

DNS Resource Records for Nimrod Routing Architecture

Status of this Memo

This document is an Internet-Draft. 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.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt' listing contained in the Internet-Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

This Internet Draft expires December 1995.

Abstract

This document describes two additional RR types for the Domain Name System[7,8] required to implement the Nimrod Routing Architecture[1]. These RRs record the Nimrod Locator and an Endpoint Identifier (EID) associated with a given Domain Name.

Introduction

Nimrod is a scalable internetwork routing architecture. The Nimrod architecture is designed to accommodate an internetwork of arbitrary size and with heterogeneous service requirements and restrictions and to admit incremental deployment throughout an internetwork. The key to Nimrod's scalability is its ability to represent and manipulate routing-related information at multiple levels of abstraction.

To do this efficiently, Nimrod separates the identification of communicating entities (Endpoints) from any topological information. Endpoint Identifiers (EIDs) are used to specify and uniquely identify entities connected to the network. Information about the topological location of an endpoint in the network is given by a Locator, which may change as the network topology changes.

During the initial deployment of the Nimrod system the mapping will be stored in the existing DNS system as two additional RRs on the Domain Name of the Endpoint. This document describes the two new RR types required to record this information.

Nimrod uses a hierarchy of abstract maps of (parts of) the network. A Locator is a topologically significant "name" for a Nimrod node, indicating where in the map hierarchy it can be found. Because it reflects location in the network, a node's Locator will change when the network topology changes. An EID is a short identifier for the endpoint of a communication (e.g. a host system) and has no structure or significance other than global uniqueness. An endpoint can retain the same EID forever, no matter where in the network it is located. Any given system has exactly one EID to identify it, but may have more than one Locator if it appears in multiple maps.

Updates of the EID and the Locator information will almost always be done through a Dynamic Update[2] protocol, triggered by normal Nimrod protocol operations. Except during testing, these will not be done by manual editing of a master file, which means that human readability is not a major concern.

1. definition of the RR types

Both of the RR types described in this document encode numbers whose structure (if any) is not meaningfully interpreted by the DNS system. Thus each is encoded as an uninterpreted string of octets. The interpretation of the values is described in the Nimrod protocol spec [[[TBD]]].

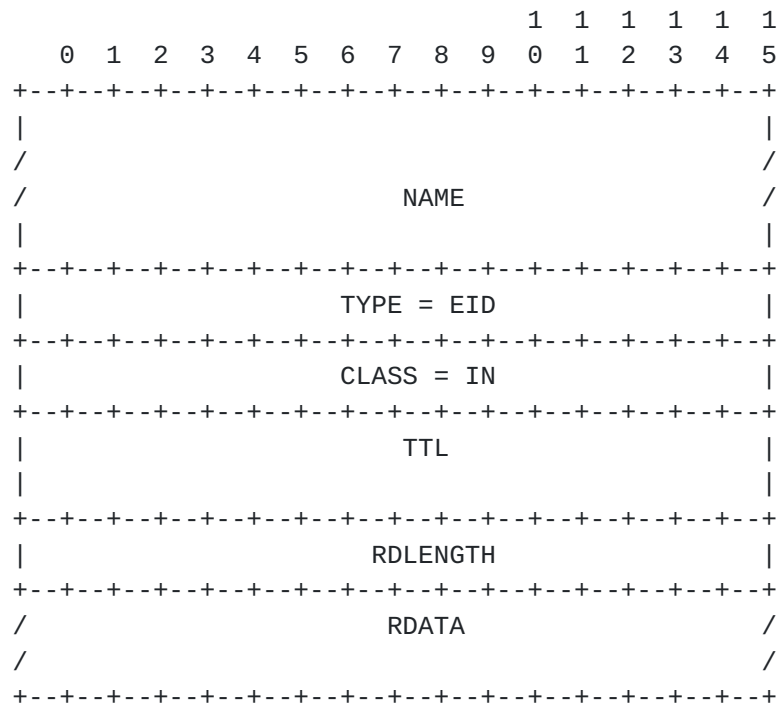
1.1. The EID (Endpoint Identifier) RR

The EID (Endpoint IDentifier) RR is defined with mnemonic "EID" and TYPE code 31 (decimal) and is used to map from domain names to EIDs. The EIDs declared in this RR may be used by any system that uses Endpoint Identifiers, but the initial use is intended for the Nimrod Routing system. EIDs are short, fixed length strings of octets whose content is meaningful to the Nimrod routing system. Since the top level RR format and semantics as defined in [Section 3.2.1 of RFC 1035](#) include a length indicator, the Domain Name System is not required to understand any internal structure.

An Endpoint can only have one unique identifier, so multiple different EID RRs at the same DNS name is an error. There are three ways to interpret such a condition when returned. If the conflict occurs when a reply is received from the authoritative server, that should be used and the existing (cached) RR should be discarded. The simplest, but less sure, way to deal with non-authoritative conflict is to ignore the RRs with the smaller

TTL and use the one with the longest remaining Time To Live. Secondly, the query can be retried at the authoritative server, with the result replacing the erroneous info as per the first item. Any caching server which is cognizant of EIDs should retain at most one EID RR as determined above, but legacy servers may not handle this requirement, so any system that needs to make use of EIDs must handle the conflict resolution described.

The format of an Endpoint Identifier (EID) RR is:



where:

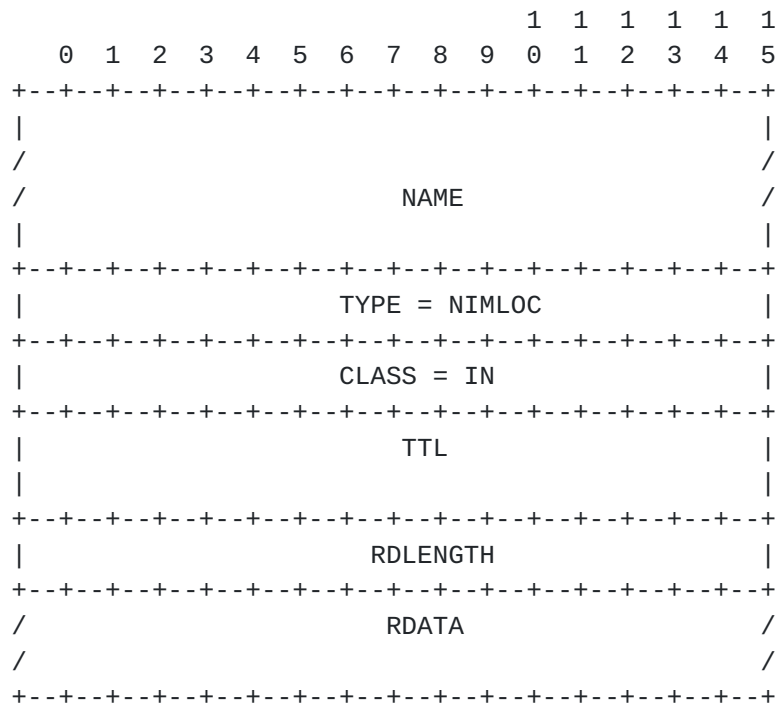
- * NAME: an owner name, i.e., the name of the DNS node to which this resource record pertains.
- * TYPE: two octets containing the EID RR TYPE code of 31 (decimal).
- * CLASS: two octets containing the RR IN CLASS code of 1.
- * TTL: a 32 bit signed integer that specifies the time interval in seconds that the resource record may be cached before the source of the information should again be consulted.
- * RDLENGTH: an unsigned 16 bit integer that specifies the length in octets of the RDATA field.
- * RDATA: a string of octets containing the Endpoint Identifier. The value is the binary encoding of the Identifier, meaningful only to the system utilizing it.

1.2. The NIMLOC (Nimrod Locator) RR

The NIMLOC (Nimrod Locator) RR is defined with mnemonic "NIMLOC" and TYPE code 32 (decimal) and is used to map from domain names to Nimrod Locators. Nimrod Locators are possibly variable length strings of octets whose content is only meaningful to the Nimrod routing system. Since the top level RR format and semantics as defined in [Section 3.2.1 of RFC 1035](#) include a length indicator, the Domain Name System is not required to understand any internal structure.

A Nimrod system may have any number of Locators associated with it. They are in this sense like A and AAAA RRs for IPv4 and IPv6 addresses. Multiple NIMLOC RRs with the same NAME, CLASS and RDATA are the same and can be merged in a cache, retaining only the highest TTL.

The format of a Nimrod Locator (NIMLOC) RR is:



where:

- * NAME: an owner name, i.e., the name of the DNS node to which this resource record pertains.
- * TYPE: two octets containing the NIMLOC RR TYPE code of 32 (decimal).
- * CLASS: two octets containing the RR IN CLASS code of 1.
- * TTL: a 32 bit signed integer that specifies the time interval in seconds that the resource record may be cached before the source of the information should again be consulted.

- * RDLENGTH: an unsigned 16 bit integer that specifies the length in octets of the RDATA field.
- * RDATA: a variable length string of octets containing the Nimrod Locator. The value is the binary encoding of the Locator specified in the Nimrod protocol[[[ref to be supplied]]].

2. Additional Section Processing

DNS servers cognizant of EID and NIMLOC type RRs should return these records in the Additional Section of any response including an A or AAAA type RR. These could be in response to either A or AAAA type queries, or some other query (e.g. NS) that specifies A and/or AAAA records in the Additional Section. This is not required for operation of the Nimrod system, as additional queries can always be made, but, in general, any time an A or AAAA RR will be used by a Nimrod agent, it will also need the EID and Locator info.

3. Master File Format

The format of NIMLOC and EID RRs follows all the rules of [RFC 1035, Section 5](#), "Master Files." The RDATA portion of both the NIMLOC and EID records contains uninterpreted binary data. The representation in the text master file is an even number of hex characters (0 to 9, a to f), case is not significant.

Example master file with NIMLOC and EID RRs (based on the example in [RFC1035](#)):

```
@   IN   SOA      VENERA      Action\.domains (
                                20      ; SERIAL
                                7200    ; REFRESH
                                600     ; RETRY
                                3600000; EXPIRE
                                60)    ; MINIMUM

      NS      A.ISI.EDU.
      NS      VENERA
      NS      VAXA
      MX      10      VENERA
      MX      20      VAXA

A      A      26.3.0.103
      EID     E32C6F78163A9348
      NIMLOC  32251A030067

VENERA A      10.1.0.52
      A      128.9.0.32
      EID     813F4B7CDAB34217
      NIMLOC  3227450A010034
      NIMLOC  75234159EAC457800920
```

VAXA	A	10.2.0.27
	A	128.9.0.33
	EID	3141592653589793
	NIMLOC	75234159EAC457800921

4. Acknowledgements

I'd like to thank the members of the DNS and Nimrod mailing lists for their comments and suggestions, and to the Nimrod Architects for the documents on which this is based. I'd also like to thank Arnt Gulbrandsen for his collected list of DNS RFCs and permission to use it as the basis for the References section and Bill Manning, the author of [RFC1348](#), for unwittingly supplying the boilerplate and diagrams I used as a basis for this document.

Specific thanks to Robert Elz, Masataka Ohta, and Martha Steenstrup for their helpful comments on early drafts.

[[[more?]]].

5. Security Considerations

Security issues are not discussed in this memo.

6. References

- [1] [draft-ietf-nimrod-routing-arch-00.txt](#): I. Castineyra, J. N. Chiappa, M. Steenstrup, "The Nimrod Routing Architecture", March 1995
- [2] [draft-ietf-dnsind-dynDNS-01.txt](#): S. Thomson, Y. Rekhter, J. Bound, "Dynamic Updates in the Domain Name System", March 1995
- [3] [RFC 1536](#): A. Kumar, J. Postel, C. Neuman, P. Danzig, S. Miller, "Common DNS Implementation Errors and Suggested Fixes.", 10/06/1993.
- [4] [RFC 1348](#): B. Manning, "DNS NSAP RRs", 07/01/1992.
- [5] [RFC 1183](#): R. Ullman, P. Mockapetris, L. Mamakos, C. Everhart, "New DNS RR Definitions", 10/08/1990.
- [6] [RFC 1101](#): P. Mockapetris, "DNS encoding of network names and other types", 04/01/1989.
- [7] [RFC 1035](#): P. Mockapetris, "Domain names - implementation and specification", 11/01/1987.

- [8] [RFC 1034](#): P. Mockapetris, "Domain names - concepts and facilities", 11/01/1987.
- [9] [RFC 1033](#): M. Lottor, "Domain administrators operations guide", 11/01/1987.
- [10] [RFC 1032](#): M. Stahl, "Domain administrators guide", 11/01/1987.
- [11] [RFC 974](#): C. Partridge, "Mail routing and the domain system", 01/01/1986.

[7.](#) Authors' Address:

Michael A. Patton
Bolt Beranek and Newman
10 Moulton Street
Cambridge, MA, 02138

Phone: (617) 873 2737
FAX: (617) 873 3457
Email: map@bbn.com

This Internet Draft expires December 1995