IS-IS Flooding over TCP

(draft-hsmit-lsr-isis-flooding-over-tcp-00)

Henk Smit

Gunter Van de Velde

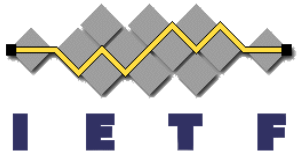

IETF 103, 6 November 2018
Bangkok, Thailand

What is it about?

- Proposal to use TCP for IS-IS flooding
  - Remove from IS-IS the task of packet pacing and reliable flooding
- New TLV in point-to-point IIHs
  - Signal the intent of a router to do flooding over TCP
- New small header to encapsulate IS-IS PDUs in the TCP byte-stream (similar to BGP)

Why?

- General scaling factors of IS-IS
  - Number of adjacencies (little can be done)
  - Flooding of LSPs (improving this has positive impact upon IS-IS scaling)
  - Dijkstra's SPF (can not be improved upon easily)
- Current scaling limitations of IS-IS flooding
  - Packet pacing and throughput
    - Historical 30msec between updates (if 1000 nodes -> 30sec on updates!!!!)
  - Reliable flooding on point-to-point interfaces
  - Unreliability of CSNPs
    - 91 LSPs in one CSNP (if miss/drop one, then 91 LSPs are exchanged)

Solution: Reliable Flooding using TCP

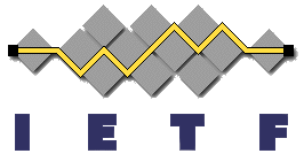

- Remove packet pacing and reliability from base IS-IS
- Initially flooding over TCP has to be negotiated
- This means introduction of new TLV in IIH
- TLV must be supported/introduced at each side
- Only for p2p interface
- TLV contain:
  - TLV type 1: TCP port#
  - TLV type 2: IPv4 address
  - TLV type 3: IPv6 address
  - (initially that is enough)
- TCP uses bytestream. In BGP a to keep sync a marker is used, and so is a marker proposed for ISISoTCP

New behavior: Establish IS-IS TCP Flooding

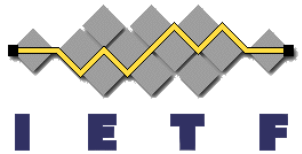

- Initially router look for new TLV in IIH
- If both router find
  - then lowest System-id setup TCP session
  - Both router keep sending IIH over p2p link
- After TCP is established we verify the TCP session
  - Each router will send initial IIH over TCP session
  - Used to identify/authenticate remote router
  - Set SRM bits
  - Send CSNPs, process and exchange LSP's
  - No more retransmission of LSPs or verification using SNP's because it is in hands of TCP

New behavior: During IS-IS TCP Flooding

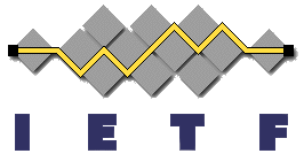

- When router receive LSP, verify if it is new or not
  - If same version: router does nothing
  - If older version: router set SRM-bit and return newer LSP to peer
  - If newer version: set SRM-bit on all interfaces, except the one it received original LSP upon
  - Receiving router does not set SSN-bit and does not send an acknowledgement (PSNP)
- Periodically, or event driven, the router will check its LSDB for LSPs with the SRM-bit set
  - it will send as many of those LSPs to neighbors, via TCP
  - no packet-pacing, TCP take care of flow-control, no retransmission by IS-IS, no PSNP, … life is simple….

Some additional thoughts

- Flooding over TCP is only done on point-to-point interfaces
- Unnumbered interfaces and reachability of the interface ip-address
- Multiple levels of hierarchy on one interface
- Downsides of using TCP for IS-IS flooding
- What to do if the TCP connection breaks
- What to do if the TCP connection can not be set up

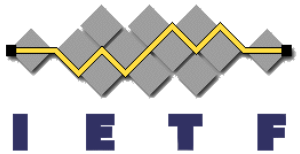

THANK YOU!