

Network Working Group
Request for Comments: 3038
Category: Standards Track

K. Nagami
Y. Katsube
Toshiba Corp.
N. Demizu
WaterSprings.ORG
H. Esaki
Univ. of Tokyo
P. Doolan
Ennovate Networks
January 2001

VCID Notification over ATM link for LDP

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

The Asynchronous Transfer Mode 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. 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 rationale 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 VPID 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 is used in this case.

- VPID notification
VCID notification is needed because the VPI at each end of the VC may not be the same. VPID notification 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 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.

2.2 VC direction

A VC has a directionality. The VCID procedure for a VC is always triggered from the upstream node of the VC, i.e., the upstream node notifies the downstream node of the VCID.

If bidirectional use of a label switched VC is allowed, the label switched VC is said to be bidirectional. In this case, two VCID procedures are taken, one for each direction.

If bidirectional use of a label switched VC is not allowed, the label switched VC is said to be unidirectional. In this case, only one VCID procedure is taken for the allowed direction.

VC directionality is communicated through LDP.

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). 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.

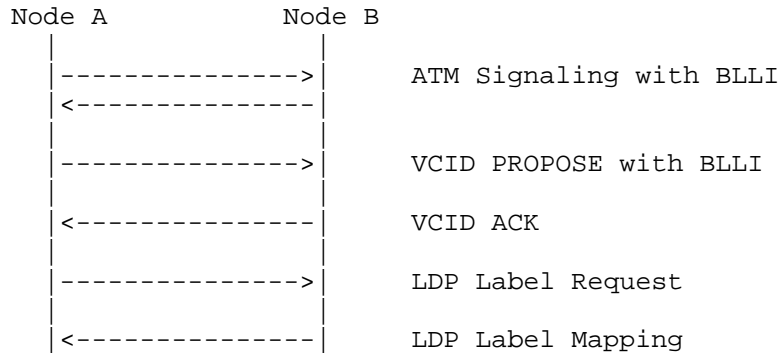
3.2 Outband Notification using a small-sized field

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". 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.

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 is sent to the node B. Then, the node B sends the LDP mapping message to the node A.



4 VPID Notification Procedure

The approach that is used for the VCID notification procedure is also applicable to share the same identifier between both ends for a VP. VPID notification procedure is defined for this purpose.

A distinct VPID notification procedure is performed for each direction of each VP.

After the VPID notification is finished for a VP, a VCID of a VC in the VP is constructed with the VPID(MSB) and VCI(LSB) of the VC. The VCID can be used by LDP without performing VCID notification procedure. The message sequence is given below.

1. An upstream node sends the VPID PROPOSE message. In the case of bidirectional label switched VC, both the upstream and downstream nodes use VCI=33. In the case of unidirectional label switched VC, the node which has larger LDP Identifier uses VCI=33 and the other node uses VCI=34. Note that VCI=32, which is used for unlabeled packet transfer, is not used for VPID notification procedure so that the same encapsulation method can be applied for both VPID procedure and inband VCID procedure.
2. The downstream node sends the VPID ACK message.
3. The upstream node sends the LDP Label Request message.
4. The downstream node sends the LDP Label Mapping message.

5 VCID Message Format

5.1 VCID Messages

An LDP VCID message consists of the LDP [LDP] fixed header followed by one or more TLV. A VCID PROPOSE inband message and a VPID PROPOSE message are 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 the VCID PROPOSE inband message and the VPID PROPOSE message are 4. Other messages are sent as TCP packets. This is the same as LDP.

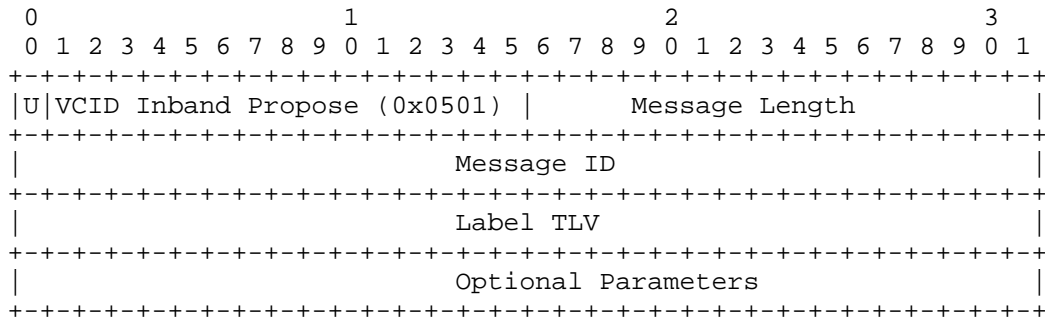
The VCID message type field is as follows:

VCID Propose inband Message	= 0x0501
VCID Propose Message	= 0x0502
VCID ACK Message	= 0x0503
VCID NACK Message	= 0x0504
VPID Propose inband Message	= 0x0505
VPID ACK Message	= 0x0506
VPID NACK Message	= 0x0507

5.1.1 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.

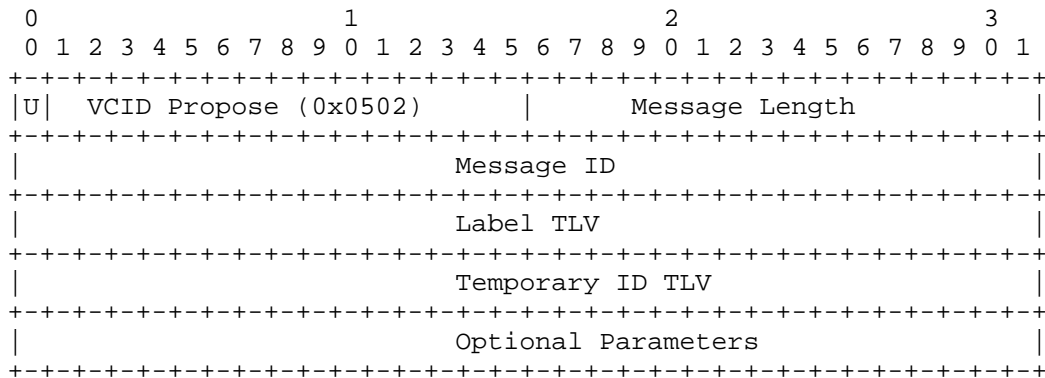


Message Id
Four octet integer used to identify this message.

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

5.1.2 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.



Message ID
Four octet integer used to identify this message.

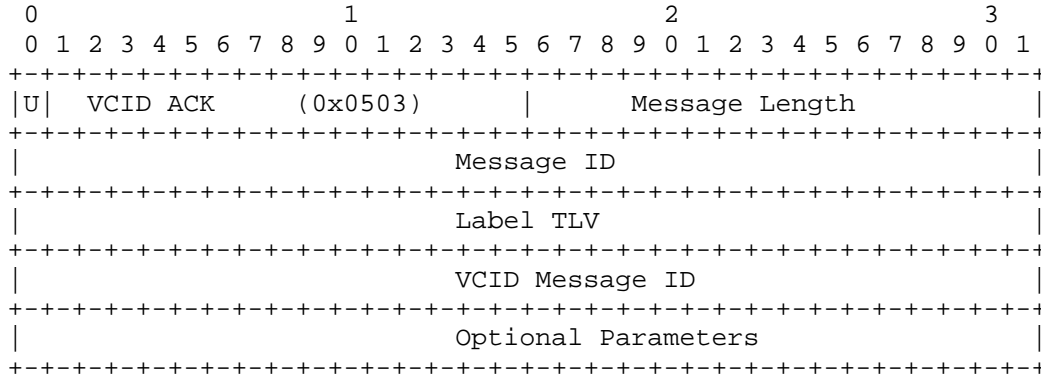
Label TLV
Label TLV contains VCID value. Type of label TLV is VCID(0x0203).

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 Type of label TLV is VCID temporary ID(0x0702).

5.1.3 VCID ACK Message

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



Message ID

Four octet integer used to identify this message.

Label TLV

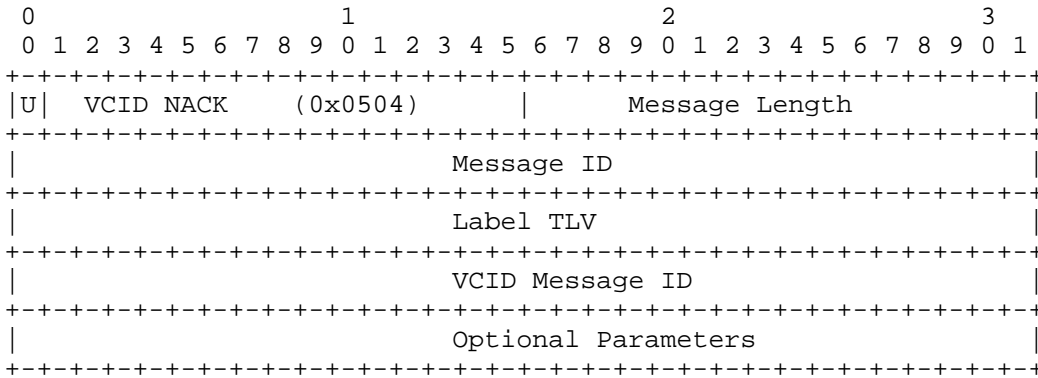
The label TLV contains the VCID value of the received VCID PROPOSE message. Type of label TLV is VCID(0x0203).

VCID Message ID

This value is the same as that of received VCID PROPOSE message.

5.1.4 VCID NACK Message

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



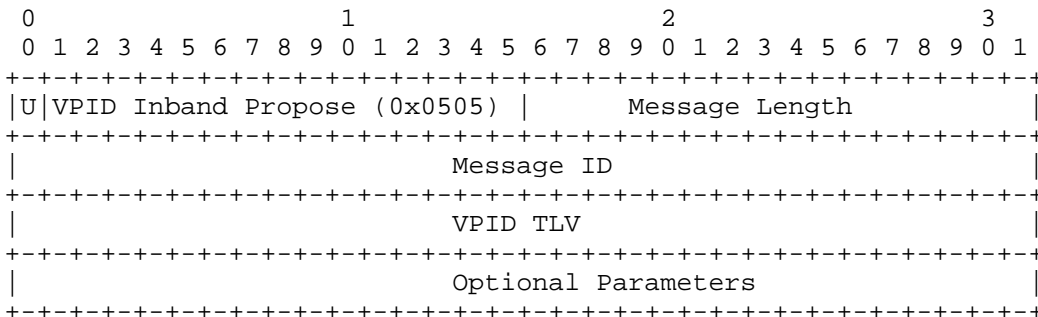
Message ID
 Four octet integer used to identify this message.

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

VCID Message ID
 This value is the same as that of received VCID PROPOSE message.

5.1.5 VPID 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 downstream node must distinguish the VPID PROPOSE message from other messages and ignore the VPID PROPOSE message when the node already received the LDP Label Request message for the VC.



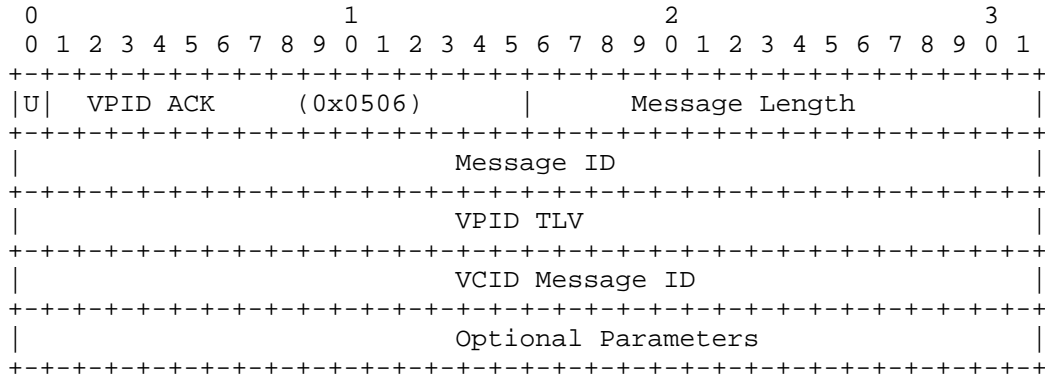
Message Id
 Four octet integer used to identify this message.

VPID TLV

VPID TLV contains VPID value. Type of label TLV is VPID(0x0703).

5.1.6 VPID ACK Message

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



Message ID

Four octet integer used to identify this message.

VPID TLV

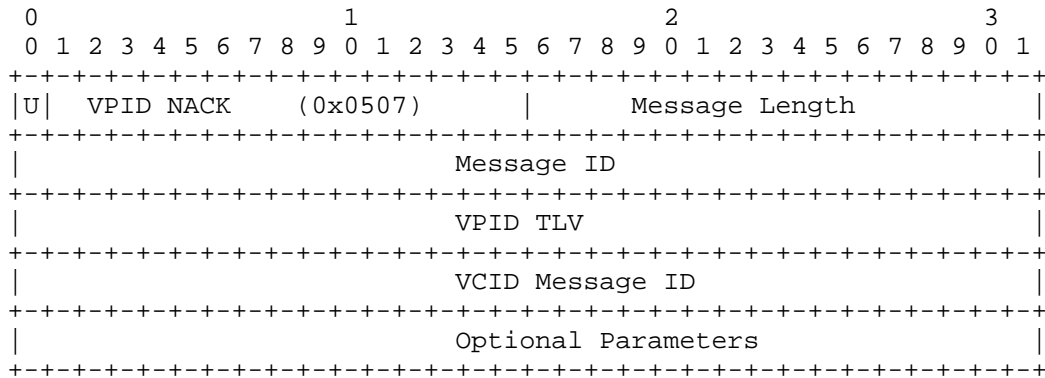
The VPID TLV contains the VPID value of the received VPID PROPOSE message.

VCID Message ID

This value is the same as that of received VCID PROPOSE message.

5.1.7 VPID NACK Message

An LSR send the VPID NACK message when the LSR accepts the VPID PROPOSE message.



Message ID
Four octet integer used to identify this message.


VPID TLV
The VPID TLV contains the VPID value of the received VPID PROPOSE message.

VCID Message ID
This value is the same as that of received VCID PROPOSE message.

5.2 Objects


5.2.1 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.



VCID
This is 4 byte VCID value.


5.2.2 VCID Message ID TLV



VCID Message ID
This is 4 byte VCID Message ID

5.2.3 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.



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

5.2.4 VPID Label TLV

An LSR uses VPID TLV for the VPID notification procedure.



VPID
This is 2 byte VPID value.

Security Considerations

This document does not introduce new security issues other than those present in the LDP and may use the same mechanisms proposed for this technology.

Acknowledgments

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

References

- [LDP] Andersson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [GIT] Suzuki, M., "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", RFC 3033, January 2001.
- [ENCAPS] Rosen, E., Viswanathan, A. and R. Callon, "MPLS Label Stack Encoding", RFC 3032, January 2001.

Authors' Addresses

Ken-ichi Nagami
Computer & Network Development Center, Toshiba Corporation,
1, Toshiba-cho, Fuchu-shi,
Tokyo, 183-8511, Japan

Phone: +81-42-333-2884
EMail: ken.nagami@toshiba.co.jp

Noritoshi Demizu
WaterSprings.ORG
1-6-11-501, Honjo, Sumida-ku, Tokyo, 130-0004, Japan

EMail: demizu@dd.iiij4u.or.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
Computer & Network Development Center, Toshiba Corporation,
1, Toshiba-cho, Fuchu-shi,
Tokyo, 183-8511, Japan

Phone: +81-42-333-2844
EMail: yasuhiko.katsube@toshiba.co.jp

Paul Doolan
Ennovate Networks
60 Codman Hill Road
Boxborough, MA

Phone: 978-263-2002 x103
EMail: pdoolan@ennovatenetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

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

