

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 16, 2016

S. Previdi, Ed.
C. Filsfils, Ed.
Cisco Systems, Inc.
B. Decraene
S. Litkowski
Orange
M. Horneffer
Deutsche Telekom
R. Shakir
Jive Communications
December 14, 2015

**SPRING Problem Statement and Requirements
draft-ietf-spring-problem-statement-06**

Abstract

The ability for a node to specify a forwarding path, other than the normal shortest path, that a particular packet will traverse, benefits a number of network functions. Source-based routing mechanisms have previously been specified for network protocols, but have not seen widespread adoption. In this context, the term 'source' means 'the point at which the explicit route is imposed' and therefore it is not limited to the originator of the packet (i.e.: the node imposing the explicit route may be the ingress node of an operator's network).

This document outlines various use cases, with their requirements, that need to be taken into account by the Source Packet Routing in Networking (SPRING) architecture for unicast traffic. Multicast use-cases and requirements are out of scope of this document.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 June 16, 2016.

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.	Dataplanes	4
3.	SPRING Use Cases	4
3.1.	IGP-based MPLS Tunneling	4
3.1.1.	Example of IGP-based MPLS Tunnels	4
3.2.	Fast Reroute (FRR)	5
3.3.	Traffic Engineering	5
3.3.1.	Examples of Traffic Engineering Use Cases	6
3.4.	Interoperability with non-SPRING nodes	12
4.	Security Considerations	12
5.	IANA Considerations	13
6.	Manageability Considerations	13
7.	Contributors	13
8.	Acknowledgements	13
9.	References	13
9.1.	Normative References	13
9.2.	Informative References	14
	Authors' Addresses	15

[1.](#) Introduction

The ability for a node to specify a unicast forwarding path, other than the normal shortest path, that a particular packet will traverse, benefits a number of network functions, for example:

Some types of network virtualization, including multi-topology networks and the partitioning of network resources for VPNs

Network, link, path and node protection such as fast re-route

Network programmability

OAM techniques

Simplification and reduction of network signaling components

Load balancing and traffic engineering

Source-based routing mechanisms have previously been specified for network protocols, but have not seen widespread adoption other than in MPLS traffic engineering.

These network functions may require greater flexibility and per packet source imposed routing than can be achieved through the use of the previously defined methods. In the context of this document, the term 'source' means 'the point at which the explicit route is imposed' and therefore it is not limited to the originator of the packet (i.e.: the node imposing the explicit route may be the ingress node of an operator's network). Throughout this document we refer to this definition of 'source'.

In this context, Source Packet Routing in Networking (SPRING) architecture is being defined in order to address the use cases and requirements described in this document.

The SPRING architecture SHOULD allow incremental and selective deployment without any requirement of flag day or massive upgrade of all network elements.

The SPRING architecture SHOULD allow optimal virtualization: put policy state in the packet header and not in the intermediate nodes along the path. Hence, the policy is completely virtualized away from midpoints and tail-ends.

The SPRING architecture objective is not to replace existing source routing and traffic engineering mechanisms but rather complement them and address use cases where removal of signaling and path state in the core is a requirement.

2. Dataplanes

The SPRING architecture SHOULD be general in order to ease its applicability to different dataplanes.

The SPRING architecture SHOULD leverage the existing MPLS dataplane without any modification and leverage IPv6 dataplane with a new IPv6 Routing Header Type (IPv6 Routing Header is defined in [[RFC2460](#)]).

3. SPRING Use Cases

3.1. IGP-based MPLS Tunneling

The source-based routing model, applied to the MPLS dataplane, offers the ability to tunnel services like VPN ([RFC4364](#)), VPLS ([RFC4761](#), [RFC4762](#)) and VPWS ([RFC6624](#)), from an ingress PE to an egress PE, with or without the expression of an explicit path and without requiring forwarding plane or control plane state in intermediate nodes. Point-to-multipoint and multipoint-to-multipoint tunnels are outside of the scope of this document.

3.1.1. Example of IGP-based MPLS Tunnels

This section illustrates an example use-case.




Figure 1: IGP-based MPLS Tunneling

In Figure 1 above, the four nodes A, CE1, CE2 and Z are part of the same VPN. CE2 advertises to PE2 a route to Z. PE2 binds a local label LZ to that route and propagates the route and its label via MPBGP to PE1 with nhop 192.0.2.2 (i.e.: the local address of PE2). PE1 installs the VPN prefix Z in the appropriate VRF and resolves the next-hop onto the node segment associated with PE2.

In order to cope with the reality of current deployments, the SPRING architecture SHOULD allow PE to PE forwarding according to the IGP shortest path without the addition of any other signaling protocol. The packet each PE forwards across the network will contain (within their label stack) the necessary information derived from the topology database in order to deliver the packet to the remote PE.

3.2. Fast Reroute (FRR)

FRR technologies have been deployed by network operators in order to cope with link or node failures through pre-computation of backup paths.

The SPRING architecture SHOULD address the following requirements:

- o support of Fast Reroute (FRR) on any topology
- o pre-computation and setup of backup path without any additional signaling (other than the regular IGP/BGP protocols)
- o support of shared risk constraints
- o support of node and link protection
- o support of microloop avoidance

Further illustrations of the problem statement for FRR are to be found in [[I-D.ietf-spring-resiliency-use-cases](#)].

3.3. Traffic Engineering

Traffic Engineering has been addressed using IGP protocol extensions (for resources information propagation) and RSVP-TE for signaling explicit paths ([[RFC5305](#)], [[RFC3630](#)], [[RFC3209](#)]). Different contexts and modes have been defined (single vs. multiple domains, with or without bandwidth admission control, centralized vs. distributed path computation, etc).

In all cases, one of the major components of the TE architecture is the soft state based signaling protocol (RSVP-TE) which is used in order to signal and establish the explicit path. Each path, once computed, need to be signaled and state for each path must be present in each node traversed by the path. This incurs a scalability problem especially in the context of SDN where traffic differentiation may be done at a finer granularity (e.g.: application specific). Also the amount of state needed to be maintained and periodically refreshed in all involved nodes contributes significantly to complexity and the number of failures cases, and thus increases operational effort while decreasing overall network reliability.

The source-based routing model allows traffic engineering to be implemented without the need of a signaling component.

The SPRING architecture SHOULD support traffic engineering, including:

- o loose or strict options
- o bandwidth admission control
- o distributed vs. centralized model (PCE [[I-D.ietf-pce-stateful-pce](#)], SDN Controller)
- o disjointness in dual-plane networks
- o egress peering traffic engineering
- o load-balancing among non-parallel links (i.e.: links connected to different adjacent neighbors).
- o Limiting (scalable, preferably zero) per-service state and signaling on midpoint and tail-end routers.
- o ECMP-awareness
- o node resiliency property (i.e.: the traffic-engineering policy is not anchored to a specific core node whose failure could impact the service).

3.3.1. Examples of Traffic Engineering Use Cases

Here follows the description of two sets of use cases:

- o Traffic Engineering without Admission Control
- o Traffic Engineering with Admission Control

3.3.1.1. Traffic Engineering without Bandwidth Admission Control

In this section, we describe Traffic Engineering use-cases without bandwidth admission control.

3.3.1.1.1. Disjointness in dual-plane networks

Many networks are built according to the dual-plane design, as illustrated in Figure 2:

Each aggregation region k is connected to the core by two C routers $C1k$ and $C2k$ where k refers to the region.

C1k is part of plane 1 and aggregation region k

C2k is part of plane 2 and aggregation region k

$C1_k$ has a link to $C2_j$ iff $k = j$.

The core nodes of a given region are directly connected.
Inter-region links only connect core nodes of the same plane.

$$\{C1k \text{ has a link to } C1j\} \text{ iff } \{C2k \text{ has a link to } C2j\}.$$

The distribution of these links depends on the topological properties of the core of the AS. The design rule presented above specifies that these links appear in both core planes.

We assume a common design rule found in such deployments: the inter-plane link costs ($C_{ik-C_{jk}}$ where $i < j$) are set such that the route to an edge destination from a given plane stays within the plane unless the plane is partitioned.




Figure 2: Dual-Plane Network and Disjointness

In this scenario, the operator requires the ability to deploy different strategies. For example, Edge Router A should be able to use the three following options:

- o the traffic is load-balanced across any ECMP path through the network
- o the traffic is load-balanced across any ECMP path within the Plane1 of the network
- o the traffic is load-balanced across any ECMP path within the Plane2 of the network

Most of the data traffic from A to Z would use the first option, such as to exploit the capacity efficiently. The operator would use the two other choices for specific premium traffic that has requested disjoint transport.

The SPRING architecture SHOULD support this use case with the following requirements:

- o Zero per-service state and signaling on midpoint and tail-end routers.
- o ECMP-awareness.
- o Node resiliency property: the traffic-engineering policy is not anchored to a specific core node whose failure could impact the service.

3.3.1.1.2. Egress Peering Traffic Engineering




Figure 3: Egress peering traffic engineering

Let us assume, in the network depicted in Figure 3, that:

C in AS1 learns about destination Z of AS 4 via two BGP paths (AS2, AS4) and (AS3, AS4).

C may or may not be configured so to enforce next-hop-self behavior before propagating the paths within AS1.

C may propagate all the paths to Z within AS1 (BGP add-paths, [[I-D.ietf-idr-add-paths](#)]).

C may install in its FIB only the route via AS2, or only the route via AS3, or both.

In that context, the SPRING architecture SHOULD allow the operator of AS1 to apply a traffic-engineering policy such as the following one, regardless the configured behavior of next-hop-self:

Steer 60% of the Z-destined traffic received at A via AS2 and 40% via AS3.

Steer 80% of the Z-destined traffic received at B via AS2 and 20% via AS3.

The SPRING architecture SHOULD allow an ingress node (i.e., an explicit route source node) to select the exit point of a packet as any combination of an egress node, an egress interface, a peering neighbor, and a peering AS.

The use cases and requirements for Egress Peer Engineering are described in [[I-D.ietf-spring-segment-routing-central-epe](#)].

3.3.1.1.3. Load-balancing among non-parallel links

The SPRING architecture SHOULD allow a given node to load share traffic across multiple non parallel links (i.e.: links connected to different adjacent routers) even if these lead to different neighbors. This may be useful to support traffic engineering policies.




Figure 4: Multiple (non-parallel) Adjacencies

In the above example, the operator requires PE1 to load-balance its PE2-destined traffic between the ABCDE and ABFE equal-cost paths in a controlled way where the operator SHOULD be allowed to distribute

traffic unevenly between paths (Weighted Equal Cost Multiplath, WECP).

3.3.1.2. Traffic Engineering with Bandwidth Admission Control

The implementation of bandwidth admission control within a network (and its possible routing consequence which consists in routing along explicit paths where the bandwidth is available) requires a capacity planning process.

The spreading of load among ECMP paths is a key attribute of the capacity planning processes applied to packet-based networks.

3.3.1.2.1. Capacity Planning Process

Capacity Planning anticipates the routing of the traffic matrix onto the network topology, for a set of expected traffic and topology variations. The heart of the process consists in simulating the placement of the traffic along ECMP-aware shortest-paths and accounting for the resulting bandwidth usage.

The bandwidth accounting of a demand along its shortest-path is a basic capability of any planning tool or PCE server.

For example, in the network topology described below, and assuming a default IGP metric of 1 and IGP metric of 2 for link GF, a 1600Mbps A-to-Z flow is accounted as consuming 1600Mbps on links AB and FZ, 800Mbps on links BC, BG and GF, and 400Mbps on links CD, DF, CE and EF.




Figure 5: Capacity Planning an ECMP-based demand

ECMP is extremely frequent in SP, Enterprise and DC architectures and it is not rare to see as much as 128 different ECMP paths between a source and a destination within a single network domain. It is a key efficiency objective to spread the traffic among as many ECMP paths as possible.

This is illustrated in the below network diagram which consists of a subset of a network where already 5 ECMP paths are observed from A to M.




Figure 6: ECMP Topology Example

When the capacity planning process detects that a traffic growth scenario and topology variation would lead to congestion, a capacity increase is triggered and if it cannot be deployed in due time, a traffic engineering solution is activated within the network.

A basic traffic engineering objective consists of finding the smallest set of demands that need to be routed off their shortest path to eliminate the congestion, then to compute an explicit path for each of them and instantiating these traffic-engineered policies in the network.

The SPRING architecture SHOULD offer a simple support for ECMP-based shortest path placement as well as for explicit path policy without incurring additional signaling in the domain. This includes:

- o the ability to steer a packet across a set of ECMP paths
- o the ability to diverge from a set of ECMP shortest paths to one or more paths not in the set of shortest paths

3.3.1.2.2. SDN/SR use-case

The SDN use-case lies in the SDN controller, (e.g.: Stateful PCE as described in [[I-D.ietf-pce-stateful-pce](#)]).

The SDN controller is responsible to control the evolution of the traffic matrix and topology. It accepts or denies the addition of new traffic into the network. It decides how to route the accepted traffic. It monitors the topology and upon topological change, determines the minimum traffic that should be rerouted on an alternate path to alleviate a bandwidth congestion issue.

The algorithms supporting this behavior are a local matter of the SDN controller and are outside the scope of this document.

The means of collecting traffic and topology information are the same as what would be used with other SDN-based traffic-engineering solutions.

The means of instantiating policy information at a traffic-engineering head-end are the same as what would be used with other SDN-based traffic-engineering solutions.

In the context of Centralized-Based Optimization and the SDN use-case, here are the functionalities that the SPRING architecture SHOULD deliver:

Explicit routing capability with or without ECMP-awareness.

No signaling hop-by-hop through the network.

State is only maintained at the policy head-end. No state is maintained at mid-points and tail-ends.

Automated guaranteed FRR for any topology.

Optimum virtualization: the policy state is in the packet header and not in the intermediate nodes along the path. The policy is completely virtualized away from midpoints and tail-ends.

Highly responsive to change: the SDN Controller only needs to apply a policy change at the head-end. No delay is introduced due to programming the midpoints and tail-end along the path.

3.4. Interoperability with non-SPRING nodes

SPRING nodes MUST inter-operate with non-SPRING nodes and in both MPLS and IPv6 dataplanes in order to allow a gradual deployment of SPRING on existing MPLS and IPv6 networks.

4. Security Considerations

There is an assumed trust model such that the source imposing an explicit route on a packet is assumed to be allowed to do so. It is assumed that the default behavior is to strip any internal routing information from the packet before the packet is forwarded outside the domain. In such context trust boundaries SHOULD strip explicit routes from a packet.

For each data plane technology that SPRING specifies, a security analysis MUST be provided showing how protection is provided against an attacker disrupting the network by for example, maliciously injecting SPRING packets.

5. IANA Considerations

This document does not request any IANA allocations.

6. Manageability Considerations

The SPRING WG SHOULD provide OAM and the management needed to manage SPRING enabled networks. The SPRING procedures MAY also be used as a tool for OAM in SPRING enabled networks.

OAM use cases and requirements are described in [\[I-D.geib-spring-oam-usecase\]](#) and [\[I-D.kumar-spring-sr-oam-requirement\]](#).

7. Contributors

The following individuals substantially contributed to the content of this documents:

Ruediger Geib
Deutsche Telekom
Heinrich Hertz Str. 3-7
Darmstadt 64295
DE
Email: Ruediger.Geib@telekom.de

Robert Raszuk
Mirantis Inc.
615 National Ave.
94043 Mt View, CA
US
Email: robert@raszuk.net

8. Acknowledgements

The authors would like to thank Yakov Rekhter for his contribution to this document.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), DOI 10.17487/RFC4364, February 2006, <<http://www.rfc-editor.org/info/rfc4364>>.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<http://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<http://www.rfc-editor.org/info/rfc4762>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", [RFC 6624](#), DOI 10.17487/RFC6624, May 2012, <<http://www.rfc-editor.org/info/rfc6624>>.

9.2. Informative References

- [I-D.geib-spring-oam-usecase]
Geib, R., Filsfils, C., Pignataro, C., and N. Kumar, "Use case for a scalable and topology aware MPLS data plane monitoring system", [draft-geib-spring-oam-usecase-06](#) (work in progress), July 2015.

[I-D.ietf-idr-add-paths]

Walton, D., Retana, A., Chen, E., and J. Scudder,
"Advertisement of Multiple Paths in BGP", [draft-ietf-idr-add-paths-13](#) (work in progress), December 2015.

[I-D.ietf-pce-stateful-pce]

Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP
Extensions for Stateful PCE", [draft-ietf-pce-stateful-pce-13](#) (work in progress), December 2015.

[I-D.ietf-spring-resiliency-use-cases]

Francois, P., Filsfils, C., Decraene, B., and r.
rjs@rob.sh, "Use-cases for Resiliency in SPRING", [draft-ietf-spring-resiliency-use-cases-02](#) (work in progress),
December 2015.

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

Filsfils, C., Previdi, S., Ginsburg, D., and D. Afanasiev,
"Segment Routing Centralized Egress Peer Engineering",
[draft-ietf-spring-segment-routing-central-epe-00](#) (work in
progress), October 2015.

[I-D.kumar-spring-sr-oam-requirement]

Kumar, N., Pignataro, C., Akiya, N., Geib, R., Mirsky, G.,
and S. Litkowski, "OAM Requirements for Segment Routing
Network", [draft-kumar-spring-sr-oam-requirement-03](#) (work
in progress), March 2015.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico, 200
Rome 00142
Italy

Email: sprevidi@cisco.com

Clarence Filsfils (editor)
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com

Bruno Decraene
Orange
FR

Email: bruno.decraene@orange.com

Stephane Litkowski
Orange
FR

Email: stephane.litkowski@orange.com

Martin Horneffer
Deutsche Telekom
Hammer Str. 216-226
Muenster 48153
DE

Email: Martin.Horneffer@telekom.de

Rob Shakir
Jive Communications, Inc.
1275 West 1600 North, Suite 100
Orem, UT 84057

Email: rjs@rob.sh

