Network Working Group
Internet Draft
Category: Standards Track

Fatai Zhang
Dan Li
Huawei
D. Ceccarelli
D. Caviglia
Ericsson
Guoying Zhang
CATR
P.Grandi
S.Belotti
Alcatel-Lucent

Expires: April 2012

October 13, 2011

Link Management Protocol (LMP) extensions for G.709 Optical Transport Networks

draft-zhang-ccamp-gmpls-g709-lmp-discovery-05.txt

Status of this Memo

Abstract

Recent progress of the Optical Transport Network (OTN) has introduced new signal types (i.e., ODU0, ODU4, ODU2e and ODUflex) and new Tributary Slot granularity (1.25Gbps).

Since equipments deployed prior to recently defined ITU-T recommendations only support 2.5 Gbps Tributary Slot granularity and ODU1, ODU2 and ODU3 containers, the compatibility problem should be considered. In addition, a Higher Order ODU (HO ODU) link may not support all the types of Lower Order ODU (LO ODU) signals defined by the new OTN standard because of the limitation of the devices at the two ends of a link. In these cases, the control plane is required to run the capability discovering functions for the evolutive OTN.

This document describes the extensions to the Link Management Protocol (LMP) needed to discover the capability of HO ODU link, including the granularity of Tributary Slot to be used and the LO ODU signal types that the link can support.

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].

Table of Contents

1. Introduction

The Link Management Protocol (LMP) defined in [RFC4204] is being developed as part of the Generalized MPLS (GMPLS) protocol suite to manage Traffic Engineering (TE) links.

Recently, great progress has been made for the Optical Transport Networking (OTN) technologies in ITU-T. New ODU containers (i.e., ODU0, ODU4, ODU2e and ODUflex) and a new Tributary Slot (TS) granularity (1.25Gbps) have been introduced by the [G709-V3], enhancing the flexibility of OTNs.

With the evolution and deployment of G.709 technology, the backward compatibility problem requires to be considered. In data plane, the equipment supporting 1.25Gbps TS can combine the specific Tributary Slots together (e.g., combination of TS#i and TS#i+4 on a HO ODU2 link) so that it can interwork with other equipments which support 2.5Gbps TS. From the control plane point of view, it is necessary to discover which type of TS is supported at both ends of a link, so that it can choose and reserve the TS resources correctly in this link for the connection.

Additionally, the requirement of discovering the signal types of Lower Order ODU (LO ODU) that can be supported by a Higher Order ODU (HO ODU) should be taken into account. Equipment at one end of a HO ODU link may not support to transport some types of LO ODU signals (e.g., may not support the ODUflex). In this case, this HO ODU link should not be selected for those types of LO ODU connections.

From the perspective of control plane, it is necessary to discover the capability of a HO ODUk or OTUk link including the granularity of TS to be used and the LO ODU signal types that the link can support. Note that this capability information can be, in principle, discovered by routing. Since in certain case, routing is not present (e.g. UNI case) we need to extend link management protocol capabilities to cover this aspect. Obviously, in case of routing presence, the discovering procedure by LMP could also be optional.

This document extends the LMP and describes the solution of discovering HO ODU link capability.

2. Terminology

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 of the Evolutive G.709

The traditional OTN standard [ITUT-G709] describes the optical transport hierarchy (OTH) and introduces three ODU signal types (i.e., ODU1, ODU2 and ODU3). The ODUj can be mapped into one or more Tributary Slots (with a granularity of 2.5Gbps) of OPUk where j<k. The ODUj can also be mapped into OTUj (j=1, 2 or 3) directly.

Recent revisions of ITU-T Recommendation G.709 have introduced new features for the evolutive Optical Transport Networks (OTN). New ODU signals, including ODU0, ODU4, ODU2e and ODUflex, are described in [G709-V3]. This document also defines the new multiplexing hierarchy for the evolutive OTN. In this multiplexing hierarchy, LO ODUj can be mapped into an OTUj, or multiplexed into a HO ODUk (where j<k) by occupying several tributary slots.

In case of LO ODUj mapping into OTUj, the following mappings are defined:

- ODU1 into OTU1 mapping
- ODU2 into OTU2 mapping
- ODU3 into OTU3 mapping
- ODU4 into OTU4 mapping

In case of LO ODUj multiplexing into HO ODUk, a new Tributary Slot granularity (i.e., 1.25Gbps) is introduced in [G709-V3]. For the evolutive OTN, the multiplexing of ODUj (j = 0, 1, 2, 2e, 3, flex) into an ODUk (k > j) signal can be depicted as follows:

- ODU0 into ODU1 multiplexing (with 1,25Gbps TS granularity)
- ODU0, ODU1, ODUflex into ODU2 multiplexing (with 1.25Gbps TS granularity)
- ODU1 into ODU2 multiplexing (with 2.5Gbps TS granularity)
- ODU0, ODU1, ODU2, ODU2e and ODUflex into ODU3 multiplexing (with 1.25Gbps TS granularity)
- ODU1, ODU2 into ODU3 multiplexing (with 2.5Gbps TS granularity)
- ODU0, ODU1, ODU2, ODU2e, ODU3 and ODUflex into ODU4 multiplexing (with 1.25Gbps TS granularity)

3.1. Data Plane Backward Compatibility

Equipment supporting a 1.25Gbps TS structure for OPU2 or OPU3 must be backward compatible with equipment which supports only the 2.5G TS structure. Specific Tributary Slots must be combined together (e.g., combination of TS#i and TS#i+4 on a HO ODU2 link) for the LO ODU at one end of the HO ODU link which supports the 1.25Gbps TS structure, so that the LO ODU can be carried on the HO ODU link correctly.

In the following example, suppose that the two ends of an ODU2 or ODU3 link support different TS structure, where node A supports the 1.25Gbps TS structure, while node B supports the 2.5Gbps TS, as shown in the figure below:

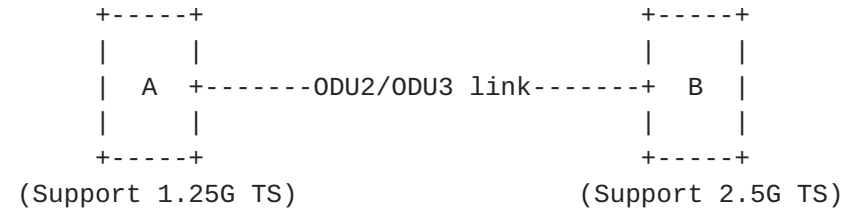

- In case of ODU1 multiplexing into ODU2, node A maps the ODU1 into the TS#i and TS#i+4 (where i<=4) (with the granularity of 1.25Gbps) of OPU2, so that node B can retrieve the ODU1 from the TS#i (with the granularity of 2.5Gbps) of the OPU2, and vice versa.

- In case of ODU1 multiplexing into ODU3, node A maps the ODU1 into the TS#i and TS#i+16 (where i<=16) (with the granularity of 1.25Gbps) of OPU3, so that node B can retrieve the ODU1 from the TS#i (with the granularity of 2.5Gbps) of the OPU3, and vice versa.

- In case of ODU2 multiplexing into ODU3, node A maps the ODU2 into the TS#a/TS#a+16, TS#b/TS#b+16, TS#c/TS#c+16 and TS#d/TS#d+16 (where a<b<c<d<=16) (with the granularity of 1.25Gbps) of OPU3, so that node B can retrieve the ODU2 from the TS#a, TS#b, TS#c and TS#d (with the granularity of 2.5Gbps) of the OPU3, and vice versa.

4. Link Capability Discovery Requirements

4.1. Discovering the Granularity of the TS

As described in section 3.1, if the two ends of a link use different granularities of TS, The LO ODU must be mapped into specific combined Tributary Slots in the end of link with TS of 1.25Gbps.

From the perspective of control plane, when creating a LO ODU connection, the node MUST select and reserve specific TS for the connection if the two ends of a link use different granularities of TS. For example, for an ODU2 link, we suppose that node A only supports the 2.5Gbps TS while node B supports the 1.25Gbps TS. When node B receives a Path message from node A requesting an ODU1 connection, node B MUST reserve the TS#i and TS#i+4 (where i<=4) (with the granularity of 1.25Gbps) and tell node A via the label carried in the Resv message that the TS#i (with the granularity of 2.5Gbps) among the 4 slots has been reserved for the ODU1 connection. Otherwise, the reservation procedure will fail.

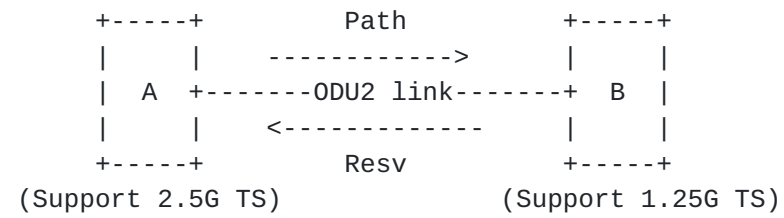

Therefore, for an ODU2 or ODU3 link, in order to reserve TS resources correctly for a LO ODU connection, the control plane of the two ends MUST know which granularity the other end can support before creating the LO ODU connection.

4.2. Discovering the Supported LO ODU Signal Types

Many new ODU signal types are introduced by [G709-V3], such as ODU0, ODU4, ODU2e and ODUflex. It is possible that equipment does not always support all the LO ODU signal types introduced by [G709-V3]. If one end of a HO ODU link can not support a certain LO ODU signal type and there is no HO ODU FA LSP able to support this LO ODU signal, the HO ODU link/FA LSP can not be selected to carry such type of LO ODU connection.

For example, in the following figure, if the interfaces IF1, IF2, IF8, IF7, IF5 and IF6 can support ODUflex signals, while the interfaces IF 3 and IF4 can not support ODUflex signals. In this case, if one ODUflex connection from A to C is requested, and there is no HO ODU FA LSP from node A to C through node B, link #1 and #2 should be excluded, link #3 and link #4 are the candidates (the possible path could be A-D-C through link #3 and link #4).

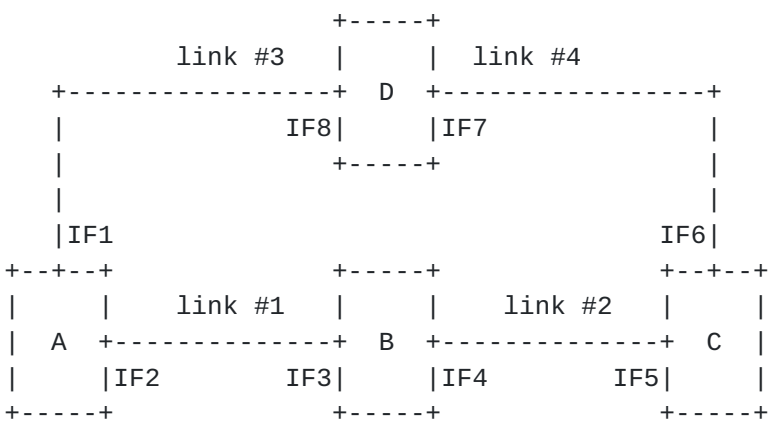

Therefore, it is necessary for the two ends of a HO ODU link to discover which types of LO ODU can be supported by the HO ODU link. After discovering, the capability information can be flooded by IGP, so that the correct path for an ODU connection can be calculated.

5. Extensions: LMP Link Summary Message

[RFC4204] defines the Link Management Protocol (LMP) which consists of four main procedures: control channel management, link property correlation, link connectivity verification, and fault management. As part of LMP, the link property correlation is used to verify the consistency of the TE and data link information on both sides of a link. This document extends the link property correlation procedure to discover the capability of both sides of a HO ODU link.

The designated HO ODU overhead bytes (e.g., the GCC1 and GCC2 overhead bytes) can be used as the control channel to carry the LMP message after the HO ODU link is created. The out-of-band Data Communication Network (DCN) can also be used.

5.1. Message Extension

Three messages are used for link property correlation: LinkSummary, LinkSummaryAck and LinkSummaryNack Message. This document does not change the basic procedure of LMP but just add a new subobject (HO ODU Link Capability Subobject) in the DATA_LINK object to carry the capability of one end of a HO ODU link.

The formats of LinkSummary, LinkSummaryAck and LinkSummaryNack messages are defined in [RFC4204].

5.1.1. LinkSummary Message

The local end of a TE link can send a LinkSummary message to the remote end to start the negotiation about the capability that the TE link can support.

One new Subobject named HO ODU Link Capability Subobject in the DATA_LINK object is introduced by this document. This new subobect is used to tell the remote end of the HO ODU link which TS granularity and which LO ODU signal types that the local end can support. When the DATA_LINK object carries the new HO ODU Link Capability Subobject, the N flag SHOULD be set to 1 which means that the subobject is negotiable.

5.1.2. LinkSummaryAck Message

The LinkSummaryAck message is used to tell the remote end that it has the same capability as the remote end after the LinkSummary message is received by the local end.

5.1.3. LinkSummaryNack Message

The LinkSummaryNack message is used to tell the remote end that it has different capability from the remote end after the LinkSummary message is received by the local end. The LinkSummaryNack message also carries the HO ODU Link Capability Subobject in the DATA_LINK object to tell the remote end the exact capability of the HO ODU link after negotiation, i.e., the granularity of TS and the types of LO ODU that both side of the HO ODU link can support.

5.2. Object Definitions

A new HO ODU Link Capability subobject type is introduced to the DATA LINK object to carry the HO ODU link capability information. The format of the new subobject is defined as follow:

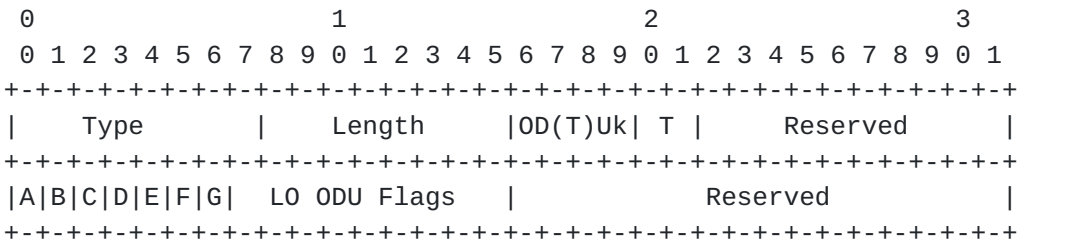

Type (8 bits):

The value of this subobject type is TBD.

Length (8 bits):

The Length field contains the total length of the subobject in bytes, including the Type and Length fields. As for RFC 4204, the Length MUST be at least 4, and MUST be a multiple of 4. Value of this field is 8.

OD(T)Uk (4 bits):

This field is used to indicate the HO ODU link type (in case of LO ODUj multiplexing into HO ODUk, wherein j<k) or the OTU link type (in case of LO ODUk mapping into OTUk).

OD(T)Uk field	Signal type of HO ODUk or OTUk
0	Reserved (for future use)
1	HO ODU1 or OTU1
2	HO ODU2 or OTU2
3	HO ODU3 or OTU3
4	HO ODU4 or OTU4
5-15	Reserved (for future use)

T (2 bits):

The T bits are used to indicate the granularity of the TS of the HO ODU link.

T field	TS type
00	Meaningless
01	1.25Gbps TS granularity
10	2.5Gbps TS granularity
11	Reserved (for future use)

In case that an OTUk link only support ODUj (j=k) into OTUk mapping and does not support any ODUj into ODUk (j<k) multiplexing, then the T field is not meaningful and MUST be filled with 0 and be ignored on receipt.

LO ODU flags (A|B|C|D|E|F|G) (16 bits):

These flags are used to indicate which LO ODU signal types that one end or the both end can support. The flags will be set to 1 if the corresponding LO ODU signal types are supported to be mapped or multiplexed into the OTUk or HO ODUk link.

This rule imposes that:

- At least one flag is set to 1.

- When the ODUj (j=k) flag corresponding to the signal type HO ODUk/OTUk is set to 1, then the signal type OD(T)Uk has to be intended as LO ODUk and direct mapping over OTUk is supported.

  * Furthermore, if only the ODUj(j=k) flag is set to 1, it means that the HO ODUk/OTUk link only supports ODUj(j=k) into OTUk mapping. In other words, the link does not support any ODUj into ODUk (j<k) multiplexing (i.e., payload type != 20/21), but may support carrying various non-ODU client signals listed in Table 15-8 of [G709-V3].

- When an ODUj (j<k) flag not corresponding to the signal type HO ODUk/OTUk is set to 1 then the signal type OD(T)Uk has to be intended as HO ODUk and multiplexing of LO ODUj over HO ODUk is supported.

    Flag A: indicates whether LO ODU0 is supported.

    Flag B: indicates whether LO ODU1 is supported.

    Flag C: indicates whether LO ODU2 is supported.

    Flag D: indicates whether LO ODU3 is supported.

    Flag E: indicates whether LO ODU4 is supported.

    Flag F: indicates whether LO ODU2e is supported.

    Flag G: indicates whether LO ODUflex is supported.

For example, if one end of an OTU2 link supports LO ODU0, LO ODU1, LO ODUflex into HO ODU2 multiplexing and supports LO ODU2 into OTU2 mapping, the flags A, B, C, and G will be set to 1.

As a further example, if one end of an OTU2 link supports only LO ODU2 into OTU2 mapping but no multiplexing, only flag C will be set to 1.

The remaining flags are reserved for future use and MUST be set to 0.

5.3. Procedures

The Link Summary messages used for capability discovery for HO ODUk or OTUk link are sent between adjacent nodes after the HO ODU link is created or driven by some events (e.g., an operator command). The procedure is described below:

- The local end of the HO ODU link sends a LinkSummary message including one or more DATA_LINK objects, each of which contains the Local_Interface_Id, the Remote_Interface_Id, and the HO ODU link capability subobject. This subobject carries the capability that the local end can support, i.e., the granularity of TS and the set of LO ODU signal types that the local end can support. The LinkSummary message is sent to the remote end.

- On receipt of the LinkSummary message, the remote end of the HO ODU link firstly determines whether the local/remote Interface_Id mappings match those that are stored locally as described in [RFC4204], and then obtains the HO ODU link capability subobject and determines the capability of the HO ODU link that both ends can support. The detail procedures are as follow:

  - Only if both ends support the 1.25Gbps TS, the remote end would choose the 1.25Gbps as the negotiated granularity for the HO ODU link. In other cases, the 2.5Gbps TS MUST be used (e.g., if the local end can support 1.25Gbps, and the remote end can support 2.5Gbps, and then the local end should imitate 2.5Gbps).

  - The remote end compares the two sets of LO ODU signal types that the local end and the remote end can support, and calculates the intersection of them, i.e., extracts all the LO ODU signal types that both two ends can support. This intersection is the set of LO ODU signal types that the HO ODU link can support.

- If both the two ends support the same capability, i.e., they support the same granularity of TS and the same LO ODU signal types, the remote end replies a LinkSummaryAck message to the local end. So the both ends know what capability the HO ODU link can support.

- If the two ends support different capabilities, i.e., they support different granularities of TS or different LO ODU signal types, the remote end replies a LinkSummaryNack message to the local end. The LinkSummaryNack message carries an ERROR_CODE object and one or more DATA_LINK objects. The ERROR_CODE "Renegotiate LINK_SUMMARY parameters" (see [RFC4204]) indicates that the two ends of the HO ODU link support different capabilities, and the DATA_LINK object carries the HO ODU link capability subobject which contains the negotiated granularity of TS and the set of LO ODU signal types that both ends can support. The local end can learn the HO ODU link capability after receiving the LinkSummaryNack message.

- If the remote end does not support the HO ODU link capability negotiation procedure, the LinkSummaryNack message MUST be responded with an ERROR_CODE "Not support of HO ODU Link Capability subobject" (TBA) indicating the reason of rejection.

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.

[RFC4204] J. Lang, Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.

[OTN-frwk] Fatai Zhang et al, "Framework for GMPLS and PCE Control of G.709 Optical Transport Networks", draft-ietf-ccamp-gmpls-g709-framework-05.txt, September 9, 2011.

[ITUT-G709] ITU-T, "Interface for the Optical Transport Network (OTN)", G.709 Recommendation, March 2003.

[G709-V3] ITU-T, "Interfaces for the Optical Transport Network (OTN)", G.709 Recommendation, December 2009.

9.2. Informative References

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4328] D. Papadimitriou, Ed. "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", RFC 4328, Jan 2006.

10. 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

Dan Li
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129 P.R.China

Phone: +86-755-28972910
Email: huawei.danli@huawei.com

Daniele Ceccarelli
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: daniele.ceccarelli@ericsson.com

Diego Caviglia
Ericsson
Via A. Negrone 1/A
Genova - Sestri Ponente
Italy

Email: diego.caviglia@ericsson.com

Guoying Zhang
China Academy of Telecommunication Research of MII
11 Yue Tan Nan Jie Beijing, P.R.China

Phone: +86-10-68094272
Email: zhangguoying@mail.ritt.com.cn

Pietro Grandi
Alcatel-Lucent
Optics CTO
Via Trento 30 20059 Vimercate (Milano) Italy
+39 039 6864930

Email: pietro_vittorio.grandi@alcatel-lucent.it

Sergio Belotti
Alcatel-Lucent
Optics CTO
Via Trento 30 20059 Vimercate (Milano) Italy
+39 039 6863033

Email: sergio.belotti@alcatel-lucent.it

11. Contributors

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

Disclaimer of Validity

Full Copyright Statement