

spring
Internet-Draft
Intended status: Informational
Expires: January 26, 2018

R. Geib, Ed.
Deutsche Telekom
C. Filsfils
C. Pignataro, Ed.
N. Kumar
Cisco Systems, Inc.
July 25, 2017

A Scalable and Topology-Aware MPLS Dataplane Monitoring System draft-ietf-spring-oam-usecase-09

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 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 January 26, 2018.

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.	Introduction	2
2.	Terminology and Acronyms	5
2.1.	Terminology	5
2.2.	Acronyms	5
3.	An MPLS Topology-Aware Path Monitoring System	6
4.	SR-based Path Monitoring Use Case Illustration	7
4.1.	Use Case 1 - LSP Dataplane Monitoring	7
4.2.	Use Case 2 - Monitoring a Remote Bundle	10
4.3.	Use Case 3 - Fault Localization	11
5.	Path Trace and Failure Notification	11
6.	Applying SR to Monitoring non-SR based LSPs (LDP and possibly RSVP-TE)	12
7.	PMS Monitoring of Different Segment ID Types	13
8.	Connectivity Verification Using PMS	13
9.	IANA Considerations	14
10.	Security Considerations	14
11.	Acknowledgements	15
12.	References	16
12.1.	Normative References	16
12.2.	Informative References	16
	Authors' Addresses	17

[1.](#) Introduction

Network operator need to be able to monitor the forwarding paths used to transport 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 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 non Segment Routing Label Switched Paths (LSP's) like LDP as well as to monitoring of Segment Routed LSP's ([section 7](#) offers some more information). As compared to non Segment Routing approaches, Segment Routing is expected to

simplify such a monitoring system by enabling MPLS topology detection based on IGP signaled segments. The MPLS topology should be detected and correlated with the IGP topology, which is too detected by IGP signaling. Thus a centralized and MPLS topology aware monitoring unit can be realized in a Segment Routed domain. This topology awareness can be used for Operation, Administration, and Maintenance (OAM) purposes as described by this document.

Benefits offered by the system:

- o The system described here allows to set up an SR domain wide centralized connectivity validation. Many operators operators of large networks regard centralized monitoring system as useful..
- o The MPLS Ping (or continuity check) packets never leave the MPLS user data plane.
- o SR allows to transport MPLS path trace or connectivity validation packets for every Label Switched Path to all nodes of an SR domain. This use case doesn't describe new path trace features. The system described here allows to set up an SR domain wide centralized connectivity validation, which is useful in large network operator domains.
- o The system sending the monitoring packet is also receiving it. The payload of the monitoring packet may be chosen freely. This allows sending probing packets which represent customer traffic, possibly from multiple services (e.g., small Voice over IP packet, larger HTTP packets) and embedding of useful monitoring data (e.g., accurate time stamps since both sender and receiver have the same clock, sequence numbers to ease the measurement...).
- o Set up of a flexible MPLS monitoring system in terms of deployment: from one single centralized one to a set of distributed systems (e.g., on a per region or service base), and in terms of redundancy from 1+1 to N+1.

In addition to monitoring paths, problem localization is required. Topology awareness is an important feature of link state IGPs deployed by operators of large networks. MPLS topology awareness combined with IGP topology awareness enables a simple and scalable data plane based monitoring mechanism. Faults can be localized:

- o by capturing the Interior Gateway Protocol (IGP) topology and analyzing 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.

MPLS OAM offers flexible traceroute (connectivity verification) features to detect and execute data paths of an MPLS domain. By utilizing the Equal Cost Multipath (ECMP) related tool set offered, e.g., by [RFC 8029](#) [[RFC8029](#)], 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, which can be 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 from a variety of protocols can be used to check continuity. These include Internet Control Message Protocol [[RFC0792](#)] [[RFC4443](#)] [[RFC4884](#)] [[RFC4950](#)], Bidirectional Forwarding Detection (BFD) [[RFC5884](#)], Seamless Bidirectional Forwarding Detection (S-BFD) [[RFC7880](#)] [[RFC7881](#)] (see [Section 3.4 of \[RFC7882\]](#)), and MPLS LSP Ping [[RFC8029](#)]. They can also 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 monitoring requires pushing a large label stack, a software based implementation is usually more flexible than an hardware based one. Hence router label stack depth and label composition limitations don't limit MPLS OAM choices.

[I-D.ietf-mpls-spring-lsp-ping] discusses SR OAM applicability and MPLS traceroute enhancements adding functionality to the use cases described by this document.

2. Terminology and Acronyms

2.1. Terminology

Continuity Check

is defined in [Section 2.2.7 of RFC 7276](#) [[RFC7276](#)].

Connectivity Verification

is defined in [Section 2.2.7 of RFC 7276](#) [[RFC7276](#)].

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 utilize 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](#)].

2.2. 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

3. 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 the 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 correspondence 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.

4. SR-based Path Monitoring Use Case Illustration

4.1. 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 [RFC 8029](#) [[RFC8029](#)] (using the Downstream (Detailed) Mapping information resulting from a path trace). 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 [RFC 8029](#) [[RFC8029](#)] 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 8029](#) [[RFC8029](#)] 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, all Adj-SIDs along that path should be part of the label stack.

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 MPLS LSP Ping [RFC 8029](#) [[RFC8029](#)]. They may also incorporate the additional enhancements defined in [[I-D.ietf-mpls-spring-lsp-ping](#)] to incorporate further MPLS trace route features. ICMP Ping based continuity checks don't require router control plane activity. Prior to monitoring a path, MPLS OAM may be used to detect ECMP dependent 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.

[4.2.](#) Use Case 2 - Monitoring a Remote Bundle



SR based probing of all the links of a remote bundle

Figure 2

In the figure, R1 addresses Link "x" Lx by the Adjacency SID 99x, while R2 addresses Link 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. 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.

4.3. 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 problem is that for some (possibly unknown) reason SR packets to be forwarded from R2 via the interface identified by Adjacency SID 663 are lost.

The example above only illustrates a method to localize 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.

5. Path Trace and Failure Notification

Sometimes forwarding along a single path indeed doesn't work, while the control plane information is healthy. Such a situation may occur after maintenance work within a domain. An operator may perform on demand-tests, but execution of automated PMS path trace checks may be set up too (scope may be limited to a subset of important end-to-end paths crossing the router or network section after completion of the maintenance work there). Upon detection of a path which can't be used, the operator needs to be notified. A check ensuring that re-routing event is differed from a path facing whose forwarding behavior doesn't correspond to the control plane information is necessary (but out of scope of this document).

Adding an automated problem solution to the PMS features only makes sense, if the root cause of the symptom appears often, can be assumed to be non-ambiguous by its symptoms, can be solved by a pre-determined chain of commands and the automated PMS reaction not doing any collateral damage. A closer analysis is out of scope of this document.

The PMS is expected to check control plane liveness after a path repair effort was executed. It doesn't matter whether the path repair was triggered manually or by an automated system.

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

The MPLS path monitoring system described by this document can be realized with non-Segment Routing (SR) based technology. Making such a non-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 tLDP 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. Enhancing IGPs to exchange of MPLS topology information as done by SR significantly simplifies and stabilizes 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 example given in figure 1, 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 standard conformant Measurement Agents ([[RFC6808](#)] connected to an MPLS domain at three different sites). The delay measurements of the PMS and the IPPM Measurement Agents were compared based on a statistical test in [[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.


7. PMS Monitoring of Different Segment ID Types

MPLS SR topology awareness should allow the PMS 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 [\[I-D.ietf-mpls-spring-lsp-ping\]](#).

To match control plane information with data plane information for all relevant types of Segment IDs, [\[I-D.ietf-mpls-spring-lsp-ping\]](#) enhances MPLS OAM functions defined by [RFC 8029](#) [\[RFC8029\]](#).

8. 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.



Connectivity verification with a PMS

Figure 3

Let's assign the following Node SIDs to the nodes of the figure: PMS = 10, LER i = 20, LER j = 30, LER m = 40. PMS is intended to validate the path between LER m and LER j. In order to validate this

path, PMS will send the probe packet with label stack of (top to bottom): {40} {30} {10}. Imagine any of the below forwarding entry misprogrammed situation:

- o LSRa receiving any packet with top label 40 will POP and forwards to LSR1 instead of LER m.
- o LSR1 receiving any packet with top label 30 will pop and forward to LER i instead of LER j.

In any of these above situation, the probe packet will be delivered back to PMS leading to a falsified path liveliness indication by the PMS.

Connectivity Verification functions helps us to verify if the probe is taking the expected path. For example, PMS can intermittently send the probe packet with label stack of (top to bottom): {40;ttl=255} {30;ttl=1} {10;ttl=255}. The probe packet may carry information about LER m which could be carried in Target FEC Stack in case of MPLS Echo Request or Discriminator in case of Seamless BFD. When LER m receives the packet, it will punt due to TTL expiry and sends a positive response. In the above mentioned misprogramming situation, LSRa will forwards to LSR1 which will send a negative response to PMS as the information in probe does not match the local node. PMS can do the same for bottom label as well. This will help perform connectivity verification and ensure that the path between LER m and LER j is working as expected.

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

The PMS builds packets with intent of performing OAM tasks. It uses address information based on topology information, rather than a protocol.

The PMS allows to insert traffic into non-SR domains. This may be required in the case of an LDP domain attached to the SR domain, but it can be used to compromise security in the case of external IP domains and MPLS based VPNs.

To avoid a PMS to insert traffic into an MPLS VPN domain, one or more sets of label ranges may be reserved for service labels within an SR domain. The PMS should be configured to reject usage of these service label values. In the same way, misuse of IP destination

addresses is blocked if only IP-destination address values conforming to [RFC 8029](#) [[RFC8029](#)] are settable by the PMS.

To limit potential misuse, access to a PMS needs to be authorized and should be logged. OAM supported by a PMS requires skilled personal and hence only experts requiring PMS access should be allowed to access such a system. It is recommended to directly attach a PMS to an SR domain. Connecting a PMS to an SR domain is technically possible, but adds further security issues. A tunnel based access of a PMS to an SR domain is not recommended.

Use of stale MPLS or IGP routing information could cause a PMS monitoring packet to leave the domain where it originated. PMS monitoring packets should not be sent using stale MPLS or IGP routing information. As it is necessary to know that the information is stale in order to follow the instruction, as is the case with for example convergence events that may be ongoing at the time of diagnostic measurement.

Traffic insertion by a PMS may be unintended, especially if the IGP or MPLS topology stored locally are in stale state. As soon as the PMS has an indication, that its IGP or MPLS topology are stale, it should stop operations involving network sections whose topology may not be accurate. Note however that it is a task of an OAM system to discover and locate network sections having where forwarding behavior is not matching control plane state. As soon as a PMS or an operator of a PMS has the impression, that the PMS topology information is stale, measures need to be taken to refresh the topology information. These measures should be part of the PMS design. Matching forwarding and control plane state by periodically automated execution of [RFC 8029](#) [[RFC8029](#)] mechanisms may be such a feature. Whenever network maintenance tasks are performed by operators, the PMS topology discovery should be started asynchronously after network maintenance has been finished.

A PMS loosing network connectivity or crashing must remove all IGP and MPLS topology information prior to restarting operation.

A PMS may operate routine measurements. If these are automated, care must be taken to avoid unintended traffic insertion. On the other hand, large scale operation based on operator configuration itself is a source of unintended misconfigurations and should be avoided.

[11.](#) Acknowledgements

The authors would like to thank Nobo Akiya for his contribution. Raik Leipnitz 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. Finally, Bruno Decraene's shepherd review led to a clarified document.

12. References

12.1. Normative References

- [I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S.,
and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-12](#) (work in progress), June 2017.
- [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>>.

12.2. Informative References

- [I-D.ietf-mpls-spring-lsp-ping]
Kumar, N., Swallow, G., Pignataro, C., Akiya, N., Kini,
S., Gredler, H., and M. Chen, "Label Switched Path (LSP)
Ping/Traceroute for Segment Routing Networks with MPLS
Dataplane", [draft-ietf-mpls-spring-lsp-ping-03](#) (work in
progress), June 2017.
- [I-D.leipnitz-spring-pms-implementation-report]
Leipnitz, R. and R. Geib, "A scalable and topology aware
MPLS data plane monitoring system", [draft-leipnitz-spring-pms-implementation-report-00](#) (work in progress), June
2016.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5,
[RFC 792](#), DOI 10.17487/RFC0792, September 1981,
<<http://www.rfc-editor.org/info/rfc792>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet
Control Message Protocol (ICMPv6) for the Internet
Protocol Version 6 (IPv6) Specification", STD 89,
[RFC 4443](#), DOI 10.17487/RFC4443, March 2006,
<<http://www.rfc-editor.org/info/rfc4443>>.
- [RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro,
"Extended ICMP to Support Multi-Part Messages", [RFC 4884](#),
DOI 10.17487/RFC4884, April 2007,
<<http://www.rfc-editor.org/info/rfc4884>>.

- [RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP Extensions for Multiprotocol Label Switching", [RFC 4950](#), DOI 10.17487/RFC4950, August 2007, <<http://www.rfc-editor.org/info/rfc4950>>.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", [RFC 5884](#), DOI 10.17487/RFC5884, June 2010, <<http://www.rfc-editor.org/info/rfc5884>>.
- [RFC6576] Geib, R., Ed., Morton, A., Fardid, R., and A. Steinmitz, "IP Performance Metrics (IPPM) Standard Advancement Testing", [BCP 176](#), [RFC 6576](#), DOI 10.17487/RFC6576, March 2012, <<http://www.rfc-editor.org/info/rfc6576>>.
- [RFC6808] Ciavattone, L., Geib, R., Morton, A., and M. Wieser, "Test Plan and Results Supporting Advancement of [RFC 2679](#) on the Standards Track", [RFC 6808](#), DOI 10.17487/RFC6808, December 2012, <<http://www.rfc-editor.org/info/rfc6808>>.
- [RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", [RFC 7880](#), DOI 10.17487/RFC7880, July 2016, <<http://www.rfc-editor.org/info/rfc7880>>.
- [RFC7881] Pignataro, C., Ward, D., and N. Akiya, "Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS", [RFC 7881](#), DOI 10.17487/RFC7881, July 2016, <<http://www.rfc-editor.org/info/rfc7881>>.
- [RFC7882] Aldrin, S., Pignataro, C., Mirsky, G., and N. Kumar, "Seamless Bidirectional Forwarding Detection (S-BFD) Use Cases", [RFC 7882](#), DOI 10.17487/RFC7882, July 2016, <<http://www.rfc-editor.org/info/rfc7882>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", [RFC 8029](#), DOI 10.17487/RFC8029, March 2017, <<http://www.rfc-editor.org/info/rfc8029>>.

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

