-TBA2 and IPv6-TBA3 are covered by a single IPv6 prefix, IPv6-PREFIX-TBA.

The addresses specified for Omniscient AS112 Servers are deliberately different from those assigned to Existing AS112 Servers for reasons discussed in [Section 4](#).

6. Updates to [RFC 6304](#)

6.1. Changes to [Section 3.4](#), Routing Software

Omniscient AS112 Nodes with IPv4 connectivity should originate the IPv4 service prefix associated with Existing AS112 Nodes, 192.175.48.0/24, and also the IPv4 service prefix associated with Omniscient AS112 Nodes, IPv4-PREFIX.

Omniscient AS112 Nodes with IPv6 connectivity should originate the IPv6 service prefix IPv6-PREFIX-TBA.

Applying this direction to the "bgpd.conf" file included as an example in this section results in the configuration shown in Figure 1.

```
! bgpd.conf
!
hostname as112-bgpd
password <something>
enable password <supersomething>
!
! Note that all AS112 nodes use the local Autonomous System
! Number 112, and originate IPv4 and IPv6 prefixes (where IPv4
! and IPv6 connectivity is available) as follows:
!
!   IPv4:   192.175.48.0/24
!           IPv4-PREFIX-TBA
!
!   IPv6:   IPv6-PREFIX-TBA
!
! All other addresses shown below are illustrative, and
! actual numbers will depend on local circumstances.
!
router bgp 112
  bgp router-id 203.0.113.1
  !
  address-family ipv4
    network 192.175.48.0
    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
    network 192.175.48.0/24
    network IPv4-PREFIX-TBA
  !
  address-family ipv6 unicast
    neighbor 2001:db8::1 remote-as 64496
    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 64497
    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 IPv6-PREFIX-TBA
```



```
!  
ip prefix-list AS112-v4 permit 192.175.48.0/24  
ip prefix-list AS112-v4 permit IPv4-PREFIX-TBA  
!  
ipv6 prefix-list AS112-v6 permit IPv6-PREFIX-TBA  
!  
ip as-path access-list 1 permit ^$
```

Figure 1

6.2. Changes to [Section 3.5](#), DNS Software

Omniscient AS112 Servers with IPv4 connectivity should include DNS software configured to listen on the addresses IPv4-TBA1, IPv4-TBA2 and IPv4-TBA3 in addition to the addresses used by Existing AS112 Servers.

Omniscient AS112 Servers with IPv6 connectivity should include DNS software configured to listen on the addresses IPv6-TBA1, IPv6-TBA2 and IPv6-TBA3.

Omniscient AS112 Servers serve an empty, unsigned root zone instead of explicitly serving the zones specified in [[RFC6304](#)].

Applying this direction to the "named.conf" file included as an example in this section results in the configuration fragment shown in Figure 2.

```
options {  
    // The following configuration stanza is for Omniscient AS112  
    // Servers with IPv4 connectivity  
  
    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 correspond to Existing AS112 Server  
        // addresses  
  
        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 addresses are required by Omniscient AS112 Servers
```



```
    IPv4-TBA1;          // A.AS112.NET
    IPv4-TBA2;          // B.AS112.NET
    IPv4-TBA3;          // C.AS112.NET
};

// The following configuration stanza is for Omniscient AS112
// Servers with IPv6 connectivity

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

    IPv6-TBA1;          // A.AS112.NET
    IPv6-TBA2;          // B.AS112.NET
    IPv6-TBA3;          // C.AS112.NET
};

directory "/var/named";
recursion no;          // authoritative-only server
query-source address *;
};

// 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
// 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; };
};

// Substantially empty root zone (replaces explicit zone
// configuration specified in RFC 6304 for Existing AS112 Servers)

zone "." {
    type master;
    file "db.empty";
};

// Also answer authoritatively for the HOSTNAME.AS112.NET zone,
// which contains data of operational relevance.

zone "hostname.as112.net" {
    type master;
```



```
file "db.hostname.as112.net";

// No other zones should be hosted on this name server.
};
```

Figure 2

The "db.empty" file is updated to include references to nameservers used by Omniscient AS112 Servers, as shown in Figure 3.

```
; db.empty
;
; Empty zone for AS112 server.
;
$TTL      1W
@ IN SOA  A.AS112.NET. hostmaster.root-servers.org. (
                                1          ; serial number
                                1W         ; refresh
                                1M         ; retry
                                1W         ; expire
                                1W )      ; negative caching TTL
;
    NS     B.AS112.NET.
    NS     C.AS112.NET.
;
; There should be no other resource records included in this zone.
;
```

Figure 3

6.3. Changes to [Section 3.6](#), Testing a Newly Installed Node

Testing should include all configured service addresses for an Omniscient AS112 Server (IPv4 or IPv6 or both, as appropriate). Note that the IPv4 service addresses include those described in [[RFC6304](#)] for Existing AS112 Servers.

7. IANA Considerations

This document describes infrastructure which could be used in the future to direct the IANA to delegate or redelegate infrastructure zones under its administrative control.

However, this document makes no request of the IANA.

8. Security Considerations

The contents of the Security Considerations section of [[RFC6304](#)] should be reviewed, since that discussion is pertinent to the operation of Omniscient AS112 Servers as well as Existing AS112 Servers.

The deployment of Omniscient AS112 Servers enables new delegations to AS112 Servers.

Queries received by an AS112 Server might reveal operational data for which there is an expectation of privacy. For example, leaked queries for an organisation's internal DNS names which are sent to an AS112 Server might reveal the existence of those names to the AS112 Server operator. The delegation of new zones to AS112 Servers has the potential to increase opportunities for such unintentional information leakage.

The delegation of new zones to AS112 Servers has the potential to increase the traffic received by those servers. AS112 Server operators are encouraged to monitor traffic levels, and to take appropriate steps if traffic levels threaten the stability of their networks.

9. Acknowledgements

The authors thank and acknowledge the contributions of Dr Paul Vixie, Shane Kerr and Bill Manning in the preparation of this document.

10. References

10.1. 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.
- [RFC6304] Abley, J. and W. Maton, "AS112 Nameserver Operations", [RFC 6304](#), July 2011.

10.2. Informative References

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

[Appendix A](#). Document Notes

This section (and sub-sections) contain information useful for development and review of this document, and should be removed prior to publication.

[A.1](#). Venue

This document is an individual submission, and is not the product of an IETF working group. However, a suitable venue for discussion is the dnsop working group mailing list.

[A.2](#). Textual Substitutions

The strings "IPv4-TBA1", "IPv4-TBA2" and "IPv4-TBA3" should be replaced in this document should be replaced with IPv4 addresses assigned for the purpose described. The covering IPv4 prefix for all three addresses should replace the string "IPv4-PREFIX-TBA".

Similarly, the strings "IPv6-TBA1", "IPv6-TBA2", "IPv6-TBA3" and "IPv6-PREFIX-TBA" should be substituted in the text with assigned production values.

[A.3](#). Open Questions

1. Where to get IPv4 and IPv6 assignments from? There has already been an assignment to DNS-OARC by ARIN for v6 service for AS112 servers.

[A.4](#). Change History

[A.4.1](#). [draft-wkumari-dnsop-omniscient-as112-00](#)

Document title changed to include the dnsop keyword, so that IETF document automation can send courtesy notifications of document actions to the dnsop working group.

Abstract and introduction expanded.

[RFC2119](#) requirements notation removed, since this is an informational document and any normative language would be toothless.

Discussion broken out into Protocol Considerations, Operational Considerations and Addressing Considerations.

Detailed updates to [[RFC6304](#)] added.

[A.4.2. draft-wkumari-omniscient-as112-00](#)

Initial draft, circulated privately, not submitted.

Authors' Addresses

Warren Kumari
Google
1600 Ampitheatre Parkway
Mountain View, CA 94043
USA

Email: warren@kumari.net

William F. Maton Sotomayor
National Research Council of Canada
1200 Montreal Road
Ottawa, ON K1A 0R6
Canada

Phone: +1 613 993 0880
Email: wfms@ryouko.imsb.nrc.ca

Joe Abley
ICANN
12025 Waterfront Drive, Suite 300
Los Angeles, CA 90094-2536
USA

Phone: +1 519 670 9327
Email: joe.abley@icann.org

