Scaling Workload Mobility with the Cisco Unified Fabric

To help maximize compute resources, server virtualization technology has become a common strategy across many organizations for optimizing compute resources. Customers are now starting to take advantage of one of the major flexibilities provided by a virtualized environment: workload mobility. Since virtual machines are no longer tied to physical hardware, they can be moved to address several business and IT challenges.

Why Workload Mobility Is Important

Workload mobility helps improve application availability and the overall efficiency of data centers. By enabling workload mobility, you can move workloads to different locations within or across data centers to avoid outages associated by maintenance downtime or pending disruptive events such as hurricanes. Workload mobility is also a key tool for resource pooling and load balancing across data centers to maximize available compute resources.

Workload Mobility Network Challenges

How to move workloads within a VLAN or IP subnet is well understood today. However, since virtual machines retain their IP addresses when they move, moving workloads across a Layer 3 boundary (across a data center or across geographical dispersed data centers) requires innovative network technologies. that can extend VLANs (IP subnets) beyond layer 3 boundaries. Essentially, for workload mobility to be transparent, the workload has to move within the same IP subnet or VLAN.

To enable scalable workload mobility, Cisco® Unified Fabric provides several network innovations that extend VLANs within a data center and across data centers, allowing transparent workload mobility so that business objectives can be met.

The Workload Mobility Technologies in Cisco Unified Fabric

The Cisco Unified Fabric offers numerous innovative technologies that enable transparent workload mobility within and across data centers.

The Cisco Nexus® 1000V Series Switch is a software switch that integrates directly into the hypervisor. It provides a Virtual Ethernet Module (VEM) at the hypervisor level that acts like a physical Ethernet card. Each virtual machine (VM) is plugged into its own VEM, allowing for differentiated policy per VM. As VMs are moved, the VM policies (including the VLAN and quality of service (QoS) features, and security policy) automatically migrate with the VM. This level of VM policy "stickiness" is critical for transparent VM mobility.

Figure 1. Cisco Nexus 1000V Series Architecture

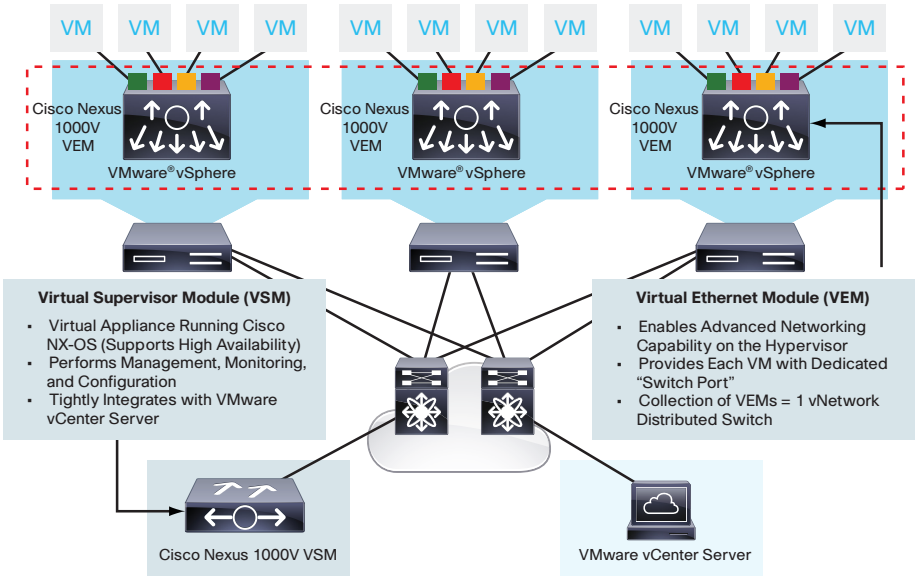

Cisco FabricPath provides a highly scalable solution to extend Layer 2 networks (VLANs) across data centers, a key requirement for workload mobility. Today, deploying Layer 2 domains is hampered by Spanning Tree Protocol scalability, which limits workload mobility to small Layer 2 domains.

Cisco FabricPath overcomes the limitations of Spanning Tree Protocol by combining the simplicity of Layer 2 Ethernet with the scalability and resiliency of Layer 3 routing. It eliminates the need for spanning tree, allowing you to build highly scalable Layer 2 Ethernet networks with multipathing capabilities. With Cisco FabricPath, you can build scalable network designs with more than 12,000 10 Gigabit Ethernet servers in single domain.

Virtual Extensible LAN (VXLAN) is another Cisco Unified Fabric technology that enables scalable workload mobility in the data center. VXLAN is an overlay technology that enables Layer 2 VLANs to extend over existing data center Layer 3 networks, making it possible to move workloads transparently across the data center beyond Layer 3 boundaries. **Cisco Overlay Transport Virtualization (OTV)** extends Layer 2 Ethernet VLANs between data centers over any network that supports Internet Protocol (IP). It allows you to extend VLAN domains between geographically distant data centers, enabling transparent workload mobility beyond data center walls.

Figure 2. OTV: Transparent Workload Mobility Across Any Data Center Anywhere

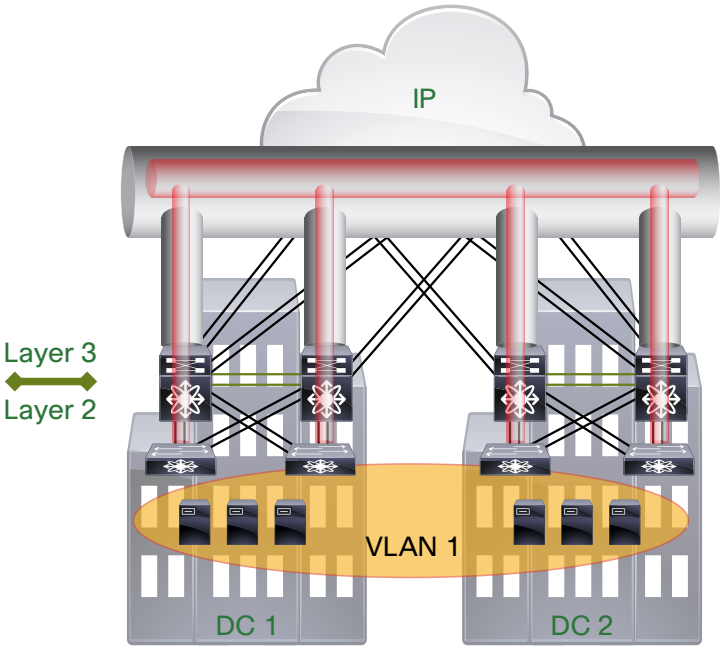

OTV is designed to be used as an overlay technology over existing networks, making it easy to deploy without requiring a network redesign.

Cisco Locator/ID Separation Protocol (LISP) is a new IP routing architecture, developed by Cisco and the Internet Engineering Task Force (IETF). LISP splits the device identity and its location into two different numbering spaces. This separation enables global IP address portability. For workload mobility, LISP means that a virtual machine can to be moved across data centers or across the country while preserving its IP address. By eliminating the need to renumber IP addresses, LISP allows customers to move entire workloads across Layer 3 boundaries to different IP subnets while still maintaining connectivity. With the deployment of LISP routing architecture, the network can truly deliver any application to any location–at any time and at any scale.

Why Cisco?

Cisco Unified Fabric offers the industry's most innovative features to enable workload mobility at any scale, within or across data centers. Cisco's innovative products and solutions, together with a broad range of service programs to accelerate customer success, are delivered through a combination of people, processes, tools, and partners, and result in high levels of customer satisfaction. Cisco continues to lead the industry with innovative solutions and technologies that are the product of years of networking experience.

C45-711050-00 06/12