

Network Working Group
Request for Comments: 5085
Category: Standards Track

T. Nadeau, Ed.
C. Pignataro, Ed.
Cisco Systems, Inc.
December 2007

Pseudowire Virtual Circuit Connectivity Verification (VCCV):
A Control Channel for Pseudowires

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.

This document describes Virtual Circuit Connectivity Verification (VCCV), which provides a control channel that is associated with a pseudowire (PW), as well as the corresponding operations and management functions (such as connectivity verification) to be used over that control channel. VCCV applies to all supported access circuit and transport types currently defined for PWs.

Table of Contents

1.	Introduction	3
1.1.	Specification of Requirements	5
2.	Abbreviations	5
3.	Overview of VCCV	6
4.	CC Types and CV Types	8
5.	VCCV Control Channel for MPLS PWs	10
5.1.	VCCV Control Channel Types for MPLS	10
5.1.1.	In-Band VCCV (Type 1)	11
5.1.2.	Out-of-Band VCCV (Type 2)	12
5.1.3.	TTL Expiry VCCV (Type 3)	12
5.2.	VCCV Connectivity Verification Types for MPLS	13
5.2.1.	ICMP Ping	13
5.2.2.	MPLS LSP Ping	13
5.3.	VCCV Capability Advertisement for MPLS PWs	13
5.3.1.	VCCV Capability Advertisement LDP Sub-TLV	14
6.	VCCV Control Channel for L2TPv3/IP PWs	15
6.1.	VCCV Control Channel Type for L2TPv3	16
6.2.	VCCV Connectivity Verification Type for L2TPv3	17
6.2.1.	L2TPv3 VCCV using ICMP Ping	17
6.3.	L2TPv3 VCCV Capability Advertisement for L2TPv3	17
6.3.1.	L2TPv3 VCCV Capability AVP	17
7.	Capability Advertisement Selection	19
8.	IANA Considerations	19
8.1.	VCCV Interface Parameters Sub-TLV	19
8.1.1.	MPLS VCCV Control Channel (CC) Types	19
8.1.2.	MPLS VCCV Connectivity Verification (CV) Types	20
8.2.	PW Associated Channel Type	21
8.3.	L2TPv3 Assignments	21
8.3.1.	Control Message Attribute Value Pairs (AVPs)	21
8.3.2.	Default L2-Specific Sublayer Bits	21
8.3.3.	ATM-Specific Sublayer Bits	21
8.3.4.	VCCV Capability AVP Values	22
9.	Congestion Considerations	23
10.	Security Considerations	24
11.	Acknowledgements	25
12.	References	26
12.1.	Normative References	26
12.2.	Informative References	26

1. Introduction

There is a need for fault detection and diagnostic mechanisms that can be used for end-to-end fault detection and diagnostics for a Pseudowire, as a means of determining the PW's true operational state. Operators have indicated in [RFC4377] and [RFC3916] that such a tool is required for PW operation and maintenance. This document defines a protocol called Virtual Circuit Connectivity Verification (VCCV) that satisfies these requirements. VCCV is, in its simplest description, a control channel between a pseudowire's ingress and egress points over which connectivity verification messages can be sent.

The Pseudowire Edge-to-Edge Emulation (PWE3) Working Group defines a mechanism that emulates the essential attributes of a telecommunications service (such as a T1 leased line or Frame Relay) over a variety of Packet Switched Network (PSN) types [RFC3985]. PWE3 is intended to provide only the minimum necessary functionality to emulate the service with the required degree of faithfulness for the given service definition. The required functions of PWs include encapsulating service-specific bit streams, cells, or PDUs arriving at an ingress port and carrying them across an IP path or MPLS tunnel. In some cases, it is necessary to perform other operations, such as managing their timing and order, to emulate the behavior and characteristics of the service to the required degree of faithfulness.

From the perspective of Customer Edge (CE) devices, the PW is characterized as an unshared link or circuit of the chosen service. In some cases, there may be deficiencies in the PW emulation that impact the traffic carried over a PW and therefore limit the applicability of this technology. These limitations must be fully described in the appropriate service-specific documentation.

For each service type, there will be one default mode of operation that all PEs offering that service type must support. However, optional modes have been defined to improve the faithfulness of the emulated service, as well as to offer a means by which older implementations may support these services.

Figure 1 depicts the architecture of a pseudowire as defined in [RFC3985]. It further depicts where the VCCV control channel resides within this architecture, which will be discussed in detail shortly.




Figure 1: PWE3 VCCV Operation Reference Model

From Figure 1, Customer Edge (CE) routers CE1 and CE2 are attached to the emulated service via Attachment Circuits (ACs), and to each of the Provider Edge (PE) routers (PE1 and PE2, respectively). An AC can be a Frame Relay Data Link Connection Identifier (DLCI), an ATM Virtual Path Identifier / Virtual Channel Identifier (VPI/VCI), an Ethernet port, etc. The PE devices provide pseudowire emulation, enabling the CEs to communicate over the PSN. A pseudowire exists between these PEs traversing the provider network. VCCV provides several means of creating a control channel over the PW, between the PE routers that attach the PW.

Figure 2 depicts how the VCCV control channel is associated with the pseudowire protocol stack.

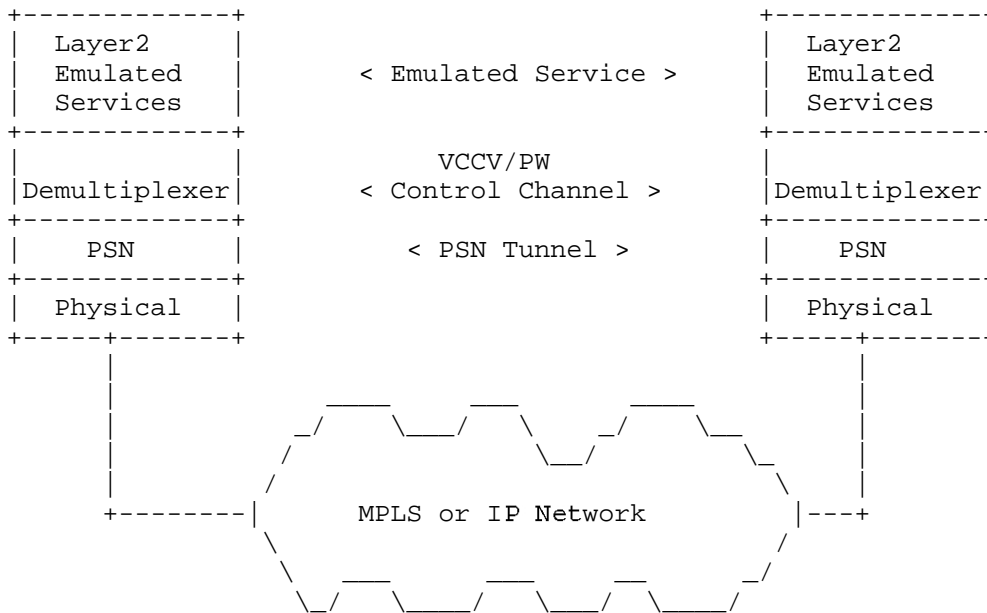


Figure 2: PWE3 Protocol Stack Reference Model including the VCCV Control Channel

VCCV messages are encapsulated using the PWE3 encapsulation as described in Sections 5 and 6, so that they are handled and processed in the same manner (or in some cases, a similar manner) as the PW PDUs for which they provide a control channel. These VCCV messages are exchanged only after the capability (expressed as two VCCV type spaces, namely the VCCV Control Channel and Connectivity Verification Types) and desire to exchange such traffic has been advertised between the PEs (see Sections 5.3 and 6.3), and VCCV types chosen.

1.1. Specification of Requirements

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. Abbreviations

- AC Attachment Circuit [RFC3985].
- AVP Attribute Value Pair [RFC3931].
- CC Control Channel (used as CC Type).

CE Customer Edge.

CV Connectivity Verification (used as CV Type).

CW Control Word [RFC3985].

L2SS L2-Specific Sublayer [RFC3931].

LCCE L2TP Control Connection Endpoint [RFC3931].

OAM Operation and Maintenance.

PE Provider Edge.

PSN Packet Switched Network [RFC3985].

PW Pseudowire [RFC3985].

PW-ACH PW Associated Channel Header [RFC4385].

VCCV Virtual Circuit Connectivity Verification.

3. Overview of VCCV

The goal of VCCV is to verify and further diagnose the pseudowire forwarding path. To this end, VCCV is comprised of different components:

- o a means of signaling VCCV capabilities to a peer PE,
- o an encapsulation for the VCCV control channel messages that allows the receiving PE to intercept, interpret, and process them locally as OAM messages, and
- o specifications for the operation of the various VCCV operational modes transmitted within the VCCV messages.

When a pseudowire is first signaled using the Label Distribution Protocol (LDP) [RFC4447] or the Layer Two Tunneling Protocol version 3 (L2TPv3) [RFC3931], a message is sent from the initiating PE to the receiving PE requesting that a pseudowire be set up. This message has been extended to include VCCV capability information (see Section 4). The VCCV capability information indicates to the receiving PE which combinations of Control Channel (CC) and Connectivity Verification (CV) Types it is capable of receiving. If the receiving PE agrees to establish the PW, it will return its capabilities in the subsequent signaling message to indicate which CC

and CV Types it is capable of processing. Precedence rules for which CC and CV Type to choose in cases where more than one is specified in this message are defined in Section 7 of this document.

Once the PW is signaled, data for the PW will flow between the PEs terminating the PW. At this time, the PEs can begin transmitting VCCV messages based on the CC and CV Type combinations just discussed. To this end, VCCV defines an encapsulation for these messages that identifies them as belonging to the control channel for the PW. This encapsulation is designed to both allow the control channel to be processed functionally in the same manner as the data traffic for the PW in order to faithfully test the data plane for the PE, and allow the PE to intercept and process these VCCV messages instead of forwarding them out of the AC towards the CE as if they were data traffic. In this way, the most basic function of the VCCV control channel is to verify connectivity of the pseudowire and the data plane used to transport the data path for the pseudowire. It should be noted that because of the number of combinations of optional and mandatory data-plane encapsulations for PW data traffic, VCCV defines a number of Control Channel (CC) and Connectivity Verification (CV) types in order to support as many of these as possible. While designed to support most of the existing combinations (both mandatory and optional), VCCV does define a default CC and CV Type combination for each PW Demultiplexer type, as will be described in detail later in this document.

VCCV can be used both as a fault detection and/or a diagnostic tool for pseudowires. For example, an operator can periodically invoke VCCV on a timed, on-going basis for proactive connectivity verification on an active pseudowire, or on an ad hoc or as-needed basis as a means of manual connectivity verification. When invoking VCCV, the operator triggers a combination of one of its various CC Types and one of its various CV Types. The CV Types include LSP Ping [RFC4379] for MPLS PWs, and ICMP Ping [RFC0792] [RFC4443] for both MPLS and L2TPv3 PWs. We define a matrix of acceptable CC and CV Type combinations further in this specification.

The control channel maintained by VCCV can additionally carry fault detection status between the endpoints of the pseudowire. Furthermore, this information can then be translated into the native OAM status codes used by the native access technologies, such as ATM, Frame-Relay or Ethernet. The specific details of such status interworking is out of the scope of this document, and is only noted here to illustrate the utility of VCCV for such purposes. Complete details can be found in [MSG-MAP] and [RFC4447].

4. CC Types and CV Types

The VCCV Control Channel (CC) Type defines several possible types of control channel that VCCV can support. These control channels can in turn carry several types of protocols defined by the Connectivity Verification (CV) Type. VCCV potentially supports multiple CV Types concurrently, but it only supports the use of a single CC Type. The specific type or types of VCCV packets that can be accepted and sent by a router are indicated during capability advertisement as described in Sections 5.3 and 6.3. The various VCCV CV Types supported are used only when they apply to the context of the PW demultiplexer in use. For example, the LSP Ping CV Type should only be used when MPLS Labels are utilized as PW Demultiplexer.

Once a set of VCCV capabilities is received and advertised, a CC Type and CV Type(s) that match both the received and transmitted capabilities can be selected. That is, a PE router needs to only allow Types that are both received and advertised to be selected, performing a logical AND between the received and transmitted bitflag fields. The specific CC Type and CV Type(s) are then chosen within the constraints and rules specified in Section 7. Once a specific CC Type has been chosen (i.e., it matches both the transmitted and received VCCV CC capability), transmitted and replied to, this CC Type MUST be the only one used until such time as the pseudowire is re-signaled. In addition, based on these rules and the procedures defined in Section 5.2 of [RFC4447], the pseudowire MUST be re-signaled if a different set of capabilities types is desired. The relevant portion of Section 5.2 of [RFC4447] is:

Interface Parameter Sub-TLV

Note that as the "interface parameter sub-TLV" is part of the FEC, the rules of LDP make it impossible to change the interface parameters once the pseudowire has been set up.

The CC and CV Type indicator fields are defined as 8-bit bitmasks used to indicate the specific CC or CV Type or Types (i.e., none, one, or more) of control channel packets that may be sent on the VCCV control channel. These values represent the numerical value corresponding to the actual bit being set in the bitfield. The definition of each CC and CV Type is dependent on the PW type context, either MPLS or L2TPv3, within which it is defined.

以上内容仅为本文档的试下载部分，为可阅读页数的一半内容。如要下载或阅读全文，请访问：<https://d.book118.com/688107051111006101>