Integrating SDN into the Data Center

Connecting islands of resources with the MX Series Universal SDN Gateway

Table of Contents

Executive Summary

The need to support increasingly more complex and virtualized applications with greater agility and at lower cost is driving a wave of change in data networking.

Over the years, applications have evolved from simply presenting data to being sources of timely information for business decision making. Applications have advanced from one per server, to virtualized multitier models on multiple servers, to highly distributed models on racks of servers, to natively virtualized, dynamic applications that can span entire data centers. At the same time, the user base has gone from a few specialists to many business managers to the entire organization, eventually including the supply chain and even customers using PCs and mobile devices from anywhere.

Network architectures must adapt to serve these new application models by becoming easier to deploy and manage. This requirement is prompting a transition from traditional networking to software-defined networking or SDN.

Introduction

As organizations plan to roll out new virtualized applications, they will need to consider how to transition to software-defined networking (SDN) and the impact it will have on the way they build their networks. As the number and type of applications that need to be managed continue to grow, there is an increasing need to connect disparate resources—some virtualized and some on legacy infrastructure, some local and some remote, and as SDN is adopted, there will be a need to connect to this environment.

Making the transition to SDN will require the ability to convert from one connection type to another and from one environment to another. Juniper Networks is building network infrastructure that will enable organizations to effectively make the transition to virtualized networks, while maintaining access to existing data and applications and preserving investments in network hardware infrastructure. This paper will examine the development of and challenges with application architectures, and it will show how Juniper Networks® MX Series 3D Universal Edge Routers can support new application deployment models, providing a bridge between new SDN and legacy infrastructure and resources.

Evolving Application Architectures

Applications have been advancing to better serve businesses with timely, critical information. As a result, these applications have become increasingly distributed and are accessed by more people throughout the organization and its supply chain from many more locations. This evolution is putting tremendous stress on the network, which needs to evolve to accommodate these demands. Enterprise organizations have many thousands of applications running in their data centers that were built over the years and present various challenges to the network.

In the early days of networking, everything was relatively simple. Applications ran on mainframes or minicomputers in the data center and were accessed from directly connected terminals, mainly to provide a few specialists with access to data records. With the arrival of client/server computing, organizations started deploying more specialized applications on servers attached to a LAN. Users accessed these applications locally with PCs or over leased lines from a branch office. There was only one application per server, but data sets and transaction volumes were growing. The demands on the network were more complex, as there was a need for switching and remote access.

The mid 1990s saw the emergence of service-oriented architecture (SOA), where applications contained more timely and critical business information used to drive the business. Employees in branch offices accessed these applications over the WAN or the Internet. The supply chain got connected so that inventories could be managed and correlated with sales. These applications were built in a three-tier model with a Web front end, a business logic tier, and a backend data store, all running on racks of servers. In this model, there was more traffic between tiers, but it flowed primarily in a north-south direction from the users to the application. The demands on the network grew with the greater need to connect over the WAN and the Internet and a greater need for network segmentation. The multitiered nature of the applications required additional network services such as server load balancers for the Web front end, a WAN optimization controller to increase throughput, and firewalls between the application tiers. Setting up the network became a major operation; provisioning was slow and cumbersome, as it was a manual process requiring every device to be touched and configured.

More recently, organizations have begun using cloud-based applications to run the business in near real time. These applications are distributed and modular, with racks of servers required for each business process such as inventory, sales transactions, and fulfillment. They create a vast amount of inter-application traffic and require hundreds or even thousands of network connections to stitch them together. Users everywhere—employees, customers, supply chain—connect to them using a variety of devices. The advent of cloud computing also saw the rise of process automation tools for Dev/Ops. Traditional networking, with its manual processes, continued to be a burden on the organization, slowing deployment and limiting business agility.

Since 2010, applications have become even more distributed, and might even be load-balanced across data centers. Everyone has access to them from a variety of devices and from just about anywhere. These applications are dynamic and can spawn new processes or move them on the fly, adding considerable complexity to network configurations. They are also being built on open source infrastructure, creating the need to integrate new environments.

At the same time, Big Data analytics—where data from multiple sources is streamed to applications in near real time and used for on-the-fly analysis to help drive the business—are also on the rise.

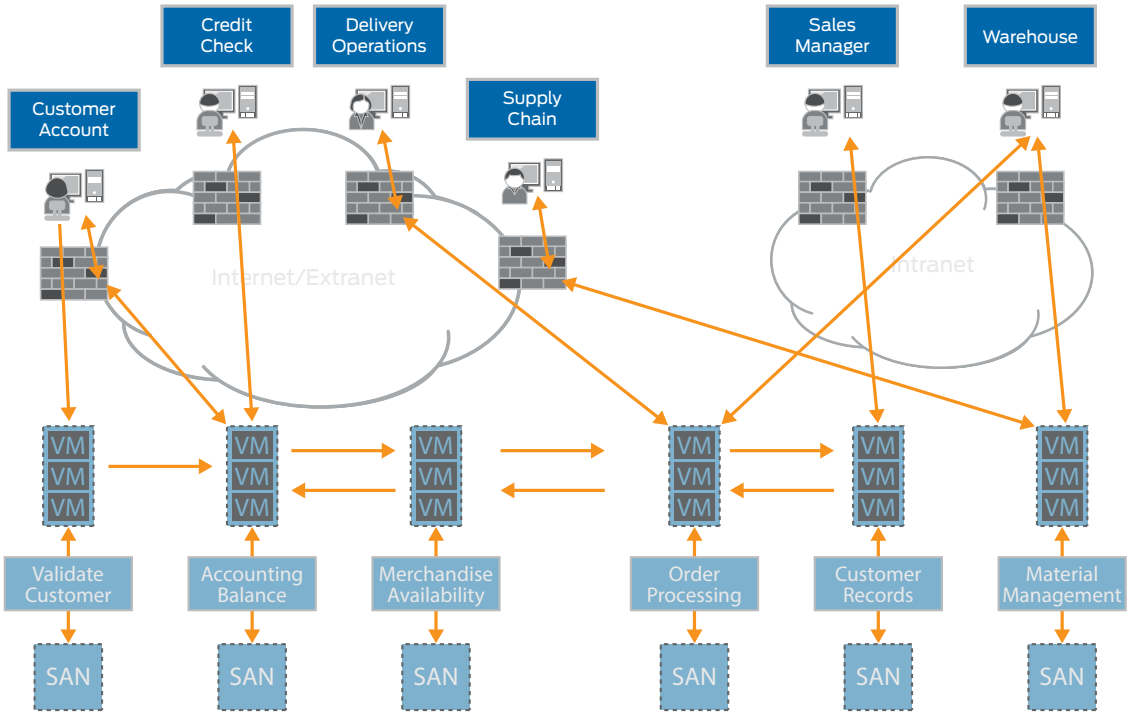

Figure 1: Typical distributed application architecture for a large enterprise

The Move to Virtualization

Virtualization has emerged as a way for IT infrastructures to deal with the increasing number of applications and their changing architectures. Server virtualization came first, developed to deal with issues such as resource utilization and application availability. While server virtualization solved a number of problems, it added a layer of complexity to the network. Abstracting servers from the hardware created a need to connect physical networks to virtual servers, resulting in the creation of virtual switches, virtual routers, and a host of other virtualized devices to serve in this environment.

The Impact of Server Virtualization

The move to server virtualization was led primarily by VMware; today, enterprise organizations have thousands of applications running on VMware infrastructure. For the most part, client/server applications can be virtualized. There are some limitations; for instance, some applications cannot change their IP addresses. Legacy data stores are usually not virtualized, so they need to be connected via bare metal servers, often over Layer 2.

Server virtualization has had an even greater impact on enterprise applications built in the SOA or three-tier model. Virtualized SOA applications require large numbers of connections between the virtual machines containing the Web tier and the business logic tier, the bare metal servers containing the data stores, and the physical network itself. There are also a number of network services devices that need to be connected. These include firewalls deployed in front of the various tiers. There may also be server load balancers in front of the Web tier. All of these devices must connect to the network and many of them exist as virtual appliances. To further complicate matters, the virtual machines (VMs) these applications run on might be moved around to balance workloads.

Newer web-based applications are highly distributed and have a number of processes that each run on their own banks of servers. These might be VMware environments, but they can also use a variety of hypervisors and are often designed to run in open source environments—for instance, they might run natively in Linux. They are intelligent and have dynamic behaviors, such as spinning-up instances in response to traffic loads. As a result, the network needs to respond to them

on the fly. They create the need for a huge mesh of interconnections to the various modules that might include multiple data stores of different types, a Web front end, and a communications layer. A high number of transactions will flow through the system, so network performance is a considerable issue. They also require access to the network services infrastructure, including server load balancers and security, and they might be in both physical and virtual formats.

The Need for Network Virtualization

The data center has reached a tipping point. Over the years, complexity has crept into the network in an attempt to accommodate the growth of applications; now, a solution is needed if real progress is to be made.

Organizations deploying new applications are looking to increase business agility and gain management insights. Rolling out new applications to realize this objective requires configuring the network device by device—an incredibly time-consuming task that increases administrative overhead and slows application deployment. Since these new applications touch the customer and enable commerce, speed of deployment means revenue to the organization. Since applications are natively virtualized and traffic flows mainly from VM to VM through virtual switches and routers, they can be connected and managed at the virtual network layer. As a result of this "software-defined network," systems have been developed and organizations are deploying them to manage the connections to their applications. These systems are managed via a central controller that uses tunneling protocols to connect to the various components of the application, including the virtual machines, virtual switches and routers, and virtualized versions of network services appliances such as load balancers, firewalls and intrusion detection/prevention gateways.

These software overlay networks promise to solve a number of issues. Some virtualized applications require Layer 2 adjacency to their storage resources and it is required for VM live migration. This is forcing broadcast domains to grow, pushing the limits of VLANs. Overlay networks alleviate the workload mobility problem by providing Layer 2 connectivity, independent of physical locality or underlying network design. They also have mechanisms to overcome VLAN limitations. Since they are connected to the applications when resources move, the virtual networks move with them. Virtual overlay networks promise to speed up provisioning and accelerate application rollout by providing agility and elasticity in the way network connections are deployed. However, the virtual overlay network isn't a panacea—there are other issues to consider, such as connecting to non-virtualized resources, connecting to remote resources, and connecting between disparate SDN resources.

Need for Universal SDN Gateway Services

As organizations virtualize their networks, they will find that they have islands of resources that need to be connected. This will be based on a number of factors: whether the resources are virtualized or bare metal; whether they are local or remote; and whether they are SDN connected. Connections must be made between these resources whether they are coming over Layer 2, Layer 3, or from any type of VPN, or from legacy systems or SDN environments. It will be necessary to bridge between environments, as they can have dissimilar control plane mechanisms or dissimilar tunnel types that need to be stitched together.

Islands of Legacy and Virtualized Resources

As organizations roll out new applications, they will need to connect them to the virtual network infrastructure. The new applications might be web-based, fully virtualized and highly distributed, even across data centers. Organizations will have older SOA-type applications with Web front ends, and they will have many legacy client/server or home-grown applications that cannot be virtualized. There will be islands of data used by older applications and a variety of data stores, some small and some in large clusters, all of which will need to be interconnected. The gateway should be able to connect from an SDN environment using a tunnel protocol, convert to native L2 or L3, and then connect to a hypervisor using a vSwitch, to a bare metal server that might be hosting a database, to a non-x86 compute platform, or to IP storage.

Islands of Resources by Location

A number of other resources need to be connected: branch office users, data centers, customers, and supply chain. The gateway must be able to connect to resources in an SDN environment in one location using a protocol like virtual private LAN service (VPLS) or MPLS Ethernet VPN (EVPN), or a Layer 3 protocol. It then needs to connect over the WAN to a remote location, and then convert back to an SDN-compatible tunnel type, or connect via Layer 2 or Layer 3 to a bare metal server or a hypervisor. The gateway should be able to connect from an SDN system via Layer 3 over the WAN to a remote location and then to another SDN system. It should be able to connect from an SDN tunnel protocol and then to a remote location via Layer 3 to another SDN system. It should also be able to connect from an SDN system and then over the WAN via generic routing encapsulation (GRE) or via an L3 VPN to a branch office. Finally, the gateway should be able to connect from an SDN environment and then via Layer 3 stitching to an integrated routing and bridging (IRB) interface and then to the Internet.

Islands of SDN-Enabled Applications

Organizations might want to use multiple SDN controllers from multiple vendors. Cloud service providers might use an open source controller such as OpenStack for one application, but maintain a different system for a customer. It is also possible they might need to migrate from one system to another, or connect to resources from one system to the other.

Organizations might operate multiple instances of a particular controller for various reasons, such as to exercise control where one system is for development and the other for production. They might want to put the systems on different networks to avoid risk in the event of a network failure. Significant operations are required to implement an SDN-to-SDN gateway where it is necessary to peer with multiple controllers. The gateway will need to connect from one SDN system to another that might have different control plane mechanisms. Or they may use different SDN controllers that can't exchange routes. For this to work, a universal gateway is required to perform route exchange via the control plane and make the necessary connections.

The MX Series Universal SDN Gateway Solution

Figure 2: Juniper Networks MX Series 3D Universal Edge Routers with Universal SDN Gateway capabilities

Juniper has developed a set of Universal SDN Gateway capabilities that run on the MX Series routers to connect all devices and resources in the data center and across the WAN. These gateway functions build upon the rich Layer 2 and Layer 3 VPN capabilities already in the MX Series routers, including standards-based protocols for Data Center Interconnect (DCI) such as EVPN, VPLS, and MPLS. There are four functions that the MX Series Universal SDN Gateway performs: L2 SDN Gateway, L3 Gateway, SDN-to-WAN Gateway, and SDN-to-SDN Gateway.

Types of SDN Gateway Services

Layer 2 SDN Gateway

The Layer 2 SDN Gateway provides SDN-to-non-SDN translation services for connections over Layer 2 on the same IP subnet, using bridging to maintain the same addressing. This functionality enables SDN controllers to communicate over Layer 2 to non-SDN VMs, bare metal servers, and L4-L7 network services which can be physical devices or running as virtualized services on x86 servers.

Layer 3 SDN Gateway

The Layer 3 SDN Gateway provides SDN-to-non-SDN translation services for resources on different IP subnets. This functionality enables the VM in the overlay environment to communicate over Layer 3 to legacy environments, to the Internet or other L3 destinations, to VMs in non-SDN environments, to bare metal servers, and to L4-L7 services.

SDN-to-WAN Gateway

The SDN-to-WAN gateway function provides SDN-to-WAN translation for devices that are on the same or different IP subnets, with the same or different encapsulation—for example, from Virtual Extensible LAN (VXLAN) to EVPN, MPLSoverGRE to EVPN, VXLAN to GRE, or VXLAN to an L3VPN. These functions are performed on the MX Series routers in Juniper's programmable silicon known as the Trio ASIC using multiple label operations on the same chip at the same time.

SDN-to-SDN Gateway

Connecting multiple hypervisors controlled by their own SDN environment requires a gateway for tunnel translation and control plane information sharing. The SDN-to-SDN Gateway translates between SDN controller types on the same or different IP subnets, which might be local or remote, with the same or different tunnel encapsulations such as VXLAN to VXLAN, VXLAN to MPLSoverGRE, or VXLAN to Network Virtualization using GRE (NVGRE). These functions are performed on Juniper's programmable Trio ASIC, which has the ability to peer and exchange routes with multiple SDN controllers simultaneously, avoiding the need to insert a second router or a virtual appliance into the architecture. SDN systems have their own control plane methods such as OVSBD or BGP; the MX Series device can take learned control plane information from OVSDB on one interface and BGP on another, then share the information, encapsulating the data plane information in the appropriate tunnel type based on the learning. Without this capability, many of these functions would have to be performed in software, possibly by virtual appliances running on x86 servers, at much lower throughputs.

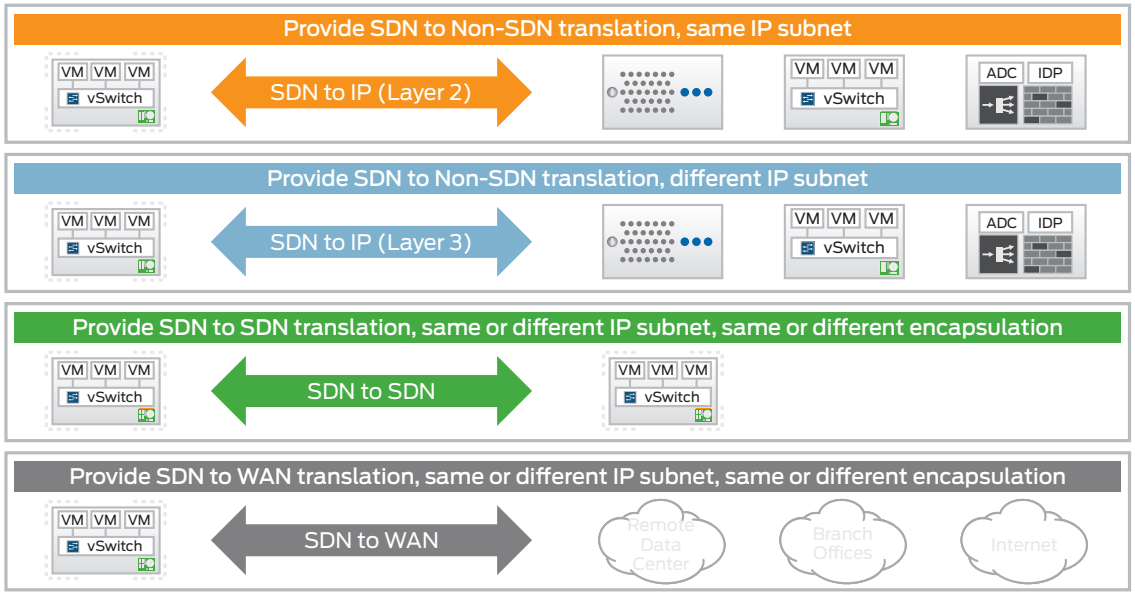

Figure 3: Four universal SDN gateway types: L2, L3, WAN and SDN

Additional Gateway Capabilities

This solution provides a high degree of multitenancy. For example, it is possible to create an architecture using a Layer 2 overlay that provides separation for each tenant, with one instance connecting to legacy resources and the other connecting to another overlay. The MX Series platform can, for instance, connect an SDN to an L2 virtual routing and forwarding (VRF) table, which can connect to a VLAN port or connect at the data center edge and stitch the connection into a tunnel. It can also bridge from a VXLAN to a VLAN and from other VPNs to a VLAN, using tunnel type tags for the separation. Tunnels such as VXLAN can be identified and mapped to a VRF; this feature supports VM mobility between non-overlay and overlay systems. If a tunnel conversion is not conducted, such as from VXLAN to an EPVN, then MPLS TAG lookup and conversion is required. This requires not only translating from a VXLAN header into a native IP header, but simultaneously doing a L3 lookup as well. Since Juniper uses programmable silicon in the MX Series, it can do simultaneous, high-performance lookups in hardware to solve this problem.

Use Cases for the MX Series Universal SDN Gateway

Enterprise Use Cases

Virtualized Applications

Many enterprises deploying virtualized applications to connect their employees, supply chain, and customers are using SDN overlay to get these applications up and running faster. These businesses will likely have other types of applications in their data center—home grown, client/server, SOA type—and will need to interconnect these disparate environments to share data or access particular application components, or they may be migrating applications from one environment to another. They can use the MX Series Universal SDN Gateway to convert from one connection type to another to enable access. The MX Series device can also be used to connect SDN environments to legacy infrastructure, such as IP storage or non-x86 compute platforms. It can also connect SDN environments, or non-SDN hypervisors to SDN environments.

Disaster Avoidance

Many organizations are seeking ways to avoid disaster-related outages, and to rapidly recover from such disasters should they be unavoidable. As a result, these organizations are implementing schemes for workload mobility from their primary data center to another active or backup data center. This may entail moving VMs from one data center to another over the WAN, and one way to do this is through live VM migration. However, in this scenario, the addressing scheme cannot change; this requires stretching L2 over the WAN. EVPNs, supported on the MX Series, offer a more robust way to accomplish the same objective.

Other organizations are using VXLANs to overcome VLAN limitations and need a way to extend them over the WAN. EVPNs are preferred over the WAN, but many top-of-rack devices cannot support direct termination of EVPNs. The Juniper solution solves this problem with EVPN-to-VXLAN stitching to provide the necessary L2 stretch. This solution is enhanced by the ability to do location-aware forwarding to avoid routing errors and the so-called “trombone effect.”

For Layer 3 connections, L3VPNs provide the established set of capabilities used for disaster recovery between data centers, and L2VPNs are fully supported on the MX Series. The benefits of EVPN include active/active multihoming, fast convergence, improved administrative and policy control, and an easier migration and interworking path to enable new services. The MX Series supports location-aware forwarding for long-distance VM mobility for the most efficient network path, and it provides support for replication of broadcast, unknown unicast, and multicast (BUM) traffic in hardware.

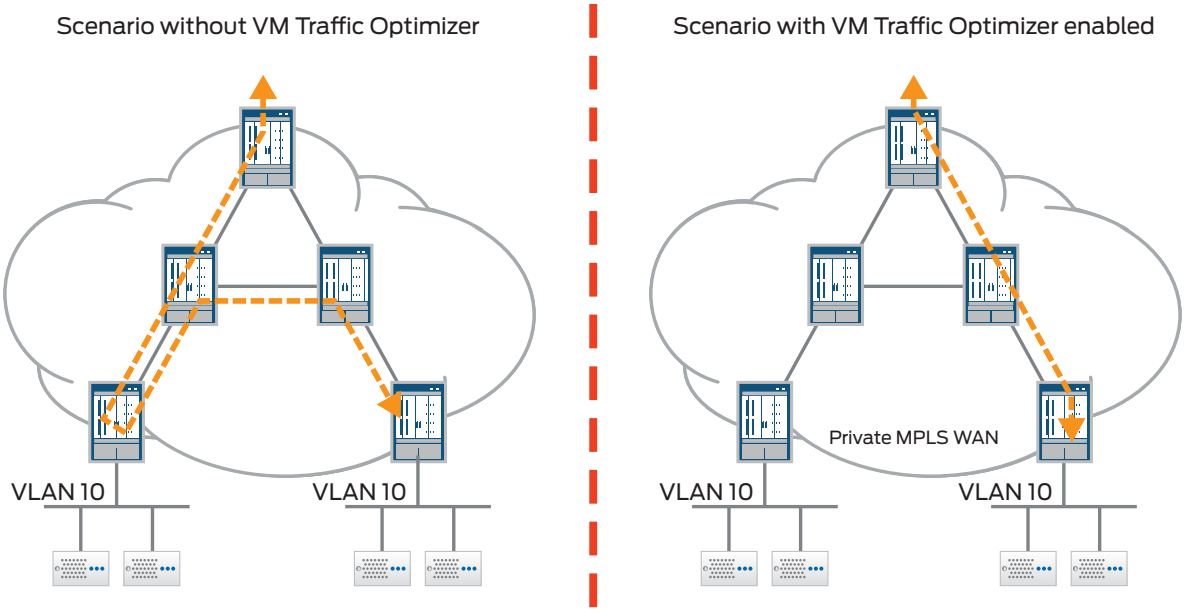

Figure 4: VM traffic optimizer ensures proper ingress and egress routing, eliminating the “trombone” effect

Remote VPN Services

The enterprise needs network connectivity from branch office locations to data centers. This connectivity can be a variety of IPVPN connection types such as IPsec or L3 IPVPNs. The MX Series offers all of the standard VPN services and is widely deployed for this purpose. With the gateway capabilities, VPN services on the MX Series now include the ability to connect branch offices to SDN environments in the data center via VXLAN and over the WAN via GRE to an MX Series device in a branch office or via an L3 VPN.

The MX Series can also provide more scalable connections from the branch office to the Internet. For example, connecting from an SDN controller via VXLAN and then via L3 stitching to the IRB to Internet, or from an SDN controller via L3VPNs to the Internet. Since many network operators only want to use L3 access, they will use IPsec or an IPVPN to connect sessions over IP over the WAN. The MX Series provides the stitching from the WAN VPN to the preferred LAN connection type. It is also possible to use the MX Series to combine connections at L3, using the same IRB and route for multiple connections. This ability can be used to combine LAN and WAN, L2 and L3 connections, overlays and VLANs in any combination. Connections can be made on the same box, at the same time, which is useful if you want to migrate from legacy protocols such as from VPLS to MPLS EVPN, where VPLS can be stitched to EVPN for DCI and VM migration.

Service Provider Use Cases

Data Center Interconnect

Organizations with multiple data centers often need to connect them for data replication or to access applications and resources. VPNs allow service providers to offer services to an enterprise by interconnecting their data centers for them. Juniper recently announced EVPNs to complement these services. The service provider can set up the data center interconnects for the enterprise's account and put them into a Layer 2 domain for VM mobility using EVPNs, or put them into an L3 domain using L3VPN.

This outsourced service provides the enterprise with an integrated and managed VPN solution that complements and interworks with VPLS and 2547bis service offerings at no added cost, technology, or operational risk. The benefits of EVPN as well as L3VPN include active/active multihoming, fast convergence, improved administrative and policy control, and an easier migration and interworking path to enable new services in the cloud and VM mobility optimizations. Since these data centers will contain a variety of resources and connection types, the SDN Gateway services can be used to handle the conversions, whatever they may be. It is important to base these services on open standards for interoperability between devices and networks.

Hybrid Cloud/Private Cloud

Enterprise organizations are looking to optimize their compute resources to ensure application performance and deliver a superior user experience. They want to build out their data center capacity only to support average workloads and use capacity from cloud service providers for the overflow. Data center operators offer hybrid cloud services where they provide hosted infrastructure to the enterprise, which the enterprise can use to host applications, spin up virtual machines, and move virtual machines as needed to support workloads that might need increased capacity. The enterprise may also have confidential and mission-critical workloads that it doesn't want to risk putting on the public cloud and prefers to host in the private cloud infrastructure. The service provider can connect the enterprise to their cloud hosting centers using L3VPN services on the MX Series to enable workload mobility using L2 stretch capabilities provided by the VxLAN-to-EVPN stitching capability on the MX Series. The MX Series Universal SDN Gateway can extend a customer's L2 or L3 network in a seamless fashion to the cloud infrastructure by manipulating the tunnel tags, preserving their extended L2 or L3 space. If there is an over-lapping tag for a VLAN, for example, it can be modified so that the extension is transparent to the tenant.

Application Hosting

Data center operators hosting applications for the enterprise or offering Software as a Service (SaaS) applications can use MX Series Universal SDN Gateway capabilities to connect branch office locations over the WAN using a variety of VPN services. They can also use the MX Series devices to span SaaS instances or connect to storage resources. If the service providers are managing their hosted or SaaS application using SDN systems, they can use the MX Series to bridge between instances of SDN systems to allow access to resources between applications or to migrate from one environment to the other. It is also possible to connect an enterprise application that is managed by an SDN to various backend resources via L2 or L3 VPNs or native IP, and to connect from virtualized resources to bare metal resources.

Big Data as a Service

With the proliferation of data from all sorts of sources, such as order management systems, online transaction systems, point of sales systems, mobile devices and social media, there is a wealth of information that needs to be analyzed and acted upon. Much of this information comes in near real time and can be used to steer sales and marketing efforts or to guide inventory management and just-in-time production systems. Large enterprise organizations are setting up systems to manage Big Data analytics. Some cloud service providers are hosting this infrastructure for the enterprise. This means that many data stores need to be hosted and often need to be replicated over the WAN from one data center to another.

This in turn requires various methods of access. Some of the applications are managed in an SDN environment; these organizations can make use of the VPN capabilities for DCI, SDN-to-SDN conversions, and L2 and L3 access methods for data stores and high-speed networking to move data quickly and enable rapid processing for analytics.

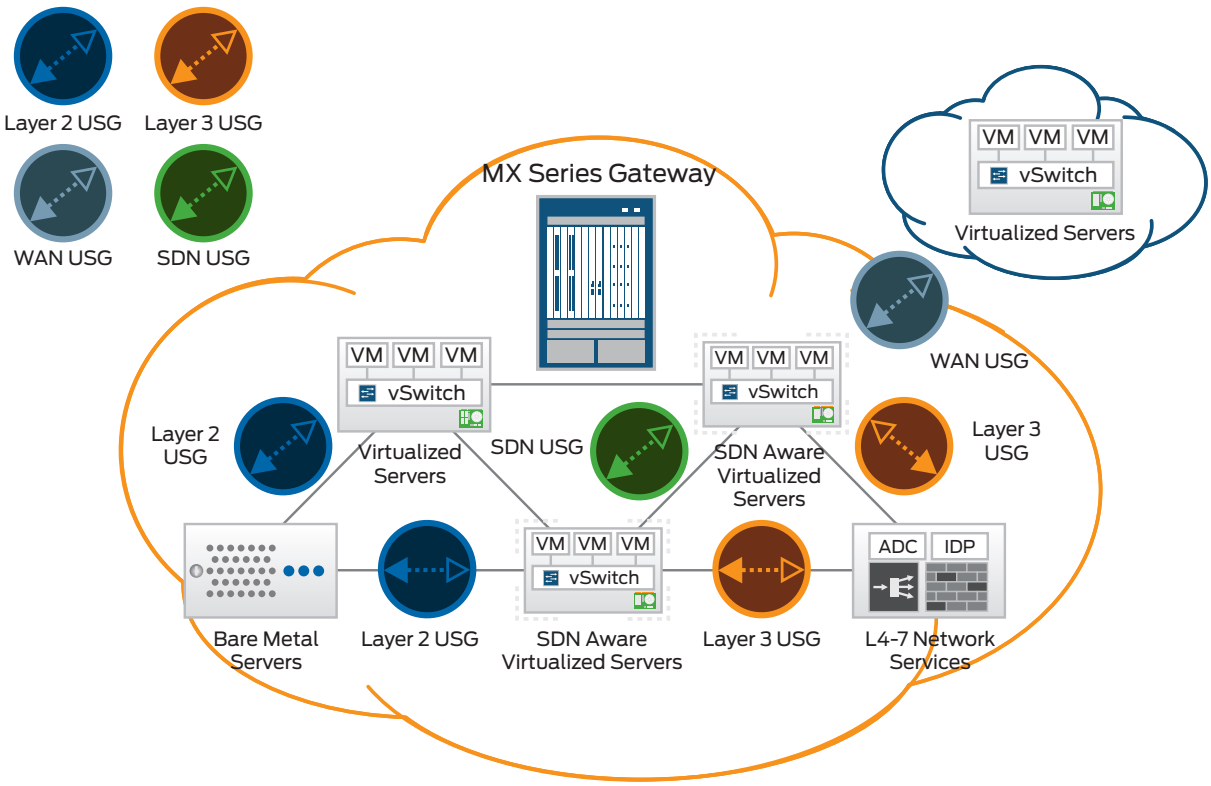

Figure 5: Universal Service Gateways in the network

Conclusion

Organizations investigating overay networks today need to think about their network infrastructure and how it will serve them in the future. Investing in the right network architecture will be critical to their ability to move to SDN while supporting their current environment.

As businesses continue to deploy new applications and move to virtualized environments, they will need to connect them to legacy applications and data stores. Businesses should base their decisions on network architectures that deliver the features and capabilities they need now as well as in the future. They will want a network architecture that's easy to integrate with SDN controllers. They will also want open standards and interoperability between their environments and network devices. They will want to take advantage of automation tools to increase network agility and the rest of their IT infrastructure. Finally, they will want to integrate with their choice of hypervisor and tie the network to it. To accomplish this, businesses will need to work with a network vendor that understands how to deliver agility and won't lock them in to a particular infrastructure.

With robust support for IP networking, a complete set of L2 and L3 VPN services and integration with SDN controllers, the Juniper Networks MX Series Router with Universal SDN Gateway capabilities provides a complete solution for transitioning networks to SDN, as well as the business agility that brings. This is accomplished while maintaining interoperability with existing resources and protecting investment in existing network infrastructure.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701

2000542-001-EN Sept 2015