Network Working Group
Internet-Draft
Updates: 3122 (if approved)
Intended status: Standards Track
Expires: April 26, 2009

P. Thubert, Ed.
E. Levy-Abegnoli
Cisco

October 23, 2008

IPv6 Inverse Neighbor Discovery Update
draft-thubert-3122bis-00.txt

Status of this Memo

Abstract

This draft updates the Inverse Discovery Specification [RFC3122] to provide Secure Neighbor Discovery. The behaviour of the protocol is slightly amended to enable an easier management of the addresses on a link and enable Secure ND.

Table of Contents

1. Introduction

This draft updates the Inverse Neighbor Discovery Specification [RFC3122]. Any behaviour or format that is not explicitely changed is preserved.

The behaviour of the protocol is slightly amended :

- Secure Inverse Neighbor Discovery is added for unicast addresses with a 64bit interface ID. This specification provides the additional options that are required to sign Inverse ND messages with the properties that are defined in [RFC3971] and details how they may be used to prove the ownership of advertised addresses.

- Fragmentation of ND messages is accepted but not required. [RFC3122] requires the use of multiple Advertisement messages when the Target Address List does not fit within MTU. With this specification, it is acceptable to fragment a message, but it is still possible to use multiple Advertisement messages, which can be necessary in particular in the context of Secure Inverse Neighbor Discovery.

- [RFC3122] does not allow a crisp management of all Addresses that a peer may use on an interface. When multiple Advertisement messages are used, it is possible to miss one and thus miss some information. With this specification, Address Management is improved in such a way that it is possible to advertise the addition or the deletion of a single address and to get the exhaustive list of all the addresses that a neighbor might use to source packets on an interface.

- With IPv4, Inverse ARP is traditionally applied to Point to Multipoint networks only. [RFC3122] claims to apply to "Frame Relay networks", and "also apply to other networks with similar behavior". This specification extends the domain of applicability of Inverse Neighbor Discovery and provides some additional

information on how and why Inverse ND MAY be used on all types of interface.

- Typos such as the length field in the Source/Target Address List are corrected.

The concept of transaction is introduced to group multiple messages into a single set to enable the advertisement of the complete list of all addresses used to source packets on on an interface. Whenever possible, a node should use one message per transaction.

This is problematic when:

- The list of addresses is so large that it causes the message to be larger than MTU and the node can not fragment.

- Secure Inverse ND is applied and not all of the addresses are based on a same CGA modifier (see [RFC3972]).

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].

2. Inverse Neighbor Discovery Messages

This section updates the Inverse Neighbor Discovery Messages defined in [RFC3122] section 2.

A new field from the ICMP reserved part is used to indicate the version and preserve backward compatibility. Version 0 is [RFC3122]. Version 1 is this specification. A node proposes a version in the Inverse Neighbor Discovery Solicitation Message and responds with the smallest of its own preferred version and the received proposed version in an Advertisement.

Another new field from the ICMP reserved part is used to indicate the Transaction ID of a Neighbor Discovery Solicitation Message, in order to correlate multiple Advertisement messages that may result from one Solicitation Message. A sequence number and a 'more' flag are also

added to enable the soliciting node to check that it actually received all the Advertisement messages for a given Transaction ID.

2.1. Inverse Neighbor Discovery Solicitation Message

The Inverse Neighbor Discovery Solicitation Message can be used to obtain the full list of IPv6 addresses from the remote end of a Point to Point link such as a PPP link, a tunnel or a Virtual Channel.

The Inverse Neighbor Discovery Solicitation can also be used as a heartbeat mechanism to verify whether a Point to Point link such as a tunnel is still up when there is no signal from the lower layers to indicate a failure.

The Inverse Neighbor Discovery Solicitation Message is changed as follows:

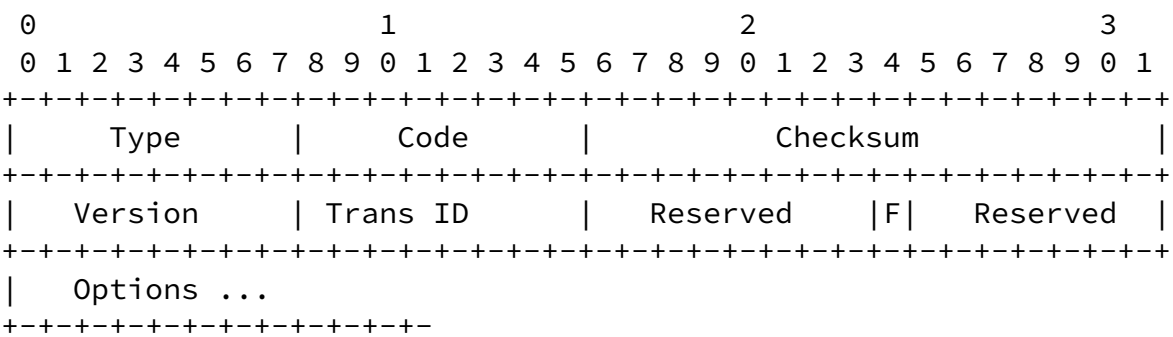

Modified Inverse Neighbor Discovery Solicitation Message

This specification adds the following fields:

Version: 8bit field. Version of 0 indicates the support of [RFC3122] only. Version of 1 indicates the desire to follow this specification and the backward compatibility with version 0.

Transaction ID: 8bit field. The Transaction ID is incremented with each Inverse Neighbor Discovery Solicitation Message sent to the

same neighbor. The transaction ID can not be set to zero so it starts at 1 and wraps from 255 directly to 1.

F: 1bit field. The "F" flag indicates a request of the Full List of addresses on the peer side of the Link.

2.2. Inverse Neighbor Discovery Advertisement Message

The Inverse Neighbor Discovery Advertisement Message can be used as a response to an Inverse Neighbor Discovery Solicitation Message or asynchronously at any point of time once a first Inverse Neighbor Discovery Solicitation Message was received indicating that the remote peer supports this specification.

Multiple Inverse Neighbor Discovery Advertisement Messages might be needed to report the full list of addresses. Those messages are correlated by a same transaction ID and sequenced. All Messages but the last have the More bit set to indicate that further Messages are to be expected for that transaction.

An unsolicited Advertisement is used to advertise the addition or the deletion of a single address and is contained in a single Inverse Neighbor Discovery Advertisement Message, with a transaction ID of 0.

The Inverse Neighbor Discovery Advertisement Message is changed as follows:

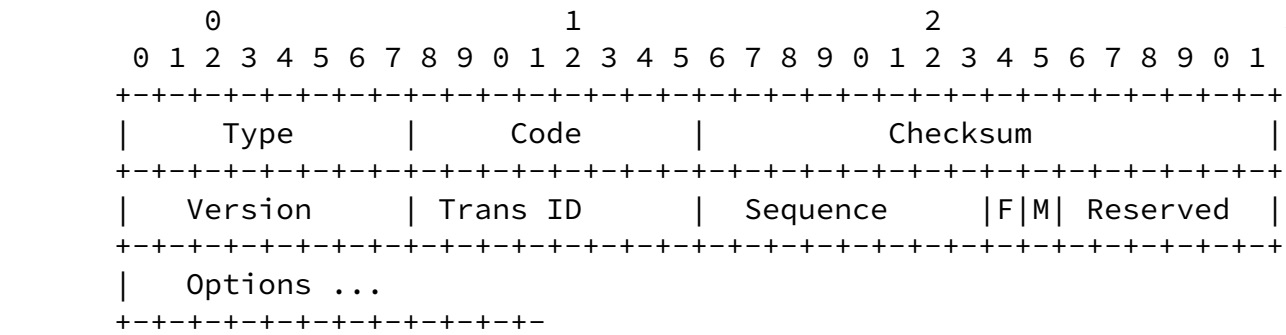

Modified Inverse Neighbor Discovery Advertisement Message

This specification adds the following fields:

Sequence: 8bit field. The sequence echoes that of the last received Inverse Neighbor Discovery Solicitation Message.

Version: 8bit field. Version of 0 indicates the support of [RFC3122] only. Version of 1 indicates the desire to follow this specification. Version can only be set to 1 if the Version in the Solicitation Message was 1 or above.

Transaction ID: 8bit field. The Transaction ID echoes that of the Inverse Neighbor Discovery Solicitation Message that this Message is responding to. The transaction ID zero is used for unsolicited Advertisements.

Sequence: 8bit field. The sequence is reset to zero for a new transaction ID and incremented with each Advertisement Message sent to a same Neighbor for a same Transaction ID. It is used to detect the loss of one Inverse Neighbor Discovery Advertisement Message in a Transaction that involved multiple ones.

M: 1bit field. The More "M" flag indicates that there are more messages for that transaction ID.

F: 1bit field. The "F" flag indicates that the Full List of addresses will be provided for that transaction.

Upon a request by the remote peer of the Full List of addresses, this node SHOULD answer with all the addresses that can be used to reach it over this link in the modified Target Address List.

If the IND Solicitation does not request the full list then his node MAY answer with all the addresses that can be used to reach it over this link in the modified Target Address List.

3. Inverse Neighbor Discovery Options

This section updates the Inverse Neighbor Discovery Options defined in [RFC3122] section 3.

3.1. Source/Target Address List

With this specification, the Source/Target Address List may be

partial or full. It can be used to indicate addition or deletion of indivisual addresses. This new support can only be used once the first Inverse Neighbor Discovery Solicitation Message is received from the remote peer indicating the support of this specification.

Until the Version that is supported by the peer is known, the only Inverse Neighbor Discovery Messages that this node should send are Solicitations, and this option if present can only be a Source Address List option with a list of addresses that can be used to reach this node over this link.

This specification uses a Length field of 8 bits, assuming that most implementations of [RFC3122] also use a Length field of 8 bits and that the misalignment in section 3.1 is commonly undestood as an undetected typo. An error in reading the Length field can be detected when confronting the length of the IPv6 packet and the expected length of this option.

The Inverse Neighbor Discovery Advertisement Message is changed as follows:

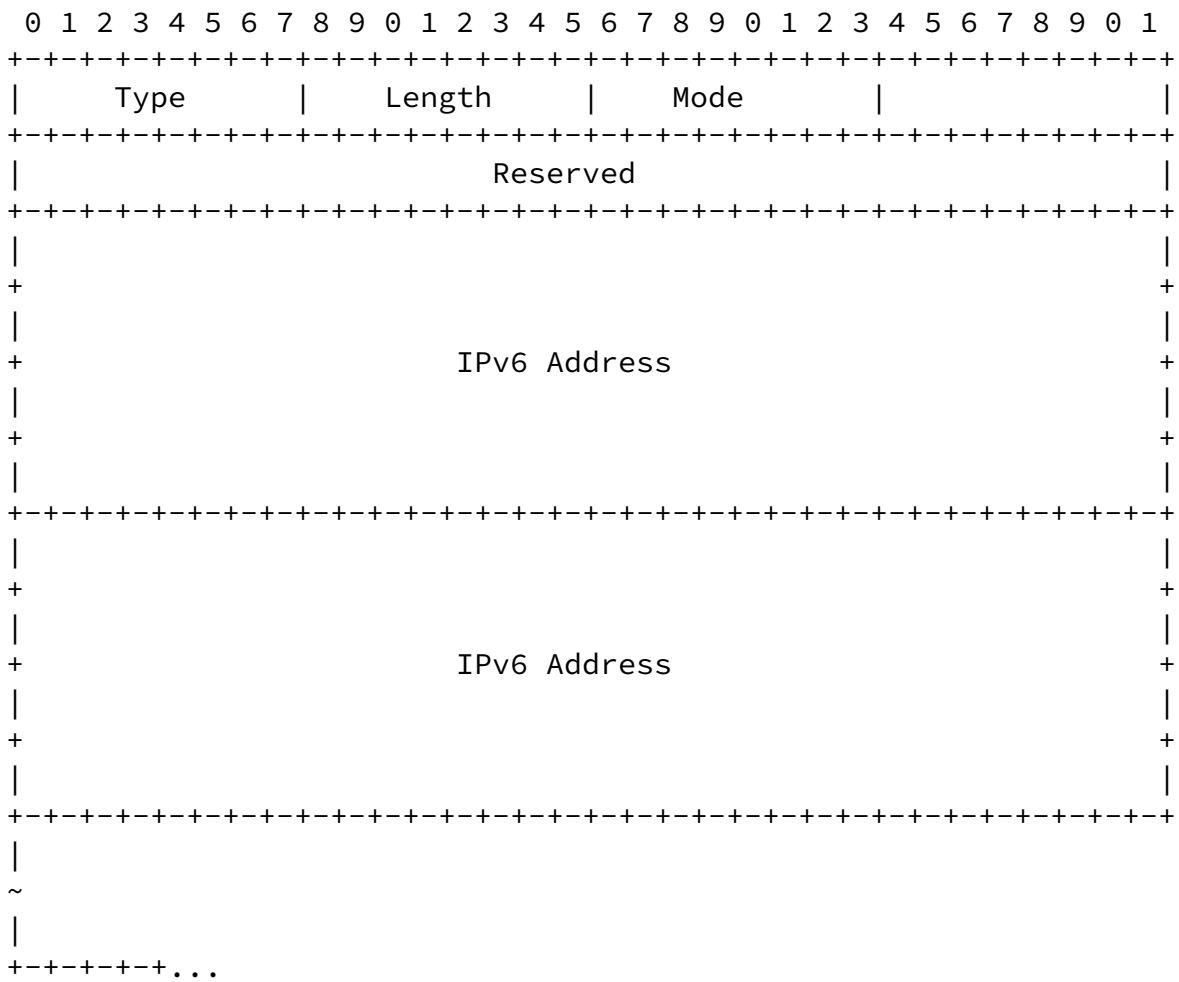

Modified Source/Target Address List option

This specification adds the following fields:

Mode 8bit enumeration

RFC3122: Mode of 0 indicates that the list is built conforming to [RFC3122]. All the addresses listed are usable but the list might not be complete.

Full: Mode of 1 indicates that the list might contain addresses that are defined on another interface but may be used to source or receive packets over this interface. This is the mode that is used to in reply to a Solicitation with the "F" bit set.

Added: Mode of 2 indicates that the list is a list of recently added addresses but not necessarily part of a full report.

Deleted: Mode of 3 indicates that the list is a list of recently removed addresses that may no more be used on this link.

Upon receiving an Address List, a node should verify that the addresses in the list don't collide with any of its own address. In case it does, the duplicate address received in the list will be ignored.

When Secure IND is being used, all the addresses listed in the Target Address List option in one Inverse Neighbor Discovery Advertisement Message must be based on the same CGA modifier. If multiple modifiers are used or some addresses were not built based on CGA, then they must be split in multiple Inverse Neighbor Discovery Advertisement Messages.

4. Secure Inverse Neighbor Discovery

The list of addresses provided in Source/Target address list can be defended using the CGA and the RSA options specified in [RFC3972] and [RFC3971]. However, in the case of Secure Inverse Neighbor Discovery, several addresses announced in one message (IND Solicitation or Advertisement) are defended by a single CGA option and a single RSA option. That mandates that all addresses in the list are based on CGA and were computed with the same modifier, the same collision count, and the same security parameter sec. In this case, the CGA option is as following:

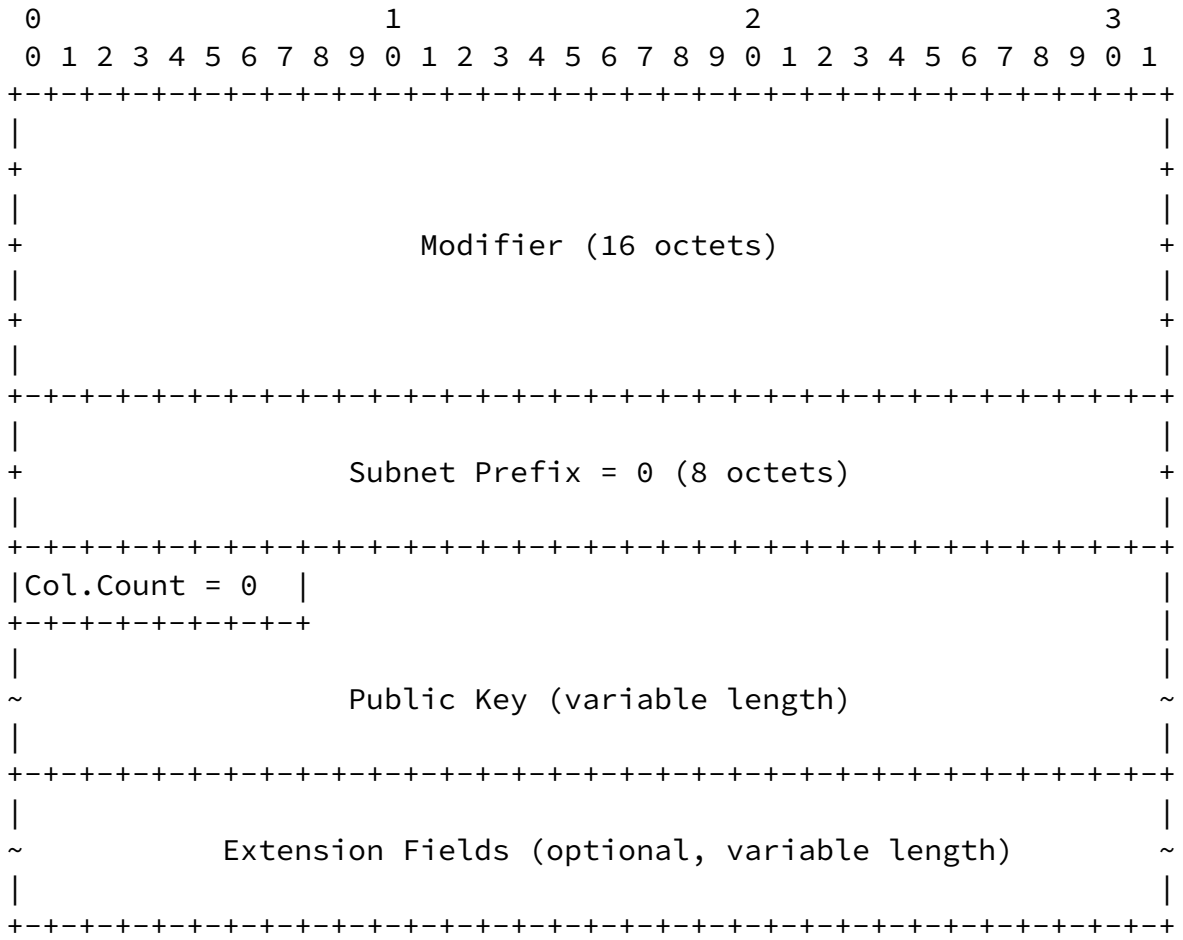

Since each address in the list is going to carry its own prefix (/64 if CGA is used), the Subnet Prefix in the CGA option is set to zero. Therefore, it should not be checked by the receiver against the prefix (/64 bits) of each CGA address found in the Address List.

The collision count is always zero, since no Duplicate Address Detection is performed, other than the node ignoring peer addresses colliding with one of its own.

The remaining of the CGA algorithm, as described in [RFC3972] applies. The 16*Sec leftmost bits of Hash2 should equal zero. Hash1 is computed separately for each address in the list, using the first 64 bits as the Subnet Prefix, and the interface identifier should

match Hash1 (except for bits 0, 1, 2, 6, and 7, which encode the collision count and the "u" and "g" bits).

The RSA option must also be provided in the message, and the signature must verify with the public key provided in the CGA option

In order to protect against replays, Timestamp and Nonce options, should also be used in the message, with similar rules as one described in [RFC3971]. When the message is a solicitation (INS), it

should have a nonce option. When the message is sollicited (INA as a response to INS), it should repeat the nonce value seen in the solicitation. As far as unsolicited message and solicitation, the timestamp option is required.

5. Interface Type and usages

Because of IPv4 and the ARP legacy, Inverse Neighbor Discovery is usually associated to Point to Multipoint (P2MP) or Non-Broadcast Multi-Access (NBMA) networks. And certainly, this specification is usable on such networks as Frame Relay or Multidrop tunnels.

But the similarity with IPv4 is limited and this specification enables a lot more:

5.1. Point to Multipoint Networks

IPv6 Secure Inverse Neighbor Discovery can be applied to P2MP and NBMA networks to prevent the theft an address by another Node.

5.2. Point-to-Point Links

As opposed to IPv4, using Inverse Neighbor Discovery makes a lot of sense on Point-to-Point link such as PPP or tunnels:

This specification inherits from [RFC3122] the support of the authentication header to authenticate the remote peer on the link. For P2P links, this might prove more relevant than Secure ND itself.

Because IPv6 operates purely at layer 3, the PPP Network Control Protocol for IPv6 defined in [RFC5072] provides a way to negotiate a

unique, 64bit interface identifier to be used for the address autoconfiguration but does not enable to advertise an IPv6 address. This would not fit anyway since IPv6 might use many addresses of various lifetimes on a same interface. This specification provides the means to create and maintain the list of addresses that can be used to reach the remote peer at any point of time.

A number of Denial of Service attacks are documented when using [RFC4861] by sending packets to addresses that are not assigned but belong to a prefix that is associated to the P2P link. On those links, Inverse Neighbor Discovery enables a proactive model that defeats those attacks. Any packet that is received for a destination that is not in the ND table is simply dropped.

5.3. Broadcast Networks

A multihomed host attached to a broadcast network might use an address that belongs to another interface on another subnet to source a packet. This makes the validation of source addresses very problematic. With this specification, a router may solicit the full list of all addreses that this host might use to source packets on that interface, and prove the ownership using SeND. The router might then accept packets that are sourced off-link and may install a host route to that address.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

This draft improves the security model in [RFC3122] by adding the capability to use Secure Neighbor Discovery

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC3122] Conta, A., "Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification", RFC 3122, June 2001.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

8.2. Informative References

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC5072] S.Varada, Haskin, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, September 2007.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com

Eric Levy-Abegnoli
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 20
Email: elevyabe@cisco.com

Full Copyright Statement

THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.