

Network Working Group
Internet-Draft
Expires: September 1, 2003

J. Arkko
Ericsson
March 3, 2003

Effects of ICMPv6 on IKE
draft-arkko-icmpv6-ike-effects-02.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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/1id-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 September 1, 2003.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

The ICMPv6 protocol provides many functions which in IPv4 were either non-existent or provided by lower layers. IPv6 architecture also makes it possible to secure all IP packets using IPsec, even ICMPv6 messages. IPsec architecture has a Security Policy Database that specifies which traffic is protected, and how. It turns out that the specification of policies in the presence of ICMPv6 traffic is hard, particularly with ICMPv6 packets related to Neighbor Discovery. Sound looking policies may easily lead to loops: The establishment of security requires Neighbor Discovery messages which can not be sent since security has not been established yet. The purpose of this draft is to inform system administrators and IPsec implementors in which manner they can handle the ICMPv6 messages. Common

understanding of the way that these messages are handled is also necessary for interoperability, in case vendors hardcode such rules in to products.

Table of Contents

- [1.](#) Introduction [3](#)
- [2.](#) Terminology [4](#)
- [3.](#) Neighbor Discovery and ICMPv6 Tasks [5](#)
 - [3.1](#) Path MTU Discovery [5](#)
 - [3.2](#) Error Notification [5](#)
 - [3.3](#) Informational Notifications [5](#)
 - [3.4](#) Router and Prefix Discovery [5](#)
 - [3.5](#) Address Autoconfiguration [6](#)
 - [3.6](#) Duplicate Address Detection [6](#)
 - [3.7](#) Address Resolution [6](#)
 - [3.8](#) Neighbor Reachability Detection [6](#)
 - [3.9](#) Redirect [7](#)
 - [3.10](#) Router Renumbering [7](#)
- [4.](#) Factors Affecting the Policy Rules [8](#)
 - [4.1](#) Nature of the Addresses [8](#)
 - [4.2](#) Network Topology [8](#)
 - [4.3](#) Role in Estaliblishing Communications [9](#)
 - [4.4](#) Protecting the Infrastructure versus Communications . [10](#)
- [5.](#) Analysis of the ICMPv6 Messages [11](#)
 - [5.1](#) Destination Unreachable [11](#)
 - [5.2](#) Packet Too Big [11](#)
 - [5.3](#) Time Exceeded [11](#)
 - [5.4](#) Parameter Problem [11](#)
 - [5.5](#) Echo Request [11](#)
 - [5.6](#) Echo Reply [12](#)
 - [5.7](#) Redirect [12](#)
 - [5.8](#) Router Solicitation [12](#)
 - [5.9](#) Router Advertisement [12](#)
 - [5.10](#) Neighbour Solicitation [12](#)
 - [5.11](#) Neighbour Advertisement [12](#)
 - [5.12](#) Router Renumbering [13](#)
- [6.](#) Summary [14](#)
- [7.](#) Further Work [16](#)
 - Normative References [17](#)
 - Informative References [18](#)
 - Author's Address [18](#)
- [A.](#) Acknowledgements [19](#)
 - Intellectual Property and Copyright Statements [20](#)

Arkko

Expires September 1, 2003

[Page 2]

1. Introduction

The ICMPv6 [8] and IPv6 Neighbor Discovery [6] protocols provide many functions which in IPv4 were either non-existent or provided by lower layers. For instance, IPv6 implements address resolution using an IP packet, ICMPv6 Neighbour Solicitation message. In contrast, IPv4 uses an ARP message at a lower layer.

IPv6 architecture makes it possible to secure all IP packets using IPsec [4], even ICMPv6 and Neighbor Discovery messages and even to multicast addresses. IPsec architecture has a Security Policy Database that specifies which traffic is protected, and how. It turns out that the specification of policies in the presence of Neighbor Discovery traffic is not easy. For instance, a simple policy of protecting all traffic between two hosts on the same network would trap even address resolution messages, leading to a situation where IKE ca not establish a Security Association since in order to send the IKE UDP packets one would have had to send the Neighbour Solicitation Message, which would have required an SA.

The purpose of this draft is to inform system administrators and IPsec implementors in which manner they can handle the Neighbor Discovery messages. System administrators do not want to study the IPv6 specifications in order to understand how they shall configure their routers. IPsec implementors want to understand what kind of policies they can offer with respect to the Neighbor Discovery messages.

Common understanding of the way that these messages are handled is also very much necessary for interoperability, as some vendors may be hardcoding some of the low-level policy operations in their products. If the rules between two vendors' products are incompatible for a particular message we may end with the sender sending cleartext and the receiver requiring IPsec, causing the packet to be dropped and possibly all connectivity between the two nodes lost.

This document does not imply any changes to the ICMPv6, Neighbor Discovery, IPsec, or IKE specifications. It is merely provided for configuration guidance.

2. Terminology

The keywords "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](#) [2].

3. Neighbor Discovery and ICMPv6 Tasks

In IPv6, ICMP has several tasks, and many of these tasks are overloaded on a few central message types such as the Neighbour Discovery message. In this chapter we explain the tasks and their effects in order to understand better how the messages should be treated.

3.1 Path MTU Discovery

Path MTUs are dynamically determined by IPv6 in order to optimize the size of the packets sent to a particular destination [1].

The ICMPv6 Packet Too Big messages [8] are used as a part of the Path MTU Discovery procedure.

3.2 Error Notification

ICMPv6 handles basic error situations of the IP layer, such as finding out that a particular destination is not available.

The Destination Unreachable, Packet Too Big, Parameter Problem, and Time Exceeded messages are a part of the error handling procedure [8]. Note that the Packet Too Big message also plays a role in the Path MTU Discovery procedure.

3.3 Informational Notifications

For debugging and network analysis purposes, ICMPv6 includes informational messages [8]. These message are necessary also in IPsec contexts and over IPsec tunnels due to the complex nature of some tunnel setups.

The Echo Request and Echo Reply messages are used solely for this purpose.

3.4 Router and Prefix Discovery

Router and prefix discovery is a part of the Neighbour Discovery protocol [6], which in turn is a part of the ICMPv6. The main purpose of the router discovery is to find neighboring routers that are willing to forward packets on the behalf of hosts. Prefix discovery involves determining which destinations are local for an attached link. This information is used both by the address autoconfiguration process, and routing. Typically, address autoconfiguration and other tasks can not proceed at all until the router discovery process has run.

The Router Solicitation and Router Advertisement messages are used for this and only this purpose.

[3.5](#) Address Autoconfiguration

Address autoconfiguration is another part of the Neighbour Discovery protocol [6]. Its purpose is to automatically assign addresses to interfaces. It comes in two variants, stateless and statefull. In this document we consider only the stateless autoconfiguration aspects. Obviously, no higher layer traffic can be sent until all participating nodes have addresses. This includes also IKE UDP traffic.

The Neighbour Solicitation and Advertisement messages are used for this purpose, among other things. Furthermore, Router and Prefix Discovery and Duplicate Address Detection have an effect to the Address Autoconfiguration tasks.

[3.6](#) Duplicate Address Detection

As a part of the stateless address autoconfiguration procedure, nodes check for duplicate addresses prior to assigning an address to an interface [7]. This procedure uses the same messages as the Neighbour Discovery protocol. Since the rules outlined in RFC 2462 [7] forbid the use of an address for both sending and receiving packets until it has been found unique, no higher layer traffic is possible until this procedure has completed.

The Neighbour Solicitation and Advertisement messages are used also for this purpose.

[3.7](#) Address Resolution

In address resolution, nodes determine the link-layer address of a local destination given only the destination's IP address [6]. Again, no higher level traffic can proceed until the sender knows the hardware address of the destination or the next hop router.

The Neighbour Solicitation and Advertisement messages are used also for this purpose.

[3.8](#) Neighbor Reachability Detection

Hosts monitor the reachability of local destinations and routers in the Neighbour Unreachability procedure, which is a part of the Neighbour Discovery protocol [6]. No higher level traffic can proceed if this procedure flushes out neighbour cache entries after (perhaps incorrectly) determining that the peer is not reachable.

The Neighbour Solicitation and Advertisement messages are used also for this purpose.

[3.9](#) Redirect

In the Redirect procedure, a router informs a host of a better first-hop node to reach a particular destination [6]. It is a part of the Neighbour Discovery protocol. As routers forward packets regardless of them being sent first to the wrong place, communications can still be established without the ability to process Redirect messages.

The Redirect message is used solely for the Redirect procedure.

[3.10](#) Router Renumbering

This procedure [9] allows address prefixes on routers to be configured and reconfigured in the similar manner as Neighbor Discovery and Address Autoconfiguration works for hosts. Incorrect processing or blocking of messages related to this procedure may render a node's address sets invalid, thereby preventing further communications.

The Router Renumbering message is used solely for the Router Renumbering procedure.

4. Factors Affecting the Policy Rules

4.1 Nature of the Addresses

Neighbor Discovery messages are sent using various kinds of source and destination address types. The nature of the destination address is of relevance here, as the destination address is used to find the right security association. The destination address can be either a well known multicast address, a computed multicast address, such as the solicited-node multicast address, or a unicast address. Many Neighbor Discovery messages use multicast addresses in most cases. Some messages can also be sent to unicast addresses in certain situations. For instance, the Neighbor Solicitation messages are usually sent to multicast addresses, but the Neighbor Advertisement messages are also sent to unicast addresses when sent as a response to a node that has an address.

ICMPv6 messages are sent using various kinds of source and destination address types. The source address is usually a unicast address, but during address autoconfiguration message exchanges, the unspecified address `::` is also used as a source address [7]. The destination address can be either a well known multicast address, a generated multicast address such as the solicited-node multicast address, or a unicast address. While many ICMPv6 messages use multicast addresses most of the time, some also use unicast addresses sometimes. For instance, the Neighbour Solicitation messages are usually sent to multicast addresses, but the Neighbour Advertisement messages are also sent to unicast addresses when sent as a response to a node that has an address.

IPsec [4] can be used for the protection of both unicast and multicast traffic. However, in order to automatically negotiate mutually acceptable security associations and to refresh keys, IKE [5] needs to be used. IKE is only capable of negotiating SAs for unicast communications.

Obviously, policies MUST be configured so that multicast traffic does not require dynamic SAs. However, while this is a necessary condition it is not sufficient to make sure that that IKE works. The policies MUST also exclude unicast traffic which contains ICMPv6 messages required before UDP can work between the two nodes.

4.2 Network Topology

ICMP traffic has different implications for hosts and security gateways. In general, security gateways SHOULD carry all ICMP traffic related to the protected traffic in the same tunnel as the traffic itself. For instance, when an ICMPv6 Packet Too Big message

is generated on the unprotected segment of a packet's path, that message should be relayed through the tunnel to ensure that the sender recognizes the MTU problem.

Between hosts similar rules apply. However, messages related to the establishment of communication between the hosts - such as for address resolution - MUST NOT be passed through the tunnel at least when the tunnel does not exist yet and IKE would be needed to establish it.

Note that the distinctions in network topology are more due to the actual network architecture than the selected IPsec mode, be it tunnel or transport.

ICMPv6 messages can be classified according to whether they are meant for end-to-end communications or communications within a link. There are also messages that we classify as 'any-to-end', which can be sent from any point within a path back to the source, typically to announce an error in processing the original packet. For instance, the address resolution messages are solely for local communications [6], whereas the Destination Unreachable messages are any-to-end in nature. End-to-end and any-to-end messages MUST always be passed through tunnels. Local messages may be passed through IPsec process under certain conditions.

4.3 Role in Establishing Communications

ICMPv6 messages can also be classified according to their role for establishing communications between two nodes. For the purposes of this discussion, the relevant issue is whether or not the messages must be passed through before IKE can use UDP packets to negotiate SAs. For instance, address autoconfiguration, duplicate address detection, and address resolution obviously MUST be completed before UDP packets can be passed.

Neighbour reachability detection is also capable of disrupting IKE communications. The reference [6] states the following:

In some cases (e.g., UDP-based protocols and routers forwarding packets to hosts) such reachability information may not be readily available from upper-layer protocols. When no hints are available and a node is sending packets to a neighbor, the node actively probes the neighbor using unicast Neighbor Solicitation messages to verify that the forward path is still working.

This means that unless the IKE implementation explicitly handles forward progress notifications towards the IPv6 stack, the stack can

not know about the reachability towards the other host. Since the hosts may be using tunnel mode and other address in the inner packets than the regular addresses on the hosts, the stack can not learn of forward progress through regular IPsec AH or ESP packets.

Therefore, neighbour reachability **MUST** also be allowed to work independent of IKE SA establishment.

As IKE messages may contain certificates, it is quite possible that an MTU limit may be exceeded somewhere within the network. If this is possible in a given network, the policies **MUST** allow ICMP Packet Too Big messages to be received. Note that these messages may well be received either in the clear, on manually configured SAs, or on dynamic SAs. If the router generating the Packet Too Big message does not yet have an SA with the original host, it can initiate IKE negotiations to create one. In case that this new negotiation fails due to reaching another MTU limit, other routers may be involved along the way. But ultimately the process reaches the closest router to which the MTU is known and will not cause any ICMP error messages.

[4.4](#) Protecting the Infrastructure versus Communications

IPsec can be used to protect the end-to-end communications or the underlying control messages (such as ICMPv6). It can even be used to protect both. Since many of the control messages are sent to multicast addresses, if IPsec is used then manual SA configuration **MUST** be performed instead of IKE-based SA negotiation.

As we have talked about some messages in some situations having to be independent of IKE, it does not necessarily imply that they have to be passed through in the clear. Instead, systems **MAY** use manually configured IPsec SAs to protect e.g. all ICMPv6 communications within one network. (Note that setting these manual SAs up requires some care as discussed in [\[13\]](#).)

A plausible security policy configuration could therefore be one where all ICMPv6 messages within the local network must be protected by manual SAs, and all other communications must be protected by IKE-negotiated SAs.

5. Analysis of the ICMPv6 Messages

5.1 Destination Unreachable

This message is always sent between unicast addresses [8]. It is an end-to-end message. Destination Unreachable is never a relevant message for establishing dynamic SAs, unless advanced failover schemes rely on the knowledge to quickly determine unreachable IKE peers.

5.2 Packet Too Big

This message is also always sent between unicast addresses [8] even if might be sent as a response to a multicast message. It is an end-to-end message.

Packet Too Big has, however, a role in establishing communications. End-to-end communications, that is. In order to pass through long IKE packets, Packet Too Big responses from the network MUST be considered. Therefore, it MUST be possible for policies to be configured so that such messages can be received. Note that as dicussed previously, the Packet Too Big messages themselves can be protected in various ways.

5.3 Time Exceeded

This message is also always sent between unicast addresses [8] and is an end-to-end message. Like Packet Too Big, it too has a role in establishing end-to-end communications under certain special situations.

5.4 Parameter Problem

This message is similar to Packet Too Big in the sense that it uses only unicast messages even if it could be sent as a response to a multicast packet. It's role is also end-to-end. While in theory its role in establishing communications is similar to Packet Too Big and Time Exceeded, in practise it is hard to see the kind of IKE and IPv6 stack version problem that could result in this message being sent.

5.5 Echo Request

Echo Request uses unicast addresses as source addresses, but may be sent to any legal IPv6 address, even multicast and anycast addresses [8]. Echo Requests run end-to-end but never have a role in establishing communications.

[5.6](#) Echo Reply

Echo Reply is similar to Echo Request in other respects, but uses only unicast addresses.

[5.7](#) Redirect

The Redirect message is always sent between unicast addresses [6]. It is only used for local purposes, not for end-to-end communications. It is not strictly necessary in order to establish communications. Nevertheless, it can be viewed as a logical add-on to the Neighbour Discovery messages such as Router Advertisement, and as such SHOULD be treated in a similar manner.

[5.8](#) Router Solicitation

This message uses either the unspecified address or an unicast address as a source address. The destination address is typically a multicast address. This message is always used only local. Since address autoconfiguration and routing depend on the ability of the routers and address prefixes to be found, this message is required before any communications can be established. Therefore, this message MUST be allowed to work independent of IKE SA establishment.

[5.9](#) Router Advertisement

This message has always a unicast source address, but the destination address can be either a unicast or a multicast address. Like the solicitation message, the advertisement is also link local only and required for establishing any communications. Therefore, this message MUST be allowed to work independent of IKE SA establishment.

[5.10](#) Neighbour Solicitation

The source address of this message is either a unicast address or (if Duplicate Address Detection is in progress) the unspecified address [6, 8]. The destination is either a multicast address, unicast address, or an anycast address. Neighbour Solicitation and Advertisement messages are used for multiple purposes: address autoconfiguration, duplicate address detection, and reachability detection. In all these roles they act only locally on the link, and getting them through is required before any communications can be established. Therefore, this message MUST be allowed to work independent of IKE SA establishment.

[5.11](#) Neighbour Advertisement

The source address of this message is a unicast address, and the

destination is either a unicast or a multicast address. Like the solicitation message, this message is link local only and is required before any communications can be established. Therefore, this message MUST be allowed to work independent of IKE SA establishment.

5.12 Router Renumbering

These messages are sent from a unicast address to either a multicast or a unicast address. The message are not solely link local, they are used for end-to-end purposes such as having a central management station renumber all routers in a corporate network. As a result of the RR procedure, automatically configured addresses and prefixes may be changed. However, it is expected that a transition period exists where both addresses are still acceptable, making it possible to still proceed with IKE negotiations to create SAs for the RR procedure. We can therefore assume that the procedure MAY use manual or dynamic SAs as desired by the system administrators.

6. Summary

Based on the above, the ICMPv6 messages can be classified as follows:

MESSAGE	ROLE	USE IKE?
Dest Unreachable	Any-to-End	MAY(1,2)
Packet Too Big	Any-to-End	MAY(1,3)
Time Exceeded	Any-to-End	MAY(1,3)
Parameter Problem	End-to-End	MAY(4)
Echo Request	End-to-End	MAY(4)
Echo Reply	End-to-End	MAY(4)
Redirect	Link Local	SHOULD NOT(5)
Router Solicit	Link Local	MUST NOT(6)
Router Advert	Link Local	MUST NOT(6)
Neighbour Solicit	Link Local	MUST NOT(6)
Neighbour Advert	Link Local	MUST NOT(6)
Router Renumbering	End-to-End	MAY(4)

Explanations:

- (1) These error messages have an end-to-end nature but may be generated by intermediate routers as well.
- (2) This MAY have to be considered by implementations that wish to base failover decisions on the Unreachable message.
- (3) These messages have an impact on the success of IKE messages e.g. when certificates are passed in IKE packets. It MUST be possible for policies to be configured so that these messages can be received while the IKE negotiations are still ongoing. Different security policy configurations MUST be supported, including trusting cleartext messages or protecting the messages from intermediate nodes using other, new dynamic SA negotiations.

- (4) These messages MAY be treated using regular IPsec and/or IKE processing.
- (5) This message SHOULD NOT use IKE in order to make their treatment equal with the rest of the link local messages, but in theory Redirect MAY be handled differently, e.g. using dynamic SAs.
- (6) These messages MUST NOT use dynamic SAs.

These policy rules may be expressed in various ways on a particular host or a router. It is necessary to use the ICMPv6 type in making the policy decisions. As [9] states, "This is consistent with, although not mentioned by, the Security Architecture specification". Only the following requirement for all implementations is stated here. Products that provide hardcoded security policies for ICMPv6 messages SHOULD enable user specified policies to be expressed at a higher priority level so that a possibility is still retained for modifying the rules due to e.g. interoperability problems.

7. Further Work

This draft discusses the use of IPsec on ICMPv6 messages on a principle level. It does not take a stand on how the policies are expressed, for instance whether IPsec products need to have hardcoded rules for handling these messages, or whether the Security Policy Databases should be general enough to make it possible to express the policies in them even for the ICMPv6 messages.

This draft does not address stateful address autoconfiguration aspects of IPv6.

This draft does not address the use of dynamic security associations in the context of multicast traffic. Now that the multicast key management working group has been founded in the IETF, a question eventually arises whether or not the results of that work can be used to protect the infrastructure multicast messages.

Normative References

- [1] McCann, J., Deering, S. and J. Mogul, "Path MTU Discovery for IP version 6", [RFC 1981](#), August 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [3] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), July 1998.
- [4] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [5] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998.
- [6] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [7] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [8] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 2463](#), December 1998.
- [9] Crawford, M., "Router Renumbering for IPv6", [RFC 2894](#), August 2000.
- [10] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 3041](#), January 2001.

Informative References

- [11] Arkko, J., Kempf, J., Sommerfeld, B. and B. Zill, "SEcure Neighbor Discovery (SEND) Protocol", [draft-ietf-send-ipsec-00.txt](#) (work in progress), February 2003.
- [12] Nikander, P., "IPv6 Neighbor Discovery trust models and threats", [draft-ietf-send-psreq-00](#) (work in progress), October 2002.
- [13] Arkko, J., "Manual SA Configuration for IPv6 Link Local Messages", [draft-arkko-manual-icmpv6-sas-01](#) (work in progress), June 2002.
- [14] Nikander, P., "Denial-of-Service, Address Ownership, and Early Authentication in the IPv6 World", Proceedings of the Cambridge Security Protocols Workshop, April 2001.

Author's Address

Jari Arkko
Ericsson
Jorvas 02420
Finland

E-Mail: jari.arkko@ericsson.com

[Appendix A](#). Acknowledgements

The author would like to thank Pekka Nikander, Markku Rossi, Tero Kivinen, Michael Richardson, Erik Nordmark, and James Kempf for interesting discussions in this problem space.

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.