MPLS Working Group
Internet Draft
Intended status: Informational

A. Fulignoli (Ed.)
Ericsson

S. Boutros (Ed.)
Cisco Systems, Inc

M.Vigoureux (Ed.)
Alcatel-Lucent

Expires: January 2010 July 7, 2009

MPLS-TP BFD for Proactive CC-CV and RDI
draft-fulignoli-mpls-tp-bfd-cv-proactive-and-rdi-01.txt

Status of this Memo

Copyright Statement

Abstract

Several documents on BFD based OAM for MPLS-TP has been put forward and the dependencies between those drafts are not yet fully sorted out; this document is one of these drafts. It is published in now to make ideas, motivations and approaches available. However we expect the final BFD based solution for MPLS-TP will be a cooperation of the parties between the existing drafts and that the BFD based OAM solution for MPLS-TP will merge into an agreed set of drafts approved by the MEAD team.

This document specifies the BFD extension and behaviour to meet the requirements for MPLS-TP proactive Continuity Check and Connectivity Verification functionality and the RDI functionality as defined in [3].

Table of Contents

1. Introduction

Several documents on BFD based OAM for MPLS-TP has been put forward and the dependencies between those drafts are not yet fully sorted out; this document is one of these drafts. It is published in now to make ideas, motivations and approaches available.

However we expect the final BFD based solution for MPLS-TP will be a cooperation of the parties between the existing drafts and that the BFD based OAM solution for MPLS-TP will merge into an agreed set of drafts approved by the MEAD team.

This document specifies the BFD extension and behaviour to meet the requirements for MPLS-TP proactive Continuity Check and Connectivity Verification functionality and the RDI functionality as defined in [3].

As recommended in [4], the BFD tool needs to be extended for the CV functionality by the addition of a unique identifier in order to meet the requirements.

As described in [5], the Proactive Continuity Check (CC) and Continuity Verification (CV) function are used together to detect loss of continuity (LOC), unintended connectivity between two MEs (e.g. mismerging or misconnection) as well as unintended connectivity within the ME with an unexpected MEP. It MUST operate both in bidirectional and unidirectional p2p and unidirectional p2mp connection.

The mechanism MUST foresee the configuration of the transmit frequency.

The mechanism MUST be the same for LSP, (MS-)PW and Section as well as for LSP Tandem Connection and PW Tandem Connection.

1.1. Terminology

LME	LSP Maintenance Entity
LTCME	LSP Tandem Connection Maintenance Entity
ME	Maintenance Entity
MEP	Maintenance End Point
MIP	Maintenance Intermediate Point
PME	PW Maintenance Entity
PTCME	PW Tandem Connection Maintenance Entity
SME	Section Maintenance Entity
TLV	Type Length Value

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 [1].

2. Trail Termination Source Identifier (TTSI) TLV Object

The MPLS Generic Associated Channel specification (see[2] section 3) describes the ACH TLV structure that can be used to provide additional context information to the G-ACh packet.

In this section the TLV Objects for providing the MEP Identifier information and the ME Identifier information as required by[5] are described.

[Editor's note - Some ACH TLV objects defined in this section can be moved in future versions of [10] and referenced by future versions of this draft]

Note: in order to simply implementations (e.g. planning processing resources), especially when BFD implementation is hardware-assisted, it would be desirable to define the maximum possible length for the all ACH TLV objects.

The TTSI TLV Object in figure below consists of the MEP-ID followed by the Unique ME identifier 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    TTSIAchTlvType = TBD       |           Length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       MEP ID value            |Reserved ( fixed to all ZEROs) !
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
~                  Unique ME ID TLV                             ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Length

2 octets field; it specifies the number of octets which follows the Length field.

MEP ID value

13-bit integer value field, identifying the transmitting MEP within the ME. The three MSBs of the first octet are not used and set to ZERO.

Unique ME ID

This is the ME Identifier TLV Object as described in section 2.1.

2.1. Unique ME Identifier

The globally unique ME Identifier ACH TLV objects have the following structure :

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    ME ID Type                 |           Length              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   ~                     ME ID Value                               ~
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

ME ID Type

2 octet field; it identifies the format of the ME Identifier;

Length

2 octets field; it identifies the length in octets of the ME ID Section that follows the length field.

ME ID payload

value of the ME identifier; its semantic depends on the format.

The ME Identifier Type transmitted and expected MUST be the same at both MEPs. For statically provisioned connections, the ME Identifier transmitted and expected is statically configured at both MEPs. For dynamically established connections, the ME Identifier transmitted and expected is signaled via the control plane. The extension of ME identifier signaling is outside the scope of this document.

Some possible ME Identifier formats are reported in the following sections.

2.1.1. LSP ME ID IPv4 Source/Destination Address Format

This ME ID format MAY be used to identify an LME (as defined in [5]) where IPv4 addresses are used to identify the LERs terminating the LSP.

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    ME ID Type                 |           Length              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                   IPv4 source address                         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                   IPv4 destination address                    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |         Tunnel Index          | TunnelInstance  Index         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

ME ID Type

2 octet field; it identifies the specific format, value = TBD;

Length

2 octets field; set to 12 (octets);

IPv4 source address

4 octets field; set to the IPv4 address of the LSP source port/node;

IPv4 destination address

4 octets field; set to the IPv4 address of the LSP destination port/node;

TunnelIndex

2 octets field as defined in RFC 3812

TunnelInstance Index

2 octets field as defined in RFC 3812

2.1.2. LSP ME ID IPv6 Source/Destination Address Format

This ME ID format MAY be used to identify an LME (as defined in [5]) where IPv6 addresses are used to identify the LERs terminating the LSP.

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    ME ID Type                 |           Length              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                      IPv6 source address                      |
   ~                         (16 bytes)                            ~
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                      IPv6 destination address                 |
   ~                         (16 bytes)                            ~
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |         Tunnel Index          | TunnelInstance  Index         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

ME ID Type

2 octet field; it identifies the specific format, value = TBD;

Length

2 octets field; set to 36 (octets);

IPv6 source address

4 octets field; set to the IPv6 address of the LSP source port/node;

IPv6 destination address

4 octets field; set to the IPv6 address of the LSP destination port/node;

TunnelIndex

2 octets field as defined in RFC 3812

TunnelInstance Index

2 octets field as defined in RFC 3812

2.1.3. Type FEC128PWv4 Format

This TLV is defined in [10]. It contains a PW ID that terminates on a PE identified by an IPv4 address.

This ME ID format MAY be used to identify a PME (as defined in [5]) where IPv4 addresses are used to identify the T-PEs terminating the PW and FEC128 is used to identify the PW.

2.1.4. Type FEC128PWv6 Format

This TLV is defined in [10]. It contains a PW ID that terminates on a PE identified by an IPv6 address.

This ME ID format MAY be used to identify a PME (as defined in [5]) where IPv6 addresses are used to identify the T-PEs terminating the PW and FEC128 is used to identify the PW.

Editor's note: implementation impacts of FEC129PWv4 and FEC129PWv6 (as defined in [10])when used as ME Identifier in a cc-cv proactive tool needs further study (see note in section 2 regarding length of ME ID).

2.1.5. ICC-based Format

This ME ID format MAY be used to identify SME, LME, LTCME, PME and PTCME(as defined in [5]) independently on LER/T-PE addressing schemes as well as of the FECs used to identify the PW.

Editor's note: for these reasons it is suggested to have this format as a MUST and the other format as optional in the MPLS-TP environment.

EndofEditornote

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |        ME ID Type             |           Length              |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   +                                                               +
   |                        MEG ID                                 |
   +                      (13 bytes)                               +
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |               |0 0 0 0 0 0 0 0|0 0 0 0 0 0 0 0|0 0 0 0 0 0 0 0|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

ME ID Type

2 octet field; it identifies the specific format, value = TBD;

Length

2 octets field; set to 16;

MEG ID value

13-octet field. Refer to Annex A of ITU-T Recommendation Y.1731 for the format used for the MEG ID field with ICC-based format.

Editor's note: to decide if insert all ICC-based MEG ID format (as for Figure A-2 of Y.1731 or only the MEG-ID value as reported now in above figure EndofEditornote

3. Two different ACH encapsulation of BFD tool

Among the solutions analyzed in [11], the one based on two separate tools, running with two different ACH encapsulations (i.e. two different ACH channel types):

- the current BFD with only CC functionality;

- a new tool based on current BFD that meet all the OAM MPLS-TP requirement.

is proposed in this document as the solution to be standardized.

As all analyzed solutions reported in [11], even this one implies extension of CV types, foreseen by [6] and yet extended by[7], in order to include the MPLS-TP OAM mechanism too for PW Fault Detection only. This is due to the fact that VCCV also includes mechanisms for negotiating the control channel and connectivity verification (i.e. OAM functions) between PEs.

This version of the draft is focused on fault detection on the transport path.

3.1. Current BFD with only CC functionality

The current BFD, with only CC functionality, is encapsulated in the G-ACH using as Channel type code point the 0x0007 value as described in [7]. This mechanism can be even extended to SME, LME, LTCME and PTCME.

Figure below shows G-ACH encapsulation of current BFD as in [7]

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 1|0 0 0 0|0 0 0 0 0 0 0 0|0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                      BFD control packet                       |
+                                                               +
:                               ...                             :
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

3.2. New tool based on current BFD

The new tool is obtained introducing TTSI TLV , described in section 2. between the ACH and the current BFD (defined in [8]) as detailed below.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 1|0 0 0 0|0 0 0 0 0 0 0 0|  MPLS-TP CC-CV proactive      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    ACH TLV Header                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                    TTSI  TLV                                  
:                              ...                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
~                    BFD Control packet                         ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The first four bytes represent the G-ACH ([2]).

- first nibble: set to 0001b to indicate a channel associated with a PW, a LSP or a Section;

- Version and Reserved fields are set to 0, as specified in[2] [2];

- G-ACH Channel Type field with a new TBD code point meaning "MPLS-TP CC-CV proactive" indicating that the message is an MPLS-TP OAM CC-CV proactive message. The value MUST be assigned

The G-ACH is followed by the ACH TLV Header as defined in Section 3 of [2] and by one and only one TTSI ACH TLV object carrying the MEP Identifier and the Unique ME Identifier Information. (see section 2.)

The IP-based addressing ACH TLV objects (see [10]) that allow the usage if this tool also within an IP/MPLS environment can be even present. In any case the TTSI ACH TLV object MUST be always carried in the last ACH TLV object (i.e just before the BFD packet).

The BFD control packet is the base BFD as described [8].

The benefit of this solution is to maintain the base behavior and protocol version of BFD as defined in[8] and in [9]; however, it's necessary to define some rules that specify the BFD profile for MPLS-TP application as detailed in the next sections.

The proposed solutions needs even some considerations on the optional Authentication Section how described in section 10.

From now on this solution is referred as extended BFD.

4. Backward compatibility

For backward compatibility, it is still possible to run the current BFD that supports only CC functionality on some transport paths and the new tool that supports CC and CV functionality on other transport paths. In any case, only one tool for OAM instance at time, configurable by operator, can run.

A MEP that is configured to support CC and CV functionality, as required by MPLS-TP, MUST be capable to receive existing BFD packets (encapsulated with GAL/G-ACH or PW-ACH) that supports only CC functionality and MUST consider them as an unexpected packet, i.e. detect a misconnection defect.

5. BFD behavior specification for MPLS-TP OAM proactive CC&CV

In this section some rules that specify the MPLS-TP application of the extended BFD is detailed. These rules apply even for the base BFD, limited to the CC functionality only, when an MPLS-TP node interoperates with MPLS legacy nodes or when only the CC functionality is required .

The BFD control packet is the base BFD as described in [8].

The extended BFD MUST operate both in bidirectional and unidirectional p2p and unidirectional p2mp connection.

This version of the draft is focused on unidirectional and bidirectional p2p connection of the MPLS-TP CC-CV proactive tool.

Unidirectional p2mp connections are even reported but it requires further analysis.

In this document the BFD operating mode is always Asynchronous; in this mode the systems send CC/CV BFD packets, carrying a unique ME identifier and sender MEP identifier, at regular configurable timing rate. If a LOC defect or a mis-connectivity defect or a period mis-configuration defect occurs, the BFD session is declared down. For proactive CC-CV functionality description please refers to [5].

5.1. BFD State Machine

The BFD State Machine is described in [8] for bidirectional p2p connection and in [9] for p2mp unidirectional connection.

The main goals in defining the state machine for MPLS-TP extended BFD are:

1. **to interoperate with the state machines defined in [8] and in** [9];

2. **to satisfy the CC-CV proactive monitoring functionality and even the RDI functionality as specified in the MPLS-TP OAM framework** [5].

The diagram in Figure 1 provides an overview of the state machine for MEP configured on p2p bidirectional transport path, as defined in [8].

```
                               +--+
                               |  | UP, ADMIN DOWN, Defects,
                               |  V
                       DOWN  +------+  INIT
                +------------|      |------------+
                |            | DOWN |            |
                |  +-------->|      |<--------+  |
                |  |         +------+         |  |
                |  |                          |  |
                |  |               ADMIN DOWN,|  |
                |  |ADMIN DOWN,          DOWN,|  |
                |  |Defects           Defects |  |
                V  |                          |  V
              +------+                      +------+
         +----|      |                      |      |----+
     DOWN|    | INIT |--------------------->|  UP  |    |INIT, UP
         +--->|      | INIT, UP             |      |<---+
              +------+                      +------+
```

Figure 1: State Machine for p2p bidirectional connection

The diagram in Figure 2 provides an overview of the state machine for MEP Sink of unidirectional p2p unidirectional transport path based on the state machine defined in [9];

```
                               (DOWN), ADMIN DOWN,
                    +------+   Defects            +------+
               +----|      |<---------------------|      |----+
         (DOWN)|    | DOWN |                      |  UP  |    |UP
    ADMIN DOWN,+--->|      |--------------------->|      |<---+
         Defects    +------+          UP          +------+
```

Figure 2: State Machine for p2p unidirectional connection

State transitions on MEP Source on unidirectional p2p path are administratively driven.

In both diagram, each arc represents the state of the remote system (as received in the State field in the BFD Control packet) or indicates the expiration of the Detection Timer, here extended to general defect raising and clearing conditions as reported in Table 1.

As reported in [8], another state (AdminDown) exists so that a session can be administratively put down indefinitely. In the above diagram Transitions involving AdminDown state are deleted for clarity; further considerations are reported in section 5.4.

Editor's note:

. It is suggested to introduce the following new bfd.SessionType in [9]:

o MPLS-TP bidirectional p2p: Transport Profile of the Classic point-to-point BFD .

o MPLS-TP unidirectional Sink

o MPLS-TP unidirectional Source

. In [9] DOWN State in received BFD is even reported. It is suggested to remove it from [9] as well as from Figure 3 as State transitions on MEP Source on unidirectional p2p and p2mp path are administratively driven .

. If in Figure 1 we assume that

o a MPLS-TP node transits from DOWN to UP State even when receiving UP State and not only INIT State and that

- o a Source MEP of a unidirectional path only transmits BFD packets in UP state or AdminDown State, the diagram in Figure 2 becomes a ''sub-branch'' of state machine reported in Figure 1

EndofEditor's note

. The ''Defects Raise'' and ''Defects Clear'' Event in Figure 1 and Figure 2 occurs when:

- o Defect Raise Event = dLOC or dUNME or dUNM or dUNP;

- o Defect Clear Event = (not dLOC)and(not dUNME)and(not dUNM)and (not dUNP);

- o See section 5.2. for defect conditions description.

. When on a configured bidirectional MEP the proactive CC-CV monitoring is enabled, the MEP sends the CC/CV BFD packet with frequency of the configured transmission period and it also expects to receive the CC/CV BFD packets from its peer MEP with the same transmission period (see [5]).

. In a unidirectional (point-to-point or point-to-multipoint) transport path , where the proactive CC-CV monitoring is enabled, only the Source MEP is enabled to generate CC/CV BFD packets with frequency of the configured transmission period and always with UP State information . This MEP does not expect to receive any CC/CV BFD packets from its peer MEP in the ME

. a MEP Sink, configured on a unidirectional transport path where the proactive CC-CV monitoring is enabled, expects to receive the CC/CV BFD packets from its peer MEP at the configured period; the defects detection procedure is the same as the bidirectional MEP; no CC/CV BFD packets are sent on the ME.

It is worth noticing that CC-CV proactive monitoring can be enabled/disabled by an operator on a configured ME and this MUST be not traffic affecting. Please refers to [5] for hitless enable/disable CC-CV proactive monitoring procedure.

When a BFD session transits from one state to another, the traffic MUST not be affected. The blocking of traffic as consequent action MUST be driven only by a defect's consequent action as specified in the consequent action section 5.3.

When the CC-CV proactive monitoring is disabled on a ME no CC/CV BFD packets are sent on the ME neither defects are monitored; however, the MEP must terminate all OAM packets it receives and in case of CC-CV proactive packets it must recognize a mis-connectivity even if the CC-CV proactive monitoring is disabled on the MEP.

5.2. Defect conditions

Defect triggered at the MPLS-TP layer by the CC-CV proactive tool are reported in this section.

5.2.1. Loss of Continuity defect(dLOC)

If no CC/CV BFD packets from a peer MEP are received within the interval equal to K times (see Table 1) the receiving MEP's configured CC-CV transmission period, loss of continuity with peer MEP is detected.

5.2.2. Mis-connectivity defect (dUNME)

If an extended BFD with a ME ID different than the configured expected one is received , mis-connectivity defect is detected.

5.2.3. Unexpected MEP ID defect (dUNM)

If a CC-CV BFD packet with expected ME ID but with an incorrect MEP ID, including the receiving MEP's own MEP ID, is received, Unexpected MEP is detected.

5.2.4. Unexpected Period Defect (dUNP)

If a CC-CV BFD packet is received with a correct ME ID, a correct MEP ID, but with a period field value different than the receiving MEP's own CC-CV transmission period, Unexpected Period is detected.

5.2.5. RDI Defect (dRDI)

The Remote Defect Indicator defect monitors the presence of an RDI maintenance signal. A MEP detects RDI defect when it receives a CC/CV BFD packet with the RDI field set.

The following Table summarizes the defect raising and clearing conditions

Defect	Raising Condition	Clearing Condition
dLOC	#Rx BFD==0 (K*CV_Period)	#Rx BFD >= N (K*CV_Period)
dUNME	Unexpected ME Identifier	#UnexpMEId==0 (K*CV_Period)
dUNM	Unexpected MEP Id	#UnexpMEPId==0 (K*CV_Period)
dUNP	Unexpected Period	#UnexpPeriod==0 (K*CV_Period)
dRDI	Rx RDI == 1	Rx RDI ==0

Table 1: Table of Raising Clearing Defect Conditions

In Table 1, N and K are protocol constants to be defined; the CV_Period corresponds to the configured transmission period;

K value is stored in bfd.DetectMult variable while the CV_Period is stored bfd.DesiredMinTxInterval variable and carried in the relative BFD packet fields.

5.3. Consequent action

- If a MEP detects an unexpected ME Identifier, or an unexpected MEP, it MUST block (BLK) all the traffic (including also the user data packets) that it receives from the misconnected transport path.

- If a MEP detects LOC defect and the CC-CV monitoring is enabled it MUST block (BLK) all the traffic (including also the user data packets) that it receives from the misconnected transport path. This consequent action is configurable (see [5])

- If a MEP detects an unexpected ME Identifier, or an unexpected MEP it MUST declare a Transport Path Fail (TSF).

- If a MEP detects LOC defect and the CC-CV monitoring is enabled it MUST declare a Transport Path Fail (TSF).

- If a MEP declare a Transport Path Fail it MUST insert an RDI information towards its peer MEP in a bidirectional connection

Consequent actions are identified by abbreviation prefixed with an 'a': aXXX, e.g. aTSF. The consequent actions detailed above can be so summarized :

aBLK <-- (dUNME or dUNM) or (dLOC and CC-CV_Monitoring_Enabled and LOC_cons_action_Enabled)

aTSF <-- (dLOC and CC-CV_Monitoring_Enabled)or dUNME or dUNM or Server_SF

aRDI <-- aTSF

5.4. Administrative Down State

As reported in[8]: ''a session MAY be kept administratively down by entering the AdminDown state and sending an explanatory diagnostic code in the Diagnostic field. AdminDown state means that the session is being held administratively down. This causes the remote system to enter Down state, and remain there until the local system exits AdminDown state. AdminDown state has no semantic implications for the availability of the forwarding path.''

Please not that the AdminDown state semantic MUST be not confused with the LOCK functionality as reported in [3]. In this document the AdminDown state semantic is equivalent to disabling on a MEP the CC-CV proactive monitoring; in this case the source MEP SHOULD send BFD Control packets in AdminDown state for a period equal to(bfd.DesiredMinTxInterval * bfd.DetectMult) in order to ensure that the remote system is aware of the state change.

An MPLS-TP MEP receiving a BFD packet with AdminDown State MUST transit to the DOWN State and report the event to the operator. Editor's note:

In this case the operator should disable the functionality even on the other node. Should the node still report the LOC defect ? If Yes the AdminDown State received can be a simple report/warning to the operator.

End editor's note

5.5. Timer Negotiation

The negotiation of time value foreseen by the base BFD (see [8]) MUST be disabled on the MPLS-TP extended BFD. The timer negotiation should be even disabled on an MPLS-TP node that runs only the CC functionality as well as on a node that interoperates with current BFD. It's up to the operator configure the BFD transmission period on the MPLS-TP node in a way that is suitable for the MPLS legacy node.

The configured BFD packet transmission period MUST be stored into the bfd.DesiredMinTxInterval variable and carried into the ''Desired Min TX Interval field'' of the transmitted BFD

For a bidirectional point-to-point transport path the same bfd.DesiredMinTxInterval value MUST be stored even in bfd.RequiredMinRxInterval; source MEP of unidirectional session MUST set the bfd.RequiredMinRxInterval to 0.

The bfd.DesiredMinRxInterval value is carried into the ''Required Min RX Interval field '' of the transmitted BFD.

If BFD packets are received with the ''Desired Min TX Interval field '' different from than expected and stored in the local bfd.DesiredMinTxInterval variable the Unexpected Period defect is detected.

Please note that this behavior doesn't preclude the base BFD session running on a MPLS legacy node to adapt to the BFD timers transmitted by an MPLS-TP node.

The configured transmission period MUST be one of the following values, both for the extended BFD and current BFD running on a MPLS-TP node:

. 3.33 ms: Default transmission period for protection switching application (transmission rate of 300 packets/second);

. 10 ms: (Transmission rate is 100 packets/second);

. 100 ms: Default transmission period for performance monitoring application (transmission rate of 10 packets/second);

. 1 s: Default transmission period for fault management application (transmission rate of 1 packets/second).

5.6. Demultiplexing and the Discriminator Fields

The context of MPLS-TP OAM packets is based on MPLS label and G-ACH, eliminating in the BFD the need to exchange Discriminator values.

An MPLS-TP node that interoperates with a base BFD can apply the same discriminator field semantic as described in [8] or:

. It MUST set the My discriminator field to a nonzero value (it can be a fixed value);

. It MUST reflect back the received value of My discriminator field into the transmitted Your discriminator field, or set it to zero if that value is unknown.

6. Remote Defect Indication

The Remote Defect Indication (RDI) is an indicator that is transmitted by a MEP to communicate to its peer MEP that a signal fail condition exists. RDI is only used for bidirectional connections and is associated with proactive CC & CV packet generation.[5]

Editor's note: in last version of MPLS-TP OAM requirement ([3])the RDI functionality is now even for unidirectional connection if a return path exist. However this implies a further analysis

End of Editor's note

The Diagnostic (Diag) field of base BFD ([8]) can be used for this functionality. However, there isn't a total correspondence among the values foreseen by [8] and the defect conditions detected by the proactive CC-CV tool that require the RDI function. A solution could be that any defect that requires the RDI information being sent to the peer MEP is encoded in the Diagnostic (Diag) field with the value 1 (corresponding to the ''Control Detection Time Expired'' in [8]). The value 0 indicates RDI condition has been cleared.

This applies for the extended BFD and for the base BFD when a MPLS-TP node runs only the CC functionality

7. BFD Control Packets field

As defined in [8], the mandatory Section of a BFD Control packet has the following format:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Vers |  Diag   |Sta|P|F|C|A|D|M|  Detect Mult  |    Length     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       My Discriminator                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Your Discriminator                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Desired Min TX Interval                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Required Min RX Interval                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 Required Min Echo RX Interval                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The contents of transmitted and expected BFD Control packets on a MPLS-TP node for the mechanism proposed in this document are detailed below both for extended BFD as well as current BFD:

Version

Set to the current version number (1). If version number is not correct the packet is discarded.

Diagnostic (Diag)

The following value are requested to be supported among the value reported in [8]:

0 -- No Diagnostic

1 -- Remote Defects Indication

3 -- Neighbor Signaled Session Down

7 -- Administratively Down

Different values are ignored in receipt.

State (Sta)

The current BFD session state as seen by the transmitting system. Values are:

0 -- AdminDown

1 -- Down

2 -- Init

3 -- Up

Poll (P)

If set, the transmitting system is requesting verification of connectivity, or of a parameter change, and is expecting a packet with the Final (F) bit in reply. If clear, the transmitting system is not requesting verification.

Final (F)

If set, the transmitting system is responding to a received BFD Control packet that had the Poll (P) bit set. If clear, the transmitting system is not responding to a Poll.

Control Plane Independent (C)

Is not required to be supported on a MPLS-TP node. Set to 1 on transmit and ignored in receipt.

Authentication Present (A)

Set to 1 if authentication is in use on this session (bfd.AuthType is nonzero), or 0 if not.

Demand (D)

Set to 0 in asynchronous mode.

Multipoint (M)

Set to 1 if bfd.SessionType is MultipointHead (Source MEP of a point to multipoint connection). Otherwise it is set to 0.

Detect Mult

Detection time multiplier. The configured transmit interval, multiplied by this value, provides the Detection Time for the transmitting system in Asynchronous mode. This is the K protocol constants as defined in Table 1.

Length

Set to the appropriate length, in bytes, of the only BFD Control packet, based on the fixed header length (24) plus any Authentication Section.

My Discriminator

A unique, nonzero discriminator value generated by the transmitting system. See section 5.6.

Your Discriminator

The discriminator received from the corresponding remote system. This field reflects back the received value of My Discriminator, or is zero if that value is unknown. See section 5.6.

Desired Min TX Interval

This is the configured BFD packet transmission period, in microseconds and stored into the bfd.DesiredMinTxInterval variable. See section 5.5.

Required Min RX Interval

This is the configured BFD packet period, in microseconds and stored into the bfd.RequiredMinRxInterval variable. Source MEP of unidirectional session MUST set it to 0 See section 5.5.

Required Min Echo RX Interval

Is not required to be supported on a MPLS-TP node; set to 0.

An optional Authentication Section MAY be present. For details see section 10.

8. Interoperability with current BFD

There's no interoperability issue among a MPLS-TP node running the current BFD, with the rules defined in previous sections, and legacy MPLS node at PW layer network level as the use of the CC Type 1 was previously defined and limited to PWs. (See [7])

Instead, any Network Partitioning scenario with LSP Stitching present interoperability issues as LSP Ping is designated to bootstrap the BFD session in an MPLS environment and the session BFD messages for MPLS are transmitted using a IP/UDP encapsulation (see [12] and [4]); the IWF requires further analysis.

This section will be completed in the next version of the draft.

9. Point to Multipoint transport paths

Solution described in section 3.2. is also valid for p2mp connections. The extended multipoint BFD packets are explicitly marked as such, via the setting of the M bit (see [9]).

The diagrams reported Figure 2 provides also an overview of the state machine for MEP Sink configured on each tail of a p2mp path.

MEP Source, head of unidirectional p2mp, will never receive packets and have no Detection Timer, and as such all state transitions are administratively driven.

The MEP Source sends the extended multipoint BFD packet in the multicast tree at a configured transmission period.

A unidirectional point-to-multipoint connection containing n end-points contains (n-1) MEs, each one independently monitored by MEP Sink configured on each tail as described in section 5.

Editor's note:

. It is suggested to introduce the following new bfd.SessionType in [9]:

oMPLS-TP MultipointHead: Transport Profile of MultipointHead oMPLS-TP MultipointTail: Transport Profile of MultipointTail

EndofEditor's note

This section requires further analysis.

10. Acknowledgments

<Add any acknowledgements>

This document was prepared using 2-Word-v2.0.template.dot.

11. Contributors

12. IANA Considerations

<Add any IANA considerations>

13. Security Considerations

Base BFD [8] foresees an optional authentication section; that can be extended even to the tool proposed in this document.

Authentication methods that require checksum calculation on the outgoing packet must extend the checksum even on the ME Identifier Section. This is possible but seems uncorrelated with the solution proposed in this document: it could be better to use the simple password authentication method.

It is also worth noticing that the interactions between authentication and connectivity verification need further analysis.

14. References

14.1. Normative References

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

[2] Bocci, et al., " MPLS Generic Associated Channel ", RFC 5586 , June 2009

[3] Vigoureux, M., Betts, M., Ward, D., "Requirements for OAM in MPLS Transport Networks", draft-ietf-mpls-tp-oam-requirements-01 (work in progress), March 2009

[4] Sprecher, N., Nadeau, T., van Helvoort, H., Weingarten, Y., " MPLS-TP OAM Analysis", draft-sprecher-mpls-tp-oam-analysis-04 (work in progress), May 2009

[5] Busi,I., Niven-Jenkins, B. "MPLS-TP OAM Framework and Overview", draft-ietf-mpls-tp-oam-framework-00(work in progress), March 2009

[6] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007

[7] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)",ID draft-ietf-pwe3-vccv-bfd-04.txt, Work in Progress, May 2009

[8] Katz, D. and D. Ward, "Bidirectional Forwarding Detection", draft-ietf-bfd-base-09.txt (work in progress), February 2009.

[9] Katz, D. and D. Ward, "BFD for Multipoint Networks", ID draft-katz-ward-bfd-multipoint-02.txt, February 2009

[10] S. Boutros, et. al., "Definition of ACH TLV Structure", draft-bryant-mpls-tp-ach-tlv-00.txt, Work in Progress, January 2009.

[11] A. Fulignoli, et. al.,''MPLS-TP Proactive Continuity and Connectivity Verification'', draft-fhbs-mpls-tp-cv-proactive-00.txt (work in progress), February 2009

[12] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "BFD For MPLS LSPs", draft-ietf-bfd-mpls-07 (work in progress), June 2008.

[13] Martini, L.and M.Morrow, ''Pseudo Wire (PW) OAM Message Mapping'', draft-eitd-pwe3-oam-msg-map-10 (work in progress), April 2009

14.2. Informative References

Authors' Addresses

Annamaria Fulignoli (Editor)
Ericsson
Email: annamaria.fulignoli@ericsson.com

Sami Boutros
Cisco Systems, Inc.
Email: sboutros@cisco.com

Martin Vigoureux
Alcatel-Lucent
Email: martin.vigoureux@alcatel-lucent.com