Scalable Cloud Network with Cisco Nexus 1000V Series Switches and VXLAN

What You Will Learn

Many customers are building private or public clouds. Intrinsic to cloud computing is having multiple tenants with numerous applications using the cloud infrastructure. Each of these tenants and applications needs to be logically isolated from each other, even at the networking level. For example, a three-tier application can have multiple virtual machines requiring logically isolated networks between the virtual machines. Traditional network isolation techniques such as IEEE 802.1Q VLAN provide 4096 LAN segments (via a 12-bit VLAN identifier) and may not provide enough segments for large cloud deployments. Cisco and a group of industry vendors are working together to address new requirements of scalable LAN segmentation as well as transporting virtual machines across a broader diameter. The underlying technology, referred to as Virtual Extensible LAN (or VXLAN), defines a 24-bit LAN segment identifier to provide segmentation at cloud scale. In addition, VXLAN provides an architecture for customers to grow their cloud deployments with repeatable pods in different subnets. VXLAN can also enable virtual machines to be migrated between servers in different subnets. With Cisco Nexus® 1000V Series Switches supporting VXLAN, customers can quickly and confidently deploy their applications to the cloud.

Cloud Computing Demands More Logical Networks

Traditional servers have unique network addresses to help ensure proper communication. Network isolation techniques, such as VLANs, typically are used to isolate different logical parts of the network, such as a management VLAN, production VLAN, or DMZ VLAN.

In a cloud environment, each tenant requires a logical network isolated from all other tenants. Furthermore, each application from a tenant demands its own logical network, to isolate itself from other applications. To provide instant provisioning, cloud management tools, such as VMware vCloud Director, even duplicate the application’s virtual machines, including the virtual machines’ network addresses, with the result that a logical network is required for each instance of the application.

Challenges with 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, there are numerous switches and tools that provide robust network troubleshooting and monitoring capabilities, enabling mission-critical applications to depend on the network. Unfortunately, the IEEE 802.1Q standard specifies a 12-bit VLAN identifier, which hinders 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 scale. Unfortunately, these techniques cannot make use of all the links in a port channel, which is often found in the data center network or in some cases do not behave well with Network Address Translation (NAT). In addition, because of the encapsulation, monitoring capabilities are lost, preventing troubleshooting and monitoring. Hence, customers are no longer confident in deploying Tier 1 applications or applications requiring regulatory compliance in the cloud.

VXLAN Solution

VXLAN solves these challenges with a MAC in User Datagram Protocol (MAC-in-UDP) encapsulation technique. VXLAN uses a 24-bit segment identifier to scale (Figure 1). In addition, the UDP encapsulation enables the logical network to be extended to different subnets and helps ensure high utilization of port channel links (Figure 2). Instead of broadcasting a frame as in a case of unknown unicast, the UDP packet is multicasted to the set of servers that have virtual machines on the same segment. Within each segment, traditional switching takes place and can therefore provide a much larger number of logical networks.

Figure 1. Figure 1 VXLAN Format

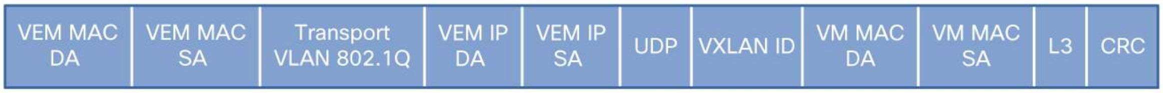

VEM MAC DA	VEM MAC SA	Transport VLAN 802.1Q	VEM IP DA	VEM IP SA	UDP	VXLAN ID	VM MAC DA	VM MAC SA	L3	CRC

As shown in Figure 2, the Cisco® VXLAN solution enables:

- Logical networks to be extended among virtual machines placed in different subnets
- Flexible, scalable cloud architecture in which new servers can be added in different subnets
- Migration of virtual machines between servers in different subnets

Figure 2. Figure 2 Scalability with VXLAN

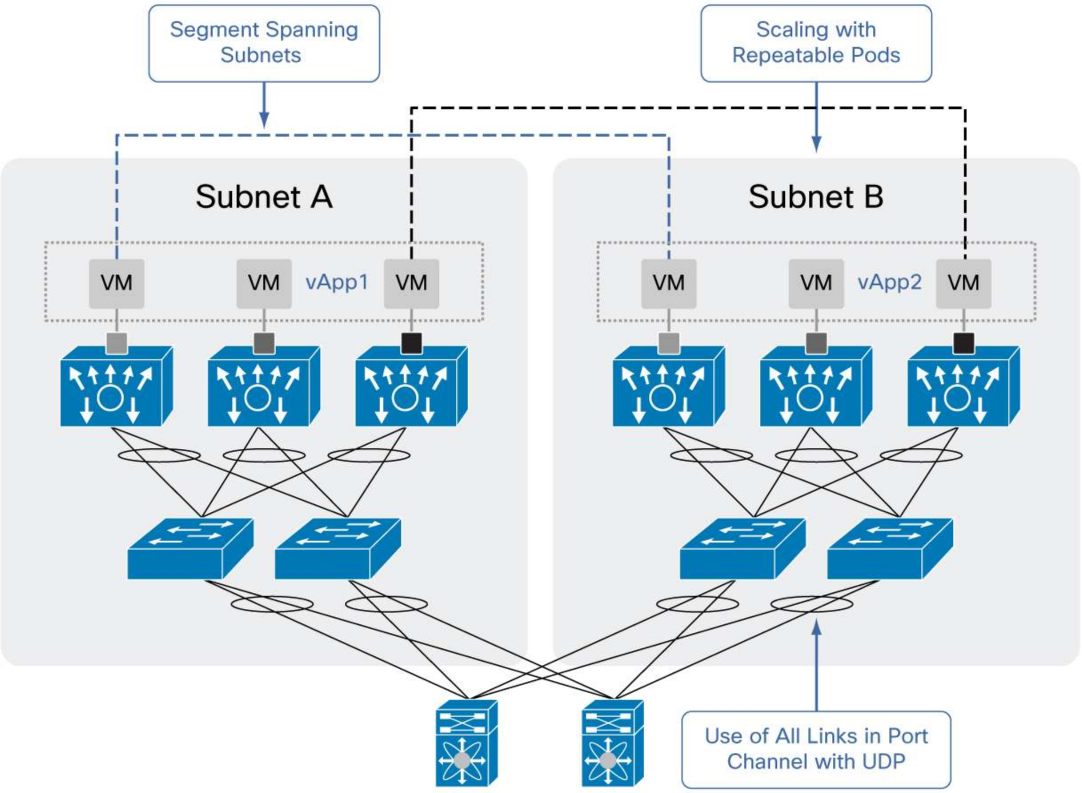
Segment Spanning Subnets
Scaling with Repeatable Pods
Subnet A
Subnet B
VM
VM
vApp1
VM
VM
VM
vApp2
VM
Use of All Links in Port Channel with UDP

Cisco Nexus 1000V Series with VXLAN

The Cisco Nexus 1000V Series supports VXLAN and provides significant benefits beyond VXLAN's baseline capabilities:

- Fully supports VMware vCloud Director 1.5
  - The Cisco Nexus 1000V Series version 1.5 [4.2(1)SV1(5)] is fully integrated into VMware vCloud Director, providing on-demand provisioning of the network.
- Supports VMware vSphere 4.1 and 5.0
  - The Cisco Nexus 1000V Series provides VXLAN capabilities in VMware vSphere 4.1 and 5.0.
- Extends existing operational model to the cloud
  - The Cisco Nexus 1000V Series offers a non-disruptive operational model for network and server administrators. With the Cisco Nexus 1000V Series supporting VXLAN, the same operational model can now be extended to the cloud without disrupting the existing operational model, accelerating cloud deployment.
- Supports Cisco vPath technology for virtualized network services
  - The Cisco Nexus 1000V Series supports Cisco virtual service datapath or vPath, which is an architecture that supports a variety of virtualized network services, such as Cisco Virtual Security Gateway (VSG) and Virtual Wide Area Application Services (vWAAS).
- Supports traditional monitoring tools for troubleshooting and regulatory compliance
  - The Cisco Nexus 1000V Series provides port statistics, Cisco NetFlow Version 9, and Encapsulated Remote Switched Port Analyzer (ERSPAN), enabling applications deployed on VXLAN to be monitored in the same way as on physical servers. Therefore, customers can confidently deploy mission-critical applications and those requiring regulatory compliance.
- Supports tenant-specific networking policy
  - The Cisco Nexus 1000V Series allows customers to provide a different network policy for each tenant, further enabling the cloud service provider to differentiate.
- Supports differentiated service-level agreements (SLAs)
  - The Cisco Nexus 1000V Series already supports advanced quality of service (QoS) for Layer 2 (IEEE 802.1p) and Layer 3 differentiated services code point (DSCP). With VXLAN's UDP encapsulation, DSCP marking can be used to provide a different SLA for each logical network. Hence, cloud service providers can provide differentiated service offerings.
- Provides 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.

Working with OTV and LISP

VXLAN is intended for creating more logical networks in a cloud environment. Overlay Transport Virtualization (OTV) while using similar frame format as VXLAN, is a data center interconnect technology extending Layer 2 domains to different data centers over Layer-3. Unlike VXLAN, OTV has simpler deployment requirements since it does not mandate multicast-enabled transport network. Locator ID Separation Protocol (LISP) goes a step further by providing IP address mobility between data centers with dynamic routing updates. While VXLAN, OTV, and LISP may share similar frame format, they serve very different networking purposes and are hence complimentary to each other.

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 a scalable cloud architecture with replicated server pods in different subnets. Because of the Layer 3 approach of UDP, virtual machine migration extends even to different subnets. Cisco Nexus 1000V Series switch with VXLAN support provides numerous advantages for customers, enabling customers to use LAN segments in a robust and customizable way without disrupting existing operational models. The unique capabilities of the Cisco Nexus 1000V Series with VXLAN help 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 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

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.**

Printed in USA C11-685115-00 09/11