Personal
Internet Draft
Technologies
Document: draft-oneill-mip-multicast-00.txt
Expires: Dec 2002

A. O'Neill
Flarion

5 July 2002

Mobility Management and IP Multicast
<draft-oneill-mip-multicast-00.txt>

Status of this Memo

Copyright Notice

Abstract

Mobile IP provides a mobile node, that visits a foreign subnet, the ability to continue to use an address from its home subnet (the home address) as a source address. This is achieved through the allocation of a Care of Address on the foreign subnet that is used as the end-point of a redirection tunnel from a home agent on the home subnet. Mobile IP in RFC 3220 states that when the mobile node originates multicast traffic intended for the foreign multicast system, it can only do so by first obtaining an IP address from the foreign subnet (a Collocated Care of Address) and then using this address as the multicast source address. This is to ensure that the source address will pass multicast routing reverse path forwarding checks.

This foreign multicast model is however extremely restrictive, and still very problematic to multicast routing and applications when the mobile node

regularly changes foreign subnets, as is common in wireless systems. This is because the source address continues to evolve which must be tracked by source specific multicast application and routing signalling. Using the home multicast system, again described above, is also non-optimal because the mobile node receiver is then serviced by packets that must be tunnelled from its home agent which, removes any multicast routing benefits (ie network based tree building). This draft therefore describes modifications to the foreign multicast interface between mobile IP and multicast routing that enable the mobile node to use its persistent home address as a multicast source address.

INDEX

Abstract

1. Introduction

Mobile IP provides a mobile node, which visits a foreign subnet, with the ability to continue to use an address from its home subnet (the home address) as a source address. This is achieved through the allocation of a Care of Address on the foreign subnet that is used as the end-point of a tunnel from a Home Agent on the home subnet. Mobile IP in RFC 3220 [MIPv4] and in [MIPv6] states that when the mobile node originates multicast traffic intended for the foreign multicast system, it can only do so by first obtaining an IP address from the foreign subnet (a Collocated Care of Address) and then using

this address as the multicast source address. This is to ensure that the source address will pass multicast routing reverse path forwarding checks as mentioned, for example, in the MIPv4 RFC text which is repeated overleaf.

From RFC 3220 section 4.4. Multicast Datagram Routing, page 66

A mobile node that wishes to send datagrams to a multicast group also has two options: (1) send directly on the visited network; or (2) send via a tunnel to its home agent. Because multicast routing in

general depends upon the IP source address, a mobile node which sends multicast datagrams directly on the visited network MUST use a co-located care-of address as the IP source address. Similarly, a mobile node which tunnels a multicast datagram to its home agent MUST use its home address as the IP source address of both the (inner) multicast datagram and the (outer) encapsulating datagram. This second option assumes that the home agent is a multicast router.

This foreign multicast model is however extremely restrictive, and still very problematic to multicast routing and applications when the mobile node regularly changes foreign subnets, as is common in wireless systems. This is because the source address continues to evolve which must be tracked by source specific multicast application and routing signalling. Using the home multicast system, again described above, is also non-optimal because the mobile node receiver is then serviced by packets that must be tunnelled from its home agent which, removes any multicast routing benefits (ie network based tree building). This draft therefore describes modifications to the foreign multicast interface between mobile IP and multicast routing that enable the mobile node to use its persistent home address as a multicast source address. It concentrates primarily on MIPv4, but mentions related MIPv6 issues and opportunities, which are brought together in section 8 along with a detailed description of the present MIPv6 foreign multicast scheme.

2. Conventions used in this document

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].

3. Terminology used in this document

Much of the terminology used in this document borrows from Mobile IPv4 [MIPv4], MIP Reverse Tunnelling [RevTun] and IP multicast RFCs and drafts. This draft introduces the following additional terminology.

Home multicast	Multicast via the home IGMP/MLD signalling.
Foreign multicast	Multicast via the IGMP/MLD signalling on the visited Subnet.
Hybrid Multicast	Foreign multicast reception, home multicast origination.
Designated Router	The DR is the multicast router/forwarder for a subnet.
OldDR / NewDR	Sender DRs as part of hand-off between subnets.
RPF Redirection	Redirecting the RPF check to point to the FA/DR and not the multicast source address.
Cross-over Router	The furthest router from the senders oldDR on a source

	tree that has the sender newDR on a different RPF interface.
Hand-Off Router	The router that issues an explicit Join towards the newDR and is the closest router from the old DR that has the newDR on the same RPF interface.
RPF Header	An IPv6 routing header indicating the preferred multicast RPF point for (S,G) packets.

4. Motivation

The motivation for this work is to enable a mobile node to have the option of using the more efficient foreign multicast delivery system. This requires the typical mobile node in wireless systems to use its home address as a foreign multicast source address, rather than a Collocated Care of Address, and yet still pass multicast routing RPF checks. This is to enable source specific multicast application and routing state to survive mobile node hand-offs between access routers that would typically not survive when using a Collocated Care of Address. Changes in these multicast source Collocated addresses would otherwise require multicast receiver application and routing signalling to be kept informed of each new source address change and to modify application and routing state in sympathy with such changes. These changes would unavoidably lead to lost packets and/or excessive signalling.
An associated motivation for this work is to avoid a mobile node, that wishes to source multicast traffic, from having to acquire a Collocated Care of Address from each foreign subnet, which is particularly expensive in MIPv4.

5. Limitations of MIP Multicast

5.1. Commercial Considerations

It is clear that MIP typically has a closely coupled policy layer that enables the home and foreign operators to control MN capabilities and packet routing when on the foreign subnet. In many cases the home operator wishes all packets to be routed to and from the home network for security, cost or customer control reasons. Similarly, the foreign operator also wishes to protect its own services and users from being affected by the presence of the roaming MN. In contrast, the foreign operator could alternatively require that the MN makes use of its own services whilst in the foreign domain and supporting this is probably a desire by the home operator to divert commodity traffic flows away from its home network and instead be delivered more efficiently by the foreign operator. These network and service control tensions are addressed by the policy layer. They need to be resolved by the AAA exchanges that occur during the request to connect to the foreign subnet.
This draft does not overly consider which of these commercial models is more important for MIP multicast and simply aims to make all practical options available to the parties involved. A discussion of the type of AAA support required for the specific suggestions in this draft are however outlined in section 7.

5.2. Foreign Multicast System in RFC3220

The foreign subnet has an IGMP Querier, a multicast designated router (DR) and at least one Foreign Agent (FA). The IGMP Querier issues Queries to the all-multicast-systems IP broadcast address, and the multicast DR and other receive GMRs from the MNs addressed to the multicast group address of interest. These IGMP messages are transmitted unencapsulated over the foreign subnet. The foreign multicast system uses native multicast routing from multicast senders in the Internet, down the multicast distribution tree towards those DRs that have joined to the group on behalf of their attached MNs. The DRs then transmit the multicast packets unencapsulated over the access link to the MN members that are members of the multicast group identified by the group destination address.

It can be seen that when multiple DRs in an access network are members of the same multicast group, and each DR has multiple MNs on that group, that the use of native multicast forwarding results in a single copy of each packet from the core of the network out across the access network and over the access link to the MN members. This represents the potential for significant bandwidth savings when compared to the home multicast system.

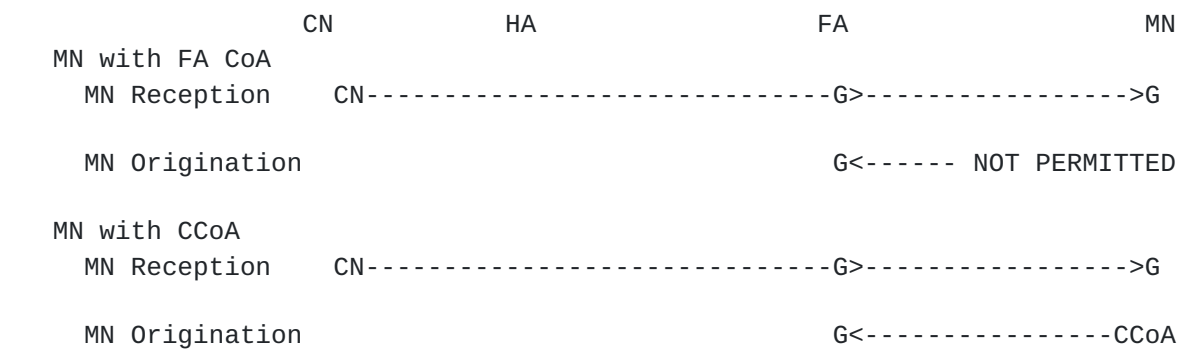

Figure 1. Forward and reverse foreign network multicast in RFC 3220

MN origination is not permitted when the MN has a FA CoA and hence such MNs can only receive multicast content. This prevents MIPv4 MNs, which typically cannot afford to be given a unique CCoA at each FA, nor afford the delay of continually updating this CCoA on hand-off, from taking part in bi-directional multicast flows that are typical with RTP sessions, and common in other multicast data applications.

MN origination is permitted when the MN has a CCoA, with that CCoA used as the source address. Once again multicast packets are sent unencapsulated over the access link, this time from the MN to the multicast DR on the subnet. This router forwards the packets into the multicast tree for the group contained in the destination address field of the packet. It can be seen here that whilst the multicast forwarding is bandwidth efficient, through the use of native multicast, it is limited to MNs with CCoAs.

5.3. Home Multicast System in RFC3220

The home multicast system in RFC 3220 uses a bi-directional tunnel between the HA and the MN CoA. The MN can have either a FA CoA or a CCoA from the FA and the resulting forwarding and encapsulations are shown graphically in figure 1. The MN should set the 'B' bit in the MIP RREQ to request the HA to forward to the MN, amongst other broadcast traffic, IGMP Queries and possibly IGMP Group Membership Reports to the MN. The MN can then issue solicited or unsolicited GMRs for the groups and group senders of interest to that MN,

and
the HA can then keep MN specific IGMP state to enable it to make appropriate forwarding decisions for multicast traffic arriving to the home subnet. Note that the HA must also export the MN GMRs to the home subnet, so that it can be seen by the IGMP Querier, received by the multicast DR on the home subnet for injection into the multicast tree building protocol, and also be seen by other MNs on the subnet to suppress their own GMRs. The IGMP Queries and GMRs
must be sent encapsulated over the foreign subnet to avoid them being confused with foreign subnet IGMP signalling, with the encapsulation being the same as that used for multicast content.

Figure 2. Forward and reverse home network multicast in RFC 3220

The major limitation here is that MN reception of multicast content is via a unicast tunnel from its HA. This tunnel is required to hide the multicast content from the foreign multicast system and to identify the target MN (the HoA/CCoA is otherwise missing due to the multicast packet having a group destination address). There is then potential for significant replication load being placed on the HA (and associated loss of bandwidth efficiency) when significant numbers of registered MNs at that HA are members of the same multicast group. In addition, when multiple MNs on the same foreign subnet are members of the same multicast group then multiple copies of the same content must be delivered to that foreign subnet and delivered over the air-interface in wireless systems. Only in the case that neither the HA nor the FA has multiple members of the same group (low membership coherence) is their no gain to be had from using multicast (network tree building and replication) between the HA and the FA. In all other cases, the absence of a multicast delivery tree potentially results in significant inefficiencies. When comparing the delivery costs (encapsulation processing and overhead) of multicast and unicast content from the HA in this model, it is evident that it is potentially better to use a multicast to unicast gateway on the home subnet and delivery any content using unicast, instead of incurring the additional cost and complexity of the unicast encapsulation and associated multicast signalling.

MN origination of content is via a unicast tunnel from the MN to the HA, using the HoA as a multicast source address. The unicast tunnel is less of an issue here because source specific branches from senders are common in multicast tree building and therefore the unicast tunnel does not result in

significant multicast inefficiencies. The tunnel is required simply to hide the multicast content from the foreign multicast system.

5.4. Non-Member Multicast Senders

It is permitted in the multicast architecture for a host to send traffic towards a multicast group of which it is not a member. When a host sends an IGMP GMR for group G, it is specifically asking to be a receiver of the group but may also wish to send to that group. A non-member sender is not a receiver on the group and is likely only a transient sender. This is

typically used to support early media flows in parallel with IGMP processing,
for transient senders that do not wish to disturb the multicast routing fabric, and for supporting sensor devices that are clearly not interested in the traffic from other senders. A non-member sender simply originates packets
to the group G and the multicast designated router forwards these into the multicast receiver tree, without initiating any receiver tree building activity in the multicast routing protocol. Non-member sender traffic is still however exposed to RPF checks. Non-member senders can be supported in either home or foreign multicast systems and therefore IGMP (or MLD) signalling may not occur before MN originated traffic flows.

5.5. Reverse Tunnelling Enhancements from RFC 3024

Reverse Tunnelling was developed to provide topologically correct tunnels back to the HA. When the MN has a FA CoA and wishes to tunnel traffic, including MN originated multicast, back to the HA then in RFC3220 as shown in
figure 2, this is achieved by a MN to HA tunnel using the HoA as a source address. Unfortunately, the HoA is not a topologically correct source address
and hence risks being dropped in the Internet in routers deploying source address checking. RFC 3024 instead adds the ability for the tunnel to instead
be initiated from the FA towards the HA and hence being topologically correct. Reverse tunnelling is requested by setting the 'T' bit in the MIP RREQ from the MN to the FA and onto the HA. There are then two modes for forwarding between and the MN to FA. In the default Direct Delivery Style (DDS), the MN sends the packets unencapsulated to the FA which then tunnels all received packets to the HA in the reverse tunnel. Essentially, all MN originated packets are viewed as being home network packets and foreign multicast is not permitted. Unfortunately, this method does not work however for multicast packets on a broadcast foreign subnet (not point to point) because these home broadcast packets will be confused with foreign broadcast packets by other MNs on the subnet and can be incorrectly received. RFC 3024 therefore mandates the second form of reverse tunnel forwarding known as Encapsulating Delivery Style (EDS). In EDS, the MN can selectively reverse tunnel packets to the HA through the use of an encapsulation between the MN and the FA. EDS mode is selected by the MN in the MIP RREQ by including an EDS extension as well as setting the 'T' bit.

The EDS extension is not forwarded to the HA and is viewed as purely a local matter between the MN and the FA. Packets which are encapsulated in a tunnel from the MN HoA to the FA are switched into a MIP tunnel from the FA CoA to the HA address. The HA then decapsulates them, and then forwards them onto the home subnet. The encapsulation on the foreign subnet also means that home
multicast packets are hidden from the foreign broadcast subnet and are

therefore not confusing to other MNs on the foreign subnet. Packets that are not to be reverse tunnelled are sent natively by the MN to the FA, which then forwards them normally, which in the case of foreign multicast is down the multicast tree. Essentially, EDS mode enables a MN to potentially partake in both home and foreign network multicast at the same time with the encapsulations over the foreign subnet being used to separate out IGMP and multicast content from/to the home and foreign multicast systems. Clearly, a MN would not wish to join the same multicast groups via both systems and so the MN, FA and HA need to have some configuration or AAA policy to decide which multicast systems the MN can participate in, and its limitations on that system (receiver, sender, receiver and sender, multicast group scope). Additionally, as has been described in 5.1, it is only a CCoA that enables a MN to originate traffic towards the foreign multicast system.

5.6 The Problem with CCoAs

Referring back to section 4.4 in RFC3220, the claim is made that MN origination into the foreign multicast system must use a CCoA. This is because the multicast source address must be topologically correct to pass the multicast reverse path forwarding check. This process, which is common to almost all multicast forwarding engines, is used to build source trees and to prevent routing loops. This is achieved by having the multicast forwarding engine in each multicast router, look-up the unicast source address within each multicast packet, as a unicast destination address in the unicast forwarding table. If the multicast packet arrives on an incoming interface, which is also the outgoing interface that would be used to forward unicast packets to that unicast destination address, then the RPF check has succeeded and the multicast packet may for replicated and forwarded. Putting aside transient routing effects, this RPF check will generally succeed for MN originated packets when the topologically correct CCoA is used as a source address. However, the RPF check will not generally succeed if the MN HoA is used as the source address because the HoA is topologically incorrect (belonging to the home subnet) and hence the multicast packets will not be received over the interface used to forward unicast packets towards the HoA. RFC 3220 is therefore correct in its statements and decisions in this regard.

However, RFC 3220 fails to mention that there is an additional problem with CCoAs. This is that the CCoA is a transient address and must change on each hand-off if the MN keeps moving. Each such address changes has a damaging affect on multicast applications and routing. For example, IGMPv3 enables a host to undertake source specific membership of a group, specifically enabling a MN to ignore content or receive content only from specific senders to that group. Protocol Independent Multicast-Sparse Mode (PIM-SM) also has source specific JOIN and PRUNE mechanisms that act in sympathy with sender specific group membership signalling to ensure only requested content is delivered down the multicast tree. In addition, PIM-SM supports both shared and source specific trees with the source specific PIM JOINs creating (S, G) state to over-rule the (*,G) shared tree state.

Source Specific Multicast (SSM) is also under development in the IETF in which the multicast destination group address as well as the senders unicast address identifies the multicast 'channel', and multicast routers keep (S+G) state. Finally, multicast transport and session layers applications typically use the multicast source address to 'demultiplex' content into sender specific feeds. This is because at a simplistic level, many to many network multicast is simply a superposition of multiple one to many transport flows.

In all these cases, a change in the multicast source address will create significant problems, requiring the address change to be communicated down the tree in advance of the CCoA update, so that new source specific routing state can be installed for (S2,G) or (S2+G) instead of (S1,G) or (S1+G). It must also be known to the host so that receiver applications can update transport, session and application state, to avoid application confusion and data corruption or loss. The scale of the update (all router and host sender specific state for that sender) coupled with the likely speed of hand-offs (and hence CCoA changes), makes the choice of the CCoA as a source address extremely problematic. Essentially, it completely prevents the MN from using the foreign multicast system and it must instead use the less efficient home multicast system. In solving the RPF problem, and preserving the packets of a single MN originator, it is clear that RFC 3220 creates an even bigger

problem, with wider implications on multicast routing stability. This is because the mobility is being directly exposed to the global multicast routing system through the address change, but is not being exposed to the unicast system (MIP instead deals with the latter). Note that the use of the HoA in the home multicast system coupled with the unicast tunnelling back to home subnet is one obvious way for multicast and MIP to collaborate in getting the job done. This however misses the efficiency of the foreign multicast delivery tree. The only way to correct this problem is to use the MN HoA as a multicast source address for the foreign multicast system (aligning somewhat home and foreign multicast processing on the hosts) and then find scalable means for MIP and the foreign multicast routing to work together to preserve the senders packets through the RPF check process. A range of techniques are next described in section 6, with the different techniques potentially forming an evolution and interoperability capability, as MIP and multicast technologies and standards evolve. They are described in overview to stimulate discussion between mobility and multicast researchers, so that standards activity can be commenced to address this opportunity.

6. HoA based MIP Multicast

The aim, in summary, is to enable a MN to originate IP multicast traffic using the HoA as a source address and have those packets correctly delivered by the foreign multicast system by specifically bypassing or satisfying the multicast RPF checks. This needs to work when the MN has requested EDS reverse tunnelling ('T' bit set plus EDS extension) or when no reverse tunnelling has been requested ('T' bit unset'). With DDS reverse tunnelling, it is clear that foreign multicast is by definition prevented and is not discussed further. An additional requirement is that the MN should not need to be aware of how the local FA is addressing this problem so that the MN can simply be made aware that foreign multicast origination is possible and then undertake home and foreign multicast as befits its configuration, incoming signalling, and the policy exchanged between the HA and FA. This idealised foreign multicast system is shown in figure 3 where the MN believes the FA (specifically the multicast designated router on the foreign subnet if different from the FA) is able to inject the multicast packets into the foreign multicast system and the multicast system will safely deliver them through the Internet to the multitude of CN multicast receivers on that group. This distribution should at all times be limited by the appropriate scope of the multicast group. Note that in the idealised system, the MN processing is the same for both a FA CoA and a CCoA.

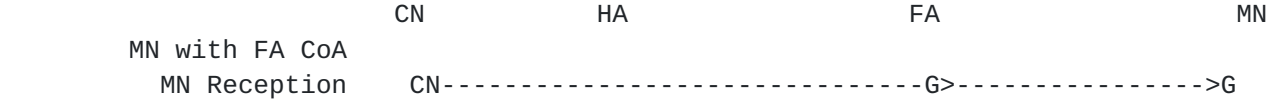

Figure 3. Idealised Foreign Multicast System

Before discussing the alternative solutions to this problem, it is important to point out that these solutions are covered at a relatively high-level and significant work in standards and subsequent engineering may be required to turn these suggestions into commercial reality. For now, they should simply be considered as examples to justify the potentially reality of MIP foreign multicast, using the HoA as a source address.

6.1 Hybrid Multicast System

The first (early) solution to this problem is to combine the best of the home and foreign multicast systems in satisfying the problem, thereby creating a hybrid multicast system. For simplicity, we will assume that the FA is also the multicast designated router for the foreign subnet. We can also assume that unencapsulated IGMP and multicast packets with a HoA source address are intended for the foreign multicast system, whether or not the 'B' bit is set or EDS RevTun has been requested and the 'T' bit is set. Essentially, we care greatly about being able to receive multicast via the foreign multicast system to accrue the bandwidth efficiencies, but care less about the path of the sender specific MN originated traffic.

The MN may or may not be a member of the multicast group G, whose scope must encompass both the home and foreign subnets (global scope only). If the MN is a member of group G then it will have sent an IGMP GMR for group G to the FA/DR and the FA/DR will have tracked IGMP state and initiated multicast tree building to add the FA/DR onto the receiver trees for the groups of interest to its MNs. This MN, and other MNs on that foreign subnet, will then receive multicast from the FA/DR in an unencapsulated form, and via a bandwidth efficient foreign multicast tree. We will now discuss how this is complemented with MN originated traffic.

6.1.1. MNs with FA CoA

The MN originates traffic to the group G by sending unencapsulated packets onto the foreign subnet with its HoA as a source address and a destination address of G. On a broadcast subnet, other members of group G on that subnet will also receive the packet. The FA also receives these packets but instead of injecting them into the foreign multicast system, it instead reverse tunnels them to the HA that matches the senders HoA in the visitors list, with the non-local HoA address acting as the trigger for this redirection. The HA knows the FA CoA from the registration and should therefore be happy to receive encapsulated packets from the registered FA CoA to the HA. Specifically, it needs to be happy to receive multicast packets via the FA CoA when the 'T' bit is not set. It will of course already be happy if the 'T' bit is set from RFC3024.

The HA has no IGMP state for such packets and treats them as non-member sender packets by injecting them into the home multicast system without initiating any receiver tree building. These MN originated packets are topologically correct at the HA and hence will satisfy any subsequent RPF checks except under transient routing situations. Note that the FA must first undertake its own multicast RPF check on the multicast packet using MIP visitor list state instead of the unicast routing state, before forwarding the packet to the HA. In addition, the FA needs to deliver the MN originated packets to other receivers of that group on that FA if the MNs are using point to point links.

This is because the multicast routing in the FA must itself not forward the packets again onto the foreign subnet when received down the multicast distribution tree, if they contain source addresses matching HoAs in the visitor list. This is necessary to avoid duplicate delivery on broadcast foreign subnets and is commonly achieved by the DR installing a source specific routing entry to discard packets arriving via a unidirectional shared tree that were originated from that DR.

6.1.2. MNs with CCoAs

When a MN has a CCoA it may or not have registered via the DR, which has an FA in MIPv4 or an attendant in MIPv6. If the MN has not registered via the DR then the DR cannot support hybrid multicast by forwarding the MN originated multicast packets to the HA because it has no state to do so. More specifically, the DR in this case cannot support MN originated foreign multicast traffic at all because any packets with a HoA source address, including IGMP GMRs, are topologically incorrect and hence will be discarded during ingress filtering in the absence of MIP visitor list state.

If the MN has registered via the DR then the DR will know the MN/CCoA/HA binding but in addition needs to know the MN/HoA binding to enable it to pass MN originated packets during ingress filtering and to then be able to tunnel them to the HA from the DR address. There are clearly ways that the MN could provide this information to the DR and the HA via a MIP extension, so that the tunnelling is both possible and acceptable, but this clearly requires the MN to be aware of the need to add such an extension. Thankfully, this extension is also required to enable the FA to natively forward unicast packets from the HoA and hence does not imply that the MN needs to know about the foreign multicast mechanism. The FA/DR receives the unencapsulated MN originated multicast packets and forwards them to the HA using the FA/DR address as a source address. The HA then receives packets from an address that is not equal to the registered CoA, but is equal to the source address of the received registration via that FA/DR and hence should be accepted because the HA and FA share an SA. The use of the HoA as a source address for foreign multicast can therefore only be permitted if the MN has registered via the FA/attendent and has informed that node of the MN HoA for ingress filtering purposes.

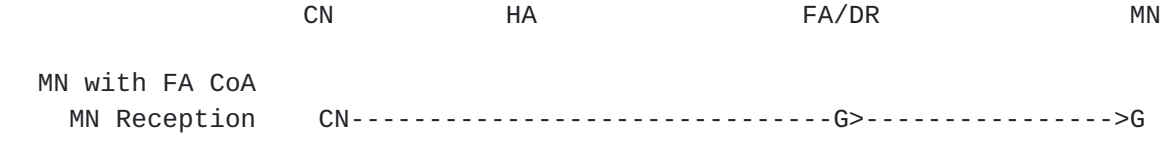

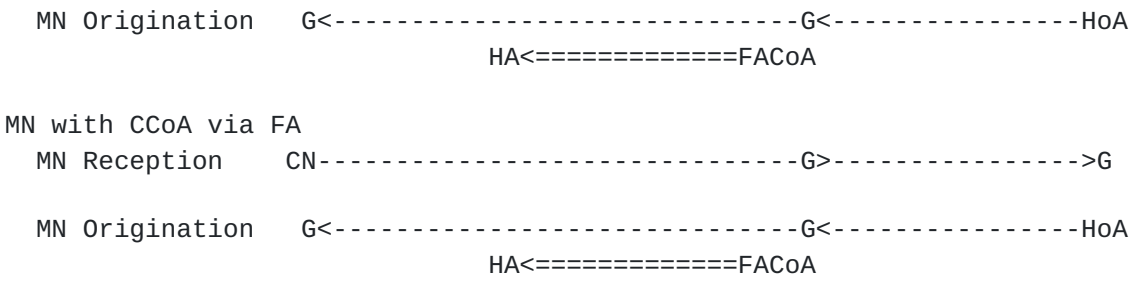

Figure 3. Hybrid Multicast System

Note that the use of the HoA as a multicast source address implies very different processing on the MN than the existing use of the CCoA. This is clearly a significant concern because MIPv6 relies solely on a MN CCoA, and its use for foreign multicast is potentially broken as discussed in section 8. This effectively means that MNs should be prevented from initiating foreign multicast content from the CCoA except in the specific case that the MN is sure that the CCoA will not be changing during the lifetime of the multicast session. This clearly excludes cellular mobility environments.

6.2 Shared Tree Solution

Some multicast routing protocols, such as PIM-SM, use a shared tree from a root node out to all receivers. The RPF check in this shared receiver tree is not made towards the senders unicast address in the multicast packet, but is instead made towards the root node whose address is distributed to all the multicast routers. Therefore the RPF problem with a non-local HoA address only needs to be solved between the senders designated router and the root node for such shared trees. This can be achieved by using a unicast tunnel from the DR to the root node, and then have the root node forward the senders packets down the receiver tree.

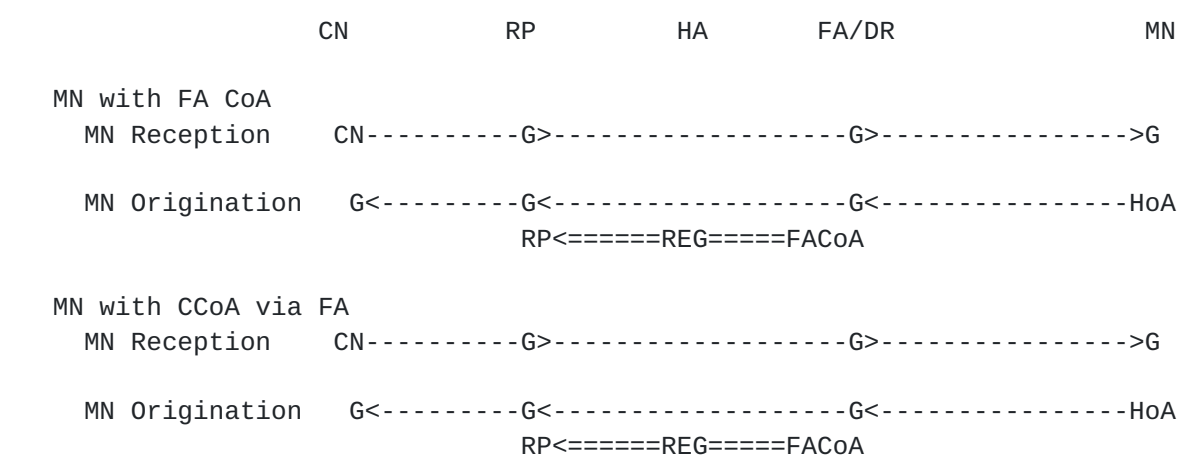

Figure 4. PIM-SM RP MIP Solution

In the specific case of PIM-SM, shown in figure 4, the root node is called the Rendevouz Point (RP) and PIM-SM already has mechanisms to enable the DR to tunnel packets directly to the RP using a Register message encapsulation. In the case of member senders, the RP is able to then try to PIM JOIN back to the sender via the senders DR and transmit periodic Register Stop messages to the senders DR. The PIM JOIN is intended to enable the senders DR to send packets natively to the RP via a source specific branch whilst the Register

Stop is intended to prevent the parallel encapsulation of multicast packets from the senders DR to the RP in Register messages. In addition to the source specific branch to the RP from the senders DR, any receiver DR is also allowed to try to build a source specific branch towards the senders DR and in so doing bypass the RP and the shared tree.

In the case of foreign multicast, both of these source specific branches will be directed towards the HoA in the source address of the multicast packet (and hence the HA subnet), and not towards the senders DR, which is the FA. These source specific PIM JOINs will therefore install state that will not be exercised by packets unless they cross part of the shared tree. The existence of this state is not a problem however as it will not cause problems and will eventually safely time out due to the soft state refresh model in PIM. The Register message could be extended to inform the RP not to bother attempting a PIM JOIN for this (S,G) and hence avoid wasted PIM JOIN and Register Stop signalling. In addition, the DR or RP could also periodically send a new PIM message down the multicast tree to the receiver DRs to also instruct them not to undertake source specific JOINs for this (S,G).

Comparing this model to the hybrid approach of 6.1, we can see a number of advantages that accrue when the multicast protocol uses a shared tree and supports unicast encapsulation to the root node. Firstly, in the hybrid approach it is clear that senders packets first go to the HA and then potentially onto the root node within a shared tree protocol. Therefore, sending the packets directly to the root node removes an additional encapsulated hop which is specifically useful when the HA and RP and far apart. In addition, we can see that by leaving all forwarding to the FA/DR we remove any uncertainties about the scope of group G, that otherwise limits the hybrid approach to global scope only. Finally, the most widely deployed multicast protocol in the Internet is PIM-SM and therefore the availability of a shared tree protocol with encapsulation to the root is generally assured. This solution does not however address the problem associated with CCoAs, nor does it deal with other types of multicast routing protocols with MOSPF a particular concern. MOSPF domains can of course still rely on the Hybrid solution of 6.2.

6.3 MIPv4 FA Multicast Encapsulation and MIPv6 RPF Redirect Option

The first two solutions rely on a unicast encapsulation to a point at which the HoA source address can pass subsequent RPF checks. An alternative solution is to encapsulate throughout the tree using an MIP multicast encapsulation. The FA encapsulates packets with a non-local HoA source address that have passed MIP aware ingress checking, using its own address as the source address and the inner destination group G as the outer destination address. This packet will then have a topologically correct source address and can be correctly forwarded by any multicast protocol that builds on-

demand source trees to the receivers of G via (anyFA,G) state. The receivers will then decapsulate such packets to reveal the original multicast packet with the HoA source address, which will then be checked against the source specific host membership state before being passed up to the transport layer.

Essentially the host is prepared to receive from any FA and therefore does not need to be kept informed of FA CoA changes. Any source specific tree building triggered by the receiver DR state should be suppressed when the data is received encapsulated like this. This approach again bypasses the HA with all the associated benefits, and this time is applicable to multicast protocols other than PIM-SM which would continue to use the Register encapsualtion. This however comes at the cost of host (and potentially DR) complexity, and the bandwidth inefficiency of the multicast encapsulation. In addition, this clearly does not work directly for explicit join protocols, the and in general limits any (S,G) pruning to the granularity of the FA address, such that multiple senders at the same FA cannot be selected/ pruned.

The cost of the encapsulation can be reduced in MIPv6 through the use of a new hop by hop option header. This 'RPF Redirect' option would be added by the MN and checked by all multicast routers, and includes the CCoA. The RPF redirect header therefore provides the opposite binding information to routers to that provided to the host by the Home Address Option (HAO) The RPF Redirect option would then be used in multicast routers to redirect RPF checks towards the senders DR instead of the HoA. The RPF Redirect Option or the source address of the FA encapsulation technique can be used by hosts to redirect source specific joins and prunes towards the newDR address, rather than towards the multicast source address of the original multicast packet. These redirections are clearly however affected by FA changes and this issue will be left for discussion in section 6.4.4.

6.4 Multicast Signalling Extensions - RPF Redirection

Section 6.3 handles the RPF problem in the data plane, but an alternative is to handle it in the multicast signalling plane. This is a longer-term solution in that changes to multicast signalling need to be widely deployed throughout a domain before they can be used, and because of the time to design, standardise and implement changes to deployed protocols. Essentially, the multicast RPF mechanism and the signalling in each routing protocol needs to be extended to support an arbitrary RPF point for the (S,G) in question. This is known henceforth in this document as RPF redirection and is analogous to the aims of the RPF Redirect option previously mentioned.

During hand-off, the CoA address is changing at the senders end and therefore each sender DR needs to advertise the new RPF point for that sender. This is achieved by injecting an RPF Redirect message into the multicast routing system using hop by hop multicast protocol signalling that is sent down the present tree or broadcast to multicast neighbours (protocol specific). In the latter case, each router passes the current RPF point to any subsequent joiners so that the sender DR only needs to undertake a periodic refresh of the RPF point in the absence of mobility. In the former case, the RPF point needs to be flooded rapidly by routers that detect that the old and new RPF points are on different interfaces so that the rapid flood is limited to those routers that are affected by the change. Note that the change in the senders DR might take the MN to a newDR that has no other members of group G, and also leave the oldDR with no other members of group G. Therefore some of the above RPF redirection signalling can be coupled with tree building activity (join/prune/membership flood) at the old and newDRs.

When the MN originator undertakes a hand-off, the oldDR and the newDR need to collaborate to update the RPF point. There is a cross-over router in the vicinity of both the old and new DRs that is the last router that would discard packets from the MN sent via the new DR, when using the oldDR as an RPF point. The newDR and all intermediate routers to the cross-over router need to be on the sender multicast tree for the group of interest, and must have the RPF point set to the newDR, in advance of multicast data being sent, to avoid that data being discarded. Once on the tree, then the newDR needs to send periodic RPF Redirects from the newDR to maintain the RPF point and the tree. At the same time, the receiver tree must also be updated. The precise mechanisms are of course multicast protocol specific due to the wide range of protocol mechanisms such as explicit join (PIM), member report broadcast (MOSPF), and data flood and prune (DVMRP) as examples.

6.4.1 PIM-SM Example (SSM also)

PIM uses explicit join/prune messaging to build either a shared tree from the RP, or source specific tree using the source address contained in multicast packets received on the shared tree. A source tree or Register encapsulation can be used between the DR and the RP in the case of a shared tree. During a hand-off on a shared tree with a source specific branch to the RP, the oldDR first needs to send an RPF Redirect down the multicast tree to first seek out the cross-over router. The RPF Redirect message contains the old DR address, the destination group address, the HoA source address, and the newDR address.

The cross-over router is identified by the next hop router towards the RP, termed the Hand-Off (HO) router, which detects that the RPF Redirect message does not affect the local RPF check because the oldDR and newDR interfaces are already the same. This router therefore issues an immediate JOIN towards the newDR to create (S,G,<newDR) state where the third <field indicates the RPF point and is a specific extension to the PIM Join message. This combines with the existing (S,G,<oldDR) state in the cross-over router to preserve packets from either DR during the hand-off, and is then passed up through the intermediate routers to the newDR. The newDR can then forward packets down the (S,G,<newDR) tree. Whilst waiting for this join, the multicast packets can be sent via the oldDR (ie via the old link or via the new link and a reverse tunnel to the oldDR) down the (S,G,<oldDR) tree, or can alternatively use a Register encapsulation to the RP. The HO router also passes the RPF Redirect down the tree to redirect oldDR state towards the newDR, setting a flag bit to indicate that the HO router has already been passed to suppress any immediate Join signalling in subsequent nodes. Note that this subsequent signalling between the HO and the RP can proceed relatively slowly compared to the signalling between the oldDR and the HO and between the HO and the newDR. This HO signal can also trigger the RP to commence issuing Register Stop messages to the newDR in the absence of the reception of Register encapsulated multicast data from the newDR. The cross-over router can issue a prune back towards the oldDR when it receives the join from the HO but should probably only do so once packets are received from the newDR branch over the new incoming interface. In the absence of such a prune normal PIM soft-state timers will still succeed in removing the oldDR branch in a less timely manner.

It is clear that, between the RP and the receiver DRs on the shared tree, the routers keep (*,G,<RP) state and hence the RPF Redirect message can be suppressed at this point. The Redirect message should be propagated however

if the operator wishes to suppress or redirect receiver DRs joining to a source tree. The redirect message informs the receiver DRs that the source specific join state should be (S,G,<newDR) and not (S,G,<S) and configuration in the receiver DRs or a flag in the Redirect message from the RP can control whether or not the source tree is permitted. Once again the source specific Join message needs to include the newDR as the RPF point, as well as the source address, and the message must be routed towards the newDR and not the source address.

If the source tree is permitted, then the source join will be directed towards a last DR that may not be the newDR due to a subsequent hand-off. The last DR should therefore send an immediate Redirect pointing to the newDR, which requires the last DR to keep hand-off state for a small number of seconds. This should ensure that the source join can rapidly catch-up with the movement of the MN.

6.4.2 MOSPF Example

MOSPF uses the flooding of membership reports within OSPF group-membership-LSA messages to build an on-demand shortest path source specific multicast tree. The shortest path calculation uses an RPF check from the sender network towards the networks that have members of the host group. The RPF point is determined by mapping the multicast source address to the source network in the link state database. MOSPF therefore needs to be extended to enable a source network to distribute an arbitrary mapping between source addresses and source networks so that the RPF check can be undertaken towards the source network that contains the FA and not the source network containing the HA. One approach would be to include any non-local sender addresses in the group-membership-LSA from the DR for the foreign subnet, which would cause MOSPF routers to learn this association that would over-rule the default mapping based on prefix matching. Clearly, this only works for member senders, and non-local multicast senders need to wait for the flood to complete before sending multicast packets into the domain to otherwise avoid having packets fail the default RPF check. Alternatively, the DR could use the FA/DR encapsulation of sect 6.3 whilst awaiting completion of the flood.

During hand-off, a controlled flood of the newRPF point is required within the region encompassing the old DR, newDR and cross-over routers to Redirect the RPF point for the non-local multicast source address. This can be achieved by the newDR advertising the MN HoA in an updated group-membership-LSA, which it floods into the domain. The group-membership-LSA will reach all routers in the domain but will likely reach the cross-over router very quickly. The oldDR continues to also distribute the MN HoA in its group-membership-LSA during the hand-off and only when this is complete does it discard IGMP membership state and update its group-membership-LSA to omit this MN HoA. During the hand-off, the MOSPF routers can use either the old or new DRs as RPF points. After the hand-off, the newDR will be the sole RPF point but in all routers beyond the local mobility region, a mixture of old and newDR RPFs will exist until the flood has completed. This is fine because the RPF interface will be the same for the old and newDRs outside of the local mobility region by definition. It may be necessary for a flag to be added to the group-membership-LSA to indicate whether or not it should be forwarded beyond the local mobility region. This is to enable the group-membership-LSAs to be generated sufficiently fast for good hand-off performance without impacting the signalling overhead on all the links in the domain. It is presently assumed however that aggregation through the presence of multiple non-local source addresses per group-membership-LSA, the extended

MIP hand-off period provided by make before break cellular technologies and the use of dual RPF points during that hand-off should be sufficient to avoid the excessive transfer of group-membership-LSAs over domain links.

6.4.3 DVMRP / PIM-DM

These protocols use a data flood followed by a receiver prune to efficiently support dense multicast membership. Both protocols use an RPF check towards the multicast source address to control the data flood but the lack of membership signalling in advance of data flow potentially renders the RPF redirection model inappropriate. DVMRP does however maintain its own multicast routing tables by propagating prefix routes and so there is potentially an opportunity for host specific routes for non-local senders to be included in this route distribution at the cost of routing stability. Therefore these protocols should instead use FA/DR encapsulation or the

RPF Redirect Option until they are extended with membership signalling phases, a development which is not considered to be likely.

6.4.4 RPF Redirect Option and FA Encapsulation

The use of these techniques from section 6.3 to bypass the RPF check also enable the host to learn the DR address on the foreign subnet as well as the non-local source address. These can then be used in source specific Joins as the destination address for tree building, instead of the non-local senders address that is already included in the PIM Join for example. Once the source specific Join is received back at the newDR then the RPF header or encapsulation can be dropped. These can then be re-applied on each hand-off whilst the routing protocol messaging in the domain adjusts the source specific tree and hence enables the impact of fast hand-off on the routing protocol to be minimise. A receiver and its receiver DR should however use the rate of such sender DR changes to decide whether it is appropriate to inject source specific Joins into the domain because clearly if the routing protocol has insufficient time to converge then the RPF header or encapsulation should be permanently used.

6.4.5 Incremental Deployment

Clearly, it is possible that each domain as a whole can be upgraded to support RPF redirection in the signalling plane as discussed in section 6.4. It is not however practical to expect that all such domains in the Internet will be upgraded at the same time and it is also not possible for a sender DR to know whether all domains on the path between the sender and all receivers have been so upgraded. Waiting until all domains upgrade before an Internet-wide flag day is not practical and creates a chicken and egg problem in that no-one has a motivation to upgrade until everyone else has. Therefore, this draft recommends that Multicast Border Routers (MBR) be aware of whether or not the next hop domain supports RPF redirection. If supported, then all is well and the multicast packet can be natively injected into that domain. However, if it is not supported then the MBR needs to apply one of the encapsulation techniques of 6.1 through 6.3, with 6.3 being mandated as the default mechanism for on-demand tree building protocols and RP Register encapsulation for PIM domains. This however needs to be assessed as part of the inter-domain multicast routing architecture (BGMP et al).

6.4.6 Topological Leaps

The limited propagation of the RPF Redirections rely on the fact that the MN is taking incremental steps across the edge of the topology such that the number of routers whose RPF checks need to be updated is always localised and hence limited. However, a MN could undertake a hand-off that represents a gigantic leap across the Internet topology (corporate to public network,

public cellular to home ADSL etc) and hence create a massive number of routers who are not on the tree and an additional number of routers whose RPF points are invalidated. There are a number of solutions to this, which should be triggered by the MN and the access router. Firstly, a MN clearly understands its movement through changes in the NAI and/or prefixes that are advertised to it. The MN can then expect that any multicast traffic is likely to be interrupted as the tree is rebuilt towards that new access router, and specifically that MN originated multicast might be overly affected. Secondly, the new access router can determine from the previous access router address (via PFANE or similar mechanism) that a topological leap has occurred. The

access router can then use, for MN originated traffic, either a unicast tunnel to the HA / root node, or alternatively a multicast tunnel to the group, until the source tree is built towards the senders DR. At this point, the current DR, can commence native forwarding and RPF redirection.

6.4.7. CCoA v HoA Summary

It is clear that one of the aims of using the HoA as a source address was to avoid the mobility dynamics, as the CCoA changes, from being exposed to the source specific state in the multicast routing protocol. RPF Redirection itself though also involves exposing the multicast routing to the mobility dynamics by propagating the evolving RPF point. This is however much less problematic than exposing the CCoA for the following reasons. This is because the RPF changes only need to be propagated in the local area where the movement results in changes to the required incoming interface. In the wider Internet the existing incoming interface (HoA,G,<anyremoteDR) is still valid, and hence forwarding will continue to succeed for the existing state. In contrast, a CCoA change must be propagated throughout the Internet routing system to update (CCoA1,G) to (CCoA2,G) which is clearly impractical.

This problem could be avoided in IPv6 by using universal interfaceIds (eg EUI-64) only, rather than the full 128 bit address (prefix + EUI) for the source specific multicast forwarding state, leading to the use of (EUI,G) state. A MN that is then moving between subnets will have an unchanged EUI but an evolving prefix, the latter though being masked from the multicast forwarding state that is now valid Internet wide. One can also imagine using prefix1 + universal interfaceID in the multicast state, and mask out sub-prefix1 bits from the Internet multicast system, where prefix1 represents all the networks through which the mobile can roam and still use this multicast address, and sub-prefix 1 masks away from the multicast protocol the bits that will change as a MN moves through the sub-networks under prefix1. This requires the multicast routing protocol to distribute and maintain a source mask as well as the source address in the routing state, but this is clearly overly complex. Finally, note that at the cost of a loss of resolution it is possible to use similar techniques in IPv4 by using a source network specific tree, rather than source specific tree, defined by a prefix2 that covers the networks under which the tree is valid and through which a sender to that group is allowed to roam. The least significant bits are masked out so that any CCoA allocated to the MN under that prefix2 will still be valid in the global multicast state. The cost here though is that the loss of resolution will be problematic to some multicast routing protocols in the local area close to the movement of the sender and probably makes this impractical without significant multicast protocol redesign. It is also invalid in the SSM model where all the source bits define the channel.

In summary, CCoAs can be masked from the multicast routing protocol by

limiting such protocols to (*,G) state and shared trees. Existing deployed protocols (DVMRP, PIM) however already build source trees using (S,G) state, and SSM specifically mandates such state. Other protocols build source trees on-demand using the RPF check and hence do not keep (S,G) state but this is likely to change due to the SSM requirements. The partial solution is to use a mask to limit tree building based on invariant source bits with IPv6 specifically supporting this via the universal interface ID. The only complete solution however, is to migrate to using the HoA for the source state, and to modify the multicast protocols to support an arbitrary RPF point and RPF redirection.

7. AAA Support for MIP Multicast

The Register encapsulation, FA/DR encapsulation and the RPF redirection techniques do not need additional assistance from the AAA layer other than controlling whether or not its MN can exploit foreign multicast and in what roles (sender/receiver etc). However, the hybrid multicast technique of section 6.1 may need the assistance of the AAA layer if it is necessary for the FA to know in advance from the HA if this capability is supported. This could be undertaken through MIP RREQ/RREP (BU) extensions between the HA to the FA but could alternatively be included in the state exchanged between the home and foreign AAA systems. The latter is particularly useful if it enables the home AAA to then dynamically assign a HA to the MN that is able to support the hybrid function.

8. IPv6 Considerations

The present proposal for foreign multicast in MIPv6 is to use the CCoA as the source address and use multicast BUs via the HA to populate the CNs with the binding between the senders CCoA and the HoA. The MN then sends multicast packets using the CCoA as the source address but including the Home Address Option (HAO). The CNs can then isolate the transport and higher layers from the CCoA changes as a MN moves by swapping the CCoA with the HAO. In addition, the CNs can undertake source specific joins for ASM or SSM by referring to the binding in the binding cache and sending these joins to the CCoA at the foreign subnet rather than to the HoA on the home subnet.

Whilst architecturally nice given its similarity with the unicast model of route optimisation, there is a critical difference. In the unicast model, the routing state in the Internet is already in place for reach CCoA and therefore the routing fabric is not affected by the senders mobility. Only the binding cache is exposed to the mobility. In the multicast case, the senders movement leads to the need to create a new source specific multicast tree for (newCCoA, G) for the existing population of receivers, in all routers between those receivers and the mobile sender. This not only causes synchronised tree building activity from that receiver population, but also has to be undertaken in some reasonable fraction of the hand-off interval to be at all useful. In cellular systems, this interval can be a small number of seconds and it is therefore clear that CCoA based multicast as mention in the MIPv6 design is also flawed.

The techniques in this draft are therefore still clearly necessary and

applicable, and in some cases enhanced by the IPv6 mechanisms, especially in regard to the RPF Redirect option. However, it is equally clear that the absence of MIPv6 FA CoAs or multicast protocol standards severely limits what can be achieved in regard to MIP multicast for a while. However, this equally provides a fresher starting point from which to developed MIP foreign multicast systems which include RPF redirection.

9. Security Considerations

Securing MIP and multicast router message exchanges are obvious minimal requirements that always need to be observed. More detailed analysis of the problem space, and of the solutions mention in this draft amongst other

candidate proposals, for supporting HoA originated foreign multicast, is necessary before concrete statements can be made about possible new threats and required security mechanisms. It is however clear that a security review will need to be undertaken specifically with the RPF Redirect option and RPF redirection in general. In addition, controls need to be placed on what nodes can tunnel multicast data to nodes such as the root node and the HA especially at multicast border routers. In addition, the multicast tunnelling to receivers by the FA/DR may be problematic.

10. Notice Regarding Intellectual Property Rights

Flarion's submissions will conform with RFC 2026. Flarion may seek patent protection on some or all of the technical information submitted by its employees in connection with the IETF's standards process. If part(s) of a submission by Flarion is (are) included in a standard and Flarion owns patent(s) and/or pending patent application(s) that are essential to implementation of such included part(s) in said standard, Flarion is prepared to grant a license on fair, reasonable, reciprocal (license back) and non-discriminatory terms on such included part(s).

11. References

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

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

[MIPv4] C.E. Perkins, ``IP Mobility Support for IPv4," RFC3220, January 2002.

[MIPv6] D. Johnson, C. Perkins, ``Mobility Support in IPv6", Internet-Draft, draft-ietf-mobileip-ipv6-18.txt (work in progress), 1 June 2002.

[IGMPv3] B. Cain et al, " Internet Group Management Protocol Version 3", Internet-draft, draft-ietf-idmr-igmp-v3-05.txt, November 2000.

[MLDv2] R. Vida, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", Internet-draft, draft-vida-mld-v2-03.txt, June 2002.

[PIM-SM] B. Fenner et al, "Protocol Independent Multicast - Sparse Mode", Internet-draft, draft-ietf-pim-sm-v2-new-05.txt, 1 March 2002

[SSM] H. Holbrook, B. Cain, "Source-Specific Multicast for IP", Internet-draft, draft-ietf-ssm-arch-00.txt, 21 November 2001.

[PIM-DM] A. Adams et al, "Protocol Independent Multicast - Dense Mode", Internet-draft, draft-ietf-pim-dm-new-v2-01.txt, February 15, 2002.

[DVMRP] T. Pusateri, "Distance Vector Multicast Routing Protocol", Internet-draft, draft-ietf-idmr-dvmrp-v3-10, August 2000.

[MOSPF] J. Moy, "Multicast Extensions to OSPF", RFC 1584, March 1994.

Author's Addresses

Alan O'Neill
Flarion Technologies
Phone: +1 908 947 7033
Email: oneill@flarion.com