

Internet Engineering Task Force (IETF)
Request for Comments: 7535
Category: Informational
ISSN: 2070-1721

J. Abley
Dyn, Inc.
B. Dickson
Twitter, Inc.
W. Kumari
Google
G. Michaelson
APNIC
May 2015

AS112 Redirection Using DNAME

Abstract

AS112 provides a mechanism for handling reverse lookups on IP addresses that are not unique (e.g., RFC 1918 addresses). This document describes modifications to the deployment and use of AS112 infrastructure that will allow zones to be added and dropped much more easily, using DNAME resource records.

This approach makes it possible for any DNS zone administrator to sink traffic relating to parts of the global DNS namespace under their control to the AS112 infrastructure without coordination with the operators of AS112 infrastructure.

Status of This Memo

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

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are 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/rfc7535>.

Copyright Notice

Copyright (c) 2015 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. Design Overview	4
3. AS112 Operations	5
3.1. Extensions to Support DNAME Redirection	5
3.2. Redirection of Query Traffic to AS112 Servers	5
4. Continuity of AS112 Operations	6
5. Candidate Zones for AS112 Redirection	6
6. DNAME Deployment Considerations	7
7. IAB Statement Regarding This .ARPA Request	8
8. IANA Considerations	8
8.1. Address Assignment	8
8.2. Hosting of AS112.ARPA	10
8.3. Delegation of AS112.ARPA	10
9. Security Considerations	10
10. References	11
10.1. Normative References	11
10.2. Informative References	11
Appendix A. Assessing Support for DNAME in the Real World	13
A.1. Methodology	13
A.2. Results	15
Acknowledgements	16
Authors' Addresses	16

1. 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) 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 IN-ADDR.ARPA authoritative servers. The AS112 project is named after the Autonomous System Number (ASN) that was assigned to it.

Prior to implementation of this technique, the AS112 project did not accommodate the addition and removal of DNS zones elegantly. Since additional zones of definitively local significance are known to exist, this presents a problem. This document describes modifications to the deployment and use of AS112 infrastructure that will allow zones to be added and dropped much more easily.

The AS112 project is described in detail in [RFC7534].

The AS112 nameservers (PRISONER.IANA.ORG, BLACKHOLE-1.IANA.ORG, and BLACKHOLE-2.IANA.ORG) are required to answer authoritatively for each and every zone that is delegated to them. If a zone is delegated to AS112 nameservers without those nameservers being configured ahead of time to answer authoritatively for that zone, there is a detrimental impact on clients following referrals for queries within that zone. This misconfiguration is colloquially known as a "lame delegation".

AS112 nameserver operators are only loosely coordinated, and hence adding support for a new zone (or, correspondingly, removing support for a zone that is no longer delegated to the AS112 nameservers) is difficult to accomplish with accuracy. Testing AS112 nameservers remotely to see whether they are configured to answer authoritatively for a particular zone is similarly challenging, since AS112 nodes are distributed using anycast [RFC4786].

This document defines a more flexible approach for sinking queries on AS112 infrastructure that can be deployed alongside unmodified, existing AS112 nodes. Instead of delegating additional zones directly to AS112 nameservers, DNAME [RFC6672] redirection is used. This approach has the advantage that query traffic for arbitrary parts of the namespace can be directed to AS112 servers without those servers having to be reconfigured every time a zone is added or removed.

This approach makes it possible for any DNS zone administrator to sink traffic relating to parts of the global DNS namespace under their control to the AS112 infrastructure without coordination with the operators of AS112 infrastructure.

2. Design Overview

A new zone, EMPTY.AS112.ARPA, is delegated to a single nameserver BLACKHOLE.AS112.ARPA (IPv4 address 192.31.196.1, IPv6 address 2001:4:112::1).

The IPv4 address 192.31.196.1 has been selected from the prefix assigned by the IANA such that the address is coverable by a single IPv4 /24 prefix, and that no other address covered by that prefix is in use. The IPv6 address 2001:4:112::1 has been similarly assigned such that no other address within a covering /48 is in use. This addressing plan accommodates the anycast distribution of the BLACKHOLE.AS112.ARPA service using a single IPv4 service prefix and a single IPv6 service prefix. See [RFC4786] for more discussion of anycast service distribution; see Section 8 for the specific actions completed by IANA per this document.

Some or all of the existing AS112 nodes should be extended to support these new nameserver addresses and to host the EMPTY.AS112.ARPA zone. See [RFC7534] for revised guidance to AS112 server operators.

Each part of the DNS namespace for which it is desirable to sink queries at AS112 nameservers should be redirected to the EMPTY.AS112.ARPA zone using DNAME [RFC6672]. See Section 3.2 for guidance to zone administrators.

3. AS112 Operations

3.1. Extensions to Support DNAME Redirection

Guidance to operators of AS112 nodes is extended to include configuration of the 192.31.196.1 and 2001:4:112::1 addresses, and the corresponding announcement of covering routes for those addresses, and to host the EMPTY.AS112.ARPA zone.

IPv4-only AS112 nodes should only configure the 192.31.196.1 nameserver address; IPv6-only AS112 nodes should only configure the 2001:4:112::1 nameserver address.

It is only necessary for a single AS112 server operator to implement these extensions for this mechanism to function as intended. It is beneficial if many more than one AS112 server operator makes these changes, however, since that provides for greater distribution and capacity for the nameservers serving the EMPTY.AS112.ARPA zone. It is not necessary for all AS112 server operators to make these changes for the mechanism to be viable.

Detailed instructions for the implementation of these extensions are included in [RFC7534].

3.2. Redirection of Query Traffic to AS112 Servers

Once the EMPTY.AS112.ARPA zone has been deployed using the nameservers described in Section 3.1, redirections may be installed in the DNS namespace for queries that are intended to be answered by the AS112 infrastructure.

For example, reverse queries corresponding to TEST-NET-1 (192.0.2.0/24) [RFC5737] could be redirected to AS112 nameservers by installing a DNAME resource record in the 192.IN-ADDR.ARPA zone, as illustrated in Figure 1.

```
$ORIGIN 192.IN-ADDR.ARPA.  
...  
2.0      IN      DNAME    EMPTY.AS112.ARPA.  
...
```

Figure 1

There is no practical limit to the number of redirections that can be configured in this fashion. Redirection of a particular part of the namespace to EMPTY.AS112.ARPA can be removed at any time, under the control of the administrators of the corresponding part of the DNS namespace. No changes to deployed AS112 nodes incorporating the

extensions described in this document are required to support additional redirections. A list of possible candidates for AS112 redirection can be found in Section 5.

DNAME resource records deployed for this purpose can be signed with DNSSEC [RFC4033], providing a secure means of authenticating the legitimacy of each redirection.

4. Continuity of AS112 Operations

Existing guidance to AS112 server operators to accept and respond to queries directed at the PRISONER.IANA.ORG, BLACKHOLE-1.IANA.ORG, and BLACKHOLE-2.IANA.ORG nameservers should continue to be followed, and no changes to the delegation of existing zones hosted on AS112 servers should occur. These measures are intended to provide continuity of operations for zones currently delegated to AS112 servers and avoid any accidental client impact due to the changes proposed in this document.

Once it has become empirically and quantitatively clear that the EMPTY.AS112.ARPA zone is well hosted to the extent that it is thought that the existing, unmodified AS112 servers host 10.IN-ADDR.ARPA, the decision might be made to replace the delegation of those [RFC1918] zones with DNAME redirection. Once implemented, the PRISONER.IANA.ORG, BLACKHOLE-1.IANA.ORG, and BLACKHOLE-2.IANA.ORG nameservers could be retired. This document gives no such direction to the IANA, however.

5. Candidate Zones for AS112 Redirection

All zones listed in [RFC6303] are candidates for AS112 redirection.

Since no pre-provisioning is required on the part of AS112 operators to facilitate sinking of any name in the DNS namespace by AS112 infrastructure, this mechanism supports AS112 redirection by any zone owner in the DNS.

This document is simply concerned with provision of the AS112 redirection service and does not specify that any particular AS112 redirection be put in place.

6. DNAME Deployment Considerations

DNAME was specified years after the original implementations of [RFC1035], and hence universal deployment cannot be expected. [RFC6672] specifies a fallback mechanism that makes use of synthesised CNAME RRSets for this reason. The expectation that design choices in the DNAME specification ought to mitigate any lack of deployment is reviewed below. Experimental validation of those expectations is included in Appendix A.

It is a fundamental design requirement of AS112 service that responses be cached. We can safely declare DNAME support on the authoritative server to be a prerequisite for DNAME redirection, but the cases where individual elements in resolver chains do not support DNAME processing deserve closer examination.

The expected behaviour when a DNAME response is supplied to a resolver that does not support DNAME is that the accompanying, synthesised CNAME will be accepted and cached. Re-query frequency will be determined by the TTLs (Time to Live) returned by the DNAME-responding authoritative servers.

Resolution of the CNAME target is straightforward and functions exactly as the AS112 project has operated since it was deployed. The negative caching [RFC2308] of the CNAME target follows the parameters defined in the target zone, EMPTY.AS112.ARPA. This has the side effects that all redirected names ultimately landing on an AS112 node will be negatively cached with the same parameters, but this lack of flexibility seems non-controversial; the effect of reducing the negative cache TTL would be increased query volume on the AS112 node operator concerned, and hence controls seem well aligned with operation.

Validating resolvers (i.e., those requesting and processing DNSSEC [RFC4033] metadata) are required to implement DNAME and hence should not make use of synthesised CNAME RRs. The lack of signature over a received CNAME RR should hence not limit the ability to sign the (DNAME) redirection point, and for those (DNAME) signatures to be validated.

In the case where a recursive server implements DNAME but DNAME is not implemented in a stub resolver, CNAME synthesis will again provide a viable path.

DNAME support on AS112 nodes themselves is never required under this proposal.

7. IAB Statement Regarding This .ARPA Request

With the publication of this document, the IAB approves of the delegation of 'AS112' in the ARPA domain. Under [RFC3172], the IAB has requested that IANA delegate and provision "AS112.ARPA" as specified in this specification. However, the IAB does not take any architectural or technical position about this specification.

8. IANA Considerations

8.1. Address Assignment

Per this document, IANA has assigned IPv4 and IPv6 number resources in conformance with Section 4 of [RFC2860].

The IANA has assigned one IPv4 /24 netblock and registered its use in the "IANA IPv4 Special-Purpose Address Registry" [RFC6890] as follows:

Name	Value
Address Block	192.31.196.0/24
Name	AS112-v4
RFC	RFC 7535
Allocation Date	2014-12
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

IANA has assigned one IPv6 /48 netblock and registered its use in the "IANA IPv6 Special-Purpose Address Registry" [RFC6890] as follows:

Name	Value
Address Block	2001:4:112::/48
Name	AS112-v6
RFC	RFC 7535
Allocation Date	2014-12
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

8.2. Hosting of AS112.ARPA

The IANA hosts and signs the zone AS112.ARPA using nameservers and DNSSEC signing infrastructure of their choosing, as shown in Figure 2. SOA RDATA may be adjusted by the IANA to suit their operational requirements.

```
$ORIGIN AS112.ARPA.
$TTL 3600
```

```
@      IN      SOA      BLACKHOLE.AS112.ARPA. NOC.DNS.ICANN.ORG. (
                                1              ; serial
                                10800           ; refresh
                                3600            ; retry
                                1209600        ; expire
                                3600 )          ; negative cache TTL

                                NS      A.IANA-SERVERS.NET.
                                NS      B.IANA-SERVERS.NET.
                                NS      C.IANA-SERVERS.NET.

BLACKHOLE      A      192.31.196.1
                AAAA   2001:4:112::1

HOSTNAME        NS      BLACKHOLE

EMPTY           NS      BLACKHOLE
```

Figure 2

8.3. Delegation of AS112.ARPA

The IANA has arranged delegation from the ARPA zone according to normal IANA procedure for ARPA zone management, to the nameservers used in carrying out the direction in Section 8.2. The whois contact information for the new record is specified by the IAB under [RFC3172].

9. Security Considerations

This document presents no known additional security concerns to the Internet.

For security considerations relating to AS112 service in general, see [RFC7534].

10. References

10.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, DOI 10.17487/RFC2308, March 1998, <<http://www.rfc-editor.org/info/rfc2308>>.
- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", RFC 6672, DOI 10.17487/RFC6672, June 2012, <<http://www.rfc-editor.org/info/rfc6672>>.
- [RFC7534] Abley, J. and W. Sotomayor, "AS112 Nameserver Operations", RFC 7534, DOI 10.17487/RFC7534, May 2015, <<http://www.rfc-editor.org/info/rfc7534>>.

10.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<http://www.rfc-editor.org/info/rfc1918>>.
- [RFC2860] Carpenter, B., Baker, F., and M. Roberts, "Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority", RFC 2860, DOI 10.17487/RFC2860, June 2000, <<http://www.rfc-editor.org/info/rfc2860>>.
- [RFC3172] Huston, G., Ed., "Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa")", BCP 52, RFC 3172, DOI 10.17487/RFC3172, September 2001, <<http://www.rfc-editor.org/info/rfc3172>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<http://www.rfc-editor.org/info/rfc4033>>.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, DOI 10.17487/RFC4786, December 2006, <<http://www.rfc-editor.org/info/rfc4786>>.

- [RFC5737] Arkko, J., Cotton, M., and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation", RFC 5737, DOI 10.17487/RFC5737, January 2010, <<http://www.rfc-editor.org/info/rfc5737>>.
- [RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC 6303, DOI 10.17487/RFC6303, July 2011, <<http://www.rfc-editor.org/info/rfc6303>>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<http://www.rfc-editor.org/info/rfc6890>>.

Appendix A. Assessing Support for DNAME in the Real World

To measure the extent to which the DNAME construct is supported in the Internet, we have used an experimental technique to test the DNS resolvers used by end hosts and derive from the test a measurement of DNAME support within the Internet.

A.1. Methodology

The test was conducted by loading a user's browser with four URLs to retrieve. The first three comprise the test setup, while the final URL communicates the result to the experiment controller. The URLs are:

- A `http://a.<unique_string>.dname.example.com/1x1.png?`
`a.<unique_string>.dname`
- B `http://b.dname.example.com/1x1.png?`
`b.<unique_string>.dname`
- C `http://c.<unique_string>.target.example.net/1x1.png?`
`c.<unique_string>.target`
- D `http://results.recorder.example.net/1x1.png?`
`results.<unique_string>?za=<a_result>&zb=<b_result>&zc=<c_result>`

The A URL is designed to test the end user's capability to resolve a name that has never been seen before, so that the resolution of this domain name will reliably result in a query at the authoritative nameserver. This is intended to test the use of domain names where there is a dynamic component that also uses the DNAME construct.

The B URL is deliberately designed to be cached by caching resolvers that are used in the process of resolving the domain name.

The C URL is a control URL. This is a unique URL, similar to A, but does not refer to a DNAME structure.

The D URL uses a static cacheable domain name.

The <unique_string> value is common to the four URLs used in each individual instance of this test but varies from test to test. The result is that each end user is presented with a unique string.

The contents of the EXAMPLE.COM, TARGET.EXAMPLE.NET, and RECORDER.EXAMPLE.NET zones are shown in Figure 3.

```
$ORIGIN EXAMPLE.COM.
...
DNAME.                IN  DNAME  TARGET.EXAMPLE.NET.
...

$ORIGIN TARGET.EXAMPLE.NET.
...
B                      IN  A      192.0.2.0
*                      IN  A      192.0.2.0
...

$ORIGIN RECORDER.EXAMPLE.NET.
...
RESULTS                IN  A      192.0.2.0
...
```

Figure 3

The first three URLs (A, B, and C) are loaded as tasks into the user's browser upon execution of the test's script. The script starts a timer with each of these URLs to measure the elapsed time to fetch the URL. The script then waits for the three fetches to complete, or 10 seconds, whichever occurs first. The script then loads the results of the three timers into the GET arguments of the D URL and performs a fetch to pass these results back to the experiment's server.

Logs on the web server reached at RESULTS.RECORDER.EXAMPLE.NET will include entries of the form shown in Figure 4. If any of the URLs fail to load within 10 seconds, the D URL will report the failure as a "null" timer value.

```
GET /1x1.png?results.<unique_string>?za=1822&zb=1674&zc=1582
GET /1x1.png?results.<unique_string>?za=null&zb=null&zc=161
```

Figure 4

The script has been encoded in Adobe Flash with a simple image in the form of an online advertisement. An online advertisement network has been used to distribute the script. The script is invoked when the advertisement is presented in the end user's browser or application and does not require the user to click on the supplied image in any way. The advertisement placement parameters were set to the broadest possible scope to sample users from across the entire Internet.

A.2. Results

The test was loaded into an advertisement distributed on 2013-10-10 and 2013-10-11.

	Count	Percentage
Recorded Results:	338,478	
A or B Loaded:	331,896	98.1%
A Fail and B Fail:	6,492	1.9%
A Fail and B Load:	4,249	1.3%
A Load and B Fail:	1,624	0.5%
C Fail:	9,355	2.8%

Table 1

These results indicate that at most 1.9% of tested clients use DNS resolvers that fail to resolve a domain name that contains a DNAME redirection. However, the failure rate of slightly lower than 3% for the control URL indicates that the failure rate for the DNAME construct lies within the bounds of error within the experimental framework. We conclude that there is no evidence of a consistent failure on the part of deployed DNS resolvers to correctly resolve a DNAME construct.

This experiment was conducted by Geoff Huston and George Michaelson.

Acknowledgements

The authors acknowledge the valuable contributions of Bob Harold and other participants in the DNSOP working group in the preparation of this document.

Authors' Addresses

Joe Abley
Dyn, Inc.
103-186 Albert Street
London, ON N6A 1M1
Canada

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

Brian Dickson
Twitter, Inc.

EMail: bdickson@twitter.com

Warren Kumari
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States

EMail: warren@kumari.net

George Michaelson
APNIC

EMail: ggm@apnic.net

