l1vpn WG

Hamid Ould-Brahim
Don Fedyk
Nortel

Internet Draft
Expiration Date: September 2006

Yakov Rekhter
Juniper Networks

March 2006

BGP-based Auto-Discovery for L1VPNs

draft-ouldbrahim-l1vpn-bgp-auto-discovery-01.txt

Status of this Memo

Abstract

The purpose of this draft is to define a BGP-based auto-discovery mechanism for layer-1 VPNs. The auto-discovery mechanism for l1vpns allows the provider network devices to dynamically discover the set of PEs having ports attached to CEs member of the same VPN. That information is necessary for completing the signaling phase. One main objective of l1vpn auto-discovery mechanism is to support "single-end provisioning" model, where addition of a new port to a given l1vpn would involve configuration changes only on the PE that has this port and on the CE that is connected to the PE via this port.

1. Introduction

The purpose of this draft is to define a BGP-based auto-discovery mechanism for layer-1 VPNs. The auto-discovery mechanism for l1vpns allows the provider network devices to dynamically discover the set of PEs having ports attached to CEs member of the same VPN. That information is necessary for completing the signaling phase. One main objective of l1vpn auto-discovery mechanism is to support "single-end provisioning" model, where addition of a new port to a given l1vpn would involve configuration changes only on the PE that has this port and on the CE that is connected to the PE via this port.

The auto-discovery mechanism proceeds by having a PE advertises to other PEs, at a minimum, its own IP address and the list of <private address, provider address> tuples local to that PE. Once that information is received, the remote PEs will identify the list of VPN members they have in common with the advertising PE, and use the information carried within the discovery mechanism to perform address resolution during signaling phase.

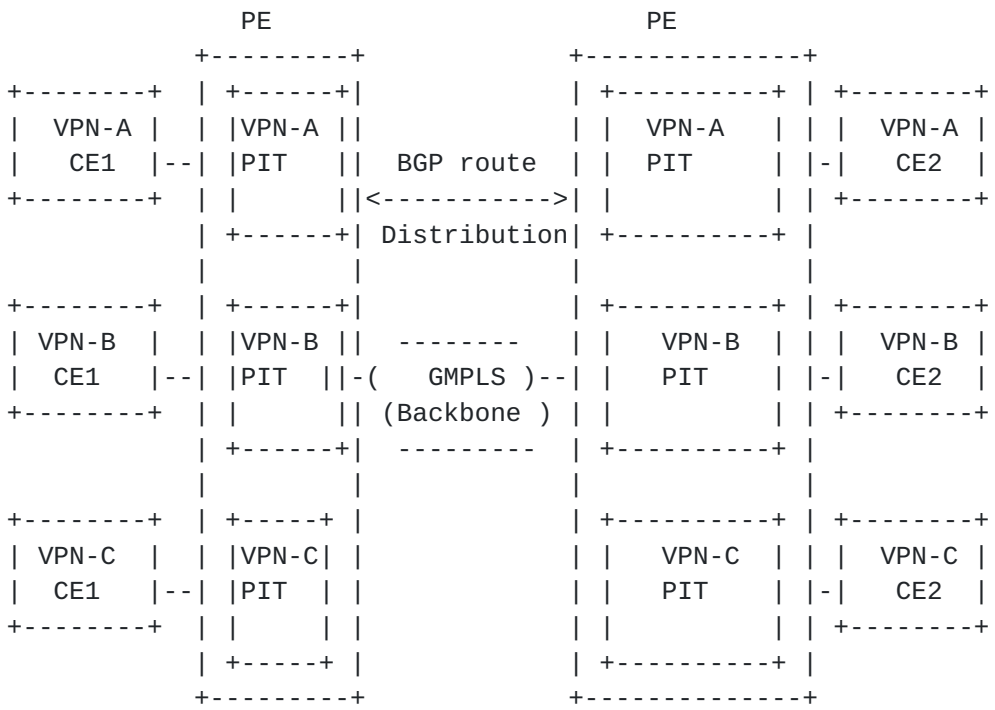

Figure 1 BGP auto-discovery for l1vpn

This version of the draft focuses on describing an auto-discovery mechanism for the basic mode only. Details for the

enhanced mode will be described in future revised version of this draft.

2. Procedures

In the context of l1vpns, a CE is connected to a PE via one or more ports, where each port may consists of one or more channels or sub-channels. Each port on a CE that connects the CE to a PE has an identifier that is unique within that l1vpn (but need not be unique across several l1vpn). We refer to this identifier as the customer port identifier (CPI). Each port on a PE has as well an identifier that is unique within that provider network. We refer to this identifier as the provider port identifier (PPI). Note that IP addresses used for CPIs, PPIs could be either IPv4 or IPv6 addresses.

A PE maintains for each l1vpn configured on that PE a port information tables (PIT) associated with each l1vpn that has at least one port configured on a PE. A PIT contains a list of <CPI, PPI> tuples for all the ports within its l1vpn. Note that a PIT may as well hold routing information (for example when CPIs are learnt using a routing protocol).

A PIT on a given PE is populated from two sources: the information related to the CEs ports attached to the ports on that PE (this information could be optionally received from the CEs), and the information received from other PEs through the auto-discovery mechanism. We ll refer to the former as the "local" information, and to the latter as the "remote" information.

Propagation of local information to other PEs is accomplished by using BGP multiprotocol extensions as specified in [BGP-VPN-AUTODISCOVERY]. To restrict the flow of this information to only the PITs within a given l1vpn, we use BGP route filtering based on the Route Target Extended Community [BGP-COMM], as follows.

Each PIT on a PE is configured with one or more Route Target Communities, called "export Route Targets", that are used for tagging the local information when it is exported into provider s BGP. The granularity of such tagging could be as fine as a single <CPI, PPI> pair. In addition, each PIT on a PE is configured with one or more Route Target Communities, called "import Route Targets", that restrict the set of routes that could be imported from provider s BGP into the PIT to only the routes that have at least of these Communities.

When a service provider adds a new l1vpn port to a particular PE, this port is associated at provisioning time with a PIT on that PE, and this PIT is associated (again at provisioning time) with that l1vpn.

Note that since the protocol used to populate a PIT with remote information is BGP, since BGP works across multiple routing domains, it follows that the mechanisms described in this document could support l1vpns that span multiple routing domains.

3. Carrying l1vpn information in BGP

The <CPI, PPI> mapping is carried using the Multiprotocol Extensions BGP [RFC2858]. [RFC2858] defines the format of two BGP attributes, MP_REACH_NLRI and MP_UNREACH_NLRI that can be used to announce and withdraw the announcement of reachability information. We introduce a new a new subsequent address family identifier (to be assigned by the IANA), and also a new NLRI format for carrying the CPI and PPI information.

One or more <PPI, CPI> tuples could be carried in the above mentioned BGP attributes.

The format of encoding a single <PPI, CPI> tuple is shown in Figure 2 below:

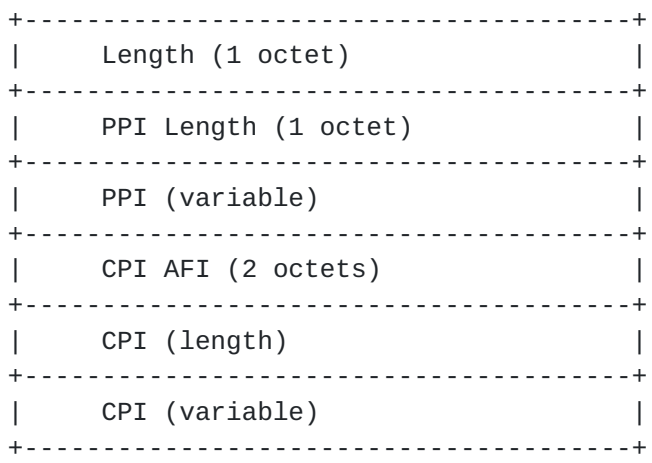

Figure 2: NLRI BGP encoding

The use and meaning of these fields are as follows:

Length:

A one octet field whose value indicates the length of the <PPI, CPI> Information tuple in octets.

PPI Length:

A one octet field whose value indicates the length of of the PPI field

PPI field:

A variable length field that contains the value of

the PPI (either an address or <port index, address> tuple

CPI AFI field:

A two octets field whose value indicates address family of the CPI.

CPI Length:

A once octet field whose value indicates the length of the CPI field.

CPI (variable):

A variable length field that contains the CPI value (either an address or <port index, address> tuple.

4. Security Considerations

TBD.

5. References

[BGP-VPN-AUTODISCOVERY] Ould-Brahim, H., Rosen, E., Rekhter, Y., "Using BGP as an Auto-Discovery Mechanism for Layer-3 and Layer-2 VPNs", draft-ietf-l3vpn-bgpvpn-auto-05.txt, work in progress

[GVPN] Ould-Brahim, H., Rekhter, Y., et al., "Generalized VPNs using BGP and GMPLS toolkit", work in progress, August 2005.

[BGP-COMM] Ramachandra, Tappan, et al., "BGP Extended Communities Attribute", draft-ietf-idr-bgp-ext-communities-08.txt, August 2005, work in progress.

[BGP-MP] Bates, Chandra, Katz, and Rekhter, "Multiprotocol Extensions for BGP4", February 1998, RFC 2283.

[L1VPN-FRMK] Tomonori Takeda, et al., " Framework and Requirements for Layer 1 Virtual Private Networks", draft-ietf-l1vpn-framework-00.txt, August 2005, work in progress.

6. Author's Addresses

Hamid Ould-Brahim

Nortel
P O Box 3511 Station C

Ottawa ON K1Y 4H7 Canada
Phone: +1 (613) 765 3418
Email: hbrahim@nortel.com

Yakov Rekhter
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089
Email: yakov@juniper.net

Don Fedyk
Nortel
600 Technology Park
Billerica, Massachusetts
01821 U.S.A
Phone: +1 (978) 288 3041
Email: dwfedyk@nortel.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of and Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement