Network Working Group
Request for Comments: 4606
Obsoletes: 3946
Category: Standards Track
E. Mannie
Perceval
D. Papadimitriou
Alcatel
August 2006

Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control

Status of This Memo

Copyright Notice

Abstract

This document provides minor clarification to RFC 3946.

This document is a companion to the Generalized Multi-protocol Label Switching (GMPLS) signaling. It defines the Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) technology-specific information needed when GMPLS signaling is used.

Table of Contents

1. Introduction

As described in [RFC3945], Generalized MPLS (GMPLS) extends MPLS from supporting packet (Packet Switching Capable, or PSC) interfaces and switching to include support of four new classes of interfaces and switching: Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC) and Fiber-Switch Capable (FSC). A functional description of the extensions to MPLS signaling needed to support the new classes of interfaces and switching is provided in [RFC3471]. [RFC3473] describes RSVP-TE-specific formats and mechanisms needed to support all five classes of interfaces, and CR-LDP extensions can be found in [RFC3472].

This document presents details that are specific to Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH). Per [RFC3471], SONET/SDH-specific parameters are carried in the signaling protocol in traffic parameter specific objects.

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].

Moreover, the reader is assumed to be familiar with the terminology in American National Standards Institute (ANSI) [T1.105] and ITU-T [G.707], as well as with that in [RFC3471], [RFC3472], and [RFC3473]. The following abbreviations are used in this document:

DCC: Data Communications Channel.
LOVC: Lower-Order Virtual Container
HOVC: Higher-Order Virtual Container
MS: Multiplex Section.
MSOH: Multiplex Section overhead.
POH: Path overhead.
RS: Regenerator Section.
RSOH: Regenerator Section overhead.
SDH: Synchronous digital hierarchy.
SOH: Section overhead.
SONET: Synchronous Optical Network.
SPE: Synchronous Payload Envelope.
STM(-N): Synchronous Transport Module (-N) (SDH).
STS(-N): Synchronous Transport Signal-Level N (SONET).
VC-n: Virtual Container-n (SDH).
VTn: Virtual Tributary-n (SONET).

2. SONET and SDH Traffic Parameters

This section defines the GMPLS traffic parameters for SONET/SDH. The protocol-specific formats, for the SONET/SDH-specific RSVP-TE objects and CR-LDP TLVs, are described in Sections 2.2 and 2.3, respectively.

These traffic parameters specify a base set of capabilities for SONET ANSI [T1.105] and SDH ITU-T [G.707], such as concatenation and transparency. Other documents may further enhance this set of capabilities in the future. For instance, signaling for SDH over PDH ITU-T G.832 or sub-STM-0 ITU-T G.708 interfaces could be defined.

The traffic parameters defined hereafter (see Section 2.1) MUST be used when the label is encoded as SUKLM as defined in this memo (see Section 3). They MUST also be used when requesting one of Section/RS or Line/MS overhead transparent STS-1/STM-0, STS-3*N/STM-N (N=1, 4, 16, 64, 256) signals.

The traffic parameters and label encoding defined in [RFC3471], Section 3.2, MUST be used for fully transparent STS-1/STM-0, STS-3*N/STM-N (N=1, 4, 16, 64, 256) signal requests. A fully transparent signal is one for which all overhead is left unmodified by intermediate nodes; i.e., when all defined Transparency (T) bits would be set if the traffic parameters defined in Section 2.1 were used.

2.1. SONET/SDH Traffic Parameters

The traffic parameters for SONET/SDH are organized 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Signal Type  |      RCC      |              NCC              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              NVC              |        Multiplier (MT)        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Transparency (T)                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Profile (P)                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Annex 1 lists examples of SONET and SDH signal coding.

o) Signal Type (ST): 8 bits

This field indicates the type of Elementary Signal that constitutes the requested Label Switched Path (LSP). Several transforms can be applied successively on the Elementary Signal to build the Final Signal actually being requested for the LSP.

Each transform application is optional and must be ignored if zero, except the Multiplier (MT), which cannot be zero and is ignored if equal to one.

Transforms must be applied strictly in the following order:

- First, contiguous concatenation (by using the RCC and NCC fields) can be optionally applied on the Elementary Signal, resulting in a contiguously concatenated signal.

- Second, virtual concatenation (by using the NVC field) can be optionally applied on the Elementary Signal, resulting in a virtually concatenated signal.

- Third, some transparency (by using the Transparency field) can be optionally specified when a frame is requested as signal rather than an SPE- or VC-based signal.

- Fourth, a multiplication (by using the Multiplier field) can be optionally applied directly on the Elementary Signal, on the contiguously concatenated signal obtained from the first phase, on the virtually concatenated signal obtained from the second phase, or on these signals combined with some transparency.

Permitted Signal Type values for SONET/SDH are

```
Value  Type (Elementary Signal)
-----  ------------------------
  1     VT1.5  SPE / VC-11
  2     VT2    SPE / VC-12
  3     VT3    SPE
  4     VT6    SPE / VC-2
  5     STS-1  SPE / VC-3
  6     STS-3c SPE / VC-4
  7     STS-1      / STM-0   (only when transparency is requested)
  8     STS-3      / STM-1   (only when transparency is requested)
  9     STS-12     / STM-4   (only when transparency is requested)
  10    STS-48     / STM-16  (only when transparency is requested)
  11    STS-192    / STM-64  (only when transparency is requested)
  12    STS-768    / STM-256 (only when transparency is requested)
```

A dedicated signal type is assigned to a SONET STS-3c SPE instead of being coded as a contiguous concatenation of three STS-1 SPEs. This is done in order to provide easy interworking between SONET and SDH signaling.

Appendix 1 adds one signal type (optional) to the above values.

o) Requested Contiguous Concatenation (RCC): 8 bits

This field is used to request the optional SONET/SDH contiguous concatenation of the Elementary Signal.

This field is a vector of flags. Each flag indicates the support of a particular type of contiguous concatenation. Several flags can be set at the same time to indicate a choice.

These flags allow an upstream node to indicate to a downstream node the different types of contiguous concatenation that it supports. However, the downstream node decides which one to use according to its own rules.

A downstream node receiving simultaneously more than one flag chooses a particular type of contiguous concatenation, if any is supported, and according to criteria that are out of this document’s scope. A downstream node that doesn’t support any of the concatenation types indicated by the field must refuse the LSP request. In particular, it must refuse the LSP request if it doesn’t support contiguous concatenation at all.

When several flags have been set, the upstream node retrieves the (single) type of contiguous concatenation the downstream node has selected by looking at the position indicated by the first label and the number of labels as returned by the downstream node (see also Section 3).

The entire field is set to zero to indicate that no contiguous concatenation is requested at all (default value). A non-zero field indicates that some contiguous concatenation is requested.

The following flag is defined:

Flag 1 (bit 1): Standard contiguous concatenation.

Flag 1 indicates that the standard SONET/SDH contiguous concatenation, as defined in [T1.105]/[G.707], is supported. Note that bit 1 is the low-order bit. Other flags are reserved for extensions; if not used, they must be set to zero when sent and should be ignored when received.

See note 1 in the section on the NCC about the SONET contiguous concatenation of STS-1 SPEs when the number of components is a multiple of three.

o) Number of Contiguous Components (NCC): 16 bits

This field indicates the number of identical SONET SPEs/SDH VCs (i.e., Elementary Signal) that are requested to be concatenated, as specified in the RCC field.

Note 1: When a SONET STS-Nc SPE with N=3*X is requested, the Elementary Signal to be used must always be an STS-3c_SPE signal type, and the value of NCC must always be equal to X. This allows facilitating the interworking between SONET and SDH. In particular, it means that the contiguous concatenation of three STS-1 SPEs cannot be requested, as according to this specification this type of signal must be coded using the STS-3c SPE signal type.

Note 2: When a transparent STS-N/STM-N signal is requested that is limited to a single contiguously concatenated STS-Nc_SPE/VC-4-Nc, the signal type must be STS-N/STM-N, RCC with flag 1, NCC set to 1.

The NCC value must be consistent with the type of contiguous concatenation being requested in the RCC field. In particular, this field is irrelevant if no contiguous concatenation is requested (RCC = 0). In that case, it must be set to zero when sent and should be ignored when received. A RCC value different from 0 implies a number of contiguous components greater than or equal to 1.

Note 3: Following these rules, when a VC-4 signal is requested, the RCC and the NCC values SHOULD be set to 0, whereas for an STS-3c SPE signal, the RCC and the NCC values SHOULD be set 1. However, if local conditions allow, since the setting of the RCC and NCC values is locally driven, the requesting upstream node MAY set the RCC and NCC values to either SDH or SONET settings without impacting the function. Moreover, the downstream node SHOULD accept the requested values if local conditions allow. If these values cannot be supported, the receiver downstream node SHOULD generate a PathErr/NOTIFICATION message (see Sections 2.2 and 2.3, respectively).

o) Number of Virtual Components (NVC): 16 bits

This field indicates the number of signals that are requested to be virtually concatenated. These signals are all of the same type by definition. They are Elementary Signal SPEs/VCs for which signal types are defined in this document; i.e., VT1.5_SPE/VC-11,

VT2_SPE/VC-12, VT3_SPE, VT6_SPE/VC-2, STS-1_SPE/VC-3, or STS-3c_SPE/VC-4.

This field is set to 0 (default value) to indicate that no virtual concatenation is requested.

o) Multiplier (MT): 16 bits

This field indicates the number of identical signals that are requested for the LSP; i.e., that form the Final Signal. These signals can be identical Elementary Signals, identical contiguously concatenated signals, or identical virtually concatenated signals. Note that all of these signals thus belong to the same LSP.

The distinction between the components of multiple virtually concatenated signals is done via the order of the labels that are specified in the signaling. The first set of labels must describe the first component (set of individual signals belonging to the first virtual concatenated signal), the second set must describe the second component (set of individual signals belonging to the second virtual concatenated signal), and so on.

This field is set to one (default value) to indicate that exactly one instance of a signal is being requested. Intermediate and egress nodes MUST verify that the node itself and the interfaces on which the LSP will be established can support the requested multiplier value. If the requested values cannot be supported, the receiver node MUST generate a PathErr/NOTIFICATION message (see Sections 2.2 and 2.3, respectively).

Zero is an invalid value. If a zero is received, the node MUST generate a PathErr/NOTIFICATION message (see Sections 2.2 and 2.3, respectively).

Note 1: When a transparent STS-N/STM-N signal is requested that is limited to a single contiguously concatenated STS-Nc-SPE/VC-4-Nc, the multiplier field MUST be equal to 1 (only valid value).

o) Transparency (T): 32 bits

This field is a vector of flags that indicates the type of transparency being requested. Several flags can be combined to provide different types of transparency. Not all combinations are necessarily valid. The default value for this field is zero, i.e., no transparency is requested.

Transparency, as defined from the point of view of this signaling specification, is only applicable to the fields in the SONET/SDH frame overheads. In the SONET case, these are the fields in the Section Overhead (SOH) and the Line Overhead (LOH). In the SDH case, these are the fields in the Regenerator Section Overhead (RSOH), the Multiplex Section overhead (MSOH), and the pointer fields between the two. With SONET, the pointer fields are part of the LOH.

Note also that transparency is only applicable when the following signal types are used: STS-1/STM-0, STS-3/STM-1, STS-12/STM-4, STS-48/STM-16, STS-192/STM-64, and STS-768/STM-256. At least one transparency type must be specified when such a signal type is requested.

Transparency indicates precisely which fields in these overheads must be delivered unmodified at the other end of the LSP. An ingress Label Switching Router (LSR) requesting transparency will pass these overhead fields that must be delivered to the egress LSR without any change. From the ingress and egress LSRs point of views, these fields must be seen as being unmodified.

Transparency is applied not at the interfaces with the initiating and terminating LSRs but only between intermediate LSRs. The transparency field is used to request an LSP that supports the requested transparency type; it may also be used to set up the transparency process to be applied at each intermediate LSR.

The different transparency flags are as follows:

    Flag 1 (bit 1): Section/Regenerator Section layer
    Flag 2 (bit 2): Line/Multiplex Section layer

where bit 1 is the low-order bit. Other flags are reserved; they should be set to zero when sent and ignored when received. A flag is set to one to indicate that the corresponding transparency is requested.

Intermediate and egress nodes MUST verify that the node itself and the interfaces on which the LSP will be established can support the requested transparency. If the requested flags cannot be supported, the receiver node MUST generate a PathErr/NOTIFICATION message (see Sections 2.2 and 2.3, respectively).

Section/Regenerator Section layer transparency means that the entire frames must be delivered unmodified. This implies that pointers cannot be adjusted. When Section/Regenerator Section layer transparency is used all other flags MUST be ignored.

Line/Multiplex Section layer transparency means that the LOH/MSOH must be delivered unmodified. This implies that pointers cannot be adjusted.

o) Profile (P): 32 bits

This field is intended to indicate particular capabilities that must be supported for the LSP; for example, monitoring capabilities.

No standard profile is currently defined, and this field SHOULD be set to zero when transmitted and ignored when received.

In the future, TLV-based extensions may be created.

2.2. RSVP-TE Details

For RSVP-TE, the SONET/SDH traffic parameters are carried in the SONET/SDH SENDER_TSPEC and FLOWSPEC objects. The same format is used both for the SENDER_TSPEC object and for FLOWSPEC objects. The content of the objects is defined above, in Section 2.1. The objects have the following class and type for SONET ANSI T1.105 and SDH ITU-T G.707:

    SONET/SDH SENDER_TSPEC object: Class = 12, C-Type = 4
    SONET/SDH FLOWSPEC object: Class = 9, C-Type = 4

There is no Adspec associated with the SONET/SDH SENDER_TSPEC. Either the Adspec is omitted, or an int-serv Adspec with the Default General Characterization Parameters and Guaranteed Service fragment is used; see [RFC2210].

For a particular sender in a session, the contents of the FLOWSPEC object received in a Resv message SHOULD be identical to the contents of the SENDER_TSPEC object received in the corresponding Path message. If the objects do not match, a ResvErr message with a "Traffic Control Error/Bad Flowspec value" error SHOULD be generated.

Intermediate and egress nodes MUST verify that the node itself and the interfaces on which the LSP will be established can support the requested Signal Type, RCC, NCC, NVC and Multiplier (as defined in Section 2.1). If the requested value(s) can not be supported, the receiver node MUST generate a PathErr message with a "Traffic Control Error/ Service unsupported" indication (see [RFC2205]).

In addition, if the MT field is received with a zero value, the node MUST generate a PathErr message with a "Traffic Control Error/Bad Tspec value" indication (see [RFC2205]).

Intermediate nodes MUST also verify that the node itself and the interfaces on which the LSP will be established can support the requested Transparency (as defined in Section 2.1). If the requested value(s) cannot be supported, the receiver node MUST generate a PathErr message with a "Traffic Control Error/Service unsupported" indication (see [RFC2205]).

2.3. CR-LDP Details

For CR-LDP, the SONET/SDH traffic parameters are carried in the SONET/SDH Traffic Parameters TLV. The content of the TLV is defined above, in Section 2.1. The header of the TLV 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|U|F|          Type             |      Length                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The type field for the SONET/SDH Traffic Parameters TLV is 0x0838.

Intermediate and egress nodes MUST verify that the node itself and the interfaces on which the LSP will be established can support the requested Signal Type, RCC, NCC, NVC, and Multiplier (as defined in Section 2.1). If the requested value(s) cannot be supported, the receiver node MUST generate a NOTIFICATION message with a "Resource Unavailable" status code (see [RFC3212]).

In addition, if the MT field is received with a zero value, the node MUST generate a NOTIFICATION message with a "Resource Unavailable" status code (see [RFC3212]).

Intermediate nodes MUST also verify that the node itself and the interfaces on which the LSP will be established can support the requested Transparency (as defined in Section 2.1). If the requested value(s) cannot be supported, the receiver node MUST generate a NOTIFICATION message with a "Resource Unavailable" status code (see [RFC3212]).

3. SONET and SDH Labels

SONET and SDH each define a multiplexing structure. Both structures are trees whose roots are, respectively, an STS-N or an STM-N and whose leaves are the signals that can be transported via the time-slots and switched between time-slots within an ingress port and time-slots within an egress port; i.e., a VTx SPE, an STS-x SPE, or a VC-x. A SONET/SDH label will identify the exact position (i.e., first time-slot) of a particular VTx SPE, STS-x SPE, or VC-x signal in a multiplexing structure. SONET and SDH labels are carried in the Generalized Label per [RFC3473] and [RFC3472].

Note that by time-slots we mean the time-slots as they appear logically and sequentially in the multiplex, not as they appear after any possible interleaving.

These multiplexing structures will be used as naming trees to create unique multiplex entry names or labels. The same format of label is used for SONET and SDH. As explained in [RFC3471], a label does not identify the "class" to which the label belongs. This is implicitly determined by the link on which the label is used.

In case of signal concatenation or multiplication, a list of labels can appear in the Label field of a Generalized Label.

In case of contiguous concatenation, only one label appears in the Label field. This unique label is encoded as a single 32-bit label value (as defined in this section) of the Generalized Label object (Class-Num = 16, C-Type = 2)/TLV (0x0825). This label identifies the lowest time-slot occupied by the contiguously concatenated signal. By lowest time-slot, we mean the one having the lowest label (value) when compared as an integer value; i.e., the time-slot occupied by the first component signal of the concatenated signal encountered descending the tree.

In case of virtual concatenation, the explicit ordered list of all labels in the concatenation is given. This ordered list of labels is encoded as a sequence of 32-bit label values (as defined in this section) of the Generalized Label object (Class-Num = 16, C-Type = 2)/TLV (0x0825). Each label indicates the first time-slot occupied by a component of the virtually concatenated signal. The order of the labels must reflect the order of the payloads to concatenate (not the physical order of time-slots). The above representation limits virtual concatenation to remain within a single (component) link; it imposes, as such, a restriction compared to the ANSI [T1.105]/ ITU-T [G.707] recommendations. The standard definition for virtual concatenation allows each virtual concatenation components to travel over diverse paths. Within GMPLS, virtual concatenation components must travel over the same (component) link if they are part of the same LSP. This is due to the way that labels are bound to a (component) link. Note, however, that the routing of components on different paths is indeed equivalent to establishing different LSPs, each one having its own route. Several LSPs can be initiated and terminated between the same nodes, and their corresponding components can then be associated together (i.e., virtually concatenated).

In case of multiplication (i.e., using the multiplier transform), the explicit ordered list of all labels that take part in the Final Signal is given. This ordered list of labels is encoded as a sequence of 32-bit label values (as defined in this section) of the Generalized Label object (Class-Num = 16, C-Type = 2)/TLV (0x0825). In case of multiplication of virtually concatenated signals, the explicit ordered list of the set of labels that take part in the Final Signal is given. The first set of labels indicates the time-slots occupied by the first virtually concatenated signal, the second set of labels indicates the time-slots occupied by the second virtually concatenated signal, and so on. The above representation limits multiplication to remain within a single (component) link.

The format of the label for SONET and/or SDH TDM-LSR link 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               S               |   U   |   K   |   L   |   M   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

This is an extension of the numbering scheme defined in [G.707], Sections 7.3.7 through 7.3.13; i.e., the (K, L, M) numbering. Note that the higher order numbering scheme defined in [G.707], Sections 7.3.1 through 7.3.6, is not used here.

Each letter indicates a possible branch number starting at the parent node in the multiplex structure. Branches are considered as being numbered in increasing order, starting from the top of the multiplexing structure. The numbering starts at 1; zero is used to indicate a non-significant or ignored field.

When a field is not significant or ignored in a particular context, it MUST be set to zero when transmitted and ignored when received.

When a hierarchy of SONET/SDH LSPs is used, a higher-order LSP with a given bandwidth can be used to carry lower-order LSPs. Remember that a higher-order LSP is established through a SONET/SDH higher-order path layer network, and a lower-order LSP through a SONET/SDH lower-order path layer network (see also ITU-T G.803, Section 3, for the corresponding definitions). In this context, the higher-order SONET/SDH LSP behaves as a "virtual link" with a given bandwidth (e.g., VC-3); it may also be used as a Forwarding Adjacency. A lower-order SONET/SDH LSP can be established through that higher-order LSP. Since a label is local to a (virtual) link, the highest part of that label (i.e., the S, U, and K fields) is non-significant and is set to zero; i.e., the label is "0,0,0,L,M". Similarly, if the structure of the lower-order LSP is unknown or not relevant, the lowest part of that label (i.e., the L and M fields) is non-significant and is set to zero; i.e., the label is "S,U,K,0,0".

For instance, a VC-3 LSP can be used to carry lower-order LSPs. In that case, the labels allocated between the two ends of the VC-3 LSP for the lower-order LSPs will have S, U, and K set to zero (i.e., non-significant) while L and M will be used to indicate the signal allocated in that VC-3.

In case of tunneling, such as VC-4 containing VC-3 containing VC-12/VC-11, where the SUKLM structure is not adequate to represent the full signal structure, a hierarchical approach must be used; i.e., per layer network signaling.

The possible values of S, U, K, L, and M are defined as follows:

1. S=1->N is the index of a particular STS-3/AUG-1 inside an STS-N/STM-N multiplex. S is only significant for SONET STS-N (N>1) and SDH STM-N (N>0). S must be 0 and ignored for STS-1 and STM-0.

2. U=1->3 is the index of a particular STS-1_SPE/VC-3 within an STS-3/AUG-1. U is only significant for SONET STS-N (N>1) and SDH STM-N (N>0). U must be 0 and ignored for STS-1 and STM-0.

3. K=1->3 is the index of a particular TUG-3 within a VC-4. K is only significant for an SDH VC-4 structured in TUG-3s. K must be 0 and ignored in all other cases.

4. L=1->7 is the index of a particular VT_Group/TUG-2 within an STS-1_SPE/TUG-3 or VC-3. L must be 0 and ignored in all other cases.

5. M is the index of a particular VT1.5_SPE/VC-11, VT2_SPE/VC-12, or VT3_SPE within a VT_Group/TUG-2. M=1->2 indicates a specific VT3 SPE inside the corresponding VT Group; these values MUST NOT be used for SDH, since there is no equivalent of VT3 with SDH. M=3->5 indicates a specific VT2_SPE/VC-12 inside the corresponding VT_Group/TUG-2. M=6->9 indicates a specific VT1.5_SPE/VC-11 inside the corresponding VT_Group/TUG-2.

Note that a label always has to be interpreted according the SONET/SDH traffic parameters; i.e., a label by itself does not allow knowing which signal is being requested (a label is context sensitive).

The label format defined in this section, referred to as SUKLM, MUST be used for any SONET/SDH signal requests that are not transparent; i.e., when all Transparency (T) bits defined in Section 2.1 are set to zero. Any transparent STS-1/STM-0/STS-3*N/STM-N (N=1, 4, 16, 64, 256) signal request MUST use a label format as defined in [RFC3471].

The S encoding is summarized in the following table:

S	SDH	SONET
0	other	other
1	1st AUG-1	1st STS-3
2	2nd AUG-1	2nd STS-3
3	3rd AUG-1	3rd STS-3
4	4rd AUG-1	4rd STS-3
:	:	:
N	Nth AUG-1	Nth STS-3

The U encoding is summarized in the following table:

U	SDH AUG-1	SONET STS-3
0	other	other
1	1st VC-3	1st STS-1 SPE
2	2nd VC-3	2nd STS-1 SPE
3	3rd VC-3	3rd STS-1 SPE

The K encoding is summarized in the following table:

K	SDH VC-4
0	other
1	1st TUG-3
2	2nd TUG-3
3	3rd TUG-3

The L encoding is summarized in the following table:

L	SDH TUG-3	SDH VC-3	SONET STS-1 SPE
0	other	other	other
1	1st TUG-2	1st TUG-2	1st VTG
2	2nd TUG-2	2nd TUG-2	2nd VTG
3	3rd TUG-2	3rd TUG-2	3rd VTG
4	4th TUG-2	4th TUG-2	4th VTG
5	5th TUG-2	5th TUG-2	5th VTG
6	6th TUG-2	6th TUG-2	6th VTG
7	7th TUG-2	7th TUG-2	7th VTG

The M encoding is summarized in the following table:

M	SDH TUG-2	SONET VTG
0	other	other
1	-	1st VT3 SPE
2	-	2nd VT3 SPE
3	1st VC-12	1st VT2 SPE
4	2nd VC-12	2nd VT2 SPE
5	3rd VC-12	3rd VT2 SPE
6	1st VC-11	1st VT1.5 SPE
7	2nd VC-11	2nd VT1.5 SPE
8	3rd VC-11	3rd VT1.5 SPE
9	4th VC-11	4th VT1.5 SPE

Examples of Labels

Example 1: the label for the STS-3c_SPE/VC-4 in the Sth STS-3/AUG-1 is: S>0, U=0, K=0, L=0, M=0.

Example 2: the label for the VC-3 within the Kth-1 TUG-3 within the VC-4 in the Sth AUG-1 is: S>0, U=0, K>0, L=0, M=0.

Example 3: the label for the Uth-1 STS-1_SPE/VC-3 within the Sth STS-3/AUG-1 is: S>0, U>0, K=0, L=0, M=0.

Example 4: the label for the VT6/VC-2 in the Lth-1 VT Group/TUG-2 in the Uth-1 STS-1_SPE/VC-3 within the Sth STS-3/AUG-1 is: S>0, U>0, K=0, L>0, M=0.

Example 5: the label for the 3rd VT1.5_SPE/VC-11 in the Lth-1 VT Group/TUG-2 within the Uth-1 STS-1_SPE/VC-3 within the Sth STS-3/AUG-1 is: S>0, U>0, K=0, L>0, M=8.

Example 6: the label for the STS-12c SPE/VC-4-4c which uses the 9th STS-3/AUG-1 as its first timeslot is: S=9, U=0, K=0, L=0, M=0.

In case of contiguous concatenation, the label that is used is the lowest label (value) of the contiguously concatenated signal, as explained before. The higher part of the label indicates where the signal starts, and the lowest part is not significant.

In case of STM-0/STS-1, the values of S, U, and K must be equal to zero, according to the field coding rules. For instance, when a VC-3 in an STM-0 is requested, the label is S=0, U=0, K=0, L=0, M=0. When a VC-11 in a VC-3 in an STM-0 is requested, the label is S=0, U=0, K=0, L>0, M=6..9.

Note: when a Section/RS or Line/MS transparent STS-1/STM-0/ STS-3*N/STM-N (N=1, 4, 16, 64, 256) signal is requested, the SUKLM label format and encoding is not applicable, and the label encoding MUST follow the rules defined in [RFC3471], Section 3.2.

4. Acknowledgements

Valuable comments and input were received from the CCAMP mailing list, where outstanding discussions took place.

The authors would like to thank Richard Rabbat for his valuable input, which lead to this revision.

5. Security Considerations

This document introduces no new security considerations to either [RFC3473] or [RFC3472]. GMPLS security is described in Section 11 of [RFC3471] and refers to [RFC3209] for RSVP-TE and to [RFC3212] for CR-LDP.

6. IANA Considerations

Three values defined by IANA for RFC 3946 now apply to this document.

Two RSVP C-Types in registry:
http://www.iana.org/assignments/rsvp-parameters

- A SONET/SDH SENDER_TSPEC object: Class = 12, C-Type = 4 (see Section 2.2).

- A SONET/SDH FLOWSPEC object: Class = 9, C-Type = 4 (see Section 2.2).

One LDP TLV Type in registry:
http://www.iana.org/assignments/ldp-namespaces

- A type field for the SONET/SDH Traffic Parameters TLV (see Section 2.3).

Contributors

Contributors are listed in alphabetical order:

Stefan Ansorge (Alcatel)
Lorenzstrasse 10
70435 Stuttgart, Germany
EMail: stefan.ansorge@alcatel.de

Peter Ashwood-Smith (Nortel)
PO. Box 3511 Station C,
Ottawa, ON K1Y 4H7, Canada
EMail:petera@nortelnetworks.com

Ayan Banerjee (Calient)
5853 Rue Ferrari
San Jose, CA 95138, USA
EMail: abanerjee@calient.net

Lou Berger (Movaz)
7926 Jones Branch Drive
McLean, VA 22102, USA
EMail: lberger@movaz.com

Greg Bernstein (Ciena)
10480 Ridgeview Court
Cupertino, CA 94014, USA
EMail: greg@ciena.com

Angela Chiu (Celion)
One Sheila Drive, Suite 2
Tinton Falls, NJ 07724-2658
EMail: angela.chiu@celion.com

John Drake (Calient)
5853 Rue Ferrari
San Jose, CA 95138, USA
EMail: jdrake@calient.net

Yanhe Fan (Axiowave)
100 Nickerson Road
Marlborough, MA 01752, USA
EMail: yfan@axiowave.com

Michele Fontana (Alcatel)
Via Trento 30,
I-20059 Vimercate, Italy
EMail: michele.fontana@alcatel.it

Gert Grammel (Alcatel)
Lorenzstrasse, 10
70435 Stuttgart, Germany
EMail: gert.grammel@alcatel.de

Juergen Heiles (Siemens)
Hofmannstr. 51
D-81379 Munich, Germany
EMail: juergen.heiles@siemens.com

Suresh Katukam (Cisco)
1450 N. McDowell Blvd,
Petaluma, CA 94954-6515, USA
EMail: suresh.katukam@cisco.com

Kireeti Kompella (Juniper)
1194 N. Mathilda Ave.
Sunnyvale, CA 94089, USA
EMail: kireeti@juniper.net

Jonathan P. Lang (Calient)
25 Castilian
Goleta, CA 93117, USA
EMail: jplang@calient.net

Fong Liaw (Solas Research)
EMail: fongliaw@yahoo.com

Zhi-Wei Lin (Lucent)
101 Crawfords Corner Rd
Holmdel, NJ 07733-3030, USA
EMail: zwlin@lucent.com

Ben Mack-Crane (Tellabs)
EMail: ben.mack-crane@tellabs.com

Dimitrios Pendarakis (Tellium)
2 Crescent Place, P.O. Box 901
Oceanport, NJ 07757-0901, USA
EMail: dpendarakis@tellium.com

Mike Raftelis (White Rock)
18111 Preston Road
Dallas, TX 75252, USA

Bala Rajagopalan (Tellium)
2 Crescent Place, P.O. Box 901
Oceanport, NJ 07757-0901, USA
EMail: braja@tellium.com

Yakov Rekhter (Juniper)
1194 N. Mathilda Ave.
Sunnyvale, CA 94089, USA
EMail: yakov@juniper.net

Debanjan Saha (Tellium)
2 Crescent Place, P.O. Box 901
Oceanport, NJ 07757-0901, USA
EMail: dsaha@tellium.com

Vishal Sharma (Metanoia)
335 Elan Village Lane
San Jose, CA 95134, USA
EMail: vsharma87@yahoo.com

George Swallow (Cisco)
250 Apollo Drive
Chelmsford, MA 01824, USA
EMail: swallow@cisco.com

Z. Bo Tang (Tellium)
2 Crescent Place, P.O. Box 901
Oceanport, NJ 07757-0901, USA
EMail: btang@tellium.com

Eve Varma (Lucent)
101 Crawfords Corner Rd
Holmdel, NJ 07733-3030, USA
EMail: evarma@lucent.com

Yangguang Xu (Lucent)
21-2A41, 1600 Osgood Street
North Andover, MA 01845, USA
EMail: xuyg@lucent.com

Appendix 1. Signal Type Values Extension for VC-3

This appendix defines the following optional additional Signal Type value for the Signal Type field of Section 2.1:

Value	Type
20	"VC-3 via AU-3 at the end"

According to the ITU-T [G.707] recommendation, a VC-3 in the TU-3/TUG-3/VC-4/AU-4 branch of the SDH multiplex cannot be structured in TUG-2s; however, a VC-3 in the AU-3 branch can be. In addition, a VC-3 could be switched between the two branches, if required.

A VC-3 circuit could be terminated on an ingress interface of an LSR (e.g., forming a VC-3 forwarding adjacency). This LSR could then want to demultiplex this VC-3 and switch internal low-order LSPs. For implementation reasons, this could be only possible if the LSR receives the VC-3 in the AU-3 branch. For example, for an LSR not able to switch internally from a TU-3 branch to an AU-3 branch on its incoming interface before demultiplexing and then switching the content with its switch fabric.

In that case, it is useful to indicate that the VC-3 LSP must be terminated at the end in the AU-3 branch instead of the TU-3 branch.

This is achieved by using the "VC-3 via AU-3 at the end" signal type. This information can be used, for instance, by the penultimate LSR to switch an incoming VC-3 received in any branch to the AU-3 branch on the outgoing interface to the destination LSR.

The "VC-3 via AU-3 at the end" signal type does not imply that the VC-3 must be switched via the AU-3 branch at some other places in the network. The VC-3 signal type just indicates that a VC-3 in any branch is suitable.

Annex 1. Examples

This annex defines examples of SONET and SDH signal coding. The objective is to help the reader to understand how the traffic parameter coding works and not to give examples of typical SONET or SDH signals.

As stated above, signal types are Elementary Signals to which successive concatenation, multiplication, and transparency transforms can be applied to obtain Final Signals.

1. A VC-4 signal is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with value 0 to a VC-4 Elementary Signal.

2. A VC-4-7v signal is formed by the application of RCC with value 0, NCC with value 0, NVC with value 7 (virtual concatenation of 7 components), MT with value 1, and T with value 0 to a VC-4 Elementary Signal.

3. A VC-4-16c signal is formed by the application of RCC with value 1 (standard contiguous concatenation), NCC with value 16, NVC with value 0, MT with value 1, and T with value 0 to a VC-4 Elementary Signal.

4. An STM-16 signal with Multiplex Section layer transparency is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with flag 2 to an STM-16 Elementary Signal.

5. An STM-4 signal with Multiplex Section layer transparency is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with flag 2 applied to an STM-4 Elementary Signal.

6. An STM-256 signal with Multiplex Section layer transparency is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with flag 2 applied to an STM-256 Elementary Signal.

7. An STS-1 SPE signal is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with value 0 to an STS-1 SPE Elementary Signal.

8. An STS-3c SPE signal is formed by the application of RCC with value 1 (standard contiguous concatenation), NCC with value 1, NVC with value 0, MT with value 1, and T with value 0 to an STS-3c SPE Elementary Signal.

9. An STS-48c SPE signal is formed by the application of RCC with value 1 (standard contiguous concatenation), NCC with value 16, NVC with value 0, MT with value 1, and T with value 0 to an STS-3c SPE Elementary Signal.

10. An STS-1-3v SPE signal is formed by the application of RCC with value 0, NVC with value 3 (virtual concatenation of 3 components), MT with value 1, and T with value 0 to an STS-1 SPE Elementary Signal.

11. An STS-3c-9v SPE signal is formed by the application of RCC with value 1, NCC with value 1, NVC with value 9 (virtual concatenation of 9 STS-3c), MT with value 1, and T with value 0 to an STS-3c SPE Elementary Signal.

12. An STS-12 signal with Section layer (full) transparency is formed by the application of RCC with value 0, NCC with value 0, NVC with value 0, MT with value 1, and T with flag 1 to an STS-12 Elementary Signal.

13. A 3 x STS-768c SPE signal is formed by the application of RCC with value 1, NCC with value 256, NVC with value 0, MT with value 3, and T with value 0 to an STS-3c SPE Elementary Signal.

14. A 5 x VC-4-13v composed signal is formed by the application of RCC with value 0, NVC with value 13, MT with value 5, and T with value 0 to a VC-4 Elementary Signal.

The encoding of these examples is summarized in the following table:

Signal	ST	RCC	NCC	NVC	MT	T
VC-4	6	0	0	0	1	0
VC-4-7v	6	0	0	7	1	0
VC-4-16c	6	1	16	0	1	0
STM-16 MS transparent	10	0	0	0	1	2
STM-4 MS transparent	9	0	0	0	1	2
STM-256 MS transparent	12	0	0	0	1	2
STS-1 SPE	5	0	0	0	1	0
STS-3c SPE	6	1	1	0	1	0
STS-48c SPE	6	1	16	0	1	0
STS-1-3v SPE	5	0	0	3	1	0
STS-3c-9v SPE	6	1	1	9	1	0
STS-12 Section transparent	9	0	0	0	1	1
3 x STS-768c SPE	6	1	256	0	3	0
5 x VC-4-13v	6	0	0	13	5	0

Normative References

[G.707] ITU-T Recommendation G.707, "Network Node Interface for the Synchronous Digital Hierarchy", October 2000.

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

[RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[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.

[RFC3212] Jamoussi, B., Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heinanen, J., Kilty, T., and A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, January 2002.

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3472] Ashwood-Smith, P. and L. Berger, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Constraint-based Routed Label Distribution Protocol (CR-LDP) Extensions", RFC 3472, January 2003.

[RFC3473] Berger, L., "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.

[T1.105] "Synchronous Optical Network (SONET): Basic Description Including Multiplex Structure, Rates, and Formats", ANSI T1.105, October 2000.

Authors' Addresses

Eric Mannie
Perceval
Rue Tenbosch, 9
1000 Brussels
Belgium

Phone: +32-2-6409194
EMail: eric.mannie@perceval.net

Dimitri Papadimitriou
Alcatel
Copernicuslaan 50
B-2018 Antwerpen, Belgium

Phone: +32 3 240-8491
EMail: dimitri.papadimitriou@alcatel.be

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).