

Network Working Group  
Internet Draft  
<draft-adamson-elasticmcast-00>  
Intended status: Experimental  
Expires: May 2014

B. Adamson  
Naval Research Laboratory  
C. Danilov  
Boeing Corporation  
J. Macker  
Naval Research Laboratory  
November 4, 2013

# Elastic Multicast Routing Protocol

## draft-adamson-elasticmcast-00

### Status of this Memo

CONFORMANCE UNDEFINED.

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 in May 2014.

### Abstract

This document describes an Internet Protocol (IP) multicast routing protocol suitable for dynamic network environments including mobile wireless. To handle high dynamics, this routing mechanism uses redundant forwarding, based upon the Simplified Multicast Forwarding (SMF) approach of [RFC6621], while converging to regular multicast distribution trees where or when the network becomes relatively stable. The rationale is that intermittent connectivity directly affects the ability of routers to synchronize on their view of the network, thus making it difficult to converge on efficient distribution trees, while network wide broadcast may be prohibitively expensive for relatively sparse groups. A hybrid approach, called Elastic Multicast, is specified which dynamically switches between limited scope broadcast and tree path forwarding independently at each node. The trees created during stable periods and portions of the network are pruned from the SMF efficient flooding mesh.

## Table of Contents

<b>1 Introduction</b> .....	<b>3</b>
<b>2 Terminology</b> .....	<b>4</b>
<b>3 Applicability Statement</b> .....	<b>5</b>
<b>4 Overview</b> .....	<b>6</b>
<b>5 Control Message Formats</b> .....	<b>7</b>
5.1 Elastic Multicast Acknowledgment (EM-ACK).....	7
5.2 Elastic Multicast Advertisement (EM-ADV).....	7
<b>6 Detailed Protocol Operation</b> .....	<b>9</b>
6.1 Elastic Multicast Timers.....	10
<b>7 SMF Relay Set Algorithm Considerations</b> .....	<b>11</b>
<b>8 IGMP and MLD Considerations</b> .....	<b>12</b>
<b>9 Border Gateway Considerations</b> .....	<b>13</b>
<b>10 Security Considerations</b> .....	<b>14</b>
<b>11 IANA Considerations</b> .....	<b>15</b>
<b>12 Acknowledgments</b> .....	<b>16</b>
<b>13 References</b> .....	<b>17</b>
13.1 Normative References.....	17
13.2 Informative References.....	17
<b>Authors' Addresses</b> .....	<b>19</b>
<b>Intellectual Property and Copyright Statements</b> .....	<b>20</b>

## 1. Introduction

IP Multicast is an efficient way to distribute data to multiple receivers in networks. And, in wireless networks, the often broadcast nature of the communication media makes this even more advantageous and often important, given capacity limits. In relatively stable networks, multicast forwarding trees can be formed by protocols such as Protocol Independent Multicast (PIM) Dense Mode (DM)[RFC3973] or PIM Sparse Mode (SM)[RFC4601] to route and replicate packets from a source to multiple destinations, thus reducing the overhead of multiple point-to-point forwarding paths run in parallel, one for each receiver of the multicast traffic. Building and maintaining multicast distribution trees requires some coordination between the network routers in order to avoid loops and to make sure that the dissemination paths are valid. In highly dynamic environments, the cost of maintaining the distribution trees can become prohibitive and the routers may be unable to synchronize their view of the network, leading to dropped traffic or transmission over invalid paths (*TBD-reference?*). Additionally, these existing IP multicast routing protocols do not provide for the case in wireless networks where a received packet must be retransmitted via the same interface as which it was received to reach adjacent hosts or routers that were not in reception range of the previous-hop transmitter. This is a fundamental difference from wired network connectivity.

A practical solution that proves very effective in relatively small networks with dense groups is Simplified Multicast Forwarding (SMF) [RFC6621] which can efficiently flood IP multicast traffic to the entire network, regardless of group membership. The SMF approach can provide efficient flooding when a subset of well-connected nodes that form a Connected Dominating Set (CDS) is selected to broadcast the data and relay it to the entire network. The SMF specification describes distributed relay set selection algorithms that can form CDS with local (2-hop neighborhood) information that can be collected via the Neighborhood Discovery Protocol (NHDP) of [RFC6130]. SMF utilizes Duplicate Packet Detection (DPD) forwarding that can be applied to the wireless interface case as needed instead of the typical reverse path forwarding checks used by other IP multicast routing protocols. The approach specified herein, called Elastic Multicast[CH2012], builds off the group forwarding mesh established by SMF with distributed relay set selection and dynamically prunes the mesh as any network stability permits. This hybrid approach allows for highly efficient, group-based multicast datagram distribution when and where the network is relatively stable and provides for SMF-based flooding in times and network portions where unstable connectivity and high dynamics occur. The pruning is implemented more as a "grafting" process in that acknowledgments from neighboring forwarders keep relevant portions of the SMF forwarding mesh active. In dynamic environments, a richer portion of the forwarding mesh is kept active as the acknowledgments are

Starting with the SMF mesh and driven by IP multicast datagram transmissions, the Elastic Multicast protocol forms multicast distribution trees whenever possible to minimize overhead of flooding to unnecessary parts of the topology and adaptively re-expands the forwarding base to additional redundant nodes when needed, when and where these trees become unstable. The data-driven aspect is based on active presence of IP multicast traffic flows among the routers where the flows are timed out due to lack of transmission activity. This mechanism of automatically expanding the forwarding base and then reducing it when not needed is based on acknowledgments in response to detected IP multicast traffic activity. The acknowledgments are aggregated and repeated at the intermediate forwarding nodes for active multicast traffic flows. The pruned SMF forwarding state is periodically timed out for active flows and flooding is reinstated. Thus, at a low duty cycle, active traffic flows are flooded (efficiently per the configured SMF algorithm) throughout the SMF domain to excite routers for Elastic Multicast acknowledgments. An OPTIONAL mechanism is described where periodic data traffic flooding may be economized by use of a control plane message that is periodically hop-by-hop disseminated throughout the network to advertise the aggregate set of active, but unacknowledged, IP multicast flows instead of the actual user traffic. The default flooding behavior of Elastic Multicast traffic is configurable on a per-flow (protocol, source, destination, traffic class) basis.

Elastic Multicast routing trades off pure efficiency in favor of robust datagram delivery in these types of network environments. However, it should be noted that in wireless network systems, the impact of Layer 2 Media Access Control (MAC) mechanisms and radio broadcast contention can extend multiple hops and, in limited scope wireless networks, SMF flooding with CDS relay sets is not necessarily much less efficient than group-specific forwarding in many cases, particularly in dynamic topologies. Since Elastic Multicast builds upon this SMF basis, it is expected to perform similarly and provide additional utility for cases of sparse group membership and stable connectivity.

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The terms introduced in [RFC5444], including "packet", "message", "TLV Block", "TLV", and "address" are to be interpreted as described therein.

The following abbreviations are used throughout this document:

<b>Abbreviation</b>	<b>Definition</b>
MANET	Mobile Ad hoc Network
SMF	Simplified Multicast Forwarding
CDS	Connected Dominating Set
NHDP	Neighborhood Discovery Protocol
DPD	Duplicate Packet Detection
FIB	Forwarding Information Base
TLV	type-length-value encoding

### 3. Applicability Statement

Within dynamic wireless routing topologies, maintaining traditional forwarding trees to support a multicast routing protocol is often not as effective as in wired networks due to the reduced reliability and increased dynamics of mesh topologies [MGL04] [GM99].

The SMF [RFC6621] specification provides for efficient forwarding of non-group-specific, IP multicast datagrams in wireless networks. This specification extends upon SMF to dynamically establish group-specific forwarding trees that are pared down from the base flooding mesh when and where network conditions of stability permit.

Elastic Multicast is compatible with different relay set selection algorithms and forwarding heuristics as described in the SMF specification [RFC6621] appendices. Elastic Multicast is also compatible with both the Any Source Multicast (ASM) [RFC1112]

- 

and Single Source Multicast (SSM) [RFC4607] models of multicast group membership.

Elastic Multicast deployments are able to connect and interoperate with existing standard multicast protocols operating within more conventional Internet infrastructures. To this end, a multicast border router or proxy mechanism **MUST** be used when deployed alongside more fixed-infrastructure IP multicast routing such as the PIM variants [RFC3973] and [RFC4601].

This document does not presently support forwarding of packets with directed broadcast addresses as a destination [RFC2644].

## 4. Overview

The Elastic Multicast protocol extends upon the SMF specification by providing for controlled forwarding of group-specific multicast traffic flows. Multicast flows are identified by the source address, group destination address, and optionally, protocol, port, and/or traffic class information. By default, newly detected flows are flooded per SMF to the entire network, but with a preset token bucket limited traffic shaping. Elastic multicast acknowledgment messages (EM-ACK) for flows of interest sent in response to the initially flooded traffic keep relevant portions of the SMF mesh "active" where token bucket limitations are lifted and other portions of the SMF mesh are "deactivated" providing only low rate forwarding until EM-ACK activity potentially reactivates it. As a more efficient alternative to gratuitous forwarding of the user data traffic, an Elastic Multicast advertisement message (EM-ADV) is specified that can allow for a "bundled" listing of active flow descriptions that is flooded within the routing area. The use of EM-ADV messages or low duty cycle flooding of the user traffic SHOULD be configurable on a per-flow basis. Initial forwarding/flooding of user traffic may be preferable depending upon the application and network use case.

Elastic Multicast monitors group membership of local hosts (collocated applications, directly attached hosts/devices, or neighboring non-routing hosts) to determine its transmission of EM-ACK messages and forwarding behaviors. When a local host is joined to a group, the Elastic Multicast router relieves any token bucket restrictions for the given group and provides periodic EM-ACK to neighboring routers upon receipt of applicable multicast traffic (or equivalent EM-ADV messages).

An Elastic Multicast router responds to IGMP joins and leaves issued by directly connected hosts. Based on such joins and leaves from connected hosts, an Elastic Multicast router may subscribe to different multicast groups. A router will be a subscriber for a multicast group as long as it has at least one directly connected host that joins that group.

The Elastic Multicast protocol may be deployed in a stand-alone MANET, or can be interfaced with PIM-DM enabled networks through Elastic Multicast gateways.

## 5. Control Message Formats

There is one principal and one optional control message used by the Elastic Multicast protocol. The first message is the Elastic Multicast Acknowledgment (EM-ACK) that is sent in response to received, non-duplicative, multicast user traffic from neighboring "upstream" routers. This message is directed to the given "upstream" router, informing it to make the flow "active" with respect to forwarding and similarly signal (via EM-ACK messages) its neighboring "upstream" routers that the flow is actively subscribed. The second, optional message is the Elastic Multicast Advertisement (EM-ADV) message that can be used as a surrogate instead of forwarding user traffic for unacknowledged flows. Multiple flow descriptions can be "bundled" into individual EM-ADV messages that can reduce the flooding overhead of advertising active multicast flows.

For both of these message types, standard MANET packet building block [RFC5444] formats will be defined. Note, as such, these messages may be sent independently or possibly opportunistically piggy-backed with other MANET link-local protocol messages such as NHDP [RFC6130].

### 5.1 Elastic Multicast Acknowledgment (EM-ACK)

The principal control message used in the Elastic Multicast protocol is acknowledgment message (EM-ACK) that is used to convey subscription interest in active multicast groups among routers in response to detected (transmitted) multicast traffic or the equivalent advertisement (EM-ADV) as described below. The EM-ACK message is directed to the specific neighboring router from which the local router received non-duplicative packets for groups of interest (i.e. joined by locally associate applications or hosts, or which are active due to directed EM-ACK messages received from other neighboring routers). Each EM-ACK message contains the following fields:

- router-id: A unique, addressable identifier for the intended "upstream" router. For example, in a broadcast wireless medium it can be the source Ethernet address of the original data multicast packet that is acknowledged.
- source-addr: The source IP address of the multicast packet being acknowledged
- group-addr: The destination IP multicast address of the packet being acknowledged

A standard MANET packet building block [RFC5444] message format for the EM-ACK message will be specified in a future revision of this document.

### 5.2 Elastic Multicast Advertisement (EM-ADV)

The Elastic Multicast Advertisement (EM-ADV) message provides an alternative means to convey the set of active, but unacknowledged, traffic flows of which the router is aware. Information regarding multiple flows can be bundled into a single EM-ADV message, thus reducing the overhead of advertising active traffic flows for which EM-ACK responses might be generated. Note that the use of EM-ADV message should be optional and configurable on a per-flow basis as, for some use cases, it may be preferable to flood the user traffic directly instead. The concept of using the EM-ADV message to proactively establish Elastic Multicast routing state prior to traffic transmission is being explored and may be described in a future revision of this draft.

Each EM-ADV message consists of a list of active multicast flow descriptions, each with the following fields:

- group-addr: IP multicast group destination address of active flow being advertised
- Source address: IP source address of flow
- Protocol: IP protocol id of flow
- Traffic class: IPv4 TOS / IPv6 traffic class of flow
- Token bucket parameters: bucket depth and rate
- Tagger identifier (Tagger-ID): address of router originating this flow description
- Duplicate packet detection identifier (DPD-ID): DPD-ID of packet triggering this flow description

The Tagger-ID and DPD-ID fields are set by the originator of the given flow description. The Tagger-ID is a unique router identifier of the Elastic Multicast router originating the EM-ADV flow description and the DPD-ID is the duplicate packet detection identifier for the packet that triggered the EM-ADV flow description generation. A standard MANET packet building block [RFC5444] message format for the EM-ADV message will be specified

in a future revision of this document. Additionally, provisions for "gateway address" and/or "wildcard" address descriptors may be provided to potentially facilitate border gateway and routing area group membership collection in future revisions.

## 6. Detailed Protocol Operation

Elastic Multicast works as an extension of SMF. Building on the distributed relay set selection and robust, efficient flooding, the goal of Elastic Multicast is to keep multicast forwarding active only for the subset of SMF relays needed to reach multicast group subscribers. The mechanism described here provides for this relay set pruning in stable parts of the network while still relying on the SMF flooding behavior in the highly dynamic parts of the network where it is hard or impossible to determine which nodes are on the critical path to the destination(s). By eliminating from the set of forwarding nodes those that do not move multicast traffic closer to the subscriber destinations, the number of retransmissions, and consequently the total overhead of the protocol will decrease.

By default, the SMF protocol does not maintain or use group membership information. In its simplest form, called *classic flooding* (CF), all SMF-enabled nodes rebroadcast all multicast packets, with the end result that all connected nodes get all multicast data. Duplicate Packet Detection (DPD), based on a unique packet identifiers and/or a hash of the packet content, is used to make sure that the same packet is not broadcasted more than once by a node, thus eliminating forwarding loops. In order to reduce overhead, SMF can use inputs from an external protocol, such as [OSPF-MDR](#) [RFC5614], [OLSR](#) [RFC3626], or [NHDP](#) [RFC6130] to enable forwarding by a selected Connected Dominating Set (CDS) of relays through which the entire MANET can be reached. This mechanism works well in relatively dense networks, significantly reducing the number of retransmissions. However, it still propagates multicast data throughout the entire MANET to all nodes, regardless of whether they subscribe to the specific multicast groups or not. Elastic Multicast works to leverage the distributed nature and robustness of SMF, but allow for improved efficiency by attempting to minimize the set of nodes forwarding traffic where it can for specific multicast groups.

To achieve this, Elastic Multicast treats low and high data rate flows of multicast traffic differently. Low data rate flows are flooded throughout the MANET area per SMF as usual. High data rate flows are limited to low data rate forwarding unless Elastic Multicast acknowledgment (EM-ACK) messages are received to "activate" high data rate forwarding of the given flow. Consistent acknowledgment under stable topology conditions to a consistent subset of the SMF relays serves to prune the SMF mesh towards a more efficient distribution tree for the flow while more dynamic acknowledgements to a varied (with topology changes) set of forwarders under less stable topologies serves to activate a larger, more redundant portion of the SMF mesh. The threshold that differentiates low and high data rate flows can be configurable. But, for an example, let's consider that a low data rate threshold is set at 1 packet per second. Thus, flows of less than 1 packet per second are disseminated throughout the entire MANET, while the protocol attempts to prune the distribution only to subscriber nodes for flows of more than 1 packet per second.

Elastic Multicast routers only maintain a local, temporary subscription for the multicast groups for which they are supposed to rebroadcast data packets, and do not maintain any information about the global network topology or global node membership, thus sharing the scalability benefits and most of the simplicity of SMF. However, routers that use Elastic Multicast need to be aware of the multicast groups to which locally attached devices subscribe. The standard protocols used in practice today for multicast join and leave operation are [IGMP](#) [RFC3376] and [MLD](#) [RFC3810]. The Elastic Multicast protocol does not change the functionality of the IGMP or MLD joins and leaves; it simply uses them to maintain the local membership of the attached devices or collocated processes. The considerations for IGMP and MLD operation are described further in [IGMP \(Section 8\)](#).

Elastic Multicast routers maintain a traffic shaping token bucket for each multicast data flow they observe. A multicast data flow is defined as a <sourceAddress:groupAddress> pair, although other parameters such as protocol, source/destination ports, or traffic class value may also be considered. The token bucket limits the rate at which multicast data for a certain flow is broadcasted throughout the network, in effect acting as a low rate enforcer per flow for SMF network broadcast. The bucket depth allows for some initial limited, gratuitous flooding of traffic for new flows and for bursty, but low average rate, traffic sources. When a flow is active, because of received local group membership or forwarding acknowledgments from adjacent Elastic Multicast routers, the token bucket enforcement is relaxed until the group membership is dropped and/or the acknowledgment activity for the flow has timed out.

The Elastic Multicast group-specific behavior is initiated by IP multicast group membership (e.g. IGMP join messages) by locally attached subscribers. For example, suppose an Elastic Multicast router becomes subscribed

to a certain group as a result of one of its attached devices joining that group. Upon receiving a non-duplicative multicast data packet addressed to the subscribed group, the router sends an Elastic Multicast Acknowledgement (EM-ACK) to the upstream router from which it received the multicast packet. The EM-ACK contains the source IP and the destination group of the original multicast packet. We call an upstream router that receives such a EM-ACK for a flow a forwarding router for that flow. A forwarding router disables the token bucket for a preconfigured period of time, and for that time sends unicast EM-ACKs upstream when it receives multicast packets for that flow. The difference between a forwarding router for a flow and a subscribing router for a multicast group due to some of its locally attached devices joining that group is that the forwarding router only sends EM-ACKs for the flows for which it is receiving EM-ACKs, while the end subscribing router sends EM-ACKs for all flows addressed to that group, regardless of source of the flows. Note, in the case of Single Source Multicast (SSM), the group identifier includes the source address and thus MAY coincide with the flow identifier.

In a static network, trees of temporary multicast forwarders for higher rate traffic is formed from the sources to the subscribers of the groups through the mechanism described above. Forwarding routers send EM-ACKs either periodically or after a certain number of packets for that flow was received, whichever occurs first, in order to refresh the forwarding state at upstream routers. If a forwarding router does not receive EM-ACKs for a multicast flow during a certain timeout, it simply re-enables the rate-limiting token bucket and no longer participates in the high rate forwarding or sending EM-ACKs upstream for that flow. It continues, however, to forward data at the lower rate, as limited by the token bucket. Other than EM-ACKs, no other control messages need to be used in the Elastic Multicast mechanism. However, the optional EM-ADV message MAY be used as a surrogate for low rate forwarding and a single EM-ADV can represent multiple traffic flows with a single message. Elastic Multicast routers use EM-ADV messaging should have its use configurable on a per-flow basis.

When the network topology becomes dynamic, or in parts of the network that become dynamic, more routers will become high load forwarders, as they receive EM-ACKs from downstream routers. Note that downstream routers send unicast EM-ACKs to the upstream router from which they received multicast data. If a forwarder router moves out of range, the next closest router will still forward data at a lower rate and will become activated when an EM-ACK is received. This mechanism expands the forwarding base and add redundancy in the dynamic parts of the network where more routers will be forwarding at a high rate, duplicating traffic, while in the stable parts of the network the protocol will form single paths or trees.

The Elastic Multicast protocol is similar to PIM Dense Mode (PIM-DM), in that it does not maintain topology information, it functions completely decentralized, and defaults in sending multicast data to all nodes in the network at a low rate. However, we believe that Elastic Multicast is more suitable to dynamic networks than PIM-DM mainly because Elastic Multicast routers do not depend on a unicast routing protocol for checking reverse paths in order to avoid loops, nor maintain point-to-point relationships with their neighbors. This allows them to broadcast multicast data, instead of streaming it point-to-point to downstream routers, and because of this the end to end packet delivery is less likely to be affected when network connectivity changes.

## 6.1 Elastic Multicast Timers

*TBD - describe the timers that control EM-ACK transmission and per-flow high rate forwarding activation / deactivation.*

## 7. SMF Relay Set Algorithm Considerations

SMF has the ability to select only some of the nodes to participate in the forwarding mechanism, such that in relatively dense areas, if only a subset of the nodes act as relays, the data forwarded still reaches all nodes in the MANET. Such a subset of forwarder nodes is called a Connected Dominant Set (CDS). Several CDS selection algorithms are currently implemented, and in particular SMF has been proven to perform well with the source-based multipoint relay (S-MPR) and essential connected dominating set (ECDS) algorithms described in [RFC6621].

For CF and ECDS operation where the SMF forwarding rules are simple, the Elastic Multicast operation is straightforward where EM-ACK messages are directed to the previous hop forwarder whenever a non-duplicative multicast packet is received for a flow and its activation is due to be refreshed due to timeout. Other SMF relay set forwarding rules may require special consideration. For example, the forwarding rules of the S-MPR algorithm necessitate slightly different handling. In the case of S-MPR, the EM-ACK message should be sent to the upstream router only when the received packet is non-duplicative *and* received from an MPR selector for the local router. It is expected that different relay set selection algorithms and selection criteria (e.g. metrics) will have impact on the utility of Elastic Multicast traffic for different application purposes. It is also expected that augmented relay set selection algorithms for efficient multi-interface operation will also impact the requirements for EM-ACK generation for correct protocol operation, but it is anticipated that Elastic Multicast can be adapted to these changes. Further experimentation and analysis should be conducted to determine the tradeoffs for network deployment and operations.

## 8. IGMP and MLD Considerations

The Elastic Multicast nodes act as regular multicast routers for local hosts subscribing to multicast groups. A local host can be:

1. an application residing on the same network node as the router,
2. a device attached directly to one of the non-MANET router interfaces, or
3. a separate neighboring host in the wireless network that does not have multicast routing capability but rather relies on the Elastic Multicast routers to forward their data appropriately.

In the case #1, the Elastic Multicast router may directly receive group membership join and leave notifications via operating system or implementation-specific mechanisms. In the case #2, the Elastic Multicast router should perform IGMP or MLD query functions per are [IGMP](#) [RFC3376] and [MLD](#) [RFC3810], respectively. For the case #3, standard IGMP or MLD operation is not sufficient for operation on MANET interfaces since adjacencies on wireless interfaces are incongruent for different routers due to radio propagation. An alternative form of IGMP/MLD query and response might be implemented where all routers send queries without observing the Querier Election provisions and MANET hosts would select (e.g., based on querier address or link quality) the router to which to direct their responses (i.e., address the reply to a specific router). This issue needs further study and may be addressed in a later version of this draft or other documents.

In any case, once an Elastic Multicast router has determined local host subscription(s), it will provide EM-ACK messages for subscribed groups and perform multicast forwarding per this specification.

## 9. Border Gateway Considerations

Elastic Multicast can be compatible with other multicast routing protocols (e.g. PIM) if border gateway provisions are met. Given the similarities to PIM-DM, this allows a relatively straight forward interconnection with PIM-DM domains. Elastic Multicast gateways can forward multicast traffic received from their PIM-DM neighbors towards the MANET and send PIM join and leave messages to the upstream PIM routers based on the subscription of other routers in the MANET. Similarly, on the egress side Elastic Multicast gateways can act as subscribers to groups for which they receive PIM join messages from their downstream PIM routers. In the case where the Elastic Multicast routing area is a stub network, the adjacent gateway routing protocol membership information can be used to manage Elastic Multicast behavior in the same manner as host group membership for traffic that egresses the Elastic Multicast area. Group membership subscription needs to be aggregated from within the Elastic Multicast area and provided to the gateway router so that appropriate multicast traffic from external networks will be routed to the Elastic Multicast area. An approach to do this is described in [CDHM07]. Additionally, an extended approach that allows for multiple gateways is describe in [DHS08]. Further consider of border gateway operation, including cases where the Elastic Multicast area is a transit domain, may be addressed in future revisions of this draft.

## 10. Security Considerations

*(TBD - We can refer to the SMF Security Considerations as applicable, the PacketBB-Sec document, etc. Some security for the EM-ACK messages should be provided via PacketBB-Sec)*

## 11. IANA Considerations

*(TBD - Call out any IANA assignments that need to be made. Most notably this will be the Elastic Multicast message types in the PacketBB message type namespace, a registry for any Elastic Multicast message TLVs, and those specific TLVs.)*

## 12. Acknowledgments

*(TBD)*

Some core contributors/authors in alphabetical order:

Tom Henderson, et al

The RFC text was produced using Marshall Rose's xml2rfc tool and Bill Fenner's XMLmind add-ons.

## 13. References

### 13.1 Normative References

- [RFC0791] Postel, J., "[Internet Protocol](#)", STD 5, RFC 791, September 1981.
- [RFC1112] Deering, S., "[Host extensions for IP multicasting](#)", STD 5, RFC 1112, August 1989.
- [RFC2119] Bradner, S., "[Key words for use in RFCs to Indicate Requirement Levels](#)", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S.E. and R.M. Hinden, "[Internet Protocol, Version 6 \(IPv6\) Specification](#)", RFC 2460, December 1998.
- [RFC2644] Senie, D., "[Changing the Default for Directed Broadcasts in Routers](#)", BCP 34, RFC 2644, August 1999.
- [RFC2780] Bradner, S. and V. Paxson, "[IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers](#)", BCP 37, RFC 2780, March 2000.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "[Internet Group Management Protocol, Version 3](#)", RFC 3376, October 2002.
- [RFC3626] Clausen, T. and P. Jacquet, "[Optimized Link State Routing Protocol \(OLSR\)](#)", RFC 3626, October 2003.
- [RFC3810] Vida, R. and L. Costa, "[Multicast Listener Discovery Version 2 \(MLDv2\) for IPv6](#)", RFC 3810, June 2004.
- [RFC4607] Holbrook, H. and B. Cain, "[Source-Specific Multicast for IP](#)", RFC 4607, August 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "[Guidelines for Writing an IANA Considerations Section in RFCs](#)", BCP 26, RFC 5226, May 2008.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "[Generalized Mobile Ad Hoc Network \(MANET\) Packet/Message Format](#)", RFC 5444, February 2009.
- [RFC5614] Ogier, R. and P. Spagnolo, "[Mobile Ad Hoc Network \(MANET\) Extension of OSPF Using Connected Dominating Set \(CDS\) Flooding](#)", RFC 5614, August 2009.
- [RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "[IANA Guidelines for IPv4 Multicast Address Assignments](#)", BCP 51, RFC 5771, March 2010.
- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "[Mobile Ad Hoc Network \(MANET\) Neighborhood Discovery Protocol \(NHDP\)](#)", RFC 6130, April 2011.
- [RFC6621] Macker, J., "[Simplified Multicast Forwarding](#)", RFC 6621, May 2012.

### 13.2 Informative References

- [CDHM07] Chakeres, I., Danilov, C., and T. Henderson, "Connecting MANET Multicast", IEEE MILCOM 2007 Proceedings, 2007.
- [CH2012] Danilov, C., Henderson, T., Brewer, O., Kim, J., Macker, J., and B. Adamson, "Elastic Multicast for Tactical Networks", IEEE MILCOM 2012, October 2012.
- [DHS08] Danilov, C., Henderson, T., and T. Spagnolo, "MANET Multicast with Multiple Gateways", IEEE MILCOM 2008 Proceedings, 2008.
- [GM99] Garcia-Luna-Aceves, JJ and E. Madruga, "The core-assisted mesh protocol", Selected Areas in Communications, IEEE Journal on Volume 17, Issue 8, August 1999.
- [JLMV02] Jacquet, P., Laouiti, V., Minet, P., and L. Viennot, "Performance of multipoint relaying in ad hoc mobile routing protocols", Networking, 2002.
- [MDC04] Macker, J., Dean, J., and W. Chao, "Simplified Multicast Forwarding in Mobile Ad hoc Networks", IEEE MILCOM 2004 Proceedings, 2004.

- [MDDA07] Macker, J., Downard, I., Dean, J., and R.B. Adamson, "Evaluation of distributed cover set algorithms in mobile ad hoc network for simplified multicast forwarding", ACM SIGMOBILE Mobile Computing and Communications Review Volume 11 , Issue 3, July 2007.
- [MGL04] Mohapatra, P., Gui, C., and J. Li, "Group Communications in Mobile Ad hoc Networks", IEEE Computer Vol. 37, No. 2, February 2004.
- [RFC2501] Corson, M.S. and J. Macker, "[Mobile Ad hoc Networking \(MANET\): Routing Protocol Performance Issues and Evaluation Considerations](#)", RFC 2501, January 1999.
- [RFC3973] Adams, A., Nicholas, J., and W. Siadak, "[Protocol Independent Multicast - Dense Mode \(PIM-DM\): Protocol Specification \(Revised\)](#)", RFC 3973, January 2005.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "[Protocol Independent Multicast - Sparse Mode \(PIM-SM\): Protocol Specification \(Revised\)](#)", RFC 4601, August 2006.

## Authors' Addresses

**Brian Adamson**

Naval Research Laboratory

Washington, DC 20375

USA

E-Mail: [brian.adamson@nrl.navy.mil](mailto:brian.adamson@nrl.navy.mil)

**Claudiu Danilov**

Boeing Corporation

Seattle, WA

USA

E-Mail: [Claudiu.B.Danilov@boeing.com](mailto:Claudiu.B.Danilov@boeing.com)

**Joseph Macker**

Naval Research Laboratory

Washington, DC 20375

USA

E-Mail: [joseph.macker@nrl.navy.mil](mailto:joseph.macker@nrl.navy.mil)

## Full Copyright Statement

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](mailto:ietf-ipr@ietf.org)<sup>1</sup>.

<sup>1</sup> <mailto:ietf-ipr@ietf.org>