                  Embedding the Rendezvous Point (RP) Address
                        in an IPv6 Multicast Address

Status of this Memo

Copyright Notice

Abstract

   This memo defines an address allocation policy in which the address
   of the Rendezvous Point (RP) is encoded in an IPv6 multicast group
   address.  For Protocol Independent Multicast - Sparse Mode (PIM-SM),
   this can be seen as a specification of a group-to-RP mapping
   mechanism.  This allows an easy deployment of scalable inter-domain
   multicast and simplifies the intra-domain multicast configuration as
   well.  This memo updates the addressing format presented in RFC 3306.

Table of Contents

1.  Introduction

1.1.  Background

   As has been noticed [V6MISSUES], there exists a deployment problem
   with global, interdomain IPv6 multicast: PIM-SM [PIM-SM] RPs have no
   way of communicating the information about (active) multicast sources
   to other multicast domains, as Multicast Source Discovery Protocol
   (MSDP) [MSDP] has deliberately not been specified for IPv6.
   Therefore the whole interdomain Any Source Multicast (ASM) model is
   rendered unusable; Source-Specific Multicast (SSM) [SSM] avoids these
   problems but is not a complete solution for several reasons, as noted
   below.

   Further, it has been noted that there are some problems with the
   support and deployment of mechanisms SSM would require [V6MISSUES]:
   it seems unlikely that SSM could be usable as the only interdomain
   multicast routing mechanism in the short term.

1.2.  Solution

   This memo describes a multicast address allocation policy in which
   the address of the RP is encoded in the IPv6 multicast group address,
   and specifies a PIM-SM group-to-RP mapping to use the encoding,
   leveraging, and extending unicast-prefix-based addressing [RFC3306].

   This mechanism not only provides a simple solution for IPv6
   interdomain Any Source Multicast but can be used as a simple solution
   for IPv6 intra-domain ASM with scoped multicast addresses as well.

It can also be used as an automatic RP discovery mechanism in those
deployment scenarios that would have previously used the Bootstrap
Router protocol (BSR) [BSR].

The solution consists of three elements:

o  A specification of a subrange of [RFC3306] IPv6 multicast group
   addresses defined by setting one previously unused bit of the
   Flags field to "1",

o  a specification of the mapping by which such a group address
   encodes the RP address that is to be used with this group, and

o  a description of operational procedures to operate ASM with PIM-SM
   on these IPv6 multicast groups.

Addresses in the subrange will be called embedded-RP addresses.

This scheme obviates the need for MSDP, and the routers are not
required to include any multicast configuration, except when they act
as an RP.

This memo updates the addressing format presented in RFC 3306.

Some design tradeoffs are discussed in Appendix A.

1.3.  Assumptions and Scope

A 128-bit RP address can't be embedded into a 128-bit group address
with space left to carry the group identity itself. An appropriate
form of encoding is thus defined by requiring that the Interface-IDs
of RPs in the embedded-RP range can be assigned to be a specific
value.

If these assumptions can't be followed, operational procedures and
configuration must be slightly changed, or this mechanism can't be
used.

The assignment of multicast addresses is outside the scope of this
document; it is up to the RP and applications to ensure that group
addresses are unique by using some unspecified method.  However, the
mechanisms are probably similar to those used with [RFC3306].

Similarly, RP failure management methods, such as Anycast-RP, are out
of scope for this document.  These do not work without additional
specification or deployment.  This is covered briefly in Section 6.1.

1.4.  Terminology

   Embedded-RP behaves as if all the members of the group were intra-
   domain to the information distribution. However, as it gives a
   solution for the global IPv6 multicast Internet, spanning multiple
   administrative domains, we say it is a solution for inter-domain
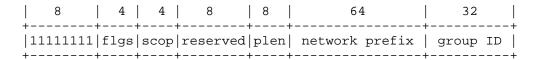   multicast.

   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 [RFC2119].

1.5.  Abbreviations

      ASM     Any Source Multicast
      BSR     Bootstrap Router
      DR      Designated Router
      IGP     Interior Gateway Protocol
      MLD     Multicast Listener Discovery
      MSDP    Multicast Source Discovery Protocol
      PIM     Protocol Independent Multicast
      PIM-SM  Protocol Independent Multicast - Sparse Mode
      RIID    RP Interface ID (as specified in this memo)
      RP      Rendezvous Point
      RPF     Reverse Path Forwarding
      SPT     Shortest Path Tree
      SSM     Source-Specific Multicast

2.  Unicast-Prefix-based Address Format

   As described in [RFC3306], the multicast address format is as
   follows:

      |   8    | 4  | 4  |   8    | 8  |       64       |    32    |
      +--------+----+----+--------+----+----------------+----------+
      |11111111|flgs|scop|reserved|plen| network prefix | group ID |
      +--------+----+----+--------+----+----------------+----------+

   Where flgs are "0011".  (The first two bits are as yet undefined,
   sent as zero and ignored on receipt.)

3.  Modified Unicast-Prefix-based Address Format

    This memo specifies a modification to the unicast-prefix-based
    address format by specifying the second high-order bit ("R-bit") as
    follows:

```
     |   8    | 4  | 4  | 4  | 4  | 8  |       64       |    32    |
     +--------+----+----+----+----+----+----------------+----------+
     |11111111|flgs|scop|rsvd|RIID|plen| network prefix | group ID |
     +--------+----+----+----+----+----+----------------+----------+
                                    +-+-+-+-+
     flgs is a set of four flags:   |0|R|P|T|
                                    +-+-+-+-+
```

    When the highest-order bit is 0, R = 1 indicates a multicast address
    that embeds the address on the RP.  Then P MUST be set to 1, and
    consequently T MUST be set to 1, as specified in [RFC3306].  In
    effect, this implies the prefix FF70::/12.  In this case, the last 4
    bits of the previously reserved field are interpreted as embedding
    the RP interface ID, as specified in this memo.

    The behavior is unspecified if P or T is not set to 1, as then the
    prefix would not be FF70::/12.  Likewise, the encoding and the
    protocol mode used when the two high-order bits in "flgs" are set to
    11 ("FFF0::/12") is intentionally unspecified until such time that
    the highest-order bit is defined.  Without further IETF
    specification, implementations SHOULD NOT treat the FFF0::/12 range
    as Embedded-RP.

    R = 0 indicates a multicast address that does not embed the address
    of the RP and follows the semantics defined in [ADDRARCH] and
    [RFC3306].  In this context, the value of "RIID" MUST be sent as zero
    and MUST be ignored on receipt.

4.  Embedding the Address of the RP in the Multicast Address

    The address of the RP can only be embedded in unicast-prefix-based
    ASM addresses.

    That is, to identify whether it is a multicast address as specified
    in this memo and to be processed any further, an address must satisfy
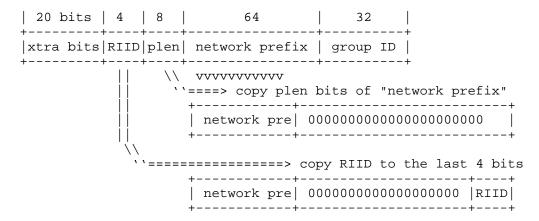    all of the following:

    o It MUST be a multicast address with "flgs" set to 0111, that is, to
      be of the prefix FF70::/12,

    o  "plen" MUST NOT be 0 (i.e., not SSM), and

   o  "plen" MUST NOT be greater than 64.

   The address of the RP can be obtained from a multicast address
   satisfying the above criteria by taking the following two steps:

   1. Copy the first "plen" bits of the "network prefix" to a zeroed
      128-bit address structure, and

   2. replace the last 4 bits with the contents of "RIID".

   These two steps could be illustrated as follows:

       | 20 bits | 4  | 8  |       64       |    32    |
       +---------+----+----+----------------+----------+
       |xtra bits|RIID|plen| network prefix | group ID |
       +---------+----+----+----------------+----------+
                 ||     \\  vvvvvvvvvvv
                 ||       ``====> copy plen bits of "network prefix"
                 ||           +-----------+-------------------------+
                 ||           | network pre| 00000000000000000000000   |
                 ||           +-----------+-------------------------+
                  \\
                   ``================> copy RIID to the last 4 bits
                             +-----------+---------------------+----+
                             | network pre| 0000000000000000000 |RIID|
                             +-----------+---------------------+----+

   One should note that there are several operational scenarios (see
   Example 3 below) when the [RFC3306] statement "all non-significant
   bits of the network prefix field SHOULD be zero" is ignored.  This is
   to allow multicast group address allocations to be consistent with
   unicast prefixes; the multicast addresses would still use the RP
   associated with the network prefix.

   "plen" higher than 64 MUST NOT be used, as that would overlap with
   the high-order bits of multicast group-id.

   When processing an encoding to get the RP address, the multicast
   routers MUST perform at least the same address validity checks to the
   calculated RP address as to one received via other means (like BSR
   [BSR] or MSDP for IPv4).  At least fe80::/10, ::/16, and ff00::/8
   MUST be excluded.  This is particularly important, as the information
   is obtained from an untrusted source, i.e., any Internet user's
   input.

   One should note that the 4 bits reserved for "RIID" set the upper
   bound for RPs for the combination of scope, network prefix, and group
   ID -- without varying any of these, one can have 2^4-1 = 15 different

RPs (as RIID=0 is reserved, see section 6.3).  However, each of these
is an IPv6 group address of its own (i.e., there can be only one RP
per multicast address).

5.  Examples

Four examples of multicast address allocation and resulting group-
to-RP mappings are described here to better illustrate the
possibilities provided by the encoding.

5.1.  Example 1

The network administrator of 2001:DB8::/32 wants to set up an RP for
the network and all the customers, by placing it on an existing
subnet, e.g., 2001:DB8:BEEF:FEED::/64.

In that case, the group addresses would be something like
"FF7x:y40:2001:DB8:BEEF:FEED::/96", and then their RP address would
be "2001:DB8:BEEF:FEED::y".  There are still 32 bits of multicast
group-ids to assign to customers and self ("y" could be anything from
1 to F, as 0 must not be used).

5.2.  Example 2

As in Example 1, the network administrator of 2001:DB8::/32 wants to
set up the RP but, to make it more flexible, wants to place it on a
specifically routed subnet and wants to keep larger address space for
group allocations.  That is, the administrator selects the least
specific part of the unicast prefix, with plen=32, and the group
addresses will be from the multicast prefix:

    FF7x:y20:2001:DB8::/64

where "x" is the multicast scope, "y" is the interface ID of the RP
address, and there are 64 bits for group-ids or assignments.  In this
case, the address of the RP would be:

    2001:DB8::y

The address 2001:DB8::y/128 is assigned to a router as a loopback
address and is injected into the routing system; if the network
administrator sets up only one or two RPs (and, e.g., not one RP per
subnet), this approach may be preferable to the one described in
Example 1.

5.3.  Example 3

   As in Example 2, the network administrator can also assign multicast
   prefixes such as "FF7x:y20:2001:DB8:DEAD::/80" to some of customers.
   In this case the RP address would still be "2001:DB8::y".  (Note that
   this is just a more specific subcase of Example 2, where the
   administrator assigns a multicast prefix, not just individual group-
   ids.)

   Note the second rule of deriving the RP address: the "plen" field in
   the multicast address, 0x20 = 32, refers to the length of "network
   prefix" field considered when obtaining the RP address.  In this
   case, only the first 32 bits of the network prefix field, "2001:DB8",
   are preserved: the value of "plen" takes no stance on actual
   unicast/multicast prefix lengths allocated or used in the networks,
   here from 2001:DB8:DEAD::/48.

   In short, this distinction allows more flexible RP address
   configuration in the scenarios where it is desirable to have the
   group addresses be consistent with the unicast prefix allocations.

5.4.  Example 4

   In the network of Examples 1, 2, and 3, the network admin sets up
   addresses for use by customers, but an organization wants to have its
   own PIM-SM domain.  The organization can pick multicast addresses
   such as "FF7x:y30:2001:DB8:BEEF::/80", and then the RP address would
   be "2001:DB8:BEEF::y".

6.  Operational Considerations

   This section describes the major operational considerations for those
   deploying this mechanism.

6.1.  RP Redundancy

   A technique called "Anycast RP" is used within a PIM-SM domain to
   share an address and multicast state information between a set of RPs
   mainly for redundancy purposes.  Typically, MSDP has been used for
   this as well [ANYCASTRP].  There are also other approaches, such as
   using PIM for sharing this information [ANYPIMRP].

   The most feasible candidate for RP failover is using PIM for Anycast
   RP or "anycasting" (i.e., the shared-unicast model [ANYCAST]) the RP
   address in the Interior Gateway Protocol (IGP) without state sharing
   (although depending on the redundancy requirements, this may or may
   not be enough).  However, the redundancy mechanisms are outside of
   the scope of this memo.

6.2.  RP Deployment

   As there is no need to share inter-domain state with MSDP, each
   Designated Router connecting multicast sources could act as an RP
   without scalability concerns about setting up and maintaining MSDP
   sessions.

   This might be particularly attractive when one is concerned about RP
   redundancy.  In the case where the DR close to a major source for a
   group acts as the RP, a certain amount of fate-sharing properties can
   be obtained without using any RP failover mechanisms: if the DR goes
   down, the multicast transmission may not work anymore in any case.

   Along the same lines, its may also be desirable to distribute the RP
   responsibilities to multiple RPs.  As long as different RPs serve
   different groups, this is trivial: each group could map to a
   different RP (or sufficiently many different RPs that the load on one
   RP is not a problem).  However, load sharing challenges one group
   faces are similar to those of Anycast-RP.

6.3.  Guidelines for Assigning IPv6 Addresses to RPs

   With this mechanism, the RP can be given basically any unicast
   network prefix up to /64. The interface identifier will have to be
   manually configured to match "RIID".

   RIID = 0 must not be used, as using it would cause ambiguity with the
   Subnet-Router Anycast Address [ADDRARCH].

   If an administrator wishes to use an RP address that does not conform
   to the addressing topology but is still from the network provider's
   unicast prefix (e.g., an additional loopback address assigned on a
   router, as described in Example 2 in Section 5.1), that address can
   be injected into the routing system via a host route.

6.4.  Use as a Substitute for BSR

   With embedded-RP, use of BSR or other RP configuration mechanisms
   throughout the PIM domain is not necessary, as each group address
   specifies the RP to be used.

6.5.  Controlling the Use of RPs

   Compared to the MSDP inter-domain ASM model, the control and
   management of who can use an RP, and how, changes slightly and
   deserves explicit discussion.

MSDP advertisement filtering typically includes at least two
capabilities: filtering who is able to create a global session
("source filtering") and filtering which groups should be globally
accessible ("group filtering").  These are done to prevent local
groups from being advertised to the outside or unauthorized senders
from creating global groups.

However, such controls do not yet block the outsiders from using such
groups, as they could join the groups even without Source Active
advertisement with a (Source, Group) or (S,G) Join by
guessing/learning the source and/or the group address.  For proper
protection, one should set up, for example, PIM multicast scoping
borders at the border routers.  Therefore, embedded-RP has by default
a roughly equivalent level of "protection" as MSDP with SA filtering.

A new issue with control is that nodes in a "foreign domain" may
register to an RP, or send PIM Join to an RP.  (These have been
possible in the past as well, to a degree, but only through willful
attempts or purposeful RP configuration at DRs.)  The main threat in
this case is that an outsider may illegitimately use the RP to host
his/hers own group(s).  This can be mitigated to an extent by
filtering which groups or group ranges are allowed at the RP; more
specific controls are beyond the scope of this memo.  Note that this
does not seem to be a serious threat in the first place, as anyone
with a /64 unicast prefix can create their own RP without having to
illegitimately get it from someone else.

7.  The Embedded-RP Group-to-RP Mapping Mechanism

   This section specifies the group-to-RP mapping mechanism for Embedded
   RP.

7.1.  PIM-SM Group-to-RP Mapping

   The only PIM-SM modification required is implementing this mechanism
   as one group-to-RP mapping method.

   The implementation will have to recognize the address format and
   derive and use the RP address by using the rules in Section 4.  This
   information is used at least when performing Reverse Path Forwarding
   (RPF) lookups, when processing Join/Prune messages, or performing
   Register-encapsulation.

   To avoid loops and inconsistencies, for addresses in the range
   FF70::/12, the Embedded-RP mapping MUST be considered the longest
   possible match and higher priority than any other mechanism.

It is worth noting that compared to the other group-to-RP mapping
mechanisms, which can be precomputed, the embedded-RP mapping must be
redone for every new IPv6 group address that would map to a different
RP.  For efficiency, the results may be cached in an implementation-
specific manner, to avoid computation for every embedded-RP packet.

This group-to-RP mapping mechanism must be supported by the RP, the
DR adjacent to the senders, and any router on the path from any
receiver to the RP.  Paths for Shortest Path Tree (SPT) formation and
Register-Stop do not require the support, as those are accomplished
with an (S,G) Join.

7.2.  Overview of the Model

   This section gives a high-level, non-normative overview of how
   Embedded RP operates, as specified in the previous section.

   The steps when a receiver wishes to join a group are as follows:

   1. A receiver finds out a group address by some means (e.g., SDR or a
      web page).

   2. The receiver issues an Multicast Listener Discovery (MLD) Report,
      joining the group.

   3. The receiver's DR will initiate the PIM-SM Join process towards
      the RP encoded in the multicast address, irrespective of whether
      it is in the "local" or "remote" PIM domain.

   The steps when a sender wishes to send to a group are as follows:

   1. A sender finds out a group address by using an unspecified method
      (e.g., by contacting the administrator for group assignment or
      using a multicast address assignment protocol).

   2. The sender sends to the group.

   3. The sender's DR will send the packets unicast-encapsulated in
      PIM-SM Register-messages to the RP address encoded in the
      multicast address (in the special case that DR is the RP, such
      sending is only conceptual).

   In fact, all the messages go as specified in [PIM-SM]; embedded-RP
   just acts as a group-to-RP mapping mechanism.  Instead of obtaining
   the address of the RP from local configuration or configuration
   protocols (e.g., BSR), the algorithm derives it transparently from
   the encoded multicast address.

8.  Scalability Analysis

   Interdomain MSDP model for connecting PIM-SM domains is mostly
   hierarchical in configuration and deployment, but flat with regard to
   information distribution.  The embedded-RP inter-domain model behaves
   as if every group formed its own Internet-wide PIM domain, with the
   group mapping to a single RP, wherever the receivers or senders are
   located.  Hence, the inter-domain multicast becomes a flat, RP-
   centered topology.  The scaling issues are described below.

   Previously, foreign sources sent the unicast-encapsulated data to
   their "local" RP; now they are sent to the "foreign" RP responsible
   for the specific group.  This is especially important with large
   multicast groups where there are a lot of heavy senders --
   particularly if implementations do not handle unicast-decapsulation
   well.

   With IPv4 ASM multicast, there are roughly two kinds of Internet-wide
   state: MSDP (propagated everywhere), and multicast routing state (on
   the receiver or sender branches).  The former is eliminated, but the
   backbone routers might end up with (*, G) and (S, G, rpt) state
   between receivers (and past receivers, for PIM Prunes) and the RP, in
   addition to (S, G) states between the receivers and senders, if SPT
   is used.  However, the total amount of state is smaller.

   In both inter-domain and intra-domain cases, the embedded-RP model is
   practically identical to the traditional PIM-SM in intra-domain.  On
   the other hand, PIM-SM has been deployed (in IPv4) in inter-domain
   using MSDP; compared to that inter-domain model, this specification
   simplifies the tree construction (i.e., multicast routing) by
   removing the RP for senders and receivers in foreign domains and
   eliminating the MSDP information distribution.

   As the address of the RP is tied to the multicast address, the RP
   failure management becomes more difficult, as the deployed failover
   or redundancy mechanisms (e.g., BSR, Anycast-RP with MSDP) cannot be
   used as-is.  However, Anycast-RP using PIM provides equal redundancy;
   this described briefly in Section 6.1.

   The PIM-SM specification states, "Any RP address configured or
   learned MUST be a domain-wide reachable address".  What "reachable"
   precisely means is not clear, even without embedded-RP.  This
   statement cannot be proven, especially with the foreign RPs, as one
   cannot even guarantee that the RP exists.  Instead of manually
   configuring RPs and DRs (configuring a non-existent RP was possible,
   though rare), with this specification the hosts and users using
   multicast indirectly specify the RP themselves, lowering the
   expectancy of the RP reachability.  This is a relatively significant

problem but not much different from the current multicast deployment:
e.g., MLDv2 (S,G) joins, whether ASM or SSM, yield the same result
[PIMSEC].

Being able to join/send to remote RPs raises security concerns that
are considered separately, but it has an advantage too: every group
has a "responsible RP" that is able to control (to some extent) who
is able to send to the group.

A more extensive description and comparison of the inter-domain
multicast routing models (traditional ASM with MSDP, embedded-RP,
SSM) and their security properties has been described in [PIMSEC].

9.  Acknowledgements

Jerome Durand commented on an early version of this memo.  Marshall
Eubanks noted an issue regarding short plen values.  Tom Pusateri
noted problems with an earlier SPT-join approach.  Rami Lehtonen
pointed out issues with the scope of SA-state and provided extensive
commentary.  Nidhi Bhaskar gave the document a thorough review.
Toerless Eckert, Hugh Holbrook, and Dave Meyer provided very
extensive feedback.  In particular, Pavlin Radoslavov, Dino
Farinacci, Nidhi Bhaskar, and Jerome Durand provided good comments
during and after WG last call.  Mark Allman, Bill Fenner, Thomas
Narten, and Alex Zinin provided substantive comments during the IESG
evaluation.  The whole MboneD working group is also acknowledged for
continued support and comments.

10.  Security Considerations

The addresses of RPs are encoded in the multicast addresses, thus
becoming more visible as single points of failure.  Even though this
does not significantly affect the multicast routing security, it may
expose the RP to other kinds of attacks.  The operators are
encouraged to pay special attention to securing these routers.  See
Section 6.1 for considerations regarding failover and Section 6.2 for
placement of RPs leading to a degree of fate-sharing properties.

As any RP will have to accept PIM-SM Join/Prune/Register messages
from any DR, this might cause a potential Denial of Service attack
scenario.  However, this can be mitigated, as the RP can discard all
such messages for all multicast addresses that do not encode the
address of the RP.  Both the sender- and receiver-based attacks are
described at greater length in [PIMSEC].

Additionally, the implementation SHOULD also allow manual
configuration of which multicast prefixes are allowed to be used.
This can be used to limit the use of the RP to designated groups
only.  In some cases, being able to restrict (at the RP) which
unicast addresses are allowed to send or join to a group is
desirable.  (However, note that Join/Prune messages would still leave
state in the network, and Register messages can be spoofed [PIMSEC].)
Obviously, these controls are only possible at the RP, not at the
intermediate routers or the DR.

It is RECOMMENDED that routers supporting this specification do not
act as RPs unless explicitly configured to do so, as becoming an RP
does not require any advertisement (e.g., through BSR or manually).
Otherwise, any router could potentially become an RP (and be abused
as such).  Further, multicast groups or group ranges to-be-served MAY
need to be explicitly configured at the RPs, to protect them from
being used unwillingly.  Note that the more specific controls (e.g.,
"insider-must-create" or "invite-outsiders" models) as to who is
allowed to use the groups are beyond the scope of this memo.

Excluding internal-only groups from MSDP advertisements does not
protect the groups from outsiders but only offers security by
obscurity; embedded-RP offers similar level of protection.  When real
protection is desired, PIM scoping for example, should be set up at
the borders. This is described at more length in Section 6.5.

One should observe that the embedded-RP threat model is actually
rather similar to SSM; both mechanisms significantly reduce the
threats at the sender side.  On the receiver side, the threats are
somewhat comparable, as an attacker could do an MLDv2 (S,G) join
towards a non-existent source, which the local RP could not block
based on the MSDP information.

The implementation MUST perform at least the same address validity
checks to the embedded-RP address as it would to one received via
other means; at least fe80::/10, ::/16, and ff00::/8 should be
excluded.  This is particularly important, as the information is
derived from the untrusted source (i.e., any user in the Internet),
not from the local configuration.

A more extensive description and comparison of the inter-domain
multicast routing models (traditional ASM with MSDP, embedded-RP,
SSM) and their security properties has been done separately in
[PIMSEC].

11.  References

11.1.  Normative References

   [ADDRARCH]   Hinden, R. and S. Deering, "Internet Protocol Version 6
                (IPv6) Addressing Architecture", RFC 3513, April 2003.

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

   [RFC3306]    Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6
                Multicast Addresses", RFC 3306, August 2002.

11.2.  Informative References

   [ANYCAST]    Hagino, J. and K. Ettikan, "An analysis of IPv6 anycast",
                Work in Progress, June 2003.

   [ANYCASTRP]  Kim, D., Meyer, D., Kilmer, H., and D. Farinacci,
                "Anycast Rendevous Point (RP) mechanism using Protocol
                Independent Multicast (PIM) and Multicast Source
                Discovery Protocol (MSDP)", RFC 3446, January 2003.

   [ANYPIMRP]   Farinacci, D. and Y. Cai, "Anycast-RP using PIM", Work in
                Progress, June 2004.

   [BSR]        Fenner, B., et al., "Bootstrap Router (BSR) Mechanism for
                PIM Sparse Mode", Work in Progress, July 2004.

   [MSDP]       Fenner, B. and D. Meyer, "Multicast Source Discovery
                Protocol (MSDP)", RFC 3618, October 2003.

   [PIMSEC]     Savola, P., Lehtonen, R., and D. Meyer, "PIM-SM Multicast
                Routing Security Issues and Enhancements", Work in
                Progress, October 2004.

   [PIM-SM]     Fenner, B. et al, "Protocol Independent Multicast -
                Sparse Mode (PIM-SM): Protocol Specification (Revised)",
                Work in Progress, July 2004.

   [SSM]        Holbrook, H. et al, "Source-Specific Multicast for IP",
                Work in Progress, September 2004.

   [V6MISSUES]  Savola, P., "IPv6 Multicast Deployment Issues", Work in
                Progress, September 2004.

A.  Discussion about Design Tradeoffs

   The document only specifies FF70::/12 for now; if/when the upper-most
   bit is used, one must specify how FFF0::/12 applies to Embedded-RP.
   For example, a different mode of PIM or another protocol might use
   that range, in contrast to FF70::/12, as currently specified, being
   for PIM-SM only.

   Instead of using flags bits ("FF70::/12"), one could have used the
   leftmost reserved bits instead ("FF3x:8000::/17").

   It has been argued that instead of allowing the operator to specify
   RIID, the value could be pre-determined (e.g., "1").  However, this
   has not been adopted, as this eliminates address assignment
   flexibility from the operator.

   Values 64 < "plen" < 96 would overlap with upper bits of the
   multicast group-id; due to this restriction, "plen" must not exceed
   64 bits.  This is in line with RFC 3306.

   The embedded-RP addressing could be used to convey other information
   (other than RP address) as well, for example, what should be the RPT
   threshold for PIM-SM.  These could be, whether feasible or not,
   encoded in the RP address somehow, or in the multicast group address.
   In any case, such modifications are beyond the scope of this memo.

   For the cases where the RPs do not exist or are unreachable, or too
   much state is being generated to reach in a resource exhaustion
   Denial of Service attack, some forms of rate-limiting or other
   mechanisms could be deployed to mitigate the threats while trying not
   to disturb the legitimate usage.  However, as the threats are
   generic, they are considered out of scope and discussed separately in
   [PIMSEC].

Authors' Addresses

    Pekka Savola
    CSC/FUNET
    Espoo, Finland

    EMail: psavola@funet.fi


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

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