

Provider Provisioned VPN  
Internet Draft  
[draft-fedyk-bgpvpon-auto-00.txt](#)  
Expiration Date: September 2001

Don Fedyk  
Hamid Ould-Brahim  
Peter Ashwood-Smith  
Nortel Networks

Yakov Rekhter  
Juniper Networks

Eric C. Rosen  
Cisco Systems

March 2001

## BGP based Auto-Discovery mechanism for Optical VPNs

### Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [[RFC-2026](#)], except that the right to produce derivative works is not granted.

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


Consider a service provider network that offers Optical Virtual

Private Network (OVPN) service. An important goal in the OVPN service is the ability to support what is known as \_single end

provisioning\_, where addition of a new port to a given OVPN would involve configuration/provisioning changes only on the devices connected to that port. Another important goal in the OVPN service is the ability to establish/terminate an optical connection between a pair of (existing) ports within an OVPN without involving configuration/provisioning changes in any of the provider devices. In this document we describe a set of mechanisms that accomplishes these goals.

## 1. Optical VPN Reference Model

Consider a service provider network that consists of devices such as Optical Network Element (ONE) which may be Optical Cross Connects (OXC's). We partition these devices into P (provider) ONEs and PE (provider edge) ONEs. The P ONEs are connected only to the ONEs within the provider's network. The PE ONEs are connected to the ONEs within the provider network, as well as to the devices outside of the provider network. We'll refer to such other devices as Client Devices (CDs). An example of a CD would be a router, or a SONET/SDH cross-connect.






Figure 1 Optical VPN Reference Model

A CD is connected to a PE ONE via one or more ports, where each port may consists of one or more channels or sub-channels (e.g., wavelength or wavelength and timeslot respectively). For purpose of this discussion we assume that all the channels within a given port have shared similar characteristics (e.g., bandwidth, encoding, etc\_), and can be interchanged from the CDs point of view. Channels

on different ports of a CD need not have the same characteristics. There may be more than one port between a given CD PE pair. A CD may be connected to more than one PE ONE (with at least one port per each PE ONE). And, of course, a PE ONE may have more than one CD connected to it.

A pair of CDs could be connected through the service provider network via an optical connection. It is precisely this optical connection that forms the basic unit of service that the service provider network offers. If a port by which a CD is connected to a PE ONE consists of multiple channels (e.g., multiple wavelengths), the CD could establish optical connection to multiple other CDs over this single port.

In the context of this document we'll refer to an Optical VPN (OVPN) as a collection of ports that connect the CDs owned by the same organization to the service provider network. A given service provider network could support multiple OVPNs. Moreover, not all ports on a given PE ONE that connect that PE ONE to CDs must belong to the same OVPN.

An important goal in the OVPN service is the ability to support what is known as \_single end provisioning\_, where addition of a new port to a given OVPN would involve configuration/prov-isioning changes only on the PE ONE that has this port and on the CD that is connected to the PE ONE via this port. Another important goal in the OVPN service is the ability to establish/terminate an optical connection between a pair of (existing) ports within an OVPN without involving configuration/provisioning changes in any of the provider's ONEs. In this document we describe a set of mechanisms that accomplishes these goals.

The actual connectivity among ports within a given OVPN is controlled by the OVPN itself, and is outside the scope of this document. It is up to the CDs to create their own topology. The CDs have control on CD-to-CD connectivity.

This allows creation of the OVPN. The learning of the complete OVPN members is outside of the scope auto-discovery mechanism described in this document.

Since this model involves minimal provisioning changes when changing the connectivity among the ports within a OVPN on the providers network and the OVPNs themselves are controlled by the CDs, the tariff structure may be on a port basis or alternatively tariffs could be triggered of signaling mechanisms.

Finally, it is assumed that CD-to-CD optical connectivity is based on GMPLS [[GMPLS](#)] and typically provides UNI [[OIF-UNI](#)] and/or NNI capability. This signaling is not covered in this document.

## [2.](#) Overview of operations

This document assumes that within a given OVPN each port has an identifier that is unique within that OVPN (but need not be unique across several OVPNs). One way to accomplish this is to assign each port an IP address that is unique within a given OVPN, and use this address as a port identifier. Another way to accomplish this is to assigned each port an interface index that is unique within a given CD, assign each CD an IP address that is unique within a given OVPN, and then use a tuple <interface index, CD IP address> acts as a port identifier.

This document assumes that within a service provider network, each port on a PE ONE has an identifier that is unique within that network. One way to accomplish this would be to assign each port on a PE ONE an interface index, assign each PE ONE an IP address that is unique within the service provider network, and then use a tuple <interface index, PE ONE IP address> as a port identifier within the provider network.

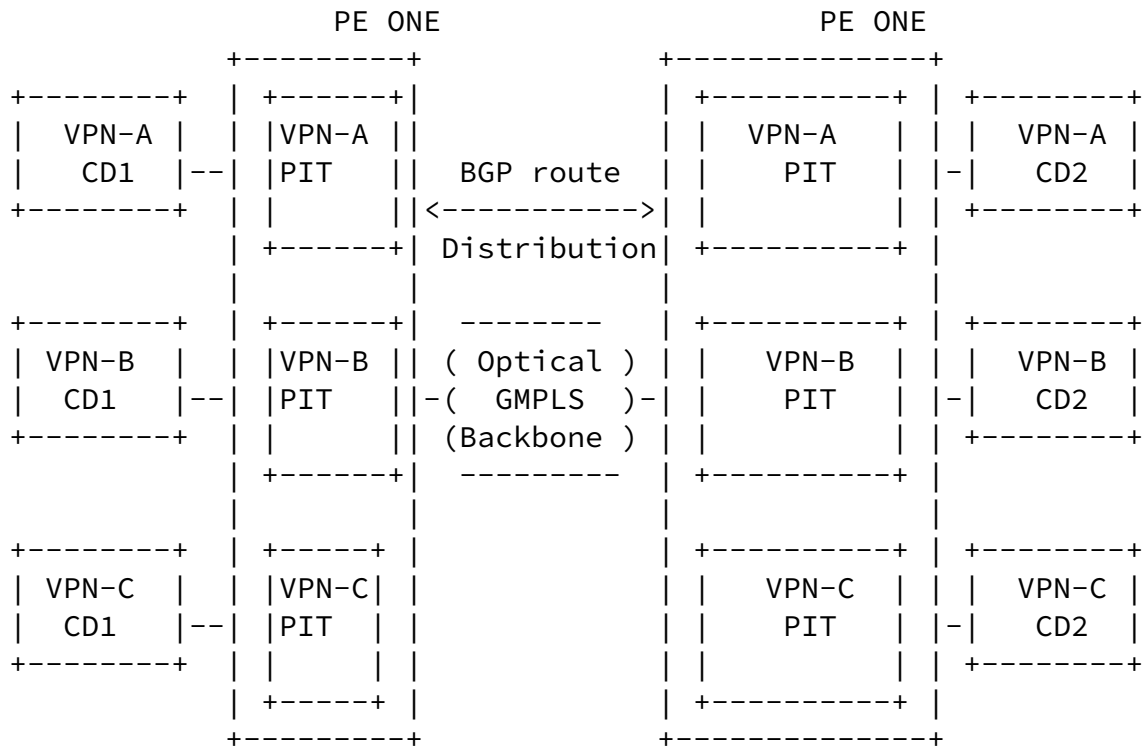


Figure 2 OVPN Components

As a result, each port has an identifier that is unique within a given OVPN, and (another) identifier that is unique within the

service provider network. We'll refer to the former as the customer port identifier (CPI), and to the latter as the provider port identifier (PPI).

Each PE ONE maintains a Port Information Table (PIT) for each OVPN that has at least one port on that PE ONE. A PIT contains a list of <CPI, PPI> tuples for all the ports within its OVPN.

A PIT on a given PE ONE is populated from two sources: the information received from the CDs attached to the ports on that PE ONEs, and the information received from other PE ONEs. We'll refer to the former as the `_local_` information, and to the latter as the `_remote_` information.

The local information is propagated to other PE ONEs by using BGP with multi-protocol extensions. To restrict the flow of this information to only the PITs within a given OVPN, we use BGP route filtering based on the Route Target Extended Community, as follows.

When a service provider adds a new OVPN, the service provider allocates a new BGP Route Target Extended Community that will be used for the purpose of this OVPN. Each PIT on a PE ONE is configured with the Route Target associated with the OVPN of that PIT, and that creates an association between a PIT and an OVPN. When exporting local information into provider's BGP, this information is tagged with the Route Target Community. When importing remote information into a particular PIT, only the information with the Route Target Community equal to the one configured for the PIT could be imported into the PIT.

When a service provider adds a new OVPN port to a particular PE ONE, this port is associated at provisioning time with a PIT on that PE ONE, and this PIT is associated (again at provisioning time) with that OVPN.

Once a port is configured on the PE ONE, the CD that is attached via this port to the PE ONE passes to the PE ONE the CPI information of that port. This document assumes that this is accomplished by using a (subset of) GMPLS signaling. This information, combined with the PPI information available to the PE ONE, enables the PE ONE to create a tuple <CPI, PPI> for such port, and then use this tuple to populate the PIT of the OVPN associated with that port.

A PE ONE uses the information in its PITs to provide CDs connected to that PE ONE with the information about CPIs of other ports within the same OVPN, which we'll refer to as `_target ports_`. This document assumes that this is accomplished by using a (subset of) GMPLS signaling. Once a CD has this information, the CD uses a (subset of) GMPLS signaling to request the provider network to establish an optical connection to a target port. The request originated by the CD contains the CPI of the port on the CD that CD wants to use for the optical connection, and the CPI of the target port. When the PE

ONE attached to the CD that originated the request receives the request, the PE ONE identifies the appropriate PIT, and then uses the information in that PIT to find out the PPI associated with the CPI of the target port carried in the request. The PPI should be

sufficient for the PE ONE to establish an optical connection. Ultimately the request reaches the CD associated with the target CPI (note that the request still carries the CPI of the CD that originated the request). If the CD associated with the target CPI accepts the request, the optical connection is established.

Note that a CD need not establish an optical connection to every target port that CD knows about \_ it is a local to the CD matter to select a subset of target ports to which the CD will try to establish optical connections.

A port, in addition to its CPI and PPI may also have other information associated with it that describes characteristics of the channels within that port, such as encoding supported by the channels, bandwidth of a channel, total unreserved bandwidth within the port, etc\_ This information is used to ensure that ports at each end of an optical connection have compatible characteristics, and that there are sufficient unallocated resources to establish an optical connection. Distribution of this information (including the mechanisms for distributing this information) is identical to the distribution of the CPI information. Distributing changes to this information due to establishing/terminating of optical connections is identical to the distribution of the CPI information, except that thresholds should be used to contain the volume of control traffic caused by such distribution.

It may happen that for a given pair of ports within an OVPN, each of the CDs connected to these ports would concurrently try to establish an optical connection to the other CD. If having a pair of optical connections between a pair of ports is viewed as undesirable, the a way to resolve this is have CD with the lower value of CPI is required to terminate the optical connection originated by the CD. This option could be controlled by configuration on the CD devices.

### [3](#) Encoding

This section specifies encoding of various information defined in this document

#### [3.1](#) Encoding of CPI and channel characteristics in GMPLS Signaling

[TBD]

#### [3.2](#) Encoding of CPI, PPI, and channel characteristics in BGP

[TBD]

### [4](#) Other issues

While the above text assumes that the service provider network consists of ONEs and ports are connected via optical connections, the mechanisms described in this document could be applied in an environment, where the service provider network consists of SONET/SDH cross connects and ports are connected via SONET/SDH sub-channels with each other.

Since the protocol used to populate a PIT with remote information is BGP, since BGP works across multiple routing domains, and since GMPLS signaling isn't restricted to a single routing domain, it follows that the mechanisms described in this document could support an environment that consists of multiple routing domains.

## [5.](#) Security Considerations

[TBD]

## [7.](#) References

- [BGP-COMM] Ramachandra, Tappan, "BGP Extended Communities Attribute", February 2000, work in progress.
- [RFC-2283] Bates, Chandra, Katz, and Rekhter, "Multiprotocol Extensions for BGP4", [RFC2283](#), February 1998.
- [BGP-MPLS] Rekhter Y, Rosen E., "Carrying Label Information in BGP4", January 2001, work in progress.
- [RFC-3031] Rosen, Viswanathan, and Callon, "Multiprotocol Label Switching Architecture", [RFC3031](#), January 2001.
- [RFC-3032] Rosen, Rekhter, Tappan, Farinacci, Fedorkow, Li, and Conta, "MPLS Label Stack Encoding", [RFC3032](#), January 2001.
- [RFC-2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC2026](#), October 1996.
- [RFC-2685] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [VPN-BGP] Ould-Brahim H., Gleeson B., Ashwood-Smith P., Rosen E., Rekhter Y., "Using BGP as an Auto-Discovery Mechanism for Network-based VPNs", work in progress.



[GMPLS] Ashwood-Smith, P., Berger, L. et al., \_Generalized MPLS - Signaling Functional Description\_, November 2000, work in progress.

Fedyk, et al.

Septembre 2000

[Page 7]

---

Internet-Draft

[draft-fedyk-bgpvpon-auto-00.txt](#)

March 2001

[Framework] Rajagopalan, B. et al., \_IP over Optical Networks: A Framework \_, November 2000, work in progress.

[OIF-UNI] Optical Networking Forum., \_User Network Interface (UNI) 1.0 Signaling Specification\_, work in progress.

## 8. Acknowledgments.

The authors would like to thank Osama Aboul-Magd for reviewing the draft and providing comments.

---

## 8. Author's Addresses

Don Fedyk

Nortel Networks  
600 Technology Park  
Billerica, Massachusetts  
01821 U.S.A.  
Phone: +1 (978) 288 3041  
Email: dwfedyk@nortelnetworks.com

Hamid Ould-Brahim  
Nortel Networks  
P O Box 3511 Station C  
Ottawa ON K1Y 4H7 Canada  
Phone: +1 (613) 765 3418  
Email: hbrahim@nortelnetworks.com

Peter Ashwood-Smith  
Nortel Networks  
P.O. Box 3511 Station C,  
Ottawa, ON K1Y 4H7, Canada  
Phone: +1 613 763 4534  
Email: petera@nortelnetworks.com

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

Eric C. Rosen

Cisco Systems, Inc.  
250 Apollo drive  
Chelmsford, MA, 01824  
E-mail: [erosen@cisco.com](mailto:erosen@cisco.com)

Ould-Brahim, et. al

[draft-fedyk-bgpvpn-auto-00.txt](#)

9

March 2001

#### Full Copyright Statement

Copyright (C) The Internet Society (2000). 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.

