

Network Working Group
Request for Comments: 4209
Category: Standards Track

A. Fredette, Ed.
Hatteras Networks
J. Lang, Ed.
Sonos Inc.
October 2005

Link Management Protocol (LMP) for
Dense Wavelength Division Multiplexing (DWDM) Optical Line Systems

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The Link Management Protocol (LMP) is defined to manage traffic engineering (TE) links. In its present form, LMP focuses on peer nodes, i.e., nodes that peer in signaling and/or routing. This document proposes extensions to LMP to allow it to be used between a peer node and an adjacent optical line system (OLS). These extensions are intended to satisfy the "Optical Link Interface Requirements" described in a companion document.

1. Introduction

Networks are being developed with routers, switches, optical cross-connects (OXCs), dense wavelength division multiplexing (DWDM) optical line systems (OLSeS), and add-drop multiplexors (ADMs) that use a common control plane (e.g., Generalized MPLS (GMPLS)) to dynamically provision resources and to provide network survivability using protection and restoration techniques.

The Link Management Protocol (LMP) is being developed as part of the GMPLS protocol suite to manage traffic engineering (TE) links [RFC4204]. In its present form, LMP focuses on peer nodes, i.e., nodes that peer in signaling and/or routing (e.g., OXC-to-OXC, as illustrated in Figure 1). In this document, extensions to LMP are proposed to allow it to be used between a peer node and an adjacent optical line system (OLS). These extensions are intended to satisfy

the "Optical Link Interface Requirements" described in [OLI]. It is assumed that the reader is familiar with LMP, as defined in [RFC4204].




Figure 1: LMP Model

Consider two peer nodes (e.g., two OXCs) interconnected by a wavelength-multiplexed link, i.e., a DWDM optical link (see Figure 1 above). Information about the configuration of this link and its current state is known by the two OLSes (OLS1 and OLS2). Allowing them to communicate this information to the corresponding peer nodes (OXC1 and OXC2) via LMP can improve network usability by reducing required manual configuration and by enhancing fault detection and recovery.

Information about the state of LSPs using the DWDM optical link is known by the peer nodes (OXC1 and OXC2), and allowing them to communicate this information to the corresponding OLSes (OLS1 and OLS2) is useful for alarm management and link monitoring. Alarm management is important because the administrative state of an LSP, known to the peer nodes (e.g., via the Admin Status object of GMPLS signaling [RFC3471]), can be used to suppress spurious alarm reporting from the OLSes.

The model for extending LMP to OLSes is shown in Figure 2.




Figure 2: Extended LMP Model

In this model, a peer node may have LMP sessions with adjacent OLSes, as well as adjacent peer nodes. In Figure 2, for example, the OXC1-OXC2 LMP session can be used to build traffic-engineering (TE) links for GMPLS signaling and routing, as described in [RFC4204]. The OXC1-OLS1 and the OXC2-OLS2 LMP sessions are used to exchange information about the configuration of the DWDM optical link and its current state and information about the state of LSPs using that link.

The latter type of LMP sessions is discussed in this document. It is important to note that a peer node may have LMP sessions with one or more OLSes and an OLS may have LMP sessions with one or more peer nodes.

Although there are many similarities between an LMP session between two peer nodes and an LMP session between a peer node and an OLS, there are some differences as well. The former type of LMP session is used to provide the basis for GMPLS signaling and routing. The latter type of LMP session is used to augment knowledge about the links between peer nodes.

A peer node maintains its peer node-to-OLS LMP sessions and its peer node-to-peer node LMP sessions independently. This means that it MUST be possible for LMP sessions to come up in any order. In particular, it MUST be possible for a peer node-to-peer node LMP session to come up in the absence of any peer node-to-OLS LMP sessions, and vice versa.

1.1. 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].

The reader is assumed to be familiar with the terminology in [RFC4204].

DWDM: Dense wavelength division multiplexing

OLS: Optical line system

Opaque:

A device is called X-opaque if it examines or modifies the X aspect of the signal while forwarding an incoming signal from input to output.

OXC: Optical cross-connect

Transparent:

As defined in [RFC4204], a device is called X-transparent if it forwards incoming signals from input to output without examining or modifying the X aspect of the signal. For example, a Frame Relay switch is network-layer transparent; an all-optical switch is electrically transparent.

1.2. Scope of LMP-WDM Protocol

This document focuses on extensions required for use with opaque OLSes. In particular, this document is intended for use with OLSes having SONET, SDH, and Ethernet user ports.

At the time of this writing, work is ongoing in the area of fully transparent wavelength routing; however, it is premature to identify the necessary information to be exchanged between a peer node and an OLS in this context. Nevertheless, the protocol described in this document provides the necessary framework in which to exchange additional information that is deemed appropriate.

2. LMP Extensions for Optical Line Systems

LMP currently consists of four main procedures, of which the first two are mandatory and the last two are optional:

1. Control channel management
2. Link property correlation
3. Link verification
4. Fault management

All four functions are supported in LMP-WDM.

2.1. Control Channel Management

As in [RFC4204], we do not specify the exact implementation of the control channel; it could be, for example, a separate wavelength, fiber, Ethernet link, an IP tunnel routed over a separate management network, a multi-hop IP network, or the overhead bytes of a data link.

The control channel management for a peer node-to-OLS link is the same as for a peer node-to-peer node link, as described in [RFC4204].

To distinguish between a peer node-to-OLS LMP session and a peer node-to-peer node LMP session, a new LMP-WDM CONFIG object is defined (C-Type = 2). The format of the CONFIG object is as follows:

Class = 6

o C-Type = 2, LMP-WDM_CONFIG

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|W|O|               (Reserved)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Reserved field should be sent as zero and ignored on receipt.

WDM: 1 bit

This bit indicates support for the LMP-WDM extensions defined in this document.

OLS: 1 bit

If set, this bit indicates that the sender is an optical line system (OLS). If clear, this bit indicates that the sender is a peer node.

The LMP-WDM extensions are designed for peer node-to-OLS LMP sessions. The OLS bit allows a node to identify itself as an OLS or a peer node. This is used to detect misconfiguration of a peer node-to-OLS LMP session between two peer nodes or a peer node-to-peer node LMP session between a peer node and an OLS.

If the node does not support the LMP-WDM extensions, it MUST reply to the Config message with a ConfigAck message.

If a peer node that is configured to run LMP-WDM receives a Config message with the OLS bit clear in LMP-WDM_CONFIG object, it MUST reply to the Config message with a ConfigAck message.

2.2. Link Verification

The Test procedure used with OLSes is the same as described in [RFC4204]. The VerifyTransportMechanism (included in the BeginVerify and BeginVerifyAck messages) is used to allow nodes to negotiate a link verification method and is essential for line systems that have access to overhead bytes rather than the payload. The VerifyId (provided by the remote node in the BeginVerifyAck message and used in all subsequent Test messages) is used to differentiate Test messages from different LMP Link Verification procedures. In

addition to the Test procedure described in [RFC4204], the trace monitoring function of [RFC4207] may be used for link verification when the OLS user ports are SONET or SDH.

In a combined LMP and LMP-WDM context, there is an interplay between the data links being managed by peer node-to-peer node LMP sessions and peer node-to-OLS LMP sessions. For example, in Figure 2, the OXC1-OLS1 LMP session manages the data links between OXC1 and OLS1, and the OXC2-OLS2 LMP session manages the data links between OXC2 and OLS2. However, the OXC1-OXC2 LMP session manages the data links between OXC1 and OXC2, which are actually a concatenation of the data links between OXC1 and OLS1, the DWDM span between OLS1 and OLS2, and the data links between OXC2 and OLS2. It is these concatenated links that comprise the TE links that are advertised in the GMPLS TE link state database.

The implication of this is that when the data links between OXC1 and OXC2 are being verified, using the LMP link verification procedure, OLS1 and OLS2 need to make themselves transparent with respect to these concatenated data links. The coordination of verification of OXC1-OLS1 and OXC2-OLS2 data links to ensure this transparency is the responsibility of the peer nodes, OXC1 and OXC2.

It is also necessary for these peer nodes to understand the mappings between the data links of the peer node - OLS LMP session and the concatenated data links of the peer node - peer node LMP session.

2.3. Link Summarization

As in [RFC4204], the LinkSummary message is used to synchronize the Interface_Ids and correlate the properties of the TE link. (Note that the term "TE link" originated from routing/signaling applications of LMP, and this concept does not necessarily apply to an OLS. However, the term is used in this document to remain consistent with LMP terminology.) The LinkSummary message includes one or more DATA_LINK objects. The contents of the DATA_LINK object consist of a series of variable-length data items called Data Link sub-objects describing the capabilities of the data links.

In this document, several additional Data Link sub-objects are defined to describe additional link characteristics. The link characteristics are, in general, those needed by the CSPF to select the path for a particular LSP. These link characteristics describe the specified peer node-to-OLS data link, as well as the associated DWDM span between the two OLSes.

The format of the Data Link sub-objects follows the format described in [RFC4204] and is shown below for readability:

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |   Length   |   (Sub-object contents)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 8 bits

The Type indicates the type of contents of the sub-object.

Length: 8 bits

The Length field contains the total length of the sub-object in bytes, including the Type and Length fields. The Length MUST be at least 4, and MUST be a multiple of 4.

The following link characteristics are exchanged on a per data link basis.

2.3.1. Link Group ID


The main purpose of the Link Group ID is to reduce control traffic during failures that affect many data links. A local ID may be assigned to a group of data links. This ID can be used to reduce the control traffic in the event of a failure by enabling a single ChannelStatus message with the LINK GROUP CHANNEL_STATUS object (see Section 2.4.1) to be used for a group of data links instead of individual ChannelStatus messages for each data link. A data link may be a member of multiple groups. This is achieved by including multiple Link Group ID sub-objects in the LinkSummary message.

The Link Group ID feature allows Link Groups to be assigned based on the types of fault correlation and aggregation supported by a given OLS. From a practical perspective, the Link Group ID is used to map (or group) data links into "failable entities" known primarily to the OLS. If one of those failable entities fails, all associated data links are failed and the peer node is notified with a single message.

For example, an OLS could create a Link Group for each laser in the OLS. The data links associated with each laser would then each be assigned the Link Group ID for that laser. If a laser fails, the OLS would then report a single failure affecting all of the data links with a Link Group ID of the failed laser. The peer node that receives the single failure notification then knows which data links are affected. Similarly, an OLS could create a Link Group ID for a

fiber, to report a failure affecting all of the data links associated with that fiber if a loss-of-signal (LOS) is detected for that fiber.

The format of the Link Group ID sub-object (Type = 3, Length = 8) is as follows:



The Reserved field should be sent as zero and ignored on receipt.


Link Group ID: 32 bits

Link Group ID 0xFFFFFFFF is reserved and indicates all data links in a TE link. All data links are members of Link Group 0xFFFFFFFF by default.

2.3.2. Shared Risk Link Group (SRLG) Identifier

This identifies the SRLGs of which the data link is a member. This information may be configured on an OLS by the user and used for diverse path computation (see [RFC4202]).

The format of the SRLG sub-object (Type = 4, Length = (N+1)*4 where N is the number of SRLG values) is as follows:



The Reserved field should be sent as zero and ignored on receipt.

Shared Risk Link Group Value: 32 bits

See [RFC4202]. List as many SRLGs as apply.

2.3.3. Bit Error Rate (BER) Estimate

This object provides an estimate of the BER for the data link.

The Bit Error Rate (BER) is the proportion of bits that have errors relative to the total number of bits received in a transmission, usually expressed as ten to a negative power. For example, a transmission might have a BER of "10 to the minus 13", meaning that, out of every 10,000,000,000,000 bits transmitted, one bit may be in error. The BER is an indication of overall signal quality.

The format of the BER Estimate sub-object (Type = 5; Length = 4) 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	2	3	4	5	6	7	8	9
Type										Length										BER										(Reserved)									

The Reserved field should be sent as zero and ignored on receipt.

BER: 8 bits

The exponent from the BER representation described above. That is, if the BER is 10 to the minus X, the BER field is set to X.

2.3.4. Optical Protection

This indicates whether the link is protected by the OLS. This information can be used as a measure of link capability. It may be advertised by routing and used by signaling as a selection criterion, as described in [RFC3471].

The format of the Optical Protection sub-object (Type = 6; Length = 4) 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	2	3	4	5	6	7	8	9
Type										Length										(Reserved)										Link Flags									

The Reserved field should be sent as zero and ignored on receipt.


Link Flags: 6 bits

Encoding for Link Flags is defined in Section 7 of [RFC3471].

2.3.5. Total Span Length

This indicates the total distance of fiber in the OLS. This may be used as a routing metric or to estimate delay.

The format of the Total Span Length sub-object (Type = 7, Length = 8) is as follows:



The Reserved field should be sent as zero and ignored on receipt.


Span Length: 32 bits

This value represents the total length of the WDM span in meters, expressed as an unsigned (long) integer.

2.3.6. Administrative Group (Color)

The administrative group (or Color) to which the data link belongs.

The format of the Administrative Group sub-object (Type = 8, Length = 8) is as follows:



The Reserved field should be sent as zero and ignored on receipt.

Administrative Group: 32 bits

A 32-bit value, as defined in [RFC3630].

2.4. Fault Management

The Fault Management procedure used between a peer and an OLS follows the procedures described in [RFC4204]; some further extensions are defined in this section. The information learned from the OLS-peer fault management procedures may be used to trigger peer-peer LMP fault management, or may be used to trigger GMPLS signaling/routing procedures directly.

Fault management consists of three major functions:

1. Fault Detection
2. Fault Localization
3. Fault Notification

The fault detection mechanisms are the responsibility of the individual nodes and are not specified as part of this protocol.

Fault detection mechanisms may include a Bit Error Rate (BER) exceeding a threshold, and loss-of-signal (LOS) and SONET/SDH-level errors. It is the responsibility of the OLS to translate these failures into (Signal) OK, Signal Failure (SF), or Signal Degrade (SD), as described in [RFC4204].

That is, an OLS uses the messages defined in the LMP fault localization procedures (ChannelStatus, ChannelStatusAck, ChannelStatusRequest, and ChannelStatusResponse messages) to inform the adjacent peer node of failures it has detected, in order to initiate the LMP fault localization procedures between peer nodes, but it does not participate in those procedures.

The OLS may also execute its own fault localization process to allow it to determine the location of the fault along the DWDM span. For example, the OLS may be able to pinpoint the fault to a particular amplifier in a span of thousands of kilometers in length.

To report data link failures and recovery conditions, LMP-WDM uses the ChannelStatus, ChannelStatusAck, ChannelStatusRequest, and ChannelStatusResponse messages defined in [RFC4204].

Each data link is identified by an Interface_ID. In addition, a Link Group ID may be assigned to a group of data links (see Section 2.3.1). The Link Group ID may be used to reduce the control traffic by providing channel status information for a group of data links. A new LINK GROUP CHANNEL_STATUS object is defined below for this purpose. This object may be used in place of the CHANNEL_STATUS objects described in [RFC4204] in the ChannelStatus message.

2.4.1. LINK_GROUP CHANNEL_STATUS Object

The LINK_GROUP CHANNEL_STATUS object is used to indicate the status of the data links belonging to a particular Link Group. The correlation of data links to Group ID is made with the Link Group ID sub-object of the DATA_LINK object.

The format of the LINK_GROUP CHANNEL_STATUS object is as follows (Class = 13, C-Type = 4):

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Link Group ID                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|A|D|                               Channel Status                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               :                               |
//                               :                               //
|                               :                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Link Group ID                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|A|D|                               Channel Status                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Link Group ID: 32 bits

The Link Group ID 0xFFFFFFFF is reserved and indicates all data links in a TE link. All data links are members of the Link Group 0xFFFFFFFF by default.

Channel Status: 32 bits

The values for the Channel Status field are defined in [RFC4204].

This object is non-negotiable.

3. 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. As such, this document introduces no other new security considerations not covered in [RFC4204].

4. 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 Object Class Types:

- o under CONFIG class name (as defined in [RFC4204])
 - LMP-WDM_CONFIG (C-Type = 2)
- o under CHANNEL_STATUS class name (as defined in [RFC4204])
 - LINK_GROUP (C-Type = 4)

LMP Sub-Object Class names:

- o under DATA_LINK Class name (as defined in [RFC4204])
 - Link_GroupId (sub-object Type = 3)
 - SRLG (sub-object Type = 4)
 - BER_Estimate (sub-object Type = 5)
 - Optical_Protection (sub-object Type = 6)
 - Total_Span_Length (sub-object Type = 7)
 - Administrative_Group (sub-object Type = 8)

5. Contributors

The authors would like to acknowledge Osama S. Aboul-Magd, Stuart Brorson, Sudheer Dharanikota, John Drake, David Drysdale, W. L. Edwards, Adrian Farrel, Andre Fredette, Rohit Goyal, Hirokazu Ishimatsu, Monika Jaeger, Ram Krishnan, Jonathan P. Lang, Raghu Mannam, Eric Mannie, Dimitri Papadimitriou, Jagan Shantigram, Ed Snyder, George Swallow, Gopala Tumuluri, Yong Xue, Lucy Yong, and John Yu.

6. References

6.1. Normative References

- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, September 2005.
- [RFC4204] Lang, J., Ed., "The Link Management Protocol (LMP)", RFC 4204, September 2005.
- [RFC4207] Lang, J., and D. Papadimitriou, "Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) Encoding for Link Management Protocol (LMP) Test Messages", RFC 4207, September 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

6.2. Informative References

- [OLI] Fredette, A., Editor, "Optical Link Interface Requirements", Work in Progress.

Editors' Addresses

Andre Fredette
Hatteras Networks
P.O. Box 110025
Research Triangle Park
NC 27709-0025, USA

EMail: Afredette@HatterasNetworks.com

Jonathan P. Lang
Sonos, Inc.
223 E. De La Guerra St.
Santa Barbara, CA 93101

EMail: jplang@ieee.org

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY 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 HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF 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 this 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. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR 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 this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

