Network Working Group   Seisho Yasukawa
Internet Draft   NTT Corp
Expires: June 2006

                                               Adrian Farrel
                                               Old Dog Consulting

                                               Daniel King
                                               DNNI Ltd.

                                               Thomas D. Nadeau
                                               Cisco Systems, Inc.

                                               December 2005

        OAM Requirements for Point-to-Multipoint MPLS Networks

             draft-ietf-mpls-p2mp-oam-reqs-00.txt

Status of this Memo

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 dataplane defects is critical.

   For operators deploying services based on P2MP MPLS LSPs the
   detection and specification of how to handle those defects is
   important because such defects may not only affect the fundamental

of an MPLS network, but also because they MAY impact service level
   specification commitments for customers of their network.

   This document describes requirements for user and 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

   This document describes requirements for user and data plane
   operations and management (OAM) for point-to-multipoint (P2MP)
   Multi-Protocol Label Switching (MPLS). These requirements have been
   gathered from network operators who have extensive experience
   deploying MPLS networks and from operators who are considering

deploying P2MP MPLS networks. This draft 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 [P2MP-SIG-REQ] and [P2MP-RSVP], as
well as multicast LDP LSPs [P2MP-LDP] and [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 are
described in [MPLS-OAM] and are not repeated in this document.

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].

2.2 Terminology

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

[P2MP-SIG-REQ] 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 list of acronyms is a repeat of common acronyms defined
in many other documents, and is provided here for convenience.

CE: Customer Edge
DoS: Denial of service
ECMP: Equal Cost Multipath
LDP: Label Distribution Protocol
LSP: Label Switch Path
LSR: Label Switch Router
OAM: Operations and Management
OA&M: Operations, Administration and Maintenance.

```
        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, [LSP-PING], [RFC3812], [RFC3813], [RFC3814],
   and [RFC3815]) developed specific solutions to individual issues or
   problems. Coordination of the full OAM requirements for MPLS was
   achieved by [MPLS-OAM] 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 closely
   similar to those described for point-to-point MPLS in [MPLS-OAM].
   The subsections below do not repeat material from [MPLS-OAM], 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 [MPLS-OAM], additional text is
   supplied.

In general, it should be noted that P2MP LSPs introduce a scalability
issue 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 broken Label Switch Path
(LSP) SHOULD not require manual hop-by-hop troubleshooting of
each LSR used to switch traffic for that P2MP LSP. Any
solutions should either extend or work in close conjunction
with existing solutions developed for point-to-point MPLS, such as
those specified in [LSP-PING]. This will leverage existing software
and hardware deployments.

   Note that P2MP LSPs may introduce additional scaling concerns for
   LSP probing by tools such as [LSP-PING]. As the number of leaves
   of the P2MP LSP increases so it becomes potentially 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 introduciton of
   additional leaves may mean that the use of techniques such as
   [LSP-PING] are less appropriate for defect detection of P2MP LSPs,
   while the technique may still remain useful for defect diagnosis as
   described in the next section.

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. 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 are subject to the scalability concerns expressed in section 4.

4.3 Path Characterization

The path characterization function [MPLS-OAM] 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 interogate 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.

```
                E
               /
   A---B---C===D
               \
                F
```
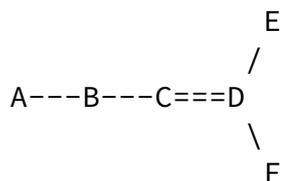
Figure 1. An example P2MP tree where the data path and control plane state branch at C, but the topology branches at D.

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 utlize P2MP LSPs. The aspects are
listed in [MPLS-OAM].

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 [MPLS-OAM], 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 for 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 [MPLS-OAM], network elements MUST provide alarm
suppression and aggregation to prevent the generation of superfluous
alarms within or across network layers. The same time constraint
issues identified in [MPLS-OAM] also exist for P2MP LSPs.

A P2MP LSP also brings the possiblity 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

   [MPLS-OAM] 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 impact upon 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 impact on the ability to detect and/or report faults.

   Where fault notification to the egress is required, there is the
   possiblity 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 [MPLS-OAM], these procedures
   will detect a broad range of defects, and SHOULD be operable where
   MPLS P2MP LSPs span multiple routing areas, or multiple Service

Yasukawa et al.   [Page 7]

Provider domains.

The same requirements as those expressed in [MPLS-OAM] 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 wide-spread deployment of MPLS requires common information modeling of management and control of OAM functionality. This is reflected in the the integration of standard MPLS-related MIBs [RFC3813], [RFC3812], [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. These

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 utlize P2MP LSPs. Extending them will facillitate 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 which 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 matricies 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 becasue 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. The Traffic flow and accounting tools MUST take this fact into account.

5. Security Considerations

   This document introduces no new security issues compared with
   [MPLS-OAM]. 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 [MPLS-OAM] 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 conseuqences since
   the traffic will be mis-delivered to multiple receivers.

   As with previous OAM function described in [MPLS-OAM], 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 extacted in bulk. This issue may be
   greater in P2MP MPLS becuase of the potential for a large number of
   receivers on a single LSP and the consequent extensive path of the
   LSP.

6. IANA Considerations

   This document creates no new requirements on IANA namespaces.

7. References

7.1 Normative References

   [MPLS-OAM]   T. Nadeau, Allan D., et al. Allan D.,
                      "OAM Requirements for MPLS Networks",
                      draft-ietf-mpls-oam-requirements-05.txt,
                      December 2004

   [RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
                      Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2 Informative References

    [LSP-PING]   Kompella, K., and Swallow, G., (Editors), "Detecting
                          MPLS Data Plane Failures", draft-ietf-mpls-lsp-ping,
                          work in progress.

    [MCAST-LDP]   Wijnands, IJ., et al., " Multicast Extensions for
                          LDP", draft-wijnands-mpls-ldp-mcast-ext, work in
                          progress.

    [P2MP-LDP]   Minei, I., et al., "Label Distribution Protocol
                          Extensions for Point-to-Multipoint Label Switched
                          Paths", draft-minei-mpls-ldp-p2mp, work in progress.

    [P2MP-LSP-PING]   Yasukawa, S., Farrel, A., Ali, Z., and Fenner, B.,
                          "Detecting Data Plane Failures in
                          Point-to-Multipoint MPLS Traffic Engineering -
                          Extensions to LSP Ping",
                          draft-yasukawa-mpls-p2mp-lsp-ping, work in progress.

    [P2MP-RSVP]   Aggarwal, R., Papadimitriou, D., and Yasukawa, S.,
                          "Extensions to RSVP-TE for Point to Multipoint TE
                          LSPs", draft-ietf-mpls-rsvp-te-p2mp, work in
                          progress.

    [P2MP-SIG-REQ]   Yasukawa, S. (Editor), "Signaling Requirements for
                          Point to Multipoint Traffic Engineered MPLS LSPs",
                          draft-ietf-mpls-p2mp-sig-requirement, work in
                          progress.

    [RFC3812]   Srinivasan, C., Viswanathan, A. and T.
                          Nadeau, "MPLS Traffic Engineering Management
                          Information Base Using SMIv2", RFC3812, June 2004.

    [RFC3813]   Srinivasan, C., Viswanathan, A. and T.
                          Nadeau, "MPLS Label Switch Router Management
                          Information Base Using SMIv2", 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.

8. Acknowledgment

    The authors wish to acknowledge and thank the following
    individuals for their valuable comments to this document:
    Dimitri Papadimitriou and Rahul Aggarwal.

9. Authors' Addresses

    Adrian Farrel
    Old Dog Consulting
    Phone: +44 (0) 1978 860944
    Email: adrian@olddog.co.uk

    Daniel King
    Darwinian Neural Network Industries Ltd.
    Phone: +44 (0)1249 665923
    Email: dan@dnni.com

    Thomas D. Nadeau
    Cisco Systems, Inc.
    1414 Massachusetts Ave.
    Boxborough, MA 01719
    Email: tnadeau@cisco.com

    Seisho Yasukawa
    NTT Corporation
    9-11, Midori-Cho 3-Chome
    Musashino-Shi, Tokyo 180-8585,
    Japan
    Phone: +81 422 59 4769
    Email: yasukawa.seisho@lab.ntt.co.jp

10. Intellectual Property Statement

    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.

11. Full Copyright Statement