

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 11, 2014

Zafar Ali
George Swallow
Matt Hartley
Clarence Filsfils
Cisco Systems

Kenji Kumaki
KDDI Corporation
July 12, 2013

Extensions to Resource ReserVation Protocol-Traffic Engineering
(RSVP-TE) for Error Notication in Generalized Multiprotocol Label
Switching (GMPLS) User-Network Interface (UNI)

draft-ali-ccamp-gmpls-uni-error-notification-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the
provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering
Task Force (IETF). Note that other groups may also distribute
working documents as Internet-Drafts. The list of current Internet-
Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months
and may be updated, replaced, or obsoleted by other documents at any
time. It is inappropriate to use Internet-Drafts as reference
material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 11, 2014.

Copyright Notice

Copyright (c) 2013 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.

Abstract

There are many scenarios in which extensions to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) are required for error notification in Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI). This document outlines these scenarios and specifies the required extensions to RSVP-TE. Although the GMPLS-UNI reference model is used to describe requirements and solutions in the document, the proposed extensions are equally applicable to other deployment scenarios such as inter-domain RSVP-TE.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction.....	3
2. Requirements.....	4
2.1. Error Node Address in ERROR_SPEC object	4
2.2. Restoration Notification.....	4
3. RSVP-TE extensions for Error Notification.....	5
3.1. Error Node Address Modification in ERROR_SPEC object ..	5
3.1.1. Error Node Address Modification Flag	5
3.1.2. Processing rules.....	5
3.1.3. Example.....	6
3.2. Restoration Notification.....	6
3.2.1. Error sub-code.....	6
3.2.2. Processing rules.....	6
4. Security Considerations.....	6
5. IANA Considerations.....	6
5.1. New ERROR_SPEC.Flags.....	6

5.2. New RSVP error sub-codes.....	7
6. Acknowledgements.....	7
7. References.....	7
7.1. Normative References.....	7
7.2. Informative References.....	8

1. Introduction

[RFC4208] provides mechanisms for GMPLS UNI signaling. Figure 1 borrows the reference network model of [RFC4208].

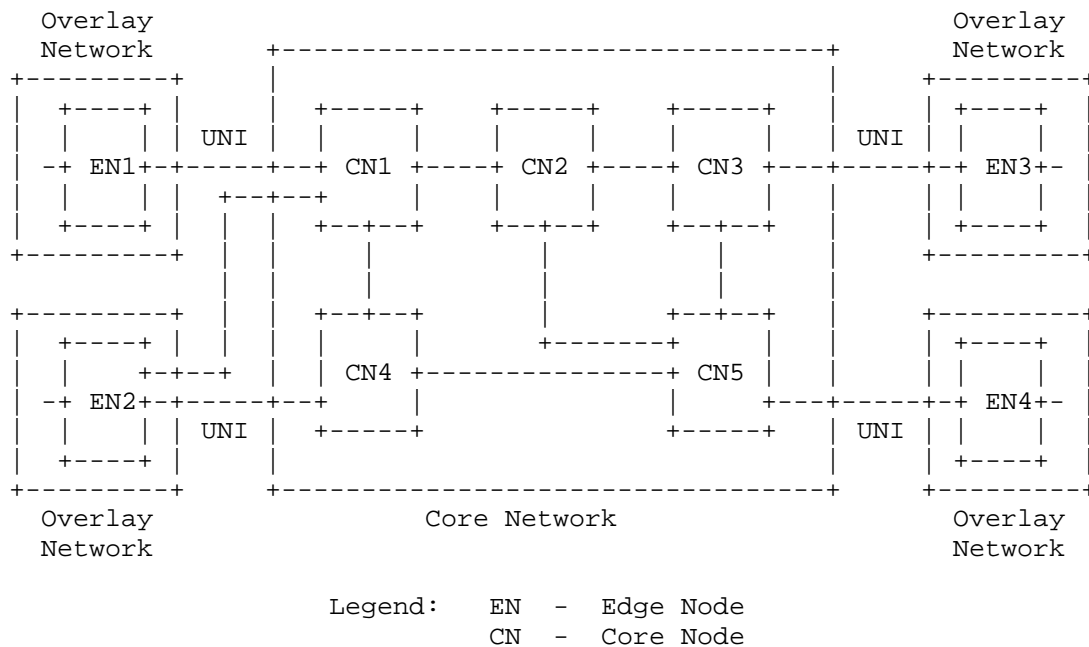


Figure 1: GMPLS UNI Reference Model

For convenience, some terms used in [RFC4208] are re-stated below:

- "source EN": the edge node which initiates the connection (e.g., EN1);
- "destination EN": the edge node where the connection is terminated (e.g., EN3);
- "ingress CN": the core node to which the source EN is attached

(e.g., CN1);

- "egress CN": the core node to which the destination EN is attached (e.g., CN3).

[RFC4208] provides mechanisms for UNI signaling whereby a single end-to-end RSVP session between source EN and destination EN is used for the user connection, just as it would be for connection creation between two core nodes. However, when considering policy considerations such as a requirement to preserve the confidentiality of topological information of the core network, additional requirements for UNI signaling exist that are not addressed by [RFC4208]. This document focuses on error notification aspects of these additional requirements.

2. Requirements

This sections outlines addition requirements related to error notification in GMPLS UNI. For the purpose of illustration, an end-to-end UNI connection that passes through EN1-CN1-CN2-CN3-EN3 is used as an example.

2.1. Error Node Address in ERROR_SPEC object

The IPv4 and IPv6 ERROR_SPEC objects are defined in [RFC2205]; both objects contain an Error Node Address of the appropriate type, defined as the IP address of the error originating node [RFC2205]. However, for confidentiality reasons, service provider of the core network may not wish to expose addresses used in the core network (other than the UNI link addresses) to edge nodes. To meet this requirement, a core node should be allowed to change the Error Node Address in the ERROR_SPEC. When such an address modification is made by a core node, the edge node should be informed that Error Node Address field in the ERROR_SPEC has been modified.

2.2. Restoration Notification

The ability to dynamically restore a Label Switched Path (LSP) is one of the fundamental features of GMPLS, including GMPLS UNI. In the context of GMPLS UNI, restoration of an LSP after a failure may be performed by the core network alone, or may be triggered by the edge nodes. There is a requirement that the two methods of restoration co-exist.

When the core network restores an LSP, in many cases there is no need to re-signal the LSP. However, in order to avoid a concurrent restoration process triggered by an edge node, it is required that the core network notify the edge network that the LSP has been restored.

3. RSVP-TE extensions for Error Notification

3.1. Error Node Address Modification in ERROR_SPEC object

3.1.1. Error Node Address Modification Flag

To allow a core node to change the Error Node Address in the ERROR_SPEC object, the following new flag value is defined for the Flags field of the IPv4 and IPv6 ERROR_SPEC objects [RFC2205], and for the IPv4 IF_ID and IPv6 IF_ID objects [RFC3477].

ERROR_SPEC.Flags = 0x08 (value to be assigned by IANA): Error Node Address Modified.

The "Error Node Address Modified" flag is applicable to all RSVP-TE messages that use the ERROR_SPEC object.

3.1.2. Processing rules

When processing an ERROR_SPEC object received in a message from another node, the processing node MAY replace the IPv4 or IPv6 address in the ERROR_SPEC object with another address relating to itself if its local policy requires it to do so.

When making an address substitution in this manner, the processing node SHOULD set the Error Node Address Modified flag in the Flags field of the ERROR_SPEC object to indicate that a change has been made.

When processing a received ERROR_SPEC object, a node SHOULD examine the ERROR_SPEC.Flags field and check for the "Error Node Address Modified" flag before processing the ERROR_SPEC.Error_Node_Address field. The processing node MAY alter its handling of the ERROR_SPEC object if this flag is set, but any such variation in handling is implementation-dependent and beyond the scope of this document.

3.1.3. Example

To illustrate usage of this flag, consider an example connection EN1-CN1-CN2-CN3-EN3 in the network specified by figure 1. In this example, the CN2 detects a failure and sends a PathErr message towards EN1, using a local address associated with the failed resource in the ERROR_SPEC.

However, CN1's local policy is to conceal the topology of the core network from edge nodes. When CN1 processes the PathErr message, it changes the address in the ERROR_SPEC to CN1's address of the EN1-CN1 UNI link. It also sets the Error Node Address Modified flag in the Flags field of the ERROR_SPEC to indicate that a change has been made. The PathErr message is then sent to EN1.

3.2. Restoration Notification

3.2.1. Error sub-code

In order to satisfy restoration notification requirement mentioned above, a new path error sub-code "LSP restored" (value to be assigned by IANA; suggested value: 16) for the error code "Notify Error" (25) is defined.

3.2.2. Processing rules

When a core node restores an existing LSP after a failure, it SHOULD send a PathErr message with the error code "Notify Error" (25) and error sub-code "LSP restored" (suggested: 16) for the LSP.

When an edge node receives a PathErr message with the error code "Notify Error" (25) and error sub-code "LSP restored" (suggested: 16), it MAY avoid triggering any other restoration process.

4. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473] and [RFC4874].

5. IANA Considerations

5.1. New ERROR_SPEC.Flags

This document defines the following new flag value for the Flags field of the IPv4 and IPv6 ERROR_SPEC object defined in [RFC2205].

ERROR_SPEC.Flags: Error Node Address Modified. Value is to be assigned by IANA (suggested value: 0x08).

5.2. New RSVP error sub-codes

IANA registry: RSVP PARAMETERS
Subsection: Error Codes and Globally-Defined Error Value Sub-Codes

For Error Code "Notify Error" (25) (see [RFC3209]) the following sub-code is defined.

Sub-code -----	Value -----
LSP Restored	To be assigned by IANA. Suggested value: 16

6. Acknowledgements

TBD.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

7.2. Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter,
"Generalized Multiprotocol Label Switching (GMPLS)
User-Network Interface (UNI): Resource ReserVation
Protocol-Traffic Engineering (RSVP-TE) Support for the
Overlay Model", RFC 4208, October 2005.
- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol
(RSVP) -- Version 1 Message Processing Rules", RFC
2209, September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS
Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: August 17, 2013

Zafar Ali
George Swallow
Clarence Filsfils
Ori Gerstel
Matt Hartley
Cisco Systems
February 18, 2013

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
Extension for Label Switched Path (LSP) Inquiry

draft-ali-ccamp-lsp-inquiry-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 15, 2013.

Copyright Notice

Copyright (c) 2012 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow

modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

RSVP-TE reoptimization procedure for Packet Switch Capable (PSC) tunnels and non-PSC tunnels has some differences. This document highlights these differences, describes how existing procedures can be used for reoptimization of non-PSC tunnels and proposes some enhancements for reoptimization of non-PSC tunnels.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction	2
1.1. Inquiry with Resource Locking.....	4
1.2. Inquiry without Resource Locking	5
2. RSVP-TE signaling procedure	6
2.1. RSVP-TE Signaling for Inquiry with Resource Locking	6
2.2. RSVP-TE Signaling for Inquiry without Resource Locking	7
3. Security Considerations	8
4. IANA Considerations	8
4.1. RSVP Attribute Bit Flags	8
5. References	9
5.1. Normative References	9
5.2. Informative References	9

1. Introduction

During reoptimization of PSC tunnel, existing and reoptimization LPSS may use independent labels and may install independent

forwarding states for each LSP. That is, existing and reoptimization LSPs may be simultaneously active in both control and data planes. However, for many non-PSC technologies label itself represents the underlying physical resource and hence cannot be shared between existing and reoptimization LSPs in the data plane. Consequently, for many non-PSC technologies, reoptimization LSPs can only be instantiated in the control plane. Furthermore, reoptimization LSP is not immediately available to carry any traffic and requires additional signaling for its activation in the data plane.

In this document, inquiry refers to a way to signal sharing of resources (e.g., labels, links) between the existing and reoptimization LSPs in the control plane without installing any forwarding states in the data plane (e.g., without installing cross-connections). In most non-PSC networks, inquiry step is required to access feasibility and characteristics of the reoptimization LSP before committing data plane resources and moving traffic to it. This is especially true of scenarios when route computation is not performed by ingress Label Edge Router (LER). These include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), e.g. inter-domain LSPs.
- . Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI) where route computation may be performed by the (server layer) core Label Switch Router (LSR) [RFC4208];

In such cases, ingress LER may like to inquire about feasibility and attributes of a better path. Cost, TE metrics and SRLG values are examples of attributes that ingress LER may like to inquire about (e.g., before making a reoptimization decision). Procedures specified in [ID.SRLG-RECORDING] and [ID.METRIC-RECORDING] may be used for this purpose.

Reoptimization is an example where inquiry procedure is needed. However, inquiry is also useful for other use cases, e.g., for probing purposes. Hence, for the generalization purposes, LSP signaled during inquiry is referred as inquiry LSP. Reoptimization LSP is an example of inquiry LSP.

Inquiry LSP may be signaled with or without resource locking in the control plane. This is detailed in the following.

1.1. Inquiry with Resource Locking

Signaling inquiry LSP with resource locking is depicted in Figure 1. Specifically, Path message (Path1) is signaled such that resource activation is Data Plane (DP) is skipped. However, inquiry LSP reserves resources in the Control Plane (CP), i.e., resources/ labels are locked/ committed in the control plane. Nonetheless, resources/ labels are still shared with the existing LSP(s) belonging to the tunnel inquiry LSP belongs to.

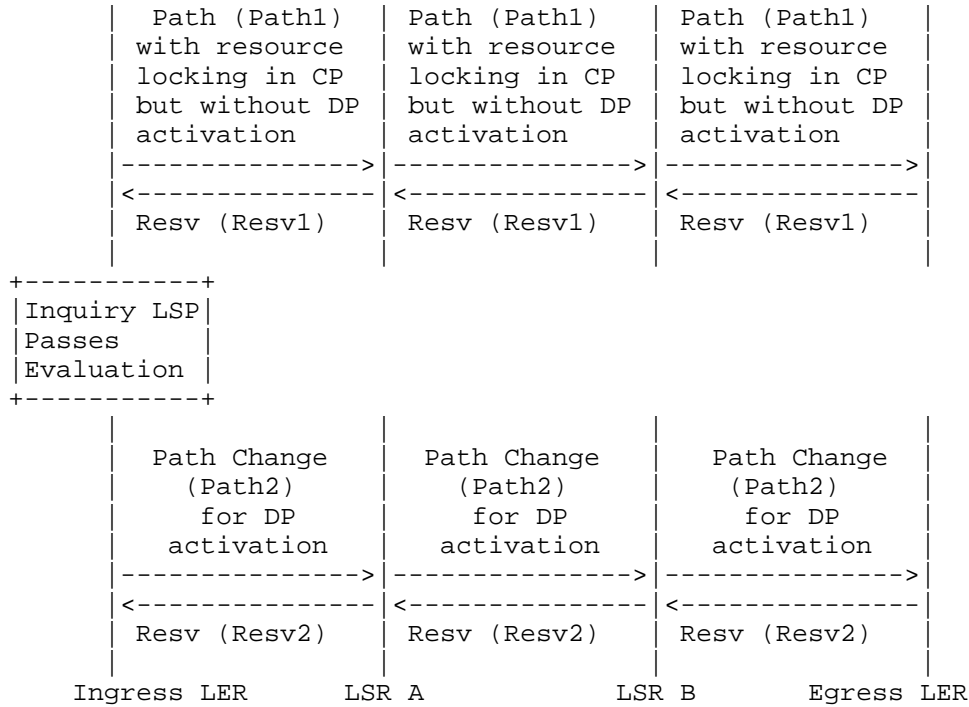


Figure 1: Inquiry LSP Signaling with Resource Locking

By the time Resv1 for the initial Path message (Path1) is received by the ingress LER, the ingress LER knows about feasibility and the requested attributes of the inquiry LSP. Please note that to ensure resource locking at the right priority, inquiry LSP needs to be signaled using the same setup and hold priority as existing LSP.

After finding feasibility and the requested attributes of the inquiry LSP, the ingress LER evaluates inquiry LSP. E.g., if the inquiry LSP is signaled for reoptimization purposes, the ingress LER

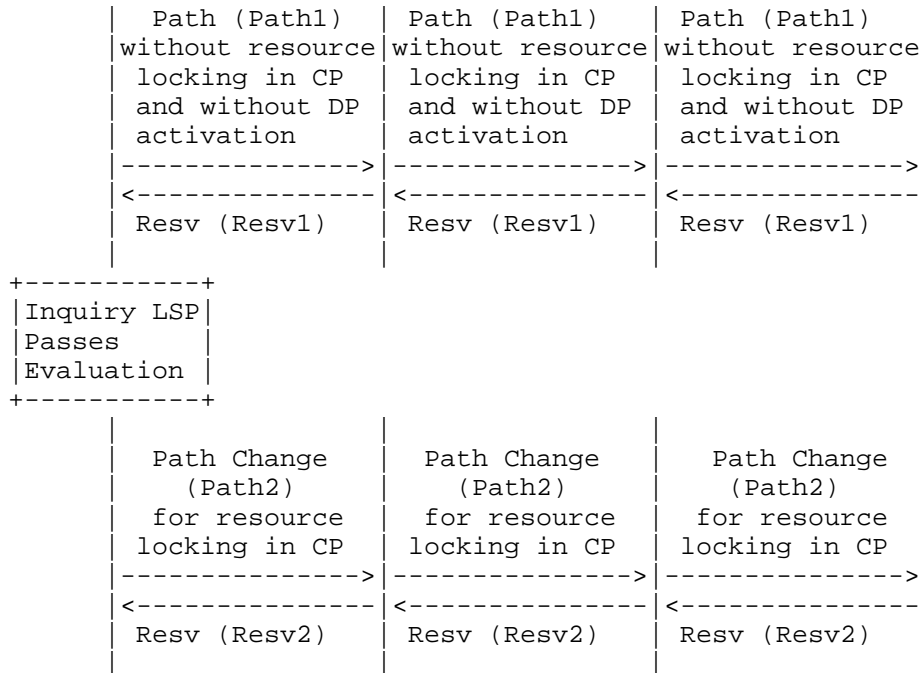
determines if it would like to reoptimize the existing LSP. If ingress LER decides not to reoptimize existing LSP, it deletes the inquiry LSP (by sending Path tear message -

- this option is not shown

in Figure 1). However, if the ingress LER decides to reoptimize the tunnel to use the inquire LSP, it initiates activation of the inquiry LSP in the data plane (as shown by Path2 and Resv2 signaling loop in Figure 1). How ingress LER moves traffic from existing LSP to inquire LSP for reoptimization purpose is beyond the scope of this document.

1.2. Inquiry without Resource Locking

When inquiry is performed with resource locking, the resources used by the inquiry LSP are locked in control plane and cannot be used by any LSP other than the LSP(s) belonging to the same tunnel (of course resource preemption based on setup and hold priority of the inquiry LSP is still possible). This limits network availability in the event of a failure, especially when inquiry is used frequently, e.g., for probing purposes. This issue can be addressed by signaling inquiry LSP without resource locking. In this case, the resources used by the inquiry LSP are not locked in control plane and can be used by any LSP, including the existing LSP(s) belonging to the same tunnel. Therefore, if inquiry LSP is signaled without resource locking, additional signaling is required to first lock the resources in the control plane, before the LSP can be activated in the data plane. This is depicted in the following Figure.



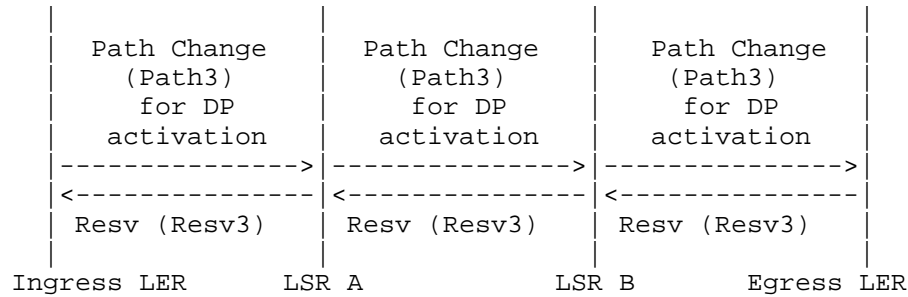


Figure 2: Inquiry LSP Signaling without Resource Locking

Initial Path message (Path1) is signaled such that resource locking in control plane and resource activation in data plane is skipped. By the time Resv1 for the initial Path message (Path1) is received by the ingress LER, the ingress LER knows about feasibility and the requested attributes of the inquiry LSP. The inquiry LSP evaluation process works in the same fashion as discussed in Section 1.1.

As Path1 is signaled without resource locking, there is no guarantee that the resources for the inquiry LSP will be available when ingress LER decides to activate the inquiry LSP. Therefore, the ingress LER first signals a path change (Path2) to lock the resource in the control plane before inquiry LSP can be activated in the data plane.

2. RSVP-TE signaling procedure

This section describes the signaling procedure for LSP inquiry with and without resource locking.

2.1. RSVP-TE Signaling for Inquiry with Resource Locking

[RFC6001] specifies procedure for signaling Soft Forwarding Adjacency (Soft FA). It defines the Pre-Planned LSP flag in the Attribute Flags TLV, which can be carried in an LSP_ATTRIBUTES object defined in [RFC5420]. The Pre-Planned LSP flag can also be used to signal inquiry LSP with resource locking.

The processing rules for the Pre-Planned LSP flag are unchanged from [RFC6001]. The procedures for the processing the Attribute Flags TLV are also unchanged from [RFC5420]. The following description is provided to help describe usage of the Pre-Planned LSP flag in the context of signaling the inquiry LSP. Example of Figure 1 is used to describe the usage.

- . Path1 message in Figure 1 is sent with the Pre-Planned LSP flag set to 1. As per [RFC6001], this enables provisioning of inquiry LSP in control plane only. In order to enable sharing of resources within the same tunnel, Path1 message is sent with shared explicit reservation style [RFC3209].
- . In order to activate inquiry LSP in the data plane, Path2 message (please see Figure 1) is sent with the Pre-Planned LSP flag set to 0.

2.2. RSVP-TE Signaling for Inquiry without Resource Locking

In order to indicate signaling of an LSP without resource provisioning in neither control plane nor data plane, a new flag in the Attribute Flags TLV, which can be carried in an LSP_ATTRIBUTES Object, is defined as follows:

- . Pre-Planned LSP Without Resource Locking flag (to be assigned by IANA, recommended bit position TBD)

The Pre-Planned LSP Without Resource Locking flag is meaningful on a Path message. The procedure for the processing the Attribute Flags TLV follows [RFC5420]. The flag is used as described in the following.

- . If the Pre-Planned LSP Without Resource Locking flag is set to 1, the transit nodes SHOULD NOT reserve resources required by the LSP in the control plane and MUST NOT install any forwarding states associated with the LSP in the data plane. However, all LERs and LSRs are required to remember the resource (labels, links, etc.) assignments and the RSVP-TE states associated with this LSP. These resources are not locked and hence can be claimed by another LSP. If the Pre-Planned LSP Without Resource Locking flag is set to 1, the Pre-Planned LSP flag MUST be ignored. In our example of Figure 2, Path1 message is sent with the Pre-Planned LSP Without Resource Locking flag set to 1.
- . If the Pre-Planned LSP Without Resource Locking flag is set to 0 and the Pre-Planned LSP flag is set to 1, the transit nodes MUST commit resources associated with the LSP in the control plane. However, if a resource is already claimed by another LSP, the processing LSR/ LER MUST send a Path Error with code: "Admission Control Failure (1)" and subcode "LSP Admission Failure (4)" and Path State Removal flag. A

processing LSR/ LER MUST NOT install any forwarding states associated with the LSP in the data plane. Path2 message in Figure 2 is sent with the Pre-Planned LSP Without Resource Locking flag set to 0 and the Pre-Planned LSP flag is set to 1. It is RECOMMENDED that no other modifications be made to other RSVP objects during this operation (signaling Path2).

- . Activation of LSP in data plane follows the procedure specific in [RFC6001]. E.g., Path2 message in Figure 2 is sent with the Pre-Planned LSP flag set to 0.

3. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], and [RFC3473] and [RFC4874].

4. IANA Considerations

4.1. RSVP Attribute Bit Flags

The IANA has created a registry and manages the space of attributes bit flags of Attribute Flags TLV as described in section 11.3 of [RFC5420]. It is requested that the IANA makes assignments from the Attribute Bit Flags defined in this document.

This document introduces the following new Attribute Bit Flag:

- Bit number: TBD
- Defining RFC: this I-D
- Name of bit: Pre-Planned LSP Without Resource Locking Flag

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiimoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.
- [RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

5.2. Informative References

- [ID.SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria, "RSVP-TE Extensions for Collecting SRLG Information", draft-ietf-ccamp-rsvp-te-srlg-collect.txt, work in progress.
- [ID.TE-METRIC-RECORDING] Ali, Z., Swallow, G., Filsfils, C., et al "RSVP-TE extension for recording TE Metric of a Label Switched Path", draft-ali-ccamp-te-metric-recording, work in progress.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

Internet Draft draft-ali-ccamp-lsp-inquiry-00.txt

[RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", RFC 2209, September 1997.

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Ori Gerstel
Cisco Systems
Email: ogerstel@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 14, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Luyuan Fang
Cisco Systems
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG
Daniele Ceccarelli
Ericsson
July 15, 2013

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
extension for signaling Objective Function and Metric Bound
draft-ali-ccamp-rc-objective-function-metric-bound-03.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the
provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering
Task Force (IETF). Note that other groups may also distribute
working documents as Internet-Drafts. The list of current
Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six
months and may be updated, replaced, or obsoleted by other
documents at any time. It is inappropriate to use Internet-Drafts
as reference material or to cite them other than as "work in
progress."

This Internet-Draft will expire on January 14, 2014.

Copyright Notice

Copyright (c) 2013 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

In particular networks such as those used by financial institutions, network performance criteria such as latency are becoming as critical to data path selection. However cost is still an important consideration. This leads to a situation where path calculation involves multiple metrics and more complex objective functions.

When using GMPLS control plane, there are many scenarios in which a node may need to request a remote node to perform path computation or expansion, like for example multi-domain LSP setup, Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI) or simply the utilization of a loose ERO in intra domain signaling. In such cases, the node requesting for the setup of an LSP needs to convey the required objective function to the remote node, to enable it to perform route computation in the desired fashion. Similarly, there are cases the ingress needs to indicate a TE metric bound for a loose segment that is expanded by a remote node.

This document defines extensions to the RSVP-TE Protocol to allow an ingress node to request the required objective function for the route computation, as well as a metric bound to influence route computation decisions at a remote node(s).

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

Copyright Notice.....	1
1. Introduction.....	3
2. RSVP-TE signaling extensions.....	4
2.1. Objective Function (OF) Subobject.....	4
2.1.1. Minimum TE Metric Cost Path Objective Function.....	6
2.1.2. Minimum IGP Metric Cost Path Objective Function.....	6
2.1.3. Minimum Latency Path Objective Function.....	6
2.1.4. Minimum Latency Variation Path Objective Function....	7
2.2. Metric subobject.....	7
2.3. Processing Rules for the OF Subobjects.....	8
2.4. Processing Rules for the Metric subobject.....	9
3. Security Considerations.....	11
4. IANA Considerations.....	11
5. Acknowledgments.....	12
6. References.....	12
6.1. Normative References.....	12
6.2. Informative References.....	12

1. Introduction

As noted in [OSPF-TE-METRIC] and [ISIS-TE-METRIC], in certain networks such as financial information networks (e.g. stock market data providers), performance criteria such as latency are becoming as critical to data path selection along with other metrics. Such networks may require selection of a path that minimizes end-to-end latency. Or a path may need to be found that minimizes some other TE metric, but subject a latency bound. Thus there is a requirement to be able to find end-to-end paths with different optimization criteria.

When the entire route for an LSP is computed at the ingress node, this requirement can be met by a local decision at that node. However, there are scenarios where partial or full route computations are performed by remote nodes. The scenarios include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), including intra-domain LSPs.
- . GMPLS-UNI where route computation may be performed by the UNI-Network (server) node [RFC 4208];
- . Multi domain LSP setup with per domain path computation;

In these scenarios, there is a need for the ingress node to convey the optimization criteria including the TE metrics (e.g., IGP metric, TE metric, hop counts, latency, etc.) to be used for the path computation to the node performing route computation or expansion. Similarly, there is a need for the ingress node to indicate a TE metric bound for the loose segment being expanded by a remote node.

[RFC5541] defines extensions to the Path Computation Element communication Protocol (PCEP) to allow a Path Computation Client (PCC) indicate in a path computation request the desired objective function. [RFC5440] defines extension to the PCEP to allow a PCC indicate in a path computation request a bound on given TE metric(s). This draft defines similar mechanisms for the RSVP-TE protocol allowing an ingress node to indicate in a Path request the desired objective function along with any associated TE metric bound(s). The nodes performing route expansion use this information to find the "best" candidate route.

2. RSVP-TE signaling extensions

This section defines RSVP-TE signaling extensions required to address the above-mentioned requirements. Two new ERO subobject types, Objective Function (OF) and Metric, are defined for this purpose. Their purpose is as follows.

- . OF subobject conveys a set of one or more specific optimization criteria that needs be followed in expanding route of a TE-LSP in MultiProtocol Label Switching (MPLS) and GMPLS networks.
- . Metric subobject indicates the bound on the path metric that needs to be observed for the loose segment to be considered as acceptable by the ingress node.

The scope of the Metric and OF subobjects is the node performing the expansion for loose ERO and the subsequent ERO subobject that identifies an abstract node. The following subsection provides the details.

2.1. Objective Function (OF) Subobject

A new ERO subobject type Objective Function (OF) is defined in order for the ingress node to indicate the required objective function on a loose hop. The ERO subobject type OF is optional. It MAY be carried within an ERO object of RSVP-TE Path message and its scope is limited to previous ERO subobject that

identifies an abstract node. For more details please refer to the Processing Rules for the OF Subobjects section.

The OF subobject 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																		
+-----+										+-----+										+-----+										+-----+																			
L										Type										Length										OF Code										Reserved									
+-----+										+-----+										+-----+										+-----+																			
										Reserved																																							
+-----+										+-----+										+-----+										+-----+																			

The fields of OF subobject are defined as follows:

L bit: The L bit MUST be set to represent a loose hop in the explicit route.

Type: The Type is to be assigned by IANA (suggested value: 66).

Length: The Length contains the total length of the subobject in bytes, including the Type field, the Length field and the length of the optional TLV(s). When there is no optional TLV, the Length is 4.

OF Code (1 byte): The identifier of the objective function. The following OF code values are suggested. These values are to be assigned by IANA.

- * OF code value 0 is reserved.

- * OF code value 1 (to be assigned by IANA) is for Minimum TE Metric Cost Path (MTMCP) OF defined in this document. See definition of MTCP OF in the following.

- * OF code value 2 (to be assigned by IANA) is for Minimum Interior Gateway Protocol (IGP) Metric Cost Path (MIMCP) OF defined in the following.

- * OF code value 3 (to be assigned by IANA) is for Minimum Load Path (MLP) OF as defined in RFC5541.

- * OF code value 4 (to be assigned by IANA) is for Maximum Residual Bandwidth Path (MBP) OF as defined in RFC5541.

- * OF code value 5 (to be assigned by IANA) is for Minimize Aggregate Bandwidth Consumption (MBC) OF as defined in RFC5541.

* OF code value 6 (to be assigned by IANA) is for Minimize the Load of the most loaded Link (MLL) OF as defined in RFC5541.

* OF code value 7 is skipped (to keep the objective function code values consistent between [RFC5541] and this draft.

* OF code value 8 (to be assigned by IANA) is for Minimum Latency Path (MLP) OF defined in this document. See definition of MLP OF in the following.

* OF code value 9 (to be assigned by IANA) is for Minimum Latency Variation Path (MLVP) OF defined in this document. See definition of MLVP OF in the following.

Other objective functions may be defined in future.

Reserved (5 bytes): This field MUST be set to zero on transmission and MUST be ignored on receipt.

2.1.1. Minimum TE Metric Cost Path Objective Function

Minimum TE Metric Cost Path (MTMCP) OF is defined as an Objective Function where a path is computed such that the sum of the TE metric of the links along the path is minimized. In the context of loose hop expansion, the ERO expanding node MUST try to find a route such that the sum of the TE metric of the links along the route is minimized.

2.1.2. Minimum IGP Metric Cost Path Objective Function

Minimum IGP Metric Cost Path (MIMCP) OF is defined as an Objective Function where a path is computed such that the sum of the IGP metric of the links along the path is minimized. In the context of loose hop expansion, the ERO expanding node MUST try to find a route such that the sum of the IGP metric of the links along the route is minimized.

2.1.3. Minimum Latency Path Objective Function

Minimum Latency Path (MLP) OF is defined as an Objective Function where a path is computed such that latency of the path is minimized. In the context of loose hop expansion, the ERO expanding node MUST try to find a route such that overall latency of the loose hop is minimized.

2.1.4. Minimum Latency Variation Path Objective Function

Minimum Latency Variation Path (MLVP) OF is defined as an Objective Function where a path is computed such that latency variation in the path is minimized. In the context of loose hop expansion, the ERO expanding node MUST try to find a route such that overall latency variation of the loose hop is minimized.

2.2. Metric Bound subobject

The ERO subobject type Metric Bound (MB) is optional. It MAY be carried within an ERO object of RSVP-TE Path message and its scope is limited to previous ERO subobject that identifies an abstract node. It is possible to identify different Metric Bound subobjects for different hops of the ERO to be expanded. For more details please refer to the Processing Rules for the Metric Bound Subobjects section.

This subobject 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
+-----+-----+-----+-----+-----+-----+-----+-----+
|L|      Type      |      Length      | metric-type |B|      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     metric-bound                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The fields of the Metric subobject are defined as follows:

L bit: The L bit is set if the subobject represents a loose hop in the explicit route. If the bit is not set, the subobject represents a strict hop in the explicit route. Please note that use of MB subobject is also applicable to strict hops, e.g., in selecting a component link within a heterogeneous bundled TE link.

Type: The Type is to be assigned by IANA (suggested value: 67).

Length: The Length is 8.

Metric-type (8 bits): Specifies the metric type associated with the partial route expended by the node processing the loose ERO. The following values are currently defined:

- * T=1: cumulative IGP cost
- * T=2: cumulative TE cost
- * T=3: Hop Counts
- * T=4: Cumulative Latency
- * T=5: Cumulative Latency Variation

B bit: Best-effort bit. When the best-effort (B) bit is set, it means that the ingress allows for the set up of an LSP that does not meeting the MB requirement. When the best-effort (B) bit is not set, it means that the MB needs to be observed.

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

Metric-bound (32 bits): The metric-bound indicates an upper bound for the path metric that MUST NOT be exceeded for the ERO expending node to consider the computed path as acceptable. The metric bound is encoded in 32 bits using IEEE floating point format as defined in [IEEE.754.1985]). When it indicates a time value (i.e. Latency or Latency Variation) it is expressed in milliseconds.

2.3. General Processing rules

A single OF subobjects SHOULD be used for each ERO object. Multiple Metric Bound subobjects MAY be indicated for each hop to be expanded and MUST be placed after each abstract node subobject. Different Metric Bounds MAY be identified for each hop expansion.

2.3.1. Processing Rules for the OF Subobjects

The basic processing rules of an ERO are not altered. Please refer to [RFC3209] for details.

The scope of the OF subobject is the previous ERO subobject that identifies an abstract node, and the subsequent ERO subobject that identifies an abstract node. Multiple OF subobjects may be present between any pair of abstract nodes. However, only first OF subobject is analyzed and others are ignored.

The following conditions SHOULD result in Path Error with error code `'Routing Problem'` and error subcode `"Bad EXPLICIT_ROUTE object"`:

- . If the first OF subobject is not preceded by an ERO subobject identifying the next hop.
- . If the OF subobject follows an ERO subobject identifying the next hop that does not have the L-bit set.

If the processing node does not understand the OF subobject, it SHOULD send a PathErr with the error code `"Routing Error"` and error value of `"Bad Explicit Route Object"` toward the sender [RFC3209].

If the processing node understands the OF subobject and the ERO passes the above mentioned sanity check and any other sanity checks associated with other ERO subobjects local to the node, the node takes the following actions:

- . If the node supports the requested OF, the node expands the loose hop using the requested OF as optimization criterion for computing the route to the next abstract node. After processing, the OF subobjects are removed from the ERO. The rest of the steps for the loose ERO processing follow procedures outlined in [RFC3209].
- . If the node understands the OF subobject but does not support the requested OF, it SHOULD send a Path Error with error code `'Routing Problem'` and a new error subcode `'Unsupported Objective Function'`. The error subcode `'Unsupported Objective Function'` for Path Error code `'Routing Problem'` is to be assigned by IANA.
- . If the OF is supported but policy does not permit applying it, the processing node SHOULD send a Path Error with error code `"Policy control failure"` (value 2) and subcode `'objective function not allowed'`. The error subcode `'objective function not allowed'` for Path Error code `'Policy control failure'` is to be assigned by IANA.

2.3.2. Processing Rules for the MB subobject

The basic processing rules of an ERO are not altered. Please refer to [RFC3209] for details.

The scope of the MB subobject is between the previous ERO subobject that identifies an abstract node, and the subsequent ERO subobject that identifies an abstract node. Multiple MB subobjects may be present between any pair of abstract nodes.

If the processing node does not understand the MB subobject, it SHOULD send a PathErr with the error code "Routing Error" and error value of "Bad Explicit Route Object" toward the sender [RFC3209].

If the processing node understands the MB subobject and the ERO passes the above mentioned sanity check and any other sanity checks associated with other ERO subobjects local to the node, the node takes the following actions:

- . For all the MB subobject(s), the node expands the ERO such that the requested metric bound(s) are met for the route between the two abstract nodes in the ERO. After processing, the Metric subobjects are removed from the ERO. The rest of the steps for the ERO processing follow procedure outlined in [RFC3209].
- . If the node understands the MB subobject but cannot find a route to the next abstract node such that the requested metric bound(s) can be satisfied and the best-effort (B) bit is not set, it SHOULD send a Path Error with error code "Routing Problem" and a new error subcode "No route available toward destination with the requested metric bounds". The error subcode "No route available toward destination with the requested metric bounds" for Path Error code "Routing Problem" is to be assigned by IANA (See IANA section for details).
- . If the node understands the Metric subobject but cannot find a route to the next abstract node such that the requested metric bound(s) can be satisfied and the best-effort (B) bit is set, it SHOULD send a Path Error message with error code "Notify Error" and a new error subcode "Route not matching the requested metric bounds" is to be assigned by IANA (See IANA section for details).
- . The ERO expanding node SHOULD respect the Metric Bound constraints in realizing any segment recovery procedure to change the route of the segment expanded by the said node. If

best-effort (B) bit is set and the new recovery segment violates the Metric Bound constraints, the ERO expanding SHOULD send a Path Error message with error code "Notify Error" and a new error subcode "Route not matching the requested metric bounds" is to be assigned by IANA (See IANA section for details).

3. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], and [RFC3473].

4. IANA Considerations

4.1. ERO Subobject

This document adds the following two new subobject of the existing entry for ERO (20, EXPLICIT_ROUTE):

Value	Description
-----	-----
TBA (suggest value: 66) subobject	Objective Function (OF)
TBA (suggest value: 67)	Metric subobject

These subobject may be present in the Explicit Route Object, but not in the Route Record Object.

OF Code values carried in OF subobject requires an IANA entry with suggested values as defined in section 2.1.

4.2. New RSVP error sub-code

For Error Code = 24 "Routing Problem" (see [RFC2205]) the following sub-code is defined.

Sub-code	Value
-----	-----
No route available toward destination with the requested metric bounds	To be assigned by IANA. Suggested Value: TBA.

For Error Code = 25 "Notify Error" (see [RFC2205]) the following sub-code is defined.

Sub-code -----	Value -----
Route not matching the requested metric bounds	To be assigned by IANA. Suggested Value: TBA.

5. Acknowledgments

Authors would like to thank Matt Hartley, Ori Gerstel, Gabriele Maria Galimberti, Luyuan Fang and Walid Wakim for their review comments.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [IEEE.754.1985] IEEE Standard 754, "Standard for Binary Floating-Point Arithmetic", August 1985.

6.2. Informative References

- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", RFC 2209, September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

George Swallow
Cisco Systems.
swallow@cisco.com

Clarence Filsfils
Cisco Systems.
cfilsfil@cisco.com

Luyuan Fang
Cisco Systems.
lufang@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Ruediger.Kunze@telekom.de

Daniele Ceccarelli
Ericsson
Email: daniele.ceccarelli@ericsson.com

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 14, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Ori Gerstel
Cisco Systems
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG
July 15, 2013

Include Routes - Extension to
Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
draft-ali-ccamp-rsvp-te-include-route-04.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2014.

Copyright Notice

Copyright (c) 2012 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.

Abstract

There are scenarios that require two or more LSPs or segments of LSPs to follow same route in the network. This document specifies methods to communicate route inclusions along the loose hops during path setup using the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) protocol.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

Copyright Notice.....	1
1. Introduction.....	2
2. RSVP-TE signaling extensions.....	4
2.1. IPv4 Point-to-Point LSP ERO subobject.....	4
2.2. IPv6 Point-to-Point LSP ERO subobject.....	6
2.3. Processing rules for LSP ERO subobjects.....	7
3. Security Considerations.....	8
4. IANA Considerations.....	9
4.1. New ERO subobject types.....	9
4.2. New RSVP error sub-codes.....	9
5. Acknowledgments.....	10
6. References.....	10
6.1. Normative References.....	10
6.2. Informative References.....	10

1. Introduction

There are scenarios that require two or more Label Switched Paths (LSPs) to follow same route in the network. E.g., many deployments require member LSPs of a bundle/ aggregated link (or Forwarding Adjacency (FA))) follow the same route. Possible reasons for two or more LSPs to follow the same end-to-end or partial route include, but are not limited to:

- . Fate sharing: an application may require that two or more LSP fail together. In the example of bundle link this would mean that if one component goes down, the entire bundle goes down.
- . Homogeneous Attributes: it is often required that two or more LSPs have the same TE metrics like latency, delay variation, etc. In the example of a bundle/ aggregated link this would meet the requirement that all component links (FAs) of a bundle should have same latency and delay variation. As noted in [OSPF-TE-METRIC] and [ISIS-TE-METRIC], in certain networks, such as financial information networks, network performance (e.g. latency and latency variation) is becoming critical and hence having bundles with component links (FAs) with homogeneous delay and delay variation is important.

Similarly, there are scenarios where two or more LSPs need to follow a given resource in the network, e.g., two partially overlapping LSPs are required. In this case, inclusion of certain abstract nodes or resources between a specific pair of abstract nodes present in an ERO is required.

The RSVP-TE specification [RFC3209] and GMPLS extensions to RSVP-TE [RFC3473] allow abstract nodes and resources to be explicitly included in a path setup, e.g., using IPv4 prefix ERO subobject [RFC3209], IPv6 prefix ERO subobject [RFC3209] and Unnumbered Interface ID ERO subobject [RFC3477], etc. However, such inclusion may not be possible in the following scenarios:

- . Inclusion of a LSP path which does not originate, terminate or traverse the source node, in which case the addresses of the path for which inclusion is required are unknown to the source node.
- . Inclusion of a LSP path which, while known at the source node of the diverse LSP, has incomplete or unavailable route information, e.g. due to confidentiality of the path attributes. In these cases, the ingress node lacks sufficient knowledge about the loose hop. The ingress node, therefore, is not able to divide a loose hop into a proper sequence of strict or a sequence of finer-grained loose hops. Inter-domain and GMPLS overlay networks may present such restrictions.

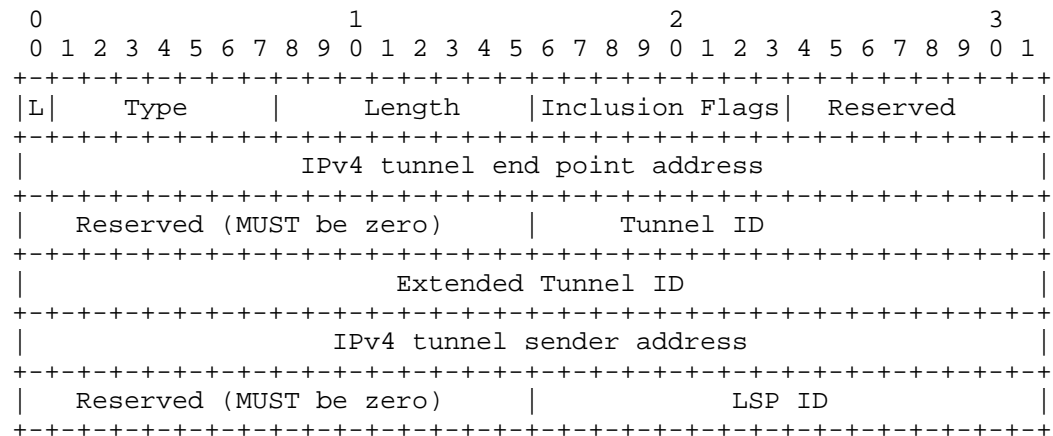
The above-mentioned use cases require relevant path inclusion requirements to be communicated to the route expanding nodes. This document addresses these requirements and defines procedures to address them.

2. RSVP-TE signaling extensions

New IPv4 and IPv6 Point-to-Point (P2P) LSP ERO subobject types are defined in this document. These ERO subobjects are used to communicate path inclusion requirements to the ERO expanding node(s). For this purpose, the subobjects carry RSVP-TE Forwarding Equivalence Class (FEC) of the reference LSP who's Path is be used to expand the loose hop of the LSP being signaled. This document only defines the use of these objects for ERO loose hops.

2.1. IPv4 Point-to-Point LSP ERO subobject

The IPv4 Point-to-Point LSP ERO subobject is defined as follows:



L

The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the ERO. If the bit is not set, the subobject represents a strict hop in the explicit route.

This document only defines the use of the subobject in loose hopes in the ERO, i.e., L bit MUST of set to 1.

Type

IPv4 Point-to-Point LSP subobject
(to be assigned by IANA; suggested value: 38).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 24.

Inclusion Flags

The Inclusion-Flags are used to communicate desirable types of inclusion. The following flags are defined.

0x01 = Mandatory inclusion

This flag is used to indicate that the route of the LSP being signaled MUST follow the path specified by the LSP subobject.

0x02 = Best-effort inclusion

This flag is used to indicate that the route of the LSP being signaled SHOULD follow the path specified by the LSP subobject.

The remaining fields are used to specify RSVP-TE FEC of the reference LSP who's Path is be used to expand the route of the LSP being signaled. Specifically,

Tunnel ID

Tunnel ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

Extended Tunnel ID

Extended Tunnel ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

IPv4 tunnel sender address

IPv4 tunnel sender address of the reference LSP who's path is be used to expand the route of the LSP being signaled.

LSP ID

LSP ID of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

2.2. IPv6 Point-to-Point LSP ERO subobject

The IPv6 Point-to-Point LSP ERO subobject is defined as follows:

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	2	3	4	5	6	7	8	9	0	1								
L										Type										Length										Inclusion Flags										Reserved									
										IPv6 tunnel end point address																																							
										IPv6 tunnel end point address (cont.)																																							
										IPv6 tunnel end point address (cont.)																																							
										IPv6 tunnel end point address (cont.)																																							
										Must Be Zero										Tunnel ID																													
										Extended Tunnel ID																																							
										Extended Tunnel ID (cont.)																																							
										Extended Tunnel ID (cont.)																																							
										Extended Tunnel ID (cont.)																																							
										IPv4 tunnel sender address																																							
										IPv4 tunnel sender address (cont.)																																							
										IPv4 tunnel sender address (cont.)																																							
										IPv4 tunnel sender address (cont.)																																							
										Reserved (MUST be zero)										LSP ID																													

L

The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the ERO.

If the bit is not set, the subobject represents a strict hop in the explicit route.

This document only defines the use of the subobject in loose hops in the ERO, i.e., L bit MUST of set to 1.

Type

IPv6 Point-to-Point LSP subobject
(to be assigned by IANA; suggested value: 39).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 48.

Inclusion Flags

The Inclusion Flags are as defined for the IPv4 Point-to-Point LSP XRO subobject.

The remaining fields are used to specific RSVP-TE FEC of the reference LSP who's Path is be used to expand the route of the LSP being signaled.

2.3. Processing rules for LSP ERO subobjects

The basic processing rules of an ERO are not altered. Please refer to [RFC3209] for details.

If an LSR strips all local subobjects from an ERO carried in a Path message (according to the procedures in [RFC3209]) and finds that the next subobject is an IPv4 P2P LSP subobject or IPv6 P2P LSP subject, it MUST attempt to resolve the LSP subobject as described in the following.

If the L bit of the LSP subobject is not set, i.e., the subobject represents a strict hop in the explicit route, the processing node MUST respond with a PathErr message with the error code "Routing Problem" (24) and the error value "Bad initial subobject" (4).

If the inclusion flags of the LSP subobject is set to "mandatory inclusion", the processing node follows the following procedure:

- If the path taken by the LSP referenced in the LSP subobject is known to the processing node and the path contains the loose abstract node in the ERO hop, the processing node MUST ensure that loose hop expansion to the next abstract node follows the referenced path.
- If the path taken by the LSP referenced in the LSP subobject is unknown to the processing node or the referenced path does not contain the loose abstract node in the ERO hop, the processing node MUST send a PathErr message with the error code "Routing Problem" (24) and the new error value "unknown or inconsistent LSP subobject" (value to be assigned by IANA) for the signaled LSP.

If the inclusion flags of the LSP subobject is set to "best-effort inclusion", the processing node follows the following procedure:

- If the path taken by the LSP referenced in the LSP subobject is known to the processing node and the path contains the loose abstract node in the ERO hop, the processing node SHOULD ensure that loose hop expansion to the next abstract node follows the referenced path.
- If the path taken by the LSP referenced in the LSP subobject is unknown to the processing node and/ or the referenced path does not contain the loose abstract node in the ERO hop, the processing node SHOULD ignore the route inclusion specified in the LSP subobject and SHOULD compute a suitable path to the loose abstract node in the ERO hop and proceed with the signaling request. After sending the Resv for the signaled LSP, the processing node SHOULD return a PathErr with the error code "Notify Error" (25) and error sub-code " unknown or inconsistent LSP subobject" (value to be assigned by IANA) for the signaled LSP.

3. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], and [RFC3473] and [RFC4874].

4. IANA Considerations

4.1. New ERO subobject types

This document adds the following new subobject of the existing entry for ERO (20, EXPLICIT_ROUTE):

Value	Description
-----	-----
TBA	IPv4 Point-to-Point LSP ERO subobject
TBA	IPv6 Point-to-Point LSP ERO subobject

These subobject may be present in the Explicit Route Object, but not in the Route Record Object.

4.2. New RSVP error sub-codes

IANA registry: RSVP PARAMETERS
Subsection: Error Codes and Globally-Defined Error Value Sub-Codes

For Error Code "Routing Problem" (24) (see [RFC3209]) the following sub-codes are defined.

Sub-code	Value
-----	-----
Unknown or inconsistent LSP subobject	To be assigned by IANA.

For Error Code "Notify Error" (25) (see [RFC3209]) the following sub-codes are defined.

Sub-code	Value
-----	-----
Unknown or inconsistent LSP subobject	To be assigned by IANA.

5. Acknowledgments

Authors would like to thank Gabriele Maria Galimberti, Luyuan Fang and Walid Wakim for their review comments.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

6.2. Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiimoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", RFC 2209, September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

George Swallow
Cisco Systems, Inc.
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

Ori Gerstel
Cisco Systems
ogerstel@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Ruediger.Kunze@telekom.de

CCAMP Working Group
Internet Draft
Intended status: Standards Track

Vishnu Pavan Beeram (Ed)
Juniper Networks
Igor Bryskin (Ed)
ADVA Optical Networking

Expires: January 15, 2013

July 15, 2013

Mutually Exclusive Link Group (MELG)
draft-beeram-ccamp-melg-01.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on January 15, 2013.

Copyright Notice

Copyright (c) 2013 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.

Abstract

This document introduces the concept of MELG ("Mutually Exclusive Link Group") and discusses its usage in the context of mutually exclusive Virtual TE Links.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Table of Contents


1. Introduction.....	2
2. Mutually Exclusive Virtual TE Links.....	3
3. Mutually Exclusive Link Group.....	6
4. Protocol Extensions.....	6
4.1. OSPF.....	6
4.2. ISIS.....	7
5. Security Considerations.....	9
6. IANA Considerations.....	9
6.1. OSPF.....	9
6.2. ISIS.....	9
7. Normative References.....	9
8. Acknowledgments.....	9

1. Introduction

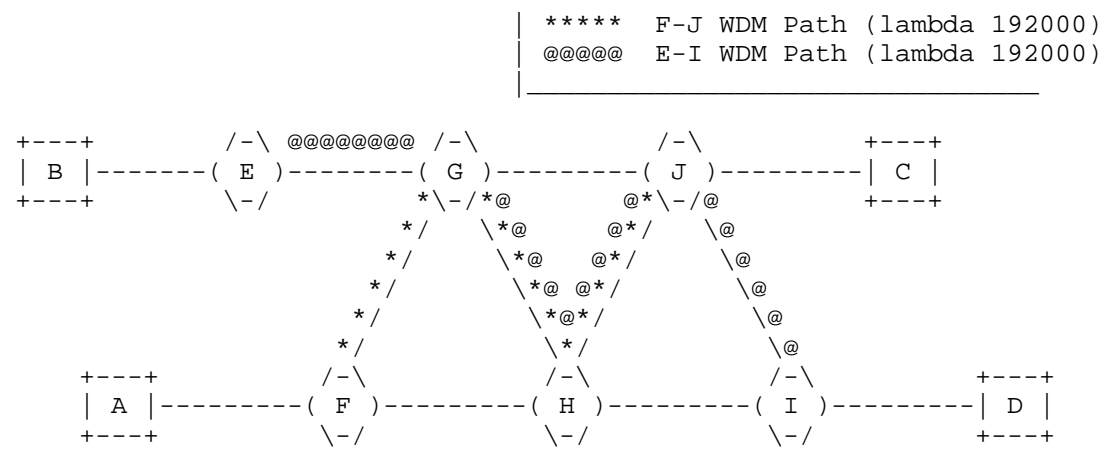
A Virtual TE Link (as defined in [RFC6001]) advertised into a Client Network Domain represents a potentiality to setup an LSP in the Server Network Domain to support the advertised TE link. The Virtual TE Link gets advertised like any other TE link and follows exactly the same rules that are defined for the advertising, processing and use of regular TE links [RFC4202]. However, "mutual exclusivity" is one attribute that is specific to Virtual TE links. This document discusses the need to advertise this information and the means to do so.

2. Mutually Exclusive Virtual TE Links

Consider the network topology depicted in Figure 1a. This is a typical packet optical transport deployment scenario where the WDM layer network domain serves as a Server Network Domain providing transport connectivity to the packet layer network Domain (Client Network Domain).



Nodes A, B, C and D are IP routers that are connected to an Optical WDM transport network. E, F, G, H, I and J are WDM nodes that constitute the Server Network Domain. The border nodes (E, F, I and J) operate in both the server and client domains. Figure 1b depicts how the Client Network Domain TE topology looks like when there are no Client TE Links provisioned across the optical domain.



Now consider augmenting the Client TE topology by creating a couple of Virtual TE Links across the optical domain. The potential paths in the WDM network catering to these two virtual TE links are as shown in Fig 2a and the corresponding augmented Client TE topology is as illustrated in Fig 2b.

In this particular example, the potential paths in the WDM layer network supporting the Virtual TE Links not only intersect, but also require the usage of the same resource (lambda channel 192000) on the intersection. Because the Virtual TE Links depend on the same uncommitted network resource, only one of them could get activated at any given time. In other words they are mutually exclusive. This scenario is encountered when the potential paths depend on any common physical resource (e.g. transponder, regenerator, wavelength converter, etc.) that could be used by only one Server Network Domain LSP at a time.

For a Client Network Domain path computation function (especially a centralized one) that is capable of concurrent computation of multiple paths, it is important to know the existence of mutually exclusive relationship between Virtual TE Links. Absent this information, there exists the risk of yielding erroneous concurrent path computation results where only a subset of the computed paths can get successfully provisioned. This document introduces the concept of Mutually Exclusive Link Group to address this problem.

The Virtual TE Link mutual exclusivity attribute can be either (a) Static or (b) Dynamic.

In case (a), the Virtual TE Link mutual exclusivity exists permanently within a given network configuration. This type of mutual exclusivity comes into play when two or more Virtual TE Links depend on a Server Network Domain resource that could be used in its entirety by only one Virtual TE Link (when committed).

In case (b), the Virtual TE Link mutual exclusivity exists temporarily within a given network configuration. This type of mutual exclusivity comes into play when two or more Virtual TE Links depend on a Server Network Domain resource that could be shared simultaneously in some limited capacity by several Virtual TE Links (when committed). Consider, for example, a situation when three Virtual TE Links depend on a Server Domain WDM link that currently has two lambda channels available. Consider further that each Virtual TE Link (in order to be committed) requires one lambda channel to be allocated on said WDM link. Obviously, under these conditions only two out of three Virtual TE Links could be

concurrently committed. Such Virtual TE Link mutual exclusivity is dynamic and temporary because as soon as additional lambda channels become available on the WDM link, the Virtual TE Link mutual exclusivity will cease to exist - it will become possible to commit all three Virtual TE Links concurrently.

This revision of the draft discusses only the Static Virtual TE Link mutual exclusivity. Dynamic Virtual TE Link mutual exclusivity will be addressed in later revisions.

3. Mutually Exclusive Link Group

The Mutually Exclusive Link Group (MELG) construct defined in this document has 2 purposes

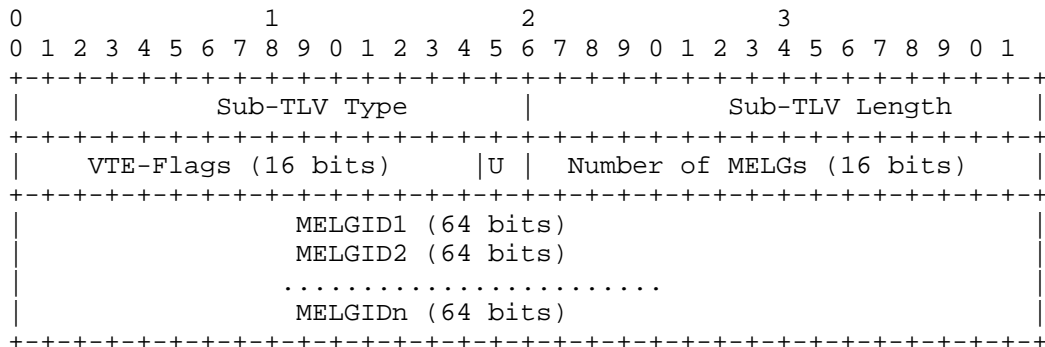
- To indicate via a separate network unique number (MELG ID) an element or a situation that makes the advertised Virtual TE Link belong to one or more Mutually Exclusive Link Groups. Path computing element will be able to decide on whether two or more Virtual TE Links are mutually exclusive or not by finding an overlap of advertised MELGs (similar to deciding on whether two or more TE links share fate or not by finding common SRLGs)
- To indicate whether the advertised Virtual TE Link is committed or not at the moment of the advertising. Such information is important for a path computation element: Committing new Virtual TE links (vs. re-using already committed ones) has a consequence of allocating more server layer resources and disabling other Virtual TE Links that have common MELGs with newly committed Virtual TE Links; Committing a new Virtual TE Link also means a longer setup time for the Client LSP and higher risk of setup-failure.

4. Protocol Extensions

4.1. OSPF

The MELG is a sub-TLV of the top level TE Link TLV. It may occur at most once within the Link TLV. The format of the MELGs sub-TLV is defined as follows:

Name: MELG
Type: TBD
Length: Variable



Number of MELGs: number of MELGs advertised for the Virtual TE Link;

VTE-Flags: Virtual TE Link specific flags;

MELGID1,MELGID2,...,MELGIDn: 64-bit network domain unique numbers associated with each of the advertised MELGs

Currently defined Virtual TE Link specific flags are:

U bit (bit 1): Uncommitted - if set, the Virtual TE Link is uncommitted at the time of the advertising (i.e. the server layer network LSP is not set up); if cleared, the Virtual TE Link is committed (i.e. the server layer LSP is fully provisioned and functioning). All other bits of the "VTE-Flags" field are reserved for future use and MUST be cleared.

Note: A Virtual TE Link advertisement MAY include MELGs sub-TLV with zero MELGs for the purpose of communicating to the TE domain whether the Virtual TE Link is currently committed or not.

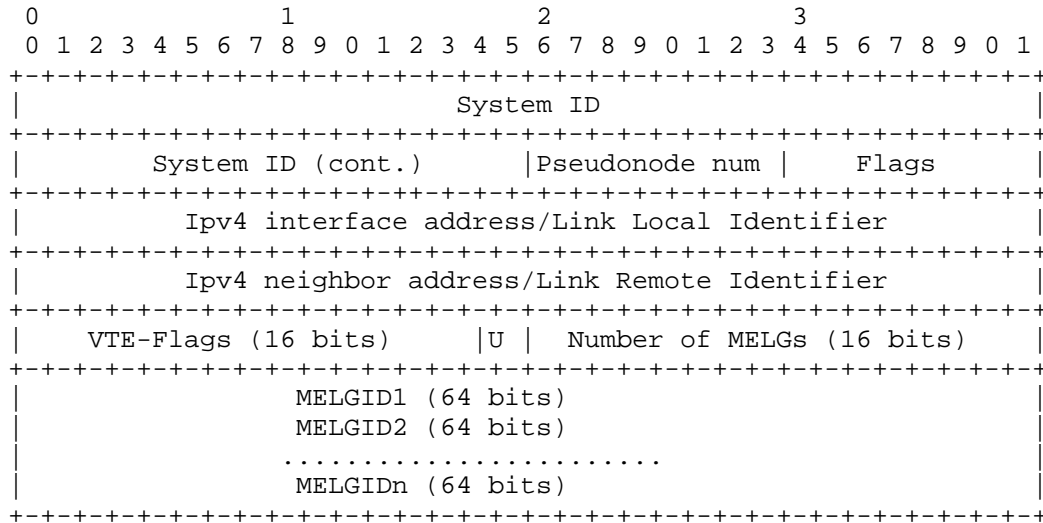
4.2. ISIS

The MELG TLV (of type TBD) contains a data structure consisting of:

6	octets of System ID
1	octet of Pseudonode Number
1	octet Flag
4	octets of IPv4 interface address or 4 octets of a Link Local Identifier
4	octets of IPv4 neighbor address or 4 octets of a Link Remote Identifier
2	octets MELG-Flags

2 octets - Number of MELGs
 variable List of MELG values, where each element in the list
 has 8 octets

The following illustrates encoding of the value field of the MELG TLV.



The neighbor is identified by its System ID (6 octets), plus one octet to indicate the pseudonode number if the neighbor is on a LAN interface.

The least significant bit of the Flag octet indicates whether the interface is numbered (set to 1) or unnumbered (set to 0). All other bits are reserved and should be set to 0.

The length of the TLV is $20 + 8 * (\text{number of MELG values})$.

The semantics of "VTE-Flags", "Number of MELGs" and "MELGID Values" are the same as the ones defined under OSPF extensions.

The MELG TLV MAY occur more than once within the IS-IS Link State Protocol Data Units.

5. Security Considerations

TBD

6. IANA Considerations

6.1. OSPF

IANA is requested to allocate a new sub-TLV type for MELG (as defined in Section 4.1) under the top-level TE Link TLV.

6.2. ISIS

IANA is requested to allocate a new IS-IS TLV type for MELG (as defined in Section 4.2).

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4202] K.Kompella, Y.Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC4202, October 2005.
- [RFC6001] D.Papadimitriou, M.Vigoureaux, K.Shiomoto, D.Brungard and JL. Le Roux, "GMPLS Protocol Extensions for Multi-Layer and Multi-Region Networks", RFC 6001, October 2010.

8. Acknowledgments

Chris Bowers [cbowers@juniper.net]

Authors' Addresses

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

Igor Bryskin
ADVA Optical Networking
Email: ibryskin@advaoptical.com

John Drake
Juniper Networks
Email: jdrake@juniper.net

Gert Grammel
Juniper Networks
Email: ggrammel@juniper.net

Wes Doonan
ADVA Optical Networking
Email: wdoonan@advaoptical.com

Manuel Paul
Deutsche Telekom
Email: Manuel.Paul@telekom.de

Ruediger Kunze
Deutsche Telekom
Email: Ruediger.Kunze@telekom.de

Oscar Gonzalez de Dios
Telefonica
Email: ogondio@tid.es

Cyril Margaria
Coriant GmbH
Email: cyril.margaria@coriant.com

Friedrich Armbruster
Coriant GmbH
Email: friedrich.armbruster@coriant.com

Daniele Ceccarelli
Ericsson
Email: daniele.ceccarelli@ericsson.com

CCAMP Working Group
Internet-Draft
Intended status: Informational
Expires: January 12, 2014

D. Ceccarelli
Ericsson
O. Gonzalez de Dios
Telefonica I+D
F. Zhang
X. Zhang
Huawei Technologies
July 11, 2013

Use cases for operating networks in the overlay model context
draft-ceccadedios-ccamp-overlay-use-cases-01

Abstract

This document defines a set of use cases for operating networks in the overlay model context through the Generalized Multiprotocol Label Switching (GMPLS) overlay interfaces.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 12, 2014.

Copyright Notice

Copyright (c) 2013 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	4
2. Terminology	5
3. Variants of UNI interface	6
3.1. UNI	6
3.1.1. UNI example	7
3.2. ONI	8
3.2.1. Routing info over the ONI	9
4. Hybrid Nodes	10
5. Adding the PCEP to the UNI	11
5.1. Edge-node to core-node PCEP interface	11
5.2. PCEP and colored UNI	13
5.3. Using PCEP to obtain TE reachability info	13
6. Provisioning	13
6.1. Use case P1 - Local Trigger	14
6.2. Use case P2 - Remote Signaling	14
6.3. Use case P3 - Provisioning with requirements over the UNI	15
6.4. Use case P4 - Provisioning with requirements and collection request over the UNI	17
6.5. Use case P5 - Advertisement of collected metrics in the client layer	17
6.6. Use case P6 - Provisioning with requirements over the ONI	17
6.7. Use case P7 - Dual homing	18
7. Recovery	19
7.1. Use case R1 - Segment Recovery - Single homing	19
7.2. Use case R2 - Local recovery - Dual Homing - Single overlay node	20
7.3. Use case R3 -End to end Recovery - Dual homing - Double overlay node	20
7.4. Use case R4 - Combination of client protection and server restoration	21
8. Optimization	23
8.1. Use case O1 -	23
9. Security Considerations	23
10. IANA Considerations	23
11. Contributors	23
12. Acknowledgements	24
13. References to be added	24
14. References	24
14.1. Normative References	24
14.2. Informative References	24

Authors' Addresses	24
------------------------------	----

1. Introduction

The GMPLS overlay model [RFC 4208] specifies a client-server relationship between networks where client and server layers are managed as separate domains because of trustiness, scalability and operational issue. By means of procedures from the GMPLS protocol suite it is possible to build a topology in the client (overlay) network from Traffic Engineering paths in the server network. In this context, the UNI (User to Network Interface) is the demarcation point between networks. It is a boundary where policies, administrative and confidentiality issues apply that limit the exchange of information.

This GMPLS overlay model supports a wide variety of network scenarios. The packet over optical scenario is probably the most popular example where the overlay model applies.

In order to exploit the full potential of client/server network interworking in the overlay model, it may be desirable to know in advance whether is it feasible or not to connect two client network nodes [INTERCON-TE]. This requires to have a certain amount of TE information of the server network in the client network. This need not be the full set of TE information available within each network, but does need to express the potential of providing TE connectivity. This subset of TE information is called TE reachability information.

Reachability can be classified according to the available information:

- Non-qualified reachability: The most basic form of TE reachability is just the information of whether is it physically feasible to connect two nodes. Thus, Non-qualified reachability is the ability to be able to connect one end point to another.
- Qualified reachability: In addition to knowing of whether is it physically feasible to connect two nodes, the qualified reachability information indicates metrics of the potential connection. This metric can be a cost, a service metric (delay), bandwidth, etc.
- Reachability with associated Potential Virtual Link: In this case, the client node, in addition to have the information of the feasibility of reaching a given destination, it has the information of a possible path, that can be established at any time through UNI signalling. This possible path may have been pre-computed in advance and may contain the explicit path (e.g. ERO). Thus, client layer topology could be richened with this potential virtual TE link information.

Current standard GMPLS UNI [RFC 4208] is focused on signaling and extensions to the GMPLS UNI [RFC4208] are being proposed to gather a tighter integration between the client packet layer and the server optical one. Also, new elements, like the Path Computation element (PCE), have become more popular and may also play a role in the overlay context. In order to understand what are the requirements that need to be fulfilled, it is necessary to understand current use cases, that is, what are the network operators expecting the UNI to do.

The first group of uses cases of the overlay model is related to the automatic provisioning, by which the client (overlay) network is able to set-up a connection through the server network. The second group of use cases is related to the enhancement of the recovery mechanism by a higher coordination of overlay and server networks.

Summing up, the goal of this document is to overview the network scenarios where the overlay model applies and analyze the most interesting aspects of the use cases that fall under the above definitions, with particular focus on signaling, exchange of info, operations, path computation and L1VPN aspects.

2. Terminology

The following terms are used within the document:

- Edge node [RFC4208]: node of the client domain belonging to the overlay network, i.e. nodes with at least one interface connected to the server domain.
- Core node [RFC4208]: node of the server domain.
- Access link: link between core node and edge node. It is the link where the UNI is usually implemented.
- Remote node: node in the client domain which has no direct access to the server domain but can reach it through an edge node in its same administrative domain.
- Local trigger: LSP setup request issued to an edge node. It triggers the setup of a client layer FA through the server domain via a UNI interface.
- Remote trigger: LSP setup request issued to a remote node. It triggers the setup of a client layer LSP which, upon reaching an edge node, will use an FA through the server domain dynamically provided via an UNI interface.

3. Variants of UNI interface

[RFC4208] defines the GMPLS UNI as an overlay interface allowing for the exchange of both routing and signaling messages between a client and a server domain. At that time only signaling extensions and procedures for the UNI interface were defined but since then a lot of RSVP-TE extensions have been defined (e.g. [LSP-DIV][TE-REC]).

In this section a taxonomy of different variants of UNI interfaces is provided.

3.1. UNI

GMPLS UNI defined in [RFC4208] allows the exchange of RSVP-TE Path messages between edge and core nodes including the Explicit Route Object (ERO) or Record Route Object (RRO). This allows for the explicit indication of the path (or part of it) to choose in the server domain or the recording of the nodes and links chosen respectively.

Subsequently new requirements have been added to the exchange of messages between edge and core nodes. In the message flow from edge to core nodes it is desirable to define a number of path computation constraints that the server domain needs to consider when providing a path that will be used as connection between two edge nodes. An example of path computation constraints consists of: SRLG diversity, max latency, max number of hops, etc. Encoding extensions for adding constraints to the server domain path computation are defined in [UNI-PLUS] and [LSP-DIV].

Similarly it is desirable that the core nodes are able to provide the edge nodes with the parameters qualifying the path provided. Such set of parameters is the same identified above and encoding extensions for its collection are defined in [TE-REC] and [SRLG-COLL].

RSVP-TE extensions are obviously inherited by the UNI interface, which has significantly been augmented with respect to the RFC4208 version.

In the overlay model applied to the packet-opto environment the UNI can be implemented over grey or colored interfaces. In the former case the transponder is placed on the RROADM and the traffic from the router to the RROADM has no G.709 framing, while in the latter the transponder is moved on the router and the traffic injected into the WDM domain is G.709 framed and managed as an alien lambda. From a procedure point of view there is no difference between the grey and the colored UNI, what is different is the type of information that

the WDM PCE need to be provided with in order to compute a path in the WDM domain but check its feasibility from router to router.

For applicability considerations regarding the GMPLS UNI please refer to [UNI-APP]

3.1.1. UNI example

This section provides an example of how SRLGs, latency and other types of parameters can be used in the edge node to core node direction as path computation constraints and in the reverse direction from core node to edge node as path qualifiers. The example is applied to a packet-opto environment.

In the reference network below, suppose that router 1 wants RROADM A to provide a path between A and G with max unidirectional latency 20ms and SRLG different from (37;52). Such parameters are passed to A via the RSVP-TE Path message.




Figure 1: Path computation constraints

Node A performs a constrained path computation in the optical domain (A-B-C-G) and provides 1 (via the RSVP-TE RESV message) with the parameters qualifying the provided path e.g. max latency 12ms and SRLG (11,55).




Figure 2: Real network topology

3.2. ONI

A different variant of GMPLS UNI interface can be implemented within the overlay model context. It foresees adding some TE (Traffic Engineering) topological information exchange between edge and core nodes. We use the term ONI (Overlay Network Interface) to identify this variant of UNI interface with such capabilities.

The ONI foresees the definition of a virtual topology in the server layer (arbitrarily configured by the network operator) which is exported by each core border node to its peering edge node by means of a routing protocol (e.g. OSPF-TE, BGP-LS). Such virtual topology comprises a mix of potential virtual TE-links [RFC4847] and virtual nodes.

A virtual TE-link, as defined in RFC4847, is a provider network Traffic Engineering (TE) link advertised to customers in routing information for purposes that include path computation. We introduce the term potential virtual TE-link to indicate those virtual TE-links whose resources are not necessarily reserved or committed in the server layer network.

A potential virtual TE link is advertised via the ONI as a real TE link and hence contributes to augmenting the client layer topology. Moreover it does not require the allocation of resources in the server layer until it is really used thus allowing the mutually exclusive sharing of server layer network resources with other potential Virtual TE Links.

On the other side a virtual node is a collection of zero or more provider network domain nodes that are collectively represented to the clients as a single node that exists in the customer layer

network and is capable of terminating of access, inter-domain and virtual links. (a virtual node can correspond 1:1 to a physical node, to a part of it or to a set of nodes).

As per the UNI, we define two variations of ONI for the packet-opto scenario, namely the ONI when we refer to grey interfaces (i.e. transponder on the RROADM) and ONI++ when referring to colored interfaces (i.e. transponder on the router).

3.2.1. Routing info over the ONI

The following reference network in an IPoWDM environment is used to depict what kind of information needs to be advertised to the edge nodes in order to augment the client layer topology.




Figure 3: Real network topology

Where:

```

+----+
| 1 | = Edge node          ***** = Access Link (UNI link)
+----+

  /-\
 ( A ) = Core node        ----- = Core domain real link
  \-/

```

Suppose that the network operator decides to let client network see two potential virtual TE-links, one between A and C and the other between D and G. Figure below show the potential virtual TE-links on

the left and corresponding real topology on the right.

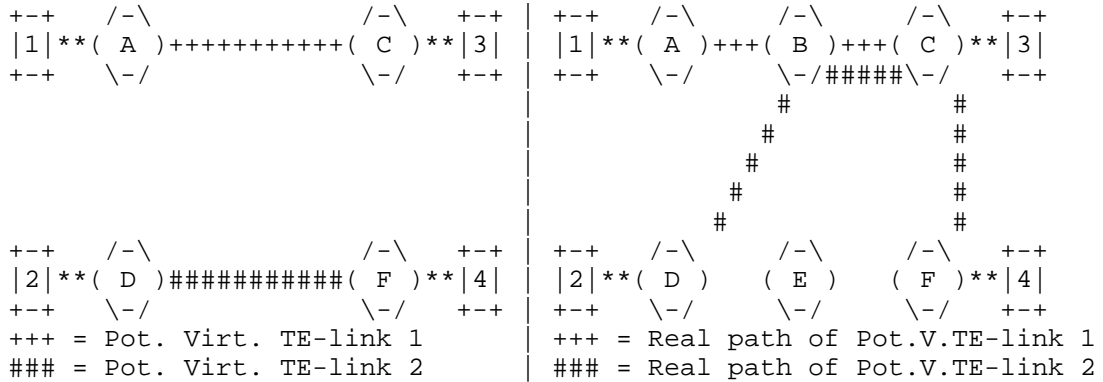


Figure 4: Virtual network topology

From the above figure it is possible to see that the edge nodes, in order to be able to compute a path, need to know the following:

- Access links: availability and TE parameters
- Potential virtual TE-links: availability, TE parameters, diversity, constraints (e.g. latency, packet loss)
- Virtual nodes: (in figure above a 1 to 1 relationship between real and virtual nodes is assumed) availability, connectivity matrix, TE parameters, diversity, constraints (in those cases where a virtual node is composed of 2 or more real nodes)

More details on potential virtual TE links diversity can be found at [MELG].

4. Hybrid Nodes

The overlay model and interfaces can be applied also to hybrid nodes [RFC5212], where it is possible to have different control plane instances in the client and server domain and keep on exploiting the advantages of managing the domains separately without losing the capability of making them tightly interwork.

5. Adding the PCEP to the UNI

In the GMPLS overlay model, the edge nodes do not participate in the routing instance of the core network. This means that the overlay network is not aware of the details (topology and TE information) of the core network. Thus, it is generally assumed that, in the case of a UNI request, the routing is performed in the core network.

The operator is able to include a set of constraints that have impact on the routing that will be performed in the core network in two different ways. One is to enrich RSVP-TE to include such constraints (please see above), and the other one is use the well-defined PCEP protocol.

The main advantage of using the PCEP interface is that the semantics for the different queries are well defined. Moreover, in case confidentiality and trustiness is an issue, the Path Key mechanism can be used in the answers to hide the details of the path. Different possible scenarios are shown bellow.

5.1. Edge-node to core-node PCEP interface

The first scenario foresees the provisioning of a PCEP interface from the core nodes to the edge nodes. Such interface is reachable from the same interfaces used for the exchange of UNI signaling messages. The path computation in the server domain can be centralized or distributed. In the former scenario the core node acts as a proxy between the PCC on the edge node and the PCE as shown in figure below.

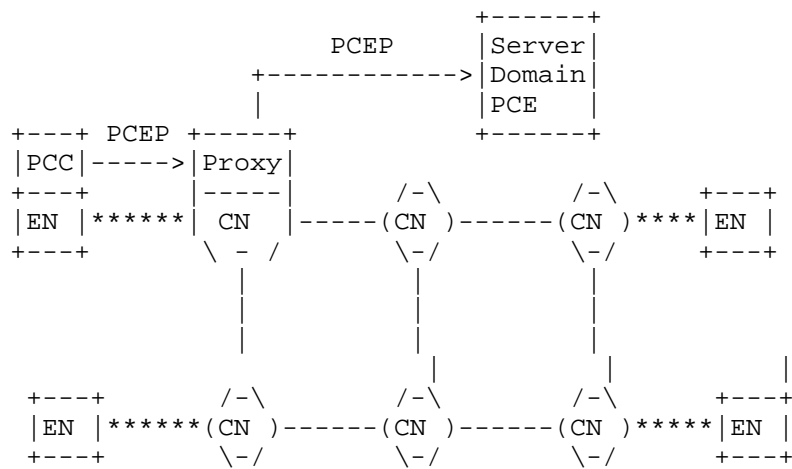


Figure 5: PCEP proxy on core node

On the other hand, when the path computation in the server domain is performed in a distributed way, the PCEP interface can be exported by the core nodes so to allow the edge node issue path computation queries. See figure below.

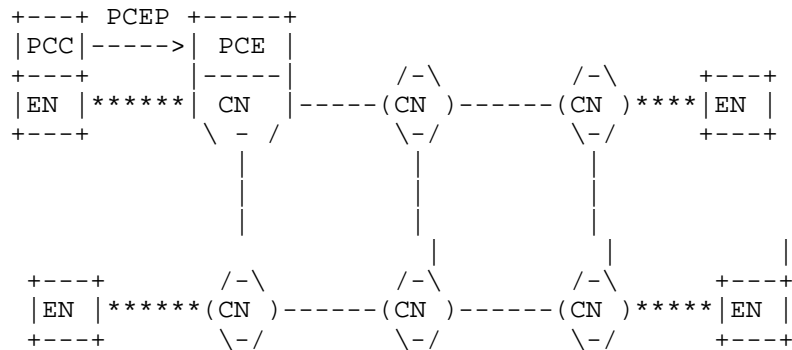


Figure 6: PCEP interface between EN and CN

In those cases where the addressing space is common between the client and the server layer, and hence with the PCE, it is also possible to have the PCC on the client layer directly speaking with the server domain PCE as shown in figure below.




Figure 7: PCE and PCC direct connection

5.2. PCEP and colored UNI

TBD

5.3. Using PCEP to obtain TE reachability info

Although the overlay network does not participate in the routing instance of the server network, RFC 4208 allows TE reachability to be exchanged between both overlay and server networks. TE reachability, as defined in [INTERCON-TE], is the ability to reach a specific address along a TE path. In the overlay context, TE reachability indicates if it is possible to reach from one node of the overlay network to another one through the server network. Also, TE reachability can indicate some TE metrics of the possible paths metrics, hop count, available bandwidth, delay, shared risk.

The PCEP interface can be used at the UNI border to query from overlay network to server network about the possible paths between nodes. The answers do not need to include a full detailed ERO, but just a collection of the desired metrics and, either a path key, or an ERO with just the end-points.

6. Provisioning

Two different use cases can be identified for path provisioning depending on where the setup trigger is issued. They will be called Local and Remote Trigger in the following. Further details regarding

different provisioning models over the UNI can be found in [UNI-APP]. What described below is in addition to the simple provisioning use case without constraints defined in [RFC4208].

6.1. Use case P1 - Local Trigger

Local Trigger is the use case defined in RFC4208, where a setup trigger is issued to one of the edge nodes belonging to the overlay network, i.e. directly connected to the UNI.




Figure 8: Local trigger

As it is possible to see in the figure above, a trigger is issued on R3 (edge node) for starting the setup request procedure over the UNI interface (R3-A). Such trigger might include an arbitrary set of constraints for the path computation in the optical domain. Once the optical LSP is setup and an adjacency in the packet layer between R3 and R5 is created, it is advertised in the rest of the packet layer and used by the signaling protocol (e.g. LDP) for setting up end-to-end (e.g. from R1 to R7) packet LSPs.

6.2. Use case P2 - Remote Signaling

A second use case consists on the utilization of a connection oriented (e.g. RSVP-TE) signaling protocol in the packet domain that allows issuing the setup trigger directly on the end nodes of the packet domain and using the RSVP-TE message for triggering the setup of the LSP in the optical domain.

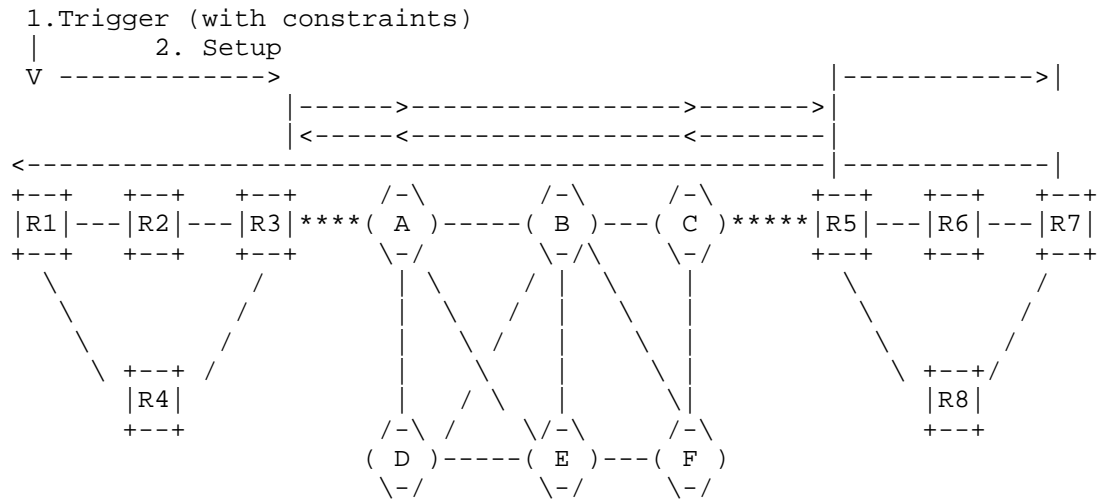


Figure 9: Local trigger

The utilization of the remote trigger allows for a strict control of the resources that will be used for the setup of the end to end path. In particular, two further approaches are possible when using the remote trigger: RSVP-TE or BGP approach. In the case of RSVP-TE it is possible to exploit the capabilities of the protocol of signaling LSPs across multiple domains and indicate in the ERO (explicit route object) the resources that need to be included in the path. It is possible to indicate a strict ERO (i.e. all the indicated resources and only those ones must be used) or a loose ERO (i.e. at least the indicated resources must be used but the PCEs of the different domains decide autonomously how to fill the gaps in the ERO).

6.3. Use case P3 - Provisioning with requirements over the UNI

This use case applies to both the local and remote trigger scenarios and describes how it is possible for the client layer to ask for a connection in the server layer with some constraints. An example of the constraints that can be applied to the path computation in the server layer consists of the following criteria:

- + Diversity: it is possible to indicate the resources must or should be avoided during the path computation by means of the Exclude Router Object (XRO) [RFC4874], the Explicit Exclusion Route Subobject (EXRS) [RFC4874] and the LSP subobject [LSP-DIV]. Such resources consists on:

```

-IPv4 prefix [RFC4874]

-IPv6 prefix [RFC4874]

-Unnumbered Interface ID [RFC4874]

-Autonomous system number [RFC4874]

-SRLG [RFC4874]

-IPv4 P2P subobject [LSP-DIV]

-IPv6 P2P subobject [LSP-DIV]

+ Latency: max delay introduced by the server layer LSP [OF-TEMB]

+ Latency variation: max latency variation allowed for the server
layer LSP [OF-TEMB]

+ Cost: max cost allowed for the server layer LSP [OF-TEMB]

```

The overlay Edge Node can include into the RSVP-TE Path message an arbitrary number of path computation constraints for the provisioning of the LSP in the server domain. In figure below and example is shown using as path computation constraints: (i) max latency should be 20ms and (ii) SRLG must be different from (A;B).

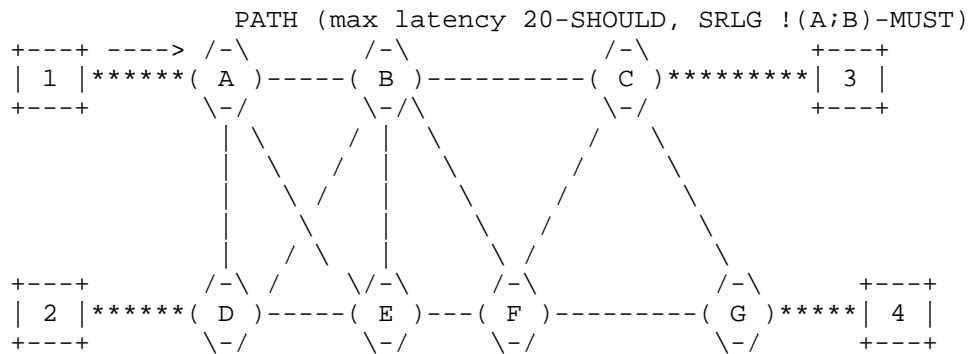


Figure 10: Use Case P3 - provisioning with requirements

If the path computation in the server domain is able to provide an LSP meeting the requirements (at least the must ones) such LSP is established and a RESV message is returned to the Edge node, otherwise an error message (PattErr) is returned.

6.4. Use case P4 - Provisioning with requirements and collection request over the UNI

This use case is an extension of Use case P3. In addition to the path request with path computation constraints, the client layer is also allowed to request for the collection in the server domain of the effective values of the parameters indicated as path computation constraints. The collection of such parameters is indicated via dedicated flags in the LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES in the Path Message. Flags defined are:

- Cost collection flag [TE-REC]
- Latency collection flag [TE-REC]
- Latency Variation collection flag [TE-REC]
- SRLG collection flag [SRLG-COLL]

In the scenario depicted in figure Figure 10a request with constraints on max latency and SRLG diversity might be issued together with the request of collecting e.g. the effective SRLGs of the provided path. Collected parameters are returned to the overlay edge node via the Record Route Object (RRO) in the RESV message.

6.5. Use case P5 - Advertisement of collected metrics in the client layer

In the case of local trigger, that is when a forwarding adjacency (FA) is created between two nodes in the client domain by means of an LPS in the server domain, it is desirable to advertise the characteristics of such FA in the client domain so that an informed path computation in the client layer can be performed.

Interior Gateway Protocol (IGP) extensions for the advertisement of the TE metrics characterizing such forwarding adjacencies have been defined for IS-IS [ISIS-TEMET] and OSPF-TE [OSPF-TEMET].

It is also possible to perform such advertisement by means of BGP-LS, where the collected data are sent to the route reflector and then forwarded to all the BGP speakers in the different client domains.

6.6. Use case P6 - Provisioning with requirements over the ONI

In the case of ONI/ONI++ each overlay edge node is provided with the overlay topology of the server domain as indicated in section Section 3.2. When a trigger is receiver by the edge node (local or remote), it is able to choose among the different available paths

depending on which one meets the constraints included into the trigger. When a path is selected, the signaling procedure occurs accordingly with a standard RSVP-TE signaling procedure including a strict (or loose) ERO. For more details please see [ENNI].

6.7. Use case P7 - Dual homing

A quite common case is the need for provisioning two (or more) optical paths with different ingress and egress nodes (i.e. different CNs) with diversity constraints. This is needed to provide two totally disjoint paths in the client layer without any kind of single point of failure (e.g. a single UNI link or a single edge node) in order to create protection schemes in the packet layer like e.g. 1+1 protection.

In this case some sort of info exchange in the client domain is needed and can be provided in two ways:

- Coordination between 1 and 3
- Coordination between 2-1 and 2-3 [UNI EXT]

In both cases node 1 is supposed to ask the optical domain to provide a path (e.g. restorable [RFC4873]) and collect its SRLGs, then pass such SRLGs info to node 3 (directly or via 2). Hence node 3 will ask the optical domain to provide a path (e.g. restorable) avoiding the given SRLGs.




Figure 11: Use Case P7 - Dual homing

As per previous use cases it is possible to either provide the path between 1 and 4, collect the SRLGs, provide the path between 3 and 6

with SRLG diversity constraints and then let the packet domain use those two packet layer adjacencies or have 2 asking 1 and 3 to respectively set up 1-4 and 3-6.

Whenever one of the paths is restored in the optical domain (e.g. B-C fails and the path is restored using A-B-E-C), an SRLG collection procedure (by 1) and exchange (to F) is needed so that a possible restoration of the protection path would avoid the new SRLGs of working path. Please see use case R4.

7. Recovery

7.1. Use case R1 - Segment Recovery - Single homing

This is one of the simplest recovery scenarios, where each overlay node has a single UNI interface available on the overlay nodes and hence only the server layer is in charge of restoring or protecting the traffic within its domains boundaries. As it is possible to see in figure below the client layer is able to recover from failures within the server domain.




Figure 12: Use case R1 - Segment Recovery - Single homing

This kind of scenario is extremely cheap in terms of resource utilization but has some limitations due to a high number of single points of failures which, in case of unavailability, prevent any type of recovery attempt. The single points of failure are: the overlay nodes (R3 and R5), the UNI links (R3-A and C-R5) and the server domain border nodes (A and C). In the following sections higher degrees of reliability are provided.

7.2. Use case R2 - Local recovery - Dual Homing - Single overlay node

In this use case it is possible to combine local recovery in the overlay nodes (e.g. Fast Rerouting) with restoration or protection in the optical domain so to be able to react to failures not only in the client domain but also affecting the server layer border nodes (A and C) and the UNI links (R3-A and R5-C). The number of single points of failure is hence significantly reduced to just the overlay nodes (R3 and R5).




Figure 13: Use case R2 - Local recovery - Dual Homing - Single overlay node

7.3. Use case R3 -End to end Recovery - Dual homing - Double overlay node

This use case focuses on recovery just in the client layer, no recovery is requested to the server layer but just the provisioning of LSPs with requirements (in this case diversity).

The provisioning of the working path occurs like the provisioning of LSP over the UNI with remote trigger and with the collection of SRLGs enabled. Once that the optical LSP supporting the working client LSP is setup (A-B-C), the SRLGs in the optical domain are collected and passed to the overlay node (R3) which forwards them to the ingress node (R1) in the reverse path of the signaling process of the working LSP (i.e. RSVP-TE Resv Message).

Once that R1 is provided with the SRLGs in the optical domains, it is able to issue the signaling of the protection LSP indicating in a loose ERO R5 and R7 as loose hops and the SRLG diversity in the ERO expansions between them. This will lead in the provisioning of a path in the photonic domain diverse from the previous one (S2-S4-S6) and hence two total diverse end to end paths. In this case the

network will be able to recover to a single failure in any part of the network without any single point of failure.




Figure 14: Use case R3 -End to end Recovery - Dual homing - Double overlay node

7.4. Use case R4 - Combination of client protection and server restoration

This is the most interesting use case when the overlay model is applied to the packet/opto scenario. The packet protection and optical restoration allows for a fast recovery of the traffic upon a failure in the network, while the restoration in the optical domain allows for always providing the packet layer with two available path against an arbitrarily high number of failure in the optical network (the number of failures depend on the meshing degree of the photonic network).

Moreover the exchange of SRLGs over the UNI interface (please see Provisioning use cases) makes sure that the optical path used for the working branch of the packet protection and the optical path used for the protection branch of the packet protection do not share any resource and hence that a single failure in the optical domain does not cause any traffic hit; i.e. traffic can be recovered with packet protection time (tens of ms) instead of photonics restoration time (tens of s).

The SRLG collection procedure is pretty simple during the provisioning phase, but more complicated when one of the paths fail and is restored.

In figure below an example is considered where a protection is provided in the client layer where the working path is setup along path (2-1-A-C-4-5) and its protection path along (2-3-F-H-G-6-5).

After the setup of the optical path between A and C (A-B-C), the SRLGs of such path are collected and the overlay edge node 3 can ask F for the setup of a path between F and H avoiding such SRLGs. This ensures that a single failure in the optical domain will not affect both the W and P.




Figure 15: Use Case R4 - Setup of paths in full diversity

In order to guarantee that different SRLGs will always be used for W and P it is necessary to perform two different steps after the provisioning of the two LSPs:

1. Each ingress node of the server domain paths (namely A and F) must be informed that in case of restoration the path computation must avoid the *actual* SRLG of the other path
2. Every time that either of the server domain paths is restored, the collection of SRLGs must be performed and communicated to the ingress node of the other LSP.




Figure 16: Use Case R4 - Restoration of paths keeping full diversity

In the example above a failure affecting W (link B-C) and a failure affecting P (link F-G) occur, but the server network is able to provide alternate paths for both of them and still provide the same level of fault resiliency to the client network.

8. Optimization

8.1. Use case O1 -

TBD

9. Security Considerations

TBD

10. IANA Considerations

TBD

11. Contributors

Diego Caviglia

Email: diego.caviglia@ericsson.com

Francesco Fondelli

Email: francesco.fondelli@ericsson.com

12. Acknowledgements

The authors would like to thank Manuela Scarella and Enrico Dutti for the comments and review.

13. References to be added

[EDITOR NOTE] : section to be fixed
[LSP-DIV] - <http://tools.ietf.org/html/draft-ietf-ccamp-lsp-diversity-01>
[TE-REC] - <http://tools.ietf.org/html/draft-ietf-ccamp-te-metric-recording-01>
[OF-TEMB] - <http://tools.ietf.org/html/draft-ali-ccamp-rc-objective-function-metric-bound-02>
[SRLG-COLL] - <http://tools.ietf.org/html/draft-ietf-ccamp-rsvp-te-srlg-collect-02>
[MELG] - <http://tools.ietf.org/html/draft-beeram-ccamp-melg-00>
[UNI-APP] - <http://tools.ietf.org/html/draft-zhang-ccamp-gmpls-uni-app-03>
[ENNI] - <http://datatracker.ietf.org/doc/draft-beeram-ccamp-gmpls-enni/>
[ISIS-TEMET] - <http://tools.ietf.org/html/draft-previdi-isis-te-metric-extensions-03>
[OSPF-TEMET] - <http://tools.ietf.org/html/draft-ietf-ospf-te-metric-extensions-04>
[INTERCON-TE] - <http://tools.ietf.org/html/draft-farrel-interconnected-te-info-exchange-00>
[UNI EXT] - <http://tools.ietf.org/html/draft-fedyk-ccamp-uni-extensions-01>

14. References

14.1. Normative References

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

14.2. Informative References

Authors' Addresses

Daniele Ceccarelli
Ericsson
Via E. Melen 77
Genova - Erzelli
Italy

Email: daniele.ceccarelli@ericsson.com

Oscar Gonzalez de Dios
Telefonica I+D
Don Ramon de la Cruz 82-84
Madrid 28045
Spain

Email: ogondio@tid.es

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Shenzhen 518129 P.R.China Bantian, Longgang District
Phone: +86-755-28972912

Email: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Shenzhen 518129 P.R.China Bantian, Longgang District
Phone: +86-755-28972913

Email: zhang.xian@huawei.com

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 14, 2013

D. Hiremagalur, Ed.
G. Grammel, Ed.
J. Drake, Ed.
Juniper
G. Galimberti, Ed.
Z. Ali, Ed.
Cisco
R. Kunze, Ed.
Deutsche Telekom
July 15, 2013

Extension to the Link Management Protocol (LMP/DWDM -rfc4209) for Dense
Wavelength Division Multiplexing (DWDM) Optical Line Systems to manage
application code of optical interface parameters in DWDM application
draft-dharinigert-ccamp-g-698-2-lmp-03

Abstract

This memo defines extensions to LMP(rfc4209) for managing Optical parameters associated with Wavelength Division Multiplexing (WDM) systems or characterized by the Optical Transport Network (OTN) in accordance with the Interface Application Code approach defined in ITU-T Recommendation G.698.2.[ITU.G698.2], G.694.1.[ITU.G694.1] and its extensions./>

Copyright Notice

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 14, 2013.

Copyright Notice

Copyright (c) 2013 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	2
2. Extensions to LMP-WDM Protocol	3
3. Black Link General Parameters - BL_General	4
4. Black Link ApplicationCode - BL_ApplicationCode	4
5. Black Link - BL_Ss	5
6. Black Link - BL_Rs	5
7. Security Considerations	6
8. IANA Considerations	6
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Authors' Addresses	8

1. Introduction

This extension is based on "draft-galikunze-ccamp-g-698-2-snmp-mib-03" and "draft-kunze-g-698-2-management-control-framework-02", for the relevant interface optical parameters described in recommendations like ITU-T G.698.2 [ITU.G698.2]. The LMP Model from RFC4902 is extended to provide link property correlation between a client and an OLS device. By using LMP, the capabilities of either end of this link are exchanged where the term 'link' refers to the attachment link between OXC and OLS (see Figure 1). By performing link property correlation, both ends of the link can agree on a common parameter window that can be supported and supervised by each device. The actual selection of a specific parameter value within the parameter window is outside the scope of LMP. In GMPLS the parameter selection (e.g. wavelength) is performed by RSVP-TE and Wavelength routing by IGP.

Figure 1 Extended LMP Model (from [RFC4209])

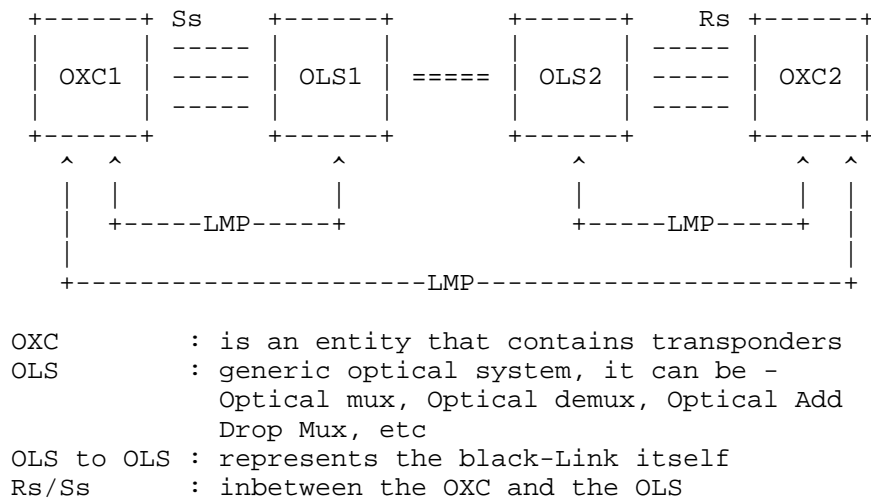


Figure 1: Extended LMP Model

2. Extensions to LMP-WDM Protocol

This document defines extensions to [RFC4209] to allow the Black Link (BL) parameters of G.698.2, as described in the draft draft-kunze-g-698-2-management-control-framework-02, to be exchanged between a router or optical switch and the optical line system to which it is attached. In particular, this document defines additional Data Link sub-objects to be carried in the LinkSummary message defined in [RFC4204] and [RFC6205]. The OXC and OLS systems may be managed by different Network management systems and hence may not know the capability and status of their peer. The intent of this draft is to enable the OXC and OLS systems to exchange this information. These messages and their usage are defined in subsequent sections of this document.

The following new messages are defined for the WDM extension for ITU-T G.698.2 [ITU.G698.2]/ITU-T G.698.1 [ITU.G698.1]/ITU-T G.959.1 [ITU.G959.1]

- BL_General (sub-object Type = TBA)
- BL_ApplicationCode (sub-object Type = TBA)
- BL_Ss (sub-object Type = TBA)
- BL_Rs (sub-object Type = TBA)

3. Black Link General Parameters - BL_General

These are the general parameters as described in [G698.2] and [G.694.1]. Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-03" for more details about these parameters and the [RFC6205] for the wavelength definition.

The general parameters are

1. Bit-Rate/line coding of optical tributary signals
2. Wavelength - (Tera Hertz) 4 bytes (see RFC6205 sec.3.2)

Figure 2: The format of the this sub-object (Type = TBA, Length = TBA) is as follows:

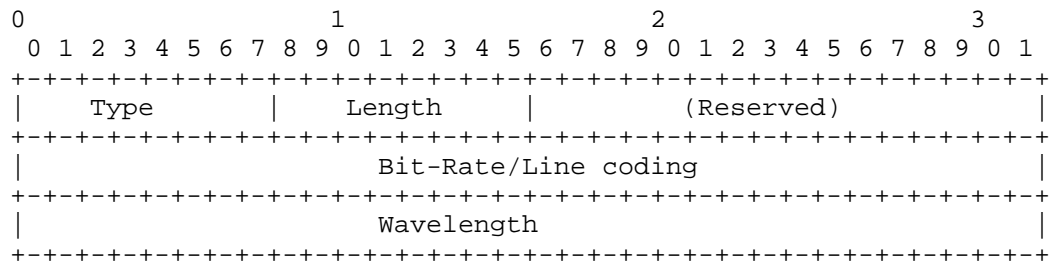


Figure 2: BL_General

4. Black Link ApplicationCode - BL_ApplicationCode

These are the general parameters as described in [G698.2]. Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-03". for more details about these parameters.

The general parameters are

1. Single-channel application codes -- 32 bytes
(from [G698.1]/[G698.2]/[G959.1])

Figure 3: The format of the this sub-object (Type = TBA, Length = TBA) is as follows:

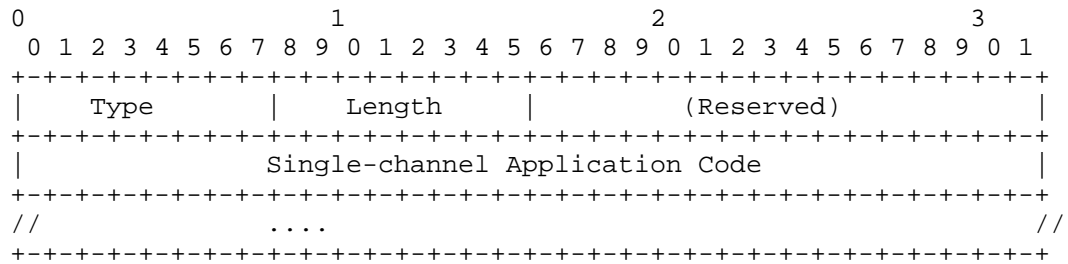


Figure 3: BL_ApplicationCode

5. Black Link - BL_Ss

These are the G.698.2 parameters at the Source(Ss reference points). Please refer to "draft-galikunze-ccamp-g-698-2-snmp-mib-03" for more details about these parameters.

1. Minimum Mean Channel Output Power -(0.1 dbm) 4 bytes
2. Maximum Mean Channel Output Power -(0.1 dbm) 4 bytes

Figure 4: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

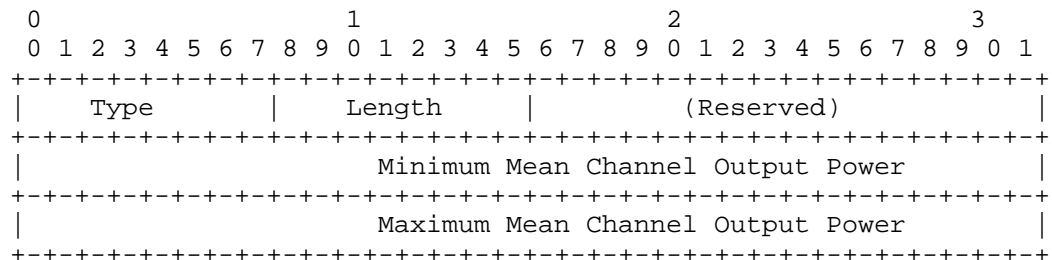


Figure 4: Black Link - BL_Ss

6. Black Link - BL_Rs

These are the G.698.2 parameters at the Sink (Rs reference points). Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-02" for more details about these parameters.

1. Minimum Mean Input Power - (0.1dbm) 4bytes
2. Maximum Mean Input Power - (0.1dbm) 4bytes
3. Minimum OSNR - (0.1dB) 4bytes
4. OSNR Tolerance - (0.1dB) 4bytes
5. Current Input Power - (0.1dbm) 4bytes

Figure 5: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

The format of the Black Link/OLS Sink sub-object (Type = TBA, Length = TBA) is as follows:

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										(Reserved)																			
Minimum Mean Input Power																																							
Maximum Mean Input Power																																							
Minimum OSNR																																							
OSNR Tolerance																																							
Current Input Power																																							

Figure 5: Black Link - BL_Rs

7. Security Considerations

LMP message security uses IPsec, as described in [RFC4204]. This document only defines new LMP objects that are carried in existing LMP messages, similar to the LMP objects in [RFC:4209]. This document does not introduce new security considerations.

8. IANA Considerations

LMP `<xref target="RFC4204"/>` defines the following name spaces and the ways in which IANA can make assignments to these namespaces:

- LMP Message Type
- LMP Object Class
- LMP Object Class type (C-Type) unique within the Object Class
- LMP Sub-object Class type (Type) unique within the Object Class

This memo introduces the following new assignments:

LMP Sub-Object Class names:

under DATA_LINK Class name (as defined in <xref target="RFC4204"/>)

- BL_General (sub-object Type = TBA)
- BL_Applicationcode (sub-object Type = TBA)
- BL_Ss (sub-object Type = TBA)
- BL_Rs (sub-object Type = TBA)

9. References

9.1. Normative References

- [RFC4204] Lang, J., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4209] Fredette, A. and J. Lang, "Link Management Protocol (LMP) for Dense Wavelength Division Multiplexing (DWDM) Optical Line Systems", RFC 4209, October 2005.
- [RFC6205] Otani, T. and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.
- [RFC4054] Strand, J. and A. Chiu, "Impairments and Other Constraints on Optical Layer Routing", RFC 4054, May 2005.
- [ITU.G698.2] International Telecommunications Union, "Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces ", ITU-T Recommendation G.698.2, November 2009.
- [ITU.G694.1] International Telecommunications Union, "Spectral grids for WDM applications: DWDM frequency grid " , ITU-T Recommendation G.698.2, February 2012.
- [ITU.G709] International Telecommunications Union, "Interface for the Optical Transport Network (OTN) ", ITU-T Recommendation G.709, March 2003.
- [ITU.G872]

International Telecommunications Union, "Architecture of optical transport networks ", ITU-T Recommendation G.872, November 2001.

9.2. Informative References

[I-D.kunze-g-698-2-management-control-framework]

Kunze, R., "A framework for Management and Control of optical interfaces supporting G.698.2", draft-kunze-g-698-2-management-control-framework-00 (work in progress), July 2011.

[I-D.galimbe-kunze-g-698-2-snmp-mib]

Kunze, R. and D. Hiremagalur, "A SNMP MIB to manage black-link optical interface parameters of DWDM applications", draft-galimbe-kunze-g-698-2-snmp-mib-02 (work in progress), March 2012.

Authors' Addresses

Dharini Hiremagalur (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: dharinih@juniper.net

Gert Grammel (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: ggrammel@juniper.net

John E. Drake (editor)
Juniper
1194 N Mathilda Avenue
HW-US, Pennsylvania
USA

Phone: +1408
Email: jdrake@juniper.net

Gabriele Galimberti (editor)
Cisco
Via Philips,12
20052 - Monza
Italy

Phone: +390392091462
Email: ggalimbe@cisco.com

Zafar Ali (editor)
Cisco
3000 Innovation Drive
KANATA
ONTARIO K2K 3E8

Email: zali@cisco.com

Ruediger Kunze (editor)
Deutsche Telekom
Dddd, xx
Berlin
Germany

Phone: +49xxxxxxxxxxx
Email: RKunze@telekom.de

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 14, 2013

D. Hiremagalur, Ed.
G. Grammel, Ed.
J. Drake, Ed.
Juniper
G. Galimberti, Ed.
Z. Ali, Ed.
Cisco
R. Kunze, Ed.
Deutsche Telekom
July 15, 2013

Extension to the Link Management Protocol (LMP/DWDM -rfc4209) for Dense
Wavelength Division Multiplexing (DWDM) Optical Line Systems to manage
optical parameters of DWDM applications
draft-dharinigert-ccamp-opt-imp-lmp-00

Abstract

This memo defines extensions to LMP(rfc4209) for managing Optical parameters associated with Wavelength Division Multiplexing (WDM) systems or characterized by the Optical Transport Network (OTN) in accordance with the parameters defined in ITU-T Recommendation G.698.2.[ITU.G698.2], G.694.1.[ITU.G694.1] and their extensions./>

Copyright Notice

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 14, 2013.

Copyright Notice

Copyright (c) 2013 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	2
2. Extensions to LMP-WDM Protocol	3
3. Optical General Parameters - OP_General	4
4. Vendor Transceiver Class - OP_VendorClass	5
5. Optical Parameters - OP_Ss	5
6. Optical Parameters - OP_SsRs	6
7. Optical Parameters - OP_Rs	7
8. DWDM Network - OLS_Status	8
9. Security Considerations	9
10. IANA Considerations	9
11. References	10
11.1. Normative References	10
11.2. Informative References	10
Authors' Addresses	11

1. Introduction

This extension is based on "draft-galikunze-ccamp-opt-imp-snmp-mib-00" and "draft-kunze-g-698-2-management-control-framework-02", for the relevant optical parameters mainly (but not restricted to) described in recommendations like ITU-T G.698.2 [ITU.G698.2] and G.694.1 [ITU.G694.1]. This document further extend the draft-dharinigert-ccamp-g-698-2-lmp-03. The LMP Model from RFC4902 is extended to provide link property correlation between a client and an OLS device. By using LMP, the capabilities of either end of this link are exchanged where the term 'link' refers to the attachment link between OXC and OLS (see Figure 1). By performing link property correlation, both ends of the link can agree on a common parameter window that can be supported and supervised by each device. The actual selection of a specific parameter value within the parameter window is outside the scope of LMP. In GMPLS the parameter selection

(e.g. wavelength) is performed by RSVP-TE and Wavelength routing by IGP. This document does not include the ITU-T G.698.2 "application-code" and few other parameters described in the draft draft-dharinigert-ccamp-g-698-2-lmp-03

Figure 1 Extended LMP Model (from [RFC4209])

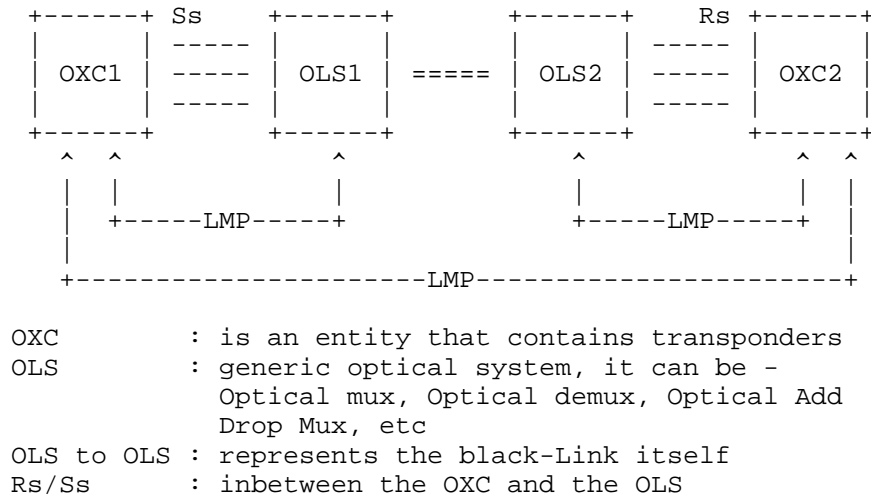


Figure 1: Extended LMP Model

2. Extensions to LMP-WDM Protocol

This document defines extensions to [RFC4209] to allow the Black Link (BL) parameters of G.698.2, as described in the draft draft-kunze-g-698-2-management-control-framework-02, to be exchanged between a router or optical switch and the optical line system to which it is attached. In particular, this document defines additional Data Link sub-objects to be carried in the LinkSummary message defined in [RFC4204] and [RFC6205]. The OXC and OLS systems may be managed by different Network management systems and hence may not know the capability and status of their peer. The intent of this draft is to enable the OXC and OLS systems to exchange this information. These messages and their usage are defined in subsequent sections of this document.

The following new messages are defined for the WDM extension for ITU-T G.698.2 [ITU.G698.2]/ITU-T G.698.1 [ITU.G698.1]/ITU-T G.959.1 [ITU.G959.1]

- BL_General (sub-object Type = TBA)
- BL_Ss (sub-object Type = TBA)
- BL_SsRs (sub-object Type = TBA)
- BL_Rs (sub-object Type = TBA)
- BL_OLS_Status (sub-object Type = TBA)

3. Optical General Parameters - OP_General

These are the general parameters as described in [G698.2] and [G.694.1]. Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-02" for more details about these parameters and the [RFC6205] for the wavelength definition.

The general parameters are

1. Min Wavelength Range - (Tera Hertz) 4 bytes (see RFC6205 sec.3.2 and 3.3 TLV): Grid/Cannel Spacing/Identifier/n
2. Max Wavelength Range - (Tera Hertz) 4 bytes (see RFC6205 sec.3.2 and 3.3 TLV): Grid/Cannel Spacing/Identifier/n
3. BER mantisa - 4 bytes
4. BER exponent - 4 bytes
5. FEC Coding - 1 byte
6. Administrative state - 1 byte
7. Operation state - 1 byte

Figure 2: The format of the this sub-object (Type = TBA, Length = TBA) is as follows:




Figure 2: OP_General

4. Vendor Transceiver Class - OP_VendorClass

These are the general parameters as described in [G698.2]. Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-02". for more details about these parameters.

The general parameters are

1. Vendor Transceiver Class -- 32 bytes
When Single-channel application code (which is defined in G.698.2) is used in the message, then the vendor transaction class need not be used (i.e. all 0) and the optional parameters except nominal central frequency need not be used. When vendor transaction class is used in the message, then the Single-channel application code need not be used (i.e. all 0) and the optional parameters needs to be used.

Figure 3: The format of the this sub-object (Type = TBA, Length = TBA) is as follows:




Figure 3: Vendor Transceiver Class

5. Optical Parameters - OP_Ss

These are the G.698.2 parameters at the Source(Ss reference points). Please refer to "draft-galikunze-ccamp-g-698-2-snmp-mib-02" for more details about these parameters.

1. Minimum Central Frequency - (THz) 4 bytes (see RFC6205 sec.3.2 and 3.3 TLV): Grid / Cannel Spacing / Identifier / n
2. Maximum Central Frequency - (THz) 4 bytes (see RFC6205 sec.3.2 and 3.3 TLV): Grid / Cannel Spacing / Identifier / n

3. Maximum Spectral Excursion - (0.1 GHz) 4 bytes
4. Maximum Tx Dispersion OSNR Penalty - (0.1 dbm) 4 bytes
5. Current Output Power - (0.1dbm) 4 bytes
6. Status of TX - Status of the Transmit link at OXC - 4 bytes

Figure 4: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

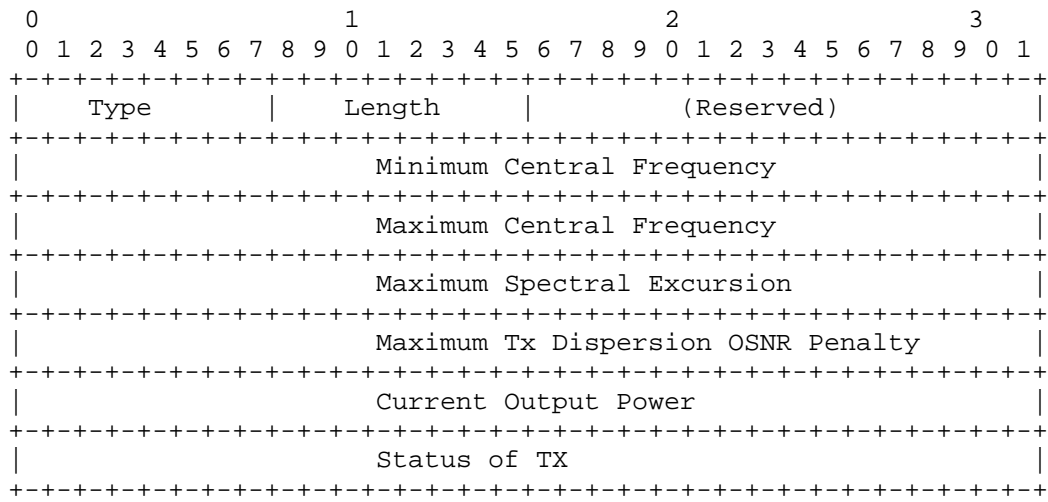


Figure 4: Optical Parameters - OP_Ss

6. Optical Paramenters - OP_SsRs

These are the G.698.2 parameters for the path (Ss-Rs). Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-02" for more details about these parameters.

1. Minimum Chromatic Dispersion - (ps/nm) 4 bytes
2. Maximum Chromatic Dispersion -(ps/nm) 4 bytes
3. Minimum Src Optical ReturnLoss -(0.1 db) 4 bytes
4. Maximum Discrete Reflectance Src To Sink - (0.1 db) 4 bytes
5. Maximum Differential Group Delay - (ps) 4 bytes
6. Maximum Polarisation Dependent Loss - (0.1 db) 4 bytes
7. Maximum Inter Channel Crosstalk - (0.1 db) 4 bytes
8. Interferometric Crosstalk - (0.1 db) 4 bytes
9. Optical Path OSNR Penalty - (0.1 db) 4 bytes
10. Fiber type - 1 byte

Figure 5: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

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										(Reserved)																			
										Minimum Chromatic Dispersion																													
										Maximum Chromatic Dispersion																													
										Minimum Src Optical ReturnLoss																													
										Maximum Discrete Reflectance Src To Sink																													
										Maximum Differential Group Delay																													
										Maximum Polarisation Dependent Loss																													
										Maximum Inter Channel Crosstalk																													
										Interferometric Crosstalk																													
										Optical Path OSNR Penalty																													
Fiber Type										Reserved																													

Figure 5: Optical Parameters - OP_SsRs

7. Optical Parameters - OP_Rs

These are the G.698.2 parameters at the Sink (Rs reference points). Please refer to the "draft-galikunze-ccamp-g-698-2-snmp-mib-02" for more details about these parameters.

1. Threshold of the input power at OLS - The power level above which the OLS will not function (0.1dbm) 4bytes
2. Current Optical OSNR (0.1dB)
3. Q factor
4. Post FEC BER Mantissa
5. Post FEC BER Exponent
6. Status of RX - Status of the Receive link at OXC - 2bytes

Figure 6: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

The format of the Black Link/OLS Sink sub-object (Type = TBA, Length = TBA) is as follows:

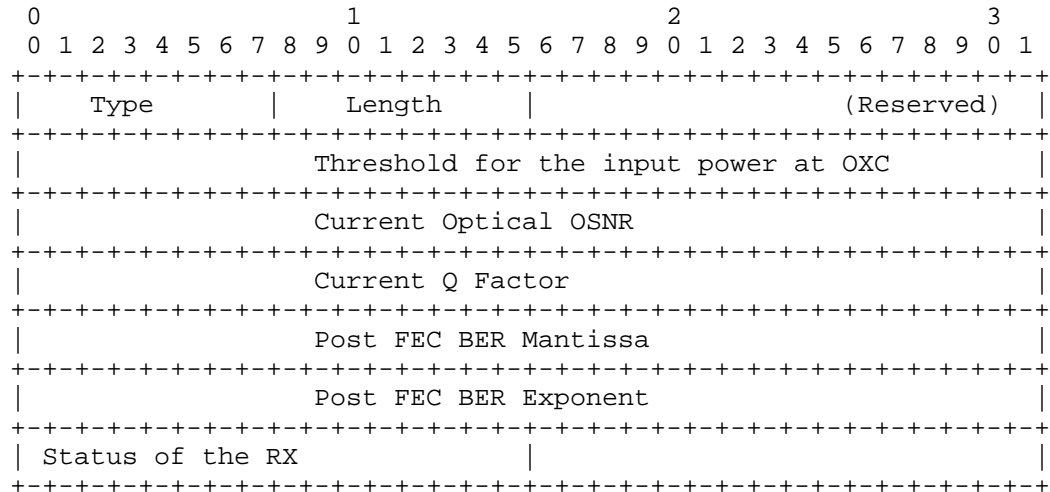


Figure 6: Optical Parameters - OP_Rs

8. DWDM Network - OLS_Status

This message is sent by the OLS to the OXC. It includes the wavelength information and the status of the OLS.

1. Wavelength - The wavelength which has been accepted by the OLS (Tera Hertz) 4 bytes. (see RFC6205 sec.3.2 and 3.3 TLV): Grid / Cannel Spacing / Identifier / n
2. Length of the Wavelength Availability Map 1 byte
3. Wavelength Availability bits - variable bits depending on the nunmer of wavelengths available (For eg 96 bits for C-band 50GHz) (Allocation is in multiples of 1byte - 96 bits - 12 bytes)
0 - wavelength is available, 1 - used - variable length
4. Current Input Power (0.1dbm) 4 bytes
- This is the current input power at OLS
5. Delta between output power at the Src(OXC)and Input Power at OLS (0.1dbm) 4 bytes
- This is the delta between the input power and the transmitted output power at the OXC (from message 2.2 BL_Src)
6. Threshold of the input power at OLS 4 bytes
- This is the power level above which the OLS will not function.
7. Current Output Power (0.1dbm) 4 bytes

- This is the transmitted output power at the OLS.
- 8. Status of Rx link at OLS 2 bytes
 - Status of the Receive link at the OLS
- 9. Status of Tx link at OLS 2 bytes
 - Status of the Transmit link at the OLS

Figure 7: The format of the Black link sub-object (Type = TBA, Length = TBA) is as follows:

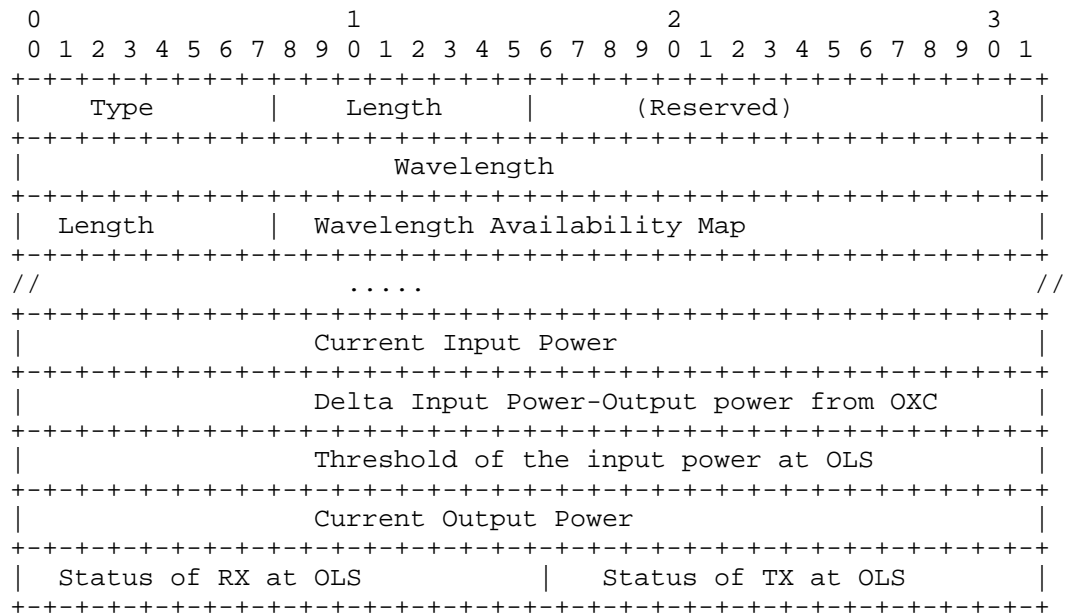


Figure 7: DWDM Network - OLS_Status

9. Security Considerations

LMP message security uses IPsec, as described in [RFC4204]. This document only defines new LMP objects that are carried in existing LMP messages, similar to the LMP objects in [RFC:4209]. This document does not introduce new security considerations.

10. IANA Considerations

LMP [RFC4204] defines the following name spaces and the ways in which IANA can make assignments to these namespaces: - LMP Message Type - LMP Object Class - LMP Object Class type (C-Type) unique within the Object Class - LMP Sub-object Class type (Type) unique within the

Object Class This memo introduces the following new assignments: LMP Sub-Object Class names: under DATA_LINK Class name (as defined in [RFC4204]) - BL_General (sub-object Type = TBA) - BL_Applicationcode (sub-object Type = TBA) - BL_Ss (sub-object Type = TBA) - BL_SsRs (sub-object Type = TBA) - BL_Rs (sub-object Type = TBA) - OLS_Status (sub-object Type = TBA)

11. References

11.1. Normative References

- [RFC4204] Lang, J., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4209] Fredette, A. and J. Lang, "Link Management Protocol (LMP) for Dense Wavelength Division Multiplexing (DWDM) Optical Line Systems", RFC 4209, October 2005.
- [RFC6205] Otani, T. and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.
- [RFC4054] Strand, J. and A. Chiu, "Impairments and Other Constraints on Optical Layer Routing", RFC 4054, May 2005.
- [ITU.G698.2] International Telecommunications Union, "Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces ", ITU-T Recommendation G.698.2, November 2009.
- [ITU.G694.1] International Telecommunications Union, "Spectral grids for WDM applications: DWDM frequency grid" ", ITU-T Recommendation G.698.2, February 2012.
- [ITU.G709] International Telecommunications Union, "Interface for the Optical Transport Network (OTN) ", ITU-T Recommendation G.709, March 2003.
- [ITU.G872] International Telecommunications Union, "Architecture of optical transport networks ", ITU-T Recommendation G.872, November 2001.

11.2. Informative References

[I-D.kunze-g-698-2-management-control-framework]

Kunze, R., "A framework for Management and Control of optical interfaces supporting G.698.2", draft-kunze-g-698-2-management-control-framework-00 (work in progress), July 2011.

[I-D.galimbe-kunze-g-698-2-snmp-mib]

Kunze, R. and D. Hiremagalur, "A SNMP MIB to manage black-link optical interface parameters of DWDM applications", draft-galimbe-kunze-g-698-2-snmp-mib-02 (work in progress), March 2012.

Authors' Addresses

Dharini Hiremagalur (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: dharinih@juniper.net

Gert Grammel (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: ggrammel@juniper.net

John E. Drake (editor)
Juniper
1194 N Mathilda Avenue
HW-US, Pennsylvania
USA

Phone: +1408
Email: jdrake@juniper.net

Gabriele Galimberti (editor)
Cisco
Via Philips,12
20052 - Monza
Italy

Phone: +390392091462
Email: ggalimbe@cisco.com

Zafar Ali (editor)
Cisco
3000 Innovation Drive
KANATA
ONTARIO K2K 3E8

Email: zali@cisco.com

Ruediger Kunze (editor)
Deutsche Telekom
Dddd, xx
Berlin
Germany

Phone: +49xxxxxxxxxxx
Email: RKunze@telekom.de

CCAMP Working Group
Internet-Draft
Intended status: Experimental
Expires: January 10, 2014

D. Dhody
U. Palle
V. Kondreddy
Huawei Technologies
R. Casellas
CTTC
July 9, 2013

Domain Subobjects for Resource ReserVation Protocol - Traffic
Engineering (RSVP-TE)
draft-dhody-ccamp-rsvp-te-domain-subobjects-02

Abstract

The Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) specification [RFC3209] and the Generalized Multiprotocol Label Switching (GMPLS) extensions to RSVP-TE [RFC3473] allow abstract nodes and resources to be explicitly included in a path setup. Further Exclude Routes extensions [RFC4874] allow abstract nodes and resources to be explicitly excluded in a path setup.

This document specifies new subobjects to include or exclude domains during path setup where domain is a collection of network elements within a common sphere of address management or path computational responsibility (such as an Interior Gateway Protocol (IGP) area or an Autonomous System (AS)). Note that the use of AS as an abstract node representing domain is already defined in [RFC3209] and [RFC4874], albeit with a 2-Byte AS number.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 10, 2014.

Copyright Notice

Copyright (c) 2013 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. Requirements Language	3
2. Terminology	3
3. Subobjects for Domains	4
3.1. Domains	4
3.2. Explicit Route Object (ERO)'s Subobjects	5
3.2.1. Autonomous system	6
3.2.2. IGP Area	6
3.2.3. Mode of Operation	7
3.3. Exclude Route Object (XRO)'s Subobjects	8
3.3.1. Autonomous system	8
3.3.2. IGP Area	9
3.3.3. Mode of Operation	10
3.4. Explicit Exclusion Route Subobject	10
4. Interaction with Path Computation Element (PCE)	11
5. Examples	11
5.1. Inter-Area LSP Path Setup	11
5.2. Inter-AS LSP Path Setup	12
5.2.1. Example 1	12
5.2.2. Example 2	13
6. IANA Considerations	15
6.1. New Subobjects	15
7. Security Considerations	15
8. Acknowledgments	15
9. References	15
9.1. Normative References	15
9.2. Informative References	16

1. Introduction

The RSVP-TE specification [RFC3209] and the GMPLS extensions to RSVP-TE [RFC3473] allow abstract nodes and resources to be explicitly included in a path setup using the Explicit Route Object (ERO).

Further Exclude Routes extensions [RFC4874] allow abstract nodes or resources to be excluded from the whole path using the Exclude Route object (XRO). To exclude certain abstract nodes or resources between a specific pair of abstract nodes present in an ERO, a Explicit Exclusion Route Subobject (EXRS) is used.

[RFC3209] already describes the notion of abstract nodes, where an abstract node is a group of nodes whose internal topology is opaque to the ingress node of the Label Switched Path (LSP). It further defines a subobject for AS, but with a 2-Byte AS number only.

This document extends the notion of abstract nodes by adding new subobjects for IGP Areas and 4-byte AS numbers (as per [RFC4893]). These subobjects MAY be included in Explicit Route Object (ERO), Exclude Route object (XRO) or Explicit Exclusion Route Subobject (EXRS).

In case of per-domain path computation [RFC5152], where the full path of an inter-domain TE LSP cannot be or is not determined at the ingress node, and signaling message may use domain identifiers. The use of these new subobjects is illustrated in Section 5.

1.1. 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 [RFC2119].

2. Terminology

The following terminology is used in this document.

AS: Autonomous System.

Domain: As per [RFC4655], any collection of network elements within a common sphere of address management or path computational responsibility. Examples of domains include Interior Gateway Protocol (IGP) areas and Autonomous Systems (ASs).

ERO: Explicit Route Object

EXRS: Explicit Exclusion Route Subobject

IGP: Interior Gateway Protocol. Either of the two routing protocols, Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS).

IS-IS: Intermediate System to Intermediate System.

OSPF: Open Shortest Path First.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Protocol.

RSVP: Resource Reservation Protocol

TE LSP: Traffic Engineering Label Switched Path.

XRO: Exclude Route Object

3. Subobjects for Domains

3.1. Domains

[RFC4726] and [RFC4655] define domain as a separate administrative or geographic environment within the network. A domain may be further defined as a zone of routing or computational ability. Under these definitions a domain might be categorized as an AS or an IGP area.

As per [RFC3209], an abstract node is a group of nodes whose internal topology is opaque to the ingress node of the LSP. Using this concept of abstraction, an explicitly routed LSP can be specified as a sequence of IP prefixes or a sequence of Autonomous Systems. In this document we extend the notion to include IGP area and 4-Byte AS number.

The sub-objects MAY appear in RSVP-TE, notably in -

- o Explicit Route Object (ERO): As per [RFC3209], an explicit route is a particular path in the network topology including abstract nodes (domains).

- o Exclude Route Object (XRO): As per [RFC4874], an exclude route identifies a list of abstract nodes (domains) that should not be traversed along the path of the LSP being established.
- o Explicit Exclusion Route Subobject (EXRS): As per [RFC4874], used to specify exclusion of certain abstract nodes between a specific pair of nodes. EXRS are a subobject carried inside the ERO. These subobjects are used to specify the domains that must be excluded between two abstract nodes.

3.2. Explicit Route Object (ERO)'s Subobjects

As stated in [RFC3209], an explicit route is a particular path in the network topology. In addition to the ability to identify specific nodes along the path, an explicit route can identify a group of nodes (abstract nodes) that must be traversed along the path.

Some subobjects are defined in [RFC3209], [RFC3473], [RFC3477], [RFC4874] and [RFC5553] but new subobjects related to domains are needed.

The following subobject types are used in ERO.

Type	Subobject
1	IPv4 prefix
2	IPv6 prefix
3	Label
4	Unnumbered Interface ID
32	Autonomous system number (2 Byte)
33	Explicit Exclusion (EXRS)
34	SRLG
64	IPv4 Path Key
65	IPv6 Path Key


This document extends the above list to support 4-Byte AS numbers and IGP Areas.

Type	Subobject
TBD	Autonomous system number (4 Byte)
TBD	OSPF Area id
TBD	ISIS Area id

3.2.1. Autonomous system

[RFC3209] already defines 2-Byte AS number.

To support 4-Byte AS numbers as per [RFC4893], the following subobject is defined:



L: The L bit is an attribute of the subobject as defined in [RFC3209].

Type: (TBA by IANA) indicating a 4-Byte AS Number.

Length: 8 (Total length of the subobject in bytes).


Reserved: Zero at transmission, ignored at receipt.

AS-ID: The 4-Byte AS Number. Note that if 2-Byte AS numbers are in use, the low order bits (16 through 31) should be used and the high order bits (0 through 15) should be set to zero.

3.2.2. IGP Area

Since the length and format of Area-id is different for OSPF and ISIS, the following two subobjects are defined:

For OSPF, the area-id is a 32 bit number. The subobject is encoded as follows:



L: The L bit is an attribute of the subobject as defined in [RFC3209].

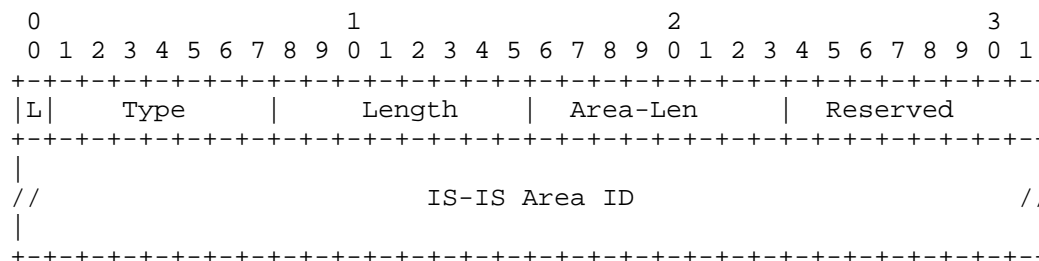
Type: (TBA by IANA) indicating a 4-Byte OSPF Area ID.

Length: 8 (Total length of the subobject in bytes).

Reserved: Zero at transmission, ignored at receipt.

OSPF Area Id: The 4-Byte OSPF Area ID.

For IS-IS, the area-id is of variable length and thus the length of the subobject is variable. The Area-id is as described in IS-IS by ISO standard [ISO 10589]. The subobject is encoded as follows:



L: The L bit is an attribute of the subobject as defined in [RFC3209].

Type: (TBA by IANA) indicating IS-IS Area ID.

Length: Variable. As per [RFC3209], the total length of the subobject in bytes, including the L, Type and Length fields. The Length MUST be at least 4, and MUST be a multiple of 4.

Area-Len: Variable (Length of the actual (non-padded) IS-IS Area Identifier in octets; Valid values are from 2 to 11 inclusive).

Reserved: Zero at transmission, ignored at receipt.

IS-IS Area Id: The variable-length IS-IS area identifier. Padded with trailing zeroes to a four-byte boundary.

3.2.3. Mode of Operation

The new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY also be used in the ERO to specify an abstract node (a group of nodes whose internal topology is opaque to the ingress node of the

LSP).

All the rules of processing (for example Next Hop Selection, L bit processing, unrecognized subobjects etc) are as per the [RFC3209].

3.3. Exclude Route Object (XRO)'s Subobjects

As stated in [RFC4874], the exclude route identifies a list of abstract nodes that should not be traversed along the path of the LSP being established.

Some subobjects are defined in [RFC3209], [RFC3477], [RFC4874] and [RFC6001] but new subobjects related to domains are needed.

The following subobject types are used in XRO.

Type	Subobject
1	IPv4 prefix
2	IPv6 prefix
3	Label
4	Unnumbered Interface ID
32	Autonomous system number (2 Byte)
34	SRLG


This document extends the above list to support 4-Byte AS numbers and IGP Areas.

Type	Subobject
TBD	Autonomous system number (4 Byte)
TBD	OSPF Area id
TBD	ISIS Area id

3.3.1. Autonomous system

[RFC3209] and [RFC4874] already define a 2-Byte AS number.

To support 4-Byte AS numbers as per [RFC4893], the following subobject is defined:



The meaning of the L bit, similar to [RFC4874], is as follows:

0: indicates that the abstract node (AS) specified MUST be excluded.


1: indicates that the abstract node (AS) specified SHOULD be avoided.

The meaning of all the other elements (Type, Length, Reserved and 4-Byte AS Id) is same as explained above in Section 3.2.1.

3.3.2. IGP Area

Since the length and format of Area-id is different for OSPF and ISIS, the following two subobjects are defined:

For OSPF, the area-id is a 32 bit number. The subobject is encoded as follows:



The meaning of the L bit, similar to [RFC4874], is as follows:

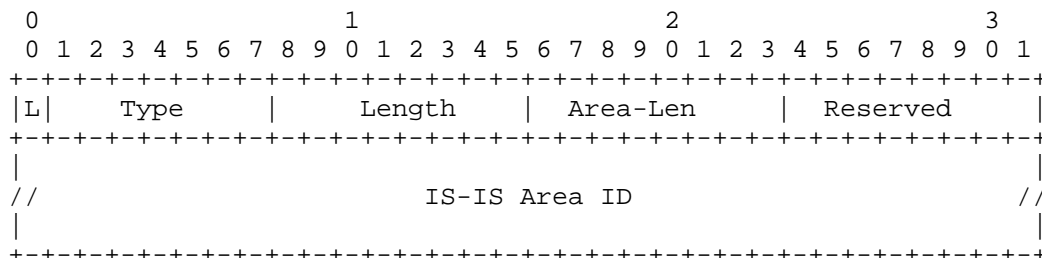
0: indicates that the abstract node (OSPF Area) specified MUST be excluded.

1: indicates that the abstract node (OSPF Area) specified SHOULD be avoided.

The meaning of all the other elements (Type, Length, Reserved and OSPF Area Id) is same as explained above in Section 3.2.2.

For IS-IS, the area-id is of variable length and thus the length of the subobject is variable. The Area-id is as described in IS-IS by

ISO standard [ISO 10589]. The subobject is encoded as follows:



The meaning of the L bit, similar to [RFC4874], is as follows:

0: indicates that the abstract node (IS-IS Area) specified MUST be excluded.

1: indicates that the abstract node (IS-IS Area) specified SHOULD be avoided.

The meaning of all the other elements (Type, Length, Area-Len, Reserved and IS-IS Area Id) is same as explained above in Section 3.2.2.

3.3.3. Mode of Operation

The new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY also be used in the XRO to specify exclusion of an abstract node (a group of nodes whose internal topology is opaque to the ingress node of the LSP).

All the rules of processing are as per the [RFC4874].

3.4. Explicit Exclusion Route Subobject

As per [RFC4874], the Explicit Exclusion Route defines abstract nodes or resources that must not or should not be used on the path between two inclusive abstract nodes or resources in the explicit route. EXRS is an ERO subobject that contains one or more subobjects of its own, called EXRS subobjects.

The EXRS subobject may carry any of the subobjects defined for XRO, thus the new subobjects to support 4-Byte AS and IGP (OSPF / ISIS) Area MAY also be used in the EXRS. The meanings of the fields of the new XRO subobjects are unchanged when the subobjects are included in an EXRS, except that scope of the exclusion is limited to the single hop between the previous and subsequent elements in the ERO.

All the rules of processing are as per the [RFC4874].


4. Interaction with Path Computation Element (PCE)

The domain subobjects to be used in Path Computation Element Protocol (PCEP) are referred to in [PCE-DOMAIN]. Note that the new domain subobjects follow the principle that subobjects used in PCEP [RFC5440] are identical to the subobjects used in RSVP-TE.

5. Examples

5.1. Inter-Area LSP Path Setup

In an inter-area LSP path setup where the ingress and the egress belong to different IGP areas within the same AS, the domain subobjects MAY be represented using an ordered list of IGP area subobjects in an ERO. The AS number MAY be skipped, as area information is enough to uniquely identify a domain.



* All IGP Area in one AS (AS 100)

Figure 1: Domain Corresponding to IGP Area

As per Figure 1, the signaling at Ingress MAY be -

ERO:(A1, ABF1, Area B, Area C, Egress); or


ERO:(A1, ABF1, AS 100, Area B, AS 100, Area C, Egress).

The AS subobject is optional and it MAY be skipped. An RSVP-TE implementation should be able to understand both notations and there is no change in the processing rules as mentioned in [RFC3209].

5.2. Inter-AS LSP Path Setup

5.2.1. Example 1

In an inter-AS LSP path setup where the ingress and the egress belong to different AS, the domain subobjects MAY be represented using an ordered list of AS subobjects in an ERO.



* All AS have one area (area 0)

Figure 2: Domain Corresponding to AS

As per Figure 2, the signaling at Ingress MAY be -

ERO:(A1, A2, AS B, AS C, Egress); or

ERO:(A1, A2, AS B, Area 0, AS C, Area 0, Egress).

Each AS has a single IGP area (area 0), Area subobject is optional and it MAY be skipped as AS is enough to uniquely identify a domain. An RSVP-TE implementation should be able to understand both notations and there is no change in the processing rules as mentioned in [RFC3209].

Note that to get a domain disjoint path, the ingress may also signal the backup path with -

XRO:(AS B)

5.2.2. Example 2

As shown in Figure 3, where AS 200 is made up of multiple areas, the signaling MAY include both AS and Area subobject to uniquely identify a domain.




Figure 3: Domain Corresponding to AS and Area

As per Figure 3, the signaling at Ingress MAY be -

ERO:(X1, AS 200, Area D, Area B, Area C, Egress).

The combination of both an AS and an Area uniquely identifies a domain, note that an Area domain identifier always belongs to the previous AS that appears before it or, if no AS subobjects are present, it is assumed to be the current AS. Also note that there are no changes in the processing rules as mentioned in [RFC3209] with

respect to subobjects.

6. IANA Considerations

6.1. New Subobjects

IANA registry: RSVP PARAMETERS

Subsection: Class Names, Class Numbers, and Class Types

IANA is requested to add further subobjects to the existing entry for:

20	EXPLICIT_ROUTE
232	EXCLUDE_ROUTE

Subobject	Type	Reference
TBA	4-Byte AS number	[This I.D.]
TBA	OSPF Area ID	[This I.D.]
TBA	IS-IS Area ID	[This I.D.]

7. Security Considerations

Security considerations for MPLS-TE and GMPLS signaling are covered in [RFC3209] and [RFC3473]. This document does not introduce any new messages or any substantive new processing, and so those security considerations continue to apply.

The route exclusion security consideration are covered in [RFC4874] and continue to apply.

8. Acknowledgments

We would like to thank Lou Berger, George Swallow, Chirag Shah, Reeja Paul and Sandeep Boina for their useful comments and suggestions.

9. References

9.1. Normative References

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

9.2. Informative References

- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4726] Farrel, A., Vasseur, J., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.
- [RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", RFC 4893, May 2007.
- [RFC5152] Vasseur, JP., Ayyangar, A., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.
- [RFC5553] Farrel, A., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, May 2009.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiimoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.

- [PCE-DOMAIN] Dhody, D., Palle, U., and R. Casellas, "Standard Representation Of Domain Sequence. (draft-ietf-pce-pcep-domain-sequence)", July 2013.
- [ISO 10589] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002.

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: dhruv.dhody@huawei.com

Udayasree Palle
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: udayasree.palle@huawei.com

Venugopal Reddy Kondreddy
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
INDIA

EMail: venugopalreddyk@huawei.com

Ramon Casellas
CTTC
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
SPAIN

EMail: ramon.casellas@cttc.es

Network Working Group
Internet Draft
Intended status: Standards Track

D. Fedyk
D. Beller
Lieven Levrau
Alcatel-Lucent
D. Ceccarelli
Ericsson
F. Zhang
Huawei Technologies
Y. Tochio
Fujitsu
X. Fu
ZTE

Expires: January 16, 2014

July 15, 2013

UNI Extensions for Diversity and Latency Support
draft-fedyk-ccamp-uni-extensions-02.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright Notice

Copyright (c) 2013 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.

Abstract

This document builds on the GMPLS overlay model [RFC4208] and defines extensions to the GMPLS User-Network Interface (UNI) to support route diversity within the core network for sets of LSPs initiated by edge nodes. A particular example where route diversity within the core network is desired, are dual-homed edge nodes. The document also defines GMPLS UNI extensions to deal with latency requirements for edge node initiated LSPs.

This document uses a VPN model that is based on the same premise as L1VPN framework [RFC4847] but may also be applied to other technologies. The extensions are applicable both to VPN and non VPN environments. These extensions move the UNI from basic connectivity to enhanced mode connectivity by including additional constraints while minimizing the exchange of CE to PE information. These extensions are applicable to the overlay extension service model. Route Diversity for customer LSPs are a common requirement applicable to L1VPNs. The UNI mechanisms described in this document are L1VPN compatible and can be applied to achieve diversity for sets of customer LSPs.

The UNI extensions in support of latency constraints can also be applied to the extended overlay service model in order for the customer LSPs to meet certain latency requirements.

Table of Contents

1. Introduction	3
2. Conventions used in this document	4
3. Contributors	4
4. LSP Diversity in the Overlay Extension Service Model	4
4.1. LSP diversity for dual-homed customer edge (CE) devices	5
4.1.1. Exchanging SRLG information between the PEs via the CE device	8
4.1.1.1. Operational Procedures	8
4.1.1.2. Error Handling Procedures	9
4.1.2. Using Path Affinity Set Extension	10
4.1.2.1. Operational Procedures	13
4.1.2.2. Error Handling Procedures	13
4.1.2.3. Distribution of the Path Affinity Set Information	14
5. Latency Signaling Extensions	15
5.1. RSVP-TE Extensions	16
5.2. Operational Procedures	16
5.3. Error Handling Procedures	16
6. Security Considerations	16
7. IANA Considerations	17
8. References	17
8.1. Normative References	17
8.2. Informative References	17
Authors' Addresses	19

1. Introduction

This document builds on the GMPLS overlay model [RFC4208] and defines extensions to the GMPLS User-Network Interface (UNI) to support route diversity within the core network for sets of LSPs initiated by edge nodes. In the following, the term customer edge (CE) device is used synonymously for the term edge node (EN) as in [RFC4208].

Moreover, the VPN terminology (CE and PE) [RFC4026] is used below when the core network is a VPN but is also applicable to UNI interfaces [RFC4208].

This document uses a VPN model that is based on the same premise as L1VPN framework [RFC4847] but may also be applied to other technologies. The extensions are applicable both to VPN and non VPN environments. These extensions move the UNI from basic connectivity to enhanced mode connectivity by including additional constraints while minimizing the exchange of CE to PE information. These extensions are applicable to the overlay extension service model.

The overlay model assumes a UNI interface between the edge nodes of the respective transport domains. Route diversity for LSPs from single homed CE and dual-home CEs is a common requirement in optical transport networks. This document describes two signaling variations that may be used for supporting LSP diversity within the overlay extension service model considering dual-homing. Dual-homing is typically used to avoid a single point of failure (UNI link, PE) or if two disjoint connections are forming a protection group in the CE device, e.g., 1+1 protection. While both methods are similar in that they utilize common mechanisms in the PE network to achieve diversity, they are distinguished according to whether the CE is permitted to retrieve provider SRLG diversity information for an LSP from a PE1 and pass it on to a PE2 (SRLG information is shared with the CE), or whether a new attribute is used that allows the PE2 that receives this attribute to derive the SRLG information for an LSP based on the attribute value. Figure 1 below is depicting the scenario.

The extended overlay service model can support other extensions for VPN signaling, for example, those related to latency. When requesting diverse LSPs, latency may also be an additional requirement.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

3. Contributors

The Authors would like to thank Eve Varma and Sergio Belotti for their review and contributions to this document.

4. LSP Diversity in the Overlay Extension Service Model

The L1VPN Framework [RFC4847] (Enhanced Mode) describes the overlay extension service model, which builds upon the UNI Overlay [RFC4208] serving as the interface between the CE edge node and the PE edge node. In this service model, a CE receives a list of CE-PE TE link addresses to which it can request a L1VPN connection (i.e., membership information) and may include additional information concerning these TE links. This document further builds on the overlay extension service model by adding shared constraint information for path diversity in the optical transport network.

While the L1VPN for optical transport is an example specific VPN technology the term VPN is used generically since the extensions can apply to GMPLS UNIs and VPNs for other technologies.

Two signaling variations are outlined here that may be used for supporting LSP diversity within the overlay extension service model considering dual-homing. While both methods utilize common mechanisms in the PE network to achieve diversity, they are distinguished according to whether the CE is permitted to retrieve provider SRLG diversity information for an LSP from a PE1 and pass it on to a PE2 (SRLG information is shared with the CE or whether a new attribute is used that allows the PE2 that receives this attribute to derive the SRLG information for an LSP based on this attribute value. The selection between these methods is governed by both PE-network specific policies and approaches taken (i.e., in terms of how the provider chooses to perform routing internal to their network).

The first method (see 3.1.1) assumes that provider Shared Resource Link Group (SRLG) Identifier information is both available and shareable (policy decision) with the CE. Since SRLG IDs can then be used (passed transparently between PEs via the dual-homed CE) as signaled information on a UNI message, a mechanism supporting LSP diversity for the overlay extension service model can be provided via straightforward signaling extensions.

The second method (see 3.1.2) assumes that provider SRLG IDs are either not available or not shareable (based on provider network operator policy) with the CE. For this case, a mechanism is provided where information signaled to the PE on UNI messages does not require shared knowledge of provider SRLG IDs to support LSP diversity for the overlay extension model.

While both methods could be implemented in the same PE network, it is likely that a GMPLS VPN CE network would use only one mechanism at a time.

4.1. LSP diversity for dual-homed customer edge (CE) devices

Single-homed CE devices are connected to a single PE device via a single UNI link (could be a bundle of parallel links which are typically using the same fiber cable). This single UNI link may constitute a single point of failure. Such a single point of failure can be avoided when the CE device is connected to two PE devices via two UNI interfaces as depicted for CE1 in Figure 1 below.

For the dual-homing case, it is possible to establish two connections from the source CE device to the same destination CE device where one connection is using one UNI link to, for example, PE1 and the other

connection is using the UNI link to PE2. In order to avoid single points of failure within the provider network, it is necessary to also ensure path (LSP) diversity within the provider network in order to achieve end-to-end diversity for the two LSPs between the two CE devices. This document describes how it is possible to enable such path diversity to be achieved within the provider network (which is subject to additional routing constraints). [RFC4202] defines SRLG information that can be used to allow GMPLS to provide path diversity in a GMPLS controlled transport network. As the two connections are entering the provider network at different PE devices, the PE device that receives the connection request for the second connection needs to be capable of determining the additional path computation constraints such that the path of the second LSP is disjoint with respect to the already established first connection entering the network at a different PE device. The methods described in this document allow a PE device to determine the SRLG information for a connection in the provider network that is entering the network on a different PE device.

PE SRLG information can be used directly by a CE if the CE understands the context, and the CE view is limited to its VPN context. In this case, there is a dependency on the provider information and there is a need to be able to query the SRLG in the provider network.

It may, on the other hand, be preferable to avoid this dependency and to decouple the SRLG identifier space used in the provider network from the SRLG space used in the client network. This is possible with both methods detailed below. Even for the method where provider SRLG information is passing through the CE device (note the CE device does not need to process and decode this information) the two SRLG identifier spaces can remain fully decoupled and the operator of the client network is free to assign SRLG identifiers from the client SRLG identifier space to the CE to CE connection that is passing through the provider network.

Referring to Figure 1, the UNI signaling mechanism must support at least one of the two mechanisms described in this document for CE dual homing to achieve LSP diversity in the provider network.

The described mechanisms can also be applied to a scenario where two CE devices are connected to two different PE devices. In this case, the additional information that is exchanged across the UNI interfaces also needs to be exchanged between the two CE devices in order to achieve the desired diversity in the provider network.

This information may be configured or exchanged by some automated

mechanism not described in this document.

In the dual-homing example, CE1 can locally correlate the LSP requests. For the slightly more complicated example involving CE2 and CE3, both requiring a path that shall be diverse to a connection initiated by the other CE device, CE2 and CE3 need to have a common view of the SRLG information to be signaled. In this document, we detail the required diversity information and the signaling of this diversity information; however, the means for distributing this information within the PE domain or the CE domain is out of scope.




Figure 1 Overlay Reference Diagram

In an overlay model, the information exchanged between the CE and the PE is kept to a minimum.

How diversity is achieved, in terms of configuration, distribution and usage in each part of the transport networks should be kept independent and separate from how diversity is signaled at the UNI between the two transport networks.

Signaling parameters discussed in this document are:

- o SRLG information (see [RFC4202])
- o Path Affinity Set

4.1.1.1. Exchanging SRLG information between the PEs via the CE device

SRLG information is defined in [RFC4202] and if the SRLG information of an LSP is known, it can be used to calculate a path for another LSP that is SRLG diverse with respect to an existing LSP. SRLG information is an unordered list of SRLGs. SRLG information is normally not shared between the transport network and the client network; i.e., not shared with the CEs of a VPN in the VPN context. However, this becomes more challenging when a CE is dual-homed. For example, CE1 in Figure 1 may have requested an LSP1 from CE1 to CE2 via PE1 and PE3. CE1 could subsequently request an LSP2 to CE2 via PE2 and PE3 with the requirement that it should be maximally SRLG disjoint with respect to LSP1. Since PE2 does not have any information about LSP1, PE2 would need to know the SRLG information associated with LSP1. If CE1 could request the SRLG information of LSP1 from PE1, it could then transparently pass this information to PE2 as part of the LSP2 setup request, and PE2 would now be capable of calculating a path for LSP2 that is SRLG disjoint with respect to LSP1.

The exchange of SRLG information is achieved on a per VPN LSP basis using the existing RSVP-TE signaling procedures. It can be exchanged in the PATH (exclusion information) or RESV message in the original request or it can be requested by the CE at any time the path is active.

It shall be noted that SRLG information is an unordered list of SRLG identifiers and the encoding of SRLG information for RSVP signaling is already defined in [SRLG_info]. Even if SRLG information is known for several LSPs it is not possible for the CEs to derive the provider network topology from this information.

4.1.1.1.1. Operational Procedures

Retrieving SRLG information from a PE for an existing LSP:

When a dual-homed CE device intends to establish an LSP to the same destination CE device via another PE node, it can request the SRLG information for an already established LSP by setting the SRLG information flag in the LSP attributes sub-object of the RSVP PATH message (IANA to assign the new SRLG flag). As long as the SRLG information flag is set in the PATH message, the PE node inserts the

SRLG sub-object as defined in [SRLG_info] into the RSVP RESV message that contains the current SRLG information for the LSP. If the provider network's policy has been configured so as not to share SRLG information with the client network, the SRLG sub-object is not inserted in the RESV message even if the SRLG information flag was set in the received PATH message. Note that the SRLG information is expected to be always up-to-date.

Establishment of a new LSP with SRLG diversity constraints:

When a dual-homed CE device sends an LSP setup requests to a PE device for a new LSP that is required to be SRLG diverse with respect to an existing LSP that is entering the network via another PE device, the CE device sets the SRLG diversity flag (note: IANA to assign the new SRLG diversity flag) in the LSP attributes sub-object of the PATH message that initiates the setup of this new LSP. When the PE device receives this request it calculates a path to the given destination and uses the received SRLG information as path computation constraints.

4.1.1.2. Error Handling Procedures

When the CE device receives a RSVP PATH message with the SRLG information flag set and if the provider's network policy does not permit sharing of SRLG information, the PE device shall notify the CE device by sending a RSVP PathErr with a Notify error code (error code to be defined) "Retrieval of SRLG information not permitted". As described above, the PE device must not include the SRLG sub-object with the SRLG information for the LSP in the RSVP RESV message.

If the PE device receives a RSVP PATH message for a new LSP with the SRLG diversity flag set and SRLG information in the SRLG sub-object, the PE device tries to calculate a route to the given destination that is SRLG diverse with respect to the provided SRLG information. If no route can be found, a RSVP PathErr message with an error code (error code to be defined) "No SRLG diverse route available toward destination".

If the PE device receives a RSVP PATH message for a new LSP with the SRLG diversity flag set and SRLG information in the SRLG sub-object and if the PE device does not support the SRLG sub-object, the PE device shall send a PathErr message to the CE device, indicating an "Unknown object class".

Further error handling cases will be added in the next revision of

this document.

4.1.2. Using Path Affinity Set Extension

The Path Affinity Set (PAS) is used to signal diversity in a pure CE context by abstracting SRLG information. There are two types of diversity information in the PAS. The first type of information is a single PAS identifier. The Second part is the optional PATH information, in the form of Source and Destination addresses of an exclude path or set of paths that MAY be specified. The motive behind the PAS information is to have as little exchange of diversity information as possible between the VPN CE and PE elements.

Rather than a detailed CE or PE SRLG list, the Path Affinity Set contains an abstract SRLG identifier that associates the given path as diverse. Logically the identifier is in a VPN context and therefore only unique with respect to a particular VPN.

How the CE determines the PAS identifier is a local matter for the CE administrator. A CE may signal the PAS identifier as a diversity object in the PATH message. This identifier is a suggested identifier and may be overridden by a PE under some conditions.

For example, a PAS identifier can be used with no prior exchange of PAS information between the CE and the PE. Upon reception of the PAS identifier information the PE can infer the CE's requirements. The actual PAS identifier used will be returned in the RESV message. Optionally an empty PAS identifier allows the PE to pick the PAS identifier.

Similar to the section 4.1.1 on SRLG information, a PE can return PAS identifier as the response to a Query allowing flexibility.

A PE interprets the specific PAS identifier, for example, "123" as meaning to exclude the PE SRLG information (or equivalent) that has been allocated by LSPs associated with this Path Affinity Set identifier "123", for any LSPs associated with the resources assigned to the VPN. For example, if a Path exists for the LSP with the identifier "123", the PE would use local knowledge of the PE SRLGs associated with the "123" LSPs and exclude those SRLGs in the path request. In other words, two LSPs that need to be diverse both signal "123" and the PEs interpret this as meaning not to use shared resources. Alternatively, a PE could use the PAS identifier to select from already established LSPs. Once the path is established it becomes the "123" identifier or optionally another PAS identifier for that VPN that replaces "123".

The optional PAS Source and Destination Address tuple represents one or more source addresses and destination addresses associated with the CE Path Affinity Set identifier. These associated address tuples represent paths that use resources that should be excluded for the establishment of the current LSP. The address tuple information gives both finer grain details on the path diversity request and serves as an alternative identifier in the case when the PAS identifier is not known by the PE. The address tuples used in signaling is within a CE context and its interpretation is local to a PE that receives a Path request from a CE. The PE can use the address information to relate to PE Addresses and PE SRLG information. When a PE satisfies a connection setup for a (SRLG) diverse signaled path, the PE may optionally record the PE SRLG information for that connection in terms of PE based parameters and associate that with the CE addresses in the Path message.

Specifically for L1VPNs, Port Information table (PIT) [RFC5251] can be leveraged to translate between CE based addresses and PE based addresses. The Path Affinity Set and associated PE addresses with PE SRLG information can be distributed via the IGP in the provider transport network (or by other means such as configuration); they can be utilized by other PEs when other CE Paths are setup that would require path/connection diversity. This information is distributed on a VPN basis and contains a PAS identifier, PE addresses and SRLG information.

If diversity is not signaled, the assumption is that no diversity is required and the Provider network is free to route the LSP to optimize traffic. No Path affinity set information needs to be recorded for these LSPs. If a diversity object is included in the connection request, the PE in the Provider Network should be able to look-up the existing Provider SRLG information from the provider network and choose an LSP that is maximally diverse from other LSPs.

The mechanisms to achieve this are outside the scope of this document.

A new VPN Diverse LSP LABEL object is specified:

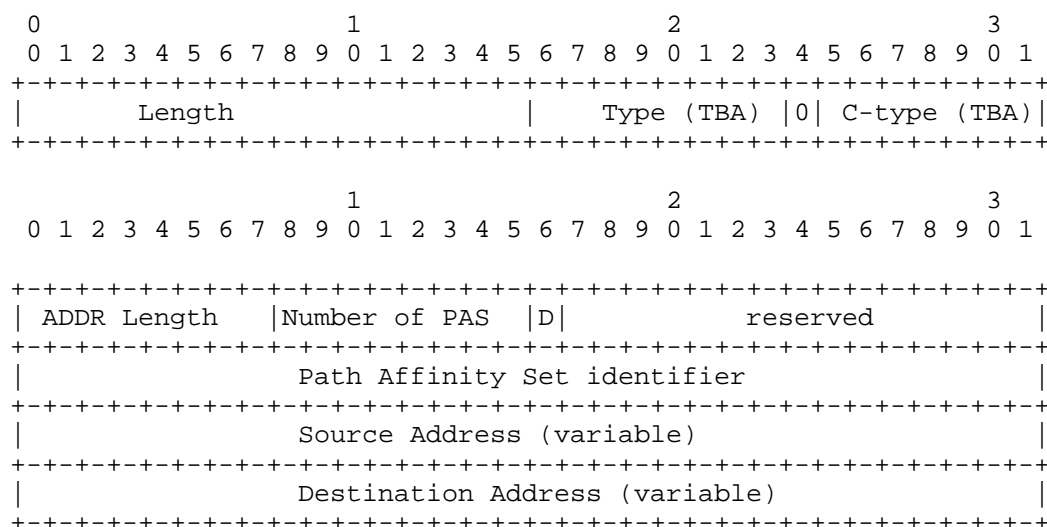


Figure 2 Diverse LSP information

1. The Address Length field (8 bits) is the number of bytes for both the source address and destination address. The address may be in any format from 1 to 32 bytes but the key point is the customers can maintain their existing addresses. A value of zero indicates there are no addresses included.
2. The Number of Path Affinity (8 bits) sets is included in the object. This is typically 1. Addition of other sets is for further study.
3. The Path affinity Set identifier (4 bytes) is a single number that represents a summarized SRLG for this path. Paths with that same Path Affinity set should be set up with diverse paths and associated with the path affinity set. A value of all zeros allows the PE to pick a PAS identifier to return. A PAS identifier of an established path may be different than the requested path identifier.
4. The diversity Bit (D) (one Bit) indicates if the diversity must be satisfied when set as a one. If a PE finds an established path with a Path Affinity set matching the signaled Path Affinity Set or the signaled Address tuple it should attempt find a diverse path.

5. The Diverse Path Source address/destination address tuple is that of an established LSP in the PE network that belongs to the same Path Affinity Set identifier. If the path for these addresses is not established or cannot be determined by the PE edge processing the PATH request then the path is established only with the Path Affinity identifier. If the path(s) for these address tuples are known by the PE the PE uses the SRLG information associated with these addresses. If in any case a diverse path cannot be setup then the Diverse bit controls whether a path is established anyway. The PE must use the PIT to translate CE Addresses into provider addresses when correlating with provider SRLG information. How SRLG information and network address tuples are distributed is for future study.

4.1.2.1. Operational Procedures

When a CE constructs a PATH message it may optionally specify and insert a Path Affinity Set in the PATH message. This Path Affinity Set may optionally include the address of an LSP that that could belong to the same Path Affinity Set. The Path Affinity Set identifier is a value (0 through $2^{32}-255$) that is independent of the mechanism the CE or the PE use for diversity. The Path Affinity Set is a single identifier that can be used to request diversity and associate diversity.

When processing a CE PATH message in a VPN Overlay, the PE first looks up the PE based addresses in the Provider Index Table (PIT). If the Path Affinity Set is included in the PATH message, the PE must look up the SRLG information (or equivalent) in the PE network that has been allocated by LSPs associated with a Path Affinity Set and exclude those resources from the path computation for this LSP if it is a new path. The PE may alternatively choose from an existing path with a disjoint set of resources. If a path that is disjoint cannot be found, the value of the PAS diversity bit determines whether a path should be setup anyway. If the PAS diversity bit is clear, one can still attempt to setup the LSP. A PE should still attempt to minimize shared resources but that is an implementation issue, and is outside the scope of this document.

Optionally the CE may use a value of all zeros in the PAS identifier allowing the PE to select an appropriate PAS identifier. Also the PE may to override the PAS identifier allowing the PE to re-assign the identifier if required. A CE should not assume that the PAS identifier used for setup is the actual PAS identifier.

4.1.2.2. Error Handling Procedures

The PAS object must be understood by the PE device. Otherwise, the CE should not use the PAS object. Path Message processing of the PAS object SHOULD follow CTYPE 0. An Error code of IANA (TBD) indicates that the PAS object is not understood.

When a PAS identifier is not recognized by a PE it must assume this LSP defines that PAS identifier however the PE may override PAS identifier under certain conditions.

If the identifier is recognized but the Source Address-Destination address pair(s) are not recognized, this LSP must be set up using the PAS identifier only.

If the identifier is recognized and the Source Address-Destination address pair(s) are also recognized, then the PE SHOULD use the PE SRLG information associated with the LSPs identified by the address pairs to select a disjoint path.

The Following are the additional error codes:

1. Route Blocked by Exclude Route Value IANA (TBA).

4.1.2.3. Distribution of the Path Affinity Set Information

Information about SRLG is already available in the IGP TE database. A PE network can be designed to have additional opaque records for Provider paths that distribute PE paths and SRLG on a VPN basis. When a PE path is setup, the following information allows a PE to lookup the PE diversity information:

- o L1 VPN Identifier 8 bytes
- o Path Affinity Set Identifier
- o Source PE Address
- o Destination PE Address
- o List of PE SRLG (variable)

The source PE address and destination PE address are the same addresses in the VPN PIT and correspond to the respective CE address identifiers.

Note that all of the information is local to the PE context and is not shared with the CE. The VPN Identifier is associated with a CE. The only value that is signaled from the CE is the Path Affinity Set and optionally the addresses of an existing LSP. The PE stores source and destination PE addresses of the LSP in their native format along with the SRLG information. This information is internal to the PE network and is always known.

PE paths may be setup on demand or they may be pre-established. When paths are pre-established, the Path Affinity Set is set to unassigned 0x0000 and is ignored. When a CE uses a pre-established path the PE may set the Path SRLG Path Affinity Set value if the CE signals one otherwise the Path Affinity Set remains unassigned 0x0000.

5. Latency Signaling Extensions

Some network applications are sensitive to latency (sometimes also called delay) while other applications are sensitive to latency variation (sometimes also called delay variation). Specifically, real time applications typically do have certain latency requirements. It shall be noted that latency variation is typically not an issue for TDM networks including the WDM layer. For these technologies the latency is constant and there is no latency variation added. Latency variation is typically caused in packet networks or when packet based services are encapsulated into a constant bit rate server layer signal, which requires buffering of the arriving packets that may arrive in bursts. An example is an Ethernet VLAN service that is mapped into a constant bit rate server layer such as an ODUk or ODUflex OTN signal.

The GMPLS UNI as defined in [RFC4208] does not support latency as a signaling parameter that would allow a CE device to signal to the PE device that latency and/or latency variation constraints need to be met when a path is calculated for the requested LSP. The path computation function does typically calculate a route to the given destination that has the least TE metric (least cost routing). However, if a CE device requests an LSP via the UNI interface for an application that is sensitive to latency/latency variation, it should be possible to signal to the PE device that the objective function should rather take latency into account instead of the TE metric.

In order to support latency/latency variation as path computation constraint, the network has to support latency/latency variation as TE metric extension as defined in [DRAFT OSPF TE METRIC EXT] - note that [DRAFT OSPF TE METRIC EXT] is using the terms delay/delay variation instead of latency/latency variation.

A latency requirement can be added to signaling in the form of a

constraint [DRAFT OBJECTIVE FUNCTION]. The constraint can take the form of:

- o Minimal latency
- o Maximum acceptable latency (upper bound)
- o Minimal latency variation
- o Maximum acceptable latency variation (upper bound), if applicable

While some systems may be able to compute routes based on delay metrics it is usual that minimizing the accumulated TE link metric (link cost) or the number of hops subject to bandwidth reservation are satisfied as the object function and delay is not considered. When considering diversity latency falls after diversity constraints have been satisfied.

Recording the latency of existing paths [DRAFT_TE_METRIC RECORD] to ensure they meet a maximum acceptable latency can be utilized to ensure latency constraint is met.

When a low latency path is required, the minimize latency subject to other constraints criteria should be signaled. A CE device can use the recorded latency to ensure that the maximum acceptable latency has been met.

5.1. RSVP-TE Extensions

At the UNI, the RSVP-TE extensions as defined in [DRAFT OBJECTIVE FUNCTION] SHALL be used for signaling the PE device whether a path with minimal latency is requested or whether certain latency/latency variation upper bound constraints shall be met for the end-to-end connection, i.e., from the source CE device to the destination CE device. The following objective function (OF) code point SHALL be used in the OF sub-object of the ERO to indicate that latency/latency variation constraints SHALL be taken into account when the path computation function that is invoked by the PE node that expands the route from the PE device to the destination CE device:

- o OF code value 8 (to be assigned by IANA) is for the Minimum Latency Path (MLP) OF
- o OF code value 9 (to be assigned by IANA) is for Minimum Latency Variation Path (MLVP) OF

Additionally, an optional OF metric-bound sub-object MAY be carried within an ERO object of the RSVP-TE Path message. The two metric-

bound sub-objects defined in [DRAFT OBJECTIVE FUNCTION] that are corresponding to the two OFs above are:

- o metric bound sub-object of Type T=4: Cumulative Latency
- o metric bound sub-object of Type T=5: Cumulative Latency Variation

The metric-bound indicates an upper bound for the path metric that MUST NOT be exceeded for the ERO expending node to consider the computed path as acceptable. It shall be noted that the metric bound included in the RSVP-TE Path message at the UNI has end-to-end significance, which means that the upper bound metric constraint MUST be met for the path from the source CE device to the destination CE device.

5.2. Operational Procedures

The processing rules as defined in [DRAFT OBJECTIVE FUNCTION] for the OF sub-object and the optional OF metric-bound sub-object SHALL be applied at the ingress PE device when the source CE device requests an LSP (It shall be noted that [DRAFT OBJECTIVE FUNCTION] has a wider scope and may also apply to inter-domain interfaces, i.e., when the provider network is composed of multiple separate domains.).

5.3. Error Handling Procedures

The error handling rules as defined in [DRAFT OBJECTIVE FUNCTION] for the OF sub-object and the optional OF metric-bound sub-object SHALL be applied.

6. Security Considerations

Security for L1VPNs is covered in [RFC4847], [RFC5251] and [RFC5253]. In this document, the model follows a generic GMPLS VPN based on the L1VPN control plane model where CE addresses are completely distinct from the PE addresses.

The use of a private network assumes that entities outside the network cannot spoof or modify control plane communications between CE and PE. Furthermore, all entities in the private network are assumed to be trusted. Thus, no security mechanisms are required by the protocol exchanges described in this document.

However, an operator that is concerned about the security of their private control plane network may use the authentication and

integrity functions available in RSVP-TE [RFC3473] or utilize IPsec ([RFC4301], [RFC4302], [RFC4835], [RFC5996], and [RFC6071]) for the point-to-point signaling between PE and CE. See [RFC5920] for a full discussion of the security options available for the GMPLS control plane.

7. IANA Considerations

TBD

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4202] Kompella, K., Rekhter, Y., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC5251] Fedyk, D., Rekhter, Y., Editors "Layer 1 VPN Basic Mode", RFC 5251, July 2008.
- [SRLG_info] Zhang, F., Li, D., Gonzalez de Dios, O., Margaria, C., Hartley, M., "RSVP-TE Extensions for Collecting SRLG Information", draft-ietf-ccamp-rsvp-te-srlg-collect-02.txt, February 2013.
- [DRAFT OBJECTIVE FUNCTION] Ali, Z., Swallow, G., Filsfils, C., Fang, L., Kumaki, K., Kunze, R., "Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) extension for signaling Objective Function and Metric Bound", draft-ali-ccamp-rc-

objective-function-metric-bound-02.txt, July 2012.

8.2. Informative References

- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
- [RFC6071] Frankel, S. and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", RFC 6071, February 2011.
- [RFC3473] Berger, L. (editor), "Generalized MPLS Signaling - RSVP-TE Extensions", RFC 3473, January 2003.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC4835] Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4835, April 2007.
- [RFC4847] Takeda, T., Editor "Framework and Requirements for Layer Virtual Private Networks", RFC 4847, April 2007.
- [RFC5253] Takeda, T., Ed., "Applicability Statement for Layer 1 Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July 2008.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[DRAFT_TE_METRIC RECORD] Ali, Z., Swallow, G., Filsfils, C., Hartley, M., Kumaki, K., Kunze, R., "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) extension for recording TE Metric of a Label Switched Path", draft-ietf-ccamp-te-metric-recording-02.txt, July 2013.

[DRAFT_OSPF TE METRIC EXT] Giacalone, S., Ward, D., Drake, J., Atlas, A., Previdi, S., "OSPF Traffic Engineering (TE) Metric Extensions", draft-ietf-ospf-te-metric-extensions-04.txt, June 2013.

Copyright (c) 2013 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in Section 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

Authors' Addresses

Don Fedyk
Alcatel-Lucent
Groton, MA, 01450
Email: donald.fedyk@alcatel-lucent.com

Dieter Beller
Alcatel-Lucent
Email: Dieter.Beller@alcatel-lucent.com

Lieven Levrau
Alcatel-Lucent
Email: Lieven.Levrau@alcatel-lucent.com

Daniele Ceccarelli
Ericsson
Email: Daniele.Ceccarelli@ericsson.com

Fatai Zhang
Huawei Technologies
Email: zhangfatai@huawei.com

Yuji Tochio
Fujitsu
Email: tochio@jp.fujitsu.com

Xihua Fu
ZTE
Email: fu.xihua@zte.com.cn

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

. G.Galimberti, Ed.
Cisco
. R.Kunze, Ed.
Deutsche Telekom
. Kam Lam, Ed.
Alcatel-Lucent
. D. Hiremagalur, Ed.
Juniper
July 15, 2013

An SNMP MIB extension to RFC3591 to manage optical interface parameters
of DWDM applications
draft-galikunze-ccamp-g-698-2-snmp-mib-03

Abstract

This memo defines a module of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) in TCP/IP- based internet. In particular, it defines objects for managing Optical parameters associated with Dense Wavelength Division Multiplexing (DWDM) interfaces. This is an extension of the RFC3591 to support the optical parameters described in ITU-T G.698.2. [ITU.G698.2]

The MIB module defined in this memo can be used for Optical Parameters monitoring and/or configuration of the endpoints of Black Links.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 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	2
2. The Internet-Standard Management Framework	4
3. Conventions	4
4. Overview	4
4.1. Optical Parameters Description	5
4.1.1. Rs-Ss Configuration	6
4.1.2. Table of Application Codes	7
4.2. Use of ifTable	8
4.2.1. Use of ifTable for OPS Layer	10
4.2.2. Use of ifTable for OCh Layer	11
4.2.3. Use of ifStackTable	11
5. Structure of the MIB Module	12
6. Object Definitions	12
7. Relationship to Other MIB Modules	19
7.1. Relationship to the [TEMPLATE TODO] MIB	19
7.2. MIB modules required for IMPORTS	19
8. Definitions	19
9. Security Considerations	19
10. IANA Considerations	20
11. Contributors	21
12. References	21
12.1. Normative References	22
12.2. Informative References	24
Appendix A. Change Log	24
Appendix B. Open Issues	24
Authors' Addresses	24

1. Introduction

This memo defines a portion of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) in TCP/IP- based internets. In particular, it defines objects for managing Optical parameters associated with Wavelength Division Multiplexing (WDM) systems in accordance with the optical interface defined in G.698.2 [ITU.G698.2]

Black Link approach allows supporting an optical transmitter/receiver pair of one vendor to inject a DWDM channel and run it over an optical network composed of amplifiers, filters, add-drop multiplexers from a different vendor. From architectural point of view, the "Black Link" is a set of pre-configured/qualified network connections between the G.698.2 reference points S and R. The black links will be managed at the edges (i.e. the transmitters and receivers attached to the S and R reference points respectively) for the relevant parameters specified in G.698.2 [ITU.G698.2], G.798 [ITU.G798], G.874 [ITU.G874], and the performance parameters specified G.7710/Y.1701 [ITU-T G.7710] and G.874.1 [ITU.G874.1].

The G.698.2 [ITU.G698.2] provides optical parameter values for physical layer interfaces of Dense Wavelength Division Multiplexing (DWDM) systems primarily intended for metro applications which include optical amplifiers. Applications are defined in G.698.2 [ITU.G698.2] using optical interface parameters at the single-channel connection points between optical transmitters and the optical multiplexer, as well as between optical receivers and the optical demultiplexer in the DWDM system. This Recommendation uses a methodology which does not specify the details of the optical link, e.g. the maximum fibre length, explicitly. The Recommendation currently includes unidirectional DWDM applications at 2.5 and 10 Gbit/s (with 100 GHz and 50 GHz channel frequency spacing). Work is still underway for 40 and 100 Gbit/s interfaces. There is possibility for extensions to a lower channel frequency spacing. This document specifically refers to the "application code" defined in the G.698.2 [ITU.G698.2] plus few optical parameter not included in the application code definition.

This draft refers and supports also the draft-kunze-g-698-2-management-control-framework

The building of an SNMP MIB describing the optical parameters defined in G.698.2 [ITU.G698.2] G.798 [ITU.G798], G.874 [ITU.G874], parameters specified G.7710/Y.1701 [ITU-T G.7710] allows the different vendors and operator to retrieve, provision and exchange information related to Optical black links in a standardized way. This facilitates interworking in case of using optical interfaces from different vendors at the end of the link.

The MIB, reporting the Optical parameters and their values, characterizes the features and the performances of the optical components and allow a reliable black link design in case of multivendor optical networks.

Although RFC 3591 [RFC3591] describes and defines the SNMP MIB of a number of key optical parameters, alarms and Performance Monitoring, a more complete description of optical parameters and processes can be found in the ITU-T Recommendations. Appendix A of this document provides an overview about the extensive ITU-T documentation in this area. The same considerations can be applied to the RFC 4054 [RFC4054]

2. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].


Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies a MIB module that is compliant to the SMIV2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

3. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119]. In the description of OIDs the convention: Set (S) Get (G) and Trap (T) conventions will describe the action allowed by the parameter.

4. Overview

Figure 1 shows a set of reference points, for the linear "black link" approach, for single-channel connection (Ss and Rs) between transmitters (Tx) and receivers (Rx). Here the DWDM network elements include an OM and an OD (which are used as a pair with the opposing element), one or more optical amplifiers and may also include one or more OADMs.



Ss = reference point at the DWDM network element tributary output

Rs = reference point at the DWDM network element tributary input

Lx = Lambda x

OM = Optical Mux

OD = Optical Demux

OADM = Optical Add Drop Mux

from Fig. 5.1/G.698.2

Figure 1: Linear Black Link

G.698.2 [ITU.G698.2] defines also Ring Black Link configurations [Fig. 5.2/G.698.2] and Bidirectional Black Link configurations [Fig. 5.3/G.698.2]

4.1. Optical Parameters Description

The black links are managed at the edges, i.e. at the transmitters (Tx) and receivers (Rx) attached to the S and R reference points respectively. The parameters that could be managed at the black link edges are specified in G.698.2 [ITU.G698.2] section 5.3 referring the "application code" notation

The definitions of the optical parameters are provided below to increase the readability of the document, where the definition is ended by (G) the parameter can be retrieve with a GET, when (S) it can be provisioned by a SET, (G,S) can be either GET and SET.

To support the management of these parameters, the SNMP MIB in RFC 3591 [RFC3591] is extended with a new MIB module defined in section 6 of this document. This new MIB module includes the definition of new configuration table of the OCh Layer for the parameters at Tx (S) and Rx (R).

4.1.1. Rs-Ss Configuration

The Rs-Ss configuration table allows configuration of Wavelength, Power and Application codes as described in [ITU.G698.2] and G.694.1 [ITU.G694.1]

This parameter report the current Transceiver Output power, it can be either a setting and measured value (G, S) NEED TO DISCUSS ON THIS.

Wavelength Value (see G.694.1 Table 1):

This parameter indicates the wavelength value that Ss and Rs will be set to work (in THz) se in particular Section 6/G.694.1 (G, S).

Vendor Transceiver Class:

Other than specifying all the Transceiver parameter, it might be convenient for the vendors to summarize a set of parameters in a single proprietary parameter: the Class of transceiver. The Transceiver classification will be based on the Vendor Name and the main TX and RX parameters (i.e. Trunk Mode, Framing, Bit rate, Trunk Type, Channel Band, Channel Grid, Modulation Format, Channel Modulation Format, FEC Coding, Electrical Signal Framing at Tx, Minimum maximum Chromatic Disperion (CD) at Rx, Maximum Polarization Mode Dispersion (PMD) at Rx, Maximum differential group delay at Rx, Loopbacks, TDC, Pre-FEC BER, Q-factor, Q-margin,etc.). If this parameter is used, the MIB parameters specifying the Transceiver characteristics may not be significant and the vendor will be responsible to specify the Class contents and values. The Vendor can publish the parameters of its Classes or declare to be compatible with published Classes.(G) Optional for compliance. (not mentioned in G.698)

Number of Vendor Transceiver Class Supported

This parameter indicates the number of Vendor Transceiver codes supported by this interface (G).

Single-channel application codes (see G.698.2):

This parameter indicates the transceiver application code at Ss and Rs as defined in [ITU.G698.2] Chapter 5.4 - this parameter can be called Optical Interface Identifier OII as per [draft-martinelli-wson-interface-class] (G).

Number of Single-channel application codes Supported

This parameter indicates the number of Single-channel application codes supported by this interface (G).

Current Laser Output power:

This parameter report the current Transceiver Output power, it can be either a setting and measured value (G, S).

Minimum Laser Output power:

This parameter report the minimum Transceiver Output power supported by this interface (G).

Maximum Laser Output power:

This parameter report the maximum Transceiver Output power supported by this interface (G).

Current Laser Input power:

This parameter report the current Transceiver Input power (G).

Minimum Laser Input power:

This parameter report the minimum Transceiver Input power supported by this interface (G).

Maximum Laser Input power:

This parameter report the maximum Transceiver Input power supported by this interface (G).

PARAMETERS	Get/Set	Reference
Wavelength Value	G,S	G.694.1 S.6
Vendor Transceiver Class	G	N.A.
Number of Vendor Transceiver Class Supported	G	N.A.
Single-channel application codes	G	G.698.2 S.5.3
Number of Single-channel application codes Supported	G	N.A.
Current Output Power	G,S	N.A.
Minimum Output Power	G	N.A.
Maximum Output Power	G	N.A.
Current Input Power	G	N.A.
Minimum Input Power	G	N.A.
Maximum Input Power	G	N.A.

Table 1: Rs-Ss Configuration

4.1.2. Table of Application Codes

This table has a list of Application codes supported by this interface at point R are defined in G.698.2.

4.1.2.1.

Maximum and minimum mean input power:

The maximum and minimum values of the average received power (in dBm) at point Rs. (G)

PARAMETERS	Get/Set	Reference
MAX and min mean input power	G	G.698.2 S.7.4.1

Table 2: mandatory parameters

4.2. Use of ifTable

This section specifies how the MIB II interfaces group, as defined in RFC 2863 [RFC2863], is used for the link ends of a black link. Only the ifGeneralInformationGroup will be supported for the ifTable and the ifStackTable to maintain the relationship between the OCh and OPS layers. The OCh and OPS layers are managed in the ifTable using IfEntries that correlate to the layers depicted in Figure 1.

For example, a device with TX and/or RX will have an Optical Physical Section (OPS) layer, and an Optical Channel (OCh) layer. There is a one to n relationship between the OPS and OCh layers.

EDITOR NOTE: Reason for changing from OChr to OCh: Work on revised G.872 in the SG15 December 2011 meeting agreed to remove OChr from the architecture and to update G.709 to account for this architectural change. The meeting also agreed to consent the revised text of G.872 and G.709 at the September 2012 SG15 meeting.

Figure 2 In the following figures, opticalChannel and opticalPhysicalSection are abbreviated as och and ops respectively.




Figure 2: OTN Layers for OPS and OCh

Each opticalChannel IfEntry is mapped to one of the m opticalPhysicalSection IfEntries, where m is greater than or equal to 1. Conversely, each opticalTransPhysicalSection port entry is mapped to one of the n opticalChannel IfEntries, where n is greater than or equal to 1.

The design of the Optical Interface MIB provides the option to model an interface either as a single bidirectional object containing both sink and source functions or as a pair of unidirectional objects, one containing sink functions and the other containing source functions.

If the sink and source for a given protocol layer are to be modelled as separate objects, then there need to be two ifTable entries, one that corresponds to the sink and one that corresponds to the source, where the directionality information is provided in the configuration tables for that layer via the associated Directionality objects. The agent is expected to maintain consistent directionality values between ifStackTable layers (e.g., a sink must not be stacked in a 1:1 manner on top of a source, or vice-versa), and all protocol layers that are represented by a given ifTable entry are expected to have the same directionality.

When separate ifTable entries are used for the source and sink functions of a given physical interface, association between the two uni-directional ifTable entries (one for the source function and the other for the sink functions) should be provided. It is recommended

that identical ifName values are used for the two ifTable entries to indicate such association. An implementation shall explicitly state what mechanism is used to indicate the association, if ifName is not used.

4.2.1. Use of ifTable for OPS Layer

Only the ifGeneralInformationGroup needs to be supported.

ifTable Object	Use for OTN OPS Layer
=====	
ifIndex	The interface index.
ifDescr	Optical Transport Network (OTN) Optical Physical Section (OPS)
ifType	opticalPhysicalSection (xxx)
<<<Editor Note: Need new IANA registration value for xxx. >>>	
ifSpeed	Actual bandwidth of the interface in bits per second. If the bandwidth of the interface is greater than the maximum value of 4,294,967,295, then the maximum value is reported and ifHighSpeed must be used to report the interface's speed.
ifPhysAddress	An octet string with zero length. (There is no specific address associated with the interface.)
ifAdminStatus	The desired administrative state of the interface. Supports read-only access.
ifOperStatus	The operational state of the interface. The value lowerLayerDown(7) is not used, since there is no lower layer interface. This object is set to notPresent(6) if a component is missing, otherwise it is set to down(2) if either of the objects optIfOPSnCurrentStatus indicates that any defect is present.
ifLastChange	The value of sysUpTime at the last change in ifOperStatus.
ifName	Enterprise-specific convention (e.g., TL-1 AID)

to identify the physical or data entity associated with this interface or an OCTET STRING of zero length. The enterprise-specific convention is intended to provide the means to reference one or more enterprise-specific tables.

ifLinkUpDownTrapEnable Default value is enabled(1). Supports read-only access.

ifHighSpeed Actual bandwidth of the interface in Mega-bits per second. A value of n represents a range of 'n-0.5' to 'n+0.499999'.

ifConnectorPresent Set to true(1).

ifAlias The (non-volatile) alias name for this interface as assigned by the network manager.

4.2.2. Use of ifTable for OCh Layer

Use of ifTable for OCh Layer See RFC 3591 [RFC3591] section 2.4

4.2.3. Use of ifStackTable

Use of the ifStackTable and ifInvStackTable to associate the opticalPhysicalSection and opticalChannel interface entries is best illustrated by the example shown in Figure 3. The example assumes an ops interface with ifIndex i that carries two multiplexed OCh interfaces with ifIndex values of j and k, respectively. The example shows that j and k are stacked above (i.e., multiplexed into) i. Furthermore, it shows that there is no layer lower than i and no layer higher than j and/or k.

Figure 3

HigherLayer	LowerLayer
0	j
0	k
j	i
k	i
i	0

Figure 3: Use of ifStackTable for an OTN port

For the inverse stack table, it provides the same information as the interface stack table, with the order of the Higher and Lower layer interfaces reversed.

5. Structure of the MIB Module

EDITOR NOTE: text will be provided based on the MIB module in Section 6

6. Object Definitions

EDITOR NOTE: Once the scope in Section 1 and the parameters in Section 4 are finalized, a MIB module will be defined. It could be an extension to the OPT-IF-MIB module of RFC 3591. >>>

```
OPT-IF-698-MIB DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    MODULE-IDENTITY,
    OBJECT-TYPE,
    Gauge32,
    Integer32,
    Unsigned32,
    Counter64,
    transmission,
    NOTIFICATION-TYPE
    FROM SNMPv2-SMI
    TEXTUAL-CONVENTION,
    RowPointer,
    RowStatus,
    TruthValue,
    DisplayString,
    DateAndTime
```

```
FROM SNMPv2-TC
SnmpAdminString
FROM SNMP-FRAMEWORK-MIB
MODULE-COMPLIANCE, OBJECT-GROUP
FROM SNMPv2-CONF
ifIndex
FROM IF-MIB
optIfMibModule
FROM OPT-IF-MIB;
```

```
-- This is the MIB module for the optical parameters -
-- Application codes associated with the black link end points.
```

```
optIfXcvrMibModule MODULE-IDENTITY
    LAST-UPDATED "201204250000Z"
    ORGANIZATION "IETF Ops/Camp MIB Working Group"
    CONTACT-INFO
        "WG charter:
         http://www.ietf.org/html.charters/

        Mailing Lists:
        Editor: Gabriele Galimberti
        Email: ggalimbe@cisco.com"
    DESCRIPTION
        "The MIB module to describe Black Link transceiver
        characteristics to rfc3591.
        Copyright (C) The Internet Society (2012). This version
        of this MIB module is part of ; see the RFC
        itself for full legal notices."
    REVISION "201305050000Z"
    DESCRIPTION
        "Draft version 1.0"
    REVISION "201305050000Z"
    DESCRIPTION
        "Draft version 2.0"
    REVISION "201302270000Z"
    DESCRIPTION
        "Draft version 3.0"
    REVISION "201307020000Z"
    DESCRIPTION
        "Mib has in application code/vendor transceivercode G.698."
    ::= { optIfMibModule 4 }
```

```
-- Addition to the RFC 3591 objects
optIfOchSsRsGroup OBJECT IDENTIFIER ::= { optIfXcvrMibModule 1 }
```

```
-- OCh Ss/Rs config table
-- The application code/vendor transceiver class for the Black Link
-- Ss-Rs will be added to the OchConfigTable
```

```
optIfOchSsRsConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOchSsRsConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A table of Och General config extension parameters"
    ::= { optIfOchSsRsGroup 1 }
```

```
optIfOchSsRsConfigEntry OBJECT-TYPE
    SYNTAX OptIfOchSsRsConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual row that contains G.698 parameters for an
        interface."
    INDEX { ifIndex }
    ::= { optIfOchSsRsConfigTable 1 }
```

```
OptIfOchSsRsConfigEntry ::=
    SEQUENCE {
        optIfOchWavelengthn                Unsigned32,
        optIfOchInterfaceVendorTransceiverClass DisplayString,
        optIfOchNumberVendorClassesSupported Unsigned32,
        optIfOchInterfaceApplicationCode    DisplayString,
        optIfOchNumberApplicationCodesSupported Unsigned32,
        optIfOchOutputPower                  Integer32,
        optIfOchMinOutputPower                Integer32,
        optIfOchMaxOutputPower                Integer32,
        optIfOchInputPower                    Integer32,
        optIfOchMinInputPower                  Integer32,
        optIfOchMaxInputPower                  Integer32
```



```

    }

    optIfOchWavelengthn OBJECT-TYPE
        SYNTAX Unsigned32
        MAX-ACCESS read-write
        STATUS current
        DESCRIPTION
            " This parameter indicate minimum wavelength spectrum - n, in
              a definite wavelength Band (L, C and S) as represented in
              [RFC6205] by the formula -
              Wavelength (nm ) = 1471nm + n* optIfOchMimumumChannelSpacing
                                   (converted to nm)
              Eg - optIfOchMimumumChannelSpacing in nm
              'Wavelength (nm ) = 1471nm + n* 20nm (20nm is the spacing
                for CWDM)'
            "
        ::= { optIfOchSsRsConfigEntry 1 }

    optIfOchInterfaceVendorTransceiverClass OBJECT-TYPE
        SYNTAX DisplayString
        MAX-ACCESS read-write
        STATUS current
        DESCRIPTION
            "As defined in G.698
              Vendors can summarize a set of parameters in a
              single proprietary parameter: the Class of transceiver. The
              Transceiver classification will be based on the Vendor Name
              and the main TX and RX parameters (i.e. Trunk Mode, Framing,
              Bit rate, Trunk Type etc).
              This defines the tranceiver class that is/should be used by
              this interface. The optIfOchSrcVendorTranscieverClassTable
              has all the vendor classes supported by this interface."
        ::= { optIfOchSsRsConfigEntry 2 }

    optIfOchNumberVendorClassesSupported OBJECT-TYPE
        SYNTAX Unsigned32
        MAX-ACCESS read-only
        STATUS current
        DESCRIPTION
            " Number of Vedor classes supported by this interface."
        ::= { optIfOchSsRsConfigEntry 3 }

    optIfOchInterfaceApplicationCode OBJECT-TYPE
        SYNTAX DisplayString
        MAX-ACCESS read-write
        STATUS current
        DESCRIPTION

```

```

        "This parameter indicates the transceiver application code at
        Ss and Rs as defined in [ITU.G698.2] Chapter 5.3, that
        is/should be used by this interface. The
        optIfOChSrcApplicationCodeTable has all the application
        codes supported by this interface. "
 ::= { optIfOChSsRsConfigEntry 4 }

optIfOChNumberApplicationCodesSupported OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " Number of Application codes supported by this interface."
 ::= { optIfOChSsRsConfigEntry 5 }

optIfOChOutputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.01dbm"
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        " The output power for this interface in .01 dbm "
 ::= { optIfOChSsRsConfigEntry 6 }

optIfOChMinOutputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.01dbm"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " The minimum output power for this interface in .01 dbm "
 ::= { optIfOChSsRsConfigEntry 7 }

optIfOChInputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.01dbm"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " The input power for this interface in .01 dbm "
 ::= { optIfOChSsRsConfigEntry 8 }

optIfOChMinInputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.01dbm"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
```

```

        " The minimum input power for this interface in .01 dbm "
        ::= { optIfOChSsRsConfigEntry 9 }

optIfOChMaxInputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.01dbm"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " The maximum input power for this interface in .01 dbm "
        ::= { optIfOChSsRsConfigEntry 10 }

-- Table of Application codes supported by the interface
-- OptIfOChSrcApplicationCodeEntry

optIfOChSrcApplicationCodeTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOChSrcApplicationCodeEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Table of Application codes supported by this interface."
        ::= { optIfOChSsRsGroup 2 }

optIfOChSrcApplicationCodeEntry OBJECT-TYPE
    SYNTAX OptIfOChSrcApplicationCodeEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual row that contains the Application code for this
        interface."
    INDEX { ifIndex, optIfOChApplicationCodeNumber }
    ::= { optIfOChSrcApplicationCodeTable 1 }

OptIfOChSrcApplicationCodeEntry ::=
    SEQUENCE {
        optIfOChApplicationCodeNumber Integer32,
        optIfOChApplicationCode DisplayString
    }

optIfOChApplicationCodeNumber OBJECT-TYPE
    SYNTAX Integer32 (1..255)
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        " The number of the application code supported at this
        interface. The interface can support more than one
        application codes.

```

```

    "
    ::= { optIfOchSrcApplicationCodeEntry 1}

optIfOchApplicationCode OBJECT-TYPE
    SYNTAX DisplayString
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " The application code supported by this interface DWDM
        link."
    ::= { optIfOchSrcApplicationCodeEntry 2}

-- Table of Vendor Transceiver class supported by the interface
-- OptIfOchSrcVendorTransceiverClassEntry

optIfOchSrcVendorTransceiverClassTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOchSrcVendorTransceiverClassEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A table of OCh Src (Ss) transceiver classes supported by
        this interface."
    ::= { optIfOchSsRsGroup 3 }

optIfOchSrcVendorTransceiverClassEntry OBJECT-TYPE
    SYNTAX OptIfOchSrcVendorTransceiverClassEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual row that contains the transceiver classes
        supported by this interface."
    INDEX { ifIndex, optIfOchTransceiverClassNumber }
    ::= { optIfOchSrcVendorTransceiverClassTable 1 }

OptIfOchSrcVendorTransceiverClassEntry ::=
    SEQUENCE {
        optIfOchTransceiverClassNumber      Integer32,
        optIfOchTransceiverClass            DisplayString
    }

optIfOchTransceiverClassNumber OBJECT-TYPE
    SYNTAX Integer32 (1..255)
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        " The number of the application code supported at this
        interface. The interface can support more than one
        application codes.

```

```
"
 ::= { optIfOChSrcVendorTransceiverClassEntry 1}

optIfOChTransceiverClass OBJECT-TYPE
    SYNTAX DisplayString
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        " Vendor transceiver class supported by this interface."
    ::= { optIfOChSrcVendorTransceiverClassEntry 2}

END
```

7. Relationship to Other MIB Modules

7.1. Relationship to the [TEMPLATE TODO] MIB

7.2. MIB modules required for IMPORTS

8. Definitions

[TEMPLATE TODO]: put your valid MIB module here.
A list of tools that can help automate the process of checking MIB definitions can be found at <http://www.ops.ietf.org/mib-review-tools.html>

9. Security Considerations

There are a number of management objects defined in this MIB module with a MAX-ACCESS clause of read-write and/or read-create. Such objects may be considered sensitive or vulnerable in some network environments. The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations. These are the tables and objects and their sensitivity/vulnerability:

Some of the readable objects in this MIB module (i.e., objects with a MAX-ACCESS other than not-accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to control even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP.

SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example by using IPsec),

even then, there is no control as to who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in this MIB module.

It is RECOMMENDED that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Further, deployment of SNMP versions prior to SNMPv3 is NOT RECOMMENDED. Instead, it is RECOMMENDED to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator responsibility to ensure that the SNMP entity giving access to an instance of this MIB module is properly configured to give access to the objects only to those principals (users) that have legitimate rights to indeed GET or SET (change/create/delete) them.

10. IANA Considerations

Option #1:

The MIB module in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
-----	-----
sampleMIB	{ mib-2 XXX }

Option #2:

Editor's Note (to be removed prior to publication): the IANA is requested to assign a value for "XXX" under the 'mib-2' subtree and to record the assignment in the SMI Numbers registry. When the assignment has been made, the RFC Editor is asked to replace "XXX" (here and in the MIB module) with the assigned value and to remove this note.

Note well: prior to official assignment by the IANA, an internet draft MUST use placeholders (such as "XXX" above) rather than actual numbers. See RFC4181 Section 4.5 for an example of how this is done in an internet draft MIB module.

Option #3:

This memo includes no request to IANA.

11. Contributors

Arnold Mattheus
Deutsche Telekom
Darmstadt
Germany
email a.mattheus@telekom.de

Manuel Paul
Deutsche Telekom
Berlin
Germany
email Manuel.Paul@telekom.de

Frank Luennemann
Deutsche Telekom
Munster
Germany
email Frank.Luennemann@telekom.de

Scott Mansfield
Ericsson Inc.
email scott.mansfield@ericsson.com

Najam Saquib
Cisco
Ludwig-Erhard-Strasse 3
ESCHBORN, HESSEN 65760
GERMANY
email nasaquib@cisco.com

Walid Wakim
Cisco
9501 Technology Blvd
ROSEMONT, ILLINOIS 60018
UNITED STATES
email wwakim@cisco.com

Ori Gerstel
Cisco
32 HaMelacha St., (HaSharon Bldg)
SOUTH NETANYA, HAMERKAZ 42504
ISRAEL
email ogerstel@cisco.com

12. References

12.1. Normative References

- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.
- [RFC3591] Lam, H-K., Stewart, M., and A. Huynh, "Definitions of Managed Objects for the Optical Interface Type", RFC 3591, September 2003.
- [RFC6205] Otani, T. and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.
- [ITU.G698.2] International Telecommunications Union, "Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces ", ITU-T Recommendation G.698.2, November 2009.
- [ITU.G709] International Telecommunications Union, "Interface for the Optical Transport Network (OTN) ", ITU-T Recommendation G.709, March 2003.
- [ITU.G872] International Telecommunications Union, "Architecture of optical transport networks ", ITU-T Recommendation G.872, November 2001.
- [ITU.G798]

International Telecommunications Union, "Characteristics of optical transport network hierarchy equipment functional blocks ", ITU-T Recommendation G.798, October 2010.

[ITU.G874]

International Telecommunications Union, "Management aspects of optical transport network elements ", ITU-T Recommendation G.874, July 2010.

[ITU.G874.1]

International Telecommunications Union, "Optical transport network (OTN): Protocol-neutral management information model for the network element view ", ITU-T Recommendation G.874.1, January 2002.

[ITU.G959.1]

International Telecommunications Union, "Optical transport network physical layer interfaces ", ITU-T Recommendation G.959.1, November 2009.

[ITU.G826]

International Telecommunications Union, "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections ", ITU-T Recommendation G.826, November 2009.

[ITU.G8201]

International Telecommunications Union, "Error performance parameters and objectives for multi-operator international paths within the Optical Transport Network (OTN) ", ITU-T Recommendation G.8201, April 2011.

[ITU.G694.1]

International Telecommunications Union, "Spectral grids for WDM applications: DWDM frequency grid ", ITU-T Recommendation G.694.1, June 2002.

[ITU.G7710]

International Telecommunications Union, "Common equipment management function requirements ", ITU-T Recommendation G.7710, May 2008.

12.2. Informative References

- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart,
"Introduction and Applicability Statements for Internet-
Standard Management Framework", RFC 3410, December 2002.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629,
June 1999.
- [RFC4181] Heard, C., "Guidelines for Authors and Reviewers of MIB
Documents", BCP 111, RFC 4181, September 2005.
- [I-D.kunze-g-698-2-management-control-framework]
Kunze, R., "A framework for Management and Control of
optical interfaces supporting G.698.2", draft-
kunze-g-698-2-management-control-framework-00 (work in
progress), July 2011.
- [RFC4054] Strand, J. and A. Chiu, "Impairments and Other Constraints
on Optical Layer Routing", RFC 4054, May 2005.

Appendix A. Change Log

This optional section should be removed before the internet draft is submitted to the IESG for publication as an RFC.

Note to RFC Editor: please remove this appendix before publication as an RFC.

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Authors' Addresses

Gabriele Galimberti (editor)
Cisco
Via Philips,12
20052 - Monza
Italy

Phone: +390392091462
Email: ggalimbe@cisco.com

Ruediger Kunze (editor)
Deutsche Telekom
Dddd, xx
Berlin
Germany

Phone: +49xxxxxxxxxxx
Email: RKunze@telekom.de

Hing-Kam Lam (editor)
Alcatel-Lucent
600-700 Mountain Avenue, Murray Hill
New Jersey, 07974
USA

Phone: +17323313476
Email: kam.lam@alcatel-lucent.com

Dharini Hiremagalur (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: dharinih@juniper.net

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

. G.Galimberti, Ed.
Cisco
. R.Kunze, Ed.
Deutsche Telekom
. Kam Lam, Ed.
Alcatel-Lucent
. D. Hiremagalur, Ed.
Juniper
July 15, 2013

An SNMP MIB extension to RFC3591 to manage optical interface parameters
of DWDM applications
draft-galikunze-ccamp-opt-imp-snmp-mib-00

Abstract

This memo defines a module of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) in TCP/IP- based internets. In particular, it defines objects for managing Optical parameters associated with Dense Wavelength Division Multiplexing (DWDM) interfaces or characterized by the Optical Transport Network (OTN). This is an extension of the RFC3591 to support the optical parameters mainly but not only described in recommendations like ITU-T G.698.2. [ITU.G698.2]

The MIB module defined in this memo can be used for Optical Parameters monitoring and/or configuration of the endpoints of Black Links.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 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
2. The Internet-Standard Management Framework	4
3. Conventions	4
4. Overview	5
4.1. Optical Parameters Description	6
4.1.1. General	6
4.1.2. Parameters at Ss	8
4.1.3. Optical path from point Ss to Rs	9
4.1.4. Interface at point Rs	11
4.1.5. Alarms and Threshold definition	12
4.1.6. Performance Monitoring (PM) description	13
4.1.7. Generic Parameter description	15
4.2. Use of ifTable	15
4.2.1. Use of ifTable for OPS Layer	17
4.2.2. Use of ifTable for OCh Layer	18
4.2.3. Use of ifStackTable	18
5. Structure of the MIB Module	19
6. Object Definitions	19
7. Relationship to Other MIB Modules	77
7.1. Relationship to the [TEMPLATE TODO] MIB	77
7.2. MIB modules required for IMPORTS	77
8. Definitions	77
9. Security Considerations	77
10. IANA Considerations	78

11. Contributors	78
12. References	79
12.1. Normative References	79
12.2. Informative References	81
Appendix A. Change Log	82
Appendix B. Open Issues	82
Authors' Addresses	82

1. Introduction

This memo defines a portion of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) in TCP/IP- based internets. In particular, it defines objects for managing Optical parameters associated with Wavelength Division Multiplexing (WDM) systems or characterized by the Optical Transport Network (OTN) in accordance with but not limited to the optical interface defined in G.698.2 [ITU.G698.2]

This document further extend the draft-galikunze-ccamp-g-698-2-snmp-mib-03 where the ITU-T G.698.2 "application-code" and few other parameters described.

Black Link approach allows supporting an optical transmitter/receiver pair of one vendor to inject a DWDM channel and run it over an optical network composed of amplifiers, filters, add-drop multiplexers from a different vendor. From architectural point of view, the "Black Link" is a set of pre-configured/qualified network connections between the G.698.2 reference points S and R. The black links will be managed at the edges (i.e. the transmitters and receivers attached to the S and R reference points respectively) for the relevant parameters specified in G.698.2 [ITU.G698.2], G.798 [ITU.G798], G.874 [ITU.G874], and the performance parameters specified G.7710/Y.1701 [ITU-T G.7710] and and G.874.1 [ITU.G874.1].

The G.698.2 [ITU.G698.2] provides optical parameter values for physical layer interfaces of Dense Wavelength Division Multiplexing (DWDM) systems primarily intended for metro applications which include optical amplifiers. Applications are defined in G.698.2 [ITU.G698.2] using optical interface parameters at the single-channel connection points between optical transmitters and the optical multiplexer, as well as between optical receivers and the optical demultiplexer in the DWDM system. This Recommendation uses a methodology which does not specify the details of the optical link, e.g. the maximum fibre length, explicitly. The Recommendation currently includes unidirectional DWDM applications at 2.5 and 10 Gbit/s (with 100 GHz and 50 GHz channel frequency spacing). Work is still underway for 40 and 100 Gbit/s interfaces. There is possibility for extensions to a lower channel frequency spacing.

This draft refers and supports also the draft-kunze-g-698-2-management-control-framework

The building of an SNMP MIB describing the optical parameters defined in G.698.2 [ITU.G698.2] G.798 [ITU.G798], G.874 [ITU.G874], parameters specified G.7710/Y.1701 [ITU-T G.7710] allows the different vendors and operator to retrieve, provision and exchange information related to Optical black links in a standardized way. This facilitates interworking in case of using optical interfaces from different vendors at the end of the link.

The MIB, reporting the Optical parameters and their values, characterizes the features and the performances of the optical components and allow a reliable black link design in case of multivendor optical networks.

Although RFC 3591 [RFC3591] describes and defines the SNMP MIB of a number of key optical parameters, alarms and Performance Monitoring, a more complete description of optical parameters and processes can be found in the ITU-T Recommendations. Appendix A of this document provides an overview about the extensive ITU-T documentation in this area. The same considerations can be applied to the RFC 4054 [RFC4054]

2. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies a MIB module that is compliant to the SMIV2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].


3. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119] In the description of OIDs the convention: Set (S) Get (G) and Trap (T) conventions will describe the action allowed by the parameter.

4. Overview

In this document, the term OTN (Optical Transport Network) system is used to describe devices that are compliant with the requirements specified in the ITU-T Recommendations G.872 [ITU.G872], G.709 [ITU.G709], G.798 [ITU.G798], G.874 [ITU.G874], and G.874.1 [ITU.G874.1] while refers to G.698.2 [ITU.G698.2] for the Black Link and DWDM parameter description.

Figure 1 shows a set of reference points, for the linear "black link" approach, for single-channel connection (Ss and Rs) between transmitters (Tx) and receivers (Rx). Here the DWDM network elements include an OM and an OD (which are used as a pair with the opposing element), one or more optical amplifiers and may also include one or more OADMs.



Ss = reference point at the DWDM network element tributary output
 Rs = reference point at the DWDM network element tributary input
 Lx = Lambda x
 OM = Optical Mux
 OD = Optical Demux
 OADM = Optical Add Drop Mux

from Fig. 5.1/G.698.2

Figure 1: Linear Black Link

G.698.2 [ITU.G698.2] defines also Ring Black Link configurations [Fig. 5.2/G.698.2] and Bidirectional Black Link configurations [Fig. 5.3/G.698.2]

4.1. Optical Parameters Description

The black links are managed at the edges, i.e. at the transmitters (Tx) and receivers (Rx) attached to the S and R reference points respectively. The parameters that could be managed at the black link edges are specified in G.698.2 [ITU.G698.2] for the optical interface, in G.798 [ITU.G798] for the equipment aspect, and in G.7710/Y.1701 [ITU.G7710] and G.874 [ITU.G874] for fault management and performance monitoring.

The definitions of the optical parameters are provided below to increase the readability of the document, where the definition is ended by (G) the parameter can be retrieve with a GET, when (S) it can be provisioned by a SET, (G,S) can be either GET and SET.

To support the management of these parameters, the SNMP MIB in RFC 3591 [RFC3591] is extended with a new MIB module defined in section 6 of this document. This new MIB module includes the definition of new configuration table of the OCh Layer for the parameters at Tx (S) and Rx (R).

4.1.1. General

The following general parameters from G.698.2 [ITU.G698.2] and G.694.1 [ITU.G694.1] provide general information at the optical interface reference points.

Minimum channel spacing:

This is the minimum nominal difference in frequency (in GHz) between two adjacent channels (G).

Bit rate/line coding of optical tributary signals:

Optical tributary signal class NRZ 2.5G (from nominally 622 Mbit/s to nominally 2.67 Gbit/s) or NRZ 10G nominally 2.4 Gbit/s to nominally 10.71 Gbit/s. (nominally 2.4 Gbit/s to nominally 10.71 Gbit/s). 40Gbit/s and 100Gbit/s are under study (G, S).

FEC Coding:

This parameter indicate what Forward Error Correction (FEC) code is used at Ss and Rs (G, S) (not mentioned in G.698). EDITOR NOTE: Need to check whether this parameter is to be put in "vendor specific" parameter or can be a standard parameter as defined in G.698.2. Is this the various adaptations (FEC encoding types) specified in G.798 clauses 12.3.1.1 (with FEC), 12.3.1.2 (without FEC), and 12.3.1.5 (vendor-specific FEC) .

Maximum bit error ratio (BER):

This parameter indicate the maximum Bit error rate can be supported by the application at the Receiver. In case of FEC applications it is intended after the FEC correction (G) .

Fiber type:

Fiber type as per fibre types are chosen from those defined in ITU-T Recs G.652, G.653, G.654 and G.655 (G,S) .

Wavelength Range (see G.694.1): [ITU.G694.1]

This parameter indicate minimum and maximum wavelength spectrum (G) in a definite wavelength Band (L, C and S).

Wavelength Value (see G.694.1 Table 1):

This parameter indicates the wavelength value that Ss and Rs will be set to work (in THz) se in particular Section 6/G.694.1 (G, S).

Vendor Transceiver Class:

Other than specifying all the Transceiver parameter, it might be convenient for the vendors to summarize a set of parameters in a single proprietary parameter: the Class of transceiver. The Transceiver classification will be based on the Vendor Name and the main TX and RX parameters (i.e. Trunk Mode, Framing, Bit rate, Trunk Type, Channel Band, Channel Grid, Modulation Format, Channel Modulation Format, FEC Coding, Electrical Signal Framing at Tx, Minimum maximum Chromatic Disperion (CD) at Rx, Maximum Polarization Mode Dispersion (PMD) at Rx, Maximum differential group delay at Rx, Loopbacks, TDC, Pre-FEC BER, Q-factor, Q-margin,etc.). If this parameter is used, the MIB parameters specifying the Transceiver characteristics may not be significant and the vendor will be responsible to specify the Class contents and values. The Vendor can publish the parameters of its Classes or declare to be compatible with published Classes.(G) Optional for compliance. (not mentioned in G.698)

Single-channel application codes (see G.698.2):

This parameter indicates the transceiver application code at Ss and Rs as defined in [ITU.G698.2] Chapter 5.4 - this parameter can be called Optical Interface Identifier OII as per [draft-martinelli-wson-interface-class] (G).

PARAMETERS	Get/Set	Reference
Minimum channel spacing	G	G.698.2 S.7.1.1
Bit rate/line coding of opt. trib. signals	G,S	G.698.2 S.7.1.2
FEC Coding	G,S	G.975
Maximum bit error ratio (BER)	G	G.698.2 S.7.1.3
Fiber type	G,S	G.698.2 S.7.1.4
Wavelength Range	G	G.694.1 S.6
Wavelength Value	G,S	G.694.1 S.6
Vendor Transceiver Class	G	N.A.
Single-channel application codes	G	G.698.2 S.5.3

Table 1: General parameters

4.1.2. Parameters at Ss

The following parameters for the interface at point S are defined in G.698.2 [ITU.G698.2].

Maximum and minimum mean channel output power:

The mean launched power at Ss is the average power (in dBm) of a pseudo-random data sequence coupled into the DWDM link. It is defined as the rrange (Max and Min) of the parameter (G, S)

Minimum and maximum central frequency:

The central frequency is the nominal single-channel frequency (in THz) on which the digital coded information of the particular optical channel is modulated by use of the NRZ line code. The central frequencies of all channels within an application lie on the frequency grid for the minimum channel spacing of the application given in ITU-T Rec. G.694.1. This parameter give the Maximum and minimum frequency interval the channel must be modulated (G)

Maximum spectral excursion:

This is the maximum acceptable difference between the nominal central frequency (in GHz) of the channel and the minus 15 dB points of the transmitter spectrum furthest from the nominal central frequency measured at point Ss. (G)

Maximum transmitter (residual) dispersion OSNR penalty (B.3/G.959.1) [ITU.G959.1]

Defines a reference receiver that this penalty is measured with.
Lowest OSNR at Ss with worst case (residual) dispersion minus the
Lowest OSNR at Ss with no dispersion. Lowest OSNR at Ss with no
dispersion (G)

Minimum side mode suppression ratio, Minimum channel extinction
ratio, Eye mask:

Although are defined in G.698.2 are not supported by this draft
(G).

Current Laser Output power:

This parameter report the current Transceiver Output power, it can
be either a setting and measured value (G, S) NEED TO DISCUSS ON
THIS.

PARAMETERS	Get/Set	Reference
MAX and min mean channel output power	G,S	G.698.2 S.7.2.1
Min and MAX central frequency	G	G.698.2 S.7.2.2
MAX spectral excursion	G	G.698.2 S.7.2.3
MAX transmitter (residual) disper. OSNR penalty	G	G.698.2 S.7.2.7
MAX side mode suppression ratio, min channel extinction ratio, Eye mask	G	G.698.2 S.7.2.6
Current Laser Output power	G,S	N.A.

Table 2: parameters at Ss

4.1.3. Optical path from point Ss to Rs

The following parameters for the optical path from point S and R are defined in G.698.2 [ITU.G698.2].

Maximum and minimum (residual) chromatic dispersion:

These parameters define the maximum and minimum value of the
optical path "end to end chromatic dispersion" (in ps/nm) that the
system shall be able to tolerate. (G)

Minimum optical return loss at Ss:

These parameter defines minimum optical return loss (in dB) of the
cable plant at the source reference point (Ss), including any
connectors (G)

Maximum discrete reflectance between Ss and Rs:

Optical reflectance is defined to be the ratio of the reflected optical power present at a point, to the optical power incident to that point. Control of reflections is discussed extensively in ITU-T Rec. G.957 (G)

Maximum differential group delay:

Differential group delay (DGD) is the time difference between the fractions of a pulse that are transmitted in the two principal states of polarization of an optical signal. For distances greater than several kilometres, and assuming random (strong) polarization mode coupling, DGD in a fibre can be statistically modelled as having a Maxwellian distribution. (G)

Maximum polarization dependent loss:

The polarization dependent loss (PDL) is the difference (in dB) between the maximum and minimum values of the channel insertion loss (or gain) of the black link from point SS to RS due to a variation of the state of polarization (SOP) over all SOPs. (G)

Maximum inter-channel crosstalk:

Inter-channel crosstalk is defined as the ratio of total power in all of the disturbing channels to that in the wanted channel, where the wanted and disturbing channels are at different wavelengths. The parameter specify the isolation of a link conforming to the "black link" approach such that under the worst-case operating conditions the inter-channel crosstalk at any reference point RS is less than the maximum inter-channel crosstalk value (G)

Maximum interferometric crosstalk:

This parameter places a requirement on the isolation of a link conforming to the "black link" approach such that under the worst case operating conditions the interferometric crosstalk at any reference point RS is less than the maximum interferometric crosstalk value. (G)

Maximum optical path OSNR penalty:

The optical path OSNR penalty is defined as the difference between the Lowest OSNR at Rs and Lowest OSNR at Ss that meets the BER requirement (G)

Maximum ripple:

Although is defined in G.698.2, this parameter is not supported by this draft.

PARAMETERS	Get/Set	Reference
------------	---------	-----------

MAX and min (residual) chromatic dispersion	G	G.698.2 S.7.3.2
Min optical return loss at Ss	G	G.698.2 S.7.3.3
MAX discrete reflectance between Ss and Rs	G	G.698.2 S.7.3.4
MAX differential group delay	G	G.698.2 S.7.3.5
MAX polarization dependent loss	G	G.698.2 S.7.3.6
MAX inter-channel crosstalk	G	G.698.2 S.7.3.7
MAX interferometric crosstalk	G	G.698.2 S.7.3.8
MAX optical path OSNR penalty	G	G.698.2 S.7.3.9
MAX ripple	G	G.698.2 S.7.3.1

Table 3: parameters between Ss and Rs

4.1.4. Interface at point Rs

The following parameters for the interface at point R are defined in G.698.2.

4.1.4.1. Mandatory parameters

Maximum and minimum mean input power:

The maximum and minimum values of the average received power (in dBm) at point Rs. (G)

Minimum optical signal-to-noise ratio (OSNR):

The minimum optical signal-to-noise ratio (OSNR) is the minimum value of the ratio of the signal power in the wanted channel to the highest noise power density in the range of the central frequency plus and minus the maximum spectral excursion (G)

Receiver OSNR tolerance:

The receiver OSNR tolerance is defined as the minimum value of OSNR at point Rs that can be tolerated while maintaining the maximum BER of the application. (G)

Maximum reflectance at receiver:

Although is defined in G.698.2, this parameter is not supported by this draft (G).

PARAMETERS	Get/Set	Reference
MAX and min mean input power	G	G.698.2 S.7.4.1
Min optical signal-to-noise ratio (OSNR)	G	G.698.2 S.7.4.2
Receiver OSNR tolerance	G	G.698.2 S.7.4.3

MAX reflectance at receiver	G	G.698.2 S.7.4.4	
-----	-----	-----	-----

Table 4: mandatory parameters

4.1.4.2. Optional parameters

Current Chromatic Dispersion (CD):

Residual Chromatic Dispersion measuread at Rx Transceiver port (G).

Current Optical Signal to Noise Ratio (OSNR):

Current Optical Signal to Noise Ratio (OSNR) estimated at Rx Transceiver port (G).

Current Quality factor (Q):

"Q" factor estimated at Rx Transceiver port (G).

PARAMETERS	Get/Set	Reference
Current Chromatic Dispersion (CD)	G	N.A.
Current Opt. Signal to Noise Ratio (OSNR)	G	N.A.
Current Quality factor (Q)	G	N.A.

Table 5: optional parameters

4.1.5. Alarms and Threshold definition

This section describes the Alarms and the Thresholds at Ss and Rs points according to ITU-T Recommendations G.798 [ITU.G798], G.874 [ITU.G874], and G.874.1 [ITU.G874.1].

OTN alarms defined in RFC3591:

Threshold Crossing Alert (TCA Alarm)

LOW-TXPOWER

HIGH-TXPOWER

LOW-RXPOWER

HIGH-RXPOWER

Loss of Signal (LOS)

Loss of Frame (LOF)

Server Signal Failure-P (SSF-P)

Loss of Multiframe (LOM)

OTN Thresholds (for TCA) defined in RFC3591

LOW-TXPOWER

HIGH-TXPOWER

LOW-RXPOWER

HIGH-RXPOWER

As the above parameters/alarms are already defined in RFC3591, they are out of scope of this document and the RFC3591 will continue to be the only reference for them

The list below reports the new Alarms and Thresholds not managed in RFC3591

4.1.6. Performance Monitoring (PM) description

This section describes the Performance Monitoring parameters and their thresholds at Ss and Rs points (Near -End and Far-End) according to ITU-T Recommendations G.826 [ITU.G826], G.8201 [ITU.G8201], G.709 [ITU.G709], G.798 [ITU.G798], G.874 [ITU.G874], and G.874.1 [ITU.G874.1].

BIP8 :

Number of BIP8's occurred in an observation period (G). BIP-8 consists of a parity byte calculated bit-wise across a large number of bytes in a transmission transport frame. BIP-8 bits are set such that the overall data stream, including the BIP-8 byte, has even parity.

Errored Second (es) :

It is a one-second period in which there is one or more errored blocks or during which a defect (e.g. Loss of Signal (LOS)) is detected. The number of errored seconds is summed over 15-minute and 24-hour intervals. (G)

Severely Errored Seconds (ses) :

It is a one-second period in which the errored block ratio exceeds a threshold or during which a defect is detected. See ITU-T Recommendation G.8201 Table 7-1 for details. The number of

severely errored seconds is summed over 15-minute and 24-hour intervals.(G)

Unavailable Seconds (uas) :

It is a one-second period in the unavailable time. A period of unavailable time begins at the onset of ten consecutive SES events. These ten seconds are considered to be part of unavailable time. A new period of available time begins at the onset of ten consecutive non-SES events. These ten seconds are considered to be part of available time. (G)

Background Block Errors (bbe) :

An errored block not occurring as part of an SES(G)

Error Seconds Ratio (esr) :

The ratio of ES in available time to total seconds in available time during a fixed measurement interval(G)

Severely Errored Seconds Ratio (sesr) :

The ratio of SES in available time to total seconds in available time during a fixed measurement interval(G)

Background Block Errored Seconds Ratio (bber) :

The ratio of Background Block Errors (BBE) to total blocks in available time during a fixed measurement interval. The count of total blocks excludes all blocks during SESs.(G)

FEC corrected Bit Error (FECcorrErr):

The number of bits corrected by the FEC are counted over one second (G)

FEC un-corrected Bit Error :

The number of bits un-corrected by the FEC are counted over one second (G)

Pre-FEC Bit Error :

The number of Errored bits at receiving side before the FEC function counted over one second (G)

OTN Valid Intervals :

The number of contiguous 15 minute intervals for which valid OTN performance monitoring data is available for the particular interface (G)

FEC Valid Intervals :

The number of contiguous 15 minute intervals for which valid FEC PM data is available for the particular interface.(G)

4.1.7. Generic Parameter description

This section describes the Generic Parameters at Ss and Rs points according to ITU-T Recommendations G.872 [ITU.G872], G.709 [ITU.G709], G.798 [ITU.G798], G.874 [ITU.G874], and G.874.1 [ITU.G874.1].

Interface Admin Status :

The Administrative Status of an Interface: Up/Down - In Service/
Out of Service (can be Automatic in Service) (G/S)

Interface Operational Status :

The Operational Status of an Interface: Up/Down - In Service/Out
of Service (G)

4.2. Use of ifTable

This section specifies how the MIB II interfaces group, as defined in RFC 2863 [RFC2863], is used for the link ends of a black link. Only the ifGeneralInformationGroup will be supported for the ifTable and the ifStackTable to maintain the relationship between the OCh and OPS layers. The OCh and OPS layers are managed in the ifTable using IfEntries that correlate to the layers depicted in Figure 1.

For example, a device with TX and/or RX will have an Optical Physical Section (OPS) layer, and an Optical Channel (OCh) layer. There is a one to n relationship between the OPS and OCh layers.

EDITOR NOTE: Reason for changing from OChr to OCh: Work on revised G.872 in the SG15 December 2011 meeting agreed to remove OChr from the architecture and to update G.709 to account for this architectural change. The meeting also agreed to consent the revised text of G.872 and G.709 at the September 2012 SG15 meeting.

Figure 2 In the following figures, opticalChannel and opticalPhysicalSection are abbreviated as och and ops respectively.




Figure 2: OTN Layers for OPS and OCh

Each opticalChannel IfEntry is mapped to one of the m opticalPhysicalSection IfEntries, where m is greater than or equal to 1. Conversely, each opticalTransPhysicalSection port entry is mapped to one of the n opticalChannel IfEntries, where n is greater than or equal to 1.

The design of the Optical Interface MIB provides the option to model an interface either as a single bidirectional object containing both sink and source functions or as a pair of unidirectional objects, one containing sink functions and the other containing source functions.

If the sink and source for a given protocol layer are to be modelled as separate objects, then there need to be two ifTable entries, one that corresponds to the sink and one that corresponds to the source, where the directionality information is provided in the configuration tables for that layer via the associated Directionality objects. The agent is expected to maintain consistent directionality values between ifStackTable layers (e.g., a sink must not be stacked in a 1:1 manner on top of a source, or vice-versa), and all protocol layers that are represented by a given ifTable entry are expected to have the same directionality.

When separate ifTable entries are used for the source and sink functions of a given physical interface, association between the two uni-directional ifTable entries (one for the source function and the other for the sink functions) should be provided. It is recommended

that identical ifName values are used for the two ifTable entries to indicate such association. An implementation shall explicitly state what mechanism is used to indicate the association, if ifName is not used.

4.2.1. Use of ifTable for OPS Layer

Only the ifGeneralInformationGroup needs to be supported.

ifTable Object	Use for OTN OPS Layer
=====	
ifIndex	The interface index.
ifDescr	Optical Transport Network (OTN) Optical Physical Section (OPS)
ifType	opticalPhysicalSection (xxx)
<<<Editor Note: Need new IANA registration value for xxx. >>>	
ifSpeed	Actual bandwidth of the interface in bits per second. If the bandwidth of the interface is greater than the maximum value of 4,294,967,295, then the maximum value is reported and ifHighSpeed must be used to report the interface's speed.
ifPhysAddress	An octet string with zero length. (There is no specific address associated with the interface.)
ifAdminStatus	The desired administrative state of the interface. Supports read-only access.
ifOperStatus	The operational state of the interface. The value lowerLayerDown(7) is not used, since there is no lower layer interface. This object is set to notPresent(6) if a component is missing, otherwise it is set to down(2) if either of the objects optIfOPSnCurrentStatus indicates that any defect is present.
ifLastChange	The value of sysUpTime at the last change in ifOperStatus.
ifName	Enterprise-specific convention (e.g., TL-1 AID)

to identify the physical or data entity associated with this interface or an OCTET STRING of zero length. The enterprise-specific convention is intended to provide the means to reference one or more enterprise-specific tables.

ifLinkUpDownTrapEnable Default value is enabled(1). Supports read-only access.

ifHighSpeed Actual bandwidth of the interface in Mega-bits per second. A value of n represents a range of 'n-0.5' to 'n+0.499999'.

ifConnectorPresent Set to true(1).

ifAlias The (non-volatile) alias name for this interface as assigned by the network manager.

4.2.2. Use of ifTable for OCh Layer

Use of ifTable for OCh Layer See RFC 3591 [RFC3591] section 2.4

4.2.3. Use of ifStackTable

Use of the ifStackTable and ifInvStackTable to associate the opticalPhysicalSection and opticalChannel interface entries is best illustrated by the example shown in Figure 3. The example assumes an ops interface with ifIndex i that carries two multiplexed och interfaces with ifIndex values of j and k, respectively. The example shows that j and k are stacked above (i.e., multiplexed into) i. Furthermore, it shows that there is no layer lower than i and no layer higher than j and/or k.

Figure 3

HigherLayer	LowerLayer
0	j
0	k
j	i
k	i
i	0

Figure 3: Use of ifStackTable for an OTN port

For the inverse stack table, it provides the same information as the interface stack table, with the order of the Higher and Lower layer interfaces reversed.

5. Structure of the MIB Module

EDITOR NOTE: text will be provided based on the MIB module in Section 6

6. Object Definitions

EDITOR NOTE: Once the scope in Section 1 and the parameters in Section 4 are finalized, a MIB module will be defined. It could be an extension to the OPT-IF-MIB module of RFC 3591. >>>

```
PT-IF-EXT-MIB DEFINITIONS ::= BEGIN
```

```
IMPORTS
```

```
    MODULE-IDENTITY,
    OBJECT-TYPE,
    Gauge32,
    Integer32,
    Unsigned32,
    Counter64,
    transmission,
    NOTIFICATION-TYPE
    FROM SNMPv2-SMI
    TEXTUAL-CONVENTION,
    RowPointer,
    RowStatus,
    TruthValue,
    DisplayString,
    DateAndTime
```

```

        FROM SNMPv2-TC
    SnmpAdminString
        FROM SNMP-FRAMEWORK-MIB
    MODULE-COMPLIANCE, OBJECT-GROUP
        FROM SNMPv2-CONF
    ifIndex
        FROM IF-MIB
    OptIfDirectionality,
    optIfOChConfigEntry,
    optIfOChSinkCurrentEntry,
    optIfMibModule
        FROM OPT-IF-MIB;

-
-- This is the MIB module for the optical parameters associated with
-- the black link end points
-- extension for Alarms and PM

optIfExtMibModule MODULE-IDENTITY
    LAST-UPDATED "201307020000Z"
    ORGANIZATION "IETF Ops/Camp MIB Working Group"
    CONTACT-INFO
        "WG charter:
         http://www.ietf.org/html.charters/

        Mailing Lists:
        Editor: Gabriele Galimberti
        Email: ggalimbe@cisco.com"
    DESCRIPTION
        "The MIB module to describe Black Link extension to rfc3591.

        Copyright (C) The Internet Society (2013). This version
        of this MIB module is part of ; see the RFC
        itself for full legal notices."
    REVISION "201307020000Z"
    DESCRIPTION
        " Draft version 1.0
        Mib has all transceiver capabilities as described in G.698,
        alarm and PM's."
    ::= { optIfMibModule 3 }
```

OptIfChannelSpacing ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Channel spacing

- 1 - 100 Ghz
- 2 - 50GHz
- 3 - 25GHz
- 4 - 12.5GHz
- 5 - 6.25Ghz

"

SYNTAX INTEGER {
 spacing100Ghz(1),
 spacing50Ghz(2),
 spacing25Ghz(3),
 spacing12point5Ghz(4),
 spacing6point25Ghz(5)
}

OptIfBitRateLineCoding ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Optical tributary signal class

- 1 - NRZ 2.5G from nominally 622 Mbit/s to nominal 2.67 Gbit/s)
- 2 - NRZ 10G nominally 2.4 Gbit/s to nominally 10.71 Gbit/s.
- 3 - 40Gbits/s
- 4 - 100Gbits/s
- 5 - 400Gbits/s

40Gbits/s and above are under study. "

SYNTAX INTEGER {
 rate2point5G(1),
 rate10G(2),
 rate40G(3),
 rate100G(4),
 rate400G(5)
}

OptIfFiberTypeRecommendation ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

" Fiber Types - ITU-T Recs G.652, G.653, G.654 and G.655

One for recommendation and one for category.

- G.652 A, B, C, D
- G.653 A, B
- G.654 A, B, C
- G.655 C, D, E
- G.656
- G.657 A, B


```
"
SYNTAX INTEGER {
    g652(1),
    g653(2),
    g654(3),
    g655(4),
    g656(5),
    g657(6)
}

OptIfFiberTypeCategory ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        " Fiber Types - ITU-T Recs G.652, G.653, G.654 and G.655
          G.652 A, B, C, D
          G.653 A, B
          G.654 A, B, C
          G.655 C, D, E
          G.656
          G.657 A, B
          Categories - A, B, C, D and E
        "
    SYNTAX INTEGER {
        categoryA(1),
        categoryB(2),
        categoryC(3),
        categoryD(4),
        categoryE(5)
    }

OptIfGridTypes ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        " The types of Grid as defined by ITU-T
        "
    SYNTAX INTEGER {
        gridReserved(0),
        gridITUDWDM(1),
        gridITUCWDM(2),
        gridITUFlex(4)
    }

OptIfDataType ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        "
        This parameter indicates the parameters for the table are for
        the Near End or Far End data.
```

```

    1 - Near End
    2 - Far End
"
SYNTAX INTEGER {
    nearEnd(1),
    farEnd(2)
}

OptIfOTNLayer ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        "
        This parameter indicates the parameters for the table are for
        OTUk, ODUk, TCM performance data.
        1 - OTUk
        2 - ODUk
        3 - TCM
        The ODUk layer and TCM sublayer PM is not related to the black
        link PM management, but since this could be a common PM model
        for the ODUk layer and TCM sublayers, they are included here so
        it may be used for simple scenarios where only lower order ODUk
        or higher order ODUk is present. For scenarios where both lower
        order ODUk and higher order ODUk are present, further extension
        to the MIBmodel is required, in particular for the indexing for
        these layers.
        "
    SYNTAX INTEGER {
        optIfOTUkLayer(1),
        optIfODUkLayer(2),
        optIfTCMSubLayer(3)
    }

optIfOTNAlarmSeverity ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        " Severity of the notification.
        "
    SYNTAX INTEGER {
        optIfCritical(1),
        optIfMajor(2),
        optIfMinor(3),
        optIfInfo(4)
    }
--
-- Alarm for the OCh and OTUk layer
--
OptIfOTNChAlarms ::= TEXTUAL-CONVENTION
    STATUS current
```

DESCRIPTION

" This is the possible alarms from the OCh and OTUk layer."

```
SYNTAX INTEGER {
    -- OTN Loss of signal alarm
    optIfOtnLosAlarm(1),
    -- OTN Loss of Frame alarm
    optIfOtnLofAlarm(2),
    -- OTN Loss of Multi Frame alarm
    optIfOtnLomAlarm(3),
    -- OTN Server Signal Failure alarm
    optIfOtnOtuSsfAlarm(4),
    -- OTN OTU Backward Defect Indicator alarm
    optIfOtnOtuBdiAlarm(5),
    -- OTN OTU Trail Trace Identifier Mismatch alarm
    optIfOtnOtuTimAlarm(6),
    -- OTN OTU Degraded alarm,
    optIfOtnOtuDegAlarm(7),
    -- OTN OTU Fec Excessive Errors alarm
    optIfOptIfOtnOtuFecExcessiveErrsAlarm(8),
    -- OTN OTU BBE Thresholdalarm
    optIf15MinThreshBBETCA(9),
    -- OTN OTU ES Thresholdalarm
    optIf15MinThreshESTCA(10),
    -- OTN OTU SES Threshold alarm
    optIf15MinThreshSESTCA(11),
    -- OTN OTU UAS Threshold alarm
    optIf15MinThreshUASTCA(12),
    -- OTN OTU Bip8 Threshold alarm
    optIf15MinThreshBip8TCA(13),
    -- OTN FEC uncorrectedwords TCA
    optIf15MinThreshFECUnCorrectedWordsTCA(14),
    -- OTN Pre FEC BER TCA
    optIf15MinThreshPreFECBERTCA(15)
}
```

OptIfOTNODUkTcmAlarms ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

" This is the alarms from the ODUk and TCM layer."

```
SYNTAX INTEGER {
    -- OTN ODU/TCM Open Connection Indicator
    optIfOTNOdukTcmOciAlarm(1),
    -- OTN ODU/TCM LCK
    optIfOTNOdukTcmLckAlarm(2),
    -- OTN ODU/TCM Backward Defect Indicator
    optIfOTNOdukTcmBdiAlarm(3),
    -- OTN ODU/TCM Trail Trace Identifier Mismatch
    optIfOTNOdukTcmTimAlarm(4),
}
```

```

-- OTN ODU/TCM Degraded
optIfOTNoduKtcmDegAlarm(5),
-- OTN ODU/TCM LTC Loss of Tandem connection
optIfOTNoduKtcmLtcAlarm(6),
-- OTN ODU/TCM CSF - Client Signal Failure
optIfOTNoduKtcmCSfAlarm(7),
-- OTN ODU/TCM Server Signal Failure
optIfOTNoduKtcmSSfAlarm(8),
-- OTN OTU BBE Threshold
optIfOTNoduKtcm15MinThreshBBETCA(9),
-- OTN OTU ES Threshold
optIfOTNoduKtcm15MinThreshESTCA(10),
-- OTN OTU SES Threshold
optIfOTNoduKtcm15MinThreshSESTCA(11),
-- OTN OTU UAS Threshold
optIfOTNoduKtcm15MinThreshUASTCA(12)
}

-- Addition to the RFC 3591 objects
optIfOTNNotifications OBJECT IDENTIFIER ::= { optIfExtMibModule 0 }
optIfOPSmEntry          OBJECT IDENTIFIER ::= { optIfExtMibModule 1 }
optIfOChXcvrGroup       OBJECT IDENTIFIER ::= { optIfExtMibModule 2 }
optIfOTNPMGroup         OBJECT IDENTIFIER ::= { optIfExtMibModule 3 }
optIfOTNAlarm           OBJECT IDENTIFIER ::= { optIfExtMibModule 4 }

-- OPS - Optical Physical Section
optIfOPSmConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOPSmConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A table of OPS General config parameters."
    ::= { optIfOPSmEntry 1 }

optIfOPSmConfigEntry OBJECT-TYPE
    SYNTAX OptIfOPSmConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "An conceptual row of OPS General config parameters."
    INDEX { ifIndex }
    ::= { optIfOPSmConfigTable 1 }

OptIfOPSmConfigEntry ::=

```

```
SEQUENCE {
    optIfOPSmDirectionality
                                OptIfDirectionality,
    optIfOPSmFiberTypeRecommendation
                                OptIfFiberTypeRecommendation,
    optIfOPSmFiberTypeCategory
                                OptIfFiberTypeCategory
}

optIfOPSmDirectionality OBJECT-TYPE
    SYNTAX      OptIfDirectionality
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Indicates the directionality of the entity."
    ::= { optIfOPSmConfigEntry 1 }

optIfOPSmFiberTypeRecommendation OBJECT-TYPE
    SYNTAX      OptIfFiberTypeRecommendation
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Fiber type as per fibre types are chosen from those defined
         in ITU-T Recs G.652, G.653, G.654, G.655, G.656 and G.657."
    ::= { optIfOPSmConfigEntry 2 }

optIfOPSmFiberTypeCategory OBJECT-TYPE
    SYNTAX      OptIfFiberTypeCategory
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Fiber type as per fibre types are chosen from those defined
         in ITU-T Recs G.652, G.653, and G.655.
         The categories are A, B, C, D and E."
    ::= { optIfOPSmConfigEntry 3 }

-- Tranceiver general parameters table
-- General parameters for the Black Link Ss-Rs

optIfXcvrTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A table of Och General config extension parameters"
    ::= { optIfOChXcvrGroup 1 }
```

```

optIfXcvrEntry OBJECT-TYPE
    SYNTAX      OptIfXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual row that contains OCh configuration extension
        information of an interface."
    AUGMENTS { optIfOChConfigEntry }
    ::= { optIfXcvrTable 1 }

OptIfXcvrEntry ::=
    SEQUENCE {
        optIfOChMimumumChannelSpacing      OptIfChannelSpacing,
        optIfOChBitRateLineCoding          OptIfBitRateLineCoding,
        optIfOChFEC                        Unsigned32,
        optIfOChSinkMaximumBERMantissa      Unsigned32,
        optIfOChSinkMaximumBERExponent      Unsigned32,
        optIfOChMinWavelength               Unsigned32,
        optIfOChMaxWavelength               Unsigned32
    }

optIfOChMimumumChannelSpacing OBJECT-TYPE
    SYNTAX      OptIfChannelSpacing
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "A minimum nominal difference in frequency (GHz) between two
        adjacent channels."
    ::= { optIfXcvrEntry 1 }

optIfOChBitRateLineCoding OBJECT-TYPE
    SYNTAX      OptIfBitRateLineCoding
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Optical tributary signal class
        NRZ 2.5G (from nominally 622 Mbit/s to nominally 2.67
        Gbit/s) NRZ 10G (nominally 2.4 Gbit/s to nominally 10.71
        Gbit/s) "
    ::= { optIfXcvrEntry 2 }

optIfOChFEC OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION
        "This parameter indicates what Forward Error Correction

```

(FEC) code is used at Source and Sink.
GFEC (from G709) and the I.x EFEC's
(G.975 - Table I.1 super FEC).

- 1 - No FEC
- 2 - GFEC
- 3 - I.2 EFEC
- 4 - I.3 EFEC
- 5 - I.4 EFEC
- 6 - I.5 EFEC
- 7 - I.6 EFEC
- 8 - I.7 EFEC
- 9 - I.8 EFEC
- 10 - I.9 EFEC
- 11 - 100G FEC (for new applications)
- 12 - 100G EFEC (for new applications)
- 99 - Vendor Specific "

::= { optIfXcvrEntry 3 }

optIfOchSinkMaximumBERMantissa OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This parameter indicate the maximum Bit(mantissa) error
rate can be supported by the application at the Receiver.
In case of FEC applications it is intended after the FEC
correction."

::= { optIfXcvrEntry 4 }

optIfOchSinkMaximumBERExponent OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This parameter indicate the maximum Bit(exponent) error
rate can be supported by the application at the Receiver.
In case of FEC applications it is intended after the FEC
correction."

::= { optIfXcvrEntry 5 }

optIfOchMinWavelength OBJECT-TYPE

SYNTAX Unsigned32

UNITS "0.01 nm"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This parameter indicate minimum wavelength spectrum in a
definite wavelength Band (L, C and S) "

```

 ::= { optIfXcvrEntry 6 }

optIfOchMaxWavelength OBJECT-TYPE
    SYNTAX      Unsigned32
    UNITS       "0.01 nm"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This parameter indicate maximum wavelength spectrum in a
        definite wavelength Band (L, C and S) "
    ::= { optIfXcvrEntry 7 }

-- Tranceiver (xcvr) Parameters at OCh Src (Ss)
-- OptIfOchSrcXcvrEntry

optIfOchSrcXcvrTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfOchSrcXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A configuration table of OCh Src (Ss) parameters."
    ::= { optIfOchXcvrGroup 2 }

optIfOchSrcXcvrEntry OBJECT-TYPE
    SYNTAX      OptIfOchSrcXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual row that contains the Src (Ss) configuration
        parameters for a given interface."
    INDEX       { ifIndex }
    ::= { optIfOchSrcXcvrTable 1 }

OptIfOchSrcXcvrEntry ::=
    SEQUENCE {
        optIfOchMinimumMeanChannelOutputPower      Integer32,
        optIfOchMaximumMeanChannelOutputPower      Integer32,
        optIfOchMinimumCentralFrequency             Unsigned32,
        optIfOchMaximumCentralFrequency             Unsigned32,
        optIfOchMaximumSpectralExcursion            Unsigned32,
        optIfOchMaximumTxDispersionOSNRPenalty      Integer32
    }

optIfOchMinimumMeanChannelOutputPower OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.01 dbm"
    MAX-ACCESS  read-only
    STATUS      current

```



```
DESCRIPTION
    "The minimum mean launched power at Ss is the average power
    (in dbm) of a pseudo-random data sequence coupled into the
    DWDM link."
 ::= { optIfOChSrcXcvrEntry 1}

optIfOChMaximumMeanChannelOutputPower OBJECT-TYPE
    SYNTAX Integer32
    UNITS      "0.01 dbm"
    MAX-ACCESS read-only
    STATUS     current
    DESCRIPTION
        "The maximum mean launched power at Ss is the average power
        (in dbm) of a pseudo-random data sequence coupled into the
        DWDM link."
    ::= { optIfOChSrcXcvrEntry 2}

optIfOChMinimumCentralFrequency OBJECT-TYPE
    SYNTAX Unsigned32
    UNITS      "0.01 THz"
    MAX-ACCESS read-only
    STATUS     current
    DESCRIPTION
        "The minimum central frequency is the nominal single-channel
        frequency (in THz) on which the digital coded information of
        the particular optical channel is modulated by use of the
        NRZ line code. Eg 191.5THz will be represented as 19150 "
    ::= { optIfOChSrcXcvrEntry 3}

optIfOChMaximumCentralFrequency OBJECT-TYPE
    SYNTAX Unsigned32
    UNITS      "0.01 THz"
    MAX-ACCESS read-only
    STATUS     current
    DESCRIPTION
        "The maximum central frequency is the nominal single-channel
        frequency(in THz) on which the digital coded information of
        the particular optical channel is modulated by use of the
        NRZ line code. Eg 191.5THz will be represented as 19150 "
    ::= { optIfOChSrcXcvrEntry 4}

optIfOChMaximumSpectralExcursion OBJECT-TYPE
    SYNTAX Unsigned32
    UNITS      "0.1 GHz"
    MAX-ACCESS read-only
    STATUS     current
    DESCRIPTION
        "This is the maximum acceptable difference between the
```

```

        nominal central frequency (in GHz) of the channel and the
        minus 15 dB points of the transmitter spectrum furthest
        from the nominal central frequency measured at point Ss."
 ::= { optIfOChSrcXcvrEntry 5}

optIfOChMaximumTxDispersionOSNRPenalty OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.1 dB"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Defines a reference receiver that this penalty is measured
        with. Lowest OSNR at Ss with worst case (residual)
        dispersion minus the lowest OSNR at Ss with no dispersion.
        Lowest OSNR at Ss with no dispersion. "
    ::= { optIfOChSrcXcvrEntry 6}

-- Optical Path from Point Src (Ss) to Sink (Rs)
-- Alternatively this can be optIfOChSsRsTable

optIfOChSrcSinkXcvrTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfOChSrcSinkXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A table of paramters for the optical path from Src to Sink
        (Ss to Rs). "
    ::= { optIfOChXcvrGroup 3 }

optIfOChSrcSinkXcvrEntry OBJECT-TYPE
    SYNTAX      OptIfOChSrcSinkXcvrEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual row that contains the optical path Src-Sink
        (Ss-Rs) configuration parameters for a given interface."
    INDEX       { ifIndex }
    ::= { optIfOChSrcSinkXcvrTable 1 }

OptIfOChSrcSinkXcvrEntry ::=
    SEQUENCE {
        optIfOChSrcSinkMinimumChromaticDispersion
                                                    Integer32,
        optIfOChSrcSinkMaximumChromaticDispersion
                                                    Integer32,
        optIfOChSrcSinkMinimumSrcOpticalReturnLoss
                                                    Integer32,
        optIfOChSrcSinkMaximumDiscreteReflectanceSrcToSink
    }

```

```

        Integer32,
    optIfOChSrcSinkMaximumDifferentialGroupDelay
        Integer32,
    optIfOChSrcSinkMaximumPolarisationDependentLoss
        Integer32,
    optIfOChSrcSinkMaximumInterChannelCrosstalk
        Integer32,
    optIfOChSrcSinkInterFerometricCrosstalk
        Integer32,
    optIfOChSrcSinkOpticalPathOSNRPenalty
        Integer32
}

optIfOChSrcSinkMinimumChromaticDispersion OBJECT-TYPE
SYNTAX Integer32
UNITS "ps/nm"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "These parameters define the minimum value of the
    optical path 'end to end chromatic dispersion' (in ps/nm)
    that the system shall be able to tolerate."
 ::= { optIfOChSrcSinkXcvrEntry 1}

optIfOChSrcSinkMaximumChromaticDispersion OBJECT-TYPE
SYNTAX Integer32
UNITS "ps/nm"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "These parameters define the maximum value of the
    optical path 'end to end chromatic dispersion' (in ps/nm)
    that the system shall be able to tolerate."
 ::= { optIfOChSrcSinkXcvrEntry 2 }

optIfOChSrcSinkMinimumSrcOpticalReturnLoss OBJECT-TYPE
SYNTAX Integer32
UNITS "0.1 dB"
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "These parameter defines minimum optical return loss (in dB)
    of the cable plant at the source reference point (Src/Ss),
    including any connectors."
 ::= { optIfOChSrcSinkXcvrEntry 3 }

optIfOChSrcSinkMaximumDiscreteReflectanceSrcToSink OBJECT-TYPE
SYNTAX Integer32

```

```
UNITS          "0.1 dB"
MAX-ACCESS    read-only
STATUS        current
DESCRIPTION
    "Optical reflectance is defined to be the ratio of the
     reflected optical power present at a point, to the optical
     power incident to that point. Control of reflections is
     discussed extensively in ITU-T Rec. G.957."
 ::= { optIfOChSrcSinkXcvrEntry 4}

optIfOChSrcSinkMaximumDifferentialGroupDelay  OBJECT-TYPE
SYNTAX  Integer32
UNITS   "ps"
MAX-ACCESS  read-only
STATUS    current
DESCRIPTION
    "Differential group delay (DGD) is the time difference
     between the fractions of a pulse that are transmitted in
     the two principal states of polarization of an optical
     signal. For distances greater than several kilometres,
     and assuming random (strong) polarization mode coupling,
     DGD in a fibre can be statistically modelled as having a
     Maxwellian distribution."
 ::= { optIfOChSrcSinkXcvrEntry 5}

optIfOChSrcSinkMaximumPolarisationDependentLoss  OBJECT-TYPE
SYNTAX  Integer32
UNITS   "0.1 dB"
MAX-ACCESS  read-only
STATUS    current
DESCRIPTION
    "The polarisation dependent loss (PDL) is the difference
     (in dB) between the maximum and minimum values of the
     channel insertion loss (or gain) of the black-link from
     point SS to RS due to a variation of the state of
     polarization (SOP) over all SOPs."
 ::= { optIfOChSrcSinkXcvrEntry 6}

optIfOChSrcSinkMaximumInterChannelCrosstalk  OBJECT-TYPE
SYNTAX  Integer32
UNITS   "0.1 dB"
MAX-ACCESS  read-only
STATUS    current
DESCRIPTION
    "Inter-channel crosstalk is defined as the ratio of total
     power in all of the disturbing channels to that in the
     wanted channel, where the wanted and disturbing channels
     are at different wavelengths. The parameter specify the
```

```

        isolation of a link conforming to the 'black-link' approach
        such that under the worst-case operating conditions the
        inter-channel crosstalk at any reference point RS is less
        than the maximum inter-channel crosstalk value."
 ::= { optIfOChSrcSinkXcvrEntry 7}

optIfOChSrcSinkInterFerometricCrosstalk OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.1 dB"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "This parameter places a requirement on the isolation of a
        link conforming to the 'black-link' approach such that
        under the worst case operating conditions the
        interferometric crosstalk at any reference point RS is less
        than the maximum interferometric crosstalk value.."
 ::= { optIfOChSrcSinkXcvrEntry 8}

optIfOChSrcSinkOpticalPathOSNRPenalty OBJECT-TYPE
    SYNTAX Integer32
    UNITS "0.1 dB"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The optical path OSNR penalty is defined as the difference
        between the Lowest OSNR at Rs and Lowest OSNR at Ss that
        meets the BER requirement."
 ::= { optIfOChSrcSinkXcvrEntry 9}

-- Parameters at Sink (Rs)
-- optIfOChSinkXcvrTable
optIfOChSinkXcvrTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOChSinkXcvrEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A table of OCh Sink (Rs) configuration parameters."
 ::= { optIfOChXcvrGroup 4 }

optIfOChSinkXcvrEntry OBJECT-TYPE
    SYNTAX OptIfOChSinkXcvrEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual row that contains the Sink (Rs) configuration
        parameters for a given interface."
    INDEX { ifIndex }

```

```

 ::= { optIfOChSinkXcvrTable 1 }

OptIfOChSinkXcvrEntry ::=
    SEQUENCE {
        optIfOChSinkMinimumMeanIntputPower      Integer32,
        optIfOChSinkMaximumMeanIntputPower      Integer32,
        optIfOChSinkMinimumOSNR                 Integer32,
        optIfOChSinkOSNRTolerance               Integer32
    }

optIfOChSinkMinimumMeanIntputPower OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.01 dbm"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        " The minimum values of the average received power (in dbm
          at point the Sink (Rs))."
    ::= { optIfOChSinkXcvrEntry 1}

optIfOChSinkMaximumMeanIntputPower OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.01 dbm"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The maximum values of the average received power (in dbm)
          at point the Sink (Rs)."
```

```

    ::= { optIfOChSinkXcvrEntry 2}

optIfOChSinkMinimumOSNR OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.1 dB"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The minimum optical signal-to-noise ratio (OSNR) is the
          minimum value of the ratio of the signal power in the wanted
          channel to the highest noise power density in the range of
          the central frequency plus and minus the maximum spectral
          excursion."
```

```

    ::= { optIfOChSinkXcvrEntry 3}

optIfOChSinkOSNRTolerance OBJECT-TYPE
    SYNTAX      Integer32
    UNITS       "0.1 dB"
    MAX-ACCESS  read-only
    STATUS      current

```

DESCRIPTION

"The receiver OSNR tolerance is defined as the minimum value of OSNR at point Sink (Rs) that can be tolerated while maintaining the maximum BER of the application. Sink (Rs)."

::= { optIfOChSinkXcvrEntry 4 }

-- The OptIfOChSinkCurrentExtEntry table is an extension to the
 -- optIfOChSinkCurrentExtEntry
 -- following optional parameters for current status
 -- OptIfOChSinkCurrentExtEntry

optIfOChSinkCurrentExtTable OBJECT-TYPE

SYNTAX SEQUENCE OF OptIfOChSinkCurrentExtEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A table of OCh sink extension to the performance monitoring information for the current 15-minute interval."

::= { optIfOTNPMGroup 1 }

optIfOChSinkCurrentExtEntry OBJECT-TYPE

SYNTAX OptIfOChSinkCurrentExtEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A conceptual row that contains OCh sink performance monitoring information for an interface for the current 15-minute interval."

AUGMENTS { optIfOChSinkCurrentEntry }

::= { optIfOChSinkCurrentExtTable 1 }

OptIfOChSinkCurrentExtEntry ::=

SEQUENCE {	
optIfOChSinkCurrentChromaticDispersion	Integer32,
optIfOChSinkCurrentOSNR	Integer32,
optIfOChSinkCurrentQ	Integer32
}	

optIfOChSinkCurrentChromaticDispersion OBJECT-TYPE

SYNTAX Integer32

UNITS "ps/nm"

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Residual Chromatic Dispersion measured at Rx Transceiver port."

```

 ::= { optIfOChSinkCurrentExtEntry 1}

optIfOChSinkCurrentOSNR OBJECT-TYPE
    SYNTAX      Integer32
    UNITS        "0.1 dB"
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "Current Optical Signal to Noise Ratio (OSNR) estimated at
         Rx Transceiver port ."
    ::= { optIfOChSinkCurrentExtEntry 2}

optIfOChSinkCurrentQ OBJECT-TYPE
    SYNTAX      Integer32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "'Q' factor estimated at Rx Transceiver port."
    ::= { optIfOChSinkCurrentExtEntry 3}

-- Performance Monitoring
-- OTN PM Config Table
--

optIfOTNPMConfigTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfOTNPMConfigEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A table of performance monitoring configuration for the type
         'optIfOTNPMConfigLayer' layer."
    ::= { optIfOTNPMGroup 2 }

optIfOTNPMConfigEntry OBJECT-TYPE
    SYNTAX      OptIfOTNPMConfigEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A conceptual entry in the performance monitoring
         configuration for the type 'optIfOTNPMConfigLayer'
         layer."
    INDEX       { ifIndex, optIfOTNPMConfigType, optIfOTNPMConfigLayer,
                  optIfOTNPMConfigTCMLLevel }
    ::= { optIfOTNPMConfigTable 1 }

OptIfOTNPMConfigEntry ::=
    SEQUENCE {

```



```

    optIfOTNPMConfigType          OptIfOTNType,
    optIfOTNPMConfigLayer         OptIfOTNLayer,
    optIfOTNPMConfigTCMLLevel     Unsigned32,
    optIfOTNPMESRInterval         Unsigned32,
    optIfOTNPMSESRInterval        Unsigned32,
    optIfOTNPMValidIntervals      Unsigned32,
    optIfOTNPM15MinBip8Threshold  Unsigned32,
    optIfOTNPM15MinESsThreshold   Unsigned32,
    optIfOTNPM15MinSESSsThreshold Unsigned32,
    optIfOTNPM15MinUASsThreshold  Unsigned32,
    optIfOTNPM15MinBBEsThreshold  Unsigned32,
    optIfOTNPM24HourBip8Threshold Unsigned32,
    optIfOTNPM24HourESsThreshold  Unsigned32,
    optIfOTNPM24HourSESSsThreshold Unsigned32,
    optIfOTNPM24HourUASsThreshold Unsigned32,
    optIfOTNPM24HourBBEsThreshold Unsigned32
}

optIfOTNPMConfigType OBJECT-TYPE
    SYNTAX  OptIfOTNType
    MAX-ACCESS not-accessible
    STATUS  current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for the Near End or Far End performance data.
        1 - Near End
        2 - Far End "
    ::= { optIfOTNPMConfigEntry 1}

optIfOTNPMConfigLayer OBJECT-TYPE
    SYNTAX  OptIfOTNLayer
    MAX-ACCESS not-accessible
    STATUS  current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for OTUk, ODUk, TCMn performance data.
        1 - OTUk
        2 - ODUk
        3 - TCM
        The ODUk/TCM sublayer PM is not related to the black link PM
        management, but since this is a common PM model for the
        ODU/TCM layer, we may include it here."
    ::= { optIfOTNPMConfigEntry 2}

optIfOTNPMConfigTCMLLevel OBJECT-TYPE
    SYNTAX  Unsigned32 (0..6)
    MAX-ACCESS not-accessible
    STATUS  current

```

DESCRIPTION

"This parameter indicates the TCM level (1-6)
if the PM is of the type TCM. This will be 0 for OTUK/ODUK."

::= { optIfOTNPMConfigEntry 3 }

optIfOTNPMESRInterval OBJECT-TYPE

SYNTAX Unsigned32 (1..96)

UNITS "seconds"

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"This parameter indicates the measurement interval
for error seconds ratio."

::= { optIfOTNPMConfigEntry 4 }

optIfOTNPMSESInterval OBJECT-TYPE

SYNTAX Unsigned32

UNITS "seconds"

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"This parameter indicates the measurement interval
for severely error seconds ratio."

::= { optIfOTNPMConfigEntry 5 }

optIfOTNPMValidIntervals OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The number of contiguous 15 minute intervals for which valid
PM data is available for the particular interface."

::= { optIfOTNPMConfigEntry 6 }

optIfOTNPM15MinBip8Threshold OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"The number of Bip8 encountered by the interface within any
given 15 minutes performance data collection period, which
causes the SNMP agent to send optIf15MinThreshBip8TCA. One
notification will be sent per interval per interface. A
value of '0' will disable the notification."

::= { optIfOTNPMConfigEntry 7 }

optIfOTNPM15MinESsThreshold OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of ES encountered by the interface within any
 given 15 minutes performance data collection period, which
 causes the SNMP agent to send optIf15MinThreshEstCA. One
 notification will be sent per interval per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 8 }

optIfOTNPM15MinSESsThreshold OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of SES encountered by the interface within any
 given 15 minutes performance data collection period, which
 causes the SNMP agent to send optIf15MinThreshSESTCA. One
 notification will be sent per interval per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 9 }

optIfOTNPM15MinUASsThreshold OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of UAS encountered by the interface within any
 given 15 minutes performance data collection period, which
 causes the SNMP agent to send optIf15MinThreshUASTCA. One
 notification will be sent per interval per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 10 }

optIfOTNPM15MinBBEsThreshold OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of UAS encountered by the interface within any
 given 15 minutes performance data collection period, which
 causes the SNMP agent to send optIf15MinThreshBBETCA. One
 notification will be sent per interval per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 11 }

optIfOTNPM24HourBip8Threshold OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of Bip8 encountered by the interface within any
 given 24 Hour performance data collection period, which
 causes the SNMP agent to send optIf15MinThreshBip8TCA. One
 notification will be sent per interval per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 12 }

optIfOTNPM24HourESsThreshold OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of ES encountered by the interface within any
 given 24 hour performance data collection period, which
 causes the SNMP agent to send optIf24HourThreshEsTCA. One
 notification will be sent per 24 hour per interface. A
 value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 13 }

optIfOTNPM24HourSESsThreshold OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of SES encountered by the interface within any
 given 24 hour performance data collection period, which
 causes the SNMP agent to send optIf24HourThreshSESsTCA. One
 notification will be sent per 24 hour per interface. A value
 of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 14 }

optIfOTNPM24HourUASsThreshold OBJECT-TYPE

SYNTAX Unsigned32
MAX-ACCESS read-write
STATUS current
DESCRIPTION
 "The number of SES encountered by the interface within any
 given 24 hour performance data collection period, which
 causes the SNMP agent to send optIf24HourThreshUASsTCA. One
 notification will be sent per 24 hour per interface. A value
 of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 15 }

optIfOTNPM24HourBBEsThreshold OBJECT-TYPE

```

SYNTAX      Unsigned32
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The number of BBE encountered by the interface within any
    given 24 hour performance data collection period, which
    causes the SNMP agent to send optIf24HourThreshBBEsTCA. One
    notification will be sent per 24 hour per interface. A
    value of '0' will disable the notification."
 ::= { optIfOTNPMConfigEntry 16 }

--
-- PM Current Entry at either the OTU/ODUk/TCM
--
optIfOTNPMCurrentTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfOTNPMCurrentEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A table for the Performance monitoring Current Table."
    ::= { optIfOTNPMGroup 3}

optIfOTNPMCurrentEntry OBJECT-TYPE
    SYNTAX      OptIfOTNPMCurrentEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring Current table for the type
        'optIfOTNPMCurrentLayer' layer."
    INDEX { ifIndex, optIfOTNPMCurrentType ,
            optIfOTNPMCurrentLayer, optIfOTNPMCurrentTCMLLevel  }
    ::= { optIfOTNPMCurrentTable 1 }

OptIfOTNPMCurrentEntry ::=
    SEQUENCE {
        optIfOTNPMCurrentType                OptIfOTNType,
        optIfOTNPMCurrentLayer                OptIfOTNLayer,
        optIfOTNPMCurrentTCMLLevel            Unsigned32,
        optIfOTNPMCurrentSuspectedFlag        TruthValue,
        optIfOTNPMCurrentBip8                 Unsigned32,
        optIfOTNPMCurrentESS                  Unsigned32,
        optIfOTNPMCurrentSESS                 Unsigned32,
        optIfOTNPMCurrentUASs                 Unsigned32,
        optIfOTNPMCurrentBBEs                 Unsigned32,
        optIfOTNPMCurrentESR                  Unsigned32,
        optIfOTNPMCurrentSESR                 Unsigned32,
        optIfOTNPMCurrentBBER                 Unsigned32,

```

```

        optIfOTNPMCurrentBIP8                Unsigned32,
        optIfOTNPMCurrentElapsedTime          Unsigned32
    }

optIfOTNPMCurrentType OBJECT-TYPE
    SYNTAX      OptIfOTNType
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for the Near End or Far End performance data.
        1 - Near End
        2 - Far End "
    ::= { optIfOTNPMCurrentEntry 1}

optIfOTNPMCurrentLayer OBJECT-TYPE
    SYNTAX      OptIfOTNLayer
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for OTUk, ODUk, TCMn performance data.
        1 - OTUk (Och which is used for the black link)
        2 - ODUk
        3 - TCM
        The ODUk/TCM sublayer PM is not related to the black link PM
        management, but since this is a common PM model for the
        ODU/TCM layer, we may include it here."
    ::= { optIfOTNPMCurrentEntry 2}

optIfOTNPMCurrentTCMLevel OBJECT-TYPE
    SYNTAX      Unsigned32 (0..6)
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This parameter indicates the TCM level (1-6)
        if the PM is of the type TCM. This will be 0 for OTUK/ODUK."
    ::= { optIfOTNPMCurrentEntry 3}

optIfOTNPMCurrentSuspectedFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "If true, the data in this entry may be unreliable."
    ::= { optIfOTNPMCurrentEntry 4}

```

```
optIfOTNPMCurrentBip8    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "Number of Failures occurred in an observation period."
    ::= { optIfOTNPMCurrentEntry 5}

optIfOTNPMCurrentESSs    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "This is the number of seconds in which one or more bits are
         in error or during which Loss of Signal (LOS) or Alarm
         Indication Signal (AIS) is detected."
    ::= { optIfOTNPMCurrentEntry 6}

optIfOTNPMCurrentSESSs    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The number of seconds which have a severe error.
         This is the number of seconds in which the bit-error ratio =
         1x10Eminus3 or during which Loss of Signal (LOS) or Alarm
         Indication Signal (AIS) is detected."
    ::= { optIfOTNPMCurrentEntry 7}

optIfOTNPMCurrentUASSs    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "It is the number of unavailable seconds.
         A period of unavailable time begins at the onset of ten
         consecutive SES events. These ten seconds are considered
         to be part of unavailable time. A new period of available
         time begins at the onset of ten consecutive non-SES events.
         These ten seconds are considered to be part of available
         time."
    ::= { optIfOTNPMCurrentEntry 8}

optIfOTNPMCurrentBBEs    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
```

```
        "An errored block not occurring as part of an SES."
 ::= { optIfOTNPMCurrentEntry 9}

optIfOTNPMCurrentESR OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The ratio of ES in available time to total seconds in
         available time during a fixed measurement interval."
 ::= { optIfOTNPMCurrentEntry 10}

optIfOTNPMCurrentSESR OBJECT-TYPE
    SYNTAX      Unsigned32
    UNITS       ".001"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The ratio of SES in available time to total seconds in
         available time during a fixed measurement interval."
 ::= { optIfOTNPMCurrentEntry 11}

optIfOTNPMCurrentBBER OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The ratio of BER in available time to total seconds in
         available time during a fixed measurement interval."
 ::= { optIfOTNPMCurrentEntry 12 }

optIfOTNPMCurrentBIP8 OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The BIP8 count for this period."
 ::= { optIfOTNPMCurrentEntry 13 }

optIfOTNPMCurrentElapsedTime OBJECT-TYPE
    SYNTAX      Unsigned32
    UNITS       "seconds"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Time elapsed for this 15 minute interval"
```



```

 ::= { optIfOTNPMCurrentEntry 14 }

--
-- OTN PM Interval Table
-- Upto 96 15-minute intervals
--
optIfOTNPMIntervalTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMIntervalEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring Interval Table."
    ::= { optIfOTNPMGroup 4 }

optIfOTNPMIntervalEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMIntervalEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring Interval table for the type
        'optIfOTNPMIntervalLayer' layer."
    INDEX { ifIndex, optIfOTNPMIntervalType,
            optIfOTNPMIntervalLayer, optIfOTNPMIntervalTCMLLevel,
            optIfOTNPMIntervalNumber }
    ::= { optIfOTNPMIntervalTable 1 }

OptIfOTNPMIntervalEntry ::=
    SEQUENCE {
        optIfOTNPMIntervalType OptIfOTNType,
        optIfOTNPMIntervalLayer OptIfOTNLayer,
        optIfOTNPMIntervalTCMLLevel Unsigned32,
        optIfOTNPMIntervalNumber Unsigned32,
        optIfOTNPMIntervalSuspectedFlag TruthValue,
        optIfOTNPMIntervalBip8 Unsigned32,
        optIfOTNPMIntervalESS Unsigned32,
        optIfOTNPMIntervalSESS Unsigned32,
        optIfOTNPMIntervalUASS Unsigned32,
        optIfOTNPMIntervalBBES Unsigned32,
        optIfOTNPMIntervalESR Unsigned32,
        optIfOTNPMIntervalSESR Unsigned32,
        optIfOTNPMIntervalBBER Unsigned32,
        optIfOTNPMIntervalBIP8 Unsigned32,
        optIfOTNPMIntervalTimeStamp DateAndTime
    }

optIfOTNPMIntervalType OBJECT-TYPE
    SYNTAX OptIfOTNType

```

```
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "This parameter indicates the parameters for the table are
    for the Near End or Far End performance data.
    1 - Near End
    2 - Far End "
 ::= { optIfOTNPMIntervalEntry 1}

optIfOTNPMIntervalLayer OBJECT-TYPE
SYNTAX OptIfOTNLayer
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "This parameter indicates the parameters for the table are
    for OTUk, ODUk, TCMn performance data.
    1 - OTUk
    2 - ODUk
    3 - TCM
    The ODUk/TCM sublayer PM is not related to the black link PM
    management, but since this is a common PM model for the
    ODU/TCM layer, we may include it here."
 ::= { optIfOTNPMIntervalEntry 2}

optIfOTNPMIntervalTCMLevel OBJECT-TYPE
SYNTAX Unsigned32 (0..6)
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "This parameter indicates the TCM level (1-6)
    if the PM is of the type TCM. This will be 0 for OTUK/ODUK."
 ::= { optIfOTNPMIntervalEntry 3}

optIfOTNPMIntervalNumber OBJECT-TYPE
SYNTAX Unsigned32 (1..96)
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "A number between 1 and 96, where 1 is the most
    recently completed 15 minute interval and 96 is
    the 15 minutes interval completed 23 hours and 45
    minutes prior to interval 1."
 ::= { optIfOTNPMIntervalEntry 4}

optIfOTNPMIntervalSuspectedFlag OBJECT-TYPE
SYNTAX TruthValue
MAX-ACCESS read-only
STATUS current
```

DESCRIPTION

"If true, the data in this entry may be unreliable."

::= { optIfOTNPMIntervalEntry 5 }

optIfOTNPMIntervalBip8 OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Number of Failures occurred in an observation period."

::= { optIfOTNPMIntervalEntry 6 }

optIfOTNPMIntervalESS OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"It is a one-second period in which one or more bits are in error or during which Loss of Signal (LOS) or Alarm Indication Signal (AIS) is detected."

::= { optIfOTNPMIntervalEntry 7 }

optIfOTNPMIntervalSESS OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The number of seconds which have a severe error.

It is a one-second period which has a bit-error ratio =

1×10^E minus 3 or during which Loss of Signal (LOS) or Alarm

Indication Signal (AIS) is detected."

::= { optIfOTNPMIntervalEntry 8 }

optIfOTNPMIntervalUASS OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"It is the number of unavailable seconds in this 15 minute interval. A period of unavailable time begins at the onset of ten consecutive SES events. These ten seconds are considered to be part of unavailable time. A new period of available time begins at the onset of ten consecutive non-SES events. These ten seconds are considered to be part of available time."

::= { optIfOTNPMIntervalEntry 9 }

```
optIfOTNPMIntervalBBEs    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "An errored block not occurring as part of an SES."
    ::= { optIfOTNPMIntervalEntry 10}

optIfOTNPMIntervalESR    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The ratio of ES in available time to total seconds in
         available time during a fixed measurement interval."
    ::= { optIfOTNPMIntervalEntry 11}

optIfOTNPMIntervalSESR    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The ratio of SES in available time to total seconds in
         available time during a fixed measurement interval."
    ::= { optIfOTNPMIntervalEntry 12}

optIfOTNPMIntervalBBER    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The ratio of BBE in available time to total seconds in
         available time during a fixed measurement interval."
    ::= { optIfOTNPMIntervalEntry 13}

optIfOTNPMIntervalBIP8    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "BIP8 for this period."
    ::= { optIfOTNPMIntervalEntry 14}

optIfOTNPMIntervalTimeStamp OBJECT-TYPE
    SYNTAX  DateAndTime
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
```

```

        "Time stamp of this interval."
 ::= { optIfOTNPMIntervalEntry 15}

--
-- PM Current Day Entry
--
optIfOTNPMCurrentDayTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMCurrentDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring Current Day Table."
    ::= { optIfOTNPMGroup 5 }

optIfOTNPMCurrentDayEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMCurrentDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring Current day table for the type
        'optIfOTNPMCurrentDayLayer' layer."
    INDEX { ifIndex, optIfOTNPMCurrentDayType,
            optIfOTNPMCurrentDayLayer,
            optIfOTNPMCurrentDayTCMLLevel }
    ::= { optIfOTNPMCurrentDayTable 1 }

OptIfOTNPMCurrentDayEntry ::=
    SEQUENCE {
        optIfOTNPMCurrentDayType                OptIfOTNType,
        optIfOTNPMCurrentDayLayer                OptIfOTNLayer,
        optIfOTNPMCurrentDayTCMLLevel            Unsigned32,
        optIfOTNPMCurrentDaySuspectedFlag        TruthValue,
        optIfOTNPMCurrentDayBip8                 Unsigned32,
        optIfOTNPMCurrentDayESS                  Unsigned32,
        optIfOTNPMCurrentDaySESS                 Unsigned32,
        optIfOTNPMCurrentDayUASs                 Unsigned32,
        optIfOTNPMCurrentDayBBEs                 Unsigned32,
        optIfOTNPMCurrentDayESR                  Unsigned32,
        optIfOTNPMCurrentDaySESR                 Unsigned32,
        optIfOTNPMCurrentDayBBER                 Unsigned32,
        optIfOTNPMCurrentDayBIP8                 Unsigned32,
        optIfOTNPMCurrentDayElapsedTime          Unsigned32
    }

optIfOTNPMCurrentDayType OBJECT-TYPE
    SYNTAX OptIfOTNType

```

```
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "This parameter indicates the parameters for the table are
    for the Near End or Far End performance data.
    1 - Near End
    2 - Far End "
 ::= { optIfOTNPMCurrentDayEntry 1}

optIfOTNPMCurrentDayLayer OBJECT-TYPE
    SYNTAX OptIfOTNLayer
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for OTUk, ODUk, TCMn performance data.
        1 - OTUk
        2 - ODUk
        3 - TCM
        The ODUk/TCM sublayer PM is not related to the black link PM
        management, but since this is a common PM model for the
        ODU/TCM layer, we may include it here."
 ::= { optIfOTNPMCurrentDayEntry 2}

optIfOTNPMCurrentDayTCMLevel OBJECT-TYPE
    SYNTAX Unsigned32 (0..6)
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "This parameter indicates the TCM level (1-6)
        if the PM is of the type TCM. This will be 0 for OTUK/ODUK."
 ::= { optIfOTNPMCurrentDayEntry 3}

optIfOTNPMCurrentDaySuspectedFlag OBJECT-TYPE
    SYNTAX TruthValue
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "If true, the data in this entry may be unreliable."
 ::= { optIfOTNPMCurrentDayEntry 4}

optIfOTNPMCurrentDayBip8 OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "Number of Failures occurred in an observation period."
```

```
::= { optIfOTNPMCurrentDayEntry 5}
```

```
optIfOTNPMCurrentDayESS OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of seconds which have an error.
        It is a one-second period in which one or more bits are in
        error or during which Loss of Signal (LOS) or Alarm
        Indication Signal (AIS) is detected."
    ::= { optIfOTNPMCurrentDayEntry 6}
```

```
optIfOTNPMCurrentDaySESS OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of seconds which have a severe error.
        It is a one-second period which has a bit-error ratio =
        1x10Eminus3 or during which Loss of Signal (LOS) or Alarm
        Indication Signal (AIS) is detected."
    ::= { optIfOTNPMCurrentDayEntry 7}
```

```
optIfOTNPMCurrentDayUASS OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "It is the number of unavailable seconds in the cunrrent day.
        A period of unavailable time begins at the onset of ten
        consecutive SES events. These ten seconds are considered to
        be part of unavailable time. A new period of available time
        begins at the onset of ten consecutive non-SES events. These
        ten seconds are considered to be part of available time."
    ::= { optIfOTNPMCurrentDayEntry 8}
```

```
optIfOTNPMCurrentDayBBES OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "An errored block not occurring as part of an SES."
    ::= { optIfOTNPMCurrentDayEntry 9}
```

```
optIfOTNPMCurrentDayESR OBJECT-TYPE
```

```
SYNTAX  Unsigned32
UNITS    ".001"
MAX-ACCESS  read-only
STATUS    current
DESCRIPTION
    "The ratio of ES in available time to total seconds in
    available time during a fixed measurement interval."
 ::= { optIfOTNPMCurrentDayEntry 10}

optIfOTNPMCurrentDaySESR  OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS    current
    DESCRIPTION
        "The ratio of SES in available time to total seconds in
        available time during a fixed measurement interval."
    ::= { optIfOTNPMCurrentDayEntry 11}

optIfOTNPMCurrentDayBBER  OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS    current
    DESCRIPTION
        "The ratio of BBE in available time to total seconds in
        available time during a fixed measurement interval."
    ::= { optIfOTNPMCurrentDayEntry 12}

optIfOTNPMCurrentDayBIP8  OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS    current
    DESCRIPTION
        "BIP8 for this period."
    ::= { optIfOTNPMCurrentDayEntry 13}

optIfOTNPMCurrentDayElapsedTime  OBJECT-TYPE
    SYNTAX      Unsigned32
    UNITS        "seconds"
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "Time elapsed for current day"
    ::= { optIfOTNPMCurrentDayEntry 14 }

--
-- PM Prev Day Entry
--
optIfOTNMPPrevDayTable OBJECT-TYPE
```



```

SYNTAX SEQUENCE OF OptIfOTNPMPPrevDayEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "A Performance monitoring Previous Day Table."
 ::= { optIfOTNPMPGroup 6 }

```

```

optIfOTNPMPPrevDayEntry OBJECT-TYPE
SYNTAX      OptIfOTNPMPPrevDayEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "A conceptual entry in the Near end or Far End performance
     monitoring previous day table for the type
     'optIfOTNPMPPrevDayLayer' layer."
INDEX { ifIndex, optIfOTNPMPPrevDayType      ,
         optIfOTNPMPPrevDayLayer, optIfOTNPMPPrevDayTCMLLevel  }
 ::= { optIfOTNPMPPrevDayTable 1 }

```

```

OptIfOTNPMPPrevDayEntry ::=
SEQUENCE {
    optIfOTNPMPPrevDayType      OptIfOTNType,
    optIfOTNPMPPrevDayLayer     OptIfOTNLayer,
    optIfOTNPMPPrevDayTCMLLevel Unsigned32,
    optIfOTNPMPPrevDaySuspectedFlag TruthValue,
    optIfOTNPMPPrevDayBip8      Unsigned32,
    optIfOTNPMPPrevDayESs       Unsigned32,
    optIfOTNPMPPrevDaySESS      Unsigned32,
    optIfOTNPMPPrevDayUASS      Unsigned32,
    optIfOTNPMPPrevDayBBES      Unsigned32,
    optIfOTNPMPPrevDayESR       Unsigned32,
    optIfOTNPMPPrevDaySESR      Unsigned32,
    optIfOTNPMPPrevDayBBER      Unsigned32,
    optIfOTNPMPPrevDayBIP8      Unsigned32,
    optIfOTNPMPPrevDayTimeStamp DateAndTime
}

```

```

optIfOTNPMPPrevDayType      OBJECT-TYPE
SYNTAX OptIfOTNType
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
    "This parameter indicates the parameters for the table are
     for the Near End or Far End performance data.
     1 - Near End
     2 - Far End "
 ::= { optIfOTNPMPPrevDayEntry 1}

```

```
optIfOTNPMPPrevDayLayer OBJECT-TYPE
    SYNTAX  OptIfOTNLayer
    MAX-ACCESS  not-accessible
    STATUS  current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for OTUk, ODUk, TCMn performance data.
        1 - OTUk
        2 - ODUk
        3 - TCM
        The ODUk/TCM sublayer PM is not related to the black link PM
        ODU/TCM management, but since this is a common PM model for
        the layer, we may include it here."
    ::= { optIfOTNPMPPrevDayEntry 2}

optIfOTNPMPPrevDayTCMLevel OBJECT-TYPE
    SYNTAX  Unsigned32 (0..6)
    MAX-ACCESS  not-accessible
    STATUS  current
    DESCRIPTION
        "This parameter indicates the TCM level (1-6)
        if the PM is of the type TCM."
    ::= { optIfOTNPMPPrevDayEntry 3}

optIfOTNPMPPrevDaySuspectedFlag OBJECT-TYPE
    SYNTAX  TruthValue
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "If true, the data in this entry may be unreliable."
    ::= { optIfOTNPMPPrevDayEntry 4}

optIfOTNPMPPrevDayBip8 OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "Number of pre FEC failures occurred in an observation
        period."
    ::= { optIfOTNPMPPrevDayEntry 5}

optIfOTNPMPPrevDayESs OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The number of seconds which have an error."
```

It is a one-second period in which one or more bits are in error or during which Loss of Signal (LOS) or Alarm Indication Signal(AIS) is detected."

::= { optIfOTNPMPPrevDayEntry 6}

optIfOTNPMPPrevDaySESS OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The number of seconds which have a severe error. A severely errored second, is a one-second period which has a bit-error ratio = $1 \times 10^{\text{Eminus}3}$ or during which Loss of Signal (LOS) or Alarm Indication Signal (AIS) is detected."

::= { optIfOTNPMPPrevDayEntry 7}

optIfOTNPMPPrevDayUASS OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"It is the number of unavailable seconds in the previous day. A period of unavailable time begins at the onset of ten consecutive SES events. These ten seconds are considered to be part of unavailable time. A new period of available time begins at the onset of ten consecutive non-SES events. These ten seconds are considered to be part of available time."

::= { optIfOTNPMPPrevDayEntry 8}

optIfOTNPMPPrevDayBBES OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"An errored block not occurring as part of an SES."

::= { optIfOTNPMPPrevDayEntry 9}

optIfOTNPMPPrevDayESR OBJECT-TYPE

SYNTAX Unsigned32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The ratio of ES in available time to total seconds in available time during a fixed measurement interval."

::= { optIfOTNPMPPrevDayEntry 10}

optIfOTNPMPPrevDaySESR OBJECT-TYPE

```
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "The ratio of SES in available time to total seconds in
    available time during a fixed measurement interval."
 ::= { optIfOTNPMPrevDayEntry 11}

optIfOTNPMPrevDayBBER OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The ratio of BBE in available time to total seconds in
        available time during a fixed measurement interval."
    ::= { optIfOTNPMPrevDayEntry 12}

optIfOTNPMPrevDayBIP8 OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "BIP8 for this period."
    ::= { optIfOTNPMPrevDayEntry 13}

optIfOTNPMPrevDayTimeStamp OBJECT-TYPE
    SYNTAX DateAndTime
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "Time stamp of this interval."
    ::= { optIfOTNPMPrevDayEntry 14}

--
-- OTN FEC PM Config Table
--
optIfOTNPMFECConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMFECConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A table of performance monitoring FEC configuration."
    ::= { optIfOTNPMGroup 7 }

optIfOTNPMFECConfigEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMFECConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
```

DESCRIPTION

"A conceptual entry in the performance monitoring FEC configuration layer."

INDEX { ifIndex, optIfOTNPMFECConfigType }
 ::= { optIfOTNPMFECConfigTable 1 }

OptIfOTNPMFECConfigEntry ::=

```
SEQUENCE {
    optIfOTNPMFECConfigType          OptIfOTNType,
    optIfOTNPMFECValidIntervals      Unsigned32,
    optIfOTNPM15MinFECUnCorrectedWordsThreshold Unsigned32,
    optIfOTNPM15MinPreFECBERMantissaThreshold Unsigned32,
    optIfOTNPM15MinPreFECBERExponentThreshold Unsigned32
}
```

optIfOTNPMFECConfigType OBJECT-TYPE

SYNTAX OptIfOTNType
 MAX-ACCESS not-accessible
 STATUS current
 DESCRIPTION

"This parameter indicates the parameters for the table are for the Near End or Far End performance data.

1 - Near End
 2 - Far End "

::= { optIfOTNPMFECConfigEntry 1}

optIfOTNPMFECValidIntervals OBJECT-TYPE

SYNTAX Unsigned32
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The number of contiguous 15 minute intervals for which valid FEC PM data is available for the particular interface."

::= {optIfOTNPMFECConfigEntry 2}

optIfOTNPM15MinFECUnCorrectedWordsThreshold OBJECT-TYPE

SYNTAX Unsigned32
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The number of Uncorrected words encountered by the interface within any given 15 minutes performance data collection period, which causes the SNMP agent to send optIf15MinThreshFECUnCorrectedWordsTCA. One notification will be sent per interval per interface. A value of '0' will disable the notification."

::= {optIfOTNPMFECConfigEntry 3}

```

optIfOTNPM15MinPreFECBERMantissaThreshold OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "BER (mantissa) by the interface within any
        given 15 minutes performance data collection period, which
        causes the SNMP agent to send optIf15MinThreshPreFECBERTCA.
        One notification will be sent per interval per interface.
        A value of '0' will disable the notification."
    ::= {optIfOTNPMFECConfigEntry 4}

optIfOTNPM15MinPreFECBERExponentThreshold OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The Pre FEC BER (exponent) by the interface within any
        given 15 minutes performance data collection period, which
        causes the SNMP agent to send optIf15MinThreshPreFECBERTCA.
        One notification will be sent per interval per interface. A
        value of '0' will disable the notification."
    ::= {optIfOTNPMFECConfigEntry 5}

--
-- FEC PM Table
--
optIfOTNPMFECCurrentTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMFECCurrentEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring FEC Current Table."
    ::= { optIfOTNPMGroup 8 }

optIfOTNPMFECCurrentEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMFECCurrentEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring FEC current table."
    INDEX { ifIndex, optIfOTNPMFECCurrentType }
    ::= { optIfOTNPMFECCurrentTable 1 }

OptIfOTNPMFECCurrentEntry ::=
    SEQUENCE {

```

optIfOTNPMFECCurrentType	OptIfOTNType,
optIfOTNPMFECCurrentSuspectedFlag	TruthValue,
optIfOTNPMCurrentFECCorrectedErr	Counter64,
optIfOTNPMCurrentFECUncorrectedWords	Counter64,
optIfOTNPMCurrentFECBERMantissa	Unsigned32,
optIfOTNPMCurrentFECBERExponent	Unsigned32,
optIfOTNPMCurrentFECMinBERMantissa	Unsigned32,
optIfOTNPMCurrentFECMinBERExponent	Unsigned32,
optIfOTNPMCurrentFECMaxBERMantissa	Unsigned32,
optIfOTNPMCurrentFECMaxBERExponent	Unsigned32,
optIfOTNPMCurrentFECAvgBERMantissa	Unsigned32,
optIfOTNPMCurrentFECAvgBERExponent	Unsigned32,
optIfOTNPMCurrentFECElapsedTime	Unsigned32

}

optIfOTNPMFECCurrentType OBJECT-TYPE
 SYNTAX OptIfOTNType
 MAX-ACCESS not-accessible
 STATUS current
 DESCRIPTION
 "This parameter indicates the parameters for the table are
 for the Near End or Far End performance data.
 1 - Near End
 2 - Far End "
 ::= { optIfOTNPMFECCurrentEntry 1}

optIfOTNPMFECCurrentSuspectedFlag OBJECT-TYPE
 SYNTAX TruthValue
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION
 "If true, the data in this entry may be unreliable."
 ::= { optIfOTNPMFECCurrentEntry 2}

optIfOTNPMCurrentFECCorrectedErr OBJECT-TYPE
 SYNTAX Counter64
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION
 "The number of bits corrected by the FEC are counted in the
 interval."
 ::= { optIfOTNPMFECCurrentEntry 3}

optIfOTNPMCurrentFECUncorrectedWords OBJECT-TYPE
 SYNTAX Counter64
 MAX-ACCESS read-only
 STATUS current

DESCRIPTION
 "The number of un-corrected words by the FEC are counted over
 the interval."
 ::= { optIfOTNPMFECCurrentEntry 4}

optIfOTNPMCurrentFECBERMantissa OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The number of Errored bits at receiving side before the FEC
 function counted over one second .. mantissa."
 ::= { optIfOTNPMFECCurrentEntry 5}

optIfOTNPMCurrentFECBERExponent OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The number of Errored bits at receiving side before the FEC
 function counted over one second .. exponent."
 ::= { optIfOTNPMFECCurrentEntry 6}

optIfOTNPMCurrentFECMinBERMantissa OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The minimum number of Errored bits at receiving side before
 the FEC function counted over one second .. mantissa."
 ::= { optIfOTNPMFECCurrentEntry 7}

optIfOTNPMCurrentFECMinBERExponent OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The minimum number of Errored bits at receiving side before
 the FEC function counted over one second.. exponent."
 ::= { optIfOTNPMFECCurrentEntry 8}

optIfOTNPMCurrentFECMaxBERMantissa OBJECT-TYPE
SYNTAX Unsigned32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
 "The maximum number of Errored bits at receiving side before
 the FECfunction counted over one second .. mantissa."


```
 ::= { optIfOTNPMFECCurrentEntry 9}

optIfOTNPMCurrentFECMaxBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The maximum number of Errored bits at receiving side before
        the FEC function counted over one second .. exponent."
    ::= { optIfOTNPMFECCurrentEntry 10}

optIfOTNPMCurrentFECAvgBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average number of Errored bits at receiving side before
        the FEC function counted over one second .. mantissa."
    ::= { optIfOTNPMFECCurrentEntry 11}

optIfOTNPMCurrentFECAvgBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average number of Errored bits at receiving side before
        the FEC function counted over one second .. exponent."
    ::= { optIfOTNPMFECCurrentEntry 12}

optIfOTNPMCurrentFECElapsedTime OBJECT-TYPE
    SYNTAX Unsigned32
    UNITS "seconds"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "Time elapsed for this 15 minute interval."
    ::= { optIfOTNPMFECCurrentEntry 13 }

--
-- FEC PM Interval Table
--
optIfOTNPMFECIntervalTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMFECIntervalEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring FEC Interval Table."
    ::= { optIfOTNPMGroup 9 }
```

```

optIfOTNPMFECIntervalEntry OBJECT-TYPE
    SYNTAX      OptIfOTNPMFECIntervalEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring FEC interval table."
    INDEX { ifIndex, optIfOTNPMFECIntervalType,
            optIfOTNPMFECIntervalNumber }
    ::= { optIfOTNPMFECIntervalTable 1 }

OptIfOTNPMFECIntervalEntry ::=
    SEQUENCE {
        optIfOTNPMFECIntervalType          OptIfOTNType,
        optIfOTNPMFECIntervalNumber        Unsigned32,
        optIfOTNPMFECIntervalSuspectedFlag TruthValue,
        optIfOTNPMIntervalFECCorrectedErr  Counter64,
        optIfOTNPMIntervalFECUncorrectedWords Counter64,
        optIfOTNPMIntervalMinFECBERMantissa Unsigned32,
        optIfOTNPMIntervalMinFECBERExponent Unsigned32,
        optIfOTNPMIntervalMaxFECBERMantissa Unsigned32,
        optIfOTNPMIntervalMaxFECBERExponent Unsigned32,
        optIfOTNPMIntervalAvgFECBERMantissa Unsigned32,
        optIfOTNPMIntervalAvgFECBERExponent Unsigned32,
        optIfOTNPMFECIntervalTimeStamp     DateAndTime
    }

optIfOTNPMFECIntervalType          OBJECT-TYPE
    SYNTAX      OptIfOTNType
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for the Near End or Far End performance data.
        1 - Near End
        2 - Far End "
    ::= { optIfOTNPMFECIntervalEntry 1}

optIfOTNPMFECIntervalNumber        OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A number between 1 and 96, where 1 is the most
        recently completed 15 minute interval and 96 is
        the 15 minutes interval completed 23 hours and 45
        minutes prior to interval 1."
    ::= { optIfOTNPMFECIntervalEntry 2}

```

```
optIfOTNPMFECIntervalSuspectedFlag    OBJECT-TYPE
    SYNTAX  TruthValue
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "If true, the data in this entry may be unreliable."
    ::= { optIfOTNPMFECIntervalEntry 3}

optIfOTNPMIntervalFECCorrectedErr    OBJECT-TYPE
    SYNTAX  Counter64
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The number of bits corrected by the FEC are counted in the
        interval."
    ::= { optIfOTNPMFECIntervalEntry 4}

optIfOTNPMIntervalFECUncorrectedWords    OBJECT-TYPE
    SYNTAX  Counter64
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The number of words un-corrected words by the FEC are
        counted over the interval."
    ::= { optIfOTNPMFECIntervalEntry 5}

optIfOTNPMIntervalMinFECBERMantissa    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The minimum bit error rate at receiving side before the FEC
        function counted over one second .. mantissa. This is the
        minimum Pre FEC BER in the current 24hour period."
    ::= { optIfOTNPMFECIntervalEntry 6}

optIfOTNPMIntervalMinFECBERExponent    OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The minimum bit error rate at receiving side before the FEC
        function counted over one second .. exponent. This is the
        minimum Pre FEC BER in the current 24hour period."
    ::= { optIfOTNPMFECIntervalEntry 7}

optIfOTNPMIntervalMaxFECBERMantissa    OBJECT-TYPE
    SYNTAX  Unsigned32
```

```
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "The maximum bit error rate at receiving side before the FEC
    function counted over one second .. mantissa. This is the
    maximum Pre FEC BER in the current 24hour period."
 ::= { optIfOTNPMFECIntervalEntry 8}

optIfOTNPMIntervalMaxFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The maximum bit error rate at receiving side before the FEC
        function counted over one second .. exponent. This is the
        maximum Pre FEC BER in the current 24hour period."
    ::= { optIfOTNPMFECIntervalEntry 9}

optIfOTNPMIntervalAvgFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. mantissa. This is the
        average Pre FEC BER in the current 24hour period."
    ::= { optIfOTNPMFECIntervalEntry 10}

optIfOTNPMIntervalAvgFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. exponent. This is the
        average Pre FEC BER in the current 24hour period."
    ::= { optIfOTNPMFECIntervalEntry 11}

optIfOTNPMFECIntervalTimeStamp OBJECT-TYPE
    SYNTAX DateAndTime
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "Time stamp of this interval."
    ::= { optIfOTNPMFECIntervalEntry 12 }

--
-- FEC PM Current Day day Table
```

```

--
optIfOTNPMFECCurrentDayTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMFECCurrentDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring FEC current day table."
    ::= { optIfOTNPMGroup 10 }

optIfOTNPMFECCurrentDayEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMFECCurrentDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring FEC current day table."
    INDEX { ifIndex, optIfOTNPMFECCurrentDayType }
    ::= { optIfOTNPMFECCurrentDayTable 1 }

OptIfOTNPMFECCurrentDayEntry ::=
    SEQUENCE {
        optIfOTNPMFECCurrentDayType                               OptIfOTNType,
        optIfOTNPMFECCurrentDaySuspectedFlag                     TruthValue,
        optIfOTNPMCurrentDayFECCorrectedErr                      Counter64,
        optIfOTNPMCurrentDayFECUncorrectedWords                  Counter64,
        optIfOTNPMCurrentDayMinFECBERMantissa                    Unsigned32,
        optIfOTNPMCurrentDayMinFECBERExponent                    Unsigned32,
        optIfOTNPMCurrentDayMaxFECBERMantissa                    Unsigned32,
        optIfOTNPMCurrentDayMaxFECBERExponent                    Unsigned32,
        optIfOTNPMCurrentDayAvgFECBERMantissa                    Unsigned32,
        optIfOTNPMCurrentDayAvgFECBERExponent                    Unsigned32,
        optIfOTNPMFECCurrentDayElapsedTime                       Unsigned32
    }

optIfOTNPMFECCurrentDayType OBJECT-TYPE
    SYNTAX OptIfOTNType
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for the Near End or Far End performance data.
        1 - Near End
        2 - Far End "
    ::= { optIfOTNPMFECCurrentDayEntry 1}

optIfOTNPMFECCurrentDaySuspectedFlag OBJECT-TYPE
    SYNTAX TruthValue

```

```
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "If true, the data in this entry may be unreliable."
 ::= { optIfOTNPMFECCurrentDayEntry 2}

optIfOTNPMCurrentDayFECCorrectedErr OBJECT-TYPE
    SYNTAX Counter64
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of bits corrected by the FEC are counted in the
         interval."
    ::= { optIfOTNPMFECCurrentDayEntry 3}

optIfOTNPMCurrentDayFECUncorrectedWords OBJECT-TYPE
    SYNTAX Counter64
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of words un-corrected by the FEC are counted over
         the Day."
    ::= { optIfOTNPMFECCurrentDayEntry 4}

optIfOTNPMCurrentDayMinFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The minimum bit error rate at receiving side before the FEC
         function counted over one second .. mantissa. This is the
         minimum PreFEC BER in the current 24hour period."
    ::= { optIfOTNPMFECCurrentDayEntry 5}

optIfOTNPMCurrentDayMinFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The minimum bit error rate at receiving side before the FEC
         function counted over one second .. exponent. This is the
         minimum PreFEC BER in the current 24hour period."
    ::= { optIfOTNPMFECCurrentDayEntry 6}

optIfOTNPMCurrentDayMaxFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
```

```
DESCRIPTION
    "The maximun bit error rate at receiving side before the FEC
    function counted over one second .. mantissa. This is the
    maximum PreFEC BER in the current 24hour period."
 ::= { optIfOTNPMFECCurrentDayEntry 7}

optIfOTNPMCurrentDayMaxFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The maximun bit error rate at receiving side before the FEC
        function counted over one second .. exponent. This is the
        maximum PreFEC BER in the current 24hour period."
    ::= { optIfOTNPMFECCurrentDayEntry 8}

optIfOTNPMCurrentDayAvgFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. mantissa. This is the
        average PreFEC BER in the current 24hour period."
    ::= { optIfOTNPMFECCurrentDayEntry 9}

optIfOTNPMCurrentDayAvgFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. exponent. This is the
        average PreFEC BER in the current 24hour period."
    ::= { optIfOTNPMFECCurrentDayEntry 10}

optIfOTNPMFECCurrentDayElapsedTime OBJECT-TYPE
    SYNTAX Unsigned32
    UNITS "seconds"
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "Time elapsed for current day."
    ::= { optIfOTNPMFECCurrentDayEntry 11}

--
-- FEC PM Prev day Table
--
```

```

optIfOTNPMFECPrevDayTable OBJECT-TYPE
    SYNTAX SEQUENCE OF OptIfOTNPMFECPrevDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A Performance monitoring FEC previous day table."
    ::= { optIfOTNPMGroup 11 }

optIfOTNPMFECPrevDayEntry OBJECT-TYPE
    SYNTAX OptIfOTNPMFECPrevDayEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "A conceptual entry in the Near end or Far End performance
        monitoring FEC previous day table."
    INDEX { ifIndex, optIfOTNPMFECPrevDayType }
    ::= { optIfOTNPMFECPrevDayTable 1 }

OptIfOTNPMFECPrevDayEntry ::=
    SEQUENCE {
        optIfOTNPMFECPrevDayType                OptIfOTNType,
        optIfOTNPMFECPrevDaySuspectedFlag        TruthValue,
        optIfOTNPMPrevDayFECCorrectedErr         Counter64,
        optIfOTNPMPrevDayFECUncorrectedWords     Counter64,
        optIfOTNPMPrevDayMinFECBERMantissa       Unsigned32,
        optIfOTNPMPrevDayMinFECBERExponent       Unsigned32,
        optIfOTNPMPrevDayMaxFECBERMantissa       Unsigned32,
        optIfOTNPMPrevDayMaxFECBERExponent       Unsigned32,
        optIfOTNPMPrevDayAvgFECBERMantissa       Unsigned32,
        optIfOTNPMPrevDayAvgFECBERExponent       Unsigned32,
        optIfOTNPMFECPrevDayTimeStamp            DateAndTime
    }

optIfOTNPMFECPrevDayType OBJECT-TYPE
    SYNTAX OptIfOTNType
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "This parameter indicates the parameters for the table are
        for the Near End or Far End performance data.
        1 - Near End
        2 - Far End "
    ::= { optIfOTNPMFECPrevDayEntry 1}

optIfOTNPMFECPrevDaySuspectedFlag OBJECT-TYPE
    SYNTAX TruthValue
    MAX-ACCESS read-only

```



```
STATUS current
DESCRIPTION
    "If true, the data in this entry may be unreliable."
 ::= { optIfOTNPMFECPrevDayEntry 2}

optIfOTNPMPrevDayFECCorrectedErr OBJECT-TYPE
    SYNTAX Counter64
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of bits corrected by the FEC are counted in the
         previous day."
    ::= { optIfOTNPMFECPrevDayEntry 3}

optIfOTNPMPrevDayFECUncorrectedWords OBJECT-TYPE
    SYNTAX Counter64
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The number of un-corrected words by the FEC are counted
         over the previous Day."
    ::= { optIfOTNPMFECPrevDayEntry 4}

optIfOTNPMPrevDayMinFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The maximun bit error rate at receiving side before the FEC
         function counted over one second .. mantissa. This is the
         maximum Pre FEC BER in the previous 24hour period."
    ::= { optIfOTNPMFECPrevDayEntry 5}

optIfOTNPMPrevDayMinFECBERExponent OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The minimun bit error rate at receiving side before the FEC
         function counted over one second .. exponent. This is the
         maximum Pre FEC BER in the previous 24hour period."
    ::= { optIfOTNPMFECPrevDayEntry 6}

optIfOTNPMPrevDayMaxFECBERMantissa OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
```

```

        "The maximun bit error rate at receiving side before the FEC
        function counted over one second .. mantissa. This is the
        maximum Pre FEC BER in the previous 24hour
        period (mantissa)."
```

```
 ::= { optIfOTNPMFECPrevDayEntry 7}
```

```
optIfOTNPMPrevDayMaxFECBERExponent OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The maximun bit error rate at receiving side before the FEC
        function counted over one second .. exponent (eg -3).
        This is the maximum Pre FEC BER in the previous 24hour
        period."
```

```
 ::= { optIfOTNPMFECPrevDayEntry 8}
```

```
optIfOTNPMPrevDayAvgFECBERMantissa OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. mantissa. This is the
        average Pre FEC BER during the previous 24hour
        period (mantissa)."
```

```
 ::= { optIfOTNPMFECPrevDayEntry 9}
```

```
optIfOTNPMPrevDayAvgFECBERExponent OBJECT-TYPE
    SYNTAX  Unsigned32
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "The average bit error rate at receiving side before the FEC
        function counted over one second .. exponent (eg -3).
        This is the average Pre FEC BER during the previous 24hour
        period."
```

```
 ::= { optIfOTNPMFECPrevDayEntry 10}
```

```
optIfOTNPMFECPrevDayTimeStamp OBJECT-TYPE
    SYNTAX  DateAndTime
    MAX-ACCESS  read-only
    STATUS  current
    DESCRIPTION
        "Time stamp for the Prev day."
```

```
 ::= { optIfOTNPMFECPrevDayEntry 11}
```

```
--
```

```

-- OTN OTUk Alarm Table
--
optIfOTNOChOTUkAlarmTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF OptIfOTNOChOTUkAlarmEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A table of OCh/OTUk alarm entries."

    ::= { optIfOTNAlarm 1 }

optIfOTNOChOTUkAlarmEntry OBJECT-TYPE
    SYNTAX      OptIfOTNOChOTUkAlarmEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual entry in the OCh/OTUk alarm table."
    INDEX { ifIndex }
    ::= { optIfOTNOChOTUkAlarmTable 1 }

OptIfOTNOChOTUkAlarmEntry ::= SEQUENCE {
    optIfOTNOChOTUkAlarmLocation      OptIfOTNType,
    optIfOTNOChOTUkAlarmDirection    OptIfDirectionality,
    optIfOTNOChOTUkAlarmLayer        OptIfOTNLayer,
    optIfOTNOChOTUkAlarmType          OptIfOTNOChAlarms,
    optIfOTNOChOTUkAlarmSeverity      OptIfOTNAlarmSeverity,
    optIfOTNOChOTUkAlarmDate          DateAndTime
}

optIfOTNOChOTUkAlarmLocation OBJECT-TYPE
    SYNTAX      OptIfOTNType
    MAX-ACCESS  accessible-for-notify
    STATUS      current
    DESCRIPTION
        "The object identifies indicates if this entry was for
        Near end/Far end."
    ::= { optIfOTNOChOTUkAlarmEntry 1 }

optIfOTNOChOTUkAlarmDirection OBJECT-TYPE
    SYNTAX      OptIfDirectionality
    MAX-ACCESS  accessible-for-notify
    STATUS      current
    DESCRIPTION
        "The object identifies indicates if this entry was for
        for the Tx/Rx or both."
    ::= { optIfOTNOChOTUkAlarmEntry 2 }

optIfOTNOChOTUkAlarmLayer OBJECT-TYPE

```

```
SYNTAX      OptIfOTNLayer
MAX-ACCESS  accessible-for-notify
STATUS      current
DESCRIPTION
    "This specifies which sublayer this alarm is for."
 ::= { optIfOTNOChOTUkAlarmEntry 3 }

optIfOTNOChOTUkAlarmType OBJECT-TYPE
SYNTAX      OptIfOTNOChAlarms
MAX-ACCESS  accessible-for-notify
STATUS      current
DESCRIPTION
    "This specifies the type of alarm of the sublayer
     'optIfOTNAlarmLayer' ."
 ::= { optIfOTNOChOTUkAlarmEntry 4 }

optIfOTNOChOTUkAlarmSeverity OBJECT-TYPE
SYNTAX      OptIfOTNAlarmSeverity
MAX-ACCESS  accessible-for-notify
STATUS      current
DESCRIPTION
    "The object identifies the severity of the last alarm/alert
     that most recently was set or cleared."
 ::= { optIfOTNOChOTUkAlarmEntry 5 }

optIfOTNOChOTUkAlarmDate OBJECT-TYPE
SYNTAX      DateAndTime
MAX-ACCESS  accessible-for-notify
STATUS      current
DESCRIPTION
    "This specifies the date and time when this alarm occurred."
 ::= { optIfOTNOChOTUkAlarmEntry 6 }

--
-- OTN ODUkTcm Alarm Table
--
optIfOTNODUkTcmAlarmTable OBJECT-TYPE
SYNTAX      SEQUENCE OF OptIfOTNODUkTcmAlarmEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "A table of ODUk/Tcm alarm entries."

 ::= { optIfOTNAlarm 2 }

optIfOTNODUkTcmAlarmEntry OBJECT-TYPE
SYNTAX      OptIfOTNODUkTcmAlarmEntry
MAX-ACCESS  not-accessible
```

```

    STATUS      current
    DESCRIPTION
        "A conceptual entry in the ODUk/Tcm alarm table."
    INDEX { ifIndex }
    ::= { optIfOTNODUkTcmAlarmTable 1 }

OptIfOTNODUkTcmAlarmEntry ::= SEQUENCE {
    optIfOTNODUkTcmAlarmLocation      OptIfOTNType,
    optIfOTNODUkTcmAlarmDirection    OptIfDirectionality,
    optIfOTNODUkTcmAlarmLayer         OptIfOTNLayer,
    optIfOTNODUkTcmAlarmTCMLLevel     Unsigned32,
    optIfOTNODUkTcmAlarmType          OptIfOTNODUkTcmAlarms,
    optIfOTNODUkTcmAlarmSeverity      OptIfOTNAlarmSeverity,
    optIfOTNODUkTcmAlarmDate          DateAndTime
}

optIfOTNODUkTcmAlarmLocation OBJECT-TYPE
    SYNTAX      OptIfOTNType
    MAX-ACCESS   accessible-for-notify
    STATUS      current
    DESCRIPTION
        "The object identifies indicates if this entry was for
        Near end/Far end."
    ::= { optIfOTNODUkTcmAlarmEntry 1 }

optIfOTNODUkTcmAlarmDirection OBJECT-TYPE
    SYNTAX      OptIfDirectionality
    MAX-ACCESS   accessible-for-notify
    STATUS      current
    DESCRIPTION
        "The object identifies indicates if this entry was for
        for the Tx/Rx or both."
    ::= { optIfOTNODUkTcmAlarmEntry 2 }

optIfOTNODUkTcmAlarmLayer OBJECT-TYPE
    SYNTAX      OptIfOTNLayer
    MAX-ACCESS   accessible-for-notify
    STATUS      current
    DESCRIPTION
        "This specifies which sublayer this alarm is for ODUk/TCM."
    ::= { optIfOTNODUkTcmAlarmEntry 3 }

optIfOTNODUkTcmAlarmTCMLLevel OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   accessible-for-notify
    STATUS      current
    DESCRIPTION
        "TCM level 1-6 of the alarm. It will be 0 if alarm sublayer

```

```

        is ODUk."
 ::= { optIfOTNODUkTcmAlarmEntry 4 }

optIfOTNODUkTcmAlarmType OBJECT-TYPE
    SYNTAX      OptIfOTNODUkTcmAlarms
    MAX-ACCESS   accessible-for-notify
    STATUS       current
    DESCRIPTION
        "This specifies the type of alarm 'optIfOTNODUkTcmAlarms'
         of the sublayer ."
 ::= { optIfOTNODUkTcmAlarmEntry 5 }

optIfOTNODUkTcmAlarmSeverity OBJECT-TYPE
    SYNTAX      OptIfOTNAlarmSeverity
    MAX-ACCESS   accessible-for-notify
    STATUS       current
    DESCRIPTION
        "The object identifies the severity of the last alarm/alert
         that most recently was set or cleared."
 ::= { optIfOTNODUkTcmAlarmEntry 6 }

optIfOTNODUkTcmAlarmDate OBJECT-TYPE
    SYNTAX      DateAndTime
    MAX-ACCESS   accessible-for-notify
    STATUS       current
    DESCRIPTION
        "This specifies the date and time when this alarm occurred."
 ::= { optIfOTNODUkTcmAlarmEntry 7 }

--
-- OTN Notifications
--

optIfOTNOChOTUkAlarmSet NOTIFICATION-TYPE
    OBJECTS { optIfOTNOChOTUkAlarmLocation,
               optIfOTNOChOTUkAlarmDirection,
               optIfOTNOChOTUkAlarmLayer,
               optIfOTNOChOTUkAlarmType,
               optIfOTNOChOTUkAlarmSeverity,
               optIfOTNOChOTUkAlarmDate }
    STATUS       current
    DESCRIPTION
        "Notification of a recently set OTN alarm of OCh/OTUk Layer."
 ::= { optIfOTNNotifications 1 }

optIfOTNOChOTUkAlarmClear NOTIFICATION-TYPE
    OBJECTS { optIfOTNOChOTUkAlarmLocation,

```

```
        optIfOTNOChOTUkAlarmDirection,
        optIfOTNOChOTUkAlarmLayer,
        optIfOTNOChOTUkAlarmType,
        optIfOTNOChOTUkAlarmSeverity,
        optIfOTNOChOTUkAlarmDate }
STATUS current
DESCRIPTION
    "Notification of a recently clear OTN alarm of OCh/OTUk
    Layer. "
::= { optIfOTNNotifications 2 }

optIfOTNODUkTcmAlarmSet NOTIFICATION-TYPE
OBJECTS { optIfOTNODUkTcmAlarmLocation,
          optIfOTNODUkTcmAlarmDirection,
          optIfOTNODUkTcmAlarmLayer,
          optIfOTNODUkTcmAlarmTCMLevel,
          optIfOTNODUkTcmAlarmType,
          optIfOTNODUkTcmAlarmSeverity,
          optIfOTNODUkTcmAlarmDate }
STATUS current
DESCRIPTION
    "Notification of a recently set OTN alarm of OTUk/Tcm
    Layer. "
::= { optIfOTNNotifications 3 }

optIfOTNODUkTcmAlarmClear NOTIFICATION-TYPE
OBJECTS { optIfOTNODUkTcmAlarmLocation,
          optIfOTNODUkTcmAlarmDirection,
          optIfOTNODUkTcmAlarmLayer,
          optIfOTNODUkTcmAlarmTCMLevel,
          optIfOTNODUkTcmAlarmType,
          optIfOTNODUkTcmAlarmSeverity,
          optIfOTNODUkTcmAlarmDate }
STATUS current
DESCRIPTION
    "Notification of a recently cleared OTN alarm of OTUk/Tcm
    Layer. "
::= { optIfOTNNotifications 4 }

END
```

7. Relationship to Other MIB Modules

7.1. Relationship to the [TEMPLATE TODO] MIB

7.2. MIB modules required for IMPORTS

8. Definitions

[TEMPLATE TODO]: put your valid MIB module here.
A list of tools that can help automate the process of checking MIB definitions can be found at <http://www.ops.ietf.org/mib-review-tools.html>

9. Security Considerations

There are a number of management objects defined in this MIB module with a MAX-ACCESS clause of read-write and/or read-create. Such objects may be considered sensitive or vulnerable in some network environments. The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations. These are the tables and objects and their sensitivity/vulnerability:

Some of the readable objects in this MIB module (i.e., objects with a MAX-ACCESS other than not-accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to control even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP.

SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example by using IPsec), even then, there is no control as to who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in this MIB module.

It is RECOMMENDED that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Further, deployment of SNMP versions prior to SNMPv3 is NOT RECOMMENDED. Instead, it is RECOMMENDED to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator responsibility to ensure that the SNMP entity giving access to an instance of this MIB module is properly configured to give access to the objects only to those principals (users) that have legitimate rights to indeed GET or SET (change/create/delete) them.

10. IANA Considerations

Option #1:

The MIB module in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
-----	-----
sampleMIB	{ mib-2 XXX }

Option #2:

Editor's Note (to be removed prior to publication): the IANA is requested to assign a value for "XXX" under the 'mib-2' subtree and to record the assignment in the SMI Numbers registry. When the assignment has been made, the RFC Editor is asked to replace "XXX" (here and in the MIB module) with the assigned value and to remove this note.

Note well: prior to official assignment by the IANA, an internet draft MUST use placeholders (such as "XXX" above) rather than actual numbers. See RFC4181 Section 4.5 for an example of how this is done in an internet draft MIB module.

Option #3:

This memo includes no request to IANA.

11. Contributors

Arnold Mattheus
Deutsche Telekom
Darmstadt
Germany
email a.mattheus@telekom.de

Manuel Paul

Deutsche Telekom
Berlin
Germany
email Manuel.Paul@telekom.de

Frank Luennemann
Deutsche Telekom
Munster
Germany
email Frank.Luennemann@telekom.de

Scott Mansfield
Ericsson Inc.
email scott.mansfield@ericsson.com

Najam Saquib
Cisco
Ludwig-Erhard-Strasse 3
ESCHBORN, HESSEN 65760
GERMANY
email nasaquib@cisco.com

Walid Wakim
Cisco
9501 Technology Blvd
ROSEMONT, ILLINOIS 60018
UNITED STATES
email wwakim@cisco.com

Ori Gerstel
Cisco
32 HaMelacha St., (HaSharon Bldg)
SOUTH NETANYA, HAMERKAZ 42504
ISRAEL
email ogerstel@cisco.com

12. References

12.1. Normative References

- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.
- [RFC3591] Lam, H-K., Stewart, M., and A. Huynh, "Definitions of Managed Objects for the Optical Interface Type", RFC 3591, September 2003.
- [RFC6205] Otani, T. and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.
- [ITU.G698.2] International Telecommunications Union, "Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces ", ITU-T Recommendation G.698.2, November 2009.
- [ITU.G709] International Telecommunications Union, "Interface for the Optical Transport Network (OTN) ", ITU-T Recommendation G.709, March 2003.
- [ITU.G872] International Telecommunications Union, "Architecture of optical transport networks ", ITU-T Recommendation G.872, November 2001.
- [ITU.G798] International Telecommunications Union, "Characteristics of optical transport network hierarchy equipment functional blocks ", ITU-T Recommendation G.798, October 2010.
- [ITU.G874] International Telecommunications Union, "Management aspects of optical transport network elements ", ITU-T Recommendation G.874, July 2010.
- [ITU.G874.1]

International Telecommunications Union, "Optical transport network (OTN): Protocol-neutral management information model for the network element view ", ITU-T Recommendation G.874.1, January 2002.

[ITU.G959.1]

International Telecommunications Union, "Optical transport network physical layer interfaces ", ITU-T Recommendation G.959.1, November 2009.

[ITU.G826]

International Telecommunications Union, "End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections ", ITU-T Recommendation G.826, November 2009.

[ITU.G8201]

International Telecommunications Union, "Error performance parameters and objectives for multi-operator international paths within the Optical Transport Network (OTN) ", ITU-T Recommendation G.8201, April 2011.

[ITU.G694.1]

International Telecommunications Union, "Spectral grids for WDM applications: DWDM frequency grid ", ITU-T Recommendation G.694.1, June 2002.

[ITU.G7710]

International Telecommunications Union, "Common equipment management function requirements ", ITU-T Recommendation G.7710, May 2008.

12.2. Informative References

[RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC4181] Heard, C., "Guidelines for Authors and Reviewers of MIB Documents", BCP 111, RFC 4181, September 2005.

[I-D.kunze-g-698-2-management-control-framework]

Kunze, R., "A framework for Management and Control of optical interfaces supporting G.698.2", draft-kunze-g-698-2-management-control-framework-00 (work in progress), July 2011.

[RFC4054] Strand, J. and A. Chiu, "Impairments and Other Constraints on Optical Layer Routing", RFC 4054, May 2005.

Appendix A. Change Log

This optional section should be removed before the internet draft is submitted to the IESG for publication as an RFC.

Note to RFC Editor: please remove this appendix before publication as an RFC.

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Authors' Addresses

Gabriele Galimberti (editor)
Cisco
Via Philips,12
20052 - Monza
Italy

Phone: +390392091462
Email: ggalimbe@cisco.com

Ruediger Kunze (editor)
Deutsche Telekom
Dddd, xx
Berlin
Germany

Phone: +49xxxxxxxxxx
Email: RKunze@telekom.de

Hing-Kam Lam (editor)
Alcatel-Lucent
600-700 Mountain Avenue, Murray Hill
New Jersey, 07974
USA

Phone: +17323313476
Email: kam.lam@alcatel-lucent.com

Dharini Hiremagalur (editor)
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 California
USA

Phone: +1408
Email: dharinih@juniper.net

CCAMP Working Group
Internet-Draft
Intended status: BCP
Expires: January 16, 2014

Rakesh Gandhi
Zafar Ali
Gabriele Maria Galimberti
Cisco Systems, Inc.
July 15, 2013

RSVP-TE Extensions For Signaling GMPLS Restoration LSP
draft-gandhi-ccamp-gmpls-restoration-lsp-01

Abstract

In transport networks, there are requirements where Generalized Multi-Protocol Label Switching (GMPLS) recovery scheme need to employ restoration LSP while keeping resources for the working and/ or protecting LSPs reserved in the network. Existing GMPLS recovery procedures do not address these requirements. This document describes best common practice for using RSVP-TE for GMPLS recovery with restoration LSP.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 31, 2013.

Copyright Notice

Copyright (c) 2013 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
2. Conventions used in this document	4
3. Restoration LSP Signaling Extensions	5
3.1. Signaling Procedure	5
4. IANA Considerations	5
5. Security Considerations	5
6. Acknowledgement	5
7. References	6
7.1. Normative references	6
7.2. Informative References	7
Authors' Addresses	7

1. Introduction

Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for different switching technologies [RFC3471] [RFC3473]. These switching technologies provide several protection schemes [RFC4426][RFC4427] (e.g. 1+1, 1:N and M:N). GMPLS RSVP-TE signaling has been extended to support various recovery schemes to establish Label Switched Paths (LSPs) [RFC4872][RFC4873], typically working LSP and protecting LSP. [RFC4427] Section 7 specifies various schemes for GMPLS restoration.

In GMPLS recovery schemes currently considered, restoration LSP is signaled after the failure has been detected and notified on the working LSP. These schemes assume that working LSP is removed from the network before restoration LSP is signaled. In transport networks, as working LSPs are typically signaled over a nominal path, there are many scenarios where service providers would like to keep resources associated with the working LSPs reserved. This is to make sure that the service (working LSP) can use the nominal path when the failure is repaired. Consequently, in transport networks one can employ a recovery scheme where a new restoration LSP is signaled while working LSP and/ or protecting LSP are not torn down in control plane due to a failure. Restoration LSP differs from a secondary LSP in the way that secondary LSP does not reserve resources in the data plane and is not able to carry any traffic until it is refreshed whereas restoration LSP does reserve resources and is able to carry traffic.

One example of the recovery scheme considered in this draft is 1+R recovery. The 1+R recovery is exemplified in Figure 1. In this example, working LSP on path A-B-C-Z is pre-established. Typically after a failure detection and notification on the working LSP, a second LSP on path A-H-I-J-Z is established as a restoration LSP. Unlike protection LSP, restoration LSP is signaled on as needed basis.




Figure 1: An example of 1+R recovery scheme

During failure with 1+R recovery scheme, in general, working LSP resources are not released and working and restoration LSPs coexist in the network. Nonetheless, working and restoration LSPs can share

network resources. Typically when failure is recovered on the working LSP, restoration LSP is no longer required and torn down (e.g. revertive mode).

Another example of the recovery scheme considered in this draft is 1+1+R. In 1+1+R, a restoration LSP is signaled for the working LSP and/ or the protecting LSP after the failure has been detected and notified on the working LSP or the protecting LSP. The 1+1+R recovery is exemplified in Figure 2. In this example, working LSP on path A-B-C-Z and protecting LSP on path A-D-E-F-Z are pre-established. After a failure detection and notification on a working LSP or protecting LSP, a third LSP on path A-H-I-J-Z is established as a restoration LSP. The restoration LSP in this case provides protection against a second order failure. Restoration LSP is torn down when the failure on the working or protecting LSP is repaired.




Figure 2: An example of 1+1+R recovery scheme

[RFC4872] Section 14 defines PROTECTION object for GMPLS recovery signaling. The PROTECTION object is used to identify primary and secondary LSPs using S bit and protecting and working LSPs using P bit. However, the PROTECTION object does not have a way to identify restoration LSP. [RFC4872] and [RFC6689] define the usage of ASSOCIATION object for further associating GMPLS working and protecting LSPs for the case where restoration LSP is signaled for GMPLS recovery after the working or protecting LSPs are removed.

This draft outlines the best common practice for identifying the restoration LSP for GMPLS recovery where working and protecting LSP resources are kept reserved.

2. 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 [RFC2119].

3. Restoration LSP Signaling Extensions

3.1. Signaling Procedure

Where GMPLS recovery scheme need to employ restoration LSP while keeping resources for the working and/ or protecting LSPs reserved in the network, restoration LSP is signaled with ASSOCIATION object with the association ID set to the LSP ID of the LSP it is restoring. For example, when a restoration LSP is signaled for a working LSP, the ASSOCIATION object in the restoration LSP contains the association ID set to the LSP ID of the working LSP. Similarly, when a restoration LSP is signaled for a protecting LSP, the ASSOCIATION object in the restoration LSP contains the association ID set to the LSP ID of the protecting LSP.

The procedure for signaling the PROTECTION object is specified in [RFC4872][RFC4873] and does not change. Restoration LSP for the working LSP is signaled with P bit cleared and restoration LSP for the protecting LSP is signaled with P bit set.

When using a GMPLS recovery mode, where working LSP is destroyed, and the restoration LSP is promoted to be the new working LSP, restoration LSP RSVP Path message MUST be refreshed by using the ASSOCIATION_OBJECT.LSP_ID from the destroyed working LSP ASSOCIATION_OBJECT.LSP_ID.

When using a GMPLS recovery mode, where a protecting LSP is destroyed, and the restoration LSP is promoted to be the new protecting LSP, restoration LSP RSVP Path message MUST be refreshed by using the ASSOCIATION_OBJECT.LSP_ID from the destroyed protecting LSP ASSOCIATION_OBJECT.LSP_ID.

4. IANA Considerations

This document makes no request for IANA action.

5. Security Considerations

This document introduces no additional security considerations. For a general discussion on MPLS and GMPLS related security issues, see the MPLS/GMPLS security framework [RFC5920]. In addition, the considerations specified in [RFC4872] and [RFC4873] will apply.

6. Acknowledgement

The authors would like to thank George Swallow for the discussion on the GMPLS restoration.

7. References

7.1. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [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., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4872] Lang, J., Rekhter, Y., and D. Papadimitriou, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC6689] Berger, L., "Usage of the RSVP ASSOCIATION Object", RFC 6689, July 2012.

7.2. Informative References

- [RFC4426] Lang, J., Rajagopalan B., and D.Papadimitriou, Editors,
 "Generalized Multiprotocol Label Switching (GMPLS)
 Recovery Functional Specification", RFC 4426, March 2006.
- [RFC4427] Mannie, E., Ed. and D. Papadimitriou, Ed., "Recovery
 (Protection and Restoration) Terminology for Generalized
 Multi-Protocol Label Switching, RFC 4427, March 2006.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS
 Networks", RFC 5920, July 2010.

Authors' Addresses

Rakesh Gandhi
Cisco Systems, Inc.

Email: rgandhi@cisco.com

Zafar Ali
Cisco Systems, Inc.

Email: zali@cisco.com

Gabriele Maria Galimberti
Cisco Systems, Inc.

Email: ggalimbe@cisco.com

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 15, 2014

G. Galimberti, Ed.
G. Martinelli, Ed.
Cisco
D. Hiremagalur
G. Grammel
Juniper
July 14, 2013

A SNMP MIB to manage GMPLS with General Constraints support
draft-gmngm-ccamp-gencons-snmp-mib-02

Abstract

This memo defines a portion of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) for GMPLS based networks.

In particular in the context Wavelength Switching Optical Network (WSO) two sets of information were defined: a general constraints set (reusable by other technologies) and a WSO specific set. This document defines a MIB module for supporting general constraint information.

Copyright Notice

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 15, 2014.

Copyright Notice

Copyright (c) 2013 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	2
2. The Internet-Standard Management Framework	3
3. Conventions	3
4. Overview	3
5. Structure of the MIB Module	4
5.1. gmplsGenConsAvailableLabelsTable	4
5.2. gmplsGenConsSharedBackupLabelsTable	5
5.3. gmplsGenConsConnMatrixTable	5
5.4. gmplsGenConsPortLabelRestrictionTable	5
6. Relationship to Other MIB Modules	6
6.1. Relationship to the [TEMPLATE TODO] MIB	6
6.2. MIB modules required for IMPORTS	6
7. Definitions	6
8. Security Considerations	12
9. IANA Considerations	13
10. Contributors	13
11. References	14
11.1. Normative References	14
11.2. Informative References	14
Appendix A. Change Log	15
Appendix B. Open Issues	15
Authors' Addresses	15

1. Introduction

This memo defines a portion of the Management Information Base (MIB) used by Simple Network Management Protocol (SNMP) in GMPLS networks in particular for Wavelength Switched Optical Networks (WSON) as defined in [RFC6163].

Those extensions were divided in two parts: generic constrains (as they can be easily applied to other technologies) and WSON specific constraints. This document aim to defines MIBs extention to conver

only the generic constrain part. The WSON specific MIB extentions will be covered by a separate document [I-D.gmngm-ccamp-wson-snmp-mib].

As such, document [I-D.ietf-ccamp-general-constraint-encode] defines specific TLVs while [RFC6825] implement OSPF-TE related extentions. This MIB document aim to cover information defined in those general constrain drafts.

[EDITOR NOTE] Very early draft to start MIB activity on GMPSL-WSON related extentions and collect feedback from working group.

2. The Internet-Standard Management Framework

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to section 7 of RFC 3410 [RFC3410].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies a MIB module that is compliant to the SMIV2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

3. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119]. In the description of OIDs the convention: Set (S) Get (G) and Trap (T) conventions will describe the action allowed by the parameter.

4. Overview

Regarding existing GMPLS MIBs modules, since the TED module [RFC6825] already provide and extension to previous GMPLS modules, we provide here a direct extension to it. Additional GMPLS MIB modules this document uses are [RFC4802] and [RFC4803].

Current GMPLS MIBs are covered by several documents. The most important to reference here are the [RFC4802] and [RFC4803]. Most recent works on GMPLS MIBs is in [RFC6825], whenever possible this document will reuse the same approach.

General constrains can be classified in two broad categories: link information (as other GMPLS TED information) and Node information (this is different from what currently available). For link information the most similar definitions are the ones from [RFC4803] where the label table is defined. For node information however, new specific information has to be defined.

5. Structure of the MIB Module

Modules defined here provide additional information to existing GMPLS MIBS in order to represent the general constrains information as reported in [I-D.ietf-ccamp-general-constraint-encode]. This module is organized into two tables as reported in the following sub sections.

5.1. gmplsGenConsAvailableLabelsTable

This object represent the Labels availability as defined by [I-D.ietf-ccamp-general-constraint-encode] section 2.3. This information may be introduced by specific technologies but is represented in a general form. An application example is [draft-ietf-ccamp-wson-signal-compatibility-ospf] that advertise such information under its specific ISC values.

The table entry is composed by:

- o gmplsGenConsLabelIndex (type: Unsigned 32)
- o gmplsGenConsLabelInterface (type: typesInterfaceIndexOrZero)
- o gmplsGenConsISCD (type: IANA GmplsSwitchingTypeTC needs to be properly extended).
[EDITOR NOTE1: could we assume that ISCD=151 for WSON hence the associated label follows RFC6205. Likely flexgrid will have the same constrain.]
[EDITOR NOTE2: we could probably add an index reference to the tedSwCapTable so we have all switching capability info through one pointer.]
- o gmplsGenConsLabelValue (type: a value that might be greater than 32 bits, not sure about max length). Value format may vary upon the ISCD however the generic label format defined in the past still apply. A type might be reused could be (gmplsLabelPortWavelength Unsigned32) defined in GMPLS LSR MIB. However we need to support labels bigger than 32 bits.

Depending on switching capability available on an interface, this table MUST initialized with all possible available labels values. As

long an tunnels are created and labels are used, the existing `gmplsLabelTable` defined in [RFC4803], is filled up while entries are removed from this `gmplsGenConsAvailableLabelsTable`.

5.2. `gmplsGenConsSharedBackupLabelsTable`

The purpose of shared backup labels is defined in [I-D.ietf-ccamp-general-constraint-encode]. As in the previous case the information is advertised through the specific ISCD.

The entry for this table is equal to definitions in Section 5.1 hence we have the same table format.

5.3. `gmplsGenConsConnMatrixTable`

The Connectivity Matrix indicates the Node constraints introduced by [I-D.ietf-ccamp-general-constraint-encode] as additional constrains compared to link/label constrains.

The table entry shall have the following information:

- o An identifier of the local node. E.g. `tedLocalRouterId` (type: `TedRouterIdTC`).
- o Connectivity Matrix Type (1: RWA, 2: Optical Impairments)
- o Connectivity Matrix ID: the unique identifier for the current connectivity matrix. Type: integer.
- o Link Ingress: (Type: `TedLinkIndexTC`) the link identifier for the ingress link (ingress and egress link identify a possible connectivity for the node).
- o Link Egress: (Type: `TedLinkIndexTC`) the link identifier for the egress link (ingress and egress link indentify a possible connectivity for the node).

5.4. `gmplsGenConsPortLabelRestrictionTable`

This table the port label constraints introduced by [I-D.ietf-ccamp-general-constraint-encode] as an additional constrains on the port or vs permitted labels. This constrains related to the connectivity matrix (from previous section) and are advertised through [I-D.ietf-ccamp-gmpls-general-constraints-ospf-te].

The entry for this table has the following elements:

- o An identifier for the local node. E.g. tedLocalRouterId (type: TedRouterIdTC).
- o Connectivity matrix id: an index in the previous table
- o Restriction Type (Values: 0 Simple Label, 1 Channel Count, 2 Label Rangel, 3 Simple Label and Channel Count, 4 Link Label Exclusivity)
- o Label Set: used depending on restriction type
- o Max Channel Count / Max Label Set (Integer): used depending on restriction type
- o Link Set: used depending on restrictio type

6. Relationship to Other MIB Modules

6.1. Relationship to the [TEMPLATE TODO] MIB

6.2. MIB modules required for IMPORTS

7. Definitions

TED-GENCONS-MIB DEFINITIONS ::= BEGIN

IMPORTS

MODULE-IDENTITY, OBJECT-TYPE, TimeTicks, NOTIFICATION-TYPE,
 Unsigned32, Counter32, Integer32
 FROM SNMPv2-SMI
 DateAndTime, TEXTUAL-CONVENTION, RowStatus, TruthValue
 FROM SNMPv2-TC
 IANAGmplsSwitchingTypeTC
 FROM IANA-GMPLS-TC-MIB
 GmplsFreeformLabelTC
 FROM GMPLS-TC-STD-MIB
 ifIndex, ifDescr
 FROM IF-MIB;

tedGenConsGMPLSMibModule MODULE-IDENTITY

LAST-UPDATED

"201307070000Z" -- Thu Jul 7 10:00:00 PST 2013

ORGANIZATION

" IETF Common Control And Measurement Plane (CCAMP) Working
 Group "

CONTACT-INFO

"WG charter:
<http://www.ietf.org/html.charters/>

Mailing Lists:
 Editor: Gabriele Galimberti
 Email: ggalimbe@cisco.com"

DESCRIPTION

"This MIB module defines objects used for managing the
 the generic constraints for switched networks in GMPLS
 networks ."

REVISION "201201270000Z"

DESCRIPTION

"Draft version 1.0"

::= { tedGenConsGMPLSMibRoot 1 }

--

-- Textual Conventions

--

TedGenConsInterfaceIndexOrZero ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION "interface index from 0 ..2147483647"

SYNTAX INTEGER (0..2147483647)

tedGenConsGmpls OBJECT IDENTIFIER ::= { tedGenConsGMPLSMibModule 1 }

gmplsGenConsAvailableLabelsTable OBJECT-TYPE

SYNTAX SEQUENCE OF GmplsGenConsAvailableLabelsEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"Information about the shared backup Labels
 availability as defined by
 [I-D.ietf-ccamp-general-constraint-encode]
 section 2.3. "

::= { tedGenConsGmpls 1 }

gmplsGenConsAvailableLabelsEntry OBJECT-TYPE

SYNTAX GmplsGenConsAvailableLabelsEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A conceptual row availability labels Table."

INDEX { gmplsGenConsLabelIndex }

::= { gmplsGenConsAvailableLabelsTable 1 }

GmplsGenConsAvailableLabelsEntry ::=

SEQUENCE {
 gmplsGenConsLabelIndex

```

        Unsigned32,
gmplsGenConsLabelInterface
        TedGenConsInterfaceIndexOrZero,
gmplsGenConsISCD
        IANAGmplsSwitchingTypeTC,
gmplsGenConsLabelValue
        Unsigned32
    }

gmplsGenConsLabelIndex OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        " Label Index for this table"
    ::= { gmplsGenConsAvailableLabelsEntry 1 }

gmplsGenConsLabelInterface OBJECT-TYPE
    SYNTAX      TedGenConsInterfaceIndexOrZero
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The interface Index of this interface"
    ::= { gmplsGenConsAvailableLabelsEntry 2 }

gmplsGenConsISCD OBJECT-TYPE
    SYNTAX      IANAGmplsSwitchingTypeTC
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The interface switching type as defined in rfc4802.
        The type needs to be extended for eg flex grid"
    ::= { gmplsGenConsAvailableLabelsEntry 3 }

gmplsGenConsLabelValue OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The Label value."
    ::= { gmplsGenConsAvailableLabelsEntry 4 }

--
--
-- This purpose of the shared backup tabble is defined by
-- [I-D.ietf-ccamp-general-constraint-encode].
-- As if the available table the information is advertised through

```

```

-- the specific ISCD
--
gmplsGenConsSharedBackupLabelsTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF GmplsGenConsSharedBackupLabelsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "Information about the shared backup Labels
        availability as defined by
        [I-D.ietf-ccamp-general-constraint-encode]
        section 2.3. "
    ::= { tedGenConsGmpls 2 }

gmplsGenConsSharedBackupLabelsEntry OBJECT-TYPE
    SYNTAX      GmplsGenConsSharedBackupLabelsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "A conceptual row  availability labels Table."
    INDEX      { gmplsGenConsBackupLabelIndex }
    ::= { gmplsGenConsSharedBackupLabelsTable 1 }

GmplsGenConsSharedBackupLabelsEntry ::=
    SEQUENCE {
        gmplsGenConsBackupLabelIndex
            Unsigned32,
        gmplsGenConsBackupLabelInterface
            TedGenConsInterfaceIndexOrZero,
        gmplsGenConsBackupISCD
            IANAGmplsSwitchingTypeTC,
        gmplsGenConsBackupLabelValue
            Unsigned32
    }

gmplsGenConsBackupLabelIndex OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        " Label Index for this table"
    ::= { gmplsGenConsSharedBackupLabelsEntry 1 }

gmplsGenConsBackupLabelInterface OBJECT-TYPE
    SYNTAX      TedGenConsInterfaceIndexOrZero
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION

```

```

        "The interface Index of this interface"
 ::= { gmplsGenConsSharedBackupLabelsEntry 2 }

gmplsGenConsBackupISCD OBJECT-TYPE
    SYNTAX      IANAGmplsSwitchingTypeTC
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The interface switching type as defined in rfc4802.
         The type needs to be extended for eg flex grid"
 ::= { gmplsGenConsSharedBackupLabelsEntry 3 }

gmplsGenConsBackupLabelValue OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The Label value."
 ::= { gmplsGenConsSharedBackupLabelsEntry 4 }

--
-- gmplsGenConsPortNodeTable
-- This table indicates the Node constraints introduced by
-- [I-D.ietf-ccamp-general-constraint-encode] as additional
-- constraints compared to link/label constraints reported above

gmplsGenConsPortNodeTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF GmplsGeniPortConsNodeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "Information about the Port-label Node constraints."
 ::= { tedGenConsGmpls 3 }

gmplsGenConsPortNodeEntry OBJECT-TYPE
    SYNTAX      GmplsGenConsPortNodeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A conceptual row in the Node Table."
    INDEX      { gmplsGenConsPortNodeIndex }
 ::= { gmplsGenConsPortNodeTable 1 }

GmplsGenConsPortNodeEntry ::=
    SEQUENCE {
        gmplsGenConsPortNodeIndex
            Unsigned32,
```

```

        gmplsGenConsPortLabelRestriction
            Unsigned32
    }

    gmplsGenConsPortNodeIndex OBJECT-TYPE
        SYNTAX      Unsigned32
        MAX-ACCESS   not-accessible
        STATUS       current
        DESCRIPTION
            " Port node restristion Index for this table"
        ::= { gmplsGenConsPortNodeEntry 1 }

    gmplsGenConsBackupLabelInterface OBJECT-TYPE
        SYNTAX      Unsigned32
        MAX-ACCESS   read-write
        STATUS       current
        DESCRIPTION
            "This information represent a constain on ports vs
            labels (i.e. some ports may not support all
            wavelenghts)."
        ::= { gmplsGenConsPortNodeEntry 2 }

-- gmplsGenConsConnectvityNodeTable
-- This table indicates the Node's connectivity matrix
-- ie some internal constraints in terms of connectivity
gmplsGenConsiConnectvityNodeTable  OBJECT-TYPE
    SYNTAX      SEQUENCE OF GmplsGeniConnectvityConsNodeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "Information about the constraints in terms of
        connectivity for the node. "
    ::= { tedGenConsGmpls 4 }

gmplsGenConsConnectvityNodeEntry OBJECT-TYPE
    SYNTAX      GmplsGenConsConnectvityNodeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A conceptual row in the Connectivity Node Table."
    INDEX       { gmplsGenConsConnectvityNodeIndex }
    ::= { gmplsGenConsPortConnectvityTable 1 }

GmplsGenConsConnectvityNodeEntry ::=
    SEQUENCE {
        gmplsGenConsConnectvityNodeIndex
            Unsigned32,

```



```
        gmplsGenConsConnectvityMatrix
            Unsigned32
    }

    gmplsGenConsConnectvityNodeIndex OBJECT-TYPE
        SYNTAX             Unsigned32
        MAX-ACCESS          not-accessible
        STATUS              current
        DESCRIPTION
            " Connectivity matrix node Index for this table"
        ::= { gmplsGenConsConnectvityNodeEntry 1 }

    gmplsGenConsConnectvityMatrix OBJECT-TYPE
        SYNTAX             Unsigned32
        MAX-ACCESS          read-write
        STATUS              current
        DESCRIPTION
            "This information represent some node internal
             constraint in term of connectivity."
        ::= { gmplsGenConsConnectvityNodeEntry 2 }

END
```

8. Security Considerations

There are a number of management objects defined in this MIB module with a MAX-ACCESS clause of read-write and/or read-create. Such objects may be considered sensitive or vulnerable in some network environments. The support for SET operations in a non-secure environment without proper protection can have a negative effect on network operations. These are the tables and objects and their sensitivity/vulnerability:

Some of the readable objects in this MIB module (i.e., objects with a MAX-ACCESS other than not-accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to control even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP.

SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example by using IPsec), even then, there is no control as to who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in this MIB module.

It is RECOMMENDED that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Further, deployment of SNMP versions prior to SNMPv3 is NOT RECOMMENDED. Instead, it is RECOMMENDED to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator responsibility to ensure that the SNMP entity giving access to an instance of this MIB module is properly configured to give access to the objects only to those principals (users) that have legitimate rights to indeed GET or SET (change/create/delete) them.

9. IANA Considerations

Option #1:

The MIB module in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
-----	-----
sampleMIB	{ mib-2 XXX }

Option #2:

Editor's Note (to be removed prior to publication): the IANA is requested to assign a value for "XXX" under the 'mib-2' subtree and to record the assignment in the SMI Numbers registry. When the assignment has been made, the RFC Editor is asked to replace "XXX" (here and in the MIB module) with the assigned value and to remove this note.

Note well: prior to official assignment by the IANA, an internet draft MUST use placeholders (such as "XXX" above) rather than actual numbers. See RFC4181 Section 4.5 for an example of how this is done in an internet draft MIB module.

Option #3:

This memo includes no request to IANA.

10. Contributors

to be added.

11. References

11.1. Normative References

- [RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.
- [RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.
- [RFC6825] Miyazawa, M., Otani, T., Kumaki, K., and T. Nadeau, "Traffic Engineering Database Management Information Base in Support of MPLS-TE/GMPLS", RFC 6825, January 2013.
- [I-D.ietf-ccamp-general-constraint-encode] Bernstein, G., Lee, Y., Li, D., and W. Imajuku, "General Network Element Constraint Encoding for GMPLS Controlled Networks", draft-ietf-ccamp-general-constraint-encode-11 (work in progress), May 2013.
- [I-D.ietf-ccamp-gmpls-general-constraints-ospf-te] Zhang, F., Lee, Y., Han, J., Bernstein, G., and Y. Xu, "OSPF-TE Extensions for General Network Element Constraints", draft-ietf-ccamp-gmpls-general-constraints-ospf-te-05 (work in progress), June 2013.

11.2. Informative References

- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

- [RFC4181] Heard, C., "Guidelines for Authors and Reviewers of MIB Documents", BCP 111, RFC 4181, September 2005.
- [RFC4802] Nadeau, T. and A. Farrel, "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007.
- [RFC4803] Nadeau, T. and A. Farrel, "Generalized Multiprotocol Label Switching (GMPLS) Label Switching Router (LSR) Management Information Base", RFC 4803, February 2007.
- [RFC6163] Lee, Y., Bernstein, G., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, April 2011.
- [I-D.ietf-ccamp-rwa-info]
Lee, Y., Bernstein, G., Li, D., and W. Imajuku, "Routing and Wavelength Assignment Information Model for Wavelength Switched Optical Networks", draft-ietf-ccamp-rwa-info-18 (work in progress), May 2013.
- [I-D.gmggm-ccamp-wson-snmp-mib]
Galimberti, G., Martinelli, G., Hiremagalur, D., and G. Grammel, "A SNMP MIB to manage GMPLS TED with WSON specific support", draft-gmggm-ccamp-wson-snmp-mib-00 (work in progress), July 2012.

Appendix A. Change Log

This optional section should be removed before the internet draft is submitted to the IESG for publication as an RFC.

Note to RFC Editor: please remove this appendix before publication as an RFC.

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Authors' Addresses

Gabriele M. Galimberti (editor)
Cisco
Via Philips,12
20900 - Monza
Italy

Phone: +390392091462
Email: ggalimbe@cisco.com

Giovanni Martinelli (editor)
Cisco
Via Philips,12
20900 - Monza
Italy

Email: giomarti@cisco.com

Dharini Hiremagalur
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 CA
USA

Email: dharinih@juniper.net

Gert Grammel
Juniper
1194 N Mathilda Avenue
Sunnyvale - 94089 CA
USA

Email: ggrammel@juniper.net

CCAMP WG
Internet Draft
Updates: 4427
Intended status: Standards track
Expires: January 2014

Huub van Helvoort
Huawei Technologies
July 11, 2013

Update Forced Switch Priority
draft-helvoort-ccamp-fs-priority-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on January 11, 2014.

Copyright Notice

Copyright (c) 2013 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.

Abstract

This document clarifies the definitions related to Manual Switch and Forced Switch. This document updates RFC 4427.

Table of Contents

1. Introduction	2
2. Manual Switch and Forced Switch.....	3
3. Security Considerations	3
4. IANA Considerations	3
5. References	4
5.1. Normative References.....	4
5.2. Informative References.....	4
6. Acknowledgments	4

1. Introduction

The external commands, Manual Switch and Forced Switch, provide an operator the ability to control the recovery schemes. The definitions in [RFC4427] provide an informative description but do not provide enough description to distinguish the processing of these commands for usage with priority treatment. The current description has led to the open question of how the commands are to be processed for relative priority treatment.

This document provides clarification of the terms relative to [G.808.1].

2. Manual Switch and Forced Switch

This section updates [RFC4427]. Section 4.13 of [RFC4427] contains the following text regarding the definitions:

D. Forced switch-over for normal traffic:

A switch-over action, initiated externally, that switches normal traffic to the recovery LSP/span, unless an equal or higher priority switch-over command is in effect.

E. Manual switch-over for normal traffic:

A switch-over action, initiated externally, that switches normal traffic to the recovery LSP/span, unless a fault condition exists on other LSPs/spans (including the recovery LSP/span) or an equal or higher priority switch-over command is in effect.

This definition does not provide enough detail with respect to their usage relative to priorities. In order to avoid mis-interpretation, this document adds the following Note as clarification:

Note:

For 1+1 protection schemes (which do not use a communication channel), a Forced Switch over for normal traffic will have the highest priority, it will not be overridden by a signal fail on the protection channel. For other protection schemes (which do use a communication channel), a Forced Switch over for normal traffic will be overridden if a signal fail is present on the protection channel. Definitions of ITU-T terminology in this section are intended to aid understanding of the concepts. For the full definition of these terms and their use, the reader is referred to the appropriate ITU-T Recommendations and RFCs.

3. Security Considerations

This document clarifies usage of terms defined in [RFC4427]. No new information is conveyed; therefore no additional considerations are included here.

4. IANA Considerations

There are no items for IANA to consider.

5. References

5.1. Normative References

[RFC4427] Mannie E. and Papadimitriou D., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS) ", RFC 4427, March 2006.

5.2. Informative References

[G.808.1] ITU-T Recommendation G.808.1, "Generic protection switching - Linear trail and subnetwork protection", February 2010.

6. Acknowledgments

<Add any acknowledgements>

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

Authors' Addresses

Huub van Helvoort

Huawei Technologies

Karspeldreef 4

Amsterdam 1011 CJ

The Netherlands

Phone: +31 20 4300832

Email: Huub.van.Helvoort@huawei.com

CCAMP
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

C. Margaria, Ed.
Coriant GmbH
G. Martinelli
Cisco
S. Balls
B. Wright
Metaswitch
July 15, 2013

LSP Attribute in ERO
draft-ietf-ccamp-lsp-attribute-ro-02

Abstract

LSP attributes can be specified or recorded for whole path, but they cannot be targeted to a specific hop. This document proposes alternative ways to extend the semantic for RSVP ERO object to target LSP attributes to a specific hop.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 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	2
1.1. Contributing Authors	2
1.2. Requirements Language	3
2. Requirements	3
3. ERO Attribute	3
3.1. ERO_ATTRIBUTE subobject	3
3.2. HOP Attributes TLVs	4
3.3. Procedures	5
4. Recording Hop attribute per LSP	5
4.1. RRO Hop Attributes subobject	5
4.2. Procedures	6
4.2.1. Subobject presence rule	6
4.2.2. Reporting Compliance with ERO Hop Attributes	6
5. IANA Considerations	7
6. Security Considerations	7
7. Acknowledgments	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	8

1. Introduction

Generalized MPLS (GMPLS) Traffic Engineering (TE) Label Switched Paths (LSPs) can be route-constrained by making use of the Explicit Route (ERO) object and related sub-objects as defined in [RFC3209], [RFC3473], [RFC3477], [RFC4873], [RFC4874], [RFC5520] and [RFC5553]. Those route constraints are extended by a number of documents, including element defined in [RFC6163], [I-D.ietf-ccamp-wson-signaling], [I-D.dong-ccamp-rsvp-te-mpls-tp-li-lb] or [I-D.ali-ccamp-rc-objective-function-metric-bound], for example the WSON_SIGNALING object, Metric and Objective Function subobjects.

RSVP already supports generic extension of LSP attributes in [RFC5420]. In order to support current and future ERO constraint extensions this document defines a mechanism to target generic attributes at a specific hop.

1.1. Contributing Authors

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 [RFC2119].

2. Requirements

The requirement is to provide a generic mechanism to carry information related to specific nodes when signaling an LSP. This document does not restrict what that information can be used for. A mechanism similar to LSP attribute defined [RFC5420] should be expressed in ERO and SERO objects. A new ERO sub-object is defined, containing a list of generic Hop attributes. The mechanism defined in this document limits itself to single HOP attributes, and does not address attributes valid for a LSP section [[This can be revised based on feedback]]

3. ERO Attribute

The ERO Attributes subobject may be carried in the ERO or SERO object if they are present. The subobject uses the standard format of an ERO subobject.

3.1. ERO_ATTRIBUTE subobject

The length is variable and content MUST be the same as for the LSP_ATTRIBUTE object with Attributes TLVs. The size of the ERO sub-object limits the size of the LSP Attribute TLV to 250 bytes. The typical size of currently defined and forthcoming LSP_ATTRIBUTE TLVs applicable to a specific hop (WSON_SIGNALING, OF and Metric) is not foreseen to exceed this limit.

The ERO attribute subobject is defined as follows:

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|   Type   |   Length   |   Reserved   |R|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                         |
//                               Attributes TLVs                               //
|                                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

See [RFC3209] for a description of L parameters. The attributes TLV are encoded as defined in section Section 3.2.

Type x TBD by IANA.

Length The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length MUST be always divisible by 4.

Reserved Reserved, must be set to 0 when the subobject is inserted in the ERO, MUST NOT be changed when a node process the ERO and must be ignored on the node addressed by the preceding ERO subobjects.


R This bit reflects the LSP_REQUIRED_ATTRIBUTE and LSP_ATTRIBUTE semantic defined in [RFC5420]. When set indicates required LSP attributes to be processed by the node, when cleared the LSP attributes are not required as described in Section 3.3.

Attributes TLVs as defined in Section 3.2 .

3.2. HOP Attributes TLVs

ERO Attributes carried by the new objects defined in this document are encoded within TLVs. One or more TLVs may be present in each object. There are no ordering rules for TLVs, and no interpretation should be placed on the order in which TLVs are received.

Each TLV is encoded as follows.



Type The identifier of the TLV..

Length Indicates the total length of the TLV in octets. That is, the combined length of the Type, Length, and Value fields, i.e., four plus the length of the Value field in octets. The entire TLV MUST be padded with between zero and three trailing zeros to make

it four-octet aligned. The Length field does not count any padding.

Value The data carried in the TLV.

3.3. Procedures

As described in [RFC3209] and [RFC3473] the ERO is managed as a list where each hop information starts with a subobject identifying an abstract node or link. The ERO attribute subobject must be appended after the existing subobjects defined in [RFC3209], [RFC3473], [RFC3477], [RFC4873], [RFC4874], [RFC5520] and [RFC5553]. Several ERO attribute subobject MAY be present, for each hop.

If a node is processing an ERO attribute subobject and does not support handling of the subobject it will behave as described in [RFC3209] when an unrecognized ERO subobject is encountered. This node will return a PathErr with error code "Routing Error" and error value "Bad EXPLICIT_ROUTE object" with the EXPLICIT_ROUTE object included, truncated (on the left) to the offending unrecognized subobject.

When the R bit is set a node MUST examine the attribute TLV present in the subobject following the rules described in [RFC5420] section 5.2. When the R bit is not set a node MUST examine the attribute TLV present in the subobject following the rules described in [RFC5420] section 4.2.

A node processing an ERO attribute subobject with an HOP attribute TLV longer than the ERO subobject SHOULD return a PathErr with error code "Routing Error" and error value "Bad EXPLICIT_ROUTE object" with the EXPLICIT_ROUTE object included, truncated (on the left) to the offending malformed subobject. The processing of the Hop attribute TLVs should be described in the documents defining them.

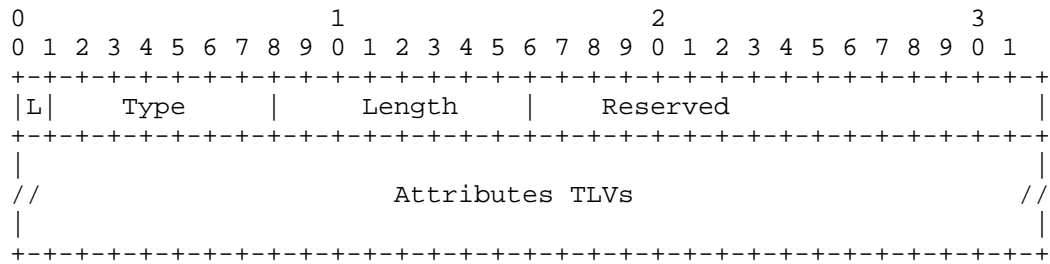
4. Recording Hop attribute per LSP

In some cases it is important to determine if an optional Hop attribute has been processed by a node.

4.1. RRO Hop Attributes subobject

The RRO Hop Attributes subobject may be carried in the RECORD_ROUTE object if it is present. The subobject uses the standard format of an RRO subobject.

The RRO Hop attribute subobject is defined as follows:



See [RFC3209] for a description of L parameters. The attributes TLV are encoded as defined in section Section 3.2.

Type x TBD by IANA.

Length The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length MUST be always divisible by 4.

Reserved Reserved, must be set to 0 when the subobject is inserted in the RRO, MUST NOT be changed when a node process the RRO and must be ignored on the node addressed by the preceding RRO subobjects.

Attributes TLVs The processed or addition HOP attributes, using the format defined in Section 3.2 .

4.2. Procedures

4.2.1. Subobject presence rule

The RRO rules defined in [RFC3209] are not changed. The RRO Hop attribute subobject must be pushed after the RRO attribute subobject (if present) defined in in [RFC5420]. The RRO Hop attribute subobject MAX be present between a pair of subobject identifying LSR or links. All such subobjects MUST be forwarded unmodified by transit LSRs.

4.2.2. Reporting Compliance with ERO Hop Attributes

To report that an ERO Hop attribute has been considered, or to report an additional attribute, an LSR MAY add a RRO Hop Attributes subobject with the HOP Attribute used. The requirement to report compliance MUST be specified in the document that defines the usage of an Hop attribute. [[This is not the most efficient encoding, a more efficient encoding would use a bit field ala RFC5420]]

5. IANA Considerations

TBD once a final approach has been chosen.

6. Security Considerations

None.

7. Acknowledgments

The authors would like to thanks Lou Berger for his directions and Attila Takacs for inspiring this [I-D.kern-ccamp-rsvpte-hop-attributes]. The authors also thanks Dirk Schroetter for his contribution to the initial versions of the documents (version -00 up to -02).

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

- [RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.
- [RFC5553] Farrel, A., Bradford, R., and JP. Vasseur, "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC 5553, May 2009.

8.2. Informative References

- [I-D.ali-ccamp-rc-objective-function-metric-bound]
Ali, Z., Swallow, G., Filsfils, C., Fang, L., Kumaki, K., and R. Kunze, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) extension for signaling Objective Function and Metric Bound", draft-ali-ccamp-rc-objective-function-metric-bound-02 (work in progress), July 2012.
- [I-D.dong-ccamp-rsvp-te-mpls-tp-li-lb]
Dong, J., Chen, M., and Z. Li, "GMPLS RSVP-TE Extensions for Lock Instruct and Loopback", draft-dong-ccamp-rsvp-te-mpls-tp-li-lb-05 (work in progress), December 2012.
- [I-D.ietf-ccamp-wson-signaling]
Bernstein, G., Xu, S., Lee, Y., Martinelli, G., and H. Harai, "Signaling Extensions for Wavelength Switched Optical Networks", draft-ietf-ccamp-wson-signaling-06 (work in progress), July 2013.
- [I-D.kern-ccamp-rsvpte-hop-attributes]
Kern, A. and A. Takacs, "Encoding of Attributes of LSP intermediate hops using RSVP-TE", draft-kern-ccamp-rsvpte-hop-attributes-00 (work in progress), October 2009.
- [RFC6163] Lee, Y., Bernstein, G., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSOs)", RFC 6163, April 2011.

Authors' Addresses

Cyril Margaria (editor)
Coriant GmbH
St Martin Strasse 76
Munich 81541
Germany

Phone: +49 89 5159 16934
Email: cyril.margaria@coriant.com

Giovanni Martinelli
Cisco
via Philips 12
Monza 20900
IT

Phone: +39 039 209 2044
Email: giomarti@cisco.com

Steve Balls
Metaswitch
100 Church Street
Enfield EN2 6BQ
UJ

Phone: +44 208 366 1177
Email: steve.balls@metaswitch.com

Ben Wright
Metaswitch
100 Church Street
Enfield EN2 6BQ
UJ

Phone: +44 208 366 1177
Email: Ben.Wright@metaswitch.com

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 14, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Ori Gerstel
Gabriele Maria Galimberti
Cisco Systems
Kenji Kumaki
KDDI Corporation
Ruediger Kunze
Deutsche Telekom AG
Julien Meuric
France Telecom Orange
July 15, 2013

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Path
Diversity using Exclude Routes

draft-ietf-ccamp-lsp-diversity-02.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 13, 2014.

Copyright Notice

Copyright (c) 2013 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.

Ali, Swallow, Filsfils

Expires August 2013

[Page 1]

Abstract

RFC 4874 specifies methods by which route exclusions may be communicated during RSVP-TE signaling in networks where precise explicit paths are not computed by the LSP source node. This document specifies signaling for additional route exclusions based on Paths currently existing or expected to exist within the network.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction.....	2
2. RSVP-TE signaling extensions.....	4
2.1. Terminology.....	5
2.2. Path XRO Subobjects.....	5
2.2.1. IPv4 Point-to-Point Path subobject.....	5
2.2.2. IPv6 Point-to-Point Path subobject.....	8
2.3. Processing rules for the Path XRO subobjects.....	9
2.4. Path EXRS Subobject.....	12
3. Security Considerations.....	13
4. IANA Considerations.....	13
4.1. New XRO subobject types.....	13
4.2. New EXRS subobject types.....	14
4.3. New RSVP error sub-codes.....	14
5. Acknowledgements.....	14
6. References.....	14
6.1. Normative References.....	14
6.2. Informative References.....	15

1. Introduction

Path diversity is a well-known requirement from Service Providers. Such diversity is required to ensure Label-Switched Paths (LSPs) may be established without sharing resources, thus greatly reducing the probability of simultaneous connection failures.

When route computation for paths that need to be diverse is performed at the LSP's source node, this requirement can be met by a local decision at that node. However, there are scenarios when route computations are performed by remote nodes, thus there is a need for relevant diversity requirements to be communicated to those nodes. These include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), e.g. inter-domain LSPs;
- . Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI) where route computation may be performed by the (server layer) core node [RFC4208].

[RFC4874] introduced a means of specifying nodes and resources to be excluded from a route, using the eXclude Route Object (XRO) and Explicit Exclusion Route Subobject (EXRS).

[RFC4874] facilitates the calculation of diverse routes for LSPs based on known properties of those paths including addresses of links and nodes traversed, and Shared Risk Link Groups (SRLGs) of traversed links. This requires that these properties of the path(s) from which diversity is required be known to the source node which initiates signaling. However, there are circumstances under which this may not be possible or desirable, including (but not limited to):

- . Exclusion of a path which does not originate, terminate or traverse the source node signaling the diverse LSP, in which case the addresses and SRLGs of the path from which diversity is required are unknown to the source node.
- . Exclusion of a path which, while known at the source node of the diverse LSP, has incomplete or unavailable route information, e.g. due to confidentiality of the path attributes. In other words, the scenario in which the reference path is hosted by the source / requesting node but the properties required to construct an XRO object are not known to source / requesting node. Inter-domain and GMPLS overlay networks may present such restrictions.
- . If the source node knows the route of the reference path from which diversity is required, it can use this information to construct an XRO and send it in the path message during the signaling of a diverse LSP. However, if the route of the excluded path changes (e.g. due to re-optimization or failure in the network), the source node would need to change the

diverse path to ensure that it remains diverse from the excluded path. It is preferable to have this decision made by the node that performed the path-calculation for the diverse path. For example, in the case of GMPLS-UNI, it is better to have such responsibility at the server layer as opposed to at the client layer so that the diversity requirements are transparent to the client layer. Furthermore, in all networking scenarios, if the node performing the route computation/expansion is aware of the diversity requirements of the two paths, it may consider joint re-optimization of the diverse paths.

This document addresses such scenarios and defines procedures that may be used to exclude the route taken by a particular LSP, or the routes taken by all LSPs belonging to a single tunnel. Note that this diversity requirement is different from the diversity requirements of path protection where both the reference and diverse LSPs belong to the same tunnel. The diversity requirements considered in this document do not require that the paths in question belonging to the same tunnel or share the same source or destination node.

The means by which the node calculating or expanding the route of the signaled LSP discovers the route of the path(s) from which the signaled LSP requires diversity are beyond the scope of this document.

This document addresses only the exclusion of point-to-point paths; point-to-multipoint paths will be addressed in a future version.

If mutually diverse routes are desired for two LSPs belonging to different tunnels, it is recommended that they be signaled with XRO LSP subobjects referencing each other. The processing rules specified in this document cover this case.

2. RSVP-TE signaling extensions

This section describes the signaling extensions required to address the aforementioned requirements. Specifically, this document defines a new LSP subobject to be signaled in the EXCLUDE_ROUTE object (XRO) and/ or Explicit Exclusion Route Subobject (EXRS) defined in [RFC4874]. Inclusion of the LSP subobject in any other RSVP object is not defined.

2.1. Terminology

In this document, the following terminology is adopted:

Excluded path: the path from which diversity is required.

Diverse LSP: the LSP being signaled with XRO/ EXRS containing the path subobject referencing the excluded path(s).

Processing node: the node performing a path-calculation involving an exclusion specified in an XRO or EXRS.

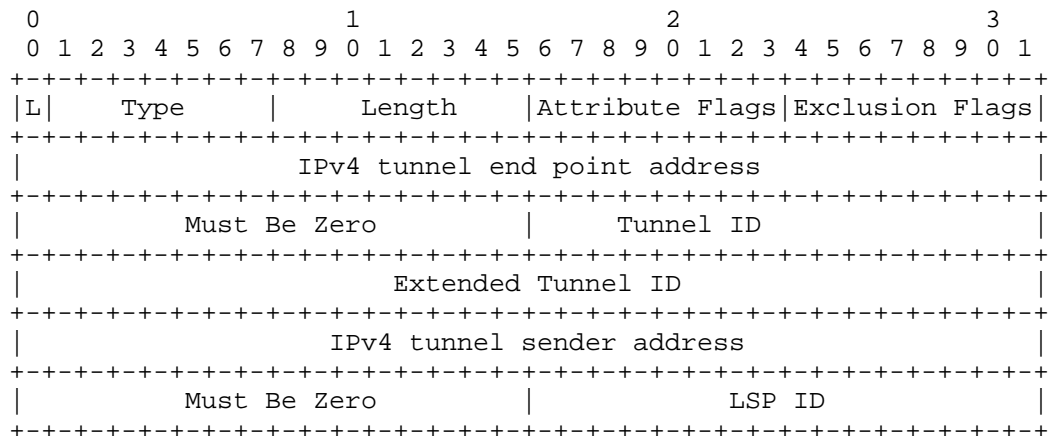
Destination node: in the context of an XRO, this is the destination of the LSP being signaled. In the context of an EXRS, the destination node is the last explicit node to which the loose hop is expanded.

Penultimate node: in the context of an XRO, this is the penultimate hop of the LSP being signaled. In the context of an EXRS, the penultimate node is the penultimate node of the loose hop undergoing expansion.

2.2. Path XRO Subobjects

New IPv4 and IPv6 Point-to-Point (P2P) Path XRO subobjects are defined by this document as follows.

2.2.1. IPv4 Point-to-Point Path subobject



L

The L-flag is used as for the other XRO subobjects defined in [RFC4874].

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

Type

IPv4 Point-to-Point Path subobject
(to be assigned by IANA; suggested value: 36).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 24.

Attribute Flags

The Attribute Flags are used to communicate desirable attributes of the LSP being signaled. The following flags are defined. None, all or multiple attribute flags MAY be set within the same subobject.

0x01 = LSP ID to be ignored

This flag is used to indicate tunnel level exclusion. Specifically, this flag is used to indicate that the lsp-id field of the subobject is to be ignored and the exclusion applies to any LSP matching the rest of the supplied FEC.

0x02 = Destination node exception

This flag is used to indicate that the destination node of the LSP being signaled MAY be shared with the excluded path even when this violates the exclusion flags.

0x04 = Processing node exception

This flag is used to indicate that the processing node MAY be shared with the excluded path even when this violates the exclusion flags.

0x08 = Penultimate node exception

This flag is used to indicate that the penultimate node of the LSP being signaled MAY be shared with the excluded path even when this violates the exclusion flags.

Exclusion Flags

The Exclusion-Flags are used to communicate desirable types of exclusion. The following flags are defined.

0x01 = SRLG exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be SRLG diverse from the excluded path specified by the LSP subobject.

0x02 = Node exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be node diverse from the excluded path specified by the LSP subobject.

(Note: the meaning of this flag may be modified by the value of the Attribute-flags.)

0x04 = Link exclusion

This flag is used to indicate that the route of the LSP being signaled is requested to be link diverse from the path specified by the LSP subobject.

The remaining fields are as defined in [RFC3209].

1 indicates that the attribute specified SHOULD be avoided.

Type

IPv6 Point-to-Point Path subobject
(to be assigned by IANA; suggested value: 37).

Length

The length contains the total length of the subobject in bytes, including the type and length fields. The length is always 48.

The Attribute Flags and Exclusion Flags are as defined for the IPv4 Point-to-Point LSP XRO subobject.

The remaining fields are as defined in [RFC3209].

2.3. Processing rules for the Path XRO subobjects

XRO processing as described in [RFC4874] is unchanged.

If the processing node is the destination for the LSP being signaled, it SHOULD NOT process a Path XRO subobject.

If the L-flag is not set, the processing node follows the following procedure:

- The processing node MUST ensure that any route calculated for the signaled LSP respects the requested exclusion flags with respect to the excluded path referenced by the subobject, including local resources.
- If the processing node fails to find a route that meets the requested constraint, the processing node MUST return a PathErr with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67).
- If the excluded path referenced in the LSP subobject is unknown to the processing node, the processing node SHOULD ignore the LSP subobject in the XRO and SHOULD proceed with the signaling request. After sending the Resv for the signaled LSP, the

processing node SHOULD return a PathErr with the error code "Notify Error" (25) and error sub-code "Route of XRO path unknown" (value to be assigned by IANA, suggested value: 13) for the signaled LSP.

If the L-flag is set, the processing node follows the following procedure:

- The processing node SHOULD respect the requested exclusion flags with respect to the excluded path as far as possible.
- If the processing node fails to find a route that meets the requested constraint, it SHOULD proceed with signaling using a suitable route that meets the constraint as far as possible. After sending the Resv for the signaled LSP, it SHOULD return a PathErr message with error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14) to the source node.
- If the excluded path referenced in the LSP subobject is unknown to the processing node, the processing node SHOULD ignore the LSP subobject in the XRO and SHOULD proceed with the signaling request. After sending the Resv for signaled LSP, the processing node SHOULD return a PathErr message with the error code "Notify Error" (25) and error sub-code "Route of XRO path unknown" for the signaled LSP.

If, subsequent to the initial signaling of a diverse LSP:

- an excluded path referenced in the diverse LSP's XRO subobject becomes known to the processing node (e.g. when the excluded path is signaled), or
- A change in the excluded path becomes known to the processing node,

the processing node SHOULD re-evaluate the exclusion and diversity constraints requested by the diverse LSP to determine whether they are still satisfied.

- If the requested exclusion constraints for the diverse LSP are no longer satisfied and an alternative route for the diverse LSP that can satisfy those constraints exists, the processing node SHOULD send a PathErr message for the diverse LSP with the error code "Notify Error" (25) and error sub-code "Preferable path exists" (6). A source node receiving a PathErr message

with this error code and sub-code combination MAY try to reoptimize the diverse tunnel to the new compliant path.

- If the requested exclusion constraints for the diverse LSP are no longer satisfied and no alternative path for the diverse LSP that can satisfy those constraints exists, then:
 - o If the L-flag was not set in the original exclusion, the processing node MUST send a PathErr message for the diverse LSP with the error code "Routing Problem" (24) and error sub-code "Route blocked by Exclude Route" (67). The PSR flag SHOULD NOT be set.
 - o If the L-flag was set in the original exclusion, the processing node SHOULD send a PathErr message for the diverse LSP with the error code error code "Notify Error" (25) and error sub-code "Failed to respect Exclude Route" (value: to be assigned by IANA, suggest value: 14).

The following rules apply whether or not the L-flag is set:

- An XRO object MAY contain multiple path subobjects.
- As specified in [RFC4874], a node receiving a Path message carrying an XRO MAY reject the message if the XRO is too large or complicated for the local implementation or the rules of local policy. In this case, the node MUST send a PathErr message with the error code "Routing Error" (24) and error sub-code "XRO Too Complex" (68). A source node receiving this error code/sub-code combination MAY reduce the complexity of the XRO or route around the node that rejected the XRO.
- A source node receiving a PathErr message with the error code "Notify Error" (25) and error sub-codes "Route of XRO path unknown" or "Failed to respect Exclude Route" MAY take no action.
- The attribute-flags affect the processing of the XRO subobject as follows:
 - o When the "LSP ID to be ignored" flag is set, the processing node MUST calculate a route based on exclusions from the routes of all known LSPs matching the tunnel-id, source, destination and extended tunnel-id specified in the subobject. When this flag is not set, the lsp-id is not ignored and the exclusion applies only to the specified LSP (i.e., LSP level exclusion).

- o When the "destination node exception" flag is not set, the exclusion flags SHOULD also be respected for the destination node.
- o When the "processing node exception" flag is not set, the exclusion flags SHOULD also be respected for the processing node.
- o When the "penultimate node exception" flag is not set, the exclusion flags SHOULD also be respected for the penultimate node.

2.4. Path EXRS Subobject

[RFC4874] defines the EXRS ERO subobject. An EXRS is used to identify abstract nodes or resources that must not or should not be used on the path between two inclusive abstract nodes or resources in the explicit route. An EXRS contains one or more subobjects of its own, called EXRS subobjects [RFC4874].

An EXRS MAY include an IPv4 Point-to-Point (P2P) Path subobject as specified in section 2.2.1. In this case, the EXRS format would be as follows:

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	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																																												
L Type Length Reserved																L Type Length Attribute Flags Exclusion Flags																IPv4 tunnel end point address																Must Be Zero Tunnel ID																Extended Tunnel ID																IPv4 tunnel sender address																Must Be Zero LSP ID															

The meaning of respective fields in EXRS header are as defined in [RFC4874]. The meaning of respective fields in IPv4 P2P Path subobject is as defined earlier in this document.

The processing rules for the EXRS object are unchanged from [RFC4874]. When the EXRS contains one or more Path subobject(s), the processing rules specified in Section 2.3 apply to the node processing the ERO with the EXRS subobject.

If a loose-hop expansion results in the creation of another loose-hop in the outgoing ERO, the processing node MAY include the EXRS in the newly-created loose hop for further processing by downstream nodes.

The processing node exception for the EXRS subobject applies to the node processing the ERO.

The destination node exception for the EXRS subobject applies to the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

The penultimate node exception for the EXRS subobject applies to the node before the explicit node identified by the ERO subobject that identifies the next abstract node. This flag is only processed if the L bit is set in the ERO subobject that identifies the next abstract node.

3. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC2205], [RFC3209], [RFC3473] and [RFC4874].

4. IANA Considerations

4.1. New XRO subobject types

IANA registry: RSVP PARAMETERS

Subsection: Class Names, Class Numbers, and Class Types

This document introduces two new subobjects for the EXCLUDE_ROUTE object [RFC4874], C-Type 1.

Subobject Type	Subobject Description
-----	-----
To be assigned by IANA (suggested value: 36)	IPv4 P2P Path subobject

- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

6.2. Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC2205] Braden, R. (Ed.), Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReserVation Protocol -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

Ori Gerstel
Cisco Systems
ogerstel@cisco.com

Gabriele Maria Galimberti
Cisco Systems
ggalimbe@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Ruediger.Kunze@telekom.de

Julien Meuric
France Telecom Orange
Email: julien.meuric@orange.com

CCAMP Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 19, 2013

F. Zhang, Ed.
ZTE
R. Jing
China Telecom
R. Gandhi
Cisco Systems
March 18, 2013

RSVP-TE Extensions for Associated Bidirectional LSPs
draft-ietf-ccamp-mpls-tp-rsvpte-ext-associated-lsp-06

Abstract

The MPLS Transport Profile (MPLS-TP) requirements document [RFC5654], describes that MPLS-TP MUST support associated bidirectional point-to-point LSPs.

This document provides a method to bind two unidirectional Label Switched Paths (LSPs) into an associated bidirectional LSP. The association is achieved by defining the new Association Types in the Extended ASSOCIATION object.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 19, 2013.

Copyright Notice

Copyright (c) 2013 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
2. Conventions used in this document	3
3. Overview	4
3.1. Provisioning Model	4
3.2. Signaling Procedure	4
3.2.1. Single Sided Provisioning Model	5
3.2.2. Double Sided Provisioning Model	5
3.2.3. Asymmetric Bandwidth LSPs	5
3.2.4. Recovery Considerations	6
3.2.5. Associated Bidirectional LSPs and LSP Recovery	6
3.2.6. Associated Bidirectional LSPs and TE Mesh-Groups	7
3.2.7. MPLS-TP Associated Bidirectional LSP Identifiers	7
3.2.8. Teardown of Associated Bidirectional LSPs	7
4. Association of LSPs	8
4.1. Extended ASSOCIATION Object	8
4.1.1 Signaling of the Extended Association Object	9
4.2. REVERSE_LSP Object	9
4.2.1. Format	9
4.2.1.1. Subobjects	10
4.2.2. LSP Control	10
4.2.3. Updated RSVP Message Formats	10
4.2.4. Compatibility	11
5. IANA Considerations	11
5.1. Association Type	11
5.2. REVERSE_LSP Object	12
6. Security Considerations	12
7. Acknowledgement	12
8. References	13
8.1. Normative references	13
8.2. Informative References	14
Authors' Addresses	15

1. Introduction

The MPLS Transport Profile (MPLS-TP) requirements document [RFC5654] describes that MPLS-TP MUST support associated bidirectional point-to-point LSPs. Furthermore, an associated bidirectional LSP is useful for protection switching, for Operations, Administrations and Maintenance (OAM) messages that require a reply path.

The requirements described in [RFC5654] are specifically mentioned in Section 2.1. (General Requirements), and are repeated below:

7. MPLS-TP MUST support associated bidirectional point-to-point LSPs.

11. The end points of an associated bidirectional LSP MUST be aware of the pairing relationship of the forward and reverse LSPs used to support the bidirectional service.

12. Nodes on the LSP of an associated bidirectional LSP where both the forward and backward directions transit the same node in the same (sub)layer as the LSP SHOULD be aware of the pairing relationship of the forward and the backward directions of the LSP.

14. MPLS-TP MUST support bidirectional LSPs with asymmetric bandwidth requirements, i.e., the amount of reserved bandwidth differs between the forward and backward directions.

50. The MPLS-TP control plane MUST support establishing associated bidirectional P2P LSP including configuration of protection functions and any associated maintenance functions.

The above requirements are also repeated in [RFC6373].

The notion of association, as well as the corresponding Resource reSerVation Protocol (RSVP) ASSOCIATION object, is defined in [RFC4872], [RFC4873] and [RFC6689]. In that context, the object is used to associate recovery LSPs with the LSP they are protecting. This object also has broader applicability as a mechanism to associate RSVP state, and [RFC6780] defines the Extended ASSOCIATION object that can be more generally applied.

This document provides a method to bind two reverse unidirectional Label Switched Paths (LSPs) into an associated bidirectional LSP. The association is achieved by defining the new Association Types in the Extended ASSOCIATION object.

2. 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 [RFC2119].

3. Overview

3.1. Provisioning Model

The associated bidirectional LSP's forward and backward directions are set up, monitored, and protected independently as required by [RFC5654]. Configuration information regarding the LSPs can be sent to one end or both ends of the LSP. Depending on the method chosen, there are two models of signaling associated bidirectional LSP. The first model is the single sided provisioning, the second model is the double sided provisioning.

For the single sided provisioning, the configurations are sent to one end. Firstly, a unidirectional tunnel is configured on this end, then a LSP under this tunnel is initiated with the Extended ASSOCIATION object carried in the Path message to trigger the peer end to set up the corresponding reverse TE tunnel and LSP.

For the double sided provisioning, the two unidirectional TE tunnels are configured independently, then the LSPs under the tunnels are signaled with the Extended ASSOCIATION objects carried in the Path message to indicate each other to associate the two LSPs together to be an associated bidirectional LSP.

A number of scenarios exist for binding LSPs together to be an associated bidirectional LSP. These include: (1) both of them do not exist; (2) both of them exist; (3) one LSP exists, but the other one need to be established. In all scenarios described, the provisioning models discussed above are applicable.

3.2. Signaling Procedure

This section describes the signaling procedures for associating bidirectional LSPs.

Consider the topology described in Figure 1. (An example of associated bidirectional LSP). The LSP1 [via nodes A,D,B] (from A to B) and LSP2 [via nodes B,D,C,A] (from B to A) are being established or have been established, which can form an associated bidirectional LSP between node A and node B.

LSP1 and LSP2 are referenced at the data plane level by the

identifiers: A-Node_ID::A-Tunnel_Num::A-LSP_Num::B-Node_ID and
B-Node_ID::B-Tunnel_Num::B-LSP_Num::A-Node_ID, respectively
[RFC6370].




Figure 1: An example of associated bidirectional LSP

3.2.1. Single Sided Provisioning Model

For the single sided provisioning model, LSP1 is triggered by LSP2 or LSP2 is triggered by LSP1. When LSP2 is triggered by LSP1, LSP1 is initialized or refreshed (if LSP1 already exists) at node A with the Extended ASSOCIATION object inserted in the Path message, the Association Type must be set to "Single Sided Associated Bidirectional LSPs". Terminating node B is triggered to set up LSP2 by the received Extended ASSOCIATION object with the Association Type set to the value "Single Sided Associated Bidirectional LSPs", the Extended ASSOCIATION object inserted in LSP2's Path message is the same as in LSP1's Path message.

When LSP1 is triggered by LSP2, the same rules are applicable. Based on the same values of the Extended ASSOCIATION objects in the two LSPs' Path messages, the two LSPs can be bound together to be an associated bidirectional LSP.

3.2.2. Double Sided Provisioning Model

For the double sided provisioning model, the Association Type must be set to "Double Sided Associated Bidirectional LSPs".

Identification of the LSPs as being Associated Bidirectional LSPs occurs based on the identical contents in the LSPs' Extended ASSOCIATION objects.

3.2.3. Asymmetric Bandwidth LSPs

A variety of applications, such as Internet services and the return paths of OAM messages, exist and which MAY have different bandwidth requirements for each direction. Additional [RFC5654] also specifies an asymmetric bandwidth requirement. This requirement is

specifically mentioned in Section 2.1. (General Requirements), and is repeated below:

14. MPLS-TP MUST support bidirectional LSPs with asymmetric bandwidth requirements, i.e., the amount of reserved bandwidth differs between the forward and backward directions.

The approach for supporting asymmetric bandwidth co-routed bidirectional LSPs is defined in [RFC6387]. As to the asymmetric bandwidth associated bidirectional LSPs, the existing SENDER_TSPEC object must be carried in the REVERSE_LSP object as a sub-object in the initialized LSP's Path message to specify the reverse LSP's traffic parameters in case that single sided provisioning model is adopted. Consider the topology described in Figure 1 in the context of asymmetric associated bidirectional LSP, and take LSP2 triggered by LSP1 as an example. Node B is triggered to set up the reverse LSP2 with the corresponding asymmetric bandwidth by the Extended ASSOCIATION object with Association Type "Single Sided Associated Bidirectional LSPs" and the SENDER_TSPEC sub-object in LSP1's Path message, and the SENDER_TSPEC object in the LSP2' Path message is the same as the the SENDER_TSPEC sub-object in LSP1's Path message. When double sided provisioning model is used, the two opposite LSPs with asymmetric bandwidths are concurrently initialized, and this requirement will be satisfied simultaneously.

3.2.4. Recovery Considerations

Consider the topology described in Figure 1, LSP1 and LSP2 form the associated bidirectional LSP. Under the scenario of recovery, a third LSP (LSP3) may be used to protect LSP1. LSP3 can be established before or after the failure occurs, it can share the same TE tunnel with LSP1.

When node A detects that LSP1 is broken or needs to be reoptimized, LSP3 will be initialized or refreshed with the Extended ASSOCIATION object inherited from LSP1's Path message. Furthermore, if LSP3 is the protecting LSP [RFC4872], the ASSOCIATION object and PROTECTION object [RFC4872] need to be inherited from the LSP1 also. In this way, based on the same Extended ASSOCIATION object, LSP2 and LSP3 will compose the new associated bidirectional LSPs.

3.2.5. Associated Bidirectional LSPs and LSP Recovery

LSP recovery as defined in [RFC4872], [RFC4873] and [RFC4090] is not impacted by this document. The recovery mechanisms defined in [RFC4872] and [RFC4873] rely on the use of ASSOCIATION Objects, but use a different Association Type field value than defined in this

document so should not be impacted. The mechanisms defined in [RFC4090] does not rely on the use of ASSOCIATION Objects and is therefore also not impacted by the mechanisms defined in this document.

3.2.6. Associated Bidirectional LSPs and TE Mesh-Groups

TE mesh-groups is defined in [RFC4972]. A node supporting both Associated Bidirectional LSPs and TE mesh-groups, MAY include an ASSOCIATION object as defined in this document in Path messages of LSPs used to support the mesh-group. To enable unambiguous identification of the mesh-group's associated bidirectional LSPs, the information carried in the ASSOCIATION object, including the contents of the Association Source and Identifier fields MUST be provisioned.

3.2.7. MPLS-TP Associated Bidirectional LSP Identifiers

[RFC6370] defines the LSP associated identifiers based on the signaling parameters of each unidirectional LSP. The combination of each unidirectional LSP's parameters is used to identify the Associated Bidirectional LSP. Using the mechanisms defined in this document, any node that is along the path of both unidirectional LSPs can identify which pair of unidirectional LSPs support an Associated Bidirectional LSP as well as the signaling parameters required by [RFC6370]. Note that the LSP end-points will always be the path of both unidirectional LSPs.

3.2.8. Teardown of Associated Bidirectional LSPs

Associated bidirectional LSPs teardown also follows standard procedures defined in [RFC3209] and [RFC3473] either without or with the administrative status. Note that teardown procedures of the associated bidirectional LSPs are independent of each other, so it is possible that while one LSP1 follows graceful teardown with administrative status, the other LSP2 is torn down without administrative status (using PathTear/ResvTear/PathErr with state removal). However, for the double sided associated bidirectional LSPs, the teardown of LSP1 does not mean that LSP2 must be deleted, which depends on the local policy. While for the single sided associated bidirectional LSPs, the teardown of the initialized LSP should induce the teardown of the trigger-established LSP, but the teardown of the trigger-established LSP (using PathErr with state removal) should not induce the teardown of the initialized LSP (which depends on the local policy).

4. Association of LSPs

4.1. Extended ASSOCIATION Object

The Extended ASSOCIATION object is defined in [RFC6780], which enables MPLS-TP required LSP identification. The Extended ASSOCIATION object is used as follows for associated bidirectional LSPs.

Association Types:

In order to bind two reverse unidirectional LSPs to be an associated bidirectional LSP, new Association Types are defined in this document:

Value	Type
-----	-----
4 (TBD)	Double Sided Associated Bidirectional LSPs (D)
5 (TBD)	Single Sided Associated Bidirectional LSPs (A)

Association ID: 16 bits

For both double sided and single sided provisioning, Association ID is a value assigned by the node that originates the association.

Association Source: 4 or 16 bytes

Same as for IPv4 and IPv6 ASSOCIATION objects, see [RFC4872].

For double sided provisioning, Association Source is set to an address selected by the node that originates the association (which may be a management entity.)

For single sided provisioning, Association Source is set to an address assigned to the node that originates the LSP.

Global Association Source: 4 bytes

Same as for IPv4 and IPv6 Extended ASSOCIATION objects defined in [RFC6780].

For both double sided and single sided provisioning, Global Association Source, when used, is set to the Global_ID [RFC6370] of the node that originates association.

Extended Association ID: 4 or 16 bytes

This field contains data that is additional information to support

unique identification.

For both double sided and single sided provisioning, Extended Association ID, when used, is selected by the node that originates the association.

If either Global Association Source or Extended Association Address is required, an Extended ASSOCIATION object [RFC6780] is used. Otherwise an ASSOCIATION object [RFC4872] is used.

4.1.1 Signaling of the Extended Association Object

As described in [RFC6780], association is always done based on matching Path state or Resv state. Upstream initialized association is represented in Extended ASSOCIATION objects carried in Path message and downstream initialized association is represented in Extended ASSOCIATION objects carried in Resv messages. The new Association Types defined in this document are only used in upstream initialized association. Thus they can only appear in Extended ASSOCIATION objects signaled in Path message.

The rules associated with the processing of the Extended ASSOCIATION objects in RSVP message are discussed in [RFC6780]. It said that in the absence of Association Type-specific rules for identifying association, the included Extended ASSOCIATION objects MUST be identical. This document adds no specific rules, the association will always operate based on the same Extended ASSOCIATION objects.

4.2. REVERSE_LSP Object

Path Computation Element (PCE)-based approaches, see [RFC4655], may be used for path computation of a GMPLS LSP, and consequently an associated bidirectional LSP, across domains and in a single domain. The ingress Label Switching Router (LSR), maybe serve as a PCE or Path Computation Client (PCC), has more information about the reverse LSP. When the forward LSP is signaled, the reverse LSP's traffic parameters, explicit route, LSP attributes, etc, can be carried in the REVERSE_LSP object of the forward LSP's Path message. The egress LSR can be triggered to establish the reverse LSP according to the control information.

4.2.1. Format

The information of the reverse LSP is specified via the REVERSE_LSP object, which is optional with class numbers in the form 11bbbbbbb has

the following format:

Class = TBD (of the form 11bbbbbb), C_Type = 1 (TBD)

```

      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
      +-----+-----+-----+-----+-----+-----+-----+-----+
      |
      //                               (Subobjects)                               //
      |
      +-----+-----+-----+-----+-----+-----+-----+-----+

```

This object MUST NOT be used when the Extended ASSOCIATION object do not exist or exist but the Association Type is not "Single Sided Associated Bidirectional LSPs".

4.2.1.1. Subobjects

The contents of a REVERSE_LSP object are a series of variable-length data items called subobjects, which can be SENDER_TSPCE, EXPLICIT_ROUTE object (ERO), Session Attribute Object, Admin Status Object, LSP_ATTRIBUTES Object, LSP_REQUIRED_ATTRIBUTES Object, PROTECTION Object, ASSOCIATION Object, Extended ASSOCIATION Objects, etc.

4.2.2. LSP Control

The signaling procedure without the REVERSE_LSP object carried in the LSP1's Path message is described in section 3.2.1, which is the default option. A node includes a REVERSE_LSP object and Extended ASSOCIATION object with an "Single Sided Associated Bidirectional LSPs" Association Type in an outgoing Path message when it wishes to control the reverse LSP, and the receiver node B MUST convert the subobjects of the REVERSE_LSP object into the corresponding objects that carried in LSP2's Path message. The case of a non-supporting egress node is outside of this document. If node A want to tear down the associated bidirectional LSP, a PathTear message will be sent out and Node B is triggered to tear down LSP2.

4.2.3. Updated RSVP Message Formats

This section presents the RSVP message-related formats as modified by this document. Unmodified RSVP message formats are not listed.

The format of a Path message is as follows:

```
<Path Message> ::= <Common Header> [ <INTEGRITY> ]
```

```

[ [ <MESSAGE_ID_ACK> | <MESSAGE_ID_NACK> ] ... ]
[ <MESSAGE_ID> ]
<SESSION> <RSVP_HOP>
<TIME_VALUES>
[ <EXPLICIT_ROUTE> ]
<LABEL_REQUEST>
[ <PROTECTION> ]
[ <LABEL_SET> ... ]
[ <SESSION_ATTRIBUTE> ]
[ <NOTIFY_REQUEST> ... ]
[ <ADMIN_STATUS> ]
[ <EXTENDED_ASSOCIATION> ... ]
[ <REVERSE_LSP> ]
[ <POLICY_DATA> ... ]
<sender descriptor>

```

The format of the <sender descriptor> is not modified by the present document.

4.2.4. Compatibility

The REVERSE_LSP object is defined with class numbers in the form 11bbbbbb, which ensures compatibility with non-supporting nodes. Per [RFC2205], nodes not supporting this extension will ignore the object but forward it, unexamined and unmodified, in all messages resulting from this message. Especially, this object received in PathTear, or PathErr messages should be forwarded immediately in the same message, but should be saved with the corresponding state and forwarded in any refresh message resulting from that state when received in Path message.

5. IANA Considerations

IANA is requested to administer assignment of new values for namespace defined in this document and summarized in this section.

5.1. Association Type

Within the current document, two new Association Types are defined in the Extended ASSOCIATION object.

Value	Type
-----	-----
4 (TBD)	Double Sided Associated Bidirectional LSPs (D)
5 (TBD)	Single Sided Associated Bidirectional LSPs (A)

5.2. REVERSE_LSP Object

A new class named REVERSE_LSP has been created in the 11bbbbbb range (TBD) with the following definition:

Class Types or C-types (1, TBD):

There are no other IANA considerations introduced by this document.

6. Security Considerations

This document introduces two new Association Types, and except this, there are no security issues about the Extended ASSOCIATION object are introduced here.

The procedures defined in this document result in an increase in the amount of topology information carried in signaling messages since the presence of the REVERSE_LSP object necessarily means that there is more information about associated bidirectional LSPs. Thus, in the event of the interception of a signaling message, slightly more could be deduced about the state of the network than was previously the case, but this is judged to be a very minor security risk as this information is already available via routing.

Otherwise, this document introduces no additional security considerations. For a general discussion on MPLS and GMPLS related security issues, see the MPLS/GMPLS security framework [RFC5920].

7. Acknowledgement

The authors would like to thank Lou Berger for his great guidance in this work, George Swallow and Jie Dong for the discussion of recovery, Lamberto Sterling for his valuable comments on the section of asymmetric bandwidths, Daniel King for the review of the document, Attila Takacs for the discussion of the provisioning model. At the same time, the authors would also like to acknowledge the contributions of Bo Wu, Xihua Fu, Lizhong Jin for the initial discussions, and Wenjuan He for the prototype implementation. The authors would also like to thank Siva Sivabalan, Eric Osborne and Robert Sawaya for the discussion on the association object.

8. References

8.1. Normative references

- [RFC6780] Berger, L., Le Faucheur, F., and A. Narayanan, "RSVP Association Object Extensions", RFC 6780, October 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4090] Pan, P., Swallow, G., Atlas, A., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC4872] Lang, J., Rekhter, Y., and D. Papadimitriou, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC4972] Vasseur, JP., Leroux, JL., Yasukawa, S., Previdi, S., Psenak, P., Mabbey, P., "Routing Extensions for Discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) Mesh Membership", RFC 4972, July 2007.

8.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.
- [RFC6373] Andersson, L., Berger, L., Fang, L., Bitar, N., and E. Gray, "MPLS Transport Profile (MPLS-TP) Control Plane Framework", RFC 6373, September 2011.
- [RFC6387] Takacs, A., Berger, L., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", RFC 6387, September 2011.
- [RFC6689] Berger, L., "Usage of The RSVP Association Object", RFC 6689, July 2012.

Authors' Addresses

Fei Zhang (editor)
ZTE

Email: zhang.fei3@zte.com.cn

Ruiquan Jing
China Telecom

Email: jingrq@ctbri.com.cn

Fan Yang
ZTE

Email: yang.fan5@zte.com.cn

Weilian Jiang
ZTE

Email: jiang.weilian@zte.com.cn

Rakesh Gandhi
Cisco Systems

Email: rgandhi@cisco.com

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: January 13, 2014

Zafar Ali
George Swallow
Clarence Filsfils
Matt Hartley
Cisco Systems

Kenji Kumaki
KDDI Corporation

Ruediger Kunze
Deutsche Telekom AG
July 14, 2013

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
extension for recording TE Metric of a Label Switched Path
draft-ietf-ccamp-te-metric-recording-02.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 13, 2014.

Copyright Notice

Copyright (c) 2013 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.

Abstract

There are many scenarios in which Traffic Engineering (TE) metrics such as cost, latency and latency variation associated with a Forwarding Adjacency (FA) or Routing Adjacency (RA) Label Switched Path (LSP) are not available to the ingress and egress nodes. This draft provides extensions for the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) for the support of the discovery of cost, latency and latency variation of an LSP.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

Copyright Notice.....	1
1. Introduction.....	3
2. RSVP-TE Requirement.....	3
2.1. Cost, Latency and Latency Variation Collection Indication....	4
2.2. Cost, Latency and Latency Variation Collection.....	4
2.3. Cost, Latency and Latency Variation Update.....	4
3. RSVP-TE signaling extensions.....	4
3.1. Cost, Latency and Latency Variation Collection Flags.....	4
3.2. Cost Subobject.....	5
3.3. Latency Subobject.....	6
3.4. Latency Variation Subobject.....	7
3.5. Signaling Procedures.....	8
4. Security Considerations.....	9
5. IANA Considerations.....	9
5.1. RSVP Attribute Bit Flags.....	9

Ali, Swallow, Filmsils Expires January 2014 [Page 2]

5.2. New RSVP error sub-code.....	10
6. Acknowledgments.....	11
7. References.....	11
7.1. Normative References.....	11
7.2. Informative References.....	12

1. Introduction

There are many scenarios in packet and optical networks where the route information of an LSP may not be provided to the ingress node for confidentiality reasons and/or the ingress node may not run the same routing instance as the intermediate nodes traversed by the path. In such scenarios, the ingress node cannot determine the cost, latency and latency variation properties of the LSP's route. Similarly, in Generalized Multi-Protocol Label Switching (GMPLS) networks signaling bidirectional LSP, the egress node cannot determine the cost, latency and latency variation properties of the LSP route. A multi-domain or multi-layer network is an example of such networks. Similarly, a GMPLS User-Network Interface (UNI) [RFC4208] is also an example of such networks.

In certain networks, such as financial information networks, network performance information (e.g. latency, latency variation) is becoming as critical to data path selection as other metrics [DRAFT-OSPF-TE-METRIC], [DRAFT-ISIS-TE-METRIC]. If cost, latency or latency variation associated with an FA or an RA LSP is not available to the ingress or egress node, it cannot be advertised as an attribute of the FA or RA. One possible way to address this issue is to configure cost, latency and latency variation values manually. However, in the event of an LSP being rerouted (e.g. due to re-optimization), such configuration information may become invalid. Consequently, in case where that an LSP is advertised as a TE-Link, the ingress and/or egress nodes cannot provide the correct latency, latency variation and cost attribute associated with the TE-Link automatically.

In summary, there is a requirement for the ingress and egress nodes to learn the cost, latency and latency variation attributes of an FA or RA LSP. This draft provides extensions to the Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) for the support of the automatic discovery of these attributes.

2. RSVP-TE Requirement

This section outlines RSVP-TE requirements for the support of the automatic discovery of cost, latency and latency variation attributes of an LSP. These requirements are very similar to the requirement of discovering the Shared Risk Link Groups (SRLGs) associated with the route taken by an LSP [DRAFT-SRLG-RECORDING].

2.1. Cost, Latency and Latency Variation Collection Indication

The ingress node of the LSP must be capable of indicating whether the cost, latency and latency variation attributes of the LSP should be collected during the signaling procedure of setting up the LSP. No cost, latency or latency variation information is collected without an explicit request being made by the ingress node.

2.2. Cost, Latency and Latency Variation Collection

If requested, cost, latency and latency variation is collected during the setup of an LSP. The endpoints of the LSP may use the collected information for routing, flooding and TE link configuration and other purposes.

2.3. Cost, Latency and Latency Variation Update

When the cost, latency and latency variation property of a TE link along the route of a LSP for which that property was collected changes, e.g., if the administrator changes cost of a TE link, the node where the change occurred needs to be capable of updating the cost, latency and latency variation information of the path and signaling this to the end-points. Similarly, if a path segment of the LSP is rerouted, the endpoints of the rerouted segment need to be capable of updating the cost, latency and latency variation information of the path. Any node, which adds cost, latency or latency variation information to an LSP during initial setup, needs to signal changes to these values to both endpoints.

3. RSVP-TE signaling extensions

3.1. Cost, Latency and Latency Variation Collection Flags

Three Attribute flags are defined in the Attribute Flags TLV, which can be set and carried in either the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Objects.

- Cost Collection flag (to be assigned by IANA)
- Latency Collection flag (to be assigned by IANA)
- Latency Variation Collection flag (to be assigned by IANA)

These flags are meaningful in a Path message. If the Cost Collection flag is set to 1, the transit nodes SHOULD report the cost information in the Record Route Objects (RRO) of both the Path and Resv messages.

If the Cost Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.


If the Latency Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.

If the Latency Variation Collection flag is set to 1, the transit nodes SHOULD report latency variation information in the Record Route Objects (RRO) of both the Path and Resv messages.

The procedure for the processing the Attribute Flags TLV follows [RFC5420].

3.2. Cost Subobject

The cost subobject is defined for the RRO to record the cost information of the LSP. Its format is similar to the RRO subobjects (ROUTE_RECORD sub-object) defined in [RFC3209].



Type: TBA1 - Cost subobject (to be assigned by IANA).

Length: The Length value is set to 8 or 12 depending on the presence of Upstream Cost information.

Reserved: This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

Downstream Cost: Cost of the local link along the route of the LSP in the direction of the tail-end node, encoded as a 32-bit integer. Based on the policy at the recording node, the cost value can be set to the Interior Gateway Protocol (IGP) metric or TE metric of the link in question. This approach has been taken to avoid defining a flag for each cost type in the Attribute-Flags TLV. It is assumed that, based on policy, all nodes report the same cost-type and that

the ingress and egress nodes know the cost type reported in the RRO.

Upstream Cost: Cost of the local link along the route of the LSP in the direction of the head-end node, encoded as a 32-bit integer. Based on the policy at the recording node, the cost value can be set to the Interior Gateway Protocol (IGP) metric or TE metric of the link in question. This approach has been taken to avoid defining a flag for each cost type in the Attribute-Flags TLV. It is assumed that, based on policy, all nodes report the same cost-type and that the ingress and egress nodes know the cost type reported in the RRO.

3.3. Latency Subobject

The Latency subobject is defined for RRO to record the latency information of the LSP. Its format is similar the RRO subobjects defined in [RFC3209].

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								Reserved (must be zero)															
Reserved								Downstream Delay																							
Reserved								Upstream Delay																							

Type: TBA2 - Latency subobject (to be assigned by IANA).

Length: 8 or 12 depending on the presence of Upstream Cost information.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay respectively, as defined in [DRAFT-OSPF-TE-METRIC].

Reserved: These fields are reserved for future use. They MUST be set to 0 when sent and MUST be ignored when received.

Downstream Delay: Delay of the local link along the route of the LSP in the direction of the tail-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

Upstream Delay: Delay of the local link along the route of the LSP in the direction of the head-end node, encoded as 24-

bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

3.4. Latency Variation Subobject

The Latency Variation subobject is defined for RRO to record the Latency Variation information of the LSP. Its format is similar to the RRO subobjects defined in [RFC3209].

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								Reserved (must be zero)															
A Reserved								Downstream Delay Variation																							
A Reserved								Upstream Delay Variation																							

Type: TBA3 - Latency Variation subobject (to be assigned by IANA).

Length: 8 or 12 depending on the presence of Upstream Latency Variation information.

A-bit: These fields represent the Anomalous (A) bit associated with the Downstream and Upstream Delay respectively, as defined in [DRAFT-OSPF-TE-METRIC].

Reserved: These fields are reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

Downstream Delay Variation: Delay Variation of the local link along the route of the LSP in the direction of the tail-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

Upstream Delay Variation: Delay Variation of the local link along the route of the LSP in the direction of the head-end node, encoded as 24-bit integer. When set to 0, it has not been measured. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

3.5. Signaling Procedures

Typically, the ingress node learns the route of an LSP by adding a RRO in the Path message. If an ingress node also desires cost, latency and/or latency variation recording, it sets the appropriate flag(s) in the Attribute Flags TLV of the LSP_ATTRIBUTES (if recording is desired but not mandatory) or LSP_REQUIRED_ATTRIBUTES (if recording is mandatory) Object. None, all or any of the Cost Collection, Latency Collection or Latency Variation Collection flags may be set in the Attribute Flags TLV of the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Object. The rules for processing the LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES Objects and RRO are not changed. The corresponding sub-objects MUST be included in the RRO, with the Downstream (only) information filled in.

When a node receives a Path message which carries an LSP_REQUIRED_ATTRIBUTES Object and the Cost, Latency and/or Latency Variation Collection Flag(s) is (are) set, if local policy disallows providing the requested information to the endpoints, the node MUST return a Path Error message with error code "Policy Control Failure (2)" and one of the following error subcodes:

- . "Cost Recoding Rejected" (value to be assigned by IANA, suggested value 105) if Cost Collection Flag is set.
- . "Latency Recording Rejected" (value to be assigned by IANA, suggested value 106) if Latency Collection Flag is set.
- . "Latency Variation Recording Rejected" (value to be assigned by IANA, suggested value 107) if Latency Variation Collection Flag is set.

When a node receives a Path message which carries an LSP_ATTRIBUTES Object and the Cost, Latency and/or Latency Variation Collection Flag(s) is (are) set, if local policy disallows providing the requested information to the endpoints, the Path message SHOULD NOT be rejected due to Metric recording restriction and the Path message is forwarded without the appropriate sub-object(s) in the Path RRO.

If local policy permits the recording of the requested information, the processing node SHOULD add the requested subobject(s) with the cost, latency and/or latency variation metric value(s) associated with the local hop to the Path RRO. If the LSP being setup is bidirectional, both Downstream and Upstream information SHOULD be included. If the LSP is unidirectional, only Downstream information SHOULD be included.

the local hop in the Path RRO. When the egress node receives the Path message, it can calculate the end-to-end cost, latency and/or latency variation properties of the LSP.

Before the Resv message is sent to the upstream node, the egress node adds the requested subobject(s) with the downstream cost, latency and/or latency variation metric value(s) associated with the local hop to the Resv RRO in a similar manner to that specified above for the addition of Path RRO sub-objects by transit nodes.

Similarly, the intermediate nodes of the LSP provide the requested metric value(s) associated with the local hop in the Resv RRO. When the ingress node receives the Resv message, it can calculate the end-to-end cost, latency and/or latency variation properties of the LSP.

Typically, cost and latency are additive metrics, but latency variation is not an additive metric. The means by which the ingress and egress nodes compute the end-to-end cost, latency and latency variation metric from information recorded in the RRO is beyond the scope of this document.

Based on the local policy, the ingress and egress nodes can advertise the calculated end-to-end cost, latency and/or latency variation properties of the FA or RA LSP in TE link advertisement to the routing instance based on the procedure described in [DRAFT-OSPF-TE-METRIC], [DRAFT-ISIS-TE-METRIC].

Based on the local policy, a transit node (e.g. the edge node of a domain) may edit a Path or Resv RRO to remove route information (e.g. node or interface identifier information) before forwarding it. A node that does this SHOULD summarize the cost, latency and latency variation data removed as a single value for each for the loose hop that is summarized by the transit node. How a transit node calculates the cost, latency and/or latency variation metric for the segment summarized by the transit node is beyond the scope of this document.

4. Security Considerations

This document does not introduce any additional security issues above those identified in [RFC5920], [RFC5420], [RFC2205], [RFC3209], and [RFC3473].

5. IANA Considerations

5.1. RSVP Attribute Bit Flags

The IANA has created a registry and manages the space of attributes bit flags of Attribute Flags TLV as described in section 11.3 of [RFC5420]. It is requested that the IANA makes

Ali, Swallow, Filsfils Expires January 2014 [Page 9]

Internet-Draft draft-ietf-ccamp-te-metric-recording-02.txt

assignments from the Attribute Bit Flags defined in this document.

This document introduces the following three new Attribute Bit Flag:

- Bit number: TBD (recommended bit position 11)
- Defining RFC: this I-D
- Name of bit: Cost Collection Flag

- Bit number: TBD (recommended bit position 12)
- Defining RFC: this I-D
- Name of bit: Latency Collection Flag

- Bit number: TBD (recommended bit position 13)
- Defining RFC: this I-D
- Name of bit: Latency Variation Flag

5.2. ROUTE_RECORD subobject

This document introduces the following three new RRO subobject:

Type	Name	Reference
-----	-----	-----
TBD (35)	Cost subobject	This I-D
TBD (36)	Latency subobject	This I-D
TBD (37)	Latency Variation subobject	This I-D

5.2. New RSVP error sub-code

For Error Code = 2 "Policy Control Failure" (see [RFC2205]) the following sub-code is defined.

Sub-code	Value
-----	-----
Ali, Swallow, Filsfils	Expires January 2014

Internet-Draft draft-ietf-ccamp-te-metric-recording-02.txt

Cost Recoding Rejected	To be assigned by IANA. Suggested Value: 105.
Latency Recoding Rejected	To be assigned by IANA. Suggested Value: 106.
Latency Variation Recoding Rejected	To be assigned by IANA. Suggested Value: 107.

6. Acknowledgments

Authors would like to thank Ori Gerstel, Gabriele Maria Galimberti, Luyuan Fang and Walid Wakim for their review comments.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [DRAFT-OSPF-TE-METRIC] S. Giacalone, D. Ward, J. Drake, A. Atlas, S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", draft-ietf-ospf-te-metric-extensions, work in progress.
- [DRAFT-ISIS-TE-METRIC] S. Previdi, S. Giacalone, D. Ward, J. Drake, A. Atlas, C. Filsfils, "IS-IS Traffic Engineering (TE) Metric Extensions", draft-ietf-isis-te-metric-extensions, work in progress.
- [DRAFT-SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria,, "RSVP-TE Extensions for Collecting SRLG Information", draft-ietf-ccamp-rsvp-te-srlg-collect.txt, work in progress.

Internet-Draft draft-ietf-ccamp-te-metric-recording-02.txt

7.2. Informative References

- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", RFC 2209, September 1997.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

George Swallow
Cisco Systems, Inc.
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

Kenji Kumaki
KDDI Corporation
Email: ke-kumaki@kddi.com

Rudiger Kunze
Deutsche Telekom AG
Ruediger.Kunze@telekom.de

Network Working Group
Internet Draft
Intended status: Standards Track

H. Long
M.Ye
Huawei Technologies Co., Ltd
G. Mirsky
Ericsson
A Alessandro
Telecom Italia S.p.A
July 3, 2013

Expires: January 2014

OSPF Routing Extension for links with variable discrete bandwidth
draft-long-ccamp-ospf-availability-extension-00.txt

Abstract

Packet switching network may contain links with variable discrete bandwidth, e.g., copper, radio, etc. The bandwidth of such link may change discretely in reaction to changing external environment. Availability is typically used for describing such links during network planning. This document describes an extension for OSPF routing for route computation in a Packet Switched Network (PSN) which contains link with variable discrete bandwidth by introducing an optional availability sub-TLV.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on January 7, 2009.

Copyright Notice

Copyright (c) 2013 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
2. Overview	3
3. Extension to OSPF Routing Protocol	4
3.1. Interface Switching Capacity Descriptor.....	4
3.2. ISCD Availability sub-TLV.....	4
3.3. Signaling Process.....	5
4. Security Considerations.....	5
5. IANA Considerations	5
6. References	5
6.1. Normative References.....	5
6.2. Informative References.....	6
7. Acknowledgments	6

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

The following acronyms are used in this draft:

OSPF	Open Shortest Path First
PSN	Packet Switched Network
SNR	Signal-to-noise Ratio
LSP	Label Switched Path
ISCD	Interface Switching Capacity Descriptor

PE Provider Edge

LSA Link State Advertisement

1. Introduction

There are some data communication technologies that allow seamless change of maximum physical bandwidth. For example, in mobile backhaul network, microwave links are very popular for providing connection of last hops. In case of heavy rain, to maintain the link connectivity, the microwave link will lower the modulation level since demodulating lower modulation level need lower signal-to-noise ratio (SNR). This is called adaptive modulation technology [EN 302 217]. However, lower modulation level also means lower link bandwidth. When link bandwidth reduces by modulation down-shifting, high priority traffic can be maintained, while lower priority traffic is dropped. Similarly the cooper links may change their effective link bandwidth due to external interference.

The parameter, availability [G.827, F.1703, P.530], is often used to describe the link capacity during network planning. Assigning different availability classes to different types of service over such kind of links provides more efficient planning of link capacity. To set up an LSP across these links, availability information is required for the nodes to verify bandwidth satisfaction and make bandwidth reservation. The availability information should be inherited from the availability requirements of the services expected to be carried on the LSP, voice service usually needs "five nines" availability, while non-real time data packets may needs four or three nines availability.

For the route computation, the availability information should be provided along with bandwidth resource information. In this document, an extension on Interface Switching Capacity Descriptor (ISCD) [RFC4202] for availability information support in routing signaling. The extension reuses the reserved field in the ISCD and also introduces an optional availability sub-TLV.

2. Overview

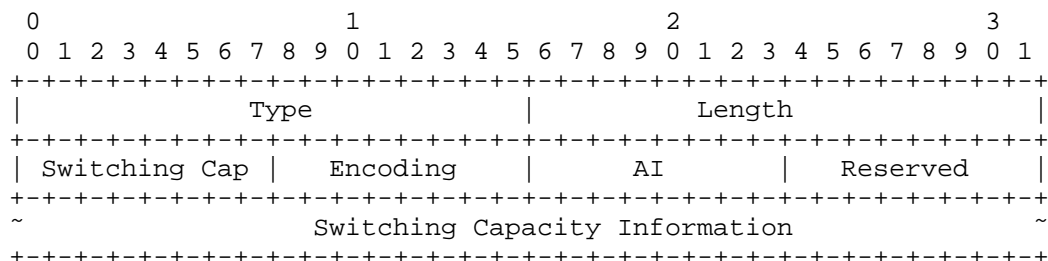
A node which has link(s) with variable bandwidth attached should contain a <bandwidth, availability> information list in its OSPF TE LSA messages. The list provides the information that how much bandwidth a link can support for a specified availability. This information is used for path calculation by the PE node(s).

To setup a label switching path (LSP), a PE node may collect link information which is spread in OSPF TE LSA message by network nodes to get know about the network topology, and calculate out a LSP route based on the network topology, and send the calculated LSP route to signaling to initiate a PATH/RESV message for setting up the LSP.

3. Extension to OSPF Routing Protocol

3.1. Interface Switching Capacity Descriptor

The Interface Switching Capacity Descriptor (ISCD) sub-TLV [RFC 4203] has the following format:



Type: TBD, 16 bits;

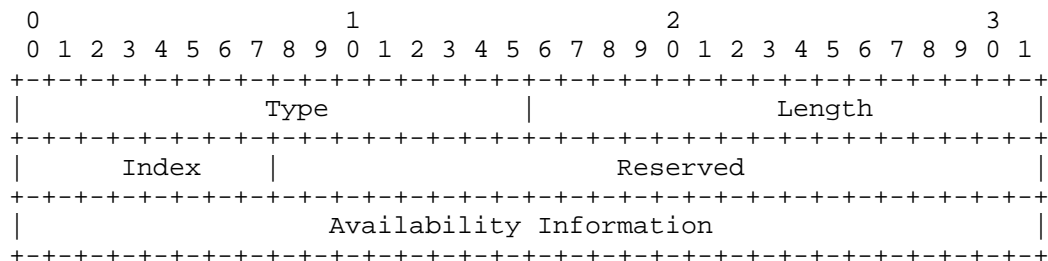
Length: 16 bits;

AI: ISCD Availability sub-TLV index, 8 bits

This field is the index of availability sub-TLV for this ISCD sub-TLV.

3.2. ISCD Availability sub-TLV

The availability sub-TLV has the following format:



Type: TBD, 16 bits;

Length: 16 bits;

Index: 8 bits

This field is the index of this availability sub-TLV, referred by the AI field of the ISCD sub-TLV.

Availability Information: 32 bits

This field is a 32-bit IEEE floating point number which describes the availability guarantee of the switching capacity in the ISCD object which has the AI value equal to Index of this sub-TLV. The value must be less than 1.

3.3. Signaling Process

A node which has link(s) with variable bandwidth attached should contain one or more ISCD Availability sub-TLVs in its OSPF TE LSA messages. Each ISCD Availability sub-TLV provides the information that how much bandwidth a link can support for a specified availability. This information is used for path calculation by the PE node(s).

4. Security Considerations

This document does not introduce new security considerations to the existing OSPF protocol.

5. IANA Considerations

TBD

6. References

6.1. Normative References

[RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.

[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.

- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4202] Kompella, K. and Rekhter, Y. (Editors), "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [G.827] ITU-T Recommendation, "Availability performance parameters and objectives for end-to-end international constant bit-rate digital paths", September, 2003.
- [F.1703] ITU-R Recommendation, "Availability objectives for real digital fixed wireless links used in 27 500 km hypothetical reference paths and connections", January, 2005.
- [P.530] ITU-R Recommendation, "Propagation data and prediction methods required for the design of terrestrial line-of-sight systems", February, 2012
- [EN 302 217] ETSI standard, "Fixed Radio Systems; Characteristics and requirements for point-to-point equipment and antennas", April, 2009

6.2. Informative References

- [MCOS] Minei, I., Gan, D., Kompella, K., and X. Li, "Extensions for Differentiated Services-aware Traffic Engineered LSPs", Work in Progress, June 2006.

7. Acknowledgments

Authors' Addresses

Hao Long
Huawei Technologies Co., Ltd.
No.1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731, P.R.China

Phone: +86-18615778750
Email: longhao@huawei.com

Min Ye
Huawei Technologies Co., Ltd.
No.1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731, P.R.China

Email: amy.yemin@huawei.com

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Alessandro D'Alessandro
Telecom Italia S.p.A

Email: alessandro.dalessandro@telecomitalia.it

Network Working Group
Internet Draft
Intended status: Standards Track

H. Long
M.Ye
Huawei Technologies Co., Ltd
G. Mirsky
Ericsson
A Alessandro
Telecom Italia S.p.A
July 3, 2013

Expires: January 2014

RSVP-TE Signaling Extension for Bandwidth availability
draft-long-ccamp-rsvp-te-bandwidth-availability-01.txt

Abstract

Packet switching network usually contains links with variable bandwidth, e.g., copper, radio, etc. The bandwidth of such link is sensitive to external environment. Availability is typically used for describing the link during network planning. This document describes an extension for RSVP-TE signaling for setting up a label switching path (LSP) in a Packet Switched Network (PSN) network which contains variable bandwidth link by introducing an optional availability field in RSVP-TE signaling.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on January 7, 2009.

Copyright Notice

Copyright (c) 2013 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
2. Overview	4
3. Extension to RSVP-TE Signaling.....	4
3.1. SENDER_TSPEC Object.....	4
3.1.1. Bandwidth Profile TLV.....	5
3.2. FLOWSPEC Object.....	6
3.3. Signaling Process.....	6
4. Security Considerations.....	7
5. IANA Considerations	7
6. References	7
6.1. Normative References.....	7
6.2. Informative References.....	8
7. Acknowledgments	8

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

The following acronyms are used in this draft:

RSVP-TE	Resource Reservation Protocol-Traffic Engineering
LSP	Label Switched Path
PSN	Packet Switched Network
SNR	Signal-to-noise Ratio

TLV	Type Length Value
PE	Provider Edge
LSA	Link State Advertisement

1. Introduction

The RSVP-TE specification [RFC3209] and GMPLS extensions [RFC3473] specify the signaling message including the bandwidth request for setting up a label switching path in a PSN network.

There are some data communication technologies that allow seamless change of maximum physical bandwidth. For example, in mobile backhaul network, microwave links are very popular for providing connection of last hops. In case of heavy rain, to maintain the link connectivity, the microwave link will lower the modulation level since demodulating lower modulation level need lower signal-to-noise ratio (SNR). This is called adaptive modulation technology [EN 302 217]. However, lower modulation level also means lower link bandwidth. When link bandwidth reduces by modulation down-shifting, high priority traffic can be maintained, while lower priority traffic is dropped. Similarly the cooper links may change their link bandwidth due to external interference.

The parameter, availability [G.827, F.1703, P.530], is often used to describe the link capacity during network planning. Assigning different availability classes to different types of service over such kind of links provides more efficient planning of link capacity. To set up a LSP across these links, availability information is required for the nodes to verify bandwidth satisfaction and make bandwidth reservation. The availability information should be inherited from the availability requirements of the services expected to be carried on the LSP, voice service usually needs "five nines" availability, while non-real time data packets may needs four or three nines availability. Since different service types may need different availabilities guarantee, multiple <availability, bandwidth> pairs may be required when signaling.

To fulfill LSP setup by signaling in these scenarios, this document specifies the following extension:

- o A new SENDER_TSPEC object is defined which includes multiple bandwidth profiles with different availability. This object is an extension on the Ethernet SENDER_TSPEC defined by [RFC6003] which support multiple bandwidth profile TLVs, but limited in the scope of Ethernet. The extension uses the object generically, and amends availability information in the bandwidth profile TLV.

2. Overview

A PSN tunnel may span one or more links in a network. To setup a label switching path (LSP), a PE node may collect link information which is spread in routing message, e.g., OSPF TE LSA message, by network nodes to get know about the network topology, and calculate out a LSP route based on the network topology, and send the calculated LSP route to signaling to initiate a PATH/RESV message for setting up the LSP.


In case that there is(are) link(s) with variable bandwidth in a network, a <bandwidth, availability> requirement list should be specified for a LSP. Each <bandwidth, availability> pair in the list means a bandwidth with specified availability is required. The list could be inherited from the result of service planning for the LSP.

When a PE node initiates a PATH/RESV signaling for setting up the LSP, the PATH message should carry the <bandwidth, availability> requirement list as bandwidth request, and the intermediate node(s) will allocate the bandwidth resource for each availability requirement from the remaining bandwidth with corresponding availability. An error message may be returned if any <bandwidth, availability> request cannot be satisfied.

3. Extension to RSVP-TE Signaling

3.1. SENDER_TSPEC Object

The SENDER_TSPEC object (Class-Num = 12) has the following format:



Class-Specific Information: 32 bits

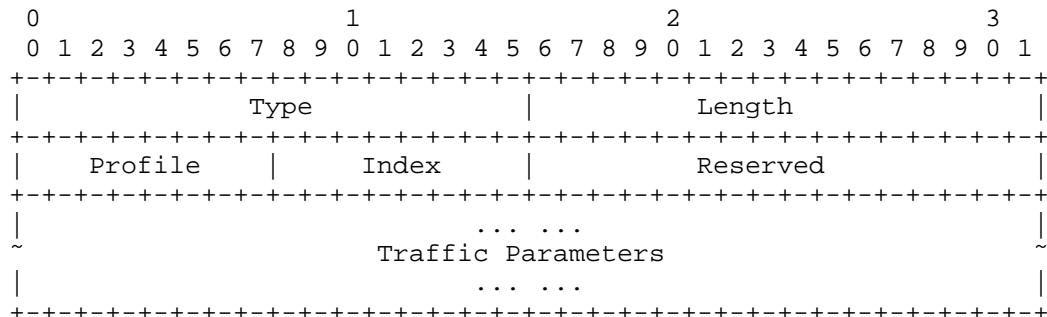
This field indicates the specific information for each C-Type.

TLV (Type-Length-Value):

The SENDER_TSPEC object MUST include at least one TLV and MAY include more than one TLV.

3.1.1.1. Bandwidth Profile TLV

The Bandwidth Profile TLV has the following format.



Type: TBD, 16 bits;


Length: 16 bits;

Profile: 8 bits

This field is defined as a bit vector of binary flags. The following flags are defined:

Flag 3 (bit 2): Availability Flag (AF)

When The Flag 3 is set to value 1, there is an availability sub-TLV included in this Bandwidth Profile TLV. The availability sub-TLV has the following format:



Type (2 octets): TBD

Length (2 octets): 4

Availability (4 octets): a 32-bit floating number describes availability requirement for this bandwidth request. The value must be less than 1.

Index: 8 bits

See [RFC6003] section 4.1.

Traffic Parameters:

This field includes the traffic parameters information. The format is different for different C-Type.

C-Type = IntServ: See [RFC2210];

C-Type = Ethernet: See [RFC6003];

3.2. FLOWSPEC Object

The FLOWSPEC object (Class-Num = 9, Class-Type = TBD) has the same format as the Ethernet SENDER_TSPEC object.

3.3. Signaling Process

The source node initiates PATH messages including one or more Bandwidth Profile TLVs with different availability value in the SENDER_TSPEC object. Each Bandwidth Profile TLV specifies the portion of bandwidth request with referred availability requirement.

The destination nodes check whether it can satisfy the bandwidth requirement by comparing each bandwidth requirement inside the SENDER_TSPEC objects with the remaining link sub-bandwidth resource with respective availability guarantee when received the PATH message.

- o If all bandwidth requirements can be satisfied, it should reserve the bandwidth resource from each remaining sub-bandwidth portion to set up this LSP. Optionally, the higher availability bandwidth can be allocated to lower availability request when the lower availability bandwidth cannot satisfy the request.
- o If at least one bandwidth requirement cannot be satisfied, it should generate PathErr message with the error code "Traffic Control Error" and the error value "Bad Tspec value" (see [RFC2205]).

4. Security Considerations

This document does not introduce new security considerations to the existing RSVP-TE signaling protocol.

5. IANA Considerations

TBD

6. References

6.1. Normative References

- [RFC2210] Wroclawski, J., ''The Use of RSVP with IETF Integrated Services'', RFC 2210, September 1997.
- [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.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC6003] Papadimitriou, D. ''Ethernet Traffic Parameters'', RFC 6003, October 2010.

- [G.827] ITU-T Recommendation, ''Availability performance parameters and objectives for end-to-end international constant bit-rate digital paths'', September, 2003.
- [F.1703] ITU-R Recommendation, ''Availability objectives for real digital fixed wireless links used in 27 500 km hypothetical reference paths and connections'', January, 2005.
- [P.530] ITU-R Recommendation, '' Propagation data and prediction methods required for the design of terrestrial line-of-sight systems'', February, 2012
- [EN 302 217] ETSI standard, ''Fixed Radio Systems; Characteristics and requirements for point-to-point equipment and antennas'', April, 2009

6.2. Informative References

- [MCOS] Minei, I., Gan, D., Kompella, K., and X. Li, "Extensions for Differentiated Services-aware Traffic Engineered LSPs", Work in Progress, June 2006.

7. Acknowledgments

Authors' Addresses

Hao Long
Huawei Technologies Co., Ltd.
No.1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731, P.R.China

Phone: +86-18615778750
Email: longhao@huawei.com

Min Ye
Huawei Technologies Co., Ltd.
No.1899, Xiyuan Avenue, Hi-tech Western District
Chengdu 611731, P.R.China

Email: amy.yemin@huawei.com

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Alessandro D'Alessandro
Telecom Italia S.p.A

Email: alessandro.dalessandro@telecomitalia.it

CCAMP
Internet-Draft
Intended status: Standards Track
Expires: January 14, 2014

G. Martinelli, Ed.
Cisco
A. Zanardi, Ed.
CREATE-NET
X. Zhang, Ed.
Huawei Technologies
G. Galimberti
Cisco
D. Siracusa
CREATE-NET
July 13, 2013

Information Encoding for WSON with Impairments Validation
draft-martinelli-ccamp-wson-iv-encode-02

Abstract

Impairment-Aware (IA) Routing and Wavelength Assignment (RWA) function might be required in Wavelength Switched Optical Networks (WSON) that already support RWA. This document defines proper encoding to support this operation. It goes in addition to the available impairment-free WSON encoding and it is fully compatible with it.

As the information model, the encoding is independent from control plane architectures and protocol implementations. Its definitions can be used in related protocol extensions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2014.

Copyright Notice

Copyright (c) 2013 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	2
1.1. Requirements Language	2
2. Encoding	3
2.1. Optical Parameter	3
2.2. Impairment Vector	4
2.3. Impairment Matrix	5
3. Acknowledgements	6
4. Contributing Authors	6
5. IANA Considerations	7
6. Security Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	7
Authors' Addresses	8

1. Introduction

In case of WSON where optical impairments play a significant role, the framework document [RFC6566] defines related control plane architectural options for Impairment Aware Routing and Wavelength Assignment (IA-RWA). This document provides a suitable encoding for the related WSON impairment information model as defined [I-D.martinelli-ccamp-wson-iv-info].

This document directly refers to ITU recommendations [ITU.G680] and [ITU.G697] as already detailed in the information model.


1.1. 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 [RFC2119].

2. Encoding

2.1. Optical Parameter

The OPTICAL_PARAM is defined as a sub TLV object.



The following flag is defined:

S. Standard bit. S=1 identifies a set of parameters standardized by ITU; while S=0 identifies a non-standardized set of parameters.

With the flag S=1 the following parameters are defined:

ParamSource. Where this parameter is defined. Currently only [ITU.G697] has defined this with value 1.

ParamID. Parameter identifier according to the source. [ITU.G697] table V.3 defines the following identifiers:


1. Total Power (dBm)
2. Channel Power (dBm)
3. Reserved (Defined in [ITU.G697] but not used)
4. Reserved (Defined in [ITU.G697] but not used)
5. OSNR (db)
6. Q Factor (a pure number)
7. PMD (ps)
8. Residual Chromatic Dispersion (ps/nm)

Value. Value for the parameter. As defined by [ITU.G697], it is a 32 bit IEEE floating point number.

Variance. Variance for the parameter, a 32 bit IEEE floating point number.

2.2. Impairment Vector

This sub-TLV is a list of optical parameters and they MAY have a wavelength dependency information.

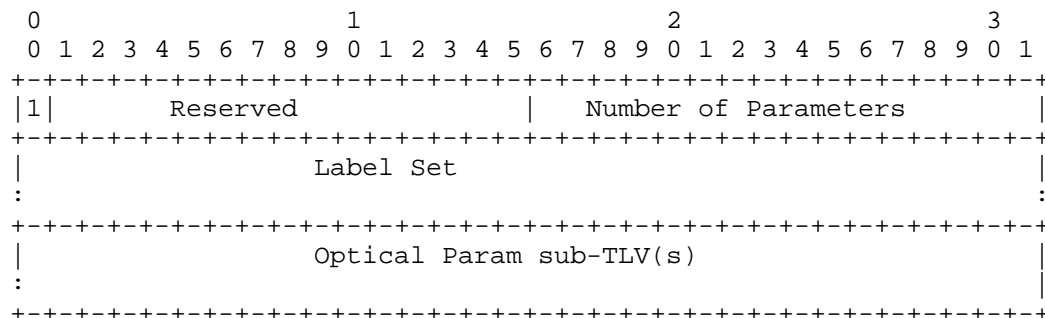


Where:

W = 0. Wavelength Dependency flag. There is no wavelength dependency.

Number of Parameters contained in this vector.

Optical Param sub-TLV(s) present a list of Object as defined in Section 2.1.



Where:

W = 1. Wavelength Dependency flag. There is wavelength dependency.

The Label Set object is defined in [I-D.ietf-ccamp-general-constraint-encode] Section 2.1. Likely an

inclusive range will be the only option required by the Action defined in the Label Set.

2.3. Impairment Matrix

As defined by the [I-D.martinelli-ccamp-wson-iv-info], the impairment matrix follows the same structure as the connectivity matrix.

[Xian's note]: (1) Similar problem as mentioned above applies, the structure differs when N is set to different values; so they should be described separately; (2) I would prefer the "W" stay with the previous structure. (3) Since the format of impairment matrix does not follow exactly as the connectivity matrix. i would suggest revising the sentence above to reflect this.

```

      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
+-----+-----+-----+-----+
| Connectivity | MatrixID | Reserved | 0 |
+-----+-----+-----+-----+
|               Link Set A #1               |
|               :               |
+-----+-----+-----+-----+
|               Link Set B #1               |
|               :               |
+-----+-----+-----+-----+
| Impairment Vector sub-TLV(s) |
|               :               |
+-----+-----+-----+-----+
| Additional Link Set pairs and Impairment Vector(s) |
|               :               |
+-----+-----+-----+-----+

      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
+-----+-----+-----+-----+
| Connectivity | MatrixID | Reserved | 1 |
+-----+-----+-----+-----+
| Impairment Vector sub-TLV(s) |
|               :               |
+-----+-----+-----+-----+

```

Where:

Connectivity: value MUST be 2 for the impairment matrix (Values 0 and 1 are already defined by [I-D.ietf-ccamp-general-constraint-encode]).

MatrixID: matrix identifier, the scope of this integer number is shared with [I-D.ietf-ccamp-rwa-info].

N: Node scope flag. With this flag set there's no Link Set information but only a list of optical parameters TLVs that apply to the whole optical node.

The usage of multiple matrixes with connectivity type equal to 2 (Impairment Matrix) MIGHT be used to group optical parameters by connectivity. For example, if a subset of parameters apply to the whole node, a unique matrix with flag N=1 is used. At the same time another subset of parameters applies only to some LinkSet pairs, a specific Impairment Matrix will be added.

3. Acknowledgements

TBD

4. Contributing Authors

This document was the collective work of several authors. The text and content of this document was contributed by the editors and the co-authors listed below (the contact information for the editors appears in appropriate section and is not repeated below):

Moustafa Kattan
Cisco
DUBAI, 500321
UNITED ARAB EMIRATES

Email: mkattan@cisco.com

Young Lee
Huawei
1700 Alma Drive, Suite 100
Plano, TX 75075
USA

Phone: +1 972 509 5599 x2240
Fax: +1 469 229 5397
Email: ylee@huawei.com

Fatai Zhang
Huawei
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
P.R. China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

5. IANA Considerations

This document does not contain any IANA request.

6. Security Considerations

This document defines an protocol-neutral encoding for an information model describing impairments in optical networks and it does not introduce any security issues. If such a encoding is put into use within a network it will by its nature contain details of the physical characteristics of an optical network. Such information would need to be protected from intentional or unintentional disclosure.

7. References

7.1. Normative References

[ITU.G680] International Telecommunications Union, "Physical transfer functions of optical network elements ", ITU-T Recommendation G.680, July 2007.

[ITU.G697] International Telecommunications Union, "Optical monitoring for dense wavelength division multiplexing systems ", ITU-T Recommendation G.697, February 2012.

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

7.2. Informative References

[I-D.ietf-ccamp-general-constraint-encode]

Bernstein, G., Lee, Y., Li, D., and W. Imajuku, "General Network Element Constraint Encoding for GMPLS Controlled Networks", draft-ietf-ccamp-general-constraint-encode-11 (work in progress), May 2013.

[I-D.ietf-ccamp-rwa-info]

Lee, Y., Bernstein, G., Li, D., and W. Imajuku, "Routing and Wavelength Assignment Information Model for Wavelength Switched Optical Networks", draft-ietf-ccamp-rwa-info-18 (work in progress), May 2013.

[I-D.martinelli-ccamp-wson-iv-info]

Martinelli, G., Kattan, M., Galimberti, G., and A. Zanardi, "Information Model for Wavelength Switched Optical Networks (WSON) with Optical Impairments Validation.", draft-martinelli-ccamp-wson-iv-info-01 (work in progress), February 2013.

[I-D.narten-iana-considerations-rfc2434bis]

Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", draft-narten-iana-considerations-rfc2434bis-09 (work in progress), March 2008.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC6566] Lee, Y., Bernstein, G., Li, D., and G. Martinelli, "A Framework for the Control of Wavelength Switched Optical Networks (WSONs) with Impairments", RFC 6566, March 2012.

Authors' Addresses

Giovanni Martinelli (editor)
Cisco
via Philips 12
Monza 20900
Italy

Phone: +39 039 2092044
Email: giomarti@cisco.com

Andrea Zanardi (editor)
CREATE-NET
via alla Cascata 56 C, Povo
Trento 38100
Italy

Email: andrea.zanardi@create-net.org

Xian Zhang (editor)
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129
P.R. China

Phone: +86 755 28972913
Email: zhang.xian@huawei.com

Gabriele M. Galimberti
Cisco
Via Philips,12
Monza 20900
Italy

Phone: +39 039 2091462
Email: ggalimbe@cisco.com

Domenico Siracusa
CREATE-NET
via alla Cascata 56 C, Povo
Trento 38100
Italy

Email: domenico.siracusa@create-net.org

CCAMP
Internet-Draft
Intended status: Informational
Expires: January 14, 2014

G. Martinelli, Ed.
Cisco
X. Zhang, Ed.
Huawei Technologies
G. Galimberti
Cisco
A. Zanardi
D. Siracusa
CREATE-NET
July 13, 2013

Information Model for Wavelength Switched Optical Networks (WSONs) with
Impairments Validation
draft-martinelli-ccamp-wson-iv-info-02

Abstract

This document defines an information model to support Impairment-Aware (IA) Routing and Wavelength Assignment (RWA) function. This operation might be required in Wavelength Switched Optical Networks (WSON) that already support RWA and the information model defined here goes in addition and it is fully compatible with the already defined information model for impairment-free RWA process in WSON.

This information model shall support all control plane architectural options defined for WSON with impairment validation.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 14, 2014.

Copyright Notice

Copyright (c) 2013 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	2
1.1. Applicability	3
2. Properties of an Impairment Information Model	3
3. Background from WSON Information Model	5
4. Optical Impairment Information Model	6
4.1. The Optical Impairment Vector	7
4.2. Node Information	7
4.3. Link Information	9
4.4. Path Information	9
5. Encoding Considerations	10
6. Control Plane Architectures	10
6.1. IV-Centralized	11
6.2. IV-Distributed	11
7. Acknowledgements	11
8. Contributing Authors	11
9. IANA Considerations	13
10. Security Considerations	13
11. References	13
11.1. Normative References	13
11.2. Informative References	13
Appendix A. G.680 Essential information	14
Authors' Addresses	14

1. Introduction

In the context of Wavelength Switched Optical Network (WSON), [RFC6163] describes the basic framework for a GMPLS and PCE-based Routing and Wavelength Assignment (RWA) control plane. The associated information model [I-D.ietf-ccamp-rwa-info] defines all information/parameters required by an RWA process.

There are cases of WSON where optical impairments plays a significant role and are considered as important constraints. The framework

document [RFC6566] defines problem scope and related control plane architectural options for the Impairment Aware Routing and Wavelength Assignment (IA-RWA) operation. Options include different combinations of Impairment Validation (IV) and RWA functions in term of different combination of control plane functions (i.e., PCE, Routing, Signaling).

This document provides an information model for the impairment aware case to allow the impairment validation function implemented in the control plane or enabled by control plane available information. This model goes in addition to [I-D.ietf-ccamp-rwa-info] and it shall support any control plane architectural option described by the framework document (see sections 4.2 and 4.3 of [RFC6566]) where a set of control plane combinations of control plane functions vs. IV functions is provided.

1.1. Applicability

This document targets at Scenario C defined in [RFC6566] section 4.1.1., i.e., approximate impairment estimation. The usage of "Approximate" concept helps to clarify that a path resulting from a IV-RWA process are not guaranteed to work from an optical perspective however, has better chance to be feasible than other(s). Setting up the lighth-path, there is still a possibility to fail but this is an acceptable in term of GMPLS control plane.

The information model defined in this document strictly relates to a control plane information model hence it does not define any new Optical Computational Model. Optical Computational Models are defined by ITU-T, [ITU.G680] defines transfer functions for some linear optical parameters while there's no general optical computational model for non-linear effects. This information model uses the available optical computational model as a reference, and identify GMPLS WSON control plane extensions to allow IV functions to perform its task.

Following here above considerations, scope of this information model is to provide a generic mechanism so will easily support additional impairment parameters and/or additional optical computational models.

2. Properties of an Impairment Information Model

An information model may have several attributes or properties that need to be defined for each optical parameter made available to the control plane. The properties will help to determine how the control plane can deal with a specific impairment parameter, depending on architectural options chosen within the overall impairment framework [RFC6566]. In some case, properties value will help to identify the level of approximation supported by the IV process.

- o Time Dependency

This identifies how an impairment parameter may vary with time. There could be cases where there is no time dependency, while in other cases there may be need of re-evaluation after a certain time. In this category, variations in impairments due to environmental factors such as those discussed in [G.sup47] are considered. In some cases, an impairment parameter that has time dependency may be considered as a constant for approximation. In this information model, we do neglect this property.

- o Wavelength Dependency

This property identifies if an impairment parameter can be considered as constant over all the wavelength spectrum of interest or not. Also in this case a detailed impairment evaluation might lead to consider the exact value while an approximation IV might take a constant value for all wavelengths. In this information model, we consider both case: dependency / no dependency on a specific wavelength. This property appears directly in the information model definitions and related encoding.

- o Linearity

As impairments are representation of physical effects, there are some that have a linear behavior while other are non-linear. Linear approximation is in scope of Scenario C of [RFC6566]. During the impairment validation process, this property implies that the optical effect (or quantity) satisfies the superposition principle, thus a final result can be calculated by the sum of each component. The linearity implies the additivity of optical quantities considered during an impairment validation process. The non-linear effects in general does not satisfy this property. The information model presented in this document however, easily allow introduction of non-linear optical effects with a linear approximated contribution to the linear ones.

- o Multi-Channel

There are cases where a channel's impairments take different values depending on the aside wavelengths already in place, this is mostly due to non-linear impairments. The result would be a dependency among different LSPs sharing the same path. This information model do not consider this kind of property.

The following table summarize the above considerations where in the first column reports the list of properties to be considered for each optical parameter, while the second column states if this property is taken into account or not by this information model.

Property	Info Model Awareness
Time Dependency	no
Wavelength Dependency	yes
Linearity	yes
Multi-channel	no

Table 1: Optical Impairment Properties

3. Background from WSOON Information Model

In this section we report terms already defined for the WSOON-RWA (impairment free) as in [I-D.ietf-ccamp-rwa-info] and [I-D.ietf-ccamp-general-constraint-encode]. The purpose is to provide essential information that will be reused or extended for the impairment case.

In particular [I-D.ietf-ccamp-rwa-info] defines the connectivity matrix as the following:

```
ConnectivityMatrix ::= <MatrixID> <ConnType> <Matrix>
```

According to [I-D.ietf-ccamp-general-constraint-encode], this definition is further detailed as:

```
ConnectivityMatrix ::=
  <MatrixID> <ConnType> ((<LinkSet> <LinkSet>) ...)
```

This second formula highlights how the connectivity matrix is built by pairs of LinkSet objects identifying the internal connectivity capability due to internal optical node constraint(s). It's essentially binary information and tell if a wavelength or a set of wavelengths can go from an input port to an output port.

As an additional note, connectivity matrix belongs to node information and is purely static. Dynamic information related to the actual usage of the connections is available through specific extension to link information.

4. Optical Impairment Information Model

The idea behind this information model is to categorize the impairment parameters into three types and extend the information model already defined for impairment-free WSONs. The three categories are:

- o Node Information. The concept of connectivity matrix is reused and extended to introduce an impairment matrix.
- o Link Information representing impairment information related to a specific link or hop.
- o Path Information representing the impairment information related to the whole path.

All the above three categories will make use of a generic container, the Impairment Vector, to transport optical impairment information.

In term of optical parameters, this document is not to rephrase content from [ITU.G680] but only provide necessary building blocks that allow the IA-RWA process to apply the specific Optical Computational Model (i.e. transfer functions). Section 9 of [ITU.G680] defines the optical computational models and provides information to calculate the following optical parameters:

- o OSNR. Section 9.1
- o Optical Power. As per Section 9.1, required by Optical Computation Model for OSNR calculation.
- o Chromatic Dispersion (CD). Section 9.2
- o Polarization Mode Dispersion (PMD). Section 9.3
- o Polarization Dependent Loss (PDL). Section 9.3

This information model however will allow however to add additional parameters beyond the one defined by [ITU.G680] in order to support additional computational models. This mechanism could eventually be applicable to both linear and non-linear parameters.

This information model makes the assumption that the each optical node in the network is able to provide the control plane protocols with its own parameter values however, no assumption is made on how the optical node gets those value information (e.g. internally computed, provisioned by a network management system, etc.). To this extent, the information model intentionally ignores all internal detailed parameters that are used by the formulas of the Optical Computational Model (i.e., "transfer function") and simply provides the object containers to carry results of the formulas.

4.1. The Optical Impairment Vector

Optical Impairment Vector (OIV) is defined as a list of optical parameters to be associated to a WSON node or a WSON link. It is defined as:

```
<OIV> ::= ([<LabelSet>] <OPTICAL_PARAM>) ...
```

The optional LabelSet object enables wavelength dependency property as per Table 1. LabelSet has its definition in [I-D.ietf-ccamp-general-constraint-encode].

OPTICAL_PARAM. This object represents an optical parameter. The Impairment vector can contain a set of parameters as identified by [ITU.G697] since those parameters match the terms of the linear impairments computational models provided by [ITU.G680]. This information model does not speculate about the set of parameters (since defined elsewhere, e.g. ITU-T), however it does not preclude extensions by adding new parameters.

4.2. Node Information

Impairment matrix describes a list of the optical parameters that applies to a network element as a whole or ingress/egress port pairs of a network element. Wavelength dependency property of optical parameters is also considered.

```
ImpairmentMatrix ::= <MatrixID> <ConnType>  
    ((<LinkSet> <LinkSet> <OIV>) ...)
```

Where:


MatrixID. This ID is a unique identifier for the matrix. It shall be unique in scope among connectivity matrices defined in [I-D.ietf-ccamp-rwa-info] and impairment matrices defined here.

ConnType. This number identifies the type of matrix and it shall be unique in scope with other values defined by impairment-free WSON documents.

LinkSet. Same object definition and usage as [I-D.ietf-ccamp-general-constraint-encode]. The pairs of LinkSet identify one or more internal node constrain.

OIV. The Optical Impairment Vector defined above.

The model can be represented as a multidimensional matrix shown in the following picture



The connectivity matrix from [I-D.ietf-ccamp-general-constraint-encode] is only a two dimensional matrix, containing only binary information, through the LinkSet pairs. In this model, a third dimension is added by generalizing the binary information through the Optical Impairment Vector associated with each LinkSet pair. Optical parameters in the picture are reported just as examples while details go into specific encoding draft [I-D.martinelli-ccamp-wson-iv-encode].

This representation shows the most general case however, the total amount of information transported by control plane protocols can be greatly reduced by proper encoding when the same set of values apply to all LinkSet pairs.

[EDITOR NODE: first run of the information model does looks for generality not for optimizing the quantity of information. We'll deal with optimization in a further step.]

4.3. Link Information

For the list of optical parameters associated to the link, the same approach used for the node-specific impairment information can be applied. The link-specific impairment information is extended from [I-D.ietf-ccamp-rwa-info] as the following:

```
<DynamicLinkInfo> ::=  <LinkID> <AvailableLabels>
                        [<SharedBackupLabels>] [<OIV>]
```

DynamicLinkInfo is already defined in [I-D.ietf-ccamp-rwa-info] while OIV is the Optical Impairment Vector is defined in the previous section.

4.4. Path Information

There are cases where the optical impariments can only be described as a contrains on the overall end to end path. In such case, the optical impariment and/or parameter, cannot be derived (using a simple function) from the set of node / link contributions.

An equivalent case is the option reported by [RFC6566] on IV-Candidate paths where, the control plane knows a list of optically feasible paths so a new path setup can be selected among that list. Independent from the protocols and functions combination (i.e. RWA vs. Routing vs. PCE), the IV-Candidates imply a path property stating that a path is optically feasible.

<PathInfo> ::= <OIV>

[EDITOR NOTE: section to be completed, especially to evaluate protocol implications. Likely resemble to RSVP ADSPEC].

5. Encoding Considerations

Details about encoding will be defined in a separate document [I-D.martinelli-ccamp-wson-iv-encode] however worth remembering that, within [ITU.G697] Appending V, ITU already provides a guideline for encoding some optical parameters.

In particular [ITU.G697] indicates that each parameter shall be represented by a 32 bit floating point number.

As an additional consideration, actual values for each optical parameter are provided by each optical node and it could provide by direct measurement or from some internal computation starting from indirect measurement. In any case the encoding shall provide the possibility to associate a variance with the parameter. This information will enable the function implementing IA-RWA process to make some additional considerations on wavelength feasibility. [RFC6566] Section 4.1.3 reports some considerations regarding this degree of confidence during the impairment validation process.

6. Control Plane Architectures

This section briefly describes how the definitions contained in this information model will match the architectural options described by [RFC6566].

The first assumption is that the WSON GMPLS extensions are available and operational. To such extent, the WSON-RWA will provide the following information through its path computation (and RWA process):

- o The wavelengths connectivity, considering also the connectivity constraints limited by reconfigurable optics, and wavelengths availability.
- o The interface compatibility at the physical level.
- o The Optical-Electro-Optical (OEO) availability within the network (and related physical interface compatibility). As already stated by the framework this information it's very important for impairment validation:

- a. If the IV functions fail (path optically infeasible), the path computation function may use an available OEO point to find a feasible path. In normally operated networks OEO are mainly used to support optically unfeasible path than mere wavelength conversion.
- b. The OEO points reset the optical impairment information since a new light is generated.

6.1. IV-Centralized

Centralized IV process is performed by a single entity (e.g., a PCE). Given sufficient impairment information, it can either be used to provide a list of paths between two nodes, which are valid in terms of optical impairments. Alternatively, it can help validate whether a particular selected path and wavelength is feasible or not. This requires distribution of impairment information to the entity performing the IV process.

[EDITOR NOTE: to be completed]

6.2. IV-Distributed

For the distributed IV process, common computational models are needed together with the information model defined in this document. Computational models for the optical impairments are defined by ITU standard body. The currently available computation models are reported in [ITU.G680] and only cover the linear impairment case. This does not require the distribution of impairment information since they can be collected hop-by-hop using a control plane signaling protocol.

[EDITOR NOTE: to be completed]

7. Acknowledgements

TBD

8. Contributing Authors

This document was the collective work of several authors. The text and content of this document was contributed by the editors and the co-authors listed below (the contact information for the editors appears in appropriate section and is not repeated below):

Moustafa Kattan

Cisco
DUBAI, 500321
UNITED ARAB EMIRATES

Email: mkattan@cisco.com

Tim Gibbon
Department of Physics
Nelson Mandela Metropolitan University
SOUTH AFRICA

Email: Tim.Gibbon@nmmu.ac.za

Young Lee
Huawei
1700 Alma Drive, Suite 100
Plano, TX 75075
USA

Greg M. Bernstein
Grotto Networking
Fremont, CA
USA

Phone: +1 510 573 2237
Email: gregb@grotto-networking.com

Phone: +1 972 509 5599 x2240
Fax: +1 469 229 5397
Email: ylee@huawei.com

Fatai Zhang
Huawei
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
P.R. China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

9. IANA Considerations

This document does not contain any IANA requirement.

10. Security Considerations

This document defines an information model for impairments in optical networks. If such a model is put into use within a network it will by its nature contain details of the physical characteristics of an optical network. Such information would need to be protected from intentional or unintentional disclosure.

11. References

11.1. Normative References

[ITU.G680]

International Telecommunications Union, "Physical transfer functions of optical network elements ", ITU-T Recommendation G.680, July 2007.

[ITU.G697]

International Telecommunications Union, "Optical monitoring for dense wavelength division multiplexing systems ", ITU-T Recommendation G.697, February 2012.

[RFC2119]

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

11.2. Informative References

[I-D.ietf-ccamp-general-constraint-encode]

Bernstein, G., Lee, Y., Li, D., and W. Imajuku, "General Network Element Constraint Encoding for GMPLS Controlled Networks", draft-ietf-ccamp-general-constraint-encode-11 (work in progress), May 2013.

[I-D.ietf-ccamp-rwa-info]

Lee, Y., Bernstein, G., Li, D., and W. Imajuku, "Routing and Wavelength Assignment Information Model for Wavelength Switched Optical Networks", draft-ietf-ccamp-rwa-info-18 (work in progress), May 2013.

[I-D.martinelli-ccamp-wson-iv-encode]

Martinelli, G., Kattan, M., Galimberti, G., and A. Zanardi, "Encoding for WSON Information Model with Impairments Validation.", draft-martinelli-ccamp-wson-iv-encode-01 (work in progress), February 2013.

- [I-D.narten-iana-considerations-rfc2434bis]
Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", draft-narten-iana-considerations-rfc2434bis-09 (work in progress), March 2008.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [RFC6163] Lee, Y., Bernstein, G., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSOONS)", RFC 6163, April 2011.
- [RFC6566] Lee, Y., Bernstein, G., Li, D., and G. Martinelli, "A Framework for the Control of Wavelength Switched Optical Networks (WSOONS) with Impairments", RFC 6566, March 2012.

Appendix A. G.680 Essential information

TBD if we need some info instead of reading [ITU.G680]

[Xian's note]: I thought about listing parameters to show different categories. but it seems to contradict to the idea of providing a general enough information model. Any suggestions?

Authors' Addresses

Giovanni Martinelli (editor)
Cisco
via Philips 12
Monza 20900
Italy

Phone: +39 039 2092044
Email: giomarti@cisco.com

Xian Zhang (editor)
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129
P.R. China

Phone: +86 755 28972913
Email: zhang.xian@huawei.com

Gabriele M. Galimberti
Cisco
Via Philips,12
Monza 20900
Italy

Phone: +39 039 2091462
Email: ggalimbe@cisco.com

Andrea Zanardi
CREATE-NET
via alla Cascata 56 C, Povo
Trento 38100
Italy

Email: andrea.zanardi@create-net.org

Domenico Siracusa
CREATE-NET
via alla Cascata 56 C, Povo
Trento 38100
Italy

Email: domenico.siracusa@create-net.org

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 07, 2013

O. Gonzalez de Dios, Ed.
Telefonica I+D
R. Casellas, Ed.
CTTC
F. Zhang
Huawei
X. Fu
ZTE
D. Ceccarelli
Ericsson
I. Hussain
Infinera
June 05, 2013

Framework and Requirements for GMPLS based control of Flexi-grid DWDM
networks
draft-ogrcetal-ccamp-flexi-grid-fwk-03

Abstract

This document defines a framework and the associated control plane requirements for the GMPLS based control of flexi-grid DWDM networks. To allow efficient allocation of optical spectral bandwidth for high bit-rate systems, the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) has extended the recommendations [G.694.1] and [G.872] to include the concept of flexible grid. A new DWDM grid has been developed within the ITU-T Study Group 15 by defining a set of nominal central frequencies, channel spacings and the concept of "frequency slot". In such environment, a data plane connection is switched based on allocated, variable-sized frequency ranges within the optical spectrum.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 07, 2013.

Copyright Notice

Copyright (c) 2013 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. Requirements Language	3
2. Introduction	3
3. Acronyms	4
4. Flexi-grid Networks	4
4.1. Flexi-grid in the context of OTN	4
4.2. Terminology	5
4.2.1. Frequency Slots	5
4.2.2. Media Channels	7
4.2.3. Media Layer Elements	7
4.2.4. Optical Tributary Signals	8
4.3. Flexi-grid layered network model	8
4.3.1. Hierarchy in the Media Layer	9
4.3.2. DWDM flexi-grid enabled network element models	10
5. GMPLS applicability	10
5.1. General considerations	11
5.2. Considerations on TE Links	11
5.3. Considerations on Labeled Switched Path (LSP) in Flexi-grid	13
5.4. Control Plane modeling of Network elements	17
5.5. Media Layer Resource Allocation considerations	17
5.6. Neighbor Discovery and Link Property Correlation	21
5.7. Path Computation / Routing and Spectrum Assignment (RSA)	21
5.7.1. Architectural Approaches to RSA	22
5.8. Routing / Topology dissemination	23
5.8.1. Available Frequency Ranges/slots of DWDM Links	23
5.8.2. Available Slot Width Ranges of DWDM Links	23
5.8.3. Spectrum Management	23
5.8.4. Information Model	24
6. Control Plane Requirements	25

6.1. Functional requirements	25
6.2. Routing/Topology Dissemination requirements	25
6.3. Signaling requirements	25
7. Security Considerations	25
8. Contributing Authors	25
9. Acknowledgments	27
10. References	28
10.1. Normative References	28
10.2. Informative References	29
Authors' Addresses	29

1. 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 [RFC2119].

2. Introduction

The term "Flexible grid" (flexi-grid for short) as defined by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) Study Group 15 in the latest version of [G.694.1], refers to the updated set of nominal central frequencies (a frequency grid), channel spacing and optical spectrum management/allocation considerations that have been defined in order to allow an efficient and flexible allocation and configuration of optical spectral bandwidth for high bit-rate systems.

A key concept of flexi-grid is the "frequency slot"; a variable-sized optical frequency range that can be allocated to a data connection. As detailed later in the document, a frequency slot is characterized by its nominal central frequency and its slot width which, as per [G.694.1], is constrained to be a multiple of a given slot width granularity.

Compared to a traditional fixed grid network, which uses fixed size optical spectrum frequency ranges or "frequency slots" with typical channel separations of 50 GHz, a flexible grid network can select its media channels with a more flexible choice of slot widths, allocating as much optical spectrum as required, allowing high bit rate signals (e.g., 400G, 1T or higher) that do not fit in the fixed grid.

From a networking perspective, a flexible grid network is assumed to be a layered network [G.872][G.800] in which the media layer is the server layer and the optical signal layer is the client layer. In the media layer, switching is based on a frequency slot, and the size of a media channel is given by the properties of the associated frequency slot. In this layered network, the media channel transports an Optical Tributary Signal.

A Wavelength Switched Optical Network (WSON), addressed in [RFC6163], is a term commonly used to refer to the application/deployment of a Generalized Multi-Protocol Label Switching (GMPLS)-based control plane for the control (provisioning/recovery, etc) of a fixed grid WDM network in which media (spectrum) and signal are jointly considered

This document defines the framework for a GMPLS-based control of flexi-grid enabled DWDM networks (in the scope defined by ITU-T layered Optical Transport Networks [G.872]), as well as a set of associated control plane requirements. An important design consideration relates to the decoupling of the management of the optical spectrum resource and the client signals to be transported.

3. Acronyms

EFS: Effective Frequency Slot

FS: Frequency Slot

NCF: Nominal Central Frequency

OCh: Optical Channel

OCh-P: Optical Channel Payload

OTS: Optical Tributary Signal

OCC: Optical Channel Carrier

SWG: Slot Width Granularity

4. Flexi-grid Networks

4.1. Flexi-grid in the context of OTN

[G.872] describes from a network level the functional architecture of Optical Transport Networks (OTN). The OTN is decomposed into independent layer networks with client/layer relationships among them. A simplified view of the OTN layers is shown in Figure 1.




Figure 1: Generic OTN overview

In the OTN layering context, the media layer is the server layer of the optical signal layer. The optical signal is guided to its destination by the media layer by means of a network media channel. In the media layer, switching is based on a frequency slot, and the size of a media channel is given by the properties of the associated frequency slot.

In this scope, this document uses the term flexi-grid enabled DWDM network to refer to a network in which switching is based on frequency slots defined using the flexible grid, and covers mainly the Media Layer as well as the required adaptations from the Signal layer. The present document is thus focused on the control and management of the media layer.

4.2. Terminology

This section presents the definition of the terms used in flexi-grid networks. These terms are included in the ITU-T recommendations [G.694.1], [G.872]), [G.870], [G.8080] and [G.959.1-2013].

Where appropriate, this documents also uses terminology and lexicography from [RFC4397].

4.2.1. Frequency Slots

This subsection is focused on the frequency slot related terms.

- o Frequency Slot [G.694.1]: The frequency range allocated to a slot within the flexible grid and unavailable to other slots. A frequency slot is defined by its nominal central frequency and its slot width.

Nominal Central Frequency: each of the allowed frequencies as per the definition of flexible DWDM grid in [G.694.1]. The set of nominal central frequencies can be built using the following expression $f = 193.1 \text{ THz} + n \times 0.00625 \text{ THz}$, where 193.1 THz is ITU-T "anchor frequency" for transmission over the C band, n is a positive or negative integer including 0.




Figure 2: Anchor frequency and set of nominal central frequencies

Nominal Central Frequency Granularity: It is the spacing between allowed nominal central frequencies and it is set to 6.25 GHz (note: sometimes referred to as 0.00625 THz).

Slot Width Granularity: 12.5 GHz, as defined in [G.694.1].

Slot Width: The slot width determines the "amount" of optical spectrum regardless of its actual "position" in the frequency axis. A slot width is constrained to be $m \times \text{SWG}$ (that is, $m \times 12.5 \text{ GHz}$), where m is an integer greater than or equal to 1.

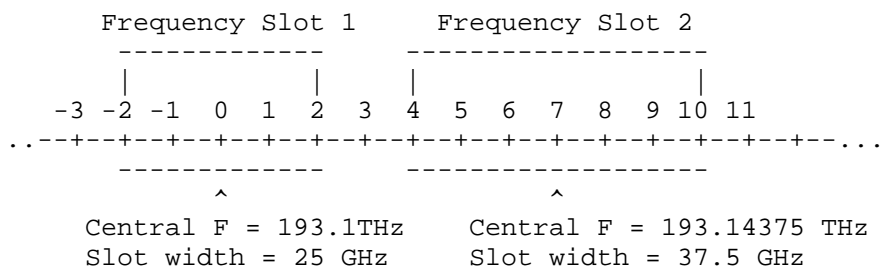


Figure 3: Example Frequency slots

- o The symbol '+' represents the allowed nominal central frequencies, the '--' represents the nominal central frequency granularity, and the '^' represents the slot nominal central frequency. The number on the top of the '+' symbol represents the 'n' in the frequency calculation formula. The nominal central frequency is 193.1 THz when n equals zero.

Effective Frequency Slot: the effective frequency slot of a media channel is the common part of the frequency slots along the media channel through a particular path through the optical network. It is a logical construct derived from the (intersection of) frequency slots allocated to each device in the path. The effective frequency slot is an attribute of a media channel and, being a frequency slot, it is described by its nominal central frequency and slot width, according to the already described rules.




Figure 4: Effective Frequency Slot

4.2.2. Media Channels

Media Channel: a media association that represents both the topology (i.e., path through the media) and the resource (frequency slot) that it occupies. As a topological construct, it represents a (effective) frequency slot supported by a concatenation of media elements (fibers, amplifiers, filters, switching matrices...). This term is used to identify the end-to-end physical layer entity with its corresponding (one or more) frequency slots local at each link filters.

Network Media Channel: It is a media channel that transports an Optical Tributary Signal [Editor's note: this definition goes beyond current G.870 definition, which is still tightened to a particular case of OTS, the OCh-P]

4.2.3. Media Layer Elements

Media Element: a media element only directs the optical signal or affects the properties of an optical signal, it does not modify the properties of the information that has been modulated to produce the optical signal [G.870]. Examples of media elements include fibers, amplifiers, filters and switching matrices.

Media Channel Matrixes: the media channel matrix provides flexible connectivity for the media channels. That is, it represents a point

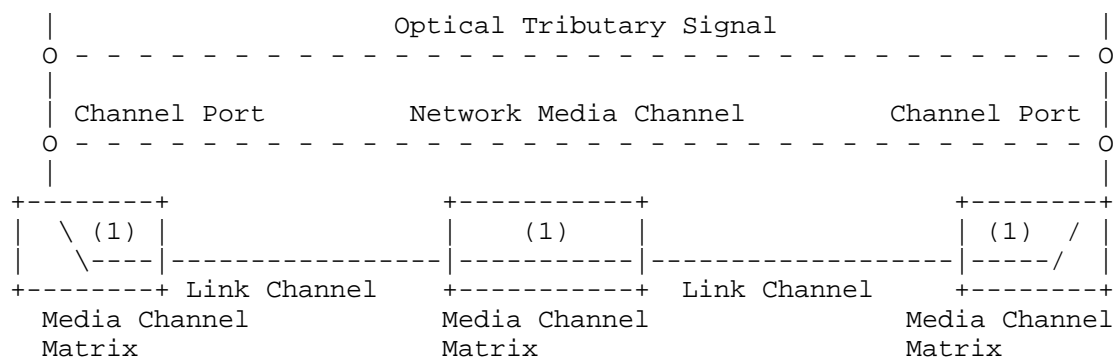
of flexibility where relationships between the media ports at the edge of a media channel matrix may be created and broken. The relationship between these ports is called a matrix channel. (Network) Media Channels are switched in a Media Channel Matrix.

4.2.4. Optical Tributary Signals

Optical Tributary Signal [G.959.1-2013]: The optical signal that is placed within a network media channel for transport across the optical network. This may consist of a single modulated optical carrier or a group of modulated optical carriers or subcarriers. One particular example of Optical Tributary Signal is an Optical Channel Payload (OCh-P) [G.872].

4.3. Flexi-grid layered network model


In the OTN layered network, the network media channel transports a single Optical Tributary Signal (see Figure 5)



(1) - Matrix Channel

Figure 5: Simplified Layered Network Model

A particular example of Optical Tributary Signal is the OCh-P. Figure 6 shows the example of the layered network model particularized for the OCH-P case, as defined in G.805.



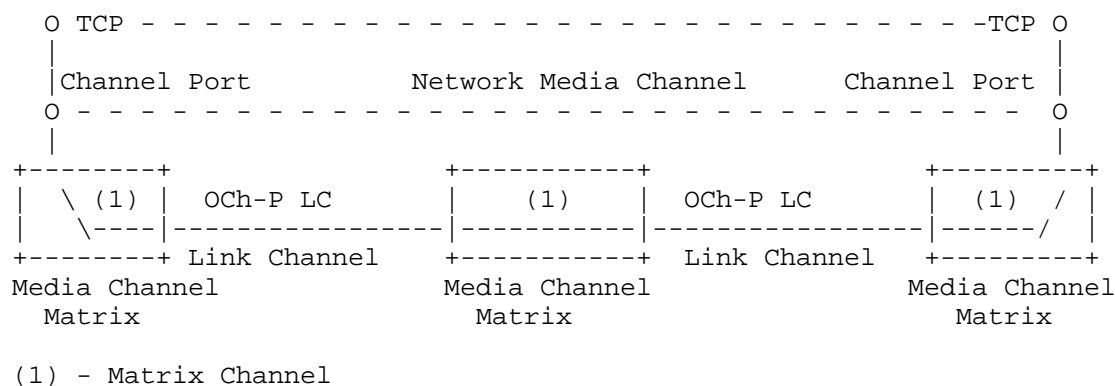



Figure 6: Layered Network Model according to G.805

By definition a network media channel only supports a single Optical Tributary signal. How several Optical Tributary signals are bound together is out of the scope of the present document and is a matter of the signal layer.

4.3.1. Hierarchy in the Media Layer

In summary, the concept of frequency slot is a logical abstraction that represents a frequency range while the media layer represents the underlying media support. Media Channels are media associations, characterized by their (effective) frequency slot, respectively; and media channels are switched in media channel matrixes. From the control and management perspective, a media channel can be logically splitted in other media channels.

In Figure 7 , a Media Channel has been configured and dimensioned to support two network media channels, each of them carrying one optical tributary signal.



```

      <- Network Media Channel->      <- Network Media Channel->

<----- Media Channel ----->

X - Frequency Slot Central Frequency

o - signal central frequency

```

Figure 7: Example of Media Channel / Network Media Channels and associated frequency slots

4.3.2. DWDM flexi-grid enabled network element models

Similar to fixed grid networks, a flexible grid network is also constructed from subsystems that include Wavelength Division Multiplexing (WDM) links, tunable transmitters and receivers, i.e., media elements including media layer switching elements (media matrices), as well as electro-optical network elements, all of them with flexible grid characteristics.

As stated in [G.694.1] the flexible DWDM grid defined in Clause 7 has a nominal central frequency granularity of 6.25 GHz and a slot width granularity of 12.5 GHz. However, devices or applications that make use of the flexible grid may not be capable of supporting every possible slot width or position. In other words, applications may be defined where only a subset of the possible slot widths and positions are required to be supported. For example, an application could be defined where the nominal central frequency granularity is 12.5 GHz (by only requiring values of n that are even) and that only requires slot widths as a multiple of 25 GHz (by only requiring values of m that are even).

5. GMPLS applicability

The goal of this section is to provide an insight of the application of GMPLS to control flexi-grid networks, while specific requirements are covered in the next section. The present framework is aimed at controlling the media layer within the Optical Transport Network (OTN) hierarchy and the required adaptations of the signal layer. This document also defines the term SSON (Spectrum-Switched Optical Network) to refer to a Flexi-grid enabled DWDM network that is controlled by a GMPLS/PCE control plane.

This section provides a mapping of the ITU-T G.872 architectural aspects to GMPLS/Control plane terms, and considers the relationship between the architectural concept/construct of media channel and its control plane representations (e.g. as a TE link).

5.1. General considerations

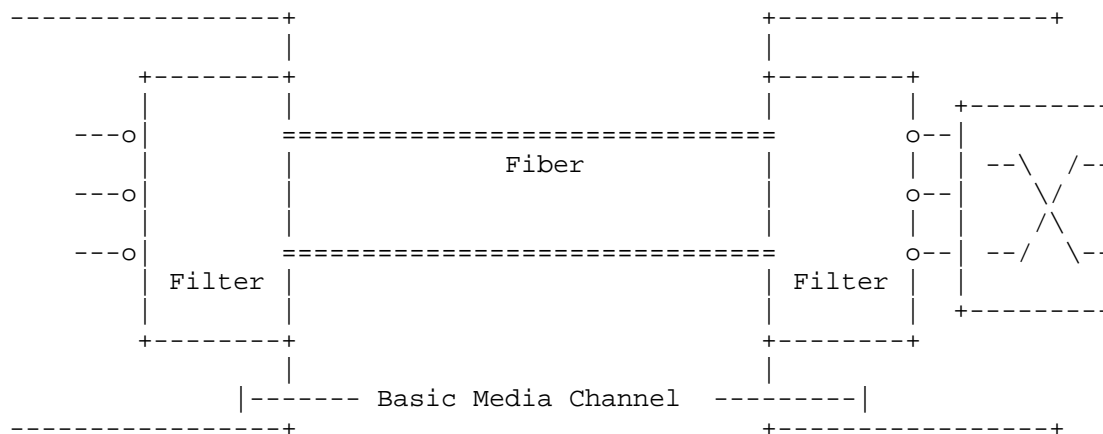
The GMPLS control of the media layer deals with the establishment of media channels, which are switched in media channel matrixes. GMPLS labels locally represent the media channel and its associated frequency slot. Network media channels are considered a particular case of media channels when the end points are transceivers (that is, source and destination of an Optical Tributary Signal)

5.2. Considerations on TE Links

From a theoretical / abstract point of view, a fiber can be modeled as having a frequency slot that ranges from $(-\infty, +\infty)$. This representation helps understand the relationship between frequency slots / ranges.

The frequency slot is a local concept that applies locally to a component / element. When applied to a media channel, we are referring to its effective frequency slot as defined in [G.872].

The association of a filter, a fiber and a filter is a media channel in its most basic form, which from the control plane perspective may be modeled as a (physical) TE-link with a contiguous optical spectrum at start of day. A means to represent this is that the portion of spectrum available at time t_0 depends on which filters are placed at the ends of the fiber and how they have been configured. Once filters are placed we have the one hop media channel. In practical terms, associating a fiber with the terminating filters determines the usable optical spectrum.






Figure 8: (Basic) Media channel and TE link

Additionally, when a cross-connect for a specific frequency slot is considered, the underlying media support is still a media channel, augmented, so to speak, with a bigger association of media elements and a resulting effective slot. When this media channel is the result of the association of basic media channels and media layer matrix cross-connects, this architectural construct can be represented as / corresponds to a Label Switched Path (LSP) from a control plane perspective. In other words, It is possible to "concatenate" several media channels (e.g. Patch on intermediate nodes) to create a single media channel.




Figure 9: Extended Media Channel

Additionally, if appropriate, it can also be represented as a TE link or Forwarding Adjacency (FA), augmenting the control plane network model.




Figure 10: Extended Media Channel / TE Link / FA

5.3. Considerations on Labeled Switched Path (LSP) in Flexi-grid

The flexi-grid LSP is seen as a control plane representation of a media channel. Since network media channels are media channels, an LSP may also be the control plane representation of a network media channel, in a particular context. From a control plane perspective, the main difference (regardless of the actual effective frequency slot which may be dimensioned arbitrarily) is that the LSP that represents a network media channel also includes the endpoints (transceivers), including the cross-connects at the ingress / egress nodes. The ports towards the client can still be represented as interfaces from the control plane perspective.

Figure 11 describes an LSP routed along 3 nodes. The LSP is terminated before the optical matrix of the ingress and egress nodes and can represent a Media Channel. This case does NOT (and cannot) represent a network media channel as it does not include (and cannot include) the transceivers.

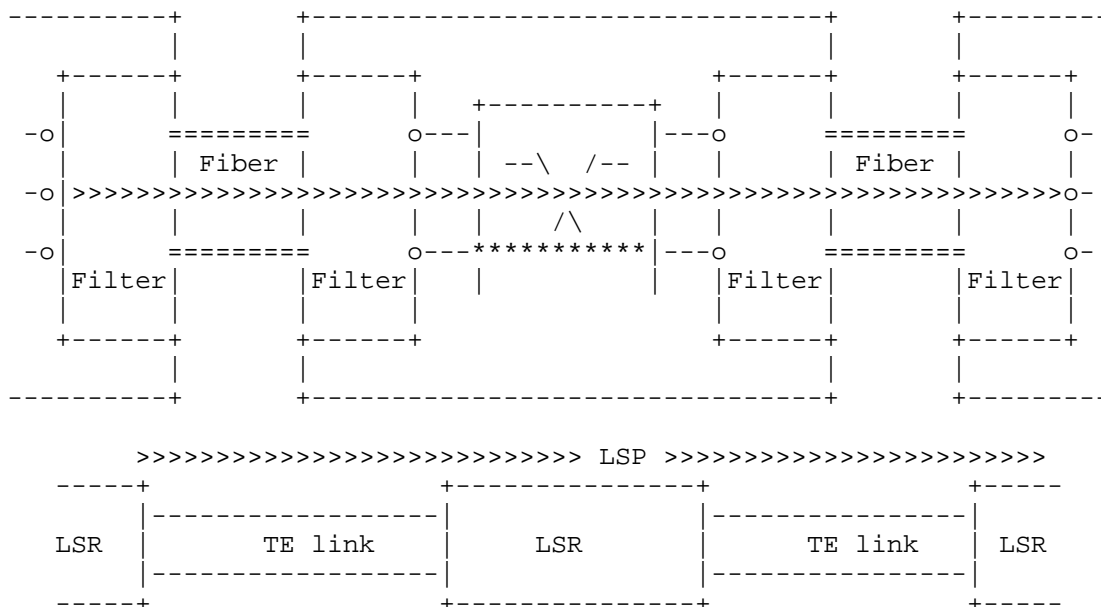
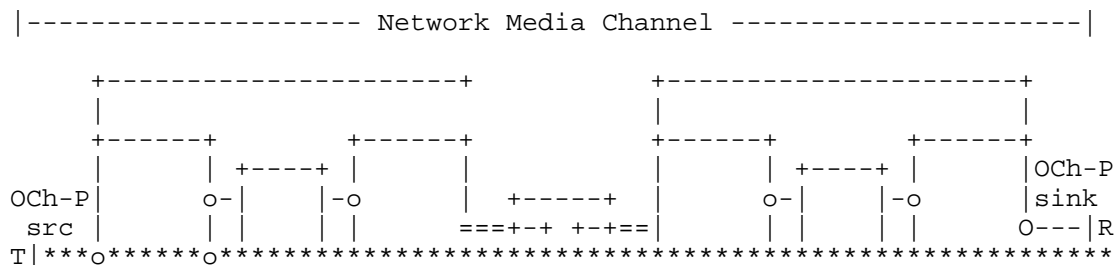


Figure 11: Flex-grid LSP representing a media channel that starts at the filter of the outgoing interface of the ingress LSR and ends at the filter of the incoming interface of the egress LSR

In Figure 12 a Network Media Channel is represented as terminated at the DWDM side of the transponder, this is commonly named as OCh-trail connection.







Figure 12: LSP representing a network media channel (OCh-Trail)

In a third case, a Network Media Channel terminated on the Filter ports of the Ingress and Egress nodes. This is named in G.872 as OCh-NC (we need to discuss the implications, if any, once modeled at the control plane level of models B and C).






Figure 13: LSP representing a network media channel (OCh-P NC)

[Note: not clear the difference, from a control plane perspective, of figs Figure 12 and Figure 13.]

Applying the notion of hierarchy at the media layer, by using the LSP as a FA, the media channel created can support multiple (sub) media channels. [Editot note : a specific behavior related to Hierarchies will be verified at a later point in time].




Figure 14: MRN/MLN topology view with TE link / FA

Note that there is only one media layer switch matrix (one implementation is FlexGrid ROADM) in SSON, while "signal layer LSP is mainly for the purpose of management and control of individual optical signal". Signal layer LSPs (OChs) with the same attributions (such as source and destination) could be grouped into one media-layer LSP (media channel), which has advantages in spectral efficiency (reduce guard band between adjacent OChs in one FSC) and LSP management. However, assuming some network elements indeed perform signal layer switch in SSON, there must be enough guard band between adjacent OChs in one media channel, in order to compensate filter concatenation effect and other effects caused by signal layer switching elements. In such condition, the separation of signal layer from media layer cannot bring any benefit in spectral efficiency and in other aspects, but make the network switch and

control more complex. If two OChs must switch to different ports, it is better to carry them by different FSCs and the media layer switch is enough in this scenario.

5.4. Control Plane modeling of Network elements

Optical transmitters/receivers may have different tunability constraints, and media channel matrixes may have switching restrictions. Additionally, a key feature of their implementation is their highly asymmetric switching capability which is described in [RFC6163] in detail. Media matrices include line side ports which are connected to DWDM links and tributary side input/output ports which can be connected to transmitters/receivers.

A set of common constraints can be defined:

- o The minimum and maximum slot width.
- o Granularity: the optical hardware may not be able to select parameters with the lowest granularity (e.g. 6.25 GHz for nominal central frequencies or 12.5 GHz for slot width granularity).
- o Available frequency ranges: the set or union of frequency ranges that are not allocated (i.e. available). The relative grouping and distribution of available frequency ranges in a fiber is usually referred to as "fragmentation".
- o Available slot width ranges: the set or union of slot width ranges supported by media matrices. It includes the following information.
 - * Slot width threshold: the minimum and maximum Slot Width supported by the media matrix. For example, the slot width can be from 50GHz to 200GHz.
 - * Step granularity: the minimum step by which the optical filter bandwidth of the media matrix can be increased or decreased. This parameter is typically equal to slot width granularity (i.e. 12.5GHz) or integer multiples of 12.5GHz.

[Editor's note: different configurations such as C/CD/CDC will be added later. This section should state specifics to media channel matrices, ROADM models need to be moved to an appendix].

5.5. Media Layer Resource Allocation considerations

A media channel has an associated effective frequency slot. From the perspective of network control and management, this effective slot is

seen as the "usable" frequency slot end to end. The establishment of an LSP related the establishment of the media channel and effective frequency slot.

In this context, when used unqualified, the frequency slot is a local term, which applies at each hop. An effective frequency slot applies at the media channel (LSP) level

A "service" request is characterized as a minimum, by its required effective slot width. This does not preclude that the request may add additional constraints such as imposing also the nominal central frequency. A given frequency slot is requested for the media channel say, with the Path message. Regardless of the actual encoding, the Path message sender descriptor sender_tspec shall specify a minimum frequency slot width that needs to be fulfilled.

In order to allocate a proper effective frequency slot for a LSP, the signaling should specify its required slot width.

An effective frequency slot must equally be described in terms of a central nominal frequency and its slot width (in terms of usable spectrum of the effective frequency slot). That is, one must be able to obtain an end-to-end equivalent n and m parameters. We refer to this as the "effective frequency slot of the media channel/LSP must be valid".

In GMPLS the requested effective frequency slot is represented to the TSpec and the effective frequency slot is mapped to the FlowSpec.

The switched element corresponds in GMPLS to the 'label'. As in flexi-grid the switched element is a frequency slot, the label represents a frequency slot. Consequently, the label in flexi-grid must convey the necessary information to obtain the frequency slot characteristics (i.e, center and width, the n and m parameters). The frequency slot is locally identified by the label

The local frequency slot may change at each hop, typically given hardware constraints (e.g. a given node cannot support the finest granularity). Locally n and m may change. As long as a given downstream node allocates enough optical spectrum, m can be different along the path. This covers the issue where concrete media matrices can have different slot width granularities. Such "local" m will appear in the allocated label that encodes the frequency slot as well as the flow descriptor flowspec.

Different modes are considered: RSA with explicit label control, and for R+DSA, the GMPLS signaling procedure is similar to the one described in section 4.1.3 of [RFC6163] except that the label set

should specify the available nominal central frequencies that meet the slot width requirement of the LSP. The intermediate nodes can collect the acceptable central frequencies that meet the slot width requirement hop by hop. The tail-end node also needs to know the slot width of a LSP to assign the proper frequency resource. Compared with [RFC6163], except identifying the resource (i.e., fixed wavelength for WSON and frequency resource for flexible grids), the other signaling requirements (e.g., unidirectional or bidirectional, with or without converters) are the same as WSON described in the section 6.1 of [RFC6163].

Regarding how a GMPLS control plane can assign n and m , different cases can apply:

- a) n and m can both change. It is the effective slot what matters. Some entity needs to make sure the effective frequency slot remains valid.
- b) m can change; n needs to be the same along the path. This ensures that the nominal central frequency stays the same.
- c) n and m need to be the same.
- d) n can change, m needs to be the same.

In consequence, an entity such as a PCE can make sure that the n and m stay the same along the path. Any constraint (including frequency slot and width granularities) is taken into account during path computation. alternatively, A PCE (or a source node) can compute a path and the actual frequency slot assignment is done, for example, with a distributed (signaling) procedure:

Each downstream node ensures that m is \geq requested _{m} .

Since a downstream node cannot foresee what an upstream node will allocate in turn, a way we can ensure that the effective frequency slot is valid is then by ensuring that the same " n " is allocated. By forcing the same n , we avoid cases where the effective frequency slot of the media channel is invalid (that is, the resulting frequency slot cannot be described by its n and m parameters).

Maybe this is a too hard restriction, since a node (or even a centralized/combined RSA entity) can make sure that the resulting end to end (effective) frequency slot is valid, even if n is different locally. That means, the effective (end to end) frequency slot that characterizes the media channel is one and determined by its n and m , but are logical, in the sense that they

are the result of the intersection of local (filters) freq slots which may have different freq. slots

For Figure Figure 15 the effective slot is valid by ensuring that the minimum m is greater than the requested m. The effective slot (intersection) is the lowest m (bottleneck).

For Figure Figure 16 the effective slot is valid by ensuring that it is valid at each hop in the upstream direction. The intersection needs to be computed. Invalid slots could result otherwise.

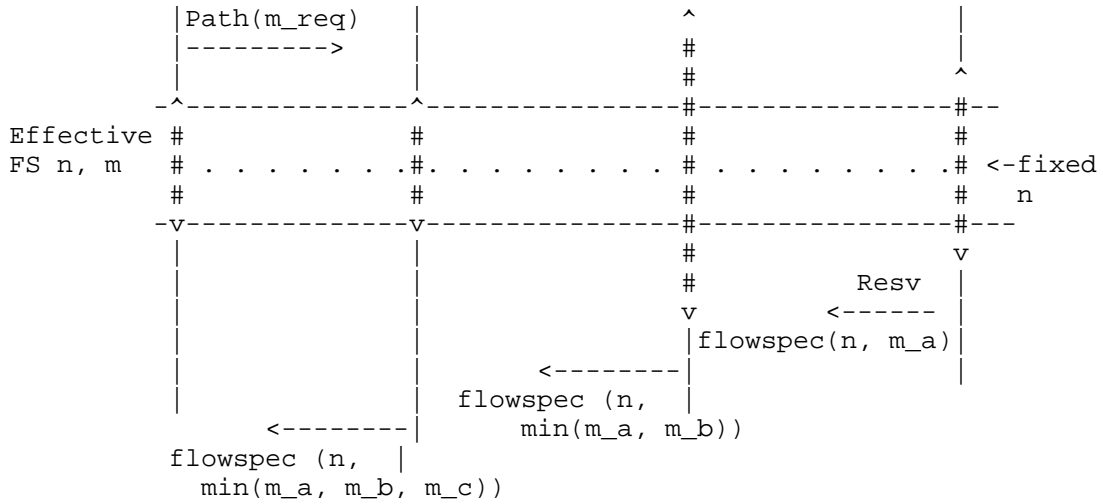



Figure 15: Distributed allocation with different m and same n



```

<-----|
flowspec ([intersect] FSa,FSb,FSc)

```

Figure 16: Distributed allocation with different m and different n

Note, when a media channel is bound to one OCh-P (i.e is a Network media channel), the EFS must be the one of the OCh-P. The media channel setup by the LSP may contains the EFS of the network media channel EFS. This is an endpoint property, the egress and ingress SHOULD constrain the EFS to OCh-P EFS .

5.6. Neighbor Discovery and Link Property Correlation

Potential interworking problems between fixed-grid DWDM and flexible-grid DWDM nodes, may appear. Additionally, even two flexible-grid optical nodes may have different grid properties, leading to link property conflict.

Devices or applications that make use of the flexible-grid may not be able to support every possible slot width. In other words, applications may be defined where different grid granularity can be supported. Taking node F as an example, an application could be defined where the nominal central frequency granularity is 12.5 GHz requiring slot widths being multiple of 25 GHz. Therefore the link between two optical nodes with different grid granularity must be configured to align with the larger of both granularities. Besides, different nodes may have different slot width tuning ranges.

In summary, in a DWDM Link between two nodes, at least the following properties should be negotiated:

Grid capability (channel spacing) - Between fixed-grid and flexible-grid nodes.

Grid granularity - Between two flexible-grid nodes.

Slot width tuning range - Between two flexible-grid nodes.

5.7. Path Computation / Routing and Spectrum Assignment (RSA)

Much like in WSON, in which if there is no (available) wavelength converters in an optical network, an LSP is subject to the "wavelength continuity constraint" (see section 4 of [RFC6163]), if the capability of shifting or converting an allocated frequency slot, the LSP is subject to the Optical "Spectrum Continuity Constraint".

Because of the limited availability of wavelength/spectrum converters (sparse translucent optical network) the wavelength/spectrum continuity constraint should always be considered. When available, information regarding spectrum conversion capabilities at the optical nodes may be used by RSA (Routing and Spectrum Assignment) mechanisms.

The RSA process determines a route and frequency slot for a LSP. Hence, when a route is computed the spectrum assignment process (SA) should determine the central frequency and slot width based on the slot width and available central frequencies information of the transmitter and receiver, and the available frequency ranges information and available slot width ranges of the links that the route traverses.

5.7.1. Architectural Approaches to RSA

Similar to RWA for fixed grids, different ways of performing RSA in conjunction with the control plane can be considered. The approaches included in this document are provided for reference purposes only; other possible options could also be deployed.

5.7.1.1. Combined RSA (R&SA)

In this case, a computation entity performs both routing and frequency slot assignment. The computation entity should have the detailed network information, e.g. connectivity topology constructed by nodes/links information, available frequency ranges on each link, node capabilities, etc.

The computation entity could reside either on a PCE or the ingress node.

5.7.1.2. Separated RSA (R+SA)

In this case, routing computation and frequency slot assignment are performed by different entities. The first entity computes the routes and provides them to the second entity; the second entity assigns the frequency slot.

The first entity should get the connectivity topology to compute the proper routes; the second entity should get the available frequency ranges of the links and nodes' capabilities information to assign the spectrum.

5.7.1.3. Routing and Distributed SA (R+DSA)

In this case, one entity computes the route but the frequency slot assignment is performed hop-by-hop in a distributed way along the route. The available central frequencies which meet the spectrum continuity constraint should be collected hop by hop along the route. This procedure can be implemented by the GMPLS signaling protocol.

5.8. Routing / Topology dissemination

In the case of combined RSA architecture, the computation entity needs to get the detailed network information, i.e. connectivity topology, node capabilities and available frequency ranges of the links. Route computation is performed based on the connectivity topology and node capabilities; spectrum assignment is performed based on the available frequency ranges of the links. The computation entity may get the detailed network information by the GMPLS routing protocol. Compared with [RFC6163], except wavelength-specific availability information, the connectivity topology and node capabilities are the same as WSON, which can be advertised by GMPLS routing protocol (refer to section 6.2 of [RFC6163]). This section analyses the necessary changes on link information brought by flexible grids.

5.8.1. Available Frequency Ranges/slots of DWDM Links

In the case of flexible grids, channel central frequencies span from 193.1 THz towards both ends of the C band spectrum with 6.25 GHz granularity. Different LSPs could make use of different slot widths on the same link. Hence, the available frequency ranges should be advertised.

5.8.2. Available Slot Width Ranges of DWDM Links

The available slot width ranges needs to be advertised, in combination with the Available frequency ranges, in order to verify whether a LSP with a given slot width can be set up or not; this is constrained by the available slot width ranges of the media matrix. Depending on the availability of the slot width ranges, it is possible to allocate more spectrum than strictly needed by the LSP.

5.8.3. Spectrum Management

[Editors' note: the part on the hierarchy of the optical spectrum could be confusing, we can discuss it]. The total available spectrum on a fiber could be described as a resource that can be divided by a media device into a set of Frequency Slots. In terms of managing spectrum, it is necessary to be able to speak about different granularities of managed spectrum. For example, a part of the spectrum could be assigned to a third party to manage. This need to

partition creates the impression that spectrum is a hierarchy in view of Management and Control Plane. The hierarchy is created within a management system, and it is an access right hierarchy only. It is a management hierarchy without any actual resource hierarchy within fiber. The end of fiber is a link end and presents a fiber port which represents all of spectrum available on the fiber. Each spectrum allocation appears as Link Channel Port (i.e., frequency slot port) within fiber.

5.8.4. Information Model

Fixed DM grids can also be described via suitable choices of slots in a flexible DWDM grid. However, devices or applications that make use of the flexible grid may not be capable of supporting every possible slot width or central frequency position. Following is the definition of information model, not intended to limit any IGP encoding implementation. For example, information required for routing/path selection may be the set of available nominal central frequencies from which a frequency slot of the required width can be allocated. A convenient encoding for this information (may be as a frequency slot or sets of contiguous slices) is further study in IGP encoding document.

```

<Available Spectrum in Fiber for frequency slot> ::=
  <Available Frequency Range-List>
  <Available Central Frequency Granularity >
  <Available Slot Width Granularity>
  <Minimal Slot Width>
  <Maximal Slot Width>

<Available Frequency Range-List> ::=
  <Available Frequency Range >[< Available Frequency Range-List>]

<Available Frequency Range > ::=
  <Start Spectrum Position><End Spectrum Position> |
  <Sets of contiguous slices>

<Available Central Frequency Granularity> ::= n x 6.25GHz,
  where n is positive integer, such as 6.25GHz, 12.5GHz, 25GHz, 50GHz
  or 100GHz

<Available Slot Width Granularity> ::= m x 12.5GHz,
  where m is positive integer

<Minimal Slot Width> ::= j x 12.5GHz,
  j is a positive integer

<Maximal Slot Width> ::= k x 12.5GHz,

```

k is a positive integer ($k \geq j$)

Figure 17: Routing Information model

6. Control Plane Requirements

This section provides a high level view of the requirements for GMPLS /PCE flexi-grid control plane. A detailed list of requirements will be provided in the next version of the document

6.1. Functional requirements

- o It must be able to dynamically set up media channels
- o It must be able to dynamically set up network media channels
- o It must must be able to dynamically set up a set of co-routed network media channels, and associate them logically

6.2. Routing/Topology Dissemination requirements

The computation entity needs to get the detailed network information: connectivity topology, node capabilities and available frequency ranges of the links

6.3. Signaling requirements

- o The signaling must be able to configure the minimum width (m) of an LSP.
- o The signaling must be able to configure the nominal central frequency (n) of an LSP.
- o It must be possible to collect the local frequency slot assigned at each link along the path

7. Security Considerations

TBD

8. Contributing Authors

Qilei Wang
ZTE
Ruanjian Avenue, Nanjing, China
wang.qilei@zte.com.cn

Malcolm Betts
ZTE
malcolm.betts@zte.com.cn

Xian Zhang
Huawei
zhang.xian@huawei.com

Cyril Margaria
Nokia Siemens Networks
St Martin Strasse 76, Munich, 81541, Germany
+49 89 5159 16934
cyril.margaria@nsn.com

Sergio Belotti
Alcatel Lucent
Optics CTO
Via Trento 30 20059 Vimercate (Milano) Italy
+39 039 6863033
sergio.belotti@alcatel-lucent.com

Yao Li
Nanjing University
wsliguotou@hotmail.com

Fei Zhang
ZTE
Zijinghua Road, Nanjing, China
zhang.fei3@zte.com.cn

Lei Wang
ZTE
East Huayuan Road, Haidian district, Beijing, China
wang.lei131@zte.com.cn

Guoying Zhang
China Academy of Telecom Research
No.52 Huayuan Bei Road, Beijing, China
zhangguoying@rictt.cn

Takehiro Tsuritani
KDDI R&D Laboratories Inc.
2-1-15 Ohara, Fujimino, Saitama, Japan
tsuri@kddilabs.jp

Lei Liu
KDDI R&D Laboratories Inc.
2-1-15 Ohara, Fujimino, Saitama, Japan

le-liu@kddilabs.jp

Eve Varma
Alcatel-Lucent
+1 732 239 7656
eve.varma@alcatel-lucent.com

Young Lee
Huawei

Jianrui Han
Huawei

Sharfuddin Syed
Infinera

Rajan Rao
Infinera

Marco Sosa
Infinera

Biao Lu
Infinera

Abinder Dhillon
Infinera

Felipe Jimenez Arribas
Telefonica I+D

Andrew G. Malis
Verizon

Adrian Farrel
Old Dog Consulting

Daniel King
Old Dog Consulting

Huub van Helvoort

9. Acknowledgments

The authors would like to thank Pete Anslow for his insights and clarifications.

10. References

10.1. Normative References

- [G.694.1] International Telecommunications Union, "ITU-T Recommendation G.694.1, Spectral grids for WDM applications: DWDM frequency grid", November 2012.
- [G.709] International Telecommunications Union, "ITU-T Recommendation G.709, Interfaces for the Optical Transport Network (OTN). ", March 2009.
- [G.800] International Telecommunications Union, "ITU-T Recommendation G.800, Unified functional architecture of transport networks.", February 2012.
- [G.805] International Telecommunications Union, "ITU-T Recommendation G.805, Generic functional architecture of transport networks.", March 2000.
- [G.8080] International Telecommunications Union, "ITU-T Recommendation G.8080/Y.1304, Architecture for the automatically switched optical network", 2012.
- [G.870] International Telecommunications Union, "ITU-T Recommendation G.870/Y.1352, Terms and definitions for optical transport networks", November 2012.
- [G.872] International Telecommunications Union, "ITU-T Recommendation G.872, Architecture of optical transport networks, draft v0.16 2012/09 (for discussion)", 2012.
- [G.959.1-2013] International Telecommunications Union, "Update of ITU-T Recommendation G.959.1, Optical transport network physical layer interfaces (to appear in July 2013)", 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008.
- [RFC6163] Lee, Y., Bernstein, G., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, April 2011.

10.2. Informative References

- [RFC4397] Bryskin, I. and A. Farrel, "A Lexicography for the Interpretation of Generalized Multiprotocol Label Switching (GMPLS) Terminology within the Context of the ITU-T's Automatically Switched Optical Network (ASON) Architecture", RFC 4397, February 2006.

Authors' Addresses

Oscar Gonzalez de Dios (editor)
Telefonica I+D
Don Ramon de la Cruz 82-84
Madrid 28045
Spain

Phone: +34913128832
Email: ogondio@tid.es

Ramon Casellas (editor)
CTTC
Av. Carl Friedrich Gauss n.7
Castelldefels Barcelona
Spain

Phone: +34 93 645 29 00
Email: ramon.casellas@cttc.es

Fatai Zhang
Huawei
Huawei Base, Bantian, Longgang District
Shenzhen 518129
China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Xihua Fu
ZTE
Ruanjian Avenue
Nanjing
China

Email: fu.xihua@zte.com.cn

Daniele Ceccarelli
Ericsson
Via Calda 5
Genova
Italy

Phone: +39 010 600 2512
Email: daniele.ceccarelli@ericsson.com

Iftekhhar Hussain
Infinera
140 Caspian Ct.
Sunnyvale 94089
USA

Phone: 408-572-5233
Email: ihussain@infinera.com

INTERNET-DRAFT
Intended Status: Standard Track
Expires: January 11, 2014

Khuzema Pithewan
Rajan Rao
Infinera
July 10, 2013

OSPF-TE extensions for MLNMRN based on OTN
draft-rao-ccamp-mlnmrn-otn-ospfte-ext-02.txt

Abstract

This document specifies OSPF extensions for multi-layer/multi-region where one of the regions is OTN.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2013 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
2.	OTN Layer Identification	3
3.	Interface Adjustment Capability Extensions for OTN	4
4.	SONET/SDH Layer Identification	5
5.	Interface Adjustment Capability Extensions for SONET/SDH	5
6	Procedure	6
7	Examples	6
8	IANA Considerations	8
9	Security Considerations	8
10	References	8
11	Authors' Addresses	9

1 Introduction

This document specifies the OSPF extensions required to work in multi-region networks involving OTN. The specification is based on the requirement as specified in RFC 5212. As per the said RFC, ISCD characterizes the information associated to one or more network layers. Same RFC also says that the information about the adjustment capabilities of the nodes in the network allow the path computation process to select an end-to-end multi-layer or multi-region path that includes links with different switching capabilities joined by LSRs that can adapt (i.e., adjust) the signal between the links. By inference, information about the adjustment capabilities should be able to identify a layer in ISCD, if ISCD specifies more than one layer.

RFC6001 specifies how to advertise adjustment capabilities between two switching regions. IACD definition has provision to extend it for a specific technology through Adjustment Capability Specific information (ACSI) field, if required. ACSI field can be used to identify a layer in the multi-layer ISCD. OTN being defined as multi-layer ISCD, the corresponding IACD needs to be extended to be able to carry layer identification so as to enable multi-layer/multi-region path computation. similar to OTN, legacy SONET/SDH also has fixed multiplexing hierarchy. Adaptation layer in SONET/SDH hierarchy requires identification information in ACSI, in order to adapt SONET/SDH signal to OTN and vice versa.

If the interface does not support de-multiplexing, then layer identification is not required. Hence, Layer ID sub-TLV is optional.

2. OTN Layer Identification

[GMPLS-OTN-OSPF] defines attributes that identifies a layer in multi-layer OTN ISCD. These attributes are part of Bandwidth sub-TLV in Switch capability specific information of ISCD. These attributes are reproduced here for completeness sake.

- * Signal Type: Layer for which bandwidth is being advertised.
- * Hierarchy : also called as multiplexing branch that specifies all the layers between server layer and signal type.
- * TSG : Time Slot Granularity

Absence of this sub-TLV for OTN means that the OTN ISCD doesn't support multiplexing.

3. Interface Adjustment Capability Extensions for OTN

RFC6001 defines IACD sub-TLV as follows. Please refer to the RFC for definition of individual fields of the sub-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								
Lower SC										Lower Encoding										Upper SC										Upper Encoding									
Max LSP Bandwidth at priority 0																																							
Max LSP Bandwidth at priority 1																																							
Max LSP Bandwidth at priority 2																																							
Max LSP Bandwidth at priority 3																																							
Max LSP Bandwidth at priority 4																																							
Max LSP Bandwidth at priority 5																																							
Max LSP Bandwidth at priority 6																																							
Max LSP Bandwidth at priority 7																																							
Adjustment Capability-specific information																																							
(variable)																																							

Adjustment Capability-specific information abbreviated as ACSI henceforth for OTN G.709v3 carries LayerID Sub-TLV which is defined as follows

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 = 1										R	Reserved										Length																		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Signal type										Num of stages										TSG	Res					Stage#1													
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
Stage#2											...										Stage#N										Padding								
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+										+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									

The definition & meaning of fields used in the above sub-TLV is same

as in bandwidth sub-TLV of ISCD as defined in [GMPLS-OTN-OSPF]. This LayerID sub-TLV is applicable only when one of the regions is OTN, which means either lower or upper SC and Encoding type MUST have Switch Cap as OTN-TDM and encoding type as G.709 ODUk.

R bit is used to make sense whether the Layer ID is for Lower region or upper region. 1 means upper region and 0 means lower.

The 8 priorities of the BW as defined in main IACD structure, is adjustment capability between the two regions where one of the region is identifies by LayerID sub-TLV.


4. SONET/SDH Layer Identification

G.707 defines the structure of SDH multiplexing hierarchy and RFC 4606 defines generalized label structure needed to fully specify SONET/SDH multiplexing hierarchy. This Label structure also referred as SUKLM structure identifies all the layers of the multiplexing hierarchy along with time slots. For the purpose of this draft, only layer identification is needed, hence each layer can be identified by a bit. Bit value 1 signifies presence of the layer and 0, its absence. 5 Bits, each representing one layer is sufficient to fully identify the SONET/SDH multiplexing hierarchy.

Absence of sub-TLV means that the SONET/SDH ISCD doesn't support multiplexing and needs only transparent mapping to other Interface.

5. Interface Adjustment Capability Extensions for SONET/SDH

Layer ID sub TLV for SONET/SDH is defined as follows



SUKLM bits signifies the presence of SONET/SDH layers and these bits together fully specifies the multiplexing hierarchy. Refer to Section 3 of RFC 4606 for full specification of SUKLM bits.

6 Procedure

A node advertising IACD for the bandwidth between regions where one or both of them are hierarchical i.e. OTN or SONET/SDH, MUST include the Layer ID sub-TLV as part of ACSI as defined above.

For multi-region path computation, the path computing node MUST look at the LayerID sub-TLV (in ACSI part of IACD) if lower/upper {SC,Enc] is {OTN-TDM,G.709ODUK} or {TDM,SONET/SDH} to identify the layer for correct layer for BW check.

7 Examples

This section exemplifies TLV values for various technology region combinations, where one of the region is OTN

a. Ethernet and OTN When upper region is Ethernet and lower region is OTN

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									
PSC-1										Ethernet										OTN-TDM										G.709 ODUk										
										Max LSP Bandwidth at priority 0																														
										/ / / / / / / / / / / / / / / /																														
										Max LSP Bandwidth at priority 7																														
Type = 1										0	Reserved										Length																			
Signal type										Num of stages					TSG					Res					Stage#1															
Stage#2										...					Stage#N										Padding															

b. OTN and FlexChannel

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																		
OTN-TDM										G.709 ODUk										SCSC										Lambda																			
										Max LSP Bandwidth at priority 0																																							
										/ / / / / / / / / / / / / / / /																																							
										Max LSP Bandwidth at priority 7																																							
Type = 1										1 Reserved										Length																													
Signal type										Num of stages										TSG										Res										Stage#1									
Stage#2										...										Stage#N										Padding																			

c. OTN and SONET/SDH

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																		
OTN-TDM										G.709 ODUk										TDM										Sonet/SDH																			
										Max LSP Bandwidth at priority 0																																							
										/ / / / / / / / / / / / / / / /																																							
										Max LSP Bandwidth at priority 7																																							
Type = 1										1 Reserved										Length																													
Signal type										Num of stages										TSG Res										Stage#1																			
Stage#2										...										Stage#N										Padding																			
Type = 2										0 Reserved										Length																													
S U K L M										Padding																																							

d. OTN and OTN

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										
OTN-TDM										G.709 ODUk										OTN-TDM										G.709 ODUk											
										Max LSP Bandwidth at priority 0																															
										/ / / / / / / / / / / / / / / /																															
										Max LSP Bandwidth at priority 7																															
Type = 1										0		Reserved										Length																			
Signal type										Num of stages										TSG		Res					Stage#1														
Stage#2										...												Stage#N					Padding														
Type = 1										1		Reserved										Length																			
Signal type										Num of stages										TSG		Res					Stage#1														
Stage#2										...												Stage#N					Padding														

8 IANA Considerations

TBD

9 Security Considerations

TBD

10 References

- [RFC5212] K. Shiomoto, Papadimitriou, D., JL. Le Roux, Vigoureux, M., Brungard, D., "Requirements for GMPLS-Based Multi-Layer and Multi-Region Networks (MLN/ MRN)", RFC 5212, July 2008.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.
- [RFC4606] E. Mannie, Perceval, D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, Aug 2006

[GMPLS-OTN-OSPF] Traffic Engineering Extensions to OSPF for
Generalized MPLS (GMPLS)

11 Authors' Addresses

Khuzema Pithewan
Infinera
140 Caspian Ct., Sunnyvale, CA 94089
Email: kpithewan@infinera.com

Rajan Rao
Infinera
140 Caspian Ct., Sunnyvale, CA 94089
Email: rrao@infinera.com

Network Working Group
Internet-Draft
Intended status: Informational

Fatai Zhang
Xian Zhang
Huawei
Ramon Casellas
CTTC
O. Gonzalez de Dios
Telefonica
D. Ceccarelli
Ericsson
June 29, 2013

Expires: December 28, 2013

GMPLS OSPF-TE Extensions in support of Flexible Grid DWDM Networks

draft-zhang-ccamp-flexible-grid-ospf-ext-02.txt

Abstract

This memo describes the OSPF-TE extensions in support of GMPLS control of networks that include devices that use the new flexible optical grid.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on December 28, 2013.

Copyright Notice

Copyright (c) 2013 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	2
2. Terminology	3
2.1. Conventions Used in this Document	3
3. Requirements for Flexi-grid Routing	3
3.1. Available Frequency Ranges.....	4
3.2. Application Compliance Considerations	5
3.3. Comparison with Fixed-grid DWDM Links	5
4. Extensions	6
4.1. Available Labels Set Sub-TLV.....	6
4.1.1. Inclusive/Exclusive Label Range	7
4.1.2. Inclusive/Exclusive Label Lists	7
4.1.3. Bitmap	7
4.2. Extensions to Port Label Restriction sub-TLV	7
4.3. Examples for Available Label Set Sub-TLV	8
5. IANA Considerations	9
6. Security Considerations.....	9
7. References	10
7.1. Normative References.....	10
7.2. Informative References.....	10
8. Authors' Addresses	11
9. Contributors' Addresses.....	11

1. Introduction

[G.694.1] defines the Dense Wavelength Division Multiplexing (DWDM) frequency grids for Wavelength Division Multiplexing (WDM) applications. A frequency grid is a reference set of frequencies used to denote allowed nominal central frequencies that may be used for defining applications. The channel spacing is the frequency spacing between two allowed nominal central frequencies. All of the

wavelengths on a fiber should use different central frequencies and occupy a fixed bandwidth of frequency.

Fixed grid channel spacing is selected from 12.5 GHz, 25 GHz, 50 GHz, 100 GHz and integer multiples of 100 GHz. But [G.694.1] also defines "flexible grids", also known as "flexi-grid". The terms "frequency slot" (i.e., the frequency range allocated to a specific channel and unavailable to other channels within a flexible grid) and "slot width" (i.e., the full width of a frequency slot in a flexible grid) are used to define a flexible grid.

[FLEX-FWK] defines a framework and the associated control plane requirements for the GMPLS based control of flexi-grid DWDM networks.

[RFC6163] provides a framework for GMPLS and Path Computation Element (PCE) control of Wavelength Switched Optical Networks (WSONs), and [WSON-OSPF] defines the requirements and OSPF-TE extensions in support of GMPLS control of a WSON.

[FLEX-SIG] describes requirements and protocol extensions for signaling to set up LSPs in networks that support the flexi-grid, and this document complements [FLEX-SIG] by describing the requirement and extensions for OSPF-TE routing in a flexi-grid network.

2. Terminology

For terminology related to flexi-grid, please consult [FLEX-FWK] and [G.694.1].

2.1. Conventions Used in this Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

3. Requirements for Flexi-grid Routing

The architecture for establishing LSPs in a Spectrum Switched optical Network (SSON) is described in [FLEX-FWK].

An flexi-LSP occupies a specific frequency slot, i.e. a range of frequencies. The process of computing a route and the allocation of a frequency slot is referred to as RSA (Routing and Spectrum Assignment). [FLEX-FWK] describes three types of architectural approaches to RSA: combined RSA; separated RSA; and distributed SA. The first two approaches among them could be called "centralized SA"

Figure 1 shows two adjacent frequency slots on a link. The highest frequency of frequency slot 1 denoted by $n=2$ is the lowest frequency of slot 2. In this example, it means that the frequency range from $n=-2$ to $n=10$ is occupied and is unavailable to other flexi-LSPs.

Hence, in order to clearly show which LSPs can be supported and what frequency slots are unavailable, the available frequency ranges should be advertised by the routing protocol for the flexi-grid DWDM links. A set of non-overlapping available frequency ranges should be disseminated in order to allow efficient resource management of flexi-grid DWDM links and RSA procedures which are described in section 5.7 of [FLEX-FWK].

3.2. Application Compliance Considerations

As described in [G.694.1], devices or applications that make use of the flexi-grid may not be capable of supporting every possible slot width or position (i.e., central frequency). In other words, applications or implementations may be defined where only a subset of the possible slot widths and positions are required to be supported.

For example, an application could be defined where the nominal central frequency granularity is 12.5 GHz (by only requiring values of n that are even) and that only requires slot widths as a multiple of 25 GHz (by only requiring values of m that are even).

Hence, in order to support all possible applications and implementations the following information should be advertised for a flexi-grid DWDM link:

- o Central frequency granularity: a multiplier of 6.25 GHz.
- o Slot width granularity: a multiplier of 12.5 GHz.
- o Slot width range: two multipliers of 12.5GHz, each indicate the minimal and maximal slot width supported by a port respectively.

The combination of slot width range and slot width granularity can be used to determine the slot widths set supported by a port.

3.3. Comparison with Fixed-grid DWDM Links

In the case of fixed-grid DWDM links, each wavelength has a pre-defined central frequency and each wavelength has the same frequency range (i.e., there is a uniform channel spacing). Hence all the wavelengths on a DWDM link can be identified uniquely simply by

giving it an identifier (such as the central wavelength [RFC6205]), and the status of the wavelengths (available or not) can be advertised through a routing protocol.

Figure 2 shows a link that supports a fixed-grid with 50 GHz channel spacing. The central frequencies of the wavelengths are pre-defined by values of 'n' and each wavelength occupies a fixed 50 GHz frequency range as described in [G.694.1].




Figure 2 - A Link Supports Fixed Wavelengths with 50 GHz Channel Spacing

Unlike the fixed-grid DWDM links, on a flexi-grid DWDM link the slot width of the frequency slot are flexible as described in section 2.1. That is, the value of m in the formula is uncertain before a frequency slot is actually allocated. For this reason, the available frequency slot/ranges need to be advertised for a flexi-grid DWDM link instead of the specific "wavelengths" that are sufficient for a fixed-grid link.

4. Extensions

As described in [FLEX-FWK], the network connectivity topology constructed by the links/nodes and node capabilities are the same as for WSON, and can be advertised by the GMPLS routing protocols (refer to section 6.2 of [RFC6163]). In the flexi-grid case, the available frequency ranges instead of the specific "wavelengths" are advertised for the link. This section defines the GMPLS OSPF-TE extensions in support of advertising the available frequency ranges for flexi-grid DWDM links.

4.1. Available Labels Set Sub-TLV

As described in section 3.1, the available frequency ranges other than the available frequency slots should be advertised for the flexi-grid DWDM links. The label encoding defined in [FLEX-LBL] is

used to encode the label field in Available Labels Set sub-TLV [GEN-Encode].

4.1.1. Inclusive/Exclusive Label Range

The inclusive/exclusive label ranges format of the Available Labels Set sub-TLV defined in [GEN-ENCODE] can be used for specifying the frequency ranges of the flexi-grid DWDM links.

Note that multiple Available Labels Set sub-TLVs may be needed if there are multiple discontinuous frequency ranges on a link.

4.1.2. Inclusive/Exclusive Label Lists

The inclusive/exclusive label lists format of Available Labels Set sub-TLV defined in [GEN-ENCODE] can be used for specifying the available central frequencies of flexi-grid DWDM links.

4.1.3. Bitmap

The bitmap format of Available Labels Set sub-TLV defined in [GEN-ENCODE] can be used for specifying the available central frequencies of the flexi-grid DWDM links.

Each bit in the bit map represents a particular central frequency with a value of 1/0 indicating whether the central frequency is in the set or not. Bit position zero represents the lowest central frequency and corresponds to the base label, while each succeeding bit position represents the next central frequency logically above the previous.

4.2. Extensions to Port Label Restriction sub-TLV

As described in Section 3.2, a port that supports flexi-grid may support only a restricted subset of the full flexible grid. The Port Label Restriction sub-TLV is defined in [GEN-ENCODE] and [GEN-OSPF]. It can be used to describe the label restrictions on a port. A new restriction type, the flexi-grid Restriction Type, is defined here to specify the restrictions on a port to support flexi-grid.

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
MatrixID								RstType = TBA								Reserved															
C.F.G								S.W.G								Min Width								Max Width							

MatrixID (8 bits): As defined in [GEN-ENCODE].

RstType (Restriction Type, 8 bits): Takes the value (TBD) to indicate the restrictions on a port to support flexi-grid.

C.F.G (Central Frequency Granularity, 8 bits): A positive integer. Its value indicates the multiple of 6.25 GHz in terms of central frequency granularity.

S.W.G (Slot Width Granularity, 8 bits): A positive integer. Its value indicates the multiple of 12.5 GHz in terms of slot width granularity.

Min Width (8 bits): A positive integer. Its value indicates the multiple of 12.5 GHz in terms of the supported minimal slot width.

Max Width (8 bits): A positive integer. Its value indicates the multiple of 12.5 GHz in terms of the supported maximal slot width.

4.3. Examples for Available Label Set Sub-TLV

Figure 3 shows an example of available frequency range of a flexi-grid DWDM link.




Figure 3 - Flexi-grid DWDM Link

The symbol '+' represents the allowed nominal central frequency. The symbol "--" represents a 6.25 GHz frequency unit. The number on the top of the line represents the 'n' in the frequency calculation formula $(193.1 + n * 0.00625)$. The nominal central frequency is 193.1 THz when n equals zero.

Assume that the central frequency granularity is 6.25GHz, the label set can be encoded as follows:

Inclusive Label Range:

- o Start Slot = -2;
- o End Slot = 8.

The available central frequencies (-1, 0, 1, 2, 3, 4, 5, 6, 7) can be deduced by the Inclusive Label Range, because the Central Frequency Granularity is 6.25 GHz.

Inclusive Label Lists:

- o List Entry 1 = slot -1;
- o List Entry 2 = slot 0;
- o List Entry 3 = slot 1;
- o List Entry 4 = slot 2;
- o List Entry 5 = slot 3;
- o List Entry 6 = slot 4;
- o List Entry 7 = slot 5;
- o List Entry 8 = slot 6;
- o List Entry 9 = slot 7.

Bitmap:

- o Base Slot = -1;
- o Bitmap = 1111111111(padded out to a full multiple of 32 bits)

5. IANA Considerations

[GEN-OSPF] defines the Port label Restriction sub-TLV of OSPF TE Link TLV. It also creates a registry of values of the Restriction Type field of that sub-TLV

IANA is requested to assign a new value from that registry as follows:

Value	Meaning	Reference
TBD	Flexi-grid restriction	[This.I-D]

6. Security Considerations

This document does not introduce any further security issues other than those discussed in [RFC3630], [RFC4203].

7. References

7.1. Normative References

- [RFC2119] S. Bradner, "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.
- [G.694.1] ITU-T Recommendation G.694.1 (revision 2), "Spectral grids for WDM applications: DWDM frequency grid", February 2012.
- [GEN-ENCODE] Bernstein, G., Lee, Y., Li, D., and W. Imajuku, "General Network Element Constraint Encoding for GMPLS Controlled Networks", draft-ietf-ccamp-general-constraint-encode, work in progress.
- [GEN-OSPF] Fatai Zhang, Y. Lee, Jianrui Han, G. Bernstein and Yunbin Xu, " OSPF-TE Extensions for General Network Element Constraints ", draft-ietf-ccamp-gmpls-general-constraints-ospf-te, work in progress.
- [RFC6205] T. Otani and D. Li, "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, March 2011.
- [FLEX-LBL] King, D., Farrel, A. and Y. Li, "Generalized Labels for the Flexi-Grid in Lambda Switch Capable (LSC) Label Switching Routers", draft-farrkingel-ccamp-flexigrid-lambda-label, work in progress.

7.2. Informative References

- [RFC6163] Y. Lee, G. Bernstein and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, April 2011.
- [FLEX-SIG] F.Zhang et al, "RSVP-TE Signaling Extensions in support of Flexible-grid", draft-zhang-ccamp-flexible-grid-rsvp-te-ext, work in progress.
- [FLEX-FWK] Gonzalez de Dios, O., Casellas R., Zhang, F., Fu, X., Ceccarelli, D., and I. Hussain, "Framework and Requirements for GMPLS based control of Flexi-grid DWDM networks", draft-ogrcetal-camp-flexi-grid-fwk, work in progress.

[WSON-OSPF] Y. Lee and G. Bernstein, "GMPLS OSPF Enhancement for Signal and Network Element Compatibility for Wavelength Switched Optical Networks ", draft-ietf-ccamp-wson-signal-compatibility-ospf, work in progress.

8. Authors' Addresses

Fatai Zhang
Huawei Technologies
Email: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
Email: zhang.xian@huawei.com

Ramon Casellas, Ph.D.
CTTC
Spain
Phone: +34 936452916
Email: ramon.casellas@cttc.es

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain
Phone: +34 913374013
Email: ogondio@tid.es

Daniele Ceccarelli
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy
Email: daniele.ceccarelli@ericsson.com

9. Contributors' Addresses

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Network Working Group
Internet-Draft
Intended status: Standards Track

Fatai Zhang
Xian Zhang
Huawei
Adrian Farrel
Old Dog Consulting
Oscar Gonzalez de Dios
Telefonica
D. Ceccarelli
Ericsson
June 25, 2013

Expires: December 24, 2013

RSVP-TE Signaling Extensions in support of Flexible Grid

draft-zhang-ccamp-flexible-grid-rsvp-te-ext-02.txt

Abstract

This memo describes the extensions to RSVP-TE signaling to support Label Switched Paths in a GMPLS-controlled network that includes devices using the new flexible optical grid.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on December 24, 2013.

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. Conventions used in this document	3
3. Requirements for Flexible Grid Signaling	3
3.1. Slot Width	4
3.2. Frequency Slot	4
4. Protocol Extensions	5
4.1. Traffic Parameters.....	5
4.1.1. Applicability to Fixed Grid Networks	6
4.2. Generalized Label.....	6
4.3. Signaling Procedures.....	6
5. IANA Considerations	7
5.1. RSVP Objects Class Types.....	7
6. Manageability Considerations.....	7
7. Security Considerations.....	7
8. References	7
8.1. Normative References.....	7
8.2. Informative References.....	8
9. Contributors' Address.....	8
10. Authors' Addresses	9

1. Introduction

[G.694.1] defines the Dense Wavelength Division Multiplexing (DWDM) frequency grids for Wavelength Division Multiplexing (WDM) applications. A frequency grid is a reference set of frequencies used to denote allowed nominal central frequencies that may be used for defining applications. The channel spacing is the frequency spacing between two allowed nominal central frequencies. All of the wavelengths on a fiber use different central frequencies and occupy a fixed bandwidth of frequency.

Fixed grid channel spacing is selected from 12.5 GHz, 25 GHz, 50 GHz, 100 GHz and integer multiples of 100 GHz. But [G.694.1] also defines "flexible grids", known as "flexi-grid". The terms "frequency slot" (i.e. the frequency range allocated to a specific channel and unavailable to other channels within a flexible grid)" and "slot width" (i.e. the full width of a frequency slot in a flexible grid) are introduced to define a flexible grid.

[FLEX-FWK] defines a framework and the associated control plane requirements for the GMPLS based control of flexi-grid DWDM networks.

[RFC6163] provides a framework for GMPLS and Path Computation Element (PCE) control of Wavelength Switched Optical Networks (WSONs), and [WSON-SIG] describes the requirements and protocol extensions for signaling to set up Label Switched Paths (LSPs) in WSONs.

This document describes the additional requirements and protocol extensions for signaling to set up LSPs in networks that support the flexi-grid.

2. Terminology

For terminology related to flexi-grid, please refer to [FLEX-FWK] and [G.694.1].

2.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

3. Requirements for Flexible Grid Signaling

The architecture for establishing LSPs in a flexi-grid network is described in [FLEX-FWK].

A optical spectrum LSP occupies a specific frequency slot, i.e. a range of frequencies. The process of computing a route and the allocation of a frequency slot is referred to as RSA (Routing and Spectrum Assignment). [FLEX-FWK] describes three types of architecture approaches to RSA: combined RSA, separated RSA and distributed SA. The first two approaches are referred to as ''centralized SA'', because both routing and spectrum (frequency slot) assignment are performed by centralized entity before the signaling procedure.

In the case of centralized SA, the assigned frequency slot is specified in the Path message during LSP setup. In the case of distributed SA, the slot width of the flexi-grid LSP is specified in the Path message, allowing the involved network elements to select the frequency slot to be used.

If the capability of switching or converting the whole optical spectrum allocated to an optical spectrum LSP is not available at nodes along the path of the LSP, the LSP is subject to the Optical ''Spectrum Continuity Constraint'', as described in [FLEX-FWK].

The remainder of this section states the additional requirements for signaling in a flexi-grid network.

3.1. Slot Width

The slot width is an end-to-end parameter representing how much frequency resource is requested for a flexi-grid LSP. It is equivalent of optical bandwidth, although the amount of bandwidth associated with a slot width will depend on the encoding.

Different LSPs may request different amounts of frequency resource in flexible grid networks, so the slot width needs to be carried in the signaling message during LSP establishment. This enables the nodes along the LSP to know how much frequency resource has been requested (in a Path message) and has been allocated (by a Resv message) for the LSP.

3.2. Frequency Slot

The frequency slot information identifies which part of the frequency spectrum is allocated on each link for a flexi-LSP.

This information is required in Resv message to indicate, hop-by-hop, the central frequency of the allocated resource. In combination with the slot width indicated in a Resv message (see Section 3.1) the central frequency carried in a Resv message identifies the resources reserved for the LSP (known as the frequency slot).

The frequency slot can be represented by the two parameters as follows:

$$\text{Frequency slot} = [(\text{central frequency}) - (\text{slot width})/2] \sim [(\text{central frequency}) + (\text{slot width})/2]$$

As is common with other resource identifiers (i.e., labels) in GMPLS signaling, it must be possible for the head-end LSP to suggest or require the central frequency to be used for the LSP. Furthermore, for bidirectional LSPs, the Path message must be able to specify the central frequency to be used for reverse direction traffic.

As described in [G.694.1], the allowed frequency slots for the flexible DWDM grid have a nominal central frequency (in THz) defined by:

$$193.1 + n * 0.00625$$

where n is zero or a positive or negative integer.

The slot width (in GHz) is defined as:

$$12.5 * m$$

where m is a positive integer.

It is possible that implementing a subset of the possible slot widths and central frequencies are supported. For example, an implementation could be built where the nominal central frequency granularity is 12.5 GHz (by only requiring values of n that are even) and that only supports slot widths as a multiple of 25 GHz (by only allowing values of m that are even).

Further details can be found in [FLEX-FWK].

4. Protocol Extensions

This section defines the extensions to RSVP-TE signaling for GMPLS [RFC3473] to support flexible grid networks.

4.1. Traffic Parameters

In RSVP-TE, the SENDER_TSPEC object in the Path message indicates the requested resource reservation. The FLOWSPEC object in the Resv message indicates the actual resource reservation.

As described in Section 3.1, the slot width represents how much frequency resource is requested for a flexi-grid LSP. That is, it describes the end-to-end traffic profile of the LSP. Therefore, the traffic parameters for a flexi-grid LSP encode the slot width.

This document defines new C-Types for the SENDER_TSPEC and FLOWSPEC objects to carry Spectrum Switched Optical Network (SSON) traffic parameters:

SSON SENDER_TSPEC: Class = 12, C-Type = TBD1.

SSON FLOWSPEC: Class = 9, C-Type = TBD2.

The SSON traffic parameters carried in both objects have the same format as shown in Figure 1.




Figure 1: The SSON Traffic Parameters

m (8 bits): the slot width is specified by $m \times 12.5$ GHz.

The Reserved bits MUST be set to zero and ignored upon receipt.

4.1.1. Applicability to Fixed Grid Networks

Note that the slot width (i.e., traffic parameters) of a fixed grid defined in [G.694.1] can also be specified by using the SSON traffic parameters. The fixed grid channel spacings (12.5 GHz, 25 GHz, 50 GHz, 100 GHz and integer multiples of 100 GHz) are also the multiple of 12.5 GHz, so the m parameter can be used to represent these slot widths.

Therefore, it is possible to consider using the new traffic parameter object types in common signaling messages for flexi-grid and legacy DWDM networks.

4.2. Generalized Label

In the case of a flexible grid network, the labels that have been requested or allocated as signaled in the RSVP-TE objects are encoded as described in [FLEX-LBL]. This new label encoding can appear in any RSVP-TE object or sub-object that can carry a label.

As noted in Section 4.2 of [FLEX-LBL], the m parameter forms part of the label as well as part of the traffic parameters.

4.3. Signaling Procedures

There are no differences between the signaling procedure described for LSP control in [FLEX-FWK] and those required for use in a fixed-grid network [WSN-SIG]. Obviously, the TSpec, FlowSpec and label formats described in Section 3 are used. The signaling procedures for distributed SA and centralized SA can be applied.

5. IANA Considerations

5.1. RSVP Objects Class Types

This document introduces two new Class Types for existing RSVP objects. IANA is requested to make allocations from the "Resource ReSerVation Protocol (RSVP) Parameters" registry using the "Class Names, Class Numbers, and Class Types" sub-registry.

Class Number	Class Name	Reference
-----	-----	-----
9	FLOWSPEC	[RFC2205]

Class Type (C-Type):

(TBD2) SSON FLOWSPEC [This.I-D]

Class Number	Class Name	Reference
-----	-----	-----
12	SENDER_TSPEC	[RFC2205]

Class Type (C-Type):

(TBD1) SSON SENDER_TSPEC [This.I-D]

IANA is requested to assign the same value for TBD1 and TBD2, and a value of 8 is suggested.

6. Manageability Considerations

This document makes minor modifications to GMPLS signaling, but does not change the manageability considerations for such networks. Clearly, protocol analysis tools and other diagnostic aids (including logging systems and MIB modules) will need to be enhanced to support the new traffic parameters and label formats.

7. Security Considerations

This document introduces no new security considerations to [RFC3473].

8. References

8.1. Normative References

[RFC2119] S. Bradner, "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.

- [RFC3473] L. Berger, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [G.694.1] ITU-T Recommendation G.694.1 (revision 2), ''Spectral grids for WDM applications: DWDM frequency grid'', February 2012.
- [FLEX-LBL] King, D., Farrel, A. and Y. Li, ''Generalized Labels for the Flexi-Grid in Lambda Switched Capable (LSC) Label Switching Routers'', draft-farrkingel-ccamp-flexigrid-lambda-label, work in progress.

8.2. Informative References

- [RFC2205] Braden, R., Zhang L., Berson, S., Herzog, S. and S. Jamin, ''Resource ReSerVation Protocol (RSVP) - Version 1, Functional Specification'', RFC2205, September 1997.
- [RFC6163] Y. Lee, G. Bernstein and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, April 2011.
- [FLEX-FWK] Gonzalez de Dios, O., Casellas R., Zhang, F., Fu, X., Ceccarelli, D., and I. Hussain, ''Framework and Requirements for GMPLS based control of Flexi-grid DWDM networks'', draft-ogrcetal-cammp-flexi-grid-fwk, work in progress.
- [WSON-SIG] G. Bernstein, Sugang Xu, Y. Lee, G. Martinelli and Hiroaki Harai, "Signaling Extensions for Wavelength Switched Optical Networks", draft-ietf-ccamp-wson-signaling, work in progress.

9. Contributors' Address

Ramon Casellas
CTTC
Av. Carl Friedrich Gauss n7
Castelldefels, Barcelona 08860
Spain

Email: ramon.casellas@cttc.es

Felipe Jimenez Arribas
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain
Email: felipej@tid.es

Yi Lin
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972914
Email: yi.lin@huawei.com

10. Authors' Addresses

Fatai Zhang
Huawei Technologies
Email: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
Email: zhang.xian@huawei.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain

Phone: +34 913374013
Email: ogondio@tid.es

Daniele Ceccarelli
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy
Email: daniele.ceccarelli@ericsson.com

Intellectual Property

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions.

For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Full Copyright Statement

Copyright (c) 2013 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.

Network Working Group
Internet Draft
Category: Standards Track

Fatai Zhang
Xian Zhang
Huawei
O. Gonzalez de Dios
Telefonica Investigacion y Desarrollo
C. Margaria. C
Coriant
July 11, 2013

Expires: January 10, 2014

GMPLS-based Hierarchy LSP Creation
in Multi-Region and Multi-Layer Networks

draft-zhang-ccamp-gmpls-h-lsp-mln-05.txt

Abstract

This specification describes the hierarchical LSP creation models in the Multi-Region and Multi-Layer Networks (MRN/MLN), and provides the extensions to the existing protocol mechanisms described in [RFC4206], [RFC6107] and [RFC6001] to create a hierarchical LSP in multiple layer networks.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 10, 2014.

Table of Contents

1. Introduction	2
1.1. Conventions used in this document	3
2. Provisioning of FA-LSP in Server Layer Network	3
2.1. Selection of Switching Layers.....	3
2.2. Selection of Switching Granularity Levels	4
2.3. Selection of Adaptation Capabilities	6
3. Signaling Requirements for Server Layer Selection	7
3.1. Model 1: Pre-provisioning of FA-LSP	8
3.2. Model 2: Signaling triggered server layer path computation and setup	9
3.3. Model 3: Signaling triggered server layer path, with explicit server path	9
4. Signaling Extensions ERO Sub-Object	10
4.1. SERVER_LAYER_INFO ERO Subobject	10
4.2. Processing of SERVER_LAYER_INFO sub-object	12
4.3. Alternative Encoding Solutions	12
5. Security Considerations.....	13
6. IANA Considerations	13
7. Acknowledgments	13
8. References	13
8.1. Normative References.....	13
8.2. Informative Reference.....	14
9. Authors' Addresses	15

1. Introduction

Networks may comprise multiple layers which have different switching technologies or different switching granularity levels. The GMPLS technology is required to support control of such network.

[RFC5212] defines the concept of MRN/MLN and describes the framework and requirements of GMPLS controlled MRN/MLN. The GMPLS extension for MRN/MLN, including routing and signaling aspects, is described in [RFC6001].

[RFC4206] and [RFC6107] describe how to set up a hierarchical LSP passing through multi-layer networks and how to advertise the forwarding adjacency LSP (FA-LSP) created in the server layer network as a TE link via GMPLS signaling and routing protocols.

Based on these existing standards, this document further describes the provisioning of a FA-LSP when the region-edge nodes support

multiple interface switching capabilities and/or multiple switching granularities and/or adaptation functions, and then provides the extensions to the RSVP-TE protocol in order to set up a hierarchical LSP according to the modes of hierarchical LSP provisioning.

1.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Provisioning of FA-LSP in Server Layer Network

2.1. Selection of Switching Layers

As described in [RFC5212], the edge node of a region always has multiple Interface Switching Capabilities (ISCs), i.e., it contains multiple matrices which may be connected to each other by internal links. Nodes with multiple ISCs are further classified as "simplex" or "hybrid" nodes by [RFC5212] and [RFC5339], where the simplex node advertises several TE links each with a single ISC value carried in its ISCD sub-TLV, while the hybrid node advertises a single TE link containing more than one ISCD each with a different ISC value. An example of a hybrid node with a link having multiple ISCs is shown in Figure 1, copied from [RFC5339].




Figure 1 - Hybrid node (Copied from [RFC5339])

In the case where a edge node of a region is a hybrid node, selection of which server layer to create the FA-LSP is necessary.

Figure 2 shows an multi-layer network, where node B and C are region edge nodes having three switching matrices which support, for instance, PSC, TDM and WDM switching, respectively. The three switching matrices are connected to each other by the internal links. Both the link between B and E and the link between E and C support TDM and WDM switching capabilities.




Figure 2 - MLN with multiple ISCs at edge node

As can be seen in Figure 2, there are two choices when providing FA in the PSC layer network between node B and C: one is creating a FA-LSP with TDM switching matrix through node B, E and C, the other is creating a FA-LSP with WDM switching matrix through node B, E and C.

[RFC6001] introduces a new SC (Switching Capability) sub-object into the XRO (ref. to [RFC4874]). This sub-object is used to indicate which switching capability is not expected to be used. When one of the switching capabilities is selected, the SC sub-object can be included in the message to exclude all other SCs.

2.2. Selection of Switching Granularity Levels

Even in the case where the edge node only has one switching capability in the server layer, there may be still multiple choices for the server layer network to set up a FA-LSP to provide new FA in the client layer network. This is because the server layer network may have the capability of providing different switching granularity levels for the FA-LSP.




Figure 3a - Multiple switching granularities in server layer

Figure 3a shows an example multi-region network, where the edge node B and C have PSC and TDM switching matrices, and where the TDM switching matrix supports ODU1, ODU2 and ODU3 switching levels. Therefore, when an FA between node B and C in the PSC layer network is needed, either of ODU1, ODU2 or ODU3 connection (FA-LSP) can be created in the TDM layer network.




Figure 3b - TDM nested LSP provisioning

Figure 3b is another example multi-layer network within the same region. When there is a need to set up an FA between node B and C for the client layer ODU0 connection, the server layer has multiple

choices, e.g., ODU1 or ODU2 or ODU3, for the FA-LSP if the multi-stage multiplexing is supported at node B and C.




Figure 3c - PSC nested LSP provisioning

Figure 3c is a third example showing an LSP nesting scenario in a PSC signal-layer network (e.g., an MPLS-TP network). A PSC tunnel passing through node B, E and C is requested to carry the client layer LSP. There are multiple choices of the bandwidth of the tunnel, on the premise that the bandwidth of the FA-LSP is equal to or larger than the client layer LSP.

The selection of server layer switching matrix and switching granularity is based on both policy and bandwidth resources. The selection can be performed by a planning tool and/or NMS/PCE/VNTM (Virtual Network Topology Manager, see [RFC5623]) and/or the network node.

2.3. Selection of Adaptation Capabilities

Adaptation function also needs to be selected when creating the server layer connection. This is because the edge nodes may support multiple adaptation functions.




Figure 4 - Selection of adaptation function

For example, in Figure 4, edge node B supports two adaptation functions, i.e., adaptation_function_A and adaptation_function_B, while edge node C only supports adaptation_function_A. In this case, only adaptation_function_A can be used for the server layer connection.

The Call procedure ([RFC4974]) may be used between edge node B and C to negotiate and determine the adaptation function for the server layer if the Call function is supported.

3. Signaling Requirements for Server Layer Selection

[RFC5623], the framework of PCE-based MLN, provides the models of cross-layer LSP path computation and creation, which are listed below:

- Inter-Layer Path Computation Models:
 - o Single PCE
 - o Multiple PCE with inter-PCE
 - o Multiple PCE without inter-PCE
- Inter-Layer Path Control Models:
 - o PCE-VNTM cooperation

- o Higher-layer signaling trigger
- o NMS-VNTM cooperation (integrated flavor)
- o NMS-VNTM cooperation (separate flavor)

This section keeps alignment with [RFC5623] except that the restriction of using a PCE for path computation is not necessary (i.e., other element, such as a network node, may also have path computation capability).

In this document, those models in [RFC4206] are mapped into 3 models on the viewpoint of signaling:

- Model 1: Pre-provisioning of FA-LSP
- Model 2: Signaling triggered server layer path computation and setup
- Model 3: Signaling triggered server layer path, with explicit server path.

3.1. Model 1: Pre-provisioning of FA-LSP

In this model, the FA-LSP in the server layer is created before initiating the signaling of the client layer LSP. Two typical scenarios using this model are:

- Network planning and building at the stage of client network initialization.
- NMS/VNTM triggering the creation of FA-LSP when computing the path of client layer LSP. The path control models of PCE-VNTM cooperation and NMS-VNTM cooperation (both integrated and separate flavor) in [RFC5623] belong to this scenario.

In such case, the server layer selection and path computation is performed by planning tool or NMS/PCE/VNTM or the edge node. The signaling of client layer LSP and server layer FA-LSP are separated. The normal LSP creation procedures ([RFC3471] and [RFC3473]) are followed to set up these two LSPs and no new extension is required.

3.2. Model 2: Signaling triggered server layer path computation and setup

In this model, the source node of client layer LSP only computes the route within its own layer network. When the signaling of the client layer LSP reaches at the region edge node, the edge node performs server layer FA-LSP path computation and then creates the FA-LSP. When a PCE is introduced to perform path computation in each layer of the multi-layer network, this model is the same as the model of "higher-layer signaling trigger with Multiple PCE without inter-PCE" in [RFC5623].

In such case, the edge node will receive the client layer PATH message with a loose ERO indicating an FA is requested, and may perform the server layer selection (e.g., through the server layer PCE or the VNTM) and then compute and set up the FA-LSP. The signaling procedure of client layer LSP and server layer FA-LSP is described in detail in [RFC4206] and [RFC6107].

It's possible that the source node of the client layer LSP selects the server layer SC and/or granularity and/or adaptation function when performing path computation in the client layer, and requests or suggests the edge node to use an appointed server layer to create the FA-LSP.

In this case, the XRO including SC sub-object ([RFC6001]) is adopted for the server layer SC exclusion, which can be used indirectly to select server layer SC. Such solution is not straightforward enough. Furthermore it cannot be used for the selection of server layer granularity and adaptation function. Therefore, new extensions for the selection of server layer SC, switching granularity and adaptation function are required.

3.3. Model 3: Signaling triggered server layer path, with explicit server path

In this model, the source node of the client layer LSP performs a full path computation including the client layer and the server layer routes. The server layer FA-LSP creation is triggered at the edge node by the client layer LSP signaling. When a PCE is introduced to perform path computation in the multi-layer network, this model is the same as the model of "Higher-layer signaling trigger with Single PCE" or "Higher-layer signaling trigger with Multiple PCE with inter-PCE" in [RFC5623].

In such case, the server layer selection and server layer path computation is performed at the source node of the client layer LSP (e.g., through VNTM or PCE), but not at the edge node.

In [RFC4206], the ERO which contains the list of nodes and links (including the client layer and server layer) along the path is used in the client layer PATH message. The edge node can find out the tail end of the FA-LSP based on the switching capability of the node using the IGP database (see session 6.2 of [RFC 4206]).

Similar to the problem of model 2, the edge node is not aware of which switching granularity and which adaptation function to be selected for the FA-LSP because the ERO and/or XRO do not contain such information. Therefore, the edge node may not be able to create the FA-LSP, or may select another switching granularity by itself which is different from the one selected previously at the source node, which makes the creation of hierarchy LSP out of control.

Therefore, new extensions for the selection of server layer SC, switching granularity and adaptation function are also required in this model.

4. Signaling Extensions ERO Sub-Object

4.1. SERVER_LAYER_INFO ERO Subobject

In order to solve the problems described in the previous sections, a new sub-object named SERVER_LAYER_INFO sub-object is introduced in this document, which is carried in the ERO and is used to explicitly indicate which server layer to create the FA-LSP.

The SERVER_LAYER_INFO sub-object is put immediately after the node or link (interface) address sub-object, indicating the related node is a region edge node on the LSP in the ERO.

The format of the SERVER_LAYER_INFO sub-object is shown 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
+-----+-----+-----+-----+-----+-----+-----+-----+
|L|      Type      |      Length      |M|      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
| LSP Enc. Type |Switching Type |      G-PID      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Traffic Spec Length      | TSpec Type      | Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Traffic Parameters                                     |
~                                                                 ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

[Editor's note: the encoding is still under discussion.]

- L bit: MUST be zero and MUST be ignored when received.
- Type: The SERVER_LAYER_INFO sub-object has a type of xx (TBD).
- Length: The total length of the sub-object in bytes, including the Type and Length fields. The value of this field is always a multiple of 4.
- M (Mandatory) bit: When set, it means the edge node MUST set up the FA-LSP in the appointed server layer; otherwise, the appointed server layer is suggested and the edge node may select other server layer by local policy.
- LSP Encoding Type, Switching Type and G-PID: These 3 fields are used to point out which switching layer is requested to set up the FA-LSP. The values of these 3 fields are inherited from the Generalized Label Request Object in GMPLS signaling, referring to [RFC3471], [RFC3473] and other related standards and drafts. Note that G-PID can be used to indicate the payload type of the server layer (i.e., the client signal) as well as the adaptation function for adapting the client signal into the server layer FA-LSP.
- Traffic Spec Length, TSpec Type, Traffic Parameters: The traffic parameters field is used to indicate the switching granularity of the FA-LSP. The format of this field depends on the TSpec Type Traffic Spec Length and is consistent with the existing standards and drafts. For example, the traffic parameters of Ethernet, SONET/SDH and OTN are defined in [RFC6003], [RFC4606] and [OTN-ctrl] respectively.

4.2. Processing of SERVER_LAYER_INFO sub-object

As described in RFC3209 and RFC3473 the ERO is managed as a sub-object list. The SERVER_LAYER_INFO sub-object MUST be appended after the existing sub-object defined in [RFC3209], [RFC3473], [RFC3477], [RFC4873], [RFC4874], [RFC5520] and [RFC5553] TBD:extensions.

When a node receives a PATH message containing ERO and finds that there is a SERVER_LAYER_INFO sub-object immediately after the node or link address sub-object related to itself, the node determines that it's a region edge node. Then, the edge node finds out the server layer selection information from the sub-object:

- Determine the switching layer by the LSP Encoding Type and Switching Type fields;
- Determine the switching granularity of the FA-LSP by the Traffic Parameters field;
- Determine the adaptation function for adapting the client signal into the server layer FA-LSP by the G-PID field.

The edge node MUST then determine the other edge of the region, i.e., the tail end of the FA-LSP, with respect to the subsequence of hops of the ERO. The node that satisfies the following conditions will be treated as the tail end of the FA-LSP:

- There is a SERVER_LAYER_INFO sub-object that immediately follows the node or link address sub-object which is related to that node;
- The LSP Encoding Type, Switching Type, G-PID and the Traffic Parameters fields of this SERVER_LAYER_INFO sub-object is the same as the SERVER_LAYER_INFO sub-object corresponding to the head end;
- The node is the first one that satisfies the two conditions above in the subsequence of hops of the ERO.

If a match of tail end is found, the head end now has the clear server layer information of the FA-LSP and then initiates an RSVP-TE session to create the FA-LSP in the appointed server layer between the head end and the tail end.

4.3. Alternative Encoding Solutions

[Editor's note: the section is still under discussion.]

A first alternative solution is to use the mechanism defined in [LSP-RO], i.e., create an ERO HOP attribute TLV.

The content and procedure are not changed from the previous section.

5. A second alternative solution aims to simplify the SERVER_LAYER_INFO processing by using the SERO mechanisms. This can be a new requirements to the SERO or to the ERO Hop attribute. This alternative is not further described here but mentioned for discussions.

6. Security Considerations

TBD.

7. IANA Considerations

TBD.

8. Acknowledgments

TBD.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC3209] D. Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, December 2001.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] L. Berger, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

- [RFC5212] K. Shiimoto et al, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC5212, July 2008.
- [RFC5339] JL. Le Roux et al, "Evaluation of Existing GMPLS Protocols against Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC5339, September 2008.
- [RFC4206] K. Kompella et al, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC4206, October 2005.
- [RFC6107] K. Shiimoto, A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", RFC6107, February 2011.
- [RFC6001] Dimitri Papadimitriou et al, "Generalized Multi-Protocol Label Switching (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC6001, October, 2010.

9.2. Informative Reference

- [RFC4974] D. Papadimitriou and A. Farrel, "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC4974, August 2007.
- [RFC5623] E. Oki et al, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, September 2009.
- [RFC4606] E. Mannie, D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, August 2006.
- [OTN-ctrl] Fatai Zhang et al, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for the evolving G.709 Optical Transport Networks Control", draft-ietf-ccamp-gmpls-signaling-g709v3-08.txt, April, 2013.
- [RFC6003] D. Papadimitriou, "Ethernet Traffic Parameters", RFC6003, October, 2010.
- [LSP-RO] Margaria, C., Giovanni, G., et al, "draft-ietf-ccamp-lsp-attribute-ro", draft-ietf-ccamp-lsp-attribute-ro-01.txt, work I progress;

10. Authors' Addresses

Fatai Zhang
Huawei Technologies
F3-1B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972603
Email: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
F3-1B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972645
Email: huawei.danli@huawei.com

Yi Lin
Huawei Technologies Co., Ltd.
F3-1B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972597
Email: yi.lin@huawei.com

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045 Spain

Phone: +34 913374013
Email: ogondio@tid.es

Cyril Margaria
Coriant GmbH
St Martin Strasse 76
Munich, 81541
Germany

Phone: +49 89 5159 16934
Email: cyril.margaria@coriant.com

Intellectual Property

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions.

For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of

the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Notice

Copyright (c) 2013 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.

Network Working Group
Internet Draft
Category: Informational

Fatai Zhang
Huawei
O. Gonzalez de Dios
Telefonica Investigacion y Desarrollo
A. Farrel
Old Dog Consulting
Xian Zhang
Huawei
D. Ceccarelli
Ericsson
July 09, 2013

Expires: January 08, 2014

Applicability of Generalized Multiprotocol Label Switching (GMPLS)
User-Network Interface (UNI)

draft-zhang-ccamp-gmpls-uni-app-04.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 09, 2014.

Abstract

Generalized Multiprotocol Label Switching (GMPLS) defines a set of protocols for the creation of Label Switched Paths (LSPs) in various switching technologies. The GMPLS User-Network Interface (UNI) was developed in RFC4208 in order to be applied to an overlay network architectural model.

This document examines a number of GMPLS UNI application scenarios. It shows how techniques developed after the GMPLS UNI can be applied to automate or enable critical processes for these applications. This document also suggests simple extensions that could be made to existing technologies to further enable the UNI and points out some unresolved issues.

Table of Contents

1. Introduction	3
2. UNI Addressing	5
3. UNI Auto Discovery	6
4. UNI Path Computation.....	7
4.1. UNI Link Selection.....	8
5. Additional Parameters across UNI.....	10
5.1. Constrained Path Computation.....	10
5.2. Collection Requests over UNI.....	11
6. UNI Path Provisioning Models.....	11
6.1. Flat Model	12
6.2. Stitching Model.....	12
6.3. Session Shuffling Model.....	13
6.4. Hierarchal Model.....	13
7. UNI Recovery	14
7.1. End-to-end Recovery.....	15
7.1.1. Serial Provisioning of Working and Protection Paths	15
7.1.2. Concurrent Computation of Working and Protection Path	16
7.2. Segment Recovery.....	17
8. UNI Call	17
8.1. Exchange of UNI Link Information.....	18
8.2. Control of Call Route.....	18
9. UNI Multicast	19
9.1. UNI Multicast Connection Model	19
9.2. UNI Multicast Connection Provisioning	21
10. Security Considerations.....	22
11. IANA Considerations.....	22
12. Acknowledgments	22
13. References	23
13.1. Normative References.....	23
13.2. Informative References.....	25
14. Contributors' Address.....	26
15. Authors' Addresses	27

1. Introduction

Generalized Multiprotocol Label Switching (GMPLS) [RFC3945] defines a set of protocols, including Open Shortest Path First - Traffic Engineering (OSPF-TE) [RFC4203] and Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) [RFC3473], which can be used to create Label Switched Paths (LSPs) in a number of deployment scenarios with various transport technologies.

The User-Network Interface (UNI) reference point is defined in the Automatically Switched Optical Network (ASON) [G.8080]. According to [G.8080], the UNI may be implemented as a peering between a client-side entity (UNI-C) and a network-side entity (UNI-N). End-to-end connectivity between UNI-C nodes is achieved across the core network by three components: a UNI request from source UNI-C to source UNI-N; a core network connection from source UNI-N to destination UNI-N; and a UNI request from destination UNI-N to destination UNI-C.

The GMPLS overlay model, as per [RFC4208], can be applied at the UNI, as shown in Figure 1.




Figure 1 - Applying GMPLS overlay model at UNI

In Figure 1, assume that there is an end-to-end UNI connection passing through EN1-CN1-CN2-CN3-EN3. For convenience, some terms used in this document are defined below:

- "source EN" refers to the edge-node which initiates the connection (i.e., EN1);
- "destination EN" refers to the edge-node where the connection is terminated (i.e., EN3);
- "ingress CN" refers to the core-node to which the source EN is attached (i.e., CN1);
- "egress CN" refers to the core-node to which the destination EN is attached (i.e., CN3).

[RFC4208] provides mechanisms for UNI signaling, which are compatible with GMPLS RSVP-TE signaling ([RFC3471] and [RFC3473]). A single end-to-end RSVP session between source EN and destination EN is used for the user connection, just as it would be for connection creation between two core nodes. However, when considering the isolation of topology information between the core network and the overlay network, additional processing of the RSVP-TE Explicit Route Object (ERO) and Record Route Object (RRO) is required. For example, the ingress CN should verify the ERO it receives against its topology database and may enhance it with additional path information before forwarding the PATH message. And the ingress/egress CN may edit or remove the RRO in order to hide the path segment used inside the core network from the EN.

The GMPLS UNI can be used in many application scenarios. For example, in a multi-layer network [RFC6001] the interface between client layer node and server layer node can be seen as a UNI. Or, when deploying VPN services such as Layer One Virtual Private Networks (L1VPNs) [RFC4847], [RFC5253], users can connect to a service provider network via a UNI.

This document examines a number of current and future GMPLS application scenarios. It shows how techniques developed after the GMPLS UNI can be used to automate or enable critical aspects of these application scenarios. It points out some potential technology extensions that could improve UNI operation, and highlights some unresolved issues.

2. UNI Addressing

In [RFC4208], the GMPLS overlay model is applied at the UNI reference point, and it is required that the edge-node and its attached core-node of the overlay network share the same address space that is used by GMPLS to signal between the edge-nodes across the core network. Under this condition, the user connection can be created using a single end-to-end RSVP session, which is consistent with the RSVP model. Therefore, RSVP-TE defined in [RFC3473] can be used for support GMPLS UNI without any extensions.

However, in some deployments of the GMPLS UNI, it is not practical for the EN and its attached CN to share the same address space. This can arise if the core and overlay networks were designed and deployed separately or belong to different carriers. For example, the core network may use IPv6 addresses, while the overlay network uses IPv4 addresses. Or, since the core network is a closed system, the assignment of the IP addresses of the CNs may be independent of other IP addresses outside the core network. This implies that the nodes in the core network may use addresses which could collide with the edge nodes in the overlay network.

[RFC4208] does not state how to ensure that an edge-node and its attached core-node share the same address space. This document analyses the addressing deployment scenarios as follows:

1. Overlay network and core network share a common addressing policy. This might be quite feasible in a multi-layer network operated by a single carrier.

In this scenario, end-to-end UNI connectivity may use a single RSVP session, and the core routing information (assuming it is shared and not stripped for confidentiality reasons) will be meaningful to the ENs. Note, however, that the overlay model examined by this document assumes that there is some separation between the overlay and core networks, and this might mean that the overlay network is not able to see the topology or routing information of the core network even when they share a common address space.

2. ENs have visibility into the core network, but overlay and core networks have different address spaces. This is the more common model envisaged by [RFC4208] and for basic mode L1VPN deployments [RFC5251]. The previous scenario can be seen to be a special case of this scenario where the two address spaces are complementary. In this deployment the ENs each have two addresses: one in the overlay network and one in the core network. The source EN is

aware of the addresses for itself, the ingress CN, the egress CN, and the destination EN in the address space of the core network. It may also have full visibility into the core network, but this is not a requirement.

In this scenario, the ENs are responsible for performing address mapping between the overlay network's addresses for the ENs, and the core network's addresses for the same nodes and/or its TE links. A typical deployment may assign addresses in the core network address space for the EN and/or its TE links at the EN side, so that EN can use these addresses to communicate with the core network for UNI connection provisioning.

In this deployment, a single end-to-end RSVP-TE session can still be utilized from the source EN to the destination EN using addressing and naming from the core network's address space.

3. ENs do not have any knowledge of the core address space, or do not support the address space the core network uses (e.g., ENs do not support IPv6 that is used by the core network). ENs will have no visibility into the core network.

In this scenario, the ingress CN is responsible for mapping addresses to the core address space and filling in any additional routing information. A typical deployment is to assign addresses in the overlay address space for the ingress CN and/or its TE links at the CN side, so that the EN can use overlay addresses to reach the ingress CN and to identify the destination EN.

In this deployment the end-to-end connectivity must be created either using "session stitching" (see Section 6.2) or "session shuffling" (see Section 6.3).

3. UNI Auto Discovery

When the end-to-end connection is set up across the core network, it must be targeted at the destination CN so that it can be extended to the destination EN. This means that either the source EN must know the identity of the destination CN to which the destination EN is attached, or the source CN must know this information. This requires some form of "discovery" (possibly including configuration), and depending on the addressing scheme in use (see Section 2), address mapping needs to be performed by the source EN or the source CN.

The discovery problem may be exacerbated when a variety of services are requested since the source EN will need to know the capabilities and available resources on the link between the destination CN and the destination EN. It could discover this by attempting to set up a connection and by drawing conclusions from connection setup failures, but this is not efficient. Furthermore, in the case of a dual-homed destination EN (such as EN2 in Figure 1), a choice of destination CN must be made, and that choice may be influenced by the capabilities and available resources on the CN-EN links leading to the destination EN.

If the UNI is applied in an L1VPN scenario, two mechanisms for auto discovery have been defined. Auto discovery of UNI using OSPFv2 is provided in [RFC5252] using an L1VPN LSA to advertise the L1VPN information via the L1VPN info TLV and the TE information of the CE-PE link (in the language of UNI, it's the EN-CN link) via the TE link TLV. Auto discovery of UNI using BGP is provided in [RFC5195] by having each edge CN advertise to other edge CN the following information, at a minimum: its own IP address and the list of <private address, provider address> tuples local to that PE. Once that information is received, the remote PEs will identify the list of VPN members they have in common with the advertising PE, and use the information carried within the discovery mechanism to perform address resolution during the signaling phase of Layer-1 VPN connections.

4. UNI Path Computation

End-to-end UNI path computation includes three parts: the selection of the source UNI link, the path computation inside the core network and the selection of the destination UNI link.

The selection of UNI links may not be necessary in all scenarios. One example is in the case of single-homing with only one UNI link between EN and CN. Another example is manual selection of the UNI link when the service is requested (i.e., as a function of the service request such as the port mapping used in a L1VPN). In such cases, the CN to which the source EN is attached, or the path Computation Element (PCE) ([RFC4655]) which is responsible for the core network, can perform the path computation across the core network when the UNI signaling request is sent from the source EN to the source CN.

4.1. UNI Link Selection

This document is specific to the overlay architectural model, and that means that the source EN does not have the topology and TE information of the core network. Therefore, in the case of multi-homing (i.e., the source EN is connected to more than one CN), the source EN does not have enough information to make a correct choice among all the UNI links between itself and the core network for an optimal end-to-end connection.

In this case, a PCE whose computation domain covers both the core network and the ENs attached to it can be used. Note that the GMPLS UNI predates PCE and hence a PCE was not available in early GMPLS UNI deployments. A PCE that has the topology and TE information of the core network can use the UNI discovery mechanism described in Section 3 to learn the EN-CN relationship and the TE information of the UNI links, and therefore has the ability to select the optimal UNI link for the connection.

Figure 2 shows the procedures for UNI path computation using a single PCE with visibility into the core network and information about all of the CN-EN links. When the UNI path computation request is received, the PCE can help the source EN to compute the end-to-end route of the UNI connection based on routing information it has access to, so that the source EN can create the UNI connection using the optimal UNI link. As shown in Figure 2, the following steps are carried out:

Step 1: EN1 requests a path from EN1 to EN2 by sending a PCReq message to the PCE;

Step 2: The PCE computes a path based on its view of the core network and knowledge of all the EN-CN links. In this case, it returns the path EN1-CN4-CN5-CN6-EN2 to the EN1 node;

Step 3: EN1 starts the signaling process to set up the LSP by using a standard RSVP signaling process, using the path information as computed.

If confidentiality of the topology within the core network needs to be preserved, the Path Key Subobject (PKS) can be used for either approach outlined here (see [RFC5520] and [RFC5553]). In the PCRep message returned to EN1, the Confidential Path Segment (CPS) (i.e., CN4-CN5-CN6) is encoded as a PKS by the PCE. Therefore, EN1 only learns the selected UNI link from the PCE. When CN4 receives the UNI signaling message from EN1 carrying the PKS, CN4 asks the PCE to decode the PKS and then continues to signal the LSP.




Figure 2 - Procedure using a PCE for UNI path computation

Note that the case described in this section, the PCE needs to be visible to the ENs, and there also needs to be a control channel between the PCE and the ENs for the exchange of PCE Protocol (PCEP) messages. An alternative implementation could be that a PCE is located inside each CN to which the source EN is attached, so that the source EN can use the UNI control channel to send and receive the PCEP messages.

The node requesting for a LSP, crossing UNI, may not be an EN node, as depicted in Figure 3. The procedure described above still applies.

In this case, if an explicit route is desired there is an additional requirement that the PCE needs to have visibility into the overlay networks. Otherwise, the PCE can only provide the route between two EN nodes as illustrated in Figure 3.

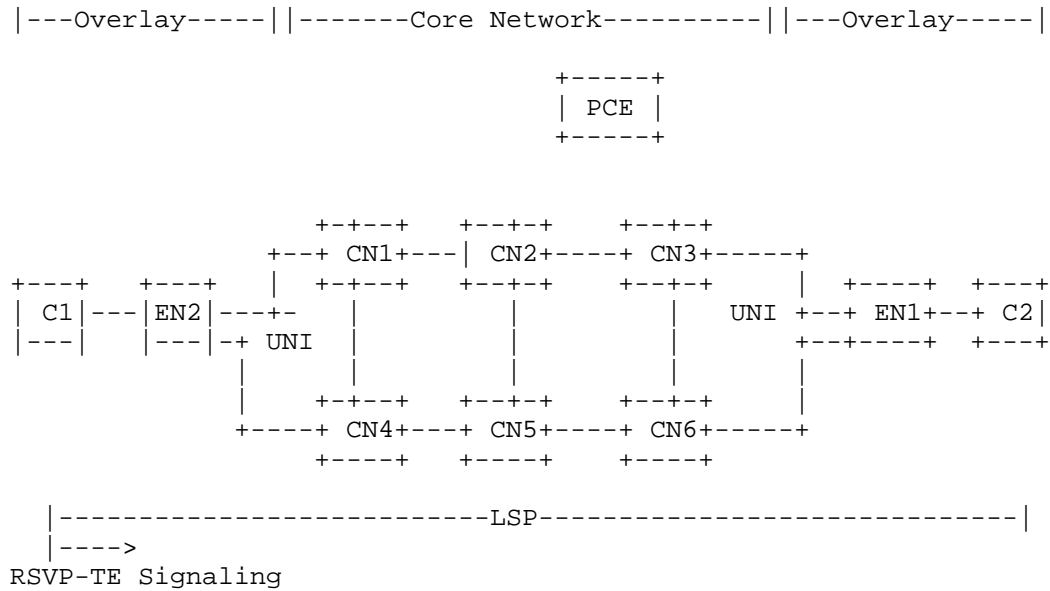


Figure 3 - Procedure of non-EN Node in the Overlay Path Computation

5. Additional Parameters across UNI

With new extensions currently proposed for RSVP-TE protocol, new parameters/functions can also be applicable to UNI.

5.1. Constrained Path Computation

Constraints that can be applied to the path computation in the core network are:

+ Diversity: it is possible to indicate the resources must or should be avoided during the path computation by means of the Exclude Router Object (XRO) [RFC4874], the Explicit Exclusion Route Subobject (EXRS) [RFC4874] and the LSP subobject [LSP-DIV]. Such resources can consist of:

```
-IPv4 prefix, IPv6 prefix, Unnumbered Interface ID, AS
Number and SRLG [RFC4874]
```

-IPv4 P2P subobject and IPv6 P2P subobject [LSP-DIV]

+ Latency, Latency Variation and Cost: max delay/delay variation and cost allowed by the server layer LSP [UNI PLUS]

The overlay Edge Node can include into the RSVP-TE Path message an arbitrary number of path computation constraints for the provisioning of the LSP in the server domain. For example, in Figure 2, EN1 can request a path with a constraint: max latency should be 200ms.

If the path computation in the core network is able to provide an LSP meeting the requirements (at least those requirements which must be met) such LSP is established and a RESV message is returned to the Edge node; otherwise an error message (PathErr) is returned.

Use cases described in Section 7 can be viewed as a special use case of diversity.

5.2. Collection Requests over UNI

In addition to the path request with path computation constraints, the overlay nodes can also request for the collection in the core network of the effective values of the parameters indicated as path computation constraints. The collection of such parameters is indicated via dedicated flags in the LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES in Path Message. Flags defined are:

- Cost collection flag [TE-REC]
- Latency collection flag [TE-REC]
- Latency Variation collection flag [TE-REC]
- SRLG collection flag [SRLG-FA]

In the scenario depicted in Figure 2 a request with constraints on max latency might be issued together with the request of collecting e.g. the effective SRLGs of the provided path, in order to set up a SRLG-disjoint recovery path, as explained in Section 7. Collected parameters are returned to the overlay edge node via the Record Route Object (RRO) in the RESV message.

6. UNI Path Provisioning Models

The basic GMPLS UNI application is to provide end-to-end connections between edge-nodes through a core network via the overlay model. This section briefly describes four ways in which the end-to-end LSP can be created and operated across the core network.




Figure 5 - The Stitching Model

This model allows the core network a degree of independence so that the S-LSP can be set up and modified without the knowledge of the overlay network. Remember that stitching is a data plane function, so that the EN-CN LSP segments are cross-connected to the S-LSP at the edge CNs. This means that, just as in Section 6.1, the overlay and core networks must have the same switching capabilities. However, the control plane for the stitching model operates just as the hierarchical model described in Section 6.4, so the S-LSP appears as a single hop in the overlay network.

6.3. Session Shuffling Model

The session shuffling approach ([RFC5251]) is a modification of the flat model described in Section 6.1. In this approach a single end-to-end session is established, but as the signaling messages pass through the ingress and egress CNs, address mapping is performed on all addresses carried by the messages to replace the addresses with values from the correct address space. The ERO and RRO are stripped from the messages as previously discussed, so there is no need for the CNs to examine those objects to map addresses. However, all other addresses must be mapped including the important session identifiers (the source and destination addresses). Viewed from the outside (perhaps through an NMS) this gives the impression of session stitching because the session has different identifiers as it crosses the core network. An NMS might, therefore, present the shuffling model as the stitching model, or it might operate the same address shuffling/mapping as is used by CNs.

6.4. Hierarchical Model

If the ENs and CNs have the same switching capability, a tunnel between the ingress and egress core-nodes can be provisioned to carry the end-to-end connection. The tunnel may have a larger capacity than

the end-to-end UNI connection, depending on the policies configured at the ingress CN of the core network. The end-to-end connection can be nested into a tunnel, which forms the LSP hierarchy [RFC4206] as shown in Figure 6. If the tunnel has a larger capacity, other LSPs can also be nested within the same tunnel.

Alternatively, if the ENs and CNs have different switching capabilities the LSP hierarchical model can also be used exactly as described in [RFC4206].

In the hierarchal model, the end-to-end connection can be divided into three hops: one for each UNI link and one hop across the core network. The core network tunnel can be pre-provisioned via network planning, or triggered by the UNI signalling. For the latter case, [RFC5212], [RFC6001] and other multi-layer network related specifications can be used to create the hierarchical LSP.




Figure 6 - The Hierarchal Model

7. UNI Recovery

One of the significant uses of GMPLS is to provide recovery mechanisms for connections. Recovery and protection mechanisms are also needed in many UNI scenarios, and the relationship between the overlay and core network provide obvious places at which to operate the recovery techniques.

7.1. End-to-end Recovery

In the case of multi-homing, UNI end-to-end recovery is possible. As shown in Figure 7, the working path (W) and the protection path (P) are disjoint from each other not only inside the core network, but also at both the source and destination sides of the UNI. Mechanisms need to be provided to ensure the selection of disjoint working and backup paths as discussed in the following subsections.

It should be noted that end-to-end recovery can be operated even when the ENs are single-homed. However, obviously, in this case there is no protection against the failure of an EN-CN link, or of the edge CN itself.




Figure 7 - UNI End-to-End Recovery

7.1.1. Serial Provisioning of Working and Protection Paths

In serial provisioning, one path is computed before another and the associated LSP may even be set up before the second path is computed. In the case where the working path is computed and created before the protection path, path computation for the protection path needs to select a (maximally) disjoint path given this existing working path.

If the EN is allowed to see details of the core network, the EN can use the RRO to collect the route of the working path. It can then use the Exclude Route Object (XRO) to exclude the working path when signaling the protection path, as described in [RFC4874].

But in most cases, in order to preserve the confidentiality of topology within the core network, the route of the working path as it

traverses the core network will be hidden from the EN. In such cases, the RRO and XRO mechanism cannot be used. Alternative includes:

- Only collect the Shared Risk Group (SRG) information, but not the full path information [SLRG-FA]. This is because the SRG information is normally less confidential than the information of node ID and link ID.

- Another possible solution is encrypted the SRG information and provide it to the EN nodes, so that the EN nodes can using this information to convey the diversity constraint, as the method specified in [UNIExt].

- In an application scenario where a PCE is involved inside the core network, then the Path Key mechanism can be used. The confidential path segment, i.e., the route of the working path as it traverses the core network, is encoded as a PKS by the PCE when computing the working path [RFC5520]. This PKS can be used by the EN when it requests the PCE to compute a protection path, to exclude the nodes and links used by the working path. As previously described, the PKS is also used in signaling [RFC5553] so that the EN can indicate to the CN what path to use across the core network.

In order to specify the diversity requirement, it is required that the PKS should be carried in the XRO in both PCEP message and RSVP-TE signaling.

7.1.2. Concurrent Computation of Working and Protection Path

The working and protection path can be computed at the same time (e.g., by PCE or by one of the CNs to which the source EN is attached).

[PCE-GMPLS] adds support for an end node to request a protected service using the protection types defined in [RFC4872]. Therefore, it's possible that the source EN requests the edge CN or PCE to compute both the working and the protection path at the same time. At this time, the disjunction requirement can be resolved inside the path computation server.

Same as described in the previous section, the path segment traversing the core network can be encoded as a PKS if confidentiality is requested.

7.2. Segment Recovery

The UNI connection may request protection only inside the core network, especially in case of single-homing. A UNI segment protection example is shown in Figure 8. In this case, the core network provides a "recovery domain".




Figure 8 - UNI Segment Recovery

[RFC4873] provides a mechanism for segment recovery, in which the PROTECTION Object is extended to indicate segment recovery, and the Secondary ERO (SERO) is introduced for the explicit control of the protection LSP between the branch node and the merge node.

However, in the overlay model, the mechanisms of segment recovery described in [RFC4873] may not be appropriate. In particular, the source EN might not know the CN to which the destination EN is attached. That means that the source EN knows the branch for the protection segment, but does not know the merge node.

But the model shown in Figure 8 is particularly important because it places the responsibility for service delivery with the edge CNs. This will be a common operational model in overlay networks. Fortunately the stitching model (Section 6.2) and the hierarchical model (Section 6.4) are good at providing the necessary protection within the core network without the ENs having to be aware of the paths in the core network.

8. UNI Call

The Call is a fundamental component of the ASON model [G.8080]. It is used to maintain the association between one or more user

applications and the network, and to control the set-up, release, modification, and maintenance of sets of Connections (LSPs). In simple cases, the Call and Connection can be established at the same time and in a strict one-to-one ratio. In this case, Call signaling requires only minor extensions to connection signaling. However, if Calls are handled separately from Connections, or if more than one Connection can be associated with a single Call, additional Call signaling is required.

The GMPLS Call, defined in [RFC4974], provides a mechanism to negotiate agreement between endpoints possibly in cooperation with the nodes that provide access to the network. Typically the GMPLS Call can be applied in the UNI scenario for access link capability exchange, policy, authorization, security, and so on.

8.1. Exchange of UNI Link Information

It is possible that the TE attributes of the access link (i.e., the UNI link) are not shared across the core network. So the source EN may not have the TE information of the destination access link as well as the capability of the destination EN. For example, in case of TDM network, the Virtual Concatenation (VCAT) and Link Capacity Adjustment Scheme (LCAS) capability of the destination EN may not be known.

In this case, the source EN can raise a Call carrying the LINK_CAPABILITY object to have a capability exchange with the destination EN, as described in [RFC4974].

8.2. Control of Call Route

When applying the Call, it's possible that there are multiple core network domains between the source EN (Call initiator) and the destination EN (Call terminator), or there is more than one Call manager in the core network (e.g., in the multi-homing scenario where the CNS to which the ENs are attached act as the Call managers).

In the both cases, when establishing the Call, there may be multiple alternative routes for the Call message to reach the destination EN. One can simply use the hop-by-hop manner (i.e., each Call manager determines the next Call manager to which the Call message will be sent by itself) to control the path of the Call.

However, in the practical deployment of UNI Call, commercial and policy motivations normally play an important role in selecting the Call route, especially in the multi-domain scenario. In this case, the hop-by-hop manner is not practical because the route of the Call

needs to be pre-determined in consideration of commercial and policy factors before establishing the Call.

Therefore, it is desirable to allow full control of the Call by the source EN. That is, the source EN can identify the full Call route and signal it explicitly, so that the Call message can be forwarded along the desired route. Moreover, the management plane needs to be able to identify the Call route explicitly as an instruction to the source EN.

9. UNI Multicast

Data plane multicasting is supported in existing Traffic-Engineering networks. GMPLS provides extensions to RSVP-TE to support provisioning of point-to-multipoint (P2MP) TE LSPs via the control plane, as described in [RFC4461] and [RFC4875].

In the scenarios where P2MP is supported using the overlay architectural model, it is a requirement to transport signals from one source EN to multiple destination ENs. One could create a point-to-point (P2P) connection between the source EN and each destination EN, but it will likely be a waste of bandwidth resource both of the UNI link and in the core network.

Therefore, there are some scenarios required to support point-to-multipoint (P2MP) TE LSPs from one source EN to multiple leaf ENs.

9.1. UNI Multicast Connection Model

There are two cases for the UNI multicast. For the first case, only the ingress and egress CNs in the core network support P2MP. The core network has to provide multiple P2P connections between ingress CN and each egress CN for the end-to-end UNI multicast, as shown in Figure 9. This relieves the pressure on the source UNI link, but does not help the over use of the core links such as CN1-CN2.




Figure 10 - All CNs support multicast

For example, in the Ethernet over OTN scenario, if the core network can support ODU0 multicast, then an ODU0 P2MP LSP can be created inside the core network to carry the client Gigabit Ethernet (GE) signal for the ENs.

Note that the branching of the P2MP connection could also happen at the source EN if the EN is multi-homed. In this case, each branch from the source EN uses a separate UNI link connecting the source EN to the core network. For each UNI branch, the connection model inside the core network is the same as described in this section.

9.2. UNI Multicast Connection Provisioning

The four UNI connection provisioning models, as described in Section 5, should also be applied in the UNI multicast scenario.

For the flat model, one end-to-end P2MP session as described in [RFC4875] can be used to create the P2MP LSP from source EN to leaf ENs.

For the stitching model, multiple P2P LSP segments or one P2MP LSP segment between the ingress CN and each egress CNs needs to be created and then stitched to the UNI P2MP LSP. GMPLS UNI signaling should have the capability to convey the multicast information by using stitching model.

For the session shuffling model, one end-to-end P2MP session can be used to create the P2MP LSP, with an address mapping performed at both ingress and egress CNs.

For the hierarchical model, multiple P2P LSP tunnels or one P2MP LSP tunnel between the ingress CN and each egress CNs needs be triggered by the UNI signaling for creating the P2MP LSP. GMPLS UNI signaling should have the capability to convey the multicast information by using the hierarchical model.

10. Security Considerations

[RFC5920] provides an overview of security vulnerabilities and protection mechanisms for the GMPLS control plane, which is applicable to this document.

The details of the specific security measures of the overlay network architectural model are provided in [RFC4208], which permits the core network to filter out specific RSVP objects to hide its topology from the EN.

Furthermore, if PCE is used, the security issues described in [RFC4655] should also be considered.

Additionally, when the PKS mechanism is applied, the security issues can be dealt with using [RFC5520] and [RFC5553].

11. IANA Considerations

This informational document does not make any requests for IANA action.

12. Acknowledgments

The authors would like to thank Zafar Ali for his comments.

13. References

13.1. Normative References

- [RFC3209] D. Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, December 2001.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] L. Berger, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4203] Kompella, K., and Rekhter, Y., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [RFC4206] K. Kompella et al, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC4206, October 2005.
- [RFC4208] G. Swallow et al, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC4208, October 2005.
- [RFC4655] A. Farrel et al, "A Path Computation Element (PCE)-Based Architecture", RFC4655, August 2006.
- [RFC4847] T. Takeda, Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC4847, April 2007.
- [RFC4872] J.P. Lang et al, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC4872, May 2007.
- [RFC4873] L. Berger et al, "GMPLS Segment Recovery", RFC4873, May 2007.

- [RFC4874] CY. Lee et al, "Exclude Routes - Extension to Resource Reservation Protocol-Traffic Engineering (RSVP-TE)", RFC4874, April 2007.
- [RFC4875] R. Aggarwal et al, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC4875, May 2007.
- [RFC4974] D. Papadimitriou and A. Farrel, Ed., "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC4974, August 2007.
- [RFC5150] A. Ayyangar et al, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC5150, February 2008.
- [RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, June 2008.
- [RFC5251] D. Fedyk and Y. Rekhter, Ed., "Layer 1 VPN Basic Mode", RFC5251, July 2008.
- [RFC5252] I. Bryskin and L. Berger Ed., "OSPF-Based Layer 1 VPN Auto-Discovery", RFC5252, July 2008.
- [RFC5520] R. Bradford, Ed., "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC5520, April 2009.
- [RFC5553] A. Farrel, Ed., "Resource Reservation Protocol (RSVP) Extensions for Path Key Support", RFC5553, May 2009.
- [RFC6001] Dimitri Papadimitriou et al, "Generalized Multi-Protocol Label Switching (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC6001, October, 2010.
- [RFC6107] K. Shiimoto, A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", RFC6107, February 2011.
- [G.8080] ITU-T Rec. G.8080/Y.1304, "Architecture for the Automatically Switched Optical Network (ASON)," June 2006 (and Amend.2, September 2010).

13.2. Informative References

- [RFC4461] S. Yasukawa, Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC4461, April 2006.
- [RFC5212] K. Shiomoto et al, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC5212, July 2008.
- [RFC5253] T. Takeda, Ed., "Applicability Statement for Layer 1 Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July 2008.
- [RFC5339] J.L. Le Roux et al, "Evaluation of Existing GMPLS Protocols against Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC5339, September 2008.
- [RFC5441] J.P. Vasseur et al, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC5441, April 2009.
- [RFC5623] Oki, E., Takeda, T., Le Roux, J.L., and Farrel, A., "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", RFC 5623, September 2009.
- [RFC5920] L. Fang, Ed., "Security Framework for MPLS and GMPLS Networks", RFC5920, July 2010.
- [Call-ext] Fatai Zhang et al, "RSVP-TE extensions to GMPLS Calls", draft-zhang-ccamp-gmpls-call-extensions-01.txt, July 08, 2009.
- [PCE-GMPLS] C. Margaria et al, "PCEP extensions for GMPLS", draft-ietf-pce-gmpls-pcep-extensions-07.txt, October 21, 2012.
- [SRLG-FA] Fatai Zhang et al, "RSVP-TE Extensions for Configuration SRLG of an FA", draft-ietf-ccamp-rsvp-te-srlg-collect-02.txt, work in progress.
- [RFC6344] G. Bernstein et al, "Operating Virtual Concatenation (VCAT) and the Link Capacity Adjustment Scheme (LCAS) with Generalized Multi-Protocol Label Switching (GMPLS)", RFC6344, August 2011.

- [UNIEExt] D. Fedyk, D. Beller, Lieven Levrau, D. Ceccarelli, F. Zhang, et al, "UNI Extensions for Diversity and Latency Support", draft-fedyk-ccamp-uni-extensions-00.txt, Feb. 2013;
- [LSP-DIV] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Path Diversity using Exclude Routes", draft-ietf-ccamp-lsp-diversity-01.txt, work in progress;
- [UNI-PLUS] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) extension for signaling Objective Function and Metric Bound", draft-ali-ccamp-rc-objective-function-metric-bound-02.txt, work in progress;
- [TE-REC] A., Zafar, G., Swallow et al, "Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)extension for recording TE Metric of a Label Switched Path", draft-ietf-ccamp-te-metric-recording-01.txt, work in progress;

14. Contributors' Address

Yi Lin
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972914
Email: yi.lin@huawei.com

Young Lee
Huawei Technologies
1700 Alma Drive, Suite 100
Plano, TX 75075
USA

Phone: (972) 509-5599 (x2240)
Email: leeyoung@huawei.com

Dan Li
Huawei Technologies

F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28973237
Email: huawei.danli@huawei.com

15. Authors' Addresses

Fatai Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972912
Email: zhangfatai@huawei.com

Oscar Gonzalez de Dios
Telefonica Investigacion y Desarrollo
Emilio Vargas 6
Madrid, 28045
Spain

Phone: +34 913374013
Email: ogondio@tid.es

Adrian Farrel
Old Dog Consulting

EMail: adrian@olddog.co.uk

Xian Zhang
Huawei Technologies
F3-5-B R&D Center, Huawei Base
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972913
Email: zhang.xian@huawei.com

Daniele Ceccarelli

Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: daniele.ceccarelli@ericsson.com

Intellectual Property

The IETF Trust 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 any IETF 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.

Copies of Intellectual Property 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 any standard or specification contained in an IETF Document. Please address the information to the IETF at ietf-ipr@ietf.org.

The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions.

For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms,

conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.

Disclaimer of Validity

All IETF Documents and the information contained therein 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 THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Notice

Copyright (c) 2013 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.

RSVP-TE Extensions for Bit Error Rate (BER) Measurement
draft-zhang-ccamp-rsvpte-ber-measure-00

Abstract

In the mobile backhaul network, the mobile service is sensitive to Bit Error Rate (BER). When the BER value of the service exceeds the threshold, the Base Station will stop working and the User Equipments (UEs) cannot obtain voice and data services anymore. Now the mobile backhaul tends to be IP/MPLS network and MPLS TE LSP is used to bear the mobile service which may be encapsulated in PW or L3VPN end to end. Then the ingress Label Switched Router (LSR) of the MPLS TE LSP needs to get information on BER along the path of the LSP. This document proposes new extensions of RSVP-TE to advertise the BER measurement requirement of the specific LSP to all of the transit LSRs and the egress LSR, and to report the BER measurement result from any transit or egress LSR towards the ingress LSR.

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 [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 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	2
2. Terminology	3
3. BER_REQUEST TLV	3
3.1. Format of BER_REQUEST	3
3.2. Operations for BER_REQUEST TLV	5
4. Bit Error Indication Report	5
4.1. Error Code for BER measurement report	6
4.2. Operations for BER Error Code	6
5. IANA Considerations	6
6. Security Considerations	6
7. Normative References	6
Authors' Addresses	7

1. Introduction

Bit Error Rate (BER) is a significant parameter for the mobile service, which can cause the Base Station to stop working when its value exceeds the threshold of the service. In IP/MPLS based mobile backhaul network, PW and L3VPN are adopted to bear the mobile service end-to-end, and MPLS TE LSP is adopted as the transport tunnel for which Hot-standby (MPLS TE HSB) or fast reroute (MPLS TE FRR) technologies is used to meet the SLA(Service Level Agreement). There are different kinds of failure detection methods, such as BFD or MPLS OAM, to trigger MPLS TE HSB or FRR to switch traffic fast when failure happens. But as to BER, even if the BER value exceeds the threshold, the detection mechanisms cannot detect the failure to trigger traffic switch to the backup path. In this document, we propose new extensions of RSVP-TE to advertise the BER measurement requirement of the LSP to its transit LSRs and the egress LSR, and to report the BER measurement result from any transit LSR or the egress LSR towards the ingress LSR.

There are two types of BER measurement requirements: one is the single-point BER measurement and the other is the multi-point BER measurement. The first one is to measure if the BER value of one point of the LSP path has exceeded the threshold of the service. The second one is to measure if the sum of BER value of multiple points of the LSP exceeds the threshold of the service. In this document, we just focus on the single-point BER measurement. The multi-point BER measurement will be described in the future version.

For the single-point BER measurement, there are two new extensions of RSVP-TE protocol. One extension is to advertise the BER measurement requirement to all of the transit LSRs and the egress LSR, then these LSRs along the path will start BER measurement for the LSP. The other extension is to report the BER measurement result from any transit LSR or the egress LSR towards the ingress LSR.

2. Terminology

BER: Bit Error Rate

RAN: Radio Access Network

LSR: Label Switch Router

LSP: Label Switch Path

3. BER_REQUEST TLV

3.1. Format of BER_REQUEST

Path Message of RSVP-TE is used to signal the BER measurement requirement of the LSP, and the LSP_ATTRIBUTES object will be included in the Path Message. The LSP_ATTRIBUTES object which is defined in [RFC5420] is used to signal attributes required in support of an LSP, or to indicate the nature or use of an LSP. The LSP_ATTRIBUTES object format is as below (refer to[RFC5420]):

LSP_ATTRIBUTES class = 197, C-Type = 1




Figure 1: LSP_ATTRIBUTES object

The LSP_ATTRIBUTES object class is 197 of the form 11bbbbbb. This C-Num value (see [RFC2205], Section 3.10) ensures that LSRs that do not recognize the object pass it on transparently. One C-Type is defined, C-Type = 1 for LSP Attributes.

The Attributes TLVs are encoded as below:




Figure 2: Attributes TLVs format

Here, we define the BER_REQUEST TLV, which is a new type of Attribute TLV, to indicate the BER measurement requirement of the LSP. The format of BER_REQUEST TLV is as below:

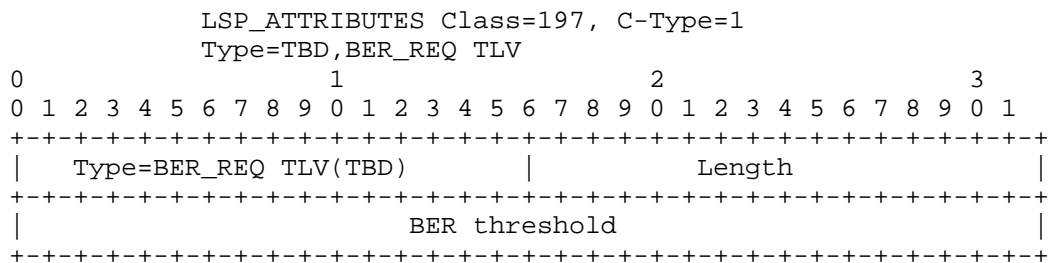


Figure 3: BER_REQ TLV

Type

The identifier of the BER_REQUEST TLV which should be allocated by IANA.

Length

Indicates the total length of the TLV in two octets.

Value

The BER threshold of the service, which is a 32-bit IEEE floating point number. Positive infinity is represented as an IEEE single-precision floating-point number with an exponent of all ones (255) and a sign and mantissa of all zeros. The format of IEEE floating-point numbers is further summarized in [RFC1832]

3.2. Operations for BER_REQUEST TLV

BER_REQUEST TLV is one type of attribute TLVs of the LSP_ATTRIBUTE object, which is optional and may be placed in Path messages to advertise the BER measurement requirement of the LSP. The process of the LSP_ATTRIBUTE object can refer to section 4.2 in [RFC5420].

When a RSVP-TE LSP requires the BER measurement of the path, the ingress LSR MUST send a Path Message with BER_REQUEST TLV in which the BER threshold value is set.

When a LSR receives a Path Message with the BER_REQUEST TLV, the LSR SHOULD start the BER measurement for the LSP. The LSR MUST pass the Path Message with BER_REQUEST TLV unchanged to the next LSR. If the measured BER value exceeds the BER threshold value set in the BER_REQUEST TLV, the LSR MUST report the bit error result towards the ingress LSR of the LSP.

If a LSR cannot support the BER_REQUEST TLV, the LSR SHOULD ignore this TLV and pass the Path Message with BER_REQUEST TLV unchanged to the next LSR.

For a LSR which has started the BER measurement on receiving the Path Message with the BER_REQUEST TLV, if the LSR receives the updated Path Message without BER_REQUEST TLV, it MUST stop the BER measurement for this LSP, and pass the Path Message to next LSRs unchanged.

4. Bit Error Indication Report

For the single-point BER measurement, the LSR should report the BER measurement result towards the ingress LSR of the LSP when the BER measurement value exceeds the threshold of the service. This document we propose a new type of Error Code and its Error Value of the ERROR_SPEC object, which is defined in [RFC2205] and [RFC3209], to report the BER measurement result within PathErr Message.

4.1. Error Code for BER measurement report

The ERROR_SPEC object is defined in [RFC2205] and [RFC3209]. Here we define a new BER Error Code as below (The Error Code for "BER measurement report" is to be defined) :

Error code	Error value	Description
TBD	0	Bit Error Elimination
	1	Bit Error Indication

4.2. Operations for BER Error Code

The BER measurement result is reported through a new Error Code and corresponding Error Value of the ERROR_SPEC object, which is placed in PathErr Message.

For a LSR which has started the BER measurement for the specific LSP, if the BER measurement value exceeds the threshold of the service, a PathErr Message MUST be triggered to send towards the ingress LSR of this LSP. The PathErr Message MUST include BER Error code with Error Value 1 for Bit Error Indication. When the BER measurement value becomes less than the BER threshold value, the LSR MUST send a PathErr Message with a value of 0 for Bit Error Elimination.

5. IANA Considerations

IANA should allocate the type value of the BER_REQUEST TLV and the BER Error Code, which are defined in this document.

6. Security Considerations

The extensions of RSVP TE for BER in this document do not introduce any new security issues, and the reader is referred to the security considerations expressed in [RFC2205], [RFC3209], and [RFC5420].

7. Normative References

[RFC1832] Srinivasan, R., "XDR: External Data Representation Standard", RFC 1832, August 1995.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [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.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Li Zhang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: monica.zhangli@huawei.com