Network Working Group
Request for Comments: 5252
Category: Standards Track

I. Bryskin
ADVA Optical Networking
L. Berger
LabN Consulting, LLC

July 2008

OSPF-Based Layer 1 VPN Auto-Discovery

Status of This Memo

Abstract

This document defines an Open Shortest Path First (OSPF) based Layer 1 Virtual Private Network (L1VPN) auto-discovery mechanism. This mechanism enables provider edge (PE) devices using OSPF to dynamically learn about the existence of each other, and attributes of configured customer edge (CE) links and their associations with L1VPNs. This document builds on the L1VPN framework and requirements and provides a L1VPN basic mode auto-discovery mechanism.

Table of Contents

1. Introduction

1.1. Overview

The framework for Layer 1 VPNs is described in [RFC4847]. Basic mode operation is further defined in [RFC5251]. The L1VPN Basic Mode (L1VPN-BM) document [RFC5251] identifies the information that is necessary to map customer information (ports identifiers) to provider information (identifiers). It also states that this mapping information may be provided via provisioning or via an auto-discovery mechanism. This document provides such an auto-discovery mechanism using Open Shortest Path First (OSPF) version 2. Use of OSPF version 3 and support for IPv6 are out of scope of this document and will be defined separately.

Figure 1 shows the L1VPN basic service being supported using OSPF-based L1VPN auto-discovery. This figure shows two PE routers interconnected over a GMPLS backbone. Each PE is attached to three CE devices belonging to three different L1VPN connections. In this network, OSPF is used to provide the VPN membership, port mapping, and related information required to support basic mode operation.

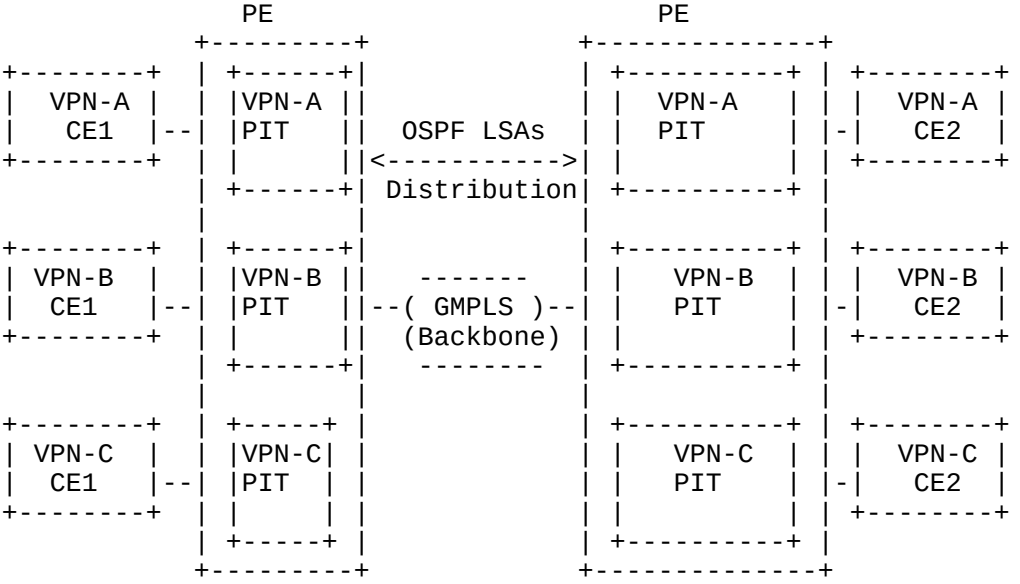

Figure 1: OSPF Auto-Discovery for L1VPNs

See [RFC5195] for a parallel L1VPN auto-discovery that uses BGP. The OSPF approach described in this document is particularly useful in networks where BGP is not typically used.

The approach used in this document to provide OSPF-based L1VPN auto-discovery uses a new type of Opaque Link State Advertisement (LSA) that is referred to as an L1VPN LSA. The L1VPN LSA carries information in TLV (type, length, value) structures. An L1VPN-specific TLV is defined below to propagate VPN membership and port information. This TLV is referred to as the L1VPN Info TLV. The L1VPN LSA may also carry Traffic Engineering (TE) TLVs; see [RFC3630] and [RFC4203].

1.2. Terminology

The reader of this document should be familiar with the terms used in [RFC4847] and [RFC5251]. The reader of this document should also be familiar with [RFC2328], [RFC5250], and [RFC3630]. In particular, the following terms:

L1VPN - Layer 1 Virtual Private Network

CE - Customer (edge) network element directly connected to the provider network (terminates one or more links to one or more PEs); it is also connected to one or more Cs and/or other CEs

C - Customer network element that is not connected to the provider network but is connected to one or more other Cs and/or CEs

PE - Provider (edge) network element directly connected to one or more customer networks (terminates one or more links to one or more CEs associated with the same or different L1VPNs); it is also connected to one or more Ps and/or other PEs

P - Provider (core) network element that is not directly connected to any customer networks; P is connected to one or more other Ps and/or PEs

LSA - OSPF link State Advertisement

LSDB - Link State Database: a data structure supported by an IGP speaker

PIT - Port Information Table

CPI - Customer Port Identifier

PPI - Provider Port Identifier

1.3. 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 [RFC2119].

2. L1VPN LSA and Its TLVs

This section defines the L1VPN LSA and its TLVs.

2.1. L1VPN LSA

The format of a L1VPN LSA is as follows:

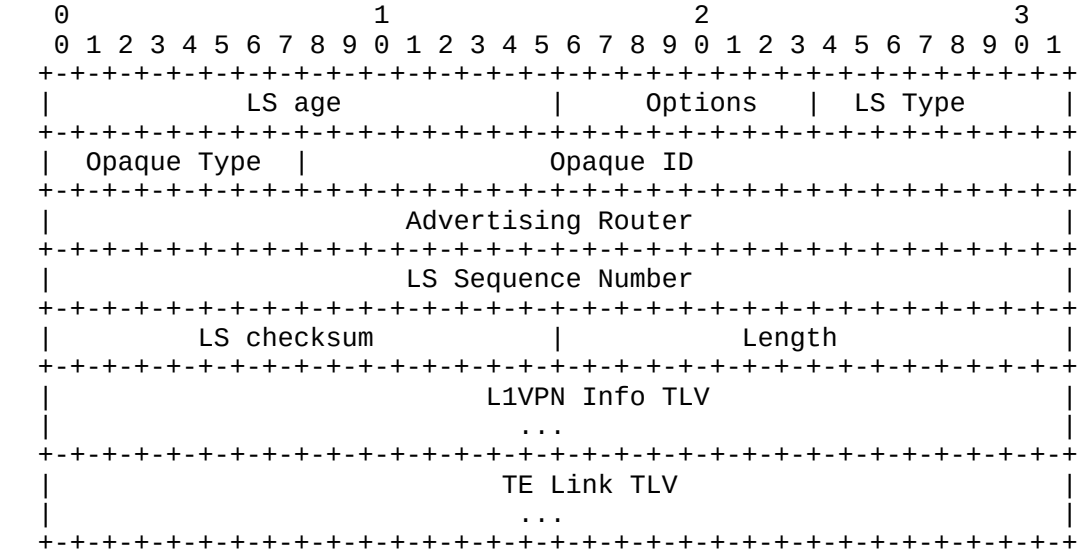

LS age
As defined in [RFC2328].

Options
As defined in [RFC2328].

LS Type
This field MUST be set to 11, i.e., an Autonomous System (AS) scoped Opaque LSA [RFC5250].

Opaque Type
The value of this field MUST be set to 5.

Opaque ID
As defined in [RFC5250].

Advertising Router
As defined in [RFC2328].

LS Sequence Number
As defined in [RFC2328].

LS checksum
As defined in [RFC2328].

Length
As defined in [RFC2328].

L1VPN Info TLV
A single TLV, as defined in Section 3.2, MUST be present. If more than one L1VPN Info TLV is present, only the first TLV is processed and the others MUST be ignored on receipt.

TE Link TLV
A single TE Link TLV (as defined in [RFC3630] and [RFC4203]) MAY be included in a L1VPN LSA.

2.2. L1VPN INFO TLV

The following TLV is introduced:

Name: L1VPN IPv4 Info
Type: 1
Length: Variable

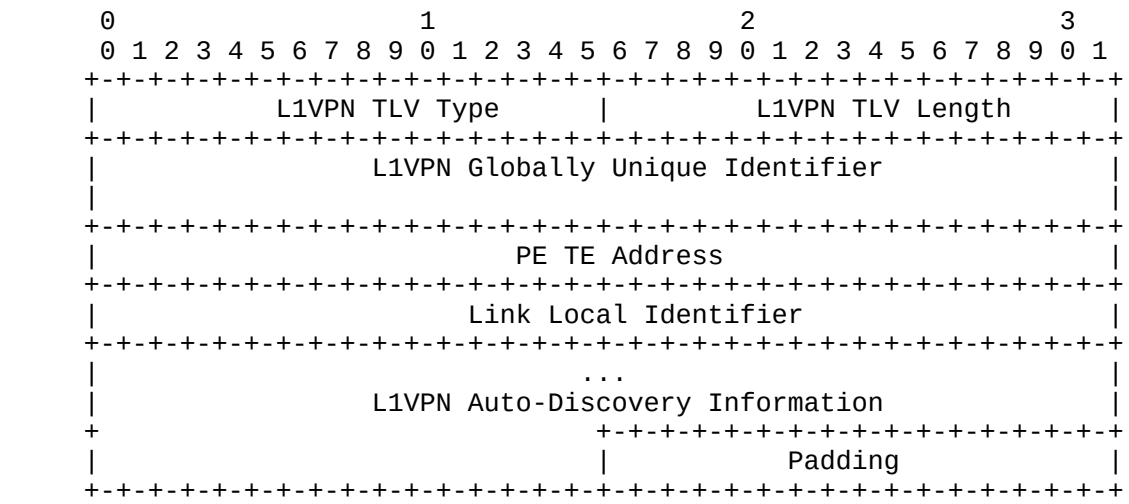

L1VPN TLV Type
The type of the TLV.

TLV Length
The length of the TLV in bytes, excluding the 4 bytes of the TLV header and, if present, the length of the Padding field.

L1VPN Globally Unique Identifier
As defined in [RFC5251].

PE TE Address
This field MUST carry an address that has been advertised by the LSA originator per [RFC3630] and is either the Router Address TLV or Local interface IP address link sub-TLV. It will typically carry the TE Router Address.

Link Local Identifier
This field is used to support unnumbered links. When an unnumbered PE TE link is represented, this field MUST contain a value advertised by the LSA originator per [RFC4203] in a Link Local/Remote Identifiers link sub-TLV. When a numbered link is represented, this field MUST be set to 0.

L1VPN Auto-discovery information
As defined in [RFC5251].

Padding
A field of variable length and of sufficient size to ensure that the TLV is aligned on a 4-byte boundary. This field is only required when the L1VPN Auto-discovery information field is not 4-byte aligned. This field MUST be less than 4 bytes long, and MUST NOT be present when the size of the L1VPN Auto-discovery information field is 4-byte aligned.

3. L1VPN LSA Advertising and Processing

PEs advertise local <CPI, PPI> tuples in L1VPN LSAs containing L1VPN Info TLVs. Each PE MUST originate a separate L1VPN LSA with AS flooding scope for each local CE-to-PE link. The LSA MUST be originated each time a PE restarts and every time there is a change in the PIT entry associated with a local CE-to-PE link. The LSA MUST include a single L1VPN Info TLV and MAY include a single TE Link TLV as per [RFC3630] and [RFC4203]. The TE Link TLV carries TE attributes of the associated CE-to-PE link. Note that because CEs are outside of the provider TE domain, the attributes of CE-to-PE links are not advertised via normal OSPF-TE procedures as described in [RFC3630] and [RFC4203]. If more than one L1VPN Info TLVs and/or TE Link TLVs are found in the LSA, the subsequent TLVs SHOULD be ignored by the receiving PEs.

L1VPN LSAs are of AS-scope (LS type is set to 11) and therefore are flooded to all PEs within the AS according to [RFC5250]. Every time a PE receives a new, removed, or modified L1VPN LSA, the PE MUST check whether it maintains a PIT associated with the L1VPN specified in the L1VPN globally unique identifier field. If this is the case (the appropriate PIT will be found if one or more local CE-to-PE links that belong to the L1VPN are configured), the PE SHOULD add, remove, or modify the PIT entry associated with each of the advertised CE-to-PE links accordingly. (An implementation MAY choose to not remove or modify the PIT according to local policy or management directives.) Thus, in the normal steady-state case, all PEs associated with a particular L1VPN will have identical local PITs for an L1VPN.

3.1. Discussion and Example

The L1VPN auto-discovery mechanism described in this document does not prevent a PE from applying any local policy with respect to PIT management. An example of such a local policy would be the ability to configure permanent (static) PIT entries. Another example would

be the ability to ignore information carried in L1VPN LSAs advertised by a specific TE.

The reason why it is required that the value specified in the PE TE Address field of the L1VPN Info TLV matches a valid PE TE Router ID or numbered TE Link ID is to ensure that CEs attached to this PE can be resolved to the PE as it is known to the Traffic Engineering Database (TED) and hence TE paths toward the CEs across the provider domain can be computed.

Let us consider the example presented in Figure 2.

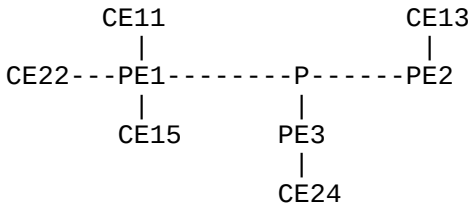

Figure 2: Single Area Configuration

Let us assume that PE1 is connected to CE11 and CE15 in L1VPN1 and to CE22 in L1VPN2; PE2 is connected to CE13 in L1VPN1; PE3 is connected to CE24 in L1VPN2. In this configuration PE1 manages two PITs: PIT1 for L1VPN1 and PIT2 for L1VPN2; PE2 manages only PIT1; and PE3 manages only PIT2. PE1 originates three L1VPN LSAs, each containing a L1VPN Info TLV advertising links PE1-CE11, PE1-CE22, and PE1-CE15, respectively. PE2 originates a single L1VPN LSA for link PE2-CE13, and PE3 originates a single L1VPN LSA for link PE3-CE24. In steady state, the PIT1 on PE1 and PE3 will contain information on links PE1-CE11, PE1-CE15, and PE2-CE13; PIT2 on PE1 and PE2 will contain entries for links PE1-CE22 and PE3-CE24. Thus, all PEs will learn about all remote PE-to-CE links for all L1VPNs supported by PEs.

Note that P in this configuration does not have links connecting it to any L1VPNs. It neither originates L1VPN LSAs nor maintains any PITs. However, it does participate in the flooding of all of the L1VPN LSAs and hence maintains the LSAs in its LSDB. This is a cause for scalability concerns and could prove to be problematic in large networks.

4. Backward Compatibility

Neither the TLV nor the LSA introduced in this document present any interoperability issues. Per [RFC5250], OSPF speakers that do not support the L1VPN auto-discovery application (Ps for example) just

participate in the L1VPN LSAs flooding process but should ignore the LSAs contents.

5. Security Considerations

The approach presented in this document describes how PEs dynamically learn L1VPN-specific information. Mechanisms to deliver the VPN membership information to CEs are explicitly out of scope of this document. Therefore, the security issues raised in this document are limited to within the OSPF domain.

This defined approach reuses mechanisms defined in [RFC2328] and [RFC5250]. Therefore, the same security approaches and considerations apply to this approach. OSPF provides several security mechanisms that can be applied. Specifically, OSPF supports multiple types of authentication, limits the frequency of LSA origination and acceptance, and provides techniques to avoid and limit impact database overflow. In cases where end-to-end authentication is desired, OSPF's neighbor-to-neighbor authentication approach can be augmented with an experimental extension to OSPF; see [RFC2154], which supports the signing and authentication of LSAs.

6. IANA Considerations

This document requests the assignment of an OSPF Opaque LSA type. IANA has made the assignment in the form:

Value	Opaque Type	Reference
5	L1VPN LSA	[RFC5252]

7. Acknowledgments

We would like to thank Adrian Farrel and Anton Smirnov for their useful comments.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC5250] Berger, L., Bryskin, I., and A. Zinin, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

[RFC5251] Fedyk, D., Ed., Rekhter, Y., Ed., Papadimitriou, D., Rabbat, R., and L. Berger, "Layer 1 VPN Basic Mode", RFC 5251, July 2008.

8.2. Informative References

[RFC2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

[RFC4847] Takeda, T., Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC 4847, April 2007.

[RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, June 2008.

Authors' Addresses

Igor Bryskin
ADVA Optical Networking Inc
7926 Jones Branch Drive
Suite 615
McLean, VA 22102
EMail: ibryskin@advaoptical.com

Lou Berger
LabN Consulting, LLC
EMail: lberger@labn.net

Full Copyright Statement

Intellectual Property