Network Working Group
Request for Comments: 4381
Category: Informational

M. Behringer
Cisco Systems Inc

February 2006

Analysis of the Security of 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

IESG Note

The content of this RFC was at one time considered by the IETF, and therefore it may resemble a current IETF work in progress or a published IETF work. This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose, and in particular notes that the decision to publish is not based on IETF review for such things as security, congestion control or inappropriate interaction with deployed protocols. The RFC Editor has chosen to publish this document at its discretion. Readers of this RFC should exercise caution in evaluating its value for implementation and deployment. See RFC 3932 for more information.

Abstract

This document analyses the security of the BGP/MPLS IP virtual private network (VPN) architecture that is described in RFC 4364, for the benefit of service providers and VPN users.

The analysis shows that BGP/MPLS IP VPN networks can be as secure as traditional layer-2 VPN services using Asynchronous Transfer Mode (ATM) or Frame Relay.

Table of Contents

1. Scope and Introduction

As Multiprotocol Label Switching (MPLS) is becoming a more widespread technology for providing IP virtual private network (VPN) services, the security of the BGP/MPLS IP VPN architecture is of increasing concern to service providers and VPN customers. This document gives an overview of the security of the BGP/MPLS IP VPN architecture that is described in RFC 4364 [1], and compares it with the security of traditional layer-2 services such as ATM or Frame Relay.

The term "MPLS core" is defined for this document as the set of Provider Edge (PE) and provider (P) routers that provide a BGP/MPLS IP VPN service, typically under the control of a single service provider (SP). This document assumes that the MPLS core network is trusted and secure. Thus, it does not address basic security concerns such as securing the network elements against unauthorised access, misconfigurations of the core, or attacks internal to the core. A customer that does not wish to trust the service provider network must use additional security mechanisms such as IPsec over the MPLS infrastructure.

This document analyses only the security features of BGP/MPLS IP VPNs, not the security of routing protocols in general. IPsec technology is also not covered, except to highlight the combination of MPLS VPNs with IPsec.

The overall security of a system has three aspects: the architecture, the implementation, and the operation of the system. Security issues can exist in any of these aspects. This document analyses only the architectural security of BGP/MPLS IP VPNs, not implementation or operational security issues.

This document is targeted at technical staff of service providers and enterprises. Knowledge of the basic BGP/MPLS IP VPN architecture as described in RFC 4364 [1] is required to understand this document. For specific Layer 3 VPN terminology and reference models refer to [11].

Section 2 of this document specifies the typical VPN requirements a VPN user might have, and section 3 analyses how RFC 4364 [1] addresses these requirements. Section 4 discusses specific security issues of multi-AS (Autonomous System) MPLS architectures, and section 5 lists security features that are not covered by this architecture and therefore need to be addressed separately. Section 6 highlights potential security issues on layer 2 that might impact the overall security of a BGP/MPLS IP VPN service. The findings of this document are summarized in section 7.

2. Security Requirements of VPN Networks

Both service providers offering any type of VPN services and customers using them have specific demands for security. Mostly, they compare MPLS-based solutions with traditional layer 2-based VPN solutions such as Frame Relay and ATM, since these are widely deployed and accepted. This section outlines the typical security requirements for VPN networks. The following section discusses if and how BGP/MPLS IP VPNs address these requirements, for both the MPLS core and the connected VPNs.

2.1. Address Space, Routing, and Traffic Separation

Non-intersecting layer 3 VPNs of the same VPN service are assumed to have independent address spaces. For example, two non-intersecting VPNs may each use the same 10/8 network addresses without conflict. In addition, traffic from one VPN must never enter another VPN. This implies separation of routing protocol information, so that routing tables must also be separate per VPN. Specifically:

- Any VPN must be able to use the same address space as any other VPN.
- Any VPN must be able to use the same address space as the MPLS core.
- Traffic, including routing traffic, from one VPN must never flow to another VPN.
- Routing information, as well as distribution and processing of that information, for one VPN instance must be independent from any other VPN instance.
- Routing information, as well as distribution and processing of that information, for one VPN instance must be independent from the core.

From a security point of view, the basic requirement is to prevent packets destined to a host a.b.c.d within a given VPN reaching a host with the same address in another VPN or in the core, and to prevent routing packets to another VPN even if it does not contain that destination address.

Confidentiality, as defined in the L3VPN Security Framework [11], is a requirement that goes beyond simple isolation of VPNs and provides protection against eavesdropping on any transmission medium. Encryption is the mechanism used to provide confidentiality. This document considers confidentiality an optional VPN requirement, since many existing VPN deployments do not encrypt transit traffic.

2.2. Hiding the Core Infrastructure

The internal structure of the core network (MPLS PE and P elements) should not be externally visible. Whilst breaking this requirement is not a security problem in itself, many service providers believe it is advantageous if the internal addresses and network structure are hidden from the outside world. An argument is that denial-of-service (DoS) attacks against a core router are much easier to carry out if an attacker knows the router addresses. Addresses can always be guessed, but attacks are more difficult if addresses are not known. The core should be as invisible to the outside world as a comparable layer 2 infrastructure (e.g., Frame Relay, ATM). Core network elements should also not be accessible from within a VPN.

Security should never rely entirely on obscurity, i.e., the hiding of information. Services should be equally secure if the implementation is known. However, there is a strong market perception that hiding of details is advantageous. This point addresses that market perception.

2.3. Resistance to Attacks

There are two basic types of attacks: DoS attacks, where resources become unavailable to authorised users, and intrusions, where resources become available to unauthorised users. BGP/MPLS IP VPN networks must provide at least the same level of protection against both forms of attack as current layer 2 networks.

For intrusions, there are two fundamental ways to protect the network: first, to harden protocols that could be abused (e.g., Telnet into a router), and second, to make the network as inaccessible as possible. This is achieved by a combination of packet filtering / firewalling and address hiding, as discussed above.

DoS attacks are easier to execute, since a single known IP address might be enough information to attack a machine. This can be done using normal "permitted" traffic, but using higher than normal packet rates, so that other users cannot access the targeted machine. The only way to be invulnerable to this kind of attack is to make sure that machines are not reachable, again by packet filtering and optionally by address hiding.

This document concentrates on protecting the core network against attacks from the "outside", i.e., the Internet and connected VPNs. Protection against attacks from the "inside", i.e., an attacker who has logical or physical access to the core network, is not discussed here.

2.4. Impossibility of Label Spoofing

Assuming the address and traffic separation discussed above, an attacker might try to access other VPNs by inserting packets with a label that he does not "own". This could be done from the outside, i.e., another Customer Edge (CE) router or from the Internet, or from within the MPLS core. The latter case (from within the core) will not be discussed, since we assume that the core network is provided securely. Should protection against an insecure core be required, it is necessary to use security protocols such as IPsec across the MPLS infrastructure, at least from CE to CE, since the PEs belong to the core.

Depending on the way that CE routers are connected to PE routers, it might be possible to intrude into a VPN that is connected to the same PE, using layer 2 attack mechanisms such as 802.1Q-label spoofing or ATM VPI/VCI spoofing. Layer 2 security issues will be discussed in section 6.

It is required that VPNs cannot abuse the MPLS label mechanisms or protocols to gain unauthorised access to other VPNs or the core.

3. Analysis of BGP/MPLS IP VPN Security

In this section, the BGP/MPLS IP VPN architecture is analysed with respect to the security requirements listed above.

3.1. Address Space, Routing, and Traffic Separation

BGP/MPLS allows distinct IP VPNs to use the same address space, which can also be private address space (RFC 1918 [2]). This is achieved by adding a 64-bit Route Distinguisher (RD) to each IPv4 route, making VPN-unique addresses also unique in the MPLS core. This "extended" address is also called a "VPN-IPv4 address". Thus, customers of a BGP/MPLS IP VPN service do not need to change their current addressing plan.

Each PE router maintains a separate Virtual Routing and Forwarding instance (VRF) for each connected VPN. A VRF includes the addresses of that VPN as well as the addresses of the PE routers with which the CE routers are peering. All addresses of a VRF, including these PE addresses, belong logically to the VPN and are accessible from the VPN. The fact that PE addresses are accessible to the VPN is not an issue if static routing is used between the PE and CE routers, since packet filters can be deployed to block access to all addresses of the VRF on the PE router. If dynamic routing protocols are used, the CE routers need to have the address of the peer PE router in the core configured. In an environment where the service provider manages the

CE routers as CPE, this can be invisible to the customer. The address space on the CE-PE link (including the peering PE address) is considered part of the VPN address space. Since address space can overlap between VPNs, the CE-PE link addresses can overlap between VPNs. For practical management considerations, SPs typically address CE-PE links from a global pool, maintaining uniqueness across the core.

Routing separation between VPNs can also be achieved. Each VRF is populated with routes from one VPN through statically configured routes or through routing protocols that run between the PE and CE router. Since each VPN is associated with a separate VRF there is no interference between VPNs on the PE router.

Across the core to the other PE routers separation is maintained with unique VPN identifiers in multiprotocol BGP, the Route Distinguishers (RDs). VPN routes including the RD are exclusively exchanged between PE routers by Multi-Protocol BGP (MP-BGP, RFC 2858 [8]) across the core. These BGP routing updates are not re-distributed into the core, but only to the other PE routers, where the information is kept again in VPN-specific VRFs. Thus, routing across a BGP/MPLS network is separate per VPN.

On the data plane, traffic separation is achieved by the ingress PE pre-pending a VPN-specific label to the packets. The packets with the VPN labels are sent through the core to the egress PE, where the VPN label is used to select the egress VRF.

Given the addressing, routing, and traffic separation across an BGP/MPLS IP VPN core network, it can be assumed that this architecture offers in this respect the same security as a layer-2 VPN. It is not possible to intrude from a VPN or the core into another VPN unless this has been explicitly configured.

If and when confidentiality is required, it can be achieved in BGP/MPLS IP VPNs by overlaying encryption services over the network. However, encryption is not a standard service on BGP/MPLS IP VPNs. See also section 5.2.

3.2. Hiding of the BGP/MPLS IP VPN Core Infrastructure

Service providers and end-customers do not normally want their network topology revealed to the outside. This makes attacks more difficult to execute: If an attacker doesn't know the address of a victim, he can only guess the IP addresses to attack. Since most DoS attacks don't provide direct feedback to the attacker it would be difficult to attack the network. It has to be mentioned specifically

that information hiding as such does not provide security. However, in the market this is a perceived requirement.

With a known IP address, a potential attacker can launch a DoS attack more easily against that device. Therefore, the ideal is to not reveal any information about the internal network to the outside world. This applies to the customer network and the core. A number of additional security measures also have to be taken: most of all, extensive packet filtering.

For security reasons, it is recommended for any core network to filter packets from the "outside" (Internet or connected VPNs) destined to the core infrastructure. This makes it very hard to attack the core, although some functionality such as pinging core routers will be lost. Traceroute across the core will still work, since it addresses a destination outside the core.

MPLS does not reveal unnecessary information to the outside, not even to customer VPNs. The addressing of the core can be done with private addresses (RFC 1918 [2]) or public addresses. Since the interface to the VPNs as well as the Internet is BGP, there is no need to reveal any internal information. The only information required in the case of a routing protocol between PE and CE is the address of the PE router. If no dynamic routing is required, static routing on unnumbered interfaces can be configured between the PE and CE. With this measure, the BGP/MPLS IP VPN core can be kept completely hidden.

Customer VPNs must advertise their routes to the BGP/MPLS IP VPN core (dynamically or statically), to ensure reachability across their VPN. In some cases, VPN users prefer that the service provider have no visibility of the addressing plan of the VPN. The following has to be noted: First, the information known to the core is not about specific hosts, but networks (routes); this offers a degree of abstraction. Second, in a VPN-only BGP/MPLS IP VPN network (no Internet access) this is equal to existing layer-2 models, where the customer has to trust the service provider. Also, in a Frame Relay or ATM network, routing and addressing information about the VPNs can be seen on the core network.

In a VPN service with shared Internet access, the service provider will typically announce the routes of customers who wish to use the Internet to his upstream or peer providers. This can be done directly if the VPN customer uses public address space, or via Network Address Translation (NAT) to obscure the addressing information of the customers' networks. In either case, the customer does not reveal more information than would be revealed by a general Internet service. Core information will not be revealed, except for

the peering address(es) of the PE router(s) that hold(s) the peering with the Internet. These addresses must be secured as in a traditional IP backbone.

In summary, in a pure MPLS-VPN service, where no Internet access is provided, information hiding is as good as on a comparable FR or ATM network. No addressing information is revealed to third parties or the Internet. If a customer chooses to access the Internet via the BGP/MPLS IP VPN core, he will have to reveal the same information as required for a normal Internet service. NAT can be used for further obscurity. Being reachable from the Internet automatically exposes a customer network to additional security threats. Appropriate security mechanisms have to be deployed such as firewalls and intrusion detection systems. This is true for any Internet access, over MPLS or direct.

A BGP/MPLS IP VPN network with no interconnections to the Internet has security equal to that of FR or ATM VPN networks. With an Internet access from the MPLS cloud, the service provider has to reveal at least one IP address (of the peering PE router) to the next provider, and thus to the outside world.

3.3. Resistance to Attacks

Section 3.1 shows that it is impossible to directly intrude into other VPNs. Another possibility is to attack the MPLS core and try to attack other VPNs from there. As shown above, it is impossible to address a P router directly. The only addresses reachable from a VPN or the Internet are the peering addresses of the PE routers. Thus, there are two basic ways that the BGP/MPLS IP VPN core can be attacked:

1. By attacking the PE routers directly.
2. By attacking the signaling mechanisms of MPLS (mostly routing).

To attack an element of a BGP/MPLS IP VPN network, it is first necessary to know the address of the element. As discussed in section 3.2, the addressing structure of the BGP/MPLS IP VPN core is hidden from the outside world. Thus, an attacker cannot know the IP address of any router in the core to attack. The attacker could guess addresses and send packets to these addresses. However, due to the address separation of MPLS each incoming packet will be treated as belonging to the address space of the customer. Thus, it is impossible to reach an internal router, even by guessing IP addresses. There is only one exception to this rule, which is the peer interface of the PE router. This address of the PE is the only attack point from the outside (a VPN or Internet).

The routing between a VPN and the BGP/MPLS IP VPN core can be configured two ways:

1. Static: In this case, the PE routers are configured with static routes to the networks behind each CE, and the CEs are configured to statically point to the PE router for any network in other parts of the VPN (mostly a default route). There are two sub-cases: The static route can point to the IP address of the PE router or to an interface of the CE router (e.g., serial0).
2. Dynamic: A routing protocol (e.g., Routing Information Protocol (RIP), OSPF, BGP) is used to exchange routing information between the CE and PE at each peering point.

In the case of a static route that points to an interface, the CE router doesn't need to know any IP addresses of the core network or even of the PE router. This has the disadvantage of needing a more extensive (static) configuration, but is the most secure option. In this case, it is also possible to configure packet filters on the PE interface to deny any packet to the PE interface. This protects the router and the whole core from attack.

In all other cases, each CE router needs to know at least the router ID (RID, i.e., peer IP address) of the PE router in the core, and thus has a potential destination for an attack. One could imagine various attacks on various services running on a router. In practice, access to the PE router over the CE-PE interface can be limited to the required routing protocol by using access control lists (ACLs). This limits the point of attack to one routing protocol, for example, BGP. A potential attack could be to send an extensive number of routes, or to flood the PE router with routing updates. Both could lead to a DoS, however, not to unauthorised access.

To reduce this risk, it is necessary to configure the routing protocol on the PE router to operate as securely as possible. This can be done in various ways:

- By accepting only routing protocol packets, and only from the CE router. The inbound ACL on each CE interface of the PE router should allow only routing protocol packets from the CE to the PE.
- By configuring MD5 authentication for routing protocols. This is available for BGP (RFC 2385 [6]), OSPF (RFC 2154 [4]), and RIP2 (RFC 2082 [3]), for example. This avoids packets being spoofed from other parts of the customer network than the CE router. It requires the service provider and customer to agree on a shared secret between all CE and PE routers. It is necessary to do this for all VPN customers. It is not sufficient to do this only for the customer with the highest security requirements.

- By configuring parameters of the routing protocol to further secure this communication. For example, the rate of routing updates should be restricted where possible (in BGP through damping); a maximum number of routes accepted per VRF and per routing neighbor should be configured where possible; and the Generalized TTL Security Mechanism (GTSM; RFC 3682 [10]) should be used for all supported protocols.

In summary, it is not possible to intrude from one VPN into other VPNs, or the core. However, it is theoretically possible to attack the routing protocol port to execute a DoS attack against the PE router. This in turn might have a negative impact on other VPNs on this PE router. For this reason, PE routers must be extremely well secured, especially on their interfaces to CE routers. ACLs must be configured to limit access only to the port(s) of the routing protocol, and only from the CE router. Further routing protocols' security mechanisms such as MD5 authentication, maximum prefix limits, and Time to Live (TTL) security mechanisms should be used on all PE-CE peerings. With all these security measures, the only possible attack is a DoS attack against the routing protocol itself. BGP has a number of countermeasures such as prefix filtering and damping built into the protocol, to assist with stability. It is also easy to track the source of such a potential DoS attack. Without dynamic routing between CEs and PEs, the security is equivalent to the security of ATM or Frame Relay networks.

3.4. Label Spoofing

Similar to IP spoofing attacks, where an attacker fakes the source IP address of a packet, it is also theoretically possible to spoof the label of an MPLS packet. In the first section, the assumption was made that the core network is trusted. If this assumption cannot be made, IPsec must be run over the MPLS cloud. Thus in this section the emphasis is on whether it is possible to insert packets with spoofed labels into the MPLS network from the outside, i.e., from a VPN (CE router) or from the Internet.

The interface between a CE router and its peering PE router is an IP interface, i.e., without labels. The CE router is unaware of the MPLS core, and thinks it is sending IP packets to another router. The "intelligence" is done in the PE device, where, based on the configuration, the label is chosen and pre-pended to the packet. This is the case for all PE routers, towards CE routers as well as the upstream service provider. All interfaces into the MPLS cloud only require IP packets, without labels.

For security reasons, a PE router should never accept a packet with a label from a CE router. RFC 3031 [9] specifies: "Therefore, when a labeled packet is received with an invalid incoming label, it MUST be discarded, UNLESS it is determined by some means (not within the scope of the current document) that forwarding it unlabeled cannot cause any harm." Since accepting labels on the CE interface would potentially allow passing packets to other VPNs it is not permitted by the RFC.

Thus, it is impossible for an outside attacker to send labeled packets into the BGP/MPLS IP VPN core.

There remains the possibility to spoof the IP address of a packet being sent to the MPLS core. Since there is strict address separation within the PE router, and each VPN has its own VRF, this can only harm the VPN the spoofed packet originated from; that is, a VPN customer can attack only himself. MPLS doesn't add any security risk here.

The Inter-AS and Carrier's Carrier cases are special cases, since on the interfaces between providers typically packets with labels are exchanged. See section 4 for an analysis of these architectures.

3.5. Comparison with ATM/FR VPNs

ATM and FR VPN services enjoy a very high reputation in terms of security. Although ATM and FR VPNs can be provided in a secure manner, it has been reported that these technologies also can have security vulnerabilities [14]. In ATM/FR as in any other networking technology, the security depends on the configuration of the network being secure, and errors can also lead to security problems.

4. Security of Advanced BGP/MPLS IP VPN Architectures

The BGP/MPLS IP VPN architecture described in RFC 2547 [7] defines the PE-CE interface as the only external interface seen from the service provider network. In this case, the PE treats the CE as untrusted and only accepts IP packets from the CE. The IP address range is treated as belonging to the VPN of the CE, so the PE maintains full control over VPN separation.

RFC 4364 [1] has subsequently defined a more complex architecture, with more open interfaces. These interfaces allow the exchange of label information and labeled packets to and from devices outside the control of the service provider. This section discusses the security implications of this advanced architecture.

4.1. Carriers' Carrier

In the Carriers' Carrier (CsC) architecture, the CE is linked to a VRF on the PE. The CE may send labeled packets to the PE. The label has been previously assigned by the PE to the CE, and represents the label switched path (LSP) from this CE to the remote CE via the carrier's network.

RFC 4364 [1] specifies for this case: "When the PE receives a labeled packet from a CE, it must verify that the top label is one that was distributed to that CE." This ensures that the CE can only use labels that the PE correctly associates with the corresponding VPN. Packets with incorrect labels will be discarded, and thus label spoofing is impossible.

The use of label maps on the PE leaves the control of the label information entirely with the PE, so that this has no impact on the security of the solution.

The packet underneath the top label will -- as in standard RFC 2547 [7] networks -- remain local to the customer carrier's VPN and not be inspected in the carriers' carrier core. Potential spoofing of subsequent labels or IP addresses remains local to the carrier's VPN; it has no implication on the carriers' carrier core nor on other VPNs in that core. This is specifically stated in section 6 of RFC 4364 [1].

Note that if the PE and CE are interconnected using a shared layer 2 infrastructure such as a switch, attacks are possible on layer 2, which might enable a third party on the shared layer 2 network to intrude into a VPN on that PE router. RFC 4364 [1] specifies therefore that either all devices on a shared layer 2 network have to be part of the same VPN, or the layer 2 network must be split logically to avoid this issue. This will be discussed in more detail in section 6.

In the CsC architecture, the customer carrier needs to trust the carriers' carrier for correct configuration and operation. The customer of the carrier thus implicitly needs to trust both his carrier and the carriers' carrier.

In summary, a correctly configured carriers' carrier network provides the same level of security as comparable layer 2 networks or traditional RFC 2547 [7] networks.

4.2. Inter-Provider Backbones

RFC 4364 [1] specifies three sub-cases for the inter-provider backbone (Inter-AS) case.

a) VRF-to-VRF connections at the autonomous system border routers (ASBRs).

In this case, each PE sees and treats the other PE as a CE; each will not accept labeled packets, and there is no signaling between the PEs other than inside the VRFs on both sides. Thus, the separation of the VPNs on both sides and the security of those are the same as on a single AS RFC 2547 [7] network. This has already been shown to have the same security properties as traditional layer 2 VPNs.

This solution has potential scalability issues in that the ASBRs need to maintain a VRF per VPN, and all of the VRFs need to hold all routes of the specific VPNs. Thus, an ASBR can run into memory problems affecting all VPNs if one single VRF contains too many routes. Thus, the service providers needs to ensure that the ASBRs are properly dimensioned and apply appropriate security measures such as limiting the number of prefixes per VRF.

The two service providers connecting their VPNs in this way must trust each other. Since the VPNs are separated on different (sub-)interfaces, all signaling between ASBRs remains within a given VPN. This means that dynamic cross-VPN security breaches are impossible. It is conceivable that a service provider connects a specific VPN to the wrong interface, thus interconnecting two VPNs that should not be connected. This must be controlled operationally.

b) EBGP redistribution of labeled VPN-IPv4 routes from AS to neighboring AS.

In this case, ASBRs on both sides hold full routing information for all shared VPNs on both sides. This is not held in separate VRFs, but in the BGP database. (This is typically limited to the Inter-AS VPNs through filtering.) The separation inside the PE is maintained through the use of VPN-IPv4 addresses. The control plane between the ASBRs uses Multi-Protocol BGP (MP-BGP, RFC 2858 [8]). It exchanges VPN routes as VPN-IPv4 addresses, the ASBR addresses as BGP next-hop IPv4 addresses, and labels to be used in the data plane.

The data plane is separated through the use of a single label, representing a VRF or a subset thereof. RFC 4364 [1] states that an ASBR should only accept packets with a label that it has assigned to this router. This prevents the insertion of packets with unknown labels, but it is possible for a service provider to use any label

that the ASBR of the other provider has passed on. This allows one provider to insert packets into any VPN of the other provider for which it has a label.

This solution also needs to consider the security on layer 2 at the interconnection. The RFC states that this type of interconnection should only be implemented on private interconnection points. See section 6 for more details.

RFC 4364 [1] states that a trust relationship between the two connecting ASes must exist for this model to work securely. Effectively, all ASes interconnected in this way form a single zone of trust. The VPN customer needs to trust all the service providers involved in the provisioning of his VPN on this architecture.

c) PEs exchange labeled VPN-IPv4 routes, ASBRs only exchange loopbacks of PEs with labels.

In this solution, there are effectively two control connections between ASes. The route reflectors (RRs) exchange the VPN-IPv4 routes via multihop eBGP. The ASBRs only exchange the labeled addresses of those PE routers that hold VPN routes that are shared between those ASes. This maintains scalability for the ASBRs, since they do not need to know the VPN-IPv4 routes.

In this solution, the top label specifies an LSP to an egress PE router, and the second label specifies a VPN connected to this egress PE. The security of the ASBR connection has the same constraints as in solution b): An ASBR should only accept packets with top labels that it has assigned to the other router, thus verifying that the packet is addressed to a valid PE router. Any label, which was assigned to the other ASBR, will be accepted. It is impossible for an ASBR to distinguish between different egress PEs or between different VPNs on those PEs. A malicious service provider of one AS could introduce packets into any VPN on a PE of the other AS; it only needs a valid LSP on its ASBR and PEs to the corresponding PE on the other AS. The VPN label can be statistically guessed from the theoretical label space, which allows unidirectional traffic into a VPN.

This means that such an ASBR-ASBR connection can only be made with a trusted party over a private interface, as described in b).

In addition, this solution exchanges labeled VPN-IPv4 addresses between route reflectors (RRs) via MP-eBGP. The control plane itself can be protected via routing authentication (RFC 2385 [6]), which ensures that the routing information has been originated by the expected RR and has not been modified in transit. The received VPN

information cannot be verified, as in the previous case. Thus, a service provider can introduce bogus routes for any shared VPN. The ASes need to trust each other to configure their respective networks correctly. All ASes involved in this design form one trusted zone. The customer needs to trust all service providers involved.

The difference between case b) and case c) is that in b) the ASBRs act as iBGP next-hops for their AS; thus, each SP needs to know of the other SP's core only the addresses of the ASBRs. In case c), the SPs exchange the loopback addresses of their PE routers; thus, each SP reveals information to the other about its PE routers, and these routers must be accessible from the other AS. As stated above, accessibility does not necessarily mean insecurity, and networks should never rely on "security through obscurity". This should not be an issue if the PE routers are appropriately secured. However, there is an increasing perception that network devices should generally not be accessible.

In addition, there are scalability considerations for case c). A number of BGP peerings have to be made for the overall network including all ASes linked this way. SPs on both sides need to work together in defining a scalable architecture, probably with route reflectors.

In summary, all of these Inter-AS solutions logically merge several provider networks. For all cases of Inter-AS configuration, all ASes form a single zone of trust and service providers need to trust each other. For the VPN customer, the security of the overall solution is equal to the security of traditional RFC 2547 [7] networks, but the customer needs to trust all service providers involved in the provisioning of this Inter-AS solution.

5. What BGP/MPLS IP VPNs Do Not Provide

5.1. Protection against Misconfigurations of the Core and Attacks 'within' the Core

The security mechanisms discussed here assume correct configuration of the network elements of the core network (PE and P routers). Deliberate or inadvertent misconfiguration may result in severe security leaks.

Note that this paragraph specifically refers to the core network, i.e., the PE and P elements. Misconfigurations of any of the customer side elements such as the CE router are covered by the security mechanisms above. This means that a potential attacker must have access to either PE or P routers to gain advantage from misconfigurations. If an attacker has access to core elements, or is

able to insert into the core additional equipment, he will be able to attack both the core network and the connected VPNs. Thus, the following is important:

- To avoid the risk of misconfigurations, it is important that the equipment is easy to configure and that SP staff have the appropriate training and experience when configuring the network. Proper tools are required to configure the core network.
- To minimise the risk of "internal" attacks, the core network must be properly secured. This includes network element security, management security, physical security of the service provider infrastructure, access control to service provider installations, and other standard SP security mechanisms.

BGP/MPLS IP VPNs can only provide a secure service if the core network is provided in a secure fashion. This document assumes this to be the case.

There are various approaches to control the security of a core if the VPN customer cannot or does not want to trust the service provider. IPsec from customer-controlled devices is one of them. The document "CE-to-CE Member Verification for Layer 3 VPNs" [13] proposes a CE-based authentication scheme using tokens, aimed at detecting misconfigurations in the MPLS core. The document "MPLS VPN Import/Export Verification" [12] proposes a similar scheme based on using the MD5 routing authentication. Both schemes aim to detect and prevent misconfigurations in the core.

5.2. Data Encryption, Integrity, and Origin Authentication

BGP/MPLS IP VPNs themselves do not provide encryption, integrity, or authentication service. If these are required, IPsec should be used over the MPLS infrastructure. The same applies to ATM and Frame Relay: IPsec can provide these missing services.

5.3. Customer Network Security

BGP/MPLS IP VPNs can be secured so that they are comparable with other VPN services. However, the security of the core network is only one factor for the overall security of a customer’s network. Threats in today’s networks do not come only from an "outside" connection, but also from the "inside" and from other entry points (modems, for example). To reach a good security level for a customer network in a BGP/MPLS infrastructure, MPLS security is necessary but not sufficient. The same applies to other VPN technologies like ATM or Frame Relay. See also RFC 2196 [5] for more information on how to secure a network.

6. Layer 2 Security Considerations

In most cases of Inter-AS or Carrier's Carrier solutions, a network will be interconnected to other networks via a point-to-point private connection. This connection cannot be interfered with by third parties. It is important to understand that the use of any shared-medium layer 2 technology for such interconnections, such as Ethernet switches, may carry additional security risks.

There are two types of risks with layer 2 infrastructure:

a) Attacks against layer 2 protocols or mechanisms

Risks in a layer 2 environment include many different forms of Address Resolution Protocol (ARP) attacks, VLAN trunking attacks, or Content Addressable Memory (CAM) overflow attacks. For example, ARP spoofing allows an attacker to redirect traffic between two routers through his device, gaining access to all packets between those two routers.

These attacks can be prevented by appropriate security measures, but often these security concerns are overlooked. It is of the utmost importance that if a shared medium (such as a switch) is used in the above scenarios, that all available layer 2 security mechanisms are used to prevent layer 2 based attacks.

b) Traffic insertion attacks

Where many routers share a common layer 2 network (for example, at an Internet exchange point), it is possible for a third party to introduce packets into a network. This has been abused in the past on traditional exchange points when some service providers have defaulted to another provider on this exchange point. In effect, they are sending all their traffic into the other SP's network even though the control plane (routing) might not allow that.

For this reason, routers on exchange points (or other shared layer 2 connections) should only accept non-labeled IP packets into the global routing table. Any labeled packet must be discarded. This maintains the security of connected networks.

Some of the above designs require the exchange of labeled packets. This would make it possible for a third party to introduce labeled packets, which if correctly crafted might be associated with certain VPNs on an BGP/MPLS IP VPN network, effectively introducing false packets into a VPN.

The current recommendation is therefore to discard labeled packets on generic shared-medium layer 2 networks such as Internet exchange points (IXPs). Where labeled packets need to be exchanged, it is strongly recommended to use private connections.

7. Summary and Conclusions

BGP/MPLS IP VPNs provide full address and traffic separation as in traditional layer-2 VPN services. It hides addressing structures of the core and other VPNs, and it is not possible to intrude into other VPNs abusing the BGP/MPLS mechanisms. It is also impossible to intrude into the MPLS core if this is properly secured. However, there is a significant difference between BGP/MPLS-based IP VPNs and, for example, FR- or ATM-based VPNs: The control structure of the core is layer 3 in the case of MPLS. This caused significant skepticism in the industry towards MPLS, since this might open the architecture to DoS attacks from other VPNs or the Internet (if connected).

As shown in this document, it is possible to secure a BGP/MPLS IP VPN infrastructure to the same level of security as a comparable ATM or FR service. It is also possible to offer Internet connectivity to MPLS VPNs in a secure manner, and to interconnect different VPNs via firewalls. Although ATM and FR services have a strong reputation with regard to security, it has been shown that also in these networks security problems can exist [14].

As far as attacks from within the MPLS core are concerned, all VPN classes (BGP/MPLS, FR, ATM) have the same problem: If an attacker can install a sniffer, he can read information in all VPNs, and if the attacker has access to the core devices, he can execute a large number of attacks, from packet spoofing to introducing new peer routers. There are a number of precautionary measures outlined above that a service provider can use to tighten security of the core, but the security of the BGP/MPLS IP VPN architecture depends on the security of the service provider. If the service provider is not trusted, the only way to fully secure a VPN against attacks from the "inside" of the VPN service is to run IPsec on top, from the CE devices or beyond.

This document discussed many aspects of BGP/MPLS IP VPN security. It has to be noted that the overall security of this architecture depends on all components and is determined by the security of the weakest part of the solution. For example, a perfectly secured static BGP/MPLS IP VPN network with secured Internet access and secure management is still open to many attacks if there is a weak remote access solution in place.

8. Security Considerations

The entire document is discussing security considerations of the RFC 4364 [1] architecture.

9. Acknowledgements

The author would like to thank everybody who has provided input to this document. Specific thanks go to Yakov Rekhter, for his continued strong support, and Eric Rosen, Loa Andersson, Alexander Renner, Jim Guichard, Monique Morrow, Eric Vyncke, and Steve Simlo, for their extended feedback and support.

10. Normative References

[1] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

11. Informative References

[2] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[3] Baker, F., Atkinson, R., and G. Malkin, "RIP-2 MD5 Authentication", RFC 2082, January 1997.

[4] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.

[5] Fraser, B., "Site Security Handbook", RFC 2196, September 1997.

[6] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

[7] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.

[8] Bates, T., Rekhter, Y., Chandra, R., and D. Katz, "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000.

[9] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[10] Gill, V., Heasley, J., and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)", RFC 3682, February 2004.

[11] Fang, L., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, July 2005.

[12] Behringer, M., Guichard, J., and P. Marques, "MPLS VPN Import/Export Verification", Work in Progress, June 2004.

[13] Bonica, R. and Y. Rekhter, "CE-to-CE Member Verification for Layer 3 VPNs", Work in Progress, September 2003.

[14] DataComm, "Data Communications Report, Vol 15, No 4: Frame Relay and ATM: Are they really secure?", February 2000.

Author's Address

Michael H. Behringer
Cisco Systems Inc
Village d'Entreprises Green Side
400, Avenue Roumanille, Batiment T 3
Biot - Sophia Antipolis 06410
France

EMail: mbehring@cisco.com
URI: http://www.cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgement

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