

Network Working Group
Request for Comments: 1788
Category: Experimental

W. Simpson
Daydreamer
April 1995

ICMP Domain Name Messages

Status of this Memo

This document defines an Experimental Protocol for the Internet community. This does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

IESG Note:

An Internet Engineering Steering Group comment from the co-Area Director for IPng: Please note well that this memo is an individual product of the author. It presents one view of the IN-ADDR mechanism, motivated by discussion in the IPNG WG of the difficulty of secure, dynamic update of the reverse tree. Other IETF discussion and ongoing standards work on this area will be found in the IP Next Generation (ipngwg), DNS IXFR, Notification, and Dynamic Update (dnsind), DNS Security (dnssec) working groups.

Abstract

This document specifies ICMP messages for learning the Fully Qualified Domain Name associated with an IP address.

Table of Contents

1.	Introduction	2
1.1	Direct Query	3
1.2	Multicast	3
1.3	Domain Names	3
1.4	Messages	4
2.	Domain Name Request	4
3.	Domain Name Reply	5
	SECURITY CONSIDERATIONS	6
	REFERENCES	6
	ACKNOWLEDGEMENTS	7
	AUTHOR'S ADDRESS	7

1. Introduction

The Domain Name System (DNS) is described in [RFC-1034]. The IN-ADDR domain of the DNS is specified [RFC-1035] to perform address to domain name resolution, and to facilitate queries to locate all gateways (routers) on a particular network in the Internet.

Neither function has been remarkably successful. The IN-ADDR domain is not reliably populated.

As multiple routers were used at boundaries and within networks, the IN-ADDR mechanism was found to be inadequate. The location of routers by hosts is now performed using "ICMP Router Discovery Messages" [RFC-1256].

As network numbers migrated to "classless" routing and aggregation, the IN-ADDR delegation granularity has fragmented, and requires overlapping administration. The "reverse" IN-ADDR administration frequently does not follow the same delegation as the "forward" domain name tree. This structure is not amenable to cooperative secure updating of the DNS.

As application servers have appeared which require the Domain Name for user interaction and security logging, the IN-ADDR servers have been inundated with queries. This produces long user visible pauses at the initiation of sessions.

1.1. Direct Query

This document proposes that each unicast address be queried directly for its corresponding Domain Name. This has the advantages that the naming is under the same administration as the address assignment, and the queries are distributed in the same fashion as IP routing. In effect, the routing is used to index the database.

1.2. Multicast

Only a few well-known multicast addresses are populated in the IN-ADDR domain. The ephemeral nature of most multicast addresses is not conducive to cooperative secure updating of the DNS.

However, the technique described here is not useful for multicast addresses. A query to a multicast address could result in a storm of replies. Most multicast groups are not named, or the member nodes are not configured with the name.

The IN-ADDR method SHOULD continue to be used for reverse lookup of well-known multicast addresses in the range 224.0.0.0 to 224.0.255.255. Other multicast addresses are an issue for further study.

1.3. Domain Names

Each Domain Name is expressed as a sequence of labels. Each label is represented as a one octet length field, followed by that number of octets. Since every Domain Name ends with the null label of the root, a Domain Name is terminated by a length byte of zero. The most significant two bits of every length octet must be '00', and the remaining six bits of the length field limit the label to 63 octets or less.

When the most significant two bits of the length octet are '11', the length is interpreted as a 2 octet sequence, indicating an offset from the beginning of the message (Type field). Further details are described in [RFC-1035] "Message Compression".

To simplify implementations, the total length of a Domain Name (including label octets and label length octets) is restricted to 255 octets or less.

1.4. Messages

The datagram format and basic facilities are already defined for ICMP [RFC-792].

Up-to-date values of the ICMP Type field are specified in the most recent "Assigned Numbers" [RFC-1700]. This document concerns the following values:

- 37 Domain Name Request
- 38 Domain Name Reply

2. Domain Name Request

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Code      |      Checksum      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Identifier      |      Sequence Number      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Type 37

Code 0

Checksum The ICMP Checksum.

Identifier If Code is zero, a value to aid in matching requests and replies. For example, it might be used like a port in TCP or UDP to identify a session. May be zero.

Sequence Number If Code is zero, a value to aid in matching requests and replies. For example, the number might be incremented on each request sent. May be zero.

A separate Domain Name Request is used for each IP Destination queried.

An ICMP Domain Name Request received with a broadcast or multicast Destination MUST be silently discarded.

On receipt of an ICMP error message, the implementations MAY attempt to resolve the Domain Name using the IN-ADDR method.

3. Domain Name Reply

```

+-----+-----+-----+-----+-----+-----+-----+-----+
|   Type   |   Code   |           Checksum           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Identifier          |           Sequence Number    |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Time-To-Live                      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|   Names ...                   |
+-----+-----+-----+-----+

```

Type 38

Code 0

Checksum The ICMP Checksum.

Identifier Copied from the request.

Sequence Number Copied from the request.

Time-To-Live The number of seconds that the name may be cached.
For historic reasons, this value is a signed 2s-
complement number.

Names zero or more Fully Qualified Domain Names. The
length of this field is determined from the total
length of the datagram.

When no names are known, the field is eliminated
(zero length), but the Reply is sent as an
authoritative indication that no name is known.

When more than one name is known, all such names
SHOULD be listed.

Any name which cannot entirely fit within the Reply
MTU is not sent.

The IP Source in a Reply MUST be the same as the IP Destination of
the corresponding Request message.

Every host and router MUST implement an ICMP Domain Name server
function that receives Domain Name Requests and sends corresponding
Domain Name Replies.

A host SHOULD also implement an application- layer interface for sending a Domain Name Request and receiving a Domain Name Reply, for diagnostic purposes.

Security Considerations

A primary purpose of this specification is to provide a mechanism for address to name resolution which is more secure than the IN-ADDR reverse tree. This mechanism is amenable to use of the IP Security Protocols for authentication and privacy.

Although the routing infrastructure to the Destination does not provide security in and of itself, it is as least as reliable as delivery of correspondence for the other sessions with the same peer.

A DNS cryptographic signature, located by using the reply in the forward DNS direction, can be used to verify the reply itself.

References

[RFC-792]

Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, USC/Information Sciences Institute, September 1981.

[RFC-1034]

Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, USC/Information Sciences Institute, November 1987.

[RFC-1035]

Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, USC/Information Sciences Institute, November 1987.

[RFC-1256]

Deering, S., Editor, "ICMP Router Discovery Messages", RFC 1256, Xerox PARC, September 1991.

[RFC-1700]

Reynolds, J., and J. Postel, "ASSIGNED NUMBERS", STD 2, RFC 1700, USC/Information Sciences Institute, October 1994.

Acknowledgements

The DNSIND and IPng Working Groups contributed substantial amounts of discussion.

Additional comments should be submitted to the namedroppers@internic.net mailing list.

Author's Address

Questions about this memo can also be directed to:

William Allen Simpson
Daydreamer
Computer Systems Consulting Services
1384 Fontaine
Madison Heights, Michigan 48071

Bill.Simpson@um.cc.umich.edu
bsimpson@MorningStar.com

