

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 30, 2017

A. D'Alessandro
Telecom Italia
L. Andersson
Huawei Technologies
S. Ueno
NTT Communications
K. Arai
Y. Koike
NTT
July 29, 2016

Hitless path segment monitoring
draft-ietf-mpls-tp-temporal-hitless-psm-11.txt

Abstract

One of the most important OAM capabilities for transport network operation is fault localisation. An in-service, on-demand segment monitoring function of a transport path is indispensable, particularly when the service monitoring function is activated only between end points. However, the current segment monitoring approach defined for MPLS (including the transport profile (MPLS-TP)) in [RFC 6371](#) [[RFC6371](#)] has drawbacks. This document provides an analysis of the existing MPLS-TP OAM mechanisms for the path segment monitoring and provides requirements to guide the development of new OAM tools to support a Hitless Path Segment Monitoring (HPSM).

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 30, 2017.

Copyright Notice

Copyright (c) 2016 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.	Conventions used in this document	3
2.1.	Terminology	3
2.2.	Definitions	4
3.	Problem Statement	4
4.	Requirements for Hitless Path Segment Monitoring	7
4.1.	Backward compatibility	7
4.2.	Non-intrusive segment monitoring	8
4.3.	Multiple segments monitoring	8
4.4.	Single and multiple level monitoring	8
4.5.	HPSM and end-to-end proactive monitoring independence . .	9
4.6.	Arbitrary segment monitoring	10
4.7.	Fault while HPSM is operational	11
4.8.	HPSM Manageability	12
4.9.	Supported OAM functions	13
5.	Summary	13
6.	Security Considerations	14
7.	IANA Considerations	14
8.	Contributors	14
9.	Acknowledgements	14
10.	References	15
10.1.	Normative References	15
10.2.	Informative References	15
	Authors' Addresses	15

[1.](#) Introduction

According to the MPLS-TP OAM requirements [RFC 5860](#) [[RFC5860](#)], mechanisms MUST be available for alerting service providers of faults or defects that affects their services. In addition, to ensure that faults or service degradation can be localized, operators need a

function to diagnose the detected problem. Using end-to-end monitoring for this purpose is insufficient in that an operator will not be able to localize a fault or service degradation accurately.

Thus, a segment monitoring function that can focus on a specific segment of a transport path and can provide a detailed analysis is indispensable to promptly and accurately localize the fault. For MPLS-TP, a path segment monitoring function has been defined to perform this task. However, as noted in the MPLS-TP OAM Framework [RFC 6371](#) [[RFC6371](#)], the current method for segment monitoring of a transport path has implications that hinder the usage in an operator network.

This document, after elaborating on the problem statement for the path segment monitoring function as it is currently defined, provides requirements for an on-demand segment monitoring function without traffic disruption.

2. Conventions used in this document

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

2.1. Terminology

ATM - Asynchronous Transfer Mode

HPSM - Hitless Path Segment Monitoring

LSP - Label Switched Path

LSR - Label Switching Router

ME - Maintenance Entity

MEG - Maintenance Entity Group

MEP - Maintenance Entity Group End Point

MIP - Maintenance Entity Group Intermediate Point

OTN - Optical Transport Network

TCM - Tandem connection monitoring

SPME - Sub-path Maintenance Element

2.2. Definitions

None.

3. Problem Statement

To monitor (and to protect and/or manage) MPLS-TP network segments a Sub-Path Maintenance Element (SPME) function has been defined in [RFC 5921](#) [[RFC5921](#)]. The SPME is defined between the edges of the segment of a transport path that needs to be monitored, protected, or managed. SPME is created by stacking the shim header (MPLS header) according to [RFC 3031](#) [[RFC3031](#)] and it is defined as the segment where the header is stacked. OAM messages can be initiated at the edge of the SPME and sent to the peer edge of the SPME or to a MIP along the SPME by setting the TTL value of the label stack entry (LSE) and interface identifier value at the corresponding hierarchical LSP level in case of a per-node model.

MPLS-TP segment monitoring should satisfy two network objectives according to [section 3.8 of RFC 6371](#) [[RFC6371](#)]:

(N1) The monitoring and maintenance of current transport paths has to be conducted in-service without traffic disruption.

(N2) Segment monitoring must not modify the forwarding of the segment portion of the transport path.

The SPME function that has been defined in [RFC 5921](#) [[RFC5921](#)] has the following drawbacks:

(P1) It increases network management complexity, because a new sublayer and new MEPS and MIPs have to be configured for the SPME.

(P2) Original conditions of the path change.

(P3) The client traffic over a transport path is disrupted if the SPME is configured on-demand.

Problem (P1) is related to the management of each additional sub-layer required for segment monitoring in a MPLS-TP network. When an SPME is applied to administer on-demand OAM functions in MPLS-TP networks, a rule for operationally differentiating those SPME will be required at least within an administrative domain. This forces operators to implement at least an additional layer into the management systems that will only be used for on-demand path segment monitoring. From the perspective of operation, increasing the number of managed layers and managed addresses/identifiers is not desirable in view of keeping the management systems as simple as possible.

Moreover, using the currently defined methods, on-demand setting of SPMEs causes problems (P2) and (P3) due to additional label stacking.

Problem (P2) arises from the fact that MPLS exposed label value and MPLS frames length changes. The monitoring function should monitor the status without changing any conditions of the targeted, to be monitored, segment or transport path. Changing the settings of the original shim header should not be allowed because this change corresponds to creating a new segment of the original transport path that differs from the original one. When the conditions of the path change, the measured values or observed data will also change and this may make the monitoring meaningless because the result of the measurement would no longer reflect the performance of the connection where the original fault or degradation occurred. As an example, setting up an on-demand SPME will result in the LSRs within the monitoring segment only looking at the added (stacked) labels and not at the labels of the original LSP. This means that problems stemming from incorrect (or unexpected) treatment of labels of the original LSP by the nodes within the monitored segment cannot be identified when setting up SPME. This might include hardware problems during label look-up, mis-configuration, etc. Therefore operators have to pay extra attention to correctly setting and checking the label values of the original LSP in the configuration. Of course, the reverse of this situation is also possible, e.g., an incorrect or unexpected treatment of SPME labels can result in false detection of a fault where no problem existed originally.

Figure 1 shows an example of SPME settings. In the figure, "X" is the label value of the original path expected at the tail-end of node D. "210" and "220" are label values allocated for SPME. The label values of the original path are modified as well as the values of the stacked labels. As shown in Figure 1, SPME changes both the length of MPLS frames and the label value(s). In particular, performance monitoring measurements (e.g. Delay Measurement and Packet Loss Measurement) are sensitive to these changes.

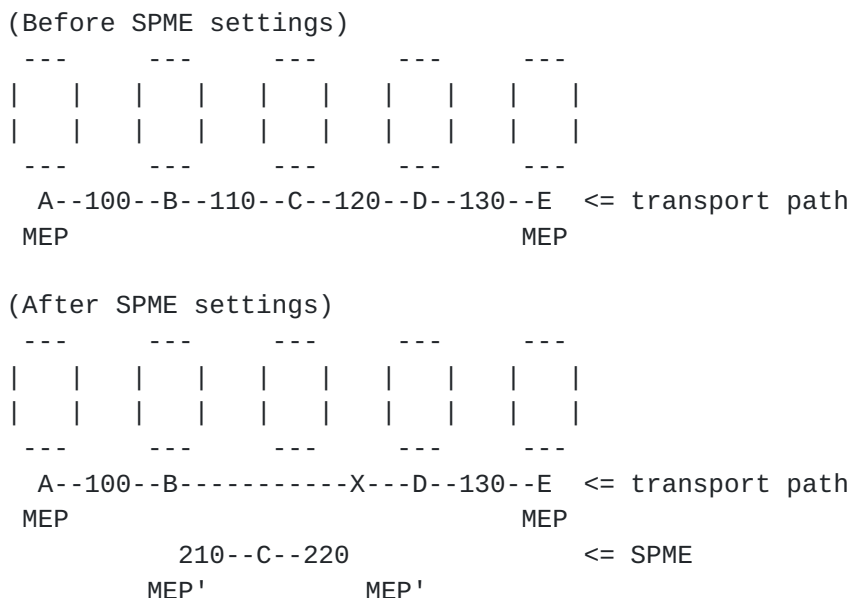


Figure 1: SPME settings example

Problem (P3) can be avoided if the operator sets SPMEs in advance and maintains them until the end of life of a transport path. But this does not support on-demand. Furthermore SPMEs cannot be set arbitrarily because overlapping of path segments is limited to nesting relationships. As a result, possible SPME configurations of segments of an original transport path are limited due to the characteristic of the SPME shown in Figure 1, even if SPMEs are pre-configured.

Although the make-before-break procedure in the survivability document [RFC 6372](#) [RFC6372] supports configuration for monitoring according to the framework document [RFC 5921](#) [RFC5921], without traffic disruption, the configuration of an SPME is not possible without violating network objective (N2). These concerns are described in [section 3.8 of RFC 6371](#) [RFC6371].

Additionally, the make-before-break approach typically relies on a control plane and requires additional functionalities for a management system to properly support SPME creation and traffic switching from the original transport path to the SPME.

As an example, the old and new transport resources (e.g. LSP tunnels) might compete with each other for resources which they have in common. Depending on availability of resources, this competition can cause admission control to prevent the new LSP tunnel from being established as this bandwidth accounting deviates from traditional (non control plane) management system operation. While SPMEs can be

applied in any network context (single domain, multi domain, single carrier, multi carrier, etc.), the main applications are in inter-carrier or inter-domain segment monitoring where they are typically pre-configured or pre-instantiated. SPME instantiates a hierarchical path (introducing MPLS label stacking) through which OAM packets can be sent. The SPME monitoring function is also mainly important for protecting bundles of transport paths and carriers' carrier solutions within an administrative domain.

The analogy for SPME in other transport technologies is Tandem Connection Monitoring (TCM), used in Optical Transport Networks (OTN) and Ethernet transport networks, which supports on-demand but does not affect the path. For example in OTN, TCM allows the insertion and removal of performance monitoring overhead within the frame at intermediate points in the network. It is done such that their insertion and removal do not change the conditions of the path. Though as the OAM overhead is part of the frame (designated overhead bytes), it is constrained to a pre-defined number of monitoring segments.

To summarize: the problem statement is that the current sub-path maintenance based on a hierarchical LSP (SPME) is problematic for pre-configuration in terms of increasing the number of managed objects by layer stacking and identifiers/addresses. An on-demand configuration of SPME is one of the possible approaches for minimizing the impact of these issues. However, the current procedure is unfavourable because the on-demand configuration for monitoring changes the condition of the original monitored path. To avoid or minimize the impact of the drawbacks discussed above, a more efficient approach is required for the operation of an MPLS-TP transport network. A monitoring mechanism, named Hitless Path Segment Monitoring (HPSM), supporting on-demand path segment monitoring without traffic disruption is needed.

4. Requirements for Hitless Path Segment Monitoring

In the following sections, mandatory (M) and optional (O) requirements for the Hitless Path Segment Monitoring function are listed.

4.1. Backward compatibility

HPSM would be an additional OAM tool that would not replace SPME. As such:

(M1) HPSM MUST be compatible with the usage of SPME

(O1) HPSM SHOULD be applicable at the SPME layer too

(M2) HPSM MUST support both the per-node and per-interface model as specified in [RFC 6371](#) [[RFC6371](#)].

4.2. Non-intrusive segment monitoring

One of the major problems of legacy SPME highlighted in [section 3](#) is that it may not monitor the original path and it could disrupt service traffic when set-up on demand.

(M3) HPSM MUST NOT change the original conditions of transport path (e.g. must not change the length of MPLS frames, the exposed label values, etc.)

(M4) HPSM MUST support on-demand provisioning without traffic disruption.

4.3. Multiple segments monitoring

Along a transport path there may be the need to support simultaneously monitoring multiple segments

(M5) HPSM MUST support configuration of multiple monitoring segments along a transport path.




Figure 2: Multiple HPSM instances example

4.4. Single and multiple level monitoring

HPSM would apply mainly for on-demand diagnostic purposes. With the currently defined approach, the most serious problem is that there is no way to locate the degraded segment of a path without changing the conditions of the original path. Therefore, as a first step, a single level, single segment monitoring, not affecting the monitored path, is required for HPSM. A combination of multi-level and simultaneous segments monitoring is the most powerful tool for accurately diagnosing the performance of a transport path. However, in the field, a single level, multiple segments approach would be less complex for management and operations.

- (M6) HPSM MUST support single-level segment monitoring
- (O2) HPSM MAY support multi-level segment monitoring.

Figure 3 shows an example of multi-level HPSM.

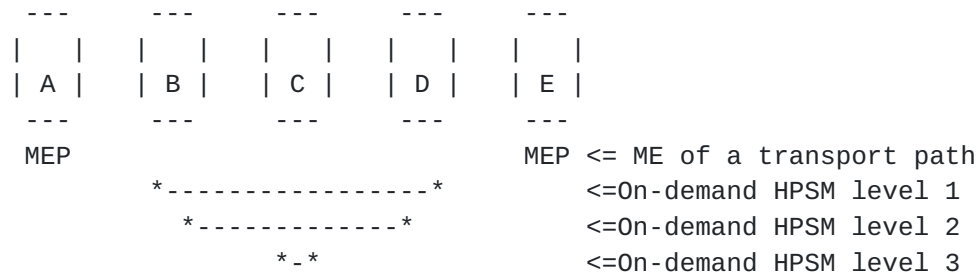


Figure 3: Multi-level HPSM example

4.5. HPSM and end-to-end proactive monitoring independence

There is a need for simultaneously using existing end-to-end proactive monitoring and on-demand path segment monitoring. Normally, the on-demand path segment monitoring is configured on a segment of a maintenance entity of a transport path. In such an environment, on-demand single-level monitoring should be performed without disrupting the pro-active monitoring of the targeted end-to-end transport path to avoid affecting user traffic performance monitoring.

- (M7) HPSM MUST support the capability of being operated concurrently to, and independently of operated of the OAM function operated on the end-to-end path




Figure 4: Independence between proactive end-to-end monitoring and on-demand HPSM

4.6. Arbitrary segment monitoring

The main objective for on-demand segment monitoring is to diagnose the fault locations. A possible realistic diagnostic procedure is to fix one end point of a segment at the MEP of the transport path under observation and change progressively the length of the segments. This example is shown in Figure 5.

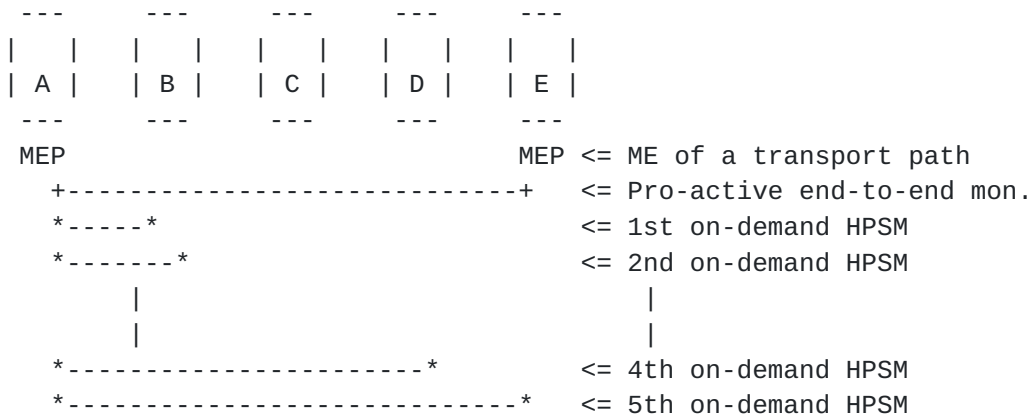


Figure 5: Localization of a defect by consecutive on-demand segment monitoring procedure

Another possible scenario is depicted in Figure 6. In this case, the operator wants to diagnose a transport path starting at a transit node, because the end nodes (A and E) are located at customer sites and consist of small boxes supporting only a subset of OAM functions. In this case, where the source entities of the diagnostic packets are limited to the position of MEPs, on-demand segment monitoring will be ineffective because not all the segments can be diagnosed (e.g. segment monitoring HPSM 3 in Figure 6 is not available and it is not possible to determine the fault location exactly).

(M8) It SHALL be possible to provision HPSM on an arbitrary segment of a transport path.

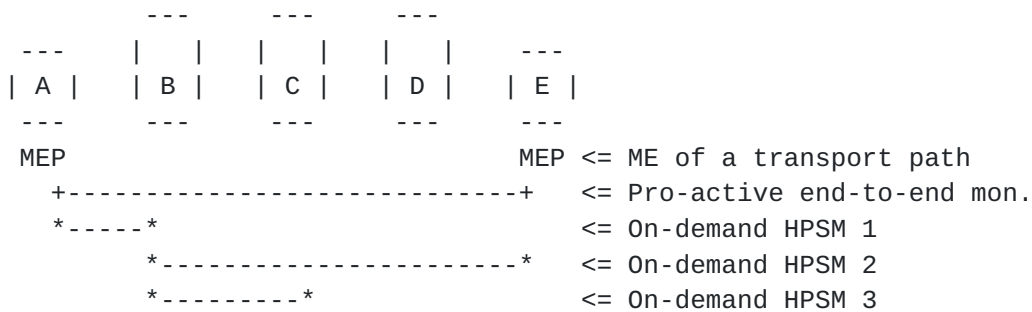


Figure 6: HPSM configuration at arbitrary segments


4.7. Fault while HPSM is operational

Node or link failures may occur while HPSM is active. In this case, if no resiliency mechanism is set-up on the subtended transport path, there is no particular requirement for HPSM. If the transport path is protected, the HPSM function should be terminated to avoid monitoring a new segment when a protection or restoration path is active.

(M9) The HPSM SHOULD avoid monitoring an unintended segment when one or more failures occur

The following examples are provided for clarification only and they are not intended to restrict any solution for meeting the requirements of HPSM.

Protection scenario A is shown in figure 7. In this scenario a working LSP and a protection LSP are set-up. HPSM is activated between nodes A and E. When a fault occurs between nodes B and C, the operation of HPSM is not affected by the protection switch and continues on the active LSP path. As a result requirement (M9) is satisfied.




Where:

- end-to-end LSP: A-B-C-D-E-F
- working LSP: A-B-C-D-E-F
- protection LSP: A-G-H-I-L-F
- HPSM: A-E

Figure 7: Protection scenario A

Protection scenario B is shown in figure 8. The difference with scenario A is that only a portion of the transport path is protected. In this case, when a fault occurs between nodes B and C on the working sub-path B-C-D, traffic will be switched to protection sub-path B-G-H-D. Assuming that OAM packet termination depends only on the TTL value of the MPLS label header, the target node of the HPSM changes from E to D due to the difference of hop counts between the working path route (A-B-C-D-E: 4 hops) and protection path route (A-B-G-H-D-E: 5 hops). As a result requirement (M9) is not satisfied.



- end-to-end LSP: A-B-C-D-E-F
- working sub-path: B-C-D
- protection sub-path: B-G-H-D
- HPSM: A-E

Figure 8: Protection scenario B

4.8. HPSM Manageability

From managing perspective, increasing the number of managed layers and managed addresses/identifiers is not desirable in view of keeping the management systems as simple as possible.

(M10)HPSM SHOULD NOT be based on additional transport layers (e.g. hierarchical LSPs)

(M11) The same identifiers used for MIPs and/or MEPs SHOULD be applied to maintenance points for the HPSM when they are instantiated in the same place along a transport path.

Anyway maintenance points for the HPSM may be different from MIPs and MEPs functional components as defined in the OAM framework document [RFC 6371](#) [RFC6371]. Investigating potential solutions for satisfying proposed HPSM requirements might lead to propose new functional components that have to be backward compatible with MPLS architecture. Solutions are outside the scope of this document.

4.9. Supported OAM functions

A maintenance point supporting the HPSM function has to be able to generate and inject OAM packets. OAM functions that may be applicable for on-demand HPSM are basically the on-demand performance monitoring functions which are defined in the OAM framework document [RFC 6371](#) [RFC6371]. The "on-demand" attribute is typically temporary for maintenance operation.

(M12) HPSM MUST support Packet Loss and Packet Delay measurement.

That because these functions are normally only supported at the end points of a transport path. If a defect occurs, it might be quite hard to locate the defect or degradation point without using the segment monitoring function. If an operator cannot locate or narrow down the cause of the fault, it is quite difficult to take prompt actions to solve the problem.

Other on-demand monitoring functions (e.g. Delay Variation measurement) are desirable but not as necessary as the functions mentioned above.

(O3) HPSM MAY support Packet Delay variation, Throughput measurement and other performance monitoring and fault management functions.

Support of out-of-service on-demand performance management functions (e.g. Throughput measurement) is not required for HPSM.

5. Summary

A new hitless path segment monitoring (HPSM) mechanism is required to provide on-demand segment monitoring without traffic disruption. It shall meet the two network objectives described in section 3.8 of [RFC 6371](#) [RFC6371] and summarized in [Section 3](#) of this document.

The mechanism should minimize the problems described in [Section 3](#), i.e. (P1), (P2) and (P3).

The solution for the on-demand segment monitoring without traffic disruption needs to cover both the per-node model and the per-interface model specified in [RFC 6371](#) [[RFC6371](#)].

The on-demand segment monitoring without traffic disruption solution needs to support on-demand Packet Loss Measurement and Packet Delay Measurement functions and optionally other performance monitoring and fault management functions (e.g. Throughput measurement, Packet Delay variation measurement, Diagnostic test, etc.).

[6.](#) Security Considerations

The security considerations defined for MPLS Transport Profile Framework in [RFC 5921](#) [[RFC5921](#)] apply to this document as well. The document provides the requirements for a new construct for performance monitoring that would make use of existing OAM tools that follow the security considerations provided in OAM Requirements for MPLS-TP in [RFC5860](#) [[RFC5860](#)].

[7.](#) IANA Considerations

There are no requests for IANA actions in this document.

Note to the RFC Editor - this section can be removed before publication.

[8.](#) Contributors

Manuel Paul

Deutsche Telekom AG

Email: manuel.paul@telekom.de

[9.](#) Acknowledgements

The authors would also like to thank Alexander Vainshtein, Dave Allan, Fei Zhang, Huub van Helvoort, Malcolm Betts, Italo Busi, Maarten Visser, Jia He and Nurit Sprecher for their comments and enhancements to the text.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), DOI 10.17487/RFC3031, January 2001, <<http://www.rfc-editor.org/info/rfc3031>>.
- [RFC5860] Vigoureux, M., Ed., Ward, D., Ed., and M. Betts, Ed., "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", [RFC 5860](#), DOI 10.17487/RFC5860, May 2010, <<http://www.rfc-editor.org/info/rfc5860>>.

10.2. Informative References

- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), DOI 10.17487/RFC5921, July 2010, <<http://www.rfc-editor.org/info/rfc5921>>.
- [RFC6371] Busi, I., Ed. and D. Allan, Ed., "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", [RFC 6371](#), DOI 10.17487/RFC6371, September 2011, <<http://www.rfc-editor.org/info/rfc6371>>.
- [RFC6372] Sprecher, N., Ed. and A. Farrel, Ed., "MPLS Transport Profile (MPLS-TP) Survivability Framework", [RFC 6372](#), DOI 10.17487/RFC6372, September 2011, <<http://www.rfc-editor.org/info/rfc6372>>.

Authors' Addresses

Alessandro D'Alessandro
Telecom Italia
Via Reiss Romoli, 274
Torino 10148
Italy

Email: alessandro.dalessandro@telecomitalia.it

Loa Andersson
Huawei Technologies

Email: loa@mail01.huawei.com

Satoshi Ueno
NTT Communications

Email: satoshi.ueno@ntt.com

Kaoru Arai
NTT

Email: arai.kaoru@lab.ntt.co.jp

Yoshinori Koike
NTT

Email: y.koike@vcd.nttbiz.com

