

BGP/MPLS VPNs

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This document describes a method by which a Service Provider with an IP backbone may provide VPNs (Virtual Private Networks) for its customers. MPLS (Multiprotocol Label Switching) is used for forwarding packets over the backbone, and BGP (Border Gateway Protocol) is used for distributing routes over the backbone. The primary goal of this method is to support the outsourcing of IP backbone services for enterprise networks. It does so in a manner which is simple for the enterprise, while still scalable and flexible for the Service Provider, and while allowing the Service Provider to add value. These techniques can also be used to provide a VPN which itself provides IP service to customers.

Table of Contents

1	Introduction	2
1.1	Virtual Private Networks	2
1.2	Edge Devices	3
1.3	VPNs with Overlapping Address Spaces	4
1.4	VPNs with Different Routes to the Same System	4
1.5	Multiple Forwarding Tables in PEs	5
1.6	SP Backbone Routers	5
1.7	Security	5
2	Sites and CEs	6
3	Per-Site Forwarding Tables in the PEs	6
3.1	Virtual Sites	8
4	VPN Route Distribution via BGP	8
4.1	The VPN-IPv4 Address Family	9
4.2	Controlling Route Distribution	10

4.2.1	The Target VPN Attribute	10
4.2.2	Route Distribution Among PEs by BGP	12
4.2.3	The VPN of Origin Attribute	13
4.2.4	Building VPNs using Target and Origin Attributes ...	14
5	Forwarding Across the Backbone	15
6	How PEs Learn Routes from CEs	16
7	How CEs learn Routes from PEs	19
8	What if the CE Supports MPLS?	19
8.1	Virtual Sites	19
8.2	Representing an ISP VPN as a Stub VPN	20
9	Security	20
9.1	Point-to-Point Security Tunnels between CE Routers .	21
9.2	Multi-Party Security Associations	21
10	Quality of Service	22
11	Scalability	22
12	Intellectual Property Considerations	23
13	Security Considerations	23
14	Acknowledgments	23
15	Authors' Addresses	24
16	References	24
17	Full Copyright Statement.....	25

1. Introduction

1.1. Virtual Private Networks

Consider a set of "sites" which are attached to a common network which we may call the "backbone". Let's apply some policy to create a number of subsets of that set, and let's impose the following rule: two sites may have IP interconnectivity over that backbone only if at least one of these subsets contains them both.

The subsets we have created are "Virtual Private Networks" (VPNs). Two sites have IP connectivity over the common backbone only if there is some VPN which contains them both. Two sites which have no VPN in common have no connectivity over that backbone.

If all the sites in a VPN are owned by the same enterprise, the VPN is a corporate "intranet". If the various sites in a VPN are owned by different enterprises, the VPN is an "extranet". A site can be in more than one VPN; e.g., in an intranet and several extranets. We regard both intranets and extranets as VPNs. In general, when we use the term VPN we will not be distinguishing between intranets and extranets.

We wish to consider the case in which the backbone is owned and operated by one or more Service Providers (SPs). The owners of the sites are the "customers" of the SPs. The policies that determine

whether a particular collection of sites is a VPN are the policies of the customers. Some customers will want the implementation of these policies to be entirely the responsibility of the SP. Other customers may want to implement these policies themselves, or to share with the SP the responsibility for implementing these policies. In this document, we are primarily discussing mechanisms that may be used to implement these policies. The mechanisms we describe are general enough to allow these policies to be implemented either by the SP alone, or by a VPN customer together with the SP. Most of the discussion is focused on the former case, however.

The mechanisms discussed in this document allow the implementation of a wide range of policies. For example, within a given VPN, we can allow every site to have a direct route to every other site ("full mesh"), or we can restrict certain pairs of sites from having direct routes to each other ("partial mesh").

In this document, we are particularly interested in the case where the common backbone offers an IP service. We are primarily concerned with the case in which an enterprise is outsourcing its backbone to a service provider, or perhaps to a set of service providers, with which it maintains contractual relationships. We are not focused on providing VPNs over the public Internet.

In the rest of this introduction, we specify some properties which VPNs should have. The remainder of this document outlines a VPN model which has all these properties. The VPN Model of this document appears to be an instance of the framework described in [4].

1.2. Edge Devices

We suppose that at each site, there are one or more Customer Edge (CE) devices, each of which is attached via some sort of data link (e.g., PPP, ATM, ethernet, Frame Relay, GRE tunnel, etc.) to one or more Provider Edge (PE) routers.

If a particular site has a single host, that host may be the CE device. If a particular site has a single subnet, that the CE device may be a switch. In general, the CE device can be expected to be a router, which we call the CE router.

We will say that a PE router is attached to a particular VPN if it is attached to a CE device which is in that VPN. Similarly, we will say that a PE router is attached to a particular site if it is attached to a CE device which is in that site.

When the CE device is a router, it is a routing peer of the PE(s) to which it is attached, but is not a routing peer of CE routers at

other sites. Routers at different sites do not directly exchange routing information with each other; in fact, they do not even need to know of each other at all (except in the case where this is necessary for security purposes, see [section 9](#)). As a consequence, very large VPNs (i.e., VPNs with a very large number of sites) are easily supported, while the routing strategy for each individual site is greatly simplified.

It is important to maintain clear administrative boundaries between the SP and its customers (cf. [4]). The PE and P routers should be administered solely by the SP, and the SP's customers should not have any management access to it. The CE devices should be administered solely by the customer (unless the customer has contracted the management services out to the SP).

1.3. VPNs with Overlapping Address Spaces

We assume that any two non-intersecting VPNs (i.e., VPNs with no sites in common) may have overlapping address spaces; the same address may be reused, for different systems, in different VPNs. As long as a given endsystem has an address which is unique within the scope of the VPNs that it belongs to, the endsystem itself does not need to know anything about VPNs.

In this model, the VPN owners do not have a backbone to administer, not even a "virtual backbone". Nor do the SPs have to administer a separate backbone or "virtual backbone" for each VPN. Site-to-site routing in the backbone is optimal (within the constraints of the policies used to form the VPNs), and is not constrained in any way by an artificial "virtual topology" of tunnels.

1.4. VPNs with Different Routes to the Same System

Although a site may be in multiple VPNs, it is not necessarily the case that the route to a given system at that site should be the same in all the VPNs. Suppose, for example, we have an intranet consisting of sites A, B, and C, and an extranet consisting of A, B, C, and the "foreign" site D. Suppose that at site A there is a server, and we want clients from B, C, or D to be able to use that server. Suppose also that at site B there is a firewall. We want all the traffic from site D to the server to pass through the firewall, so that traffic from the extranet can be access controlled. However, we don't want traffic from C to pass through the firewall on the way to the server, since this is intranet traffic.

This means that it needs to be possible to set up two routes to the server. One route, used by sites B and C, takes the traffic directly to site A. The second route, used by site D, takes the traffic

instead to the firewall at site B. If the firewall allows the traffic to pass, it then appears to be traffic coming from site B, and follows the route to site A.

1.5. Multiple Forwarding Tables in PEs

Each PE router needs to maintain a number of separate forwarding tables. Every site to which the PE is attached must be mapped to one of those forwarding tables. When a packet is received from a particular site, the forwarding table associated with that site is consulted in order to determine how to route the packet. The forwarding table associated with a particular site S is populated only with routes that lead to other sites which have at least one VPN in common with S. This prevents communication between sites which have no VPN in common, and it allows two VPNs with no site in common to use address spaces that overlap with each other.

1.6. SP Backbone Routers

The SP's backbone consists of the PE routers, as well as other routers (P routers) which do not attach to CE devices.

If every router in an SP's backbone had to maintain routing information for all the VPNs supported by the SP, this model would have severe scalability problems; the number of sites that could be supported would be limited by the amount of routing information that could be held in a single router. It is important to require therefore that the routing information about a particular VPN be present ONLY in those PE routers which attach to that VPN. In particular, the P routers should not need to have ANY per-VPN routing information whatsoever.

VPNs may span multiple service providers. We assume though that when the path between PE routers crosses a boundary between SP networks, it does so via a private peering arrangement, at which there exists mutual trust between the two providers. In particular, each provider must trust the other to pass it only correct routing information, and to pass it labeled (in the sense of MPLS [9]) packets only if those packets have been labeled by trusted sources. We also assume that it is possible for label switched paths to cross the boundary between service providers.

1.7. Security

A VPN model should, even without the use of cryptographic security measures, provide a level of security equivalent to that obtainable when a level 2 backbone (e.g., Frame Relay) is used. That is, in the absence of misconfiguration or deliberate interconnection of

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.