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

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

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:

- 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.

- 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.

- 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.