Network Working Group
Internet Draft
Category: Standards Track
Expiration Date: August 2005

George Swallow
Cisco Systems, Inc.

Mickey Spiegel
Cisco Systems, Inc.

February 2005

Soft Permanent Virtual Circuit Interworking between PWE3 and ATM

draft-swallow-pwe3-spvc-iw-02.txt

Status of this Memo

Copyright Notice

Abstract

This document defines interworking between Private Network Node Interface (PNNI) SPVC signaling and the Label Distribution Protocol

[LDP] as extended by [PWE3-CP] and [L2VPN-SIG].

Contents

1. Introduction

In current ATM deployments, Soft Permanent Virtual Connections (SPVCs) are used to provision both Asynchronous Transfer Mode (ATM) Permanent Virtual Channel Connections (PVCC) and Permanent Virtual Path Connections (PVPC) and Frame Relay (FR) PVCCs.

Pseudowires over Multiprotocol Label Switching (MPLS) and L2TPv3 PSNs are current being introduced as backbone technologies to carry these same services. Mechanisms are being developed to enable a flexible provisioning model which incorporates elements of the SPVC model in that configuration of the end service exists only at the end-points. These mechanisms are described in [PWE3-CP] and [L2VPN-SIG].

As services are migrated from ATM to PSNs, any reasonable deployment scenario mandates that there be a period of time (possibly protracted or permanent) in which services will need to be established and maintained between end-users where one end of a circuit is attached to an ATM network and the other end is attached to a PSN.

1.1. Conventions

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 RFC 2119 [KEYWORDS].

1.2. Terminology

PAE Provider ATM Edge, a customer facing node which is part of the ATM network

AE ATM Edge, a node of the ATM network which is attached to a node of the MPLS/IP network

PSN Packet Switched Network

PME Provider MPLS/IP Edge, a customer facing node which is part of the MPLS/IP PSN

ME MPLS/IP Edge, a node of the MPLS/IP PSN which is attached to a node of the ATM network

AI Attachment Identifier

AGI Attachment Group Identifier

AII	Attachment Individual Identifier
SAI	Source Attachment Identifier
TAI	Target Attachment Identifier
SAII	Source Attachment Individual Identifier
TAII	Target Attachment Individual Identifier
PNNI	Private Network-Network Interface
UNI	User Network Interface
AINI	ATM Inter-Network Interface
PVCC	Permanent Virtual Channel Connection
PVPC	Permanent Virtual Path Connection
SPVC	Soft Permanent Virtual Connection
IE	Information Element

2. Problem Statement

To facilitate the migration of ATM and Frame Relay SPVCs to a PSN carrying pseudowires, a means of interoperating ATM and LDP signaling needs to be defined. Further this interoperation must preserve the essential reasons for using SPVCs, namely, keeping configuration limited to the edge nodes supporting a particular connection and allowing the network to be able to recover in the event of the failure of any facility between those two edge nodes.

It is also useful to perform reassembly of AAL5 frames of Frame Relay connections at the boundary between the ATM network and the PSN. This serves to reduce dataplane protocol overhead in the PSN.

Finally, any solution must not preclude any existing services. In particular, Frame Relay to ATM interworking is recognized to be in wide use.

2.1. Reference Network Diagram

The diagram below shows an ATM network interfaced to a PSN. A soft PVC is to be set up between the two customer facing edge nodes, PAE, and PME. Two interconnections are shown, to indicate that the connection must be routable over either interconnection in the event of the failure of the preferred interconnection. The 'M' (for MPLS/IP PSN) was chosen over 'P' (for PSN) to allow 'P' to stand for provider.

```
               _______                          _______
              /       %                        /       %
             /         %   +-----+  +-----+   /         %
            /           %  |     |  |     |  /           %
            |           |==| AE1 |==| ME1 |==|           |
   +-----+  |           |  |     |  |     |  |           |  +-----+
   |     |  |    ATM    |  +-----+  +-----+  |  MPLS/IP  |  |     |
   | PAE |==|           |                    |    PSN    |==| PME |
   |     |  |           |  +-----+  +-----+  |           |  |     |
   +-----+  |           |  |     |  |     |  |           |  +-----+
            |           |==| AE2 |==| ME2 |==|           |
            %           /  |     |  |     |  %           /
             %         /   +-----+  +-----+   %         /
              %_______/                        %_______/
```

Key:

All node acronyms are composed of the following letters

P - Provider
M - MPLS/IP PSN
A - ATM
E - Edge

Figure 1: Reference Network

3. Requirements

The purpose of Soft Permanent Virtual Connections is two-fold. First they confine circuit specific configuration to the edge nodes where the user access circuits are terminated. Second they allow the network to automatically reroute / re-establish the circuit when failures are encountered. The requirements for this solution follow directly from this as well as some of the common uses of SPVCs.

3.1. Configuration

All per circuit configuration must be limited to the edge nodes. In particular the solution must not necessitate any per circuit configuration on the nodes that support the ATM/PSN interface.

3.2. Redundant ATM/PSN Interfaces

The solution must support multiple inter-connections between the ATM network and the PSN. Upon failure of an ATM/PSN interface, it must be possible to re-route the SPVCs that had been traversing that interface over other ATM/PSN interfaces.

3.3. Re-assembly

It must be possible to locate the AAL5 re-assembly at the ME. That is, a ME by default will perform AAL5 re-assembly for Frame Relay SPVCs. The ME's ATM/PSN interface may be configured to not perform re-assembly and leave the job to the target PME, when the target PME supports AAL5 re-assembly for Frame Relay circuits. No per circuit selection need be supported.

3.4. No Change to Existing ATM Signaling Protocols

It is highly desirable that no changes be required to existing ATM signaling protocols.

3.5. Frame Relay and ATM Circuits at the MPLS Network Edge

The solution must support both ATM and Frame Relay circuits at the Provider MPLS Edge. For the case of ATM at the PME and Frame Relay at the PAE, Frame Relay to ATM interworking is the responsibility of the ATM network. For the case of Frame Relay at both the PME and the PAE, FRF5 or FRF8.1 Frame Relay to ATM interworking capability is required at the PME (to be specific, Frame Relay to AAL5 SDU Frame encapsulation over PW) and is optional at the ME (Frame Relay over Pseudo Wire to ATM). For the case of Frame Relay at the PME and ATM at the PAE, Frame Relay to ATM interworking capability is required at the PME and is optional at the ME.

4. Non-Requirements

4.1. Frame Relay / ATM Interworking

While Frame Relay to ATM interworking is recognized as an important and pervasive service, such a service is deemed to be beyond the scope of this document. This is not to imply that such a service cannot be supported by the means specified herein. It merely implies that when Frame Relay to ATM interworking is required in the PSN network, the interworking function (IWF) is located (and configured) at the PME.

5. Background on identifiers

5.1. Pseudowire Identifiers

Pseudowires serve to connect a pair of attachment circuits (ACs). In the context of this document these ACs are either Frame Relay DLCIs or ATM VPI/VCIs. An AC is identified by an Attachment Identifier (AI) and the IP address of the egress node. AIs are defined in [1]. An AI is a logical reference for both the physical/logical port as well as the virtual circuit. That is an AC is fully identified by a node-ID (IP address) and an AI.

The AI has further structure to designate groups of identifiers and individual identifiers within a group, these are called attachment group identifiers (AGI) and attachment individual identifiers (AII), respectively. When pairs of AIs are used in signaling, they are further designated by their role in the call, with the operative terms being source and target of the call. Thus we also have the terminology, source attachment identifier (SAI), source attachment individual identifier (SAII), target attachment identifier (TAI), and target attachment individual identifier (TAII). The source and target designations are only relevant from the perspective of the pseudowire control protocol. For example, a node receiving an LDP label mapping message from a remote node will swap the SAI and TAI values when it sends a label mapping message in the reverse direction back to the originating node.

Attachment Identifiers (AIs) are carried in the Generalized ID FEC Element of LDP. Each AI encoded as two fields, the Attachment Group Identifier (AGI) and an Attachment Individual Identifier (AII), each encoded using a type-length-value (TLV) format as defined in Section 5.2.2/[PWE3-CP]. In particular:

"Note that the interpretation of a particular field as AGI, SAII, or

TAII depends on the order of its occurrence. The type field identifies the type of the AGI, SAII, or TAII.

The rules for constructing the AGI and AII are left to the specification of applications and/or models."

5.2. ATM SPVC Identifiers

In ATM, the identifying information of the attachment circuit at the destination interface consists of an ATM End-System Address (AESA) and the DLCI or VPI/VCI. The AESA identifies both the destination node and port where the end-user is attached. The DLCI or VPI/VCI are signaled in the Called party SPVC IE and are carried as literal values. The Called party SPVC IE has two formats depending on whether the service being signaled is a Frame Relay SPVC or an ATM SPVC. Furthermore, ATM SPVCCs and ATM SPVPCs are distinguished through the bearer class codepoint in the Broadband bearer capability IE.

AESAs are based on the NSAP address format. Figure 2 shows the generic format.

```
 |--AFI (Authority & Format Identifier)
 |                                   Selector
 |  |--IDP (Initial Domain Part)       |
 V  V                                  V
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| |   |    H O - D S P     |   E S I  |0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

HO-DSP   High Order Domain Specific Part
ESI      End System Identifier
```

Figure 2: Generic AESA Format

Although many formats are permitted within AESAs, all ATM Forum defined formats include a six byte End System Identifier or ESI. The ESI's role is to identify a host. Typically, the first 13 bytes of the AESA are common for all end systems attached to an ATM node. This is the default behavior in PNNI. In this case, the ESI is used to differentiate between end systems attached to the same switch. From the point of view of the egress ATM switch, the ESI maps to a physical or logical port.

Thus common practice is to use the ESI to carry the port information.

6. Proposed Solution

6.1. PSN / ATM Interface

The interface between the ATM network and the PSN can be any of the following:

ATM Forum User Network Interface (UNI)
ITU-T User Network Interface (UNI)
Private Network-Network Interface (PNNI)
ATM Inter-Network Interface (AINI)

In the case of the UNI, there must be extensions to support the Called party soft PVPC or PVCC IE defined in [PNNI]. (In this document we refer to the Called party soft PVPC or PVCC IE as simply the SPVC IE.) There may be extensions to support:

- o the Calling party soft PVPC or PVCC IE,
- o signaled VPs, using the "transparent VP service" codepoint for the bearer class in the Broadband bearer capability IE,
- o crankback indication by setting the cause location in the Cause IE to a value other than "user", and
- o frame discard indication using either the Frame Discard bits or the ATM adaptation layer parameters IE.

[Note: while the details of various versions of ATM signaling and the support for particular IEs is a subject for other Fori and thus beyond the scope of the WG (and IETF), an informative appendix appropriate references will be added if the WG so desires.]

6.2. Signaling

In ATM soft PVCs are statically defined across the UNI interfaces, but are signaled across the ATM network using PNNI signaling. For the signaling part, one edge node is configured to be active for the SPVC and the other is defined to be passive. That is, one end always initiates the call. ATM signaling creates a bidirectional circuit in a single pass. Bandwidth parameters for each direction of the circuit are carried in the setup message.

A paradigm analogous to the active/passive roles in SPVC setup above for pseudowires is known as single sided provisioning. These procedures are defined in draft-rosen-ppvpn-l2-signaling-03.txt [L2VPN-SIG] section 5.1.1.2. A small difference exists here in that the "discovery" process occurs when an ATM SETUP message arrives.

It should be noted, that the pseudowire setup remains a pair of uni-directional labels assigned by two essentially independent label requests. It is only the triggering of the reverse label request that is tied to the forward label request. Further [PWE3-CP] does not carry bandwidth parameters at all.

6.3. Mapping Identifiers

In [PWE3-CP], an IP address identifies the egress node, and the (T)AI identifies the egress port and DLCI or VPI/VCI. In ATM, an AESA identifies both the egress node and port, and the DLCI or VPI/VCI is carried as a literal in the SPVC IE.

Two issues must be addressed. First a mapping between ATM and IP addresses is needed. Second a means of translating between the ATM port and SPVC IE and the AI used in [PWE3-CP].

6.3.1. Mapping Addresses

OSI Network Service Access Point Addresses include an Authority and Format Identifier (AFI). The AFI value 35 has been assigned to the IANA. Within this format a two-octet Internet Code Point (ICP) field is available for assignment by the IANA. Currently there is one code point, 0, assigned for embedding IPv6 addresses. A format and code point assignment has been proposed by the ATM Forum in [ATM-ipaddr]. That format is shown below.

```
          |--AFI (Authority &
          |       Format Identifier)          Selector
          |  |--IDP (Initial Domain Part)       |
          V  V |<---- HO-DSP ----->|<-- ESI -->|V
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |3|0 0|I P v 4|0 0 0 0 0 0|           |0|
         |5|0 1|Address|0 0 0 0 0 0|           |0|
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

         HO-DSP   High Order Domain Specific Part
         ESI      End System Identifier
```

Figure 3: NSAP with Embedded IPv4 Address

While it is common practice to carry the port number in the ESI field, We note that there are six unused bytes in the HO-DSP as well as the Selector Byte which could be used in a situation where the ESI was not available.

The format to embed an IPv6 address in an NSAP address is defined in [RFC1888]. In this format the only unused space is the Selector Byte. This allows for the identification of 256 ports. If more ports are needed, multiple addresses must be assigned.

```
          |--AFI (Authority &
          |       Format Identifier)          Selector
          |  |--IDP (Initial Domain Part)       |
          V  V                                  V
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
         |3|0 0|    I P v 6   A d d r e s s    | |
         |5|0 0|                               | |
         +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: NSAP with Embedded IPv6 Address

While it is expected that for IPv4 the ESI will commonly be used and for IPv6 the Selector Byte must be used, the discussion below will simply refers to a generic port field.

6.3.2. Mapping Port and SPVC IEs to AIs

It is proposed that the Port and SPVC IE values be mapped into the AII. The SPVC IE has three formats, one each for Frame Relay, ATM SPVC, and ATM SPVP. Each of these will be mapped into three [to be assigned] AII types. [If there is an issue with IANA allocating

three, one, with some form of sub-typing will do.]

6.4. Configuration

PAEs: For each Permanent Virtual Connection, the PAE is configured with the target AESA and DLCI or VPI/VCI.

AEs: AEs are configured with the AESA prefix representing the set of PSN nodes reachable through its link(s) to the PSN. Multiple prefixes may be configured to allow choices of optimum nodes to reach and to allow reachability to a larger set of nodes, should some other AE or ME fail. The advertisement into the ATM network's routing protocol (e.g.PNNI) should be withdrawn if the associated link(s) have failed.

PMEs: PME require no special configuration. They are simply configured with the (T)AII of each of their ACs.

MEs: MEs must be able to encode and parse the [to be assigned] AII types. Associated with each of these AII type is an AII format (used to form TAIIs) and rules for how to extract the port from the ATM called party address. .fi

6.5. Procedures

6.5.1. Procedures within the ATM Network

In an SPVC, one end is designated as the 'owner' and is responsible for initiating the connection. In order to simplify the interworking, this solution proposes that SPVCs always be initiated from the ATM side. This obviates any need to communicate bandwidth information across the PSN to the ATM network.

The AEs as the owner of the connection, initiates PNNI signaling as it normally would. Finding a longest match associated with one or more of the AEs it performs normal PNNI routing selection and sends a SETUP message which includes the SPVC IE.

When the SETUP message arrives at the AE, it performs normal PNNI signaling processing and forwards the message across the PNNI, AINI or UNI to the ME.

6.5.2. Procedures for the ME

When the ME receives a SETUP message, performs ATM admission control. Additionally, the ME may perform additional checks to determine if it has the necessary resources to support the pseudowire connection in the forward direction. For example, in some network deployments it may determine if a PSN tunnel can be established in order to satisfy QoS or restoration requirements.

In the event that the call cannot be admitted, the ME SHOULD set an appropriate cause code, assuming that it is capable of supporting crankback procedures.

When the ME has successfully performed ATM admission control, it splices the call to a pseudowire using the signaling procedures of [L2VPN-SIG]. First, it extracts the destination IP address from the ATM called party address. It then determines if a LDP session exists with this node. If not, one is initiated. It then examines the SPVC IE to determine the type of service which is being requested. Based on the service type it selects AII type and format. It then extracts the port, VPI, VCI, and/or DLCI information as appropriate to the service and formats an AII. It formats an SAI using the same encoding rules, using the port the setup message was received on and the Connection Identifier. This AI becomes the handle which will be used to locate the context for this call. It then sends a Label Mapping message to the target node.

When it receives a Label mapping message from the target node, it uses the TAI to locate the call context and completes the ATM call.

[Note: while the details of the ATM call processing and crankback codes is a subject for other Fori and thus beyond the scope of the WG (and IETF), an informative appendix will be added if the WG so desires.]

6.5.3. Procedures for the PME

The procedures for the PME follow [L2VPN-SIG] with no changes. That is, the PME uses the TAI to identify interface and DLCI or VPI/VCI. No decoding of the TAII is necessary; the AI and AC are simply configured as in [L2VPN-SIG].

If the forward label mapping completed successfully, the PME responds with an LDP label mapping message in the reverse direction. The same encapsulation as the forward direction MUST be signaled.

6.5.4. Call Completion

When the ME receives the label mapping message, it uses the TAI to (i.e. what it sent as the SAI) locate the context for this call and then completes the ATM call by sending a CONNECT message back to the PAE.

6.6. Handling Re-assembly

When an ME detects that the requested service is Frame Relay, by default it will perform AAL5 re-assembly for Frame Relay SPVCs. In this case the AAL5 SDU frame mode encapsulation as defined in [PWE3-ENCAPS] is RECOMMENDED.

On a per ATM/PSN interface basis, an ME may be configured to not perform re-assembly for Frame Relay. No per circuit selection is provided. In this case the RECOMMENDED encapsulation is ATM N-to-1 Cell Mode.

Both ATM SPVCCs and SPVPCs are encapsulated using one of the cell-mode encapsulations. The RECOMMENDED encapsulation is ATM N-to-1 Cell Mode.

[As noted in the "Issues" section below, the Frame Relay format of the SPVC IE may not be available on some ATM equipment. Alternative means of determining that the service is Frame Relay can be achieved in many cases by examining some combination of the .

7. Issues

The Frame Relay format for the SPVC IE was added later. Some ATM equipment still uses the ATM SPVC format with the DLCI encoded in the VPI/VCI fields. To support these switches without change, and still allow AAL5 reassembly to be done at the ME, some other means of indicating that the circuit is Frame Relay needs to be established.

8. Security Considerations

This document represents a straightforward application of [L2VPN-SIG]. It poses no new security considerations over and above those identified in that document.

9. Acknowledgments

The authors would like to thank Chris Metz and Chandra Krishnamurthy for their input to this document.

10. References

10.1. Normative References

[LDP] Andersson, L. et al., "LDP Specification", RFC 3036, January 2001.

[L2VPN-SIG] Rosen, E.& Radoaca, V, "Provisioning Models and Endpoint Identifiers in L2VPN Signaling", draft-ietf-l2vpn-signaling-00.txt, September 2003.

[PWE3-CP] Martini, L. & Rosen, E., "Pseudowire Setup and Maintenance using LDP", draft-ietf-pwe3-control-protocol-04.txt, October 2003.

[PWE3-ENCAP] Martini, et al., "Encapsulation Methods for Transport of ATM Cells/Frame Over IP and MPLS Networks", draft-ietf-pwe3-atm-encap-03.txt, October 2003.

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

[PWE3-IANA] Townsley, M., "IANA Allocations for pseudo Wire Edge to Edge Emulation (PWE3)", draft-ietf-pwe3-iana-allocation-04.txt, April 2004.

[ATM-AINI] ATM Forum, "ATM Inter-Network Interface Specification Version 1.1 (AINI 1.1)", af-cs-0125.002, September 2002

[ATM-UNI] ATM Forum, "ATM User-Network Interface (UNI) Signalling

Specification Version 4.1", af-sig-0061.002, April 2002

[PNNI] ATM Forum, "Private Network-Network Interface Specification Version 1.1", af-pnni-0055.001, April 2002

[ATM-ipaddr] ATM Forum, "IP-Based Addressing Version 1.0", str-cs-ipaddr-01.01, October 2003

[RFC1888] Bound, J., et al., "OSI NSAPs and IPv6", RFC 1888, August 1996.

11. Authors' Addresses

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719

Email: swallow@cisco.com

Ethan Mickey Spiegel
Cisco Systems, Inc.
3700 Cisco Way
San Jose, CA 95134

Email: mspiegel@cisco.com

12. Full Copyright and Intellectual Property Statements

Intellectual Property

Acknowledgement

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