

Network Working Group
Internet-Draft
Updates: [6304](#) (if approved)
Intended status: Informational
Expires: December 23, 2013

W. Kumari
Google
W. Sotomayor
NRC-CNRC
J. Abley
ICANN
R. Bellis
Nominet UK
June 21, 2013

Omniscient AS112 Servers
draft-wkumari-dnsop-omniscient-as112-03

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 this project is to reduce the load of such junk queries on the authoritative name servers that would otherwise receive them, and instead direct the load to name servers operated within the AS112 project.

Due to the loosely-coordinated nature of the project, adding and dropping zones from the AS112 servers is difficult. This document proposes a mechanism by which AS112 name servers could answer authoritatively for all possible zones. This eliminates the add/drop problem, changing it to a matter of delegation within the DNS and requiring no operational changes 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."

Internet-Draft

Omniscient AS112 Servers

June 2013

This Internet-Draft will expire on December 23, 2013.

Copyright Notice

Copyright (c) 2013 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.	Overview	4
3.1.	DNSSEC	4
4.	Protocol Considerations	4
5.	Operational Considerations	6
6.	Addressing Considerations	7
7.	Updates to RFC 6304	7
7.1.	Changes to Section 3.4 , Routing Software	7
7.2.	Changes to Section 3.5 , DNS Software	8
7.3.	Changes to Section 3.6 , Testing a Newly Installed Node	9
8.	Other Approaches	9
9.	IANA Considerations	9
10.	Security Considerations	10
11.	Acknowledgements	10
12.	References	10
12.1.	Normative References	10
12.2.	Informative References	11
Appendix A.	Implementation / "Running Code"	11
Appendix B.	Document Notes	11
B.1.	Venue	11
B.2.	Textual Substitutions	11
B.3.	Open Questions	11
B.4.	Change History	12

B.4.1.	draft-wkumari-dnsop-omniscient-as112-00	13
B.4.2.	draft-wkumari-omniscient-as112-00	14
Authors' Addresses		14

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 6](#).

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.](#) Overview

An Omniscient AS112 Server is one that acts as though it is authoritative for all zones, but denies that there is any data. This allows domains to be delegated to it and to receive back an authoritative NoError / NoData response for queries for labels in that domain. This allows new zones to be added to or removed from the AS112 project by simply delegating to the Omniscient AS112 servers without needing to co-ordinate with the server operators.

It is worth noting that the term Omniscient is imperfect - it makes it sound like an Omniscient AS112 Server knows the answers to everything, whereas all it actually knows that it doesn't know anything. A thesaurus shows that a good antonym for "Omniscient" is "stupid", but we are not using this as we feel that it is too broad a term and covers most nameservers. (:-))

[3.1.](#) DNSSEC

Responses from Omniscient AS112 servers (like Existing AS112 servers) are not DNSSEC signed. This does not cause issues as responses have no meaning in the global context, and so the delegations to these are either not signed or are insecure delegations.

4. Protocol Considerations

Familiarity with [[RFC1034](#)] and [[RFC1035](#)] is assumed.

In order to safely cache the response, DNS implementations require the closest-enclosing SOA to be returned. An Omniscient AS112 server (which is not configured with a specific list of zones, and hence zone cuts) cannot necessarily know where that is. Removing labels and guessing (whether to the extreme case of removing all labels, or returning one, or anything in between) cannot be guaranteed to be appropriate, since the answers might clash with authentic answers already present in client caches. A client that has followed a referral to an Omniscient AS112 server is guaranteed not to have a cached SOA that matches the QNAME, however, so Omniscient AS112 servers use the QNAME as the SOA and owner name.

Please see [Appendix A](#) for information on an implementation ("running code") that does this.

AS112 Servers respond to AXFR (QTYPE=252) or IXFR (QTYPE=251) with RCODE=REFUSED.

A TYPE=6 (SOA) resource record for Omniscient AS112 servers contains:

- o MNAME = "a.as112.net."

- o RNAME = "hostmaster.as112.net."
- o SERIAL = 1
- o REFRESH = 604800 (7 days)
- o RETRY = 2592000 (30 days)
- o EXPIRE = 604800 (7 days)
- o MINIMUM = 604800 (7 days, negative caching TTL)

For all queries with QTYPE=2 (NS) an AS112 Server responds with an authoritative (AA=1) answer with NoError (RCODE=0), the owner name copied from the QNAME and two resource records of TYPE=2 (NS), one containing "B.AS112.NET." and the containing "C.AS112.NET."

For all queries with QTYPE=6 (SOA) an AS112 Server responds with an authoritative (AA=1) answer with NoError (RCODE=0), the owner name copied from the QNAME and one (ANCOUNT=1) resource record of TYPE=6 (SOA).

For all queries with QTYPE=255 (*, also known as ANY) an AS112 Server responds with an authoritative (AA=1) answer with NoError (RCODE=0) the owner name copied from the QNAME and three (ANCOUNT=3) resource records, one containing the SOA (as described above), and two containing NS (also as described above).

For all other queries an AS112 Server responds with an authoritative (AA=1) NoError (RCODE=0) with the owner name copied from the QNAME in the request and no answers (ANCOUNT=0). The resource record of TYPE=6 (SOA) (as described above) should be returned in the authority section. The presence of the SOA is to allow the negative cache TTL to be set(see [[RFC2308](#)]).

[5.](#) Operational Considerations

Existing AS112 Servers address the protocol considerations described in [Section 4](#) 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.

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 generate a response (as described in Section 4 ([Section 4](#))) as though they are authoritative for everything ("."). Adding or dropping a delegation for an AS112 Zone therefore imposes no operational requirements on Omniscient AS112 Server operators.

Delegation of new AS112 Zones should only be made to Omniscient AS112 Servers. Omniscient AS112 Servers, therefore, must listen on additional addresses to those used by existing AS112 Servers. Addressing is discussed in [Section 6](#).

By ensuring that Omniscient AS112 Servers listen on Existing AS112 Servers' addresses as well as the new addresses specified in [Section 6](#) 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.

[6.](#) 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 5](#).

[7](#). Updates to [RFC 6304](#)

[7.1](#). 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-TBA.

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

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

Omniscient AS112 Servers should be configured to listen on the addresses IPv6-TBA1, IPv6-TBA, IPv6-TBA3, IPv4-TBA1, IPv4-TBA2 and IPv4-TBA3 in addition to the addresses used for Existing AS112 Servers.

Omniscient AS112 Servers generate an answer as described in [Section 4](#) instead of explicitly serving the zones specified in [\[RFC6304\]](#).

As ISC BIND [\[BIND\]](#) does not provide the required functionality a custom nameserver implementation needs to be deployed, and so the example "named.conf" file in this section can be disregarded.

[7.3.](#) 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.

[8.](#) Other Approaches

[Editors note: This is provided for information only. We may want to remove this before publication... or not. If we don't, we should tidy it up, it is conversational at the moment. It is here (instead of in a "Changes / Appendix" section) so that folk read it :-P

Other solutions to this problem were considered. There are mentioned here to avoid having to have the discussion from the beginning again.]

The original version of this document suggested simply serving an empty root zone. This was tested, but found not to work. In order to prevent poisoning, recursive resolvers will only cache the response if it comes from the closest enclosing zone. This was tested in BIND both with and without 'minimal responses'. Thanks to Mark Andrews for pointing this out.

Brian Dickson suggested that DNAME could be used (instead of delegating to AS112, DNAME to "foo" and then "foo" could be the SOA. We have not yet fully tested this idea, one concern would be how widely deployed DNAME support is. This is still being considered.

Paul Vixie (and others) suggested a metazone solution. This is also worth considering more. A possible concern is stale AXFRs.

[9.](#) 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.

10. 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.

11. Acknowledgements

The authors thank and acknowledge the contributions of Dr Paul Vixie, Ed Lewis, Bill Manning, George Michaelson, Mark Andrews, Shane Kerr, Brian Dickson, S. Moonesamy, Chris Thompson, Nick Hilliard, Chris Morrow, Paul Hoffman, Kurt Erik Lindqvist, Erik Kline, Roy Arends and all the folk on the AS112 Project mailing lists in the preparation of this document.

12. References

12.1. 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.

Kumari, et al.

Expires December 23, 2013

[Page 10]

Internet-Draft

Omniscient AS112 Servers

June 2013

- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](#), March 1998.
- [RFC6304] Abley, J. and W. Maton, "AS112 Nameserver Operations", [RFC 6304](#), July 2011.

12.2. Informative References

- [BIND] Nominet UK, "Internet Systems Consortium, "BIND"", , <<http://www.isc.org/>>.
- [RFC6303] Andrews, M., "Locally Served DNS Zones", [BCP 163](#), [RFC 6303](#), July 2011.
- [evldns] Bellis, R., "evldns", , <<http://code.google.com/p/evldns/>>.

Appendix A. Implementation / "Running Code"

The "evldns" [[evldns](#)] library (written by Ray Bellis, Nominet UK) includes an Omniscient AS112 Server implementation in the file "oas112d.c"

Appendix B. Document Notes

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

B.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.

[B.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.

[B.3.](#) Open Questions

Kumari, et al.	Expires December 23, 2013	[Page 11]
----------------	---------------------------	-----------

Internet-Draft	Omniscient AS112 Servers	June 2013
----------------	--------------------------	-----------

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.

[B.4.](#) Change History

Template:

- o Initial draft
- o Initial draft, circulated privately, not submitted.

-00:

- o Rewrote much of the document (especially [Section 4](#) to explain how (and why) responses should be generated.
- o Updated "Updates to [RFC 6304](#)" section to explain the BIND does not currently implement this, and so named.conf, etc should be ignored.
- o Removed example "empty" zone.
- o Changed the addressing bit at the suggestion of SM.

-01:

- o 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.
- o Abstract and introduction expanded.
- o [RFC2119](#) requirements notation removed, since this is an informational document and any normative language would be toothless.
- o Discussion broken out into Protocol Considerations, Operational Considerations and Addressing Considerations.
- o Reverted to the custom software / synthesized answers.
- o Added in the Ray Bellis evldns stuff.

-01 to -02:

- o s/NoError/NXDomain/ -- Suggestion from Paul Vixie (and others). "Nxd says there is no such name, no matter what the type was, and

Kumari, et al.

Expires December 23, 2013

[Page 12]

Internet-Draft

Omniscient AS112 Servers

June 2013

there are no children. No data/noerror says there are either other types or children or both. We know what the truth must be and we know which implications we want the requestor to follow. Right?" -- Paul.

- o Need to retest with empty root zone, and "minimal responses". Initial test didn't seem to suppress the 'Negative Answers from Authoritative Servers' ([rfc2308](#))
- o Removed the "'Editor note: [NoError was chosen instead of NXDOMAIN because we did not think that we could reasonably return an SOA RR which clearly indicates that the QNAME does exist, and also return an NXDOMAIN.]" as we are now using NXDomain :-P
- o This version submitted by Warren, with no real discussion with co-authors. Trying to squeeze things under the -01 cutoff.

-02 to -03:

- o "AS112 servers don't respond to transfers" -> AS112 respond, but with REFUSED" -- thanks to Ed Lewis.
- o Added an overview section. Before we just leapt into protocol discussions.
- o And after much discussion on the list, and even more so in Orlando we reverted to the NoError/NoData. See Orlando DNSOP audio archives, around 50min.
- o Added [section 3.1](#) - DNSSEC
- o Added [Section 8](#) - Other Approaches. This section is conversational.

[B.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.

Kumari, et al.	Expires December 23, 2013	[Page 13]
----------------	---------------------------	-----------

Internet-Draft	Omniscient AS112 Servers	June 2013
----------------	--------------------------	-----------

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

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

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

Ray Bellis
Nominet UK
Edmund Halley Road
Oxford OX4 4DQ
United Kingdom

Phone: +44 1865 332211
Email: ray.bellis@nominet.org.uk