

Network Working Group
Internet Draft
Category: Informational

L. Yong
L. Dunbar
Huawei
M. Toy
Verizon
A. Isaac
Juniper Networks
V. Manral
Ionos Networks

Expires: June 2017

December 21, 2016

Use Cases for Data Center Network Virtualization Overlay Networks

[draft-ietf-nvo3-use-case-15](#)

Abstract

This document describes data center network virtualization overlay (NVO3) network use cases that can be deployed in various data centers and serve different data center applications.

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/1id-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 June 21, 2017.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

1.	Introduction.....	3
1.1.	Terminology.....	4
2.	Basic NV03 Networks.....	5
3.	DC NV03 Network and External Network Interconnection.....	6
3.1.	DC NV03 Network Access via the Internet.....	6
3.2.	DC NV03 Network and SP WAN VPN Interconnection.....	7
4.	DC Applications Using NV03.....	8
4.1.	Supporting Multiple Technologies.....	9
4.2.	DC Application with Multiple Virtual Networks.....	9
4.3.	Virtual Data Center (vDC).....	10
5.	Summary.....	12
6.	Security Considerations.....	12
7.	IANA Considerations.....	12
8.	Informative References.....	13
	Contributors.....	14
	Acknowledgements.....	14
	Authors' Addresses.....	14

1. Introduction

Server virtualization has changed the Information Technology (IT) industry in terms of the efficiency, cost, and speed of providing new applications and/or services such as cloud applications. However traditional data center (DC) networks have some limits in supporting cloud applications and multi tenant networks [[RFC7364](#)]. The goal of data center network virtualization overlay (NV03) networks is to decouple the communication among tenant systems from DC physical infrastructure networks and to allow one physical network infrastructure:

- o Carry many NV03 networks and isolate different NV03 network traffic on a physical network that carries NV03 network traffic.
- o Independent address spaces in individual NV03 networks such as MAC and IP.
- o Flexible Virtual Machines (VM) and/or workload placement including the ability to move them from one server to another without requiring VM address changes and physical infrastructure network configuration changes, and the ability to perform a "hot move" with no disruption to the live application running on VMs.

These characteristics of NV03 networks help address the issues that cloud applications face in data centers [[RFC7364](#)].

An NV03 network may interconnect with another NV03 network on the same physical network, or another physical network (i.e., not the physical network that the NV03 network is carried over), via a gateway. The use case examples for the latter are: 1) DCs that migrate toward an NV03 solution will be done in steps, where a portion of tenant systems in a VN is on virtualized servers while others exist on a LAN. 2) many DC applications serve to Internet users who are on physical networks; 3) some applications are CPU bound, such as Big Data analytics, and may not run on virtualized resources. Some inter-VN policies can be enforced at the gateway.

This document describes general NV03 network use cases that apply to various data centers. The use cases described here represent DC provider's interests and vision for their cloud services. The document groups the use cases into three categories from simple to advance in term of implementation. However the implementations of these use cases are outside the scope of this document. These three categories are highlighted below:

- o Basic NV03 networks ([Section 2](#)). All Tenant Systems (TS) in the network are located within the same DC. The individual networks can be either Layer 2 (L2) or Layer 3 (L3). The number of NV03 networks in a DC is much higher than what traditional VLAN based virtual networks [IEEE 802.1Q] can support. This case is often referred as to the DC East-West traffic.
- o A virtual network that spans across multiple Data Centers and/or to customer premises where NV03 networks are constructed and interconnect another virtual or physical network outside the data center. An enterprise customer may use a traditional carrier VPN or an IPsec tunnel over the Internet to communicate with its systems in the DC. This is described in [Section 3](#).
- o DC applications or services require an advanced network that contains several NV03 networks that are interconnected by the gateways. Three scenarios are described in [Section 4.1](#)) supporting multiple technologies; 2) constructing several virtual networks as a tenant network; 3) applying NV03 to a virtual Data Center (vDC).

The document uses the architecture reference model defined in [[RFC7365](#)] to describe the use cases.

1.1. Terminology

This document uses the terminologies defined in [[RFC7365](#)] and [[RFC4364](#)]. Some additional terms used in the document are listed here.

DMZ: Demilitarized Zone. A computer or small sub-network that sits between a trusted internal network, such as a corporate private LAN, and an un-trusted external network, such as the public Internet.

DNS: Domain Name Service [[RFC1035](#)]

DC Operator: A role who is responsible to construct and manage cloud service instances in their life-cycle and manage DC infrastructure that runs these cloud instances.

DC Provider: A company that uses its DC infrastructure to offer cloud services to its customers.

NAT: Network Address Translation [[RFC3022](#)]

vGW: virtual Gateway; a gateway component used for an NV03 virtual network to interconnect with another virtual/physical network.

2. Basic NV03 Networks

An NV03 network provides communications among Tenant Systems (TS) in a DC. A TS can be a physical server/device or a virtual machine (VM) on a server, i.e., end-device [[RFC7365](#)]. A DC provider often uses NV03 networks for its internal applications in which each application runs on many VMs or physical services and requires application segregation.

A Network Virtual Edge (NVE) is an NV03 architecture component [[RFC7365](#)]. It is responsible to forward and encapsulate the NV03 traffic in outbound direction; and decapsulate and forward the NV03 traffic in inbound direction [[NV03ARCH](#)]. A Network Virtualization Authority (NVA) is another NV03 architecture component [[RFC7365](#)]. An NVE obtains the reachability information of tenant systems in a NV03 network from the NVA. The tenant systems attached to the same NVE may belong to a same or different NV03 networks.

The network virtualization overlay in this context means that a virtual network is implemented with an overlay technology, i.e., within a DC, NV03 traffic is encapsulated at an NVE and carried by a tunnel to another NVE where the packet is decapsulated and sent to a target tenant system [[NV03ARCH](#)]. This architecture decouples an NV03 network construction from the DC physical network configuration, which provides the flexibility for VM placement and mobility. The architecture supports one tunnel to carry NV03 traffic belonging to different NV03 networks; thus the NV03 encapsulation header carries a virtual network identifier to differentiate NV03 traffic in a tunnel.

An NV03 network may be an L2 or L3 domain. The network provides switching (L2) or routing (L3) capability to support host (i.e. tenant systems) communications. An NV03 network may required to carry unicast traffic and/or multicast, broadcast/unknown (for L2 only) traffic from/to tenant systems. There are several ways to transport NV03 network BUM traffic [[NV03MCAST](#)].

It is worth mentioning two distinct cases regarding to NVE location. The first is where TSs and an NVE are co-located on a single end host/device, which means that the NVE can be aware of the TS's state at any time via an internal API. The second is where TSs and an NVE are not co-located, with the NVE residing on a network device; in this case, a protocol is necessary to allow the NVE to be aware of the TS's state [[NV03HYVR2NVE](#)].

One NV03 network can provide connectivity to many TSs that attach to many different NVEs in a DC. TS dynamic placement and mobility

results in frequent changes of the binding between a TS and an NVE. The TS reachability update mechanisms need be fast enough so that the updates do not cause any communication disruption/interruption. The capability of supporting many TSs in a virtual network and many more virtual networks in a DC is critical for the NV03 solution.

If a virtual network spans across multiple DC sites, one design using NV03 is to allow the network to seamlessly span across the sites without DC gateway routers' termination. In this case, the tunnel between a pair of NVEs can be carried within other intermediate tunnels over the Internet or other WANs, or an intra DC tunnel and inter DC tunnel(s) can be stitched together to form an end-to-end tunnel between the pair of NVEs that are in different DC sites. Both cases will form one virtual network across multiple DC sites.

3. DC NV03 Network and External Network Interconnection

Many customers (an enterprise or individuals) who utilize a DC provider's compute and storage resources to run their applications need to access their systems hosted in a DC through Internet or Service Providers' Wide Area Networks (WAN). A DC provider can construct a NV03 network that provides connectivity to all the resources designated for a customer and allows the customer to access the resources via a virtual gateway (vGW). This, in turn, becomes the case of interconnecting an NV03 network and the virtual private network (VPN) on the Internet or wide-area networks (WAN). Note that a VPN is not implemented by NV03 solution. Two use cases are described here.

3.1. DC NV03 Network Access via the Internet

A customer can connect to an NV03 network via the Internet in a secure way. Figure 1 illustrates an example of this case. The NV03 network has an instance at NVE1 and NVE2 and the two NVEs are connected via an IP tunnel in the Data Center. A set of tenant systems are attached to NVE1 on a server. NVE2 resides on a DC Gateway device. NVE2 terminates the tunnel and uses the VNID on the packet to pass the packet to the corresponding vGW entity on the DC GW (the vGW is the default gateway for the virtual network). A customer can access their systems, i.e., TS1 or TS_n, in the DC via the Internet by using an IPsec tunnel [[RFC4301](#)]. The IPsec tunnel is configured between the vGW and the customer gateway at the customer site. Either a static route or iBGP may be used for prefix advertisement. The vGW provides IPsec functionality such as authentication scheme and encryption; iBGP protocol traffic is carried within the IPsec tunnel. Some vGW features are listed below:

- o The vGW maintains the TS/NVE mappings and advertises the TS prefix to the customer via static route or iBGP.
- o Some vGW functions such as firewall and load balancer can be performed by locally attached network appliance devices.
- o If the NV03 network uses different address space than external users, then the vGW needs to provide the NAT function.
- o More than one IPsec tunnel can be configured for redundancy.
- o The vGW can be implemented on a server or VM. In this case, IP tunnels or IPsec tunnels can be used over the DC infrastructure.
- o DC operators need to construct a vGW for each customer.

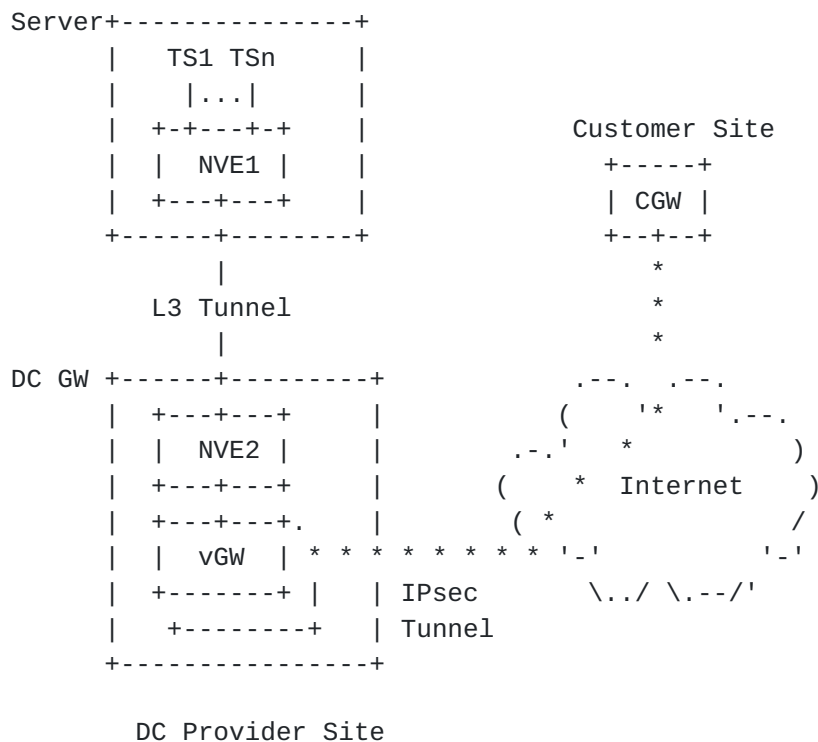


Figure 1 - DC Virtual Network Access via the Internet

3.2. DC NV03 Network and SP WAN VPN Interconnection

In this case, an Enterprise customer wants to use a Service Provider (SP) WAN VPN [RFC4364] [RFC7432] to interconnect its sites with an NV03 network in a DC site. The Service Provider constructs a VPN for

the enterprise customer. Each enterprise site peers with an SP PE. The DC Provider and VPN Service Provider can build an NV03 network and a WAN VPN independently, and then interconnect them via a local link, or a tunnel between the DC GW and WAN PE devices. The control plane interconnection options between the DC and WAN are described in [RFC4364](#) [[RFC4364](#)]. Using Option A with VRF-LITE [[VRF-LITE](#)], both ASBRs, i.e., DC GW and SP PE, maintain a routing/forwarding table (VRF). Using Option B, the DC ASBR and SP ASBR do not maintain the VRF table; they only maintain the NV03 network and VPN identifier mappings, i.e., label mapping, and swap the label on the packets in the forwarding process. Both option A and B allow the NV03 network and VPN using own identifier and two identifiers are mapped at DC GW. With option C, the VN and VPN use the same identifier and both ASBRs perform the tunnel stitching, i.e., tunnel segment mapping. Each option has pros/cons [[RFC4364](#)] and has been deployed in SP networks depending on the applications in use. BGP is used with these options for route distribution between DCs and SP WANs. Note that if the DC is the SP's Data Center, the DC GW and SP PE in this case can be merged into one device that performs the interworking of the VN and VPN within an AS.

The configurations above allow the enterprise networks to communicate with the tenant systems attached to the NV03 network in the DC without interfering with the DC provider's underlying physical networks and other NV03 networks in the DC. The enterprise can use its own address space in the NV03 network. The DC provider can manage which VM and storage elements attach to the NV03 network. The enterprise customer manages which applications run on the VMs without knowing the location of the VMs in the DC. (See [Section 4](#) for more)

Furthermore, in this use case, the DC operator can move the VMs assigned to the enterprise from one server to another in the DC without the enterprise customer being aware, i.e., with no impact on the enterprise's 'live' applications. Such advanced technologies bring DC providers great benefits in offering cloud services, but add some requirements for NV03 [[RFC7364](#)] as well.

[4. DC Applications Using NV03](#)

NV03 technology provides DC operators with the flexibility in designing and deploying different applications in an end-to-end virtualization overlay environment. The operators no longer need to worry about the constraints of the DC physical network configuration when creating VMs and configuring a network to connect them. A DC provider may use NV03 in various ways, in conjunction with other

physical networks and/or virtual networks in the DC for a reason. This section highlights some use cases for this goal.

4.1. Supporting Multiple Technologies

Servers deployed in a large data center are often installed at different times, and may have different capabilities/features. Some servers may be virtualized, while others may not; some may be equipped with virtual switches, while others may not. For the servers equipped with Hypervisor-based virtual switches, some may support a standardized NV03 encapsulation, some may not support any encapsulation, and some may support a documented encapsulation protocol (e.g. VxLAN [[RFC7348](#)], NVGRE [[RFC7637](#)]) or proprietary encapsulations. To construct a tenant network among these servers and the ToR switches, operators can construct one traditional VLAN network and two virtual networks where one uses VxLAN encapsulation and the other uses NVGRE, and interconnect these three networks via a gateway or virtual GW. The GW performs packet encapsulation/decapsulation translation between the networks.

Another case is that some software of a tenant is high CPU and memory consumption, which only makes a sense to run on metal servers; other software of the tenant may be good to run on VMs. However provider DC infrastructure is configured to use NV03 to connect to VMs and VLAN [[IEEE802.1Q](#)] connect to metal services. The tenant network requires interworking between NV03 and traditional VLAN.

4.2. DC Application with Multiple Virtual Networks

A DC application may necessarily be constructed with multi-tier zones, where each zone has different access permissions and runs different applications. For example, a three-tier zone design has a front zone (Web tier) with Web applications, a mid zone (application tier) where service applications such as credit payment or ticket booking run, and a back zone (database tier) with Data. External users are only able to communicate with the Web application in the front zone; the back zone can only receive traffic from the application zone. In this case, communications between the zones must pass through a GW/firewall. Each zone can be implemented by one NV03 network and a GW/firewall can be used to between two NV03 networks, i.e., two zones. As a result, a tunnel carrying NV03 network traffic must be terminated at the GW/firewall where the NV03 traffic is processed.

4.3. Virtual Data Center (vDC)

An enterprise data center today may deploy routers, switches, and network appliance devices to construct its internal network, DMZ, and external network access; it may have many servers and storage running various applications. With NV03 technology, a DC Provider can construct a virtual Data Center (vDC) over its physical DC infrastructure and offer a virtual Data Center service to enterprise customers. A vDC at the DC Provider site provides the same capability as the physical DC at a customer site. A customer manages its own applications running in its vDC. A DC Provider can further offer different network service functions to the customer. The network service functions may include firewall, DNS, load balancer, gateway, etc.

Figure 2 below illustrates one such scenario at service abstraction level. In this example, the vDC contains several L2 VNs (L2VNX, L2VNY, L2VNZ) to group the tenant systems together on a per-application basis, and one L3 VN (L3VNa) for the internal routing. A network firewall and gateway runs on a VM or server that connects to L3VNa and is used for inbound and outbound traffic processing. A load balancer (LB) is used in L2VNX. A VPN is also built between the gateway and enterprise router. An Enterprise customer runs Web/Mail/Voice applications on VMs within the vDC. The users at the Enterprise site access the applications running in the vDC via the VPN; Internet users access these applications via the gateway/firewall at the provider DC site.




Figure 2 - Virtual Data Center Abstraction View

The enterprise customer decides which applications should be accessible only via the intranet and which should be assessable via both the intranet and Internet, and configures the proper security policy and gateway function at the firewall/gateway. Furthermore, an

enterprise customer may want multi-zones in a vDC (See [section 4.2](#)) for the security and/or the ability to set different QoS levels for the different applications.

The vDC use case requires an NV03 solution to provide DC operators with an easy and quick way to create an NV03 network and NVEs for any vDC design, to allocate TSs and assign TSs to the corresponding NV03 network, and to illustrate vDC topology and manage/configure individual elements in the vDC in a secure way.

5. Summary

This document describes some general and potential NV03 use cases in DCs. The combination of these cases will give operators the flexibility and capability to design more sophisticated cases for various cloud applications.

DC services may vary, from infrastructure as a service (IaaS), to platform as a service (PaaS), to software as a service (SaaS). In these services, NV03 networks are just a portion of such services.

NV03 uses tunnel techniques to deliver NV03 traffic over DC physical infrastructure network. A tunnel encapsulation protocol is necessary. An NV03 tunnel may in turn be tunneled over other intermediate tunnels over the Internet or other WANs.

An NV03 network in a DC may be accessed by external users in a secure way. Many existing technologies can help achieve this.

6. Security Considerations

Security is a concern. DC operators need to provide a tenant with a secured virtual network, which means one tenant's traffic is isolated from other tenants' traffic as well as from underlay networks. DC operators also need to prevent against a tenant application attacking their underlay DC network; further, they need to protect against a tenant application attacking another tenant application via the DC infrastructure network. For example, a tenant application attempts to generate a large volume of traffic to overload the DC's underlying network. An NV03 solution has to address these issues.

7. IANA Considerations

This document does not request any action from IANA.

8. Informative References

- [IEEE802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges and Virtual Bridged Local Area", IEEE Std 802.1Q, 2011.
- [NV03HYVR2NVE] Li, Y., et al, "Hypervisor to NVE Control Plane Requirements", [draft-ietf-nvo3-hpvr2nve-cp-req-05](#), work in progress.
- [NV03ARCH] Black, D., et al, "An Architecture for Overlay Networks (NV03)", [draft-ietf-nvo3-arch-08](#), work in progress.
- [NV03MCAST] Ghanwani, A., Dunbar, L., et al, "A Framework for Multicast in Network Virtualization Overlays", [draft-ietf-nvo3-mcast-framework-05](#), work in progress.
- [RFC1035] Mockapetris, P., "DOMAIN NAMES - Implementation and Specification", [RFC1035](#), November 1987.
- [RFC3022] Srisuresh, P. and Egevang, K., "Traditional IP Network Address Translator (Traditional NAT)", [RFC3022](#), January 2001.
- [RFC4301] Kent, S., "Security Architecture for the Internet Protocol", [rfc4301](#), December 2005
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC7348] Mahalingam, M., Dutt, D., et al, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", [RFC7348](#) August 2014.
- [RFC7364] Narten, T., et al "Problem Statement: Overlays for Network Virtualization", [RFC7364](#), October 2014.
- [RFC7365] Lasserre, M., Motin, T., et al, "Framework for DC Network Virtualization", [RFC7365](#), October 2014.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A. and J. Uttaro, "BGP MPLS Based Ethernet VPN", [RFC7432](#), February 2015
- [RFC7637] Garg, P., and Wang, Y., "NVGRE: Network Virtualization using Generic Routing Encapsulation", [RFC7637](#), Sept. 2015.

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

Contributors

Vinay Bannai
PayPal
2211 N. First St,
San Jose, CA 95131
Phone: +1-408-967-7784
Email: vbannai@paypal.com

Ram Krishnan
Brocade Communications
San Jose, CA 95134
Phone: +1-408-406-7890
Email: ramk@brocade.com

Kieran Milne
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
Phone: +1-408-745-2000
Email: kmilne@juniper.net

Acknowledgements

Authors like to thank Sue Hares, Young Lee, David Black, Pedro Marques, Mike McBride, David McDysan, Randy Bush, Uma Chunduri, Eric Gray, David Allan, Joe Touch, Olufemi Komolafe, Matthew Bocci, and Alia Atlas for the review, comments, and suggestions.

Authors' Addresses

Lucy Yong
Huawei Technologies

Phone: +1-918-808-1918
Email: lucy.yong@huawei.com

Linda Dunbar
Huawei Technologies,

5340 Legacy Dr.
Plano, TX 75025 US

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

Mehmet Toy
Verizon

E-mail : mtoy054@yahoo.com

Aldrin Isaac
Juniper Networks
E-mail: aldrin.isaac@gmail.com

Vishwas Manral

Email: vishwas@ionosnetworks.com