Internet Engineering Task Force (IETF)
Request for Comments: 6658
Category: Standards Track
ISSN: 2070-1721

S. Bryant, Ed.
L. Martini
G. Swallow
Cisco Systems
A. Malis
Verizon Communications

July 2012

Packet Pseudowire Encapsulation over an MPLS PSN

Abstract

This document describes a pseudowire mechanism that is used to transport a packet service over an MPLS PSN in the case where the client Label Switching Router (LSR) and the server Provider Edge equipments are co-resident in the same equipment. This pseudowire mechanism may be used to carry all of the required layer 2 and layer 3 protocols between the pair of client LSRs.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6658.

Table of Contents

1. Introduction

There is a need to provide a method of carrying a packet service over an MPLS PSN in a way that provides isolation between the two networks. The server MPLS network may be an MPLS network or a network conforming to the MPLS Transport Profile (MPLS-TP) [RFC5317]. The client may also be either an MPLS network or a network conforming to the MPLS-TP. Considerations regarding the use of an MPLS network as a server for an MPLS-TP network are outside the scope of this document.

Where the client equipment is connected to the server equipment via a physical interface, the same data-link type must be used to attach the clients to the Provider Edge (PE) equipments, and a pseudowire (PW) of the same type as the data-link must be used [RFC3985]. The reason that interworking between different physical and data-link attachment types is specifically disallowed in the pseudowire architecture is because this is a complex task and not a simple bit-mapping exercise. The interworking is not limited to the physical and data-link interfaces and the state-machines. It also requires a compatible approach to the formation of the adjacencies between attached client network equipment. As an example, the reader should consider the differences between router adjacency formation on a point-to-point link compared to a multipoint-to-multipoint interface (e.g., Ethernet).

A further consideration is that two adjacent MPLS Label Switching Routers (LSRs) do not simply exchange MPLS packets. They exchange IP packets for adjacency formation, control, routing, label exchange, management, and monitoring purposes. In addition, they may exchange data-link packets as part of routing (e.g., IS-IS Hellos and IS-IS Link State Packets) and for Operations, Administration, and Maintenance (OAM) purposes such as the Link-Layer Discovery Protocol [IEEE.802.1AB.2009]. Thus, the two clients require an attachment mechanism that can be used to multiplex a number of protocols. In addition, it is essential to the correct operation of the network layer that all of these protocols fate share.

Where the client LSR and server PE are co-located in the same equipment, the data-link layer can be simplified to a point-to-point Ethernet used to multiplex the various data-link types onto a pseudowire. This is the method described in this document.

Appendix A provides information on alternative approaches to providing a packet PW that were considered by the PWE3 Working Group and the reasons for using the method defined in this specification.

1.1. Requirements Language

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 [RFC2119].

2. Network Reference Model

The network reference model for the packet pseudowire operating in an MPLS network is shown in Figure 1. This is an extension of Figure 3 "Pre-processing within the PWE3 Network Reference Model" from [RFC3985].

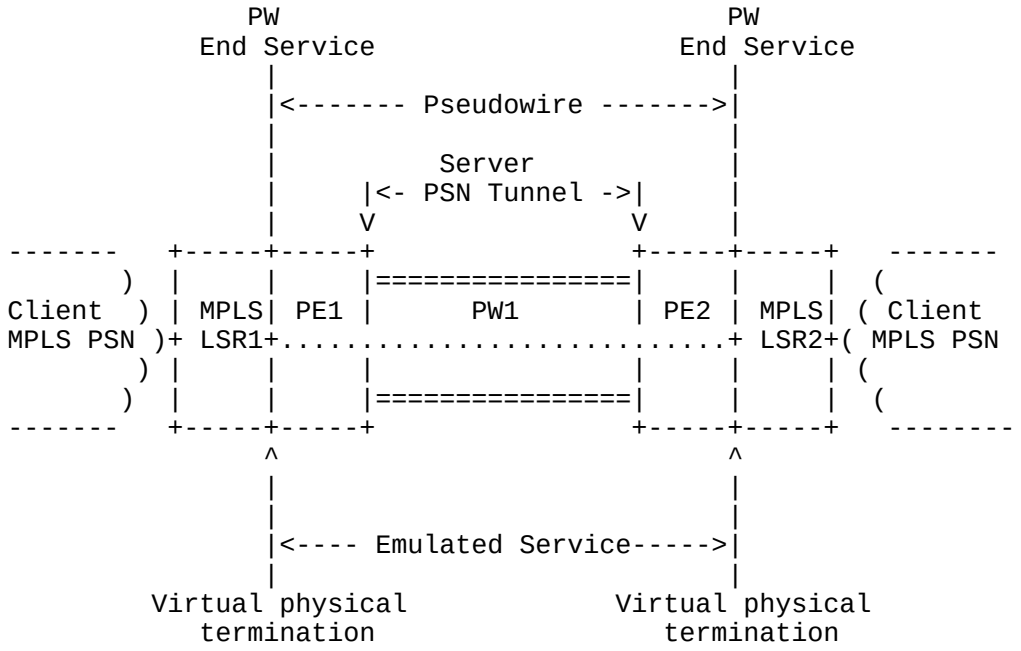

Figure 1: Packet PW Network Reference Model

In this model, the LSRs (LSR1 and LSR2) are part of the client MPLS PSN. The PEs (PE1 and PE2) are part of the server PSN that is to be used to provide connectivity between the client LSRs. The attachment circuit that is used to connect the MPLS LSRs to the PEs is a virtual interface within the equipment. A packet pseudowire is used to provide connectivity between these virtual interfaces. This packet pseudowire is used to transport all of the required layer 2 and layer 3 protocols between LSR1 and LSR2.

3. Client Network-Layer Model

The packet PW appears as a single point-to-point link to the client layer. Network-layer adjacency formation and maintenance between the client equipments will follow the normal practice needed to support the required relationship in the client layer. The assignment of metrics for this point-to-point link is a matter for the client layer. In a hop-by-hop routing network, the metrics would normally be assigned by appropriate configuration of the embedded client network-layer equipment (e.g., the embedded client LSR). Where the client was using the packet PW as part of a traffic-engineered path, it is up to the operator of the client network to ensure that the server-layer operator provides the necessary service-level agreement.

4. Forwarding Model

The packet PW forwarding model is illustrated in Figure 2. The forwarding operation can be likened to a virtual private network (VPN), in which a forwarding decision is first taken at the client layer, an encapsulation is applied, and then a second forwarding decision is taken at the server layer.

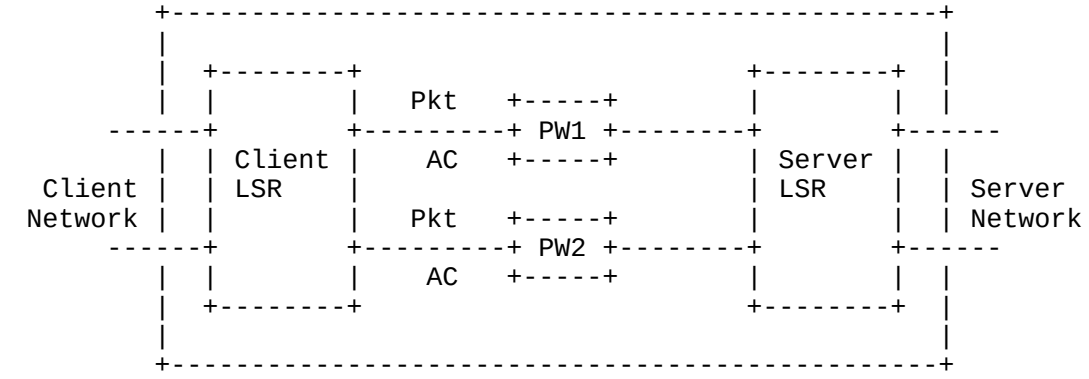

Figure 2: Packet PW Forwarding Model

A packet PW PE comprises three components: the client LSR, a PW processor, and a server LSR. Note that [RFC3985] does not formally indicate the presence of the server LSR because it does not concern itself with the server layer. However it is useful in this document to recognize that the server LSR exists.

It may be useful to first recall the operation of a layer 2 PW such as an Ethernet PW [RFC4448] within this model. The client LSR is not present, and packets arrive directly on the attachment circuit (AC) that is part of the client network. The PW function undertakes any

header processing, if configured to do so; it then optionally pushes the PW control word (CW) and finally pushes the PW label. The PW function then passes the packet to the LSR function, which pushes the label needed to reach the egress PE and forwards the packet to the next hop in the server network. At the egress PE, the packet typically arrives with the PW label at the top of the stack; the packet is thus directed to the correct PW instance. The PW instance performs any required reconstruction using, if necessary, the CW, and the packet is sent directly to the attachment circuit.

Now let us consider the case of client-layer MPLS traffic being carried over a packet PW. An LSR belonging to the client layer is embedded within the PE equipment. This is a type of native service processing element [RFC3985]. The client LSR determines the next hop in the client layer, and pushes the label needed by the next hop in the client layer. It then encapsulates the packet in an Ethernet header setting the Ethertype to MPLS, and the client LSR passes the packet to the correct PW instance. The PW instance then proceeds as defined for an Ethernet PW [RFC4448] by optionally pushing the control word, then pushing the PW label, and finally handing the packet to the server-layer LSR for delivery to the egress PE in the server layer.

At the egress PE in the server layer, the packet is first processed by the server LSR, which uses the PW label to pass the packet to the correct PW instance. This PW instance processes the packet as described in [RFC4448]. The resultant Ethernet encapsulated client packet is then passed to the egress client LSR, which then processes the packet in the normal manner.

Note that although the description above is written in terms of the behavior of an MPLS LSR, the processing model would be similar for an IP packet or any other protocol type.

Note that the semantics of the PW between the client LSRs is a point-to-point link.

5. Packet PW Encapsulation

The client network-layer packet encapsulation into a packet PW is shown in Figure 3.

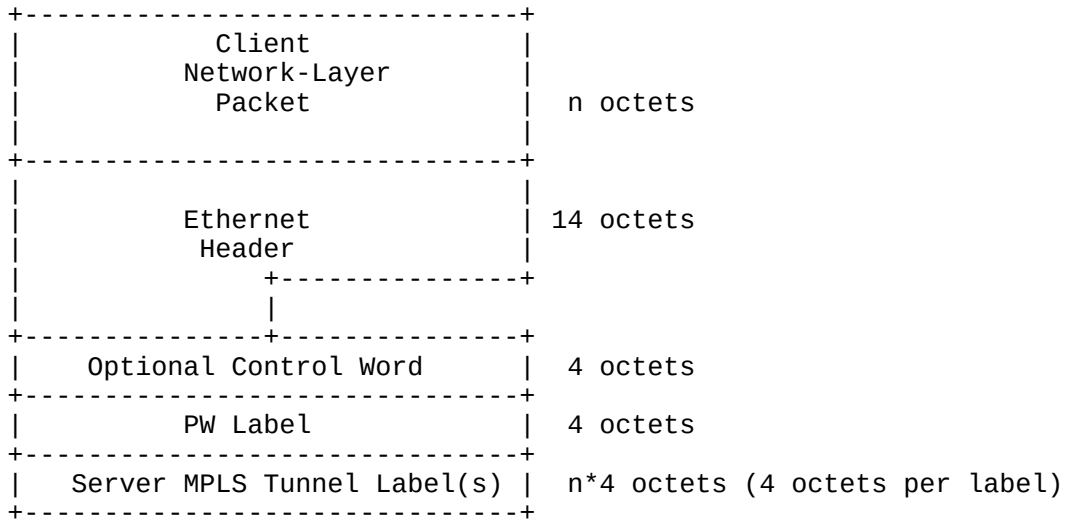

Figure 3: Packet PW Encapsulation

This conforms to the PW protocols stack as defined in [RFC4448]. The protocol stack is unremarkable except to note that the stack does not retain 32-bit alignment between the virtual Ethernet header and the PW optional control word (or the PW label when the optional components are not present in the PW header). This loss of 32 bits of alignment is necessary to preserve backwards compatibility with the Ethernet PW design [RFC4448]

Ethernet Raw Mode (PW type 5) MUST be used for the packet PW.

The PEs MAY use a local Ethernet address for the Ethernet header used to encapsulate the client network-layer packet or MAY use the special Ethernet addresses "PacketPWEthA" or "PacketPWEthB" as described below.

IANA has allocated two unicast Ethernet addresses [RFC5342] for use with this protocol, referred to as "PacketPWEthA" and "PacketPWEthB". Where [RFC4447] signaling is used to set up the PW, the LDP peers numerically compare their IP addresses. The LDP PE with the higher-value IP address will use PacketPWEthA, whilst the LDP peer with the lower-value IP address uses PacketPWEthB.

Where no signaling PW protocol is used, suitable Ethernet addresses MUST be configured at each PE.

Although this PW represents a point-to-point connection, the use of a multicast destination address in the Ethernet encapsulation is REQUIRED by some client-layer protocols. Peers MUST be prepared to handle a multicast destination address in the Ethernet encapsulation.

6. Ethernet and IEEE 802.1 Functional Restrictions

The use of Ethernet as the encapsulation mechanism for traffic between the server LSRs is a convenience based on the widespread availability of existing hardware. In this application, there is no requirement for any Ethernet feature other than its protocol multiplexing capability. Thus, for example, a server LSR is not required to implement the Ethernet OAM.

The use and applicability of VLANs, IEEE 802.1p, and IEEE 802.1Q tagging between PEs is not supported.

Point-to-multipoint and multipoint-to-multipoint operation of the virtual Ethernet is not supported.

7. Congestion Considerations

A packet pseudowire is normally used to carry IP, MPLS and their associated support protocols over an MPLS network. There are no congestion considerations beyond those that ordinarily apply to an IP or MPLS network. Where the packet protocol being carried is not IP or MPLS and the traffic volumes are greater than that ordinarily associated with the support protocols in an IP or MPLS network, the congestion considerations developed for PWs apply [RFC3985] [RFC5659].

8. Security Considerations

The virtual Ethernet approach to packet PW introduces no new security risks. A more detailed discussion of pseudowire security is given in [RFC3985], [RFC4447], and [RFC3916].

9. IANA Considerations

IANA has allocated two Ethernet unicast addresses from "IANA Unicast 48-bit MAC Addresses".

Address	Usage	Reference
00-00-5E-00-52-00	PacketPWEthA	[RFC6658]
00-00-5E-00-52-01	PacketPWEthB	[RFC6658]

10. Acknowledgements

The authors acknowledge the contributions made to this document by Sami Boutros, Giles Herron, Siva Sivabalan, and David Ward.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.

[RFC4448] Martini, L., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, April 2006.

[RFC5342] Eastlake, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", BCP 141, RFC 5342, September 2008.

11.2. Informative References

[IEEE.802.1AB.2009] Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and Metropolitan Area Networks -- Station and Media Access Control Connectivity Discovery", IEEE Standard 802.1AB, 2009.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.

[RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

[RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.

[RFC5317] Bryant, S. and L. Andersson, "Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile", RFC 5317, February 2009.

[RFC5659] Bocci, M. and S. Bryant, "An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge", RFC 5659, October 2009.

[RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.

Appendix A. Encapsulation Approaches Considered

A number of approaches to the design of a packet pseudowire (PW) were investigated by the PWE3 Working Group and were discussed in IETF meetings and on the PWE3 list. This section describes the approaches that were analyzed and the technical issues that the authors took into consideration in arriving at the approach described in the main body of this document. This appendix is provided so that engineers considering alternative optimizations can have access to the rationale for the selection of the approach described in this document.

In a typical network, there are usually no more that four network-layer protocols that need to be supported: IPv4, IPv6, MPLS, and Connectionless Network Service (CLNS). However, any solution needs to be scalable to a larger number of protocols. The approaches considered in this appendix all satisfy this minimum requirement but vary in their ability to support larger numbers of network-layer protocols.

Additionally, it is beneficial if the complete set of protocols carried over the network in support of a set of CE peers fate share. It is additionally beneficial if a single OAM session can be used to monitor the behavior of this complete set. During the investigation, various views were expressed as to where these benefits lay on the scale from absolutely required to "nice to have", but in the end, they were not a factor in reaching our conclusion.

Four candidate approaches were analyzed:

1. A protocol identifier (PID) in the PW control word (CW)

2. A PID label

3. Parallel PWs - one per protocol

4. Virtual Ethernet

A.1. A Protocol Identifier in the Control Word

In this approach, a Protocol Identifier (PID) is included in the PW control word (CW) by appending it to the generic control word [RFC4385] to make a 6-byte CW (it was thought that this approach would include 2 reserved bytes to provide 32-bit alignment, but then this was optimized out). A variant of this is just to use a 2-byte PID without a control word.

This is a simple approach and is basically a virtual PPP interface without the PPP control protocol. This has a smaller MTU than, for example, a virtual Ethernet would need; however, in forwarding terms, it is not as simple as the PID label or multiple PW approaches described next and may not be deployable on a number of existing hardware platforms.

A.2. PID Label

In this approach, the PID is indicated by including a label after the PW label that indicates the protocol type, as shown in Figure 4.

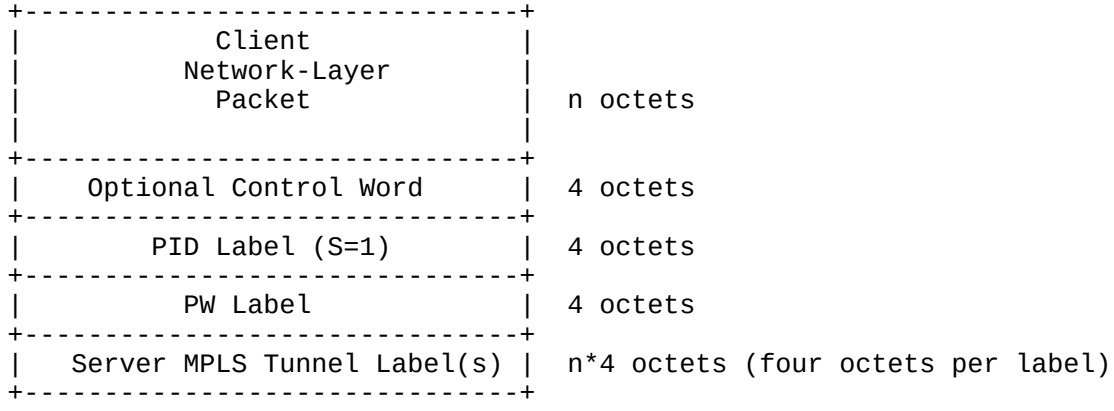

Figure 4: Encapsulation of a Pseudowire with a Pseudowire Load-Balancing Label

In the PID label approach, a new Label Distribution Protocol (LDP) Forwarding Equivalence Class (FEC) element is used to signal the mapping between protocol type and the PID label. This approach complies with [RFC3031].

A similar approach to PID label is described in Section 3.4.5 of [RFC5921]. In this case, when the client is a network-layer packet service such as IP or MPLS, a service label and demultiplexer label (which may be combined) are used to provide the necessary identifications needed to carry this traffic over an LSP.

The authors surveyed the hardware designs produced by a number of companies across the industry and concluded that whilst the approach complies with the MPLS architecture, it may conflict with a number of designers' interpretations of the existing MPLS architecture. This led to concerns that the approach may result in unexpected difficulties in the future. Specifically, there was an assumption in many designs that a forwarding decision should be made on the basis

of a single label. Whilst the approach is attractive, it cannot be supported by many commodity chip sets, and this would require new hardware, which would increase the cost of deployment and delay the introduction of a packet PW service.

A.3. Parallel PWs

In this approach, one PW is constructed for each protocol type that must be carried between the PEs. Thus, a complete packet PW would consist of a bundle of PWs. This model would be very simple and efficient from a forwarding point of view. The number of parallel PWs required would normally be relatively small. In a typical network, there are usually no more that four network-layer protocols that need to be supported: IPv4, IPv6, MPLS, and CLNS. However, any solution needs to be scalable to a larger number of protocols.

There are a number of serious downsides with this approach:

1. From an operational point of view, the lack of fate sharing between the protocol types can lead to complex faults that are difficult to diagnose.

2. There is an undesirable trade-off in the OAM related to the first point. We would have to run an OAM on each PW and bind them together, which leads to significant protocol and software complexity and does not scale well. Alternatively, we would need to run a single OAM session on one of the PWs as a proxy for the others and then diagnose any more complex failures on a case-by-case basis. To some extent, the issue of fate sharing between protocols in the bundle (for example, the assumed fate sharing between CLNS and IP in IS-IS) can be mitigated through the use of Bidirectional Forwarding Detection (BFD).

3. The need to configure, manage, and synchronize the behavior of a group of PWs as if they were a single PW leads to an increase in control-plane complexity.

The Parallel PW mechanism is therefore an approach that simplifies the forwarding plane, but only at a cost of a considerable increase in other aspects of the design, in particular, operation of the PW.

A.4. Virtual Ethernet

Using a virtual Ethernet to provide a packet PW would require PEs to include a virtual (internal) Ethernet interface and then to use an Ethernet PW [RFC4448] to carry the user traffic. This is conceptually simple and can be implemented today without any further standards action, although there are a number of applicability

considerations that it are useful to bring to the attention of the community.

Conceptually, this is a simple approach, and some deployed equipments can already do this. However, the requirement to run a complete Ethernet adjacency led us to conclude that there was a need to identify a simpler approach. The packets encapsulated in an Ethernet header have a larger MTU than the other approaches, although this is not considered to be an issue on the networks needing to carry packet PWs.

The virtual Ethernet mechanism was the first approach that the authors considered, before the merits of the other approaches appeared to make them more attractive. As we shall see below, however, the other approaches were not without issues, and it appears that the virtual Ethernet is the preferred approach to providing a packet PW.

A.5. Recommended Encapsulation

The operational complexity and the breaking of fate-sharing assumptions associated with the parallel PW approach would suggest that this is not an approach that should be further pursued.

The PID label approach gives rise to the concerns that it will break implicit behavioral and label-stack size assumptions in many implementations. Whilst those assumptions may be addressed with new hardware, this would delay the introduction of the technology to the point where it is unlikely to gain acceptance in competition with an approach that needs no new protocol design and is already supportable on many existing hardware platforms.

The PID in the CW leads to the most compact protocol stack, is simple, and requires minimal protocol work. However, it is a new forwarding design and, apart from the issue of the larger packet header and the simpler adjacency formation, offers no advantage over the virtual Ethernet.

The above considerations bring us back to the virtual Ethernet, which is a well-known protocol stack with a well-known (internal) client interface. It is already implemented in many hardware platforms and is therefore readily deployable. After considering a number of initially promising alternatives, the authors conclude that the simplicity and existing hardware make the virtual Ethernet approach to the packet PW the most attractive solution.

Authors' Addresses

Stewart Bryant (editor)
Cisco Systems
250, Longwater, Green Park,
Reading, Berks RG2 6GB
UK

EMail: stbryant@cisco.com

Luca Martini
Cisco Systems
9155 East Nichols Avenue, Suite 400
Englewood, CO 80112
USA

EMail: lmartini@cisco.com

George Swallow
Cisco Systems
1414 Massachusetts Ave
Boxborough, MA 01719
USA

EMail: swallow@cisco.com

Andrew G. Malis
Verizon Communications
60 Sylvan Rd.
Waltham, MA 02451
USA

EMail: andrew.g.malis@verizon.com