

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 5, 2012

T. Hardie
Google
March 5, 2012

The Reachability Method (RM) DNS Resource Record
draft-hardie-rm-rr-00

Abstract

This draft proposes a DNS resource record for providing adjunct reachability methods for network hosts or resources which are only accessible within limited reachability domains.

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 September 2, 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.

Internet-Draft

The RM RR

March 2012

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
2.	Applicability Statement	4
3.	DNS Consideration	4
4.	The format of the RM RR	4
4.1.	Ownername, class, and type	4
4.2.	Priority	5
4.3.	Weight	5
4.4.	Target	5
5.	RM RDATA Wire Format	5
6.	Example Use	6
7.	Acknowledgements	7
8.	IANA Considerations	7
9.	Security Considerations	7
10.	References	8
10.1.	Normative References	8
10.2.	Informative References	8
	Author's Address	8

Internet-Draft

The RM RR

March 2012

1. Introduction

For the purposes of this document, a limited reachability domain (LRD) is a network segment, overlay network, or bounded network whose nodes are not generally reachable outside an administratively controlled domain. Enterprise networks are common examples. Many limited reachability domains use private address space or unrouted space.

A reachability method in this context is a mechanism that provides access to a limited reachability domain from outside the usual administrative boundary. This document expresses reachability methods by designating the host, port, and protocol used for access, using a URI. Note that few reachability methods currently have registered URI schemes, so effective deployment may be gated by the development of an effective set of interoperable designations of these schemes.

A common method for providing access to limited reachability domain is the point-to-multipoint virtual private network. VPNs in enterprise contexts commonly provide a single tunnel end point per client, configured at set-up. Once the tunnel has successfully connected, the client gets access to the internal network's full resources, including the internal view of any split DNS. Though effective, this method generally either limits the granularity of security (to what is often called a "crunchy on the outside, soft in the middle" approach) or requires additional configuration by the network operations team to limit reachability further after the clients have traversed the tunnel. Even in the cases where such additional configuration is completed (by assigning different VPN users to different VLANs, for example), it is generally per client or user rather than by application.

An alternative approach would be to provide different limited reachability domains for different internal resources. In that approach, a client would be connected to an LRD specific to the

resource required. Because each of those LRDs could use different access methods and be available via different ingress points, it would be valuable to have a general mechanism for distributing the reachability method needed for each resource. This draft proposes a DNS Resource Record for this purpose, with the assumption that it would be made available on the public side of any split DNS.

[1.1.](#) Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

Hardie

Expires September 2, 2012

[Page 3]

Internet-Draft

The RM RR

March 2012

[2.](#) Applicability Statement

Reachability Method (RM) records can be queried directly, but it is expected that they will commonly be returned as additional data by servers relating information about hosts that are located within a limited reachability domain (e.g. with queries for the A or AAAA record associated with a host within an enterprise network). While it is possible to associate an RM RR with a service name such as those used by SRV or URI RRs, this is not generally recommended, since the reachability information may vary for each target. It would also be possible to associate a reachability method with a wildcard target (like *.internal.example.net), but this would limit the advantage over straight configuration to load-balancing entry points according to the weight and priority of the targets given in the RM RRs.

[3.](#) DNS Consideration

If the reachability method varies over time, the TTL of the RM RR will need to be managed to match and coordinated with the TTL of the resource to be reached. If the reachability method varies according to other characteristics, something akin to split DNS must be managed, with the usual conflicts with the DNS's core loose consistency model.

[4.](#) The format of the RM RR

This is the presentation format of the RM RR:

Ownername TTL Class RM Priority Weight Target

The RM RR does not cause any kind of Additional Section processing.

[4.1.](#) Ownername, class, and type

The type number for the RM record is TBD (to be assigned by IANA).

The RM resource record owner name has no special considerations.

The RM resource record is class independent.

The RM resource record has no special TTL processing requirements, though the considerations on coordination stated above apply.

[4.2.](#) Priority

The priority of the target Reachability Method in this RR. Its range is 0-65535. A client **MUST** attempt to contact the URI with the lowest-numbered priority it can reach; RMs with the same priority **SHOULD** be tried in the order defined by the weight field.

[4.3.](#) Weight

A server selection mechanism. The weight field specifies a relative weight for entries with the same priority. Larger weights **SHOULD** be given a proportionately higher probability of being selected. The range of this number is 0-65535.

[4.4.](#) Target

The target is a URI, enclosed in double-quote characters (''). Resolution of the URI is according to the definitions for the Scheme of the URI. The URI is encoded as one or more <character-string> [RFC 1035 section 3.3 \[RFC1035\]](#).

5. RM RDATA Wire Format

The RDATA for a RM RR consists of a 2 octet Priority field, a two octet Weight field, and a variable length target field.

Priority and Weight are unsigned integers in network byte order.

The Target field contains the URI of the Reachability Method (without the enclosing double- quote characters used in the presentation format), encoded as a sequence of one or more <character-string> (as specified in [section 3.3 of RFC 1035](#), where all but the last <character-string> are filled up to the maximum length of 255 octets.

The Target field can also contain an IRI, but with the additional requirements that it is in UTF-8 [RFC 3629](#) [[RFC3629](#)] and possible to convert to a URI according to [section 3.1 of RFC 3987](#) [[RFC3987](#)] and back again to an IRI according to [section 3.2](#). Other character sets than UTF-8 are not allowed. The domain name part of the IRI can be either an U-LABEL or an A-LABEL as defined in [RFC 5890](#) [[RFC5890](#)].

```

      1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Priority           |           Weight           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
/                               /
/                               /
/                               /
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

6. Example Use

Probably the easiest way to understand how this RR might be used is

by contrasting it with happens with a user without it. Imagine a user is trying to use a financial application client from a home or hotspot network. The server for that client resides within the user's enterprise network. If invoked, the client will attempt to resolve the hostname "accounting.corp.example.com". In the currently common scenario, that hostname would not be visible from a home or hotspot because of split DNS. To use it, the user must recognize that the required host was within a corporate domain and initiate a VPN in order to succeed at the DNS request. It is likely that once that VPN was up, the point-to-multipoint nature of the VPN would allow the user to connect to accounting.corp.example.com with the financial application client. Once the VPN is up, however, other resources within corp.example.com's network are likely to be reachable to processes on the user's device.

The simplest implementation of the alternate method described here would be a wrapper script for the financial application client. The wrapper script looks up the hostname accounting.corp.example.com and gets an RM record along with the A or AAAA record. The script then creates the tunnel using the method defined in the RM method, which would be something like:

```
finance.example.com IN RM 10 1 "ssh://financetunnel.example.com:/"
```

prompting the user for authentication credentials appropriate to the target if need be. The script then bind the application to the tunnel interface when it starts, so that it can only reach that portion of the network accessible from financetunnel.example.com.

Real deployments based on wrapper scripts seem unlikely, given the need to manage authentication credentials and DNSSEC validation of the records returned, but the basic series of steps would be the same: get the reachability method associated with a resource; use it

to create an overlay network or tunnel; associate the calling application with the limited reachability domain made available via the tunnel or overlay.

[7.](#) Acknowledgements

The work on which this is based was co-authored with Tom Keitel. The

text for describing both the presentation and wire formats of the priority field and the weight field of this RR are lifted wholesale from the URI RR internet-draft submitted by Patrik Faltstrom and Olaf Kolkman URI-RR [[I-D.faltstrom-uri](#)]. This document's overall organization and IANA considerations are also largely derived from that draft. Harald Alvestrand kindly provided early feedback on this draft, despite his overall impression of its wisdom..

[8.](#) IANA Considerations

This memo asks IANA to register a new Resource Record Type, adding the line below, suitably amended, to the registry named Resource Record (RR) TYPES and QTYPES as defined in BCP 42RFC 5395 [[RFC5395](#)] and located at <http://www.iana.org/assignments/dns-parameters>.

TYPE	Value and meaning	Reference
-----	-----	-----
RM	TBD a URI for a service (per the owner name)	[RFCXXXX]

[9.](#) Security Considerations

The overall goal of the RM RR is to standardize a way to nominate a monkey-in-the-middle, so using it without a DNSSEC-based assurance that the data you have received is the data placed there by the zone administrator would be deeply unwise. Placing information about the approved monkey-in-the-middle into the public DNS also makes it a potential target of attack, both for denial of service and for infiltration

Why, then, would you take this approach? If you have a series of services which need to be reachable to users but which cannot themselves be safely be made accessible to the public Internet, you can use this approach to segment the reachability to each, using unique authentication or authorization decisions for the individual overlay networks. The authentication methods for the bastion hosts can also vary and be made arbitrarily strong, something which may not be possible for services which are being used but may not be

The choice of the URI format for the target is also somewhat problematic, as few candidate reachability methods have registered URI schemes. If the registered schemes are not available, it will be tempting to mint new or local schemes which may miss critical fetures of the actual reachability methods.

10. References

10.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC3987] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", [RFC 3987](#), January 2005.
- [RFC5395] Eastlake, D., "Domain Name System (DNS) IANA Considerations", [RFC 5395](#), November 2008.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), August 2010.

10.2. Informative References

- [I-D.faltstrom-uri]
Faltstrom, P. and O. Kolkman, "The Uniform Resource Identifier (URI) DNS Resource Record",
[draft-faltstrom-uri-06](#) (work in progress), October 2010.

Internet-Draft

The RM RR

March 2012

Author's Address

Ted Hardie
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: hardie@google.com

Hardie

Expires September 2, 2012

[Page 9]