

Network Working Group
Internet Draft
Intended status: Informational
Expires: April 24, 2012

Luyuan Fang
Cisco Systems

October 24, 2011

VPN4DC Problem Statement
[draft-fang-vpn4dc-problem-statement-00.txt](#)

Abstract

Provider Provisioned IP VPNs are commonly used to interconnect multiple locations of a private network, such as an enterprise with multiple offices. Current developments in data center operations create the need to consider additional connectivity and connectivity management problems described in this document.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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 Internet-Draft will expire on April 24, 2012.

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 Simplified BSD License.

Internet Draft

VPN4DC Problem Statement

October 2011

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	VPN4DC: A problem Definition	3
4.	Private network connectivity between data centers	4
5.	Private Networks within a public data center	5
6.	Connectivity between different VPNs	5
7.	Mobile connectivity	5
8.	Security Considerations	6
9.	IANA Considerations	6
10.	Normative References	6
11.	Informative References	6
12.	Author's Address	6

Requirements Language

Although this document is not a protocol specification, 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](#) [RFC 2119].

1. Introduction

Data centers are increasingly being consolidated and outsourced in an effort both to improve the deployment time of applications as well as reduce operational costs. This coincides with an increasing requirement for compute, storage and network connectivity from applications.

The consolidation and virtualization of data centers, either public or private, has consequences in terms of network requirements. It creates several new problems for private network connectivity, which are incremental the existing L3VPN technology. It is helpful to identify and analyze these problems separately:

- Private network connectivity between different data centers, either public or private.
- Private network connectivity between different compute resources within a public data center.

- Connectivity between different private networks within or across data centers.

- Content distribution between centralized public or private data centers and enterprise branch offices.

- Private network connectivity for mobile devices.

This document defines the problems under the assumption that applications require IP unicast connectivity but no layer 2 direct adjacencies. Applications with layer 2 requirements are likely to also have assumptions of other media characteristics such as round trip time, for instance.

This document is also under the assumption that both IPv4 and IPv6 unicast are in scope, but multicast service is a topic for further discussion and outside the scope of this document.

The private network service to be provided must provide traffic isolation between different VPNs allowing the use of a common infrastructure and take into account the need to reduce operational costs.

[2.](#) Terminology

AS	Autonomous Systems
DC	Data Center
IaaS	Infrastructure as a Service
LTE	Long Term Evolution
RT	Route Target
ToR	Top-of-Rack switch
VM	Virtual Machine
VMM	Virtual Machine Manager, Hypervisor
VPN	Virtual Private Network

[3.](#) VPN4DC: A problem Definition

A VPN4DC solution needs to address the following problems that are incremental to existing IPVPN solutions:

- IP only data center: defined by a data center where VM, applications, and hypervisors require only IP connectivity

and the underlying DC infrastructure is IP only.

- Network isolation among tenants or applications sharing the same data centers.
- Hypervisors may not support BGP as a control protocol.

[Page 3]

Internet Draft

VPN4DC Problem Statement

October 2011

- Fast and secure provisioning of a VPN connectivity for a VM with low operational complexity within a data center and across data centers. This includes the ability to connect a VM to a customer VPN outside the data center, thus requiring the ability to provision the communication path within the data center to the customer VPN. It also includes interconnecting VMs within and across physical data centers in the context of a virtual data center. The customer VPN service could be provided by a BGP/MPLS VPN [[RFC 4363](#)] network service provider. The VPN connectivity provisioning is targeted to be done via in-band signaling rather than an out-of-band control infrastructure. The Software Defined Network (SDN) is addressing the latter approach. It is expected that both in-band and out-of-band provisioning control will have applicability in different environments.

[4.](#) Private network connectivity between data centers

Private data centers attach to the VPN network via a CE device, which advertises the respective IP address prefixes to the network. In this space, the requirements remain unchanged from current private networks, unless we assume the ability to migrate Virtual Machines (VMs) between different data centers.

In the case that VMs are allowed to migrate between distinct data centers, this requires that each specific IP Host prefix for a VM to be advertised to the VPN network or an "home agent" approach that can redirect traffic from one data center to another (with potential negative consequences to latency).

When private networks interconnect with public data centers, the VPN provider must interconnect with the public data center provider. In this case we are in the presence of an Inter-Provider VPN in which the VPN service provider manages part of the

connectivity and in which the data center provider provides network attachment points for multiple common customers.

As with existing Inter-AS BGP/MPLS VPN scenarios, the Route Target (RT) associated with a specific VPN (in a symmetrical VPN) must be coordinated between the two entities (service provider and data center provider). The data center provider services (e.g. the API portal to its orchestration system) must also be accessible to all the carriers VPNs.

As data center providers often have different operational procedures than network services providers it is important to identify potential solutions, from operational procedures to

[Page 4]

application APIs that can exchange the necessary information between the VPN network service provider and data center provider.

5. Private Networks within a public data center

Public data centers achieve efficiencies by executing the compute loads of many different customers over a common infrastructure for compute, storage and network.

Compute nodes are often executed as Virtual Machines, in an "Infrastructure as a Service" (IaaS) data center. The set of virtual machines corresponding to a particular customer should be constrained to a private network. L3VPN technologies have proven to be able to scale to a large number of customer routes while providing for aggregated management capability. It is important to document the applicability of BGP/MPLS L3VPN technology to VMs in a data center.

It must take into account that MPLS itself is not a common technology within data centers and as such the solution must provide for IP based forwarding. It is also important to consider whether the end-system itself can contain the routing information corresponding to the VPN overlay networks without the assistance of the Top-of-Rack (ToR) switch, which may be constrained in terms of its routing table size.

6. Connectivity between different VPNs

Within a data center, the VMs within a private network will need to

communicate with data center common services such as storage or data-base services. These services often imply high traffic volumes.

The traditional approach is to deploy stateful service appliance, between different VPNs. That may become cost prohibitive for services with high volume of traffic. It is important to consider whether pushing the desired traffic control rules to the ingress points of the network (traffic sources) may assist in addressing this operational issue.

7. Mobile connectivity

Application access is being done increasingly from clients such as cell phones or tablets that may come in via a private WiFi access point or a public WiFi or 3G/LTE access. These clients must have access to application which servers reside on a private network.

[Page 5]

Internet Draft

VPN4DC Problem Statement

October 2011

It is important to consider whether it is possible to connect applications in mobile clients to provider provisioned VPNs. For instance by using IPSec tunnels; or whether these applications are best served by content caches running in the service provider infrastructure.

The solution should assume that client, VPN provider and data center may be in different Autonomous Systems.

8. Security Considerations

The document presents the problems need to be addressed in the L3VPN for data center space. The requirements and solutions will be documented separately.

The security considerations for general requirements or individual solutions will be documented in the relevant documents.

9. IANA Considerations

This document contains no new IANA considerations.

10. Normative References

[RFC 4363] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.

11. Informative References

[RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

12. Author's Address

Luyuan Fang
Cisco Systems
111 Wood Avenue South
Iselin, NJ 08830
USA
Email: lufang@cisco.com

[Page 6]

Internet Draft

VPN4DC Problem Statement

October 2011

13. Acknowledgement

The author would like to thank Pedro Marques for his helpful comments/input.

