

Network Working Group
Request for Comments: 4687
Category: Informational

S. Yasukawa
NTT Corporation
A. Farrel
Old Dog Consulting
D. King
Aria Networks Ltd.
T. Nadeau
Cisco Systems, Inc.
September 2006

Operations and Management (OAM) Requirements
for Point-to-Multipoint MPLS Networks

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Multi-Protocol Label Switching (MPLS) has been extended to encompass point-to-multipoint (P2MP) Label Switched Paths (LSPs). As with point-to-point MPLS LSPs, the requirement to detect, handle, and diagnose control and data plane defects is critical.

For operators deploying services based on P2MP MPLS LSPs, the detection and specification of how to handle those defects are important because such defects not only may affect the fundamentals of an MPLS network, but also may impact service level specification commitments for customers of their network.

This document describes requirements for data plane operations and management for P2MP MPLS LSPs. These requirements apply to all forms of P2MP MPLS LSPs, and include P2MP Traffic Engineered (TE) LSPs and multicast LSPs.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. Conventions Used in This Document	3
2.2. Terminology	3
2.3. Acronyms	3
3. Motivations	4
4. General Requirements	4
4.1. Detection of Label Switch Path Defects	5
4.2. Diagnosis of a Broken Label Switch Path	6
4.3. Path Characterization	6
4.4. Service Level Agreement Measurement	7
4.5. Frequency of OAM Execution	8
4.6. Alarm Suppression, Aggregation, and Layer Coordination	8
4.7. Support for OAM Interworking for Fault Notification	8
4.8. Error Detection and Recovery	9
4.9. Standard Management Interfaces	9
4.10. Detection of Denial of Service Attacks	10
4.11. Per-LSP Accounting Requirements	10
5. Security Considerations	10
6. References	11
6.1. Normative References	11
6.2. Informative References	11
7. Acknowledgements	12

1. Introduction

This document describes requirements for data plane operations and management (OAM) for point-to-multipoint (P2MP) Multi-Protocol Label Switching (MPLS). This document specifies OAM requirements for P2MP MPLS, as well as for applications of P2MP MPLS.

These requirements apply to all forms of P2MP MPLS LSPs, and include P2MP Traffic Engineered (TE) LSPs [RFC4461] and [P2MP-RSVP], as well as multicast LDP LSPs [MCAST-LDP].

Note that the requirements for OAM for P2MP MPLS build heavily on the requirements for OAM for point-to-point MPLS. These latter requirements are described in [RFC4377] and are not repeated in this document.

For a generic framework for OAM in MPLS networks, refer to [RFC4378].

2. Terminology

2.1. 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].

The requirements in this document apply to OAM mechanism and protocol development, as opposed to the usual application of RFC 2119 requirements to an actual protocol, as this document does not specify a protocol.

2.2. Terminology

Definitions of key terms for MPLS OAM are found in [RFC4377] and the reader is assumed to be familiar with those definitions, which are not repeated here.

[RFC4461] includes some important definitions and terms for use within the context of P2MP MPLS. The reader should be familiar with at least the terminology section of that document.

2.3. Acronyms

The following acronyms are used in this document.

CE: Customer Edge
DoS: Denial of service
ECMP: Equal Cost Multipath

LDP: Label Distribution Protocol
LSP: Label Switched Path
LSR: Label Switching Router
OAM: Operations and Management
RSVP: Resource reSerVation Protocol
P2MP: Point-to-Multipoint
SP: Service Provider
TE: Traffic Engineering

3. Motivations

OAM for MPLS networks has been established as a fundamental requirement both through operational experience and through its documentation in numerous Internet Drafts. Many such documents (for example, [RFC4379], [RFC3812], [RFC3813], [RFC3814], and [RFC3815]) developed specific solutions to individual issues or problems. Coordination of the full OAM requirements for MPLS was achieved by [RFC4377] in recognition of the fact that the previous piecemeal approach could lead to inconsistent and inefficient applicability of OAM techniques across the MPLS architecture, and might require significant modifications to operational procedures and systems in order to provide consistent and useful OAM functionality.

This document builds on these realizations and extends the statements of MPLS OAM requirements to cover the new area of P2MP MPLS. That is, this document captures the requirements for P2MP MPLS OAM in advance of the development of specific solutions.

Nevertheless, at the time of writing, some effort had already been expended to extend existing MPLS OAM solutions to cover P2MP MPLS (for example, [P2MP-LSP-PING]). While this approach of extending existing solutions may be reasonable, in order to ensure a consistent OAM framework it is necessary to articulate the full set of requirements in a single document. This will facilitate a uniform set of MPLS OAM solutions spanning multiple MPLS deployments and concurrent applications.

4. General Requirements

The general requirements described in this section are similar to those described for point-to-point MPLS in [RFC4377]. The subsections below do not repeat material from [RFC4377], but simply give references to that document.

However, where the requirements for P2MP MPLS OAM differ from or are more extensive than those expressed in [RFC4377], additional text is supplied.

In general, it should be noted that P2MP LSPs introduce a scalability issue with respect to OAM that is not present in point-to-point MPLS. That is, an individual P2MP LSP will have more than one egress and the path to those egresses will very probably not be linear (for example, it may have a tree structure). Since the number of egresses for a single P2MP LSP is unknown and not bounded by any small number, it follows that all mechanisms defined for OAM support MUST scale well with the number of egresses and the complexity of the path of the LSP. Mechanisms that are able to deal with individual egresses will scale no worse than similar mechanisms for point-to-point LSPs, but it is desirable to develop mechanisms that are able to leverage the fact that multiple egresses are associated with a single LSP, and so achieve better scaling.

4.1. Detection of Label Switch Path Defects

The ability to detect defects in a P2MP LSP SHOULD not require manual, hop-by-hop troubleshooting of each LSR used to switch traffic for that LSP, and SHOULD rely on proactive OAM procedures (such as continuous path connectivity and Service Level Agreement (SLA) measurement mechanisms). Any solutions SHOULD either extend or work in close conjunction with existing solutions developed for point-to-point MPLS, such as those specified in [RFC4379] where this requirement is not contradicted by the other requirements in this section. This will leverage existing software and hardware deployments.

Note that P2MP LSPs may introduce additional scaling concerns for LSP probing by tools such as [RFC4379]. As the number of leaves of a P2MP LSP increases it potentially becomes more expensive to inspect the LSP to detect defects. Any tool developed for this purpose MUST be cognitive of this issue and MUST include techniques to reduce the scaling impact of an increase in the number of leaves. Nevertheless, it should also be noted that the introduction of additional leaves may mean that the use of techniques such as [RFC4379] are less appropriate for defect detection with P2MP LSPs, while the technique may still remain useful for defect diagnosis as described in the next section.

Due to the above scaling concerns, LSRs or other network resources MUST NOT be overwhelmed by the operation of normal proactive OAM procedures, and measures taken to protect LSRs and network resources against being overwhelmed MUST NOT degrade the operational value or responsiveness of proactive OAM procedures. Note that reactive OAM may violate these limits (i.e., cause visible traffic degradation) if it is necessary or useful to try to fix whatever has gone wrong.

By "overwhelmed" we mean that it MUST NOT be possible for an LSR to be so busy handling proactive OAM that it is unable to continue to process control or data plane traffic at its advertised rate. Similarly, a network resource (such as a data link) MUST NOT be carrying so much proactive OAM traffic that it is unable to carry the advertised data rate. At the same time, it is important to configure proactive OAM, if it is in use, not to raise alarms caused by the failure to receive an OAM message if the component responsible for processing the messages is unable to process because other components are consuming too many system resources -- such alarms might turn out to be false.

In practice, of course, the requirements in the previous paragraph may be met by careful specification of the anticipated data throughput of LSRs or data links. However, it should be recalled that proactive OAM procedures may be scaled linearly with the number of LSPs, and the number of LSPs is not necessarily a function of the available bandwidth in an LSR or on a data link.

4.2. Diagnosis of a Broken Label Switch Path

The ability to diagnose a broken P2MP LSP and to isolate the failed component (i.e., link or node) in the path is REQUIRED. These functions include a path connectivity test that can test all branches and leaves of a P2MP LSP for reachability, as well as a path tracing function. Note that this requirement is distinct from the requirement to detect errors or failures described in the previous section. In practice, Detection and Diagnosis/Isolation MAY be performed by separate or the same mechanisms according to the way in which the other requirements are met.

It MUST be possible for the operator (or an automated process) to stipulate a timeout after which the failure to see a response shall be flagged as an error.

Any mechanism developed to perform these functions is subject to the scalability concerns expressed in section 4.

4.3. Path Characterization

The path characterization function [RFC4377] is the ability to reveal details of LSR forwarding operations for P2MP LSPs. These details can then be compared later during subsequent testing relevant to OAM functionality. Therefore, LSRs supporting P2MP LSPs MUST provide mechanisms that allow operators to interrogate and characterize P2MP paths.

Since P2MP paths are more complex than the paths of point-to-point LSPs, the scaling concerns expressed in section 4 apply.

Note that path characterization SHOULD lead to the operator being able to determine the full tree for a P2MP LSP. That is, it is not sufficient to know the list of LSRs in the tree, but it is important to know their relative order and where the LSP branches.

Since, in some cases, the control plane state and data paths may branch at different points from the control plane and data plane topologies (for example, Figure 1), it is not sufficient to present the order of LSRs, but it is important that the branching points on that tree are clearly identified.




Figure 1. An example P2MP tree where the data path and control plane state branch at C, but the topology branches at D.

A diagnostic tool that meets the path characterization requirements SHOULD collect information that is easy to process to determine the P2MP tree for a P2MP LSP, rather than provide information that must be post-processed with some complexity.

4.4. Service Level Agreement Measurement

Mechanisms are required to measure the diverse aspects of Service Level Agreements for services that utilize P2MP LSPs. The aspects are listed in [RFC4377].

Service Level Agreements are often measured in terms of the quality and rate of data delivery. In the context of P2MP MPLS, data is delivered to multiple egress nodes. The mechanisms MUST, therefore, be capable of measuring the aspects of Service Level Agreements as they apply to each of the egress points to a P2MP LSP. At the same time, in order to diagnose issues with meeting Service Level Agreements, mechanisms SHOULD be provided to measure the aspects of the agreements at key points within the network such as at branch nodes on the P2MP tree.

4.5. Frequency of OAM Execution

As stipulated in [RFC4377], the operator MUST have the flexibility to configure OAM parameters to meet their specific operational requirements. This requirement is potentially more important in P2MP deployments where the effects of the execution of OAM functions can be potentially much greater than in a non-P2MP configuration. For example, a mechanism that causes each egress of a P2MP LSP to respond could result in a large burst of responses to a single OAM request.

Therefore, solutions produced SHOULD NOT impose any fixed limitations on the frequency of the execution of any OAM functions.

4.6. Alarm Suppression, Aggregation, and Layer Coordination

As described in [RFC4377], network elements MUST provide alarm suppression and aggregation mechanisms to prevent the generation of superfluous alarms within or across network layers. The same time constraint issues identified in [RFC4377] also apply to P2MP LSPs.

A P2MP LSP also brings the possibility of a single fault causing a larger number of alarms than for a point-to-point LSP. This can happen because there are a larger number of downstream LSRs (for example, a larger number of egresses). The resultant multiplier in the number of alarms could cause swamping of the alarm management systems to which the alarms are reported, and serves as a multiplier to the number of potentially duplicate alarms raised by the network.

Alarm aggregation or limitation techniques MUST be applied within any solution, or be available within an implementation, so that this scaling issue can be reduced. Note that this requirement introduces a second dimension to the concept of alarm aggregation. Where previously it applied to the correlation and suppression of alarms generated by different network layers, it now also applies to similar techniques applied to alarms generated by multiple downstream LSRs.

4.7. Support for OAM Interworking for Fault Notification

[RFC4377] specifies that an LSR supporting the interworking of one or more networking technologies over MPLS MUST be able to translate an MPLS defect into the native technology's error condition. This also applies to any LSR supporting P2MP LSPs. However, careful attention to the requirements for alarm suppression stipulated therein and in section 4.6 SHOULD be observed.

Note that the time constraints for fault notification and alarm propagation affect the solutions that might be applied to the scalability problem inherent in certain OAM techniques applied to

P2MP LSPs. For example, a solution to the issue of a large number of egresses all responding to some form of probe request at the same time might be to make the probes less frequent -- but this might affect the ability to detect and/or report faults.

Where fault notification to the egress is required, there is the possibility that a single fault will give rise to multiple notifications, one to each egress node of the P2MP that is downstream of the fault. Any mechanisms MUST manage this scaling issue while still continuing to deliver fault notifications in a timely manner.

Where fault notification to the ingress is required, the mechanisms MUST ensure that the notification identifies the egress nodes of the P2MP LSP that are impacted (that is, those downstream of the fault) and does not falsely imply that all egress nodes are impacted.

4.8. Error Detection and Recovery

Recovery from a fault by a network element can be facilitated by MPLS OAM procedures. As described in [RFC4377], these procedures will detect a broad range of defects, and SHOULD be operable where MPLS P2MP LSPs span multiple routing areas or multiple Service Provider domains.

The same requirements as those expressed in [RFC4377] with respect to automatic repair and operator intervention ahead of customer detection of faults apply to P2MP LSPs.

It should be observed that faults in P2MP LSPs MAY be recovered through techniques described in [P2MP-RSVP].

4.9. Standard Management Interfaces

The widespread deployment of MPLS requires common information modeling of management and control of OAM functionality. This is reflected in the integration of standard MPLS-related MIBs [RFC3812], [RFC3813], [RFC3814], [RFC3815] for fault, statistics, and configuration management. These standard interfaces provide operators with common programmatic interface access to operations and management functions and their status.

The standard MPLS-related MIB modules [RFC3812], [RFC3813], [RFC3814], and [RFC3815] SHOULD be extended wherever possible, to support P2MP LSPs, the associated OAM functions on these LSPs, and the applications that utilize P2MP LSPs. Extending them will facilitate the reuse of existing management software both in LSRs and in management systems. In cases where the existing MIB modules cannot be extended, then new MIB modules MUST be created.

4.10. Detection of Denial of Service Attacks

The ability to detect denial of service (DoS) attacks against the data or control planes that signal P2MP LSPs MUST be part of any security management related to MPLS OAM tools or techniques.

4.11. Per-LSP Accounting Requirements

In an MPLS network where P2MP LSPs are in use, Service Providers can measure traffic from an LSR to the egress of the network using some MPLS-related MIB modules (see section 4.9), for example. Other interfaces MAY exist as well and enable the creation of traffic matrices so that it is possible to know how much traffic is traveling from where to where within the network.

Analysis of traffic flows to produce a traffic matrix is more complicated where P2MP LSPs are deployed because there is no simple pairing relationship between an ingress and a single egress. Fundamental to understanding traffic flows within a network that supports P2MP LSPs will be the knowledge of where the traffic is branched for each LSP within the network, that is, where within the network the branch nodes for the LSPs are located and what their relationship is to links and other LSRs. Traffic flow and accounting tools MUST take this fact into account.

5. Security Considerations

This document introduces no new security issues compared with [RFC4377]. It is worth highlighting, however, that any tool designed to satisfy the requirements described in this document MUST include provisions to prevent its unauthorized use. Likewise, these tools MUST provide a means by which an operator can prevent denial of service attacks if those tools are used in such an attack. LSP mis-merging is described in [RFC4377] where it is pointed out that it has security implications beyond simply being a network defect. It needs to be stressed that it is in the nature of P2MP traffic flows that any erroneous delivery (such as caused by LSP mis-merging) is likely to have more far-reaching consequences since the traffic will be mis-delivered to multiple receivers.

As with the OAM functions described in [RFC4377], the performance of diagnostic functions and path characterization may involve the extraction of a significant amount of information about network construction. The network operator MAY consider this information private and wish to take steps to secure it, but further, the volume of this information may be considered as a threat to the integrity of

the network if it is extracted in bulk. This issue may be greater in P2MP MPLS because of the potential for a large number of receivers on a single LSP and the consequent extensive path of the LSP.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.

6.2. Informative References

- [MCAST-LDP] Minei, I., Ed., Kompella, K., Wijnands, I., Ed., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", Work in Progress, June 2006.
- [P2MP-LSP-PING] Yasukawa, S., Farrel, A., Ali, Z., and B. Fenner, "Detecting Data Plane Failures in Point-to-Multipoint MPLS Traffic Engineering - Extensions to LSP Ping", Work in Progress, April 2006.
- [P2MP-RSVP] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to RSVP-TE for Point to Multipoint TE LSPs", Work in Progress, July 2006.
- [RFC3812] Srinivasan, C., Viswanathan, A. and T. Nadeau, "MPLS Traffic Engineering Management Information Base Using SMIPv2", RFC3812, June 2004.
- [RFC3813] Srinivasan, C., Viswanathan, A. and T. Nadeau, "MPLS Label Switch Router Management Information Base Using SMIPv2", RFC3813, June 2004.
- [RFC3814] Nadeau, T., Srinivasan, C., and A. Viswanathan, "Multiprotocol Label Switching (MPLS) FEC-To-NHLFE (FTN) Management Information Base", RFC3814, June 2004.

- [RFC3815] Cucchiara, J., Sjostrand, H., and Luciani, J., "Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)", RFC 3815, June 2004.
- [RFC4378] Allan, D. and T. Nadeau, "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC4461] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.

7. Acknowledgements

The authors wish to acknowledge and thank the following individuals for their valuable comments on this document: Rahul Aggarwal, Neil Harrison, Ben Niven-Jenkins, and Dimitri Papadimitriou.

Authors' Addresses

Seisho Yasukawa
NTT Corporation
(R&D Strategy Department)
3-1, Otemachi 2-Chome Chiyodaku,
Tokyo 100-8116 Japan

Phone: +81 3 5205 5341
EMail: s.yasukawa@hco.ntt.co.jp

Adrian Farrel
Old Dog Consulting

Phone: +44 (0) 1978 860944
EMail: adrian@olddog.co.uk

Daniel King
Aria Networks Ltd.

Phone: +44 (0)1249 665923
EMail: daniel.king@aria-networks.com

Thomas D. Nadeau
Cisco Systems, Inc.
1414 Massachusetts Ave.
Boxborough, MA 01719

EMail: tnadeau@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).