

IETF Internet Draft

Expires: January, 2004

Document: [draft-guichard-pe-ce-addr-03.txt](#)

July, 2003

## Address Allocation for PE-CE links within a Provider Provisioned VPN Network

### 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/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

### Abstract

This document proposes to allocate a block of globally unique IPv4 addresses for the exclusive use of Service Providers that provide [L3VPN] based Services. The Service Provider may use these addresses for the provisioning of IP addresses to the links that connect Customer Edge (CE) routers with Provider Edge (PE) routers (PE-CE links), and/or for the IP addressing of attached CE routers.

The main motivation for this proposal is to simplify the Service Providers' operations in the scenario where they monitor PE-CE links, and/or CE-routers, while at the same time conserving IPv4 address space.

This addressing scheme is not intended for use by VPNs that span more than one Service Provider, unless co-operation of addressing

structure is maintained to ensure uniqueness of addresses between providers. Furthermore, although the main reference within this draft is related to services provided by [L3VPN], the recommendations

Guichard, et. al

1

formulated herein may also apply to other Layer-2 or Layer-3 VPN architectures.

## 1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC-2119](#)].

## 2. Contributing Authors

This document was the collective work of several. The editor, and co-authors listed below, contributed the text and content of this document.

Monique Morrow	Cisco Systems Inc
Jeff Aparc	Cisco Systems Inc
Jean Philippe Vasseur	Cisco Systems Inc
Yakov Rekhter	Juniper Networks Inc
Xavier Vinet	EQUANT
Vincent Parfait	EQUANT
Y. Reina Wang	AT&T Labs
Fang, Luyuan	AT&T Labs
Dr. Thomas Kuehne	Arcor AG & Co
Lars Braeunig	Arcor AG & Co

## 3. Motivation for an Additional IP Address Allocation Scheme

The [L3VPN] architecture provides a very flexible model for the deployment of layer-3 based VPN services. The customer interface to these services is typically via a CE router, and the Service Provider may manage this router, or it may be under the control of the attached customer.

The emergence of VPN services based on the [L3VPN] architecture, and the significant experience gained from the deployment of these services, has prompted the need for an additional IP address allocation scheme. The primary use for this scheme would be the IP address assignment of the links that connect CE routers with PE routers (PE-CE links), and/or the IP address assignment of the CE

routers.

The need for this scheme is driven by the explosion of [L3VPN] based services, each with many thousands of CE end-points, and a continuing trend for the migration of existing VPN customers to this service.

When the Service Provider manages the CE router, it is typical for the Service Provider to monitor this router from a central management location that is within the Service Provider's premises. The management of the CE router is useful for a number of reasons including troubleshooting, statistics collection for SLA reporting,

Guichard et. al

2

configuration and so on. Using such a centralized monitoring policy means that the Service Provider has to address each CE router with a unique IP address so that they are able to identify each CE router without any conflict with other CEs/VPNs.

Even when the Service Provider does not manage the CE router, but just monitors PE-CE links from a central management location, it still requires that the addresses assigned to all such links have to be unique across all the links of all the VPNs provided by the Service Provider.

Regardless of whether the [L3VPN] service is managed or not, there are various approaches currently available to the Service Provider when allocating IP addresses. There are various advantages and disadvantages associated with each of these approaches, each of which is explored in the following sections.

From the list of approaches, the most attractive one is discussed in [section 3.3](#) (where PE-CE links are numbered from the address space of the Service Provider registered block). However, the major drawback of this approach is that it results in consuming potentially a (very) large amount of IPv4 address space that would not be used for the purpose of connecting to the Internet. The proposal described in this document aims at preserving most of the benefits of the model described in [section 3.3](#), while at the same time eliminating its major drawbacks.

### [3.1.](#) Address space from [\[RFC-1918\]](#) for PE-CE links

Pros: There is no requirement for Service Provider registered addresses, which means that there is no need for the Service Provider to obtain these addresses from one of the addressing registries.

Cons: This relies on case-by-case discussions with all customers who use [\[RFC-1918\]](#) addresses to negotiate a target pool of [\[RFC-1918\]](#) addresses for monitoring needs, which is very time-consuming and intensive in terms of addressing management. Also, this approach is not always possible with very large customers, due to the number of [\[RFC-1918\]](#) addresses being used within the customers network.

### 3.2. Address space taken from the customer address block

Pros: There is no requirement for the Service Provider to use registered addresses, and the customer is responsible for the addressing plan. No need for the Service Provider to obtain these addresses from one of the addressing registries.

Cons: Requires co-ordination of addressing plan between the Service Provider and the customer. May be problematic if the customers' addresses are from the [\[RFC-1918\]](#) range and the Service Provider has also used [\[RFC-1918\]](#) address space, as there is the potential for an

address overlap between the Service Provider and the customer address space. Also, if the Service Provider manages CE-PE links, then this option requires the NMS used by the provider to deal with non-unique addresses.

### 3.3. Address space from Service Provider registered block

Pros: Does not require any coordination with customer IP address allocation, as no address conflict is likely. Addresses used for PE-CE links are unique across all the PE-CE links for all the VPNs supported by the Service Provider. Furthermore, each CE router is guaranteed to have a unique address for management purposes.

Cons: The Service Provider has to obtain these addresses from one of the addressing registries. This is a waste of globally unique addresses for private usage. A separate /30 or /31 subnet is required for each PE-CE connection, and/or a /32 for each CE router, which results in a very high number of wasted addresses, especially when there are VPNs with thousands of sites.

In order to save the number of required IP addresses, some specific allocation techniques have been tentatively deployed by Service Providers. However, the various solutions still present some challenges as detailed in the following sections.

### [3.3.1.](#) Unnumbered addressing for PE-CE links

Pros: Conservation of IP addresses. Only 1 IP address is required for each PE-CE link instead of a /30 or /31 subnet.

Cons: Undesirable for Network Management purposes, as the network management stations do not capture the management view of the PE-CE links. This means that a separate network management loopback interface is needed for each CE router. A further disadvantage is the requirement for an additional loopback interface on the PE router for each VRF, which may be taken from the Service Providers registered address block, or from the customer address block, or from the [RFC-1918] address block. In either case, the same set of pros and cons apply, and the use of unnumbered links is not a long-term solution for the conservation of address space due to the additional loopback requirement at the CE routers.

### [3.3.2.](#) Unique address block used for ALL VPNs

Pros: Address space can be saved as the same address block is used on the PE-CE links of ALL VPNs.

Cons: Potential for address conflict when merging one VPN with another as the same addresses may be used on the PE-CE links. Also, if the Service Provider manages CE-PE links, then this option

Guichard et. al

4

requires NMS used by the provider to deal with non-unique addresses. If these addresses are used for the CE routers also then an address conflict may occur.

### [3.3.3.](#) Unique address block used for ALL PE routers

Pros: Address space can be saved as the same address block is used for each PE router in the network. This address block is used to number all the CE-PE links of that PE router, irrespective of whether these links belong to the same or different VPNs.

Cons: Potential for address conflict as two PE-CE links belonging to the same VPN but attached to two different PE routers may have the same /30 or /31 subnet assigned. Also, if the Service Provider manages CE-PE links, then this option requires NMS used by the provider to deal with non-unique addresses. If these addresses are

used for the CE routers also then an address conflict may occur.

#### 4. Proposal

This document defines a /12 IPv4 address block that a Service Provider could use for PE-CE addressing, CE router addressing, and/or for local value-added services. The size of the IPv4 address block has been determined through analysis of existing [L3VPN] deployments and by tracking the continued trend for the migration of existing VPN customers onto this service.

The results of this analysis have shown:

- (1) assigning a /8 address block would provide a long term solution with a lower risk of address conflicts between Service Providers. However, there is no evidence at this time to suggest a need for managing millions of PE-CE links and therefore a /8 is considered too large for the purposes laid out in this proposal.
- (2) At a time when large Service Providers have already connected in excess of 40,000 CE endpoints, allocating a /16 which allows for ~16,000 connections would clearly not be sufficient

Continuing with the aim of saving IPv4 address space, the proposed approach is to request an address block to suit medium term needs, and to issue a request for a new block when and if this becomes a requirement. In this context a /12 appears to be a consensual choice.

#### 5. Operational Considerations

Reserved addresses for [L3VPN] based services facilitate address uniqueness for Service Providers within their own administrative

domain. The uniqueness of addresses is guaranteed even if the Service Provider network consists of multiple Autonomous Systems.

Overlapping of these reserved addresses between Service Providers may cause problems if a VPN client site CE router connects to two different Service Provider PE routers, and the addresses used on the PE-CE links are the same.

Private addressing [[RFC-1918](#)] is still available for use within a Service Provider and customer environment. However, the most important benefit of a dedicated address block for PE-CE connections, CE router management, and local value-added services, is the guarantee against address overlap between Service Provider and customer address spaces, as well as the guarantee that all PE-CE numbered links and CE routers will have addresses that are unique within a Service Provider. This benefit impacts the service cost for preventing address overlap and reduces the complexity in doing so.

For Service Providers who offer managed customer PE-CE connectivity, this proposal facilitates Service Provider NMS operations by guaranteeing unique addressing for the managed service thus minimizing provisioning complexity attributed to administering non-unique address space. This factor is a key benefit to Service Providers who are developing and deploying managed [L3VPN] services. One specific example of the benefit is that the Service Provider only requires a single Management VPN (the Service Provider can import to the management VPN the PE-CE address and CE router address blocks using a route-map), the number of management workstations (or instances of) is greatly reduced as there is no overlap.

## [6.](#) Security Considerations

Security issues are not addressed in this memo.

## [7.](#) IANA Considerations

IANA should allocate a /12 address block for sole use by [L3VPN] Service Providers. The actual address block assignment is TBD.

## [8.](#) References

[L3VPN], Rosen, E. et al., "BGP/MPLS VPNs", [draft-ietf-ppvpn-rfc2547bis-01](#), July, 2002.

[[RFC-1918](#)], Rekhter, Y. et al. "Address Allocation for Private Internets", [RFC 1918](#), February 1996.

## [9.](#) Authors' Address:

Guichard et. al

Jim Guichard  
Cisco Systems, Inc.  
300 Apollo Drive  
Chelmsford, MA, 01824  
Email: jguichar@cisco.com

Monique Jeanne Morrow  
Cisco Systems, Inc.  
Glatt-Com  
CH-8301, Glattzentrum, Switzerland  
Email: mmorrow@cisco.com

Jeff Apcar  
Cisco Systems, Inc  
201 Pacific Highway  
St Leonards, NSW 2068  
Australia  
Email: japcar@cisco.com

JP Vasseur  
Cisco Systems, Inc.  
300 Apollo Drive  
Chelmsford, MA, 01824  
Email: jpv@cisco.com

Yakov Rekhter  
Juniper Networks, Inc.  
1194 N. Mathilda Ave.  
Sunnyvale, CA 94089  
Email: yakov@juniper.net

Xavier VINET  
EQUANT  
9 rue du Chêne Germain - BP 80  
35512 Cesson Sevigne cedex  
FRANCE  
Email: xavier.vinet@equant.com

Vincent Parfait  
EQUANT  
1041 route des Dolines - BP 347  
06906 Sophia Antipolis Cedex  
FRANCE  
Email: Vincent.Parfait@equant.com

Y. Reina Wang  
AT&T Labs  
200 Laurel Ave  
Middletown, NJ 07748  
USA



Email: reinawang@att.com

Luyuan Fang  
AT&T Labs  
200 Laurel Avenue  
Middletown, NJ 07748  
USA  
Email: luyuanfang@att.com

Dr. Thomas Kuehne  
Arcor AG & Co.  
Alfred-Herrhausen-Allee 1  
65760 Eschborn  
GERMANY  
Email: thomas.kuehne@arcor.net

Lars Braeunig  
Arcor AG & Co.  
Alfred-Herrhausen-Allee 1  
65760 Eschborn  
GERMANY  
Email: lars.Braeunig@arcor.net

