

Internet-Draft
Expires: July 2004

Juha Heinanen (Tutpro)
W. Mark Townsley (Cisco)
Stephen Bailey (Sandburst)

Radius/L2TP Based VPLS
[draft-ietf-l2vpn-l2tp-radius-vpls-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>.

Copyright Notice

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

Abstract

This memo describes a simple mechanism to implement provider provisioned Virtual Private LAN Service (VPLS) using Radius for PE discovery and L2TP as the control and data plane protocol.

Table Of Contents

1.	Introduction	2
2.	Service Description	2
3.	Adding a Site to a VPN	3
3.1.	Configuration Actions	3
4.	Connecting a Site to a VPN	3

4.1.	Configuration Actions	3
4.2.	Protocol Actions	3
5.	Disconnecting a Site from a VPN	5
5.1.	Configuration Actions	5
5.2.	Protocol Actions	5
6.	Removing a Site from a VPN	6
6.1.	Configuration Actions	6
7.	Failure Recovery	6
8.	Exponential Back-off Behavior	7
9.	Data Plane	7
10.	Security Considerations	7
	References	8
	Authors' Addresses	8
	Full Copyright Statement	9

[1.](#) Introduction

This memo describes a simple mechanism to implement provider provisioned Virtual Private LAN Service (VPLS) [[1](#)] using Radius [[5](#)] as the PE discovery protocol and L2TPv3 [[3](#)] as the control and data plane protocol. Radius is deployed as described in [[2](#)], whereas L2TP is deployed as described in [[4](#)] with minor changes.

An advantage of a directory (such as Radius) based discovery solution for provider based VPNs is that it doesn't require BGP implementation or configuration complexity in the PE routers and can be easily deployed also in inter-AS cases where the VPN sites are attached to PEs in more than one AS. An advantage of Radius as a directory protocol is that it has been in Internet-wide use for years and can thus be deployed without a new directory infrastructure.

A similar directory based VPLS solution could be specified that uses LDP for signaling and MPLS label stack encapsulation for data transport. An L2TP based solution may, however, be preferable to providers who are already familiar with L2TP and are not deploying MPLS. An L2TP based solution may also be considered simpler to manage, because L2TP tunnels are bidirectional and because L2TP bundles control, data, and management planes in a single protocol.

[2.](#) Service Description

This memo supports VPLS service in a mode where each VPLS instance (also called VPN for short) connects one or more CEs (also called VPN sites) to a common virtual LAN. A VPN site can use either 802.1q tagged or untagged (but not both) Ethernet frames to communicate with the other sites of the VPN. In case of tagged

frames, each VPN site MUST use a single VLAN ID for the same VPN, but the VLAN ID MAY differ at each VPN site.

VPLS service MAY support Differentiated Services treatment of tagged or untagged Ethernet frames. In case of tagged frames, the desired treatment of the frame is coded in the 802.1p User Priority field. In case of untagged frames, all frames sent by a site receive a default treatment. Differentiated Services treatment as well as mapping of 802.1p User Priority values to DiffServ code points of L2TP tunnels is VPLS specific and outside the scope of this memo.

[3.](#) Adding a Site to a VPN

[3.1.](#) Configuration Actions

A site (a CE) is added to a VPN (a VPLS instance) by adding its "user name", password, and VPN identifier record, for example:

```
<SiteX@vpnY.domainZ.net, secret, vpnY.domainZ.net>
```

to Radius database as described in [\[2\]](#). After this configuration action the site can be connected to the VPN at a PE.

[4.](#) Connecting a Site to a VPN

[4.1.](#) Configuration Actions

No configuration actions are needed if a site connects to the VPN at a PE using a dynamic authentication protocol, such as 801.1x/EAP. Otherwise, the Ethernet interface of the PE to which the site is going to be connected to MUST be configured with the "user name" and password of the site, for example:

```
<providerP/SiteX@vpnY.domainZ.net, secret>
```

The provider prefix is only needed in case the site is connected to a PE of a provider that is not the administrative owner of the VPN (providerP in the above example).

The interface to which the site is connected to MAY be 802.1Q tagged or untagged. In the former case, the VLAN ID that is used to connect the site to the VPN MUST be specified.

[4.2.](#) Protocol Actions

The following protocol actions take place at the PE when a new VPN site tries to authenticate itself with the PE or when the provider

has configured a new VPN site to the PE:

- (1) The PE issues Radius Access-Request for the CE as described in [\[2\]](#). If access is granted, the PE learns the identifier of the CE's VPN and IP addresses of the VPN's PEs.
- (2) If the PE already has site(s) that belong to the same VPN as the new site, no other protocol actions take place at the PE.
- (3) Otherwise the PE establishes an L2TP Control Connection with each of the other PEs of the VPN unless one already exists. The Pseudo Wire Capabilities List AVP of the Control Connection MUST contain this and only this value:

0xTBD - Sessions without control word for connecting Ethernet VLANs are allowed

- (4) The PE establishes for this VPN an L2TP session with each of the remote PEs unless one already exists. L2TP sessions are established as defined in section 2.2 of [\[4\]](#) with the following changes and clarifications:

L2TP sessions are established as for Incoming Calls using ICRQ/ICRP/ICCN message exchange (see section 3.4.1 of [\[3\]](#)).

The Pseudo Wire Type AVP MUST have in its Attribute Value field value:

0xTBD - Ethernet VLAN

The Application ID AVP MUST have in its Application Code field value:

0xTBD - Radius/L2TP based VPLS

The End Identifier AVP MUST have in its Attribute Value field the domain name of the VPN:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| VPN Identifier (e.g. vpnY.domainZ.net) ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

The following protocol actions take place in sequence at a PE when it receives an L2TP Incoming-Call-Request from another PE for the application described in this document:

- (1) The PE checks that itself and the other PE belong to the VPN indicated by the End Identifier AVP. If the other PE is not included in the PE's current list of other PEs of the VPN, the PE issues an Access-Request request for an up to date list. If the check fails, the PE responds with a Call-Disconnect-Notify and no other protocol actions take place at the PE.

The Call-Disconnect-Notify MUST include a Result Code AVP with Error Code and Error Message fields. The Result Code MUST have the value 0x0002 (Session disconnected for the reason indicated in Error Code) and the <Error Code, Error Message> MUST have one of the two values:

```
<0xTBD, "Requesting PE does not belong to the VPN">
<0xTBD, "Requested PE does not belong to the VPN">
```

- (2) The PE checks if it already has an L2TP session with the calling PE for the VPN indicated by the End Identifier AVP. If so, the PE responds with a Call-Disconnect-Notify and no other protocol actions take place at the PE.

The Call-Disconnect-Notify MUST include a Result Code AVP with Error Code and Error Message fields. The Result Code MUST have the value 0x0002 (Session disconnected for the reason indicated in Error Code) and the <Error Code, Error Message> MUST have the value:

<0xTBD, "Session already exists for the VPN">

- (3) Otherwise the PE accepts the request with an Incoming-Call-Reply.

[5.](#) Disconnecting a Site from a VPN

[5.1.](#) Configuration Actions

When a site (CE) is to be disconnected from a VPN at a PE the "user name" and password of the site is unconfigured from the Ethernet interface to which it has been connected to.

[5.2.](#) Protocol Actions

The following protocol actions take place in sequence at the PE of the disconnected site:

- (1) The PE issues a Stop Accounting-Request as described in 5.3 of [\[2\]](#).
- (2) If the disconnected site was the last site of the VPN at the PE, the PE tears down any existing L2TP sessions for the VPN by sending each remote PE a Call-Disconnect-Notify.

The Call-Disconnect-Notify MUST include a Result Code AVP with Error Code and Error Message fields. The Result Code MUST have the value 0x0002 (Session disconnected for the reason indicated in Error Code) and the <Error Code, Error Message> MUST have the value:

<0xTBD, "Requesting PE does not anymore belong to the VPN">

When a PE receives a Call-Disconnect-Notify from another PE for the application described in this memo, no other protocol actions than normal clean up of the corresponding L2TP session are needed at the PE.

If the L2TP session that was torn down between two PEs was the last session associated with the Control Connection, either PE MAY tear down the Control Connection.

[6.](#) Removing a Site from a VPN

[6.1.](#) Configuration Actions

A site (a CE) is removed from a VPN (a VPLS instance) by removing its

<CE user name, password, VPN identifier>

record from Radius database. This configuration action MUST succeed only if Radius does not have a

<VPN identifier, PE IP address, CE user name>

record in its database where CE user name belongs to the removed CE. This is true if the site has been first disconnected from the VPN as described in [section 5](#).

[7.](#) Failure Recovery

If a PE loses its Control Connection with another PE having site(s) in a common VPN, the PE tries to re-establish the Control Connection until (a) the Control Connection gets re-established or (b) this PE or the other PE no longer have site(s) in this VPN.

Once the Control Connection gets re-established, the PE re-establishes an L2TP session with the other PE for this VPN as described in [section 4.2](#).

If an L2TP session gets teared down between two PEs and they still have site(s) in the VPN of the teared down session, the two PEs try to re-establish the session as described in [section 4.2](#) as long as the two PEs have site(s) in the VPN of the teared down session.

When a PE recovers from a crash, it adds each of the configured VPN site(s) to their respective VPN(s) as described in [section 4.2](#).

[8.](#) Exponential Back-off Behavior

If any protocol action does not succeed immediately, normal behavior is that the PE keeps on trying with exponential back-off until the action either succeeds or becomes invalid due to a change in VPN configuration. If the protocol action fails for an implementation specific prolonged period of time, the PE SHOULD notify the "owner" of the VPN about the problem via a management action.

[9.](#) Data Plane

The PEs that host the sites of a VPN act as virtual, fully connected learning bridges for the VPN.

When a PE receives a Ethernet frame from a CE for a particular VPN, it adds to it a 802.1q tag (if not already present) and sets the VLAN ID to zero. Treatment of the 802.1p User Priority field is VPLS specific and outside the scope of this memo.

When a PE needs to send an Ethernet frame to a VPN site connected to it, it either overwrites the VLAN ID with the VLAN ID used by the site for this VPN or removes the 802.1q tag if the interface of the VPN site is untagged. Treatment of the 802.1p User Priority field is VPLS specific and outside the scope of this memo.

When a PE needs to send an Ethernet frame to another PE, the PE processes the frame as described in section 3 of [\[4\]](#) using the L2TP session established for this VPLS instance. Mapping of the 802.1p User priority value to DiffServ code point of the L2TP packet is VPLS specific and outside the scope of this memo.

[10.](#) Security Considerations

Security of Radius/L2TP based VPNs depends on security of Radius and L2TP. Security of Radius is covered in [section 8](#) or [\[5\]](#) and

security of L2TP is covered in section 8 of [\[3\]](#).

11. References

- [1] Augustyn, et al., "Requirements for Virtual Private LAN Services (VPLS)". [draft-ietf-l2vpn-vpls-requirements-00.txt](#), October 2002.
- [2] Heinanen, "Using Radius for PE-Based VPN Discovery". [draft-heinanen-radius-pe-discovery-04.txt](#), June 2003.
- [3] Lau, et al., "Layer Two Tunneling Protocol (Version 3) "L2TPv3"". [draft-ietf-l2tpext-l2tp-base-11.txt](#), October 2003.
- [4] Aggarwal, et al., "Transport of Ethernet Frames over L2TPv3". [draft-ietf-l2tpext-pwe3-ethernet-01.txt](#), October 2002.
- [5] Rigney, et al., "Remote Authentication Dial In User Service (RADIUS)". [RFC 2865](#), June 2000.

Authors' Addresses

Juha Heinanen
TutPro Inc.
Utsjoki, Finland

Email: jh@tutpro.com

W. Mark Townsley
Cisco Systems
7025 Kit Creek Road
PO Box 14987
Research Triangle Park, NC 27709

Email: mark@townsley.net

Stephen Bailey
Sandburst Corporation
600 Federal Street
Andover, MA 01810 USA
USA

Phone: +1 978 689 1614
Email: steph@sandburst.com

Full Copyright Statement

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

