

MPLS-VPN



ip infusion™
application note

Introduction

Businesses today are looking to the Internet for wide area network (WAN) solutions that in the recent past they could get only by choosing Frame Relay or T1 dedicated links. To achieve the security that is required for corporate users, virtual private networks (VPNs) can be used to guarantee that traffic is securely tunneled over the Internet. Up to now, most VPNs have been provisioned using Layer 2 technologies, such as Frame Relay and asynchronous transfer mode (ATM). These technologies provided secure tunnels, were resistant to Denial-of-Service (DoS) and intrusion attacks, and provided address and routing separation. The problem with Layer 2 VPN technology is that it does not scale well. As the network grows, the number of required virtual circuits achieving optimal routing scales non-linearly. It is also difficult to provide traffic engineering using a Layer 2 VPN approach.

To solve these scaling problems, a border gateway protocol/multiprotocol label switching (BGP/MPLS) VPN standard is now being adopted to provide Layer 3 VPN solutions using BGP to carry route information over a MPLS core. This Layer 3 MPLS-VPN solution achieves all of the security of the Layer 2 approach, while adding enhanced scalability inherent in the use of Layer 3 routing technology.

The key to this approach is the use of BGP and a set of extensions, known as BGP-VPN, that allow separate route forwarding information to be maintained for each VPN client. BGP then carries this separate route forwarding information over MPLS using the label distribution protocol (LDP).

This application note discusses how IP Infusion's ZebOS™ Advanced Routing Suite (ARS) provides source code MPLS-VPN solutions for provider edge equipment manufacturers. These provider edge devices can then be used by Service Providers to provision VPN services directly to customers.

MPLS/BGP-VPN

In order to achieve the security that is necessary for VPN provisioning over the Internet using a Layer 3 approach, address and routing separation between customers is required. This is inherent in a Layer 2 approach, but must be specially designed to work in a Layer 3 based VPN solution. To solve this problem, *draft-ietf-ppvpn-rfc2547bis-00.txt* has been developed by a number of Internet experts (notably from Cisco, Juniper, ATT, Alcatel, Worldcom, and others). This draft RFC specifically defines how to provide address and routing separation using BGP, and how to send this information and the VPN traffic itself over a MPLS backbone.

The model as expressed in the draft RFC is that Service Providers (SP) own the backbone and provision VPN services from Provider Edge (PE) equipment

which communicates directly with Customer Edge (CE) equipment using standard technology such as Frame Relay, ATM, DSL, and T1. At that time, the customer would purchase VPN services directly from the SP. Then, the SP would provide the VPN service to multiple customers using a shared PE device.

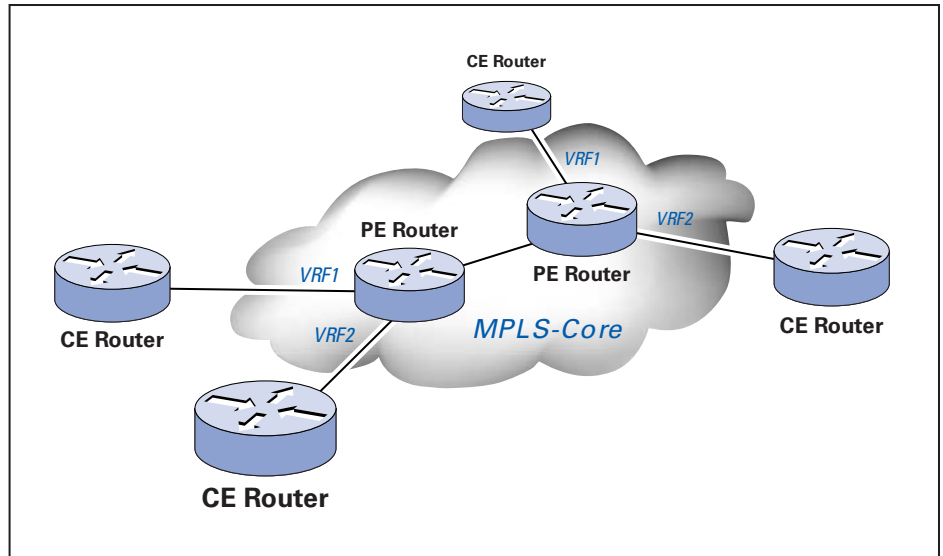
The key to providing security in the shared PE equipment is made available by the BGP-VPN extensions as defined in the draft RFC. Each PE router must maintain a number of forwarding tables, each of which map to a unique VPN class. When a packet is received from the CE equipment, the forwarding table that is mapped to that site is used to determine the routing for the data. Each VPN has its own unique forwarding table, known as a VRF (VPN Routing and Forwarding). If a PE device has multiple connections to the same site, a single VRF can be mapped to all of those connections. The BGP-VPN extensions for VRF support then allow BGP to send the specific route forwarding information to the PE router connected to the other end of the VPN. In this approach, route separation is maintained for each unique VPN customer.

In this type of architecture, only PE routers must carry the VRF information. It is not necessary that the non-edge routers on the SP backbone know anything about the VRF information. Consequently, this design greatly expands the scalability of the Layer 3 VPN approach.

In each PE router, sub-interfaces may be mapped to VRFs; the mapping is many-to-one. Two sub-interfaces may not map to the same VRF, unless they are intended to show route information, and the VRF packet's destination address is determined by the sub-interface over which it is received.

A VPN-IPv4 address concept is defined for use in distinguishing routes. A VPN-IPv4 address is a 12-byte address that begins with an 8-byte Route Distinguisher (RD) and ends with a 4-byte IPv4 address. BGP-Multi-protocol Extensions (BGP-MP) allow BGP to carry routes from this new address family. The VPN-IPv4 address family and RD ensure that if similar addresses are used in two different VPN's, then two separate routes to that address can be maintained. This is important when supporting RFC1918 private addressing.

The BGP-VPN extensions allow route distribution policies to be configured for the proper distribution of VPN route information. PE routers can also auto-discover the other PE device attached to the same VPN. This eliminates the need to reconfigure both PE devices when reconfiguring or initially configuring the VPN.



Forwarding VPN information using MPLS

As stated previously, the intermediate routers in the backbone do not need to maintain any information about the VPNs. So how are packets forwarded from one VPN to another? The answer is to use MPLS with a two-level label stack. PE routers insert 32-bit address prefixes into the Internal Gateway Protocol Routing tables of the backbone. By doing this, MPLS at each node in the SP backbone can assign a label to the corresponding route in each PE router. To certify that this is interoperable, LDP (Label Distribution Protocol) is used for setting up the label switched paths across the SP backbone.

A variety of mechanisms can be used for the CE equipment to deliver routing information to the PE router. This includes the use of static routes and BGP. BGP has many advantages for CE to PE communications. The main advantage is that it does not require multiple instances on the PE since it is explicitly designed for this function.

The ZebOS MPLS-VPN Solution

The ZebOS MPLS-VPN solution is a portable, platform-independent solution that can be integrated into PE equipment to deliver secure, peer-to-peer VPN solutions across a MPLS backbone. A number of components are necessary to provide the full MPLS-VPN solutions. The ZebOS Advanced Routing Suite provides these solutions and also contains a number of optional routing components that can be easily added to support additional IPv4 and IPv6 routing solutions.

To fully implement the MPLS-VPN solution into PE equipment, the following ZebOS modules are required:

- ZebOS Network Services Module
- ZebOS BGP Protocol Module with the optional VPN extensions
- ZebOS MPLS-LDP Module
- ZebOS MPLS Forwarder Module

Together, these modules provide the complete MPLS-VPN solution. Optional modules are available for OSPFv2, OSPFv3, RIP, and RIPng to support both IPv4 and IPv6 routing solutions along with MPLS-VPN.

The ZebOS Advanced Routing Suite is being integrated into a number of network processing environments. These implementations will allow MPLS-VPN to be supported on standard network processors and will assist equipment manufacturers in quickly getting their products to market.

The ZebOS MPLS-VPN solution supports a variety of operating systems and processors. Please refer to the individual spec sheets for more details on feature and requirements.



IP Infusion Inc.
111 W. St. John Street
Suite 910
San Jose, CA 95113
tel: 408-794-1500
fax: 408-278-0521
marketing@ipinfusion.com
www.ipinfusion.com

© Copyright 2001 IP Infusion Inc. All Rights Reserved. IP Infusion, ipinfusion and ZebOS are trademarks of IP Infusion Inc. All other brands or product names are trademarks or registered trademarks of their respective holders. All specifications within this document are subject to change without notice.

Part No. 01910001