

# Scalable Cloud Networking with Cisco Nexus 1000V Series Switches and VXLAN

## What You Will Learn

Many enterprise and service provider customers are building private or public clouds. Intrinsic to cloud computing is the presence of multiple tenants with numerous applications using the on-demand cloud infrastructure. Each of these tenants and applications needs to be logically isolated from the others, even at the network level. For example, a three-tier application can have multiple virtual machines in each tier and requires logically isolated networks between these tiers. Traditional network isolation techniques such as IEEE 802.1Q VLAN provide 4096 LAN segments (through a 12-bit VLAN identifier) and may not provide enough segments for large cloud deployments.

Cisco and a group of industry vendors, including VMware, Citrix, and Red Hat, are working together to address new requirements for scalable LAN segmentation and for transport of virtual machines across a broader network range. The underlying technology, referred to as Virtual Extensible LAN (VXLAN), defines a 24-bit LAN segment identifier that provides segmentation at cloud scale. In addition, VXLAN provides an architecture that customers can use to expand their cloud deployments with repeatable pods in different Layer 2 domains. VXLAN can also enable migration of virtual machines between servers across Layer 3 networks. With Cisco Nexus® 1000V Series Switches supporting VXLAN, customers can quickly, confidently, and securely deploy their applications in a multi-tenant cloud infrastructure.

## Cloud Computing Demands More Logical Networks

An infrastructure for a service cloud computing environment can have a large number of tenants, each with its own applications. In fact, each tenant requires a logical network isolated from all other tenants. Furthermore, each application from a tenant also requires its own logical network, to isolate it from other applications. To provide instant provisioning, cloud management tools, such as VMware vCloud Director, clone the application's virtual machines, including the virtual machines' network addresses, that demands a logical network for each instance of the application.

## Challenges of Existing Network Isolation Techniques


The VLAN has been the traditional mechanism for providing logical network isolation. Because of the ubiquity of the IEEE 802.1Q standard, numerous switches and tools are available that provide robust network troubleshooting and monitoring capabilities, enabling mission-critical applications to depend on the network. The IEEE 802.1Q standard specifies a 12-bit VLAN identifier, which limits the scalability of cloud networks beyond 4K VLANs. Some in the industry have proposed incorporation of a longer logical network identifier in a MAC-in-MAC or MAC in Generic Route Encapsulation (MAC-in-GRE) encapsulation as a way to expand scalability. Unfortunately, these techniques transport network packets inefficiently because they cannot make use of all the links in a PortChannel, which is typically implemented in data center networks.

## VXLAN Solution

VXLAN meets these challenges with a MAC in User Datagram Protocol (MAC-in-UDP) encapsulation technique and a 24-bit segment identifier in the form of a VXLAN ID (Figure 1). The larger VXLAN ID allows LAN segments to scale to 16 million in a cloud network. In addition, the UDP encapsulation allows each LAN segment to be extended across Layer 3 and helps ensure even distribution of traffic across PortChannel links (see Figure 2 later in this document).

VXLAN uses an IP multicast network to send broadcast, multicast, and unknown unicast flood frames. When a virtual machine joins a VXLAN segment, the server joins a multicast group. Broadcast traffic from the virtual machine is encapsulated and is sent using multicast to all the servers in the same multicast group. Subsequent unicast packets are encapsulated and unicast directly to the destination server without multicast. In effect, traditional switching takes place within each VXLAN segment.

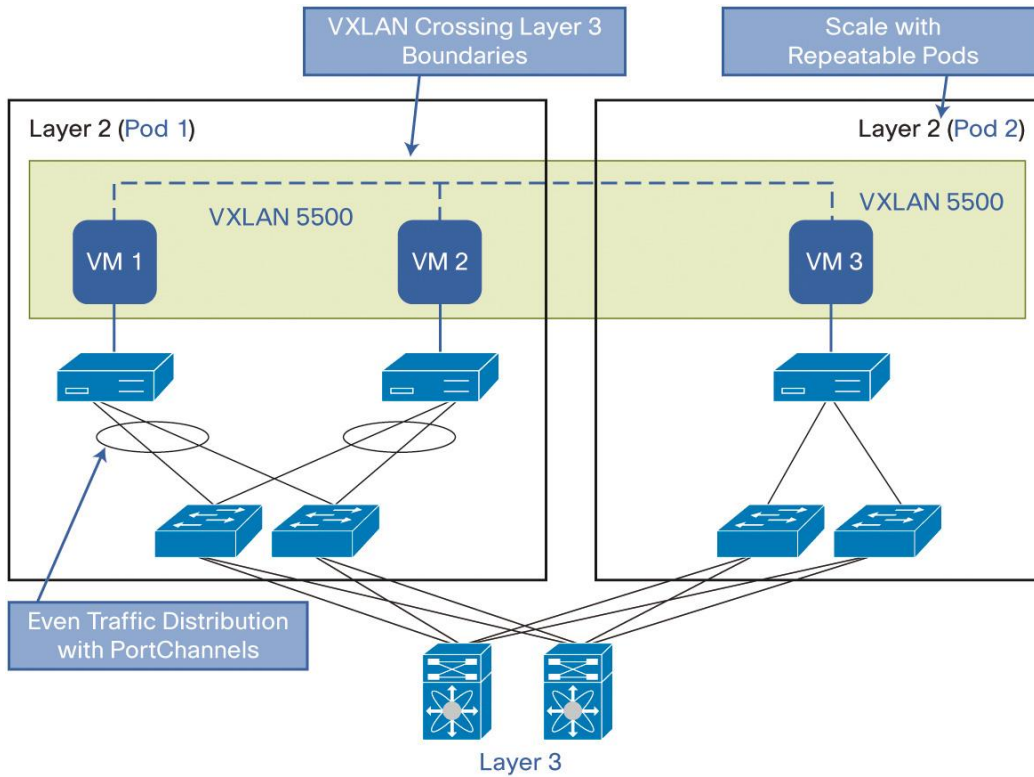
**Figure 1.** VXLAN Packet Format




The VXLAN solution provides these capabilities:

- Logical networks can be extended among virtual machines placed in different Layer 2 domains (Figure 2).
- Flexible, scalable cloud architecture enables addition of new server capacity over Layer 3 networks and accommodates elastic cloud workloads (Figure 2).
- If a virtual machine is connected only through VXLAN, then it can migrate across the Layer 3 network (gray virtual machine in Figure 3).
- If a virtual machine is connected to a VLAN (as in the case of the red virtual machine with VLAN and VXLAN connections in Figure 3), then it is restricted to migration within the Layer 2 domain. Note that a virtual machine on a VXLAN segment needs to be connected to a VLAN network in order to interact with external networks. To migrate such a virtual machine across Layer 3, you can use Overlay Transport Virtualization (OTV), discussed in more detail later in this document.

**Figure 2.** Segmentation with VXLAN




**Figure 3.** Migration Range with VLAN and VXLAN



Since VXLAN is a tunneling technique, the VXLAN gateway is required to send traffic to and from VXLAN to a traditional VLAN. In fact, for VXLAN traffic to use network services on physical devices, such as a physical firewall, the traffic needs to go through a VXLAN gateway, as shown in Figure 4. Cisco® ASA 1000V Cloud Firewall and VMware vShield Edge can all serve as VXLAN gateways.

**Figure 4.** VXLAN and Physical Network Services



### Cisco Nexus 1000V Series with VXLAN

The Cisco Nexus 1000V Series supports VXLAN and provides significant additional benefits for virtual machine network traffic in a VXLAN segment:

- Fully supports VMware vCloud Director 1.5
  - Cisco Nexus 1000V Series 1.5 [4.2(1)SV1(5.1)] is fully integrated into VMware vCloud Director, providing on-demand provisioning of the network.
- Supports tenant-specific networking policy, helping cloud service providers to differentiate their services
  - Cloud providers can provide different network policies for customized service-level agreements (SLAs).
- Supports advanced quality of service (QoS) for VXLAN
  - The Cisco Nexus 1000V Series provides Layer 3 QoS for VXLAN traffic with UDP encapsulation, helping ensure proper treatment of the packet on physical networks.
- Extends the existing operational model to the cloud
  - The Cisco Nexus 1000V Series offers a nondisruptive operational model for network and server administrators. With the Cisco Nexus 1000V Series supporting VXLAN, the existing operational model can now be extended to the cloud, accelerating cloud deployment.
- Supports Cisco vPath technology for virtualized network services
  - The Cisco Nexus 1000V Series supports Cisco virtual service data path (vPath) architecture, which supports a variety of virtualized network services, such as Cisco Virtual Security Gateway (VSG) and Virtual Wide Area Application Services (vWAAS). These virtualized network services can also be applied to traffic on VXLAN.
- Provides an XML API for customization and integration
  - The Cisco Nexus 1000V Series is based on Cisco NX-OS Software, which has a comprehensive XML API that allows customers to customize a solution and integrate with existing management tools.
- Supports VMware vSphere 4.1 and 5.0
  - Broader VMware vSphere support options.

---

## Working with OTV and LISP

VXLAN is intended for automated provisioning of logical networks in a cloud environment. OTV, using a frame format similar to that used by VXLAN, is a data center interconnect technology for extending Layer 2 domains across data centers over Layer 3. However, OTV has simpler deployment requirements than VXLAN since it does not mandate multicast-enabled transport network. OTV also contains faults within a data center and provide more robust data center interconnection. Applications in a VXLAN segment are often accessed through external networks based on VLANs, and hence OTV is required to extend such applications over Layer 3 so that both VXLAN and VLAN segments are extended.

Locator ID Separation Protocol (LISP) goes a step further by providing IP address mobility between data centers with dynamic routing updates, thus providing highly efficient routing networks for LAN segment traffic stretched across Layer 3 networks. Although VXLAN, OTV, and LISP frame formats share a similar-looking packet encapsulation structure, they serve very different networking purposes and are hence complementary to each other.

## Standardization of VXLAN

VXLAN has been submitted to IETF for standardization<sup>1</sup>, and Cisco, VMware, Citrix, and Red Hat have all jointly contributed to the standard. Hence, the networking and virtualization industries have come together to solve the scalability problem for cloud deployments with the same common standard.

## Conclusion

Cloud computing requires significantly more logical networks than traditional models. Traditional network isolation techniques such as the VLAN cannot scale adequately for the cloud. VXLAN resolves these challenges with a MAC-in-UDP approach and a 24-bit segment identifier. This solution enables scalable cloud network architecture with replicated server pods in different Layer 2 domains. Because of the UDP transport of VXLAN segments, virtual machines in a VXLAN segment can have their own LANs, but the traffic can cross Layer 3 boundaries. Cisco Nexus 1000V Series Switches with VXLAN support provide numerous advantages for customers, enabling customers to use LAN segments in a robust and customizable way without disrupting existing operational models. Cisco vPath technology enables virtualized network services to support virtual machines on VXLAN. In short, the Cisco Nexus 1000V Series with VXLAN helps ensure that customers can deploy mission-critical applications in the cloud with confidence.

## For More Information

- For more information about Cisco Nexus 1000V Series Switches, visit <http://www.cisco.com/go/1000v>.
- For more information about Cisco Virtual Security Gateway, visit <http://www.cisco.com/go/vsg>.
- For more information about Cisco Virtual Wide Area Application Services, visit <http://www.cisco.com/go/vwaas>.
- For more information about OTV, please visit <http://www.cisco.com/go/otv>.
- For more information about LISP, please visit <http://www.cisco.com/go/lisp>.
- For more information about VMware vCloud Director, visit <http://www.vmware.com/products/vcloud-director>.
- For more information about VMware vSphere, visit <http://www.vmware.com/go/vsphere>.

---

<sup>1</sup> VXLAN IETF submission: <http://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-00>




---

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)