

Network Working Group	R. Bellis	
Internet-Draft	Nominet UK	
Intended status: Standards Track	A. Bligh	
Expires: April 18, 2010	Silverscale Associates Ltd	
	W. Wijngaards	
	NLnet Labs	
	October 15, 2009	

[TOC](#)

DNS Proxy Bypass by Recursive DNS Discovery and LOCAL.ARPA draft-bellis-dns-recursive-discovery-00

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 18, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document describes a method for a DNS client resolver to discover the IP addresses of the upstream recursive DNS resolvers and hence bypass the local DNS proxy. It also directs IANA to reserve the "LOCAL.ARPA" domain name and to create a registry for well known sub-

domains of that domain name, such sub-domains being reserved for use within any network's administrative boundary.

Table of Contents

- [1.](#) Introduction
- [2.](#) Terminology
- [3.](#) LOCAL.ARPA
- [4.](#) DOMAIN.LOCAL.ARPA - Server behaviour
- [5.](#) DOMAIN.LOCAL.ARPA - Client Behaviour
- [6.](#) LOCAL.ARPA - Proxy behaviour
- [7.](#) Security Considerations
- [8.](#) IANA Considerations
- [9.](#) IAB Considerations
- [10.](#) References
 - [10.1.](#) Normative References
 - [10.2.](#) Informative References

[Appendix A.](#) Change Log

[§](#) Authors' Addresses

1. Introduction

TOC

Client DNS resolvers [\[RFC1035\] \(Mockapetris, P., "Domain names - implementation and specification," November 1987.\)](#) usually must send their queries to a recursive resolver which performs the iterative lookups on the client's behalf and returns the complete result, often from a local cache.

However, particularly in consumer and small office networks, the client resolver often does not talk directly to a recursive resolver. A very common configuration is that the client talks to a 'proxy' embedded in their local internet access gateway which acts as a simple intermediary between the client and the recursive servers. The term 'proxy' is used within this document to indicate any device to which the queries generated by the client resolver are addressed at an IP level; as such

they may include devices such as NAT devices, firewalls, and DSL gateways.

This configuration is a side-effect of the need for the DHCP server embedded in such gateways to issue DNS server addresses before those addresses have been learnt from the upstream network (see Section 5.3 of [\[RFC5625\] \(Bellis, R., "DNS Proxy Implementation Guidelines," August 2009.\)](#)).

These proxies have however been found to be generally deficient in their implementation of the DNS protocols (see [\[SAC035\] \(Bellis, R. and L. Phifer, "Test Report: DNSSEC Impact on Broadband Routers and Firewalls," September 2008.\)](#), [\[RFC5625\] \(Bellis, R., "DNS Proxy Implementation Guidelines," August 2009.\)](#)), to the extent that modern DNS extensions may fail to work correctly. Common problems include failure to deal properly with large DNS packets (including poor support for EDNS0 (ref), poor support for IP fragments, and truncation due to apparent buffer size limitations), and failure to deal with unknown RR types.

Thus, whilst the devices may appear to be functional for standard DNS protocols, they may fail to properly process all queries, in particular DNSSEC ([Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements," March 2005.](#)) [\[RFC4033\]](#) queries, which use RR types that the device concerned may not understand, and typically have larger responses. Field tests indicate that such devices are in general far more competent at passing through DNS queries addressed directly to the recursive resolvers (and the replies to such queries) than they are at processing queries addressed directly to them.

This document therefore proposes a method whereby a client resolver may discover the IP addresses of the intended recursive resolvers such that subsequent queries may be sent directly to those resolvers without passing through the gateway's DNS proxy.

To support this method IANA are directed to reserve a new domain name ("LOCAL.ARPA") which is not present in the .ARPA zone file but exists only on the recursive DNS servers local to the client stub resolver concerned. IANA are also redirected to create a registry of well known sub-domains of "LOCAL.ARPA", and this document further directs IANA to record "DOMAIN.LOCAL.ARPA" in that registry.

2. Terminology

[TOC](#)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\] \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

[TOC](#)

3. LOCAL.ARPA

This document reserves "LOCAL.ARPA" for infrastructure use within the administrative boundaries of a local network. A local network for these purposes means, in respect of a given recursive DNS server, all those hosts which are permitted to use that server to make recursive queries. The exact boundaries of a local network are implementation dependent;. It could be a corporate network, but it could also be an ISP access network including all of the customer networks connected to it.

This domain name serves as an anchor point for the discovery of well known services within a network. Whilst other technologies have been described for the discovery of services belonging to a specific domain (TODO: DNS-SD ref), the intent of "LOCAL.ARPA" is to support discovery of services on the current local network (dependent on which recursive nameserver it queries), regardless of the local domain name(s).

Note that like "SINK.ARPA" (see [\[I-D.jabley-sink-arpa\] \(Abley, J. and O. Gudmundsson, "The Eternal Non-Existence of SINK.ARPA \(and other stories\)," May 2009.\)](#)), this domain name MUST NOT actually appear in the IANA maintained zone file for .ARPA. Queries for this domain name (and by extension any sub-domain thereof) which leak beyond the local network onto the global internet MUST receive an NXDOMAIN (RCODE == 3) response.

4. DOMAIN.LOCAL.ARPA - Server behaviour

[TOC](#)

A recursive server implementing this standard, in response to an A or AAAA query for the QNAME "DOMAIN.LOCAL.ARPA" MUST do exactly one of the following (subject to the normal situations where the server is permitted to return an error):

- (a) return a response containing a list of one or more A or AAAA records for the QNAME each of which is an IP address for a recursive name server that is configured to support recursive DNS lookups by the client which sent the initial query. The server MAY algorithmically generate such records; for instance, it may return one or more of its own IP addresses, for instance the destination IP address of the query packet it received or (after appropriate sanitisation to ensure the client can query them) a list of IP addresses of its own interfaces
- (b) return a CNAME record which, if followed, will yield a list as set out in (a) above. The server MAY algorithmically generate the CNAME record, for instance to encode the querying IP address within it in an implementation dependent manner. OR
- (c) return NXDOMAIN, without performing a query to the .ARPA nameservers.

A recursive server not implementing this standard will normally recursively query the .ARPA nameservers, which will result in an NXDOMAIN, which will be cached for future queries.

An example of how a network operator running the Unbound recursive resolver might configure this is as follows:

```
local-zone: "local.arpa." static
local-data: "domain.local.arpa. 3600 IN A 192.0.2.1"
local-data: "domain.local.arpa. 3600 IN A 192.0.2.2"
local-data: "domain.local.arpa. 3600 IN AAAA 2001:db8::1"
```

indicating (for the IPv4 case) that recursive resolvers may be found at 192.0.2.1 and 192.0.2.2.

5. DOMAIN.LOCAL.ARPA - Client Behaviour

[TOC](#)

Typically when the DNS client is first started it will use DNS settings obtained via a DHCP lease which contains the IP address of the local internet access gateway in the Domain Name Server Option (6).

The DNS client resolver bootstrapped (whether as above using DHCP or otherwise) would send the query via its local DNS proxy, and receive a new list of DNS servers.

Queries for A or AAAA records will as set out above either return NXDOMAIN (in the case where the recursive server does not support this standard, or where support is present but not configured), or an error code, or a list of A records or a CNAME which when followed will yield a list of A records. The list of A records however obtained will contain one or more IP addresses corresponding to recursive name servers that are configured to support recursive DNS lookups by the client which sent the initial query.

Client resolvers supporting this standard MUST be capable of following CNAMEs and MUST follow any CNAME returned in response to a query for "DOMAIN.LOCAL.ARPA".

If the client receives an NXDOMAIN, this indicates that the recursive server concerned does not support or is not configured to support this standard. The client MAY continue to use its currently configured DNS server IP addresses. It MAY repeat the query, but it MUST NOT issue such a repeat query for "DOMAIN.LOCAL.ARPA" for [60 seconds].

If the client receives an error or no response, this may be because the proxy does not have internet connectivity. The client MAY repeat the query, but it MUST NOT issue such a repeat query for "DOMAIN.LOCAL.ARPA" for [1 second]

If the client receives a list of one or more A or AAAA records, the client MAY then use these IP addresses as the destination for subsequent recursive DNS lookup queries in preference to those issued by the local DHCP server or otherwise configured.

If the A or AAAA records have a non-zero TTL, the client SHOULD treat the records as invalid after the TTL has expired. The client MAY make repeat queries but SHOULD NOT make such repeat queries until half the TTL returned has expired.

[TOC](#)

6. LOCAL.ARPA - Proxy behaviour

Proxies MUST NOT intercept queries to "LOCAL.ARPA" or its subdomains. In particular, proxies MUST NOT return NXDOMAIN or A, AAAA or CNAME records for queries to "LOCAL.ARPA" or its subdomains unless such a response is sent to the proxy by a recursive nameserver. A proxy for this purpose does not include a device which performs a full recursive caching nameservice which is compliant with [EDNS0 \(Vixie, P., "Extension Mechanisms for DNS \(EDNS0\)," August 1999.\)](#) [RFC2671], DNSSEC, TCP support and is fully capable of handling maximum size UDP packets whether fragmented or not.

A proxy MAY return SERVFAIL if it is aware it has no connectivity to a recursive nameserver without attempting to forward the packet concerned. For instance if the proxy device is a DSL access gateway and the DSL line by which it would reach the recursive server is down.

A proxy MUST pass UDP and TCP requests (and their responses) to recursive servers transparently and unmolested. This means that proxies MUST reassemble UDP fragments. As the proxy may find it hard to detect which packets are addressed to or from the recursive nameserver, this might be achieved by applying similar considerations to all packets. It is recognised that proxies which assiduously follow this section are unlikely to be the proxies which gave rise to the need for this standard.

7. Security Considerations

[TOC](#)

TODO

8. IANA Considerations

[TOC](#)

This document directs the IANA to add the following record to the [ARPA Reserved Names Registry \(Abley, J. and O. Gudmundsson, "The Eternal Non-Existence of SINK.ARPA \(and other stories\)," May 2009.\)](#)

[I-D.jabley-sink-arpa]:

Name	Purpose	RRTypes	Reference
LOCAL.ARPA	Locally administered infrastructure	NONE	This document (RFCXXXX) S. 3

This document further directs the IANA to create a new registry as follows:

Parameter	Value
Registry Name	ARPA Reserved Local Names
Reference	This document (RFCXXXX) Section 3
Registration Procedures	IETF Standards Action

with initial contents as follows:

Name	Purpose	RRTypes	Reference
DOMAIN	Recursive DNS Discovery	A / AAAA	This document (RFCXXXX) S. 4

9. IAB Considerations

[TOC](#)

The addition of "LOCAL.ARPA" to the ARPA Reserved Names Registry requires IAB approval.

Note that addition of sub-domains of "LOCAL.ARPA" to the ARPA Reserved Local Names Registry only requires IETF Standards Action.

10. References

[TOC](#)

10.1. Normative References

[TOC](#)

[I-D.jabley-sink-arpa]	Abley, J. and O. Gudmundsson, " The Eternal Non-Existence of SINK.ARPA (and other stories) ," draft-jabley-sink-arpa-00 (work in progress), May 2009 (TXT).
[RFC1035]	Mockapetris, P., " Domain names - implementation and specification ," STD 13, RFC 1035, November 1987 (TXT).
[RFC2119]	Bradner, S., " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[RFC5625]	Bellis, R., " DNS Proxy Implementation Guidelines ," BCP 152, RFC 5625, August 2009 (TXT).

10.2. Informative References

[TOC](#)

[RFC2671]	Vixie, P. , " Extension Mechanisms for DNS (EDNS0) ," RFC 2671, August 1999 (TXT).
[RFC4033]	Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, " DNS Security Introduction and Requirements ," RFC 4033, March 2005 (TXT).
[SAC035]	Bellis, R. and L. Phifer, " Test Report: DNSSEC Impact on Broadband Routers and Firewalls ," September 2008.

Appendix A. Change Log

[TOC](#)

NB: to be removed by the RFC Editor before publication.
draft-bellis-dns-recursive-discovery-00

Initial draft

Authors' Addresses

[TOC](#)

	Ray Bellis
	Nominet UK
	Edmund Halley Road
	Oxford OX4 4DQ
	United Kingdom
Phone:	+44 1865 332211
Email:	ray.bellis@nominet.org.uk
URI:	http://www.nominet.org.uk/
	Alex Bligh
	Silverscale Associates Ltd
	15 Elsyng Road
	London SW18 2HW
	United Kingdom
Phone:	+44 20 8812 3300
Email:	alex@alex.org.uk
	Wouter Wijngaards
	NLnet Labs
	Science Park 140
	Amsterdam 1098 XG
	The Netherlands
Email:	wouter@nlnetlabs.nl