

A Framework for Computed Multicast applied to MPLS based Segment Routing

draft-allan-spring-mpls-multicast-framework-01

Dave Allan, Ericsson

Jeff Tantsura

IETF96

Progress since IETF 95

- 01 version simply includes some clarifications
- Drafts describing required IGP extensions published:
 - Draft-allan-ospf-spring-mcast-00
 - Draft-allan-isis-spring-mcast-00

Quick Recap: What is the draft about?

- The application of computation to determining the routing of multicast segments in an MPLS based SR network, and how unicast tunnels can be used as part of multicast tree construction to minimize overall network state
- The draft describes
 - Terminology
 - Overall approach
 - Loose and Specified multicast distribution trees
 - Algorithm
 - FIB installation procedures

Approach

- The draft postulates an architecture whereby multicast trees are a hybrid of roots, leaves, and replication points interconnected with unicast tunnels, with the routing of the tree and location of the nodes that need to install state is determined entirely from information in the IGP
 - Which has been augmented to add TLVs for multicast interest
- This provides multiple benefits
 - Minimized messaging to converge the network
 - Reduced dataplane state
 - Minimized bandwidth requirements
 - Unicast recovery addresses most failures
 - Re-use of the existing MPLS dataplane

Tree Pruning

- Computed trees are determined by a series of pruning steps applied to the shortest path tree from a root to the set of leaves
- Two classes of prune
 1. Those that if they fully resolve the tree are known to produce a minimum cost tree
 - This will sort out 97%+ of leaves
 2. If these do not completely resolve the tree (unique shortest path to the root from every leaf), then we start to apply “guesses”, and audit the tree for correctness at the end
 - With good “guesses” only a tiny fraction of the remainder require “fixing”

Loose and Specified Trees

- A loose tree is composed of a single multicast segment (with a SID), where only the root and the leaves have been specified in the IGP
 - The routing of the tree is wholly computed based on the current network topology
- A specified tree is composed of a concatenation of multicast segments where the roots, waypoints and leaves have been specified in the IGP
 - The routing of individual segments is still computed
 - The routing of an MDT can then be specified to an arbitrary level of granularity
 - A unique SID per segment ensures the resulting hybrid of pinned and computed components is loop free, even if not planar

IGP changes

- For both ISIS and OSPF, the approach requires the specification of 3 TLVs...
 1. A TLV that advertises compute capability and the algorithm used
 - Extensible via OUI/algorithm tuple
 2. SID/Group binding TLV
 - How membership in ASM is advertised in the IGP
 - <group><SID><Transmit/receive interest>
 3. Pinned tree descriptor TLV
 - Expressed as a list of SID cross connects
 - <source SID><cross connecting SID><destination SID>

OSPFv2 Specifics

- SRM Compute Capability TLV
 - Proposed as a TLV to the RFC 7770 Router Information LSA
- SRM SID Multicast Group Binding Sub-TLV
 - Proposed sub-TLV as part of the RFC 7684 OSPFv2 Extended Prefix TLV as a sub-TLV from OSPFv2 Extended Prefix Opaque LSA
- SRM Pinned Tree Descriptor sub-TLV
 - Proposed sub-TLV is part of the RFC 7684 OSPFv2 Extended Prefix TLV

ISIS Specifics

- SRM Compute Capability TLV
 - Proposed sub-TLV of TLV 144
- SRM SID Multicast Group Binding sub-TLV
 - Proposed as a sub-TLV from the 135, 235, 236 and 237 registry
- SRM Pinned Tree Descriptor sub-TLV
 - Proposed as a sub-TLV from the 135, 235, 236 and 237 registry

Next Steps

- Now that the IGP drafts have been published, there is a complete description of the solution available for review
 - Comments, suggestions, questions and criticism all welcome

Questions?