

MPLS Working Group
INTERNET DRAFT

Ken-ichi Nagami (Toshiba Corp.)
Noritoshi Demizu (NAIST)
Hiroshi Esaki (Univ. Tokyo)
Yasuhiro Katsube (Toshiba Corp.)
Paul Doolan (Ennovate Networks)
December 1998
Expires June 1999

VCID Notification over ATM link
<[draft-ietf-mpls-vcid-atm-02.txt](#)>

Status of this memo

This document is an Internet-Draft. 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.''

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

The ATM Label Switching Router (ATM-LSR) is one of the major applications of label switching. Because the ATM layer labels (VPI and VCI) associated with a VC rewritten with new value at every ATM switch nodes, it is not possible to use them to identify a VC in label mapping messages. The concept of Virtual Connection Identifier (VCID) is introduced to solve this problem. VCID has the same value at both ends of a VC. This document specifies the procedures for the communication of VCID values between neighboring ATM-LSRs that must occur in order to ensure this property.

1. Introduction

Several label switching schemes have been proposed to integrate Layer 2 and Layer 3. The ATM Label Switching Router (ATM-LSR) is one of the major applications of label switching.

In the case of ATM VCs, the VPI and VCI labels are, in the general case, rewritten with new values at every switch node through which

the VC passes and cannot be used to provide end to end identification of a VC.

In the context of MPLS 'stream', which are classes of packets that have some common characteristic that may be deduced by examination of the layer 3 header in the packets, are bound to layer 2 'labels'. We speak of stream being 'bound' to labels. These bindings are conveyed between peer LSRs by means of a Label Distribution Protocol [[LDP](#)].

In order to apply MPLS to ATM links, we need some way to identify ATM VCs in LDP mapping messages. In [[VCID](#)], an identifier called a Virtual Connection ID (VCID) is introduced. VCID has the same value at both ends of a VC. This document specifies the procedures for the communication of VCID values between neighboring ATM-LSRs that must occur in order to ensure this property.

[2. Overview of VCID Notification Procedures](#)

[2.1 VCID Notification procedures](#)

The ATM has several types of VCs (transparent point-to-point link/VP/PVC/SVC). A transparent point-to-point link is defined as one that has the same VPI/VCI label at both ends of a VC. For example, two nodes are directly connected (i.e., without intervening ATM switches) or are connected through a VP with the same VPI value at both ends of the VP.

There are two broad categories of VCID notification procedures; inband and outband. The categorization refers to the connection over which the messages of the VCID notification procedure are forwarded. In the case of the inband procedures, those messages are forwarded over the VC to which they refer. In contrast the out of band procedures transmit the messages over some other connection (than the VC to which they refer).

We list below the various types of link and briefly mention the VCID notification procedures employed and the rational for that choice. The procedures themselves are discussed in detail in later sections.

Transparent point-to-point link : no VCID notification

VCID notification procedure is not necessary because the label (i.e., VPI/VCI) is the same at each end of the VC.

VP : inband notification or no notification

- Inband notification

VCID notification is needed because the VPI at each end of the VC may not be the same. Inband VCID notification [[VCID](#)] is used in this case.

- No notification

If a node has only one VP to a neighboring node, VCID notification procedure is not mandatory. The VCI can be used as the VCID. This is because the VCI value is the same at each end of the VP.

PVC : inband notification

Inband VCID notification [[VCID](#)] is used in this case because the labels at each end of the VC may not be the same.

SVC : there are three possibilities

- Outband notification

If a signaling message has a field which is large enough to carry a VCID value (e.g., GIT [[GIT](#)]), then the VCID is carried directly in it.

- Outband notification using a small-sized field

If a signaling message has a field which is not large enough to carry a VCID value, this procedure is used.

- Inband notification

If a signaling message can not carry user information, this procedure is used.

When an LSP is a point-to-multipoint VC and an ATM switch in an LSR is not capable of VC merge, it may cause problems in performance and quality of service. When the LSR wants to add a new leaf to the LSP, it needs to split the active LSP temporarily to send an inband notification message.

[3. VCID Notification Procedures](#)

[3.1 Inband Notification Procedures](#)

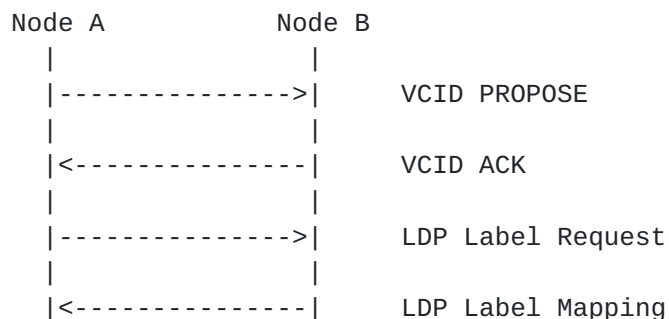
[3.1.1 Inband Notification for Point-to-point VC](#)

VCID notification is performed by transmitting a control message through the VC newly established (by signalling or management) for use as an label switched path (LSP) [[FRAME](#)]. The procedure for VCID notification between two nodes A and B is detailed below.

0. The node A establishes a VC to the destination node B. (by signalling or management)
1. The node A selects a VCID value.
2. The node A sends a VCID PROPOSE message which contains the VCID value and a message ID through the newly established VC to the node B.
3. The node A establishes an association between the outgoing label (VPI/VCI) for the VC and the VCID value.
4. The node B receives the message from the VC and establishes an

association between the VCID in the message and the incoming label(VPI/VCi) for the VC. Until the node B receives the LDP Request message, the node B discards any packet received over the VC other than the VCID PROPOSE message.

5. The node B sends an ACK message to the node A. This message contains the same VCID and message ID as specified in the received message. This message is sent through the VC for LDP.
6. When node A receives the ACK message, it checks whether the VCID and the message ID in the message are the same as the registered ones. If they are the same, node A regards that node B has established the association between the VC and VCID. Otherwise, the message is ignored. If the node A does not receive the ACK message with the expected message ID and VCID during a given period, the node A resends the VCID PROPOSE message to the node B.
7. After receiving the proposer ACK message, the node A sends an LDP REQUEST message to the node B. It contains the message ID used for VCID PROPOSE. When the node B receives the LDP REQUEST message, it regards that the node A has received the ACK correctly. The message exchange using VCID PROPOSE, VCID ACK, and LDP REQUEST messages constitutes a 3-way handshake. The 3-way handshake mechanism is required since the transmission of VCID PROPOSE message is unreliable. Once the 3-way handshake completes, the node B ignores all VCID PROPOSE messages received over the VC. The node B sends an LDP Mapping message, which contains the VCID value in the label TLV.



3.1.2 Inband notification for point-to-multipoint VC


Current LDP specification does not support multicast. But the VCID notification procedure is defined for future use. VCID notification is performed by sending a control message through the VC to be used as an LSP. The upstream node assigns the VCID value. The procedure by which it notifies the downstream node of that value is given below. The procedure is used when a new VC is created or a new leaf is added to the VC.

First, the procedure for establishing the first VC is described.

1. The upstream node assigns a VCID value for the VC. When the VCID value is already assigned to a VC, it is used for VCID.

2. The upstream node sends a message which contains the VCID value and a message ID through the VC used for an LSP. This message is transferred to all leaf nodes.

3. The upstream node establishes an association between the outgoing label for the VC and the VCID value.
4. When the downstream nodes receiving the message already received the LDP REQUEST message for the VC, the received message is discarded. Otherwise, the downstream nodes establish an association between the VCID in the message and the VC from which the message is received.
5. The downstream nodes send an ACK message to the upstream node.
6. After the upstream node receives the ACK messages, the upstream node and the downstream nodes share the VCID. The upstream node sends the LDP REQUEST message in order to make a 3-way handshake.



Second, the procedure for adding a leaf to the existing point-to-multipoint VC is described.

0. The upstream node adds the downstream node, using the ATM signaling.
1. The VCID value which already assigned to the VC is used.
2. The upstream node sends a message which contains the VCID value and a message ID through the VC used for an LSP. This message is transferred to all leaf nodes.
3. When the downstream nodes receiving the message already received the LDP REQUEST message for the VC, the received message is discarded. Otherwise, the downstream nodes establish an association between the VCID in the message and the VC from which the message is received.
4. After the upstream node receives the ACK messages, the upstream node and the downstream nodes share the VCID. The upstream node sends the LDP REQUEST message in order to make a 3-way handshake.

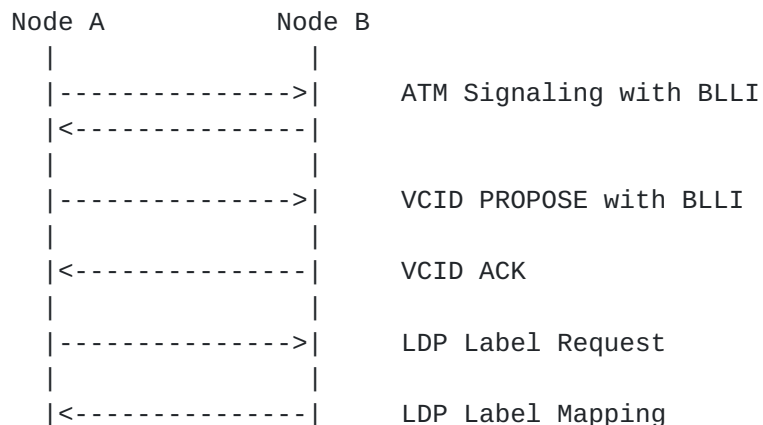
3.2 Outband Notification using a small-sized field

This method can be applied when a VC is established using an ATM signaling message and the message has a field which is not large

enough to carry a VCID value.

SETUP message of the ATM Forum UNI 3.1/4.0 has a 7-bit mandatory field for the user. This is a user specific field in the Layer 3 protocol field in the BLLI IE (Broadband Low Layer Information Information Element).

The BLLI value is used as a temporary identifier for a VC during a VCID notification procedure. This mechanism is defined as "Outband Notification using a small-sized field" described in [[VCID](#)]. The BLLI value of a new VC must not be assigned to other VCs during the procedure to avoid identifier conflict. When the association among the BLLI value, a VCID value, and the corresponding VC is established, the BLLI value can be reused for a new VC. VCID values can be assigned independently from BLLI values.



A point-to-multipoint VC can also be established using ADD_PARTY of the ATM Forum Signaling. ADD_PARTY adds a new VC leaf to an existing VC or an existing VC tree. In this procedure, the BLLI value of ADD_PARTY has to be the same value as that used to establish the first point-to-point VC of the tree. The same BLLI value can be used in different VC trees only when these VC trees can not add a leaf at the same time. As a result, the BLLI value used in the signaling must be determined by the root node of the multicast tree.

[note]

BLLI value is unique at the sender node. But BLLI value is not unique at the receiver node because multiple sender nodes may allocate the same BLLI value. So, the receiver node must recognize BLLI value and the sender address. ATM Signaling messages(SETUP and ADD_PARTY) carry both the BLLI and the sender ATM address. The receiver node can realize which node sends the BLLI message.

3.2.1 Outband notification using a small-sized field for point-to-point VC

This subsection describes procedures for establishing a VC and for notification of its VCID between neighboring LSRs for unicast traffic. VC pool [[VCPool](#)] can be applied.

The procedure employed when the upstream LSR assigns a VCID is as follows.

1. An upstream LSR establishes a VC to the downstream LSR using ATM signaling and supplies a value in the BLLI field that it is not currently using for any other (incomplete) VCID notification transaction with this peer.
2. The upstream LSR sends the VCID PROPOSE message through the VC for LDP to notify the downstream LSR of the association between the BLLI and VCID values.
3. The downstream LSR establishes the association between the VC with the BLLI value and the VCID and sends an ACK message to the upstream LSR.
4. After the upstream LSR receives the ACK message, it establishes the association between the VC and the VCID. The VC is ready to be used. At this time the BLLI value employed in this transaction is free for reuse.
5. After VCID notification, the upstream node sends the LDP REQUEST message to the downstream node. The downstream node sends the LDP mapping message, which contains the VCID value in the label TLV of LDP.

3.2.2 Outband notification using a small-sized field for point-to-multipoint VC

This subsection describes procedures for establishing the first VC for a multicast tree and for adding a new VC leaf to an existing VC tree including the notification of its VCID for a multicast stream using point-to-multipoint VCs.

In this procedure, an upstream LSR determines both the VCID and BLLI value in the multicast case. The reason that the BLLI value is determined by an upstream LSR is described above.

First, the procedure for establishing the first VC is described.

1. An upstream LSR establishes a VC by the ATM Forum Signaling between the downstream LSR with a unique BLLI value at this time.
2. The upstream LSR notifies the downstream LSR of a paired BLLI value and VCID using a message dedicated for this purpose.
3. The downstream LSR establishes the association between the VC with the BLLI value and the VCID and sends an ACK message to the upstream LSR. If the VCID is used by some other VC between the

upstream and downstream LSRs, the old VC is discarded.


4. After the upstream LSR receives the ACK message, the VC is ready to be used and the BLLI value can be used for another VC.

Second, the procedure for adding a leaf to the existing point-to-multipoint VC is described.

1. The upstream LSR establishes a VC by the ATM Forum Signaling between its downstream LSR with the BLLI value that was used during the first signaling procedure. If another VC is using the BLLI value at the same time, the upstream waits for the completion of the signaling procedure that is using this BLLI value.
2. Go to step 2 of the procedure for the first VC.

3.3 Outband notification

This method can be applied when a VC is established using a ATM signaling message and the message has a field (e.g., GIT [[GIT](#)]) which is large enough to carry a VCID value. Message format is described in [[GIT](#)]. After the VCID notification, the node A sends the LDP request message to the node B. Then, the node B sends the LDP mapping message to the node A.



4 VCID Message Format

4.1 VCID Messages

An LDP VCID message consists of the LDP [[LDP](#)] fixed header followed by one or more VCID TLV. VCID PROPOSE inband message is sent as a null encapsulation packet through a VC to be used as an LSP. There is only the label stack header before the LDP VCID PDU. A label value in the label stack entry [[ENCAPS](#)] for VCID PROPOSE inband message is 4. Other messages are sent as TCP packets. This is the same as LDP.

The VCID message type field is as follows:

VCID Propose inband Message	= 0x0501
VCID Propose Message	= 0x0502
VCID ACK Message	= 0x0503
VCID NACK Message	= 0x0504

4.1.1 VCID Propose inband Message

This message is sent as a null encapsulation packet with LDP header and label stack header through a VC to be used as an LSP. The label value is 4. The reserved label value is required because the

downstream node may receive this message after receiving the LDP Label Request message in the case of point-to-multipoint VC. The downstream node must distinguish the VCID PROPOSE message from other messages and ignore the VCID PROPOSE message when the node already received the LDP Label Request message for the VC.

```

      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|VCID Inband Propose (0x0501) |      Message Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Message ID                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Label TLV                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Optional Parameters             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Message Id

Four octet integer used to identify this message.

Label TLV

Label TLV contains VCID value. Type of label TLV is VCID(0x0203).

4.1.2 VCID Propose Message

An LSR uses the VCID PROPOSE message for the VCID notification procedure of the outband notification using a small-sized field. This message is sent through the VC for the LDP.

```

      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|  VCID Propose (0x0502)    |      Message Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Message ID                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Label TLV                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Temporary ID TLV               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Optional Parameters             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Message Id

Four octet integer used to identify this message.

Label TLV

Label TLV contains VCID value. Type of label TLV is VCID(0x0203).

Temporary ID TLV

The value carried in the user specific field in the layer 3
protocol field in the BLLI ID in the ATM Forum UNI 3.1/4.0

Nagami, et al.

[Page 9]

Type of label TLV is VCID temporary ID(0x0902).

4.1.3 VCID ACK Message

An LSR send the VCID ACK message when the LSR accepts the VCID PROPOSE message.

[illegible]

Message Id

Four octet integer used to identify this message. This value is the same as that of received VCID PROPOSE message.

Label TLV

The label TLV contains the VCID value of the received VCID PROPOSE message. Type of label TLV is VCID(0x0203).

4.1.4 VCID NACK Message

An LSR send the VCID NACK message when the LSR does not accept the VCID PROPOSE message.

[illegible]

Message Id

Four octet integer used to identify this message. This value is the same as that of received VCID PROPOSE message.

Label TLV

The label TLV contains the VCID value of the received VCID PROPOSE message. Type of label TLV is VCID(0x0203).

4.2 Objects

4.2.1 VCID Label TLV

An LSR uses VCID Label TLV to encode labels for use on the link which does not have the same data link label at both ends of a VC.

```

      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|VCID Label  (0x0203)      |              Length              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               VCID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

VCID

This is 4 byte VCID value.

4.2.2 VCID Temporary ID TLV

An LSR uses the VCID temporary ID TLV for the VCID notification procedure of the outband notification using a small-sized field.

```

      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| VCID Temporary ID (0x0601)|              Length              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Temporary ID |
+---+---+---+---+---+

```

Temporary ID:

The value carried in the user specific field in the layer 3 protocol field in the BLLI ID in the ATM Forum UNI 3.1/4.0

Security Considerations

Security issues are not discussed in this document.

Intellectual Property Considerations

Toshiba Corporation and Ennovate Networks may seek patent or other intellectual property protection for some of the aspects of the technology discussed in this document. If any standards arising from this document are or become protected by one or more patents assigned to Toshiba Corporation, Toshiba intends to license them on reasonable and non- discriminatory terms.

Acknowledgments

Nagami, et al.

[Page 11]

The authors would like to acknowledge the valuable technical comments of Shigeo Matsuzawa, Akiyoshi Mogi, Muneyoshi Suzuki, George Swallow and members of the LAST-WG of the WIDE Project.

References

- [VCID] N. Demizu, et al., "VCID: Virtual Connection Identifier",
[draft-demizu-mpls-vcid-01.txt](#), Oct. 1997
- [VCPool] N. Demizu, et al., "VC pool",
[draft-demizu-mpls-vcpool-00.txt](#), Oct. 1997
- [LDP] L. Andersson, et al., "LDP Specification",
[draft-ietf-mpls-ldp-02.txt](#), Nov. 1998
- [FRAME] R. Callon, et al., "A Framework for Multiprotocol Label
Switching", [draft-ietf-mpls-framework-02.txt](#), Nov. 1997
- [GIT] M. Suzuki, "The Assignment of the Information Field and
Protocol Identifier in the Q.2941 Generic Identifier and Q.2957
User-to-user Signaling for the Internet Protocol",
[draft-ietf-mpls-git-uus-01.txt](#), Dec. 1998
- [ENCAPS] E. Rosen, et al., "MPLS Label Stack Encoding",
[draft-ietf-mpls-label-encaps-03.txt](#), Sep. 1998

Authors Information

Ken-ichi Nagami
Information & Communication Lab., Toshiba Corporation,
3-1-1 Asahigaoka, Hino,
Tokyo, 191-8555, Japan
Phone: +81-42-585-3299
Email: ken.nagami@toshiba.co.jp

Noritoshi Demizu
Graduate School of Information Science,
Nara Institute of Science and Technology
8916-5 Takayama, Ikoma, Nara 630-0101, Japan
Phone: +81-743-72-5348
Email: nori-d@is.aist-nara.ac.jp

Hiroshi Esaki
Computer Center, University of Tokyo,
2-11-16 Yayoi, Bunkyo-ku,
Tokyo, 113-8658, Japan
Phone: +81-3-3812-1111
Email: hiroshi@wide.ad.jp

Yasuhiro Katsube
Information & Communication Lab., Toshiba Corporation,
3-1-1 Asahigaoka, Hino,

Nagami, et al.

[Page 12]

Internet Draft

[draft-ietf-mpls-vcid-atm-02.txt](#)

December 1998

Tokyo, 191-8555, Japan

Phone: +81-42-585-3299

Email: yasuhiko.katsube@toshiba.co.jp

Paul Doolan

Ennovate Networks

330 Codman Hill Road

Boxborough, MA

Phone: 978-263-2002 x103

Email: pdoolan@ennovatenetworks.com

