

Internet Engineering Task Force (IETF)
Request for Comments: 7579
Category: Standards Track
ISSN: 2070-1721

G. Bernstein, Ed.
Grotto Networking
Y. Lee, Ed.
D. Li
Huawei
W. Imajuku
NTT
J. Han
Huawei
June 2015

General Network Element Constraint Encoding for GMPLS-Controlled Networks

Abstract

Generalized Multiprotocol Label Switching (GMPLS) can be used to control a wide variety of technologies. In some of these technologies, network elements and links may impose additional routing constraints such as asymmetric switch connectivity, non-local label assignment, and label range limitations on links.

This document provides efficient, protocol-agnostic encodings for general information elements representing connectivity and label constraints as well as label availability. It is intended that protocol-specific documents will reference this memo to describe how information is carried for specific uses.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7579>.

Copyright Notice

Copyright (c) 2015 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. Node Switching Asymmetry Constraints	3
1.2. Non-local Label Assignment Constraints	4
1.3. Conventions Used in This Document	4
2. Encoding	4
2.1. Connectivity Matrix Field	5
2.2. Port Label Restrictions Field	6
2.2.1. SIMPLE_LABEL	8
2.2.2. CHANNEL_COUNT	8
2.2.3. LABEL_RANGE	9
2.2.4. SIMPLE_LABEL & CHANNEL_COUNT	9
2.2.5. LINK_LABEL_EXCLUSIVITY	10
2.3. Link Set Field	10
2.4. Available Labels Field	12
2.5. Shared Backup Labels Field	13
2.6. Label Set Field	14
3. Security Considerations	16
4. IANA Considerations	17
5. References	17
5.1. Normative References	17
5.2. Informative References	18
Appendix A. Encoding Examples	19
A.1. Link Set Field	19
A.2. Label Set Field	19
A.3. Connectivity Matrix	20
A.4. Connectivity Matrix with Bidirectional Symmetry	24
A.5. Priority Flags in Available/Shared Backup Labels	26
Contributors	27
Authors' Addresses	28

1. Introduction

Some data-plane technologies that wish to make use of a GMPLS control plane contain additional constraints on switching capability and label assignment. In addition, some of these technologies must perform non-local label assignment based on the nature of the technology, e.g., wavelength continuity constraint in Wavelength Switched Optical Networks (WSONs) [RFC6163]. Such constraints can lead to the requirement for link-by-link label availability in path computation and label assignment.

This document provides efficient encodings of information needed by the routing and label assignment process in technologies such as WSON and are potentially applicable to a wider range of technologies. Such encodings can be used to extend GMPLS signaling and routing protocols. In addition, these encodings could be used by other mechanisms to convey this same information to a path computation element (PCE).

1.1. Node Switching Asymmetry Constraints

For some network elements, the ability of a signal or packet on a particular input port to reach a particular output port may be limited. Additionally, in some network elements (e.g., a simple multiplexer), the connectivity between some input and output ports may be fixed. To take into account such constraints during path computation, we model this aspect of a network element via a connectivity matrix.

The connectivity matrix (ConnectivityMatrix) represents either the potential connectivity matrix for asymmetric switches or fixed connectivity for an asymmetric device such as a multiplexer. Note that this matrix does not represent any particular internal blocking behavior but indicates which input ports and labels (e.g., wavelengths) could possibly be connected to a particular output port and label pair. Representing internal state-dependent blocking for a node is beyond the scope of this document and, due to its highly implementation-dependent nature, would most likely not be subject to standardization in the future. The connectivity matrix is a conceptual $M \times m$ by $N \times n$ matrix where M represents the number of input ports (each with m labels) and N the number of output ports (each with n labels).

1.2. Non-local Label Assignment Constraints

If the nature of the equipment involved in a network results in a requirement for non-local label assignment, we can have constraints based on limits imposed by the ports themselves and those that are implied by the current label usage. Note that constraints such as these only become important when label assignment has a non-local character. For example, in MPLS, an LSR may have a limited range of labels available for use on an output port and a set of labels already in use on that port; these are therefore unavailable for use. This information, however, does not need to be shared unless there is some limitation on the LSR's label swapping ability. For example, if a Time Division Multiplexer (TDM) node lacks the ability to perform time-slot interchange or a WSON lacks the ability to perform wavelength conversion, then the label assignment process is not local to a single node. In this case, it may be advantageous to share the label assignment constraint information for use in path computation.

Port label restrictions (PortLabelRestriction) model the label restrictions that the network element (node) and link may impose on a port. These restrictions tell us what labels may or may not be used on a link and are intended to be relatively static. More dynamic information is contained in the information on available labels. Port label restrictions are specified relative to the port in general or to a specific connectivity matrix for increased modeling flexibility. [Switch] gives an example where both switch and fixed connectivity matrices are used and both types of constraints occur on the same port.

1.3. 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].

2. Encoding

This section provides encodings for the information elements defined in [RFC7446] that have applicability to WSON. The encodings are designed to be suitable for use in the GMPLS routing protocols OSPF [RFC4203] and IS-IS [RFC5307] and in the PCE Communication Protocol (PCEP) [RFC5440]. Note that the information distributed in [RFC4203] and [RFC5307] is arranged via the nesting of sub-TLVs within TLVs; this document defines elements to be used within such constructs. Specific constructs of sub-TLVs and the nesting of sub-TLVs of the information element defined by this document will be defined in the respective protocol enhancement documents.

2.1. Connectivity Matrix Field

The Connectivity Matrix Field represents how input ports are connected to output ports for network elements. The switch and fixed connectivity matrices can be compactly represented in terms of a minimal list of input and output port set pairs that have mutual connectivity. As described in [Switch], such a minimal list representation leads naturally to a graph representation for path computation purposes; this representation involves the fewest additional nodes and links.

The Connectivity Matrix Field is uniquely identified only by the advertising node. There may be more than one Connectivity Matrix Field associated with a node as a node can partition the switch matrix into several sub-matrices. This partitioning is primarily to limit the size of any individual information element used to represent the matrix and to enable incremental updates. When the matrix is partitioned into sub-matrices, each sub-matrix will be mutually exclusive to one another in representing which ports/labels are associated with each sub-matrix. This implies that two matrices will not have the same {src port, src label, dst port, dst label}.

Each sub-matrix is identified via a different Matrix ID that MUST represent a unique combination of {src port, src label, dst port, dst label}.

A TLV encoding of this list of link set pairs is:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
| Conn  |      MatrixID      |      Reserved      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Link Set A #1   |
|                                     :               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Link Set B #1   |
|                                     :               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Additional Link Set Pairs as Needed
|                                     to Specify Connectivity
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where:

Connectivity (Conn) (4 bits) is the device type.

- 0 - the device is fixed

- 1 - the device is switched (e.g., Reconfigurable Optical Add/Drop Multiplexer / Optical Cross-Connect (ROADM/OXC))

MatrixID represents the ID of the connectivity matrix and is an 8-bit integer. The value of 0xFF is reserved for use with port label constraints and should not be used to identify a connectivity matrix.

Link Set A #1 and Link Set B #1 together represent a pair of link sets. See Section 2.3 for a detailed description of the Link Set Field. There are two permitted combinations for the Link Set Field parameter "dir" for link set A and B pairs:

- o Link Set A dir=input, Link Set B dir=output

In this case, the meaning of the pair of link sets A and B is that any signal that inputs a link in set A can be potentially switched out of an output link in set B.

- o Link Set A dir=bidirectional, Link Set B dir=bidirectional

In this case, the meaning of the pair of link sets A and B is that any signal that inputs on the links in set A can potentially output on a link in set B and any input signal on the links in set B can potentially output on a link in set A. If link set A is an input and link set B is an output for a signal, then it implies that link set A is an output and link set B is an input for that signal.

See Appendix A for both types of encodings as applied to a ROADM example.

2.2. Port Label Restrictions Field

The Port Label Restrictions Field tells us what labels may or may not be used on a link.

The port label restrictions can be encoded as follows. More than one of these fields may be needed to fully specify a complex port constraint. When more than one of these fields is present, the resulting restriction is the union of the restrictions expressed in

each field. The use of the reserved value of 0xFF for the MatrixID indicates that a restriction applies to the port and not to a specific connectivity matrix.

```

      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 | Switching Cap | Encoding |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Additional Restriction Parameters per Restriction Type |
:                                                         :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Where:

MatrixID: either is the value in the corresponding Connectivity Matrix Field or takes the value 0xFF to indicate the restriction applies to the port regardless of any connectivity matrix.

RstType (Restriction Type) can take the following values and meanings:

- 0: SIMPLE_LABEL (Simple label selective restriction). See Section 2.2.1 for details.
- 1: CHANNEL_COUNT (Channel count restriction). See Section 2.2.2 for details.
- 2: LABEL_RANGE (Label range device with a movable center label and width). See Section 2.2.3 for details.
- 3: SIMPLE_LABEL & CHANNEL_COUNT (Combination of SIMPLE_LABEL and CHANNEL_COUNT restriction. The accompanying label set and channel count indicate labels permitted on the port and the maximum number of channels that can be simultaneously used on the port). See Section 2.2.4 for details.
- 4: LINK_LABEL_EXCLUSIVITY (A label may be used at most once amongst a set of specified ports). See Section 2.2.5 for details.

Switching Cap (Switching Capability) is defined in [RFC4203], and LSP Encoding Type is defined in [RFC3471]. The combination of these fields defines the type of labels used in specifying the port label restrictions as well as the interface type to which these restrictions apply.

The Additional Restriction Parameters per RestrictionType field is an optional field that describes additional restriction parameters for each RestrictionType pertaining to specific protocols.

2.2.1. SIMPLE_LABEL

In the case of SIMPLE_LABEL, the format is:

```

      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 = 0   | Switching Cap | Encoding   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Label Set Field
+-----+-----+-----+-----+-----+-----+-----+-----+

```

In this case, the accompanying label set indicates the labels permitted on the port/matrix.

See Section 2.6 for the definition of label set.

2.2.2. CHANNEL_COUNT

In the case of CHANNEL_COUNT, the format is:

```

      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 = 1   | Switching Cap | Encoding   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     MaxNumChannels
+-----+-----+-----+-----+-----+-----+-----+-----+

```

In this case, the accompanying MaxNumChannels indicates the maximum number of channels (labels) that can be simultaneously used on the port/matrix.

MaxNumChannels is a 32-bit integer.

2.2.3. LABEL_RANGE

In the case of LABEL_RANGE, the format is:

```

      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 = 2  | Switching Cap | Encoding  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     MaxLabelRange                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Label Set Field                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

This is a generalization of the waveband device. The MaxLabelRange indicates the maximum width of the waveband in terms of the channels spacing given in the Label Set Field. The corresponding label set is used to indicate the overall tuning range.

MaxLabelRange is a 32-bit integer.

See Section 2.6.2 for an explanation of label range.

2.2.4. SIMPLE_LABEL & CHANNEL_COUNT

In the case of SIMPLE_LABEL & CHANNEL_COUNT, the format is:

```

      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 = 3  | Switching Cap | Encoding  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     MaxNumChannels                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Label Set Field                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

In this case, the accompanying label set and MaxNumChannels indicate labels permitted on the port and the maximum number of labels that can be simultaneously used on the port.

See Section 2.6 for the definition of label set.

2.2.5. LINK_LABEL_EXCLUSIVITY

In the case of Link Label Exclusivity, the format is:

```

      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 = 4  | Switching Cap | Encoding  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Link Set Field
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

In this case, the accompanying link set indicates that a label may be used at most once among the ports in the Link Set Field.

See Section 2.3 for the definition of link set.

2.3. Link Set Field

We will frequently need to describe properties of groups of links. To do so efficiently, we can make use of a link set concept similar to the label set concept of [RFC3471]. The Link Set Field is used in the <ConnectivityMatrix>, which is defined in Section 2.1. The information carried in a link set 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Action      |Dir|  Format  |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Link Identifier 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                                                         :
:                                                         :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Link Identifier N
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Action: 8 bits

0 - Inclusive List

Indicates that one or more link identifiers are included in the link set. Each identifies a separate link that is part of the set.

1 - Inclusive Range

Indicates that the link set defines a range of links. It contains two link identifiers. The first identifier indicates the start of the range. The second identifier indicates the end of the range. All links with numeric values between the bounds are considered to be part of the set. A value of zero in either position indicates that there is no bound on the corresponding portion of the range. Note that the Action field can be set to 0x01 (Inclusive Range) only when the identifier for unnumbered link is used.

Dir: Directionality of the link set (2 bits)

0 - bidirectional

1 - input

2 - output

In optical networks, we think in terms of unidirectional and bidirectional links. For example, label restrictions or connectivity may be different for an input port than for its "companion" output port, if one exists. Note that "interfaces" such as those discussed in the Interfaces MIB [RFC2863] are assumed to be bidirectional. This also applies to the links advertised in various link state routing protocols.

Format: The format of the link identifier (6 bits)

0 - Link Local Identifier

Indicates that the links in the link set are identified by link local identifiers. All link local identifiers are supplied in the context of the advertising node.

1 - Local Interface IPv4 Address

Indicates that the links in the link set are identified by Local Interface IPv4 Address.

2 - Local Interface IPv6 Address

Indicates that the links in the link set are identified by Local Interface IPv6 Address.

Others - Reserved for future use

Note that all link identifiers in the same list must be of the same type.

Length: 16 bits


This field indicates the total length in bytes of the Link Set Field.

Link Identifier: length is dependent on the link format

The link identifier represents the port that is being described either for connectivity or for label restrictions. This can be the link local identifier of GMPLS routing [RFC4202], GMPLS OSPF routing [RFC4203], and IS-IS GMPLS routing [RFC5307]. The use of the link local identifier format can result in more compact encodings when the assignments are done in a reasonable fashion.

2.4. Available Labels Field

The Available Labels Field consists of priority flags and a single variable-length Label Set Field as follows:



Where:

PRI (Priority Flags, 8 bits): A bitmap used to indicate which priorities are being advertised. The bitmap is in ascending order, with the leftmost bit representing priority level 0 (i.e., the highest) and the rightmost bit representing priority level 7 (i.e., the lowest). A bit MUST be set (1) corresponding to each priority represented in the sub-TLV and MUST NOT be set (0) when the corresponding priority is not represented. If a label is available at priority M, it MUST be advertised available at each priority N < M. At least one priority level MUST be advertised.

The PRI field indicates the availability of the labels for use in Label Switched Path (LSP) setup and preemption as described in [RFC3209].


When a label is advertised as available for priorities 0, 1, ... M, it may be used by any LSP of priority $N \leq M$. When a label is in use by an LSP of priority M, it may be used by an LSP of priority $N < M$ if LSP preemption is supported.

When a label was initially advertised as available for priorities 0, 1, ... M and once a label is used for an LSP at a priority, say N ($N \leq M$), then this label is advertised as available for 0, ... N-1.

Note that the Label Set Field is defined in Section 2.6. See Appendix A.5 for illustrative examples.

2.5. Shared Backup Labels Field

The Shared Backup Labels Field consists of priority flags and a single variable-length Label Set Field as follows:



Where:

PRI (Priority Flags, 8 bits): A bitmap used to indicate which priorities are being advertised. The bitmap is in ascending order, with the leftmost bit representing priority level 0 (i.e., the highest) and the rightmost bit representing priority level 7 (i.e., the lowest). A bit MUST be set (1) corresponding to each priority represented in the sub-TLV and MUST NOT be set (0) when the corresponding priority is not represented. If a label is available at priority M, it MUST be advertised available at each priority $N < M$. At least one priority level MUST be advertised.

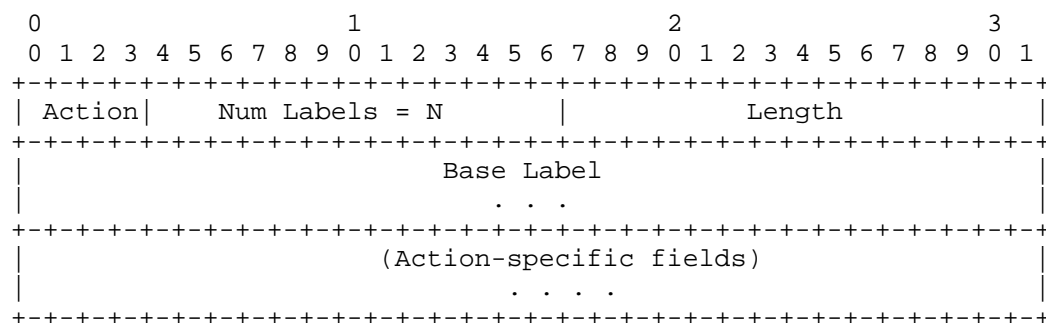
The same LSP setup and preemption rules specified in Section 2.4 apply here.

Note that Label Set Field is defined in Section 2.6. See Appendix A.5 for illustrative examples.

2.6. Label Set Field

The Label Set Field is used within the Available Labels Field or the Shared Backup Labels Field, defined in Sections 2.4 and 2.5, respectively. It is also used within SIMPLE_LABEL, LABEL_RANGE, or SIMPLE_LABEL & CHANNEL_COUNT, defined in Sections 2.2.1, 2.2.3, and 2.2.4, respectively.

The general format for a label set is given below. This format uses the Action concept from [RFC3471] with an additional Action to define a "bitmap" type of label set. Labels are variable in length. Action-specific fields are defined in Sections 2.6.1, 2.6.2, and 2.6.3.



Action:

- 0 - Inclusive List
- 1 - Exclusive List
- 2 - Inclusive Range
- 3 - Exclusive Range
- 4 - Bitmap Set

Num Labels is generally the number of labels. It has a specific meaning depending on the Action value. See Sections 2.6.1, 2.6.2, and 2.6.3 for details. Num Labels is a 12-bit integer.

Length is the length in bytes of the entire Label Set Field.

2.6.1. Inclusive/Exclusive Label Lists

For inclusive/exclusive lists (Action = 0 or 1), the wavelength set format is:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 or 1 | Num Labels = 2 |                               Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Label #1                       |
|                               . . .                         |
+-----+-----+-----+-----+-----+-----+-----+-----+
:                                                                 :
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Label #N                       |
|                               . . .                         |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Label #1 is the first label to be included/excluded, and Label #N is the last label to be included/excluded. Num Labels MUST match with N.

2.6.2. Inclusive/Exclusive Label Ranges

For inclusive/exclusive ranges (Action = 2 or 3), the label set format is:

```

      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 or 3 | Num Labels |                               Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Start Label                     |
|                               . . .                         |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               End Label                       |
|                               . . .                         |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Note that Start Label is the first label in the range to be included/excluded, and End Label is the last label in the same range. Num Labels MUST be two.

4. IANA Considerations

This document provides general protocol-independent information encodings. There is no IANA allocation request for the information elements defined in this document. IANA allocation requests will be addressed in protocol-specific documents based on the encodings defined here.

5. References

5.1. Normative References

- [G.694.1] ITU-T, "Spectral grids for WDM applications: DWDM frequency grid", ITU-T Recommendation G.694.1, February 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2863] McCloghrie, K. and F. Kastenholtz, "The Interfaces Group MIB", RFC 2863, DOI 10.17487/RFC2863, June 2000, <<http://www.rfc-editor.org/info/rfc2863>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <<http://www.rfc-editor.org/info/rfc3471>>.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, DOI 10.17487/RFC4202, October 2005, <<http://www.rfc-editor.org/info/rfc4202>>.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<http://www.rfc-editor.org/info/rfc4203>>.

- [RFC5307] Kompella, K., Ed., and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <<http://www.rfc-editor.org/info/rfc5307>>.
- [RFC6205] Otani, T., Ed., and D. Li, Ed., "Generalized Labels for Lambda-Switch-Capable (LSC) Label Switching Routers", RFC 6205, DOI 10.17487/RFC6205, March 2011, <<http://www.rfc-editor.org/info/rfc6205>>.

5.2. Informative References

- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.
- [RFC6163] Lee, Y., Ed., Bernstein, G., Ed., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", RFC 6163, DOI 10.17487/RFC6163, April 2011, <<http://www.rfc-editor.org/info/rfc6163>>.
- [RFC7446] Lee, Y., Ed., Bernstein, G., Ed., Li, D., and W. Imajuku, "Routing and Wavelength Assignment Information Model for Wavelength Switched Optical Networks", RFC 7446, DOI 10.17487/RFC7446, February 2015, <<http://www.rfc-editor.org/info/rfc7446>>.
- [Switch] Bernstein, G., Lee, Y., Gavler, A., and J. Martensson, "Modeling WDM Wavelength Switching Systems for Use in GMPLS and Automated Path Computation", Journal of Optical Communications and Networking, Volume 1, Issue 1, pp. 187-195, June 2009.

Appendix A. Encoding Examples

This appendix contains examples of the general encoding extensions applied to some simple ROADM network elements and links.

A.1. Link Set Field

Suppose that we wish to describe a set of input ports that have link local identifiers numbered 3 through 42. In the Link Set Field, we set Action = 1 to denote an inclusive range, Dir = 1 to denote input links, and Format = 0 to denote link local identifiers. Thus, we have:

```

+-----+
| Action=1      | 0 1 | 0 0 0 0 0 0 | Length = 12 |
+-----+
| Link Local Identifier = #3 |
+-----+
| Link Local Identifier = #42 |
+-----+

```

A.2. Label Set Field

In this example, we use a 40-channel C-Band Dense Wavelength Division Multiplexing (DWDM) system with 100 GHz spacing with lowest frequency 192.0 THz (1561.4 nm) and highest frequency 195.9 THz (1530.3 nm). These frequencies correspond to $n = -11$ and $n = 28$, respectively. Now suppose the following channels are available:

Frequency (THz)	n Value	bitmap position
192.0	-11	0
192.5	-6	5
193.1	0	11
193.9	8	19
194.0	9	20
195.2	21	32
195.8	27	38

Using the label format defined in [RFC6205], with the Grid value set to indicate an ITU-T A/2 [G.694.1] DWDM grid and C.S. set to indicate 100 GHz, this lambda bitmap set would then be encoded 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
+-----+-----+-----+-----+-----+-----+-----+-----+
| 4      | Num Labels = 40      | Length = 16 bytes      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = -11 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|1 0 0 0 0 1 0 0 0 0 0 1 0 0 0 0 0 0 0 1 1 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|1 0 0 0 0 0 1 0| Not used in 40 Channel system (all zeros) |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

To encode this same set as an inclusive list, we would have:

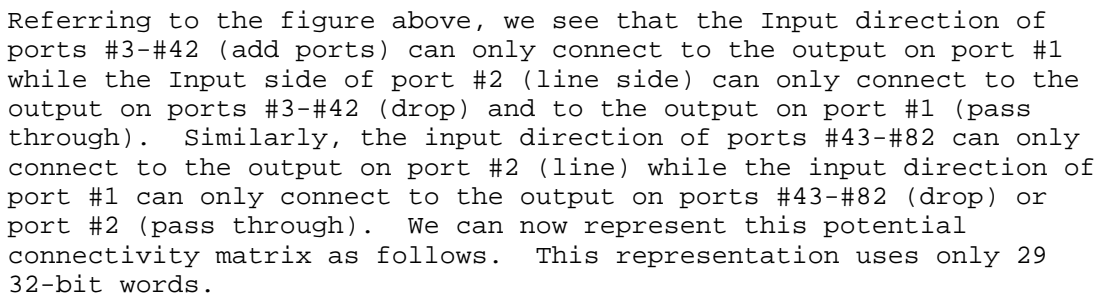
```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0      | Num Labels = 7      | Length = 32 bytes      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = -11 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = -6  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = -0  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = 8   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = 9   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = 21  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|Grid | C.S. |      Reserved      | n for lowest frequency = 27  |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

A.3. Connectivity Matrix

Suppose we have a typical 2-degree 40-channel ROADM. In addition to its two line side ports, it has 80 add and 80 drop ports. The figure below illustrates how a typical 2-degree ROADM system that works with bidirectional fiber pairs is a highly asymmetrical system composed of two unidirectional ROADM subsystems.



```

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
+-----+-----+-----+-----+
| Conn = 1 | MatrixID | Reserved |
+-----+-----+-----+-----+
Note: adds to line
+-----+-----+-----+-----+
| Action=1 | 0 1|0 0 0 0 0 0| Length = 12 |
+-----+-----+-----+-----+
| Link Local Identifier = #3 |
+-----+-----+-----+-----+
| Link Local Identifier = #42 |
+-----+-----+-----+-----+
| Action=0 | 1 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #1 |
+-----+-----+-----+-----+
Note: line to drops
+-----+-----+-----+-----+
| Action=0 | 0 1|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #2 |
+-----+-----+-----+-----+
| Action=1 | 1 0|0 0 0 0 0 0| Length = 12 |
+-----+-----+-----+-----+
| Link Local Identifier = #3 |
+-----+-----+-----+-----+
| Link Local Identifier = #42 |
+-----+-----+-----+-----+
Note: line to line
+-----+-----+-----+-----+
| Action=0 | 0 1|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #2 |
+-----+-----+-----+-----+
| Action=0 | 1 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #1 |
+-----+-----+-----+-----+
Note: adds to line
+-----+-----+-----+-----+
| Action=1 | 0 1|0 0 0 0 0 0| Length = 12 |
+-----+-----+-----+-----+
| Link Local Identifier = #43 |
+-----+-----+-----+-----+
| Link Local Identifier = #82 |
+-----+-----+-----+-----+
| Action=0 | 1 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+

```

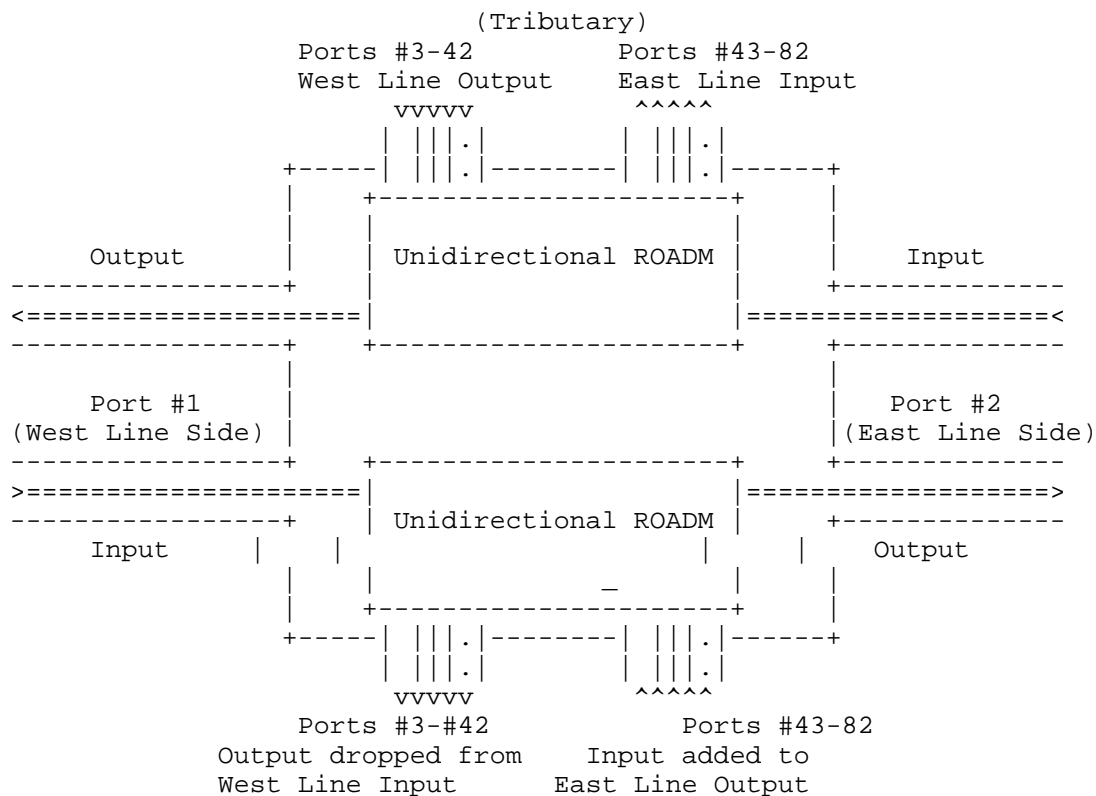
```

+++++
|                               Link Local Identifier = #2                               |
+++++
|                               Note: line to drops                               |
+++++
| Action=0      |0 1|0 0 0 0 0 0|                               Length = 8      |
+++++
|                               Link Local Identifier = #1                               |
+++++
| Action=1      |1 0|0 0 0 0 0 0|                               Length = 12     |
+++++
|                               Link Local Identifier = #43                              |
+++++
|                               Link Local Identifier = #82                              |
+++++
|                               Note: line to line                               |
+++++
| Action=0      |0 1|0 0 0 0 0 0|                               Length = 8      |
+++++
|                               Link Local Identifier = #1                               |
+++++
| Action=0      |1 0|0 0 0 0 0 0|                               Length = 8      |
+++++
|                               Link Local Identifier = #2                               |
+++++

```

A.4. Connectivity Matrix with Bidirectional Symmetry

If one has the ability to renumber the ports of the previous example as shown in the next figure, then we can take advantage of the bidirectional symmetry and use bidirectional encoding of the connectivity matrix. Note that we set `dir=bidirectional` in the `Link Set Fields`.




```

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
+-----+-----+-----+-----+
| Conn = 1 | MatrixID | Reserved |
+-----+-----+-----+-----+
Note: Add/Drop #3-42 to Line side #1
+-----+-----+-----+-----+
| Action=1 | 0 0|0 0 0 0 0 0| Length = 12 |
+-----+-----+-----+-----+
| Link Local Identifier = #3 |
+-----+-----+-----+-----+
| Link Local Identifier = #42 |
+-----+-----+-----+-----+
| Action=0 | 0 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #1 |
+-----+-----+-----+-----+
Note: line #2 to add/drops #43-82
+-----+-----+-----+-----+
| Action=0 | 0 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #2 |
+-----+-----+-----+-----+
| Action=1 | 0 0|0 0 0 0 0 0| Length = 12 |
+-----+-----+-----+-----+
| Link Local Identifier = #43 |
+-----+-----+-----+-----+
| Link Local Identifier = #82 |
+-----+-----+-----+-----+
Note: line to line
+-----+-----+-----+-----+
| Action=0 | 0 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #1 |
+-----+-----+-----+-----+
| Action=0 | 0 0|0 0 0 0 0 0| Length = 8 |
+-----+-----+-----+-----+
| Link Local Identifier = #2 |
+-----+-----+-----+-----+

```

If one wants to make a set of labels (indicated by Label Set Field #1) available only for the highest priority level (Priority Level 0) while allowing a set of labels (indicated by Label Set Field #2) to be available to all priority levels, the following encoding will express such need.

	<u>0</u>									<u>1</u>									<u>2</u>									<u>3</u>							
+	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	+		
	1	0	0	0	0	0	0	0	0		Reserved																								
+	+																																		
	Label Set Field #1																																		
:																									:										
+	+																																		
	1	1	1	1	1	1	1	1	1		Reserved																								
+	+																																		
	Label Set Field #2																																		
:																									:										
+	+																																		

Contributors

Diego Caviglia
Ericsson
Via A. Negrone 1/A 16153
Genoa
Italy
Phone: +39 010 600 3736
EMail: diego.caviglia@ericsson.com

Anders Gavler
Acreo AB
Electrum 236
SE - 164 40 Kista
Sweden
EMail: Anders.Gavler@acreo.se

Jonas Martensson
Acreo AB
Electrum 236
SE - 164 40 Kista
Sweden
EMail: Jonas.Martensson@acreo.se

Itaru Nishioka
NEC Corp.
1753 Simonumabe
Nakahara-ku, Kawasaki, Kanagawa 211-8666
Japan
Phone: +81 44 396 3287
EMail: i-nishioka@cb.jp.nec.com

Rao Rajan
Infinera
EMail: rrao@infinera.com

Giovanni Martinelli
Cisco
EMail: giomarti@cisco.com

Remi Theillaud
Marben
EMail: remi.theillaud@marben-products.com

Authors' Addresses

Greg M. Bernstein (editor)
Grotto Networking
Fremont, California
United States
Phone: (510) 573-2237
EMail: gregb@grotto-networking.com

Young Lee (editor)
Huawei Technologies
1700 Alma Drive, Suite 100
Plano, TX 75075
United States
Phone: (972) 509-5599 (x2240)
EMail: ylee@huawei.com

Dan Li
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129
China
Phone: +86-755-28973237
EMail: danli@huawei.com

Wataru Imajuku
NTT Network Innovation Labs
1-1 Hikari-no-oka, Yokosuka, Kanagawa
Japan
Phone: +81-(46) 859-4315
EMail: imajuku.wataru@lab.ntt.co.jp

Jianrui Han
Huawei Technologies Co., Ltd.
F3-5-B R&D Center, Huawei Base,
Bantian, Longgang District
Shenzhen 518129
China
Phone: +86-755-28972916
EMail: hanjianrui@huawei.com

