Internet Engineering Task Force
INTERNET DRAFT
Expires March 2003

Juha Heinanen
Song Networks
September, 2002

DNS/L2TP Based VPLS
<draft-heinanen-dns-l2tp-vpls-01.txt>

Status of this Memo

Abstract

This memo describes a simple mechanism to implement provider provisioned Virtual Private LAN Service (VPLS) using DNS and L2TP as discovery, control, and data plane protocols.

1. Introduction

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

An advantage of a directory (such as DNS) 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 DNS as a directory is that it has been in Internet-wide use for years and can thus be deployed without any new standardization effort.

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.

Although this memo proposed the use of DNS as the discovery mechanism, the described VPLS solution itself is not DNS specific. Later memos may thus specify how other protocols, such as Radius or LDAP, can replace DNS as the VPLS discovery mechanism.

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. Addition of Sites

3.1 Configuration Actions

A DNS/L2TP based VPLS is very easy to provision. Only the following two configuration actions are needed when a new site (CE) is added to a VPN (VPLS instance):

(1) If the PE does not previously connect any sites of this VPN, the IP address (A record) of the PE is added to the directory (DNS in this memo) under domain name

vpn-name.domain

The label "vpn-name" uniquely identifies the VPN within "domain", which belongs to the administrative "owner" of the VPN. An example of the domain name of a VPN is bobsVpn.serviceProvider.net.

(2) An Ethernet "interface" of the PE to which the site is connected to is configured to belong to the VPN by associating the interface with the domain name of the VPN. The interface 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.

Note that also in the case of a multi-provider VPN, the administrative "owner" of the VPN is the single body that operates the directory for the VPN zone. The "owner" of a VPN MAY choose to make all updates to the zone data of the VPN by itself or MAY allow other providers to dynamically update the zone data.

3.2 Protocol Actions

After the above configuration actions, the following protocol actions take place in sequence at the PE of the new site if the PE of the new site doesn't previously connect any site(s) of the VPN:

(1) The PE of the new site checks that its own IP address has become available in the directory under the domain name of the VPN.

(2) The PE of the new site queries the directory for IP addresses of the other (remote) PEs of the VPN.

(3) The PE of the new site establishes an L2TP Control Connection with each the remote PEs unless one already exists. The Pseudo Wire Capabilities List AVP of the Control Connection MUST contain this and only this value:

0x0004 - Sessions without control word for connecting Ethernet VLANs are allowed

The Control Connection Tie-Breaker AVP MUST be used for tie-breaking.

(4) The PE of the new site 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 6.2.1.2 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

0x0004 - Ethernet VLAN

The Application Identifier AVP MUST have in its Attribute Value field value

0xXXXX - Directory/L2TP based VPLS

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

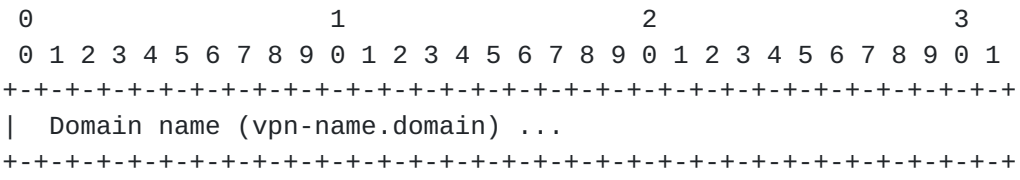

The Application Identifier AVP and End Identifier AVP thus replace the PW ID and Group ID AVPs used in [4].

The Session Tie-Breaker AVP MUST be used for tie-breaking.

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 from the directory that the other PE belongs to the VPN indicated by the End Identifier AVP and that it itself has at least one site in that VPN. If not, 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 0xXXXX (Session disconnected for the application specific reason indicated in Error Code) and the <Error Code, Error Message> MUST have one of the two values

<0x0001, "Requesting PE does not belong to the VPN">
<0x0002, "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 0xXXXX(Session disconnected for the application specific reason indicated in Error Code) and the <Error Code, Error Message> MUST have the value

<0x0003, "Session already exists for the VPN">

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

If the PE of the new site already connects site(s) of this VPN, no protocol actions take place at either the PE of the new site or at the remote PEs.

4. Removal of Sites

4.1 Configuration Actions

The following configuration actions are needed when an existing site (CE) is removed from a VPN:

(1) If the site to be removed is the last site of the VPN at the PE, the IP address of the PE is removed from the directory under the domain name of the VPN.

(2) The site is removed from the VPN by unconfiguring the VPN from the "interface" of the PE to which the site is connected to.

4.2 Protocol Actions

After the above configuration actions, the following protocol actions take place in sequence at the PE of the removed site if the removed site was the last site of the VPN at the PE:

(1) The PE checks that its IP address no longer exists in the directory under the domain name of the VPN.

(2) 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 0xXXXX (Session disconnected for the application specific reason indicated in Error Code) and the <Error Code, Error Message> MUST have the value

<0x0004, "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.

5. 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 3.2.

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 3.2 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 3.2.

6. 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.

7. 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 6.1 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.

8. Security Considerations

Security of DNS/L2TP based VPNs depends on security of DNS and L2TP. Security of DNS is covered in section 9 or [2] and security of L2TP is covered in section 9 of [3].

Acknowledgements

I would like to thank Mark Townsley of Cisco Systems for his expertise and constructive comments during the development of this memo.

References

[1] Augustyn, et al., "Requirements for Virtual Private LAN Services (VPLS)". draft-ietf-ppvpn-vpls-requirements-00.txt, March 2002.

[2] Luciani et al., "Using DNS for VPN Discovery". draft-luciani-ppvpn-vpn-discovery-02.txt, March 2002.

[3] Lau, et al., "Layer Two Tunneling Protocol (Version 3) "L2TPv3"". draft-ietf-l2tpext-l2tp-base-03.txt, June 2002.

[4] So, et al., Ethernet Pseudo Wire Emulation Edge-to-Edge (PWE3). draft-so-pwe3-ethernet-01.txt, March 2002.

Author's Address

Juha Heinanen
Song Networks, Inc.
Hallituskatu 16
33200 Tampere, Finland
Email: jh@song.fi

Full Copyright