

Internet Engineering Task Force (IETF)
Request for Comments: 6427
Category: Standards Track
ISSN: 2070-1721

G. Swallow, Ed.
Cisco Systems, Inc.
A. Fulignoli, Ed.
Ericsson
M. Vigoureux, Ed.
Alcatel-Lucent
S. Boutros
Cisco Systems, Inc.
D. Ward
Juniper Networks, Inc.
November 2011

MPLS Fault Management Operations, Administration, and Maintenance (OAM)

Abstract

This document specifies Operations, Administration, and Maintenance (OAM) messages to indicate service disruptive conditions for MPLS-based transport network Label Switched Paths. The notification mechanism employs a generic method for a service disruptive condition to be communicated to a Maintenance Entity Group End Point. This document defines an MPLS OAM channel, along with messages to communicate various types of service disruptive conditions.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6427>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. Requirements Language	5
2. MPLS Fault Management Messages	5
2.1. MPLS Alarm Indication Signal	5
2.1.1. MPLS Link Down Indication	6
2.2. MPLS Lock Report	6
2.3. Propagation of MPLS Fault Messages	7
3. MPLS Fault Management Channel	7
4. MPLS Fault Management Message Format	8
4.1. Fault Management Message TLVs	9
4.1.1. Interface Identifier TLV	10
4.1.2. Global Identifier	10
5. Sending and Receiving Fault Management Messages	10
5.1. Sending a Fault Management Message	10
5.2. Clearing a Fault Management Indication	11
5.3. Receiving a Fault Management Indication	11
6. Minimum Implementation Requirements	12
7. Security Considerations	12
8. IANA Considerations	13
8.1. Pseudowire Associated Channel Type	13
8.2. MPLS Fault OAM Message Type Registry	13
8.3. MPLS Fault OAM Flag Registry	14
8.4. MPLS Fault OAM TLV Registry	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15
10. Contributing Authors	16

1. Introduction

Proper operation of a transport network depends on the ability to quickly identify faults and focus attention on the root cause of the disruption. This document defines MPLS Fault Management Operations, Administration, and Maintenance (OAM) messages. When a fault occurs in a server (sub-)layer, Fault Management OAM messages are sent to clients of that server so that alarms, which otherwise would be generated by the subsequent disruption of the clients, may be suppressed. This prevents a storm of alarms and allows operations to focus on the actual faulty elements of the network.

In traditional transport networks, circuits such as T1 lines are typically provisioned on multiple switches. When an event that causes disruption occurs on any link or node along the path of such a transport circuit, OAM indications are generated. When received, these indications may be used to suppress alarms and/or activate a backup circuit. The MPLS-based transport network provides mechanisms equivalent to traditional transport circuits. Therefore, a Fault Management (FM) capability must be defined for MPLS. This document defines FM capabilities to meet the MPLS-TP requirements as described in RFC 5654 [1], and the MPLS-TP Operations, Administration, and Maintenance requirements as described in RFC 5860 [2]. These mechanisms are intended to be applicable to other aspects of MPLS as well. However, applicability to other types of LSPs is beyond the scope of this document.

Two broad classes of service disruptive conditions are identified.

1. **Fault:** The inability of a function to perform a required action. This does not include an inability due to preventive maintenance, lack of external resources, or planned actions.
2. **Lock:** an administrative status in which it is expected that only test traffic, if any, and OAM (dedicated to the LSP) can be sent on an LSP.

Within this document, a further term is defined: server-failure. A server-failure occurs when a fault condition or conditions have persisted long enough to consider the required service function of the server (sub-)layer to have terminated. In the case of a protected server, this would mean that the working facilities and any protection facilities have all suffered faults of the required duration.

This document specifies an MPLS OAM channel called an "MPLS-OAM Fault Management (FM)" channel. A single message format and a set of procedures are defined to communicate service disruptive conditions

from the location where they occur to the end points of LSPs that are affected by those conditions. Multiple message types and flags are used to indicate and qualify the particular condition.

Corresponding to the two classes of service disruptive conditions listed above, two messages are defined to communicate the type of condition. These are known as:

Alarm Indication Signal (AIS)

Lock Report (LKR)

1.1. Terminology

ACH: Associated Channel Header

ACH: Associated Channel

CC: Continuity Check

FM: Fault Management

GAL: Generic Associated Channel Label

LOC: Loss of Continuity

LSP: Label Switched Path

MEP: Maintenance Entity Group End Point

MPLS: Multiprotocol Label Switching

MPLS-TP: MPLS Transport Profile

MS-PW: Multi-Segment Pseudowire

OAM: Operations, Administration, and Maintenance

PHP: Penultimate Hop Pop

PW: Pseudowire

TLV: Type, Length, Value

1.2. Requirements Language

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 [3].

2. MPLS Fault Management Messages

This document defines two messages to indicate service disruptive conditions, Alarm Indication Signal and Lock Report. The semantics of the individual messages are described in subsections below. Fault OAM messages are applicable to LSPs used in the MPLS Transport Profile. Such LSPs are bound to specific server layers based upon static configuration or signaling in a client/server relationship.

Fault Management messages are carried in-band of the client LSP or MS-PW by using the Associated Channel Header (ACH). For LSPs other than PWs, the ACH is identified by the Generic Associated Channel Label (GAL) as defined in RFC 5586 [4]. To facilitate recognition and delivery of Fault Management messages, the Fault Management Channel is identified by a unique Associated Channel (ACh) code point.

Fault OAM messages are generated by intermediate nodes where a client LSP is switched. When a server (sub-)layer, e.g., a link or bidirectional LSP, used by the client LSP fails, the intermediate node sends Fault Management messages downstream towards the end point of the LSP. The messages are sent to the client MEPs by inserting them into the affected client LSPs in the direction downstream of the fault location. These messages are sent periodically until the condition is cleared.

2.1. MPLS Alarm Indication Signal

The MPLS Alarm Indication Signal (AIS) message is generated in response to detecting faults in the server (sub-)layer. The AIS message SHOULD be sent as soon as the condition is detected, but MAY be delayed owing to processing in an implementation, and MAY be suppressed if protection is achieved very rapidly. For example, an AIS message may be sent during a protection switching event and would cease being sent (or cease being forwarded by the protection switch selector) if the protection switch was successful in restoring the link. However, an implementation may instead wait to see if the protection switch is successful prior to sending any AIS messages.

The primary purpose of the AIS message is to suppress alarms in the layer network above the level at which the fault occurs. When the Link Down Indication is set, the AIS message can be used to trigger recovery mechanisms.

2.1.1. MPLS Link Down Indication

The Link Down Indication (LDI) is communicated by setting the L-Flag to 1. A node sets the L-Flag in the AIS message in response to detecting a failure in the server layer. A node **MUST NOT** set the L-Flag until the fault has been determined to be a server-failure. A node **MUST** set the L-Flag if the fault has been determined to be a server-failure. For example, during a server layer protection switching event, a node **MUST NOT** set the L-Flag. However, if the protection switch was unsuccessful in restoring the link within the expected repair time, the node **MUST** set the L-Flag.

The setting of the L-Flag can be predetermined based on the protection state. For example, if a server layer is protected and both the working and protection paths are available, the node should send AIS with the L-Flag clear upon detecting a fault condition. If the server layer is unprotected, or the server layer is protected but only the active path is available, the node should send AIS with the L-Flag set upon detecting a loss of continuity (LOC) condition. Note again that the L-Flag is not set until a server-failure has been declared. Thus, if there is any hold-off timer associated with the LOC, then the L-Flag is not set until that timer has expired.

The receipt of an AIS message with the L-Flag set **MAY** be treated as the equivalent of LOC at the client layer. The choice of treatment is related to the rate at which the Continuity Check (CC) function is running. In a normal transport environment, CC is run at a high rate in order to detect a failure within tens of milliseconds. In such an environment, the L-Flag **MAY** be ignored and the AIS message is used solely for alarm suppression.

In more general MPLS environments, the CC function may be running at a much slower rate. In this environment, the Link Down Indication enables faster switch-over upon a failure occurring along the client LSP.

2.2. MPLS Lock Report

The MPLS Lock Report (LKR) message is generated when a server (sub-)layer entity has been administratively locked. Its purpose is to communicate the locked condition to the client-layer entities. When a server layer is administratively locked, it is not available to carry client traffic. The purpose of the LKR message is to

suppress alarms in the layer network above the level at which the administrative lock occurs and to allow the clients to differentiate the lock condition from a fault condition. While the primary purpose of the LKR message is to suppress alarms, similar to AIS with the LDI (L-Flag set), the receipt of an LKR message can be treated as the equivalent of loss of continuity at the client layer.

2.3. Propagation of MPLS Fault Messages

MPLS-TP allows for a hierarchy of LSPs. When the client MEP of an LSP (that is also acting as a server layer) receives FM indications, the following rules apply. If the CC function is disabled for the server LSP, a node SHOULD generate AIS messages toward any clients when either the AIS or LKR indication is raised. Note that the L-Flag is not automatically propagated. The rules of Section 2.1.1 apply. In particular, the L-Flag is not set until a server-failure has been declared.

3. MPLS Fault Management Channel

The MPLS Fault Management channel is identified by the ACH as defined in RFC 5586 [4] with the Associated Channel Type set to the MPLS Fault Management (FM) code point = 0x0058. The FM Channel does not use ACh TLVs and MUST NOT include the ACh TLV header. The ACH with the FM ACh code point is shown below.




Figure 1: ACH Indication of the MPLS Fault Management Channel

The first three fields are defined in RFC 5586 [4].

The Fault Management Channel is 0x0058.

4. MPLS Fault Management Message Format

The format of the Fault Management message is shown below.




Figure 2: MPLS Fault OAM Message Format

Version

The Version Number is currently 1.

Reserved

This field MUST be set to zero on transmission and ignored on receipt.

Message Type

The Message Type indicates the type of condition as listed in the table below.

Msg Type	Description
0	Reserved
1	Alarm Indication Signal (AIS)
2	Lock Report (LKR)

Flags

Two flags are defined. The reserved flags in this field MUST be set to zero on transmission and ignored on receipt.




Figure 3: Flags

L-Flag

Link Down Indication. The L-Flag only has significance in the AIS message. For the LKR message, the L-Flag MUST be set to zero and ignored on receipt. See Section 2.1.1 for details on setting this bit.

R-Flag

The R-Flag is clear to indicate the presence of an FM condition and is set to one to indicate the removal of a previously sent FM condition.

Refresh Timer

The maximum time between successive FM messages specified in seconds. The range is 1 to 20. The value 0 is not permitted.

Total TLV Length

The total length in bytes of all included TLVs.

4.1. Fault Management Message TLVs

TLVs are used in Fault Management messages to carry information that may not pertain to all messages as well as to allow for extensibility. The TLVs currently defined are the IF_ID and the Global_ID.

TLVs have the following format:




Figure 4: Fault TLV Format

Type

Encodes how the Value field is to be interpreted.

Length

Specifies the length of the Value field in octets.

Value

Octet string of Length octets that encodes information to be interpreted as specified by the Type field.

4.1.1.1. Interface Identifier TLV

The Interface Identifier (IF_ID) TLV carries the IF_ID as defined in RFC 6370 [5]. The Type is 1. The length is 0x8.




Figure 5: Interface Identifier TLV Format

4.1.1.2. Global Identifier

The Global Identifier (Global_ID) TLV carries the Global_ID as defined in RFC 6370 [5]. The Type is 2. The length is 0x4.




Figure 6: Global Identifier TLV Format

5. Sending and Receiving Fault Management Messages

5.1. Sending a Fault Management Message

Service disruptive conditions are indicated by sending FM messages. The message type is set to the value corresponding to the condition. The Refresh Timer is set to the maximum time between successive FM messages. This value MUST NOT be changed on successive FM messages reporting the same incident. If the optional clearing procedures are not used, then the default value is one second. Otherwise, the default value is 20 seconds.

A Global_ID MAY be included. If the R-Flag clearing procedures are to be used, the IF_ID TLV MUST be included. Otherwise, the IF_ID TLV MAY be included.

The message is then sent. Assuming the condition persists, the message MUST be retransmitted two more times at an interval of one second. Further retransmissions are made according to the value of the Refresh Timer. Retransmissions continue until the condition is cleared.

5.2. Clearing a Fault Management Indication

When a fault is cleared, a node MUST cease sending the associated FM messages. Ceasing to send FM messages will clear the indication after 3.5 times the Refresh Timer. To clear an indication more quickly, the following procedure is used. The R-Flag of the FM message is set to one. Other fields of the FM message SHOULD NOT be modified. The message is sent immediately and then retransmitted two more times at an interval of one second. Note, however, if another fault occurs, the node MUST cease these retransmissions and generate new FM messages for the new fault.

5.3. Receiving a Fault Management Indication

When an FM message is received, a MEP examines it to ensure that it is well formed. If the message type is reserved or unknown, the message is ignored. If the version number is unknown, the message is ignored.

If the R-Flag is set to zero, the MEP checks to see if a condition matching the message type exists. If it does not, the condition specific to the message type is entered. An Expiration timer is set to 3.5 times the Refresh Timer. If the message type matches an existing condition, the message is considered a refresh and the Expiration timer is reset. In both cases, if an IF_ID TLV is present, it is recorded.

If the R-Flag is set to one, the MEP checks to see if a condition matching the message type and IF_ID exists. If it does, that condition is cleared. Otherwise, the message is ignored.

If the Expiration timer expires, the condition is cleared.

6. Minimum Implementation Requirements

At a minimum, an implementation MUST support the following:

1. Sending AIS and LKR messages at a rate of one per second.
2. Support of setting the L-Flag to indicate a server-failure.
3. Receiving AIS and LKR messages with any allowed Refresh Timer value.

The following items are OPTIONAL to implement.

1. Sending AIS and LKR messages with values of the Refresh Timer other than one second.
2. Support of receiving the L-Flag.
3. Support of setting the R-Flag to a value other than zero.
4. Support of receiving the R-Flag.
5. All TLVs.

7. Security Considerations

MPLS-TP is a subset of MPLS and so builds upon many of the aspects of the security model of MPLS. MPLS networks make the assumption that it is very hard to inject traffic into a network, and equally hard to cause traffic to be directed outside the network. The control-plane protocols utilize hop-by-hop security and assume a "chain-of-trust" model such that end-to-end control-plane security is not used. For more information on the generic aspects of MPLS security, see RFC 5920 [8].

This document describes a protocol carried in the G-ACh (RFC 5586 [4]) and so is dependent on the security of the G-ACh itself. The G-ACh is a generalization of the Associated Channel defined in RFC 4385 [6]. Thus, this document relies heavily on the security mechanisms provided for the Associated Channel as described in those two documents.

A specific concern for the G-ACh is that it can be used to provide a covert channel. This problem is wider than the scope of this document and does not need to be addressed here, but it should be noted that the channel provides end-to-end connectivity and SHOULD

NOT be policed by transit nodes. Thus, there is no simple way of preventing any traffic being carried in the G-ACh between consenting nodes.

A good discussion of the data-plane security of an Associated Channel may be found in RFC 5085 [9]. That document also describes some mitigation techniques.

It should be noted that the G-ACh is essentially connection-oriented, so injection or modification of control messages specified in this document requires the subversion of a transit node. Such subversion is generally considered hard to protect against in MPLS networks, and impossible to protect against at the protocol level. Management-level techniques are more appropriate.

Spurious fault OAM messages form a vector for a denial-of-service attack. However, since these messages are carried in a control channel, except for one case discussed below, one would have to gain access to a node providing the service in order to effect such an attack. Since transport networks are usually operated as a walled garden, such threats are less likely.

If external MPLS traffic is mapped to an LSP via a PHP forwarding operation, it is possible to insert a GAL followed by a fault OAM message. In such a situation, an operator SHOULD protect against this attack by filtering any fault OAM messages with the GAL at the top of the label stack.

8. IANA Considerations

8.1. Pseudowire Associated Channel Type

Fault OAM requires a unique Associated Channel Type that has been assigned by IANA from the Pseudowire Associated Channel Types registry.

Registry:

Value	Description	TLV Follows	Reference
0x0058	Fault OAM	No	(This Document)

8.2. MPLS Fault OAM Message Type Registry

This section details the "MPLS Fault OAM Message Type Registry", a new sub-registry of the "Multiprotocol Label Switching (MPLS) Operations, Administration, and Management (OAM) Parameters" registry. The Type space is divided into assignment ranges; the

following terms are used in describing the procedures by which IANA allocates values (as defined in RFC 5226 [7]): "Standards Action" and "Experimental Use".

MPLS Fault OAM Message Types take values in the range 0-255. Assignments in the range 0-251 are via Standards Action; values in the range 252-255 are for Experimental Use and MUST NOT be allocated.

Message Types defined in this document are:

Msg Type	Description
-----	-----
0	Reserved (not available for allocation)
1	Alarm Indication Signal (AIS)
2	Lock Report (LKR)

8.3. MPLS Fault OAM Flag Registry

This section details the "MPLS Fault OAM Flag Registry", a new sub-registry of the "Multiprotocol Label Switching (MPLS) Operations, Administration, and Management (OAM) Parameters" registry. The Flag space ranges from 0-7. All flags are allocated by "Standards Action" (as defined in RFC 5226 [7]).

Flags defined in this document are:

Bit	Hex Value	Description
---	-----	-----
0-5		Unassigned
6	0x2	L-Flag
7	0x1	R-Flag

8.4. MPLS Fault OAM TLV Registry

This sections details the "MPLS Fault OAM TLV Registry", a new sub-registry of the "Multiprotocol Label Switching (MPLS) Operations, Administration, and Management (OAM) Parameters" registry. The Type space is divided into assignment ranges; the following terms are used in describing the procedures by which IANA allocates values (as defined in RFC 5226 [7]): "Standards Action", "Specification Required", and "Experimental Use".

MPLS Fault OAM TLVs take values in the range 0-255. Assignments in the range 0-191 are via Standards Action; assignments in the range 192-247 are made via "Specification Required"; values in the range 248-255 are for Experimental Use and MUST NOT be allocated.

TLVs defined in this document are:

Value	TLV Name
-----	-----
0	Reserved (not available for allocation)
1	Interface Identifier TLV
2	Global Identifier

9. References

9.1. Normative References

- [1] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.
- [2] Vigoureux, M., Ed., Ward, D., Ed., and M. Betts, Ed., "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, May 2010.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [4] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [5] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.
- [6] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [7] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

9.2. Informative References

- [8] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [9] Nadeau, T., Ed., and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.

10. Contributing Authors

Stewart Bryant
Cisco Systems, Inc.
250, Longwater
Green Park, Reading RG2 6GB
UK

EMail: stbryant@cisco.com

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

EMail: msiva@cisco.com

Authors' Addresses

George Swallow (editor)
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, Massachusetts 01719
United States

EMail: swallow@cisco.com

Annamaria Fulignoli (editor)
Ericsson
Via Moruzzi
Pisa 56100
Italy

EMail: annamaria.fulignoli@ericsson.com

Martin Vigoureux (editor)
Alcatel-Lucent
Route de Villejust
Nozay 91620
France

EMail: martin.vigoureux@alcatel-lucent.com

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way
San Jose, California 95134
USA

EMail: sboutros@cisco.com

David Ward
Juniper Networks, Inc.

EMail: dward@juniper.net

