

Internet Engineering Task Force  
INTERNET-DRAFT

Expires January 2002

Chandrasekar Kathirvelu  
Karthik Muthukrishnan  
Tom Walsh  
Lucent Technologies

Andrew Malis  
Vivace Networks, Inc.

Fred Ammann  
COMMCARE telecommunications

July 2001

## Hierarchical VPN over MPLS Transport

<[draft-ietf-ppvnpn-hiervpn-corevpn-00.txt](#)>

### 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."

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 hierarchical Virtual Private Network (VPN) services. This approach uses Multiprotocol Label Switching (MPLS). The central vision is for the service provider to provide a virtual router service to other SPs without participating in VPNs of those SPs.

### [1.0](#). Acronyms

INTERNET-DRAFT

Hierarchical VPNs

July 2001

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
VPNID	VPN Identifier
VR	Virtual Router
VRL	Virtual Router Link
VRC	Virtual Router Console
P	Provider Device

## [2.](#) Introduction

This draft describes an approach for building Hierarchical IP VPN services out of the backbone of the SP's network. We use the VR model to describe the relationship among the VPNs, and MPLS Label stacking to explain how the data is transported in the hierarchical VPNs. An application of this technique enables the aggregation of many regional or local service Provider VPN networks across a Hierarchical VPN tunneling architecture.

The approach presented here does not require modification of any existing routing protocols.

## [3.](#) Hierarchical Relationship between VPNs

A simplified example that shows a hierarchical relationship between Virtual Routed VPNs is shown in Figure 1. NOTE: Hierarchies can be extended to more than two levels.

Hierarchical levels are designated numerically with the highest level designated as 0. Lower hierarchical levels are designated as Level 1, 2, etc. Higher level VPNs transport lower level VPNs. So:

- LEVEL 0 represents the highest hierarchical level. A Level 0 VPN transports lower level VPNs but is not itself transported by any other VPN;
- LEVEL 1 represents a VPN that is transported by a LEVEL 0 VPN

but is not itself transported across any lower or equal level VPN.

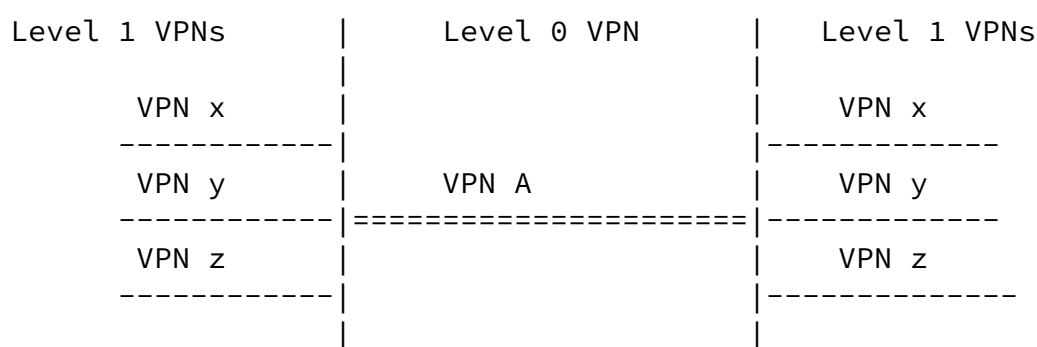


Figure 1. Hierarchical VPN Levels.

By assigning the VPNs depicted in this figure to different hierarchical levels, a hierarchical relationship between the VPNs is created. For example, the highest hierarchical level is designated as "Level 0". In this example, VPN A is a level 0 VPN. Similarly, VPNs' X, Y and Z are part of the next lowest hierarchical level, designated "Level 1". Data within a Level 1 VPN is transported transparently across the Level 0 VPN.

A possible realization of a Hierarchical VPN (similar to that depicted in Figure 1) can now be described using the VR model. This realization does not assume a single Service Provider only is involved. Specifically, in the examples which follow, SP1 and SP2 do not have to be the same Service Provider. MPLS Label stacking techniques are used to create the hierarchical levels and explain how the data is transported.

Figure 2. shows an example of a Hierarchical VPN involving two Service Providers. This example assumes that SP1 provides an international backbone network that is utilized by SP2 to connect its geographically isolated regional (or local) networks. In this example, SP2 is providing two customer VPNs, X and Y. A two level Hierarchical VPN is created to allow VPN X and VPN Y (i.e., level 1 VPNs in this hierarchy) to be transported (at level 0) across VPN A.

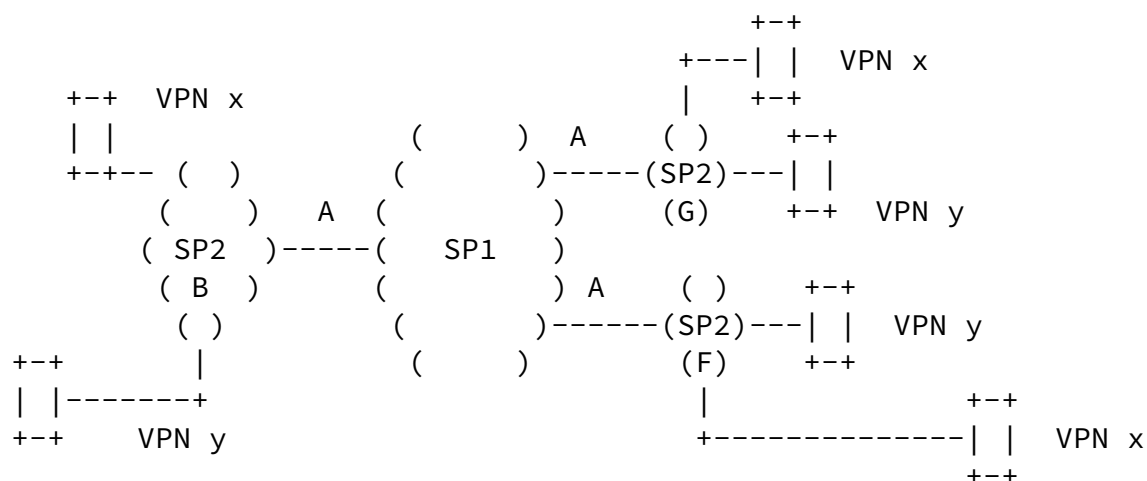




Figure 3 expands the diagram to show the relationship between SP2 and SP1. From this Figure 3, we can see that SP2 provides both end customer VPNs (i.e., VPN x and VPN y) and SP2 must also know about the backbone (i.e., VPN A) that it uses for transporting the hierarchy. On the other hand, SP1 needs to be concerned with just the Level 0 VPN A.





Kathirvelu, et al. Expires January 2002 [Page 4]

VPNs can use any label distribution protocol. The only restriction is, within a specific VPN, the same protocol should be used in all its PE devices, so that they can interwork. This is restricted by the nature of

the distribution protocol not by the VPNs.

Referring to Figure 2, SP1 provides the Level 0 VPN service (called VPN A) to SP2(B/G/F). The label distribution operates independently in each level of the VPN Hierarchy. Labels are distributed for the Level 0 VPN separately from the labels that are distributed for the Level 1 VPN. The following text describes the label distribution for each level of the hierarchical VPN.

#### Level 0 (VPN A) Label Distribution:

The PEs of SP1 share the VPN A routing information between each other. In other words, the reachability information of SP2 edge routers is exchanged. LSP tunnels are set up in VPN A between the edge routers of SP2. For example, an LSP tunnel from SP2 (edge router B) is created to SP2 (edge router G).

#### Level 1 (VPN X) Label Distribution:

The PEs of SP2 share the VPN x routing information with each other. In other words, the reachability information of the CE routers of VPN x is exchanged. LSP tunnels are set up in VPN X between the CE routers in SP2.

Usage of Penultimate Hop Popping (PHP) requires penultimate and top-most labels to be allocated from the same label space (e.g., in this case the

allocation is from VPN A's label space). This implies in the case of Hierarchical VPNs, that an additional label (i.e., the penultimate label) will be necessary between the IGP label (i.e., top-most label) for the PE and VPN destination label. This is shown in the next section on Forwarding.

In this example, it is indicated that A2 is the label for SP2-CE(G) in SP2-CE(B) and it is shown in the Forwarding section ([Section 6.](#)) how A2 is used. (see Figure 4.). This label is chosen from the VPN A Label space.

Architecturally, Level 1 VPN Y and Y are connected to Level 0 VPN A by a Virtual Router Link. Note that the edge routers of SP2 must have knowledge of all three VPNs (i.e., VPN X, VPN Y, and VPN A). When the VRL is configured for a hierarchical relationship, then the top

level VPN will allocate a label for each VRL, i.e., to each VPNs, from its label space.

## 6. Forwarding

User data from the lower level VPNs (e.g. Level 1 in Figure 4) are forwarded by the LSP tunnels of the upper level VPN (e.g. Level 0 in Figure 4). The label encoding shown in Figure 4. is explained below.

VPN x/y/z Data

```
+-----+ +-----+-----+
|Data | | Lx2|Data|
+-----+ +-----+-----+
```

```
+-----+-----+-----+
| Ax2|Lx2 | Data|
```




Figure 4 Label Encoding

1. Customer data arrives at the VPN X CE router in SP2 (B) and is encapsulated in a MPLS frame.

2. Label Lx2 is pushed on to the Label Stack. Lx2 is the peer VPN x CE label used to forward VPN X data to VPN X CE router in SP2 (G).

3. Next Label Ax2 is pushed on to the Label Stack. Ax2 is the peer VPN X attachment label with VPN A taken from VPN A's label space. This label is used by VPN A to forward data on the SP2 (G) VRL between VPN A and VPN X.

4. Finally Label A2 is pushed on to the Label Stack. This is the peer VPN A label used to forward data from the VPN A SP2 (B) PE router to the VPN A SP2 (G) PE router.

In summary, the complete LSP path therefore to move customer data on VPN X from the SP2 (B) CE to the SP2 (G) CE is as follows: a) Transport data across Level 0 (VPN A) using label A2; b) Transport data across the VRL from Level 0 to Level 1 in SP2 (G) using label Ax2 c) Transport data across Level 1 (VPN x) from SP2 (B) to SP2(G) using label Lx2

## 7. Security Considerations

Security as available in MPLS networks will be available and extended to hierarchical VPNs.

## 8. Intellectual Property Considerations

Lucent technologies may seek patent or other intellectual property protection for some of all of the technologies disclosed in this document. If any standards arising from this document are or become protected by one or more patents assigned to Lucent Technologies. Lucent



reasonable and non-discriminatory terms.

## 9. References

[Callon] Callon R., et al., "A Framework for Multiprotocol Label Switching", work in Progress

[Fox] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", [RFC 2685](#), September 1999.

[Rosen2] Rosen E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", work in progress

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

## 10. Authors' Addresses

Karthik Muthukrishnan  
Lucent Technologies  
1 Robbins Road  
Westford, MA 01886  
Phone: (978) 952-1368  
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

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

## 11. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

