

Applicability Statement for BGP/MPLS IP Virtual Private Networks (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 (2006).

Abstract

This document provides an Applicability Statement for the Virtual Private Network (VPN) solution described in RFC 4364 and other documents listed in the References section.

Table of Contents

1. Introduction	2
2. SP Provisioning Model	4
3. Supported Topologies and Traffic Types	6
4. Isolated Exchange of Data and Routing Information	7
5. Access Control and Authentication	9
6. Security Considerations	9
6.1. Protection of User Data	9
6.2. SP Security Measures	10
6.3. Security Framework Template	12
7. Addressing	18
8. Interoperability and Interworking	19
9. Network Access	19
9.1. Physical/Link Layer Topology	19
9.2. Temporary Access	19
9.3. Access Connectivity	20
10. Service Access	21
10.1. Internet Access	21
10.2. Other Services	21
11. SP Routing	22
12. Migration Impact	22
13. Scalability	23
14. QoS, SLA	26

15. Management	27
15.1. Management by the Provider	27
15.2. Management by the Customer	28
16. Acknowledgements	28
17. Normative References	29
18. Informative References	29

1. Introduction

This document provides an Applicability Statement for the Virtual Private Network (VPN) solution described in [BGP-MPLS-IP-VPN] and other documents listed in the References section. We refer to these as "BGP/MPLS IP VPNs", because Border Gateway Protocol (BGP) is used to distribute the routes, and Multiprotocol Label Switching (MPLS) is used to indicate that particular packets need to follow particular routes. The characteristics of BGP/MPLS IP VPNs are compared with the requirements specified in [L3VPN-REQS].

A VPN service is provided by a Service Provider (SP) to a customer (sometimes referred to as an enterprise). BGP/MPLS IP VPNs are intended for the situation in which:

- The customer:

- * uses the VPN only for carrying IP packets.
- * does not want to manage a routed backbone; the customer may be using routing within his sites, but wishes to outsource the inter-site routing to the SP.
- * wants the SP to make the backbone and its routing completely transparent to the customer's own routing.

If the customer has a routed infrastructure at his sites, he does not want his site routing algorithms to need to be aware of any part of the SP backbone network, other than the Provider Edge (PE) routers to which the sites are attached. In particular, the customer does not want his routers to need to be aware of either the native structure of the SP backbone or an overlay topology of tunnels through the SP backbone.

- The Service Provider:

- * has an IP backbone, with MPLS-enabled edge routers, and possibly (though not necessarily) with MPLS-enabled core routers.

- * wants to provide a service that meets the customer requirements above.
- * does not want to maintain a distinct overlay topology of tunnels for each customer.

The basic principle is to model each VPN as a self-contained "internet", where each site makes one or more access connections to an SP, sends the SP its routing information, and then relies on the SP to distribute routing information to and from the other sites in that same VPN. The service differs from Internet service, however, in that the SP strictly controls the distribution of this routing information so that routes from within a VPN are not sent outside the VPN, unless that is explicitly authorized by the customer. In fact, even within the VPN, the distribution of routes may be controlled by the SP so as to meet some policy of the customer.

The routers at a given customer site need not be routing peers of the routers at other customer sites, and indeed need not know anything about the internal structure of other customer sites. In fact, different routing protocols may run at the different sites, with each site using whatever protocol is most appropriate for that particular site.

If EBGP (the BGP procedures used between BGP speakers from different Autonomous Systems) is used on the access links that connect a Provider Edge router (PE router) to a Customer Edge router (CE router), then the SP and the customer do NOT peer in any Interior Gateway Protocol (IGP), i.e., intra-domain routing algorithm).

BGP/MPLS IP VPNs are optimized for the situation in which a customer (an enterprise) expects a service provider to operate and maintain the customer's "backbone" (i.e., the customer's inter-site routing). As such, the service provider becomes a "business partner" of the enterprise. The technical mechanisms accommodate the case in which a number of closely cooperating SPs can jointly offer the VPN service to a customer, in that the BGP-based route distribution mechanisms can operate between different SPs. If a set of SPs has sufficient agreements with respect to Quality of Service (QoS), Service Level Agreement (SLA), etc., then the customer's VPN could have sites attached to different SPs from that set.

[BGP-MPLS-IP-VPN] specifies the inter-AS (Autonomous System) mechanisms that allow a single VPN to have sites attached to different SPs. However, the design center is not an environment where a given VPN is spread among a very large number (e.g., hundreds) of SPs.

In cases where remote offices, individual telecommuters, etc., must use the public Internet to access the VPN, it is possible to "tunnel" the remote traffic to a PE router, and the PE router will treat the traffic as if it had arrived over an interface connected to the PE. Remote Point-to-Point Protocol (PPP) connections can be tunneled via Layer 2 Tunneling Protocol (L2TP) to a PE router; IPsec tunnels can also be used to tunnel traffic to a PE router across the public Internet. Of course, when the public Internet is used, issues such as QoS and SLAs must be carefully considered.

Some customers want to connect their sites over the public Internet, creating a VPN "virtual backbone", purchasing connectivity for a given site from whatever Internet Service Provider (ISP) offers the best price for connecting that site. A BGP/MPLS IP VPN is not an appropriate solution for such customers; they instead need to consider solutions (either customer-managed or provider-managed) that interconnect their sites via an overlay of secure tunnels across the Internet. (See, for example, [IPSEC-VPN].)

Some customers who do not want to connect their sites via secure site-to-site tunnels across the Internet may nevertheless want to maintain complete control over the routing in their VPN backbone. These customers will not want a "managed routing service" such as is provided by BGP/MPLS IP VPNs, since that hides all details of the backbone routing and topology from the customer. Rather, they may prefer a "virtual router" service, in which the tunnels through the SP networks are visible as links to the customer's routing algorithm. (See, for example, [VR-VPN].)

2. SP Provisioning Model

If a particular VPN attaches to a particular PE router, the SP must configure that PE router with a VPN Routing and Forwarding table (VRF), a routing table that is specific to the specified VPN. (This is known as a VPN Forwarding Instance (VFI) in the language of [L3VPN-REQS] and [L3VPN-FRMWRK].) Each interface or sub-interface at that PE that attaches to a site in the specified VPN (i.e., each local access link of that VPN) must be configured so as to be associated with that VRF. Each such interface may be unnumbered or may be assigned an address that is unique within the VPN's address space. In general, a routing algorithm needs to be run on each of these links (though static routing can be used instead). The routing algorithm can be EIGRP, or an IGP such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). (If OSPF is used, the procedures of [VPN-OSPF] MUST be implemented.) If an IGP is run on the access links, the IGP MUST be a separate IGP instance, different

than the IGP instance running among the backbone routers, and different than the IGP instance running on the access links of any other VPN. Static routing is also allowed.

The VRF is populated automatically with routes distributed from locally attached CE routers via whatever routing algorithm is run on the PE/CE links. It is also populated automatically with routes distributed from other VRFs via BGP. Standard routing decision processes are used to automatically select the proper routes. Static configuration of routes in the VRF is optional.

Each PE router must run BGP, and must be pre-configured with the identities of a small set of BGP Route Reflectors, with which it is to peer via IBGP. ("IBGP" refers to the BGP procedures used between BGP speakers from the same Autonomous System.)

In lieu of using Route Reflectors, one could configure each PE with the identities of all the other PEs, and set up a full mesh of IBGP connections. While this might be adequate for small networks, it would not scale well to large networks; the use of Route Reflectors is necessary to achieve scalability. See section 4.3.3 of [BGP-MPLS-IP-VPN] for a more complete discussion of the use of Route Reflectors, and related scalability mechanisms such as Outbound Route Filtering.

Each VRF must be configured with three parameters:

- A Route Distinguisher. This is a globally unique 8-byte value. Each VRF may have a unique Route Distinguisher (RD), or there may be a single unique RD for an entire VPN. When BGP is used to distribute VPN routing information across the SP backbone, this value is prepended to the VPN's IPv4 address prefixes, creating a new address family, the VPN-IPv4 address family. Thus, even when two VPNs have overlapping IPv4 address spaces, they have unique VPN-IPv4 address spaces.
- One or more Export Route Targets. A Route Target (RT) is a globally unique 8-byte value that BGP carries, as the Extended Communities Route Target attribute, along with routes that are exported from the VRF.
- One or more Import Route Targets. This RT is used to select routes to be imported from other VRFs into this VRF.

In the simplest cases and most common cases, the Export RT, Import RT, and RD can be identical, and all VRFs in the same VPN will distribute routes to each other (a typical intranet). In more complex cases, they can be set differently, allowing a very fine

degree of control over the distribution of routes among VRFs. This can be used to create extranets or to enforce various customer policies. In complicated cases, particular Export RTs can be assigned to particular routes using router management mechanisms. One advantage to not requiring the RD to be the same as any RT is that this may allow an RD value to be automatically determined for each VRF; RT values, on the other hand, must always be configured.

Adding a new site to a VPN is a matter of attaching the site's CE router to a PE router, configuring the interface, and, if a VRF for that VPN already exists in the PE router, associating that interface with the VRF. If a VRF for that VPN does not already exist in the PE, then one must be configured as specified above. Changes to the configuration of a PE are automatically reflected via BGP to the other PEs.

The RTs and RDs are made unique by being structured as an SP identifier followed by a number which is assigned by the identified SP. SPs may be identified by their AS numbers, or by a registered IP address owned by that SP.

Although RTs are encoded as BGP Extended Communities, the encoding itself distinguishes them from any other kind of BGP Extended Community.

3. Supported Topologies and Traffic Types

The scheme is optimized for full inter-site connectivity, in the sense that this is what the simplest configurations provide.

However, the SP has full control, through the mechanism of Route Targets, of the distribution of routing information among the set of VRFs. This enables the SP to provide hub-and-spoke or partial mesh connectivity as well as full mesh connectivity.

Note that, strictly speaking, the scheme does not create a topology, as it does not create layer 2 connections among the sites. It does, however, allow for control over the IP connectivity among the sites. It is also possible to constrain the distribution of routing in arbitrary ways, e.g., so that data from site A to site B must travel through a third site C. (In fact, if it is desired to do so, this level of control can be specified at the granularity of a single route.)

It is possible for some of the routes from a particular customer site A to be distributed to one set of remote sites, while other routes from site A are distributed to a different set of remote sites. This is done with the Route Target mechanism previously described.

Unicast IP traffic is fully supported. Customer IP packets are passed transparently.

Multicast IP traffic is optionally supported, if the SP provides the optional mechanisms of [BGP-MPLS-MCAST-VPN]. There are, however, scaling implications to the use of these mechanisms. Discussion of these implications is deferred.

Non-IP traffic is not supported. If support for non-IP traffic is necessary, either the SP must additionally provide a layer 2 tunneling service or the customer must use IP tunneling.

In general, customer routers at different sites do not become routing peers. However, a customer may, if he so desires, allow routers at different sites to be routing peers over a link that is NOT part of the VPN service. Such peering relationships are known as "IGP backdoors". To ensure the proper operation of routing when IGP backdoors are present, each VPN route that is distributed by the SP is distributed along with a corresponding routing metric. This enables the customer's IGP to compare the "backdoor routes" properly with the routes that use the SP backbone. In the particular case where a customer running OSPF within his sites wishes to have IGP backdoors, he should run OSPF on the PE/CE link, and the PEs should run the procedures of [VPN-OSPF]. (The CEs do NOT require any special OSPF procedures.)

4. Isolated Exchange of Data and Routing Information

The Route Target mechanism is used to control the distribution of routing information, so that routes from one VPN do not get sent to another. VPN routes are treated by BGP as a different address family than general Internet routes. Routes from a VRF do not get leaked to the Internet unless the VRF has been explicitly configured to allow it (and this is NOT the default).

The way in which a particular VPN is divided into sites, or the topology of any particular VPN site, is hidden from the Internet and from other VPNs. (Of course, if a particular site can receive Internet traffic, and if it responds to traceroute probes from the Internet, then any user of the Internet can learn something about the site topology. The fact that the site is in a VPN does not make this any easier or any harder.)

Similarly, Internet routes do not get leaked into the VPN, unless a VRF of that VPN is explicitly configured to import the Internet routes.

Proper configuration is essential to maintaining the isolation. In particular, each access link must be associated with the proper VRF for that access link, and each VRF must be configured with the proper set of RTs.

A number of means for exchanging reachability information between the PE and CE devices are supported: static routing, EBGp, and RIP are supported by the procedures of [BGP-MPLS-IP-VPN]. If the procedures of [VPN-OSPF] and [OSPF-2547-DNBIT] are implemented, OSPF may be used. If OSPF is used between two VPN sites that are in the same OSPF area, and if it is desired for routes over the VPN backbone to be preferred to the OSPF intra-site routes, then the "sham link" procedures of [VPN-OSPF] must be used.

The routing protocols used among the customer routers are not in any way restricted by the VPN scheme, as whatever IGP is used within the VPN, the PE/CE access links may run EBGp, or may otherwise be in a different routing domain than the site's internal links.

BGP is used for passing routing information among SPs. BGP may be authenticated by use of the TCP MD5 option, or by operating through an IPsec tunnel.

Data traveling between two customer sites is encapsulated while in transit through the backbone. The encapsulation contains sufficient information to ensure that the packet is sent to the proper PE router, and then, in conjunction with the VRF and related information at that PE, to the proper CE routers.

If two VPNs attach to the same PE, there is strict separation of forwarding at that PE, as well as strict separation of the routing information.

Isolation of traffic is similar to that provided by classical L2 VPNs which are based on Frame Relay or Asynchronous Transfer Mode (ATM). As in classical L2 VPNs, the customer must rely on the SP to properly configure the backbone network to ensure proper isolation and to maintain the security of his communications gear.

5. Access Control and Authentication

No particular means of PE/CE authentication is specified for BGP/MPLS IP VPNs. PE/CE mutual authentication may be done via any mechanism supported by the routing protocol in which the CE and PE are peers (e.g., use of the TCP MD5 authentication when the PE/CE protocol is BGP), or by any other mechanism that may be desired. With such mechanisms in place, a CE may not join a VPN until the CE authenticates itself to the Service Provider.

There is, however, no standardized method that requires a CE to authenticate itself to the customer network (rather than to the SP) before the CE is allowed to join the VPN. This is for further study.

No particular means is specified for controlling which user data packets can be forwarded by BGP/MPLS IP VPNs. BGP/MPLS IP VPNs are compatible with Access Control Lists (ACLs) and any other filtering features that are supported on the PE routers. Routing can be set up so that extranet traffic is directly through a firewall, if that is desired.

It is possible for various sorts of "tunnel interfaces" to be associated with a VRF. In this case, whatever authentication is natively used in the establishment of the tunnel interface may be used. For example, an IPsec tunnel can be used as an "access link" to attach a remote user or site to a VRF. The authentication procedure in this case is part of IPsec, not part of the VPN scheme.

Where L2TP is used, each PPP session carried in an L2TP tunnel can be associated with a VRF. The SP's Authentication, Authorization, and Accounting (AAA) server can be used to determine the VPN to which the PPP session belongs, and then the customer's AAA server can be given the opportunity to authenticate that session as well.

6. Security Considerations

6.1. Protection of User Data

No particular means of ensuring user data security is specified for BGP/MPLS IP VPNs.

The optional procedures of [MPLS/BGP-IPsec] may be used to provide authentication and/or encryption of user data as it travels from the ingress PE to the egress PE. However, the data is exposed at those two PEs, as well as on the PE/CE access links.

The customer may provide his own user data security by using IPsec tunnels that terminate within the customer sites. Such tunnels are transparent to the VPN scheme. Schemes that discover the remote tunnel endpoints automatically and then set up the tunnels automatically as needed are the best fit with this VPN technology. Note that there is no requirement in general that IPsec tunnels between customer sites terminate at CE routers.

The use of end-to-end transport mode IPsec by the customer is also transparent to the VPN scheme. In fact, the VPN scheme is compatible with any use of security by the customer, as long as a cleartext IP header is passed from CE to PE.

When data must cross the Internet to reach the ingress PE router, IPsec tunnels between the end user and the PE router can be used; the PE router must then associate each IPsec tunnel with the proper VRF. This association would have to be based on user-specific information provided by the Internet Key Exchange (IKE) protocol, such as a VPN-id.

If data is going from one SP network to another, and must cross the public Internet to get between those two networks, IPsec tunnels can be used to secure the data. This would require bilateral agreement between the two SPs. BGP connections can also be passed through an IPsec tunnel if this is deemed necessary, in order to protect user data, by a pair of SPs. QoS/SLA factors would have to be carefully considered in this case.

6.2. SP Security Measures

The SP is responsible for preventing illegitimate traffic from entering a VPN. VPN traffic is always encapsulated while traveling on the backbone, so preventing illegitimate traffic is a matter of ensuring that the PE routers to the encapsulation/decapsulation correctly and that encapsulations have not been "spoofed", i.e., that the encapsulated packets were actually encapsulated by PE routers.

This requires the SP to take various security measures. The PE and P routers must themselves be secure against break-ins (either from someone physically present or from the Internet), and neither P nor PE routers should form routing adjacencies to other P or PE routers without benefit of some kind of security. This may be authentication in the IGP, or physical security.

The PE/CE access link should be secured in some manner, though the provider may make it the responsibility of the customer to ensure that the CE is secure from compromise. If the PE/CE access link is a tunnel over the Internet, then of course some sort of authentication protocol should always be used.

Label Distribution Protocol (LDP) sessions and BGP sessions between PE and/or P routers should be authenticated. This can be done via the TCP MD5 option or by use of IPsec.

If the SP is providing the VPN service over an MPLS backbone, it should not accept MPLS packets from its external interfaces (i.e., interfaces to CE devices or to other providers' networks) unless the top label of the packet was legitimately distributed to the system from which the packet is being received. If the packet's incoming interface leads to a different SP (rather than to a customer), an appropriate trust relationship must also be present, including the trust that the other SP also provides appropriate security measures.

If the SP is providing the VPN service by using an IP (rather than an MPLS) encapsulation, or if it accepts IP-encapsulated VPN packets from other SPs, it should apply filtering at its borders so that it does not accept from other SPs or from customers any IP packets that are addressed to the PE routers, unless appropriate trust relationships are in place.

Cryptographic authentication of the encapsulated data packets is certainly advantageous when there are multiple SPs providing a single VPN.

When a dynamic routing protocol is run on the link between a CE router and a PE router, routing instability in the private network may have an effect on the PE router. For example, an unusually large number of routing updates could be sent from the CE router to the PE router, placing an unusually large processing load on the PE router. This can potentially be used as a Denial-of-Service (DoS) attack on the PE router.

This issue can be mitigated via resource partitioning in the PE, in order to limit the amount of resources (e.g., CPU and memory) that any one VPN is permitted to use in PE routers. Also, rate limits may be applied to the routing traffic sent from the CE to the PE. Alternately, when this problem is detected, the CE-to-PE interface may be shut down.

Network management traffic from the CE to the PE may be rate limited (for example, to prevent network management traffic from CE to PE to be used in a DoS attack).

6.3. Security Framework Template

Section 9 of [L2VPN-SEC-FRMWRK] provides "a brief template that may be used to evaluate and summarize how a given PPVPN [Provider-Provisioned Virtual Private Network] approach (solution) measures up against the PPVPN Security Framework". It further states "an evaluation using this template should appear in the applicability statement for each PPVPN approach". The purpose of this subsection is to provide the information in the form required by this template. Security requirements that are relevant only to L2VPNs are not applicable and are not further discussed.

- Does the approach provides complete IP address space separation for each L3VPN?

Yes.

The IP address prefixes from a particular VPN appear in their native form only in routing tables that are specific to the particular VPN. They are distributed in their native form only by routing instances that are specific to the particular VPN. When address prefixes from different VPNs are combined into a common table, or distributed by a common mechanism, the address prefixes are first prepended with a Route Distinguisher (RD). The RD is a 64-bit quantity, structured so that globally unique RD values can easily be created by an SP. As long as no two VPNs are assigned the same RD value, complete IP address space separation is provided. It is however possible for an SP to misconfigure the RD assignments.

- Does the approach provide complete IP route separation for each L3VPN?

Yes.

The distribution of routes is controlled by assigning import and export Route Targets (RTs). A route that is exported from a VRF carries an RT specified by the SP as an export RT for that VRF. The route can be imported into other VRFs only if the RT that it carries has been configured by the SP as an import RT for those other VRFs. Thus, the SP has complete control over the set of VRFs to which a route will be distributed. It is of course possible for the SP to misconfigure the RT assignments.

- Does the approach provide a means to prevent improper cross-connection of sites in separate VPNs?

This requirement is addressed in a way that is beyond the scope of the VPN mechanisms.

In BGP/MPLS IP VPNs, an SP makes a particular site part of a particular VPN by configuring the PE router's interface to that site to be associated with a particular VRF in that PE. The VRF is configured with import and export RTs, and it is the way in which VRFs are configured with RTs in the various PEs that results in a particular set of sites being connected as a VPN.

Connecting the sites properly in this way is regarded as a network management function, and the VPN scheme itself does not provide a means to prevent misconfiguration.

The VPN scheme does not provide any particular method for ensuring that a given interface from a PE leads to the CE that is expected to be there. If a routing algorithm is run on a particular PE/CE interface, any security procedures that the routing algorithm provides (e.g., MD5 authentication of BGP sessions) can be used; this is outside the scope of the VPN scheme. Also, a CE can attach to a PE via an IPsec tunnel, if this is desired, for a greater degree of security.

- Does the approach provide a means to detect improper cross-connection of sites in separate VPNs?

The base specifications for BGP/MPLS IP VPNs do not provide a means for detecting that a site has been connected to the wrong VPN. However, the optional procedure specified in [CE-VERIF] does provide such a means. Basically, each PE obtains, via protocol, a secret from each CE to which it is directly attached. When the routes from a given CE are distributed, the secret from that CE is attached as an attribute of the route. This secret will ultimately be distributed to any other CE that receives any route from the given CE. A CE that is not supposed to be part of a given VPN will not know the right secret, and if it is connected to the given VPN the other CEs in that VPN will realize that a CE that doesn't know the proper secret has been connected to the VPN.

- Does the approach protect against the introduction of unauthorized packets into each VPN?

We must look separately at the various points at which one might attempt to introduce unauthorized packets.

* Packets arriving at a PE over a PE/CE interface

If a given PE is directly connected to a given CE, the PE will accept any packets that the CE sends it. The VPN scheme has no special procedures for determining that these packets actually came from the CE. However, various means of securing the PE/CE connection can be used (for instance, the PE and CE can be connected by an IPsec tunnel) if desired. That is, this aspect of the requirement can be addressed by means that are outside the scope of the VPN specification.

Once a packet has been accepted from a CE by a PE, the packet is routed according to the VRF associated with that PE's interface to that CE. Such packets can only be sent along routes that are in that VRF. There is no way a packet from a CE can be routed to an arbitrary VPN. In particular, there is nothing a VPN user can do to cause any particular packet to be sent to the wrong VPN. So this aspect of the requirement is fully addressed.

* Packets arriving at a PE over an interface from the backbone

The optional procedures of [MPLS/BGP-IPsec] can be used to ensure that a packet that is received by a PE from the backbone will not be recognized as a VPN packet unless it actually is one. Those procedures also ensure that a received VPN packet came from a particular PE and that it carries the MPLS label that that PE put on it. These procedures protect the packet from ingress PE to egress PE, but do not protect the PE/CE interfaces.

If the optional procedures of [MPLS/BGP-IPsec] are not used, then the following considerations apply.

Undetected corruption of the routing information carried in a packet's VPN encapsulation can result in misdelivery of the packet, possibly to the wrong VPN.

If a packet enters an SP's network on an interface other than a PE/CE interface, the SP should ensure that the packet either does not look like a VPN packet or else is not routed to a PE router. This can be done in a variety of ways that are outside the scope of the VPN scheme. For example, IP packets addressed to the PE routers can be filtered, MPLS packets (or, e.g., MPLS-in-IP) from outside the SP network can be refused, etc.

In the case of a multi-provider L3VPN backbone, the SP will have to know which interfaces lead to SPs that are VPN partners, so that VPN packets can be allowed to flow on those interfaces.

If the public Internet is used as the L3VPN backbone, protection against unauthorized packets cannot be achieved by the above measures. IPsec tunnels should always be used to carry VPN traffic across the public Internet.

- Does the approach provide confidentiality (secrecy) protection, sender authentication, integrity protection, or protection against replay attacks for PPVPN user data?

If these are desired, they must be provided by mechanisms that are outside the scope of the VPN mechanisms. For instance, the users can use secure protocols on an end-to-end basis, e.g., IPsec, Secure Shell (SSH), Secure Sockets Layer (SSL), etc.

- Does the approach provide protection against unauthorized traffic pattern analysis for PPVPN user data?

Preventing an observer from obtaining traffic pattern analysis from the SP network is beyond the scope of the VPN mechanisms.

- Do the control protocol(s) used for each of the following functions provide for message integrity and peer authentication?

- * VPN membership discovery

This requirement is fully satisfied. Membership discovery is done by means of BGP. Control message integrity and peer authentication in BGP may be achieved by use of the TCP MD5 option.

- * Tunnel establishment

The answer to this question depends of course on the tunnel protocol and tunnel establishment protocol; a variety of different tunneling schemes can be used in BGP/MPLS IP VPNs. Thus, this question is out of scope.

In the common case where the tunnels are MPLS Label Switching Routers (LSRs) established by LDP, then control message integrity and peer authentication may be achieved by use of the TCP MD5 option.

- * VPN topology and reachability advertisement

With respect to PE-PE interactions, the relevant control protocol is BGP, so control message integrity and peer authentication can be achieved by use of the TCP MD5 option.

With respect to CE-PE interactions, the answer depends on the protocol used for exchanging information between PE and CE, as the security mechanisms (if any) of those protocols would need to be used. In the common case where the PE/CE protocol is BGP, the TCP MD5 option can be used.

- * VPN provisioning and management

The protocols procedures for provisioning VPNs and managing the PE routers are outside the scope of the VPN scheme.

- * VPN monitoring and attack detection and reporting

The protocols and procedures for monitoring the VPNs are outside the scope of the VPN scheme.

- What protection does the approach provide against PPVPN-specific DoS attacks (i.e., inter-trusted-zone DoS attacks)?

- * Protection of the service provider infrastructure against Data Plane or Control Plane DoS attacks originated in a private (PPVPN user) network and aimed at PPVPN mechanisms.

The PE/CE interfaces of a given VPN will generally be addressable from within that VPN. Apart from that, a user within an L3VPN has no more access to the service provider infrastructure than does any user of the Internet. Therefore, we will focus in this section on possible DoS attacks against a PE router that may occur when traffic from within a VPN is addressed to a PE router.

A user within the VPN may address traffic to a PE router and may attempt to send an excessive amount of traffic to it. Presumably, the PE routers will not accept unauthorized TCP connections or Simple Network Management Protocol (SNMP) commands, so such traffic will be thrown away; the danger is that the PE may need to use a significant proportion of its capacity to discard such traffic. However, this case is no different than the case of any SP access router that attaches to subscriber equipment. The presence of the VPN mechanisms does not make the PE any more or less vulnerable to DoS attacks from arbitrary end users.

- * Protection of the service provider infrastructure against Data Plane or Control Plane DoS attacks originated in the Internet and aimed at PPVPN mechanisms.

DoS attacks of this sort can be prevented if the PE routers are not addressable from the Internet. Alternatively, an SP can apply address filtering at its boundaries so that packets from the Internet are filtered if they are addressed to a PE router.

- * Protection of PPVPN users against Data Plane or Control Plane DoS attacks originated from the Internet or from other PPVPN users and aimed at PPVPN mechanisms.

Mechanisms already discussed prevent users in a VPN from receiving packets from the Internet, unless this is specifically allowed. In the case where it is specifically allowed, it is no different than any other situation in which a network is connected to the Internet, and there is no special vulnerability to DoS attacks due to the L3VPN mechanisms.

There is nothing to prevent a user in a VPN from mounting a DoS attack against other users in the VPN. However, the L3VPN mechanisms make this neither more nor less likely.

- Does the approach provide protection against unstable or malicious operation of a PPVPN user network?

- * Protection against high levels of, or malicious design of, routing traffic from PPVPN user networks to the service provider network.

If a dynamic routing algorithm is running on the PE/CE interface, it can be used to mount an attack on the PE router, by having the CE present the PE with an excessive number of routing events. If an end user within a VPN successfully attacks the routing algorithm of the VPN, that might also result in an excessive number of routing events being seen by the PE router. This sort of attack can be ameliorated by having the PE limit the amount of its resources that can be expended processing routing events from a particular VPN. If the PE/CE routing algorithm is BGP, then such mechanisms as route flap damping may be appropriate as well.

- * Protection against high levels of, or malicious design of, network management traffic from PPVPN user networks to the service provider network.

A user in a BGP/MPLS IP VPN has no more ability than any Internet user to send management traffic to the service provider network.

- * Protection against worms and probes originated in the PPVPN user networks, sent towards the service provider network.

A user in a BGP/MPLS IP VPN has no more ability than any Internet user to send worms or probes to the service provider network.

7. Addressing

Overlapping customer addresses are supported. There is no requirement that such addresses be in conformance with [RFC1918]. There is no requirement that customer VPN addresses be distinct from addresses in the SP network.

Any set of addresses used in the VPN can be supported, irrespective of how they are assigned, how well they aggregate, and whether they are public or private. However, the set of addresses that are reachable from a given site must be unique.

Network address translation for packets leaving/entering a VPN is possible and is transparent to the VPN scheme.

There is nothing in the architecture to preclude the mechanisms from being extended to support IPv6, provided that the appropriate IPv6-capable routing algorithms are in place. That is, PE/CE routing must support IPv6, and the PE-PE BGP must support the labeled IPv6 address family. The latter has not been specified, but its specification is obvious from the specification of the labeled IPv4 address family. The IGP used in the SP backbone need not be IPv6 capable in order to support customer IPv6 networks.

In theory, the same could be said of other network layers, but in practice a customer who has non-IP traffic to carry must expect to carry it either in site-to-site IP tunnels or using some additional service (such as a layer 2 service) from the SP.

Layer 2 addresses and identifiers are never carried across the SP backbone.

No restrictions are placed on the customer's addressing schemes or policies. Note though that the SP may place restrictions on the number of routes from a given customer site, or may charge differentially depending on the number of such routes, and such restrictions may have implications for the customer's addressing scheme. In particular, addressing schemes that facilitate route aggregation on a per-site basis will result in the most efficient use of the SP's resources, and this may be reflected in SP charging policies.

8. Interoperability and Interworking

Interoperability should be ensured by proper implementation of the published standards.

Direct PE-PE interworking over the SP backbone with other VPN solutions is not supported.

As all the different types of L3VPNs are IP networks, they can of course interwork in the same way that any two IP networks can interwork. For example, a single site can contain a CE router of one VPN scheme and a CE router of another VPN scheme, and these CE routers could be IGP peers, or they might even be the same CE router. This would result in the redistribution of routes from one type of VPN to the other, providing the necessary interworking.

9. Network Access

9.1. Physical/Link Layer Topology

The architecture and protocols do not restrict the link layer or the physical layer in any manner.

9.2. Temporary Access

Temporary access via PPP is possible, using industry standard PPP-based authentication mechanisms. For example:

- A dial-up user (or other PPP user) is authenticated by the PE, using the SP's AAA server, based on a login string or on the number dialed.
- The SP's AAA server returns a VPN-id to PE.
- The PE assigns the user to a VRF, based on that VPN-id.

- The user is then authenticated by a AAA server within the VPN (i.e., managed by the customer rather than by the SP). This AAA server would typically be addressed through the VRF (i.e., may be in VPN's private address space).
- The user gets disconnected if either authentication step is unsuccessful.

IPsec access to a VRF is also possible. In this case, the security association is between the end user and the SP.

In these ways, a user can access a BGP/MPLS IP VPN via the public Internet.

There is no explicit support for mobility, other than what is stated above.

9.3. Access Connectivity

Homing of a CE to two or more PEs is fully supported, whether or not the PEs are on the same SP network.

If a CE is connected to two or more PEs, all its PE/CE links can be used to carry traffic in both directions. In particular, traffic from different ingress PEs to a particular CE may arrive at that CE over different PE/CE links. This depends on the backbone network routing between the CE and the various ingress PEs.

If a VRF on a particular ingress PE contains several routes to a particular destination, then traffic from that ingress PE can be split among these routes. If these routes end with different PE/CE links, then traffic from that ingress PE will be split among those links.

BGP contains a multitude of knobs that allow an SP to control the traffic sent on one PE/CE link as opposed to the other. One can also make use of the Link Bandwidth extended community [BGP-EXT-COMM] to control how traffic is distributed among multiple egress PE/CE links.

The VPN scheme is of course compatible with the use of traffic engineering techniques, Resource Reservation Protocol - Traffic Engineering (RSVP-TE) based or otherwise, in the backbone network.

10. Service Access

10.1. Internet Access

Internet access and VPN access are possible from the same site. This is even possible over the same interface, as long as the VPN's internal addresses are distinct from the addresses of the systems that must be reached via the Internet. This requires only that Internet routes as well as VPN routes be imported into the VRF associated with that interface. This may be as simple as putting a default route to the Internet into that VRF.

The "route to the Internet" that is in a particular VRF need not lead directly to the Internet; it may lead to a firewall or other security device at another site of the VPN. The VPN customer can cause this to happen simply by exporting a default route from the site with the firewall. Generally, a site with a firewall will use a different virtual interface for Internet access than for VPN access, since the firewall needs to distinguish the "clean interface" from the "dirty interface".

In such a configuration, the customer would export his routes to the Internet via the firewall's dirty interface, but would export the same routes to the VPN via the clean interface. Thus, all traffic from the Internet would come through the dirty interface, then through the firewall, and possibly go to another VPN site through the clean interface. This also allows any necessary Network Address Translation (NAT) functionality to be done in the firewall.

10.2. Other Services

Any externally provided service can be accessed from the VPN, provided that it can be addressed with an address that is not otherwise in use within the VPN. Access can be firewalled or non-firewalled. If the client accessing the service does not have a globally unique IP address, and a single server provides a service to multiple VPNs, NAT will have to be applied to the client's packets before they reach the server. This can be done at a customer site, or by a VRF-specific NAT function in a PE router.

11. SP Routing

Routing through the backbone is independent of the VPN scheme and is unaffected by the presence or absence of VPNs. The only impact is that the backbone routing must carry routes to the PE routers.

The VPN routes themselves are carried in BGP as a distinct address family, different than the address family that is used to carry "ordinary" IP routes. These routes are passed from PE router to Route Reflector to PE router, and are never seen by the P routers. The Route Reflectors that carry the VPN routes can be entirely separate from the Route Reflectors that carry the "ordinary" IP routes.

The fact that two PE routers support a common VPN does not require those PE routers to form an IGP routing adjacency between themselves. The number of adjacencies in the backbone IGP is independent of and unrelated to the number of VPNs supported by any set of PE routers.

No VPN-specific protection and restoration mechanisms are needed; these are general routing considerations, and the VPN scheme is compatible with any protection and restoration mechanisms that may be available.

The SP does not manage the customer's IGP in any way, and routes are never leaked between the SP's IGP and any customer's IGP.

If the PE/CE protocol is EBGP, the SP and the customer do not ever participate in a common IGP.

12. Migration Impact

Generally, this means replacement of an existing legacy backbone with VPN backbone. The general migration mechanism would be to hook up the sites one at a time to the VPN backbone, and to start giving the routes via the VPN backbone preference to routes via the legacy backbone. Details depend on the legacy backbone's IGP. In general, one would have to manipulate the IGP metrics to provide the proper route preference.

If the legacy backbone routing protocol is OSPF, then migration is best done with OSPF as the PE/CE protocol and the PE supporting the [VPN-OSPF] procedures, OR with BGP as the PE/CE protocol, and the CE supporting the BGP/OSPF interaction specified in [VPN-OSPF].

With other legacy backbone routing protocols, the proper metrics must be set at the point (PE or CE) where the BGP routes from the SP network are being redistributed into the legacy IGP.

13. Scalability

There is no upper limit on the number of VPNs per SP network, as there is no one box in the SP network that needs to know of all VPNs. Knowledge of a particular VPN is confined to the PE routers that attach to sites in that VPN, and to the BGP Route Reflectors that receive routing data from those PEs; other systems maintain no state at all for the VPN. Note though that there is no need for any one Route Reflector to know of all VPNs.

If the SP is providing the VPN service over an MPLS backbone, then the backbone IGP must carry a host route for every Label Switched Path (LSP) egress node within the routing domain. Every PE router in the routing domain is an LSP egress node. If there are VPNs attached to PE routers that are within the routing domain, as well as PE routers that are in some second routing domain, then the border routers leading towards the second routing domain will also be LSP egress nodes. Thus, the sum of the number of PE routers plus number of border routers within a routing domain is limited by the number of routes that can be carried within the domain's IGP. This does not seem to create any practical scalability issue.

There is no upper limit on the number of site interfaces per VPN, as state for a particular interface is maintained only at the PE router to which that interface attaches. The number of site interfaces per VPN at a given PE router is limited only by the number of interfaces that that PE router can support.

The number of routes per VPN is constrained only by the number of routes that can be supported in BGP, the number of routes that can be maintained in the PEs that attach to that VPN, and the number of routes that can be maintained in the BGP Route Reflectors that hold the routes of that VPN.

The major constraint in considering scalability is the number of routes that a given PE can support. In general, a given PE can support as many VPNs as it has interfaces (including virtual interfaces or "sub-interfaces", not just physical interfaces), but it is constrained in the total number of routes it can handle. The number of routes a given PE must handle depends on the particular set of VPNs it attaches to, and the number of routes in each such VPN, and the number of "non-VPN" Internet routes (if any) that it must also handle.

The SP may need to engage in significant planning to ensure that these limits are not often reached. If these limits are reached, it may be necessary either to replace the PE with one of larger capacity or to reorganize the way in which access links lead from CEs to PEs,

in order to better concentrate the set of access links from sites that are in the same VPN. Rehoming a site to a different PE may not involve actual rewiring; if the access technology is switched, this is a matter of provisioning, but may still be a significant undertaking. If it is necessary to have downtime while performing the rehoming, the customer is impacted as well. Rehoming can also be done "virtually", by creating a layer 2 tunnel from a CE's "old" PE to its "new" PE.

An important consideration to remember is that one may have any number of INDEPENDENT BGP systems carrying VPN routes. This is unlike the case of the Internet, where the Internet BGP system must carry all the Internet routes. The difference stems from the fact that all Internet addresses must be reachable from each other, but a given VPN address is only supposed to be reachable from other addresses in the same VPN.

Scalability is also affected by the rate of changes in the reachability advertisements from CE to PE, as changes reported by a CE to its attached PE may be propagated to the other PEs. BGP mechanisms to control the rate of reported changes should be used by the SP.

Another constraint on the number of VPNs that can be supported by a particular PE router is based on the number of routing instances that the PE router can support. If the PE/CE routing is static, or is done by BGP, the number of routing protocol instances in a PE device does not depend on the number of CEs supported by the PE device. In the case of BGP, a single BGP protocol instance can support all CEs that exchange routing information using BGP. If the PE/CE router is done via RIP or OSPF, then the PE must maintain one RIP or OSPF instance per VRF. Note that the number of routing instances that can be supported may be different for different routing protocols.

Inter-AS scenarios constructed according to option (b) of section 10 of [BGP-MPLS-IP-VPN] require BGP "border routers" to hold the routes for a set of VPNs. If two SPs share in a small number of VPNs, a single border router between them provides adequate capacity. As the number of shared VPNs increases, additional border routers may be needed to handle the increased number of routes. Again, no single border router would handle all the routes from all the VPNs, so an increase in the number of VPNs can always be supported by adding more border routers.

Inter-AS scenarios constructed according to option (c) of section 10 of [BGP-MPLS-IP-VPN] eliminate the need for border routers to contain VPN routes (thus improving scalability in that dimension), but at the cost of requiring that each AS have a route to the PEs in the others.

(Inter-AS scenarios constructed according to option (a) of section 10 of [BGP-MPLS-IP-VPN] do not scale well.)

The solution of [BGP-MPLS-IP-VPN] is intended to simplify CE and site operations, by hiding the structure of the rest of the VPN from a site, and by hiding the structure of the backbone. Thus, CEs need have only a single sub-interface to the backbone, CEs at one site need not even be aware of the existence of CEs at another, and CEs at one site need not be routing peers of CEs at another. CEs are never routing peers of P routers. These factors help to scale the customer's network, but limiting the number of adjacencies each CE must see, and by limiting the total number of links that the customer's IGP must handle.

The solution of [BGP-MPLS-IP-VPN] is also intended to simplify the SP's VPN provisioning, so that potentially the SP will have to do little more than say which sites belong to which VPNs. However, as the system scales up, planning is needed to determine which PEs should home which VPNs, and which BGP RRs should take which VPNs' routing information.

P routers maintain NO per-VPN state at all; the only requirement on them is to maintain routes to the PE routers. When MPLS is used, a P router must also maintain one multipoint-to-point LSP for each such route.

However, certain VPN multicast schemes require per-multicast-group state in the P routers, summed over all VPNs. Others require only no state in the P routers at all, but will result in sending more unnecessary traffic. The complete set of tradeoffs for multicast is not that well understood yet.

Note that as the scaling of a particular PE is primarily a matter of the total number of routes that it must maintain, scalability is facilitated if the addresses are assigned in a way that permits them to be aggregated (i.e., if the customers have a sensible addressing plan).

When a dynamic routing protocol is run on the link between a CE router and a PE router, routing instability in the private network may have an effect on the PE router. For example, an unusually large number of routing updates could be sent from the CE router to the PE router, placing an unusually large processing load on the PE router.

This issue can be mitigated via resource partitioning in the PE, in order to limit the amount of resources (e.g., CPU and memory) that any one VPN is permitted to use in PE routers. Also, rate limits may be applied to the routing traffic sent from the CE to the PE.

Alternately, when this problem is detected, the CE-to-PE interface may be shut down.

14. QoS, SLA

The provision of appropriate QoS capabilities may require any combination of the following:

- QoS in the access network.
- Admission control (policing) by the PE router on the ingress access links.
- Traffic conditioning (shaping) by the PE router on the ingress access links.
- Traffic engineering in the backbone.
- Intserv/diffserv classification by the PE, for traffic arriving from the CE. Once the PE classifies the user packets, this classification needs to be preserved in the encapsulation (MPLS or IP) used to send the packet across the backbone.
- Differentiated Services Codepoint (DSCP) mapping.
- DSCP transparency.
- Random Early Discard in the backbone.

None of these features are VPN-specific. The ability to support them depends on whether the features are available on the edge and core platforms, rather than on any particular VPN scheme.

MPLS support for differentiated services is detailed in RFC 3270 [MPLS-DIFFSERV]. DSCP mapping and transparency are covered in section 2.6 of that document.

It is possible to use traffic engineering to provide, e.g., guaranteed bandwidth between two PEs for the traffic of a given VPN. The VRF entries for that VPN in each PE need to be modified so that the traffic to the other PE is directed onto the traffic-engineered path. How this is done is a local matter.

BGP/MPLS IP VPNs can support both the "hose model" and the "pipe model" of QoS. In the "pipe model", a particular quality of service (e.g., a guaranteed amount of bandwidth) would be applied to all or some of the packets traveling between a given pair of CEs. In the "hose model", a particular quality of service (e.g., a guaranteed

amount of bandwidth) would be applied to all traffic to or from a particular CE, irrespective of which other CE the traffic is going to or coming from. Since BGP/MPLS IP VPNs do not usually make use of CE-CE tunnels, the hose model is the more natural fit. Providing the pipe model would require the use of traffic engineering to explicitly create the necessary tunnels.

Many of the requirements specified in [L3VPN-REQS] stipulate that the Network Monitoring System (NMS) should support SLA monitoring and verification between the SP and the various customers by measurement of the indicators defined within the context of the SLA. The measurement of these indicators (i.e., counters) can be achieved when BGP/MPLS IP VPNs are used by employing a combination of the Management Information Base (MIB) module designed for BGP/MPLS IP VPNs [L3VPN-MIB] as well as other standard MIB modules such as the IF-MIB [IF-MIB]. Devices supporting these MIB modules can calculate SLAs based on real-time performance measurements using indicators and threshold crossing alerts. Devices can make these thresholds configurable either via a management interface such as SNMP.

15. Management

The L3VPN Requirements document [L3VPN-REQS] stipulates that the term "Provider Provisioned VPN" refers to VPNs for which the service provider participates in management and provisioning of the VPN. RFC BGP/MPLS IP VPNs can be provisioned and managed to meet these requirements. The following subsections will outline how devices supporting BGP/MPLS IP VPNs can satisfy these requirements.

15.1. Management by the Provider

The SP manages all the VPN-specific information in the PE device. This can be done using the MIB designed for BGP/MPLS IP VPNs [L3VPN-MIB], in combination with other standard MIB modules such as IF-MIB [IF-MIB], and other MPLS MIB modules [LSRMIB], [LDPMIB], [TEMIB], [FTNMIB].

Devices supporting BGP/MPLS IP VPNs that employ the management interface characteristics described above will also support the ITU-T Telecommunications Management Network Model "FCAPS" functionalities as required in the L3VPN Requirements document. These include Fault, Configuration, Accounting, Provisioning, and Security.

In BGP/MPLS IP VPNs, the SP is not required to manage the CE devices. However, if it is desired for the SP to do so, the SP may manage CE devices from a central site, provided that a route to the central site is exported into the CE's VPN, and the central site is in a VPN into which the routes to the managed CE devices have been imported.

This is a form of extranet.

If the central site is managing CE devices from several VPNs, those CE devices must have mutually unique addresses. Note that this does not enable the CE devices from different VPNs to reach each other.

The CE devices have no VPN-specific information in them. Hence the fact that they are connected together into a VPN does not require them to have any VPN-specific management MIB modules or capabilities.

15.2. Management by the Customer

CE devices may be managed from within the VPN, transparently to the SP. The CE devices have no VPN-specific information in them, and the fact that they are tied together into a VPN does not impact the customer's management of them.

Customer access to a PE device is totally at the discretion of the SP, but is not required by the solution. The PE device is a routing peer of a CE device, and can be pinged, etc.

If a customer is permitted to access the PE router for management purposes, the functions available to any particular customer need to be strictly controlled, and the use of resource partitioning may be appropriate.

Network management traffic from the CE to the PE may be rate limited (for example, to prevent network management traffic from CE to PE to be used in a DoS attack).

16. Acknowledgements

Many thanks to Jeremy De Clercq, Luyuan Fang, Dave McDysan, Ananth Nagarajan, Yakov Rekhter, and Muneyoshi Suzuki, for their comments, criticisms, and help in preparing this document. Thanks also to Thomas Nadeau for his help with the section on management, to Francois LeFaucheur for his help with the section on QoS, and to Ross Callon for his review of the document.

17. Normative References

- [BGP-EXT-COMM] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.
- [BGP-MPLS-IP-VPN] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [L3VPN-FRMWRK] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4110, July 2005.
- [L3VPN-REQS] Carugi, M. and D. McDysan, "Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks (PPVPNs)", RFC 4031, April 2005.
- [L2VPN-SEC-FRMWRK] Fang, L., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, July 2005.

18. Informative References

- [VPN-OSPF] Rosen, E., Psenak, P., and P. Pillay-Esnault, "OSPF as the PE/CE Protocol in BGP/MPLS VPNs", Work in Progress, February 2004.
- [OSPF-2547-DNBIT] Rosen, E., Psenak, P., and P. Pillay-Esnault, "Using an LSA Options Bit to Prevent Looping in BGP/MPLS IP VPNs", Work in Progress, March 2004.
- [MPLS/BGP-IPsec] Rosen, E., De Clercq, J., Paridaens, O., T'Joens, Y., and C. Sargor, "Architecture for the Use of PE-PE IPsec Tunnels in BGP/MPLS IP VPNs", Work in Progress, March 2004.
- [BGP-MPLS-MCAST-VPN] Rosen, E., Cai, Y., and IJ. Wijsnands, "Multicast in MPLS/BGP VPNs", Work in Progress, May 2004.
- [CE-VERIF] Bonica, R., Rekhter, Y., Raszuk, R., Rosen, E., and D. Tappan, "CE-to-CE Member Verification for Layer 3 VPNs", Work in Progress, September 2003.

- [FTNMIB] Nadeau, T., Srinivasan, C., and A. Viswanathan, "Multiprotocol Label Switching (MPLS) Forwarding Equivalence Class To Next Hop Label Forwarding Entry (FEC-To-NHLFE) Management Information Base (MIB)", RFC 3814, June 2004.
- [IPSEC-VPN] De Clercq, J., Paridaens, O., Krywaniuk, A., and C. Wang, "An Architecture for Provider Provisioned CE-based Virtual Private Networks using IPsec", Work in Progress, February 2004.
- [LDPMIB] Cucchiara, J., Sjostrand, H., and J. Luciani, "Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)", RFC 3815, June 2004.
- [LSRMIB] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)", RFC 3813, June 2004.
- [MPLS-DIFFSERV] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.
- [L3VPN-MIB] Nadeau, T. and H. Van Der Linde, "MPLS/BGP Virtual Private Network Management Information Base Using SMiv2", Work in Progress, August 2004.
- [IF-MIB] McCloghrie, K. and F. Kastenholtz, "The Interfaces Group MIB", RFC 2863, June 2000.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [TEMIB] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004.

[VR-VPN]

Knight, P., Ould-Brahim, H., and B. Gleeson,
"Network Based IP VPN Architecture using Virtual
Routers", Work in Progress, April 2004.

Author's Address

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719

EMail: erosen@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any 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.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

