

SD-WAN: how the control of the network can be shifted from core to edge

^{1,2}Sebastian Troia, ¹Ligia Maria Moreira Zorello and ^{1,2}Guido Maier

(1) Dipartimento di Eletttronica, Informazione e Bioingegneria, Politecnico di Milano, Milan 20133, Italy

(2) SWAN networks, Via Fabio Filzi 27, Milan 20124, Italy

Abstract—Wide Area Network (WAN) reliability has become an imperative for enterprises with Cloud-hosted applications and distributed branch offices. Many solutions and different network technologies have been proposed over the years, such as leased lines, frame relay, or Multi-Protocol Label Switching (MPLS). Those solutions offer Quality of Service (QoS) at costs that are too high for companies today. Software-Defined Wide Area Network (SD-WAN) is regarded as the promising technological solution for the next generation of enterprise networks, since it is capable of increasing network reliability and agility while reducing costs. In this paper, we provide an overview of this technology, addressing its advantages and research opportunities.

Keywords—SD-WAN, SDN, MPLS, Enterprise Networks

I. INTRODUCTION

Digitalization of services is a major process of transformation enterprises are undergoing during the latest years. This process implies to rely more and more on Cloud computing applications. Consequently, the Enterprise Network (EN) that effectively and safely interconnects all company branches to both private and public Cloud data-centers, has become a crucial asset, largely contributing to enterprises' costs and revenues. Being a geographically-distributed infrastructure, the EN of a large company includes connections across Wide Area Networks (WANs). Almost always, WAN connectivity is provided by network operators contracted by the enterprise. Cloud-computing applications require ever-increasing network performance; hence, enterprises expect greater reliability, agility, and performance from their WAN connectivity, but this often implies high costs to be paid to the network operators to obtain higher QoS.

Over the years, many different technologies have been proposed to enhance WANs in terms of cost and quality: for instance, leased lines [1], Asynchronous Transfer Mode (ATM), Frame Relay (FR), Multi-Protocol Label Switching (MPLS), and Virtual Private Networks (VPN). Among these technologies, MPLS is currently widely deployed and can provide guaranteed Quality of Service (QoS) with high efficiency. However, the cost of MPLS is quite high, and it is pushing enterprises towards the new technology named Software-Defined Wide Area Network (SD-WAN) [2].

SD-WAN promises to provide QoS at lower costs by leveraging broadband Internet [3] as WAN connectivity. As well known, Software-Defined Networking (SDN) paradigm separates the physical forwarding elements (data plane), from the network control logic (control plane), which is implemented in a logically centralized controller [4]. In the specific SD-WAN case, the controller has the function of constantly monitoring the WAN interconnections and of selecting at each instant the best performing WAN path among multiple ones that connect two remote EN sites. This ensure QoS, even if each WAN connection is a simple best-effort and inexpensive broadband-Internet path. The control logic is for instance what differentiate SD-WAN from the




Figure 1. Example of SD-WAN edge solution deployed into an EN

VPN solution, allowing providing a good degree of QoS and reliability at the same access cost of VPN.

We point out that SD-WAN can be of two types. The first type entails the implementation of SDN into the WAN, i.e. one or more software-based network controllers manage all the devices on the provider's network according to SDN-based protocols (e.g. Netconf, Openflow, etc.). It fully implements the SDN paradigm, and for the sake of clarity, we will refer to this type as *SD-WAN core*, since the control is mainly aimed at the devices of the core WAN. The second type consists of deploying SD-WAN only at the edge of the network, that is, within the Customer Premises Equipment (CPE) devices located into the branch offices, headquarters, and data-center sites of enterprises. Therefore, we will refer to this type as *SD-WAN edge*, since it is totally independent of the provider's network.

SD-WAN edge represents a revolutionary way to simplify the EN and achieves optimal application performance using centrally managed WAN virtualization. We can define it as a virtual WAN architecture that allows enterprises to leverage any combination of transport services, including MPLS, 4G/LTE and broadband internet services, and to securely connect users to applications (see Figure 1). An SD-WAN edge solution exploits a centralized control function to direct traffic across different WANs. This solution increases application performance and delivers a high-quality user experience, resulting in increased business productivity, agility, and reduced costs for IT. Nowadays, SD-WAN edge is gaining momentum as connectivity solutions in the enterprise-network market. Currently, there are around 30 SD-WAN vendors in the market with varying levels of product maturity. Just to mention the largest players, they include VeloCloud (owned by VMware), Viptela (owned by Cisco), Nuage (owned by Nokia), and Silver Peak (owned by Aruba).

On the one hand, SD-WAN edge introduces many advantages by helping enterprises optimizing their branch site connectivity; on the other, it brings numerous technical challenges, for instance: placement of the software controller (in the Cloud or at premises), reliability of the controller in case of failure, scalability of the network architecture. This short paper provides an overview of SD-WAN edge, looking at the requirements and the fundamental aspects for its implementation. We investigate the building blocks of this

technology by highlighting the research opportunities and challenges. This investigation is motivated by the commercial success of this technology; curiously enough, SD-WAN has not been widely studied in research literature.

II. SD-WAN ARCHITECTURE

The static nature of legacy WAN architecture is inapposite to cope with today's increasingly Quality of Experience (QoE) requirements of modern businesses and enterprises. The main goal of SD-WAN is to simplify networking operations in WANs by introducing innovation and flexibility compared to legacy WAN technologies. In the following, we introduce the SD-WAN architecture shown in Figure 2, composed of three planes: data, control, and orchestration.

A. Data plane

The data plane is the part of a network that carries user traffic. As previously mentioned, an SD-WAN edge solution does not provide any control of the WAN network devices (such as switches and routers) but only of those placed at the edges of the enterprise's network, called CPEs. The CPE allows the local network of a remote office to be connected to the enterprise network through the WAN connectivity. In traditional EN, CPEs were typically totally "black boxes" for the Enterprise customer, as they were fully configured and managed by the carrier that provided WAN connectivity.

In SD-WAN, the owner of the CPEs depends on the specific circumstance (it can be the customer, the carrier or the SD-WAN vendor), but in all the cases the CPEs must be made programmable by the controller. Moreover, SD-WAN CPEs can be hardware- or software-based. Two of the most popular and recent implementations of the CPEs in SD-WAN are called Virtual CPE (vCPE) and universal CPE (uCPE). The vCPE is a CPE implemented as a Virtual Network Function (VNF).

The virtual machine implementing the vCPE can be executed in a customer's private data-center or even in Cloud. In a vCPE, the hardware required at each branch office becomes very simple, e.g. a switch or a router. Usually, an Enterprise customer, beside connectivity, requires additional services at each site, such as load balancing, encryption, firewalling, antivirus, intrusion detection/prevention, etc. In the past, these services were implemented by ad-hoc hardware elements installed on-site (firewalls, load balancers, VPN servers, intrusion prevention systems, etc.). Today, these services can be implemented as VNFs at the customer's data-center or in Cloud, in the same way as the vCPE.

The uCPE is a recent enhanced version of vCPEs. It is also based on software and VNFs, but virtualization is hosted inside the remote offices themselves using local appliances with high processing capabilities, such as virtualization servers or mini-datacenters. These facilities will execute as VNFs not just the SD-WAN functions, but also the other additional network services required by the customer. Using commercial off-the-shelf hardware and a standard operating system, uCPEs can offer multi-vendor and multi-component network services in remote locations.

As shown in Figure 2, SD-WAN provides different overlay networks, *i.e.* virtual networks built on top of underlying network infrastructures connecting the vCPEs/uCPEs. The main idea of an overlay network is that some form of encapsulation is used to decouple a network service from the underlying infrastructure, for instance,

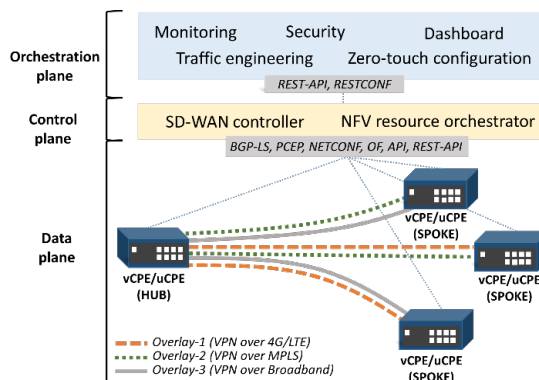


Figure 2. Typical SD-WAN architecture

MPLS, broadband internet, 4G/LTE, etc. Overlay networking uses many different networking protocols and standards built over time. Some of the protocols developed for overlay networking technology include different kinds of VPNs, such as Virtual eXtensible LAN (VXLAN), Overlay Transport Virtualization (OTV), Virtual Private LAN Service (VPLS), and Generic Routing Encapsulation (GRE). The use of overlay networking to connect multiple CPEs brings numerous advantages, such as high agility, scalability, and customization of overlay networks based on the services requested by the user.

B. Control plane

The control plane is a logical entity that receives instructions from the orchestration plane by means of the Northbound Interfaces (NBI's). It is provided by protocols such as REST API and RESTCONF. It relays them to the networking components in the data plane through the Southbound Interfaces (SBI's) provided by different alternative protocols such as BGP-LS, PCEP, NETCONF, OpenFlow (OF), custom API, and REST API. The SDN controller extracts an abstract view of the network, including statistics and event information, from the CPEs and communicates back to the applications. In the context of SD-WAN, the controller is responsible for the overlay network configuration (such as IP address management), and for pushing down policies onto the CPEs. SD-WAN providers develop their controllers together with the other components ready to be installed and used by customers. There is also the possibility of using SD-WAN controllers available open-source, as ONOS [5] and OpenDayLight [6]. Both of them have highly modular architecture offering a rich set of programmable interfaces for applications to run automatically on network systems.

C. Orchestration plane

The orchestration plane provides the management of the SD-WAN applications such as: traffic engineering, monitoring, zero-touch configuration and security. It orchestrates the applications to make them work in unison. Moreover, the SD-WAN edge provider integrates a dashboard to the customer in order to be able to manage the resources of its network in a simple and automated way. This dashboard interacts with the orchestration plane and gives the customer the opportunity to manage the applications such as traffic engineering and monitoring. In the following, we show the fundamental applications that operate within the orchestration plane that represent the backbones of an SD-WAN solution.

1) Monitoring

Monitoring has always been a crucial point for network management operations. It provides a view of the network conditions, operations, and usage statistics for traffic engineering and QoS regulations. By monitoring, we mean the collection of measurements of specific network parameters that allow having an insight into the state of links and equipment. Moreover, these parameters tell if the service provided by the network meets or not the service quality levels agreed with the customer. The most common monitored parameters are packet loss, delay, and jitter. Monitoring requires a suitable protocol on the SBI which support transmission of telemetry data from the CPEs to the controller.

2) Traffic engineering

Traffic engineering is important for network availability and reliability. The SD-WAN controller has to orchestrate enterprise traffic so that the WAN performance measured by monitoring can match the service requirements at each instant. The traffic engineering application comprises two main components:

- Service classification: it classifies the packets at the vCPE/uCPE ingress into a specific service class, such as VoIP, video streaming, etc., in order to assess the service requirements
- Policy implementation: it applies the SD-WAN policies which are associated with the service class in order to select the most suitable overlays to use at each instant

3) Security

SD-WAN edge uses authentication and encryption to fully protect end-to-end traffic [8]. IT security experts can monitor the quality of connections and ensure that all communications meet company policies for security and reliability. SD-WAN edge provides end-to-end encryption across any underlay network type, including the Internet, by exploiting technologies like IPSEC VPN, IKEv2 with certificate, end-to-end encryption using AES256, shared keys, and PKI.

III. RESEARCH OPPORTUNITIES

After giving an overview of SD-WAN and presenting its architecture, we identified different research challenges that deserve further investigations. We now briefly explain some of these potential research areas.

Control plane management. The controller is the most critical component of an SD-WAN solution. Its position can affect the performance, especially when the network is geographically extensive. The distance from the CPEs on the edge can result in delays that may slow-down the execution of controller's decisions by the CPEs. An asymmetric location of the controller w.r.t. the position of the CPEs may cause the CPEs to react asynchronously, making the routing decisions at the different customer premises inconsistent on switching transients.

Scalability and reliability. They represent an important aspect to be considered when deploying an SD-WAN edge solution. In fact, as the network grows in size (e.g., increasing the number of CPEs), the centralized SD-WAN controller becomes highly solicited (in terms of events/requests) and thus potentially overloaded (in terms of bandwidth, processing

power, and memory). Moreover, the failure of the central controller may collapse the overall network.

Traffic engineering and monitoring. The SD-WAN control plane has a global, centralized view of the network and thus can also access network statistics and properties through monitoring algorithms. With centralized traffic engineering solutions, this data can be leveraged to find globally optimal path assignments. Traditional traffic engineering mechanisms such as RSVP-TE or LDP for MPLS rely on the local, limited view of the network from the ingress router. Programmability in SD-WAN offers possibilities to implement custom, fast, and efficient adaptive routing schemes. Efficient network monitoring is required for the development of control and management applications in SD-WAN edge solutions. However, collecting the appropriate data and statistics without affecting the network performance is a challenging task. In fact, the continuous monitoring of network data and statistics may generate excessive overheads and affect the network performance, whereas the lack of monitoring may cause incorrect behavior of management applications.

Security. It is another crucial challenge for SD-WAN. The use of VPN tunnels among the CPEs and the controllers can reduce the risk associated with the most common network attacks such as DDoS. However, if a malicious CPE plugs into the network, this can put all security at risk. The development of secure authentication mechanisms of the overlay tunnels to avoid network security risks is of fundamental importance to SD-WAN.

IV. FINAL REMARKS

An enterprise WAN is a network that connects geographically spread sites of a company that could be located anywhere in the world. MPLS has been so far the main WAN technology for enterprise networking because of its high performance. Although MPLS has many advantages, SD-WAN is a new and fast-growing paradigm that could achieve similar performance but more cost-effectively. In an increasingly Cloud-centric world, this revolutionary technology is universally acclaimed as a new and unprecedented way to easily implement policies across large WANs at a fraction of the cost of traditional solutions. In this short paper, we presented an overview of this technology by showing its architecture and some of the many research challenges that remain still open and need further investigation.

REFERENCES

- [1] R. Graziani and B. Vachon, Cisco Networking Academy: Connecting Networks Companion Guide. Cisco Press, 2014.
- [2] Rangan, Raghavan Kasturi. " Trends in SD-WAN and SDN." CSI Transactions on ICT 8.1 (2020).
- [3] Press, Cisco. "SD-WAN in the Public Sector", 2020.
- [4] D. Kreutz, et al., "Software-Defined Networking: A Comprehensive Survey," in Proceedings of the IEEE, vol. 103, no. 1, pp. 14-76, 2015.
- [5] Berde, Pankaj, et al. "ONOS: towards an open, distributed SDN OS." Proceedings of the 3rd workshop on Hot topics in SDN. 2014.
- [6] Medved, Jan, et al. "Opendaylight: Towards a model-driven sdn controller architecture." Proceeding of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks 2014.
- [7] McKeown, Nick, et al. "OpenFlow: enabling innovation in campus networks." ACM SIGCOMM Computer Communication Review 2008.
- [8] Wood, M. "How to make SD-WAN secure." Network Security 2017.