

Internet Engineering Task Force
INTERNET-DRAFT
Expires April 5, 1999
<[draft-muthukrishnan-corevpn-arch-00.txt](#)>

Karthik Muthukrishnan
Andrew Malis
Ascend Communications
October 5 1998

Core IP VPN Architecture

1. Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``[id-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow Directories on [ftp.is.co.za](#) (Africa), [nic.nordu.net](#) (Europe), [munnari.oz.au](#) (Pacific Rim), [ftp.ietf.org](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

2. Acronyms

LSP	Label Switched Path
PNA	Private Network Administrator
SP	Service Provider
SPED	Service Provider Edge Device
SPNA	SP Network Administrator
VMA	VPN Multicast Address
VPNID	VPN Identifier
VR	Virtual Router
VRC	Virtual Router Console

3. Abstract

This draft presents an approach for building core VPN services in the service provider backbone, as described in [[Heinanen](#)]. This approach does not depend on MPLS running in the backbone but will benefit from it. The central vision is for the service provider to provide a virtual router service to their customers. Ease of configuration, dynamic neighbor discovery, scaling and the use of existing routing protocols as they exist today without any modifications are the keystones of this architecture.

INTERNET-DRAFT

Core VPNs

October 2, 1998

[4.](#) Introduction

This draft describes how VPN services in the backbone of the SP's network could be built. The predominant emphasis is on providing a virtual router service and every effort has been made to make the virtual router as equivalent to a physical router as possible. The aspects of a router that a virtual router needs to emulate are configuration of any combination of routing protocols, monitoring of the network and troubleshooting. Providing a logically independent routing domain to every VPN enhances the SP's ability to offer a fully flexible virtual router service that can fully serve the SP's customer without requiring physical per-VPN routers.

The approach presented here meets most of the requirements set forth in [[Heinanen](#)] but differs significantly in that we have strived to not require or depend on any modifications of any existing routing protocols. Neighbor discovery is aided by the use of an emulated LAN and is achieved by the use of ARP. This draft has made a concerted effort to draw the line between the SP and the PNA: layer 1 and layer 2 services belong and are managed by the SP while layer 3 services belong to and are managed by the PNA. By the provisioning of fully logically independent routing domains the PNA has been given the flexibility to use private and unregistered addresses. Data security is not an issue given the use of private LSPs and the use of VPNID encapsulation when forwarding on shared LSPs.

The approach espoused in this draft differs from that described in [[Jamieson](#)] in the following ways: No routing protocol is modified or used to aid in the neighbor discovery mechanism. No VPN subnet from the SP's address space is required to be allocated. No PNL to PNL direct peering is used. It is not required for the CPE gear to be also MPLS compliant, thus allowing existing enterprise routers to not have to be upgraded.

[5.](#) Objectives

1. Easy, scalable configuration of VPN endpoints in the service provider network.
2. No use of globally unique SP IP resources such as IP subnets.

3. Dynamic discovery of VRs (Virtual Routers) in the SP's cloud.
4. Virtual Routers fully configured and monitored by network administrator of the VPN.

5. Forwarding quality fully configurable; at the lowest end best effort internet LSP.
6. Differentiated services on a VPN by VPN basis based on private LSPs.
7. Security of internet routers extended to Virtual Routers.
8. Specific routing protocols not mandated between Virtual Routers.
9. No special extensions to existing routing protocols such as BGP, RIP, OSPF, ISIS etc.

[6.](#) Requirements

The service provider network must run some form of multicast routing to all nodes that will have VPN connections and to nodes that have to forward Virtual Router discovery multicast datagrams.

[7.](#) Architectural Outline

1. Every VPN is assigned a 16 bit VPNID which is unique within the SP's network. The choice of 16 bits for VPNID (rather than 32 bits) allows 65k VPNs to be built in a SP's network and simultaneously keeps this ID small enough to be transmitted in encapsulation headers.
2. The VPN service is offered in the form of a Virtual Router service. These VRs reside in the SPED and are as such confined to the edge of the SP's cloud. The VRs will use the SP's network for data and control packet forwarding but are otherwise invisible outside the SPEDs.
3. The "size" of the VR contracted to the VPN in a given SPED is

the quantity of IP resources such as routing interfaces, route filters, routing entries etc. This is entirely under the control of the SP and provides the fine granularity required to empower the SP to offer virtually infinite grades of VR service on a per-SPED level. [Example: one SPED may be the aggregating point (say headquarters of the corporation) for a given VPN and a number of other SPEDs may be access points (branch offices). In this case, the SPED connected to the headquarters may be contracted to provide a large VR while the SPEDs connected to the branch offices may house small, perhaps stub VRs].

4. One of the indicators of the size of the VPN is the number of

SPEDs in the SP's network that have connections to CPE routers. As globally unique IP resources do not have to be dedicated/assigned to VPNs, the number of SPEDs is not limited by any artificial configuration limits.

5. Layer 1 and Layer 2 entities belong to and are managed by the SP. To be specific, physical switches/routers, physical links, logical layer 2 connections (such as DLCI in Frame Relay and VPI/VCI in ATM) and LSPs (and their assignment to specific VPNs) are under the control of the SP. In the context of VPNs, it is the SP's responsibility to contract and assign layer 2 entities to specific VPNs.

6. Layer 3 entities belong to and are managed by the PNA. Examples of these entities include IP interfaces, choice of dynamic routing protocols or static routes, and routing interfaces. This provides a virtual routing domain to the PNA and empowers the PNA to design the network to achieve intranet, extranet and traffic engineering goals.

7. The PNA can manage and monitor the VPN using the methods that would have been used if physical routers rather than VRs were used. Therefore, management may be performed using SNMP or other similar methods or directly at the console, the VR console (VRC). Monitoring and troubleshooting may be performed using SNMP or similar, but may also include the use of standard tools such as ping, traceroute etc. Again, the VRC may be used for these purposes just like any physical router.

8. The VRs in the SPEDs form the VPN in the SP's network. Together, they represent a virtual routing domain. They dynamically discover each other by utilizing an emulated LAN resident in the SP's network. Each VPN in the SP's network is assigned a multicast address. Subscription to this multicast address allows a VR to discover and be discovered by other VRs.

9. Data forwarding may be done in one of several ways: hop-by-hop using some form of tunneling SPED-to-SPED, a public LSP with best-effort characteristics or a traffic engineered private LSP with differentiated characteristics. The choice of which LSP is configurable by the SP. The default is the public LSP with best-effort characteristics. The hop-by-hop mechanism is available to route packets during periods of LSP establishment and failure.

8. Scalable Configuration

A VPN is expected to have 100s to 1000s of endpoints within the SP cloud. Configuration should therefore scale at most linearly with

the number of end points. Anything worse will make this task too daunting for the service provider. To this end, all that the service provider needs to allocate/assign are physical/logical links from the private network to the service provider edge device.

9. Dynamic Neighbor Discovery

The VRs in a given VPN reside in a number of SPEDs in the network. The problem is that these VRs need to be connected together. One way to do this is to require the configuration of tunnels between these VRs in a fully meshed fashion. This is obviously not scalable from a configuration and network resource standpoint. Hence the need arises to allow these VRs to dynamically discover each other. Neighbor discovery is facilitated as follows: each VPN is given a limited emulated LAN. This emulated LAN is used in several ways:

1. Address resolution uses this LAN to resolve next-hop (private) IP addresses associated with the other VRs.
2. Routing protocols such as RIP and OSPF use this limited

emulated LAN for neighbor discovery and to send routing updates.

The per-VPN LAN is emulated using an IP multicast address. In the interest of conserving public address space and because this multicast address needs to be visible only in the SP network space, we would use an address from the Organizationally scoped multicast addresses (239.192/14) as described in [Meyer]. Each VPN is allocated an address from this range. To completely eliminate configuration in this regard, this address could be computed given the VPNID.

10. Virtual Router Configuration

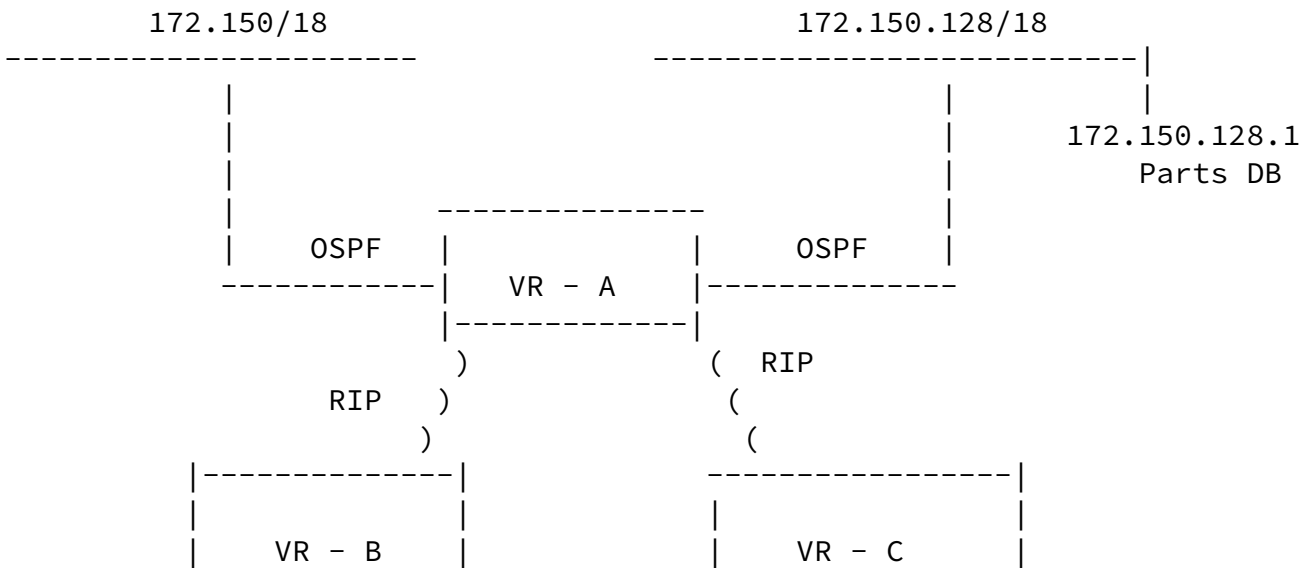




Figure 1

Each Virtual Router is configurable by the PNA as though it were a private physical router. The resources that this Virtual Router may consume is of course limited by the bounds set by the SP on a SPED by SPED basis. Each VPN has a number of physical connections (to CPE routers) and a number of logical connections (to the emulated LAN). Each of these connections is IP capable and can be configured to utilize any combination of the standard routing protocols and routing policies to achieve specific corporate network goals.

To illustrate, in Figure 1, there are 3 VRs on 3 SPEDs. VR-C and VR-B have a physical connection each to CPE equipment while VR-A has 2 physical connections. Each of the VRs has a fully IP capable logical connection to the emulated LAN. VR-A has the (physical) connections to the headquarters of the company and runs OSPF over those connections. It can therefore route packets to 172.150/18 and 172.150.128/18. VR-B runs RIP in the branch office (over the physical connection) and uses RIP (over the logical connection) to export 172.150.64/18 to VR-A. VR-A advertises a default route to VR-B over the logical connection. VR-C is the extranet connection for vendors to use to connect to the parts database at 172.150.128.1. Hence VR-C advertises a default route to VR-A over the logical connection. VR-A exports only 175.150.128.1 to VR-C. This keeps the rest of the corporate network from being subjected to a security problem.

The network administrator will configure the following:

1. OSPF connections to the 172.150/18 and 172.150.128/18 network in VR-A.
2. RIP connections to VR-B and VR-C on VR-A.
3. Route policies on VR-A to advertise only the default route to VR-B.

4. Route policies on VR-A to advertise only 172.159.128.1 to VR-C.
5. RIP on VR-B to VR-A.
6. RIP on VR-C to advertise a default route to VR-A.

11. Forwarding

As mentioned in the architectural outline, data forwarding may be done in one of four ways. The actual method in all but the first outlined here is configurable. At the high end the private LSP is preferred for data forwarding and at the other end hop-by-hop forwarding is used. The order of forwarding preference is therefore: optionally configured private LSP, best effort public LSP and lastly, hop-by-hop.

11.1 Private LSP

This LSP is optionally configured on a per-VPN basis. This LSP is usually associated with non-zero bandwidth reservation and/or a specific differentiated service or QOS class. If this LSP is available it is used for user data and for VPN private control data forwarding.

11.2 Best Effort Public LSP

VPN data packets are forwarded using this LSP if a private LSP with specified bandwidth and/or QOS characteristics is either not configured or not presently available. The LSP that is used is that destined for the egress router in VPN 0. The VPNID in the shim header is used to de-multiplex data packets from various VPNs at the egress router.

11.3 Hop-by-hop

This method of forwarding is used when no LSP is currently available to carry the traffic. This could happen when the LSP is going through a down transient. To confine the VPN routing tables to the edges of the SP network, the VPNID and the egress SPED's ID need to be carried

all the way. An approach is to tunnel the packet to the egress SPED

with the IP protocol set to IPVPN (protocol number to be allocated by IANA) and with a label pushed to represent the VPNID [TBD].

12. Differentiated Services

The configuration of private LSPs for VPNs allows the SP to offer differentiated services to paying customers. These private LSPs could be associated with any available QOS class. Multiple private LSPs with different QOS classes could be configured in a VPN with flow profiles used to sort the packets among the LSPs. This feature together with the ability to size the virtual routers allows the SP to offer truly differentiated services to the VPN customer.

13. Virtual Router Security Considerations

13.1 Data Security

This allows the SP to assure the VPN customer that data packets in one VPN never has the opportunity wander into another. From a routing standpoint, this is achieved by maintaining separate instances of routing protocols and routing tables for each virtual router. From a data forwarding standpoint, the use of VPN encapsulation headers (in the case of shared LSPs or hop-by-hop forwarding) or the use of private LSPs guarantees data privacy.

13.2 Configuration Security

Virtual routers appear as real routers to the PNA. This means that they may be configured by the PNA to achieve connectivity between offices of a corporation. Obviously, the SP has to guarantee that the PNA and the PNA's designees are the only ones to have access to the VRs on the SPEDs the private network has connections to. Since the virtual router console is functionally equivalent to a physical router, all of the authentication methods available on a physical console such as password, RADIUS, etc. are available to the PNA. By allowing only authenticated PNAs to access the VR console, the SP guarantees that the VPN is in full control of its destiny.

14. Physical Network Security

When a PNA logs in to a SPED to configure or monitor the VPN, the PNA is logged into the VR for the VPN. The PNA has layer 3 configuration and monitoring privileges for the VR. Specifically the PNA has no configuration privileges for the physical network. This provides the guarantee to the SP that a VPN administrator will not be able to inadvertently or otherwise adversely affect the SP's network.

15. Virtual Router Monitoring

All of the router monitoring features available on a physical router is available on the virtual router. This includes utilities such as "ping" and "traceroute". In addition, the ability to display private routing tables, link state databases, etc. are available.

16. Acknowledgements

Thanks to Sridhar Komandur and Peter Fetterolf of Ascend Communications for their helpful review and feedback.

17. References

[Callon] Callon R., et al, "A framework for Multiprotocol Label Switching, [draft-ietf-mpls-framework-02.txt](#)".

[Rosen] Rosen E., et al, "Multiprotocol Label Switching Architecture", [draft-ietf-mpls-arch-02.txt](#).

[Heinanen] Heinanen J., et al, "MPLS Mappings of generic VPN mechanisms", [draft-heinanen-generic-vpn-mpls-00.txt](#).

[Jamieson] Jamieson D., et al, "MPLS VPN Architecture", [draft-jamieson-mpls-vpn-00.txt](#).

[Meyer] Meyer D., "Administratively Scoped IP Multicast". [RFC 2365](#).

18. Authors' addresses

Karthik Muthukrishnan
Ascend Communications
1 Robbins Road
Phone: (978) 952-1368
Westford, MA 01886
Email: karthikm@ascend.com

Andrew Malis
Ascend Communications
1 Robbins Road
Westford, MA 01886
Phone: (978)-952-7414
Email: malis@ascend.com

Muthukrishnan, Malis

Expires March, 1999

[Page 9]