

Network Working Group  
Request for Comments: 3337  
Category: Standards Track

B. Thompson  
T. Koren  
Cisco Systems  
B. Buffam  
Seaway Networks  
December 2002

Class Extensions for PPP over  
Asynchronous Transfer Mode Adaptation Layer 2 (AAL2)

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 (2002). All Rights Reserved.

Abstract

The Point-to-Point Protocol (PPP) over Asynchronous Transfer Mode (ATM) Adaptation Layer 2 defines the encapsulation that allows a PPP session to be transported over an ATM virtual circuit using the ATM Adaptation Layer 2 (AAL2) adaptation layer. This document defines a set of class extensions to PPP over AAL2 that implement equivalent functionality to Multi Class Multi Link PPP over a single ATM virtual circuit. Instead of using Multi Link PPP as the basis for fragmentation functionality, this document uses the functionality of the Segmentation and Reassembly Service Specific Convergence Sublayer that is already required as the basic encapsulation format of PPP over AAL2.

1. Introduction

Using AAL2 as an adaptation layer for PPP transport over ATM provides a bandwidth efficient transport for IP applications that generate small packets. An example IP application that generates small packets is RTP encapsulated voice (Voice over IP).

In addition to bandwidth efficiency, real-time applications such as voice require low latency. RFC 2689 [2] describes an architecture for providing transport services for real time applications on low bit rate links. The main components of the architecture are: a real-time encapsulation format for asynchronous and synchronous low-bitrate links, a header compression architecture optimized for real-time flows, elements of negotiation protocols used between routers (or between hosts and routers), and announcement protocols used by applications to allow this negotiation to take place.

Multi Class Multi Link PPP [3] defines a fragment-oriented solution for the real-time encapsulation format part of the architecture defined in [2], i.e., for the queues-of-fragments type sender. As described in more detail in the architecture document, a real-time encapsulation format is required to guarantee low latency in the presence of large non real time packets. For example, a 1500 byte packet on a 128 kbit/s ATM virtual circuit makes this link unavailable for the transmission of real-time information for about 100 ms. This adds a worst-case delay that causes real-time applications to operate with round-trip delays that are too high for many interactive tasks. Multi Class Multi Link PPP defines a set of extensions of Multi Link PPP [4] that enable the sender to fragment the packets of various priorities into multiple classes of fragments, allowing high-priority packets to be sent between fragments of lower priorities.

This document defines a set of class extensions to PPP over AAL2 [1] that implement equivalent functionality to Multi Class Multi Link PPP over a single ATM virtual circuit. Instead of using Multi Link PPP as the basis for fragmentation functionality, this document uses the functionality of the Service Specific Segmentation and Reassembly Sublayer (SSSAR) [5] that is already required as the basic encapsulation format of PPP over AAL2.

In addition to providing fragmentation, the real time transport service must allow high priority fragments to be sent between fragments of lower priorities. This can be accomplished in PPP over AAL2 by allowing a single PPP session to span multiple AAL2 CPS [6] Channel Identifiers. Once a PPP session spans multiple AAL2 Channel IDs, the Channel ID can be used to identify the class that a fragment belongs to. Fragments belonging to a high priority class can be sent using a particular AAL2 Channel ID. Fragments of lower priority classes can be sent using different AAL2 Channel IDs. Once multiple fragment classes are identified using different AAL2 Channel IDs, the AAL2 CPS layer can be used to send fragments belonging to a high priority class between fragments of lower priorities.

The class based extensions to PPP over AAL2 use existing services of the AAL2 SSCS and CPS layers already specified in PPP over AAL2. Because of this, the extensions described in this document may be viewed as a desirable alternative to Multi Class Multi Link PPP in providing a class based transport service with PPP over AAL2.

### 1.1. Specification Language

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [7].

## 2. Requirements

This document assumes the same service requirements as defined in Multi Class Multi Link PPP [3]. The reader is referred to section 2 of Multi Class Multi Link PPP for the general requirements of a multi class fragmentation / preemption service.

## 3. Class Extensions for PPP over AAL2

PPP over AAL2 uses the Service Specific Segmentation and Reassembly Sublayer (SSSAR) [5] for the AAL type 2. The SSSAR sub-layer is used to segment PPP packets into frames that can be transported using the AAL2 CPS. The SSSAR sub-layer uses different AAL2 UUI code-points to indicate whether a segment is the last segment of a packet or not. SSSAR provides basic fragmentation functionality for all packets encapsulated using PPP over AAL2. The SSSAR layer fragments all packets into 64 byte fragments.

The AAL2 CPS layer defines a Channel ID that is used to identify multiple streams of packets within a single ATM Virtual Circuit. In this document, the AAL2 CPS Channel ID is used to identify the preemption class that a packet fragment belongs to. Since the Channel ID is used to identify different preemption classes, packet fragments from each class of traffic MUST be assigned to different Channel IDs. In addition, each PPP session MUST have at least as many Channel IDs assigned as there are different classes of preemptible traffic.

To allow PPP packets to be assigned to different preemption classes, PPP packets must be classified into multiple preemption classes as they are fragmented using SSSAR. Many classification methods may be used to determine the class that a particular PPP packet belongs to. The architecture document [2] describes possible alternatives that MAY be used to implement a real time classification scheme.

Once packets have been classified into different preemption classes, each class of traffic is then assigned a different Channel ID. Since fragments from each traffic class are now transmitted using separate Channel ID, the AAL2 CPS layer can be used to schedule fragments from the different classes. The AAL2 CPS specification [6] does not specify a method for scheduling AAL2 CPS payloads from different Channel IDs. The scheduling method required at the AAL2 CPS layer depends upon the real time requirements of applications using this service. Some real-time applications MAY require the use of a priority based CID scheduler. Other applications MAY only require a fair or weighted fair CID scheduler. Implementations of PPP over AAL2 real time transport extensions SHOULD implement AAL2 CPS CID schedulers that meet the requirements of multi-class real time applications.

#### 4. Example Implementation: Class Based Extensions for Voice Service

When PPP over AAL2 is used to transport both voice and non-voice packets over low bandwidth ATM virtual circuits, it may be necessary to preempt the transmission of a large data packet in order to transmit a voice packet with minimal delay. The example implementation described below shows an example of how the class extensions for PPP over AAL2 can be used to support a real time voice transport service over low bandwidth AAL2 virtual circuits. To guarantee low latency and loss for voice transport, the ATM virtual circuit in this example must be provisioned using a real time traffic class such as VBRnrt or VBRrt.

For the simple voice service described above, 2 classes are sufficient to guarantee low latency for voice packets. The PPP over AAL2 session in this case can be configured to run across 2 AAL2 CPS Channel IDs. One channel ID is used to transport large data packets while the other channel ID is used to transport real time voice packets.

Packets that arrive at the PPP interface must first be classified as either belonging to the real time class or belonging to the data class. A simple classifier that can be used to classify packets at this layer is packet size.

Large packets are assigned to the non-real time (or data) traffic class and small packets are assigned to the real time traffic class. The packet size used to discriminate between real time and non-real time packets may vary based on the application and transmission rate of the virtual circuit.

Once packets have been classified, they are now fragmented using the SSSAR layer of PPP over AAL2. Separate instances of the SSSAR fragmentation function run on each of the 2 Channel IDs assigned to the PPP session.

Fragments coming from the SSSAR functions are now scheduled into the AAL2 virtual circuit using the AAL2 CPS layer. Most AAL2 SAR implementations currently implement fair scheduling across multiple AAL2 Channel IDs. Since the AAL2 CPS scheduler implements fair scheduling, real time fragments will wait for at most one non-real time fragment to be transmitted on the AAL2 virtual circuit before being scheduled.

#### 5. Security Considerations

Operation of this protocol is believed to be no more and no less secure than operation of PPP over AAL2 [1].

#### 6. Acknowledgements

The authors would like to thank James Carlson for his contributions to this proposal.

#### 7. References

- [1] Thompson, B., Koren, T. and B. Buffam, "PPP Over Asynchronous Transfer Mode Adaptation Layer 2", RFC 3336, December 2002.
- [2] Bormann, C., "Providing Integrated Services over Low-bitrate Links", RFC 2689, September 1999.
- [3] Bormann, C., "The Multi-Class Extension to Multi-Link PPP", RFC 2686 September 1999.
- [4] Sklower, K., Lloyd, B., McGregor, G., Carr, D. and T. Coradetti, "The PPP Multilink Protocol (MP)", RFC 1990, August 1996.
- [5] International Telecommunications Union, "Segmentation and Reassembly Service Specific Convergence Sublayer for the AAL type 2", ITU-T Recommendation I.366.1, June 1998.
- [6] International Telecommunications Union, "BISDN ATM Adaptation layer specification: Type 2 AAL(AAL2)", ITU-T Recommendation I.363.2, September 1997.
- [7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

## 8. Authors' Addresses

Bruce Thompson  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134  
USA

Phone: +1 408 527-0446  
EMail: brucet@cisco.com

Bruce Buffam  
Seaway Networks  
One Chrysalis Way,  
Suite 300,  
Ottawa, Canada  
K2G-6P9

Phone: +1 613 723-9161  
EMail: bruce@seawaynetworks.com

Tmima Koren  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134  
USA

Phone: +1 408 527-6169  
EMail: tmima@cisco.com

## 9. Full Copyright Statement

Copyright (C) The Internet Society (2002). 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.