

draft-venaas-pim-port-pfm-00

PIM Flooding Protocol over Reliable Transport

IETF 104

Stig Venaas, stig@cisco.com

Balaji Ganesh, balagane@cisco.com

Kesavan Thiruvenkatasamy, kethiruv@cisco.com

Ramakrishnan Chokkanathapuram, ramaksun@cisco.com

The idea

- PFM relies on sending periodic updates. A first hop router would every 60s or so send a full update with all active sources.
- We want to send it often so that state can be updated in case of packet loss.
- If the signaling was reliable we would rarely need to send full updates.
- PFM is hop-by-hop. Can we utilize PIM PORT which offers reliable transport between neighbors?

Solution 1/2

- Establish PORT connections between neighbors according to PIM PORT RFC.
- We define a new PORT message for sending PFM updates. They contain a major and a minor sequence number, and the PFM message itself according to PFM RFC.
- We define a new PORT message to request an update from a neighbor. It contains a list of originators and which sequence number/version was last received from each of them. A neighbor will respond by sending everything it has that is newer from the specified originators, plus everything from other originators.

Solution 2/2

- When a router receives PFM messages newer than what it already has, it updates its own information and forwards what is new to its neighbors.
- When a PORT connection is established to a neighbor, a router would ask the neighbor if it has anything newer than what it currently has. Whenever the neighbor learns anything new after that, it is assumed that it will tell us.
- In this way we make all the routers have the latest information.

Sequence numbers

- We use a major sequence number that is increased each time a full update is sent, and a minor number that is increased each time a triggered delta is sent.
- A full update has minor sequence number 0.
- When a router receives a full update it replaces all previous information from the originator with the full update. When it receives a minor update, it stores that while keeping all existing info.
- How to ensure that the major sequence number is increased across restarts? Derive it from the clock?
 - Add a mechanism to find what is the largest major number stored in the domain? May borrow ideas from IGPs.