Network Working Group
Internet-Draft
Intended status: Experimental
Expires: February 23, 2009

Dino Farinacci
IJsbrand Wijnands
Apoorva Karan
Arjen Boers
cisco Systems
Maria Napierala
AT&T Labs

August 22, 2008

A Reliable Transport Mechanism for PIM
draft-ietf-pim-port-00.txt

Status of this Memo

Copyright Notice

Abstract

This draft describes how a reliable transport mechanism can be used by the PIM protocol to optimize CPU and bandwidth resource utilization by eliminating periodic Join/Prune message transmission. This draft proposes a modular extension to PIM to use either the TCP or SCTP transport protocol.

Table of Contents

1. Introduction

The goals of this specification are:

- To create a simple incremental mechanism to provide reliable PIM message delivery in PIM version 2.
- The reliable transport mechanism will be used for Join-Prune message transmission only.
- Can be used for link-local transmission of Join-Prune messages or multi-hop for use in a multicast VPN environments.
- When a router supports this specification, it need not use the reliable transport mechanism on every interface. That is, negotiation on per interface basis (or MDT basis) will occur.

The explicit non-goals of this specification are:

- Changes to the PIM protocol machinery as defined in [RFC4601]. The reliable transport mechanism will be used as a plugin layer so the PIM component does not know it is really there.
- Provide support for both Datagram mode and Transport mode (see Section 1.2 for definitions) on the same physical interface or MDT.

This document will specify how periodic JP message transmission can be eliminated by using TCP [RFC0761] or SCTP [RFC4960] as the reliable transport mechanism for JP messages.

This specification enables greater scalability in multicast deployment since the processing required for protocol state maintenance can be reduced. These enhancements to PIMv2 are applicable to IP multicast over routed services and VPNs [MCAST-VPN]. In addition to reduced processing on PIM enabled routers, another important feature is the reduced join and leave latency provided through a reliable transport.

In many existing and emerging networks, particularly wireless and mobile satellite systems, link degradation due to weather, interference, and other impairments can result in temporary spikes in the packet loss. In these environments, periodic PIM joining can cause join latency when messages are lost causing a retransmission only 60 seconds later. By applying a reliable transport, a lost join is retransmitted rapidly. Furthermore, when the last user leaves a multicast group, any lost prune is similarly repaired and the multicast stream is quickly removed from the wireless/satellite link.

Without a reliable transport, the multicast transmission could otherwise continue until it timed out, roughly 3 minutes later. As network resources are at a premium in many of these environments, rapid termination of the multicast stream is critical to maintaining efficient use of bandwidth.

1.1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.2. Definitions

PORT: Stands for PIM Over Reliable Transport. Which is the short form for describing the mechanism in this specification where PIM can use the TCP or SCTP transport protocol.

JP Message: An abbreviation for a Join-Prune message.

Periodic JP: A JP message sent periodically to refresh state.

Incremental JP: A JP message sent as a result of state creation or deletion events. Also known as a triggered message.

Native JP: A JP message which is carried with an IP protocol type of PIM.

Reliable JP: A JP message using TCP or SCTP for transport.

Datagram Mode: The current procedures PIM uses by encapsulating JP messages in IP packets sent either triggered or periodically.

Transport Mode: Procedures used by PIM defined in this specification for sending JP messages over the TCP or SCTP transport layer.

MDT/PMSI: Used interchangeably in this document. An MDT tunnel is one used between PE router to provide support for a Multicast VPN. The new standards term for an MDT tunnel is a Provider-Network Multicast Service Interface or PMSI.

Segmented Multi-Access LAN: A segmented (or partitioned) LAN is like a virtual overlay network using the physical LAN to realize control and data packets. Multiple overlay networks may be created using the physical LAN, much like how VLANs or PMSI overlays are configured over a multi-access phsyical LAN. The interface associated with the partitioned LAN is like an NBMA interface type so explicit tracking can be accomplished. Each partitioned or segmented LAN has it's own data-link encapsulation and link-layer multicast is still used to avoid head-end replication. This concept also applies to MDTs/PMSIs and is called "Segmented MDTs/PMSIs". A Segmented MDT/PMSI is a MDT/PMSI that has a single forwarder (i.e. a single ingress PE) for any

multicast stream.

2. Protocol Overview

PIM Over Reliable Transport (PORT) is a simple extension to PIMv2 for refresh reduction of PIM JP messages. It involves sending incremental rather than periodic JPs over a TCP/SCTP connection between PIM neighbors.

PORT can be incrementally used on an interface between PORT capable neighbors. Routers which are not PORT capable can continue to use PIM in Datagram Mode. PORT capability is detected using a new PORT Capable PIM Hello Option.

When PORT is used, only incremental JPs are sent from downstream routers to upstream routers. As such, downstream routers do not generate periodic JPs for routes which RPF to a PORT-capable neighbor.

For Joins and Prunes, which are received over a TCP/SCTP connection, the upstream router does not start or maintain timers on the outgoing interface entry. Instead, it explicitly tracks downstream routers which have expressed interest. An interface is deleted from the outgoing interface list only when all downstream routers on the interface, no longer wish to receive traffic.

Because incremental JPs are sent over a TCP/SCTP connection, no Join suppression or Prune-Override of incremental JPs is possible on multi-access LANs. As a result, upstream routers, which receive an incremental Join or Prune that creates state, explicitly track all downstream nodes. Note, for point-to-point links there is no need for explicitly tracking downstream nodes.

There is no change proposed for the PIM JP packet format. However, for JPs sent over TCP/SCTP connections, no IP Header is included. The message begins with the PIM common header, followed by the JP message. See section Section 5 for details on the common header.

3. New PIM Hello Options

3.1. PIM over the TCP Transport Protocol

Option Type: PIM-over-TCP Capable

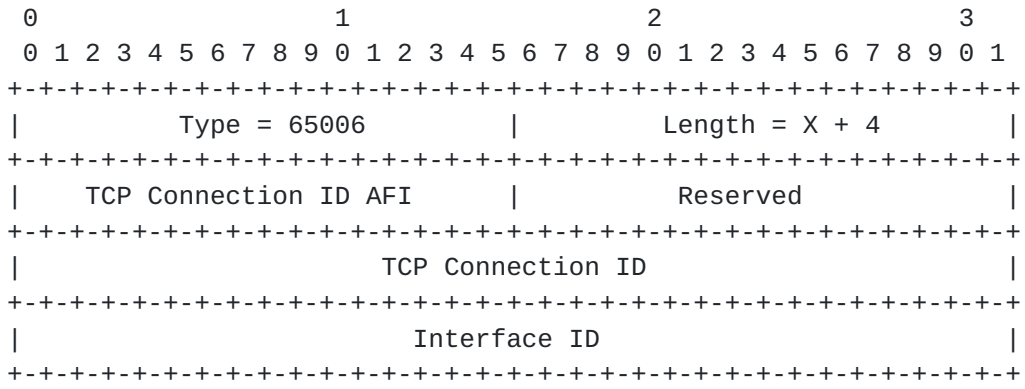

Allocated Hello Type values can be found in [HELLO-OPT].

When a router is configured to use PIM over TCP on a given interface, it MUST include the PORT Capable hello option in its Hello messages for that interface. If a router is explicitly disabled from using JP over TCP it MUST NOT include the PORT Capable hello option in its Hello messages. When the router cannot setup a TCP connection, it will refrain from including this option.

This option is only used when a physical or logical interface is a point-to-point, segmented multi-access LAN, a PMSI [MCAST-VPN], a point-to-point or point-to-multipoint GRE tunnel. In all other cases, such as multi-access LANs, Datagram Mode is used.

Implementation may provide a configuration option to enable or disable PORT functionality. We recommend that this capability be disabled by default.

Length: In bytes for the value part of the Type/Length/Value encoding. Where X is 4 bytes if IP AFI of value 1 is used and 16 bytes when IPv6 AFI of 2 is used [AFI].

TCP Connection ID AFI: The AFI value to describe the address-family of the address of the TCP Connection ID field.

Reserved: Set to zero on transmission and ignored on receipt.

TCP Connection ID: An IP or IPv6 address used to establish the TCP connection. When this field is 0, a mechanism outside the scope of this spec is used to obtain the addresses used to establish the TCP connection.

Interface ID: An Interface ID is used to associate the connection a JP message is received over with an interface which is added or removed from an oif-list. When unnumbered interfaces are used or when a single Transport connection is used for sending and receiving JP messages over multiple interfaces, the Interface ID is used convey the interface from JP message sender to JP message receiver. When a PIM router sets a locally generated value for the Interface ID in thie Hello TLV, it must send the same Interface ID value in all JP messages it is sending to the PIM neighbor.

3.2. PIM over the SCTP Transport Protocol

Option Type: PIM-over-SCTP Capable

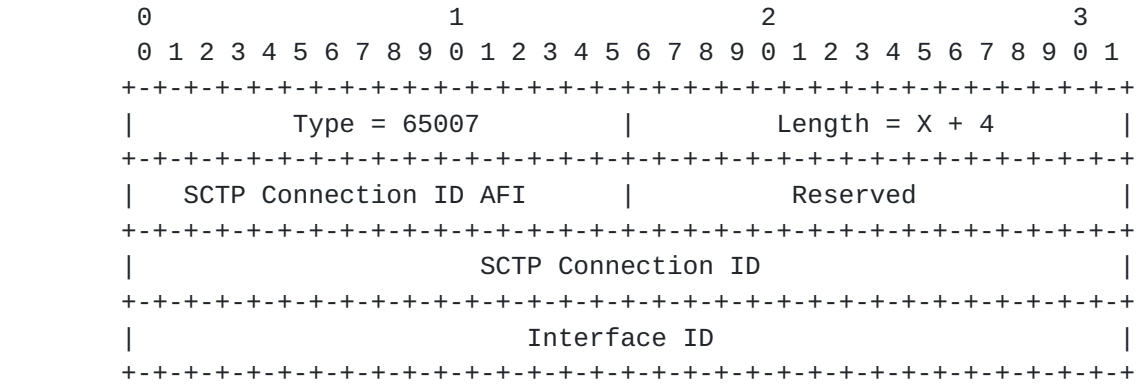

Allocated Hello Type values can be found in [HELLO-OPT].

When a router is configured to use PIM over SCTP on a given interface, it MUST include the PORT Capable hello option in its Hello messages for that interface. If a router is explicitly disabled from using JP over SCTP it MUST NOT include the PORT Capable hello option in its Hello messages. When the router cannot setup a SCTP connection, it will refrain from including this option.

This option is only used when an interface is point-to-point or when a multi-access LAN or MDT is segmented (also known as "Partitioned MDTs" in a non-broadcast multi-access (NBMA) mode. In all other

cases, such as general purpose multi-access LANs, Datagram Mode is used.

Implementation may provide a configuration option to enable or disable PORT functionality. We recommend that this capability be disabled by default.

Length: In bytes for the value part of the Type/Length/Value encoding. Where X is 4 bytes if IP AFI of value 1 is used and 16 bytes when IPv6 AFI of 2 is used [AFI].

SCTP Connection ID AFI: The AFI value to describe the address-family of the address of the SCTP Connection ID field.

Reserved: Set to zero on transmission and ignored on receipt.

SCTP Connection ID: An IP or IPv6 address used to establish the SCTP connection. When this field is 0, a mechanism outside the scope of this spec is used to obtain the addresses used to establish the SCTP connection.

Interface ID: An Interface ID is used to associate the connection a JP message is received over with an interface which is added or removed from an oif-list. When unnumbered interfaces are used or when a single Transport connection is used for sending and receiving JP messages over multiple interfaces, the Interface ID is used convey the interface from JP message sender to JP message receiver. When a PIM router sets a locally generated value for the Interface ID in thie Hello TLV, it must send the same Interface ID value in all JP messages it is sending to the PIM neighbor.

4. Establishing Transport Connections

Since this specification describes using Transport on point-to- point links or NBMA configured MDTs, a router knows when a Transport is established with the neighbor. When the Transport connection is not established, Datagram Mode is used. When the Transport connection becomes established Transport Mode is in effect where the router can suppress sending periodic JPs.

When a router receives a Hello from a neighbor it has not previously heard from, or the PORT-Capable Option is included in a Hello that was not previously included by an existing neighbor, the router will attempt to establish a Transport connection with the neighbor. When the router is using TCP it will compare the IP address it uses to send Hellos on the interface with the IP address the neighbor is using to send Hellos. The router with the lower IP address will do an active Transport open to the neighbor address. The higher IP addressed neighbor will do a passive Transport open. When the router is using SCTP, the IP address comparison not be done since the SCTP protocol can handle call collision.

The PIM router that performs the active open initiates the connection with a locally generated source transport port number and a well-known destination transport port number. The PIM router that performs the passive open listens on the well-known local transport port number and does not qualify the remote transport port number. See Section 5 for well-known port number assignment for PORT.

When a Transport connection goes down, Join or Prune state that was sent over the Transport connection is still retained. The neighbor should not be considered down until the neighbor timer has expired. This allows routers to do a control-plane switchover without disrupting the network. If a Transport connection is reestablished before the neighbor timer expires, the previous state is intact and any new JP messages sent cause state to be created or removed (depending on if it was a Join or Prune). If the neighbor timer does expire, only the upstream router, that has oif-list state, to the expired downstream neighbor will need to clear state. A downstream router, when an upstream neighboring router has expired, will simply RPF to a new neighbor where it would trigger JP messages like it would in [RFC4601]. It is required of a PIM router to clear it's neighbor table for a neighbor who has timed out due to neighbor holdtime expiration.

When a router is in Datagram Mode with a neighbor and has been sending periodic JP messages to it and then the Transport connection has been established to the neighbor, there is no requirement for the downstream router to send JP messages to the upstream neighbor. The

upstream router can keep the state maintained from the Datagram Mode creation. However when a router is in Transport Mode with a neighbor and moves to Datagram Mode because the transport connection went down (and several attempts to reestablish the transport connection fail), the router cannot be sure that all the JP data was received by the neighbor. Therefore, it is required to send a full set of JP messages to refresh or re-create state in the upstream neighbor.

An upstream neighbor does have the responsibility of removing the timer-activated timeout of an oif-list entry. When a Transport connection is established, the timer-activated timeout is disabled. When a Transport connection goes down, the timer-activated timeout for an oif-list is enabled. Both the upstream and downstream routers stay in sync based on the state of the Transport connection. If the upstream router has timer-activated timeout on oif-lists, the downstream router will be sending periodic JPs. Otherwise, the downstream router suppresses sending periodic JPs because it assumes the upstream router has disabled the timer-activated timeout of oif-list entries the downstream router has previously joined.

4.1. TCP Connection Maintenance

TCP is designed to keep connections up indefinitely during a period of network disconnection. If a PIM-over-TCP router fails, the TCP connection may stay up until the neighbor actually reboots, and even then it may continue to stay up until you actually try to send the neighbor some information. This is particularly relevant to PIM, since the flow of JPs might be in only one direction, and the downstream neighbor might never get any indication via TCP that the other end of the connection isn't really there.

Most applications using TCP want to detect when a neighbor is no longer there, so that the associated application state can be released. Also, one wants to clean up the TCP state, and not keep half-open connections around indefinitely. This is accomplished by using PIM Hellos and by not introducing an application-specific or new PIM keep-alive message. Therefore, when a GENID changes from a received PIM Hello message, and a TCP connection is established or attempting to be established, the local side will tear down the connection and attempt to reopen a new one for the new instance of the neighbor coming up.

When PORT capable routers come up and try to establish transport connections with their neighbors, but cannot for some reason, after 3 attempts to do so, the router should go into datagram mode and not advertise the PORT Hello option anymore. Operator intervention is required to restart the process after the problem is found.

4.2. Transitional Periods

There may be transitional periods when a router receives, from a given neighbor, both datagram JP messages and JP messages sent over a transport connection. When this happens, a transport connection to a particular neighbor is established, and as long as it remains established, the router MUST ignore PIM messages sent in Datagram Mode from that neighbor. Otherwise, the datagram messages could get out of order with respect to the transport messages, and the router could end up in an erroneous state of pruning joined state or joining pruned state which it is unable to recover from as long as the transport connection stays up.

4.3. On-demand versus Pre-configured Connections

Transport connections could be established when they are needed or when a router interface to other PIM neighbors has come up. The advantages of on-demand Transport connection establishment are the reduction of router resources. Especially in the case where there is no need for n^2 connections on a network interface or MDT tunnel. The disadvantages are deciding what to do when a JP message needs to be sent and a Transport connection is not established yet. An implementation can either send a Datagram Mode JP or queue the JP to be sent as a Transport Mode JP after the Transport connection is established.

If a router interface has become operational and PIM neighbors are learned from Hello messages, at that time, Transport connections may be established. The advantage is that a connection is ready to transport data by the time a JP messages needs to be sent. The disadvantage is there can be more connections established than needed. This can occur when there is a small set of RPF neighbors for the active distribution trees compared to the total number of neighbors. Even when Transport connections are pre-established before they are needed, a connection can go down and an implementation will have to deal with an on-demand situation.

Therefore, this specification recommends but does not mandate the use of on-demand Transport connection establishment.

4.4. Possible Hello Suppression Considerations

This specification indicates that a Transport connection cannot be established until a Hello message is received. One reason for this is to determine if the PIM neighbor supports this specification and the other is to determine the remote address to use to establish the Transport connection.

There are cases where it is desirable to suppress entirely the transmission of Hello messages. In this case, it is outside the scope of this document on how to determine if the PIM neighbor supports this specification as well as an out-of-band (outside of the PIM protocol) method to determine the remote address to establish the Transport connection.

4.5. Avoiding a Pair of Connections between Neighbors

To ensure there are not two connections between a pair of PIM neighbors, the following set of rules must be followed. Let A and B be two PIM neighbors where A's IP address is numerically smaller than B's IP address, and each is known to the other as having a potential PIM adjacency relationship.

At node A:

- If there is already an established TCP connection to B, on the PIM-over-TCP port, then A MUST NOT attempt to establish a new connection to B. Rather it uses the established connection to send JPs to B. (This is independent of which node initiated the connection.)

- If A has initiated a connection to B, but the connection is still in the process of being established, then A MUST refuse any connection on the PIM-over-TCP port from B.

- At any time when A does not have a connection to B which is either established or in the process of being established, A MUST accept connections from B.

At node B:

- If there is already an established TCP connection to A, on the PIM-over-TCP port, then B MUST NOT attempt to establish a new connection to A. Rather it uses the established connection to send JPs to A. (This is independent of which node initiated the connection.)

- If B has initiated a connection to A, but the connection is still in the process of being established, then if A initiates a connection to, B MUST accept the connection initiated by A and must release the connection which it (B) initiated.

5. Common Header Definition

It may be desirable for scaling purposes to include JP messages from different PIM protocol instances to be sent over the same Transport connection. Also, it may be desirable to have a set of JP messages for one address-family sent over a Transport connection that is established over a different address-family network layer.

To be able to do this we need a common header that is inserted and parsed for each PIM JP message that is sent on a Transport connection. This common header will provide both record boundary and demux points when sending over a stream protocol like Transport.

Each JP message will have in front of it the following common header in Type/Length/Value format. And multiple different TLV types can be sent over the same Transport connection.

To make sure PIM JP messages are delivered as soon as the TCP transport layer receives the JP buffer, the TCP Push flag will be set in all outgoing JP messages sent over a TCP transport connection.

PIM messages will be sent using destination TCP port number 8471. When using SCTP as the reliable transport, destination port number 8471 will be used. See Section 11 for IANA considerations.

If the buffer length of the received TLV message is less than what is encoded in the TLV Length field, the entire TLV encoded message is ignored and a error message is logged. Likewise, if the received buffer length left to process at each record parsing level, is less than the JP Message Length, the rest of the message is malformed and not processed.

Each JP message that has passed the length checks above, contained in the TLV encoding, will be error checked individually. This includes a bad PIM checksum, illegal type fields, or illegal addresses. If any parsing errors occur in a single JP message, it is skipped over and not processed but other JP message records in the TLV are still parsed and processed.

The current list of defined TLVs are:

IPv4 JP Message

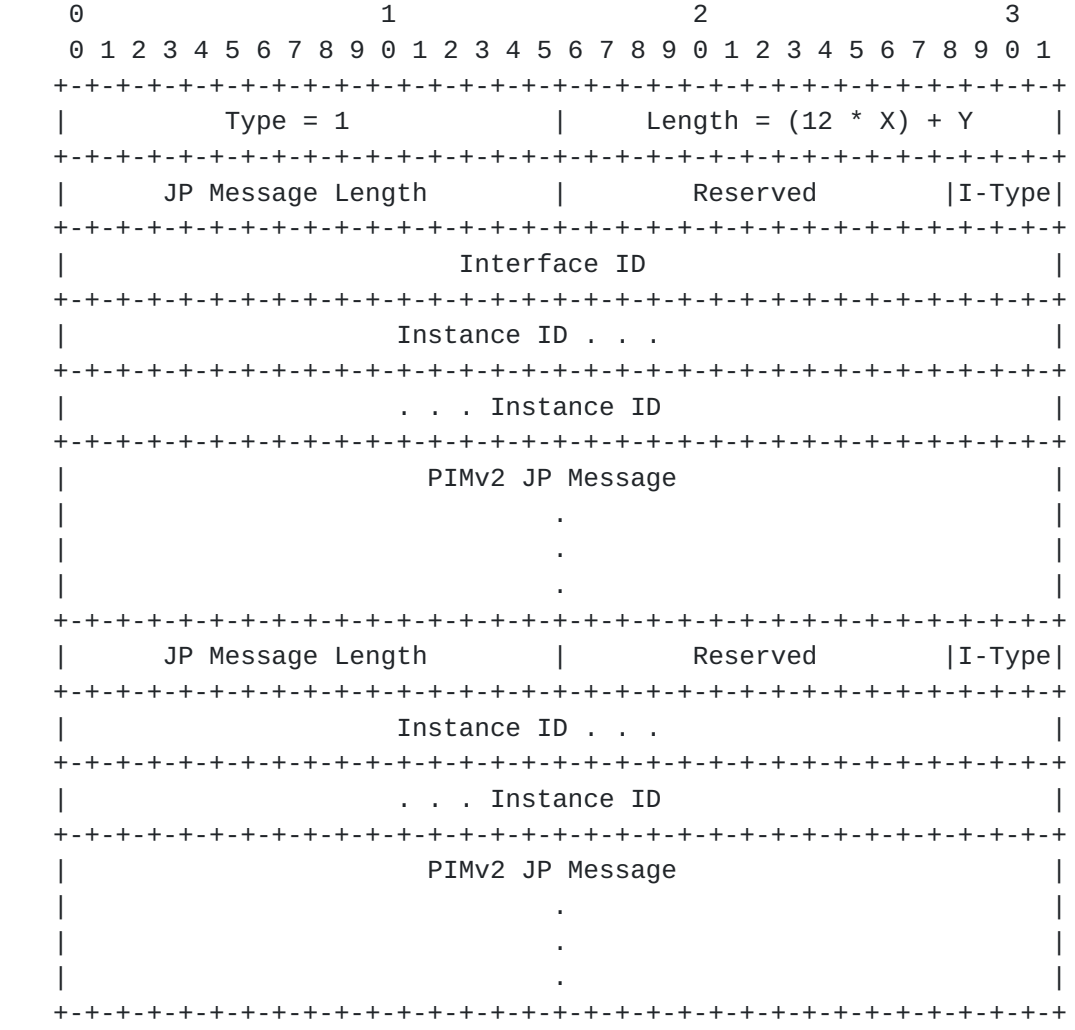

The IPv4 JP common header is used when a JP message is sent that has all IPv4 encoded addresses in the PIM payload.

Length: In bytes for the value part of the Type/Length/Value encoding. Where there are 12 bytes per JP message (where X above is the number of JP messages contained) enclosed in one transmission plus Y which is the sum of each "JP Message Length" field that appears in the transmission.

I-Type: Defines the encoding and semantics of the Instance ID field. This is not specified in this specification.

Interface ID: This is the Interface ID from the Hello TLV, defined in this specification, the PIM router is sending to the PIM neighbor. It indicates to the PIM neighbor what interface to associate the JP Join or Prune with.

Instance ID: This can be a VPN-ID. This field could also be a BGP Route Target (RT) or BGP Route Distinguisher (RD) as defined in [RFC4364]. Not specified in this specification.

Reserved: Set to zero on transmission and ignored on receipt.

JP Message Length: The number of bytes that follow which make up the PIMv2 JP message.

PIMv2 JP Message: PIMv2 Join/Prune message and payload with no IP header in front of it. As you can see from the packet format diagram, multiple JP messages can go into one TCP/SCTP stream from the same or different Instance IDs.

IPv6 JP Message

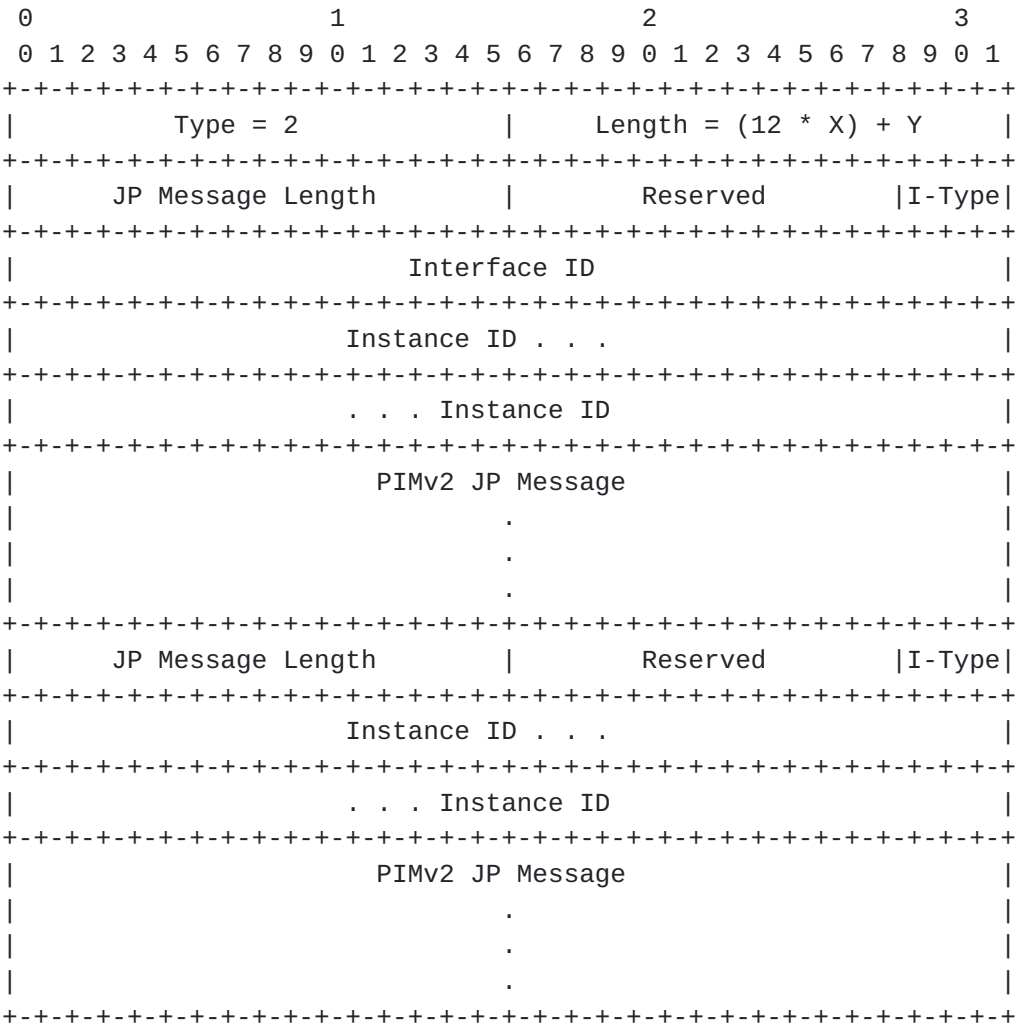

The IPv6 JP common header is used when a JP message is sent that has all IPv6 encoded addresses in the PIM payload.

Length: In bytes for the value part of the Type/Length/Value encoding. Where there are 12 bytes per JP message (where X above is the number of JP messages contained) enclosed in one transmission plus Y which is the sum of each "JP Message Length" field that appears in the transmission.

I-Type: Defines the encoding and semantics of the Instance ID field. This is not specified in this specification.

Interface ID: This is the Interface ID from the Hello TLV, defined in this specification, the PIM router is sending to the PIM neighbor. It indicates to the PIM neighbor what interface to associate the JP Join or Prune with.

Instance ID: This can be a VPN-ID, BGP Route Target (RT) or BGP Route Distinguisher (RD). Not specified in this specification.

Reserved: Set to zero on transmission and ignored on receipt.

JP Message Length: The number of bytes that follow which make up the PIMv2 JP message.

PIMv2 JP Message: PIMv2 Join/Prune message and payload with no IP header in front of it. As you can see from the packet format diagram, multiple JP messages can go into one TCP/SCTP stream from the same or different Instance IDs.

6. Join/Prune Processing

When a PORT neighbor transitions to using Transport Mode, the downstream router sends JP messages for existing routes that RPF to the neighbor over the Transport connection. In addition, periodic JP messages are stopped and only incremental JPs are sent thereafter.

A router which has a Transport connection established MUST send and receive JP messages over the Transport session to that given peer as well as accept and process native JP messages as described in [RFC4601].

When a Transport connection is established for a newly discovered neighbor, the downstream router triggers JP messages for its existing state. This is to allow the upstream router to build state it may previously not had. If state had existed due to a Native JP, the expiration timer would have been started. Now it can be stopped because the state is being sent incrementally over the Transport connection.

When a Transport connection goes down to a given neighbor, the downstream router does not have to trigger native JP messages. It can wait for its next periodic interval to send a native JP messages. When the upstream router receives the native JP message, it will start the expiration timer for the oif associated with the state from the JP message.

Note, since JP messages are sent over a Transport connection, no Prune Override or Join Suppression are possible for these messages.

7. Outgoing Interface List Explicit Tracking

Since this specification indicates the use of TCP/SCTP for PIM JP messages over point-to-point or NBMA type links, explicit tracking can be achieved by tracking only oif-list state and not per-neighbor per oif-list state. This is true for segmented LANs and in segmented MDT/PMSI environments.

By using explicit tracking of oifs, the router tracks all downstream neighbors which have expressed interest in a route on a given interface. The list of tracked routers is one of the checks used to determine whether traffic needs to be forwarded on a given interface or not.

For (*,G) and (S,G) routes, the router starts forwarding traffic on an interface when a Join is received from a neighbor on such an interface. This is tracking the oif to the neighbor. When the neighbor sends a Prune, the interface is removed and forwarding of traffic stops on the interface.

When all interfaces are removed from the oif-list, the route entry can be removed.

For (S,G,R) routes, typically is tracking Prune state on the shared tree. One at least one downstream neighbor sends a Prune over a Transport connection, the (S,G,R) state is create with a empty outgoing interface list. If a subsequent JP is received over a Transport connection which has (*,G) in the join-list and does not have (S,G,R) in the prune-list, the upstream router will add the interface the JP message was received on to the oif-list. And oif-list based explicit tracking will occur just like in the (*,G) and (S,G) route case above.

The only difference in the (S,G,R) route case, is that when the outgoing interface is pruned, the entry must stay in the route table or else forwarding will occur on the interfaces for the (*,G) entry. Therefore, explicit tracking for Prunes must be provided. Only when the (S,G,R) oif-list interfaces match the interfaces in the (*,G) can the (S,G,R) route be removed.

8. Multiple Instances and Address-Family Support

Multiple instances of the PIM protocol may be used to support multiple VPNs or within a VPN to support multiple address families. Multiple instances can cause a multiplier effect on the number of router resources consumed. To be able to have an option to use router resources more efficiently, muxing JP messages over fewer Transport connections can be performed.

There are two ways this can be accomplished, one using a common header format over a TCP connection and the other using multiple streams over a single SCTP connection.

Using the Common Header format described previously in this specification, using different TLVs, both IPv4 and IPv6 based JP messages can be encoded within a Transport connection. Likewise, within a TLV, multiple occurrences of JP messages can occur and are tagged with an instance-ID so multiple JP messages for different VPNs can use a single Transport connection.

When using SCTP multi-streaming, the common header is still used to convey instance information but an SCTP association is used, on a per-VPN basis, to send data concurrently for multiple instances. When data is sent concurrently, head of line blocking, which can occur when using TCP, is avoided.

9. Miscellany

No changes expected in processing of other PIM messages like PIM Asserts, Grafts, Graft-Acks, Registers, and Register-Stops. This goes for BSR and Auto-RP type messages as well.

This extension is applicable only to PIM-SM, PIM-SSM and Bidir-PIM. It does not take requirements for PIM-DM into consideration.

10. Security Considerations

Transport connections can be authenticated using HMACs MD5 and SHA-1 similar to use in BGP [RFC4271] and MSDP [RFC3618].

When using SCTP as the transport protocol, [RFC4895] can be used, on a per SCTP association basis to authenticate PIM data.

11. IANA Considerations

This specification requests IANA to allocate a TCP port number and a SCTP port number for the use of PIM-Over-Reliable-Transport.

12. Acknowledgments

The authors would like to give a special thank you and appreciation to Nidhi Bhaskar for her initial design and early prototype of this idea.

Appreciation goes to Randall Stewart for his authoritative review and recommendation for using SCTP.

Thanks also goes to the following for their ideas and commentary review of this specification, Mike McBride, Toerless Eckert, Yiqun Cai, Albert Tian, Suresh Boddapati, Nataraj Batchu, Daniel Voce, John Zwiebel, Yakov Rekhter, and Lenny Giuliano.

A special thank you goes to Eric Rosen for his very detailed review and commentary. Many of his comments are reflected as text in this specification.

13. References

13.1. Normative References

[RFC0761] Postel, J., "DoD standard Transmission Control Protocol", RFC 761, January 1980.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC4895] Tuexen, M., Stewart, R., Lei, P., and E. Rescorla, "Authenticated Chunks for the Stream Control Transmission Protocol (SCTP)", RFC 4895, August 2007.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

13.2. Informative References

[AFI] IANA, "Address Family Indicators (AFIs)", ADDRESS FAMILY NUMBERS http://www.iana.org/numbers.html, February 2007.

[HELLO-OPT] IANA, "PIM Hello Options", PIM-HELLO-OPTIONS per RFC4601 http://www.iana.org/assignments/pim-hello-options, March 2007.

[MCAST-VPN] Rosen and Aggarwal, "Multicast in MPLS/BGP VPNs", Internet Draft draft-ietf-l3vpn-2547bis-mcast-05.txt, July 2007.

Authors' Addresses

Dino Farinacci
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: dino@cisco.com

IJsbrand Wijnands
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: ice@cisco.com

Apoorva Karan
cisco Systems
170 Tasman Drive
San Jose, CA
USA

Email: apoorva@cisco.com

Arjen Boers
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: aboers@cisco.com

Maria Napierala
AT&T Labs
200 Laurel Drive
Middletown, New Jersey 07748>
USA

Email: mnapierala@att.com

Full Copyright Statement

Intellectual Property

Acknowledgment

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