

PWE3
Internet-Draft
Expires: January 9, 2005

Y(J) Stein
RAD Data Communications
July 11, 2004

UDP based PWs
draft-stein-pwe3-udp-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 9, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

The PWE3 charter permits three PSN types, namely "IP (both versions), L2TP, or MPLS". The majority of drafts submitted to the working group have focused on MPLS PSNs, and lesser attention has been given to L2TPv3. Only the TDMoIP drafts have explicitly specified UDP/IP as PSN, and only they stipulate the use of UDP port numbers as PW labels. This draft discusses PWs based on UDP/IP, concentrating on three issues, namely 1) the advisability of using UDP port numbers as PW labels, 2) which UDP port number to use, and 3) what label distribution protocol should be employed.

Internet-Draft

UDP-PWs

July 2004

Table of Contents

1.	Introduction	3
2.	Using UDP Port Numbers	4
3.	Source or Destination Port	4
4.	Port Number Distribution Protocol	6
5.	Summary	6
6.	Security Considerations	7
7.	IANA Considerations	7
8.	IPR Disclosure Acknowledgement	7
9.	References	8
10.	Acknowledgments	8
	Author's Address	8
	Full Copyright Statement	9

Internet-Draft

UDP-PWs

July 2004

1. Introduction

Although sanctioned as a valid PSN by the PWE3 charter, little work has been carried out concerning PWs over UDP/IP. In fact, use of UDP/IP seems to be mentioned only in two WG drafts.

The PWE3 architecture draft [[PWE3-ARCH](#)] mentions the use of UDP, and states:

PW Demultiplexing is provided by the PW label, which may take the form specified in a number of IETF protocols, e.g. an MPLS label, an L2TP session ID, or a UDP port number.

Of the service-specific drafts, only the TDMoIP drafts have explicitly specified UDP/IP as PSN, and although these drafts have stimulated much comment, their non-standard use of UDP ports has not provoked discussion.

Unlike MPLS and L2TP, UDP/IP does not provide a transport tunneling mechanism within which PWs are placed. However, UDP ports can provide addressing functionality for the PWs themselves, by populating a port field with the PW label.

The use of UDP ports as PW labels has been implemented by several vendors, and is widely deployed for TDM PWs. The use of UDP/IP as the PSN is particularly appropriate for CE to CE PWs [CE2E], since the MPLS network is presently confined to the PEs.

This draft discusses PWs based on UDP/IP, concentrating on three issues, namely

- 1) the fundamental advisability of using UDP port numbers as PW labels, (see [Section 2](#))
- 2) whether to place the PW label in source or destination port (see

[Section 3](#)), and

- 3) the label distribution protocol that should be employed (see [Section 4](#)).

Stein

Expires January 9, 2005

[Page 3]

Internet-Draft

UDP-PWs

July 2004

[2](#). Using UDP Port Numbers

The majority of PWE drafts assume an MPLS (inner) label as PW label. This is because most PWE work have focused on MPLS PSNs, as the basic PWE architecture has been designed for service-provider scenarios. MPLS is not presently available at end-user sites, and so MPLS-based PWs are not directly applicable to CE-CE scenarios (but see [CE2E]). In particular, enterprise applications where native services are converted to IP at the customer premises, require PWs based on pure IP, specifically UDP/IP.

The raison d'etre of the PW demultiplexer label is to allow multiple PWs to be carried in a single tunnel, or for the UDP/IP case, between two end systems. The PW label is situated in the packet overhead, somewhere before the PWE3 Encapsulation Layer, and employs mechanisms pre-existing in the PSN headers.

In addition to IP addresses, UDP/IP and TCP/IP employ 16-bit port numbers. The allocation of port numbers is performed by IANA, with well known port numbers in the range from 0 through 1023, registered port numbers from 1024 through 49151, and dynamic or private port numbers from 49152 through 65535.

As aforementioned, the architecture draft specifically allows the use of UDP port numbers as PW labels. Since UDP always reserves space for source and destination UDP port numbers, exploiting these fields saves inserting additional overhead to the PW packet (as would be required were L2TPv3 to be used).

3. Source or Destination Port

Unfortunately, the TCP and UDP standards themselves (RFCs 793 and 768, respectively) do not mandate methods for port number use. In particular [RFC 768](#) merely states:

Source Port is an optional field, when meaningful, it indicates the port of the sending process, and may be assumed to be the port to which a reply should be addressed in the absence of any other information. If not used, a value of zero is inserted.

Destination Port has a meaning within the context of a particular Internet destination address.

Stein

Expires January 9, 2005

[Page 4]

Internet-Draft

UDP-PWs

July 2004

While the TCP text is more descriptive, it is equally noncommittal:

To allow for many processes within a single Host to use TCP communication facilities simultaneously, the TCP provides a set of addresses or ports within each host. à

The binding of ports to processes is handled independently by each host. However, it proves useful to attach frequently used processes à to fixed sockets which are made known to the public.

These services can then be accessed through the known addresses. Establishing and learning the port addresses of other processes may involve more dynamic mechanisms.

The primary use of TCP and UDP port numbers is application demultiplexing, with requests to an IP server being forwarded to the appropriate application based on the well known or registered destination port number.

VoIP and other RTP-based protocols [[RTP](#)] use the destination port for connection demultiplexing, whereby multiple connections carrying data destined for the same application are distinguished based on their dynamic port numbers. Adopting this strategy would lead us to place

the PW label in the destination port. Such use excludes the capacity to demultiplex applications based on the content of this port. While this loss is not overly problematic, as sessions are setup in an orderly way and packets do not arrive unexpected, this use constitutes a break with tradition.

Using a port number field is superior to adding a non-standard field, since available hardware components can switch based on these fields in the IP header. Were we to desire maintaining the traditional use of destination port as application demultiplexer, we could place the PW label in the source port field. This is equally untraditional, but exploits a 16-bit field that would otherwise be unused. It has the additional advantage of being compatible with enterprise scenarios where NATs and firewalls are employed.

NAPT devices [[RFC 2391](#)] classify incoming packets based on destination IP address and port number. Firewalls [[RFC 2979](#)] may use source IP address as well. A server sitting behind a NAPT or firewall may be accessed based on a known destination port number resulting in packets being correctly routed. NAPT and firewalls that adaptively learn mappings can function properly with dynamically allocated destination port numbers, only when the connection originator sits on their local side.

[4.](#) Port Number Distribution Protocol

A typical client-server protocol, such as telnet [[TELNET](#)], has the server wait for an incoming connection by monitoring a well-known port W. The client process, before initiating a connection, selects an unused port number P. It then sends a packet to the server with destination port W and source port P. The server's responses are sent using destination port P and source port W. Thus, a bidirectional but asymmetric connections are established without any explicit setup protocol.

This idea can be extended to the more symmetric case. For example, in tftp [[TFTP](#)] both sides select identifiers, say P and Q. The initiator sends a packet with destination port set to well-known port

W and source set to P. The response has destination P and source Q, and only P and Q are used for the duration of the session.

In both of the above cases the port numbers were chosen by their source (upstream allocation). Downstream allocation is also used, for example by passive mode ftp [FTP]. Here the client issues a command to a well-known port and the server responds with a message containing the port number P to be used.

As an alternative to such implicit port number allocations, UDP/IP PWs may be set up using an explicit control protocol. An obvious possibility is the standard PWE control protocol based on targeted LDP. This has the advantage of settling on a single control protocol for all PW types except L2TPv3, and could be extended to additional cases (e.g. allocation of VLAN tags [VLAN]). It has the shortcoming of using LDP in a non-MPLS environment. Only minor changes would be required to the LDP protocol for this use, mainly related to the different label ranges (UDP port numbers are only 16 bits in length, VLAN tags are 12 bits.)

[5.](#) Summary

The main questions raised in this draft are herein summarized.

Should UDP/IP PSNs, specified in the PWE charter and the architecture draft, continue to be supported?

Should the PW label be placed in the source or destination UDP port?
Do we need registered port numbers for PWs?

What method should be used for label selection and distribution?
Should the label selection be upstream or downstream? Should an implicit or explicit distribution protocol be used? Do we want to standardize a single PWE control protocol?

[6.](#) Security Considerations

PWs do not enhance or detract from the security performance of the underlying PSN, and PWs running over UDP/IP will suffer from the same weaknesses as any other traffic on the same network.

PW labels used as UDP port numbers that are known or can be readily

surmised may pose security threats.

7. IANA Considerations

Depending on the label distribution protocol, use of UDP ports will require allocation of registered port numbers for the various PW types. 0x085E (2142) has already been assigned by IANA to TDMoIP PWs.

8. IPR Disclosure Acknowledgement

By submitting this Internet-Draft, we certify that any applicable patent or other IPR claims of which we are aware have been disclosed, and any of which we become aware will be disclosed, in accordance with [RFC 3668](#).

9. References

- [L2TPv3] [draft-ietf-l2tpext-l2tp-base-10.txt](#) (08/03) Layer Two Tunneling Protocol (L2TPv3), J. Lau et al., work in progress
- [MPLS] [RFC 3032](#) MPLS Label Stack encoding
- [PWE3-ARCH] [draft-ietf-pwe3-arch-07.txt](#) (3/04), PWE3 Architecture, Stewart Bryant et al, work in progress
- [PWCE2E] [draft-stein-pwe3-pwce2e-00.txt](#) (3/04), Pseudowire Customer Edge to Customer Edge Emulation, Y(J) Stein, work in progress
- [RTP] [RFC 3550](#) RTP: Transport Protocol for Real-Time Applications
- [TCP] [RFC 793](#) Transmission Control Protocol (TCP)
- [TELNET] [RFC 854](#) Telnet Protocol Specification
- [TFTP] [RFC 1350](#) The TFTP Protocol (Revision 2)
- [UDP] [RFC 768](#) User Datagram Protocol (UDP)
- [VLAN] IEEE 802.1Q, IEEE Standards for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks (2003)

10. Acknowledgments

The author would like to thank Sasha Vainshtein and Stewart Bryant for interesting discussion and valuable contributions to the options herein presented.

Author's Address

Yaakov (Jonathan) Stein
RAD Data Communications
24 Raoul Wallenberg St., Bldg C
Tel Aviv 69719
ISRAEL

Phone: +972 3 645-5389
EMail: yaakov_s@rad.com

Internet-Draft

UDP-PWs

July 2004

Full Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Acknowledgement

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

Stein

Expires January 9, 2005

[Page 9]