Provider Provisioned VPN Working Group
Internet Draft
draft-knight-l2vpn-lpe-ad-00.txt
Expires: April 2002

Paul Knight
Dinesh Mohan
Hamid Ould-Brahim
Vasile Radoaca
Nortel Networks, Inc.

November 2001

Logical PE Auto-Discovery Mechanism

Status of this Memo

Abstract

This document describes a lightweight protocol for VPLS information exchange between Logical PE components, consisting of the PE-Edge and PE-Core.

Table of Contents

1. Conventions used in this document

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].

2. Overview

Within the "Logical PE" architecture described in [LPE-ARCH], the functionality of the Provider-Edge devices (PEs) is distributed between low-cost Ethernet-based devices called "PE-Edges" and high-capacity core devices, labeled "PE-Cores." This distribution of functionality improves PE scalability and decreases costs to the service provider.

The logical PE (LPE) architecture supports Virtual Private LAN Services (VPLS) as described in [RFC-2764]. It can be adapted to a core network running either [L2-MART] based l2vpn architecture or [L2-KOMP] architecture, or even a core network running both architectures.

In order to support the LPE architecture, a lightweight protocol needs to exchange the VPLS-related information (membership, end-point identification, etc.) between the PE-Edges and PE-Cores. It should be simple enough to be implemented in low-cost PE-Edge devices.

The Logical PE Auto-Discovery (LPE-AD) protocol described in this draft is intended to perform these functions. LPE-AD allows a group of distributed PE-Edges and PE-Cores to exchange the VPLS information necessary to function as a unified Logical PE device within a service provider's network.

The basic functions provided by LPE-AD include:

- Notifying PE-Cores within the LPE of addition, deletion, or modification of VPLS membership on customer-facing interfaces (LPE endpoints) of the PE-Edges
- Notifying PE-Cores of the addition, deletion, or modification of PE-Edges

3. Logical PE Architecture

A detailed architecture and reference model of the "Logical PE" is presented in [LPE-ARCH]. This draft uses much of the terminology and definitions developed in that document.

The LPE model offers a way of achieving both scaling and cost objectives by distributing the PE's VPLS functions among low-cost Ethernet-based customer-facing devices and high capacity devices attached to the Service Provider's core network. The LPE model defines devices containing customer-facing interfaces, called "PE-Edges", and core provider edge devices, labeled as "PE-Cores". (The architecture does not preclude the PE-Core from containing customer-facing interfaces. In this case the PE-Core contains all the functions of a "normal" PE.)

In general, the PE-Edge manages the Service Level Agreement (SLA), including bandwidth policing, traffic classification and Quality of Service (QoS). It also manages customer Ethernet frame encapsulation and MAC learning.

The PE-Core maintains membership information. It communicates VPLS membership information to all the other PE-Core members of the same VPLS.

4. Logical PE Auto-Discovery Protocol (LPE-AD)

This section describes the requirements for LPE-AD, describes the requirements for messages used by the protocol, describes a simple LPE-AD protocol, and provides operational scenarios explaining the use of the protocol.

4.1 Terminology

The following terminology is used in this draft:

Customer Site - An entity connected to the Service Provider's network through a Customer Edge (CE) device. The purpose of Virtual Private LAN Service (VPLS) is to interconnect multiple Customer Sites associated with a single Customer. Multiple Customer Sites within one LPE should use the same CSM-ID in order to belong to a common VPLS. A single Customer Site may belong to multiple VPLSs.

Customer Site Member Identifier (CSM-ID) - A local VPN identifier within the LPE. All LPE endpoints belonging to the same VPLS within a specific LPE will be configured with the same CSM-ID. It can be exactly the same as a VPN-ID, or it may be an identifier with significance only within the LPE. A CSM-ID instance is associated with a LPE Endpoint through provisioning. More than one CSM-ID can be provisioned on a given LPE Endpoint, if VLAN tagging is used on the CE-to-LPE Endpoint link to distinguish traffic for different

VPLSs. Each CSM-ID is associated with a VPN-ID. Every instance of a specific CSM-ID MUST be associated with the same Customer VPLS (i.e., be associated with the same VPN-ID) within a single LPE. A state transition of a CSM-ID should be signaled via the LPE-AD protocol.

LPE Endpoint - Identifies a CE-facing Ethernet port in a PE-Edge or PE-Core, where VPLS instances can terminate. It acts as the demarcation to a CE device. It is associated with the actual physical interface. A state transition of the LPE Endpoint with active CSM-IDs should be signaled via the LPE-AD protocol.

PE-Edge-ID - A PE-Edge is defined within the LPE by a PE-Edge-ID, which is provisioned. The PE-Edge-ID value is an IP address.

SET domain - A Switched Ethernet Transport domain; a single Ethernet broadcast domain, within a single Service Provider's network. It may also refer to a single 802.1Q VLAN broadcast domain. The LPE-AD protocol described in this draft is designed specifically to use a SET as the transport between PE-Edge and PE-core within a LPE.

VPN-ID - VPN Identifier, is a unique identifier of a VPN within a service provider context. This can be a standard VPN-ID for the core VPN technology used by the service provider, as in [RFC 2685]. Within a LPE, a CSM-ID is mapped one-to-one to a VPN-ID.

4.2 LPE-AD - Provisioning and Distribution

LPE-AD facilitates the auto-discovery and distribution of the VPLS membership information inside the LPE. Although other mechanisms such as Edge LDP (using LDP Downstream Unsolicited label advertisements [RFC 3036]) or BGP extensions (BGP Auto-Discovery, in [BPG-AD]) have been discussed as possible methods to accomplish this, they may require functionality beyond the capability of low-cost PE-Edge devices described in the LPE model. A key requirement for LPE-AD is to be as simple as possible while fulfilling the basic functions needed, and preferably allowing for future extensions.

The purpose of Auto-Discovery in general is to minimize the requirements for device configuration by the operators of the Service Provider network. Ideally, a specific customer VPLS connection should be provisioned one time, on one device, and that device should be able to initiate the distribution of the information to all other devices that need it. The LPE-AD protocol is used to accomplish this distribution within the LPE.

There are three major models for provisioning customer connectivity within a Logical PE architecture:
o Provision a customer connection on the PE-Edges and use an LPE

Auto-Discovery mechanism to distribute the information to the PE-Core (and thence to the other VPLS members)

- o Provision the PE-Core and distribute the information to the PE-Edges through a mechanism like SNMP or LPE Auto-Discovery
- o Use a Network Management System (NMS) to define the customer connectivity and distribute the provisioning information to both the PE-Core and PE-Edges, without Auto-Discovery between the PE-Core and PE-Edges

The first model has the attraction of simplicity. The other two models place heavier demands on the PE-Core (to act almost as an NMS in provisioning other devices) or the NMS (to understand and act on the specific LPE element capabilities). In order to focus on the LPE Auto-Discovery function as a means to simplify the Service Provider's operations, in this draft we will describe the use of LPE-AD to distribute the PE-Edge provisioning information to the PE-Core, as in the first model.

The following elements will typically be provisioned by the Service Provider when establishing a Logical PE. These elements may be provisioned on the PE-Edges or the PE-Cores. They will generally be activated one time, and changed infrequently:

- o PE-Edge-IDs (one IP address per PE-Edge)
- o VPN-IDs (one per VPLS) along with mapping to CSM-IDs

The following elements will typically be provisioned by the Service Provider when creating a specific Customer VPLS connection:

- o CSM-ID (one CSM-ID for all members of a VPLS within a LPE)
- o Policy - associated with a VLAN tag or physical interface
- o LPE Endpoint - (unique ID within SET)

The LPE Auto-Discovery protocol provides the mechanism for distributing these provisioned identifiers within the LPE. Devices that are neither PE-Edges or PE-Core within the SET SHOULD NOT participate in the LPE-AD protocol. LPE-AD messages MUST NOT be propagated into Customer Sites.

4.3 One LPE-AD Model: Provision PE-Edges and Distribute to PE-Cores

The PE-Edges and PE-Cores within a LPE are interconnected via a Switched Ethernet Transport (SET) domain, as described in [LPE-ARCH]. A key consideration is to minimize the amount of traffic generated on the SET domain by LPE-AD. This is accomplished primarily by bundling VPLS Membership information per physical interface (LPE Endpoint) on the PE-Edge (i.e., including the information for all CSM-IDs on a LPE Endpoint in a single message).

When provisioning is performed at the PE-Edges, a minimal LPE-AD may consist of messages broadcast or multicast by the PE-Edges, in a given LPE. However, a more general LPE-AD protocol could include provisions for acknowledgements, or other two-way communication.

There may be multiple SET domains per PE-Core (and LPE). When an LPE is composed of more than one SET domain, then LPE-AD messages originating from a PE-Edge device must be restricted only to its SET domain. More generally, LPE-AD messages MUST NOT be propagated into Customer Sites.

Upon receiving the LPE-AD messages, the PE-Core may do the following:
o Set up the proper LPE to LPE connectivity (e.g., generate a VC-Label for each new CSM-ID)
o May send acknowledgements to the PE-Edges

Whenever there is a state change of an LPE Endpoint or CSM-ID (due to provisioning or link/port failure, etc.) the PE-Edge will send a LPE-AD Message for each LPE Endpoint involved in the change.

When a LPE-AD Message is sent after a change, it is transmitted several times, a few seconds apart. During a period with no changes, a LPE-AD Message is sent as a keep-alive every few minutes. There is no difference in the format or content of the LPE-AD Message, in either case.

The keep-alive LPE-AD Messages allow PE-Cores to learn the state of the PE-Edges in a relatively short period of time, without adding any complications to the simple protocol. For example, PE-Cores do not need to request updates, or notify the PE-Edges of their state.

4.3.1 LPE-AD Messages

The LPE-AD Message is used by a PE-Edge device to communicate the following indications:
o add a new CSM-ID or a new LPE Endpoint
o remove a CSM-ID or LPE Endpoint
o keep-alive _ this should indicate that the VPLS Membership is active

A LPE-AD Message may group the VPLS Membership information (CSM-IDs) per LPE Endpoint, in order to minimize flooding in the LPE SET and processing in the PE-Edge. In most cases, the LPE-AD Message should fit in a single packet, but since this may not be possible in every case, there should be a provision for splitting the message across multiple packets.

In its simplest form, the LPE-AD Message will have the following information:
o LPE Endpoint Identifier
o The set of CSM-IDs for that LPE Endpoint
o The PE-Edge-ID (may be implicit in the LPE-AD packet's source address)

4.3.2 PE-Core Actions

The PE-Cores in all LPEs within a Service Provider's network are connected in a fully meshed topology across the core. They perform all of the inter-PE communications. The PE-Core will perform actions described in this section, based on the information it has received in the LPE-AD Messages.

Upon receiving a LPE-AD Message, the PE-Core will perform the following actions:

- o If the LPE-AD Message is valid, the PE-Core will check the list of CSM-IDs, otherwise it should silently discard the message.
- o If there is a new LPE Endpoint/CSM-ID tuple in the message, then it will determine or generate the VC-Label(s) associated with that LPE Endpoint/CSM-ID. It may communicate with other PE-Cores. For example, it might send an LDP Mapping Message.
- o If there is a missing LPE Endpoint/CSM-ID tuple in the message (relative to the previous LPE-AD Message), then local information regarding this CSM-ID is cleared. If the VC-Label has been issued, then it may send an appropriate message (e.g., LDP Withdraw Message) to its PE-Core peers.

If no LPE-AD Message is received from a PE-Edge for a specific LPE Endpoint within the keep-alive period, then the PE-Core will consider the VPLS connections on that LPE Endpoint closed. It will remove all the CSM-ID associations with that LPE Endpoint from its internal tables, and may send withdraw messages (e.g., LDP Withdraw Messages) to other PE-Cores (or other PE peers) in the affected VPLSs, for all the VC-Labels representing the CSM-IDs for this specific LPE Endpoint. There should be a long timer to allow for multiple dropped LPE-AD Messages. It should be at least three times the keep-alive interval.

4.3.3 Summary of PE-Edge to PE-Core Auto-Discovery Protocol

To summarize the protocol in simple terms:

- o There is one type of message involved, the LPE-AD Message. The LPE-AD Message is either triggered by a change in the LPE Endpoint association with CSM-IDs (adding/removing VPLS membership, etc.) or is a keep-alive message, which is sent at a regular interval and simply repeats the current information. The message content is the same in either case.
- o The PE-Core will use one timer per LPE Endpoint. Under any failure of the PE-Edge such that no message is sent, the PE-Core relies on these timers to delete VPLS memberships for that LPE Endpoint.

5. LPE-AD Protocol Implementation Concepts

The LPE-AD protocol should be based on the UDP protocol [UDP].

The UDP Port Number should identify the UDP/LPE-AD type messages. The Port number should be configurable, and it should have a default value <TBD> assigned.

6. Security Considerations

The LPE Auto-Discovery protocol is a lightweight protocol intended for use within a network managed by a Service Provider, in a Logical PE architecture. It does not contain explicit security mechanisms. If an attacker can gain access to the links within the LPE, then LPE-AD may be vulnerable to a variety of attacks, including interception, spoofing, and replay.

7. References

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

[LPE-ARCH] Ould-Brahim, H., et al., "VPLS/LPE L2VPNs: Virtual Private LAN Services Using Logical PE Architecture", draft-ould-brahim-l2vpn-lpe-00.txt, November 2001, work in progress.

[RFC-2764] Gleeson, B., et al., "A Framework for IP Based Virtual Private Networks", RFC 2764, February 2000.

[L2-MART] Martini, L., et al., "Transport of Layer-2 Frames over MPLS", draft-martini-l2circuit-trans-mpls-07.txt, work in progress, July 2001.

[L2-KOMP] Kompella, K., et al., "MPLS based Layer-2 VPNs", draft-kompella-ppvpn-l2vpn-00.txt, June 2001, work in progress.

[RFC 2685] Fox, B. et al., "Virtual Private Networks Identifier", RFC 2685, September 1999.

[RFC 3036] Andersson, L., et al., "LDP Specification", RFC 3036, January, 2001.

[BPG-AD] Ould-Brahim, H., et al., "Using BGP as an Auto-Discovery Mechanism for Network-Based VPNs", draft-ietf-ppvpn-bgpvpn-auto-00.txt, work in progress.

[UDP] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August, 1980.

8. Acknowledgements

We would like to acknowledge the helpful comments and contributions of the following individuals: Liam Casey, Michael Chen, Hesham Elbakoury, Marc Holness, Amita Patil.

9. Intellectual Property Considerations

Nortel Networks may seek patent or other intellectual property protection for some of all of the technologies disclosed in this document. If any standards arising from this document are or become protected by one or more patents assigned to Nortel Networks, Nortel Networks intends to disclose those patents and license them on reasonable and non-discriminatory terms.

10. Authors' Addresses

Paul Knight
Nortel Networks
600 Technology Park Drive   Phone: +1 (978) 288 6414
Billerica, Ma. USA   Email: paknight@nortelnetworks.com

Dinesh Mohan
Nortel Networks
3500 Carling Avenue   Phone: +1 (613) 763 4794

Nepean, ON K2H 8E9 Canada   Email: mohand@nortelnetworks.com

Hamid Ould-Brahim
Nortel Networks
P O Box 3511 Station C Phone: +1 (613) 765 3418
Ottawa ON K1Y 4H7 Canada Email: hbrahim@nortelnetworks.com

Vasile Radoaca
Nortel Networks
600 Technology Park Phone: +1 (978) 288 6097
Billerica, MA, 01803 Email: vasile@nortelnetworks.com

Full Copyright Statement