CCAMP Working Group                                         Zafar Ali
Internet Draft                                       Roberto Cassata
Intended status: Standards Track                 Cisco Systems, Inc.
                                                      Marco Anisetti
                                                    Valerio Bellandi
                                                     Ernesto Damiani
                                                     Francesco Diana
                                                    Umberto Raimondi
                                                 University of Milan
                             T. Otani(KDDI R&D Laboratories, Inc.)

Expires: August 2008                              February 25, 2008

# Ping and Traceroute with Evidence Collection in Photonic Networks
**draft-ali-ccamp-gmpls-lsp-ping-traceroute-01.txt**

## Status of this Memo

## Copyright Notice

## Abstract

[RFC4379] describes procedures for ping and tracerouting for LSPs with PSC (packet switch capable) transit switching capability. An important implication of using transparent (non-PSC) nodes in GMPLS network is that LSP Ping solution described in [RFC4379] are not applicable to LSP with non-PSC switching capability. Another important difference between PSC and non-PSC switching technologies is the data and control plan separation in the latter case. An implication of the separation of data and control planes in GMPLS networks is that LSP traceroute procedures described in [RFC4379] are not directly applicable to GMPLS networks with separation of data and control planes.

The scope of this draft is cases where data plane does not provide the OAM functions addressed by this draft. This document is assumed that OAM mechanisms provided by the underlying data plan technology MUST be used, whenever possible. E.g., G.709 addresses the problem of trace routing in DWDM network. However, G.709 OAM mechanisms are only applicable to OEO (Optical-Electrical-Optical) capable node. This document fills in such gaps; in particular it addresses GMPLS OAM functionality in optical networks with wavelength routers, ROADMs nodes, etc. with no OEO conversion capability. For this purpose, the draft relies on control plan mechanism to provide required OAM functions. Specifically the proposed solutions are based on Link Management Protocol (LMP) [RFC4204] and RSVP-TE [RFC3209], [RFC3473] and do not require any extension to the data plan.

Conventions used in this document

In examples, "C:" and "S:" indicate lines sent by the client and server respectively.

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.

Table of Contents

## 1. Introduction

When a GMPLS LSP fails to deliver user traffic, the failure cannot always be detected by the GMPLS control plane. There is a need to provide a tool that would enable users to detect such traffic "black holes" or misrouting within a reasonable period of time, and a mechanism to isolate faults [GMPLS-OAM-REQ]. Similarly, ability to traceroute a GMPLS LSPs in networks where data and control planes are separated is a requirement [GMPLS-OAM-REQ]. This draft provides solution to these requirements.

The scope of this draft is cases where data plane does not provide the OAM functions addressed by this draft. This document is assumed that OAM mechanisms provided by the underlying data plan technology MUST be used, whenever possible. E.g., G.709 addresses the problem of trace routing in DWDM network. However, G.709 OAM mechanisms are only applicable to OEO (Optical-Electrical-Optical) capable node. This document fills in such gaps; in particular it addresses GMPLS OAM functionality in optical networks with wavelength routers, ROADMs nodes, etc. with no OEO conversion capability. For this purpose, the draft relies on control plan mechanism to provide required OAM functions.

[RFC4379] describes control plan procedures for LSP Ping for LSPs with PSC (packet switch capable) endpoint and transit switching capability devices. LSP Ping solutions described in [RFC4379], however, are not applicable to LSPs crossing or terminating non-PSC switching capable devices. This is because the solution described in RFC4379 requires all transit and end point nodes along the LSP path

to be able to intercept the MPLS OAM (Operation and Maintenance) packets and identify the Target FEC Stack being tested. Such capability is not available at nodes that are non-PSC-capable. Moreover, LSP ping mechanisms described in [RFC4379] can be inadequate even when the end points of the GMPLS LSP are PSC-capable. This is because the GMPLS LSP appears as a single hop for procedures described in [RFC4379]. In such cases, mechanisms in [RFC4379] are able to detect data plan failure in the GMPLS LSP but are still not able to isolate failures in underlying switching layers.

The Link Management Protocol (LMP) [RFC4204] fault isolation mechanism can be used to detect and isolate failures along a GMPLS LSP, but it requires the GMPLS LSP to be carrying traffic. Inability to use LSP fault isolation is a considerable limitation for operators wanting to check the health of an LSP before carrying traffic over it. This draft addresses this limitation by extending the LMP link verification procedure to check connectivity of a GMPLS LSP and extending the RSVP-TE to detect the faulty point.

For successful fault detection on a light-path, the fault isolation mechanism must be aware of all physical evidence (consisting of optical measurements such as signal power, OSNR, OCM (Optical Channel Monitor), etc.) that have effect on the light-path. The proposed technique is also suitable for optical networks that suffer of physical dysfunction due the non-ideal optical transmission medium and/or to critical situations (e.g., a fiber cut). In this scenario even if every node along the path is connected, the reachability of the end node with an acceptable signal quality is not guaranteed.

Such evidence can consist of real optical measurements or estimates computed via a prediction model. The former may require mutually exclusive access to hardware to avoid interference; therefore, evidence can also be classified as blocking or non-blocking. This draft address both type of evidence collections. Furthermore, in this draft evidence collection is performed during the phase of trace routing.

## 2. Tracerouting with Evidence Collection

Traceroute is often used for network troubleshooting. Specifically, it is used identify the LSP taken to reach a particular destination viewing the all transit nodes on the network; for that reasons it is used also to detect faulty point inside a route.

The LSP traceroute procedures described in [RFC4379] are not directly applicable to GMPLS networks with separation of data and control planes. To overcome this issue tracerouting using RSVP RRO object

[RFC4561] can be implemented. This strategy is only a control plain view. However, maintain a coherence with data channel in the sense of traversed nodes it detects a faulty point in the control channel that is largely different than finding the faulty point in the data channel.

This draft proposes a technique to address the deficiency of the use of RRO for tracerouting a GMPLS LSP. The proposed is control plan based but is able perform a traceroute with fault isolation coherent with data channel. The proposed method is able to perform tracerouting with evidence collection. It is based on the idea that for successful fault detection on an optical path, the fault isolation mechanism must be aware of all physical evidence (consisting of optical measurements such as signal power, OSNR, Optical Channel Monitor, etc.) that have effect on the light-path. Therefore measuring or estimating some physical evidences along an optical path address the actual control channel deficiency in finding the data channel coherence traceroute.

### 2.1. Optical Path Quality Evaluation

The quality of an optical path is done by collecting the physical evidences along an LSP and evaluating them (e.g. for faulty point detection). As already mention this feature integrates the tracerouting in such a way that the control channel becoming aware and coherent with data channel. The holistic analysis proposed produce also a quality of path awareness.

In this draft we extend the LSP_ATTRIBUTES to perform the evidence collection hop by hop.

Other important concept defined by this evidences collection process is that certain evidences (blocking evidence) require a mutually exclusive access. Therefore the entire LSP needs to be locked until the evidence collection process is performed. This implies that if other evidence collection process tries to retrieve evidences on the same node-resource already under Administrative Evidences Locking status, it MUST be aborted. The draft uses RSVP Admin status object to define LSP Administrative Evidences Locking status and to make sure that all nodes are ready to collect the blocking evidence.

In the following we first define Optical Evidence classification, and extension to LSP ATTRIBUTE and RSVP Admin status objects needed to perform above mentioned functionalities. The later sections details

signaling procedures with examples on how these objects are used for tracerouting with evidence collection.

## 2.2. Optical Evidence Classification and LSP Locking

Physical evidences (consisting of optical measurements such as signal power, OSNR, Optical Channel Monitor, etc.) that have effect on the light-path are classified as:

o  Blocking evidence. In general blocking evidence is a physical measurement that may require a mutually exclusive access to hardware resources while performing the measurement.

o  Non blocking evidence. Every physical values that can be probed in parallel with different RSVP-TE.

Every optical Node can be in three states related to a certain reserved resource: unlock, lock-required or lock. In fact blocking evidence MUST generate a lock on each reserved resource required for evidence reading. In general this is due to the hardware limitation of optical nodes.

In case of blocking evidence the LSP status needs to be set to "Locked". To perform this status changing we use the Admin object [RFC3471] with B bit (Blocked bit) and C bit (block Confirm) extension. In our LSP locking strategy also the R bit (Reflect bit) MUST be set since the egress node MUST return the Admin object in the Resv Message for locking confirmation or unlocking. Since we need to block an entire LSP, one node unable to measure the required blocking evidence MUST generate a lock failure (unset the C bit in the Path Admin Object). Therefore the evidence locking is considered mandatory.

The general locking procedure is defined as follow:

o  Every transit node that receives the Admin status object in the Path message with B, C and R bit set needs to check if the actual status is unlock.

o  In the case of unlock status, the node switches to lock-required state related to the required blocking evidence.

o  In the case of lock or lock-required statuses, the node forwards the Admin object message without the C bit set. This implies a lock failure.

- The Resv message performs the locking for the entire LSP in case of C and B bit set and unlocking in case of unset C bit.

- Every transit node that receives the Resv message with B and C bit set changes its status to lock.

This strategy prevents race conditions.

### 2.3. Optical Evidence Collection

Path quality evaluation is based on holistic analysis of the evidence collected inside an LSP. To determinate which evidence needs to be collected we adopt a LSP Attribute TLV sub-object.

The evidence collection is performed as follows:

- Source node sends a Path message with LSP Attribute object aimed to inform the transit nodes about the imminent evidence collection This downstream Path message also contains TLV sub-object with required evidence.

- Every transit node, when receives the message with LSP Attribute object, assembles the collected evidence (specified in TLV) inside a sub-TLV. The way an optical node gets knowledge of the evidence using information locally available at the node (e.g. via discovery of internal amplifiers, photodiode etc.) is out of the scope of this document.

- Evidence collection will be executed by the returning Resv message that collects hop-by-hop evidence objects upstream by inserting the sub-TLV inside the attached TLV. After successful forwarding of Resv message the status of transit nodes MUST be switched to unlock for preventing deadlock.

In case of blocking evidence the LSP lock MUST be performed before evidence collection.

In case of non-blocking evidence the unavailability of certain evidence in an intermediate node MUST NOT cause the request of failure (PathErr message) since the holistic evidence evaluation SHOULD be able to deal with missing non-blocking evidence.

When one transit node not in locking state receives a request for blocking evidence, an evidence collection failure (PathErr) MUST be triggered.

### 2.4. Evidence Collection Request TLV

NOTE: INFORMATION IN THIS SECTION NEED SOME CAREFUL REVISION AGAINST EXPECTED USAGE IN [RFC4420].

The proposed encoding scheme for optical evidence measurements defines a TLV associated to a particular evidence type. A TLV sub-object is encoded in an LSP_REQUIRED_ATTRIBUTES Object [RFC4420]. The TLV sub-object encoding is:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |           Length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     E-Type    |                    Reserve                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type: Collected evidence type(TBA). Can be blocking or non blocking type.

Length: length of the TLV object in bytes without the 4 byte header.

E-type (Evidence Type, 8 bits): Evidence identifier, for instance: 0 as Signal power, 1 as OSNR, 2 as Pilot Tone (as blocking evidence).

This TLV defines which type of evidence needs to be collected and specifies the evidence (signal power, OSNR, Pilot Tone, alarm etc.) in the Path message.

### 2.5. Evidence recording TLV

NOTE: INFORMATION IN THIS SECTION NEED SOME CAREFUL REVISION AGAINST EXPECTED USAGE IN [RFC4420].

For provisioned LSP, a set of evidence has to be collected through the Resv message to allow the optical quality evaluation at the ingress node. Each item of optical evidence is collected separately. Every transit node, in the Path message, finds the Evidence collection requested TLV and stores in the Evidence recording TLV (encoded in an LSP_ATTRIBUTES Object) its own measured or estimated value. Furthermore it sets the Measure Method inside the this TLV according to the kind of measured media (single lambda measurement or aggregate measurement).
This evidence collection improves the feasibility evaluation where network elements support at least only a subset of evidence.

The following TLV encode the evidence's values of the LSP associated to the evidence type defined in the Evidence Collection Request TLV.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    Type                        |         Length               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  IPv4/IPv6 Address/unnumbered                                 |
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  Evidence Value                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |Measure Method |                 Reserved                      |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type: Evidence type(TBA).

Length: length of the TLV value in bytes.

IPv4/IPv6 Address: The address of the Node that measures the evidence.

Evidence Value: Estimated or measured evidence value. For instance the Signal Optical Power as 32-bit IEEE floating point number.

Measure(ment) method: Aggregate measurement (0) or single lambda measurement (1).

## 2.6. Administrative Status Object extension

We propose and extension to Administrative status object by adding two bits for locking purpose.

Therefore the format of the extended Admin_Status Object is:

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |            Length              | Class-Num(196)|   C-Type (1)  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |R|         Reserved                                   |C|B|T|A|D|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Reflect (R): 1 bit
When set, indicates that the edge node SHOULD reflect the object/TLV back in the appropriate message.  This bit MUST NOT be set in state change request, i.e., Notify, messages.

Reserved: 25 bits. This field is reserved.  It MUST be set to zero on transmission and MUST be ignored on receipt.  These bits SHOULD be passed through unmodified by transit nodes.

Testing (T): 1 bit. When set, indicates that the local actions related to the "testing" mode should be taken.

Administratively down (A): 1 bit. When set, indicates that the local actions related to the "administratively down" state should be taken.

Deletion in progress (D): 1 bit. When set, indicates that that the local actions related to LSP teardown should be taken.  Edge nodes may use this flag to control connection teardown.

Blocking node (B): 1 bit. When set, indicates that locking procedure is ongoing.

Confirm blocking (C): 1 bit. When set, indicates that an the locking procedure is successfully ongoing.

## 2.7. Signaling Procedure for tracerouting with evidence collection

In this section we describe signaling procedures for tracerouting with evidence collection using examples. Consider a GMPLS LSP that has OXC1 as Ingress Node, OXC4 as Egress node with OXC2 and OXC3 in transit, as shown below.

```
+------+       +------+       +------+       +------+
|      | ----- |      | ===== |      | ----- |      |
| OXC1 | ===== | OXC2 | ----- | OXC3 | ----- | OXC4 |
|      | ----- |      | ----- |      | ===== |      |
+------+       +------+       +------+       +------+
```

In the following we consider three scenarios of evidence collection and describe signaling procedures associated with the evidence collection and how above mentioned extensions to LSP Attribute and admin status objects are used for this purpose.

### 2.7.1. Tracerouting with non-blocking evidence collection

The quality evaluation of an optical path is done after LSP provisioning and in case of non-blocking evidences is implemented by the following procedure:

o  OXC1 node sends a Path message with Evidence Collection Request TLV aimed to inform the transit nodes about the imminent evidence collection and about the type of evidence that needs to be collected (e.g., Signal power).

o  The transit nodes that do not support LSP_REQUIRED ATTRIBUTE object or do not support evidence request TLV will be addressed in a later version of the document.

o  Every transit node (OXC2,OXC3), when receives the Path message with Evidence Collection Request TLV, starting the internal evidence reading procedure and waits for the correspondent Resv message to forward the related Evidence recording TLV in the upstreaming flow to the ingress node OXC1. If for some reason the evidence is not available, since it is non blocking evidence, the node simply do not include the evidence measure in its own Evidence recording TLV. The holistic analysis can be performed also with a subset of the non blocking evidences.

o  Egress node OXC4 sends Resv message with Evidence Collection Request TLV containing optical evidence TLV upstream to the ingress node OXC1 and puts its own evidence value in this Evidence recording TLV.

o  Every transit node (OXC3,OXC2) inserts its own Evidences recording TLV inside Resv message in such way that ingress node collects all required evidences hop by hop using the upsteaming flow.

o  OXC1 node when receives the Resv message extract the Evidences recording TLV to perform holistic path quality analysis.

Summarizing the Evidence collection will be executed by the returning Resv message that collects hop-by-hop evidence objects upstream.

### 2.7.2. Tracerouting with blocking evidence collection and all nodes ready for evidence collection

In this scenario the locking strategy needs to be performed first to be sure that no one node in the LSP is already locked in another blocking evidences collection. Summarizing we need to be sure that all nodes along the path are ready to collect the evidence. This phase uses Admin status object in the path and Resv message to

The locking procedure is defined as follow:

o  OXC1 switches to lock-required state and sends a Path message with Admin status object with B, C and R bit set. B bit is used for

locking requirement. C bit is used for locking confirmation if set and for unlock if unset.

- Every transit node (OXC2, OXC3) that receives the Admin status object in the Path message with B, C and R bit set switches to lock-required state related to the required blocking evidence.

- Egress node OXC4 switches to lock state forward the Admin status object in the Resv message resetting the R bit.

- Every transit node (OXC3,OXC2) that receives the Resv message with B and C bit set changes its status to lock.

- Ingress node OXC1 when receive the Resv message with Admin status object with B and C bit set switches to lock states.

At the end of this procedure the entire LSP is in lock state and is ready for blocking evidence collection.

At this stage the Evidence collection can be performed as described in the Section 2.7.1 except that every transit nodes need to change its own status to unlock to prevent deadlock as described in the Evidence collection Section (2.3).

The locking strategy is performed before evidence collection to maintain a better compatibility with the future available blocking evidences kind that would require further action to be taken before starting the collection.

### 2.7.3. Tracerouting with blocking evidence collection with some node(s) blocked for evidence collection.

In this scenario the locking procedure fails since some nodes (e.g OXC3 is in locking or lock-required state over other LSP)

- OXC1 switches to lock-required state and sends a Path message with Admin status object with B, C and R bit set. B bit is used for locking requirement. C bit is used for locking confirmation if set and for unlock if unset.

- OXC2 receives the Admin status object in the Path message with B, C and R bit set switches to lock-required state related to the required blocking evidence.

o  OXC3 node when receives the Admin status object since it is already lock or lock-required over other LSP with the same resources, unset the C bit. Therefore the locking procedure will fails.

o  Egress node OXC4 since receives the Admin object without C bit set switches to unlock state and forwards the received Admin status object in the Resv message resetting the R bit.

o  Other transit nodes (OXC3, OXC2) when receive the Admin object in the Resv message with B bit set but with C bit unset, switch to unlock state.

o  The ingress node OXC1 when receives Resv message with Admin object containing B bit set and C bit unset switches to unlock.

At this stage the Locking strategy is failed since the ingress node does not receive the confirmation of successful locking (C bit set).

## 3. LSP Ping for GMPLS LSPs

Tracerouting with evidence collection described in the last section is an expensive signaling operation. Most of the time service provider's requirement is to test connectivity verification, and to perform tracerouting with evidence collection when detailed diagnostic of LSP is needed.

If the end-points of the LSP are PSC capable, LSP ping procedure in [RFC4379] can be used. However, if LSP end-points are non-PSC capable, LMP procedure described in this section can be used to provide LSP ping functionality for GMPLS LSPs. For this purpose, this draft proposes an extended LMP model as shown below.

```
          +------+       +------+       +------+       +------+
          |      | ----- |      | ===== |      | ----- |      |
          | OXC1 | ===== | OXC2 | ----- | OXC3 | ----- | OXC4 |
          |      | ----- |      | ----- |      | ===== |      |
          +------+       +------+       +------+       +------+
            ^ ^              ^  ^           ^ ^             ^  ^
            | |              |  |           | |             |  |
            | +----LMP1----+  +----LMP2---+ +-----LMP3----+  |
            |                                                |
            +----------------------LMP4----------------------+
```

Figure 1 Extended LMP Model.

In this model, non-adjacent nodes may establish and maintain LMP sessions that can be used to check the status of a GMPLS LSP. Also, the nodes continue to maintain hop-by-hop LMP sessions to build traffic-engineering (TE) links for GMPLS signaling and routing, as described in [RFC4204]. For example in Figure 1, OXC1-OXC2, OXC2-OXC3, and OXC3-OXC4 LMP sessions are used to build traffic-engineering (TE) links for GMPLS signaling and routing, while the LMP session OXC1-OXC4 (LMP4) is used to monitor the health of GMPLS LSP(s) with OXC1 and OXC4 as end-points. Note that the LMP session between LSP end-point nodes is only used for OAM purposes. Existing signaling mechanisms are used to discover remote link property.

Once an LMP session between LSP end-point nodes comes up, Link connectivity verification can be used to perform LSP connectivity verification.  This is done by sending Test messages over the GMPLS LSP and TestStatus messages back over the control channel. For this purpose, LMP connectivity verification procedure as described in [RFC4204] is used. Note that in this model the verification of a GMPLS LSP is not confined to LSPs having endpoint nodes that are PSC-capable, but effectively to LSPs of endpoint nodes that reside at any of the GMPLS switching layers.

In what follows, we outline how existing LMP and MPLS OAM procedures needs to be applied to provide tracerouting functionality in scenarios outlined above. Again recall the scope of this draft is cases where data plane does not provide the OAM functions addressed by this draft.

The control channel management for LSP ingress-node-to-egress-node is the same as described in [RFC4204]. To distinguish between a LSP ingress-node-to-egress-node LMP session and a peer node-to-peer node LMP session, a new LMP_TARGET_HELLO_CONFIG object is defined (C-Type = TBD).  The format of the CONFIG object is as follows:

Class = TBD

o     C-Type = TBD, LMP_TARGET_HELLO_CONFIG

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |T|                        (Reserved)                           |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Reserved field should be sent as zero and ignored on receipt.

T: 1 bit

This bit indicates support for the LMP-LSP-Verification extensions defined in this document.

To establish an ingress-node-to-egress-node LMP session, sender node uses the control plane's IP address of the LSP destination node, while sending the LMP_TARGET_HELLO_CONFIG message. The ConfigAck and ConfigNack messages MUST be sent to the source IP address found in the IP header of the received Config message.

### 3.1. LSP Verification Procedure

Link verification procedure described in [RFC4204] has been adapted for LSP verification. Specifically, once a control channel has been established between the ingress and egress nodes of an LSP, LSP connectivity can be verified by exchanging Test messages between nodes along the GMPLS LSP's path. Since the LSP's health can be tested along the forwarding transmit path, both endpoints nodes can (independently and simultaneously) initiate the exchange of Test messages in each direction to test for the health of bidirectional LSPs.

To initiate the link verification procedure, the Ingress (Egress) node MUST send a BeginVerify message over a control channel with the IP address of the destination (source) node of the LSP. To limit the scope of LSP Verification to a particular LSP, the local Lsp_Id assigned by the local node is used. This Lsp_Id is learned by the remote node during signaling and MUST be non-zero. If this field is zero, the verification can span multiple TE LSPs between the set of Ingress/Egress nodes involved in the verification process. The rest of the details for LSP verification are the same as described for link verification in [RFC4204].

## 4. Security Considerations

Security considerations and requirements form [RFC4204] and [RFC4379] apply equally to this document. Furthermore, there are some additional security considerations that may be induced by extended

LMP model and RSVP-TE proposed by this draft. These security considerations will be added in a later version of the draft.

## 5. Acknowledgments

Authors would like to thank Alberto Tanzi, Ferdinando Malgrati, Domenico La Fauci, Enzo Luca Passerini, Gabriele Galimberti for their valuable inputs.

## 6. **IANA Considerations**

TBA.

## 7. References

### 7.1. Normative References

[RFC4204] Lang, J., et al., "Link Management Protocol (LMP)", RFC 4204, October 2003.

[RFC4379] Kompella, K., Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209 ,December 2001.

[RFC3471] Berger, L., et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., et al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling", RFC 3473, RFC 3473, January 2003.

[RFC4561] Vasseur, J.-P.,Ali, Z., Sivabalan, S., "Definition of a Record Route Object (RRO) Node-Id Sub-Object", RFC 4561, June 2006.

[RFC4420] Farrel, A., Papadimitriou, D., Vasseur, J., and A. Ayyangar, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4420, February 2006.

### 7.2. Informative References

[GMPLS-OAM-REQ] Otani, T., et al., "OAM Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Networks", draft-ietf-ccamp-gmpls-oam-requirements-00.txt.

Author's Addresses

Zafar Ali
Cisco Systems, Inc. 200
100 South Main St. #
Ann Arbor, MI 48104
USA
Email: zali@cisco.com

Marco Anisetti
University of Milan, Department of information Technology
Via Bramante 65, 26013 Crema (CR)
Italy
Email: anisetti@dti.unimi.it

Valerio Bellandi
University of Milan, Department of information Technology
Via Bramante 65, 26013 Crema (CR)
Italy
Email: bellandi@dti.unimi.it

Roberto Cassata
Cisco Systems, Inc.
Via Philips 2, 20052 Monza (MI)
Italy
Email: rcassata@cisco.com

Ernesto Damiani
University of Milan, Department of information Technology
Via Bramante 65, 26013 Crema (CR)
Italy
Email: damiani@dti.unimi.it

Francesco Diana
University of Milan, Department of information Technology
Via Bramante 65, 26013 Crema (CR)
Italy
Email: diana@dti.unimi.it

Tomohiro Otani
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Fujimino Saitama, 356-8502. Japan
Email: otani@kddilabs.jp

Umberto Raimondi

University of Milan, Department of information Technology
Via Bramante 65, 26013 Crema (CR)
Italy
Email: uraimondi@crema.unimi.it

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.

Disclaimer of Validity

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, THE IETF TRUST 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.

Copyright Statement