

spring
Internet-Draft
Intended status: Informational
Expires: August 11, 2017

R. Geib, Ed.
Deutsche Telekom
C. Filsfils
C. Pignataro, Ed.
N. Kumar
Cisco Systems, Inc.
February 7, 2017

A Scalable and Topology-Aware MPLS Dataplane Monitoring System draft-ietf-spring-oam-usecase-05

Abstract

This document describes features of a path monitoring system and related use cases. Segment based routing enables a scalable and simple method to monitor data plane liveliness of the complete set of paths belonging to a single domain. The MPLS monitoring system adds features to the traditional MPLS ping and LSP path trace, in a very complementary way. MPLS topology awareness reduces management and control plane involvement of OAM measurements while enabling new OAM features.

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 August 7, 2017.

Copyright Notice

Copyright (c) 2017 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.	Acronyms	2
2.	Introduction	3
3.	Terminology	5
4.	An MPLS Topology-Aware Path Monitoring System	6
5.	SR-based Path Monitoring Use Case Illustration	7
5.1.	Use Case 1 - LSP Dataplane Monitoring	7
5.2.	Use Case 2 - Monitoring a Remote Bundle	10
5.3.	Use Case 3 - Fault Localization	11
6.	Failure Notification from PMS to LERi	11
7.	Applying SR to Monitoring non-SR based LSPs (LDP and possibly RSVP-TE)	12
8.	PMS Monitoring of Different Segment ID Types	13
9.	Connectivity Verification Using PMS	13
10.	Extensions of Specifications Relevant to this Use Case	13
11.	IANA Considerations	13
12.	Security Considerations	14
13.	Acknowledgements	14
14.	References	14
14.1.	Normative References	14
14.2.	Informative References	14
	Authors' Addresses	15

1. Acronyms

ECMP	Equal-Cost Multi-Path
IGP	Interior Gateway Protocol
LER	Label Edge Router
LSP	Label Switched Path
LSR	Label Switching Router
OAM	Operations, Administration, and Maintenance
PMS	Path Monitoring System
RSVP-TE	Resource ReserVation Protocol-Traffic Engineering

SID Segment Identifier

SR Segment Routing

SRGB Segment Routing Global Block

2. Introduction

It is essential for a network operator to monitor all the forwarding paths observed by the transported user packets. Monitoring packets are expected to be forwarded in dataplane in a similar way as user packets. Segment Routing enables forwarding of packets along pre-defined paths and segments and thus a Segment Routed monitoring packet can stay in dataplane while passing along one or more segments to be monitored.

This document describes illustrates a system using MPLS data plane path monitoring capabilities. The use cases introduced here are limited to a single Interior Gateway Protocol (IGP) MPLS domain.

The system applies to monitoring of pre Segment Routing LSP's (like LDP) as well as to monitoring of Segment Routed LSP's ([section 7](#) offers some more information). As compared to pre Segment Routing approaches, Segment Routing is expected to simplify such a monitoring system by enabling MPLS topology detection based on IGP signaled segments as specified by [\[I-D.ietf-isis-segment-routing-extensions\]](#), [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-idr-bgp-ls-segment-routing-ext\]](#). Thus a centralised and MPLS topology aware monitoring unit can be realized in a Segment Routed domain. This topology awareness can be used for OAM purposes as described by this document.

The system offers several benefits for network monitoring:

- o A single centralized MPLS monitoring system which is able to perform a continuity check (ping) along all Label Switched Paths of the SR domain. Monitoring packets never leave data plane.
- o The MPLS ping (or continuity check) packets never leave the MPLS data plane.
- o SR allows to transport MPLS path trace or connectivity validation packets for any existing Label Switched Path to all nodes of an SR domain. This use case doesn't describe any new path trace features, but the system described here allows to set up an SR domain wide centralised connectivity validation.

- o An MPLS monitoring system, maybe several ones if redundancy is desired, which apply SR for OAM purposes as described, offer the possibility to scale and design a flexible MPLS OAM platform as suitable for a provider.

In addition to monitoring paths, problem localization is required. Faults can be localized:

- o by capturing the Interior Gateway Protocol (IGP) topology and analysing IGP messages indicating changes of it.
- o by correlation between different SR based monitoring probes.
- o by setting up an MPLS traceroute packet for a path (or Segment) to be tested and transporting it to a node to validate path connectivity from that node on.

Topology awareness is an essential part of link state IGP. Adding MPLS topology awareness to an IGP speaking device hence enables a simple and scalable data plane based monitoring mechanism.

MPLS OAM offers flexible traceroute (connectivity verification) features to recognise and execute data paths of an MPLS domain. By utilising the ECMP related tool set offered, e.g., by [RFC 4379](#) [RFC4379], a SR based MPLS monitoring system can be enabled to:

- o detect how to route packets along different ECMP routed paths.
- o construct ping packets respectively, which can be precisely steered to paths whose connectivity is to be checked, also if ECMP is present.
- o limit the MPLS label stack of such a ping packet checking continuity of every single IGP-Segment to the maximum number of 3 labels. A smaller label stack may also be helpful, if any router interprets a limited number of packet header bytes to determine an ECMP path along which to route a packet.

Alternatively, any path may be executed by building suitable label stacks. This allows path execution without ECMP awareness.

The MPLS Path Monitoring System may be any server residing at a single interface of the domain to be monitored. The PMS doesn't need to support the complete MPLS routing or control plane. It needs to be capable to learn and maintain an accurate MPLS and IGP topology. MPLS ping and traceroute packets need to be set up and sent with the correct segment stack. The PMS further must be able to receive and decode returning ping or traceroute packets. Packets used to check

continuity could have BFD or LSP Ping format, or have any other OAM format supported by the PMS. As long as the packet used to check continuity returns back to the server while no IGP change is detected, the monitored path can be considered as validated. If the depth of label stacks to be pushed for the purpose of path monitoring is of concern for a domain, a dedicated PMS server allows to push monitoring related label stacks of arbitrary depth on this server. Hence router label stack limitations don't limit MPLS OAM choices.

Documents discussing SR OAM requirements and MPLS traceroute enhancements adding functionality to the use cases described by this document are in work within IETF, see [\[I-D.ietf-spring-sr-oam-requirement\]](#) and [\[I-D.draft-ietf-mpls-spring-lsp-ping\]](#).

3. Terminology

Continuity Check

[RFC 7276](#) [[RFC7276](#)] defines Continuity Checks to be used to verify that a destination is reachable, and are typically sent proactively, though they can be invoked on-demand as well. Segment Routing allows to realise a continuity check along any given SR domain path within data plane.

Connectivity Verification

[RFC 7276](#) [[RFC7276](#)] defines Connectivity Verification as a mechanism to check connectivity between two nodes by checking whether a path between both can be used. [RFC 4379](#) [[RFC7276](#)] specifies a Connectivity Verification for MPLS domains. As [RFC 7276](#) states, Connectivity Verification and Continuity Checks are considered complementary mechanisms and are often used in conjunction with each other. The use cases following merely treat SR based network monitoring as adding a new method to realise a Continuity Check. In special cases, the SR based Continuity Check offers limited Connectivity Verification properties. This will be in the use case descriptions, if applicable.

[RFC 7276](#) [[RFC7276](#)] defines Connectivity Verification as a mechanism to check connectivity between two nodes by checking whether a path between both can be used. [RFC 4379](#) [[RFC7276](#)] specifies a Connectivity Verification for MPLS domains. As [RFC 7276](#) states, Connectivity Verification and Continuity Checks are considered complementary mechanisms and are often used in conjunction with each other. The use cases following merely treat SR based network monitoring as adding a new method to

realise a Continuity Check. In special cases, the SR based Continuity Check offers limited Connectivity Verification properties. This will be in the use case descriptions, if applicable.

MPLS topology

The MPLS topology of an MPLS domain is the complete set of MPLS- and IP-address information and all routing and data plane information required to address and utilise every MPLS path within this domain from an MPLS Path Monitoring System attached to this MPLS domain at an arbitrary access. This document assumes availability of the MPLS topology (which can be detected with available protocols and interfaces). None of the use cases will describe how to set it up.

This document further adopts the terminology and framework described in [[I-D.ietf-spring-segment-routing](#)].

4. An MPLS Topology-Aware Path Monitoring System

Any node at least listening to the IGP of an SR domain is MPLS topology aware (the node knows all related IP addresses, SR SIDs and MPLS labels). An MPLS PMS which is able to learn the IGP LSDB (including the SID's) is able to execute arbitrary chains of label switched paths. To monitor an MPLS SR domain, a PMS needs to set up a topology data base of MPLS SR domain to be monitored. It may be used to send ping type packets to only check continuity along such a path chain based on the topology information only. In addition, the PMS can be used to trace MPLS Label Switched Path and thus verify their connectivity and correspondance between control and data plane, respectively. The PMS can direct suitable MPLS traceroute packets to any node along a path segment.

Let us describe how the PMS constructs a labels stack to transport a packet to LER i, monitor its path to LER j and then receive the packet back.

The PMS may do so by sending packets carrying the following MPLS label stack information:

- o Top Label: a path from PMS to LER i, which is expressed as Node SID of LER i.
- o Next Label: the path that needs to be monitored from LER i to LER j. If this path is a single physical interface (or a bundle of connected interfaces), it can be expressed by the related Adjacency-SID. If the shortest path from LER i to LER j is


supposed to be monitored, the Node-SID (LER j) can be used. Another option is to insert a list of segments expressing the desired path (hop by hop as an extreme case). If LER i pushes a stack of Labels based on a SR policy decision and this stack of LSPs is to be monitored, the PMS needs an interface to collect the information enabling it to address this SR created path.

- o Next Label or address: the path back to the PMS. Likely, no further segment/label is required here. Indeed, once the packet reaches LER j, the 'steering' part of the solution is done and the probe just needs to return to the PMS. This is best achieved by popping the MPLS stack and revealing a probe packet with PMS as destination address (note that in this case, the source and destination addresses could be the same). If an IP address is applied, no SID/label has to be assigned to the PMS (if it is a host/server residing in an IP subnet outside the MPLS domain).

The PMS should be physically connected to a router which is part of the SR domain. It must be able to send and receive MPLS packets via this interface. As mentioned above, routing protocol support isn't required and the PMS itself doesn't have to be involved in IGP or MPLS routing. A static route will do. Further options, like deployment of a PMS connecting to the MPLS domain by a tunnel only require more thought, as this implies security aspects. MPLS so far separates networks securely by avoiding tunnel access to MPLS domains.

5. SR-based Path Monitoring Use Case Illustration

5.1. Use Case 1 - LSP Dataplane Monitoring



Example of a PMS based LSP dataplane monitoring

Figure 1

For the sake of simplicity, let's assume that all the nodes are configured with the same SRGB [[I-D.ietf-spring-segment-routing](#)].

Let's assign the following Node SIDs to the nodes of the figure: PMS = 10, LER i = 20, LER j = 30.

The aim is to set up a continuity check of the path between LER i and LER j. As has been said, the monitoring packets are to be sent and received by the PMS. Let's assume the design aim is to be able to work with the smallest possible SR label stack. In the given topology, a fairly simple option is to perform an MPLS path trace, as specified by [RFC4379](#). The starting point for the path trace is LER i and the PMS sends the MPLS path trace packet to LER i. The MPLS echo reply of LER i should be sent to the PMS. As a result, IP destination address choices are detected, which are then used to target any one of the ECMP routed paths between LER i and LER j by the MPLS ping packets to later check path continuity. The Label stack of these ping packets doesn't need to consist of more than 3 labels. Finally, the PMS sets up and sends packets to monitor connectivity of the ECMP routed paths. The PMS does this by creating a measurement packet with the following label stack (top to bottom): 20 - 30 - 10. The ping packets reliably use the monitored path, if the IP-address information which has been detected by the MPLS trace route is used as the IP destination address (note that this IP address isn't used or required for any IP routing).

LER m forwards the packet received from the PMS to LSR1. Assuming Pen-ultimate Hop Popping to be deployed, LSR1 pops the top label and forwards the packet to LER i. There the top label has a value 30 and LER i forwards it to LER j. This will be done transmitting the packet via LSR1 or LSR2. The LSR will again pop the top label. LER j will forward the packet now carrying the top label 10 to the PMS (and it will pass a LSR and LER m).

A few observations on the example given in figure 1:

- o The path PMS to LER i must be available (i.e., a continuity check only along the path to LER i must succeed). If desired, an MPLS trace route may be used to exactly detect the data plane path taken for this MPLS Segment. It is usually sufficient to just apply any of the existing Shortest Path routed paths.
- o If ECMP is deployed, separate continuity checks monitoring all possible paths which a packet may use between LER i and LER j may be desired. This can be done by applying an MPLS trace route between LER i and LER j. Another option is to use SR routing, but this will likely require additional label information within the label stack of the ping packet. Further, if multiple links are

deployed between two nodes, SR methods to address each individual path require an Adj-SID to be assigned to each single interface. This method is based on control plane information - a connectivity verification based on MPLS traceroute seems to be a fairly good option to deal with ECMP and validation of control and data plane correlation.

- o The path LER j to PMS must be available (i.e., a continuity check only along the path from LER j to PMS must succeed). If desired, an MPLS trace route may be used to exactly detect the data plane path taken for this MPLS Segment. It is usually sufficient to just apply any of the existing Shortest Path routed paths.

Once the MPLS paths (Node-SIDs) and the required information to deal with ECMP have been detected, the path continuity between LER i and LER j can be monitored by the PMS. Path continuity monitoring by ping packets does not require [RFC4379](#) MPLS OAM functionality. All monitoring packets stay on dataplane, hence path continuity monitoring does not require control plane interaction in any LER or LSR of the domain. To ensure consistent interpretation of the results, the PMS should be aware of any changes in IGP or MPLS topology or ECMP routing. While the description given here pronouncing path connectivity checking as a simple basic application, others like checking continuity of underlying physical infrastructure or delay measurements may be desired. In both cases, a change in ECMP routing which is not caused by an IGP or MPLS topology change may not be desirable. A PMS therefore should also periodically verify connectivity of the SR paths which are monitored for continuity.


Determining a path to be executed prior to a measurement may also be done by setting up a label stack including all Node-SIDs along that path (if LSR1 has Node SID 40 in the example and it should be passed between LER i and LER j, the label stack is 20 - 40 - 30 - 10). The advantage of this method is, that it does not involve [RFC 4379](#) connectivity verification and, if there's only one physical connection between all nodes, the approach is independent of ECMP functionalities. The method still is able to monitor all link combinations of all paths of an MPLS domain. If correct forwarding along the desired paths has to be checked, or multiple physical connections exist between any two nodes, either additional information based on an MPLS trace route or additional Adj-SIDs are required to deal with ECMP.

In theory at least, a single PMS is able to monitor data plane availability of all LSPs in the domain. The PMS may be a router, but could also be dedicated monitoring system. If measurement system

reliability is an issue, more than a single PMS may be connected to the MPLS domain.

Monitoring an MPLS domain by a PMS based on SR offers the option of monitoring complete MPLS domains with limited effort and a unique possibility to scale a flexible monitoring solution as required by the operator (the number of PMS deployed is independent of the locations of the origin and destination of the monitored paths). The PMS can be enabled to send MPLS OAM packets with the label stacks and address information identical to those of the monitoring packets to any node of the MPLS domain. The routers of the monitored domain should support [RFC 4379](#) and its standardised extensions to allow for MPLS trace route. Ping based continuity checks don't require router control plane activity. Prior to monitoring a path, MPLS OAM may be used to detect ECMP dependant forwarding of a packet. A PMS may be designed to learn the IP address information required to execute a particular ECMP routed path and interfaces along that path. This allows to monitor these paths with label stacks reduced to a limited number of Node-SIDs resulting from SPF routing. The PMS does not require access to LSR / LER management- or data-plane information to do so.

5.2. Use Case 2 - Monitoring a Remote Bundle



SR based probing of all the links of a remote bundle

Figure 2

R1 addresses Lx by the Adjacency SID 99x, while R2 addresses Lx by the Adjacency SID 66(x+1).

In the above figure, the PMS needs to assess the dataplane availability of all the links within a remote bundle connected to routers R1 and R2.

The monitoring system retrieves the SID/Label information from the IGP LSDB and appends the following segment list/label stack: {72, 662, 992, 664} on its IP probe (whose source and destination addresses are the address of the PMS).

PMS sends the probe to its connected router. If the connected router is not SR compliant, a tunneling technique can be used to tunnel the probe and its MPLS stack to the first SR router. The MPLS/SR domain then forwards the probe to R2 (72 is the Node SID of R2). R2 forwards the probe to R1 over link L1 (Adjacency SID 662). R1 forwards the probe to R2 over link L2 (Adjacency SID 992). R2 forwards the probe to R1 over link L3 (Adjacency SID 664). R1 then forwards the IP probe to PMS as per classic IP forwarding.

As has been mentioned in [section 5.1](#), the PMS must be able monitor continuity of the path PMS to R2 (Node-SID 72) as well as continuity from R1 to the PMS. If both are given and packets are lost, forwarding on one of the three interfaces connecting R1 to R2 must be disturbed.

5.3. Use Case 3 - Fault Localization

In the previous example, a uni-directional fault on the middle link in direction of R2 to R1 would be localized by sending the following two probes with respective segment lists:

- o 72, 662, 992, 664
- o 72, 663, 992, 664

The first probe would succeed while the second would fail. Correlation of the measurements reveals that the only difference is using the Adjacency SID 663 of the middle link from R2 to R1 in the non successful measurement. Assuming the second probe has been routed correctly, the fault must have been occurring in R2 which didn't forward the packet to the interface identified by its Adjacency SID 663.

The example above only illustrates a method to localise a fault by correlated continuity checks. Any operational deployment requires a well designed engineering to allow for the desired non ambiguous diagnosis on the monitored section of the SR network. 'Section' here could be a path, a single physical interface, the set of all links of a bundle or an adjacency of two nodes, just to name a few. Such a design is not within scope of this document.

6. Failure Notification from PMS to LERi

PMS on detecting any failure in the path liveliness may use any out-of-band mechanism to signal the failure to LER i. This document does not propose any specific mechanism and operators can choose any existing or new approach.

Alternately, the Operator may log the failure in local monitoring system and take necessary action by manual intervention.

7. Applying SR to Monitoring non-SR based LSPs (LDP and possibly RSVP-TE)

The MPLS path monitoring system described by this document can be realised with pre-Segment Routing (SR) based technology. Making such a pre-SR MPLS monitoring system aware of a domain's complete MPLS topology requires, e.g., management plane access to the routers of the domain to be monitored or set up of a dedicated T-LDP tunnel per router to set up an LDP adjacency. To avoid the use of stale MPLS label information, the IGP must be monitored and MPLS topology must be timely aligned with IGP topology. Obviously, enhancing IGPs to exchange of MPLS topology information as done by SR significantly simplifies and stabilises such an MPLS path monitoring system.

A SR based PMS connected to a MPLS domain consisting of LER and LSR supporting SR and LDP or RSVP-TE in parallel in all nodes may use SR paths to transmit packets to and from start and end points of non-SR based LSP paths to be monitored. In the above example, the label stack top to bottom may be as follows, when sent by the PMS:

- o Top: SR based Node-SID of LER i at LER m.
- o Next: LDP or RSVP-TE label identifying the path or tunnel, respectively from LER i to LER j (at LER i).
- o Bottom: SR based Node-SID identifying the path to the PMS at LER j

While the mixed operation shown here still requires the PMS to be aware of the LER LDP-MPLS topology, the PMS may learn the SR MPLS topology by IGP and use this information.

An implementation report on a PMS operating in an LDP domain is given in [[I-D.leipnitz-spring-pms-implementation-report](#)]. In addition, this report compares delays measured with a single PMS to the results measured by three IP Performance Measurement Work Group (IPPM WG) standard conformant Measurement Agents (connected to an MPLS domain at three different sites). The delay measurements of PMS and where compared based on a statistical test published by the IPPM WG [[RFC6576](#)]. The Anderson Darling k-sample test showed that the PMS round-trip delay measurements are equal to those captured by an IPPM conformant IP measurement system for 64 Byte measurement packets with 95% confidence.

The authors are not aware of similar deployment for RSVP-TE. Identification of tunnel entry- and transit-nodes may add complexity. They are not within scope of this document.

8. PMS Monitoring of Different Segment ID Types

MPLS SR topology awareness should allow the SID to monitor liveness of SIDs related to interfaces within the SR and IGP domain, respectively. Tracing a path where an SR capable node assigns an Adj-SID for a non-SR-capable node may fail. This and other backward compatibility with non Segment Routing devices are discussed by [I-D.[draft-ietf-mpls-spring-lsp-ping](#)].

To match control plane information with data plane information, MPLS OAM functions as defined for example by [RFC4379](#) [[RFC4379](#)] are enhanced to allow collection of data relevant to check all relevant types of Segment IDs by [I-D.[draft-ietf-mpls-spring-lsp-ping](#)].

9. Connectivity Verification Using PMS

While the PMS based use cases explained in [Section 5](#) are sufficient to provide continuity check between LER i and LER j, it may not help perform connectivity verification. So in some cases like data plane programming corruption, it is possible that a transit node between LER i and LER j erroneously removes the top segment ID and forwards a monitoring packet to the PMS based on the bottom segment ID leading to a falsified path liveness indication by the PMS.

There are various method to perform basic connectivity verification like intermittently setting the TTL to 1 in bottom label so LER j selectively perform connectivity verification. Other methods are possible and may be added when requirements and solutions are specified.

10. Extensions of Specifications Relevant to this Use Case

The following activities are welcome enhancements supporting this use case, but they are not part of it:

[RFC4379](#) [[RFC4379](#)] functions should be extended to support Flow- and Entropy Label based ECMP.

11. IANA Considerations

This memo includes no request to IANA.

12. Security Considerations

As mentioned in the introduction, a PMS monitoring packet should never leave the domain where it originated. It therefore should never use stale MPLS or IGP routing information. Further, assigning different label ranges for different purposes may be useful. A well known global service level range may be excluded for utilisation within PMS measurement packets. These ideas shouldn't start a discussion. They rather should point out, that such a discussion is required when SR based OAM mechanisms like a SR are standardised.

Should the approach of a PMS connected to an SR domain by a tunnel be picked up, some fundamental MPLS security properties need to be discussed. MPLS domains so far allow to separate the MPLS network from an IP network by allowing no tunneled MPLS access to an MPLS domain.

13. Acknowledgements

The authors would like to thank Nobo Akiya for his contribution. Raik Leinertz kindly provided an editorial review. The authors would also like to thank Faisal Iqbal for an insightful review and a useful set of comments and suggestions.

14. References

14.1. Normative References

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), DOI 10.17487/RFC4379, February 2006, <<http://www.rfc-editor.org/info/rfc4379>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", [RFC 7276](#), DOI 10.17487/RFC7276, June 2014, <<http://www.rfc-editor.org/info/rfc7276>>.

14.2. Informative References

- [I-D.[draft-ietf-mpls-spring-lsp-ping](#)] IETF, "Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane", IETF, <https://datatracker.ietf.org/doc/draft-ietf-mpls-spring-lsp-ping/>, 2016.

- [I-D.ietf-idr-bgp-ls-segment-routing-ext]
IETF, "BGP Link-State extensions for Segment Routing",
IETF, <https://datatracker.ietf.org/doc/draft-ietf-idr-bgp-ls-segment-routing-ext/>, 2016.
- [I-D.ietf-isis-segment-routing-extensions]
IETF, "IS-IS Extensions for Segment Routing", IETF,
<https://datatracker.ietf.org/doc/draft-ietf-isis-segment-routing-extensions/>, 2016.
- [I-D.ietf-ospf-segment-routing-extensions]
IETF, "OSPF Extensions for Segment Routing", IETF,
<https://datatracker.ietf.org/doc/draft-ietf-ospf-segment-routing-extensions/>, 2016.
- [I-D.ietf-spring-segment-routing]
IETF, "Segment Routing Architecture", IETF,
<https://datatracker.ietf.org/doc/draft-ietf-spring-segment-routing/>, 2016.
- [I-D.ietf-spring-sr-oam-requirement]
IETF, "OAM Requirements for Segment Routing Network",
IETF, <https://datatracker.ietf.org/doc/draft-ietf-spring-sr-oam-requirement/>, 2016.
- [I-D.leipzig-spring-pms-implementation-report]
Leipzig, R. and R. Geib, "A scalable and topology aware
MPLS data plane monitoring system", IETF, [draft-leipzig-spring-pms-implementation-report-00](https://datatracker.ietf.org/doc/draft-leipzig-spring-pms-implementation-report-00/), 2016.
- [RFC6576] Geib, R., Ed., Morton, A., Fardid, R., and A. Steinmitz,
"IP Performance Metrics (IPPM) Standard Advancement
Testing", [BCP 176](https://datatracker.ietf.org/doc/BCP-176/), [RFC 6576](https://datatracker.ietf.org/doc/rfc6576/), DOI 10.17487/RFC6576, March
2012, <<http://www.rfc-editor.org/info/rfc6576>>.

Authors' Addresses

Ruediger Geib (editor)
Deutsche Telekom
Heinrich Hertz Str. 3-7
Darmstadt 64295
Germany

Phone: +49 6151 5812747
Email: Ruediger.Geib@telekom.de

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

Carlos Pignataro (editor)
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: cpignata@cisco.com

Nagendra Kumar
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: naikumar@cisco.com