

Network working group  
Internet Draft  
Category: Informational

L. Yong  
Huawei  
M. Toy  
Comcast  
A. Isaac  
Bloomberg  
V. Manral  
Hewlett-Packard  
L. Dunbar  
Huawei

Expires: February 2013

August 30, 2012

## Use Cases for DC Network Virtualization Overlays

[draft-mity-nvo3-use-case-03](#)

### 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 February, 2013.

### Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

NV03 Use Case

August 2012

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.

## Abstract

This draft describes the generalized NV03 use cases. The work intention is to help validate the NV03 framework and requirements as along with the development of the solutions.

## 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](#) [[RFC2119](#)].

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction.....</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Terminology.....</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Virtual Networks in a Data Center.....</a>	<a href="#">4</a>
<a href="#">4.</a>	<a href="#">Interconnecting DC Virtual Network and External Networks.....</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">DC Virtual Network Access via Internet.....</a>	<a href="#">6</a>
<a href="#">4.2.</a>	<a href="#">DC Virtual Network and WAN VPN Interconnection.....</a>	<a href="#">7</a>
<a href="#">5.</a>	<a href="#">DC Applications Using NV03.....</a>	<a href="#">9</a>
<a href="#">5.1.</a>	<a href="#">Supporting Multi Technologies in a Data Center.....</a>	<a href="#">9</a>
<a href="#">5.2.</a>	<a href="#">Tenant Virtual Network with Bridging/Routing.....</a>	<a href="#">10</a>
<a href="#">5.3.</a>	<a href="#">Virtual Data Center.....</a>	<a href="#">11</a>
<a href="#">6.</a>	<a href="#">OAM Considerations.....</a>	<a href="#">13</a>
<a href="#">7.</a>	<a href="#">Summary.....</a>	<a href="#">13</a>
<a href="#">8.</a>	<a href="#">Security Considerations.....</a>	<a href="#">14</a>
<a href="#">9.</a>	<a href="#">IANA Considerations.....</a>	<a href="#">14</a>
<a href="#">10.</a>	<a href="#">Acknowledgements.....</a>	<a href="#">14</a>
<a href="#">11.</a>	<a href="#">References.....</a>	<a href="#">14</a>
<a href="#">11.1.</a>	<a href="#">Normative References.....</a>	<a href="#">14</a>
<a href="#">11.2.</a>	<a href="#">Informative References.....</a>	<a href="#">15</a>
	<a href="#">Authors' Addresses.....</a>	<a href="#">16</a>

## 1. Introduction

Compute Virtualization has dramatically and quickly changed IT industry in terms of efficiency, cost, and the speed in providing a new applications and/or services. However the problems in today's data center hinder the support of an elastic cloud service and dynamic virtual tenant networks [[NVO3PRBM](#)]. The goal of DC Network Virtualization Overlays, i.e. NV03, is to decouple the end systems (VMs) and DC physical networks and to allow the network infrastructure to provide: 1) traffic isolation among one virtual network and another; 2) independent address space in each virtual network and address isolation from the infrastructure's; 3) VM placement and move from one server to another without any physical network limitation. These characteristics will help address the issues in the data centers.

Although NV03 may enable a true virtual environment where VMs and net service appliances communicate, the NV03 solution has to address how to communicate between a virtual network and a physical network. This is because 1) many traditional DCs exist and will not disappear any time soon; 2) a lot of DC applications serve to Internet and/or cooperation users; 3) some applications like Big Data analytics which are CPU bound may not want to the virtualization capability.

This document is to describe generalized NV03 use cases in various data center networks to make sure the future framework and solutions can meet the demand. Three types of the use cases are:

- o A virtual network connects many tenant end systems within a Data Center and form one L2 or L3 communication domain. A virtual network segregates its traffic from others and allows the VMs in the network moving from one server to another. The case may be used for DC internal applications that constitute the DC East-West traffic.
- o A DC provider offers a secure DC service to an enterprise customer and/or Internet users. In these cases, the enterprise customer may use a traditional VPN provided by a carrier or an

IPsec tunnel over Internet connecting to an overlay virtual network offered by a Data Center provider. This is mainly constitutes DC North-South traffic.

- o A DC provider uses NV03 to design a variety of DC applications that make use of the net service appliance, virtual compute, storage, and networking. In this case, the NV03 provides the virtual networking functions for the applications.

MITY

Expires February 2013

[Page 3]

---

Internet-Draft

NV03 Use Case

August 2012

The document uses the architecture reference model and terminologies defined in [[NV03FRWK](#)] to describe the use cases.

## [2.](#) Terminology

This document uses the terminologies defined in [[NV03FRWK](#)], [[RFC4364](#)]. Some additional terms used in the document are listed here.

CUG: Closed User Group

L2 VNI: L2 Virtual Network Instance

L3 VNI: L3 Virtual Network Instance

ARP: Address Resolution Protocol

CPE: Customer Premise Equipment

DNS: Domain Name Service

DMZ: DeMilitarized Zone

NAT: Network Address Translation

VNIF: Internal Virtual Network Interconnection Interface

## [3.](#) Virtual Networks in a Data Center

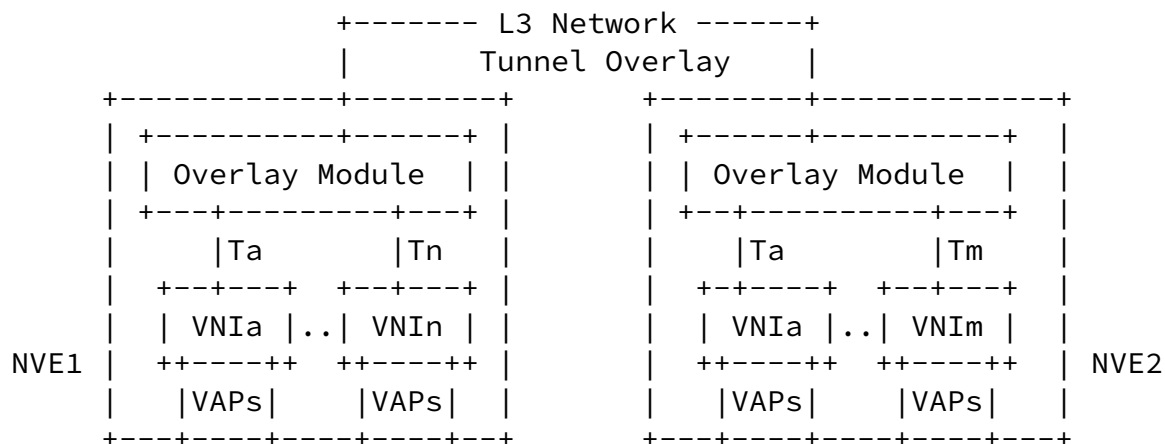
A tenant virtual network may exist within a DC. The network interconnects many tenant end systems (ESs) which can be hosted by physical servers or virtual machines (VMs). The network enables the communication among the ESs that are considered as Closed User Groups (CUG). A virtual network has a unique virtual network

identifier for switches/routers to properly differentiate them from virtual networks. The Closed User Groups are formed so that proper policies can be applied when the ESs in one CUG communicate with the ESs in other CUGs.

Figure 1 depicts this case by using the framework model. [NV03FRWK] NVE1 and NVE2 are two network virtual edges and each may exist on a server or ToR. Each NVE may be the members of multiple virtual networks that may have different topologies and run at L2 or L3 individually. In this illustration, three virtual networks with VN context Ta, Tn, and Tm are shown. The VN<sub>Ia</sub> terminates on both NVE1 and NVE2; The VN<sub>In</sub> terminates on NVE1 and the VN<sub>Im</sub> at NVE2 only. Each NVE has one overlay module to perform frame

encapsulation/decapsulation and tunneling initiation/termination. In this scenario, a tunnel between NVE1 and NVE2 is necessary for the virtual network Ta. Note that it is possible that one TES participates in more than one virtual network via one VAP for each; further if individual virtual networks use different address spaces, the TES participating in them will be configured with multiple addresses as well. A TES as a gateway is an example.

A VNI on an NVE is a forwarding table that caches and/or maintains the mapping of an end system and its attached NVE. The table entry may be updated by the control plane or data plane or the combination of both. A TES associates to one VNI via a VAP. One tenant virtual network may terminate on many NVEs and interconnect several thousands of TESs, the capability of supporting a lot of TESs per tenant instance and TES mobility is critical for NV03 solution no matter where an NVE resides.



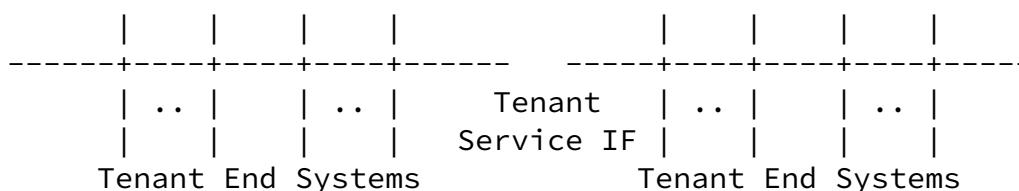


Figure 1 NV03 for Tenant End-System interconnection

Individual virtual networks may use its own address space and the space is isolated from DC infrastructure. This eliminates the route changes in the DC underlying network when VMs move. Note that the NV03 solutions still have to address VM move in the overlay network, i.e. the TES/NVE association change when a VM moves.

It is worth mentioning two scenarios regarding to the NVE location. At first an NVE resides on a server, a server manager system such as vCenter [[VMWARE](#)] is responsible to create NVE/VNs and VMs, and also responsible to assign a VM to a VN that has unique identification, the server software just makes it works properly and securely. Second an NVE resides on physical switch such as ToR, when a server manger system creates a VM and add it to a VN, the server will send a notification of TES participating in a VN to the local NVE. [[ESYS](#)][VDP] Note that if non-virtualized server is used, local configuration on NVE is necessary to attach the TES (server) to a VN. In both cases, when a local NVE notices the new attached TES in a VN, it will announce the TES to remote NVEs or to a mapping server via a control plane protocol. In the case of using mapping server, the remote NVEs can query the server for any TES location and cache it in the VNI.

If a tenant virtual network spans across multiple DC sites, one design is to allow the corresponding NV03 instance seamlessly span across those sites without DC gateway routers' termination (see [section 4.3](#)). In this case, the tunnel between a pair of NVEs may in turn be tunneled over other intermediate tunnels over the Internet or other WANs, or the intra DC and inter DC tunnels are stitched together to form an end-to-end tunnel between two NVEs.

#### 4. Interconnecting DC Virtual Network and External Networks

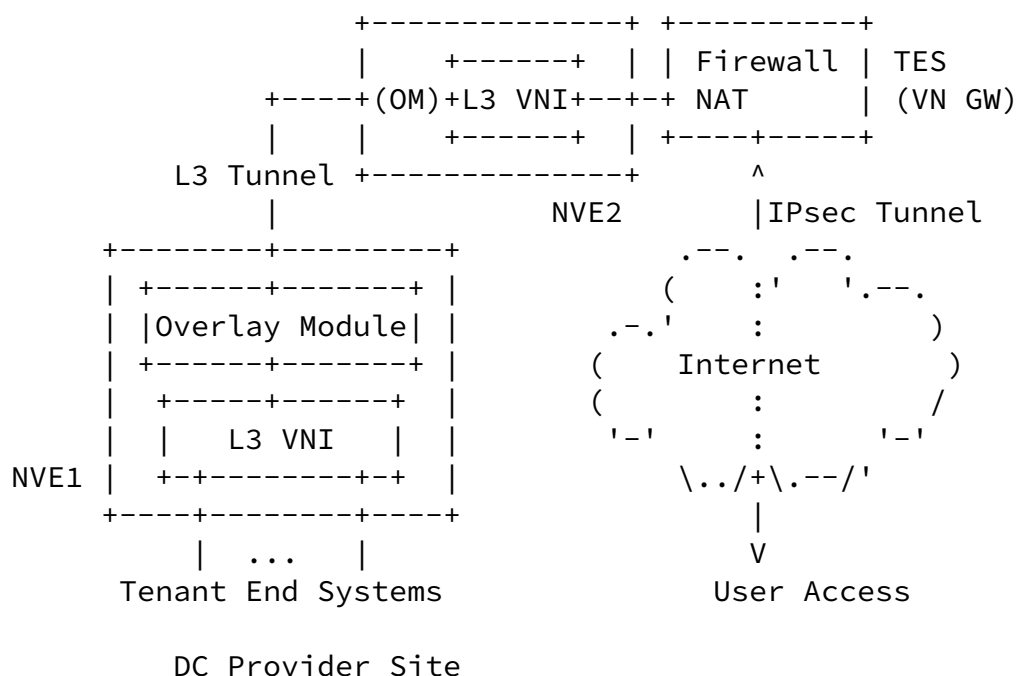
For customers (an enterprise or individuals) who want to utilize the

DC provider's compute and storage resources to run their applications, they need to access those end systems hosted in a DC through Carrier WANs or Internet. A DC provider may want to use an NV03 virtual network to connect these end systems; then it, in turn, becomes the case of interconnecting DC virtual network and external networks. Two cases are described here.

#### 4.1. DC Virtual Network Access via Internet

A user or an enterprise customer may want to connect to a DC virtual network via Internet but securely. Figure 2 illustrates this case. An L3 virtual network is configured on NVE1 and NVE2 and two NVEs are connected via an L3 tunnel in the Data Center. A set of tenant end systems attach to NVE1. The NVE2 connects to one (may be more) TES that runs the VN gateway and NAT applications (known as net service appliance). A user or customer can access the VN via Internet by using IPsec tunnel [RFC4301]. The encrypted tunnel is established between the VN GW and the user machine or CPE at enterprise location. The VN GW provides authentication scheme and

encryption. Note that VN GW function may be performed by a net service appliance or on a DC GW.



OM: Overlay Module;

Figure 2 DC Virtual Network Access via Internet

#### [4.2](#). DC Virtual Network and WAN VPN Interconnection

A DC Provider and Carrier may build a VN and VPN independently and interconnect the two at the DC GW and PE for an enterprise customer. Figure 3 depicts this case in a L3 overlay (L2 overlay is the same). The DC provider constructs an L3 VN between the NVE1 on a server and the NVE2 on the DC GW in the DC site; the carrier constructs an L3VPN between PE1 and PE2 in its IP/MPLS network. An Ethernet Interface physically connects the DC GW and PE2 devices. The local VLAN over the Ethernet interface [[VRF-LITE](#)] is configured to connect the L3VNI/NVE2 and VRF, which makes the interconnection between the L3 VN in the DC and the L3VPN in IP/MPLS network. An Ethernet Interface may be used between PE1 and CE to connect the L3VPN and enterprise physical networks.

This configuration allows the enterprise networks communicating to the L3 VN as if its own networks but not communicating with DC provider physical networks as well as not other overlay networks in the DC. The enterprise may use its own address space on the L3 VN.

MITY

Expires February 2013

[Page 7]

---

Internet-Draft

NV03 Use Case

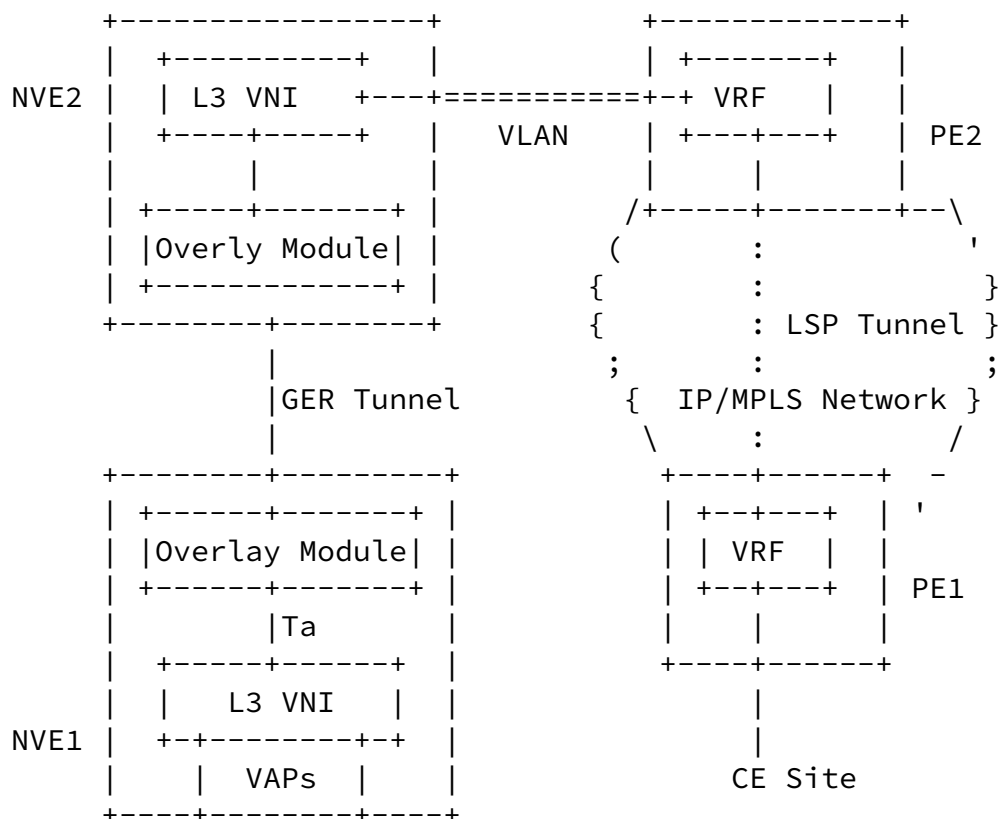
August 2012

The DC provider can manage the VM and storage assignment to the L3 VN for the enterprise customer. The enterprise customer can determine and load the applications running on the VMs. From the L3 VN perspective, an end point in the enterprise location appears as the end point associating to the NVE2. The NVE2 on the DC GW has to perform both the GRE tunnel termination [[RFC4797](#)] and the local VLAN termination and forward the packets in between. The DC provider and Carrier negotiate the local VLAN ID used on the Ethernet interface.

This configuration makes the L3VPN over the WANs only has the reachability of the L3 VN. It does not have the reachability of DC physical networks and other VNs in the DC. However, the L3VPN has the reachability of enterprise networks. Note that both the DC provider and enterprise may have multiple network locations connecting to the L3VPN.

The eBGP protocol can be used between DC GW and PE2 for the route population in between. In fact, this is like the Option A in [[RFC4364](#)]. This configuration can work with any NV03 solution. The eBGP, OSPF, or other can be used between PE1 and CE for the route

population.



| ... |  
Tenant End Systems

Enterprise Site

DC Provider Site

Figure 3 L3 VNI and L3VPN interconnection across multi networks

If an enterprise only has one location, it may use P2P VPWS [[RFC4664](#)] or L2TP [[RFC5641](#)] to connect one DC provider site. In this case, one edge connects to a physical network and another edge connects to an overlay network.

The interesting feature in this use case is that the L3 VN and compute resource are managed by the DC provider. The DC operator can place them at any location without notifying the enterprise and carrier because the DC physical network is completely isolated from the carrier and enterprise network. Furthermore, the DC operator may move the compute resources assigned to the enterprise from one place

to another in the DC without the enterprise customer awareness, i.e. no impact on the enterprise 'live' applications running these resources. Such advanced feature brings some new requirements for NV03 [[NV03PRBM](#)].

## [5.](#) DC Applications Using NV03

NV03 brings DC operators the flexibility to design different applications in a virtual environment without worry about physical network configuration in the Data Center. DC operators may build several virtual networks and interconnect them directly to form a tenant virtual network; or may allocate some VMs to run tenant applications and some to run net service applications such as Firewall, DNS for the tenant. Several use cases are given in this section.

### [5.1.](#) Supporting Multi Technologies in a Data Center

Most likely servers deployed in a large data center are rolled in at different times and may have different capacities/features. Some servers may be virtualized, some may not; some may be equipped with virtual switches, some may not. For the ones equipped with hypervisor based virtual switches, some may support VxLAN [[VXLAN](#)] encapsulation, some may support NvGRE encapsulation [[NVGRE](#)], and some may not support any types of encapsulation. To construct a tenant virtual network among these servers and the ToRs, it may use two virtual networks and a gateway to allow different implementations working together. For example, one virtual network

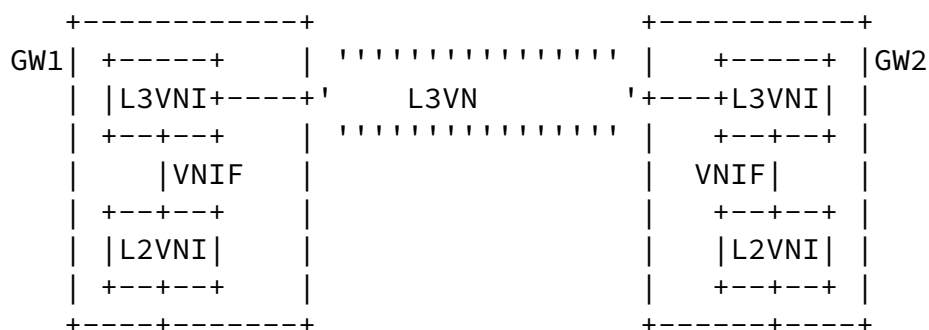
uses VxLAN encapsulation and another virtual network uses traditional VLAN.

The gateway entity, either on VMs or standalone one, participates in to both virtual networks, and maps the services and identifiers and changes the packet encapsulations.

### [5.2.](#) Tenant Virtual Network with Bridging/Routing

A tenant virtual network may span across multiple Data Centers. DC operator may want to use L2VN within a DC and L3VN outside DCs for a tenant. This is very similar to today's DC physical network configuration. L2 bridging has the simplicity and endpoint awareness while L3 routing has advantages in aggregation and scalability. For

Figure 5 depicts two DC sites. The site A constructs an L2VN that terminates on NVE1, NVE2, and GW1. An L3VN is configured between the GW1 at site A and the GW2 at site Z. An internal Virtual Network Interconnection Interface (VNIF) connects to L2VNI and L3VNI on GW1. Thus the GW1 is the members of the L2VN and L3VN. The L2VNI is the MAC/NVE mapping table and the L3VNI is IP prefix/NVE mapping table. Note that a VNI also has the mapping of TES and VAP at the local NVE. The site Z has the similar configuration. A packet coming to the GW1 from L2VN will be decapsulated and converted into an IP packet and then encapsulated and sent to the site Z. The Gateway uses ARP protocol to obtain MAC/IP mapping. Note that both the L2VN and L3VN in the figure are carried by the tunnels supported by the underlying networks which are not shown in the figure.



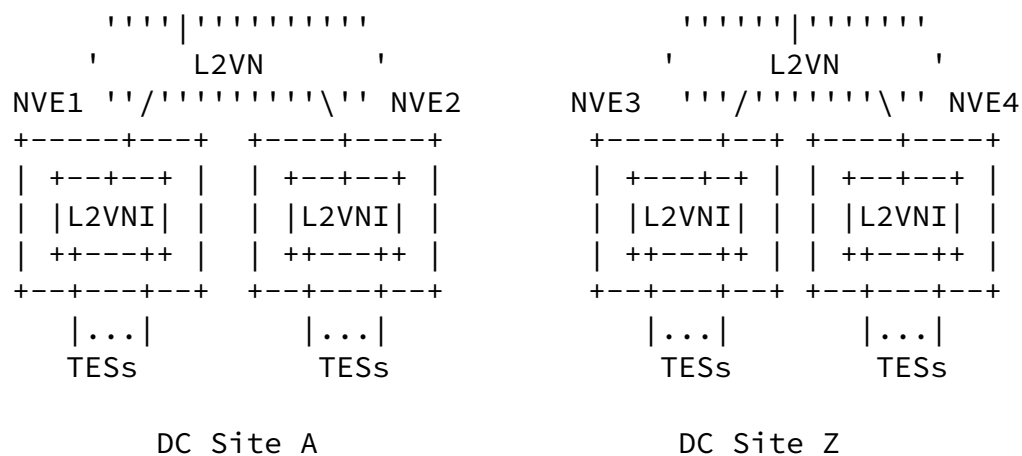


Figure 4 Tenant Virtual Network with Bridging/Routing

### 5.3. Virtual Data Center

Enterprise DC's today may often use several routers, switches, and service devices to construct its internal network, DMZ, and external network access. A DC Provider may offer a virtual DC to an enterprise customer to run enterprise applications such as website/emails. Instead of using many hardware devices, with the overlay and virtualization technology of NV03, DC operators can build them on top of a common network infrastructure for many customers and run service applications per customer basis. The service applications may include firewall, gateway, DNS, load balancer, NAT, etc.

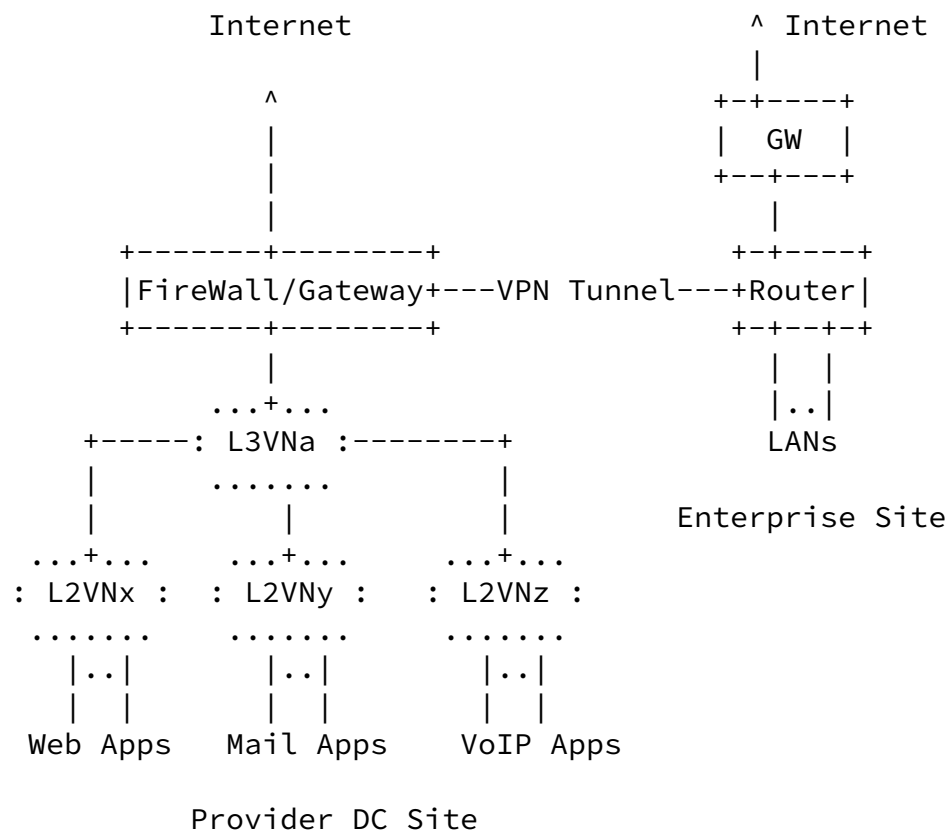
Figure 6 below illustrates this scenario. For the simple illustration, it only shows the L3VN or L2VN as virtual and overlay routers or switches. In this case, DC operators construct several L2 VNs (L2VNx, L2VNy, L2VNz in figure 6) to group the end tenant systems together per application basis, create an L3VNa for the internal routing. A server or VM runs firewall/gateway applications and connects to the L3VNa and Internet. A VPN tunnel is also built between the gateway and enterprise router. The design runs

Enterprise Web/Mail/VoIP applications at the provider DC site; lets the users at Enterprise site to access the applications via the VPN tunnel and Internet via a gateway at the Enterprise site; let Internet users access the applications via the gateway in the provider DC. The enterprise operators can also use the VPN tunnel or IPsec over Internet to access the vDC for the management purpose. The firewall/gateway provides application-level and packet-level

gateway function and/or NAT function.

The Enterprise customer decides which applications are accessed by intranet only and which by both intranet and extranet; DC operators then design and configure the proper security policy and gateway function. DC operators may further set different QoS levels for the different applications for a customer.

This application requires the NV03 solution to provide the DC operator an easy way to create NVEs and VNIs for any design and to quickly assign TESSs to a VNI, and easily configure policies on an NVE.



\* firewall/gateway may run on a server or VMs

Figure 5 Virtual Data Center by Using NV03

## 6. OAM Considerations

NV03 brings the ability for a DC provider to segregate tenant traffic. A DC provider needs to manage and maintain NV03 instances. Similarly, the tenant needs to be informed about tunnel failures impacting tenant applications.

Various OAM and SOAM tools and procedures are defined in [IEEE 802.1ag, ITU-T Y.1731, [RFC4378](#), [RFC5880](#), ITU-T Y.1564] for L2 and L3 networks, and for user, including continuity check, loopback, link trace, testing, alarms such as AIS/RDI, and on-demand and periodic measurements. These procedures may apply to tenant overlay networks and tenants not only for proactive maintenance, but also to ensure support of Service Level Agreements (SLAs).

As the tunnel traverses different networks, OAM messages need to be translated at the edge of each network to ensure end-to-end OAM.

It is important that failures at lower layers which do not affect NV03 instance are to be suppressed.

## [7](#). Summary

The document describes some basic potential use cases of NV03. The combination of these cases should give operators flexibility and power to design more sophisticated cases for various purposes.

The main differences between other overlay network technologies and NV03 is that the client edges of the NV03 network are individual and virtualized hosts, not network sites or LANs. NV03 enables these virtual hosts communicating in a true virtual environment without considering physical network configuration.

NV03 allows individual tenant virtual networks to use their own address space and isolates the space from the network infrastructure. The approach not only segregates the traffic from multi tenants on a common infrastructure but also makes VM placement and move easier.

DC applications are about providing virtual processing/storage, applications, and networking in a secured and virtualized manner, in which the NV03 is just a portion of an application. NV03 decouples the applications and DC network infrastructure configuration.

NV03's underlying network provides the tunneling between NVEs so that two NVEs appear as one hop to each other. Many tunneling technologies can serve this function. The tunneling may in turn be tunneled over other intermediate tunnels over the Internet or other WANs. It is also possible that intra DC and inter DC tunnels are stitched together to form an end-to-end tunnel between two NVEs.

A DC virtual network may be accessed via an external network in a secure way. Many existing technologies can achieve this.

The key requirements for NV03 are 1) traffic segregation; 2) supporting a large scale number of virtual networks in a common infrastructure; 3) supporting highly distributed virtual network with sparse memberships 3) VM mobility 4) auto or easy to construct a NVE and its associated TES; 5) Security 6) NV03 Management [[NV03PRBM](#)].

## [8. Security Considerations](#)

Security is a concern. DC operators need to provide a tenant a secured virtual network, which means the tenant traffic isolated from other tenant's and non-tenant VMs not placed into the tenant virtual network; they also need to prevent DC underlying network from any tenant application attacking through the tenant virtual network or one tenant application attacking another tenant application via DC networks. For example, a tenant application attempts to generate a large volume of traffic to overload DC underlying network. The NV03 solution has to address these issues.

## [9. IANA Considerations](#)

This document does not request any action from IANA.

## [10. Acknowledgements](#)

Authors like to thank Sue Hares, Young Lee, David Black, Pedro Marques, Mike McBride, David McDysan, and Randy Bush for the review, comments, and suggestions.

## [11. References](#)

### [11.1. Normative References](#)

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [IEEE 802.1ag] "Virtual Bridged Local Area Networks - Amendment 5: Connectivity Fault Management", December 2007.
- [ITU-T G.8013/Y.1731] OAM Functions and Mechanisms for Ethernet based Networks, 2011.
- [ITU-T Y.1564] "Ethernet service activation test methodology", 2011.
- [RFC4378] Allan, D., Nadeau, T., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", [RFC4378](#), February 2006
- [RFC4301] Kent, S., "Security Architecture for the Internet Protocol", [rfc4301](#), December 2005
- [RFC4664] Andersson, L., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", [rfc4664](#), September 2006
- [RFC4797] Rekhter, Y., etc, "Use of Provider Edge to Provider Edge (PE-PE) Generic Routing Encapsulation (GRE) or IP in BGP/MPLS IP Virtual Private Networks", [RFC4797](#), January 2007
- [RFC5641] McGill, N., "Layer 2 Tunneling Protocol Version 3 (L2TPv3) Extended Circuit Status Values", [rfc5641](#), April 2009.
- [RFC5880] Katz, D. and Ward, D., "Bidirectional Forwarding Detection (BFD)", [rfc5880](#), June 2010.

### 11.2. Informative References

- [ESYS] Marques, P., "End-System support for BGP-signaled IP/VPNs", [draft-marques-l3vpn-end-system-07](#), August 2012
- [NVGRE] Sridharan, M., "NVGRE: Network Virtualization using Generic Routing Encapsulation", [draft-sridharan-virtualization-nvgre-01](#), July 2012
- [NV03PRBM] Narten, T., etc "Problem Statement: Overlays for Network Virtualization", [draft-narten-nvo3-overlay-problem-statement-04](#), August 2012

Internet-Draft

NV03 Use Case

August 2012

[NV03FRWK] Lasserre, M., Motin, T., and etc, "Framework for DC Network Virtualization", [draft-lasserre-nvo3-framework-03](#), July 2012

[VDP] "IEEE P802.1Qbg Edge Virtual Bridging".

[VMWARE] VMware, "vCenter", <http://www.vmware.com>

[VRF-LITE] Cisco, "Configuring VRF-lite", <http://www.cisco.com>

[VXLAN] Mahalingam, M., Dutt, D., etc "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", [draft-mahalingam-dutt-dcops-vxlan-02.txt](#), August 2012

#### Authors' Addresses

Lucy Yong  
Huawei Technologies,  
4320 Legacy Dr.  
Plano, Tx75025 US

Phone: +1-469-277-5837  
Email: [lucy.yong@huawei.com](mailto:lucy.yong@huawei.com)

Mehmet Toy  
Comcast  
1800 Bishops Gate Blvd.,  
Mount Laurel, NJ 08054

Phone : +1-856-792-2801  
E-mail : [mehmet\\_toy@cable.comcast.com](mailto:mehmet_toy@cable.comcast.com)

Aldrin Isaac  
Bloomberg  
E-mail: [aldrin.isaac@gmail.com](mailto:aldrin.isaac@gmail.com)

Vishwas Manral  
Hewlett-Packard Corp.  
191111 Pruneridge Ave.  
Cupertino, CA 95014

Phone: 408-447-1497  
Email: vishwas.manral@hp.com

MITY

Expires February 2013

[Page 16]

---

Internet-Draft

NV03 Use Case

August 2012

Linda Dunbar  
Huawei Technologies,  
4320 Legacy Dr.  
Plano, Tx75025 US

Phone: +1-469-277-5840  
Email: linda.dunbar@huawei.com

MITY

Expires February 2013

[Page 17]