

Bootstrap Router (BSR) Mechanism for PIM

Status of this Document

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This document is a product of the IETF PIM WG. Comments should be addressed to the authors, or the WG's mailing list at pim@ietf.org.

Copyright Notice

Copyright (C) The Internet Society (2005).

This document specifies the Bootstrap Router (BSR) mechanism for the class of multicast routing protocols in the PIM (Protocol Independent Multicast) family that use the concept of a Rendezvous Point as a means for receivers to discover the sources that send to a particular multicast group. BSR is one way that a multicast router can learn the set of group-to-RP mappings required in order to function. The mechanism is dynamic, largely self-configuring, and robust to router failure.

<u>1.</u>	<u>Introduction.</u>	<u>5</u>
<u>1.1.</u>	<u>Background</u>	<u>5</u>
<u>1.2.</u>	<u>Protocol Overview.</u>	<u>7</u>
<u>1.3.</u>	<u>Administrative Scoping and BSR</u>	<u>8</u>
<u>2.</u>	<u>BSR State and Timers.</u>	<u>10</u>
<u>3.</u>	<u>Bootstrap Router Election and RP-Set</u>	
	<u>Distribution.</u>	<u>10</u>
<u>3.1.</u>	<u>Bootstrap Router Election.</u>	<u>10</u>
<u>3.1.1.</u>	<u>Per-Scope-Zone Candidate-BSR State</u>	
	<u>Machine</u>	<u>11</u>
<u>3.1.2.</u>	<u>Per-Scope-Zone State Machine for Non-</u>	
	<u>Candidate-BSR Routers</u>	<u>13</u>
<u>3.1.3.</u>	<u>Bootstrap Message Processing Checks</u>	<u>15</u>
<u>3.1.4.</u>	<u>State Machine Transition Events</u>	<u>15</u>
<u>3.1.5.</u>	<u>State Machine Actions</u>	<u>16</u>
<u>3.2.</u>	<u>Sending Candidate-RP-Advertisement Messages.</u>	<u>18</u>
<u>3.3.</u>	<u>Creating the RP-Set at the BSR</u>	<u>19</u>
<u>3.4.</u>	<u>Forwarding Bootstrap Messages.</u>	<u>22</u>
<u>3.5.</u>	<u>Unicasting Bootstrap Messages to New and</u>	
	<u>Rebooting Routers.</u>	<u>22</u>
<u>3.6.</u>	<u>Receiving and Using the RP-Set</u>	<u>23</u>
<u>4.</u>	<u>Message Formats</u>	<u>23</u>
<u>4.1.</u>	<u>Bootstrap Message Format</u>	<u>25</u>
<u>4.1.1.</u>	<u>Semantic Fragmentation of BSMs.</u>	<u>29</u>
<u>4.2.</u>	<u>Candidate-RP-Advertisement Message Format.</u>	<u>30</u>
<u>5.</u>	<u>Timers and Timer Values</u>	<u>32</u>
<u>6.</u>	<u>Security Considerations</u>	<u>35</u>
<u>6.1.</u>	<u>Possible Threats</u>	<u>35</u>
<u>6.2.</u>	<u>Limiting Third-Party DoS Attacks</u>	<u>35</u>
<u>6.3.</u>	<u>Bootstrap Message Security</u>	<u>36</u>
<u>6.3.1.</u>	<u>Rejecting Unicast Bootstrap Messages.</u>	<u>36</u>
<u>6.3.2.</u>	<u>Rejecting Bootstrap Messages from Invalid</u>	
	<u>Neighbors</u>	<u>37</u>
<u>6.4.</u>	<u>Candidate-RP-Advertisement Message Security.</u>	<u>37</u>
<u>6.4.1.</u>	<u>Non-Cryptographic Security of C-RP-Adv</u>	
	<u>Messages.</u>	<u>37</u>
<u>6.4.2.</u>	<u>Cryptographic Security of C-RP-Adv</u>	
	<u>Messages.</u>	<u>38</u>
<u>6.5.</u>	<u>Denial of Service using IPsec.</u>	<u>38</u>
<u>7.</u>	<u>Contributors.</u>	<u>39</u>
<u>8.</u>	<u>Acknowledgments</u>	<u>39</u>
<u>9.</u>	<u>IANA Considerations</u>	<u>39</u>
<u>10.</u>	<u>Normative References</u>	<u>39</u>
<u>11.</u>	<u>Informative References</u>	<u>40</u>

This document assumes some familiarity with the concepts of Protocol Independent Multicast - Sparse Mode (PIM-SM), as defined in [1], and Bi-directional Protocol Independent Multicast (BIDIR-PIM), as defined in [2], as well as with Administratively Scoped IP Multicast, as described in [3], and the IPv6 Scoped Address Architecture, described in [4].

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same Rendezvous Point (RP). The PIM specifications do not mandate the use of a single mechanism to provide routers with the information to perform this group-to-RP mapping.

This document describes the PIM Bootstrap Router (BSR) mechanism. BSR is one way that a multicast router can learn the information required to perform the group-to-RP mapping. The mechanism is dynamic, largely self-configuring, and robust to router failure.

BSR was first defined in [RFC 2362](#) [7], which has since been obsoleted. This document provides an updated specification of the BSR mechanism from [RFC 2362](#), and also extends it to cope with administratively scoped region boundaries and different flavours of routing protocols.

Throughout the document, any reference to the PIM protocol family is restricted to the subset of RP-based protocols, namely PIM-SM and BIDIR-PIM, unless stated otherwise.

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](#) [6].

[1.1](#). Background

A PIM domain is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary defined by PIM Multicast Border Routers (PMBRs). PMBRs connect each PIM domain to the rest of the internet.

Every PIM multicast group needs to be associated with the IP address of a Rendezvous Point (RP). This address is used as the root of a group-specific distribution tree whose branches extend to all nodes in the domain that want to receive traffic sent to the group. Senders inject packets into the tree in such a manner that they reach all connected receivers. How this is done and how the packets are forwarded along the distribution tree depends on the particular routing protocol.

For all senders to reach all receivers, it is crucial that all routers in the domain use the same mappings of group addresses to RP addresses.

An exception to the above is where a PIM domain has been broken up into multiple administrative scope regions. These are regions where a border has been configured so that a set of multicast groups will not be forwarded across that border. In this case, all PIM routers within the same scope region must map a particular scoped group to the same RP within that region.

In order to determine the RP for a multicast group, a PIM router maintains a collection of group-to-RP mappings, called the RP-Set. A group-to-RP mapping contains the following elements.

- o Multicast group range, expressed as an address and prefix length
- o RP priority
- o RP address
- o Hash mask length
- o SM / BIDIR flag

In general, the group ranges of these group-to-RP mappings may overlap in arbitrary ways; hence a particular multicast group may be covered by multiple group-to-RP mappings. When this is the case, the router chooses only one of the RPs by applying a deterministic algorithm so that all routers in the domain make the same choice. It is important to note that this algorithm is part of the specification of the individual routing protocols (and may differ among them), not of the BSR specification.

There are a number of ways in which such group-to-RP mappings can be established. The simplest solution is for all the routers in the domain to be statically configured with the same information. However, static configuration generally doesn't scale well, and, except when used in conjunction with Anycast-RP (see [\[8\]](#) and [\[9\]](#)), does not dynamically adapt to route around router or link failures.

The BSR mechanism provides a way in which viable group-to-RP mappings can be created and rapidly distributed to all the PIM routers in a domain. It is adaptive, in that if an RP becomes unreachable, this will be detected and the RP-Sets will be modified so that the unreachable RP is no longer used.

[1.2.](#) Protocol Overview

In this section we give an informal and non-definitive overview of the BSR mechanism. The definitive specification begins in [section 2](#).

The general idea behind the BSR mechanism is that some of the PIM routers within a PIM domain are configured to be potential RPs for the domain. These are known as Candidate-RPs (C-RPs). A subset of the C-RPs will eventually be used as the actual RPs for the domain. In addition, some of the PIM routers in the domain are configured to be candidate bootstrap routers, or Candidate-BSRs (C-BSRs). One of these C-BSRs will be elected to be the bootstrap router (BSR) for the domain, and all the PIM routers in the domain will learn the result of this election through Bootstrap messages. The C-RPs will then report their candidacy to the elected BSR, which chooses a subset of these C-RPs and distributes corresponding group-to-RP mappings to all the routers in the domain through Bootstrap messages.

In more detail, the BSR mechanism works as follows. There are four basic phases (although in practice all phases may be occurring simultaneously):

- [1.](#) BSR Election. Each Candidate-BSR originates Bootstrap messages (BSMs). Every BSM contains a BSR Priority field. Routers within the domain flood the BSMs throughout the domain. A C-BSR that hears about a higher-priority C-BSR than itself then suppresses its sending of further BSMs for some period of time. The single remaining C-BSR becomes the elected BSR, and its BSMs inform all the other routers in the domain that it is the elected BSR.
- [2.](#) C-RP Advertisement. Each Candidate-RP within a domain sends periodic Candidate-RP-Advertisement (C-RP-Adv) messages to the elected BSR. A C-RP-Adv message includes the priority of the advertising C-RP, as well as a list of group ranges for which the candidacy is advertised. In this way, the BSR learns about possible RPs that are currently up and reachable.
- [3.](#) RP-Set Formation. The BSR selects a subset of the C-RPs that it has received C-RP-Adv messages from to form the RP-Set. In general it should do this in such a way that the RP-Set is neither too large to inform all the routers in the domain about, nor too small so that load is overly concentrated on some RPs. It should also attempt to produce an RP-Set that does not change frequently.
- [4.](#) RP-Set Flooding. In future Bootstrap messages, the BSR includes the RP-Set information. Bootstrap messages are flooded through the domain, which ensures that the RP-Set rapidly reaches all the routers in the domain. BSMs are originated periodically to ensure

consistency after failure restoration.

When a PIM router receives a Bootstrap message, it adds the group-to-RP mappings contained therein to its pool of mappings obtained from other sources (e.g. static configuration). It calculates the final mappings of group addresses to RP addresses from this pool according to rules specific to the particular routing protocol and uses that information to construct multicast distribution trees.

If a PIM domain becomes partitioned, each area separated from the old BSR will elect its own BSR, which will distribute an RP-Set containing RPs that are reachable within that partition. When the partition heals, another election will occur automatically and only one of the BSRs will continue to send out Bootstrap messages. As is expected at the time of a partition or healing, some disruption in packet delivery may occur. This time will be on the order of the region's round-trip time and the BS_Timeout value.

[1.3.](#) Administrative Scoping and BSR

The mechanism described in the previous section does not work when the PIM domain is divided into administratively scoped regions. To handle this situation, we use the protocol modifications described in this section.

Administrative scoping permits a PIM domain to be divided into multiple admin-scope regions. Each admin-scope region is a convex connected set of PIM routers, and is associated with a set of group addresses. The boundary of the admin-scope region is formed by Zone Border Routers (ZBRs). ZBRs are configured not to forward traffic for any of the scoped group addresses into or out of the scoped region. It is important to note that a given scope boundary always creates at least two scoped regions: one on either side of the boundary.

In IPv4, administratively scoped regions are associated with a set of addresses given by an address and a prefix length. In IPv6, administratively scoped regions are associated with a set of addresses given by a single scope ID value. The set of addresses corresponding to a given scope ID value is defined in [5]. For example, a scope ID of 5 maps to the 16 IPv6 address ranges ff[0-f]5::/16.

There are certain topological restrictions on admin-scope regions. Firstly, the scope zone border must be complete and convex. By this we mean that there must be no path from inside the scoped zone to outside it that does not pass through a configured scope border router, and that the multicast capable path between any arbitrary pair of multicast routers in the scope zone must remain in the zone. In addition, a boundary for one scope must always be a boundary for all smaller scopes,

where a smaller scope for IPv4 is one whose address range is contained within the other address range, and for IPv6 is one whose scope ID is less than the other scope ID.

Administrative scoping complicates BSR because we do not want a PIM router within the scoped region to use an RP outside the scoped region. Thus we need to modify the basic mechanism to ensure that this doesn't happen.

This is done by running a separate copy of the basic BSR mechanism, as described in the previous section, within each admin scope region of a PIM domain. Thus a separate BSR election takes place for each admin-scope region, a C-RP typically registers to the BSR of every admin scope zone it is in, and every PIM router receives Bootstrap messages for every scope zone it is in. The Bootstrap messages sent by the BSR for a particular scope zone contain information about the RPs that should be used for the set of addresses associated with that scope zone.

Bootstrap messages are marked to indicate which scope zone they belong to. Such admin scoped Bootstrap messages are flooded in the normal way, but will not be forwarded by a ZBR across the boundary for that scope zone.

For the BSR mechanism to function correctly with admin scoping, within each admin scope region there must be at least one C-BSR, and at least one C-RP that is configured to be a C-RP for the set of group addresses associated with the scoped region.

Even when administrative scoping is used, a copy of the BSR mechanism is still used across the entire PIM domain, in order to distribute RP information for groups that are not administratively scoped. We call this copy of the mechanism Non-Scoped BSR. The copies of the mechanism run for each admin-scope region are called Scoped BSR.

Only the C-BSRs and the ZBRs need to be configured to know about the existence of the scope zones. Other routers, including the C-RPs, learn of their existence from Bootstrap messages.

All PIM routers within a PIM bootstrap domain where admin scope ranges are in use must be capable of receiving Bootstrap messages and storing the winning BSR and RP-Set for all admin scope zones that apply. Thus PIM routers that only implement [RFC 2362](#) or Non-Scoped BSR (which only allows one BSR per domain) cannot be used within the admin-scope regions of a PIM domain.

[2.](#) BSR State and Timers

A PIM router implementing BSR holds the following state.

RP-Set

Per Configured or Learned Scope Zone (Z):

At all routers:

Current Bootstrap Router's IP Address

Current Bootstrap Router's BSR Priority

Last BSM received from current BSR

Bootstrap Timer (BST(Z))

Per group-to-RP mapping (M):

Group-to-RP mapping Expiry Timer (GET(M,Z))

At a Candidate-BSR for Z:

My state: One of "Candidate-BSR", "Pending-BSR",
"Elected-BSR"

At a router that is not a Candidate-BSR for Z:

My state: One of "Accept Any", "Accept Preferred"

Scope-Zone Expiry Timer (SZT(Z))

At the current Bootstrap Router for Z only:

Per group-to-C-RP mapping (M):

Group-to-C-RP mapping Expiry Timer (CGET(M,Z))

At a C-RP only:

C-RP Advertisement Timer (CRPT)

[3. Bootstrap Router Election and RP-Set Distribution](#)

[3.1. Bootstrap Router Election](#)

For simplicity, Bootstrap messages are used in both the BSR election and the RP-Set distribution mechanisms.

Each Bootstrap message indicates the scope that it belongs to. If the Admin Scope Zone bit is set in the first group range in the Bootstrap message, the message is called a scoped BSM. If the Admin Scope Zone bit is not set in the first group range in the Bootstrap message, the message is called a non-scoped BSM.

In a scoped IPv4 BSM, the scope of the message is given by the first group range in the message, which can be any sub-range of 224/4. In a scoped IPv6 BSM, the scope of the message is given by the scope ID of the first group range in the message, which must have a mask length of at least 16. For example, a group range of ff05::/16 with the Admin Scope Zone bit set indicates that the Bootstrap message is for the scope with scope ID 5. If the mask length of the first group range in a scoped IPv6 BSM is less than 16, the message **MUST** be dropped and a warning **SHOULD** be logged.

The state machine for Bootstrap messages depends on whether or not a router has been configured to be a Candidate-BSR for a particular scope zone. The per-scope-zone state machine for a C-BSR is given below, followed by the state machine for a router that is not configured to be a C-BSR.

[3.1.1. Per-Scope-Zone Candidate-BSR State Machine](#)

When in C-BSR state				
Event	Receive Preferred BSM	Bootstrap Timer Expires	Receive Non-preferred BSM from Elected BSR	
Action	-> C-BSR state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout	-> P-BSR state Set Bootstrap Timer to BS_Rand_Override	-> P-BSR state Set Bootstrap Timer to BS_Rand_Override	

When in P-BSR state			
Event	Receive Preferred BSM	Bootstrap Timer Expires	Receive Non-preferred BSM
Action	-> C-BSR state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout	-> E-BSR state Originate BSM; Set Bootstrap Timer to BS_Period	-> P-BSR state

When in E-BSR state			
Event	Receive Preferred BSM	Bootstrap Timer Expires	Receive Non-preferred BSM
Action	-> C-BSR state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout	-> E-BSR state Originate BSM; Set Bootstrap Timer to BS_Period	-> E-BSR state Originate BSM; Set Bootstrap Timer to BS_Period

A Candidate-BSR may be in one of three states for a particular scope zone:

Candidate-BSR (C-BSR)

The router is a candidate to be the BSR for the scope zone, but currently another router is the preferred BSR.

Pending-BSR (P-BSR)

The router is a candidate to be the BSR for the scope zone. Currently no other router is the preferred BSR, but this router is not yet the elected BSR. This is a temporary state that prevents rapid thrashing of the choice of BSR during BSR election.

Elected-BSR (E-BSR)

The router is the elected BSR for the scope zone and it must perform all the BSR functions.

In addition to the three states, there is one timer:

- o The Bootstrap Timer (BST) - used to time out old bootstrap router information, and used in the election process to terminate P-BSR state.

On startup, the initial state for this configured scope zone is "Pending-BSR"; the Bootstrap Timer is initialized to BS_Rand_Override.

3.1.2. Per-Scope-Zone State Machine for Non-Candidate-BSR Routers

When in NoInfo state	
Event	Receive BSM
Action	-> AP state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout; Set SZT to SZ_Timeout

When in Accept Any state		
Event	Receive BSM	Scope-Zone Expiry Timer Expires
Action	-> AP state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout; Set SZT to SZ_Timeout	-> NoInfo state Cancel timers; Clear state

When in Accept Preferred state			
Event	Receive Preferred BSM	Bootstrap Timer Expires	Receive Non-preferred BSM
Action	-> AP state Forward BSM; Store RP-Set; Set Bootstrap Timer to BS_Timeout; Set SZT to SZ_Timeout	-> AA state Refresh RP-Set; Remove BSR state	-> AP state

A router that is not a Candidate-BSR may be in one of three states:

NoInfo

The router has no information about this scope zone. This state does not apply if the router is configured to know about this scope zone, or for the global scope zone. When in this state, no state information is held and no timers run that refer to this scope zone.

Accept Any (AA)

The router does not know of an active BSR, and will accept the first Bootstrap message it sees as giving the new BSR's identity and the RP-Set.

Accept Preferred (AP)

The router knows the identity of the current BSR, and is using the RP-Set provided by that BSR. Only Bootstrap messages from that BSR or from a C-BSR with higher weight than the current BSR will be accepted.

In addition to the three states, there are two timers:

- o The Bootstrap Timer (BST) - used to time out old bootstrap router information.
- o The Scope-Zone Expiry Timer (SZT) - used to time out the scope zone itself if Bootstrap messages specifying this scope zone stop arriving.

On startup, the initial state for this scope zone is "Accept Any" for routers that know about this scope zone, either through configuration or because the scope zone is the global scope which always exists; the Scope-Zone Expiry Timer is considered to be always running for such scope zones. For routers that do not know about a particular scope zone, the initial state is NoInfo; no timers exist for the scope zone.

[3.1.3.](#) Bootstrap Message Processing Checks

When a Bootstrap message is received, the following initial checks must be performed:

```
if ((DirectlyConnected(BSM.src_ip_address) == FALSE) OR
    (we have no Hello state for BSM.src_ip_address)) {
    drop the Bootstrap message silently
}

if (BSM.dst_ip_address == ALL-PIM-ROUTERS) {
    if (BSM.src_ip_address != RPF_neighbor(BSM.BSR_ip_address)) {
        drop the Bootstrap message silently
    }
} else if (BSM.dst_ip_address is one of my addresses) {
    if ((any previous BSM for this scope has been accepted) OR
        (more than BS_Period has elapsed since startup)) {
        #the packet was unicast, but this wasn't
        #a quick refresh on startup
        drop the Bootstrap message silently
    }
} else {
    drop the Bootstrap message silently
}

if (the interface the message arrived on is an Admin Scope
    border for the BSM.first_group_address) {
    drop the Bootstrap message silently
}
```

Basically, the packet must have come from a directly connected neighbor for which we have active Hello state. It must have been sent to the ALL-PIM-ROUTERS group by the correct upstream router towards the BSR that originated the Bootstrap message, or the router must have recently restarted, have no BSR state for that admin scope and have received the Bootstrap message by unicast. In addition it must not have arrived on an interface that is a configured admin scope border for the first group address contained in the Bootstrap message.

[3.1.4.](#) State Machine Transition Events

If the Bootstrap message passes the initial checks above without being discarded, then it may cause a state transition event in one of the above state machines. For both candidate and non-candidate BSRs, the following transition events are defined:

Receive Preferred BSM

A Bootstrap message is received from a BSR that has higher or

equal weight than the current BSR. If a router is in P-BSR state, then it uses its own weight as that of the current BSR.

A Bootstrap message is also preferred if it is from the current BSR with a lower weight than the previous BSM it sent, provided that if the router is a Candidate BSR the current BSR still has a weight higher or equal than the router itself. In this case, the "Current Bootstrap Router's BSR Priority" state must be updated. (For lower weight, see Non-preferred BSM from Elected BSR case.)

The weight of a BSR is defined to be the concatenation in fixed-precision unsigned arithmetic of the BSR Priority field from the Bootstrap message and the IP address of the BSR from the Bootstrap message (with the BSR Priority taking the most-significant bits and the IP address taking the least significant bits).

Receive Non-preferred BSM

A Bootstrap message is received from a BSR that has lower weight than the current BSR. If a router is in P-BSR state, then it uses its own weight as that of the current BSR.

Receive Non-preferred BSM from Elected BSR

A Bootstrap message is received from the elected BSR, but the BSR Priority field in the received message has changed, so that now the currently elected BSR has lower weight than the router itself.

Receive BSM

A Bootstrap message is received, regardless of BSR weight.

In addition to state machine transitions caused by the receipt of Bootstrap messages, a state machine transition takes place each time the Bootstrap Timer or Scope-Zone Expiry Timer expires.

3.1.5. State Machine Actions

The state machines specify actions that include setting the Bootstrap Timer and the Scope-Zone Expiry Timer to various values. These values are defined in [Section 5](#).

In addition to setting and cancelling the timers, the following actions may be triggered by state changes in the state machines:

Forward BSM

A Bootstrap message that passes the Bootstrap Message Processing Checks is forwarded out of all interfaces with PIM

neighbors (including the interface it is received on), except where this would cause the BSM to cross an admin-scope boundary for the scope zone indicated in the message. For details, see [section 3.4](#).

Originate BSM

A new Bootstrap message is constructed by the BSR, giving the BSR's address and BSR priority, and containing the BSR's chosen RP-Set. The message is forwarded out of all interfaces on which PIM neighbors exist, except where this would cause the BSM to cross an admin-scope boundary for the scope zone indicated in the message.

Store RP-Set

The router uses the group-to-RP mappings contained in a BSM to update its local RP-Set.

This action is skipped for an empty BSM. A BSM is empty if it contains no group ranges, or if it only contains a single group range where that group range has the Admin Scope Zone bit set (a scoped BSM) and an RP count of zero.

If a mapping does not yet exist, it is created and the associated Group-to-RP mapping Expiry Timer (GET) is initialized with the holdtime from the BSM.

If a mapping already exists, its GET is set to the holdtime from the BSM. If the holdtime is zero, the mapping is removed immediately. Note that for an existing mapping, the RP priority must be updated if changed.

Mappings for a group range are also to be immediately removed if they are not present in the received group range. This means that if there are any existing Group-to-RP mappings for a range where the respective RPs are not in the received range, then those mappings must be removed.

All RP mappings associated with the scope zone of the BSM are updated with the new hash mask length from the received BSM. This includes any RP mappings learned from the current BSR but not contained in the received BSM, as well as any RP mappings learned from any previous BSR for the scope zone.

In addition, the entire BSM is stored for use in the action Refresh RP-Set and to prime a new PIM neighbor as described below.

Refresh RP-Set

When the Bootstrap Timer expires, the router uses the copy of the last BSM that it has received to refresh its RP-Set according to the action Store RP-Set as if it had just received it. This will increase the chance that the group-to-RP mappings will not expire during the election of the new BSR.

Remove BSR state

When the Bootstrap Timer expires, all state associated with the current BSR is removed (see [section 2](#)). Note that this does not include any group-to-RP mappings.

[3.2.](#) Sending Candidate-RP-Advertisement Messages

Every C-RP periodically unicasts a C-RP-Adv message to the BSR for each scope zone for which it has state, to inform the BSR of the C-RP's willingness to function as an RP. These messages are sent with an interval of C_RP_Adv_Period, except when a new BSR is elected, see below.

When a new BSR is elected, the C-RP SHOULD send one to three C-RP-Adv messages, waiting a randomized amount of 0-3 seconds before sending each message. We recommend sending three messages because it is important that the BSR quickly learns which RPs are active, and some packet loss may occur when a new BSR is elected due to changes in the network. One way of implementing this is to set the CRPT to 0-3 seconds when the new BSR is elected, as well as setting a counter to 2. Whenever the CRPT expires, we first send a C-RP-Adv message as usual. Next, if the counter is non-zero, it is decremented and the CRPT is again set to 0-3 seconds instead of C_RP_Adv_Period.

[NOTE: Add a name for this timer and counter?]

The Priority field in these messages is used by the BSR to select which C-RPs to include in the RP-Set. Note that lower values of this field indicate higher priorities, so that a value of zero is the highest possible priority. C-RPs should by default send C-RP-Adv messages with the Priority field set to 192.

When a C-RP is being shut down, it SHOULD immediately send a C-RP-Adv message to the BSR for each scope zone for which it is currently serving as an RP; the Holdtime in this C-RP-Adv message should be zero. The BSR will then immediately time out the C-RP and generate a new Bootstrap message removing the shut down RP from the RP-Set.

[NOTE: Should a new BSM be sent immediately when a C-RP-Adv message with Holdtime of 0 is received? Need to clarify.]

A C-RP-Adv message carries a list of group address and group mask field pairs. This enables the C-RP to specify the group prefixes for which it is willing to be the RP. If the C-RP becomes an RP, it may enforce this scope acceptance when receiving Register or Join/Prune messages.

A C-RP is configured with a list of group ranges for which it should advertise itself as the C-RP. A C-RP uses the following algorithm to determine which ranges to send to a given BSR.

For each group range R in the list, the C-RP advertises that range to the scoped BSR for the smallest scope that "contains" R. For IPv6, the containing scope is determined by matching the scope identifier of the group range with the scope of the BSR. For IPv4, it is the longest-prefix match for R, amongst the known admin-scope ranges. If no scope is found to contain the group range the C-RP includes it in the C-RP-Adv sent to the non-scoped BSR. If a non-scoped BSR is not known, the range is not included in any C-RP-Adv.

In addition, for each IPv4 group range R in the list, for each scoped BSR whose scope range is strictly contained within R, the C-RP SHOULD by default advertise that BSR's scope range to that BSR. And for each IPv6 group range R in the list with prefix length < 16, the C-RP SHOULD by default advertise each sub-range of prefix length 16 to the scoped BSR with the corresponding scope ID. An implementation MAY supply a configuration option to prevent the behavior described in this paragraph, but such an option SHOULD be disabled by default.

For IPv6, the mask length of all group ranges included in the C-RP-Adv message sent to a scoped BSR MUST be ≥ 16 .

If the above algorithm determines that there are no group ranges to advertise to the BSR for a particular scope zone, a C-RP-Adv message MUST NOT be sent to that BSR. A C-RP MUST NOT send a C-RP-Adv message with no group ranges in it.

If the same router is the BSR for more than one scope zone, the C-RP-Adv messages for these scope zones MAY be combined into a single message.

If the C-RP is a ZBR for an admin scope zone, then the Admin Scope Zone bit MUST be set in the C-RP-Adv messages it sends for that scope zone; otherwise this bit MUST NOT be set. This information is currently only used for logging purposes by the BSR, but might allow for future extensions of the protocol.

[3.3.](#) Creating the RP-Set at the BSR

Upon receiving a C-RP-Adv message, the router needs to decide whether or not to accept each of the group ranges included in the message. For

each group range in the message, the router checks to see if it is the elected BSR for any scope zone that contains the group range, or if it is elected as the non-scoped BSR. If so, the group range is accepted; if not, the group range is ignored.

If the group range is accepted, a group-to-C-RP mapping is created for this group range and the RP Address from the C-RP-Adv message.

If the mapping is not already part of the C-RP-Set, it is added to the C-RP-Set and the associated Group-to-C-RP mapping Expiry Timer (CGET) is initialized to the holdtime from the C-RP-Adv message. Its priority is set to the Priority from the C-RP-Adv message.

If the mapping is already part of the C-RP-Set, it is updated with the Priority from the C-RP-Adv message and its associated CGET is reset to the holdtime from the C-RP-Adv message. If the holdtime is zero, the mapping is immediately removed from the C-RP-Set.

The hash mask length is a global property of the BSR and is therefore the same for all mappings managed by the BSR.

For compatibility with the previous version of the BSR specification, a C-RP-Adv message with no group ranges SHOULD be treated as though it contained the single group range ff00::/8 or 224/4. Therefore, according to the rule above, this group range will be accepted if and only if the router is elected as the non-scoped BSR.

When a CGET expires, the corresponding group-to-C-RP mapping is removed from the C-RP-Set.

The BSR constructs the RP-Set from the C-RP-Set. It may apply a local policy to limit the number of Candidate-RPs included in the RP-Set. The BSR may override the prefix indicated in a C-RP-Adv message unless the 'Priority' field from the C-RP-Adv message is less than 128.

For inclusion in a BSM, the RP-Set is subdivided into sets of {group-prefix, RP-Count, RP-addresses}. For each RP-address, the corresponding Holdtime is included in the "RP-Holdtime" field. The format of the Bootstrap message allows 'semantic fragmentation', if the length of the original Bootstrap message exceeds the packet maximum boundaries. However, we recommend against configuring a large number of routers as C-RPs, to reduce the semantic fragmentation required.

In general BSMs are originated at regular intervals according to the BS_Period timer. We do recommend that a BSM is also originated whenever the RP-set to be announced in the BSMs changes. This will usually happen when receiving C-RP advertisements from a new C-RP, or when a C-RP is shut down (C-RP advertisement with a holdtime of zero). There

MUST however be a minimum of 10 seconds between each time a BSM is sent. In particular, when a new BSR is elected, it will first send one BSM (which is likely to be empty since it has not yet received any C-RP advertisements), and then wait at least 10 seconds before sending a new one. During those 10 seconds, it is likely to have received C-RP advertisements from all usable C-RPs (since we say that a C-RP should send one or more advertisements with small random delays of 0-3 seconds when a new BSR is elected). For this case in particular, where routers may not have a usable RP-set, we recommend originating a BSM as soon as those 10 seconds have passed. We suggest though that a BSR can do this in general. One way of implementing this, is to decrease the Bootstrap Timer to 10 seconds whenever the RP-set changes, while not changing the timer if it is less or equal to 10.

[NOTE: Add a name for this 10s value as a function of the 0-3s random delay?]

A BSR originates separate scoped BSMS for each scope zone for which it is the elected BSR, as well as originating non-scoped BSMS if it is the elected non-scoped BSR.

Each group-to-C-RP mapping is included in precisely one of these BSM, namely the scoped BSM for the narrowest scope containing the group range of the mapping, if any, or the non-scoped BSM otherwise.

A scoped BSM MUST have at least one group range, and the first group range in a scoped BSM MUST have the "Admin Scope Zone" bit set. This group range identifies the scope of the BSM. In a scoped IPv4 BSM, the first group range is the range corresponding to the scope of the BSM. In a scoped IPv6 BSM, the first group range may be any group range subject to the general condition that all the group ranges in such a BSM MUST have a mask length of at least 16 and MUST have the same scope ID as the scope of the BSM.

RP mappings may be included in the first group range of a BSM, just as for any other group range. After this group range, other group ranges for which there are RP mappings appear in any order.

The "Admin Scope Zone" bit of all group ranges other than the first SHOULD be set to 0 on origination, and MUST be ignored on receipt.

When an elected BSR is being shut down, it should immediately originate a Bootstrap message listing its current RP-Set, but with the BSR Priority field set to the lowest priority value possible. This will cause the election of a new BSR to happen more quickly.

[3.4.](#) Forwarding Bootstrap Messages

Bootstrap messages originate at the BSR, and are hop-by-hop forwarded by intermediate routers if they pass the Bootstrap Message Processing Checks. When a Bootstrap message is forwarded, it is forwarded out of every multicast-capable interface which has PIM neighbors (including the one over which the message was received). The exception to this is if the interface is an administrative scope boundary for the admin scope zone indicated in the first group address in the Bootstrap message packet.

As an optimization, a router MAY choose not to forward a BSM out of the interface the message was received on if that interface is a point-to-point interface. On interfaces with multiple PIM neighbors, a router SHOULD forward an accepted BSM onto the interface that BSM was received on, but if the number of PIM neighbors on that interface is large, it MAY delay forwarding a BSM onto that interface by a small randomized interval to prevent message implosion. A configuration option MAY be provided to disable forwarding onto the interface a message was received on, but we recommend that the default behavior is to forward onto that interface.

Rationale: A BSM needs to be forwarded onto the interface the message was received on (in addition to the other interfaces) because the routers on a LAN may not have consistent routing information. If three routers on a LAN are A, B, and C, and at router B $RPF(BSR) == A$ and at router C $RPF(BSR) == B$, then router A originally forwards the BSM onto the LAN, but router C will only accept it when router B re-forwards the message onto the LAN. If the underlying routing protocol configuration guarantees that the routers have consistent routing information, then forwarding onto the incoming interface may safely be disabled.

A ZBR constrains all BSMs which are of equal or smaller scope than the configured boundary. That is, the BSMs are not accepted from, originated or forwarded on the interfaces on which the boundary is configured. For IPv6 the check is a comparison between the scope of the first range in the scoped BSM and the scope of the configured boundary. For IPv4, the first range in the scoped BSM is checked to see if it is contained in or is the same as the range of the configured boundary.

[3.5.](#) Unicasting Bootstrap Messages to New and Rebooting Routers

To allow new or rebooting routers to learn the RP-Set quickly, when a Hello message is received from a new neighbor, or a Hello message with a new GenID is received from an existing neighbor, one router on the LAN unicasts a stored copy of the Bootstrap message for each admin scope zone to the new or rebooting router.

The router that does this is the Designated Router (DR) on the LAN, or, if the new or rebooting router is the DR, the router that would be the DR if the new or rebooting router were excluded from the DR election process.

Before unicasting a Bootstrap message in this manner, the DR must wait until it has sent a triggered Hello message on this interface; otherwise, the new neighbor will discard the Bootstrap message.

[3.6.](#) Receiving and Using the RP-Set

The RP-Set maintained by BSR is used by RP-based multicast routing protocols like PIM-SM and BIDIR-PIM. These protocols may obtain RP-Sets from other sources as well. How the final group-to-RP mappings are obtained from these RP-Sets is not part of the BSR specification. In general, the routing protocols need to re-calculate the mappings when any of their RP-Sets change. How such a change is signalled to the routing protocol is also not part of the present specification.

Some group-to-RP mappings in the RP-Set indicate group ranges for which PIM-SM should be used; others indicate group ranges for use with BIDIR-PIM. Routers that only support one of these protocols MUST NOT ignore ranges indicated as being for the other protocol. They MUST NOT treat them as being for the protocol they support.

[4.](#) Message Formats

BSR messages are PIM messages, as defined in [\[1\]](#). The values of the PIM Message Type field for BSR messages are:

[4](#) Bootstrap

[8](#) Candidate-RP-Advertisement

As with all other PIM control messages, BSR messages have IP protocol number 103.

Candidate-RP-Advertisement messages are unicast to a BSR. Usually, Bootstrap messages are multicast with TTL 1 to the ALL-PIM-ROUTERS group, but in some circumstances (described in [section 3.5](#)) Bootstrap messages are unicast to a specific PIM neighbor.

The IP source address used for Candidate-RP-Advertisement messages is a domain-wide reachable address. The IP source address used for Bootstrap messages (regardless of whether they are being originated or forwarded) is the link-local address of the interface on which the message is being sent (that is, the same source address that the router uses for the Hello messages it sends out that interface).

All Bootstrap and Candidate-RP-Advertisement messages SHOULD carry the Router Alert IP option. See [section 6](#) for information about the way in which the Router Alert option is checked by receiving routers.

The IPv4 ALL-PIM-ROUTERS group is 224.0.0.13. The IPv6 ALL-PIM-ROUTERS group is ff02::d.

In this section we use the following terms defined in the PIM-SM specification [1]:

- ```
0 Encoded-Unicast format
0 Encoded-Group format
```

We repeat these here to aid readability.

Encoded-Unicast address

An Encoded-Unicast address takes the following format:

[illegible]

## Addr Family

The PIM address family of the 'Unicast Address' field of this address.

Values of 0-127 are as assigned by the IANA for Internet Address Families in [10]. Values 128-250 are reserved to be assigned by the IANA for PIM-specific Address Families. Values 251 through 255 are designated for private use. As there is no assignment authority for this space, collisions should be expected.

## Encoding Type

The type of encoding used within a specific Address Family. The value `0` is reserved for this field, and represents the native encoding of the Address Family.

## Unicast Address

The unicast address as represented by the given Address Family and Encoding Type.



## Encoded-Group address

Encoded-Group addresses take the following format:

```

 0 1 2 3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+
| Addr Family | Encoding Type |B| Reserved |Z| Mask Len |
+--+
|
| Group multicast Address
+--+...
```

Addr Family  
described above.

Encoding Type  
described above.

[B]IDIR bit  
When set, all BIDIR capable PIM routers will operate the protocol described in [2] for the specified group range.

Reserved  
Transmitted as zero. Ignored upon receipt.

Admin Scope [Z]one  
When set, this bit indicates that this group address range is an administratively scoped range.

Mask Len  
The Mask length field is 8 bits. The value is the number of contiguous one bits left justified used as a mask which, combined with the group address, describes a range of groups. It is less than or equal to the address length in bits for the given Address Family and Encoding Type. If the message is sent for a single group then the Mask length must equal the address length in bits for the given Address Family and Encoding Type. (e.g. 32 for IPv4 native encoding and 128 for IPv6 native encoding).

Group multicast Address  
Contains the group address.

### 4.1. Bootstrap Message Format

A Bootstrap message is divided up into 'semantic fragments' if the original message exceeds the maximum packet size boundaries. Basically, a single Bootstrap message can be sent as multiple packets (semantic



fragments), so long as the fragment tags of all the packets comprising the message is the same.

The format of a single `fragment' is given below:

```

 0 1 2 3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|PIM Ver| Type | Reserved | Checksum |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Fragment Tag | Hash Mask Len | BSR Priority |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| BSR Address (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Group Address 1 (Encoded-Group format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Count 1 | Frag RP Cnt 1 | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Address 1 (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP1 Holdtime | RP1 Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Address 2 (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP2 Holdtime | RP2 Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| . |
| . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Address m (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Rpm Holdtime | Rpm Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Group Address 2 (Encoded-Group format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| . |
| . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Group Address n (Encoded-Group format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Count n | Frag RP Cnt n | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Address 1 (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP1 Holdtime | RP1 Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP Address 2 (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RP2 Holdtime | RP2 Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| . |
| . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

| RP Address m (Encoded-Unicast format) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| RPm Holdtime | RPm Priority | Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

PIM Version, Reserved, Checksum  
Described in [\[1\]](#).

Type  
PIM Message Type. Value is 4 for a Bootstrap message.

Fragment Tag  
A randomly generated number, acts to distinguish the fragments belonging to different Bootstrap messages; fragments belonging to same Bootstrap message carry the same 'Fragment Tag'.

Hash Mask Len  
The length (in bits) of the mask to use in the hash function. For IPv4 we recommend a value of 30. For IPv6 we recommend a value of 126. This field SHOULD be the same for all fragments belonging to the same Bootstrap message.

BSR Priority  
Contains the BSR priority value of the included BSR. This field is considered as a high order byte when comparing BSR addresses. Note that for historical reasons, the highest BSR priority is 255 (the higher the better), whereas the highest RP Priority (see below) is 0 (the lower the better).

BSR Address  
The address of the bootstrap router for the domain. The format for this address is given in the Encoded-Unicast address in [\[1\]](#).

Group Address 1..n  
The group prefix (address and mask) with which the Candidate-RPs are associated. Format described in [\[1\]](#). In a fragment containing admin scope ranges, the first group address in the fragment MUST satisfy the following conditions: it MUST have the Admin Scope bit set; for IPv4 it MUST be the group range for the entire admin scope range (this is the case even if there are no RPs in the RP-Set for the entire admin scope range - in this case the sub-ranges for the RP-Set are specified later in the fragment along with their RPs); for IPv6 the Mask Len MUST be at least 16 and have the scope ID of the admin scope range.

**RP Count 1..n**

The number of Candidate-RP addresses included in the whole Bootstrap message for the corresponding group prefix. A router does not replace its old RP-Set for a given group prefix until/unless it receives 'RP-Count' addresses for that prefix; the addresses could be carried over several fragments. If only part of the RP-Set for a given group prefix was received, the router discards it, without updating that specific group prefix's RP-Set.

**Frag RP Cnt 1..m**

The number of Candidate-RP addresses included in this fragment of the Bootstrap message, for the corresponding group prefix. The 'Frag RP Cnt' field facilitates parsing of the RP-Set for a given group prefix, when carried over more than one fragment.

**RP address 1..m**

The address of the Candidate-RPs, for the corresponding group prefix. The format for these addresses is given in the Encoded-Unicast address in [\[1\]](#).

**RP1..m Holdtime**

The Holdtime (in seconds) for the corresponding RP. This field is copied from the 'Holdtime' field of the associated RP stored at the BSR.

**RP1..m Priority**

The 'Priority' of the corresponding RP and Encoded-Group Address. This field is copied from the 'Priority' field stored at the BSR when receiving a C-RP-Adv message. The highest priority is '0' (i.e. unlike BSR priority, the lower the value of the 'Priority' field, the better). Note that the priority is per RP per Group Address.

Within a Bootstrap message, the BSR Address, all the Group Addresses and all the RP Addresses MUST be of the same address family. In addition, the address family of the fields in the message MUST be the same as the IP source and destination addresses of the packet. This permits maximum implementation flexibility for dual-stack IPv4/IPv6 routers.

#### [4.1.1.](#) Semantic Fragmentation of BSMS

Bootstrap messages may be split over several PIM Bootstrap Message Fragment (BSMF) packets; this is known as semantic fragmentation. It is needed when the BSM would exceed the MTU of the link the packet will be forwarded over.

The packet would be too large because the set of group prefixes and the RPs for each group prefix are too long to fit in one packet. The BSR

will then split the BSM across several BSMF packets; each of these must be a well-formed BSMF packet in its own right.

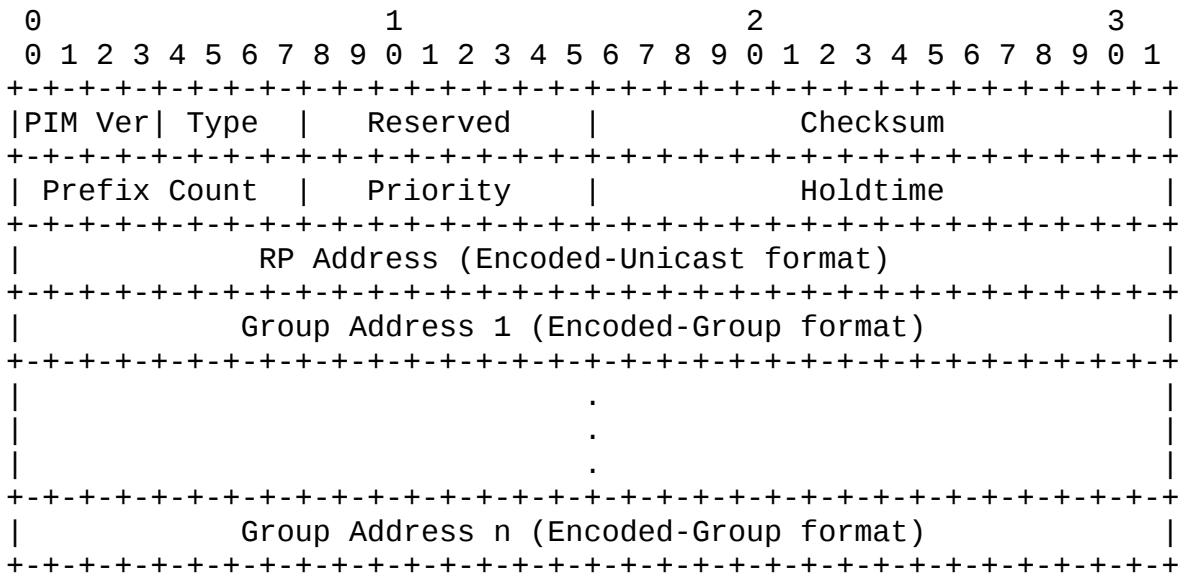
If the BSR can split up the BSM so that different group prefixes (and their RP information) can fit in different fragments, then it should do so. If one of these BSMF packets is then lost, the state from the previous BSM for the group-prefix from the missing packet will be retained. Each fragment that does arrive will update the RP information for the group-prefixes contained in that fragment, and the new group-to-RP mapping for those can be used immediately. The information from the missing fragment will be obtained when the BSM is next transmitted. In this case, whilst the Fragment Tag must be set to the same value for all BSMFs comprising a single BSM, the tag is not actually used by routers receiving the BSM.

If the list of RPs for a single group-prefix is too long to fit in a single BSMF packet, then that information must be split across multiple BSMF packets. In this case, all the BSMF packets comprising the information for that group-prefix must be received before the group-to-RP mapping in use can be modified. This is the purpose of the RP Count field - a router receiving BSMF packets from the same BSM (ie that have the same fragment tag) must wait until the BSMFs providing RP Count RPs for that group-prefix have been received before the new group-to-RP mapping can be used for that group-prefix. If a single BSMF from such a large group-prefix is lost, then that entire group-prefix will have to wait until the next BSM is originated.

Next we need to consider how a BSR would remove group-prefixes from the BSM. A router receiving a set of BSMFs cannot tell if a group-prefix is missing. If it has seen a group-prefix before, it must assume that that group-prefix still exists, and that the BSMF describing it has been lost. It should retain this information for BS\_Timeout. Thus for a BSR to remove a group-prefix from the BSR, it should include that group-prefix, but with a RP Count of zero, and it should resend this information in each BSM for BS\_Timeout.

#### 4.2. Candidate-RP-Advertisement Message Format

Candidate-RP-Advertisement messages are periodically unicast from the C-RPs to the BSR.



PIM Version, Reserved, Checksum  
Described in [1].

Type  
PIM Message Type. Value is 8 for a Candidate-RP-Advertisement message.

Prefix Count  
The number of encoded group addresses included in the message; indicating the group prefixes for which the C-RP is advertising. C-RPs MUST NOT send C-RP-Adv messages with a Prefix Count of '0'.

Priority  
The 'Priority' of the included RP, for the corresponding Encoded-Group Address (if any). The highest priority is '0' (i.e. the lower the value of the 'Priority' field, the higher the priority). This field is stored at the BSR upon receipt along with the RP address and corresponding Encoded-Group Address.

Holdtime  
The amount of time (in seconds) the advertisement is valid. This field allows advertisements to be aged out. This field should be set to 2.5 times C\_RP\_Adv\_Period.

**RP Address**

The address of the interface to advertise as a Candidate RP. The format for this address is given in the Encoded-Unicast address in [1].

**Group Address-1..n**

The group prefixes for which the C-RP is advertising. Format described in Encoded-Group-Address in [1].

Within a Candidate-RP-Advertisement message, the RP Address and all the Group Addresses MUST be of the same address family. In addition, the address family of the fields in the message MUST be the same as the IP source and destination addresses of the packet. This permits maximum implementation flexibility for dual-stack IPv4/IPv6 routers.

**5. Timers and Timer Values**

Timer Name: Bootstrap Timer (BST(Z))

| Value Name       | Value                | Explanation                                                                     |
|------------------|----------------------|---------------------------------------------------------------------------------|
| BS_Period        | Default: 60 seconds  | Periodic interval with which BSMs are normally originated                       |
| BS_Timeout       | Default: 130 seconds | Interval after which a BSR is timed out if no BSM is received from that BSR     |
| BS_Rand_Override | see below            | Randomized interval used to reduce control message overhead during BSR election |

Note that BS\_Timeout MUST be larger than BS\_Period, even if their values are changed from the defaults. We recommend that BS\_Timeout is set to 2 times BS\_Period plus 10 seconds.

BS\_Rand\_Override is calculated using the following pseudocode, in which all values are in units of seconds. The values of BS\_Rand\_Override generated by this pseudocode are between 5 and 23 seconds, with smaller values generated if the C-BSR has a high bootstrap weight, and larger values generated if the C-BSR has a low bootstrap weight.

$$\text{BS\_Rand\_Override} = 5 + \text{priorityDelay} + \text{addrDelay}$$

where priorityDelay is given by:

$$\text{priorityDelay} = 2 * \log_2(1 + \text{bestPriority} - \text{myPriority})$$

and addrDelay is given by the following for IPv4:

```
if (bestPriority == myPriority) {
 addrDelay = log_2(1 + bestAddr - myAddr) / 16
} else {
 addrDelay = 2 - (myAddr / 2^31)
}
```

and addrDelay is given by the following for IPv6:

```
if (bestPriority == myPriority) {
 addrDelay = log_2(1 + bestAddr - myAddr) / 64
} else {
 addrDelay = 2 - (myAddr / 2^127)
}
```

and bestPriority is given by:

$$\text{bestPriority} = \max(\text{storedPriority}, \text{myPriority})$$

and bestAddr is given by:

$$\text{bestAddr} = \max(\text{storedAddr}, \text{myAddr})$$

and where myAddr is the Candidate-BSR's address, storedAddr is the stored BSR's address, myPriority is the Candidate-BSR's configured priority, and storedPriority is the stored BSR's priority.



Timer Name: Scope Zone Expiry Timer (SZT(Z))

| Value Name | Value                 | Explanation                                                                              |
|------------|-----------------------|------------------------------------------------------------------------------------------|
| SZ_Timeout | Default: 1300 seconds | Interval after which a scope zone is timed out if no BSM is received for that scope zone |

Note that SZ\_Timeout MUST be larger than BS\_Timeout, even if their values are changed from the defaults. We recommend that SZ\_Timeout is set to 10 times BS\_Timeout.

Timer Name: Group-to-C-RP mapping Expiry Timer (CGET(M,Z))

| Value Name           | Value        | Explanation                    |
|----------------------|--------------|--------------------------------|
| C-RP Mapping Timeout | from message | Holdtime from C-RP-Adv message |

Timer Name: Group-to-RP mapping Expiry Timer (GET(M,Z))

| Value Name         | Value        | Explanation       |
|--------------------|--------------|-------------------|
| RP Mapping Timeout | from message | Holdtime from BSM |

Timer Name: C-RP Advertisement Timer (CRPT)

| Value Name      | Value               | Explanation                                                      |
|-----------------|---------------------|------------------------------------------------------------------|
| C_RP_Adv_Period | Default: 60 seconds | Periodic interval with which C-RP-Adv messages are sent to a BSR |

## [6.](#) Security Considerations

### [6.1.](#) Possible Threats

Threats affecting the PIM BSR mechanism are primarily of two forms: denial of service attacks, and traffic diversion attacks. An attacker that subverts the BSR mechanism can prevent multicast traffic from reaching the intended recipients, can divert multicast traffic to a place where they can monitor it, and can potentially flood third parties with traffic.

Traffic can be prevented from reaching the intended recipients by one of two mechanisms:

- o Subverting a BSM, and specifying RPs that won't actually forward traffic.
- o Registering with the BSR as a C-RP, and then not forwarding traffic.

Traffic can be diverted to a place where it can be monitored by both of the above mechanisms; in this case the RPs would forward the traffic, but are located so as to aid monitoring or man-in-the-middle attacks on the multicast traffic.

A third party can be flooded by either of the above two mechanisms by specifying the third party as the RP, and register-encapsulated traffic will then be forwarded to them.

### [6.2.](#) Limiting Third-Party DoS Attacks

The third party DoS attack above can be greatly reduced if PIM routers acting as DR do not continue to forward Register traffic to the RP in the presence of ICMP Protocol Unreachable or ICMP Host Unreachable responses. If a PIM router sending Register packets to an RP receives one of these responses to a data packet it has sent, it should rate-limit the transmission of future Register packets to that RP for a short period of time.

As this does not affect interoperability, the precise details are left to the implementor to decide. However we note that a router implementing such rate limiting must only do so if the ICMP packet correctly echoes part of a Register packet that was sent to the RP. If this check were not made, then simply sending ICMP Unreachable packets to the DR with the source address of the RP spoofed would be sufficient to cause a denial-of-service attack on the multicast traffic originating from that DR.

### [6.3.](#) Bootstrap Message Security

If a legitimate PIM router is compromised, there is little any security mechanism can do to prevent that router subverting PIM traffic in that domain. However we recommend that implementors provide a mechanism whereby a PIM router using the BSR mechanisms can be configured with the IP addresses of valid BSR routers, and that any Bootstrap message from any other BSR should then be dropped and logged as a security issue. We also recommend that this not be enabled by default, as it makes the initial configuration of a PIM domain problematic - it is the sort of feature that might be enabled once the configuration of a domain has stabilized.

The primary security requirement for BSR (as for PIM) is that it is possible to prevent hosts that are not legitimate PIM routers, either within or outside the domain, from subverting the BSR mechanism.

The Bootstrap Message Processing Checks prevent a router from accepting a Bootstrap message from outside of the PIM Domain, as the source address on Bootstrap messages must be an immediate PIM neighbor. There is however a small window of time after a reboot where a PIM router will accept a bad Bootstrap message unicast from an immediate neighbor, and it might be possible to unicast a Bootstrap message to a router during this interval from outside the domain, using the spoofed source address of a neighbor. This can be prevented if PMBRs perform source-address filtering to prevent packets entering the PIM domain with IP source addresses that are infrastructure addresses in the PIM domain.

The principal threat to Bootstrap message security comes from hosts within the PIM domain that attempt to subvert the BSR mechanism. They may be able to do this by sending PIM messages to their local router, or by unicasting a Bootstrap message to another PIM router during the brief interval after it has restarted.

#### [6.3.1.](#) Rejecting Unicast Bootstrap Messages

All Bootstrap messages SHOULD carry the Router Alert IP option. If a PIM router receives a Bootstrap message that does not carry the Router Alert option, it SHOULD drop it (a configuration option should also be provided to disable this check on a per-interface basis for backward compatibility with older PIM routers). The Router Alert option allows a PIM router to perform checks on unicast packets it would otherwise blindly forward. All PIM routers should check that packets with Router Alert that are not destined for the router itself are not PIM Bootstrap messages. Any such packets should be dropped and logged as a possible security issue - it is never acceptable for a PIM Bootstrap message to travel multiple IP hops.

### [6.3.2.](#) Rejecting Bootstrap Messages from Invalid Neighbors

Most hosts that are likely to attempt to subvert PIM BSR are likely to be located on leaf subnets. We recommend that implementors provide a configuration option that specifies an interface is a leaf subnet, and that no PIM packets are accepted on such interfaces.

On multi-access subnets with multiple PIM routers and hosts that are not trusted, we recommend that IPsec AH is used to protect communication between PIM routers, and that such routers are configured to drop and log communication attempts from any host that do not pass the authentication check. When all the PIM routers are under the same administrative control, this authentication may use a configured shared secret. The securing of interactions between PIM neighbors is discussed in more detail in the Security Considerations section of [[1](#)], and so we do not discuss the details further here. The same security mechanisms that can be used to secure PIM Join, Prune and Assert messages should also be used to secure Bootstrap messages.

### [6.4.](#) Candidate-RP-Advertisement Message Security

Even if it is not possible to subvert Bootstrap messages, an attacker might be able to perform most of the same attacks by simply sending C-RP-Adv messages to the BSR specifying the attacker's choice of RPs. Thus it is necessary to control the sending of C-RP-Adv messages in essentially the same ways that we control Bootstrap messages. However, C-RP-Adv messages are unicast and normally travel multiple hops, so controlling them is more difficult.

#### [6.4.1.](#) Non-Cryptographic Security of C-RP-Adv Messages

We specify that C-RP-Adv messages SHOULD also carry the Router Alert IP option, and that the BSR SHOULD by default drop and log C-RP-Adv messages that do not carry this option. Setting Router Alert on these packets is practical because the rate of C-RP-Adv messages should be very low, so the extra load on routers forwarding these packets will be insignificant. PIM routers forwarding such a packet may then be capable of checking whether the packet came from a valid PIM neighbor, although note that such checks are only possible if the unicast and multicast topologies in the network are congruent. If this is not the case, it is legitimate to receive a C-RP-Adv message from a router which is not a valid PIM neighbor, and therefore in this situation a PIM router MUST NOT drop C-RP-Adv messages that do not come from a valid PIM neighbor.

If the unicast and multicast topologies are known to be congruent, the following checks should be made. On interfaces that are configured to be leaf subnets, all C-RP-Adv messages should be dropped. On multi-access subnets with multiple PIM routers and hosts that are not trusted,

the router can at least check that the source MAC address is that of a valid PIM neighbor. PMBRs should ensure that no C-RP-Adv messages enter the domain from an external neighbor.

#### [6.4.2.](#) Cryptographic Security of C-RP-Adv Messages

For true security, we recommend that all C-RPs are configured to use IPsec authentication. The authentication process for a C-RP-Adv message between a C-RP and the BSR is identical to the authentication process for PIM Register messages between a DR and the relevant RP, except that there will normally be fewer C-RPs in a domain than there are DRs, so key management is a little simpler. We do not describe the details of this process further here, but refer to the Security Considerations section of [1]. Note that the use of cryptographic security for C-RP-Adv messages does not remove the need for the non-cryptographic mechanisms, as explained below.

#### [6.5.](#) Denial of Service using IPsec

An additional concern is that of Denial-of-Service attacks caused by sending high volumes of Bootstrap messages or C-RP-Adv messages with invalid IPsec authentication information. It is possible that these messages could overwhelm the CPU resources of the recipient.

The non-cryptographic security mechanisms above prevent unicast Bootstrap messages from traveling multiple hops, and constrain who can originate such messages. However, it is obviously important that PIM messages that are required to have Router Alert checked are checked for this option before the IPsec AH is checked. Thus the remaining vulnerability primarily exists for hosts on multi-access subnets containing more than one PIM router. A PIM router receiving PIM packets with Router Alert set from such a subnet should already be checking that the source MAC address is that of a valid PIM neighbor, but this is hardly strong security. In addition, we recommend that rate-limiting mechanisms can be configured, to be applied to the forwarding of unicast PIM packets containing Router Alert options. The rate-limiter MUST independently rate-limit different types of PIM packets - for example a flood of C-RP-Adv messages MUST NOT cause a rate limiter to drop low-rate Bootstrap messages. Such a rate-limiter might itself be used to cause a denial of service attack by causing valid packets to be dropped, but in practice this is more likely to constrain bad PIM messages close to their origin. In addition, the rate limiter will prevent attacks on PIM from affecting other activity on the destination router, such as unicast routing.

## 7. Contributors

Bill Fenner, Mark Handley, Roger Kermode and David Thaler have contributed greatly to this draft. They were authors of this draft up to version 03. Most of the current text is identical to 03.

## 8. Acknowledgments

PIM-SM was designed over many years by a large group of people, including ideas from Deborah Estrin, Dino Farinacci, Ahmed Helmy, Steve Deering, Van Jacobson, C. Liu, Puneet Sharma, Liming Wei, Tom Pusateri, Tony Ballardie, Scott Brim, Jon Crowcroft, Paul Francis, Joel Halpern, Horst Hodel, Polly Huang, Stephen Ostrowski, Lixia Zhang, Girish Chandranmenon, Pavlin Radoslavov, John Zwiebel, Isidor Kouvelas and Hugh Holbrook. This BSR specification draws heavily on text from [RFC 2362](#).

Many members of the PIM Working Group have contributed comments and corrections for this document, including Christopher Thomas Brown, Ardas Cilingiroglu, Murthy Esakonu, Venugopal Hemige, Prashant Jhingran, Rishabh Parekh and Katta Sambasivarao.

## 9. IANA Considerations

This document has no actions for IANA.

## 10. Normative References

- [1] W. Fenner, M. Handley, H. Holbrook, I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", Internet Draft [draft-ietf-pim-sm-v2-new-11.txt](#)
- [2] M. Handley, I. Kouvelas, T. Speakman, L. Vicisano, "Bi-directional Protocol Independent Multicast (BIDIR-PIM)", Internet Draft [draft-ietf-pim-bidir-07.txt](#)
- [3] D. Meyer, "Administratively Scoped IP Multicast", [RFC 2365](#), Jul 1998.
- [4] S. Deering, B. Haberman, T. Jinmei, E. Nordmark, B. Zill, "IPv6 Scoped Address Architecture", [RFC 4007](#), Mar 2005.
- [5] R. Hinden, S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", [RFC 3513](#), Apr 2003.
- [6] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), Mar 1997.

## 11. Informative References

- [7] D. Estrin et al., "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification", [RFC 2362](#), June 1998 (now obsolete).
- [8] D. Kim, D. Meyer, H. Kilmer, D. Farinacci, "Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", [RFC 3446](#), Jan 2003.
- [9] D. Farinacci, Y. Cai, "Anycast-RP using PIM", Internet Draft [draft-ietf-pim-anycast-rp-04.txt](#)
- [10] IANA, "Address Family Numbers", linked from <http://www.iana.org/numbers.html>

### Authors' Addresses

Nidhi Bhaskar  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
USA  
nbhaskar@cisco.com

Alexander Gall  
SWITCH  
Limmatquai 138  
P.O. Box  
CH-8021 Zurich  
Switzerland  
gall@switch.ch

James Lingard  
Data Connection Ltd  
100 Church Street  
Enfield  
EN2 6BQ  
United Kingdom  
james@lingard.com

Stig Venaas  
UNINETT  
NO-7465 Trondheim  
Norway  
venaas@uninett.no

## Copyright Statement

Copyright (C) The Internet Society (2005). 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 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).