Internet Engineering Task Force
INTERNET-DRAFT
Expires January 2002

Chandrasekar Kathirvelu
Karthik Muthukrishnan
Tom Walsh
Lucent Technologies

Andrew Malis
Vivace Networks, Inc.

Fred Ammann
COMMCARE telecommunications

July 2001

A Core MPLS IP VPN Link Broadcast And Virtual Router Discovery <draft-ietf-ppvpn-corevpn-disc-00.txt>

Status of this Memo

Abstract

An IPVPN consists of many routers, some physically discrete and some virtual, housed in a Provider Edge router. The problem that presents itself is that these virtual routers need to find each other over a virtual topology and they need to send broadcast datagrams as mandated in routing protocols [such as the neighbor discovery datagram and routing updates in OSPF, the routing updates in RIPV2 etc] and user data over this virtual topology. This memo presents an approach for solving these problems.

1. Acronyms

ARP	Address Resolution Protocol
CE	Customer Edge router
LSP	Label Switched Path
PNA	Private Network Administrator
SLA	Service Level Agreement
SP	Service Provider
SPED	Service Provider Edge Device
SPNA	SP Network Administrator
VL	Inter VR Virtual Link
VMA	VPN Multicast Address
VPNID	VPN Identifier
VR	Virtual Router
VRC	Virtual Router Console

2. Introduction

The two problems that need to be addressed are link level broadcast and inter VR discovery over a virtual topology.

Broadly we can classify the solutions as static and dynamic. The static approach calls for manually configuring the neighbor address. This, while easy to articulate and effective for small VPNs, is a configuration nightmare. Alternatively, this memo describes an approach using an IP multicast infrastructure or ARP servers for virtual routers to discover other virtual routers in a given VPN and for supporting link level broadcast over a virtual topology.

3. Static Approach

In this approach, we expect the user to configure a given VR's neighbors in that VR. The exact mechanism is static ARP. Basically, the user configures a static ARP entry for each neighboring VR. The static ARP entry provides a mapping from the neighboring VR's interface on the virtual link (VL) to the backbone visible IP address of the neighbor's hosting PE. If this methodology is used and the routing protocol(s) configured to run over the VL require link broadcast, it has to be implemented using non-native multicast forwarding paradigms; this is obviously less than optimum from a network utilization standpoint.

4. Dynamic Approach

When the VPN is large with a large number of VRs associated with it, a dynamic, automated method is preferable to static configuration. This draft presents a IP multicast based approach whereby dynamic ARP is used to discover the neighbor VR's PE address and link broadcast is achieved using encapsulation of VPN

link broadcast packets in the multicast address assigned to the VPN.

4.1. Using ARP for VR Discovery

In a physical LAN with a number of routers, when a data packet needs to be forwarded to one of the other routers, an ARP request is broadcast to the LAN. The router with the logical address queried in the ARP request responds with an ARP response with its MAC address. In identical fashion, the dynamic approach described in this draft sends an ARP request to the multicast address assigned to the VPN. The other VRs associated with this VPN receive the ARP request and the appropriate VR responds with its MAC address, i.e., its PE's backbone visible IP address. In order to achieve this, we need to have a new hardware type field in the ARP Req and Response. This enjoys the advantage that ARP is implemented in almost all the SPEDs and CPEs and it is independent of any routing protocols.

When we discuss the methods of carrying ARP, again it may depend on the SP's network configuration. We discuss few methods here.

4.1.1. ARP over Multicast

If we assume that the SP's network supports multicast then each VPN can subscribe to one multicast group as described in RFC 2917 and exchange the neighbor information.

4.2. Using Multicast For Link Broadcast

Some routing protocols, most notably IGPs, require link level multicast facilities. For example, OSPF in broadcast mode uses 224.0.0.5 to discover other OSPF routers and to send route updates, RIPv2 uses 224.0.0.9 to send route updates. If the optimizations achieved with the use of these modes is desirable over the VL, this approach calls for sending these routing

datagrams over the VPN's multicast address. This ensures that only the VRs that participate in a given VPN receive these datagrams.

5. Intellectual Property Considerations

Lucent technologies 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 Lucent Technologies. Lucent Technologies intends to disclose those patents and license them on reasonable and non-discriminatory terms.

6. References

[muthuk] K.Muthukrishnan, A.Malis "A Core MPLS IP VPN Architecture", RFC 2917 September 2000.

7. Authors' Addresses

Chandrasekar Kathirvelu
Lucent Technologies
1 Robbins Road
Westford, MA 01886
Phone: (978) 952-7116
EMail: ck32@lucent.com

Karthik Muthukrishnan
Lucent Technologies
1 Robbins Road
Westford, MA 01886
Phone: (978) 952-1368
EMail: mkarthik@lucent.com

Tom Walsh
Lucent Technologies
10 Lyberty Way
Westford, MA 01886
Phone: (978) 392-2311
EMail: tdwalsh@lucent.com

Fred Ammann

COMMCARE Telecommunications
Turmstrasse 8
CH-8952 Schlieren
Switzerland
Phone: +41 1 738 61 11
Email: fa@commcare.ch

Andrew Malis
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Phone: (408) 383-7223
EMail: Andy.Malis@vivacenetworks.com

8. Full Copyright Statement

MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.