

Network Working Group
Request for Comments: 4454
Category: Standards Track

S. Singh
M. Townsley
C. Pignataro
Cisco Systems
May 2006

Asynchronous Transfer Mode (ATM) over Layer 2 Tunneling Protocol Version 3 (L2TPv3)

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 (2006).

Abstract

The Layer 2 Tunneling Protocol, Version 3 (L2TPv3) defines an extensible tunneling protocol to transport layer 2 services over IP networks. This document describes the specifics of how to use the L2TP control plane for Asynchronous Transfer Mode (ATM) Pseudowires and provides guidelines for transporting various ATM services over an IP network.

Table of Contents

1. Introduction	2
1.1. Abbreviations	3
1.2. Specification of Requirements	3
2. Control Connection Establishment	3
3. Session Establishment and ATM Circuit Status Notification	4
3.1. L2TPv3 Session Establishment	4
3.2. L2TPv3 Session Teardown	6
3.3. L2TPv3 Session Maintenance	6
4. Encapsulation	6
4.1. ATM-Specific Sublayer	7
4.2. Sequencing	9
5. ATM Transport	9
5.1. ATM AAL5-SDU Mode	10
5.2. ATM Cell Mode	10

5.2.1. ATM VCC Cell Relay Service	11
5.2.2. ATM VPC Cell Relay Service	12
5.2.3. ATM Port Cell Relay Service	12
5.3. OAM Cell Support	12
5.3.1. VCC Switching	12
5.3.2. VPC Switching	13
6. ATM Maximum Concatenated Cells AVP	13
7. OAM Emulation Required AVP	14
8. ATM Defects Mapping and Status Notification	14
8.1. ATM Alarm Status AVP	14
9. Applicability Statement	15
9.1. ATM AAL5-SDU Mode	16
9.2. ATM Cell Relay Mode	18
10. Congestion Control	20
11. Security Considerations	21
12. IANA Considerations	21
12.1. L2-Specific Sublayer Type	21
12.2. Control Message Attribute Value Pairs (AVPs)	21
12.3. Result Code AVP Values	22
12.4. ATM Alarm Status AVP Values	22
12.5. ATM-Specific Sublayer Bits	23
13. Acknowledgements	23
14. References	23
14.1. Normative References	23
14.2. Informative References	24

1. Introduction

This document describes the specifics of how to use the Layer 2 Tunneling Protocol (L2TP) for Asynchronous Transfer Mode (ATM) Pseudowires, including encapsulation, carrying various ATM services, such as AAL5 SDU, ATM VCC/VPC/Port cell relay over L2TP, and mapping ATM defects to L2TP Set-Link-Info (SLI) messages to notify the peer L2TP Control Connection Endpoint (LCCE).

Any ATM-specific AVPs or other L2TP constructs for ATM Pseudowire (ATMPW) support are defined here as well. Support for ATM Switched Virtual Path/Connection (SVP/SVC) and Soft Permanent Virtual Path/Connection (SPVP/SPVC) are outside the scope of this document.

The reader is expected to be very familiar with the terminology and protocol constructs defined in [RFC3931].

1.1. Abbreviations

AIS	Alarm Indication Signal
ATMPW	ATM Pseudowire
AVP	Attribute Value Pair
CC	Continuity Check OAM Cell
CE	Customer Edge
HEC	Header Error Checksum
LAC	L2TP Access Concentrator (see [RFC3931])
LCCE	L2TP Control Connection Endpoint (see [RFC3931])
MSB	Most Significant Byte
OAM	Operation, Administration, and Maintenance
PE	Provider Edge
PSN	Packet Switched Network
PWE3	Pseudowire Emulation Edge to Edge
RDI	Remote Defect Indicator
SAR	Segmentation and Reassembly
SDU	Service Data Unit
SLI	Set-Link-Info, an L2TP control message
SVC	Switched Virtual Connection
SVP	Switched Virtual Path
SPVC	Soft Permanent Virtual Connection
SPVP	Soft Permanent Virtual Path
VC	Virtual Circuit
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VPC	Virtual Path Connection
VPI	Virtual Path Identifier

1.2. Specification of Requirements

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. 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].

2. Control Connection Establishment

To emulate ATM Pseudowires using L2TP, an L2TP Control Connection as described in Section 3.3 of [RFC3931] MUST be established.

The Start-Control-Connection-Request (SCCRQ) and corresponding Start-Control-Connection-Reply (SCCRP) MUST include the supported ATM Pseudowire types (see Section 3.1), in the Pseudowire Capabilities List as defined in Section 5.4.3 of [RFC3931]. This identifies the Control Connection as able to establish L2TP sessions in support of the ATM Pseudowires.

An LCCE MUST be able to uniquely identify itself in the SCCRQ and SCCRP messages via a globally unique value. By default, this is advertised via the structured Router ID AVP [RFC3931], though the unstructured Hostname AVP [RFC3931] MAY be used to identify LCCEs as well.

3. Session Establishment and ATM Circuit Status Notification

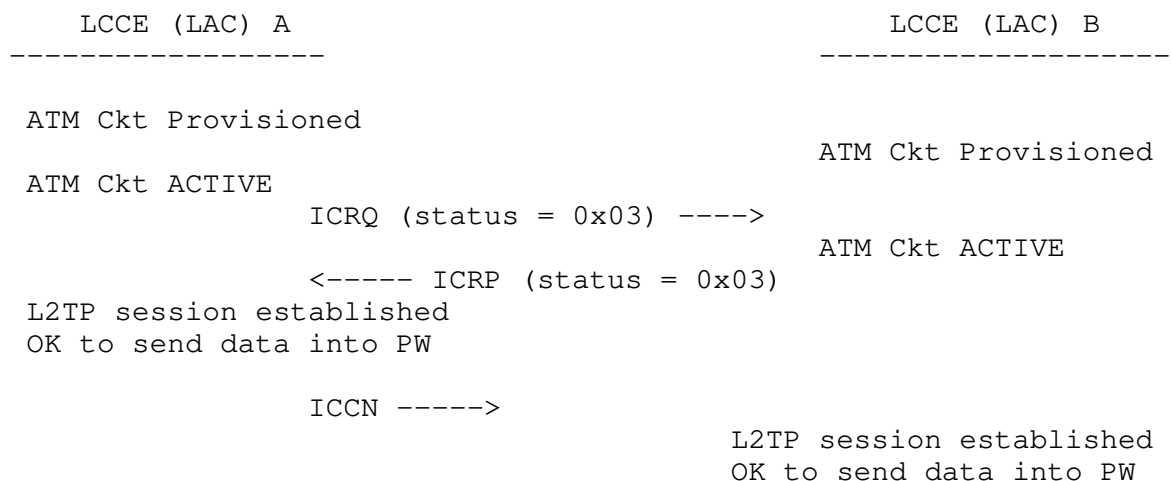
This section describes how L2TP ATMPWs or sessions are established between two LCCEs. This includes what will happen when an ATM circuit (e.g., AAL5 PVC) is created, deleted, or changes state when circuit state is in alarm.

3.1. L2TPv3 Session Establishment

ATM circuit (e.g., an AAL5 PVC) creation triggers establishment of an L2TP session using three-way handshake described in Section 3.4.1 of [RFC3931]. An LCCE MAY initiate the session immediately upon ATM circuit creation, or wait until the circuit state transitions to ACTIVE before attempting to establish a session for the ATM circuit. It MAY be preferred to wait until circuit status transitions to ACTIVE in order to delay the allocation of resources until absolutely necessary.

The Circuit Status AVP (see Section 8) MUST be present in the Incoming-Call-Request (ICRQ) and Incoming-Call-Reply (ICRP) messages, and MAY be present in the SLI message for ATMPWs.

The following figure shows how L2TP messages are exchanged to set up an ATMPW after the ATM circuit (e.g., an AAL5 PVC) becomes ACTIVE.



The following signaling elements are required for the ATMPW establishment.

- a. Pseudowire Type: One of the supported ATM-related PW types should be present in the Pseudowire Type AVP of [RFC3931].

```
0x0002  ATM AAL5 SDU VCC transport
0x0003  ATM Cell transport Port Mode
0x0009  ATM Cell transport VCC Mode
0x000A  ATM Cell transport VPC Mode
```

The above cell relay modes can also signal the ATM Maximum Concatenated Cells AVP as described in Section 6.

- b. Remote End ID: Each PW is associated with a Remote End ID akin to the VC-ID in [PWE3ATM]. Two LCCEs of a PW would have the same Remote End ID, and its format is described in Section 5.4.4 of [RFC3931].

This Remote End ID AVP MUST be present in the ICRQ in order for the remote LCCE to associate the session to the ATM circuit. The Remote End Identifier AVP defined in [RFC3931] is of opaque form, though ATMPW implementations MAY simply use a 4-octet value that is known to both LCCEs (either by direct configuration or some other means). The exact method of how this value is configured, retrieved, discovered, or otherwise determined at each LCCE is outside the scope of this document.

As with the ICRQ, the ICRP is sent only after the ATM circuit transitions to ACTIVE. If LCCE B had not been provisioned yet for the ATM circuit identified in the ICRQ, a Call-Disconnect-Notify (CDN) would have been immediately returned indicating that the circuit either was not provisioned or was not available at this LCCE. LCCE A SHOULD then exhibit a periodic retry mechanism. If so, the period and maximum number of retries MUST be configurable.

An implementation MAY send an ICRQ or ICRP before a PVC is ACTIVE, as long as the Circuit Status AVP reflects that the ATM circuit is INACTIVE and an SLI is sent when the ATM circuit becomes ACTIVE (see Section 8).

The ICCN is the final stage in the session establishment. It confirms the receipt of the ICRP with acceptable parameters to allow bidirectional traffic.

3.2. L2TPv3 Session Teardown

When an ATM circuit is unprovisioned (deleted) at either LCCE, the associated L2TP session MUST be torn down via the CDN message defined in Section 3.4.3 of [RFC3931].

3.3. L2TPv3 Session Maintenance

All sessions established by a given Control Connection utilize the L2TP Hello facility defined in Section 4.4 of [RFC3931] for session keepalive. This gives all sessions basic dead peer and path detection between LCCEs.

If the control channel utilizing the Hello message is not in-band with data traffic over the PSN, then other method MAY be used to detect the session failure, and it is left for further study.

ATMPWs over L2TP use the Set-Link-Info (SLI) control message as defined in [RFC3931] to signal ATM circuit status between LCCEs after initial session establishment. This includes ACTIVE or INACTIVE notifications of the ATM circuit, or any other parameters that may need to be shared between the LCCEs in order to provide proper PW emulation.

The SLI message MUST be sent whenever there is a status change that may be reported by any values identified in the Circuit Status AVP. The only exceptions to this are the initial ICRQ, ICRP, and CDN messages, which establish and tear down the L2TP session itself when the ATM circuit is created or deleted. The SLI message may be sent from either LCCE at any time after the first ICRQ is sent (and perhaps before an ICRP is received, requiring the peer to perform a reverse Session ID lookup).

The other application of the SLI message is to map the ATM OAM or physical layer alarms into Circuit Status AVP as described in Section 8.

4. Encapsulation

This section describes the general encapsulation format for ATM services over L2TP.

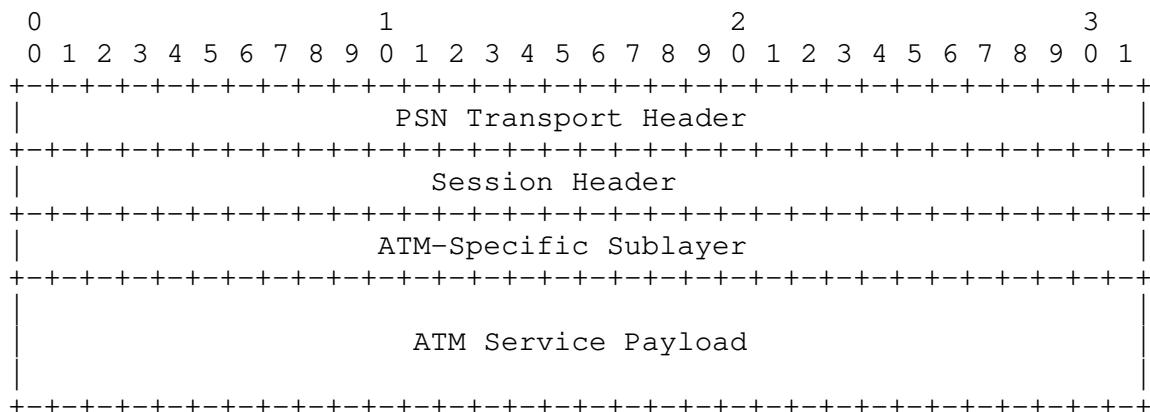


Figure 1: General Format for ATM Encapsulation over L2TPv3 over IP

The PSN Transport header is specific to IP and its underlying transport header. This header is used to transport the encapsulated ATM payload through the IP network.

The Session Header is a non-zero 32-bit Session ID with an optional Cookie up to 64-bits. This Session ID is exchanged during session setup.

The ATM-Specific Sublayer is REQUIRED for AAL5 SDU Mode and OPTIONAL for ATM Cell Mode. Please refer to Section 4.1 for more details.

4.1. ATM-Specific Sublayer

This section defines a new ATM-Specific Sublayer, an alternative to the Default L2-Specific Sublayer as mentioned in Section 4.6 of [RFC3931]. Four new flag bits (T, G, C, and U) are defined that concur with Section 8.2 of [PWE3ATM].




Figure 2: ATM-Specific Sublayer Format

The meaning of the fields of the ATM-Specific Sublayer is as follows:

* S bit

Definition of this bit is as per Section 4.6 of [RFC3931].

* B and E bits

Definitions of these bits are as per Section 5.5 of [L2TPFRAG].

If these bits are not used as per [L2TPFRAG], they MUST be set to 0 upon transmission and ignored upon reception.

* T (Transport type) bit

Bit (T) of the ATM-Specific Sublayer indicates whether the packet contains an ATM admin cell or an AAL5 payload. If T = 1, the packet contains an ATM admin cell, encapsulated according to the VCC cell relay encapsulation of Section 5.2.

If not set, the PDU contains an AAL5 payload. The ability to transport an ATM cell in the AAL5 SDU Mode is intended to provide a means of enabling administrative functionality over the AAL5 VCC (though it does not endeavor to preserve user-cell and admin-cell arrival/transport ordering, as described in Section 9.1).

* G (EFCI) Bit

The ingress LCCE device SHOULD set this bit to 1 if the Explicit Forward Congestion Indication (EFCI) bit of the final cell of the incoming AAL5 payload is set to 1, or if the EFCI bit of the single ATM cell to be transported in the packet is set to 1. Otherwise, this bit SHOULD be set to 0. The egress LCCE device SHOULD set the EFCI bit of all the outgoing cells that transport the AAL5 payload to the value contained in this field.

* C (CLP) Bit

The ingress LCCE device SHOULD set this bit to 1 if the Cell Loss Priority (CLP) bit of any of the incoming ATM cells of the AAL5 payload is set to 1, or if the CLP bit of the single ATM cell that is to be transported in the packet is set to 1. Otherwise this bit SHOULD be set to 0. The egress LCCE device SHOULD set the CLP bit of all outgoing cells that transport the AAL5 CPCS-PDU to the value contained in this field.

* U (Command/Response) Bit

When FRF.8.1 Frame Relay / ATM PVC Service Interworking (see [FRF8.1]) traffic is being transported, the CPCS-UU Least Significant Bit (LSB) of the AAL5 CPCS-PDU may contain the Frame Relay C/R bit. The ingress LCCE device SHOULD copy this bit to the U bit of the ATM-Specific Sublayer. The egress LCCE device SHOULD copy the U bit to the CPCS-UU Least Significant Bit (LSB) of the AAL5 payload.

The Sequence Number field is used in sequencing, as described in Section 4.2.

In case of a reassembly timeout, the encapsulating LCCE should discard all component cells of the AAL5 frame.

An additional enumeration is added to the L2-Specific Sublayer AVP to identify the ATM-Specific Sublayer:

- 0 - There is no L2-Specific Sublayer present.
- 1 - The Default L2-Specific Sublayer (defined in Section 4.6 of [RFC3931]) is used.
- 2 - The ATM-Specific Sublayer is used.

The first two values are already defined in the L2TPv3 base specification [RFC3931].

4.2. Sequencing

Data Packet Sequencing MAY be enabled for ATMPWs. The sequencing mechanisms described in [RFC3931] MUST be used to signal sequencing support. ATMPWs over L2TPv3 MUST request the presence of the ATM-Specific Sublayer when sequencing is enabled, and MAY request its presence at all times.

5. ATM Transport

There are two encapsulations supported for ATM transport as described below.

The ATM-Specific Sublayer is prepended to the AAL5-SDU. The other cell mode encapsulation consists of the OPTIONAL ATM-Specific Sublayer, followed by a 4-byte ATM cell header and a 48-byte ATM cell-payload.

5.1. ATM AAL5-SDU Mode

In this mode, each AAL5 VC is mapped to an L2TP session. The Ingress LCCE reassembles the AAL5 CPCS-SDU without the AAL5 trailer and any padding bytes. Incoming EFCI, CLP, and C/R (if present) are carried in an ATM-Specific Sublayer across ATPWs to the egress LCCE. The processing of these bits on ingress and egress LCCEs is defined in Section 4.1.

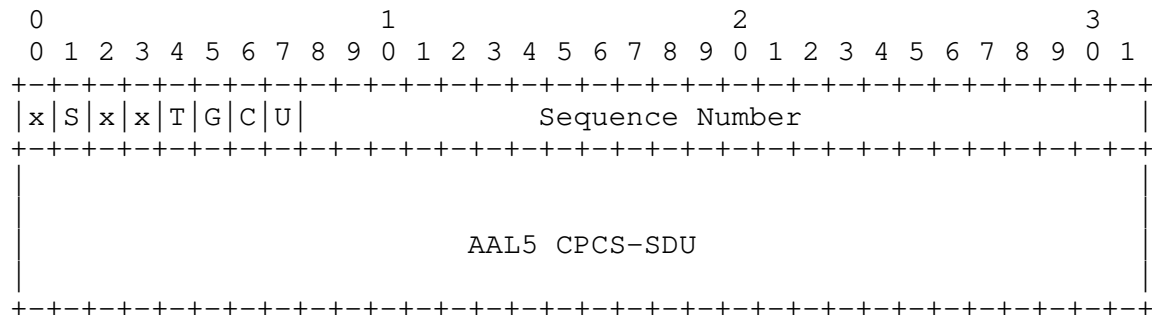


Figure 3: ATM AAL5-SDU Mode Encapsulation

If the ingress LCCE determines that an encapsulated AAL5 SDU exceeds the MTU size of the L2TPv3 session, then AAL5 SDU may be fragmented as per [L2TPFRAG] or underneath the transport layer (IP, etc.). F5 OAM cells that arrive during the reassembly of an AAL5 SDU are sent immediately on the PW followed by the AAL5 SDU payload. In this case, OAM cells' relative order with respect to user data cells is not maintained.

Performance Monitoring OAM, as specified in ITU-T 610 [I610-1], [I610-2], [I610-3] and security OAM cells as specified in [ATMSEC], should not be used in combination with AAL5 SDU Mode. These cells MAY be dropped at the ingress LCCE because cell sequence integrity is not maintained.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the ICRQ messages and MUST include the ATM AAL5 SDU VCC transport PW Type of 0x0002.

5.2. ATM Cell Mode

In this mode, ATM cells skip the reassembly process at the ingress LCCE. These cells are transported over an L2TP session, either as a single cell or as concatenated cells, into a single packet. Each ATM cell consists of a 4-byte ATM cell header and a 48-byte ATM cell-payload; the HEC is not included.

In ATM Cell Mode encapsulation, the ATM-Specific Sublayer is OPTIONAL. It can be included, if sequencing support is required. It is left to the implementation to choose to signal the Default L2-Specific Sublayer or the ATM-Specific Sublayer.

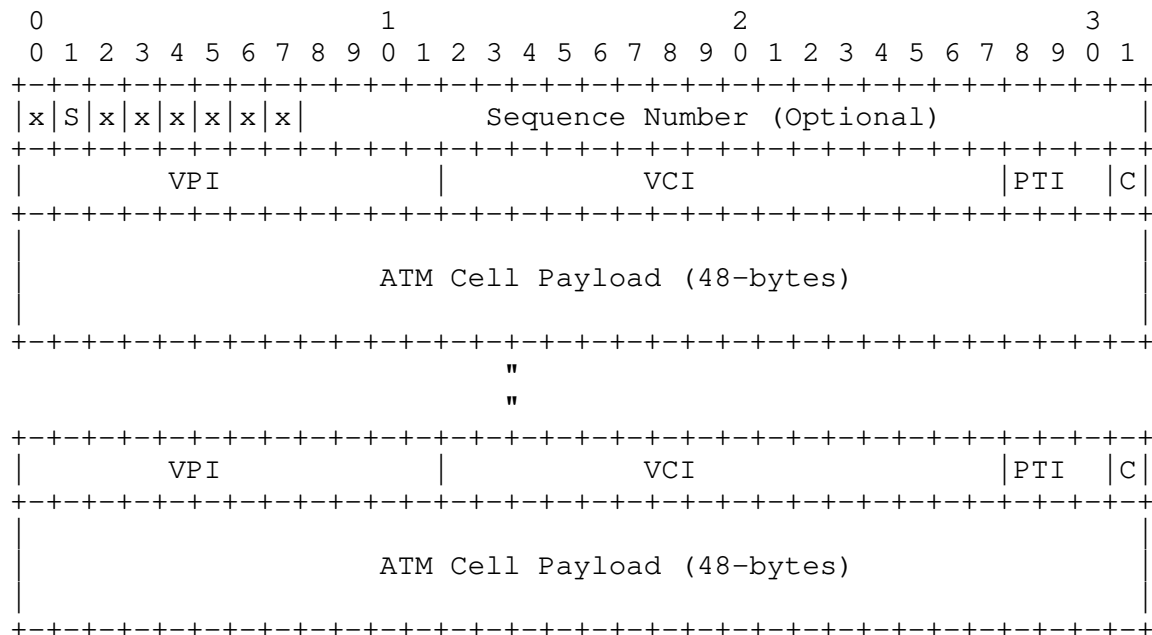


Figure 4: ATM Cell Mode Encapsulation

In the simplest case, this encapsulation can be used to transmit a single ATM cell per Pseudowire PDU. However, in order to provide better Pseudowire bandwidth efficiency, several ATM cells may be optionally encapsulated into a single Pseudowire PDU.

The maximum number of concatenated cells in a packet is limited by the MTU size of the session and also by the ability of the egress LCCE to process them. For more details about ATM Maximum Concatenated Cells, please refer to Section 6.

5.2.1. ATM VCC Cell Relay Service

A VCC cell relay service may be provided by mapping an ATM Virtual Channel Connection to a single Pseudowire using cell mode encapsulation as defined in Section 5.2.

An LCCE may map one or more VCCs to a single PW. However, a service provider may wish to provision a single VCC to a PW in order to satisfy QoS or restoration requirements.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the ICRQ messages and MUST include the ATM cell transport VCC Mode PW Type of 0x0009.

5.2.2. ATM VPC Cell Relay Service

A Virtual Path Connection cell relay service may be provided by mapping an ATM Virtual Path Connection to a single Pseudowire using cell mode encapsulation as defined in Section 5.2.

An LCCE may map one or more VPCs to a single Pseudowire.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the ICRQ messages and MUST include the ATM cell transport VPC Mode PW Type of 0x000A.

5.2.3. ATM Port Cell Relay Service

ATM port cell relay service allows an ATM port to be connected to another ATM port. All ATM cells that are received at the ingress ATM port on the LCCE are encapsulated as per Section 5.2, into Pseudowire PDU and sent to peer LCCE.

Each LCCE MUST discard any idle/unassigned cells received on an ATM port associated with ATMPWs.

The Pseudowire Type AVP defined in Section 5.4.4 of [RFC3931], Attribute Type 68, MUST be present in the ICRQ messages and MUST include the ATM Cell transport Port Mode PW Type of 0x0003.

5.3. OAM Cell Support

The OAM cells are defined in [I610-1], [I610-2], [I610-3] and [ATMSEC] can be categorized as follows:

- a. Fault Management
- b. Performance monitoring and reporting
- c. Activation/deactivation
- d. System Management (e.g., security OAM cells)

OAM Cells are always encapsulated using cell mode encapsulation, regardless of the encapsulation format used for user data.

5.3.1. VCC Switching

The LCCEs SHOULD be able to pass the F5 segment and end-to-end Fault Management, Resource Management (RM cells), Performance Management, Activation/deactivation, and System Management OAM cells.

F4 OAM cells are inserted or extracted at the VP link termination. These OAM cells are not seen at the VC link termination and are therefore not sent across the PW.

5.3.2. VPC Switching

The LCCEs MUST be able to pass the F4 segment and end-to-end Fault Management, Resource Management (RM cells), Performance Management, Activation/deactivation, and System Management OAM cells transparently according to [I610-1].

F5 OAM cells are not inserted or extracted at the VP cross-connect. The LCCEs MUST be able to pass the F5 OAM cells transparently across the PW.

6. ATM Maximum Concatenated Cells AVP

The "ATM Maximum Concatenated Cells AVP", Attribute Type 86, indicates that the egress LCCE node can process a single PDU with concatenated cells up to a specified number of cells. An LCCE node transmitting concatenated cells on this PW MUST NOT exceed the maximum number of cells as specified in this AVP. This AVP is applicable only to ATM Cell Relay PW Types (VCC, VPC, Port Cell Relay). This Attribute value may not be same in both directions of the specific PW.

The Attribute Value field for this AVP has the following format:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+--+--+--+--+--+--+--+--+--+--+--+
| ATM Maximum Concatenated Cells |
+--+--+--+--+--+--+--+--+--+--+--+

```

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP MAY be set to 0, but MAY vary (see Section 5.2 of [RFC3931]). The length (before hiding) of this AVP is 8.

This AVP is sent in an ICRQ, ICRP during session negotiation or via SLI control messages when LCCE changes the maximum number of concatenated cells configuration for a given ATM cell relay circuit.

This AVP is OPTIONAL. If the egress LCCE is configured with a maximum number of cells to be concatenated by the ingress LCCE, it SHOULD signal this value to the ingress LCCE.

7. OAM Emulation Required AVP

An "OAM Emulation Required AVP", Attribute Type 87, MAY be needed to signal OAM emulation in AAL5 SDU Mode, if an LCCE cannot support the transport of OAM cells across L2TP sessions. If OAM cell emulation is configured or detected via some other means on one side, the other LCCE MUST support OAM cell emulation as well.

This AVP is exchanged during session negotiation (in ICRQ and ICRP) or during the life of the session via SLI control messages. If the other LCCE cannot support the OAM cell emulation, the associated L2TP session MUST be torn down via CDN message with result code 22.

OAM Emulation AVP is a boolean AVP, having no Attribute Value. Its absence is FALSE and its presence is TRUE. This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see Section 5.2 of [RFC3931]). The Length (before hiding) of this AVP is 6.

8. ATM Defects Mapping and Status Notification

ATM OAM alarms or circuit status is indicated via the Circuit Status AVP as defined in Section 5.4.5 of [RFC3931]. For reference, usage of this AVP is shown below.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|           Reserved           |N|A|
+---+---+---+---+---+---+---+---+---+

```

The Value is a 16-bit mask with the two least significant bits defined, and the remaining bits are reserved for future use. Reserved bits MUST be set to 0 when sending and ignored upon receipt.

The A (Active) bit indicates whether the ATM circuit is ACTIVE (1) or INACTIVE (0).

The N (New) bit indicates whether the ATM circuit status indication is for a new ATM circuit (1) or an existing ATM circuit (0).

8.1. ATM Alarm Status AVP

An "ATM Alarm Status AVP", Attribute Type 88, indicates the reason for the ATM circuit status and specific alarm type, if any, to its peer LCCE node. This OPTIONAL AVP MAY be present in the SLI message with the Circuit Status AVP.

The Attribute Value field for this AVP has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Circuit Status Reason           |           Alarm           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Circuit Status Reason is a 2-octet unsigned integer, and the Alarm Type is also a 2-octet unsigned integer.

This AVP MAY be hidden (the H bit MAY be 0 or 1). The M bit for this AVP SHOULD be set to 0, but MAY vary (see Section 5.2 of [RFC3931]). The Length (before hiding) of this AVP is 10 octets.

This AVP is sent in the SLI message to indicate additional information about the ATM circuit status.

Circuit Status Reason values for the SLI message are as follows:

- 0 - Reserved
- 1 - No alarm or alarm cleared (default for Active Status)
- 2 - Unspecified or unknown Alarm Received (default for Inactive Status)
- 3 - ATM Circuit received F1 Alarm on ingress LCCE
- 4 - ATM Circuit received F2 Alarm on ingress LCCE
- 5 - ATM Circuit received F3 Alarm on ingress LCCE
- 6 - ATM Circuit received F4 Alarm on ingress LCCE
- 7 - ATM Circuit received F5 Alarm on ingress LCCE
- 8 - ATM Circuit down due to ATM Port shutdown on Peer LCCE
- 9 - ATM Circuit down due to loop-back timeout on ingress LCCE

The general ATM Alarm failures are encoded as below:

- 0 - Reserved
- 1 - No Alarm type specified (default)
- 2 - Alarm Indication Signal (AIS)
- 3 - Remote Defect Indicator (RDI)
- 4 - Loss of Signal (LOS)
- 5 - Loss of Pointer (LOP)
- 6 - Loss of Framer (LOF)
- 7 - Loopback cells (LB)
- 8 - Continuity Check (CC)

9. Applicability Statement

The ATM Pseudowire emulation described in this document allows for carrying various ATM services across an IP packet switched network

(PSN). These ATM services can be PVC-based, PVP-based, or port-based. In all cases, ATMPWs operate in a point-to-point deployment model.

ATMPWs support two modes of encapsulation: ATM AAL5-SDU Mode and ATM Cell Relay Mode. The following sections list their respective characteristics in relationship to the native service.

9.1. ATM AAL5-SDU Mode

ATMPWs operating in AAL5-SDU Mode only support the transport of PVC-based services. In this mode, the AAL5 CPCS-PDU from a single VCC is reassembled at the ingress LCCE, and the AAL5 CPCS-SDU (i.e., the AAL5 CPCS-PDU without CPCS-PDU Trailer or PAD octets, also referred to as AAL5 CPCS-PDU Payload) is transported over the Pseudowire. Therefore, Segmentation and Reassembly (SAR) functions are required at the LCCEs. There is a one-to-one mapping between an ATM PVC and an ATMPW operating in AAL5-SDU Mode, supporting bidirectional transport of variable length frames. With the exception of optionally transporting OAM cells, only ATM Adaptation Layer (AAL) Type 5 frames are carried in this mode, including multiprotocol over AAL5 packets [RFC2684].

The following considerations stem from ATM AAL5-SDU Mode Pseudowires not transporting the ATM cell headers and AAL5 CPCS-PDU Trailer (see Section 5.1):

- o An ATMPW operating in AAL5-SDU Mode conveys EFCI and CLP information using the G and C bits in the ATM-Specific Sublayer. In consequence, the EFCI and CLP values of individual ATM cells that constitute the AAL5 frame may be lost across the ATMPW, and CLP and EFCI transparency may not be maintained. The AAL5-SDU Mode does not preserve EFCI and CLP values for every ATM cell within the AAL5 PDU. The processing of these bits on ingress and egress is defined in Section 4.1.
- o Only the least significant bit (LSB) from the CPCS-UU (User-to-User indication) field in the CPCS-PDU Trailer is transported using the ATM-Specific Sublayer (see Section 4.1). This bit contains the Frame Relay C/R bit when FRF.8.1 Frame Relay / ATM PVC Service Interworking [FRF8.1] is used. The CPCS-UU field is not used in multiprotocol over AAL5 [RFC2684]. However, applications that transfer user to user information using the CPCS-UU octet would fail to operate.

- o The CPI (Common Part Indicator) field in the CPCS-PDU Trailer is also not transported across the ATMPW. This does not affect multiprotocol over AAL5 applications since the field is used for alignment and MUST be coded as 0x00 [RFC2684].
- o The trailing CRC field in the CPCS-PDU is stripped at the ingress LCCE and not transported over the ATMPW operating in AAL5-SDU Mode. It is in turn regenerated at the egress LCCE. Since the CRC has end-to-end significance, this means that errors introduced in the ATMPW payload during encapsulation or transit across the packet switched network may not be detected. To allow for payload integrity checking transparency on ATMPWs operating in AAL5-SDU Mode using L2TP over IP or L2TP over UDP/IP, the L2TPv3 session can utilize IPsec as specified in Section 4.1.3 of [RFC3931].

Some additional characteristics of the AAL5-SDU Mode are the following:

- o The status of the ATM PVC is signaled between LCCEs using the Circuit Status AVP. More granular cause values for the ATM circuit status and specific ATM alarm types are signaled using the ATM Alarm Status AVP (see Section 8.1). Additionally, loss of connectivity between LCCEs can be detected by the L2TPv3 keepalive mechanism (see Section 4.4 in [RFC3931]).
- o F5 OAM cells' relative order with respect to user data cells may not be maintained. F5 OAM cells that arrive during the reassembly of an AAL5 SDU are sent immediately over the PW and before the AAL5 SDU payload. At egress, these OAM cells are sent before the cells that comprise the AAL5-SDU. Therefore, applications that rely on cell sequence integrity between OAM and user data cells may not work. This includes Performance Monitoring and Security OAM cells (see Section 5.1). In addition, the AAL5-SDU service allows for OAM emulation in which OAM cells are not transported over the ATMPW (see Section 7). This is advantageous for AAL5-SDU Mode ATMPW implementations that do not support cell transport using the T-bit.
- o Fragmentation and Reassembly procedures MAY be used for managing mismatched MTUs, as specified in Section 5 of [L2TPFRAG] or in the underlying PSN (IP, etc.) between tunnel endpoints as discussed in Section 4.1.4 of [RFC3931]. Only one of these methods SHOULD be used for a given AAL5-SDU Mode ATMPW. The procedures described in [L2TPFRAG] can be used to support the maximum size of an AAL5 SDU, $2^{16} - 1$ (65535) octets. However, relying on fragmentation on the L2TP/IPv4 packet between tunnel endpoints limits the maximum size of the AAL5 SDU

that can be transported, because the maximum total length of an IPv4 datagram is already 65535 octets. In this case, the maximum AAL5 SDU that can be transported is limited to 65535 minus the encapsulating headers, 24-36 octets for L2TP-over-IPv4 or 36-48 octets for L2TP-over-UDP/IPv4. When the AAL5 payload is IPv4, an additional option is to fragment IP packets before tunnel encapsulation with L2TP/IP (see Section 4.1.4 of [RFC3931]).

- o Sequencing may be enabled on the ATMPW using the ATM-Specific Sublayer Sequence Number field, to detect lost, duplicate, or out-of-order frames on a per-session basis (see Section 4.2).
- o Quality of Service characteristics such as throughput (cell rates), burst sizes and delay variation can be provided by leveraging Quality of Service features of the LCCEs and the underlying PSN, increasing the faithfulness of ATMPWs. This includes mapping ATM service categories to a compatible PSN class of service.

9.2. ATM Cell Relay Mode

In this mode, no reassembly takes place at the ingress LCCE. There are no SAR requirements for LCCEs. Instead, ATM-layer cells are transported over the ATMPW. Consequently, all AAL types can be transported over ATMPWs operating in Cell Relay Mode. ATM Cell Relay Pseudowires can operate in three different modes (see Section 5.2): ATM VCC, ATM VPC, and ATM Port Cell Relay Services. The following are some of their characteristics:

- o The ATM cells transported over Cell Relay Mode ATMPWs consist of a 4-byte ATM cell header and a 48-byte ATM cell-payload (see Section 5.2). The ATM Service Payload of a Cell Relay Mode ATMPW is a multiple of 52 bytes. The Header Error Checksum (HEC) in the ATM cell header containing a Cyclic Redundancy Check (CRC) calculated over the first 4 bytes of the ATM cell header is not transported. Accordingly, the HEC field may not accurately reflect errors on an end-to-end basis; errors or corruption in the 4-byte ATM cell header introduced in the ATMPW payload during encapsulation or transit across the PSN may not be detected. To allow for payload integrity checking transparency on ATMPWs operating in Cell Relay Mode using L2TP over IP or L2TP over UDP/IP, the L2TPv3 session can utilize IPsec as specified in Section 4.1.3 of [RFC3931].

- o ATM PWs operating in Cell Relay Mode can transport a single ATM cell or multiple concatenated cells (see Section 6). Cell concatenation improves the bandwidth efficiency of the ATMPW (by decreasing the overhead) but introduces latency and delay variation.
- o The status of the ATM PVC is signaled between LCCEs using the Circuit Status AVP. More granular cause values for the ATM circuit status and specific ATM alarm types are signaled using the ATM Alarm Status AVP (see Section 8.1). Additionally, loss of connectivity between LCCEs can be detected by the L2TPv3 keepalive mechanism (see Section 4.4 in [RFC3931]).
- o ATM OAM cells are transported in the same fashion as user cells, and in the same order as they are received. Therefore, applications that rely on cell sequence integrity between OAM and user data cells are not adversely affected. This includes performance management and security applications that utilize OAM cells (see Section 5.3).
- o The maximum number of concatenated cells is limited by the MTU size of the session (see Section 5.2 and Section 6). Therefore, Fragmentation and Reassembly procedures are not used for Cell Relay ATMPWs. Concatenating cells to then fragment the resulting packet defeats the purpose of cell concatenation. Concatenation of cells and fragmentation act as inverse functions, with additional processing but null net effect, and should not be used together.
- o Sequencing may be enabled on the ATMPW to detect lost, duplicate, or out-of-order packets on a per-session basis (see Section 4.2).
- o Quality of Service characteristics such as throughput (cell rates), burst sizes, and delay variation can be provided by leveraging Quality of Service features of the LCCEs and the underlying PSN, increasing the faithfulness of ATMPWs. This includes mapping ATM service categories to a compatible PSN class of service, and mapping CLP and EFCI bits to PSN classes of service. For example, mapping a Constant Bit Rate (CBR) PVC to a class of service with tight loss and delay characteristics, such as an Expedited Forwarding (EF) Per-Hop Behavior (PHB) if the PSN is an IP DiffServ-enabled domain. The following characteristics of ATMPWs operating in Cell Relay Mode include additional QoS considerations:
 - ATM Cell transport VCC Pseudowires allow for mapping multiple ATM VCCs to a single ATMPW. However, a user may

wish to map a single ATM VCC per ATMPW to satisfy QoS requirements (see Section 5.2.1).

- Cell Relay ATMPWs allow for concatenating multiple cells in a single Pseudowire PDU to improve bandwidth efficiency, but may introduce latency and delay variation.

10. Congestion Control

As explained in [RFC3985], the PSN carrying the PW may be subject to congestion, with congestion characteristics depending on PSN type, network architecture, configuration, and loading. During congestion the PSN may exhibit packet loss and packet delay variation (PDV) that will impact the timing and data integrity of the ATMPW. During intervals of acute congestion, some Cell Relay ATMPWs may not be able to maintain service. The inelastic nature of some ATM services reduces the risk of congestion because the rates will not expand to consume all available bandwidth, but on the other hand, those ATM services cannot arbitrarily reduce their load on the network to eliminate congestion when it occurs.

Whenever possible, Cell Relay ATMPWs should be run over traffic-engineered PSNs providing bandwidth allocation and admission control mechanisms. IntServ-enabled domains providing the Guaranteed Service (GS) or DiffServ-enabled domains using Expedited Forwarding (EF) are examples of traffic-engineered PSNs. Such PSNs will minimize loss and delay while providing some degree of isolation of the Cell Relay ATMPW's effects from neighboring streams.

If the PSN is providing a best-effort service, then the following best-effort service congestion avoidance considerations apply: Those ATMPWs that carry constant bit rate (CBR) and variable bit rate-real time (VBR-rt) services across the PSN will most probably not behave in a TCP-friendly manner prescribed by [RFC2914]. In the presence of services that reduce transmission rate, ATMPWs carrying CBR and VBR-rt traffic SHOULD be halted when acute congestion is detected, in order to allow for other traffic or the network infrastructure itself to continue. ATMPWs carrying unspecified bit rate (UBR) traffic, which are equivalent to best-effort IP service, need not be halted during acute congestion and MAY have cells delayed or dropped by the ingress PE if necessary. ATMPWs carrying variable bit rate-non real time (VBR-nrt) services may or may not behave in a TCP-friendly manner, depending on the end user application, but are most likely safe to continue operating, since the end-user application is expected to be delay-insensitive and may also be somewhat loss-insensitive.

LCCEs SHOULD monitor for congestion (for example, by measuring packet loss or as specified in Section 6.5 of [RFC3985]) in order to ensure that the ATM service may be maintained. When severe congestion is detected (for example, when enabling sequencing and detecting that the packet loss is higher than a threshold), the ATM service SHOULD be terminated by tearing down the L2TP session via a CDN message. The PW may be restarted by manual intervention, or by automatic means after an appropriate waiting time.

11. Security Considerations

ATM over L2TPv3 is subject to the security considerations defined in [RFC3931]. There are no additional considerations specific to carrying ATM that are not present carrying other data link types.

12. IANA Considerations

The signaling mechanisms defined in this document rely upon the allocation of the following ATM Pseudowire Types (see Pseudowire Capabilities List as defined in 5.4.3 of [RFC3931] and L2TPv3 Pseudowire Types in 10.6 of [RFC3931]) by the IANA (number space created as part of publication of [RFC3931]):

Pseudowire Types

0x0002 ATM AAL5 SDU VCC transport
0x0003 ATM Cell transparent Port Mode
0x0009 ATM Cell transport VCC Mode
0x000A ATM Cell transport VPC Mode

12.1. L2-Specific Sublayer Type

This number space is created and maintained per [RFC3931].

L2-Specific Sublayer Type

2 - ATM L2-Specific Sublayer present

12.2. Control Message Attribute Value Pairs (AVPs)

This number space is managed by IANA as per [BCP0068].

A summary of the three new AVPs follows:

Control Message Attribute Value Pairs

Attribute Type	Description
-----	-----
86	ATM Maximum Concatenated Cells AVP
87	OAM Emulation Required AVP
88	ATM Alarm Status AVP

12.3. Result Code AVP Values

This number space is managed by IANA as per [BCP0068].

A new Result Code value for the CDN message is defined in Section 7. Following is a summary:

Result Code AVP (Attribute Type 1) Values

General Error Codes

- 22 - Session not established due to other LCCE
cannot support the OAM Cell Emulation

12.4. ATM Alarm Status AVP Values

This is a new registry for IANA to maintain.

New Attribute values for the ATM Alarm Status AVP in the SLI message are defined in Section 8.1. Additional values may be assigned by Expert Review [RFC2434]. Following is a summary:

ATM Alarm Status AVP (Attribute Type 88) Values

Circuit Status Reason values for the SLI message are as follows:

- 0 - Reserved
- 1 - No alarm or alarm cleared (default for Active Status)
- 2 - Unspecified or unknown Alarm Received (default for Inactive Status)
- 3 - ATM Circuit received F1 Alarm on ingress LCCE
- 4 - ATM Circuit received F2 Alarm on ingress LCCE
- 5 - ATM Circuit received F3 Alarm on ingress LCCE
- 6 - ATM Circuit received F4 Alarm on ingress LCCE
- 7 - ATM Circuit received F5 Alarm on ingress LCCE
- 8 - ATM Circuit down due to ATM Port shutdown on Peer LCCE
- 9 - ATM Circuit down due to loop-back timeout on ingress LCCE

The general ATM Alarm failures are encoded as below:

- 0 - Reserved
- 1 - No Alarm type specified (default)
- 2 - Alarm Indication Signal (AIS)
- 3 - Remote Defect Indicator (RDI)
- 4 - Loss of Signal (LOS)
- 5 - Loss of Pointer (LOP)
- 6 - Loss of Framer (LOF)
- 7 - Loopback cells (LB)
- 8 - Continuity Check (CC)

12.5. ATM-Specific Sublayer Bits

This is a new registry for IANA to maintain.

The ATM-Specific Sublayer contains 8 bits in the low-order portion of the header. Reserved bits may be assigned by IETF Consensus [RFC2434].

- Bit 0 - Reserved
- Bit 1 - S (Sequence) bit
- Bit 2 - B (Fragmentation) bit
- Bit 3 - E (Fragmentation) bit
- Bit 4 - T (Transport type) bit
- Bit 5 - G (EFCI) bit
- Bit 6 - C (CLP) bit
- Bit 7 - U (Command/Response) bit

13. Acknowledgements

Thanks for the contributions from Jed Lau, Pony Zhu, Prasad Yaditi, Durai, and Jaya Kumar.

Many thanks to Srinivas Kotamraju for editorial review.

Thanks to Shouu Yiu and Fred Shu for giving their valuable time to review this document.

14. References

14.1. Normative References

- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

14.2. Informative References

- [PWE3ATM] Martini, L., "Encapsulation Methods for Transport of ATM Over MPLS Networks", Work in Progress, September 2005.
- [L2TPFRAG] Malis, A. and M. Townsley, "PWE3 Fragmentation and Reassembly", Work in Progress, November 2005.
- [FRF8.1] "Frame Relay / ATM PVC Service Interworking Implementation Agreement (FRF 8.1)", Frame Relay Forum 2000.
- [BCP0068] Townsley, W., "Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers Authority (IANA) Considerations Update", BCP 68, RFC 3438, December 2002.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [I610-1] ITU-T Recommendation I.610 (1999): B-ISDN operation and maintenance principles and functions
- [I610-2] ITU-T Recommendation I.610, Corrigendum 1 (2000): B-ISDN operation and maintenance principles and functions (corrigendum 1)
- [I610-3] ITU-T Recommendation I.610, Amendment 1 (2000): B-ISDN operation and maintenance principles and functions (Amendment 1)
- [ATMSEC] ATM Forum Specification, af-sec-0100.002 (2001): ATM Security Specification version 1.1
- [RFC2684] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 2684, September 1999.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, September 2000.

Authors' Addresses

Sanjeev Singh
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134

EMail: sanjeevs@cisco.com

W. Mark Townsley
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

EMail: mark@townsley.net

Carlos Pignataro
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

EMail: cpignata@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

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 provided by the IETF Administrative Support Activity (IASA).

