draft-ietf-pim-multicast-lessons-learned-06

Dino, Lenny, Mike, Nils
Montreal 124

Current protocols/topics addressed

3. Lessons learned about IP Multicast **Protocols** over the last 30 years
   - 3.1. DVMRP
   - 3.2. MOSPF
   - 3.3. Shared vs Source Trees
   - 3.4. IGMP
   - 3.5. Data Driven State Creation and RPF
   - 3.6. MSDP
   - 3.7. MPLS MVPNs
   - 3.8. SD and SDR
   - 3.9. All or Nothing Problem
   - 3.10. AMT and TreeDN
   - 3.11. Network Based Source Discovery
   - 3.12. Dynamic Multicast Group Address Allocation
   - 3.13. Reliable Multicast
   - 3.14. Premature Optimization
   - 3.15. Kernel vs User Space
   - 3.16. 802.11
   - 3.17. RPT-to-SPT Switchover Thresholds

RPT-to-SPT Switchover Thresholds

A common implementation feature in PIM-SM is the ability to configure a bandwidth or packet rate threshold that triggers the switch from the shared tree (RPT) to the source tree (SPT). This is typically implemented as a configurable parameter such as "switch to SPT when exceeding X kbit/s of traffic for this (S,G)".

While the default behavior in most implementations is to switch to SPT immediately upon receiving the first data packet, this configurable threshold feature was designed to give operators control over when to build source trees for high-bandwidth flows. However, it has proven to be operationally challenging. The challenges arises from:

* The need to accurately measure traffic rates per (S,G) entry
* The state machinery required to trigger and manage the transition
* Ensuring consistent behavior across different router implementations
* The potential for flapping between RPT and SPT if thresholds are set too close to typical traffic patterns

While this feature provides granular control, many operators find the default immediate SPT-switchover behavior sufficient for their needs. The operational experience has shown that the costs of implementing and managing the threshold feature often outweigh the benefits for many deployment scenarios. Operationally, most networks set the SPT threshold to either 0 (immediate switchover) or infinity (no switchover).

This experience reinforces the broader lesson that features requiring per-flow state and measurement can introduce significant complexity with limited operational benefit. The fact that many deployments operate successfully with the simpler immediate SPT-switchover behavior suggests that the additional complexity of configurable thresholds was perhaps unnecessary for most use cases.

Next steps

- Continue to make slow and steady progress, feedback appreciated.
- What else?