

Internet Engineering Task Force (IETF)  
Request for Comments: 8403  
Category: Informational  
ISSN: 2070-1721

R. Geib, Ed.  
Deutsche Telekom  
C. Filsfils  
C. Pignataro, Ed.  
N. Kumar  
Cisco Systems, Inc.  
July 2018

## A Scalable and Topology-Aware MPLS Data-Plane Monitoring System

### Abstract

This document describes features of an MPLS 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 Label Switched Path (LSP) trace, in a very complementary way. MPLS topology awareness reduces management and control-plane involvement of Operations, Administration, and Maintenance (OAM) measurements while enabling new OAM features.

### Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8403>.

Copyright Notice

Copyright (c) 2018 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 (https://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 . . . . . 3
- 2. Terminology and Abbreviations . . . . . 5
  - 2.1. Terminology . . . . . 5
  - 2.2. Abbreviations . . . . . 6
- 3. An MPLS Topology-Aware Path Monitoring System . . . . . 6
- 4. Illustration of an SR-Based Path Monitoring Use Case . . . . . 8
  - 4.1. Use Case 1: LSP Data-Plane Monitoring . . . . . 8
  - 4.2. Use Case 2: Monitoring a Remote Bundle . . . . . 11
  - 4.3. Use Case 3: Fault Localization . . . . . 12
- 5. Path Trace and Failure Notification . . . . . 12
- 6. Applying SR to Monitoring LSPs That Are Not SR Based (LDP and Possibly RSVP-TE) . . . . . 13
- 7. PMS Monitoring of Different Segment ID Types . . . . . 14
- 8. Connectivity Verification Using PMS . . . . . 14
- 9. IANA Considerations . . . . . 15
- 10. Security Considerations . . . . . 15
- 11. References . . . . . 17
  - 11.1. Normative References . . . . . 17
  - 11.2. Informative References . . . . . 17
- Acknowledgements . . . . . 19
- Authors' Addresses . . . . . 19

## 1. Introduction

Network operators need to be able to monitor the forwarding paths used to transport user packets. Monitoring packets are expected to be forwarded in the data plane in a similar way to user packets. Segment Routing (SR) enables forwarding of packets along predefined paths and segments; thus, an SR monitoring packet can stay in the data plane while passing along one or more segments to be monitored.

This document describes a system as a functional component called (MPLS) Path Monitoring System or PMS. The PMS uses capabilities for MPLS data-plane path monitoring. The use cases introduced here are limited to a single Interior Gateway Protocol (IGP) MPLS domain. The use cases of this document refer to the PMS realized as a separate node. Although many use cases depict the PMS as a physical node, no assumption should be made, and the node could be virtual. This system is defined as a functional component abstracted to have many realizations. The terms "PMS" and "system" are used interchangeably here.

The system applies to the monitoring of non-SR LSPs like Label Distribution Protocol (LDP) as well as to the monitoring of SR LSPs (Section 7 offers some more information). As compared to non-SR approaches, SR 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 also detected by IGP signaling. Thus, a centralized and MPLS-topology-aware monitoring unit can be realized in an SR 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 ability to set up an SR-domain-wide centralized connectivity validation. Many operators of large networks regard a centralized monitoring system as useful.
- o The MPLS ping (or continuity check) packets never leave the MPLS user data plane.
- o SR allows the transport of MPLS path trace or connectivity validation packets for every LSP to all nodes of an SR domain. This use case doesn't describe new path-trace features. The system described here allows for the set up of 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 probing packets to be sent that represent customer traffic, possibly from multiple services (e.g., small Voice over IP packets, larger HTTP packets), and allows the embedding of useful monitoring data (e.g., accurate timestamps since both sender and receiver have the same clock and 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 basis), 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 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 ECMP-related tool set offered, e.g., by RFC 8029 [RFC8029], an SR-based MPLS monitoring system can be enabled to:

- o detect how to route packets along different ECMP-routed paths.
- o Construct ping packets that can be steered along a single path or ECMP towards a particular LER/LSR whose connectivity is to be checked.
- 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 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 PMS 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 of learning and maintaining 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 must further 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 (ICMP) [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 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 a hardware-based one. Hence, router label stack depth and label composition limitations don't limit MPLS OAM choices.

RFC 8287 [RFC8287] discusses SR OAM applicability and MPLS traceroute enhancements adding functionality to the use cases described by this document.

The document describes both use cases and a standalone monitoring framework. The monitoring system reuses existing IETF OAM protocols and leverage Segment Routing (Source Routing) to allow a single device to send, have exercised, and receive its own probing packets. As a consequence, there are no new interoperability considerations. A Standards Track RFC is not required; Informational status for this document is appropriate

## 2. Terminology and Abbreviations

### 2.1. Terminology

#### Continuity Check

See Section 2.2.7 of RFC 7276 [RFC7276].

#### Connectivity Verification

See 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 PMS 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 [RFC8402].

## 2.2. Abbreviations

ECMP	Equal-Cost Multipath
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 that is able to learn the IGP Link State Database (LSDB) (including the SIDs) is able to execute arbitrary chains of LSPs. To monitor an MPLS SR domain, a PMS needs to set up a topology database 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 only on the topology information. In addition, the

PMS can be used to trace MPLS LSP and, thus, verify their connectivity and correspondence between control and data planes, respectively. The PMS can direct suitable MPLS traceroute packets to any node along a path segment.

Let us describe how the PMS constructs a label 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 Adj-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 an 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 that is part of the SR domain. It must be able to send and receive MPLS packets via this interface. As mentioned above, the 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. The option to connect a PMS to an MPLS domain by a tunnel may be attractive to some operators. So far, MPLS separates networks securely by avoiding tunnel access to MPLS domains. Tunnel-based access of a PMS to an MPLS domain is out of scope of this document, as it implies additional security aspects.

4. Illustration of an SR-Based Path Monitoring Use Case

4.1. Use Case 1: LSP Data-Plane Monitoring

Figure 1 shows an example of this functional component as a system, which can be physical or virtual.




Figure 1: Example of a PMS-Based LSP Data-Plane Monitoring

For the sake of simplicity, let’s assume that all the nodes are configured with the same SRGB [RFC8402].

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, the 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 that has been



detected by the MPLS traceroute 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 Penultimate Hop Popping is 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 from PMS to LER i must be available (i.e., a continuity check along the path to LER i must succeed). If desired, an MPLS traceroute 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 that a packet may use between LER i and LER j may be desired. This can be done by applying an MPLS traceroute between LER i and LER j. Another option is to use SR, 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 correlation between control and data planes.
- 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 traceroute 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 the MPLS OAM functionality described in RFC 8029 [RFC8029]. All monitoring packets stay on the data plane; 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 this document

describes path connectivity checking as a basic application, additional monitoring (like checking continuity of underlying physical infrastructure or performing delay measurements) may be desired. A change in ECMP routing that is not caused by an IGP or MPLS topology change may not be desirable for connectivity checks and delay measurements. Therefore, a PMS should also periodically verify connectivity of the SR paths that 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 connectivity verification as specified in RFC 8029 [RFC8029] 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.

While a single PMS can detect the complete MPLS control- and data-plane topology, a reliable deployment requires two separated PMSs. Scalable permanent surveillance of a set of LSPs could require deployment of several PMSs. The PMS may be a router, but could also be a 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 PMSs 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 RFC 8287 [RFC8287] to incorporate further MPLS traceroute 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 for the monitoring of these paths with label stacks reduced to a limited number of Node-

SIDs resulting from Shortest Path First (SPF) routing. The PMS does not require access to information about LSR/LER management or data planes to do so.

4.2. Use Case 2: Monitoring a Remote Bundle




Figure 2: SR-Based Probing of All the Links of a Remote Bundle

In the figure, R1 addresses Link "x" Lx by the Adj-SID 99x, while R2 addresses Link Lx by the Adj-SID 66(x+1).

In the above figure, the PMS needs to assess the data-plane 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).

The 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 (Adj-SID 662). R1 forwards the probe to R2 over link L2 (Adj-SID 992). R2 forwards the probe to R1 over link L3 (Adj-SID 664). R1 then forwards the IP probe to the PMS as per classic IP forwarding.

As was mentioned in Section 4.1, the PMS must be able to monitor the continuity of the path PMS to R2 (Node-SID 72) as well as the 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 unidirectional 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 Adj-SID 663 of the middle link from R2 to R1 in the unsuccessful 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 Adj-SID 663 are lost.

The example above only illustrates a method to localize a fault by correlated continuity checks. Any operational deployment requires well-designed engineering to allow for the desired unambiguous 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 doesn't work, even though 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 as well (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 that can't be used, the operator needs to be notified. A check ensuring that a 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 unambiguous by its symptoms, can be solved by a predetermined chain of commands, is not collaterally damaged by the automated PMS reaction. 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 LSPs That Are Not SR Based (LDP and Possibly RSVP-TE)

The MPLS PMS described by this document can be realized with technology that is not SR based. Making such a monitoring system that is not SR MPLS based 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 aligned with IGP topology in a timely manner. Enhancing IGP to the exchange of MPLS-topology information as done by SR significantly simplifies and stabilizes such an MPLS PMS.

An SR-based PMS connected to an MPLS domain consisting of LER and LSRs supporting SR and LDP or RSVP-TE in parallel in all nodes may use SR paths to transmit packets to and from the start and endpoints of LSPs that are not SR based 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 the IGP and use this information.

An implementation report on a PMS operating in an LDP domain is given in [MPLS-PMS-REPORT]. In addition, this report compares delays measured with a single PMS to the results measured by systems that are conformant with IP Performance Metrics (IPPM) connected to the MPLS domain at three sites (see [RFC6808] for IPPM conformance). 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 node that is not SR capable may fail. This and other backward compatibility with non-SR devices are discussed by RFC 8287 [RFC8287].

To match control-plane information with data-plane information for all relevant types of Segment IDs, RFC 8287 [RFC8287] 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 checks between LER i and LER j, they may not help perform connectivity verification.




Figure 3: Connectivity Verification with a PMS

Let's assign the following Node-SIDs to the nodes of the figure: PMS = 10, LER i = 20, LER j = 30, LER m = 40. The PMS is intended to validate the path between LER m and LER j. In order to validate this path, the PMS will send the probe packet with a 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 either of the above situations, the probe packet will be delivered back to the PMS leading to a falsified path liveliness indication by the PMS.

Connectivity Verification functions help us to verify if the probe is taking the expected path. For example, the PMS can intermittently send the probe packet with a 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 the Target FEC Stack in case of an MPLS Echo Request or Discriminator in the case of Seamless BFD. When LER m receives the packet, it will punt due to Time-To-Live (TTL) expiry and send a positive response. In the above-mentioned misprogramming situation, LSRa will forward to LSR1, which will send a negative response to the PMS as the information in probe does not match the local node. The 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 document has no IANA actions.

## 10. Security Considerations

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

The PMS allows the insertion of 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 maliciously insert traffic in the case of external IP domains and MPLS-based VPNs.

To prevent a PMS from inserting 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 personnel; 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 by a tunnel 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. To carry out a desired measurement properly, the PMS must be aware of and respect the actual route changes, convergence events, as well as the assignment of Segment IDs relevant for measurements. At a minimum, the PMS must be able to listen to IGP topology changes or pull routing and segment information from routers signaling topology changes.

Traffic insertion by a PMS may be unintended, especially if the IGP or MPLS topology stored locally is 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. However, note that it is the task of an OAM system to discover and locate network sections 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 the mechanisms described in RFC 8029 [RFC8029] 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 that is losing network connectivity or crashing must remove all IGP- and MPLS-topology information prior to restarting operation.

A PMS may operate routine measurements on a large scale. Care must be taken to avoid unintended traffic insertion after topology changes that result in, e.g., changes of label assignments to routes or interfaces within a domain. If the labels concerned are part of the



label stack composed by the PMS for any measurement packet and their state is stale, the measurement initially needs to be stopped. Setup and operation of routine measurements may be automated. Secure automated PMS operation requires a working automated detection and recognition of stale routing state.

## 11. References

### 11.1. Normative References

- [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, <<https://www.rfc-editor.org/info/rfc7276>>.
- [RFC8402] Filisfilis, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

### 11.2. Informative References

- [MPLS-PMS-REPORT] Leipzig, R., Ed. and R. Geib, "A scalable and topology aware MPLS data plane monitoring system", Work in Progress, draft-leipzig-spring-pms-implementation-report-00, June 2016.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8287] Kumar, N., Ed., Pignataro, C., Ed., Swallow, G., Akiya, N., Kini, S., and M. Chen, "Label Switched Path (LSP) Ping/Traceroute for Segment Routing (SR) IGP-Prefix and IGP-Adjacency Segment Identifiers (SIDs) with MPLS Data Planes", RFC 8287, DOI 10.17487/RFC8287, December 2017, <<https://www.rfc-editor.org/info/rfc8287>>.

## 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 Document Shepherd review led to a clarified document.

## 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 Filselfil  
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  
United States of America

Email: cpignata@cisco.com

Nagendra Kumar  
Cisco Systems, Inc.  
7200 Kit Creek Road  
Research Triangle Park, NC 27709-4987  
United States of America

Email: naikumar@cisco.com