

Host Threats to Protocol Independent Multicast (PIM)

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.

Abstract

This memo complements the list of multicast infrastructure security threat analysis documents by describing Protocol Independent Multicast (PIM) threats specific to router interfaces connecting hosts.

Table of Contents

1. Introduction	2
2. Host-Interface PIM Vulnerabilities	2
2.1. Nodes May Send Illegitimate PIM Register Messages	3
2.2. Nodes May Become Illegitimate PIM Neighbors	3
2.3. Routers May Accept PIM Messages from Non-Neighbors	3
2.4. An Illegitimate Node May Be Elected as the PIM DR or DF	3
2.4.1. PIM-SM Designated Router Election	3
2.4.2. BIDIR-PIM Designated Forwarder Election	4
2.5. A Node May Become an Illegitimate PIM Asserted Forwarder	4
2.6. BIDIR-PIM Does Not Use RPF Check	4
3. On-Link Threats	5
3.1. Denial-of-Service Attack on the Link	5
3.2. Denial-of-Service Attack on the Outside	6
3.3. Confidentiality, Integrity, or Authorization Violations	6
4. Mitigation Methods	7
4.1. Passive Mode for PIM	7
4.2. Use of IPsec among PIM Routers	7
4.3. IP Filtering PIM Messages	8
4.4. Summary of Vulnerabilities and Mitigation Methods	8
5. Acknowledgements	10
6. Security Considerations	10
7. References	10
7.1. Normative References	10
7.2. Informative References	10

1. Introduction

There has been some analysis of the security threats to the multicast routing infrastructures [RFC4609], some work on implementing confidentiality, integrity, and authorization in the multicast payload [RFC3740], and also some analysis of security threats in Internet Group Management Protocol/Multicast Listener Discovery (IGMP/MLD) [DALEY-MAGMA], but no comprehensive analysis of security threats to PIM at the host-connecting (typically "Local Area Network") links.

We define these PIM host threats to include:

- o Nodes using PIM to attack or deny service to hosts on the same link,
- o Nodes using PIM to attack or deny service to valid multicast routers on the link, or
- o Nodes using PIM (Register messages) to bypass the controls of multicast routers on the link.

The attacking node is typically a host or a host acting as an illegitimate router.

A node originating multicast data can disturb existing receivers of the group on the same link, but this issue is not PIM-specific so it is out of scope. Subverting legitimate routers is out of scope. Security implications on multicast routing infrastructure are described in [RFC4609].

This document analyzes the PIM host-interface vulnerabilities, formulates a few specific threats, proposes some potential ways to mitigate these problems, and analyzes how well those methods accomplish fixing the issues.

It is assumed that the reader is familiar with the basic concepts of PIM.

Analysis of PIM-DM [RFC3973] is out of scope of this document.

2. Host-Interface PIM Vulnerabilities

This section briefly describes the main attacks against host-interface PIM signaling, before we get to the actual threats and mitigation methods in the next sections.

The attacking node may be either a malicious host or an illegitimate router.

2.1. Nodes May Send Illegitimate PIM Register Messages

PIM Register messages are sent unicast, and contain encapsulated multicast data packets. Malicious hosts or routers could also send Register messages themselves, for example, to get around rate-limits or to interfere with foreign Rendezvous Points (RPs), as described in [RFC4609].

The Register message can be targeted to any IP address, whether in or out of the local PIM domain. The source address may be spoofed, unless spoofing has been prevented [RFC3704], to create arbitrary state at the RPs.

2.2. Nodes May Become Illegitimate PIM Neighbors

When PIM has been enabled on a router's host interface, any node can also become a PIM neighbor using PIM Hello messages. Having become a PIM neighbor in this way, the node is able to send other PIM messages to the router and may use those messages to attack the router.

2.3. Routers May Accept PIM Messages from Non-Neighbors

The PIM-SM (Sparse Mode) specification recommends that PIM messages other than Hellos should not be accepted, except from valid PIM neighbors. The Bidirectional-PIM (BIDIR-PIM) specification specifies that packets from non-neighbors "SHOULD NOT" be accepted; see Section 5.2 of [RFC5015]. However, the specification does not mandate this, so some implementations may be susceptible to attack from PIM messages sent by non-neighbors.

2.4. An Illegitimate Node May Be Elected as the PIM DR or DF

2.4.1. PIM-SM Designated Router Election

In PIM-SM, the Designated Router (DR) on a Local Area Network (LAN) is responsible for Register-encapsulating data from new sources on the LAN, and for generating PIM Join/Prune messages on behalf of group members on the LAN.

A node that can become a PIM neighbor can also cause itself to be elected DR, whether or not the DR Priority option is being used in PIM Hello messages on the LAN.

2.4.2. BIDIR-PIM Designated Forwarder Election

In BIDIR-PIM [RFC5015], a Designated Forwarder (DF) is elected per link. The DF is responsible for forwarding data downstream onto the link, and also for forwarding data from its link upstream.

A node that can become a BIDIR-PIM neighbor (this is just like becoming a PIM neighbor, except that the PIM Hello messages must include the Bidirectional Capable PIM-Hello option) can cause itself to be elected DF by sending DF Offer messages with a better metric than its neighbors.

There are also some other BIDIR-PIM attacks related to DF election, including spoofing DF Offer and DF Winner messages (e.g., using a legitimate router's IP address), making all but the impersonated router believe that router is the DF. Also, an attacker might prevent the DF election from converging by sending an infinite sequence of DF Offer messages.

For further discussion of BIDIR-PIM threats, we refer to the Security Considerations section in [RFC5015].

2.5. A Node May Become an Illegitimate PIM Asserted Forwarder

With a PIM Assert message, a router can be elected to be in charge of forwarding all traffic for a particular (S,G) or (*,G) onto the LAN. This overrides DR behavior.

The specification says that Assert messages should only be accepted from known PIM neighbors, and "SHOULD" be discarded otherwise. So, either the node must be able to spoof an IP address of a current neighbor, form a PIM adjacency first, or count on these checks being disabled.

The Assert Timer, by default, is 3 minutes; the state must be refreshed or it will be removed automatically.

As noted before, it is also possible to spoof an Assert (e.g., using a legitimate router's IP address) to cause a temporary disruption on the LAN.

2.6. BIDIR-PIM Does Not Use RPF Check

PIM protocols do not perform Reverse Path Forwarding (RPF) check on the shared tree (e.g., in PIM-SM from the RP to local receivers). On the other hand, RPF check is performed, e.g., on stub host interfaces. Because all forwarding in BIDIR-PIM is based on the shared tree principle, it does not use RPF check to verify that the

forwarded packets are being received from a "topologically correct" direction. This has two immediately obvious implications:

1. A node may maintain a forwarding loop until the Time to Live (TTL) runs out by passing packets from interface A to B. This is not believed to cause significant new risk as with a similar ease such a node could generate original packets that would loop back to its other interface.
2. A node may spoof source IP addresses in multicast packets it sends. Other PIM protocols drop such packets when performing the RPF check. BIDIR-PIM accepts such packets, allowing easier Denial-of-Service (DoS) attacks on the multicast delivery tree and making the attacker less traceable.

3. On-Link Threats

The previous section described some PIM vulnerabilities; this section gives an overview of the more concrete threats exploiting those vulnerabilities.

3.1. Denial-of-Service Attack on the Link

The easiest attack is to deny the multicast service on the link. This could mean either not forwarding all (or parts of) multicast traffic from upstream onto the link, or not registering or forwarding upstream the multicast transmissions originated on the link.

These attacks can be done in multiple ways: the most typical one would be becoming the DR through becoming a neighbor with Hello messages and winning the DR election. After that, one could, for example:

- o Not send any PIM Join/Prune messages based on the IGMP reports, or
- o Not forward or register any sourced packets.

Sending PIM Prune messages may also be an effective attack vector even if the attacking node is not elected DR, since PIM Prune messages are accepted from downstream interfaces even if the router is not a DR.

An alternative mechanism is to send a PIM Assert message, spoofed to come from a valid PIM neighbor or non-spoofed if a PIM adjacency has already been formed. For the particular (S,G) or (*,G) from the Assert message, this creates the same result as getting elected as a DR. With BIDIR-PIM, similar attacks can be done by becoming the DF or by preventing the DF election from converging.

3.2. Denial-of-Service Attack on the Outside

It is also possible to perform Denial-of-Service attacks on nodes beyond the link, especially in environments where a multicast router and/or a DR is considered to be a trusted node.

In particular, if DRs perform some form of rate-limiting, for example, on new Join/Prune messages, becoming a DR and sending those messages yourself allows one to subvert these restrictions; therefore, rate-limiting functions need to be deployed at multiple layers, as described in [RFC4609].

In addition, any host can send PIM Register messages on their own, to whichever RP it wants; further, if unicast RPF (Reverse Path Forwarding) mechanisms [RFC3704] have not been applied, the packet may be spoofed. This can be done to get around rate-limits, and/or to attack remote RPs, and/or to interfere with the integrity of an ASM group. This attack is also described in [RFC4609].

Also, BIDIR-PIM does not prevent nodes from using topologically incorrect addresses (source address spoofing) making such an attack more difficult to trace.

3.3. Confidentiality, Integrity, or Authorization Violations

Contrary to unicast, any node is able to legitimately receive all multicast transmission on the link by just adjusting the appropriate link-layer multicast filters. Confidentiality (if needed) must be obtained by cryptography.

If a node can become a DR, it is able to violate the integrity of any data streams sent by sources on the LAN, by modifying (possibly in subtle, unnoticeable ways) the packets sent by the sources before Register-encapsulating them.

If a node can form a PIM neighbor adjacency or spoof the IP address of a current neighbor, then if it has external connectivity by some other means other than the LAN, the node is able to violate the integrity of any data streams sent by external sources onto the LAN. It would do this by sending an appropriate Assert message onto the LAN to prevent the genuine PIM routers forwarding the valid data, obtaining the multicast traffic via its other connection, and modifying those data packets before forwarding them onto the LAN.

In either of the above two cases, the node could operate as normal for some traffic, while violating integrity for some other traffic.

A more elaborate attack is on authorization. There are some very questionable models [HAYASHI] where the current multicast architecture is used to provide paid multicast service, and where the authorization/authentication is added to the group management protocols such as IGMP. Needless to say, if a host would be able to act as a router, it might be possible to perform all kinds of attacks: subscribe to multicast service without using IGMP (i.e., without having to pay for it), deny the service for the others on the same link, etc. In short, to be able to ensure authorization, a better architecture should be used instead (e.g., [RFC3740]).

4. Mitigation Methods

This section lists some ways to mitigate the vulnerabilities and threats listed in previous sections.

4.1. Passive Mode for PIM

The current PIM specification seems to mandate running the PIM Hello protocol on all PIM-enabled interfaces. Most implementations require PIM to be enabled on an interface in order to send PIM Register messages for data sent by sources on that interface or to do any other PIM processing.

As described in [RFC4609], running full PIM, with Hello messages and all, is unnecessary for those stub networks for which only one router is providing multicast service. 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.

4.2. Use of IPsec among PIM Routers

Instead of passive mode, or when multiple PIM routers exist on a single link, one could also use IPsec to secure the PIM messaging, to prevent anyone from subverting it. The actual procedures have been described in [RFC4601] and [LINKLOCAL].

However, it is worth noting that setting up IPsec Security Associations (SAs) manually can be a very tedious process, and the routers might not even support IPsec; further automatic key negotiation may not be feasible in these scenarios either. A Group Domain of Interpretation (GDOI) [RFC3547] server might be able to mitigate this negotiation.

4.3. IP Filtering PIM Messages

To eliminate both the unicast and multicast PIM messages, in similar scenarios to those for which PIM passive mode is applicable, it might be possible to block IP protocol 103 (all PIM messages) in an input access list. This is more effective than PIM passive mode, as this also blocks Register messages.

This is also acceptable when there is more than one PIM router on the link if IPsec is used (because the access-list processing sees the valid PIM messages as IPsec AH/ESP packets). In this case, unicast Register messages must also be protected with IPsec or the routing topology must be such that the link is never used to originate, or transit unicast Register messages.

When multiple routers exist on a link, IPsec is not required if it is possible to prevent hosts from sending PIM messages at the Ethernet switch (or equivalent) host ports. This could be accomplished in at least two ways:

1. Use IP access lists on the stub routers to allow PIM messages from the valid neighbor IP addresses only, and implement IP spoofing prevention at the Ethernet-switch-port level using proprietary mechanisms, or
2. Filter out all PIM messages at configured host ports on Ethernet switches instead of doing it on the routers.

The main benefit of this approach is that multiple stub routers can still communicate through the LAN without IPsec but hosts are not able to disturb the PIM protocol. The drawback is that Ethernet switches need to implement much finer-grained IP layer filtering, and the operational requirements of carefully maintaining these filters could be significant.

4.4. Summary of Vulnerabilities and Mitigation Methods

This section summarizes the vulnerabilities, and how well the mitigation methods are able to cope with them.

Summary of vulnerabilities and mitigations:

Sec	Vulnerability	One stub router			>1 stub routers		
		PASV	IPsec	Filt	PASV	IPsec	Filt
2.1	Hosts Registering	N	N+	Y	N	N+	Ysw
2.2	Invalid Neighbor	Y	Y	Y	*	Y	Ysw
2.3	Adjacency Not Req'd	Y	Y	Y	*	Y	Ysw
2.4	Invalid DR /DF	Y	Y	Y	*	Y	Ysw
2.5	Invalid Forwarder	Y	Y	Y	*	Y	Ysw
2.6	No RPF Check (BIDIR)	x	x	x	x	x	x

Figure 1

"*" means Yes if IPsec is used in addition; No otherwise.

"Ysw" means Yes if IPsec is used in addition or IP filtering is done on Ethernet switches on all host ports; No otherwise.

"N+" means that the use of IPsec between the on-link routers does not protect from this; IPsec would have to be used at RPs.

"x" means that, with BIDIR-PIM, IP access lists or RPF mechanisms need to be applied in stub interfaces to prevent originating packets with topologically incorrect source addresses. This needs to be done in addition to any other chosen approach.

To summarize, IP protocol filtering for all PIM messages appears to be the most complete solution when coupled with the use of IPsec between the real stub routers when there are more than one of them. However, IPsec is not required if PIM message filtering or a certain kind of IP spoofing prevention is applied on all the host ports on Ethernet switches. If hosts performing registering is not considered a serious problem, IP protocol filtering and passive-mode PIM seem to be equivalent approaches. Additionally, if BIDIR-PIM is used, ingress filtering will need to be applied in stub interfaces to multicast packets, as well as unicast, to prevent hosts using wrong source addresses.

5. Acknowledgements

Greg Daley and Gopi Durup wrote an excellent analysis of MLD security issues [DALEY-MAGMA], which gave inspiration in exploring the on-link PIM threats problem space.

Ayan Roy-Chowdhury, Beau Williamson, Bharat Joshi, Dino Farinacci, John Zwiebel, Stig Venaas, Yiqun Cai, and Eric Gray provided good feedback for this memo.

6. Security Considerations

This memo analyzes the threats to the PIM multicast routing protocol on host interfaces and proposes some possible mitigation techniques.

7. References

7.1. Normative References

- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC4609] Savola, P., Lehtonen, R., and D. Meyer, "Protocol Independent Multicast - Sparse Mode (PIM-SM) Multicast Routing Security Issues and Enhancements", RFC 4609, October 2006.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, October 2007.

7.2. Informative References

- [DALEY-MAGMA] Daley, G. and J. Combes, "Securing Neighbour Discovery Proxy Problem Statement", Work in Progress, February 2008.
- [HAYASHI] Hayashi, T., "Internet Group membership Authentication Protocol (IGAP)", Work in Progress, August 2003.
- [LINKLOCAL] Atwood, J., Islam, S., and M. Siami, "Authentication and Confidentiality in PIM-SM Link-local Messages", Work in Progress, February 2008.

- [RFC3547] Baugher, M., Weis, B., Hardjono, T., and H. Harney, "The Group Domain of Interpretation", RFC 3547, July 2003.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3740] Hardjono, T. and B. Weis, "The Multicast Group Security Architecture", RFC 3740, March 2004.
- [RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.

Authors' Addresses

Pekka Savola
CSC - Scientific Computing Ltd.
Espoo
Finland

EMail: psavola@funet.fi

James Lingard
Arastra, Inc.
P.O. Box 10905
Palo Alto, CA 94303
USA

EMail: jchl@arastra.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

