

Advertising Per-Algorithm Label Blocks

draft-bowers-spring-adv-per-algorithm-label-blocks-01

Chris Bowers cbowers@juniper.net

Pushpasis Sarkar psarkar@juniper.net

Hannes Gredler hannes@juniper.net

SPRING Working Group

IETF93 Prague

Computing locally significant labels for shortest path next-hops

$$\text{SPF_Label}(X,D) = \text{Label_Block}(X) + \text{Node_Index}(D)$$

- $\text{Label_Block}(X)$ is the label block advertised by X
- D is the destination node
- $\text{SPF_Label}(X,D)$ is the value of the label that neighbors need to apply to a packet so that X will forward the packet along the shortest path next-hop to D .

Computing locally significant labels for next-hops corresponding to other algorithms

Option 1: Per-algorithm node index

$$\text{Label}(X,D,A) = \text{Label_Block}(X) + \text{Node_Index}(D,A)$$

Option 2: Per-algorithm label block

$$\text{Label}(X,D,A) = \text{Label_Block}(X,A) + \text{Node_Index}(D)$$

- $\text{Label}(X,D,A)$ is the value of the label that neighbors need to apply to a packet so that X will forward the packet along the next-hop to D computed using destination-based forwarding algorithm A .

Option 1: Per-algorithm node index

$$\text{Label}(X,D,A) = \text{Label_Block}(X) + \text{Node_Index}(D,A)$$

- Each destination node is assigned a globally unique node index for each algorithm (a per-algorithm node index)
- The node index corresponding to a given algorithm and destination node is used to look up the corresponding locally significant label in the single label block advertised by the node X.
- Extensions specified in draft-ietf-isis-segment-routing-extensions and draft-ietf-ospf-segment-routing-extensions use option 1.

Option 2: Per-algorithm label block

$$\text{Label}(X,D,A) = \text{Label_Block}(X,A) + \text{Node_Index}(D)$$

- Each destination node is only assigned a single node index (the same node index as for the SPF algorithm)
- Each node X advertises a unique label block for each algorithm (a per-algorithm label block)
- The single node index corresponding to a destination node is used to look up the corresponding locally significant label in label block advertised by the node X for the given algorithm.
- draft-bowers-spring-adv-per-algorithm-label-blocks proposes extensions to support option 2.

Why is option 2 (per-algorithm label blocks) better than option 1 (per-algorithm node indexes)?

- Using SR to distribute labels for shortest path routes
- Advantage over LDP that any node (not just neighbors of X) can determine the FEC-label bindings distributed by node X.
 - Pro: no need for targeted LDP sessions
 - Con: need to assign and maintain tightly packed, domain-unique node index values
- Generalizing SR to distribute labels for other algorithms
- Adding a new algorithm with Option 1
 - Advertise a larger label block
 - Assign and advertise a set of node index values (unique and tightly packed wrt previously assigned node index values)
- Adding a new algorithm with Option 2
 - Advertise a new label block for the algorithm

