

Network Working Group  
Request for Comments: 4609  
Category: Informational

P. Savola  
CSC/FUNET  
R. Lehtonen  
TeliaSonera  
D. Meyer  
August 2006

Protocol Independent Multicast - Sparse Mode (PIM-SM)  
Multicast Routing Security Issues and Enhancements

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo describes security threats for the larger (intra-domain or inter-domain) multicast routing infrastructures. Only Protocol Independent Multicast - Sparse Mode (PIM-SM) is analyzed, in its three main operational modes: the traditional Any-Source Multicast (ASM) model, the source-specific multicast (SSM) model, and the ASM model enhanced by the Embedded Rendezvous Point (Embedded-RP) group-to-RP mapping mechanism. This memo also describes enhancements to the protocol operations that mitigate the identified threats.

## Table of Contents

1. Introduction .....	3
2. Terminology .....	4
3. Threats to Multicast Routing .....	4
3.1. Receiver-Based Attacks .....	5
3.1.1. Joins to Different Groups (Join Flooding) .....	5
3.2. Source-Based Attacks .....	7
3.2.1. Sending Multicast to Empty Groups (Data Flooding) ...	7
3.2.2. Disturbing Existing Group by Sending to It (Group Integrity Violation).....	8
3.3. Aggravating Factors to the Threats .....	9
3.3.1. Distant RP/Source Problem .....	9
3.3.2. No Receiver Information in PIM Joins .....	10
4. Threat Analysis .....	10
4.1. Summary of the Threats .....	10
4.2. Enhancements for Threat Mitigation .....	10
5. PIM Security Enhancements .....	11
5.1. Remote Routability Signalling .....	11
5.2. Rate-Limiting Possibilities .....	12
5.3. Specific Rate-limiting Suggestions .....	14
5.3.1. Group Management Protocol Rate-Limiter .....	14
5.3.2. Source Transmission Rate-Limiter .....	14
5.3.3. PIM Signalling Rate-Limiter .....	15
5.3.4. Unicast-Decapsulation Rate-Limiter .....	15
5.3.5. PIM Register Rate-Limiter .....	15
5.3.6. MSDP Source-Active Rate-Limiter .....	16
5.4. Passive Mode for PIM .....	16
6. Security Considerations .....	16
7. Acknowledgements .....	17
8. References .....	17
8.1. Normative References .....	17
8.2. Informative References .....	17
Appendix A. RPF Considers Interface, Not Neighbor .....	19
Appendix B. Return Routability Extensions .....	20
B.1. Sending PIM-Prune Messages Down the Tree .....	20
B.2. Analysing Multicast Group Traffic at DR .....	21
B.3. Comparison of the Above Approaches .....	21

## 1. Introduction

This document describes security threats to the Protocol Independent Multicast - Sparse Mode (PIM-SM) multicast routing infrastructures and suggests ways to make these architectures more resistant to the described threats.

Only attacks that have an effect on the multicast routing infrastructures (whether intra- or inter-domain) are considered.

"On-link" attacks where the hosts specifically target the Designated Router (DR) or other routers on the link, or where hosts disrupt other hosts on the same link, possibly using group management protocols, are discussed elsewhere (e.g., [10] and [12]). These attacks are not discussed further in this document.

Similar to unicast, the multicast payloads may need end-to-end security. Security mechanisms to provide confidentiality, authentication, and integrity are described in other documents (e.g., [9]). Attacks that these security mechanisms protect against are not discussed further in this document.

PIM builds on a model where Reverse Path Forwarding (RPF) checking is, among other things, used to ensure loop-free properties of the multicast distribution trees. As a side effect, this limits the impact of an attacker using a forged source address, which is often used as a component in unicast-based attacks. However, a host can still spoof an address within the same subnet, or spoof the source of a unicast-encapsulated PIM Register message, which a host may send on its own.

We consider PIM-SM [1] operating in the traditional Any Source Multicast (ASM) model (including the use of Multicast Source Discovery Protocol (MSDP) [2] for source discovery), in Source-Specific Multicast [3] (SSM) model, and the Embedded-RP [4] group-to-RP mapping mechanism in ASM model. Bidirectional-PIM [15] is typically deployed only in intra-domain and is similar to ASM but without register messages. Bidirectional-PIM is not finished as of this writing, and its considerations are not discussed further in this document.

## 2. Terminology

### ASM

"ASM" [6] is used to refer to the traditional Any Source Multicast model with multiple PIM domains and a signalling mechanism (MSDP) to exchange information about active sources between them.

### SSM

"SSM" [7] is used to refer to Source-Specific Multicast.

### SSM channel

SSM channel (S, G) identifies the multicast delivery tree associated with a source address S and a SSM destination address G.

### Embedded-RP

"Embedded-RP" refers to the ASM model where the Embedded-RP mapping mechanism is used to find the Rendezvous Point (RP) for a group, and MSDP is not used.

### Target Router

"Target Router" is used to refer to either the RP processing a packet (ASM or Embedded-RP) or the DR that is receiving (Source, Group) (or (S,G)) joins (in all models).

## 3. Threats to Multicast Routing

We make the broad assumption that the multicast routing networks are reasonably trusted. That is, we assume that the multicast routers themselves are well-behaved, in the same sense that unicast routers are expected to behave well. While this assumption is not entirely correct, it simplifies the analysis of threat models. The threats caused by misbehaving multicast routers (including fake multicast routers) are not considered in this memo; the generic threat model would be similar to [5]. RP discovery mechanisms like Bootstrap Router (BSR) and Auto-RP are also considered out of scope.

As the threats described in this memo are mainly Denial-of-Service (DoS) attacks, it may be useful to note that the attackers will try to find a scarce resource anywhere in the control or data plane, as described in [5].

There are multiple threats relating to the use of host-to-router signalling protocols -- such as Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) -- but these are outside the scope of this memo.

PIM-SM can be abused in the cases where RPF checks are not applicable (in particular, in the stub LAN networks), as spoofing the on-link traffic is very simple. For example, a host could get elected to become DR for the subnet, but not perform any of its functions. A host can also easily make PIM routers on the link stop forwarding multicast by sending PIM Assert messages. This implies that a willful attacker will be able to circumvent many of the potential rate-limiting functions performed at the DR (as one can always send the messages himself). The PIM-SM specification, however, states that these messages should only be accepted from known PIM neighbors; if this is performed, the hosts would first have to establish a PIM adjacency with the router. Typically, adjacencies are formed with anyone on the link, so a willful attacker would have a high probability of success in forming a protocol adjacency. These are described at some length in [1], but are also considered out of the scope of this memo.

### 3.1. Receiver-Based Attacks

These attacks are often referred to as control plane attacks, and the aim of the attacker is usually to increase the amount of multicast state information in routers above a manageable level.

#### 3.1.1. Joins to Different Groups (Join Flooding)

Join flooding occurs when a host tries to join, once or a couple of times, to a group or an SSM channel, and the DR generates a PIM Join to the Target Router. The group/SSM channel or the Target Router may or may not exist.

An example of this is a host trying to join different, non-existent groups at a very rapid pace, trying to overload the routers on the path with an excessive amount of (\*S,G) state (also referred to as "PIM State"), or the Target Router with an excessive number of packets.

Note that even if a host joins to a group multiple times, the DR only sends one PIM Join message, without waiting for any acknowledgement; the next message is only sent after the PIM Join timer expires or the state changes at the DR.

This kind of joining causes PIM state to be created, but this state is relatively short-lived (260 seconds by default, which is the default time that the state is active at DR in the absence of IGMP/MLD Reports/Leaves). Note that the host can join a number of different ASM groups or SSM channels with only one IGMPv3 [11] or MLDv2 [12] Report as the protocol allows multiple sources to be included in the same message, resulting in multiple PIM Joins from one IGMPv3/MLDv2 message.

However, even short-lived state may be harmful when the intent is to cause as much state as possible. The host can continue to send IGMP/MLD Reports to these groups to make the state attack more long-lived. This results in:

- o ASM: An (\*,G) join is sent to an intra-domain RP, causing state on that path; in turn, that RP joins to the DR of the source if the source is active. If the source address was specified by the host in the IGMPv3/MLDv2 Report, a (S,G) Join is sent directly to the DR of the source, as with SSM, below.
- o SSM: An (S,G) join is sent inter-domain to the DR of the source S, causing state on that path. If the source S does not exist, the join goes to the closest router using longest prefix matching on the path to S as possible.
- o Embedded-RP: An (\*,G) join is sent towards an inter/intra-domain RP embedded in the group G, causing state on that path. If the RP does not exist, the join goes to the router that is closest to the RP address. Similarly, an explicit (S,G) join goes to the DR, as with SSM above.

That is, SSM and Embedded-RP always enable "inter-domain" state creation. ASM defaults to intra-domain, but can be used for inter-domain state creation as well.

If the source or RP (only in case of Embedded-RP) does not exist, the multicast routing protocol does not have any means to remove the distribution tree if the joining host remains active. The worst case attack could be a host remaining active to many different groups (containing either imaginary source or RP). Please note that the imaginary RP problem is related to only Embedded-RP, where the RP address is extracted from the group address, G.

For example, if the host is able to generate 100 IGMPv3 (S,G) joins a second, each carrying 10 sources, the amount of state after 260 seconds would be 260 000 state entries -- and 100 packets per second is still a rather easily achievable number.

### 3.2. Source-Based Attacks

These attacks are often referred to as "data plane" attacks; however, with traditional ASM and MSDP, these also include an MSDP control plane threat.

#### 3.2.1. Sending Multicast to Empty Groups (Data Flooding)

Data flooding occurs when a host sends data packets to a multicast group or SSM channel for which there are no real subscribers.

Note that since register encapsulation is not subject to RPF checks, the hosts can also craft and send these packets themselves, also spoofing the source address of the register messages unless ingress filtering [13] has been deployed [14]. That is, as the initial data registering is not subject to the same RPF checks as many other multicast routing procedures, making control decisions based on that data leads to many potential threats.

Examples of this threat are a virus/worm trying to propagate to multicast addresses, an attacker trying to crash routers with excessive MSDP state, or an attacker wishing to overload the RP with encapsulated packets of different groups. This results in:

- o ASM: The DR register-encapsulates the packets in Register messages to the intra-domain RP, which may join to the source and issue a Register-Stop, but which continues to get the data. A notification about the active source is sent (unless the group or source is configured to be local) inter-domain with MSDP and propagated globally.
- o SSM: The DR receives the data, but the data does not propagate from the DR unless someone joins the (S,G) channel.
- o Embedded-RP: The DR register-encapsulates the packets to the intra/inter-domain RP, which may join to the source and issue a Register-Stop. Data continues to be encapsulated if different groups are used.

This yields many potential attacks, especially if at least parts of the multicast forwarding functions are implemented on a "slow" path or CPU in the routers:

- o The MSDP control plane traffic generated can cause a significant amount of control and data traffic, which may overload the routers receiving it. A thorough analysis of MSDP vulnerabilities can be found in [16] and is only related to the ASM. However, this is the most serious threat at the moment, because MSDP will flood the

multicast group information to all multicast domains in Internet including the multicast packet encapsulated to MSDP source-active message. This creates a lot of data and state to be shared by all multicast-enabled routers, and if the source remains active, the flooding will be repeated every 60 seconds by default.

- o As a large amount of data is forwarded on the multicast tree, if multicast forwarding is performed on CPU, it may be a serious performance bottleneck, and a way to perform DoS on the path. Similarly, the DR must always be capable of processing (and discarding, if necessary) the multicast packets received from the source. These are potentially present in every model.
- o If the encapsulation is performed on software, it may be a performance bottleneck, and a way to perform DoS on the DR. Similarly, if the decapsulation is performed on software, it may be a performance bottleneck, and a way to perform DoS on the RP. Note: the decapsulator may know (based on access configuration, a rate limit, or something else) that it doesn't need to decapsulate the packet, avoiding bottlenecks. These threats are related to ASM and Embedded-RP.

### 3.2.2. Disturbing Existing Group by Sending to It (Group Integrity Violation)

Group integrity violation occurs when a host sends packets to a group or SSM channel, which already exists, to disturb the users of the existing group/SSM channel.

The SSM service model prevents injection of packets to (S,G) channels, avoiding this problem. However, if the source address can be spoofed to be a topologically-correct address, it's possible to get the packet into the distribution tree. Typically only hosts that are on-link with the source are able to perform this, so it is not really relevant in the scope of this memo.

With ASM and Embedded-RP, sources can inject forged traffic through RPs, which provide the source discovery for the group. The RPs send the traffic over the shared tree towards receivers (routers with (\*,G) state). DR then forwards the forged traffic to receivers unless the legitimate recipients are able to filter out unwanted sources, e.g., using Multicast Source Filters (MSF) API [8]. Typically this is not used or supported by the applications using these protocols.



Note that with ASM and Embedded-RP, the RP may exert some form of control on who can send to a group, as the first packets are register-encapsulated in register packets to the RP. If the RP drops the packet based on an access list, a rate limit, or something else,

it doesn't get injected to an existing group. However, if the DR has existing (\*,G) state, the data will also be forwarded on those interfaces.

With ASM, this "source control" is distributed across all the PIM domains, which significantly decreases its applicability. Embedded-RP enables easier control because source discovery is done through a single RP per group.

As a result, in addition to possible local disturbance, the RP decapsulates the register packets and forwards them to the receivers in the multicast distribution tree, resulting in an integrity violation.

### 3.3. Aggravating Factors to the Threats

This section describes a few factors that aggravate the threats described in Sections 3.1 and 3.2. These could also be viewed as individual threats on their own.

#### 3.3.1. Distant RP/Source Problem

In the shared tree model, if the RP or a source is distant (topologically), then joins will travel to the distant RP or source and keep the state information in the path active, even if the data is being delivered locally.

Note that this problem will be exacerbated if the RP/source space is global; if a router is registering to a RP/source that is not in the local domain (say, fielded by the site's direct provider), then the routing domain is flat.

Also note that PIM assumes that the addresses used in PIM messages are valid. However, there is no way to ensure this, and using non-existent S or G in (\*,G) or (S,G) messages will cause the signalling to be set up, even though one cannot reach the address.

This will be analyzed at more length in Section 5.1.

3.3.2. No Receiver Information in PIM Joins

Only DRs, which are directly connected to receivers, know the exact receiver information (e.g., IP address). PIM does not forward that information further in the multicast distribution tree. Therefore, individual routers (e.g., domain edge routers) are not able to make policy decisions on who can be connected to the distribution tree.

4. Threat Analysis

4.1. Summary of the Threats

Trying to summarize the severity of the major classes of threats with respect to each multicast usage model, we have a matrix of resistance to different kinds of threats:

	Forged Join	Being a Source	Group Integrity
ASM	bad 1)	very bad	bad/mediocre
SSM	bad	very good	very good
Embedded-RP	bad 1),2)	good/mediocre 3)	good

Notes:

- 1) In ASM, the host can directly join also (S,G) groups with IGMPv3/MLDv2 and thus have the same characteristics as SSM (also allows inter-domain state to be created).
- 2) allows inter-domain shared state to be created.
- 3) Embedded-RP allows a host to determine the RP for a given group (or set of groups), which in turn allows that host to mount a PIM register attack. In this case, the host can mount the attack without implementing any of the PIM register machinery.

4.2. Enhancements for Threat Mitigation

There are several desirable actions ("requirements") that could be considered to mitigate these threats; these are listed below. A few more concrete suggestions are presented later in the section.

- o Inter-domain MSDP (ASM) should be retired to avoid attacks; or, if this is not reasonable, the DRs should rate-limit the register encapsulation (note that the hosts can circumvent this). More

importantly, the RPs should rate-limit the register decapsulation especially from different sources, or MSDP must rate-limit the MSDP data generation for new sources.

- o DRs should rate-limit PIM Joins and Prunes somehow; there are multiple ways this should be considered (i.e., depending on which variables are taken into consideration).
- o DRs could rate-limit register encapsulation somehow; there are multiple ways to perform this. Note that the hosts can avoid this by performing the register encapsulation themselves if so inclined.
- o RPs could rate-limit register decapsulation somehow; there are multiple ways to perform this. Note that if the source of the unicast packets is spoofed by the host, this may have an effect on how (for example) rate-limiters behave.
- o RPs should rate-limit the MSDP SA messages coming from MSDP peers.
- o RPs could limit or even disable the SA cache size. However, this could have negative effects on normal operation.
- o RPs should provide good interfaces to reject packets that are not interesting; for example, if an Embedded-RP group is not configured to be allowed in the RP, the register encapsulated packets would not even be decapsulated.
- o DRs could rate-limit the multicast traffic somehow to reduce the disturbing possibilities; there are multiple possibilities how exactly this should be considered.
- o DRs should rate-limit the number of groups/SSM channels that can be created by a given source, S.

## 5. PIM Security Enhancements

This section includes more in-depth description of the above-mentioned functions for rate-limiting, etc., as well as a description of the remote routability signalling issue.

### 5.1. Remote Routability Signalling

As described in Section 3.3.1, non-existent DRs or RPs may cause some problems when setting up multicast state. There seem to be a couple of different approaches to mitigate this, especially if rate-limiting is not extensively deployed.

With ASM and Embedded-RP, Register message delivery could be ensured somehow. For example:

- 1) At the very least, receiving an ICMP unreachable message (of any flavor) should cause the DR to stop the Register packets, as the RP will not be receiving them anyway. (However, one should note that easy spoofing of such ICMP messages could cause a DoS on legitimate traffic.)
- 2) An additional method could be implementing a timer on the DRs so that unless nothing is heard back from the RP within a defined time period, the flow of Register messages would stop. (Currently, the RPs are not required to answer back, unless they want to join to the source.)
- 3) An extreme case would be performing some form of return routability check prior to starting the register messages: first, a packet would be sent to the RP, testing its existence and willingness to serve, and also proving to the RP that the sender of the "bubble" and the sender of the registers are the same and the source address is not forged. (That is, the RP would insert a cookie in the bubble, and it would have to be present in the register message.)

It would be desirable to have some kind of state management for PIM Joins (and other messages) as well; for example, a "Join Ack" that could be used to ensure that the path to the source/RP actually exists. However, this is very difficult, if not impossible, with the current architecture: PIM messages are sent hop-by-hop, and there is not enough information to trace back the replies, for example, to notify the routers in the middle to release the corresponding state or to notify the DR that the path did not exist.

Appendix B discusses this receiver-based remote routability signalling in more detail.

## 5.2. Rate-Limiting Possibilities

There seem to be many ways to implement rate-limiting (for signalling, data encapsulation, and multicast traffic) at the DRs or RPs. The best approach likely depends on the threat model; for example, factors in the evaluation may include:

- o Whether the host is willfully malicious, uncontrolled (e.g., virus/worm), or a regular user just doing something wrong.

- o Whether the threat is aimed towards a single group, a single RP handling the group, or the (multicast) routing infrastructure in general.
- o Whether the host on a subnet is spoofing its address (but still as one that fulfills the RPF checks of the DR).
- o Whether the host may generate the PIM join (and similar) messages itself to avoid rate-limiters at the DR, if possible.
- o Whether unicast RPF checks are applied on the link (i.e., whether the host can send register-encapsulated register-messages on its own).
- o Whether blocking the misbehaving host on a subnet is allowed to also block other, legitimate hosts on the same subnet.
- o Whether these mechanisms would cause false positives on links with only properly working hosts if many of them are receivers or senders.

As should be obvious, there are many different scenarios here that seem to call for different kinds of solutions.

For example, the rate-limiting could be performed based on:

1. multicast address, or the RP where the multicast address maps to
2. source address
3. the (source address, multicast address) pair (or the RP that maps to the multicast address)
4. data rate, in case of rate-limiting the source
5. everything (multicast groups and sources would not be distinguished at all)

In the above, we assume that rate-limiting would be performed per-interface (on DRs) if a more fine-grained filter is not being used.

It should be noted that some of the rate-limiting functions can be used as a tool for DoS against legitimate multicast users. Therefore, several parameters for rate-limiting should be used to prevent such operation.

### 5.3. Specific Rate-limiting Suggestions

These suggestions take two forms: limiters designed to be run on all the edge networks, preventing or limiting an attack in the first place, and the limiters designed to be run at the border of PIM domains or at the RPs, which should provide protection in case edge-based limiting fails or was not implemented, or when additional control is required.

Almost none of the suggested rate-limiters take legitimate users into account. That is, being able to allow some hosts on a link to transmit/receive, while disallowing others, is very challenging to do right, because the attackers can easily circumvent such systems. Therefore, the intent is to limit the damage to only one link, one DR, or one RP -- and avoid the more global effects on the Internet multicast architecture.

Also, it is possible to perform white-listing of groups, sources, or (S,G) pairs from the rate-limiters so that packets related to these are not counted towards the limits. This is useful for handling an aggressive but legitimate source without modifying the limiting parameters for all the traffic, for example.

#### 5.3.1. Group Management Protocol Rate-Limiter

A Group Management Protocol rate-limiter is a token-bucket-based rate-limiter to all Group Management Protocols (IGMP, MLD) that would limit the average rate of accepted groups or sources on the specific interface, with a bucket of depth of G\_DEPTH, refilling at G\_RATE tokens per second. Example values could be G\_RATE=1 and G\_DEPTH=20. Note that, e.g., an IGMPv3 join with two included sources for one group would count as two groups/sources.

This would be the first-order defense against state-creation attacks from the hosts. However, as it cannot be guaranteed that all the routers would implement something like this, other kinds of protections would be useful as well. This harms legitimate receivers on the same link as an attacker.

#### 5.3.2. Source Transmission Rate-Limiter

A source transmission rate-limiter is a token-bucket-based rate-limiter that would limit the multicast data transmission (excluding link-local groups) on a specific interface with a bucket of depth of GSEND\_DEPTH, refilling at GSEND\_RATE tokens per second. Example values could be GSEND\_RATE=10 and GSEND\_DEPTH=20.

This would be the first-order defense against data flooding attacks. However, as it cannot be guaranteed that all routers would implement something like this, and as the RP (if SSM is not used) could be loaded from multiple senders, additional protections are needed as well. This harms legitimate senders on the same link as an attacker. This does not prevent a host from sending a lot of traffic to the same group -- an action that would harm only the DR and the RP of the group, is similar to unicast DoS attacks against one source, and is not considered critical to the overall security.

#### 5.3.3. PIM Signalling Rate-Limiter

A PIM signalling rate-limiter is a token-bucket-based rate-limiter that would limit all multicast PIM messaging, either through a specific interface or globally on the router, with a bucket of depth of PIM\_DEPTH, refilling at PIM\_RATE tokens per second. Example values could be PIM\_RATE=1000 and PIM\_DEPTH=10000.

This would be second-order defense against PIM state attacks when IGMP/MLD rate-limiters haven't been implemented or haven't been effective. This limiter might not need to be active by default, as long as the values are configurable. The main applicability for this filter would be at a border of PIM domain in case PIM state attacks are detected. This harms legitimate receivers as well.

#### 5.3.4. Unicast-Decapsulation Rate-Limiter

A unicast-decapsulation rate-limiter is a simple decapsulation rate-limiter that would protect the CPU usage in the router by limiting the packets per second (depending on the router architecture) and disregarding the source of the registers. This could also be an additional check to be used before decapsulation and checking the group to throttle the worst of the decapsulation CPU consumption. This limit should have to be quite high, and would hamper the existing legitimate sessions as well.

#### 5.3.5. PIM Register Rate-Limiter

A PIM Register rate-limiter is a token-bucket-based rate-limiter that would limit register decapsulation of PIM Register messages with a bucket of depth of REG\_DEPTH, refilling at REG\_RATE tokens per second. If the router has restarted recently, a larger initial bucket should be used. Example values could be REG\_RATE=1 and REG\_DEPTH=10 (or REG\_DEPTH=500 after restart).

This would be second-order defense against data flooding: if the DRs would not implement appropriate limiters, or if the total number of flooded groups rises too high, the RP should be able to limit the

rate with which new groups are created. This does not harm legitimate senders, as long as their groups have already been created.

#### 5.3.6. MSDP Source-Active Rate-Limiter

A MSDP source-active rate-limiter is a token-bucket-based, source-based rate-limiter, that would limit new groups per source with a bucket of depth of SAG\_DEPTH, refilling at SAG\_RATE tokens per second. Example values could be SAG\_RATE=1 and SAG\_DEPTH=10.

This would be second-order defense, at both the MSDP SA sending and receiving sites, against data flooding and MSDP vulnerabilities in particular. The specific threat being addressed here is a source (or multiple different sources) trying to "probe" (e.g., virus or worm) different multicast addresses. [16] discusses different MSDP attack prevention mechanisms at length.

#### 5.4. Passive Mode for PIM

As described in the last paragraph of Section 3, hosts are also able to form PIM adjacencies and send disrupting traffic unless great care is observed at the routers. This stems from the fact that most implementations require that stub LANs with only one PIM router must also have PIM enabled (to enable PIM processing of the sourced data, etc.) Such stub networks however do not require to actually run the PIM protocol on the link. Therefore, such implementations should provide an option to specify that the interface is "passive" with regard to PIM: no PIM packets are sent or processed (if received), but hosts can still send and receive multicast on that interface.

### 6. Security Considerations

This memo analyzes the security of PIM routing infrastructures in some detail and proposes enhancements to mitigate the observed threats.

This document does not discuss adding (strong) authentication to the multicast protocols. The PIM-SM specification [1] describes the application of IPsec for routing authentication; note that being able to authenticate the register messages and to prevent illegitimate users from establishing PIM adjacencies seem to be the two most important goals. The IGMPv3 specification [11] describes the use of IPsec for group management (IPsec for MLDv2 may be applied similarly), which is out of scope for this memo. However, note that being able to control the group memberships might reduce the receiver-based attacks.



However, one should keep in mind two caveats: authentication alone might not be sufficient, especially if the user or the host stack (consider a worm propagation scenario) cannot be expected to "behave well"; and adding such authentication likely provides new attack vectors, e.g., in the form of a CPU DoS attack with an excessive amount of cryptographic operations.

## 7. Acknowledgements

Kamil Sarac discussed "return routability" issues at length. Stig Venaas and Bharat Joshi provided feedback to improve the document quality. Bill Fenner and Russ Housley provided useful comments during the IESG evaluation.

## 8. References

### 8.1. Normative References

- [1] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [2] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.
- [3] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.
- [4] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004.
- [5] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, July 2006.

### 8.2. Informative References

- [6] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [7] Bhattacharyya, S., "An Overview of Source-Specific Multicast (SSM)", RFC 3569, July 2003.
- [8] Thaler, D., Fenner, B., and B. Quinn, "Socket Interface Extensions for Multicast Source Filters", RFC 3678, January 2004.

- [9] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC 3740, March 2004.
- [10] Daley, G. and G. Kurup, "Trust Models and Security in Multicast Listener Discovery", Work in Progress, July 2004.
- [11] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [12] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [13] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [14] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [15] Handley, M., "Bi-directional Protocol Independent Multicast (BIDIR-PIM)", Work in Progress, October 2005.
- [16] Rajvaidya, P., Ramachandran, K., and K. Almeroth, "Detection and Deflection of DoS Attacks Against the Multicast Source Discovery Protocol", UCSB Technical Report, May 2003.

#### Appendix A. RPF Considers Interface, Not Neighbor

In most current implementations, the RPF check considers only the incoming interface, and not the upstream neighbor (RPF neighbor).

This can result in accepting packets from the "wrong" RPF neighbor (the neighbor is "wrong" since, while the RPF check succeeds and the packet is forwarded, the unicast policy would not have forwarded the packet).

This is a problem in the media where more than two routers can connect to, in particular, Ethernet-based Internet Exchanges. Therefore, any neighbor on such a link could inject any PIM signalling as long as a route matching the address used in the signalling is going through the interface.

Note that for PIM signalling to be accepted, a PIM adjacency must have been established. However, typically, this does not help much against willful attackers, as PIM adjacencies are usually formed with anyone on the link. Still, the requirement is that the neighbor has enabled PIM in the concerned interface. That is, in most cases, the threat is limited to attackers within the operators in the exchange, not third parties. On the other hand, data plane forwarding has no such checks -- and having such checks would require that one look at the link-layer addresses used. That is, this checking is not as feasible as one might hope.

## Appendix B. Return Routability Extensions

The multicast state information is built from the receiver side, and it can be currently pruned only by the receiver-side DR. If the RP or the source for the group is non-existent, the state can't be pruned by the DR without return routability extensions to provide such information. There might also be a need to remove the state in some cases when there is no multicast traffic sent to that group. This section discusses the alternative ways to remove the unused state information in the routers, so that it can't be used in state-based DoS attacks. Note that rate-limiting PIM Joins gives some protection against the state attacks.

### B.1. Sending PIM-Prune Messages Down the Tree

When a router discovers the non-existence of the RP or the source, it can create a PIM-Prune message and send it back to the join originator. However, since it does not know the unicast IP address of join originator DR, it cannot directly unicast it to that router.

A possible alternative is to use a link-local multicast group address (e.g., all-pim routers local multicast address) to pass this information back toward the joining DR. Since the routers from this current router all the way back to the joining DR have forwarding state entry for the group, they can use this state information to see how to forward the PIM-Prune message back.

Each on-tree router, in addition to forwarding the PIM-Prune message, can also prune the state from its state tables. This way, the PIM-Prune message will go back to the DR by following the multicast forwarding state information created so far. In addition, if we use some sort of RPF checks during this process, we can also make it more difficult to inject such PIM-Prune messages maliciously.

A potential abuse scenario may involve an attacker that has access to a router on the direct path and can send such PIM-Prune messages down the tree branch so as to prune the branch from the tree. But such an attacker can currently achieve the same effect by sending a PIM-Prune message toward the source from the same point on the tree. So, the proposed mechanism does not really aggravate the situation.

One visible overhead in this new scenario might be that someone can send bogus join messages to create redundant PIM-Join and PIM-Prune messages in the network.

### B.2. Analyzing Multicast Group Traffic at DR

Another possible way to remove the unused state information would be to analyze individual group traffic at the DR and if there is no multicast traffic for a certain group within a certain time limit, the state should be removed. In here, if the receiver is malicious and wants to create states in the network, then it can send joins to different groups and create states on routers for each of these different groups until the DR decides that the groups are inactive and initiates the prune process. In addition, during the prune process, the routers will again process all these prune messages and therefore will be spending time.

### B.3. Comparison of the Above Approaches

Both of these solutions have the same problem of renewing the multicast state information. The DR shouldn't permanently block the state building for that group, but should restrict the PIM Joins if it notices that the receiver is abusing the system. One additional option is to block the PIM Joins to the non-existent source/RP for a certain time.

In the first approach (sending PIM-Prunes down the tree), part of the goal was to prune the states in the routers much sooner than in the second approach. (That is, the goal is to make sure that the routers will not be keeping unnecessary states for long time.)

The second approach works also for DoS attacks related to the existing source/RP addresses, could be more quickly implemented and deployed in the network, and does not have any relationship with the other deployments (no need to change all PIM routers).

## Authors' Addresses

Pekka Savola  
CSC/FUNET  
Espoo  
Finland

E-Mail: [psavola@funet.fi](mailto:psavola@funet.fi)

Rami Lehtonen  
TeliaSonera  
Hataanpaan valtatie 20  
Tampere 33100  
Finland

E-Mail: [rami.lehtonen@teliasonera.com](mailto:rami.lehtonen@teliasonera.com)

David Meyer

E-Mail: [dmm@1-4-5.net](mailto:dmm@1-4-5.net)

## Full Copyright Statement

Copyright (C) The Internet Society (2006).

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).

## Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).