

CCAMP Working Group
Internet Draft
Updates: RFC [3946](#)
Category: Standards Track
Expires: February 2007

G. Bernstein
Grotto Networking
D. Caviglia
Ericsson
R. Rabbat (ed.)
Fujitsu
H. van Helvoort
Huawei
August 4, 2006

**Operating Virtual Concatenation (VCAT) and the Link Capacity
Adjustment Scheme (LCAS) with Generalized Multi-Protocol Label
Switching (GMPLS)**
[draft-bernstein-ccamp-gmpls-vcat-lcas-04.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

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

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

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

This Internet-Draft will expire on February 4, 2007.

Abstract

This document describes requirements for, and use of, the Generalized Multi-Protocol Label Switching (GMPLS) control plane in conjunction

with the Virtual Concatenation (VCAT) layer 1 inverse multiplexing mechanism and its companion Link Capacity Adjustment Scheme (LCAS) which can be used for hitless dynamic resizing of the inverse multiplex group. These techniques apply to the Optical Transport Network (OTN), Synchronous Optical Network (SONET), Synchronous Digital Hierarchy (SDH), and Plesiochronous Digital Hierarchy (PDH) signals.

Conventions used in this document

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

Table of Contents

1. Introduction.....	3
2. Changes from draft-bernstein-ccamp-gmpls-vcat-lcas-03.....	3
3. VCAT/LCAS Scenarios and Specific Requirements.....	3
3.1. Multiple VCAT Groups per GMPLS Endpoint.....	3
3.2. Component Signal Configuration Requirements.....	4
3.3. VCAT Operation With or Without LCAS.....	4
4. GMPLS Mechanisms for Signaling VCAT/LCAS.....	4
4.1. Co-Routed Signals.....	5
4.1.1. One-shot Setup of Co-Routed Signal.....	5
4.1.2. Incremental Setup of Co-Routed Signal.....	5
4.1.3. Removing a Component Signal.....	6
4.1.4. Removing Multiple Component Signals in One Shot.....	6
4.1.5. Use of multiple LSPs for Co-Routed Signals.....	7
4.1.6. Teardown of Whole VCG.....	7
4.2. Diversely Routed Signals.....	7
4.2.1. Associating Diversely Routed Signals.....	7
4.2.2. Procedures for VCG Setup Using Diversely Routed Components.....	8
4.2.3. Procedures for VCG Reduction/Teardown Using Diversely Routed Components.....	9
4.2.4. Update of Already Established LSPs.....	9
4.2.5. One LSP per Circuit.....	9
5. IANA Considerations.....	10
6. Security Considerations.....	10
7. Contributors.....	11
8. Acknowledgments.....	11
9. References.....	12
9.1. Normative References.....	12
9.2. Informative References.....	12
Author's Addresses.....	13

Intellectual Property Statement.....	14
Disclaimer of Validity.....	14
Copyright Statement.....	14
Acknowledgment.....	15

[1. Introduction](#)

The Generalized Multi-Protocol Label Switching (GMPLS) suite of protocols allows the automated control of different switching technologies including SONET/SDH.

This document describes extensions to RSVP-TE to support the Virtual Concatenation (VCAT) layer 1 inverse multiplexing mechanism and its companion Link Capacity Adjustment Scheme (LCAS). These extensions enable the setup of diversely routed circuits that are members of the same VCAT group.

[2. Changes from \[draft-bernstein-ccamp-gmpls-vcat-lcas-03\]\(#\)](#)

- o Dropped appendices
- o Updated references
- o Fixed error w/r to sessions. The LSPs belong to different sessions
- o Many editorial changes throughout the document
- o Rewrote the introduction

[3. VCAT/LCAS Scenarios and Specific Requirements](#)

There are a number of specific requirements for the support of VCAT/LCAS in GMPLS that can be derived from the carriers' application-specific demands for the use of VCAT/LCAS and from the flexible nature of VCAT/LCAS. These are set out in the following section.

[3.1. Multiple VCAT Groups per GMPLS Endpoint](#)

In general, an LSR can be ingress/egress of one or more VCAT groups. VCAT and LCAS are interface capabilities. An LSR may have, for example, VCAT-capable interfaces that are not LCAS-capable. It may at the same time have interfaces that are neither VCAT nor LCAS-capable.

3.2. Component Signal Configuration Requirements

We list in this section the different scenarios that SHOULD be supported. Here we use the term "VCG" to refer to the entire VCAT group and the terminology "set" and "subset" to refer to the collection of potential VCAT group member signals.

- o Fixed, co-routed: A fixed bandwidth VCG, transported over a co-routed set of member signals. This is the case where the intended bandwidth of the VCG does not change and all member signals follow the same route and minimize differential delay. The intent here is the capability to allocate an amount of bandwidth close to that required at the client layer.
- o Fixed, diversely routed: A fixed bandwidth VCG, transported over at least two diversely routed subsets of member signals. In this case, the subsets are link-disjoint over at least one link of the route. The intent here is more efficient use of network resources (no unique route has the required bandwidth), and additional resilience and graceful degradation in the case of failure (note that differential delay may be a limiting factor).
- o Dynamic, co-routed: A dynamic VCG (bandwidth can be increased or decreased via the addition or removal of member signals), transported over a co-routed set of members. Intent here is dynamic sizing of bandwidth.
- o Dynamic, diversely routed: A dynamic VCAT group, transported over at least two diversely routed subsets of member signals. The intent here is dynamic resizing and resilience (but differential delay may be a limiting factor).

3.3. VCAT Operation With or Without LCAS

VCAT capabilities may be present with or without the presence of LCAS. The use of LCAS is beneficial to the provision of services, but in the absence of LCAS, VCAT is still a valid technique. Therefore GMPLS mechanisms for the operation of VCAT are REQUIRED for both the case where LCAS is available and the case where it is not available. The GMPLS procedures for the two cases SHOULD be identical.

4. GMPLS Mechanisms for Signaling VCAT/LCAS

We describe in this section the signaling mechanisms that already exist in GMPLS using RSVP-TE [[RFC3473](#)] and the extensions needed, for diversely routed paths and in support of the LCAS procedure.

[Section 4.1](#) is included for informational purposes only. It describes existing procedures and makes no changes.

[Section 4.2](#) describes new procedures to support diversely routed VCAT groups. Where possible it reuses applicable existing procedures from [section 4.1](#).

[4.1. Co-Routed Signals](#)

Note that this section is for informational purposes only.

The existing signaling protocols support co-routed signal setup using the NVC field as explained in section 2.1 of [[RFC3946bis](#)]. In this case, one single LSP is set up in support of the VCAT group.

There are two options for setting up the VCAT group, depending on hardware capability, or management preferences: one-shot setup and incremental setup.

The following sections explain the procedure based on an example of setting up a VC4-7v SDH VCAT group (corresponding to an STS-3c-7v SONET VCAT group).

[4.1.1. One-shot Setup of Co-Routed Signal](#)

An RSVP-TE Path message is used with the following parameters.

With regards to the traffic parameters, the elementary signal is chosen (6 for VC-4/STS-3c_SPE). The value of NVC is then set to 7.

A Multiplier Transform greater than 1 (say $N > 1$) is used if the operator wants to set up N VCAT groups that will belong to, and be assigned to, one LSP.

SDH or SONET labels in turn have to be assigned for each member of the VCG and concatenated to form a single Generalized Label constructed as an ordered list of 32-bit timeslot identifiers of the same format as TDM labels. [[RFC3946bis](#)] requires that the order of the labels reflect the order of the payloads to concatenate, and not the physical order of time-slots.

[4.1.2. Incremental Setup of Co-Routed Signal](#)

In some cases, it may be necessary or desirable to set up the VCG members individually, or to add group members to an existing group.

One example of this need is when the hardware that supports VCAT can only add VCAT elements one at a time or cannot automatically match the elements at the ingress and egress for the purposes of inverse multiplexing. Serial or incremental setup solves this problem.

In order to accomplish incremental setup an iterative process is used to add group members. For each iteration, NVC is incremented up to the final value required. The iteration consists of the successful completion of Path and Resv signaling. At first, NVC = 1 and the label includes just one timeslot identifier

At each of the next iterations, NVC is set to (NVC +1), one more timeslot identifier is added to the ordered list in the Generalized Label (in the Path or Resv message). A node that receives a Path message that contains changed fields will process the full Path message and, based on the new value of NVC, it will add a component signal to the VCAT group, and switch the new timeslot based on the new label information.

Following the addition of the new label to the LSP, LCAS may be used in-band to add the new label into the existing VCAT group. LCAS signaling for this function is described in [[ITU-T-G.7042](#)].

4.1.3. Removing a Component Signal

The procedure to remove a component signal is similar to that used to add components as described in [Section 4.1.2](#). The LCAS in-band signaling step is taken first to take the component out of the group. LCAS signaling is described in [[ITU-T-G.7042](#)].

In this case, the NVC value is decremented by 1 and the timeslot identifier for the dropped component is removed from the ordered list in the Generalized Label. This function is not supported without management intervention for VCAT-only interfaces as removing one component of the VCG will result in errors in the inverse-multiplexing procedure of VCAT and result in the teardown of the whole group. So, this is a feature that only LCAS-capable VCAT interfaces can support without management intervention at the end points.

4.1.4. Removing Multiple Component Signals in One Shot

The procedure is similar to 4.1.3. In this case, the NVC value is changed to the new value and all relevant timeslot identifiers for the components to be torn down are removed from the ordered list in the Generalized Label. This is also not supported for VCAT-only interfaces without management intervention as removing one component

of the VCG will tear down the whole group, but the use of LCAS can facilitate this procedure.

4.1.5. Use of multiple LSPs for Co-Routed Signals

Co-routed signals may also be supported by distinct LSPs signaled separately using exactly the techniques described for diversely routed signals in [Section 4.2](#).

4.1.6. Teardown of Whole VCG

The entire LSP is deleted in a single step (i.e., all components are removed in one go) using deletion procedures of [\[RFC3473\]](#).

4.2. Diversely Routed Signals

The initial GMPLS specification did not support diversely routed signals using the NVC construct. In fact, [\[RFC3946bis\]](#) says:

[...] 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).

Diverse routing of signals can be a useful capability but requires the extensions identified in this document.

4.2.1. Associating Diversely Routed Signals

The feature that needs to be added is the functionality to associate the components of the same VCG. For this purpose, we use the Association Object that was defined in [\[E2E-RECOVERY\]](#) to associate working and recovery LSPs.

A diversely routed VCG uses a number of routes $R \leq \text{VCG size}$, as some routes may be the same for several components. A number of LSPs, L ($\text{VCG} \geq L \geq R$) are used with each LSP establishing at least one component of the VCG, and at most all of the co-routed members of the group. For a set of c components using the same route, we set up the LSP with $\text{NVC} = c$ exactly as explained in [section 4.1.1](#). Therefore,

the association of group members or of sub-groups to form the VCG requires the association of the LSPs used to establish the group members.

To be able to distinguish the LSPs in the VCG each must have a unique identifier. LSPs are identified by the combination of Session and Sender Template fields. It is common practice when make-before-break [[RFC3209](#)] is supported to allow LSPs with the same Session, but different Sender Templates (specifically with different LSP IDs) to share resources. Since resource sharing between VCG members must not be allowed (because we want each LSP to contribute capacity to the VCG), but since we want to continue to support make-before-break for each group member, it is necessary to distinguish the LSPs in the VCG by varying the fields in the Session object. Specifically, a different Tunnel ID is used to identify each LSP in the VCG.

Thus, VCG members cannot be associated through the Session object, and the Association object is used instead.

The assignment of the Association ID is outside the scope of GMPLS but MUST be unique for each VCAT group.

Note that the use of the Association object to associate members of a VCG does not preclude the use of another instance of the object with a different Association ID to indicate the association of working and recovery LSPs, as [[E2E-RECOVERY](#)] allows the use of multiple Association objects. We differentiate between the Association objects used for the VCAT group and other Association objects through the definition of a new association type to indicate that this is a VCG association.

Association Type: 16 bits

Value	Type
-----	----
3	VCAT group

See [[E2E-RECOVERY](#)] for the definition of other fields and values of the Association object.

[4.2.2](#). Procedures for VCG Setup Using Diversely Routed Components

For every route R, use the procedure outlined in [section 4.1.1](#) or 4.1.2 depending on the capability of the equipment or local policy. The Path message MUST include the Association object with type set to 3, and each Path message MUST use the same Association ID.

Following the addition of each new LSP (i.e., once the RESV message has been received by the ingress LSR), LCAS signaling is used in-band to hitlessly add the new label into the existing group [ITU-T-G.7042].

4.2.3. Procedures for VCG Reduction/Teardown Using Diversely Routed Components

To remove the component circuits on any route, LCAS signaling is used in-band to remove the labels associated with the LSP from the group. LCAS signaling is defined in [ITU-T-G.7042].

In addition, the procedures outlined in [section 4.1.3](#) or 4.1.4 are used to tear down the unwanted LSP.

Again, this can only be done on LCAS-capable interfaces. If the procedure is attempted on VCAT-only interfaces, then the whole VCG is torn down (this is not a graceful teardown so ingress/egress initiate a Path Tear/Resv Tear) on all routes.

4.2.4. Update of Already Established LSPs

Established co-routed VCAT groups currently do not support the Association object. If a co-routed VCAT Group size is to be increased with new diversely routed members, we use the LSP modification procedure described in [RFC2205]. An Association object is added to the Path message for the existing LSP(s). That Association object can then be used to set up new diversely routed group members.

The same applies to SONET/SDH LSPs. An operator may decide to use an already cross-connected SONET/SDH LSP for diversely-routed VCAT. In this case the modification procedure described in [RFC2205] is used as well.

4.2.5. One LSP per Circuit

These procedures can support the use of as many LSPs as there are circuits in the VCG. This can be done when each circuit is separately routed, or when some of the circuits are co-routed, and each LSP will be used to set up one element of the VCG. The Association object is used to indicate the VCG association.

5. IANA Considerations

This document requests from IANA the assignment of a new Association Type within the Association object. This object was defined in [E2E-RECOVERY].

The value 3 "VCAT group" is suggested in [section 4.2.1](#).

6. Security Considerations

This document introduces a new use of the Association object for GMPLS signaling [[RFC3473](#)] to associate diversely routed VCAT group members. It does not introduce any new signaling messages, nor change the relationship between LSRs that are adjacent in the control plane. This association information in the event of an interception may indicate that there are members of the same VCAT group that take a different route and may indicate to an interceptor that the network may be more robust.

Otherwise, this document does not introduce any additional security considerations.

7. Contributors

Wataru Imajuku (NTT)
1-1 Hikari-no-oka Yokosuka Kanagawa 239-0847
Japan

Phone +81-46-859-4315
Email: imajuku.wataru@lab.ntt.co.jp

Julien Meuric
France Telecom
2, avenue Pierre Marzin
22307 Lannion Cedex
France

Phone: + 33 2 96 05 28 28
Email: julien.meuric@orange-ft.com

Lyndon Ong
Ciena
PO Box 308
Cupertino, CA 95015
United States of America

Phone: +1 408 705 2978
Email: lyong@ciena.com

8. Acknowledgments

The authors would like to thank Maarten Vissers and Adrian Farrel for extensive reviews and contributions to this draft.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1, Functional Specification", [RFC 2205](#), September 1997.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC3946bis] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", IETF draft, work in progress, December 2005.
- [E2E-RECOVERY] Lang, J.P., Rekhter, Y., and D. Papadimitriou (eds.), "RSVP-TE Extensions in support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery", IETF draft, work in progress, April 2005.

9.2. Informative References

- [ANSI-T1.105] American National Standards Institute, "Synchronous Optical Network (SONET) - Basic Description including Multiplex Structure, Rates, and Formats", ANSI T1.105-2001, May 2001.
- [ITU-T-G.7042] International Telecommunications Union, "Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenated Signals", ITU-T Recommendation G.7042, February 2004.

[ITU-T-G.7043] International Telecommunications Union, "Virtual Concatenation of Plesiochronous Digital Hierarchy (PDH) Signals", ITU-T Recommendation G.7043, July 2004.

[ITU-T-G.707] International Telecommunications Union, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)", ITU-T Recommendation G.707, December 2003.

[ITU-T-G.709] International Telecommunications Union, "Interfaces for the Optical Transport Network (OTN)", ITU-T Recommendation G.709, March 2003.

Author's Addresses

Greg Bernstein
Grotto Networking

Phone: +1-510-573-2237

Email: gregb@grotto-networking.com

Diego Caviglia
Ericsson
Via A. Negrone 1/A 16153
Genoa Italy

Phone: +39 010 600 3736

Email: diego.caviglia@marconi.com, ericsson.com)

Richard Rabbat
Fujitsu Laboratories of America
1240 East Arques Ave, MS 345
Sunnyvale, CA 94085
United States of America

Phone: +1 408-530-4537

Email: richard@us.fujitsu.com

Huub van Helvoort
Huawei Technologies, Ltd.
Kolkgriend 38, 1356 BC Almere
The Netherlands

Phone: +31 36 5315076
Email: hhelvoort@huawei.com

Intellectual Property Statement

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.

Disclaimer of Validity

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.

Copyright Statement

Copyright (C) The Internet Society (2006).

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.

Acknowledgment

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