

Network Working Group
Internet Draft
Intended Status: Informational
Expires: May 2008

Mike McBride

November 2007

PIM Refresh Reduction Problem Statement

[draft-mmcbride-pim-refresh-problem-statement-00](#)

Status of this Memo

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

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

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

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

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

Copyright Notice

Copyright (C) The IETF Trust (2007).

Internet Draft

PIM Refresh Problem Statement

November 2007

Abstract

The PIM Working Group has a charter item to provide solutions to reduce the effects of periodic join/prune processing in PIM. The L3VPN Working Group identified the periodic messaging of PIM as a potential future scaling problem for PIM based MVPNs. This document identifies the issues we are trying to solve with PIM refresh reduction.

Table of Contents

1	Introduction	2
2	History	3
3	Problem	3
4	Solutions	4
5	Security Considerations	4
6	Iana Considerations	5
7	Acknowledgments	5
8	Normative References	5
9	Informative References	5
10	Authors' Addresses	5
11	Full Copyright Statement	6
12	Intellectual Property	6

[1](#). Introduction

PIM Joins are refreshed every 60 seconds by a downstream router to keep multicast state alive at the upstream router. The more number of states, the more control traffic required for refresh.

The PIM Working Group has a charter item to provide solutions to reduce the effects of periodic join/prune processing in PIM. The L3VPN Working Group identified the periodic messaging of PIM as a potential future scaling problem for PIM based MVPNs. This document identifies the issues we are trying to solve with PIM refresh reduction.

2. History

At the November 2004 IETF, the L3VPN WG identified the effects of periodic Join/Prune processing in PIM as a potential scalability problem for PIM based MVPNs and asked that the PIM WG provide a solution. The PIM WG subsequently created a pim refresh design team to understand the problem area and provide a solution.

Although there were disagreements on certain details, the design team had rough agreement on a Join/Prune acknowledgement solution. But then the team decided to step back and ask the L3VPN WG for a requirements document to see if this is really an area for which we should provide a solution.

The PIM WG discussed the future efforts we wanted to make to PIM. The primary focus, of the working group, at the time was in getting the PIMv2 draft completed. There are many enhancements we could make to PIM. Should we refrain from such future PIM enhancements or should we leave PIM alone and close the working group? Or should we start going down the path of PIMv3?

The L3VPN did produce an MVPN requirements document which did help to understand the future scalability requirements of MVPN. But it didn't help in understanding at what point PIM scalability would cause a problem for PIM based MVPNs.

In 2007, with the PIMv2 draft complete and the PIM WG tasked with a new Charter, it was decided it would be of benefit to provide enhancements to PIM. The WG has determined there is additional work to be accomplished and is chartered to standardize extensions to [RFC 4601](#) - Protocol Independent Multicast Version 2 - Sparse Mode. These PIM extensions will include PIM refresh reduction.

3. Problem

As previously listed, the PIM WG has decided to discover ways to reduce the effects of the periodic Join/Prune processing in PIM. This could specifically help in the future scalability of PIM based MVPN, but is certainly not limited to that solution.

PIM Joins are refreshed every 60 seconds by a downstream router to keep multicast state alive at the upstream router. The more number of states, the more control traffic is required for refresh. Scaling could become an issue, especially with VPNs where there could be 1000 MVPNs with 100 mroute entries each and 10 downstream interfaces (1 million state refreshes every minute in steady state).

Route changes cause Joins to be sent to the new RPF neighbor. If the Join is lost, the disruption of traffic is for at least 60 seconds. There is no reliability in the Join/Prune exchange between downstream and upstream routers. Larger values of holdtime in the Join/Prune PDU will reduce the frequency of refreshes, but can cause larger convergence delays.

[4. Solutions](#)

This is not a solutions draft. Subsequent to this draft, there will be drafts which outline solutions to this problem. The following ideas have been discussed as possible solutions to be further specified:

- + Join/Prune Ack extension to PIM.
- + Hard state (TCP) solution.
- + A PIMv3 containing strong RPF from forwarding plane, explicit tracking of all neighbors and hard state
- + Include checksums in Hello messages rather than sending periodic JPs.
- + Use long holdtimes.

5. Security Considerations

This document is a problem statement, which describes the reduction of PIM messaging, and does not introduce security considerations by itself. Any potential solution must protect against exploiting PIM as specified in [RFC 4601](#).

6. Iana Considerations

This document does not require any action on the part of IANA.

7. Acknowledgments

We'd like to thank Dino Farinacci, Suresh Boddapati, Tom Pusateri, Marshall Eubanks, Robert Kebler, Venu Hemige, Yiqun Cai, Yakov Rekhter, Yetik Serbest, Albert Tian, for their work on the pim refresh design team and helping the PIM WG to understand a few possible solutions.

8. Normative References

[Fenner] B. Fenner, "Protocol Independent Multicast - Sparse Mode (PIM-SM)". [RFC 4601](#)

[MVPN] "Multicast in MPLS/BGP IP VPNs", Rosen, Aggarwal, July 2007,

9. Informative References

[MORIN] T. Morin, "Requirements for Multicast in L3 Provider-Provisioned VPNs", [RFC 4834](#)

10. Authors' Addresses

Mike McBride
mmcbride@cisco.com

McBride

[Page 5]

Internet Draft

PIM Refresh Problem Statement

November 2007

11. Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

12. Intellectual Property

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

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.