

mboned
Internet-Draft
Intended status: Experimental
Expires: January 12, 2009

T. Morin, Ed.
France Telecom - Orange Labs
B. Haberman
The Johns Hopkins University,
Applied Physics Laboratory
July 11, 2008

IGMP/MLD Error Feedback
draft-morin-mboned-igmpmld-error-feedback-01

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/lid-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 January 12, 2009.

Abstract

This document describes messages and procedures that can optionally be implemented in IGMP/MLD Queriers and Hosts, to provide information to multicast receivers on the reason why the IGMP/MLD Querier didn't take into account a Membership Report message.

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](#)].

Internet-Draft

IGMP/MLD Error Feedback

July 2008

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	History and problem statement	3
4.	Principle	4
5.	Procedures	4
5.1.	Procedures on the IGMP/MLD Querier	4
5.2.	Procedures on the IGMP/MLD Host	5
6.	Message encodings	6
6.1.	Feedback message	6
6.2.	Error codes	9
7.	Feedback to the application layer	9
8.	Impact on IGMP/MLD proxies and equipments doing IGMP/MLD snooping	11
8.1.	IGMP/MLD Proxies	11
8.2.	Equipments doing IGMP/MLD snooping	12
9.	IGMP/MLD Hosts stacks not implementing the Feedback mechanism	12
10.	IANA Considerations	12
11.	Security Considerations	12
12.	Acknowledgements	13
13.	References	13
13.1.	Normative References	13
13.2.	Informative References	14
Appendix A.	Protocol to carry error feedback messages	15
A.1.	ICMP	15
A.2.	IGMP/MLD	15
	Authors' Addresses	16
	Intellectual Property and Copyright Statements	17

Internet-Draft

IGMP/MLD Error Feedback

July 2008

1. Introduction

Requirements have been formulated for means to provide multicast receivers with error feedback when the IGMP/MLD Querier did not or could not process an IGMP/MLD Membership Report message ([[I-D.ietf-mboned-maccnt-req](#)], section 4). Operator's experience with IPTV deployments show that introducing such a feedback in IGMP exchanges between multicast receivers and multicast routing equipments would help provide a better service (e.g. a liaison between the IETF mboned working group and the DSLForum was made in December 2005 to discuss this issue, but didn't lead to a standardized solution).

An examples case is when an IGMP Querier refuses to take into account an IGMP Membership Report because the number of multicast channels would surpass the allowed threshold for the service. Many other examples of the case where an IGMP error feedback channel would be useful are presented in [Section 6.2](#).

This document describes new message encodings and the associated procedures, all of which being optional and preserving full backward and forward compatibility, and details the impact on the host API for multicast subscriptions.

This document doesn't state yet whether the messages should be carried over IGMP, ICMP or another protocol, but tries to document the pros and cons of the different alternatives, to guide the decision process.

2. Terminology

The terminology in this document is the terminology used in [[RFC3376](#)] and [[RFC3810](#)].

For readability, "Querier" and "Host(s)" will be used throughout this

document, in place of "IGMP or MLD Querier" and "IGMP or MLD Host(s)".

3. History and problem statement

The DSLForum expressed interest for such a feature, which was discussed [[magma-archive](#)] in a liaison with the Magma IETF Working group. The specifications described in this document try to address the comments exchanged on the magma WG mailing-list.

Morin & Haberman

Expires January 12, 2009

[Page 3]

Internet-Draft

IGMP/MLD Error Feedback

July 2008

4. Principle

The procedures described in this memo are fully optional : their only intent is to carry informative data from the Querier to the Hosts.

Most specifically, the intent is that:

- o the procedures don't change the state machine of the Querier or Host, the information carried is only meant to provide information to the application subscribed to multicast data
- o a Host implementing these specifications will behave correctly in the absence of these informations.
- o the behavior of a Querier implementing these specifications is unchanged whether or not the hosts implement these specs.

Last, these specifications are not meant to carry information about transient errors that the network is supposed to recover from, like network outages.

5. Procedures

5.1. Procedures on the IGMP/MLD Querier

The following procedures introduce a new type of message : the Feedback message. See section [Section 6](#) for details about message encodings.

Using these procedures a Querier can OPTIONALLY emit a Feedback message after receiving an IGMP or MLD Membership Report message that it can not process (see [Section 6.2](#) for example reasons on why a Querier would not process a Membership Report message).

The Feedback message carries error type/sub-type field, and information about the group to which the error pertains. Optionally, if IGMPv3/MLDv2 is used, and if the error message pertains to some specific sources, the addresses of the sources to which the error pertains are included in the message.

The address to which the Feedback message will be sent is determined as follows:

- o if IGMPv3/MLDv2 is used (and if the sender IP address is not 0.0.0.0 or 0::0), the address of the sender of the Membership Report is used

- o else, the group address specified in the Membership Report message is used

The source address MUST be the same address as the address used for Query messages, and the TTL MUST be set to 1.

If IGMPv2/MLDv1 is being used, not more than one Feedback message should be sent for a said Membership Report message.

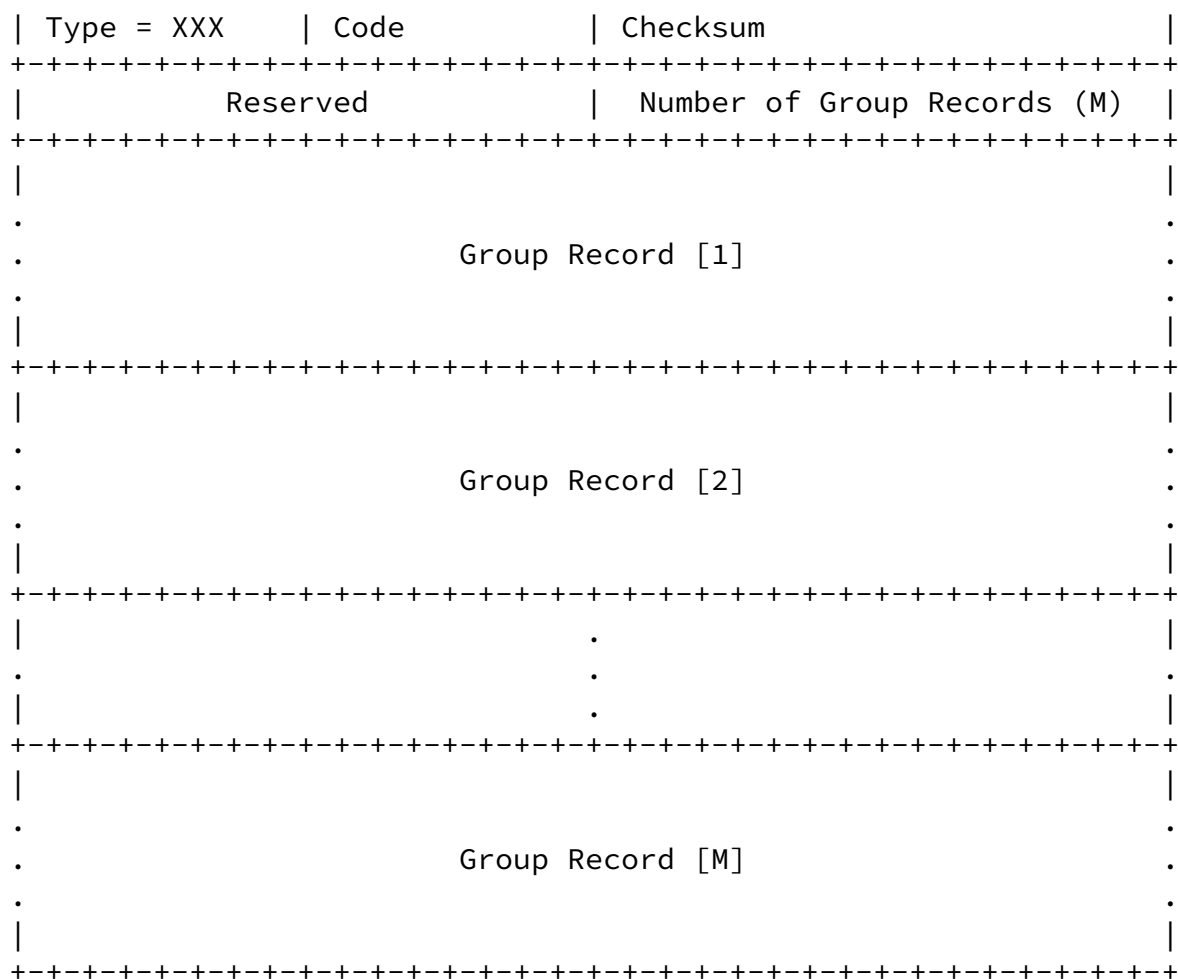
If IGMPv3/MLDv2 is being used, multiple feedback message MAY be sent if the group record of the IGMP/MLD message that triggered the error contained multiple source addresses.

In any case the amount of Feedback messages sent on a link MUST be rate-limited,

[5.2.](#) Procedures on the IGMP/MLD Host

When a Host receives an Feedback message, it MAY process it.

Processing a Feedback message consists in :



Fields:

Type: Identifies this message as a Feedback message. Currently using:

- * in the case of IPv6/MLD: 0xYY (currently using value 200 as defined in [RFC 4443](#) "private experimentation value", until another value is registered with IANA).
- * in the case of IPv4/IGMP: 0xZZ (currently using 0xF2, in the "Reserved for experimentation" range, until another value is registered with IANA)

- o the message that triggered the Feedback message is IGMPv3 or MLDv2 and the group record that triggered the error contains no source address
- o the message that triggered the Feedback message is IGMPv2 or MLDv1

[6.2.](#) Error codes

This section describes some proposed error codes:

- o improper message : the Membership Report message is improper (the group address is not in the 224/0 or FF00::/8 range, or specified sources are improper addresses)
- o IGMP or MLD version is not supported by querier
- o wildcard on an SSM group : IGMPv2 or IGMPv3/MLDv2 with an Exclude source filter mode was used in the Report, but the group address is not in the SSM range of the Querier
- o exclude source filter mode not supported by the Querier
- o group administratively prohibited
- o source(s) administratively prohibited
- o resource limit reached
- o multicast reception is disabled on the link
- o multicast routing protocol issue

[This section will later be completed to include error code numbers]

Remember that the Feedback message is NOT meant to carry information about transient errors that the network is supposed to recover from, like for instance network outages.

[7.](#) Feedback to the application layer

This section gives an example of how the information from Feedback messages is supplied to applications subscribed to multicast streams, and which expect the reception of multicast datagrams on a socket, based on Linux extensions to the POSIX [[posix](#)] network socket API.

A first requirement is full backward compatibility with applications

not supporting these specifications : an application not supporting these specifications must not be affected by a Feedback message. For instance, a wrong solution would be to return an error on a `read()` or `recv()` call.

A second requirement is to allow an application to keep receiving data on a socket, even if some error was reported through a Feedback message, for a group or channel joined on this socket. This is needed, for instance, in two cases : when a socket is used to join multiple distinct group or channels and only one of them is subject to an error ; when a socket is used to join only one multicast group, for which the Querier sends a Feedback message, but there are members for this group sending data on a directly connected link.

The proposed solution is to rely on the use of the `MSG_ERRQUEUE` flag of the `recvmsg()/recvfrom()` POSIX calls. This call allows the socket user to retrieve the network errors queued for the socket.

The MLD component receiving an MLD Feedback message containing error condition reports the error to the application via the `MSG_ERRQUEUE` flag in the `recvmsg()/recvfrom()` calls. The `MSG_ERRQUEUE` flag indicates the presence of a `sock_extended_err` data structure. When the `sock_extended_err` data structure is passed to the application, the `ee_origin` field is set to 3 (`SO_EE_ORIGIN_ICMP6`) in the case of an MLD Feedback message, and `XX` (`SO_EE_ORIGIN_YYYY`) in the case of an IGMP Feedback message [`XX` and `YYY` is to be determined in compliance with the relevant standard, 4 and `SO_EE_ORIGIN_IGMP` are proposed as interim values]. The Type and Code fields from the MLD Feedback message are copied into the `ee_type` and `ee_code` field of the `sock_extended_err` data structure.

The addresses of the multicast group and sources in error can be retrieved as follows:

- o in the IPv4 case, the group address and source address are stored, respectively, in the `ee_info` and `ee_data` fields,
- o the group address and source address can be retrieved, in all cases, by calling functions returning a `sockaddr` pointer and which

take into argument a sock_extended_err pointer (similarly as SOCK_EE_OFFENDER) and called SOCK_EE_MCAST_FEEDBACK_GRP and SOCK_EE_MCAST_FEEDBACK_SRC

If the Feedback contains multiple sources addresses, a sock_extended_err will be added to the message queue for each such sources.

An application receiving a sock_extended_err message from the MLD

component MUST NOT terminate the multicast subscription to the group or source/group address in error, except possibly if it can be ascertained that the Feedback message comes from a legitimate Querier (e.g. thanks to a mechanism like SEND [[RFC3971](#)]), and if multicast traffic for the said group or channel is not expected from any host attached to a directly-connected link.

(Another proposal would be to allow the setsockopt() and set(ipv4)sourcefilter() calls [[RFC3678](#)] to return an error. That would require the local network stack to wait for some time after sending a Membership Report message, before being able to return from the setsockopt()/set(ipv4)sourcefilter() call, and would not easily allow to carry detailed information about the specific group or channel in error. Consequently, this approach doesn't seem a viable one.)

[8.](#) Impact on IGMP/MLD proxies and equipments doing IGMP/MLD snooping

[8.1.](#) IGMP/MLD Proxies

To support this Feedback mechanism, an IGMP or MLD proxy [[RFC4605](#)] SHOULD send Feedback messages received on its Host side toward its Querier side(s) that have matching multicast memberships. The procedures for sending the Feedback messages are then the same as for a normal Querier, as specified in [Section 5](#): in particular the IGMP/MLD proxy MUST use its own address as the source address of the Feedback message.

A new member of a multicast group already forwarded by the proxy on its Querier side, will have to wait for some time before having a chance to receive a Feedback message : timers will have to trigger

before the Querier on the Host side of the proxy sends a Query, causing the proxy to send a Membership Report that may then cause the Querier on the Host side to send a Feedback message, and this Feedback message to be propagated to the new receiver.

To quickly provide Feedback messages to receivers on its Querier side, the proxy MAY cache the Feedback messages that it receives on the Host side to later match them with Membership Report messages received on its Querier side, and send relevant Feedback messages in reaction to these. If doing Feedback message caching, the proxy MUST keep only one Feedback message per (S,G) entry or (*,G) entry. It MUST also remove a Feedback message from its cache, if a same Feedback message hasn't been sent in the last <> seconds by the Querier on its Host side.

Last, an IGMP/MLD proxy MAY produce its own Feedback messages. In

that case it MUST still respect procedures of [Section 5](#).

[8.2](#). Equipments doing IGMP/MLD snooping

IGMP/MLD snooping equipments are expected to transparently interoperate with the procedures described in this document, given that [RFC 4541 section 2.2.1\(3\)](#) [[RFC4541](#)] states that "[a] switch that supports IGMP snooping must flood all unrecognized IGMP messages to all other ports".

[9](#). IGMP/MLD Hosts stacks not implementing the Feedback mechanism

To allow applications running on an IGMP/MLD Host, whose networking stack or API does not implement the Feedback mechanism described in this spec, it is proposed that IGMP/MLD Querier implementing this specification can, when configured to do so, send each Feedback message twice : once with the encoding described in these specifications, and another time encapsulated in a UDP packet.

The UDP message uses port xxx [TBD], with a payload identical to the IGMP or MLD Feedback message, except that the checksum is set to zero (the UDP message having its own checksum).

10. IANA Considerations

Request to IANA for IGMP and ICMPv6 type allocation will be needed for the messages defined in this document.

[Whether or not it is needed to define a registry for the error codes used in IGMP/MLD Feedback messages will be later determined.]

[Note to RFC Editor: this section may be removed on publication as an RFC.]

11. Security Considerations

Given that the specifications in this document do not change nor the state machine of the IGMP/MLD Querier and Host stack, nor the datagrams that will be received by an application, they are not expected to introduce security issues not already existing with IGMP/MLD or the protocol used to carry the Feedback message.

A possible issue would be to have wrong Feedback sent toward multicast applications. Such an issue could arise if spoofed Feedback messages are sent and interpreted by multicast receiver

hosts. This issue is mitigated by the fact that IGMP/MLD Hosts MUST check that the TTL of the Feedback messages is 1, and MAY also check the source IP of the Feedback message against a list of known Queriers.

Another possible issue is denial of service of the Querier function, that would be due to having the IGMP/MLD Querier be overloaded by Feedback messages to send. This is mitigated by allowing the Querier to rate-limit the flow of Feedback messages. On a LAN, such a rate-limiting would possibly result in some receivers not receiving the feedback message that they would have, which is a form of denial of service, but only on the Feedback function by itself, with no impact on the rest of the multicast group membership control protocol infrastructure. This later type of denial of service might be mitigated by doing rate-limiting based on the source address of the receivers (the source address of the Membership Report triggering the Feedback message); but such mechanism may themselves be subject to weaknesses due to Membership Report spoofing.

[12.](#) Acknowledgements

Acknowledgments go to DSLForum contributors who provided an initial proposal, to IETF participants that participated in the discussion on the magma WG list, from which guidance and inspiration was largely taken. Thank you to Bill Fenner for providing detailed information on issues related to ICMP errors in reaction to multicast datagrams.

Thanks to Toerless Eckert for his inputs and who offered a suggestion on how to handle application running on hosts not implementing the Feedback mechanism.

Message encodings are largely inspired from Report message encodings found in[RFC3376].

[13.](#) References

[13.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", [RFC 3376](#), October 2002.
- [RFC3678] Thaler, D., Fenner, B., and B. Quinn, "Socket Interface

Extensions for Multicast Source Filters", [RFC 3678](#), January 2004.

- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", [RFC 4541](#), May 2006.

- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", [RFC 4605](#), August 2006.

[13.2.](#) Informative References

- [I-D.ietf-mboned-maccnt-req]
He, H., "Requirements for Multicast AAA coordinated between Content Provider(s) and Network Service Provider(s)", [draft-ietf-mboned-maccnt-req-05](#) (work in progress), October 2007.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", [RFC 1812](#), June 1995.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.
- [fenner-icmp-mcast]
"ICMP errors in response to multicast", March 1999, <<http://www.icir.org/fenner/mcast/icmp.html>>.
- [magma-archive]
"IETF Magma WG mailing-list archives", December 2005, <<http://www1.ietf.org/mail-archive/web/magma/current/msg00815.html>>.
- [posix] "ISO/IEC 9945 Information technology, Portable Operating System Interface (POSIX), Part 1: Base Definitions", 2003.

[Appendix A.](#) Protocol to carry error feedback messages

ICMP and IGMP/MLD were possible candidates for carrying the Feedback message. This section exposes the pros/cons of both alternatives,

and ought to be removed once a decision is made on one of them.

[A.1.](#) ICMP

The Feedback message could be an ICMP message that would use a new ICMP message type (or possibly reusing existing types and codes).

Pros:

- o ICMP is already used to carry error messages between routers and hosts (e.g.. port unreachable message)
- o ICMP has an extensible format which could easily be reused for the Feedback function described in this memo
- o Implementations of network socket APIs already take into account ICMP messages

Cons:

- o ICMP has currently nothing to do with multicast today
- o some RFC explicitly forbid the sending of ICMP in reaction to receiving multicast packets, and IGMP/MLD Membership Reports are multicast packets ([\[RFC1122\] section 7.2](#) and 3.2.2, [\[RFC1812\] section 4.3.2.7](#)) (see [\[fenner-icmp-mcast\]](#))
- o ICMP messages are (currently) never sent toward multicast addresses, whereas that would be required to send Feedback messages to IGMPv2/MLDv1 hostsSo we may say that the generic argument is that the restriction for ICMP ; this has lead to practical issues to integrate this approach into existing stacks, because of the need to work around sanity checks

[A.2.](#) IGMP/MLD

The Feedback message could be an IGMP or MLD message that would use new IGMP/MLD message types.

Pros:

- o IGMP and MLD are the protocols used for all operations related to multicast subscription

Cons:

- o possibly more work to define the encodings
- o a new IANA registry might be needed to manage Feedback error codes
- o definition of how the network socket API will be used to carry the information to the applications has to be defined

Authors' Addresses

Thomas Morin (editor)
France Telecom - Orange Labs
2, avenue Pierre Marzin
Lannion 22307
France

Email: thomas.morin@orange-ftgroup.com

Brian Haberman
The Johns Hopkins University, Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel, MD 20723-6099
US

Phone: +1 443 778 1319
Email: brian@innovationslab.net

Internet-Draft

IGMP/MLD Error Feedback

July 2008

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND 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.

