# Evolution of Software Defined Networking within Cisco's VMDC

Software-Defined Networking (SDN) has the capability to revolutionize the current data center architecture and its associated networking model. This new paradigm, with its promised benefits, has the potential to create an inflection point in deploying cloud services. This paper summarizes the key characteristics of SDN as it is applied to data center virtualization, and illustrates how Cisco's Virtualized Multiservice Data Center (VMDC) solutions leverage many of these concepts today, to solve real-world customer problems.

## Challenges within the Data Center

Businesses are increasingly under pressure to respond to the ever-increasing demand from end-users and employees, who demand more from computer systems, networks, and mobile devices than ever before. As a result, service providers and enterprises are constantly exploring ways to keep up with fast evolving technology trends, business and end-user requirements, and to provide innovative applications and services with faster time to market. "Business Agility" is the watchword in this new world where providers are expected to provision and roll out services rapidly.

Software Defined Networking provides a new paradigm that attempts to respond to the new requirements of business agility and improved user experience. Many cloud-computing environments operate in an application-centric world, where virtualized applications are hosted within a public or private cloud. As a result, users can access their applications from anywhere, on any device, at any time. Users have access to more applications than ever before (on smartphones, tablets, etc.), and the user-experience of many of these applications has a dependency on the quality of the network.

The Open Network Foundation (ONF) defines Software-Defined Networking as follows:

> "The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."

Figure 1 shows layers of the SDN architecture.[1]

1. https://www.opennetworking.org/sdn-resources/sdn-definition

*Figure 1 SDN Architectural Layers*

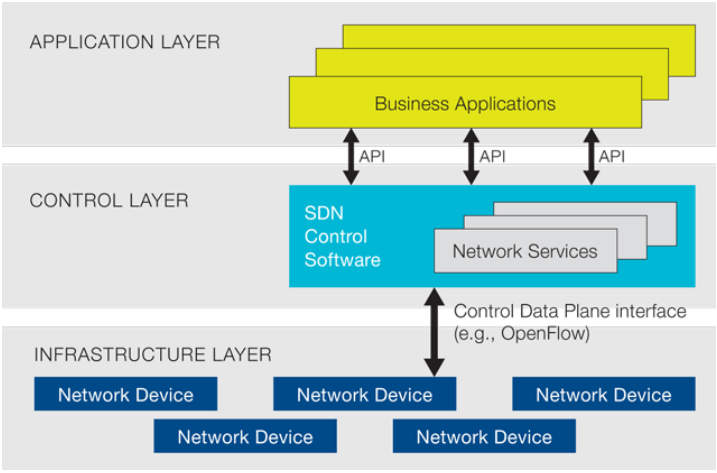


Within an SDN infrastructure, applications can request and obtain services from the underlying network infrastructure. This capability leads to the development of more proactive and dynamic applications that improve the user experience. SDN changes the way networks are designed and deployed, where the applications have more control on the configuration of the network infrastructure. SDN offers businesses the chance to build networks with increased application awareness and intelligence about Layer 4 - Layer 7 protocol attributes and delivery requirements.

Software-defined networking allows infrastructure become much more automated and therefore adaptive to the needs of the applications performing (or requesting) the automation.

# SDN Architectural Framework and Solution Characteristics

Although centralization of control is a key tenet of SDN, there are a number of other characteristics that vary with different SDN solutions. The following are certain common characteristics that are deployed in most SDN solutions today.

## Centralized Control

In conventional networks, control traffic and data traffic are tightly coupled in network devices. Additionally, in conventional networks most control functions are distributed over many devices. In contrast, the SDN paradigm attempts to pull control functions out of the network devices and consolidate them into a centralized location. With this model, once a centralized controller derives the desired forwarding behavior, forwarding instructions for packets are downloaded to the appropriate network devices. The communication between the controller and the network devices can use some form of standardized protocol such as OpenFlow to facilitate standardized network device programming.

Not all control functions can or should be centralized. Most SDN solutions still rely upon under-lying network connectivity that employs some form of a distributed routing control mechanism. Legacy, non-packet-switching networks deployed a mostly centralized control functionality that created a number of problems, primarily lack of scale, which provided the impetus for the fast adoption of a decentralized control mechanism found in today's packet switching networks. Most SDN solutions still require some sort of decentralized control functionality, and the degree of control plane centralization varies from one solution to another.

# Overlay Networks

The use of overlay networking technologies is another common characteristic to a number of SDN architectures. Overlay networks, provide a construct for the creation of logical networks that can be leveraged by edge devices and applications. Overlay tunneling technologies such as VXLAN, enable the creation of logical networks on top of the existing physical network without having to explicitly involve the underlying physical network. Some of the benefits of overlay networks are as follows:

- They can provide logical layer-2 adjacency without the need to create physically adjacent layer-2 networks. This is particularly useful for provisioning multi-datacenter environments where logical layer-2 connectivity is needed across layer-3 boundaries.
- Some tunneling technologies provide much larger numbers of layer-2 networks than VLANs, which generally are limited to 4000 segments.
- Faster and potentially simpler network provisioning and orchestration since interacting with the physical network is not required.

# SDN Solution Taxonomy

Within any networking solution, one can classify network characteristics within the following broad categories:

## Control Plane Function

In its simplest form, the control plane provides layer-2 MAC reachability and layer-3 routing information to network devices that require this information to make packet forwarding decisions. In the case of firewalls, the control plane would include stateful flow information for inspection. Control plane functionality can implemented as follows:

- Distributed - Conventional routers and switches operate using distributed protocols for control, i.e. where each device makes its own decisions about what to do, and communicate relevant information to other devices for input into their decision making process. For example, the Spanning Tree Protocol (STP), Fabric Path, and routing protocols such as IS-IS and BGP provide distributed control of packet forwarding functionality to networking devices.
- Centralized - In this case, a centralized controller provides the necessary information for a network element to make a decision. For example, these controller(s) instruct networking devices on where to forward packets by explicitly programming their MAC and FIBs.

The control plane functionality can be further classified as follows:

#### Layer-2 Reachability Control

This control mechanism provides Layer 2 MAC reachability information. It can either be implemented in a distributed manner like bridging and data-plane learning, or in a centralized manner with a controller-based device.

#### Layer-3 Reachability Control

The Layer 3 control mechanism provides Layer 3 routing and reachability information to all participating devices. Conventional routing protocols are distributed, while an SDN based system typically involves downloading controller-derived Layer 3 forwarding tables to various devices, using standardized or open protocols.

## Data Plane / Control Plane Collocation

The main function of network devices and appliances is to forward user-generated data traffic within the network infrastructure; the particular forwarding policies are dependent upon the type of device. Such network elements can be one of the following:

- **Collocated**—These are devices that use distributed control planes and which have control plan and data plane functions that are collocated, i.e., no external entity is required for the device to make decisions. These appliances can be physical or virtual. All of Cisco's physical devices, the Adaptive Security Appliance (ASA) 1000v, and Cloud Services Router (CSR) 1000v are examples of devices with collocated control plane + data plane.[1]
- **Dislocated**—The functionality of the device is distributed across multiple elements, under the control of a centralized element, i.e. the data plane and control plane of the device are dislocated. The functionality of the device is dependent on instructions coming from the centralized element. The OpenFlow enabled Cisco 3750-X and 3650-X devices, the Nexus 1000v, and the newly developed Virtual Provider Edge (vPE) solution are examples of devices with dislocated control plane and data plane. Devices that use distributed control planes may have dislocated control plane and data plane functions; devices with centralized controllers implicitly have dislocated control plane and data plane functions.

#### Services

Services such as load-balancers or firewalls can be implemented with either autonomous or dependent forwarding decision making capabilities. Examples would include a virtual-autonomous appliance like the ASA 1000v, or a virtual-dependent appliance such as an Open Virtual Switch (OVS) with a centralized firewall controller. Autonomous stateful service appliances inspect and maintain state machines for traffic flows at each device, where as dependent service appliances employ a centralized control device to externally control the service behavior. The complexity of a centralized services controller is considerable since it needs to process, store, and distribute a large number of traffic flow states associated with stateful inspection of Layer 4 - Layer 7 traffic.

Further, some service control functions need to be in the data path, e.g. an application server load balancer, which is monitoring the responsiveness of applications servers. Given this complexity, it is difficult to scale dependent services, and therefore autonomous services will still be required for functions such as stateful firewalls and application server load balancers for the foreseeable future.

#### Overlay Networks

Tenant segmentation can be provided by conventional means such as VLANs, or an overlay method such as VXLANs. As outlined previously, an overlay network within a SDN environment is a construct for the creation of logical networks that can be leveraged by edge devices and appliances. Tenant segmentation based on VLANs is normally considered a characteristic of a non-SDN systems, while overlay networks are considered a key component of an SDN-based solution.

The adoption of the above mentioned choices for L2/L3 reachability control, services, overlay networks, and data plane characteristics within a solution determines where the solution falls within the range between conventional and SDN-based solutions.

# Cisco's Virtual Multiservice Data Center

The Cisco Virtual Multiservice Data Center (VMDC) reference architecture provides a framework for building a scalable and resilient data center infrastructure. Cisco's VMDC validated designs provide guidelines that demonstrate how customers can integrate Cisco and partner technologies into a data

1. http://www.openflow.org/

center platform that supports virtualization and can lead to significant hardware consolidation. There are several variants of VMDC, each offering certain features and functionality suitable to a variety of customer needs. The following sections describe how the various versions of the VMDC architecture fall within the spectrum of conventional and SDN-based solutions.

# VMDC 2.x and 3.x

The VMDC 2.x and 3.x series of releases is built on conventional, hierarchal-based data center designs. The layers in these releases are Core (optional), Aggregation, Services, Access, and Virtual Access. Compute pods hang off the access layer, while virtual machines connect to the virtual access layer. The Core layer connects to the WAN-PE devices on the operator's IP-NGN network. Figure 2 shows the physical topology of the VMDC 2.2 release.

***Figure 2 VMDC 2.2 Physical Topology***

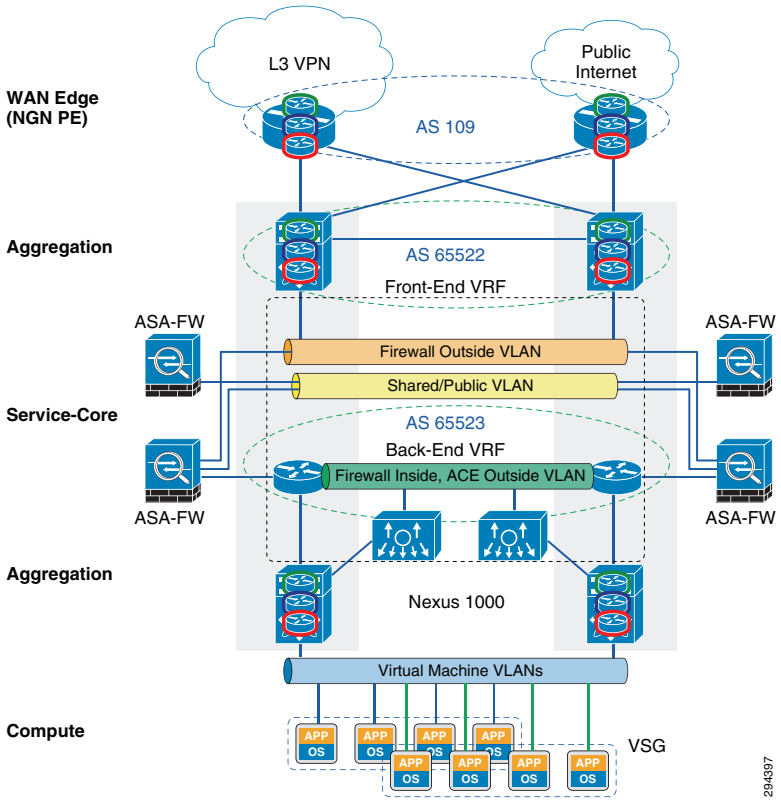


With regards to the solution's placement on the SDN spectrum, the VMDC 2.x and 3.x releases employ several relevant technologies. First, these releases leverage distributed control plane functions for Layer 3 routing, specifically BPG and OSPF. For Layer 2 bridging, the solution relies on ARP and STP, while VMDC 3.x relies on Cisco's Fabric Path technology. Both solutions rely on VLANs for tenant segmentation. These technologies are more typical of a conventional network rather than an SDN-based solution. At the virtual access layer, the solutions leverage the Nexus 1000v, which has a centralized control element (the Virtual Supervisor Module (VSM)) programming the forwarding behavior of the distributed data-plane elements (the Virtual Ethernet Modules (VEMs)). Likewise, one of the services leveraged by these releases is the Virtual Service Gateway (VSG), which has a centralized component

performing initial flow inspection via Cisco's vPath technology, and subsequently programs the VEM's action to apply to the flow. These technologies are well aligned with the SDN concepts discussed previously. Figure 3 shows the VMDC 2.x series on the conventional/SDN solution spectrum.

***Figure 3 VMDC 2.x Spectrum***

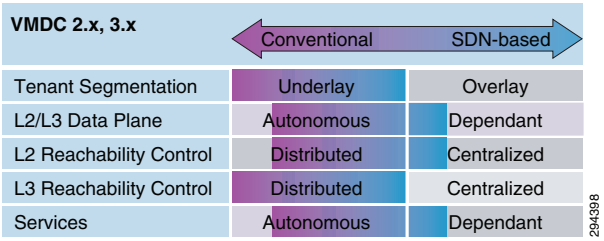


As shown in Figure 3, while the VMDC 2.x series has some relevant SDN technologies, its foundational concepts are based on conventional architectures.

# VMDC 4.x (vCE)

The VMDC 4.x series of releases focuses on Cisco's Virtual Service Architecture (VSA), where all network services are virtual, and each tenant network container leverages the Cloud Services Router (CSR) 1000v for container ingress and egress traffic. The CSR 1000v serves a number of functions, including a zone-based perimeter firewall, VPN remote access termination, and the virtual Consumer Edge (vCE) device; providing L3 routing for the container and seamlessly extending IP-NGN customers' reach into the data center. Figure 4 shows the representative physical topology of the VMDC 4.0 release.

***Figure 4 Physical Topology***

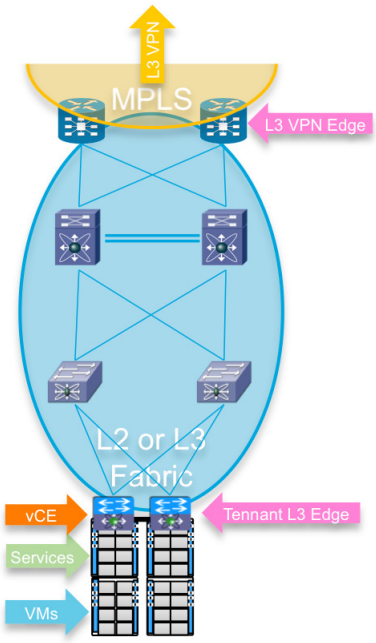

Relative to the VMDC 2.x series of releases, the VMDC 4.x series adds a significant SDN concept to the architecture by deploying VXLAN-based overlay networks for tenant segmentation, while continuing to leverage distributed control functions for switching (STP or Fabric Path) and routing (BGP). It also adds more vPath enabled virtual services like the virtual Wide Area Application Services (vWAAS) appliance and the Adaptive Security Appliance (ASA) 1000v.

Figure 5 shows the VMDC 4.x series on the conventional/SDN solution spectrum.

***Figure 5 VMDC 4.x Spectrum***

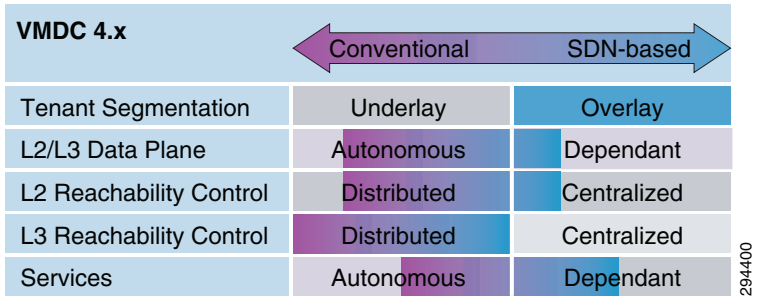


# Cisco vPE

Cisco's Virtual Provider Edge (vPE) solution leverages all of the SDN concepts [1] described earlier in this document. Figure 6 shows the high-level vPE architecture.

***Figure 6 vPE High-level Architecture***

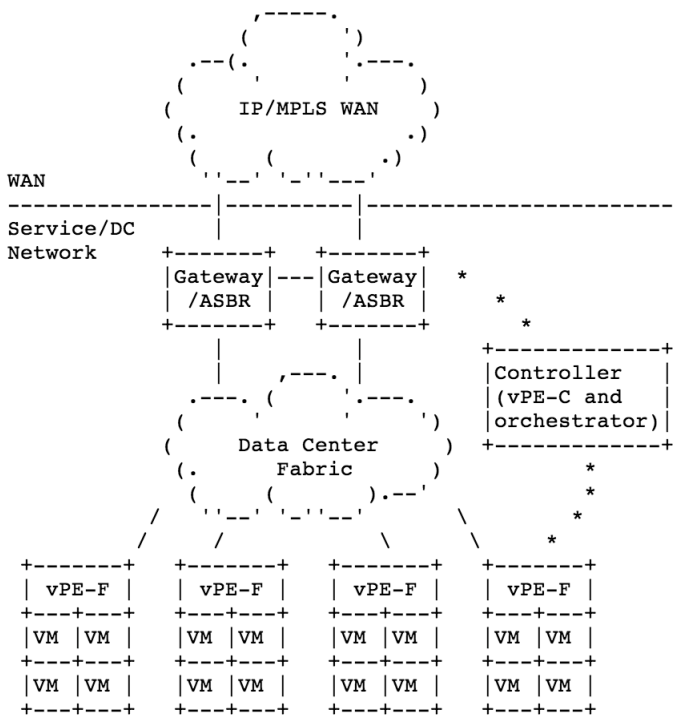


1. http://tools.ietf.org/html/draft-fang-l3vpn-virtual-pe-03

The vPE controller and the vPE forwarder represent the key tenet of SDN, i.e., the decoupling of the control-plane and data-plane. The vPE controller computes the appropriate forwarding tables for both Layer 2 and Layer 3, and subsequently programs these tables into the vPE forwarders that reside in the hypervisor hosts. Packet forwarding and service chaining are accomplished using the tables programmed by the controller, while utilizing an encapsulation, e.g., VXLAN or MPLS to create an overlay network through the data center. The vPE solution can be used over many different physical topologies, as the majority of network functions occur in the software based controllers, forwarders, and services. Similar to the vCE architecture, all the services for the vPE solution are virtual, with some being autonomous (e.g., Citrix VPX) and others dependent (VSG, ASA 1000v, vWAAS). Figure 7 shows the vPE solution on the conventional/SDN solution spectrum.

***Figure 7 vPE Spectrum***

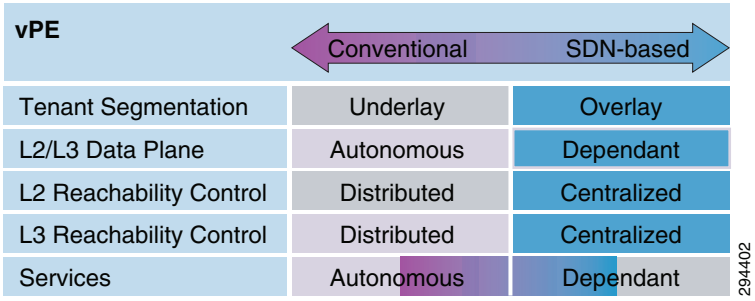


As shown in Figure 7, Cisco's vPE solution includes at least one piece of functionality from all the main SDN concepts discussed in this document. At the time of this writing, Cisco's vPE solution is scheduled to become generally available in Q1CY2014, and will become an official VMDC release shortly thereafter.

# Conclusion

Cisco sees SDN as an opportunity for the company and a benefit for customers. SDN is nothing new to Cisco-we have been delivering SDN-related technologies for some time. For example, Cisco's Nexus 1000V software switch, in production since 2009, uses separated control-data plane architecture and is currently licensed by more than 5000 customers. Cisco has also long-provided open programmatic interfaces to our operating systems to enable scalable application integration to underlying network infrastructures and access to third-party management and orchestration tools. As described in this document, the VMDC Cisco Validated Designs are proof points of Cisco's utilization of SDN concepts in data center infrastructure that exposes Cisco's best of breed hardware to operators, applications, and ultimately, end-users. Cisco is continuously evolving its technology portfolio to evolve with our customer's requirements. In some cases, those requirements will benefit from SDN technologies, and in other cases, conventional mechanisms will prove optimum. In either case, Cisco has solutions to meet those needs.

### Axel Nadimi

**Alex Nadimi, Solutions Architect, Systems Development Unit (SDU), Cisco Systems**

Alex has been with Cisco for the past 15 years and is currently working as a Solutions Architect in Cisco's Systems Development Unit. Prior to this role, he worked as a Technical Marketing Engineer in the Cisco Central Marketing Organization. He has developed solutions and technical guidance on various technologies such as security, VPN networks, WAN transport technologies, data center solutions, and virtualization. Prior to Cisco, he has worked at Hughes LAN Systems and Northern Telecom. He holds a masters of science in electrical engineering from Louisiana State University

### Brian Davis

**Brian Davis, Principal Engineer, Systems Development Unit (SDU), Cisco Systems**

Brian is a 15-year Cisco veteran, specializing in Service Provider solution architectures for the duration of his tenure. Previous to his current assignment, he specialized in DOCSIS High Speed Data, voice over IP (VoIP), and 3-Screen video architectures and deployments. Brian's latest focus in on network orchestration of Cisco's Virtual Multiservice Data Center solution, which leverages both conventional and SDN concepts and technologies. Brian graduated from Rensselaer Polytechnic Institute with a BS in Electrical and Computer Systems Engineering in 1997. When he's not thinking about networking, Brian enjoys spending time at home with his wife, 7 and 5 year old daughters, and newborn son. His latest hobby is maintaining his lawn and flowerbeds to be the envy of the neighborhood.

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.