Provider Provisioned VPN WG
Internet Draft
draft-nagarajan-ppvpn-vrbased-applicability-01.txt
Expiration Date: December 2002

Ananth Nagarajan
Sprint

Junichi Sumimoto
Muneyoshi Suzuki
NTT Corporation

Paul Knight
Nortel Networks

Benson Schliesser
SAVVIS Communications

June 2002

Applicability Statement for Virtual Router-based Layer 3 PPVPN approaches

Status of this Memo

This document is submitted to the IETF's Provider Provisioned Virtual Private Network (ppvpn) working group. Comments should be addressed to WG's mailing list at ppvpn@ppvpn.francetelecom.com. The charter may be found at http://www.ietf.org/html.charters/ppvpn-charter.html

Abstract

This document is an applicability statement for Layer 3 Provider Provisioned VPNs (L3 PPVPNs) that is based on Virtual Router (VR) approaches. This document describes how VR-based approaches meet the key requirements that are outlined in the PPVPN Applicability Statements Guideline document.

1. Summary for Sub-IP area

This document is an output of the design team formed to develop applicability statements for Layer 3 PPVPNs in the PPVPN working group. As such this work fits within the scope of the PPVPN working group. This document discusses the applicability of virtual router-based approaches for Layer 3 PPVPNs.

2. Introduction

The virtual router concept for L3 PPVPNs was first introduced in [COREVPN]. This was generalized in [PPVPNVR]. A number of autodiscovery mechanisms can be used with this approach to L3 PPVPNs, and [COREVPN] represents one such approach using IP multicast. Based on the taxonomy of PPVPNs described in [FRAMEWORK], Virtual Router based approaches are classified as PE-based Layer 3 PPVPNs.

VR-based PPVPNs are used in the following situations:

- The customer wishes to outsource the maintenance and management of inter-site VPN connectivity to the Service Provider (SP).

- The SP desires to provide VPN service without upgrading its core network to support any specific technology (e.g., MPLS), i.e., the SP can provide a Layer 3 VPN service over an existing IP routed or Layer 2 switched core network.

- The customer is not aware of the topology or mechanisms used in the SP core network and is responsible for routing between customer routers, which is independent of the routing used in the SP core. Only the customer-facing sides of the PE devices in the SP network are visible to the customer.

- The customer primarily sends IP traffic across the VPN, but may also send non-IP traffic, provided these traffic types are supported by the tunneling technologies used. It should be noted that the support of Layer 2 VPNs using VR-based mechanisms is outside the scope of this document.

This document describes how Virtual Router based approaches satisfy key requirements and metrics identified in the PPVPN Applicability Statements Guideline document [ASGUIDE]. These requirements are a subset of the requirements listed in the PPVPN Service Requirements document [REQTS]. This document is based on the guidelines specified in [ASGUIDE].

3. SP Provisioning Model

Virtual Routers (VR) have similar properties to physical routers, except that they are instantiated on a single PE device. VPNs are constructed via tunnels connecting VR pairs across the service provider backbone network. Per-VR routing protocol instantiations are run to distribute VPN reachability information. VPN membership information distribution is treated separately, and is achieved via sharing a VPN-ID, for example [RFC2685], between VRs that are members of a specific VPN. This separation of reachability and membership distribution functions is one of the key differences between the VR model and the "piggy-backing" models such as [RFC2547bis]. In [RFC2547bis], both reachability and membership information is distributed via BGP extensions between PE devices, on a per-VPN basis. The detailed VR model is described in [PPVPNVR].

3.1. Auto Discovery

In the VR-based PPVPNS, various auto discovery mechanisms are supported. VPN discovery can be achieved through directory servers, explicit configuration via a management platform, using multicast [COREVPN] or by piggybacking VPN membership and topology information via routing protocols such as BGP [VPN-BGP]. A combination of these mechanisms may also be used on a PE. For example, for some VPNs topology discovery is done only through a management platform. For others, dynamic topology discovery is achieved using existing routing protocol. BGP-based auto-discovery is described in [VPN-BGP], and is used for membership, toplogy and reachability discovery.

It is important to note that, for the VR architecture, the auto-discovery mechanism is only used to automatically exchange control VPN information between VRs. It is not intended for piggybacking VPN private reachability information onto the backbone routing instance.

4. Supported Topology and Traffic Types

VR-based PPVPNs can be constructed using either MPLS tunnels in the core network or IP tunnels (GRE, IP-in-IP, L2TP, IPSec), or Layer 2 connections such as ATM or Frame Relay. The choice of the tunneling mechanism may impact other properties of the VPN itself, including scalability, manageability, QoS, security etc. For example, the use of IPSec tunnels for encryption may impact forwarding performance as a result of sophisticated encryption mechanisms, and therefore impact the number of routes per VPN, the number of VPNs per PE, etc. Tunnels are created on a per-VPN basis. For transport across the network, a number of these tunnels may be aggregated and carried within a PE-PE tunnel. The topology of the VPN is not strictly defined. It may be any arbitrary topology, including full-mesh, and arbitrary partial-mesh.

5. Isolated exchange of routing and data information

By definition of a Virtual Private Network, the details of its addressing, topology, connectivity, and reachability as well as the data that it transports are implicitly considered to be private, and should therefore be isolated from other networks, including others that may be supported with the PPVPN infrastructure. [FRAMEWORK]

5.1. Isolation of routing information (constrained distribution of reachability information)

In a PPVPN, routing is provisioned and managed by the SP, who is responsible for maintaining isolation between networks except as explicitly intended by the VPN owner.

The VR model of PPVPNs provides for isolation by instantiating

multiple Virtual Routers (VR) on a single physical platform to support multiple VPNs. [PPVPNVR] Each VR has its own interfaces, routing tables, forwarding tables, and routing protocol instances. This provides for isolated topology, addressing, and reachability for the VPN.

Addressing and Reachability includes the assignment, discovery, and distribution of source and/or destination information for the PPVPN. The isolation of this information implies that other networks, including other VPNs and the Internet, will have no visibility into the PPVPN except as explicitly configured.

Routing information carried between VRs is carried in the same context/plane as data itself, and is therefore segregated from the underlying backbone infrastructure by the same mechanisms that segregate data between VPNs.

This model supports arbitrary routing architectures, including support for back-door links or other potentially unique routing architecture requirements. The support for arbitrary routing architectures, however, is accompanied by scalability and management issues. These issues are discussed later in this document.

In the VR approach, virtual routers are connected to the CEs through local links, and to each other across the backbone through tunneling services provided by the service provider across the backbone. All data traffic within the VR-based VPN is isolated from non-VPN traffic by these mechanisms.

5.2. Isolation of data

Data for different VPNs in the VR model is segregated through the use of different link-layer connections or tunnels over a common SP backbone. [PPVPNVR] Examples of such tunnels include GRE, L2TP, IPSec, MPLS or Layer 2 connections such as ATM or Frame Relay. It

should be noted that this isolation can be impacted by misconfiguration.

6. Access Control and Authentication

CE-PE authentication has not been specified for VR-based VPNs and is for further study. The customer must provide appropriate mechanisms for CE-PE authentication.

In order for VR-based PPVPNs to support confidentiality, integrity, authentication, and replay attack prevention, mechanisms such as IPsec may be used as tunneling mechanism or used over VPN tunnels. Even with the use of IPsec, the security level offered is dependent on the scope of the IPsec security: encrypting on a CE-to-CE basis (as in CE-based VPNs) will offer a higher level of protection than encrypting on a PE-to-PE basis (as in PE-based VPNs). Policy-based security and access control mechanisms or firewalls may be used between sites in the same VPN. These can be implemented on the PE router, or on the CE.

7. Security

7.1. Protection of user data

As described above, end-to-end (CE-to-CE) IPSec may be used to protect user data. SPs may choose to provide CE-based IPSec as a value added service. If the SP core network is also part of the public Internet, the SP may choose to provide PE-to-PE IPSec as the tunneling mechanism between VRs.

If inter-SP VPNs are to be provided, IPSec tunnels may be used. The impact on QoS and SLAs in this case will have to be studied.

In general, user data is protected via the inherent isolation provided by the inter-VR tunnels. Varying levels of security of user data may be provided based on the type of tunnel that is used.

7.2. SP Security Measures

In general, the SP should ensure that non-VPN traffic does not accidentally or maliciously enter a VPN. As such, the PE and P devices should be protected against intrusion or denial of service attacks. VR routing sessions must be authenticated. If BGP is used as an auto- discovery mechanism between VRs, it should be further authenticated using mechanisms such as TCP MD5. Filtering of data entering the PE should be performed in order to prevent the acceptance of unauthorized packets from customers or other SPs into that PE.

8. Addressing

Virtual routers may provide any or all of the services which are provided by a physical router, including Network Address Translation (NAT), packet filtering, etc. These VR capabilities can simplify the process of joining previously independent site networks, which may have overlapping address spaces. NAT can be used to satisfy intra-VPN non-unique addressing requirements. This facilitates the construction of short-term or ad-hoc VPNS. It should be noted, however, that NAT has accompanying scaling problems, and other mechanisms are needed to ensure proper routing updates, when two sites share the same routing domain.

Non-unique and private customer addresses may be supported by using encapsulation within the tunneling mechanisms employed between VR pairs (e.g., GRE, IP-in-IP etc.). As such, support for private addressing as specified in [RFC1918] allows for non-unique addresses between different VPNs.

9. Interoperability and Interworking

Interoperability and Interworking of VR-based VPNs with other L3 PPVPN mechanisms such as 2547bis is for further study. Since VRs provide all IP router functionalities, various VR-based solutions interwork and interoperate to the extent that IP networks interoperate and interwork.

10. Network Access

10.1. Physical and Link Layer Topology

VR-based mechanisms do not affect the choice of physical and link layer technologies or topologies.

10.2. Temporary Access

Temporary access for a dial-up user to a VR can be provided via PPP and AAA, using a Remote Access Server. Other access mechanisms such as IPSec can also be used. Thus, it is possible provide login and password based access to a VR-based VPN from an authorized user connected to the Internet.

10.3. Access Connectivity

Multi-homing of CEs to multiple VRs (within the same or different PEs) is supported. The PEs (and consequently the VRs) may belong to different SPs. In the case where multihoming of CEs is across different SPs, care should be taken during traffic sharing across the SPs. For example, traffic from a single ingress PE should not be split in this case.

Load sharing based on IGP or other traffic engineering mechanisms used in the SP core are naturally supported by VR-based VPNs.

11. Service Access

11.1. Internet Access

Simultaneous VPN and Internet Access can be supported via various mechanisms. A specific VR may be assigned as a default VR that is connected to the Internet. If a single VR is to be used to carry a customer's VPN as well as Internet traffic, Internet traffic can be distinguished from VPN traffic by associating a default VPN-ID with Internet traffic and pointing it to a default route to the Internet. This default route to the Internet need not be direct, but may instead point to a firewall or other security device which may use different interfaces for VPN access and Internet access.

11.2. Hosting, ASP, other services

All of the above "external" services can be supported by associating a separate address for every service that is not being used within the VPN. If a single server (for example, a web hosting server) is used to provide a particular service to all VPNs, NAT may be used to provide a unique address for clients to access that particular service. NAT can be performed either at the customer site or can be integrated into the PE. The scaling impacts of adding NAT to the PE will have to be considered.

12. SP Routing

VR-based PPVPNs do not impose any additional requirements on the IGP used in the service provider core network. However, the PE must implement the IGP used in the customer VPN. The VR-based VPNs can use the core routing protocols or may use different routing protocols between VRs than the core network.

Fault handling is a specific problem when the timers used for the VR-to-VR routing peering are shorter than the timers used for the routing peering within the service provider(s) network. In this case a single failure within a service provider network may look like a

collection of un-correlated failures in the VPN.

Moreover, since a VR doesn't really "know" what causes the failure, the VR may react to such a failure by re-routing along some other

tunnel, while this other tunnel may be also affected by the same failure. As a result, this would slow down routing convergence within the VPN.

To avoid the problems mentioned above one may consider making the timers used for the VR-to-VR peering longer than the timers used for the routing peering within the service provider network (so that failures within the service provider network would be "invisible" to the VR-VR tunnels). But that has its own set of problems. While this may be possible to accomplish within a single routing domain (one needs to appropriately set the IGP timers within the domain), doing this in a network that includes more than one routing domain may be difficult (as timers include both IGP and BGP timers, and moreover, timers include IGP timers in several routing domains). Another consequence of making the timers used for the VR-to-VR peering over the tunnels longer than the timers used for the routing peering within the service provider network is that it would increase the amount of traffic that will be "black holed" in the case of VR failures.

12.1. Core router awareness of mechanisms used

Since tunnels are established between VR pairs, the core router (P router) does not have any information of the mechanisms used to construct the VPN. If MPLS is the tunneling mechanism that is used between the VRs, the core routers may have to be MPLS enabled in order to leverage the benefits of MPLS tunnels (e.g., traffic engineering). As such, while the core routers are not aware of VPN-specific information, they should support requirements to meet relevant SLAs. (e.g., for guaranteed QoS, the core routers may need to support appropriate QoS mechanisms).

13. Migration impacts

As VR approach makes use of standard routing protocols without any extensions, any CE using the VR approach can access a PE similar to the way it would access another CE router in a private network using leased lines. Key design considerations include:

- The PEs will introduce extra router hops

- If the VR-VR backbone routing protocol differs from the sites, then IGP metric implications should be carefully evaluated. This would be particularly true for multihomed VPN sites.

Also, since the VR approach does not depend on the backbone architecture in terms of routing protocols, a VR-based L3 PPVPN can be offered on a service provider core network without the need for specific core technologies. For example, the core network does not need specific mechanisms like MPLS to be implemented on the P routers. Similarly, if the core network is a Layer 2 network based on ATM or Frame Relay, VR-based VPNs can still be constructed.

It should be noted, however, that core network mechanisms would determine the overall properties and services that may be provided over the VPN. For example, in order to support customer QoS SLAs, the core network should be robustly engineered or should support QoS mechanisms, in addition to SLA marking at the PE.

Thus, while migration impacts in the case of basic VPN functionality using VR are minimal from the customers' or providers' point of view, appropriate core mechanisms may be necessary for certain services.

14. Scalability

PE-based PPVPNs have better scalability than CE-based PPVPNs, because the total number of VPN tunnels that need to be managed are far fewer in the service provider backbone, than between CEs. Addition of a new CE in a CE-based PPVPN would require O(N^2) tunnels to be set up where N represents the total number of CEs. In comparison, addition of a new CE for a specific customer, in the case of a PE-based PPVPN, would simply require an additional connection between the new CE and the PE, because inter-PE tunnels already exist per VPN.

VR is a technology for implementing logical routing instances in a PE device. A PE device may contain more than one VR and a VR supports one VPN. Therefore, scalability of a VR and conventional physical router are basically the same, e.g., if different routing protocols are used for customer and network sides of a VR or physical router, the load is increased compared with the case when the same protocols are used. However, the scalability depends on the theoretical limits on address space, routing protocols, etc., within a single VPN.

The major factor contributing to scalability constraint in the VR approach is the number of VRs which can be supported by a PE. This is because, the number of VRs in a PE device is equal to the number of VPNs which are supported by the PE.

Resources used by a VR instance include memory and processor resources, used to support VPN tunnel mechanisms, routing protocol instances, route tables, interface management, etc. The extent to which these resources are utilized impact scalability.

Much of the resource utilization for a given VPN will be affected by the topology of the VPN. For instance, a VPN with a full-mesh topology will require that VRs have more peers for the VPN tunneling

mechanism, for routing protocol adjacencies, for security protocols, and etc., while a hub-and-spoke model will constrain the resources required for 'spoke' PE routers.

From a VR perspective, scalability also depends on whether the same routing protocols are used between VRs as in the backbone network. If the inter-VR routing protocols are different from the backbone IGP, the scaling and management impacts for configuring routing protocols on a per-VR basis may be significant. For example, it may be necessary to maintain OSPF databases for the entire customer VPN topology, as opposed to maintaining information for only directly connected customer sites. Additionally, the customer IGP may need to maintain information about the entire VR topology. Other concerns include routing loop avoidance, route dedistribution, etc. Thus, while the VR model allows separate routing protocols between customers and between VRs than the backbone IGP, this flexibility is accompanied by scalability concerns. Mechanisms such as OSPF areas may be used to circumvent such scaling issues.

15. QoS/SLA

VR-based PPVPNs support any kind of QoS that the core network and the tunneling mechanism used support.

QoS mechanisms developed for physical routers can be used with VRs, on a per-VR basis. e.g., classification, policing, drop policies, traffic shaping and scheduling/bandwidth reservation. The architecture allows separate quality of service engineering of the VPNs and the backbone.

VR-based VPNs can utilize different quality of service mechanisms. QoS mechanisms developed for physical routers can be used with VRs, on a per-VR basis. e.g. classification, policing, drop policies, traffic shaping and scheduling/bandwidth reservation. The architecture allows separate quality of service engineering of the

VPNs and the backbone. However, the tunneling mechanisms themselves should support relevant QoS mechanisms.

15.1. SLA Monitoring

VR-based VPNs can implement a variety of methods to monitor compliance with Service Level Agreements. Since the links between VRs make use of tunnels across the underlying backbone network, the SLA monitoring capabilities of the backbone network can be used to monitor the performance of the inter-VR links. Performance to SLA requirements within the PEs hosting the VRs is typically monitored via internal processes to ensure compliance from end to end. In addition, either the service provider or the VPN customer can use all existing SLA tracking tools (round trip time measurement, traceroute mapping, etc.) within the VR-based VPN.

16. Management

16.1. SP Management

[Note: This section will be completed in following versions of this draft]

16.2. Customer Management

The SP may choose to manage the customer site (i.e., the CE devices) for added revenue. If the SP uses a centralized customer management system, care should be taken to uniquely identify various CEs belonging to different VPNs, so that CE devices from different VPNs do not reach each other.

The customer may desire to have access to the PE device for monitoring purposes (e.g., ping, traceroute). Providing such access is at the discretion of the SP.

Traffic statistics in order to prove SLAs to customers may be provided on a periodic basis. Other statistics that can show

enhanced SP capabilities, including protection against Denial of Service attacks, failure etc., can be provided to the customer.

17. Security considerations

There are no additional security considerations besides those already addressed in the document.

18. Acknowledgments

The authors of this draft would like to acknowledge the suggestions and comments received from the entire Layer 3 Applicability Statement Design Team formed in the PPVPN working group. Besides the authors, the members of the design team include Marco Carugi, Eric Rosen, Jeremy De Clercq, Luyuan Fang, Dave McDysan, Cliff Wang, Olivier Paridaens, Tom Nadeau, Yakov Rekhter and Rick Wilder.

19. REFERENCES

[PPVPNVR] Ould-Brahim, H., et al., "Network based IP VPN Architecture using Virtual Routers", work in progress.

[ASGUIDE] Sumimoto, J., et al., "Guidelines of Applicability Statements for PPVPNs," work in progress.

[FRAMEWORK] R. Callon, et al., "A Framework for Layer 3 Provider Provisioned Virtual Private Networks," work in progress.

[REQTS] McDysan, D., et al., "Service requirements for Layer 3 Provider Provisioned Virtual Private Networks", work in progress.

[RFC2764] Gleeson, B., et al., "A Framework for IP Based Virtual Private Networks", RFC 2764, February 2000.

[RFC1918] Rekhter, Y. et al., "Address Allocation for Private Internets," RFC 1918, February 1996.

[RFC2685] Fox B., et al, "Virtual Private Networks Identifier", RFC

2685, September 1999.

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", RFC 2026, October 1996.

[COREVPN] Muthukrishnan, K., Malis, A., "Core MPLS IP VPN Architecture", work in progress.

[RFC2547bis] Rosen E., et al, "BGP/MPLS VPNs", work in progress.

[VPN-BGP] Ould-Brahim, H., et al, "Using BGP as an Auto-Discovery Mechanism for Network-based VPNs", work in progress.

20. Authors' Addresses

Ananth Nagarajan
Sprint
6220 Sprint Parkway
Overland Park, KS 66251
E-mail: ananth.nagarajan@mail.sprint.com

Muneyoshi Suzuki
Junichi Sumimoto
NTT Information Sharing Platform Labs.
3-9-11, Midori-cho,
Musashino-shi, Tokyo 180-8585, Japan
Email: suzuki.muneyoshi@lab.ntt.co.jp
Email: sumimoto.junichi@lab.ntt.co.jp

Paul Knight
Nortel Networks
600 Technology Park Drive
Billerica, MA 01821
E-mail: paknight@nortelnetworks.com

Benson Schliesser
SAVVIS Communications
717 Office Parkway
St. Louis, MO 63141
Phone: +1-314-468-7036
Email: bensons@savvis.net

21. Full Copyright Statement