

Internet Engineering Task Force
INTERNET-DRAFT

Expires January 2002

Andrew Malis
Vivace Networks, Inc.

Jing Ming Xiao
China UNICOM

Karthik Muthukrishnan
Chandrasekar Kathirvelu
Tom Walsh
Lucent Technologies

Fred Ammann
COMMCARE telecommunications

Junichi Sumimoto
NTT Information Sharing Platform Labs

July 2001

A Core MPLS IP VPN Architecture

<[draft-ietf-ppvpn-rfc2917bis-00.txt](#)>

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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."

This draft is not a product of any working group and was written and presented to the IETF well before the formation of any working group related to Core VPNs.

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract This memo presents an approach for building core Virtual

INTERNET-DRAFT

Core VPNs

July 2001

Private Network (VPN) services in a service provider's MPLS backbone. This approach uses Multiprotocol Label Switching (MPLS) running in the backbone to provide premium services in addition to best effort services. The central vision is for the service provider to provide a virtual router service to their customers. The keystones of this architecture are ease of configuration, user security, network security, dynamic neighbor discovery, scaling and the use of existing routing protocols as they exist today without any modifications.

[1.0.](#) Acronyms

| | |
|-------|------------------------------------|
| ARP | Address Resolution Protocol |
| CE | Customer Edge router |
| LSP | Label Switched Path |
| PNA | Private Network Administrator |
| SLA | Service Level Agreement |
| SP | Service Provider |
| PE | Service Provider Edge Device |
| SPNA | SP Network Administrator |
| VMA | VPN Multicast Address |
| VPNID | VPN Identifier |
| VR | Virtual Router |
| VRC | Virtual Router Console |
| P | Provider Device |
| DSCP | DiffServ Code Point |
| OUI | Organizationally Unique Identifier |

[2.0.](#) Introduction

This draft describes an approach for building IP VPN services out of the backbone of the SP's network. Broadly speaking, two possible approaches present themselves: the piggyback model and the virtual router approach. The piggyback model is based on overloading some semantic(s) of existing routing protocols to carry reachability information. In this document, we focus on the virtual router service.

The approach presented here does not depend on any modifications of any existing routing protocols. This draft makes a concerted effort to draw the line between the SP and the PNA: the SP owns and manages layer 1 and layer 2 services while layer 3 services belong to and are manageable 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. Due to the use of private LSPs and the use of VPNID encapsulation using label stacks over shared LSPs, data security is not an issue.

The approach espoused in this draft differs from that described in [[Rosen1](#)] in that no specific routing protocol has been overloaded to carry VPN routes. [[Rosen1](#)] specifies a way to modify BGP to carry VPN unicast routes across the SP's backbone. To carry multicast routes, further architectural work will be necessary.

[3.0](#). Virtual Routers

A virtual router is a collection of threads, either static or dynamic, in a routing device, that provides routing and forwarding services much like physical routers. A virtual router need not be a separate operating system process (although it could be); it simply has to provide the illusion that a dedicated router is available to satisfy the needs of the network(s) to which it is connected. A virtual router, like its physical counterpart, is an element in a IP domain. The other routers in this domain could be physical or virtual routers themselves. Given that the virtual router connects to a specific (logically discrete) routing domain and that a physical router can support multiple virtual routers, it follows that a physical router supports multiple (logically discrete) routing domains.

From the user (VPN customer) standpoint, it is imperative that the virtual router be as equivalent to a physical router as possible. In other words, with very minor and very few exceptions, the virtual router should appear for all purposes (configuration, management, monitoring and troubleshooting) like a dedicated physical router. The main motivation behind this requirement is to avoid upgrading or re-configuring the large installed base of routers and to avoid retraining of network administrators.

The aspects of a router that a virtual router needs to emulate are:

- Configuration of any combination of routing protocols

- Transport of Unicast and Multicast IP traffic with differentiated or with absolute Qos.
- Monitoring of the network
- Troubleshooting.

Every VPN has a logically independent routing domain. This 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. This means that the SP's "hardware" investments, namely routers and links between them, can be re-used

by multiple customers.

[4.0.](#) Objectives

4.1. Easy, scalable configuration of VPN endpoints in the service provider network. At most, one piece of configuration should be necessary when a CE is added.

4.2. No use of SP resources that are globally unique and hard to get such as IP addresses and subnets.

4.3. Dynamic discovery of VRs (Virtual Routers) in the SP's cloud. This is an optional, but extremely valuable "keep it simple" goal.

4.4. Virtual Routers should be fully configurable and monitorable by the VPN network administrator. This provides the PNA with the flexibility to either configure the VPN themselves or outsource configuration tasks to the SP.

4.5. Quality of data forwarding should be configurable on a VPN-by-VPN basis. This should translate to continuous (but perhaps discrete) grades of service. Some examples include best effort, dedicated bandwidth, QOS, and policy based forwarding services.

4.6. Differentiated services should be configurable on a VPN-by-VPN basis, perhaps based on LSPs set up for exclusive use for forwarding data traffic in the VPN.

4.7. Security of internet routers extended to virtual routers.

This means that the virtual router's data forwarding and routing functions should be as secure as a dedicated, private physical router. There should be no unintended leak of information (user data and reachability information) from one routing domain to another.

4.8. Specific routing protocols should not be mandated between virtual routers. This is critical to ensuring the VPN customer can setup the network and policies as the customer sees fit. For example, some protocols are strong in filtering, while others are strong in traffic engineering. The VPN customer might want to exploit both to achieve "best of breed" network quality.

4.9. No special extensions to existing routing protocols such as BGP, RIP, OSPF, ISIS etc. This is critical to allowing the future addition of other services such as NHRP and multicast. In addition, as advances and addenda are made to existing protocols (such as traffic engineering extensions to ISIS and OSPF), they can be easily incorporated into the VPN implementation.

[5.0.](#) Architectural Outline

5.1. Every VPN is assigned a VPNID which is unique within the SP's network.i.e., local to the SP's network.

By definition there can be different VPNIDs for the same OUI in different VPN clouds.

These different VPN Clouds are separated by one or more IP/MPLS links.

This identifier unambiguously identifies the VPN with which a packet or connection is associated. The VPNID of zero is reserved; it is associated with and represents the public internet.

5.2. The VPN service is offered in the form of a Virtual Router service. These VRs reside in the PE 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 PEs.

5.3. The "size" of the VR contracted to the VPN in a given PE is

expressed by 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 that the SP requires to offer virtually infinite grades of VR service on a per-PE level. [Example: one PE may be the aggregating point (say headquarters of the corporation) for a given VPN and a number of other PE may be access points (branch offices). In this case, the PE connected to the headquarters may be contracted to provide a large VR while the PE connected to the branch offices may house small, perhaps stub VRs]. This provision also allows the SP to design the network with an end goal of distributing the load among the routers in the network.

5.4. As per the traffic requirement we can incrementally add the PEs.

5.5. One indicator of the VPN size is the number of PE in the SP's network that have connections to CPE routers in that VPN. In this respect, a VPN with many sites that need to be connected is a "large" VPN whereas one with a few sites is a "small" VPN. Also, it is conceivable that a VPN grows or shrinks in size over time. VPNs may even merge due to corporate mergers, acquisitions and partnering agreements. These changes are easy to accommodate in this architecture, as globally unique IP resources do not have to be dedicated or assigned to VPNs. The number of PE is not limited by any artificial configuration limits.

5.6. The SP owns and manages Layer 1 and Layer 2 entities. To be specific, the SP controls physical switches or 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). In the context of VPNs, it is the SP's responsibility to contract and assign layer 2 entities to specific VPNs.

5.7. Layer 3 entities belong to and are manageable by the PNA. Examples of these entities include IP interfaces, choice of dynamic routing protocols or static routes, and routing interfaces. Note that although Layer 3 configuration logically falls under the PNA's area of responsibility, it is not necessary for the PNA to execute it. It is quite viable for the PNA to outsource the IP administration of the virtual routers to the Service Provider. Regardless of who assumes responsibility for

configuration and monitoring, this approach provides a full routing domain view to the PNA and empowers the PNA to design the network to achieve intranet, extranet and traffic engineering goals.

5.8. The VPNs can be managed as if physical routers rather than VRs were deployed. Therefore, management may be performed using SNMP or other similar methods or directly at the VR console (VRC).

5.9. Industry-standard troubleshooting tools such as 'ping,' 'traceroute,' etc., are available in a routing domain exclusively of dedicated physical routers. Therefore, monitoring and troubleshooting may be performed using SNMP or similar methods, but may also include the use of these standard tools. Again, the VRC may be used for these purposes just like any physical router.

5.10. Since the VRC is visible to the user, router specific security checks need to be put in place to make sure the VPN user is allowed access to Layer 3 resources in that VPN only and is disallowed from accessing physical resources in the router. Most routers achieve this through the use of database views.

5.11. The VRC is available to the SP as well. If configuration and monitoring has been outsourced to the SP, the SP may use the VRC to accomplish these tasks as if it were the PNA.

5.12. The VRs in the PE form the VPN in the SP's network. Together, they represent a virtual routing domain. They dynamically discover each other by utilizing variety of technics.

5.13. If SPs network supports multicast then routing protocols which uses multicast like OSPF, RIPv2 can be easily supported in VPN.

5.14. User data forwarding may be done in one of several ways:

5.14.1. An LSP with best-effort characteristics that all VPNS can use.

5.14.2. An LSP dedicated to a VPN and traffic engineered by the VPN customer.

5.14.3. A private LSP with differentiated characteristics.

5.14.4. Policy based forwarding on a dedicated L2 Virtual Circuit

The choice of the preferred method is negotiable between the SP and the VPN customer, perhaps constituting part of the SLA between them. This allows the SP to offer different grades of service to different VPN customers.

Of course, hop-by-hop forwarding is also available to forward routing packets and to forward user data packets during periods of LSP establishment and failure.

5.15. This approach does not mandate that separate operating system tasks for each of the routing protocols be run for each VR that the PE houses. Specific implementations may be tailored to the particular PE in use. Maintaining separate routing databases and forwarding tables, one per VR, is one way to get the highest performance for a given PE.

[6.0.](#) Scaleable Configuration

A typical VPN is expected to have 100s to 1000s of endpoints within the SP cloud. Therefore, configuration should scale (at most) linearly with the number of end points. To be specific, the administrator should have to add a couple of configuration items when a new customer site joins the set of VRs constituting a specific VPN. Anything worse will make this task too daunting for the service provider. In this architecture, all that the service provider needs to allocate and configure is the ingress/egress physical link (e.g. Frame Relay DLCI or ATM VPI/VCI) and the virtual connection between the VR and the Service Provider Core.

[7.0.](#) VPN IP Domain Configuration


```

#####
#                               #
#   ROUTER 'A'   #
#                               #
#####
#                               #
#                               #
#                               #
#                               #
#####
#                               #
#   ROUTER 'B'   #   #   ROUTER 'C'   #
#                               #
#                               #
#####
152.0.0.2           153.0.0.3

```

Figure 1 'Physical Routing Domain'

The physical domain in the SP's network is shown in Figure.1. In this network, physical routers A, B and C are connected together. Each of the routers has a 'public' IP address assigned to it. These addresses uniquely identify each of the routers in the SP's network.

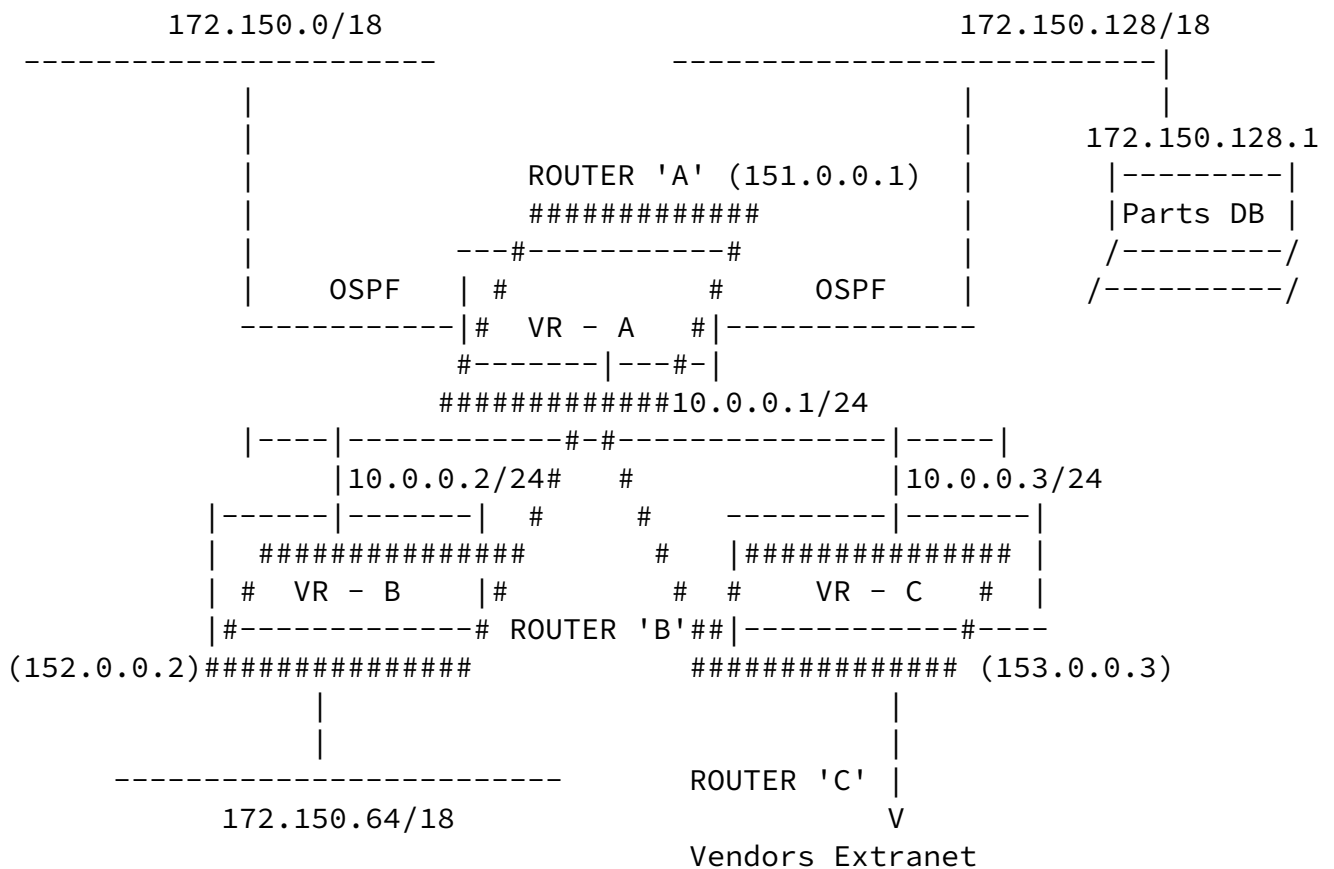


Figure 2 'Virtual Routing Domain'

Each Virtual Router is configurable by the PNA as though it were a private physical router. Of course, the SP limits the resources that this Virtual Router may consume on a PE-by-PE basis. Each VPN has a number of physical connections (to CE routers) and a number of logical connections (to the emulated LAN). The emulated LAN (shown with IP addresses 10.0.0.x/24) provides the VPN Backbone connecting the VRs in Figure 2. Each connection is IP-capable and can be configured to utilize any combination of the standard corporate network protocols and routing policies to achieve specific corporate network goals.

To illustrate, in Figure 2, 3 VRs reside on 3 PE in VPN 1. Router 'A' houses VR-A, router 'B' houses VR-B and router 'C' houses VR-C. VR-C and VR-B have a physical connection to CE 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. Therefore, it can route packets to 172.150.0/18 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. Vendors use VR-C as the extranet connection

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 a security problem.

The network administrator will configure the following:

7.1. OSPF connections to the 172.150.0/18 and 172.150.128/18 network in VR-A.

7.2. RIP connections to VR-B and VR-C on VR-A.

7.3. Route policies on VR-A to advertise only the default route to VR-B.

7.4. Route policies on VR-A to advertise only 172.150.128.1 to VR-C.

7.5. RIP on VR-B to VR-A.

7.6. RIP on VR-C to advertise a default route to VR-A.

[8.0.](#) Forwarding

As mentioned in the architectural outline, data forwarding may be done in one of several ways. In all techniques except the Hop-by-Hop technique for forwarding routing/control packets, the actual method is configurable. At the high end, policy based forwarding for quick service and at the other end best effort forwarding using public LSP is used. The order of forwarding preference is as follows:

- Policy based forwarding.
- Optionally configured private LSP.
- Best-effort public LSP.

[8.1](#) 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.

[8.2](#) 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 used is the one destined for the egress router in VPN 0.(i.e., the reserved VPNID designating the public Internet). The VPNID in the shim header is used to de-multiplex data packets from various VPNs at the egress router.

[9.0](#). Differentiated Services

Configuring private LSPs for VPNs allows the SP to offer differentiated services to paying customers. These private LSPs could be associated with any available L2 QOS class or Diff-Serv codepoints. In a VPN, multiple private LSPs with different service classes could be configured with flow profiles for sorting 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.

[10.0](#). Security Considerations

[10.1](#) Routing Security

The use of standard routing protocols such as OSPF and BGP in their unmodified form means that all the encryption and security methods (such as MD5 authentication of neighbors) are fully available in VRs. Making sure that routes are not accidentally leaked from one VPN to another is an implementation issue. One way to achieve this is to maintain separate routing and forwarding databases.

[10.2](#) Data Security

This allows the SP to assure the VPN customer that data packets in

one VPN never have the opportunity to wander into another. From a routing standpoint, this could be achieved by maintaining separate routing databases for each virtual router. From a data forwarding standpoint, the use of label stacks in the case of shared LSPs [[Rosen2](#)] [[Callon](#)] or the use of private LSPs guarantees data privacy. Packet filters may also be configured to help ease the problem.

[10.3](#) Configuration Security

Virtual routers appear as physical 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 who have access to the VRs on the PE 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

Muthukrishnan et al.

Expires January 2002

[Page 11]

INTERNET-DRAFT

Core VPNs

July 2001

physical console such as password, RADIUS, etc. are available to the PNA.

[10.4](#) Physical Network Security

When a PNA logs in to a PE to configure or monitor the VPN, the PNA is logged into the VR for the VPN. The PNA has only 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.

[11.0](#). Virtual Router Monitoring

All of the router monitoring features available on a physical router are 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.

[12.0](#). Performance Considerations

For the purposes of discussing performance and scaling issues, today's routers can be split into two planes: the routing (control) plane and the forwarding plane.

In looking at the routing plane, most modern-day routing protocols use some form of optimized calculation methodologies to calculate the shortest path(s) to end stations. For instance, OSPF and ISIS use the Dijkstra algorithm while BGP uses the "Decision Process". These algorithms are based on parsing the routing database and computing the best paths to end stations. The performance characteristics of any of these algorithms is based on either topological characteristics (ISIS and OSPF) or the number of ASs in the path to the destinations (BGP). But it is important to note that the overhead in setting up and beginning these calculations is very little for most any modern day router. This is because, although we refer to routing calculation input as "databases", these are memory resident data structures.

Therefore, the following conclusions can be drawn:

1. Beginning a routing calculation for a routing domain is nothing more than setting up some registers to point to the right database objects.
2. Based on 1, the performance of a given algorithm is not significantly worsened by the overhead required to set it up.

3. Based on 2, it follows that, when a number of routing calculations for a number of virtual routers has to be performed by a physical router, the complexity of the resulting routing calculation is nothing more than the sum of the complexities of the routing calculations of the individual virtual routers.

4. Based on 3, it follows that whether an overlay model is used or a virtual routing model is employed, the performance characteristics of a router are dependent purely on its hardware capabilities and the choice of data structures and algorithms.

To illustrate, let's say a physical router houses N VPNs, all running some routing protocol say RP. Let's also suppose that the average performance of RP's routing calculation algorithm is $f(X,Y)$ where x and y are parameters that determine performance of the algorithm for that routing protocol. As an example, for Dijkstra algorithm users such as OSPF, X could be the number of nodes in the area while Y could be the number of links. The performance of an arbitrary VPN n

is $f(X_n, Y_n)$. The performance of the (physical) router is the sum of $f(X_i, Y_i)$ for all values of i in $0 \leq i \leq N$. This conclusion is independent of the chosen VPN approach (virtual router or overlay model).

In the usual case, the forwarding plane has two inputs: the forwarding table and the packet header. The main performance parameter is the lookup algorithm. The very best performance one can get for a IP routing table lookup is by organizing the table as some form of a tree and use binary search methods to do the actual lookup. The performance of this algorithm is $O(\log n)$.


Hence, as long as the virtual routers' routing tables are distinct from each other, the lookup cost is constant for finding the routing table and $O(\log n)$ to find the entry. This is no worse or different from any router and no different from a router that employs overlay techniques to deliver VPN services. However, when the overlay router utilizes integration of multiple VPNs' routing tables, the performance is $O(\log m \times n)$ where 'm' is the number of VPNs that the routing table holds routes for.

[13.0.](#) Some Applications

Some typical applications of PPVPN are illustrated here to assist better understanding of the PPVPNs.


13.1. Example 1

World HQ wants to connect Regional HQs and small stationary outlet/storage.




13.2. Example 2:


In the similar model as said above the SP's network may differ. More than one SP will be involved in connecting the corporate user. In that case we need a SLA between SPs for the service.



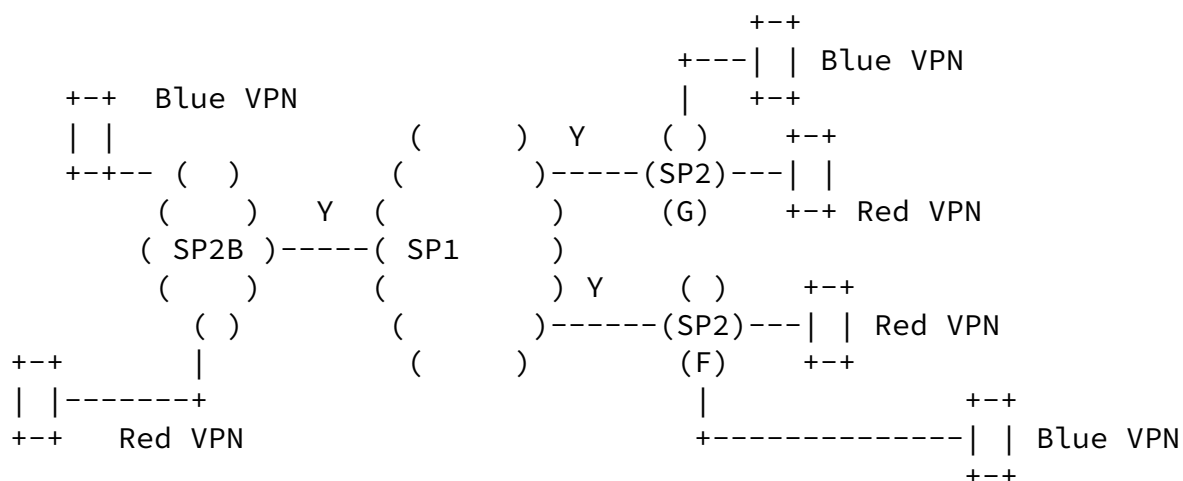
We need a policy between two SP's VPN for exchange of routes. We need VPN id conversion since the VPNs may have different ID in each SP.

13.3. Example 3: When a Mobile user wants to join the VPN, the connection can be made by an LSP tunnel or an IPSec tunnel to the nearest VPN PE.





13.4. Hierarchical VPNs



Y - Yellow VPN. Which SP2 gets from SP1.

SP2 bought Yellow VPN from SP1. SP2 has branch offices as SP2B, SP2G and SP2F in different locations. SP2 also provides VPN services and Red and Blue VPNs are SP2's customers.

SP2B CE for VPN Y towards SP1 will have the instance of Blue Red and Yellow VPNs. Relation between Yellow VPN and Blue/Red VPN is configured by policy. Yellow VPN provides separate transport for Red/Blue VPNs, it does not participate in routing, which keeps all the VPNs separate to each other, but still they can use the upper level VPN for transport. In this example Client of SP2B Red VPN can peer with RedVPN in SP2G. SP2G will peer with SP2B in Yellow VPN. Data from RedVPN in SP2B will use Yellow VPN link from SP2B through SP1 to SP2G. Data from BlueVPN will also use the same path or same LSP. In SP2G it will be demuxed to reach Red/Blue VPNs. Here YellowVPN acts like a LSP tunnel to Red/Blue VPNs.

[14.0.](#) References

[Callon] Callon R., et al, "A framework for Multiprotocol Label Switching", [draft-ietf-mpls-framework-05.txt](#).

[Fox] Fox B., et al, "Virtual Private Networks Identifier", [RFC 2685](#).

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

[Rosen1] Rosen, E., et al. [draft-rosen-rfc2547bis-02.txt](#)

[Rosen2] Rosen E., et al, "Multiprotocol Label Switching Architecture", [draft-ietf-mpls-arch-06.txt](#).

[muthuk] K.Muthukrishnan, A.Malis "A Core MPLS IP VPN Architecture", [RFC 2917](#) September 2000.

[VPN-FRAME] M.Suzuki, J.Sumimoto, "A Framework for Network-based VPNs", <[draft-suzuki-nbvpn-framework-01.txt](#)>, October 2000.

15. Authors' addresses

Karthik Muthukrishnan
Lucent Technologies
1 Robbins Road
Phone: (978) 952-1368
Westford, MA 01886
Email: mkarthik@lucent.com

Andrew Malis
Vivace Networks, Inc.
2730 Orchard Parkway
San Jose, CA 95134
Phone: (408) 383-7223
EMail: Andy.Malis@vivacenetworks.com

Chandrasekar Kathirvelu
Lucent Technologies
1 Robbins Road
Westford, MA 01886
Phone: (978) 952-7116
EMail: ck32@lucent.com

Tom Walsh
Lucent Technologies
10 Lyberty Way
Westford, MA 01886
Phone: (978) 392-2311
EMail: tdwalsh@lucent.com

INTERNET-DRAFT

Core VPNs

July 2001

Fred Ammann
COMMCARE Telecommunications
Turmstrasse 8
CH-8952 Schlieren
Switzerland
Phone: +41 1 738 61 11
Email: fa@commcare.ch

Jing Ming Xiao
China UNICOM
Data & Fixed Communication Department
6/F Office Tower 3
Henderson Center
Beijing, China
Email: unicomnet@bj.cnuninet.net

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

