

Extensions to Support Efficient Carrying of
Multicast Traffic in Layer-2 Tunneling Protocol (L2TP)

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The Layer Two Tunneling Protocol (L2TP) provides a method for tunneling PPP packets. This document describes an extension to L2TP, to make efficient use of L2TP tunnels within the context of deploying multicast services whose data will have to be conveyed by these tunnels.

Table of Contents

1.	Introduction.....	2
1.1.	Conventions Used in This Document.....	3
1.2.	Terminology.....	3
2.	Motivation for a Session-Based Solution.....	4
3.	Control Connection Establishment.....	5
3.1.	Negotiation Phase.....	5
3.2.	Multicast Capability AVP (SCCRQ, SCCRP).....	5
4.	L2TP Multicast Session Establishment Decision.....	6
4.1.	Multicast States in LNS.....	6
4.2.	Group State Determination.....	8
4.3.	Triggering.....	9
4.4.	Multicast Traffic Sent from Group Members.....	10
5.	L2TP Multicast Session Opening Process.....	11
5.1.	Multicast-Session-Request (MSRQ).....	11
5.2.	Multicast-Session-Response (MSRP).....	12
5.3.	Multicast-Session-Establishment (MSE).....	12
6.	Session Maintenance and Management.....	13
6.1.	Multicast-Session-Information (MSI).....	13
6.2.	Outgoing Sessions List Updates.....	14

6.2.1.	New Outgoing Sessions AVP (MSI).....	15
6.2.2.	New Outgoing Sessions Acknowledgement AVP (MSI).	15
6.2.3.	Withdraw Outgoing Sessions AVP (MSI).....	17
6.3.	Multicast Packets Priority AVP (MSI).....	17
6.3.1.	Global Configuration.....	18
6.3.2.	Individual Configuration.....	19
6.3.3.	Priority.....	19
7.	Multicast Session Teardown.....	19
7.1.	Operations.....	20
7.2.	Multicast-Session-End-Notify (MSEN).....	20
7.3.	Result Codes.....	21
8.	Traffic Merging.....	22
9.	IANA Considerations.....	22
10.	Security Considerations.....	23
11.	References.....	23
11.1.	Normative References.....	23
11.2.	Informative References.....	24
12.	Acknowledgements.....	24
Appendix A.	Examples of Group States Determination.....	25
Author's Address.....		27
Full Copyright Statement.....		28

1. Introduction

The deployment of IP multicast-based services may have to deal with L2TP tunnel engineering. The forwarding of multicast data within L2TP sessions may impact the throughput of L2TP tunnels because the same traffic may be sent multiple times within the same L2TP tunnel, but in different sessions. This proposal aims to reduce the impact by applying the replication mechanism of multicast traffic only when necessary.

The solution described herein provides a mechanism for transmitting multicast data only once for all the L2TP sessions that have been established in a tunnel, each multicast flow having a dedicated L2TP session.

Within the context of deploying IP multicast-based services, it is assumed that the routers of the IP network that embed a L2TP Network Server (LNS) capability may be involved in the forwarding of multicast data, toward users who access the network through an L2TP tunnel. The LNS is in charge of replicating the multicast data for each L2TP session that a receiver who has requested a multicast flow uses. In the solution described here, an LNS is able to send multicast data only once and to let the L2TP Access Concentrator (LAC) perform the traffic replication. By doing so, it is expected to spare transmission resources in the core network that supports

L2TP tunnels. This multicast extension to L2TP is designed so that it does not affect the behavior of L2TP equipment under normal conditions.

A solution whereby multicast data is carried only once in a L2TP tunnel is of interest to service providers, as edge devices are aggregating more and more users. This is particularly true for operators who are deploying xDSL (Digital Subscriber Line) services and cable infrastructures. Therefore, L2TP tunnels that may be supported by the network will have to carry multiple redundant multicast data more often. The solution described in this document applies to downstream traffic exclusively; i.e., data coming from the LNS toward end-users connected to the LAC. This downstream multicast traffic is not framed by the LNS but by the LAC, thus ensuring compatibility for all users in a common tunnel, whatever the framing scheme.

1.1. 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 [RFC2119].

1.2. Terminology

Unicast session

This term refers to the definition of "Session" as it is described in the terminology section of [RFC2661].

Multicast session

This term refers to a connection between the LAC and the LNS. Additional Control Messages and Attribute-Value-Pairs (AVPs) are defined in this document to open and maintain this connection for the particular purpose of multicast traffic transportation. This connection between the LAC and the LNS is intended to convey multicast traffic only.

Session

This term is used when there is no need to dissociate multicast from unicast sessions, and thus it designates both.

M-IGP

Designates a Multicast Interior Gateway Protocol.

Multicast flow

Designates datagrams sent to a group from a set of sources for which multicast reception is desired.

GMP

Group Management Protocol, such as:

- IGMPv1 ([RFC1112])
- IGMPv2 ([RFC2236])
- MLD ([RFC2710], [RFC3590])

SFGMP

Source Filtering Group Management Protocol, such as:

- IGMPv3 ([RFC3376])
- MLDv2 ([RFC3810])

2. Motivation for a Session-Based Solution

Multicast data have to be seen as a singular flow that may be conveyed into all the L2TP sessions that have been established in a tunnel. This means that a given L2TP session can be dedicated for the forwarding of a multicast flow that will be forwarded to multiple receivers, including those that can be reached by one or several of these L2TP sessions. A session carrying IP multicast data is independent from the underlying framing scheme and is therefore compatible with any new framing scheme that may be supported by the L2TP protocol.

Using a single L2TP session per multicast flow is motivated by the following arguments:

- The administrator of the LNS is presumably in charge of the IP multicast-based services and the related engineering aspects. As such, he must be capable of filtering multicast traffic on a multicast source basis, on a multicast group basis, and on a user basis (users who access the network using an L2TP session that terminates in this LNS).

- Having an L2TP session dedicated for a multicast flow makes it possible to enforce specific policies for multicast traffic. For instance, it is possible to change the priority treatment for multicast packets against unicast packets.
- It is not always acceptable or possible to have multicast forwarding performed within the network between the LAC and the LNS. Having the multicast traffic conveyed within an L2TP tunnel ensures a multicast service between the LNS and end-users, alleviating the need for activating multicast capabilities in the underlying network.

3. Control Connection Establishment

3.1. Negotiation Phase

The multicast extension capability is negotiated between the LAC and the LNS during the control connection establishment phase. However, establishment procedures defined in [RFC2661] remain unchanged. An LAC indicates its multicast extension capability by using a new AVP, the "Multicast Capability" AVP. There is no explicit acknowledgement sent by the LNS during the control connection establishment phase. Instead, the LNS is allowed to use multicast extension messages to open and maintain multicast sessions (see Section 5).

3.2. Multicast Capability AVP (SCCRQ, SCCRP)

In order to inform the LNS that an LAC has the ability to handle multicast sessions, the LAC sends a Multicast Capability AVP during the control connection establishment phase. This AVP is either sent in a SCCRQ or a SCCRP control message by the LAC towards the LNS.

Upon receipt of the Multicast Capability AVP, a LNS may adopt two distinct behaviors:

- 1) The LNS does not implement the L2TP multicast extension: any multicast-related information (including the Multicast Capability AVP) will be silently ignored by the LNS.
- 2) The LNS implements L2TP multicast extensions and therefore supports the Multicast Capability AVP: the LNS is allowed to send L2TP specific commands for conveying multicast traffic toward the LAC.

The multicast capability exclusively refers to the tunnel for which the AVP has been received during the control connection establishment phase. It SHOULD be possible for an LNS administrator to shut down

As a multicast router, the equipment that embeds the LNS function will keep state per group per attached network (i.e., per L2TP session). The LNS-capable equipment activating multicast extensions for L2TP will have to classify and analyze GMP and SFGMP states in order to create L2TP multicast sessions within the appropriate L2TP tunnels. This is performed in three steps:

- 1) The LNS has to compute group states for each L2TP tunnel, by using group states recorded for each L2TP session of the tunnel. Group state determination for L2TP tunnels is discussed in Section 4.2. For each L2TP tunnel, the result of this computation will issue a list of states of the form (group, filter-mode, source-list):
 - group: Denotes the multicast group.
 - filter-mode: Either INCLUDE or EXCLUDE, as defined in [RFC3376].
 - source-list: List of IP unicast addresses from which multicast reception is desired or not, depending on the filter-mode.
- 2) According to each group state, the LNS will create one or multiple replication contexts, depending on the filter-mode for the considered group and the local policy configured in the LNS.

For groups in INCLUDE mode, the LNS SHOULD implement two different policies:

- One session per (source, group) pair: the LNS creates one replication context per (source, group) pair.
- or
- One session per group: the LNS creates one replication context per (source-list, group) pair.

For groups in EXCLUDE mode, the LNS will create one replication context per (list of sources excluded by *all* the receivers, group). The list of sources represents the intersection of the sets, not the union.

- 3) For each replication context, the LNS will create one L2TP multicast session (if threshold conditions are met; see Section 4.3) and its associated Outgoing Session List (OSL). The OSL lists L2TP sessions that requested the multicast flow corresponding to the group and the associated source-filtering properties. There is one OSL per replication context; i.e., per L2TP multicast session.

For a group member running an SFGMP, it is therefore possible to receive multicast traffic from sources that have been explicitly excluded in its SFGMP membership report if other group members in the

same L2TP tunnel wish to receive packets from these sources. This behavior is comparable to the case where group members are connected to the same multi-access network. When a group is in EXCLUDE mode or in INCLUDE mode with a policy allowing one session per (group, source-list), sharing the same L2TP tunnel is equivalent to being connected to the same multi-access network in terms of multicast traffic received. For groups in INCLUDE mode with a policy allowing one L2TP multicast session per (source, group), the behavior is slightly improved because it prevents group members from receiving traffic from non-requested sources. On the other hand, this policy potentially increases the number of L2TP multicast sessions to establish and maintain. Examples are provided in Appendix A.

In order for the LAC to forward the multicast traffic received through the L2TP multicast session to group members, the LNS sends the OSL to the LAC for the related multicast session (see Section 6).

4.2. Group State Determination

Source Filtering Group Management Protocols require querier routers to keep a filter-mode per group per attached network, to condense the total desired reception state of a group to a minimum set so that all systems' memberships are satisfied.

Within the context of L2TP, each L2TP session has to be considered an attached network by GMP and SFGMP protocols. When the L2TP multicast extension is activated, each L2TP Control Connection has to be considered a pseudo attached network, as well, in order to condense group membership reports for every L2TP session in the tunnel.

Therefore, a list of group states is maintained for each L2TP Control Connection into which the membership information of each of its L2TP sessions is merged. This list of group states is a set of membership records of the form (group, filter-mode, source-list).

Each group state represents the result of a merging process applied to subscriptions on L2TP sessions of a Control Connection for a considered group. This merging process is performed in three steps:

- 1) Conversion of any GMP subscription into SFGMP subscription (IGMPv1/v2 to IGMPv3, MLDv1 to MLDv2);
- 2) Removal of subscription timers and, if filter-mode is EXCLUDE, sources with source timer > 0;
- 3) Then, resulting subscriptions are merged by using merging rules described in SFGMP specifications ([RFC3376], Section 3.2, [RFC3810], Section 4.2).

This process is also described in [PROXY]. Examples of group state determination are provided in Appendix A.

4.3. Triggering

The rules to be enforced by the LNS whereby it is decided when to open a dedicated L2TP multicast session for a multicast group SHOULD be configurable by the LNS administrator. This would typically happen whenever a threshold of MULTICAST_SESSION_THRESHOLD receivers/sessions referenced in a replication context is reached. This threshold value SHOULD be valued at 2 by default, as it is worth opening a dedicated L2TP multicast session for two group members sharing the same desired reception state (which means that two L2TP unicast sessions are concerned). In this case, the OSL will reference two distinct L2TP sessions.

The actual receipt by the LNS of multicast traffic requested by end-users can also be taken into account to decide whether the associated L2TP multicast session has to be opened.

Whenever an OSL gets empty, the LNS MUST stop sending multicast traffic over the corresponding L2TP multicast session. Then the L2TP multicast session MUST be torn down as described in Section 7.

Filter-mode changes for a group can also trigger the opening or the termination of L2TP multicast sessions in the following ways:

a) From INCLUDE Mode to EXCLUDE Mode

When a group state filter-mode switches from INCLUDE to EXCLUDE, only one replication context (and its associated L2TP multicast session) issued from this group state can exist (see Section 4.1). The LNS SHOULD keep one replication context previously created for this group state and it has to update it with:

- a new source-list that has to be excluded from forwarding
- a new OSL

The LNS MUST send an OSL update to the LAC to reflect L2TP session list changes (section 6.2), whenever appropriate. The unused L2TP multicast sessions that correspond to previously created replication contexts for the group SHOULD be terminated, either actively or passively by emptying their corresponding OSLs.

The remaining L2TP multicast session MAY also be terminated if the number of receivers is below a predefined threshold (see Section 7). To limit the duration of temporary packet loss or duplicates to receivers, the LNS has to minimize delay between OSL updates messages

sent to the LAC. Therefore, one can assume that terminating a multicast session passively gives the smoothest transition.

b) From EXCLUDE Mode to INCLUDE Mode

When a group state filter-mode switches from EXCLUDE to INCLUDE, multiple replication contexts issued by this group state may be created (see Section 4.1). The LNS SHOULD keep the replication context previously created for this group state and it has to update it accordingly with the following information:

- a new list of sources that has to be forwarded. This list has only one record if there is one replication context per (group, source)
- a new OSL

The LNS MUST send an OSL update to the LAC to reflect L2TP session list changes, whenever appropriate. If the LNS is configured to create one replication context per (group, source), L2TP multicast sessions will be opened in addition to the existing one, depending on the number of sources for the group.

If new L2TP multicast sessions have to be opened, the LNS SHOULD wait until these multicast sessions are established before updating the OSL of the original multicast session. To limit the duration of temporary packet loss or duplicates to receivers, the LNS has to minimize delay between OSL updates messages sent to the LAC.

4.4. Multicast Traffic Sent from Group Members

The present document proposes a solution to enhance the forwarding of downstream multicast traffic exclusively; i.e., data coming from the LNS toward end-users connected to the LAC. If a group member that uses an L2TP session is also a multicast source for traffic conveyed in a multicast session, datagrams may be sent back to the source. To prevent this behavior, two options can be used in the LNS:

- 1) Disable the multicast packets' forwarding capability, for those multicast datagrams sent by users connected to the network by means of an L2TP tunnel. Protocols using well-known multicast addresses MUST NOT be impacted.
- 2) Exclude from the OSL the L2TP session used by a group member that sends packets matching the replication context of this OSL. Therefore, the corresponding multicast flow is sent by the LNS over the user L2TP unicast session, using standard multicast forwarding rules.

5. L2TP Multicast Session Opening Process

The opening of an L2TP multicast session is initiated by the LNS. A three-message exchange is used to set up the session. The following is a typical sequence of events:

```

LAC                LNS
---                ---
                    (multicast session
                    triggering)

                    <- MSRQ

MSRP ->

(Ready to
 replicate)

MSE  ->

                    <- ZLB ACK

```

The ZLB ACK is sent if there are no further messages waiting in the queue for that peer.

5.1. Multicast-Session-Request (MSRQ)

Multicast-Session-Request (MSRQ) is a control message sent by the LNS to the LAC to indicate that a multicast session can be created. The LNS initiates this message according to the rules in Section 4.3. It is the first in a three-message exchange used for establishing a multicast session within an L2TP tunnel.

A LNS MUST NOT send a MSRQ control message if the remote LAC did not open the L2TP tunnel with the Multicast Capability AVP. The LAC MUST ignore MSRQ control messages sent in an L2TP tunnel, if the L2TP tunnel was not opened with control messages including a Multicast Capability AVP.

The following AVPs MUST be present in MSRQ:

```

Message Type
Assigned Session ID

```

The following AVPs MAY be present in MSRQ:

```

Random Vector
Maximum BPS

```

The Maximum BPS value is set by the LNS administrator. However, this value should be chosen in accordance with the line capabilities of the end-users. The Maximum BPS value SHOULD NOT be higher than the highest speed connection for all end-users within the L2TP tunnel.

The associated Message Type AVP is encoded with the following values:

Vendor ID = 0
Attribute Type = 0
Attribute Value = 23 (16 bits)

The M-bit MUST be set to 0, and the H-bit MUST be set to 0.

5.2. Multicast-Session-Response (MSRP)

Multicast-Session-Response (MSRP) is a control message sent by the LAC to the LNS in response to a received MSRQ message. It is the second in a three-message exchange used for establishing a multicast session within an L2TP tunnel.

MSRP is used to indicate that the MSRQ was successful and that the LAC will attempt to reserve appropriate resources to perform multicast replication for unicast sessions managed in the pertaining control connection.

The following AVPs MUST be present in MSRP:

Message Type
Assigned Session ID

The following AVP MAY be present in MSRP:

Random Vector

The associated Message Type AVP is encoded with the following values:

Vendor ID = 0
Attribute Type = 0
Attribute Value = 24 (16 bits)

The M-bit MUST be set to 0, and the H-bit MUST be set to 0.

5.3. Multicast-Session-Establishment (MSE)

Multicast-Session-Establishment (MSE) is a control message sent by the LAC to the LNS to indicate that the LAC is ready to receive necessary multicast information (Section 6) for the group using the

newly created multicast session. It is the third message in the three-message sequence used for establishing a multicast session within an L2TP tunnel.

The following AVP MUST be present in MSE:

Message Type

The following AVP MAY be present in MSE:

Sequencing Required

Sequencing will occur only from the LNS to the LAC, as a multicast session is only used to forward multicast traffic downstream.

The associated Message Type AVP is encoded with the following values:

Vendor ID = 0
Attribute Type = 0
Attribute Value = 25 (16 bits)

The M-bit MUST be set to 0, and the H-bit MUST be set to 0.

6. Session Maintenance and Management

Once the multicast session is established, the LAC has to be informed of the L2TP unicast sessions interested in receiving the traffic from the newly created multicast session, and a related optional priority parameter, defined in Section 6.3. To achieve this, a new control message type is defined: Multicast-Session-Information (MSI).

6.1. Multicast-Session-Information (MSI)

Multicast-Session-Information (MSI) control messages carry AVPs to keep the OSL synchronized between the LNS and the LAC, and to set the optional priority parameter for multicast traffic versus unicast traffic. MSI may be extended to update the multicast session with additional parameters, as needed.

Each MSI message is specific to a particular multicast session. Therefore, the control message MUST use the assigned session ID associated with the multicast session (assigned by the LAC), except for the case mentioned in 6.3.2.

The associated Message Type AVP is encoded with the following values:

Vendor ID = 0
Attribute Type = 0
Attribute Value = 26 (16 bits)

The M-bit MUST be set to 0, and the H-bit MUST be set to 0.

The following AVP MUST be present in MSI:

Message Type

The following AVPs MAY be present in MSI:

Random Vector
New Outgoing Sessions
New Outgoing Sessions Acknowledgement
Withdraw Outgoing Sessions
Multicast Packets Priority

New Outgoing Sessions, New Outgoing Sessions Acknowledgement, Withdraw Outgoing Sessions, and Multicast Packets Priority are new AVPs defined in sections 6.2 and 6.3.

6.2. Outgoing Sessions List Updates

Whenever a change occurs in the Outgoing Sessions List, the LNS MUST inform the LAC of that change. The OSL is built upon subscription reports recorded by GMP or SFGMP processes running in the LNS (Section 4.1).

The LAC maintains an OSL as a local table transmitted by the LNS. As for the LNS, the LAC has to maintain an OSL for each L2TP multicast session within an L2TP tunnel. To update the LAC OSL, the LNS sends a New Outgoing Sessions AVP for additional sessions, or sends a Withdraw Outgoing Sessions AVP to remove sessions. All sessions mentioned in these AVPs MUST be added or removed by the LAC from the relevant OSL. The Outgoing Sessions List is identified by the tunnel ID and the multicast session ID to which the updating AVP refers. To update the OSL, the following AVPs are used:

Additional session(s): New Outgoing Sessions AVP
Session(s) removal: Withdraw Outgoing Sessions AVP


These new AVPs MUST be sent in an MSI message.

6.2.1. New Outgoing Sessions AVP (MSI)

The New Outgoing Sessions AVP can only be carried within an MSI message type. This AVP piggybacks every Session ID to which the multicast traffic has to be forwarded.

The AVP has the following format:

Vendor ID = 0
Attribute = 81 (16 bits)



There can be from 1 to N Session IDs present in the New Outgoing Sessions AVP (considering the maximum value of the Length field). This AVP must be placed in an MSI message and sent after the establishment of the multicast session to indicate the initial outgoing sessions to the LAC, and must be sent at any time when one or more outgoing sessions appear during the multicast session lifetime. Upon receipt of this AVP, the LAC sends a New Outgoing Sessions Acknowledgment AVP to the LNS to notify that the LAC is ready to replicate the multicast traffic toward the indicated sessions.

Usage of this AVP is incremental; only new outgoing sessions have to be listed in the AVP.


The M-bit MUST be set to 1, and the AVP MAY be hidden (H-bit set to 0 or 1).

6.2.2. New Outgoing Sessions Acknowledgement AVP (MSI)

The New Outgoing Sessions Acknowledgement AVP can only be carried within an MSI message type. This AVP informs the LNS that the LAC is ready to replicate traffic for every Session ID listed in the AVP.

The AVP has the following format:

Vendor ID = 0
Attribute = 82 (16 bits)



This AVP must be placed in an MSI message and sent by the LAC toward the LNS to acknowledge the receipt of a New Outgoing Sessions list received in a New Outgoing Sessions AVP from the LNS.

An LNS is allowed to send multicast traffic within the L2TP multicast session as soon as a New Outgoing Sessions Acknowledgement AVP is received for the corresponding L2TP multicast session.

An LNS is allowed to stop sending packets of the corresponding multicast flow within L2TP unicast sessions only if it receives an MSI message with the New Outgoing Session Acknowledgement AVP, and only for the unicast Session IDs mentioned in this AVP. The multicast traffic can then be conveyed in L2TP unicast sessions when the L2TP multicast session goes down. From this standpoint, packets related to this multicast flow SHOULD NOT be conveyed within the L2TP unicast sessions mentioned in the AVP in order to avoid the duplication of multicast packets.

There can be from 1 to N Session IDs present in the New Outgoing Sessions Acknowledgement AVP (considering the maximum value of the Length field). Session IDs mentioned in this AVP that have not been listed in a previous New Outgoing Sessions AVP should be ignored. Non-acknowledged Session IDs MAY be listed in forthcoming New Outgoing Sessions AVPs, but multicast traffic MUST be sent to logical interfaces associated to these Session IDs as long as these Session IDs are not acknowledged for replication by the LAC.

The M-bit MUST be set to 1, and the AVP MAY be hidden (H-bit set to 0 or 1).


6.2.3. Withdraw Outgoing Sessions AVP (MSI)

The Withdraw Outgoing Sessions AVP is sent whenever there is one or more withdrawn subscriptions for the corresponding multicast flow (designated by the session ID on which the MSI is sent).

The LAC can stop forwarding packets to Session IDs mentioned in the AVP for the corresponding multicast flow as soon as it receives the MSI message embedding this Withdraw Target Session AVP.

The AVP has the following format:

Vendor ID = 0
Attribute = 83 (16 bits)



There can be from 1 to N Session IDs present in the Withdraw Outgoing Sessions AVP (considering the value of the Length field). The M-bit MUST be set to 1, and the AVP MAY be hidden (H-bit set to 0 or 1).

6.3. Multicast Packets Priority AVP (MSI)

The Multicast Packets Priority AVP is an optional AVP intended to indicate to the LAC how to process multicast traffic against unicast traffic. Even though the LAC behavior is partially described here, the nature of the traffic (layer-2 frames for unicast traffic and pure IP packets for multicast traffic) is not a criteria for enforcing a traffic prioritization policy. Traffic processing for the provisioning of a uniformly framed traffic for the final user is described in section 8.

Three different behaviors can be adopted:

- 1) Best effort: the traffic is forwarded from the LAC to the end-user in the order in which it comes from the LNS, whatever the type of traffic.


- 2) Unicast traffic priority: traffic coming down the L2TP unicast session has priority over traffic coming down the L2TP multicast session.
- 3) Multicast traffic priority: traffic coming down the L2TP multicast session has priority over traffic coming down the L2TP unicast session.

The priority is encoded as a 16-bit quantity, which can take the following values:

- 0: Best effort (default)
- 1: Unicast traffic priority
- 2: Multicast traffic priority

The AVP has the following format:

Vendor ID = 0
 Attribute = 84 (16 bits)



Note that the multicast traffic rate can reach up to Maximum BPS (as indicated in MSRQ). This rate can exceed the maximum rate allowed for a particular end-user. This means that even with a priority value of 0, the end-user may receive multicast traffic only; unicast packets might be dropped because the multicast flow overwhelms the LAC forwarding buffer(s).

The default Priority Value is 0. The M-bit MUST be set to 0, and the AVP MAY be hidden (H-bit set to 0 or 1).

There are two ways of using this AVP: global configuration and individual configuration.

6.3.1. Global Configuration

The Multicast Priority Packet AVP is sent for all L2TP unicast sessions concerned with a specific multicast flow represented by an L2TP multicast session. In this case, the AVP is sent in an L2TP MSI control message for the corresponding multicast session ID (Session ID = L2TP session for the corresponding multicast group). The

priority value applies to all L2TP unicast sessions to which the multicast group designated by the L2TP multicast session is intended, as soon as this AVP is received.

6.3.2. Individual Configuration

The Multicast Priority Packet AVP is sent for a specific L2TP unicast session that SHALL adopt a specific behavior for both unicast and multicast traffics. In this case, the AVP is sent in an L2TP MSI control message for the L2TP unicast session (Session ID = L2TP session for the concerned user). The priority value applies to the targeted session only and does not affect the other sessions. Note that in this case, all multicast packets carried in L2TP multicast sessions are treated the same way by the LAC for the concerned user.

This is the only case in which an MSI control message can be sent for an L2TP unicast session.

6.3.3. Priority

It is the responsibility of the network administrator to decide which behavior to adopt between global or individual configurations, if the AVP is sent twice (one for a multicast group and one for a specific end-user). By default, only the individual configurations SHOULD be taken into consideration in that case.

Support of the Multicast Packets Priority AVP is optional and SHOULD be configurable by the LAC administrator, if it is relevant.

7. Multicast Session Teardown

An L2TP multicast session should be torn down whenever there are no longer any users interested in receiving the corresponding multicast traffic. A multicast session becomes useless once the related OSL has fewer than a predefined number of entries, this number being defined by a threshold.

Multicast session flapping may occur when the number of OSL entries oscillates around the threshold, if the same value is used to trigger the creation or deletion of an L2TP multicast session. To avoid this behavior, two methods can be used:

- The threshold value that is used to determine whether the L2TP multicast session has to be torn down is lower than the MULTICAST_SESSION_THRESHOLD value;

- The MULTICAST_SESSION_THRESHOLD value is used to determine whether the L2TP multicast session has to be torn down. A multicast session SHOULD be killed after a period of MULTICAST_SESSION_HOLDTIME seconds if the corresponding OSL maintains fewer than a MULTICAST_SESSION_THRESHOLD number of entries. The MULTICAST_SESSION_HOLDTIME value is 10 seconds by default and SHOULD be configurable by either the LAC or the LNS administrator.

The multicast session can be torn down for multiple reasons, including specific criteria not described here (which can be vendor specific).

A multicast session teardown can be initiated by either the LAC or the LNS. However, multicast session teardown MUST be initiated by the LNS if the termination decision is motivated by the number of users interested in receiving the traffic corresponding to a multicast flow.

7.1. Operations

The actual termination of a multicast session is initiated with a new Multicast-Session-End-Notify (MSEN) control message, sent either by the LAC or by the LNS.

The following is an example of a control message exchange that terminates a multicast session:

```

LAC or LNS      LAC or LNS
-----
                  (multicast session
                  termination)

                  <- MSEN
                  (Clean up)

ZLB ACK ->
(Clean up)

```

7.2. Multicast-Session-End-Notify (MSEN)

The Multicast-Session-End-Notify (MSEN) is an L2TP control message sent by either the LAC or the LNS to request the termination of a specific multicast session within the tunnel. Its purpose is to give the peer the relevant termination information, including the reason why the termination occurred. The peer MUST clean up any associated resources and does not acknowledge the MSEN message.

As defined in [RFC2661], termination of a control connection will terminate all sessions managed within, including multicast sessions if there are any.

The MSEN message carries a Result Code AVP with an optional Error Code.

The following AVPs MUST be present in an MSEN message:

Message Type
Result Code
Assigned Session ID

The associated Message Type AVP is encoded with the following values:

Vendor ID = 0
Attribute Type = 0
Attribute Value = 27 (16 bits)

The M-bit MUST be set to 0, and the H-bit MUST be set to 0.

7.3. Result Codes

The following values are the defined result codes for MSEN control messages:

- 1 (16 bits) - No multicast traffic for the group
 - 2 (16 bits) - Session terminated for the reason indicated in the error code
 - 3 (16 bits) - No more receivers
 - 4 (16 bits) - No more receivers (filter-mode change)
- o The code 1 MAY be used when the LAC detects that no traffic is coming down the multicast session, or when the LNS doesn't receive multicast traffic to be conveyed over the L2TP multicast session during a certain period of time.
 - o The code 2 refers to General Error Codes maintained by the IANA for L2TP.
 - o The code 3 MAY be used by the LAC or the LNS when the OSL is empty.
 - o The code 4 MAY be used by the LNS when a multicast session is torn down because of a filter-mode change. This result code SHOULD also be used when the OSL becomes empty after a filter-mode change (passive termination when filter-mode changes from INCLUDE to EXCLUDE; see Section 4.3).

8. Traffic Merging

Both unicast and multicast traffics have to be merged by the LAC in order to forward properly framed data to the end-user. Multicast packets are framed by the LAC and transmitted toward the proper end-user. Methods used to achieve this function are not described here, since it is an implementation-specific issue.

All frames conveyed from the LAC to the end-users have to follow the framing scheme applied for the considered peer to which the traffic is destined (e.g., the LAC is always aware of the PPP [RFC1661] link parameters, as described in [RFC2661], Section 6.14). Note that using L2TP Multicast Extension features is not appropriate for end-users who have negotiated a sequenced layer-2 connection with the LNS. While inserting PPP-encapsulated multicast packets in a session, the LAC cannot modify PPP sequencing performed by the LNS for each PPP session.

9. IANA Considerations

This document defines:

- 5 new Message Type (Attribute Type 0) Values:
 - o Multicast-Session-Request (MSRQ) : 23
 - o Multicast-Session-Response (MSRP) : 24
 - o Multicast-Session-Establishment (MSE) : 25
 - o Multicast-Session-Information (MSI) : 26
 - o Multicast-Session-End-Notify (MSEN) : 27

- 5 new Control Message Attribute Value Pairs:
 - o Multicast Capability : 80
 - o New Outgoing Sessions : 81
 - o New Outgoing Sessions Acknowledgement : 82
 - o Withdraw Outgoing Sessions : 83
 - o Multicast Packets Priority : 84

- 4 Result Codes for the MSEN message:
 - o No multicast traffic for the group : 1
 - o Session terminated for the reason indicated in the error code : 2
 - o No more receivers : 3
 - o No more receivers (filter-mode change): 4

10. Security Considerations

It is possible for one receiver to make additional multicast traffic that has not been requested go down the link of another receiver. This can happen if a single replication context per group is used in INCLUDE mode with receivers having divergent source lists, and in EXCLUDE mode if a receiver has a source list not shared by another. This behavior can be encountered every time receivers are connected to a common multi-access network.

The extension described in this document does not introduce any additional security issues as far as the activation of the L2TP protocol is concerned.

Injecting appropriate control packets in the tunnel toward the LAC to modify Outgoing Session List and to flood end-users with unwanted multicast traffic is only possible if the control connection is hacked. As for any reception of illegitimate L2TP control messages, the following apply:

- If the spoofed control message embeds consistent sequence numbers, next messages will appear out of synch, yielding the control connection to terminate.
- If sequence numbers are inconsistent with current control connection states, the spoofed control message will be queued or discarded, as described in [RFC2661], Section 5.8.

The activation of the L2TP multicast capability on the LAC could make the equipment more sensitive to Denial of Service attacks if the control connection or the related LNS is hacked. The LAC might also be sensitive to the burden generated by the additional replication work.

As mentioned in [RFC2661], Section 9.2, securing L2TP requires that the underlying transport make encryption, integrity, and authentication services available for all L2TP traffic, including L2TP multicast traffic (control and data).

11. References

11.1. Normative References

- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.
- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC3438] Townsley, W., "Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers Authority (IANA) Considerations Update", BCP 68, RFC 3438, December 2002.
- [RFC3590] Haberman, B., "Source Address Selection for the Multicast Listener Discovery (MLD) Protocol", RFC 3590, September 2003.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

11.2. Informative References

- [PROXY] Fenner, B., He, H., Haberman, B., Sandick, H., "IGMP/MLD-based Multicast Forwarding ("IGMP/MLD Proxying")", Work in Progress.

12. Acknowledgements

Thanks to Christian Jacquenet for all the corrections done on this document and his precious advice, to Pierre Levis for his contribution about IGMP, to Francis Houllier for PPP considerations, and to Xavier Vinet for his input about thresholds. Many thanks to W. Mark Townsley, Isidor Kouvelas, and Brian Haberman for their highly valuable input on protocol definition.

Appendix A. Examples of Group States Determination

***Example 1:**

All users are managed in the same control connection.

```
Users {1, 2, 3} subscribe to (Group G1, EXCLUDE {})
Users {3, 4, 5} subscribe to (Group G2, EXCLUDE {})
```

Group states for this L2TP tunnel will be:

```
(G1, EXCLUDE, {})
(G2, EXCLUDE, {})
```

Therefore, two replication contexts will be created:

```
-RC1:
(*, G1) packets, Multicast Session MS1, OSL = 1, 2, 3
-RC2:
(*, G2) packets, Multicast Session MS2, OSL = 3, 4, 5
```

***Example 2:**

All users are managed in the same control connection.

```
Users {1, 2, 3} subscribe to (Group G1, INCLUDE {S1})
Users {4, 5, 6} subscribe to (Group G1, INCLUDE {S1,S2})
Users {7, 8, 9} subscribe to (Group G1, INCLUDE {S2})
```

The group state for this L2TP tunnel will be:

```
(G1, INCLUDE, {S1, S2})
```

If the LNS policy allows one replication context per (group, source), two replication contexts will be created:

```
-RC1:
(S1, G1) packets, Multicast Session MS1, OSL = 1, 2, 3, 4, 5, 6
-RC2:
(S2, G1) packets, Multicast Session MS2, OSL = 4, 5, 6, 7, 8, 9
```

If the LNS policy allows one replication context per (group, source-list), one replication context will be created:

```
-RC1:
({S1, S2}, G1) packets, Multicast Session MS1, OSL = [1..9]
```

*Example 3:

All users are managed in the same control connection.

Users {1, 2} subscribe to (Group G1, EXCLUDE {S1})
User {3} subscribes to (Group G1, EXCLUDE {S1, S2})

The group state for this L2TP tunnel will be:

(G1, EXCLUDE, {S1})

Therefore, one replication context will be created:

-RC1:
(*-{S1}, G1) packets, Multicast Session MS1, OSL = 1, 2, 3

Next, user {4} subscribes to (Group G1, INCLUDE {S1}). The group state for the L2TP tunnel is changed to:

(G1, EXCLUDE, {})

The replication context RC1 is changed to:

-RC1: (*, G1) packets, Multicast Session MS1, OSL = 1, 2, 3, 4

*Example 4:

All users are managed in the same control connection. The LNS policy allows one replication context per (group, source).

Users {1, 2, 3} subscribe to (Group G1, INCLUDE {S1, S2})

The group state for this L2TP tunnel will be:

(G1, INCLUDE, {S1, S2})

Therefore, two replication contexts will be created:

-RC1:
(S1, G1) packets, Multicast Session MS1, OSL = 1, 2, 3
-RC2:
(S2, G1) packets, Multicast Session MS2, OSL = 1, 2, 3

Next, user {4} subscribes to (Group G1, EXCLUDE {}), equivalent to an IGMPv2 membership report. The group state for the L2TP tunnel is changed to:

```
(G1, EXCLUDE, {})
```

The replication context RC1 is changed to:

```
-RC1: (*, G1) packets, Multicast Session MS1, OSL = 1, 2, 3, 4
```

The replication context RC2 is changed to:

```
-RC2: no packets to forward, Multicast Session MS2, OSL = {}  
(Multicast Session MS2 will be deleted)
```

When user {4} leaves G1, the group state for the L2TP tunnel goes back to:

```
(G1, INCLUDE, {S1, S2})
```

Replication contexts become:

```
-RC1:  
(S1, G1) packets, Multicast Session MS1, OSL = 1, 2, 3  
-RC2:  
(S2, G1) packets, Multicast Session MS2, OSL = 1, 2, 3  
(Multicast Session MS2 is re-established)
```

Author's Address

Gilles Bourdon
France Telecom
38-40, rue du General Leclerc
92794 Issy les Moulineaux Cedex 9 - FRANCE

Phone: +33 1 4529-4645
EMail: gilles.bourdon@francetelecom.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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.

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.

Intellectual Property

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.

Acknowledgement

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

