

Reliable PIM Registers

(draft-anish-reliable-pim-register)

Anish Peter (anishp@juniper.net)

Robert Kebler (rkebler@juniper.net)

Vikram Nagarajan (vikramna@juniper.net)

July-2015

PIM-WG, IETF-92

PIM Registers – How it is today

- When a First-Hop-Router (FHR) gets native multicast traffic, this traffic would be tunneled as PIM control packets to RP. These are PIM registers.
- PIM registers serve two purposes
 - It helps FHR to inform that it is getting traffic for a given (source, group).
 - It helps in avoiding initial packet loss.
- Each individual S, G is “ack’ed” with Register Stop
 - Register-stops prevent FHR from sending data Registers
- Subsequent to this, NULL-Registers are used to maintain the aliveness of the source

Observations

- PIM Null-Register
 - Is soft-state based
 - Packet format does not allow state refresh for multiple flows in the same message
- PIM register-stop messages inherit all the problems in Null-Register messages
- In the FHR, if Register-Stop times-out, its expected to resort to Packet-Register's (RFC defaults to 60+5s).
 - This could happen even if one RS-message gets dropped.

Reliable Registers

- Reliable-Registers would support a reliable transport between FHR and RP
- Create a “targeted” adjacency between FHR and RP
 - These routers form adjacency.
 - Sends PIM Hellos with normal Hello Options to advertise capabilities
 - Can use Anycast-RP address to find closest RP
- Use TCP/SCTP
 - Some of the same encoding as RFC 6559 (PIM PORT)
 - Reliability and Flow control
 - New messages created to notify of new active source
- FHR sends message to RP to add/remove active sources

Targeted Hellos

- As per present spec, PIM hellos are link-level
- This draft extends that to supported pim neighbors over multiple hops reached via its known unicast address
- FHR router upon learning an RP (could be anycast-RP) address would transmit targeted hellos
- RP could respond to those targeted hellos
- From these hellos RP and FHR would learn the port capability and could start with reliable-registers

Targeted hellos (cont.)

- RP when responding to targeted hello would use its unique address and would add its other address (including anycast addresses) in its secondary address TLV's.
- New TLV would be added for targeted neighbor properties/capabilities.
- Hellos will have TLV's as specified by PORT for reliable connection setup

Connection Setup

- Based on hello FHR and RP would learn its peers PORT capabilities.
- Once adjacency is formed, RP would connect to FHR to form the reliable connection.
- PORT Keep-alive could be used to maintain aliveness of session.

Hard-State Register Messages

- Stream-Register Message send by FHR
- Similar to a NULL-register
- FHR can withdraw the register when it finds doing so is appropriate (KAT trigger)
- To withdraw, set withdraw flag in the same register message

Anycast RP

- FHR would discover nearest RP by means of sending targeted hellos to anycast address.
- Reliable full mesh connection among the anycast RP-Set.
- Redistribution of source information
 - RP's would transmit stream-register messages received from FHR to all the other any-cast peers.
 - When a new anycast-RP connection is setup, an RP would send to the peers all the stream-registers it had learned from FHR.

Management Considerations

- Only configuration needed is an enable/disable knob for reliable register (No need configure peers)
- Incremental deployment is possible
- Feature support needed only on RP and FHR

Thank You

Opinions
&
Clarifications