Internet Engineering Task Force
Internet-Draft
Expires: July 22, 2004

P. Savola
CSC/FUNET
R. Lehtonen
TeliaSonera
D. Meyer
Jan 22, 2004

PIM-SM Multicast Routing Security Issues and Enhancements
draft-savola-mboned-mroutesec-00.txt

Status of this Memo

Copyright Notice

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, Source-Specific Multicast (SSM) model, and the ASM model enhanced by the Embedded RP group-to-RP mapping mechanism. This memo also describes enhancements to the protocol operations to mitigate these threats.

Table of Contents

1. Introduction

This memo 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 which have an effect on the multicast routing (whether intra- or inter-domain) are considered. For example, attacks where the hosts are specifically targeting the Designated Router (DR) or other routers of the link, or where hosts are disrupting other hosts on the same link are out of scope. Similarly, ensuring confidentiality, authentication and integrity of multicast groups and traffic is out of the scope [9].

PIM builds on a model where Reverse Path Forwarding (RPF) check is (among other things) used to ensure loop-free properties of the multicast distribution trees. As a side effect, this limits the effect of using forged source addresses, which is often 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 messages, which a host may send on its own.

We consider PIM-SM [1] operating in the traditional Any Souce 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. If Bidirectional-PIM enhancements are globally significant, and have implications, they could also be considered.

2. Terminology

ASM

Term "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

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

Embedded-RP

Embedded-RP refers to ASM model where the Embedded-RP mapping

mechanism is used to find the RP for a group, and MSDP is not needed.

Target Router

Target Router is used to refer to either the RP processing a packet (ASM or Embedded-RP), or DR close to the transmitting source (SSM).

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 behave "well", in the same sense that unicast routers are expected to behave well, and are not a significant source of abuse. This assumption is not entirely correct, but it simplifies the analysis of threat models. If seen important, the threats caused by misbehaving multicast routers (including fake multicast routers) may be considered separately.

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].

3.1 Receiver-based Attacks

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

One should note that hosts can also originate PIM messages (e.g. PIM Joins) as long as their source address passes the RPF checks. 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 yourself. The PIM-SM specification, however, states that these messages should only be accepted from known PIM neighbors [1]; if these would be implemented, the hosts would have to forge PIM Hello messages as well.

One should also 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 timer expires or the state changes at the DR.

Also, if the host uses IGMPv3 [10] or MLDv2 [11], it is able to join multiple sources for the same group and each of these joins for the

same group generates new PIM (Source, Group), or (S,G) Joins.

3.1.1 Joins to Different Groups

Description of the threat: Join Flooding. This happens when a host tries to join, once or a couple of times, to a group or a channel, and the DR generates a PIM Join towards the Target Router. The group/channel or the Targer Router may or may not exist.

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 refered to as "PIM State"), or the Target Router with an excessive number of packets.

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). It should also be noted that the host can join a number of different channels with only one IGMPv3/MLDv2 Report as the protocol allows to include multiple sources in the same message. However, even short-lived state may be harmful, if 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:

- ASM: a (*,G) join is sent towards an intra-domain RP, causing state on that path; in turn, that RP joins to the DR of the source (if it exists). If the source address was specified by the host in the IGMPv3/MLDv2 Report, a (S,G) Join is sent directly towards the specified source.

- SSM: a (S,G) join is sent inter-domain to the DR of the source S, causing state on that path. If the source does not exist, the join goes to the closest router to S as possible.

- Embedded RP: a (*,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 closest router to the RP as possible.

If the source or RP does not exist, the multicast routing protocol does not have any means to remove the distribution tree if the host remains active. Worst case attack could be a host remaining active to many different groups (containing either imaginary source or RP).

3.2 Source-based Attacks

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

3.2.1 Sending Multicast to Empty Groups

Description of the threat: Data Flooding. This happens when a host sends data packets to a multicast group or channel for which there are no real subscribers.

Note that as unicast-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 [12] has been deployed [13].

Examples of this 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 or different groups. This results in:

- ASM: the DR unicast-encapsulates the packets in Register messages to the intra-domain RP, which may join to the source and issue a Register-Stop, but 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.

- SSM: the DR receives the data, but the data does not propagate from the DR unless someone joins the (S,G) channel.

- 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. The data continues to be encapsulated.

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

- The MSDP control plane traffic generated can cause a significant amount of messages/data which may overload the routers receiving it. The thorough analysis of MSDP vulnerabilities can be found from [14]. This 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.

- As a large amount of data is forwarded on the multicast tree; if multicast forwarding is performed on software, it may be a 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.

- 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

Description of the threat: Group Integrity Violation. This happens when a host sends packets to a group or channel, which already exists, to disturb the users of the existing group/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 those hosts which are on-link with the source are able to perform this, so this is not really relevant in the scope of this memo.

With ASM and Embedded RP sources can inject bogus traffic through RPs, which provide the source discovery for the group. The RP(s) send the traffic over the shared tree towards receivers (routers with (*,G) state). DR then forwards the bogus traffic to receivers unless the legitimate recipients are able to filter out unwanted sources, e.g., using 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 unicast-encapsulated in register packets to the RP -- if the RP drops the packet based on access-list, rate-limiter or something else, it doesn't get injected to an existing group.

With ASM this "source control" is distributed across all the PIM

domains, which decreases it's applicability. Embedded RP enables easier control, because source discovery is done through single RP per group.

So, for this attack to succeed, the RP must decapsulate the packets, and join to the source.

3.3 Aggravating Factors to the Threats

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

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 also 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 would get elected to become DR for the subnet, but not perform any of its functions. These are described at some length in [1], but are also considered out of scope of this memo.

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 analysed at more length in Section 5.1.

3.3.2 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.

3.3.3 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.

3.3.4 Injecting a Bogus Route

Hosts that able to inject a bogus route can be used to "steal" PIM Joins. This prevents the correct multicast tree forming. If the injected route information changes, it causes route flapping and that could have harmful effect on multicast routing and packet delivery (depending on the group). The threat is similar to unicast case meaning that by injecting a bogus route, routing does not work correctly.

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:

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

Notes:

1) in ASM host can directly join also (S,G) groups with IGMPv3/MLDv2 and thus have same characteristics as SSM (also allows inter-domain shared state to be created).

2) allows inter-domain shared state to be created.

3) register messages can be sent to long-distance RPs with spoofed source addresses and this could create unnecessary joins towards DRs (from spoofed source address space).

4.2 Enhancements for Threat Mitigation

There are several desirable actions ("requirements") which could be considered to mitigate these threats; these are listed below. A future revision of this memo will describe the best methods and parameters at more detail.

- Inter-domain MSDP (ASM) should be retired (or not introduced) to avoid attacks; or, if this is not reasonable, the DRs should rate-limit the unicast-encapsulation (note that the hosts can avoid this) and (more importantly) the RPs should rate-limit the unicast-decapsulation especially from different sources, or MSDP must rate-limit the MSDP data generation for new sources.

- DRs should rate-limit PIM Joins and Prunes somehow; there are multiple possibilities how exactly this should be considered (i.e., which variables to take into the consideration).

- DRs could rate-limit unicast-encapsulation somehow; there are multiple ways to perform this. Note that the hosts can avoid this by performing the unicast-encapsulation themselves if so inclined.

- RPs could rate-limit unicast-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 e.g. rate-limiters behave.

- RPs should rate limit the MSDP SA messages coming from MSDP peers.

- RPs could limit or even disable the SA cache size. However, this could have negative effects on normal operation.

- RPs should provide good interfaces to reject packets which are not interesting; for example, if an Embedded RP group is not configured to be allowed in the RP, the unicast-encapsulated packets would not even be decapsulated.

- 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 that can be created by a given source, S.

5. PIM Security Enhancements

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

5.1 Remote Routability Signalling

As described in section 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.

2) An additional method could be implementing a timer on the RPs 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 (i.e., the RP would insert a cookie in the bubble, and it would have to be present in the register message.)

With all the models, PIM Joins and other state management messages could also be somehow managed. For example:

1) At the very least, receiving an ICMP unreachable message (of any flavor) should cause the DR to stop the PIM messages toward the destination, as the packets will not be received anyway.

(Currently the sending of an ICMP error message in response to a multicast packet, even though in this case the PIM message to only received by one router, is prohibited.)

2) A possible method would be limiting the number of PIM messages sent towards a destination until some response (e.g. other PIM state messages).

5.2 RPF to Check Neighbor, not Interface

As described in Section 3.3.2, especially Ethernet-based Internet Exchange Points (IXP) are susceptible to signalling attacks from any member of the IXP, as the RPF considers the Interface, not a Neighbor.

Consequently, PIM must be modified so that on non- point-to-point links, the RPF must also consider whether the neighbor is correct. Note that in case of IPv6, this requires (an already necessary) a mapping between link-local and global addresses.

5.3 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; factors in the evaluation might be e.g.:

- Whether the host is willfully malicious, uncontrolled (e.g., virus/worm), or a regular user just doing something wrong.
- Whether the threat is aimed towards a single group, a single RP handling the group, or the (multicast) routing infrastructure in general.
- Whether the host on a subnet is spoofing its address (but still as one which fulfills the RPF checks of the DR) or not.
- Whether the host may generate the PIM join (and similar) messages itself to avoid rate-limiters at the DR if possible.
- Whether unicast RPF checks are applied on the link (i.e., whether the host can send unicast-encapsulated register-messages on its own).
- Whether blocking the misbehaving host on a subnet is allowed to also block other, legitimate hosts on the same subnet.

- 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 which 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 which 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 make an assumption 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 legimate multicast users. Therefore several parameters for rate limiting should be used to prevent such operation.

The next revisions of this document (or separated in other documents, if appropriate) will include more explicit discussion of the best ways to perform rate-limiting, especially considering the effects on the legimate traffic.

6. Security Considerations

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

7. IANA Considerations

This memo is for informational purposes and does not introduce new namespaces for the IANA to manage.

Normative References

[1] Fenner, B., Handley, M., Holbrook, H. and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode PIM-SM): Protocol Specification (Revised)", draft-ietf-pim-sm-v2-new-08 (work in progress), October 2003.

[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", draft-ietf-ssm-arch-04 (work in progress), October 2003.

[4] Savola, P. and B. Haberman, "Embedding the Address of RP in IPv6 Multicast Address", draft-ietf-mboned-embeddedrp-00 (work in progress), October 2003.

[5] Barbir, A., Murphy, S. and Y. Yang, "Generic Threats to Routing Protocols", draft-ietf-rpsec-routing-threats-04 (work in progress), December 2003.

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", draft-ietf-magma-msf-api-05 (work in progress), July 2003.

[9] Hardjono, T. and B. Weis, "The Multicast Security Architecture", draft-ietf-msec-arch-05 (work in progress), January 2004.

[10] Cain, B., Deering, S., Kouvelas, I., Fenner, B. and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[11] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", draft-vida-mld-v2-08 (work in progress), December 2003.

[12] 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.

[13] Baker, F. and P. Savola, "Ingress Filtering for Multihomed

Networks", draft-savola-bcp38-multihoming-update-03 (work in progress), December 2003.

[14] Rajvaidya, P., Ramachandran, K. and K. Almeroth, "Detection and Deflection of DoS Attacks Against the Multicast Source Discovery Protocol", IEEE Infocom 2003.

Authors' Addresses

Pekka Savola
CSC/FUNET

Espoo
Finland

EMail: psavola@funet.fi

Rami Lehtonen
TeliaSonera
Hataanpaan valtatie 20
Tampere 33100
Finland

EMail: rami.lehtonen@teliasonera.com

David Meyer

EMail: dmm@1-4-5.net

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.