

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 5, 2015

C. Bowers
P. Sarkar
H. Gredler
Juniper Networks
April 3, 2015

Advertising Per-Algorithm Label Blocks
draft-bowers-spring-adv-per-algorithm-label-blocks-00

Abstract

Segment routing uses globally-known labels to accomplish destination-based forwarding along shortest paths computed using Dijkstra's algorithm with IGP metrics. This draft discusses how to use segment routing to accomplish destination-based forwarding along paths computed using other algorithms and metrics. In particular, the draft contrasts two different options for associating labels with different algorithms for computing forwarding next-hops.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 5, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Destination-based forwarding in segment routing using other algorithms	3
3.	Practical implications of using the per-algorithm node index option	4
4.	Advantages of using the per-algorithm label block option . .	5
5.	IANA Considerations	5
6.	Management Considerations	5
7.	Security Considerations	5
8.	Acknowledgements	6
9.	Informative References	6
	Authors' Addresses	6

[1.](#) Introduction

[I-D.ietf-spring-segment-routing] describes the segment routing architecture. In segment routing, LDP-like forwarding behavior along shortest paths is achieved using globally-known node labels. Globally-known node labels can be distributed in one of two ways. Each router can directly advertise a globally unique node label in the IGP. Or each router can advertise a globally unique node index value as well as a locally assigned label block, allowing any router in the IGP area to determine the mapping of locally-assigned label to globally unique node index for any other router in the area.

This draft focuses on the case where the combination of global node index and local label block is used to distribute locally-significant labels. Figure 1 shows the formula used to determine the locally-significant label values corresponding to shortest path forwarding next-hops computed using Dijkstra's shortest path first algorithm with IGP metrics, which is the default mode of operation for segment routing.

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

X is the node for which the label has local significance
D is the destination node

Figure 1: Formula for determining locally-significant shortest path labels

As a simple example, when the computing node (Y) needs to forward a packet ultimately destined for node D, Y first determines the shortest path next-hop node to reach D, which we assume to be X in this example. Y then adds the Node_Index value advertised by D to the Label_Block value advertised by X to determine the label value to apply to the packet before sending it to X.

[2.](#) Destination-based forwarding in segment routing using other algorithms

Figure 2 shows two options for generalizing the above formula, to determine locally-significant labels corresponding to forwarding next-hops computed using 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

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

X is the node for which the label has local significance

D is the destination node

A is the algorithm for computing destination-based forwarding next-hops

Figure 2: Two options for determining locally-significant labels for other algorithms

The computing router (Y) needs to forward a packet destined for node D using next-hops computed by algorithm A. Using either option, Y determines the next-hop computed by algorithm A to reach D, which we assume to be X in this example. Y then needs to figure out the correct label to apply to the packet so that X will also understand that the packet is destined for node D using forwarding next-hops

computed by algorithm A. The two options shown in Figure 2 differ in how Y determines that label value.

In Option 1 each node advertises a unique node index for each algorithm and a single label block. Y determines the label value of local significance to X to reach D using algorithm A by adding the Node_Index advertised by node D for algorithm A to the Label_Block advertised by node X. We refer to this as the per-algorithm node index option.

In Option 2 each node advertises a single node index and a unique label block for each algorithm. Y determines the label value of

local significance to X to reach D using algorithm A by adding the Node_Index advertised by node D to the Label_Block for algorithm A advertised by node X. We refer to this as the per-algorithm label block option.

The extensions currently defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-segment-routing-extensions\]](#) specify encodings for Option 1, the per-algorithm node index option. However, as illustrated in the following section, Option 2 has significant advantages over Option 1.

[3.](#) Practical implications of using the per-algorithm node index option

The following example illustrates the practical difficulties associated with using the per-algorithm node index option to support other forwarding algorithms.

Suppose an operator has a network with 100 nodes, which we will refer to as R0-R99. The operator assigns the unique node index values 0-99 to those nodes for algorithm=0, in order to accomplish shortest path routing based on IGP metrics with SR labels. Each node will need to advertise a label block of size=100.

Assume that at some future point in time, the IETF defines algorithm=1 to mean shortest path routing based on latency, and vendors implement this. Suppose that the operator wants to use latency-based SPF routes for some traffic and metric-based SPF routes for other traffic. The operator will need to define a new set of

unique node index values for algorithm=1. A reasonable choice would be to assign Node-SID values of 100-199 to R0-R99 for algorithm=1. Each node will now need to advertise a label block of size=200. So far the need for per-algorithm node index values is an annoyance, but not too difficult to deal with.

Now assume that the operator needs to add 10 new nodes to the SR domain, specifically nodes R100-R109. Each node will now need to advertise a label block of size=220. The main issue is deciding how to assign per-algorithm node index for the 10 new nodes? One option is to redo the node index numbering scheme so that R0-R109 have node index values 0-109 for algorithm=0 and node index values 110-229 for algorithm=1. However, this requires renumbering existing nodes. The other option is to avoid renumbering of nodes by assigning nodes R100-R109 node index values 200-209 for algorithm=0 and node index values 210-219 for algorithm=1. Each of these approaches has drawbacks. The first requires renumbering existing nodes, while the second is difficult to maintain since there is no obvious relationship between the node index values for different algorithms.

There are other numbering schemes for per-algorithm node index value that one can consider. However, they run into similar problems as described in the example above when trying to add more nodes or more algorithms. One could also try to mitigate these problems by anticipating growth in each dimension and pre-assigning node index values based on anticipated maximum values in each dimension. However, this can result in needing to advertise label blocks that are potentially much larger than actually needed.

[4.](#) Advantages of using the per-algorithm label block option

The use of per-algorithm label blocks avoids the problems associated with assigning and maintaining unique node index values for each forwarding algorithm.

When the SR domain is initially deployed, R0-R99 can be assigned node index values 0-99, as one would expect. When support for algorithm=1 gets added, the operator does not need to assign and configure any new node index values. Instead, the routers automate the process by advertising different label blocks for each forwarding algorithm.

When another 10 nodes are added to the SR domain, R100-R109 get

assigned node index values 100-109 as one would expect. And the router advertises a label block of size=110 for each algorithm, as one would expect. Adding new nodes in the presence of multiple forwarding algorithms is simplified significantly with the use of per-algorithm label blocks.

[5.](#) IANA Considerations

This document introduces no new IANA Considerations.

[6.](#) Management Considerations

This document proposes the use of per-algorithm label blocks to support destination-based forwarding along next-hops computed using different algorithms. The automated advertisement of per-algorithm label blocks significantly simplifies network management compared to configuration and maintenance of unique per-algorithm node indexes.

[7.](#) Security Considerations

TBD

[8.](#) Acknowledgements

TBD

[9.](#) Informative References

[I-D.ietf-isis-segment-routing-extensions]

Previdi, S., Filshie, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-03](#) (work in progress), October 2014.

[I-D.ietf-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filshie, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF

Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-04](#) (work in progress), February 2015.

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Shakir, R., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-01](#) (work in progress), February 2015.

Authors' Addresses

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: cbowers@juniper.net

Pushpasis Sarkar
Juniper Networks
Embassy Business Park
Bangalore, KA 560093
India

Email: psarkar@juniper.net

Hannes Gredler
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

