



INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.1711

(02/2004)

SERIES Y: GLOBAL INFORMATION
INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS
AND NEXT GENERATION NETWORKS

Internet protocol aspects – Operation, administration and
maintenance

Operation & Maintenance mechanism for MPLS networks

ITU-T Recommendation Y.1711

ITU-T Y-SERIES RECOMMENDATIONS

**GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT
GENERATION NETWORKS**

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation Y.1711

Operation & Maintenance mechanism for MPLS networks

Summary

This Recommendation provides mechanisms for user-plane OAM (Operation and Maintenance) functionality in MPLS networks according to the requirements and principles given in ITU-T Rec. Y.1710. This Recommendation is designed primarily to support point-to-point and multipoint-to-point explicit routed LSPs (ER-LSPs) with limited applicability to LSPs that employ penultimate hop popping (PHP).

The OAM mechanisms defined in this Recommendation assume common forwarding of the LSP payload and Y.1711 PDUs. In some situations this may not be true, such as when the LSP payload is load balanced across a plurality of parallel paths while still appearing as a single trail to the ingress and egress. LSRs introducing variations in connectivity are responsible for ensuring that the availability behaviour of Y.1711 per ingress-egress pair is preserved.

Source

ITU-T Recommendation Y.1711 was approved on 12 February 2004 by ITU-T Study Group 13 (2001-2004) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

© ITU 2004

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Page

1	Scope	1
2	References.....	1
2.1	Normative references.....	1
2.2	Informative references.....	2
3	Definitions	2
4	Symbols and abbreviations	3
5	Introduction	4
5.1	Overview of functionality.....	5
5.2	Identification of OAM packets from normal user-plane traffic	6
5.3	OAM payload	6
5.4	Handling of errored OAM packets	6
5.5	Engineering cost/risk considerations	7
5.6	Backward compatibility considerations	7
6	OAM mechanisms	7
6.1	Features common to all OAM packets	7
6.2	Connectivity Verification (CV).....	8
6.3	Fast Failure Detection (FFD).....	9
6.4	Forward Defect Indication (FDI).....	9
6.5	Backward Defect Indication (BDI)	11
6.6	On-demand diagnostic OAM mechanisms.....	11
6.7	Defect type codepoints	12
6.8	Defect type entry/exit criteria and consequent actions.....	13
7	Available and unavailable state processing	15
7.1	Short-breaks.....	16
7.2	Available/unavailable state definition	16
7.3	Near-end and far-end measurements of availability.....	17
7.4	Near-end state processing flow-chart	17
7.5	Far-end state processing flow-chart.....	20
7.6	A pictorial view of near-end and far-end state processing for a short-break and an unavailability event.....	23
8	Security aspects	23
	Appendix I – CV source and sink processing.....	25
	Appendix II – Indexing LSP Availability State Machine (ASM).....	25
	Appendix III – Different possible defect scenarios when using FFD OAM	26
	BIBLIOGRAPHY	27

... ..

ITU-T Recommendation Y.1711

Operation & Maintenance mechanism for MPLS networks

1 Scope

This Recommendation provides mechanisms for user-plane OAM (Operation and Maintenance) functionality in MPLS networks according to the requirements and principles given in ITU-T Rec. Y.1710. This Recommendation is designed primarily to support point-to-point and multipoint-to-point explicit routed LSPs (ER-LSPs) with limited applicability to LSPs that employ penultimate hop popping (PHP).

The OAM mechanisms defined in this Recommendation assume common forwarding of the LSP payload and Y.1711 PDUs. In some situations this may not be true, such as when the LSP payload is load balanced across a plurality of parallel paths while still appearing as a single trail to the ingress and egress. LSRs introducing variations in connectivity are responsible for ensuring that the availability behaviour of Y.1711 per ingress-egress pair is preserved.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

2.1 Normative references

- [1] ITU-T Recommendation I.610 (1999), *B-ISDN operation and maintenance principles and functions*.
- [2] ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks*.
- [3] ITU-T Recommendation M.20 (1992), *Maintenance philosophy for telecommunication networks*.
- [4] ITU-T Recommendation Y.1710 (2002), *Requirements for Operation & Maintenance functionality in MPLS networks*.
- [5] IETF RFC 3031 (2001), *Multiprotocol Label Switching Architecture*. Category: Standards Track.
- [6] IETF RFC 3032 (2001), *MPLS Label Stack Encoding*. Category: Standards Track.
- [7] IETF RFC 2373 (1998), *IP Version 6 Addressing Architecture*. Category: Standards Track.
- [8] IETF RFC 3270 (2002), *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*. Category: Standards Track.
- [9] IETF RFC 3209 (2001), *RSVP-TE: Extensions to RSVP for LSP Tunnels*. Category: Standards Track.
- [10] IETF RFC 3212 (2002), *Constraint-Based LSP Setup using LDP*. Category: Standards Track.

- [11] IETF RFC 3429 (2002), *Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions*.
- [12] ITU-T Recommendation Y.1712 (2004), *OAM functionality for ATM-MPLS interworking*.

2.2 Informative references

- [13] IETF RFC 1930 (1996), *Guidelines for creation, selection and registration of an Autonomous System (AS)*.

3 Definitions

This Recommendation introduces some terminology, which is required to discuss the functional network components associated with OAM. These definitions are consistent with G.805 terminology.

3.1 backward direction: The backward direction is opposite to the forward direction.

3.2 client/server (relationship between layer networks): [2] A term referring to the transparent transport of a client (i.e., higher) layer link connection by a server (i.e., lower) layer network trail.

3.3 defect: [3] Interruption of the capability of a transport entity (e.g., network connection) to transfer user or OAM information. This Recommendation defines seven types of defects (the detailed specifications of which are given later):

- **dLOCV:** Loss of Connectivity Verification defect.
- **dExcess:** Receiving excess rate of CV.
- **dPeerME:** Peer network maintenance entity defect.
- **dServer:** Server layer defect. Any server layer defect arising from a non-MPLS layer technology below the lowest MPLS layer network.
- **dTTSI_Mismatch:** Trail Termination Source Identifier Mismatch defect.
- **dTTSI_Mismerge:** Trail Termination Source Identifier Mismerge defect.
- **dUnknown:** Unknown defect in the MPLS network.

3.4 failure: [3] Termination of the capability of a transport entity to transfer user or OAM information. A failure can be caused by a persisting defect.

3.5 forward direction: The forward direction is the direction that traffic and the OAM PDUs directed to the near end trail termination packets flow on an LSP.

3.6 link connection: A partition of a layer N trail that exists between two logically adjacent switching points within the layer N network.

3.7 subnetwork: [2] A subnetwork is a contiguous topological region of a network delimited by its set of peripheral access points, and is characterized by the possible routing across the subnetwork between those access points. A network is the largest subnetwork and a node is the smallest subnetwork (at least in practical physical terms, though there are smaller subnetworks within nodes).

3.8 trail: [2] A generic transport entity at layer N which is composed of a payload field (which can carry a packet from a client higher layer N – 1 trail entity) with specific overhead added to ensure the forwarding integrity of the trail transport entity at layer N.

3.9 trail termination point: [2] A source or sink point of a trail at layer N, at which the trail overhead is added or removed respectively. A trail termination point must have a unique means of identification within the layer network.

3.10 user-plane: This refers to the set of traffic forwarding components through which traffic flows. CV OAM (or alternately FFD OAM) PDUs are periodically inserted into this traffic flow to monitor the health of those forwarding components. The user-plane is also sometimes called the data-plane (especially in IETF). Note that control-plane protocols (e.g., for signalling or routing) and management-plane protocols will require their own user-plane, and their user-plane may or may not be congruent (to varying degrees) with the traffic bearing user-plane.

4 Symbols and abbreviations

This Recommendation uses the following abbreviations:

AIS	Alarm Indication Signal
AS	Autonomous System
ASM	Availability State Machine
BDI	Backward Defect Indicator
BIP	Bit Interleaved Parity
CR-LDP	Constraint-based Routing Label Distribution Protocol
CV	Connectivity Verification
DL	Defect Location
DoS	Denial of Service
DT	Defect Type
E-LSP	EXP-Inferred-PSC LSP
FDI	Forward Defect Indicator
FFD	Fast Failure Detection
FRR	Fast Re-Route
LB	Loopback
LB-Req	Loopback Request
LB-Rsp	Loopback Response
LSP	Label Switched Path
L-LSP	Label-Only-Inferred-PSC LSP
LSR	Label Switching Router
MPLS	Multi-Protocol Label Switching
NMS	Network Management System
OAM	Operation and Maintenance
PHB	Per Hop Behaviour
PHP	Penultimate Hop Popping
PSC	PHB Scheduling Class
RSVP	Resource reSerVation Protocol
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement

TTL	Time to Live
TTSI	Trail Termination Source Identifier

5 Introduction

This Recommendation provides OAM techniques to meet the OAM requirements defined in ITU-T Rec. Y.1710. It is based on the following OAM packets that are defined in detail later, but in summary:

Connectivity Verification: Two connectivity verification probes are defined, CV and FFD. The role of CV is to instrument the LSP availability model described in clause 7. FFD is used for applications such as protection switching that require much faster response. Use of connectivity verification for applications requiring different insertion intervals is for further study.

CV: The CV flow is generated at the LSP's source LSR with a nominal frequency of 1/s and terminated at the LSP's sink LSR. The CV packet contains a network-unique identifier (TTSI) so that all types of defects can be detected.

FFD: FFD provides failure detection option for a P2P LSP independent of the CV based availability model and is not tied to the CV insertion rate. Insertion rates at 1/s or faster may also be used to instrument the availability model. Slower insertion rates will detect persistent MPLS layer defects but will not correctly instrument the availability model. Slower insertion with corresponding slower detection will impact monitored client layers as mismatches in time to detect defects that can result in additional alarms and possibly misleading alarm cycling and/or alarm persistency in the client layers. The effect of slower insertion rates on defect exit criteria (especially for mismerge and misbranch defects) is for further study.

The provisioning of FFD on an LSP will provide capability of fast failure detection (by default FFD OAM is not provisioned). It is recommended that FFD will be generated at the LSP ingress at a rate of 20 per second. This will provide failure detection on the LSP in order of 100s of milliseconds. When FFD is provisioned on an LSP CV PDU generation is disabled.

The FFD PDU contains the same information as CV. PDU processing at the egress is similar for CV and FFD with adjustments for the higher frequency of FFD PDU arrival and potential of misbranching and mismerging defects between LSPs using either FFD or CV.

It is recommended that FFD should only be used as the lowest layer LSP (Level 1), i.e., there should never be plain CV LSPs serving a FFD LSP. (Note the reason for this recommendation is that if there is a lower layer LSP running plain CV below the FFD LSP, then failures within or below the plain CV LSP may not send FDI upwards into the FFD LSP fast enough to stop the FFD thinking the true defect is in the FFD LSP. This is because it may take the plain CV LSP up to 3 s to detect a defect and insert FDI (at 1/s), whereas the FFD LSP will detect the defect in ≤ 150 ms and thus take action faster.

FDI: The FDI flow is generated in response to detecting defects (e.g., from the CV flow). Its primary purpose is to suppress alarms in layer networks above the level at which the defect occurs. It is generated at either:

- i) the LSR, which first detects a dServer/dUnknown defect; or
- ii) the LSP terminating LSR for all MPLS layer defects.

BDI: The BDI flow is injected on a return path (such as a return LSP) to inform the upstream LSR (which is the source of the forward LSP) that there is a defect at the downstream LSP's LSR sink point. BDI therefore tracks FDI in terms of its period of generation. BDI packets may be useful in 1:1/N instances of protection switching.

Performance: These are FFS. However, the intention is to have an on-demand method of determining packet/octet loss on an LSP in order to aid trouble-shooting. They are not intended to be used as a temporally permanent OAM function (unlike the CV flow), but they could be.

Note that path trace, performance monitoring and loopback functions are FFS.

BDI and loopback transactions use a return path. A return path could be:

- A dedicated return LSP;
- A shared return LSP, which is shared between many forward LSPs;
- A non-MPLS return path, such as an out-of-band IP path. This option has potential security issues. For example, the return path could be terminated on a different LSR interface, and potentially a malicious user could generate a BDI and send it to the ingress LSR. Therefore, due to the possibility of DoS attack, additional security measures must be taken. Operators should use the optional TTSI field in BDI packets in order to assure authentication of such packets so that the receivers of BDI OAM packets may verify that the sender of the packet is valid.

All OAM packets are identified within a LSP traffic stream by the use of globally well-known and reserved label codepoint (of 14). Further details on the encoding of OAM packets are provided later.

It is strongly recommended that CV OAM or FFD OAM packets be generated on each LSP (in order to detect all defects and potentially provide protection against traffic leakage both in and out of LSPs). It is also recommended that FDI OAM packets be used to suppress alarm storms. BDI packets are a useful tool for single-ended monitoring of both directions and also, in some protection-switching cases. However, these are only recommendations, and operators can choose to use some or all of the OAM packets as they see fit. Appendix I discusses some of the options for generating and processing CV flows.

OAM techniques are applied on a per LSP basis. If a segment of a given LSP at layer N is to be monitored for some reason (e.g., via a CV or P flow say), one way to do this is by creating a new server layer LSP (i.e., at layer N + 1) to cover the segment at layer N.

5.1 Overview of functionality

The OAM defect detection function is based on the periodic transmission of CV or FFD packets from ingress to egress of an LSP. CV packet generation rate is 1 packet per second whereas recommended packet generation rate for FFD packets is 20 per second. Each CV and FFD packet carries a unique TTSI (Trail Termination Source Identifier), which is composed of the source LSR identifier, and the LSP identifier.

An LSP enters a defect state when one of the defects noted in clause 3 occurs (these are defined in detail later in terms of precise entry/exit criteria and consequent actions).

In addition to the CV packet, there are other OAM packet types defined that provide consequential fault handling or performance monitoring functions. These will be defined later. All OAM packets are identified in terms of a function type by the first octet of the OAM packet payload as follows (see Table 1):