

VPN4DC  
Internet Draft  
Intended status: Information Track  
Expires: April 2012

L. Dunbar  
Huawei  
N. So  
Verizon  
October 24, 2011

VPN4DC Problem Statement  
draft-dunbar-vpn4dc-problem-statement-00.txt

## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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>

This Internet-Draft will expire on April 24, 2011.

## Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

Internet-Draft

Problem statement for VPN4DC

2011-10-05

## Abstract

VPN4DC is for extending an existing VPN to connect hosts in public data center(s) which are purchased or leased by VPN client. This draft describes the issues and problems associated with this kind of services.

## 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](#) 0.

## Table of Contents

<a href="#">1. Introduction</a>	<a href="#">2</a>
<a href="#">2. Terminology</a>	<a href="#">3</a>
<a href="#">3. Connecting hosts in public Data Centers with a VPN</a>	<a href="#">4</a>
<a href="#">4. Hosts connectivity for VPN Oriented Data Center Service</a>	<a href="#">4</a>
<a href="#">5. APIs between VPN PEs and Data Center Gateways</a>	<a href="#">6</a>
5.1. What to be communicated between VPN PEs and Data Center Gateways?	<a href="#">6</a>
<a href="#">6. Conclusion and Recommendation</a>	<a href="#">7</a>
<a href="#">7. Manageability Considerations</a>	<a href="#">7</a>
<a href="#">8. Security Considerations</a>	<a href="#">7</a>
<a href="#">9. IANA Considerations</a>	<a href="#">8</a>
<a href="#">10. Acknowledgments</a>	<a href="#">8</a>
<a href="#">11. References</a>	<a href="#">8</a>
Authors' Addresses	<a href="#">8</a>
Intellectual Property Statement	<a href="#">8</a>
Disclaimer of Validity	<a href="#">9</a>

## [1. Introduction](#)

VPN4DC lets VPN clients to connect to their leased or purchased computing resources in public data centers via their own VPNs. In another words, VPN4DC is elastic which allows a VPN to extend its connectivity to (or shrink its connectivity from) resources in one or multiple public data centers.

This kind of service is attractive to VPN customers who often do not want to use public Internet to access purchased or leased resources in public data centers.

Dunbar

Expires April 24, 2012

[Page 2]

---

Internet-Draft

Problem statement for VPN4DC

2011-10-05

For ease of description, this problem statement focuses on making today's L3VPN to seamlessly extend to client's hosts in public data centers.

## [2. Terminology](#)

Aggregation Switch: A Layer 2 switch interconnecting ToR switches

Bridge: IEEE802.1Q compliant device. In this draft, Bridge is used interchangeably with Layer 2 switch.

DC: Data Center

DA: Destination Address

EOR: End of Row switches in data center.

FDB: Filtering Database for Bridge or Layer 2 switch

Public data center: internet data center which offers hosting or computing resources to many different clients

SA: Source Address

ToR: Top of Rack Switch. It is also known as access switch

VDCS: VPN oriented data center services

VM: Virtual Machines

VPN: Virtual Private Network

VPN-o-CS: VPN oriented Computing Service

VPN4DC: Elastic VPN which can extend its VPN connectivity to, or shrink some of its connectivity's from, computing resources

in public data centers

Dunbar

Expires April 24, 2012

[Page 3]

---

Internet-Draft

Problem statement for VPN4DC

2011-10-05

### [3.](#) Connecting hosts in public Data Centers with a VPN

Hosts in data centers are managed by data center operators which are assisted by their own resource and network management systems. VPNs are managed by network service providers, which are most likely different organizations. If enterprises want to offload their applications to public data centers and connect those purchased/leased resources with their own intranet via VPN, the enterprises have to do following steps separately:

1. Contact data center operators to purchase computing resources
2. Get network configuration from the data center operators on how and where their purchased hosts are placed
3. Ask their VPN operators to add attachment circuits on the PEs which are adjacent to the data centers in which their hosts reside.

One big issue associated with this process is that the client VPN's network provider may not have PEs in close proximity to the data centers from which clients' remote hosts are purchased or leased. It can be very difficult to connect hosts in 3rd party data centers to a provider VPN. Under this scenario, the only option is to use tunnels, e.g. IPSec, between 3rd party data centers and VPN provider PEs. This approach will definitely turn away some enterprises that have paid premium for their VPNs.

When VPN service providers do have PEs co-located, or via secure links connected, with the 3rd party data centers from which clients' remote hosts are purchased, proper configuration on the PEs can be very challenging. The VPN operator has to know how their PEs are connected to the 3rd party data center gateway routers, what

protocols are supported on the connections, which network segments have the hosts belonging to a particular VPN client, etc.. To get all those information accurately, a lot of coordination is needed among VPN clients, VPN service providers, and 3rd party data center operators. This process can be very long and error prone.

#### [4.](#) Hosts connectivity for VPN Oriented Data Center Service

The VPN oriented data center service [[VDCS](#)] describes a service model where VPN clients can assume the computing resources purchased from VPN service providers are already linked with their corresponding VPNs.

Dunbar

Expires April 24, 2012

[Page 4]

---

Internet-Draft

Problem statement for VPN4DC

2011-10-05

To enable those services, VPN service providers have PEs co-located with gateways of multiple data centers, which can be their own or 3rd party data centers.

There are two cases of connectivity in VDCS service model:

1. Strictly connectivity: PE is provisioned (by its own operators) in the same fashion as today's L3VPN. When a group of hosts belonging to client X are added to Data Center Y, the PE adjacent to Y is configured properly so that Client X's VPN can exchange routes with Client X owned hosts in Y.
2. VPN attachment circuit configuration being triggered by data center gateway routers: When a group of hosts purchased by Client X are added to a Data Center Y, Y's gateway automatically triggers its adjacent PEs to add attachment circuits for the VPN which belongs to X, and then perform the Case #1 above for VPN-X's PEs to exchange routes with X's hosts in the data center Y.

The Case #1 above will require a long and deep coordination between data center management systems and VPN management systems. The Data Center management systems have to pass out at least the following attributes associated with hosts belonging to Client X:

- Which gateway routers from which Client X's hosts can be accessed
- Which physical interfaces from which Client X's hosts can be

accessed

- Which logical interfaces, e.g. subnets, VLANs, or Data Center internal VPN ID, from which Client X's hosts can be accessed.

Suppose those information can be provided by data center management systems, it is not a simple process for VPN management systems to map Client X's VPN ID with those network attributes from data center, figure out which PEs are actually connected with which Data Center Gateways and which ports on PEs are connected with DC gateway where Client X's hosts can be accessed. This process gets worse when a VPN client's hosts have to be placed in different data center locations for reasons like, regulatory, diverse locations for disaster recovery, etc.

It is well known that network is only a small portion of the overall data center infrastructure. Most likely the data center networks are managed by a smaller separate team than their computing and storage services. Majority, if not all, Data Center's overall management

systems don't have proper mechanism to get and record the information on which network segments are assigned to which clients. Therefore, it can take extremely long process to configure the PE properly for Case #1 above to work.

## [5.](#) APIs between VPN PEs and Data Center Gateways

Different data centers can have different network designs. The network segments on which a VPN client's hosts reside in different data centers can be represented by very different ways. For example, some data centers use VID to differentiate different clients, some data centers could use different subnet addresses, while other data centers could use its internal VPN IDs.

There are simply too many attributes from too many different places to be coordinated in order to configure VPN PEs manageably. There is no protocol between data center server management systems and network management systems, and there is no protocol for VPN management systems to retrieve the needed network attributes from data center management systems. In addition, data center management systems are part of computing industry which is different industry from networking.

If Data Center gateway routers can inform their adjacent VPN PEs on network attributes associated with hosts of a VPN client, the steps to get PEs configured properly will be much shorter and simpler. Even though PEs' VPN attachment circuit may not be configured directly by adjacent DC Gateway routers for security reasons, the network attributes passed from DC can be used by VPN network management system to properly configure the VPN PEs after some level of authentication.

Therefore, it is very beneficial to have some open APIs between VPN PEs and DC gateway to simplify the steps needed for VPN PE configurations.

#### 5.1. What to be communicated between VPN PEs and Data Center Gateways?

The APIs between VPN PEs and their directly connected data center gateway should at least include the request for a group of hosts in a data center to join (or leave) a specific client's VPN.

If a DC Gateway and PE are directly connected via an Ethernet interface, the network segment for a group of hosts belonging to a VPN client in data center can be easily represented by a VLAN identifier.

If a data center gateway is connected with VPN PEs via OC-n interfaces, then the data center Gateway and VPN PEs have to reach agreement on how to differentiate traffic belonging to different VPN clients. If GRE encapsulation is used on the interfaces, the data center gateway and the VPN PE has to reach agreement on which outer IP address represents which VPN client.

It is very common that Data Center network is not aware of provider VPN configuration, and vice versa. Provider VPN Management system has its own mapping between VPN client and its corresponding VPN identifier. Data Center Network management system also has its own mapping from client ID to a specific network segment, such as VLAN ID, ISID, or IP interface, etc.

When Data Center Gateway sends a request to VPN PE for a network segment to be attached to a specific client's VPN, the information it has is the client identifier. Upon receiving the request, the PE has to authenticate the request with its own management system. Upon

authenticating the request, the VPN management system has to map the client identifier, potentially a key, to the proper VPN identifier, and then configure the PE accordingly to get the VPN connected.

As of now, there aren't any available solutions to enable this in-band communication between VPN and data center gateways.

## [6.](#) Conclusion and Recommendation

VPNs represent a major industry for service providers in the enterprise (revenues at billion dollar level). It is very important for VPN Service Providers to expand its VPN services with cloud Data Center services in a secure manner. Automation and end-to-end network integration is very important for VDCS.

Therefore, we recommend IETF to investigate solutions to make it possible.

## [7.](#) Manageability Considerations

TBD

## [8.](#) Security Considerations

TBD.

Dunbar

Expires April 24, 2012

[Page 7]

---

Internet-Draft

Problem statement for VPN4DC

2011-10-05

## [9.](#) IANA Considerations

## [10.](#) Acknowledgments

We want to acknowledge the following people for their valuable inputs to this draft: K.K.Ramakrishnan.

This document was prepared using 2-Word-v2.0.template.dot.

## [11.](#) References



[VPN4DC-Req] So, et al, "Requirements of Layer 3 Virtual Private Network for Data Centers", [draft-so-VPN4DC-00](#), Oct 2011.

[VDCS] So, et al, "Requirement and Framework for VPN-Oriented Data Center Services", [draft-so-vdcs-00](#), June 2011.

#### Authors' Addresses

Linda Dunbar  
Huawei Technologies  
5340 Legacy Drive, Suite 175  
Plano, TX 75024, USA  
Phone: (469) 277 5840  
Email: [ldunbar@huawei.com](mailto:ldunbar@huawei.com)

Ning So  
Verizon Inc.  
2400 N. Glenville Ave.,  
Richardson, TX75082  
[ning.so@verizonbusiness.com](mailto:ning.so@verizonbusiness.com)

#### Intellectual Property Statement

The IETF Trust takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in any IETF Document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights.

Dunbar

Expires April 24, 2012

[Page 8]

---

Internet-Draft

Problem statement for VPN4DC

2011-10-05

Copies of Intellectual Property disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>

The IETF invites any interested party to bring to its attention any

copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement any standard or specification contained in an IETF Document. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

#### Disclaimer of Validity

All IETF Documents and the information contained therein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION THEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

#### Acknowledgment

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