



POLITECNICO DI TORINO

Master Degree course in Communications and Computer Networks
Engineering

Master Degree Thesis

**Study of applications and testing
scenarios of Citrix SD-WAN
solution**

Supervisor

Prof. Paolo GIACCONE

Candidato

Silvia PASSARO

Company Supervisor

Giuseppe CARISTIA

ACADEMIC YEAR 2020-2021

Abstract

Traditional networks are no more suitable for the increasing traffic demand of the recent years. Indeed, they were based on coupling the data and the control plane in each device of the network. This approach is old and poorly flexible because it is strictly related to hardware vendors and is too much pricey nowadays. For satisfying the high bandwidth request and the switchover to cloud-based applications, the traditional approach would imply purchasing a lot of new devices, that is costly and energy-consuming. Thus, a transition towards the Software Defined Networking (SDN) is occurring.

SDN is a new network paradigm that reverses the previous approach by decoupling the data and the control plane, implementing the first into the dummy switches of the network, and the second into a specialized hardware usually known as SDN controller or orchestrator. This new layout turns out to be more flexible, cost-effective and dynamic and makes the network control plane directly programmable.

The current work investigates the application of SDN to wide area networks (SD-WAN), where wide area networks are massive networks whose function is to connect users at the branch or campus to applications hosted on servers in the data center. Actually, the point of interest of the thesis relies on understanding what has brought enterprises to develop this new solution. In the past, wide area networks were mainly implemented through dedicated MPLS circuits, which are expensive and require long lead times for the deployment. Typically, dedicated MPLS circuits were used to help ensure security and reliable connectivity. Nowadays, the majority of applications is developed by means of Software-as-a-service platforms on the cloud, and the old MPLS is no longer up to the job.

The thesis has been realized in partnership with the Citrix vendor, with the aim to test its peculiar SD-WAN solution, its capabilities and the innovations it brings. For this purpose, many simulation scenarios have been deployed by creating a virtual network infrastructure on the VMware ESXi platform. The simulated network involves a branch office and a head quarter, both connected to two instances

of a Citrix SD-WAN appliance and communicating between them thanks to an emulator of wide area networks that mimics an MPLS circuit and an ADSL link.

The tests are based on exchanges of data among the two sites while monitoring continuously the network statistics through the graphical interface of the Citrix product. The results of the experiments have highlighted which are the benefits of the new technology, starting from the application routed traffic dynamism, to the improved operational expenditures, to the simplified and programmable management and to end up in the enhanced WAN performances.

To wrap up, the Citrix SD-WAN solution offers a competitive environment for the transition from traditional networks to a cloud-centric world. However, there are still some open issues to address: mobile users are completely transparent to SD-WAN devices, no true end to end quality of service is guaranteed, there is no on-site security functionality and there is still the need for qualified IT staff to deploy and maintain this technical solution.

Contents

1	Introduction	1
2	SDN technology: from virtualization to SD-WAN	3
2.1	Virtualization	3
2.1.1	Types of virtualization	5
2.1.2	Benefits	6
2.1.3	Towards a new approach: container deployment	7
2.2	Software Defined Networking	9
2.2.1	Main features	10
2.2.2	How SDN works	11
2.2.3	Advantages of SDN	11
2.2.4	Use-cases	12
2.2.5	Future of SDN	17
2.3	SD-WAN: Software Defined - Wide Area Networking	18
2.3.1	SD-WAN architecture	19
2.3.2	Operating principles	20
2.3.3	SD-WAN security	21
2.3.4	SD-WAN vs. SDN	23
2.3.5	Limitations	24
3	Citrix SD-WAN solution	26
3.1	Citrix SD-WAN overview	26
3.1.1	Master Control Node	27
3.1.2	Standard Edition	28
3.1.3	Wanop Edition	29
3.1.4	Premium Edition	30
3.1.5	Advanced Edition	30
3.2	Citrix SD-WAN VPX SE	30
3.2.1	Inline mode	31
3.2.2	Web Cache Communication Protocol	32
3.2.3	Virtual Inline mode	32
3.2.4	VPX usage scenarios	33

3.3	How Citrix SD-WAN VPX SE works	37
4	Simulations and results	39
4.1	Network infrastructure	39
4.2	Configuration	40
4.2.1	WAN emulator configuration	40
4.2.2	Server configuration	41
4.2.3	SD-WAN configuration	42
4.3	Citrix SD-WAN appliance	44
4.3.1	Configuration	44
4.3.2	Monitoring	44
4.3.3	Dashboard	46
4.4	WANem	47
4.5	Experimental evaluation	48
4.5.1	Scenario 1: Jperf traffic	49
4.5.2	Scenario 2: Video download and Jperf	50
4.5.3	Scenario 3: Video streaming and Jperf	52
4.5.4	Scenario 4: Download and video streaming	54
4.5.5	Scenario 5: Download with throttling profile and video streaming	55
4.5.6	Scenario 6: Network failures with video streaming and Jperf	57
4.5.7	Scenario 7: Delaying the ADSL channel	59
4.5.8	Scenario 8: Failures: severing MPLS channel	62
5	SD-WAN advantages and deployment scenarios	65
5.1	SD-WAN advantages	65
5.2	Organizations migrating to SD-WAN	67
5.2.1	Retail stores	68
5.2.2	Healthcare	69
5.2.3	Financials and Insurance	70
5.3	Deployment scenarios	70
5.3.1	Equinix: Bridging Multiple SD-WAN Deployments	71
5.3.2	Equinix: Simplification of Multi-Cloud Deployment	72
5.3.3	Lavelle Networks: SD-WAN tunnel with Internet and MPLS WANS	73
5.3.4	Lavelle Networks: SD-WAN service using Multiple Internet Service Providers	74
5.3.5	Citrix: SD-WAN and Microsoft Office 365	75
6	Conclusion	77
6.1	Future works	79
	Bibliography	80

Chapter 1

Introduction

Nowadays the world is moving towards a more robust exploitation of virtualization in every aspect of a network, from servers to network functions. Due to the increasing demand of bandwidth by users, it is necessary to guarantee a secure navigation, a fast and reliable connection with a low latency and a less costly equipment. Usually, companies rely on large Multi Protocol Label Switching (MPLS) networks for granting a reliable service to users, but they have the drawback to be very complex to implement and they require an expensive equipment. This is why virtualization had been introduced.

The current thesis investigates all-around every aspect of the transition from traditional networks to the modern architectures. The proposed innovations concern with making each stuff in a network, from network devices to network functions, software-based, that means, programmable. The new software-based network framework is called Software Defined Networking (SDN).

SDN leads to a new way of building and managing networks, because it exploits virtual environments for the implementation of the devices and manages them by means of a centralized entity, the controller. A fundamental remark regards the way the controller rules the network, which is again by software. It communicates with all the other entities in the network through interfaces and installs rules in the switches and routers that are in charge of forwarding operations.

Once said that, the real purpose of the current work is to investigate one of the most interesting and promising SDN use cases: the software defined approach applied to wide area networks (SD-WAN). SD-WAN conceives wide area networks in a new fashion, going deeper into details, it tries to replace the traditional adopted MPLS circuits with encrypted virtual tunnels virtualized through software and hypervisor platforms.

The thesis work has been realized in collaboration with NET Reply company, business unit of Sytel Reply, so it has been allowed the access to a peculiar vendor SD-WAN solution: the Citrix SD-WAN VPX SE appliance, being NET Reply a Citrix Solution Advisor. Thus, the next chapter portrays step-by-step the current need for virtualization, then how virtualization has been leveraged in the SDN paradigm and, in the end, an overall description of the SD-WAN use-case.

Then, the focus of the work has been transferred towards the experimental phase, thanks to which all the pros and cons of this new technology have been practically tested. In the following chapters it will be presented the network infrastructure that has been realized for this intent and all the tests that have been run on it.

Ultimately, the subject of the study has also covered the real scenarios that the most famous vendors have implemented and tested in order to meticulously evaluate the starting point of the future work that will improve SD-WAN architecture.

Chapter 2

SDN technology: from virtualization to SD-WAN

2.1 Virtualization

Virtualization has its roots in the late '60s, despite the fact it actually spread at the beginning of the 21st century. Indeed, companies used to rely on the so-called commodity servers, where it was possible to allow the coexistence of more users, but with the limitation of being devoted only to a single task. Later, thanks to the introduction of the hypervisor software, it began possible to think about more users sharing the same hardware resources and doing several tasks on different operating systems. Nowadays, virtualization is the key aspect of modern networks, since it allows to create a software-based representation of applications, servers, storage and networks, whose aim is to reduce IT expenses while boosting efficiency and agility [1]. It leverages a software-based approach for realizing an abstraction layer over real physical hardware that becomes the basis of several virtual computers, called virtual machines (VM). Virtual machines are virtual environments miming physical computers, each with its own operating system and the great advantage they bring is that several VMs can be created on a single physical device, increasing scalability. VMs, differently from any other digital file, can restore their state

even when restarted in a different laptop or when several instances are opened. However, virtual machines require a sort of interface for communicating with the physical hardware below them, which is called hypervisor. An hypervisor stands for a software layer that allows the coexistence of multiple (virtual) operating systems running together and sharing resources in the same machine. The hypervisor is an essential element for virtualization purposes because it distributes the computing resources of the physical hardware to the VMs in order for them not to overlap. There exist two kinds of hypervisors:

- Type 1 or bare-metal: they are directly run on the host's hardware (e.g. a server), so they interact with it without any additional effort [2];
- Type 2 or hosted: they are applications on an existing operating system (OS) and are often located in the end-users. They require an additional overhead since they need to access the OS before communicating with the underlying hardware resources [2].

But why such a huge need for virtualization? The answer can be provided with an example. If three servers are dedicated each to a single task, such as a mail server, a web server and legacy apps and each of them is exploited only for the 30% of its capacity, the overall utilization of resources is pretty low [3]. Traditional networks acted like that, they used to reserve each server for a single activity as reported in figure 2.1.

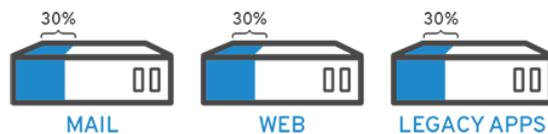


Figure 2.1. Server resource utilization in traditional networks, reproduced from [3]

Nowadays, instead, one of the key innovations introduced by virtualization stands in the possibility of devoting the same hardware to more applications; in

this example the mail server can be also employed for running legacy apps, so to save an entire dedicated server that could be employed in other activities or could be shut off with the aim of cutting down power and money. The current example clearly highlights one of the greatest advantages of virtualization, that is a better utilization of resources along with savings, graphically appreciable from figure 2.2.

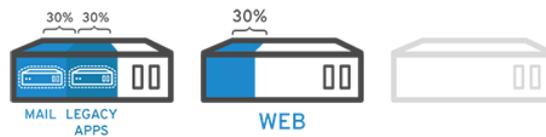


Figure 2.2. Server resource utilization with virtualization, reproduced from [3]

2.1.1 Types of virtualization

Virtualization does not mean only server virtualization as discussed up to now, but involves also different IT infrastructures, and among the most relevant ones there are: network, storage, data and cloud virtualization.

- Network virtualization: alongside Software Defined Networking (SDN) and Network Function Virtualization (NFV), this approach is based on providing an abstract whole view of the network that can be exploited by a central administrator for managing the entire structure. All the elements of a network are virtualized: from physical links to switches and routers, along with the network functions, they are all translated into software running on a hypervisor [2] [3].
- Storage virtualization: its aim is to merge all the storage devices of a network into a single shared pool of resources easily accessible from the virtual machines as necessary. Thanks to this method, the overall utilization of storage resources is improved and is easier to provision VMs with the needed resources [2] [3].

- **Data virtualization:** even the data is virtualized by specialized tools that create a software layer between applications accessing data and the system storing it. In so doing, applications can access any kind of data regardless of its format [2] [3].
- **Cloud virtualization:** by virtualizing all the resources mentioned above like storage, network and data a cloud computing model can be supported, where providers may offer several services. The main cloud-based services are: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). IaaS offers all the storage, network and server resources the user can configure; PaaS are the development environments where the user can build its cloud-based applications and SaaS involves all the software applications available on cloud [2] [3].

2.1.2 Benefits

The multiple benefits introduced by virtualization can be enumerated as follows :

- **Reduced costs:** virtualization can create significant cost savings by reducing the amount of hardware, the energy costs and streamlining maintenance and management. As a matter of fact, to run a lot of servers it is necessary to face buying, installing, upgrading and maintenance costs and by reducing the number of physical devices, these costs are decreased too.
- **Faster provisioning:** besides the considerable expenses, resource provisioning is also critical in terms of time, because after providing resources, it is necessary to wait for a long time to provide new ones or to modify their assets. In this field, virtualization comes in handy with the negligible time it takes for changing the configuration of a specific environment.
- **Fewer business disruptions:** handling a physical server failure is a quite hard task since it could take days to replace it and re-install everything on it. The

improvement introduced by virtualization in this field is that if a server fails, you simply bypass it and data keeps flowing.

- Easier backups: having full backups or snapshots of the virtual environment is pretty easy: it is a procedure run both at the system level and at the object one and it is completely transparent.
- No vendor lock-in: virtual machines reside on an abstraction level where there is no knowledge of the hardware below them. This means that you are not locked into a specific vendor for making the machine work and there is a lot a flexibility in the software implementation.
- Greater efficiency: data centers are composed by servers that need to be active 100% of the time consuming energy. However, statistical studies about their effective usage, show that the real usage is much less than the time they are switched on. So, by sharing hardware, software and infrastructure, virtualization can lead to a greater efficiency because servers are always under utilization and there is no waste.
- Head-start to the cloud: virtualization is the fundamental support towards a cloud-based approach in many fields of everyday life. A cloud-based technology would not exist without the existence of virtualization because it is strongly based on the concept of decoupling data from physical hardware and on storing them in a virtual environment.

2.1.3 Towards a new approach: container deployment

Virtualization has been the first step towards a new approach for application deployment, but there are even greater and more sophisticated innovations in this path. For better investigating the evolution in application deployment and workload support, it is worth analyzing the transition shown in figure 2.3.

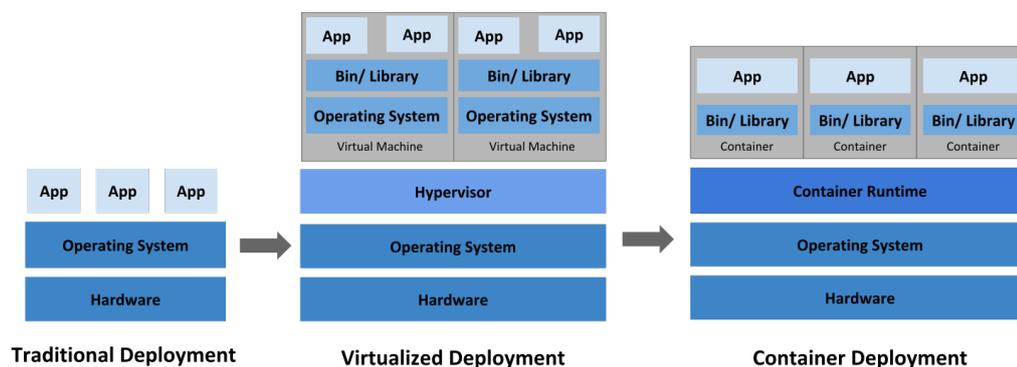


Figure 2.3. Platform deployment evolution, reproduced from [4]

The very first stage of this process is the traditional deployment era, during which many applications were hosted on the same physical server [4]. This approach is very static and leads to a poor resource utilization since it can happen that an application is more resource-hungry than another and leaves no sufficient capacity for other applications. A practical solution for this limitation is to make a single server host only one application, but this is not suitable for companies' revenue. So, the need for a more flexible baseline architecture.

The second key element of the chain is virtualization. It consists of building several virtual machines on the same physical CPU (Central Processing Unit, it is the central processor of a computer) and of running applications on top of them [4]. Each virtual machine is equipped with its own operating system and furnishes more security than the previous method, because applications are separated and cannot access other data. The introduction of virtualization has brought many advantages, like the already discussed reduction of costs and the greater scalability. However, there is an improved solution.

The last step of the evolution process leads to the introduction of containers, which are similar to virtual machines but with a lighter load. They are characterized by a lighter isolation model and they share the operating system between applications [4]. Containers are decoupled from the underlying infrastructure and

due to this, they are conceived to be portable among different cloud providers. Containers are considered the vanguard of modern technologies because of their huge advantages:

- Agile application creation and deployment: greater ease and efficiency in creating container images than using VM images;
- Consistency between development, testing and production: containers work the same way on a laptop as they do in the cloud;
- Portability between different clouds and operating systems: the same container works on Ubuntu, RHEL, CoreOS, on-premise, in the largest public clouds, and anywhere else;
- Resource isolation: application performance is predictable;
- Resource utilization: high efficiency and density.

2.2 Software Defined Networking

Before inspecting this new network framework, it is of fundamental importance to describe traditional networks. At the base of traditional networks there are: the Control plane, the Data plane and the Management plane. The Control plane is a set of distributed algorithms for topology discovery, topology tracking, route computation, installation of forwarding rules and traffic engineering. Basically, the Control plane is in charge of controlling and ruling the network in a distributed fashion. As for the Data plane, instead, it is also called Forwarding plane because it handles packet forwarding, scheduling, filtering, buffering, rate-limiting and marking. It consists of a set of local algorithms and it works at packets transmission time scale. Last, the Management plane involves local or global algorithms that deal with measurements, configuration, monitoring, protection and restoration.

Software Defined Networking (SDN) is a new network architecture whose purpose is to decouple the Control plane from the Data plane, differently from traditional networks where they coexist within each single device. SDN proposes to make the Control plane centralized and to run it on a single entity.

2.2.1 Main features

The SDN paradigm is characterized by some new pillars that make it much more flexible and groundbreaking. The novelties introduced by SDN can be scanned as follows [5]:

- Directly programmable: network control is directly programmable thanks to the logically centralized brain of the network which is the SDN orchestrator;
- Agile: abstracting control function from the forwarding plane lets administrators dynamically adjust network-wide traffic flow to meet changing needs;
- Centrally managed: network intelligence is centralized in software-based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch;
- Programmatically configured: SDN lets network managers configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs. Network managers can write programs themselves because they do not depend on proprietary software;
- Open standards-based and vendor-neutral: when implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of being provided by multiple, vendor-specific devices and protocols.

2.2.2 How SDN works

In software-defined networking, a software application controller manages the network and its activities. Instead of using hardware to support network services, SDN allows network administrators to virtualize physical network connectivity. The SDN framework is organized in three distinct layers [6]:

- **Application layer:** includes a set of applications and network functions that help improve application performance, simplify IT (Information Technology) tasks and increase security. Some relevant examples of these applications and network functions include application firewalls, application delivery controllers, load balancing and wide area networks optimization controllers. The great improvement of SDN is that it does not imply the usage of a specific hardware appliance for these applications, but it can virtualize them with a software-based approach. The current layer communicates with the control one by means of the Northbound interface.
- **Control layer:** is the centralized brain of the network and manages the Control plane through the SDN controller entity. It is responsible for managing policies, traffic flow and it processes the requirements received by the application layer by instructing the network elements in the layer below with the rules to apply.
- **Infrastructure layer:** it is equipped with physical hardware: it contains switches and routers. The mentioned network elements implement forwarding functions, data processing capabilities and collect some critical information to be sent back to the SDN controller, waiting for new instructions.

2.2.3 Advantages of SDN

The transition from traditional networks to a software-defined approach has its roots in the huge advantages it brings. The urge of enterprises of developing a new

network paradigm are mainly related to the increasing traffic demand and to the choice of moving applications and services to cloud platforms. Thus, SDN helps achieve this goal through its benefits and innovations.

Just like virtualization, the new network framework focuses its attention towards savings: instead of purchasing lots of devices, SDN relies on virtualization platforms where several virtual machines can be supported together according to the computing capabilities of the server owning the platform [7]. The result is that the amount of physical hardware is drastically reduced, and so, the costs companies are supposed to handle. However, there is something else beyond the decreased amount of appliances: the specialized hardware needed in traditional networks is much more pricey than the general-purpose servers employed in the SDN architecture.

Furthermore, the massive amount of virtual machines and, consequently, of virtualized network devices, allows to reach a greater scalability, also in real-time. As a matter of fact, the infrastructure can be enlarged and shrunken according to the real-time requirements, in order to avoid utilizing unnecessary resources, which are left free for other purposes [7]. The just described methodology is a flexible approach that helps not only save resources and money, but also adapt immediately to network changes.

Hence, the keywords of this switch-over from old networks to software-defined ones are: simplification and better resource utilization. The easier network management, in fact, does not require the intervention of network experts, whereas the flexibility of the architecture lets network managers exploit better the available resources without wastes.

2.2.4 Use-cases

The best strategy to become familiar with the SDN paradigm and to understand why it has been introduced in actual networks is to inspect its use-cases and who leverages this technology. The companies that adopted SDN first were cloud

providers, like Google, Microsoft and Facebook. Then, the SDN paradigm has been introduced in the network operators' companies in order to enhance their access networks. Only later did companies begin to adopt SDN technology [8]. What remains to be analyzed is what SDN is useful for, and the following are among the most relevant use-cases.

- Network virtualization: the most noteworthy SDN use-case is related to the virtualization of the network infrastructure. The key insight was that modern clouds required networks that could be programmatically created, managed, and torn down, without a system administrator having to manually configure them [8]. The decoupling of the Control plane from the Data plane and the centralization of the Control plane led to the exploitation of a single API (Application programming interface) entry point for the creation, modification, and deletion of virtual networks.

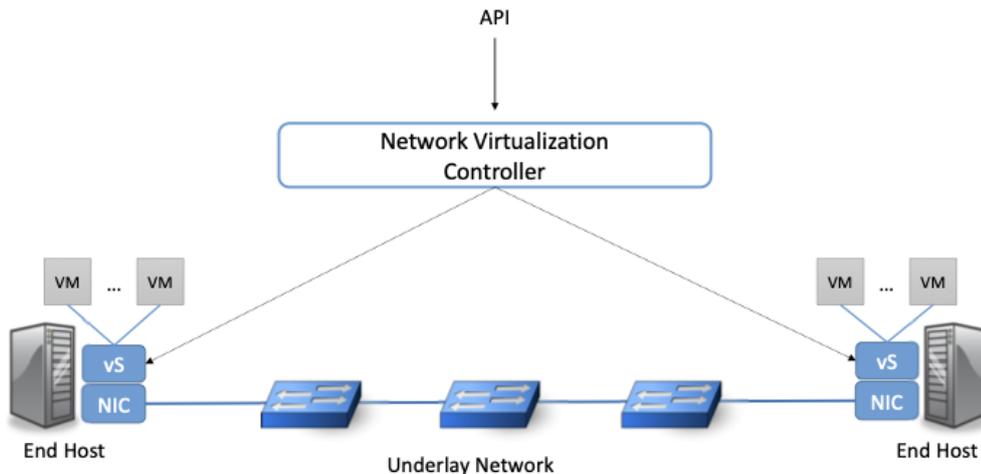


Figure 2.4. Network virtualization system (example), reproduced from [8]

Figure 2.4 shows an example of a network virtualization system, which is composed by a Network Virtualization controller able to build, control and modify networks by means of a Northbound API. It connects to switches, that can either be physical or virtual and they are programmed to forward

packets, with appropriate encapsulation, from host to host across the underlay network [8].

- Switching fabrics: however, the predominant use-case for SDN is within cloud data centers. Cloud providers have preferred to replace proprietary switches with bare-metal switches built using merchant silicon switching chips, both for monetary reasons and for increasing velocity. Cloud providers are entitled of handling the switching fabric architecture that interconnects their servers by software [8]. A leaf-spine topology is widely used for data centers, as reported in the example in figure 2.5.

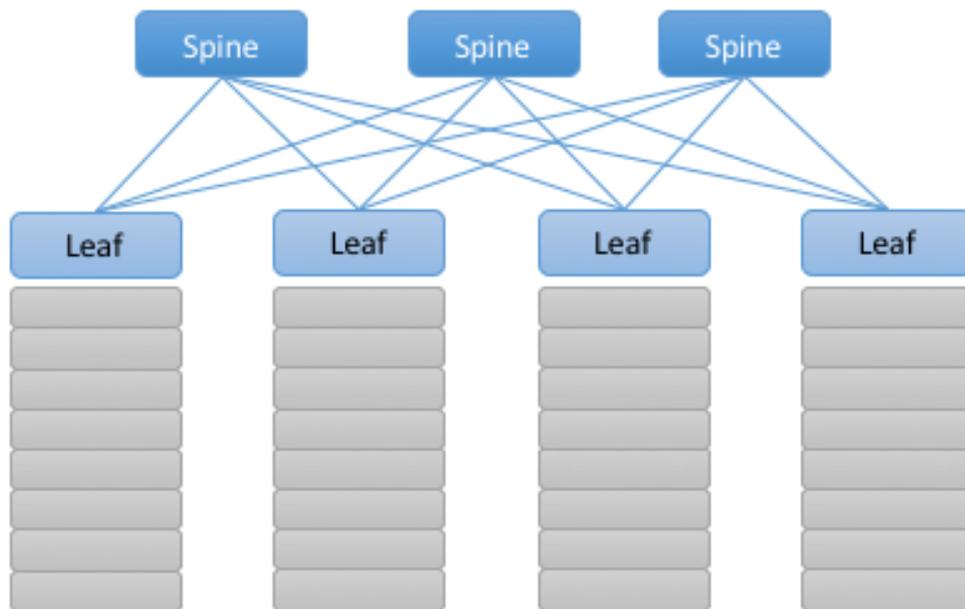


Figure 2.5. Leaf and spine switching fabric architecture, reproduced from [8]

The leaf-spine topology is usually composed by racks and each of them has a Top-of-Rack switch for connecting the switches inside the rack itself that are called leaf switches. Moreover, each leaf of the architecture must connect to a subset of available spine switches [8]. Employing SDN paradigm inside this kind of architecture means controlling by software the topology, for example at

layer 2 it is possible to implement forwarding/bridging policies while at layer 3, routing across racks.

- Access networks: SDN gives the opportunity to improve the performances of access networks that implement the last mile and connect homes, branch-offices and mobile devices to the Internet. The focus of the study is on the two most widespread access network technologies: the Passive Optical Network (PON) and the Radio Access Network (RAN) of 4G/5G standards. Nowadays, the challenge of this transition is to substitute the specific-purpose hardware that these network employ with general purpose devices to achieve a great payoff. But there is more beyond monetary reasons: building an hybrid network plenty of compute servers and access devices both controlled by software, along with the indispensable proprietary hardware, results in a very flexible and programmable solution that improves the overall performance of the access network [8]. The just described layout is called SDN-Enabled Broadband Access (SEBA) and can be appreciated in figure 2.6.

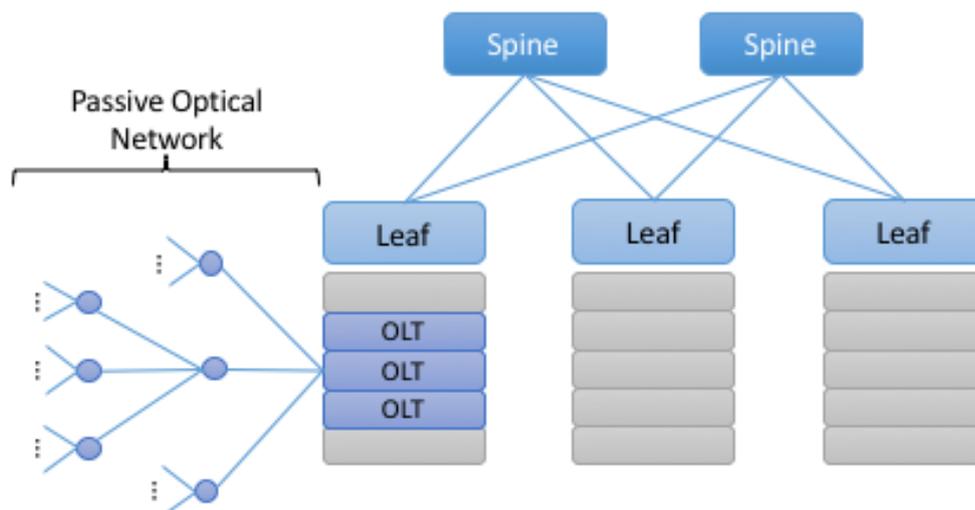


Figure 2.6. SEBA architecture, reproduced from [8]

- Network telemetry: this a pretty peculiar and new application of SDN. It

relies on the advent of programmable forwarding pipelines, the so-called In-band Network Telemetry (INT) [8].

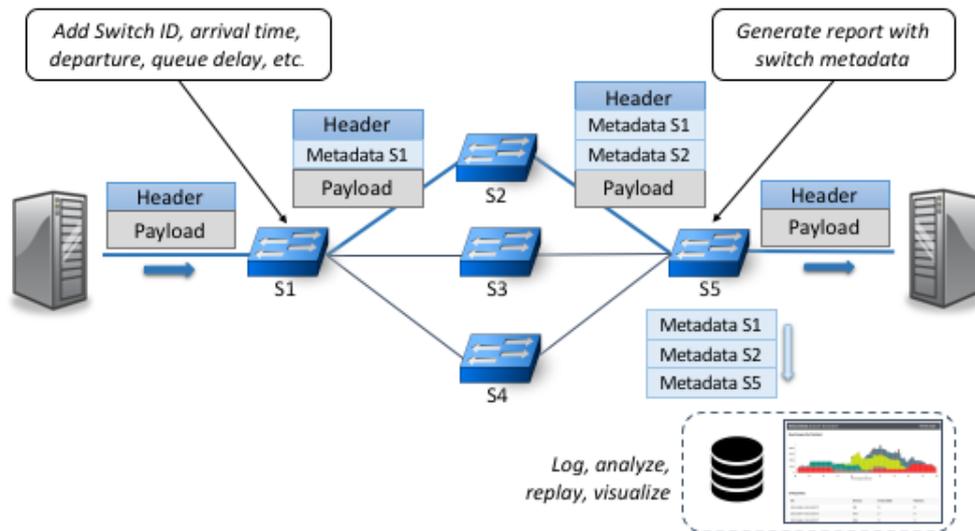


Figure 2.7. In-band Network Telemetry architecture example, reproduced from [8]

The state-of-the-art of INT is the collection of statistics about the performance and the state of the network on the fly, i.e., while packets cross the network itself, differently from the old technique done by means of the Control plane by reading fixed counters. From figure 2.7 it is possible to appreciate an example of how instructions are collected and processed. Whenever packets traverse the network each switch adds some INT metadata to the header and then traffic sinks retrieve the information about the Data plane actual state. Despite being at its early stage, this approach is potentially useful for allowing an easier troubleshooting and for simplifying the procedure of failure discovery. SDN is helpful in this context because it allows to freely experiment this innovative solution without wasting resources, something that was inconceivable years ago.

- Traffic engineering: it conveys all the practises for load balancing. Traditional routing implied to define the parameters of a network, like the weights of each

link, in order to run traffic routing algorithms and to find the best path in terms of load. This approach is old and difficult to be practically realized and, often, ended up in an optimal, but not excellent outcome. Conversely, by adopting SDN it can be reached a better result because the SDN controller has a global view of the network and is able to instruct bare-metal switches on the best way to route and balance traffic, for example splitting it in several paths.

2.2.5 Future of SDN

The subject of interest of this section is the discovery of the open challenges of SDN. For achieving this purpose, it is reasonable to identify the three phases of SDN expansion, which are depicted in figure 2.8.

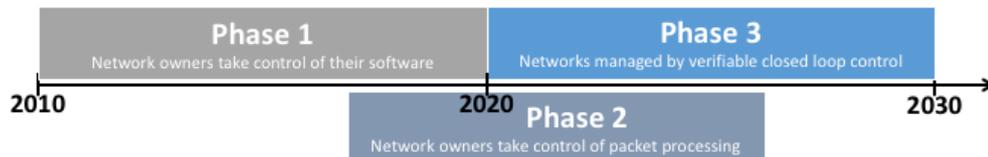


Figure 2.8. SDN challenges evolution over years, reproduced from [8]

While the first stage of the evolution regards the technology deployment and establishment, the most challenging phases are the second and the third one. Actually, the second phase is on going and is centered around the disaggregation of Control planes and on the substitution of old Data planes with P4-based ones. P4 is part of a larger innovation program called PISA, i.e. Protocol-Independent Switch Architecture used for packet processing. PISA is a ground-breaking paradigm thanks to which it is possible to provide a higher level of programmability since the SDN controller can implement some intelligence inside general-purpose switches. P4 is the abstraction employed by PISA for the Data plane programmability. With P4, switches can be configured and packets parsed and the output of the parser crosses

match-action tables to understand which is the correct action to be applied. The just described procedure turns out to be much more flexible and tries to reach one of the pillars of the second-stage SDN challenges, which is the ability of operators to fully take control of networks by software.

The third stage of the SDN evolution, instead, concerns with the top-down verification settlement. This target is about the urge of finding an effective way of verifying the network correctness. Indeed, with the advent of 5G networks that are supposed to support not only mobile devices, but also self-driving cars, drones and home appliances, the risk of cyber disasters is widely increased and must be counteracted [8]. The main idea for handling this issue is to work on a compositional network, that means, to work on small pieces of network, with the aim of facilitating the management. SDN makes this idea come true: the correctness of the network behavior is verified from top to down across each interface defined in the SDN architecture. Another remarkable element that has made this idea feasible is the exploitation of P4 language for the Data plane: it is pretty easy and it does not avail of complex frameworks like loops and pointer-based data structures that are responsible for the extreme difficulty in making analysis [8].

2.3 SD-WAN: Software Defined - Wide Area Networking

Software Defined - Wide Area Networking is the most salient SDN use-case. For most of the practical WAN implementations, enterprises have bought WAN services from the major telecommunication providers [8]. So, as a result, the majority of actual WAN networks rely on the MPLS (Multi-protocol Label Switching) solution. MPLS is a network standard intended for quality of service purposes since it is based on the idea of creating a dedicated path where packets are labeled and delivered to the destination with an established quality of service policy. Therefore, legacy

(hardware-based) WANs are extremely complex and brittle and require a significant amount of infrastructure to support remote offices [9]. Hence, software defined - wide area networks have been introduced. A software-defined wide area network is a virtual WAN architecture, in which any blend of network transport type can be virtualized and then centrally managed in software, to securely connect users to applications and desktops. [10].

2.3.1 SD-WAN architecture

SD-WAN architecture is made of three essential components, as shown in figure 2.9 [11]:

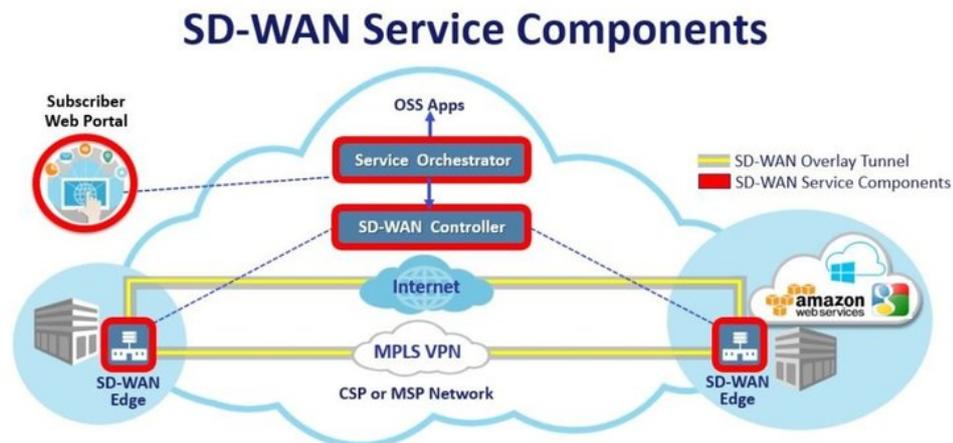


Figure 2.9. SD-WAN architecture with MPLS support, reproduced from [11]

- SD-WAN Edge: stands for the network's endpoints: in practical terms they can be branch offices, remote data centers, or cloud platforms;
- SD-WAN Controller: is the central brain of the infrastructure: it handles management, enables operators to have a global view of the network and sets policies the orchestrator has to execute;

- SD-WAN Orchestrator: is conceived as the virtualized manager of networks. Its role is about monitoring traffic and then applying policies and protocols.

Moreover, there is not only one kind of architecture, but three main ones [11]:

- On-premises SD-WAN: this type of architecture is on-site, that is, the SD-WAN device resides on physical machines. The feature that explains the name of this kind of architecture is that the network and the hardware are directly accessible from network operators, without the implication of the cloud. This is also the reason behind the usage of this architecture for private information;
- Cloud-enabled SD-WANs: implies the involvement both of the cloud and of internet. SD-WANs are built by connecting to a virtual cloud gateway over the internet. This structure provides better integration and performance with cloud-native applications;
- Cloud-Enabled with Backbone SD-WANs: the third architecture is the most complex but also complete one. In practice, it gives an extra backup option by allowing the switching from the public internet to a private connection. Thanks to this principle, the resulting network is more robust to failures and it is also more secure.

2.3.2 Operating principles

The key insight of SD-WAN is to abstract the details of networking layer giving to WANs the possibility to use several connection types, like MPLS, LTE (Long Term Evolution) or broadband internet [12].

SD-WANs are built by establishing encrypted tunnels - the overlay - between sites. Every site must involve an SD-WAN device and the traffic must cross it. Once connected to local networks, those devices automatically download custom-defined configuration and traffic policies and establish tunnels with one another or a point

of presence (PoP). Every time data crosses the appliance, the traffic is assigned to one of the existing categories and is redirected to the correct output according to what is written inside the rule tables and according to link performance status [12].

The routing and forwarding operations are managed by the three entities mentioned above. The SD-WAN devices choose the best path where to route the traffic according to statistics, which are computed instant-by-instant by the master control node, which is the most important one. Should one connection fail, the SD-WAN appliance immediately reacts by switching to another path without losing the connectivity.

Due to high level of programmability of the SD-WAN devices, the network manager can also create ad-hoc policies: according to the kind of traffic observed and its quality of service requirements, he can assign different priorities and decide how to distribute the available resources.

2.3.3 SD-WAN security

Security is one of the most critical concerns about SD-WAN deployment, because SD-WAN appliances are not always equipped with security functions. However, due to the advent of a cloud-based world, improved security features are mandatory as the risk of cyber attacks is much consistent in actual WANs. Suffice it to say that these days there are many branch offices, remote workers, campuses and Internet-of-Things (IoT) devices that require a high security level for both traditional and cloud-based applications [13]. MPLS and physical SD-WAN devices are unable to meet this requirement, so it has been proposed a new solution that involves two frameworks: the Secure Access Service Edge (SASE) and the Zero-Trust Network-Access (ZTNA) [13]. They have been developed along with the SD-WAN architecture to fully exploit its capabilities and provide a user-centric security approach.

SASE is a solution which is capable of providing core security functions like encryption, firewall and access control by making them available in the cloud points of presence (POPs) all over the world. SASE partners with ZTNA to ensure trusted access based on user identity rather than based on user's location or on its IP (Internet Protocol) address [13]. SASE providers, which are usually the same companies developing SD-WAN solutions, are entitled of building a global fabric of PoPs and peering relationships with cloud providers. The PoPs act like an on-ramp to SaaS applications and other cloud services [13]. Each POP implements a set of security functions that are activated whenever a user or a device connects in a branch or via remote access, and they can be enumerated as follows [13]:

- ZTNA;
- Secure web gateways;
- Cloud access security broker (CASB);
- Cloud-based firewalls;
- Identity services to establish both user's context and security posture.

The outcome of this brand new technology is of great benefit to the user, especially in today's world. With the emergence of smart working, employees can leverage any kind of application, from the traditional to the cloud-based ones and automatically get the appropriate level of protection, wherever they are. Indeed, ZTNA is intended for supplying user's contextual identity (such as her location and the security posture of her device) to all the different SASE security services [13]. Then, whenever a user tries to access a resource, ZTNA examines the contextual identity which dictates the policy to be applied. Therefore, the rules are executed automatically without any external human intervention, saving money and avoiding waste of time.

The current combined model also facilitates enterprises' life. In point of fact, there is only a holistic solution and infrastructure to manage security and access and to carry out tasks like troubleshooting, monitoring and deployment.

2.3.4 SD-WAN vs. SDN

Despite being an SDN use-case, SD-WAN and SDN are commonly used for different objectives. To begin, SDN is a network framework mainly adopted in data center networks for reducing the operating expenses and for enabling on-demand services in a centralized fashion, by means of a central control brain [14]. On the other hand, SD-WAN was born with a different intent: it is supposed to replace, or at least to complement, MPLS-based WAN networks also reducing deployment costs and increasing scalability.

As regards SDN, since it mainly operates inside data centers, it relies on an infrastructure where the controller is usually in close proximity to the forwarding devices, that experience a pretty stable bandwidth and latency. Wide area networks, instead, are characterized by controllers which are reasonably far from the branch offices they interconnect, so there is a huge and unpredictable variability of the performances [9]. Consequently, the Data plane in SD-WAN networks is more distributed than in SDN ones and it is separated from the centralized controller [9]. However, this is an advantage for SD-WANs, because it allows the forwarding devices at remote locations to still be operating whenever the SD-WAN controller stops working. So, remote locations handling the Data plane in SD-WAN architecture can preserve connectivity to their transport connections and still work, despite the interruption of the SD-WAN controller's service [9]. It should also be remembered that the transport connections adopted by the SD-WAN paradigm are often proprietary and not directly controlled by the SD-WAN owner. Conversely, the Data plane in data centers is handled by organizations in a tight and consistent way, since there is no need to rely on several transport connections and providers [9].

Moving on to another topic, another relevant difference in the way SDN and SD-WAN operate is about the Control plane management. The SDN Control plane is plenty of protocols which result to be fast, with constant and predictable performances and are ruled by private organizations. The statistics about the network status are collected exploiting metrics based on link-state or distance vector protocols. The same does not hold for SD-WAN Control plane: being destined to a traffic that travels across long distances, protocols are less stable and prone to significant latencies. Furthermore, the metrics leveraged in data centers for statistic collection are not suitable for WAN paths [9]. The ideal features of an SD-WAN Control plane, instead, have the aim of reducing the protocol state data exchanged between endpoints and of removing routing protocols from remote locations. The ideal replacement consists of substituting the routing protocols with Application Programming Interfaces at remote locations.

2.3.5 Limitations

The key point in using SD-WAN is the plan to replace the actual transport technology of the majority of WANs: the MPLS standard. Actually, MPLS circuits are expensive and time-consuming to connect to new locations since they are dedicated circuits. Nonetheless, as any new technology, it cannot only guarantee advantages.

First of all, since the traffic is encrypted, when exposing the branches to the internet many malwares and attacks could arise, so it is necessary to purchase some new security appliances and the costs are not completely amortised. This is due to the fact that there is not an on-site security functionality fully established yet.

It should also be taken into account that mobile users are completely transparent to SD-WAN devices, which means that a significant portion of the nowadays traffic is not manageable by this new technology.

But the most relevant issue of SD-WAN is the lack of true end-to-end quality of

service guarantees. In fact, if exploiting only SD-WAN provisioning, it is impossible to guarantee a satisfying performance of the network because without employing an hybrid infrastructure that relies also on MPLS circuits, the SD-WAN does not support yet full connectivity.

By way of conclusion, SD-WAN does not achieve all the goals it had promised. In point of fact, the communication channel among Control plane and Data plane still remains proprietary and SD-WAN networks still must rely on the existing hardware [8].

Chapter 3

Citrix SD-WAN solution

3.1 Citrix SD-WAN overview

Nowadays, many vendors are delivering their SD-WAN solutions and the one under study in the current thesis has been proposed by the Citrix vendor.

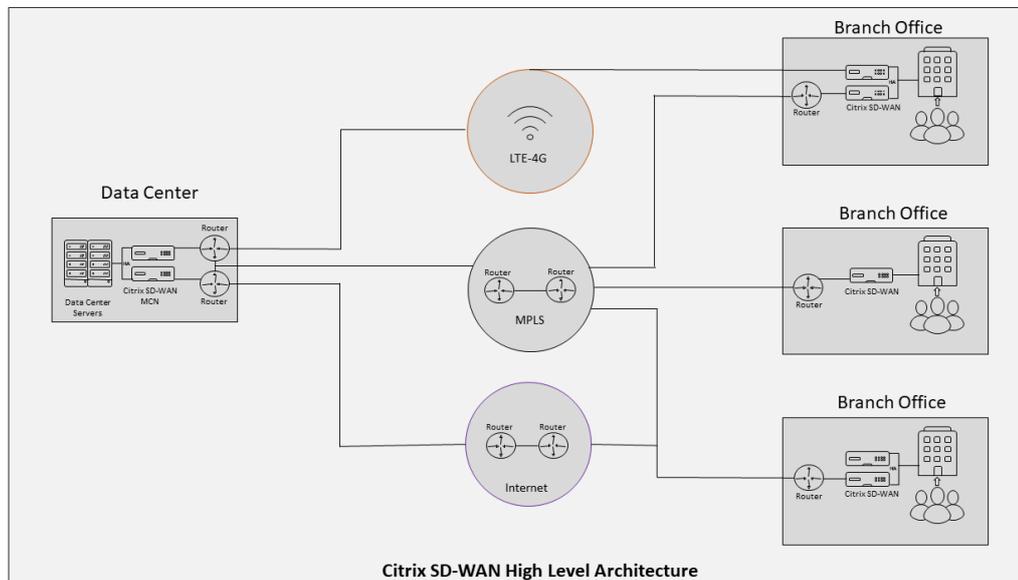


Figure 3.1. Citrix SD-WAN network architecture, reproduced from [15]

Citrix SD-WAN simplifies branch networking with a reliable and high-performance workspace experience that helps accessing SaaS applications, virtual desktops, or

traditional data centers. Citrix SD-WAN solution is typically used to virtualize wide area networks, providing connectivity between branch offices and central data centers, as displayed in figure 3.1. The solution also does provision the branch offices with additional connectivity methods to the cloud, which include tunnels to resources such as Azure virtual, enabling alternative ways to directly connect numerous branch office sites to cloud resources. Indeed, it must be taken into account that Citrix has a tight integration with cloud vendors (e.g. Microsoft Azure).

It can be deployed as the edge device where to consolidate the networking footprint or can be seamlessly integrated behind existing edge devices, also due to the lack of a stand-alone SD-WAN infrastructure.

It is fundamental to state that the Citrix SD-WAN world can be partitioned into two essential categories: the physical appliances and the virtual ones (VPX). Citrix SD-WAN physical platforms are usually branch side appliances that can be deployed in micro and small branch offices, remote sites, home offices, retail stores and temporary work-sites. A single box-in-branch solution helps to reduce the hardware footprint and eases branch deployment [16]. Therefore, the Citrix SD-WAN solution provides two possible alternatives according to the user needs, but the subject of interest that has been studied for the current thesis is the virtual appliance solution, in particular the Citrix SD-WAN VPX SE (Standard Edition). Despite the focus of the work, it is worth taking a look at the core components of the Citrix SD-WAN framework and at other editions.

3.1.1 Master Control Node

The core component of the Citrix SD-WAN architecture is the Master Control Node (MCN). The Master Control Node is the central Virtual WAN appliance. Its role is to establish and utilize virtual paths with one or more client nodes located across the LAN. It is responsible for time synchronization, routing updates and the hub for the branch devices. Moreover, it is mandatory to use a static IP address for its

configuration.

Even if there could be several MCNs, only one can be active at any given time and it monitors the entire virtual LAN. Client nodes, instead, can only monitor their local intranets along with some information for those clients to which they are connected.

3.1.2 Standard Edition

Standard edition creates a highly reliable network utilizing multiple WAN transport modes and ensures that each application takes the best path to achieve the highest application performance, which results in the optimal end-user experience.

With the real-time WAN link measurement capabilities of the solution, if individual WAN path conditions change, Citrix SD-WAN quickly detects anomalies on the path and seamlessly diverts traffic across healthier paths that are also being monitored. Condition changes are detected within 2 or 3 packets, so the appliance is able to guarantee a pretty high resilience. The experience is so smooth users won't even realize any changes have occurred.

How does packet processing work in Standard Edition? To get started, the network is segmented in three areas: the client LAN, the virtual path or LAN and the Data center network. Each packet has a source/destination IP that is destined to be changed when it crosses the virtual path. The SD-WAN device has two WAN paths available that it could choose to make the data pass through it. In practise, what it really does is to make an association of virtual IPs (VIP) at the entrance and at the exit of the virtual paths. Every single TCP/UDP traffic that travels from client LAN to server LAN gets encapsulated in the UDP tunnel between the two peers across the WAN segment and must be identical to what has been transmitted. So, the key to have an SD-WAN performing well is to have a good and reliable VIP to VIP mapping.

3.1.3 Wanop Edition

The main novelty of the Wanop edition is about TCP flow control. The common algorithms have an utilization of the link which is much lower than the actual capacity of the link. As for the Citrix SD-WAN solution, it implements an improvement of the TCP flow control such that the exploitation of the available WAN link is maximized.

The present solution, whose scheme is depicted in figure 3.2, aggressively fills the available WAN link to its measured capacity and does not perform backoff when loss or congestion is encountered. Packets are buffered and re-transmitted as needed to keep an average link usage nearly maximum capacity.



Figure 3.2. Citrix SD-WAN Wanop edition architecture, reproduced from [17]

Wanop edition also leverages Adaptive Compression Technologies. They work between appliance peers residing on opposite ends of the WAN to reduce bandwidth requirements. Wanop adopts several standard compression schemes to reduce the size of data as it traverses and propagates across the wide area network. The SD-WAN device also maintains a compression history that is shared across different sessions.

Among the novelties of the Wanop edition, it can be also reckoned the Data

Deduplication technique, thanks to which data sent earlier by one connection can be used later to optimize traffic flow over another connection. Small data streams are stored in memory for lower latency access, whereas large data streams are stored to disk. The mentioned approach is able to furnish application acceleration, data reduction and advanced flow control.

3.1.4 Premium Edition

Citrix SD-WAN Premium edition combines the features of the other editions into a single appliance solution. It takes advantage of multiple WAN links providing a larger bandwidth for application delivery. WAN overlay also enables the ability to aggregate the bandwidth of multiple WAN links. Premium edition also concerns with optimization techniques: they tokenize data payload across the tunnel, freeing up bandwidth and improving the performance of delivery. These innovations mitigate the effect of latency and end up in a turbo-charged end-user experience.

3.1.5 Advanced Edition

The last Citrix SD-WAN release (11.4) includes a fourth edition that is intended for introducing additional functionalities never explored before. As a matter of fact, Advanced Edition enables Edge Security functionality on Citrix SD-WAN appliances. It also includes the following security capabilities: web filtering, anti-malware and intrusion prevention [18].

3.2 Citrix SD-WAN VPX SE

Citrix SD-WAN VPX SE is a virtual appliance that can be hosted on several platforms like Citrix XenServer, VMware ESX, Microsoft Hyper-V or AWS. It is a virtual machine, so it is possible to deploy whatever hardware, or it can be

developed in combination with other virtual machines (servers, VPN units, or other appliances), to create a unit that precisely matches the user's needs.

An SD-WAN appliance acts like a virtual gateway, not like a router or a TCP endpoint. This means that its job is to buffer packets and put them onto the outgoing link at the right speed. The packet forwarding phase can be done in several ways: Inline mode, WCCP mode and Virtual Inline mode. If the SD-WAN appliance supports all of them, the choice is made automatically by the Ethernet or IP header. Let's investigate them further taking into account that they are defined for physical devices but they act the same also in virtual devices.

3.2.1 Inline mode

As shown in figure 3.3, in Inline mode the appliance mimics an Ethernet bridge. It transparently accelerates traffic flowing between its two (virtual) Ethernet ports and requires the most straightforward configuration.

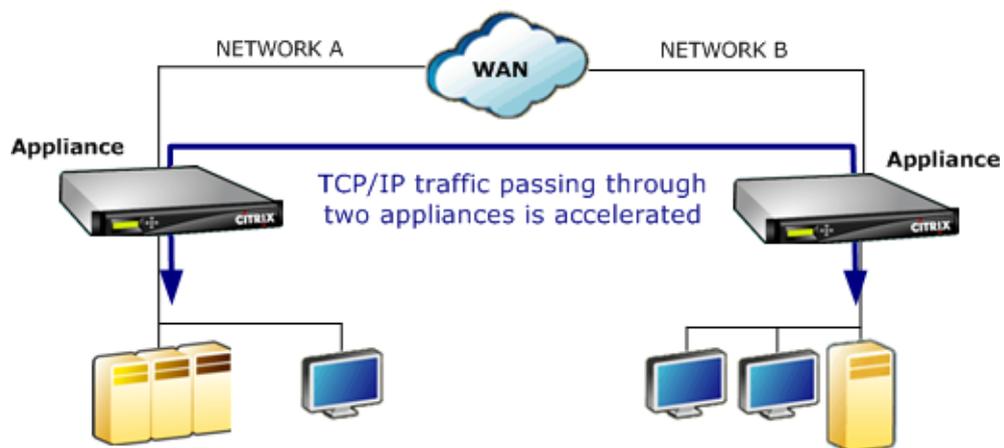


Figure 3.3. Inline deployment mode architecture, reproduced from [19]

The operating principle of the inline mode requires both the appliances to be placed between the LAN and the WAN. Traffic passes into one of the appliance's Ethernet ports and out of the other. When two sites with Inline-based appliances

communicate, every TCP connection crossing them is accelerated. All the other data is passed transparently, like there is no appliance.

3.2.2 Web Cache Communication Protocol

Web Cache Communication Protocol is a dynamic routing protocol designed by CISCO and useful whenever asymmetric routing occurs. Asymmetric routing is one of the main novelties of the SD-WAN paradigm, because it allows the data and its corresponding ACK to follow two distinct paths without being blocked by the firewall. In WCCP mode, routers divert traffic through the appliance by using the WCCP 2.0 protocol. Then, the traffic is processed within the appliance and treated as it arrived in inline mode. The WCCP appliance is connected directly to a dedicated port on the WAN router, as can be deduced from figure 3.4.

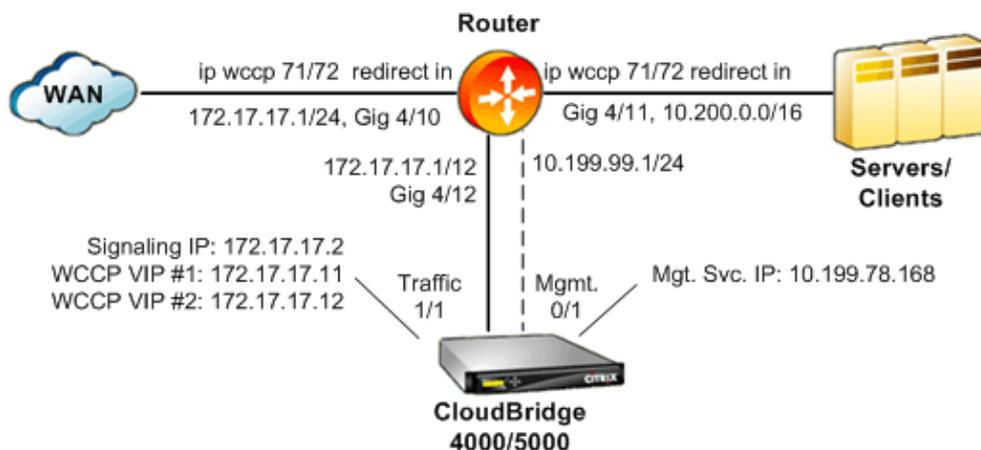


Figure 3.4. WCCP deployment mode architecture, reproduced from [20]

3.2.3 Virtual Inline mode

Virtual inline mode entails that a router sends WAN traffic to the appliance and the appliance returns it back to the router, as it is clearly highlighted in figure 3.5.

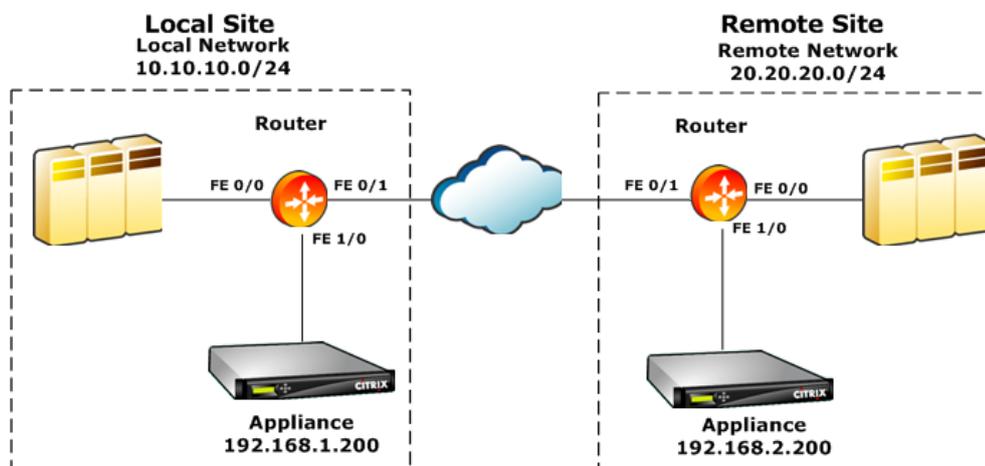


Figure 3.5. Virtual inline deployment mode architecture, reproduced from [21]

The appliance seems to act like a router, but it does not exploit routing tables. Nevertheless, Virtual Inline mode should be used only when Inline and WCCP are impractical. In this mode, router takes decisions on policy based routing, i.e., the router sends information to the appliance for acceleration and the appliances do the opposite for forwarding. Hence, it is pretty difficult to troubleshoot it. In fact, it does not contain health-checking or status-monitoring information.

3.2.4 VPX usage scenarios

As stated before, the operating principle of the Citrix SD-WAN virtual appliance is more or less the same of the physical device. Certainly, there are some differences: for example in the deployment modes, since physical devices also support Group mode and High Availability mode and, in addition, virtual appliances are not able to endorse Ethernet bypass cards and multiple accelerated bridges. Despite the lack of the mentioned functionalities, there are many usage scenarios for which the VPX release can be employed:

- Branch-office accelerator: the Citrix VPX image is installed in whatever server

involved in the infrastructure (figure 3.6) and deployed like an SD-WAN/SD-WAN appliance. Not only VPX implements all the functionalities of a physical SD-WAN device, but also it brings the advantages of virtualization;

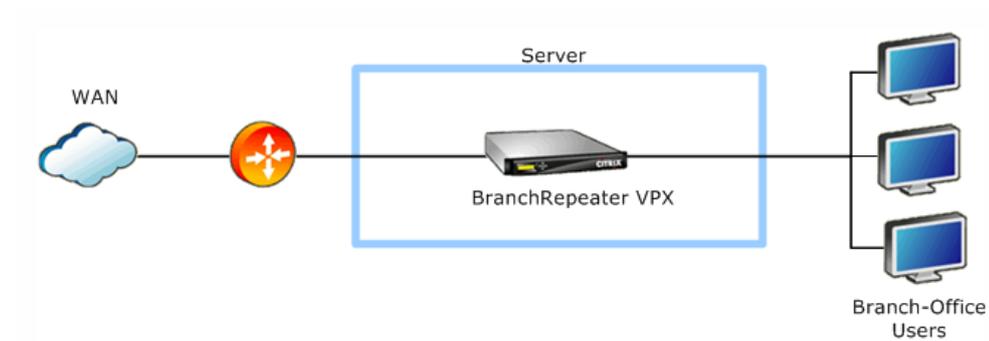


Figure 3.6. Branch-office accelerator usage scenario, reproduced from [22]

- Accelerated branch-office server: this configuration requires the insertion of a virtual branch-office server in the server portion, as can be observed in the blue box of figure 3.7.

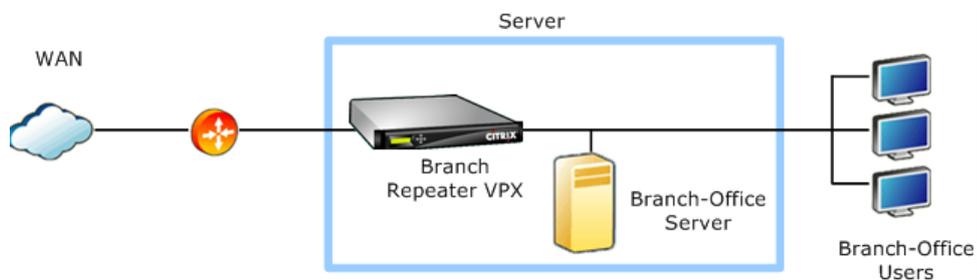


Figure 3.7. Accelerated branch-office server usage scenario, reproduced from [22]

The WAN traffic is accelerated automatically only if the virtual networks within the appliance hosting the virtual machines are assigned in such a way the path to the WAN passes through the virtual SD-WAN devices. What results is the setup of a virtual environment where the server can be configured to support any functionality. And besides, there is the possibility to implement

multiple virtual servers on the same machine, consolidating the branch-office rack as a single unit running multiple virtual machines;

- Accelerated data center servers: figure 3.8 exhibits the infrastructure of data centers employing the VPX software solution.

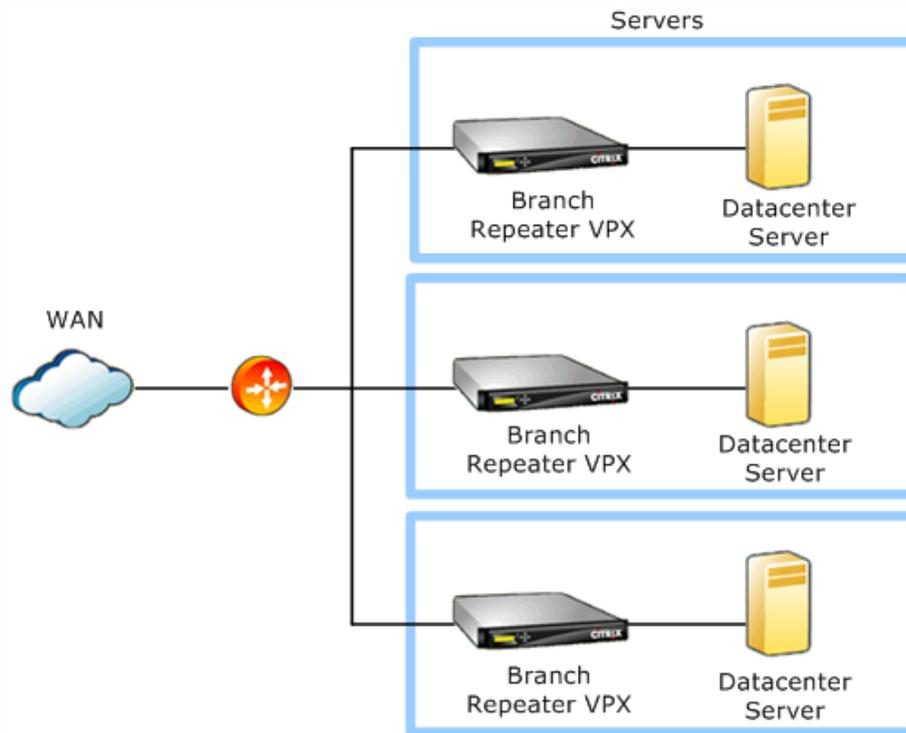


Figure 3.8. Accelerated data center servers usage scenario, reproduced from [22]

The idea behind this layout is to install VPX virtual machines on every server of a data center. This procedure ends up in a solution that scales perfectly as server capacity is added and while minimizing the number of servers by adding acceleration to them. Once propagating this principle in a significant number of servers, the aggregated acceleration exceeds any other trial that can be provided with a single appliance;

- VPN accelerator: in this scenario, by simply installing the VPN within the

VPX software, the outcome is an accelerated VPN. However, it should be kept in mind that the VPN virtual machine resides on the WAN side, whereas the VPX one is on the LAN side, as can be deduced from figure 3.9. This is due to the need for decryption and compression of the VPN traffic;

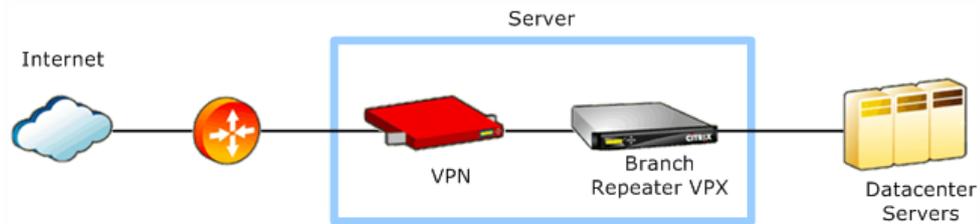


Figure 3.9. VPN accelerator usage scenario, reproduced from [22]

- Multiple VPX instances on the same server: another practical application can be the installation of multiple VPX virtual machines on the same server, as shown in figure 3.10.

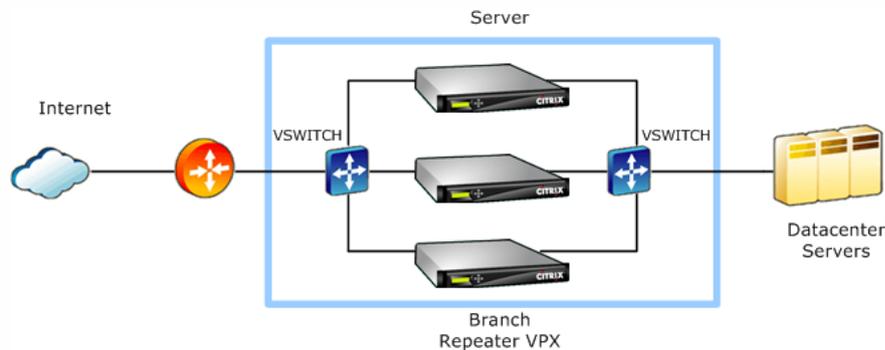


Figure 3.10. Multiple VPX instances usage scenario, reproduced from [22]

The mentioned pattern allows the creation of different types or levels of acceleration services within the same unit. Each VPX instance can be dedicated to a specific task, e.g. a critical application, or to an individual site or customer. The traffic is routed to the appropriate VPX instance by means of a VLAN switch, visible in the rightmost side of figure 3.10;

- WCCP and virtual inline deployment: due to the lack of the Ethernet bypass card, WCCP is suitable for those situations that would have required it. Actually, WCCP mode provides fault-tolerance because it implements built-in health-checking. Routers are in charge of sending traffic directly to the end point, in order to avoid crossing an unresponsive WCCP device, through the architecture depicted in figure 3.11.

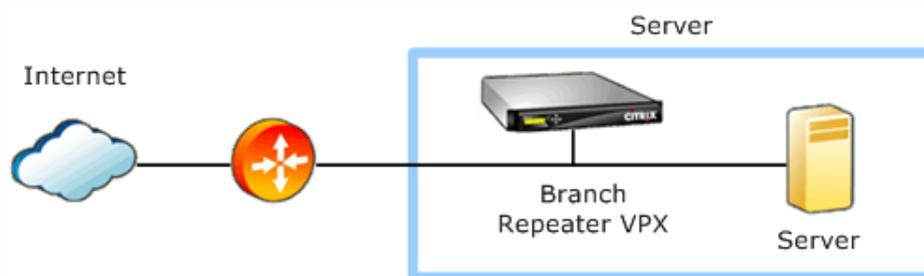


Figure 3.11. WCCP and Virtual inline deployment modes, reproduced from [22]

3.3 How Citrix SD-WAN VPX SE works

Citrix SD-WAN VPX SE solution works exactly as it was a physical appliance. Its working principle is based on the concept of building a tunnel among the connected locations. The virtual tunnel that is installed between the two connected sites is a logical path where packets are sent according to the UDP rules. Indeed, packets contain a portion dedicated to data and another one which is attached at the tail and carries the connection statistics. Thanks to them, the MCN informs every node in the topology about the best path to follow and routing is done on a software level.

The current Citrix SD-WAN virtual appliance also allows asymmetric routing. Asymmetric routing is a technique that relies on the possibility of exploiting distinct channels for sending information and receiving the corresponding ACK. Normally,

the firewall at the sender would stop and discard an ACK, coming from channel 2, of a packet sent through channel 1, due to port numbers. SD-WAN overcomes this issue and allows a more dynamic routing. Consequently, by exploiting this technique a packet sent through a fiber (e.g.) can be acknowledged from an ACK sent back through an MPLS link, without the firewall blocking this procedure.

Chapter 4

Simulations and results

4.1 Network infrastructure

The network infrastructure that has been employed in the simulation is composed by:

- a headquarter and a branch office (server 1 and server 2);
- two network emulators that are in charge of emulating two different kinds of WAN connections;
- two SD-WAN appliances, one is the Master Control Node (MCN) and the other is the client.

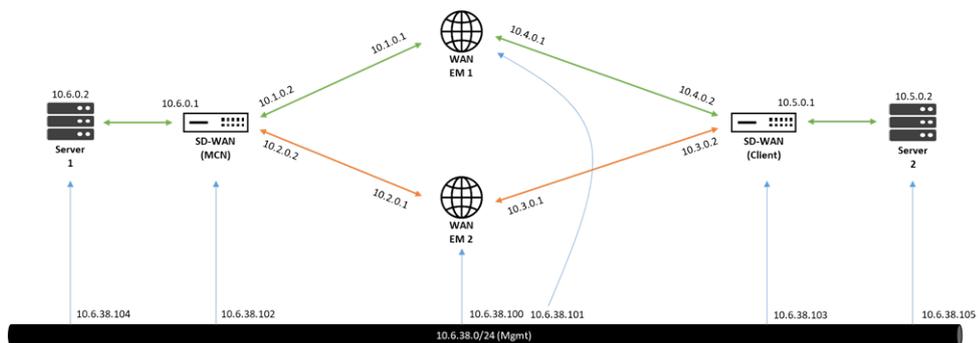


Figure 4.1. Network infrastructure for the experiments

The aim of this network is to show what are the benefits and the limitations of the SD-WAN approach. The network infrastructure is entirely constituted by virtual appliances, in the sense that they are obtained installing some software on virtual machines that stand on the VMware ESXi bare-metal hypervisor. The infrastructural scheme of the network is depicted in figure 4.1.

As highlighted through the differently coloured arrows, there are four sub-nets between the SD-WAN appliances, i.e., 10.1.0.0/24; 10.2.0.0/24; 10.3.0.0/24 and 10.4.0.0/24. The channels that span from one SD-WAN to the other are essentially two, but each of them is subdivided into two sub-paths that have the WANem as endpoint. WANem acts like a router because it is in charge of routing packets coming from one sub-net to the other. Then, there is an additional sub-net which is dedicated to the Management interface, with network address 10.6.38.0/24. The IP addresses employed are two: one is the service IP and is handy in coping with the service provided to the user, while the second one is the management address and it is exploited for all the management activities. All the devices must be equipped with an interface for interacting with the management IP.

4.2 Configuration

The baseline of the simulation relies on the configuration of each component of the network, that differs from the others. Each component of the network infrastructure is a virtual machine installed on the VMware ESXi hypervisor, which is a bare-metal hypervisor resting on VMware servers.

4.2.1 WAN emulator configuration

The WAN emulator persistent configuration consists of realizing two virtual machines, one for the MPLS circuit and one for the ADSL link. Both of them feature two interfaces, and the configuration is based on assigning the correct IP addresses

to each interface in order to let the WAN emulator communicate both with the SD-WAN appliances and with the management IP address. The WAN emulators must be endowed with a static IP address, to be added manually, because it could become a problem for the user to set a new IP address at each new connection, and this also guarantees a higher level of security.

WANem is not directly derived from the ISO file, because it would imply losing the configuration at each new boot. So, the installation is moved to the virtual disk in order to ensure a persistent setup. In addition, there is a second configuration phase that is directly done via the WANem GUI. Indeed, once connecting to the IP address of the WANem virtual machine, it is possible to properly set the connection parameters as desired by the user. In the current testing scenarios, the parameters that have been modified are: the delay, which is a fixed amount of latency added to the transfer; the jitter, that, instead, represents a variation of the delay at each transfer and the losses, which establish the probability of dropping a packet that enters the WAN emulator.

4.2.2 Server configuration

The server configuration concerns with the installation of two virtual machines provided with Windows operating system. The servers involved in the network infrastructure are treated like two computers standing one in Milan and one in Turin, with the intent of emulating a WAN connection. The server configuration follows the same steps of the WANem one, i.e., each machine has its IP address, that must belong to the already established sub-nets and must also communicate with the management IP address. After configuring each interface with the correct IP address, the server can be simply obtained by pushing an ISO file that resembles a Windows operating system. There are additional setup steps that must be followed, like the assignment of a virtual disk space and the provisioning. For the current work, it has been established to choose the "Thin provisioning" technique, that

consists of dynamically allocating the space for the machine as required. Thus, the maximum disk space is not allocated immediately, but it is destined to grow as the virtual machine starts increasing the needed capacity during its activity cycle. Moreover, thanks to the customized configuration, it has also been possible to choose the number of socket of each machine, set to 2, and the number of cores assigned to each socket, again 2.

4.2.3 SD-WAN configuration

In this phase there are many steps that must be followed. After installing the VMware vSphere Client, the very first stage of the configuration concerns with the deployment of the Citrix OVF template that has to be uploaded on the created virtual machine. According to what imposed before, the management IP is set through VMware shell and the following step is the configuration of the graphical interface. When entering the SD-WAN configuration editor, it is necessary to define the first appliance as the MCN and then to add a new configuration creating two sites, one related to the branch office (BO, set as client) and one to the headquarter (HQ). The management interface is not included in the interfaces to configure, so there is no need to add it. The necessity to have static IP regards only the main node (MCN) because every device in the network must know how to reach it, whereas remote nodes could receive any IP address.

The configuration of the SD-WAN appliance is done only on the MCN node, because later it is compacted inside a packet and then it is sent to the BO.

After the interfaces, it is fundamental to add WAN links - the ones that allow internet navigation - the MPLS link (private intranet) and the ADSL one (public internet). The WAN links' definition implies a fundamental stage: the assignment of the capacity, that must be reasonable and calibrated to the host's capacity. For the current tests, the MPLS channel and the ADSL one are both endowed with 6 Mb/s, as reported in figure 4.2. To better understand the nature of these links, it

has been decided to highlight in orange the LAN-to-WAN links and in green the WAN-to-LAN ones. It turns out from the picture that, by summing the capacity of the LAN-to-WAN and WAN-to-LAN links on the same path, the overall capacity reaches 12 Mb/s. So, the virtual paths instituted between the SD-WAN appliances, in detail, the black arcs in figure 4.2, lean on a 12 Mb/s bandwidth.

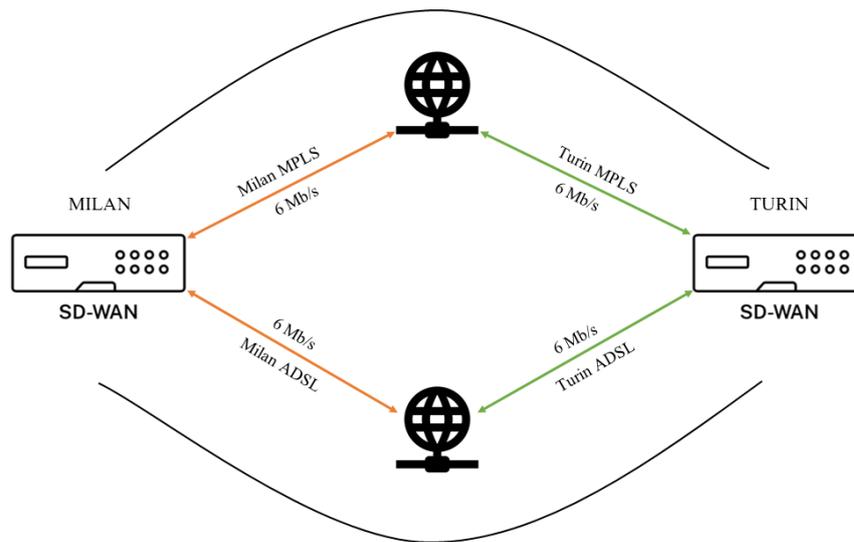


Figure 4.2. MPLS and ADSL WAN links

When adding and labeling the WAN links, it is possible to add also some features. For example the eligibility of the link, that means for which kind of traffic the link is intended. There is another important feature that must be set, the "bypass mode", that is the only one suitable for a virtual appliance. The bypass mode has the role to build a new virtual machine whenever the one that is on crashes. In terms of security, since the MPLS link is a dedicated circuit, its security mode has been set to "Trusted" because it relies on the exploitation of a firewall between the SD-WAN appliance and the WAN emulator. If employing 4G/5G links, instead, it would have been chosen the Non-trusted security mode due to the lack of the firewall.

Once the configuration has been saved, it has to be exported in a package useful for installing the configuration itself in the second virtual appliance, the one that resides on the other side of the network. The procedure implies to download the "software package" and to install a driver inside the packet that is sent in broadcast to all other nodes. The driver has the role to update the version of the appliance's software if it is different from the one of the MCN.

4.3 Citrix SD-WAN appliance

Citrix SD-WAN VPX GUI is composed by 3 sections: Dashboard, Monitoring and Configuration. The first section gives general details about the version employed, the IP addresses, the software release and the hardware support. The monitoring section, instead, is plenty of information about what is currently happening to the WAN links. Last, the configuration section allows to change the current setting which must then be propagated to all other SD-WAN nodes.

4.3.1 Configuration

The section that must be compiled first is the Configuration one. Indeed, it is here that it is possible to select which is the appliance mode: MCN or Client. The Master Control Node is the most important in the network infrastructure because it keeps the statistics of the connection and it establishes the rules to install in each device. Once the MCN mode has been selected, the SD-WAN sites with their WAN links must be properly added. It is required to specify both IP addresses of each single interface and also the sub-nets employed.

4.3.2 Monitoring

The most powerful and interesting tool that the Citrix SD-WAN appliance offers is the monitoring section. that keeps track of all the statistics at the MCN. As

depicted below in figure 4.3, the statistics part involves the analysis of each link in a detailed manner. As for starting, each link is counted twice because it is analyzed in both directions, and each link is called "Path". The logical path that connects the two sites, also known as tunnel, is then called "Virtual Path". The key point of the statistics is to detect whether the path state is good or not and also the reason behind this phenomenon. Indeed, Citrix SD-WAN creates a reliable WAN from diverse network links, continuously measuring and monitoring each link for loss, latency, jitter and congestion.

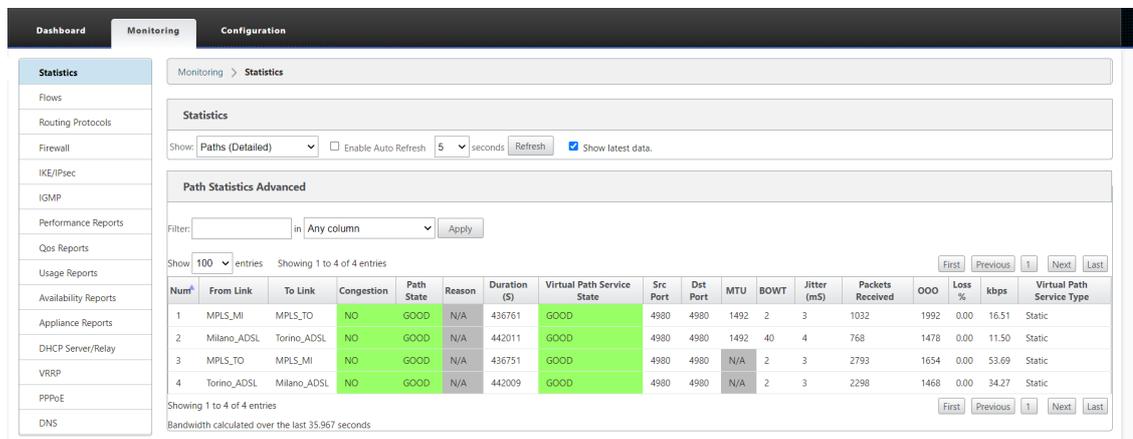


Figure 4.3. Citrix SD-WAN Monitoring graphical interface

The path state could be: GOOD, if it does not encounter any significant issue; BAD, due to losses it is assigned with a path score penalty, or DEAD, if there is no more connectivity. The reasons behind the route status change usually are: PEER, that means the remote state path changed either to Bad or Dead; SILENCE, whenever no packets are received and LOSS, when there are some missing packets.

Additionally, the statistics highlight several aspects of the connections, like the OOO, i.e. the number of out-of-order packets; the BOWT, that is the latency best one-way time, and also the MTU, which represents the maximum size of the packet that can be forwarded.

But there is more beyond the statistics portion: the Citrix SD-WAN device furnishes information about several factors of a connection. As an example, another relevant section of the SD-WAN monitoring is the QoS Report, visible in the left-most part of figure 4.3. QoS Report offers information about the kind of traffic involved in the test, and it categorizes it into four classes: bulk, interactive, real-time and control. Bulk concerns with any application that does not need a rich user experience but is more about moving data (i.e. FTP or backup/replication). Interactive is the broadest category and refers to any application that has a high degree of user interaction. Some of these applications, for example, video conferencing, are sensitive to latency and require high bandwidth. Other applications like HTTPS, may need less bandwidth but are critical to the business. real-time traffic is concerned as VoIP or VoIP like applications, such as Skype or ICA audio. In general, it refers to voice only applications that use small UDP packets, that are business critical. Control traffic, instead, is used to transfer control packets that contain routing, scheduling, and link statistics information.

Furthermore, there is the Performance Report section, that shows performance statistics at the site, virtual path, or Direction (LAN-to-WAN and WAN-to-LAN) level. With Citrix SD-WAN, it is possible to collect metrics that show the efficiency of each link in milliseconds.

The usage reports segment, instead, supplies charts about the speed of the virtual path between the sites. It allows the selection of the path direction (LAN-to-WAN or WAN-to-LAN) and the plot type.

4.3.3 Dashboard

The dashboard is the less relevant section because it does not provide extra information about the connection status. Although, it contains essential information about the product itself, like the software version edition; the appliance model, that could be either virtual or physical; the appliance mode, either MCN or client;

and also the management IP address.

4.4 WANem

WANem is a Wide Area Network Emulator, meant to provide a real experience of a Wide Area Network or Internet, during application development/testing over a LAN environment. WANem provides emulation of Wide Area Network characteristics and thus allows data/voice applications to be tested in a realistic WAN environment before they are moved into production at an affordable cost. [23]. Thanks to WANem, it is possible to emulate the behavior of different kinds of WAN connections. The two connections employed for the current simulation are: MPLS and ADSL. MPLS is a layer 2.5 protocol according to which telco providers realize a dedicated circuit where nothing except user data travels. The ADSL link, instead, merges together the user data with the traffic following the same set of routers and the level of protection is extremely lower than the MPLS one. Also the quality of service perceives this impact: data is exposed to losses, collision and physical links' failures much more than what happens in an MPLS scenario. By the way, since working in a virtual environment, the conceptual differences among the two technologies do not hold and they basically behave equally. However, thanks to the usage of the graphical interface reported below in figure 4.4, it is possible to change settings as desired by the user and to differentiate between the two WAN links. In real applications, MPLS connection is much reliable since it is based on a dedicated circuit where only the data intended for the user are allowed to pass. So, its performances are stable and they can never overcome a lower bound. ADSL connections, instead, are pretty much unpredictable and less secure, since they let data pass according to many possible paths which are established by routers along with their routing tables.

WANem commands successfully created
WANem is running

Interface: eth2		Packet Limit 1000 (Default=1000)				Symmetrical Network: Yes		
Bandwidth	Choose BW	Other				Other: Specify BW(Kbps)		
Delay		Loss		Duplication		Packet reordering		
Delay time(ms)	40	Loss(%)	80	Duplication(%)	0	Reordering(%)	0	
Jitter(ms)	10	Correlation(%)	0	Correlation(%)	0	Correlation(%)	0	
Correlation(%)	0					Correlation(%)	0	
Distribution	-NA-					Gap(packets)	0	
Idle timer Disconnect	Type: none	Idle Timer				Disconnect Timer		
Random Disconnect	Type: none	MTTF Low		MTTF High		MTTR Low	MTTR High	
Random connection Disconnect	Type: none	MTTF Low		MTTF High		MTTR Low	MTTR High	
IP source address	any	IP source subnet		IP dest address	any	IP dest subnet		
							Application port if any	any

Display commands only, do not execute them

Figure 4.4. WAN emulator graphical interface

4.5 Experimental evaluation

The infrastructure adopted for simulation purposes involves two different virtual channels: the first one emulates an MPLS connection and the second one an ADSL one. Practically speaking, they are perfectly equal since they are not physical wires but only virtual emulators. However, thanks to the WANem GUI, the parameters can be tuned according to user purposes and also for observing how the SD-WAN appliances react to changes in the network status.

The tests that have been run are focused on two essential kinds of traffic: the traffic employed for downloading a huge video file thanks to the HTTP protocol and for reproducing it through VLC platform, and another kind of traffic destined to statistics and generated by the Jperf tool by means of iperf commands.

The aim of the subsequent simulations is to gather some interesting results about the capabilities of the SD-WAN devices in separate situations, each one different from the other in terms of level of criticism. For this reason, several scenarios have been analyzed changing every time the tasks and the parameters' setting.

4.5.1 Scenario 1: Jperf traffic

The very first trial involves the installation of the Jperf tool on both sites, where Jperf is a software intended for measuring the actual throughput of the network by injecting some bulk traffic inside it. This is the easiest scenario: the Milan site acts like a client, which by default sends data, and the Turin one is the server, that, instead, receives the traffic from the client. The network status is unaffected, the 6 Mb/s of the ADSL channel and the 6 Mb/s of the MPLS one are totally available, so it can be expected to observe a throughput that settles around 12 Mb/s overall. Moreover, no extra traffic interferes with the Jperf one, so it should grab the whole bandwidth.



Figure 4.5. Simplest scenario: only Jperf traffic

As shown in figure 4.5, after the initial ramp-up period, the throughput becomes more or less stationary and it settles around 11.5 Mb/s, as it can be stated by clicking on the chart in the graphical interface. In other words, when the entire bandwidth is at the disposal of the Jperf tool, it fully leverages it and no other kind of traffic is generated. In fact, according to the Citrix dashboard, the traffic

injected in the network by the Jperf tool is classified as bulk. In addition, bulk class has the lowest priority and can be considered a scrap class.

4.5.2 Scenario 2: Video download and Jperf

The second test is trickier since it implies sharing resources among two distinct applications: one is Jperf and the other one is the downloading of the video file from the browser by means of HTTP protocol. By typing on the HTTP bar "10.6.0.2/down.msu.zip" which is the server address with the zip file containing the video, the downloading procedure starts. This peculiar kind of HTTP traffic is classified as real-time, and it is the one with the highest priority. Later, after letting the browser complete its transfer, it has been tried to open a Jperf session. What surprises is that the Jperf test is not even able to start, so thanks to Citrix Dashboard (fig 4.6) it has been carried out why. Indeed, by looking at the bandwidth that the downloading procedure grabs, it has been discovered that it steals all the available capacity, giving no room for the rest.

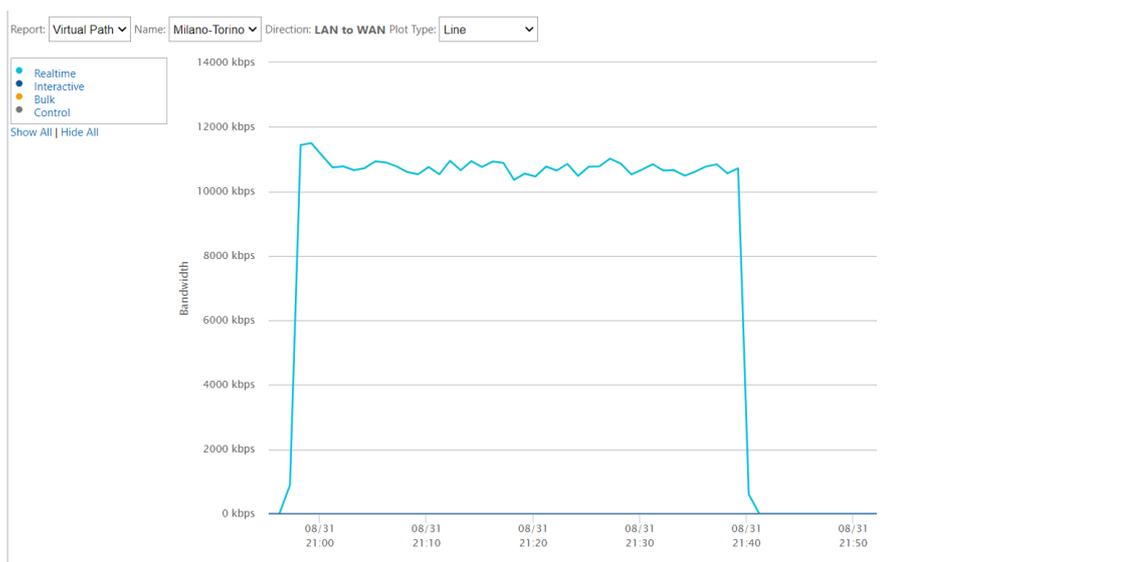


Figure 4.6. Video downloading bandwidth exploitation

After finding out what are the reason behind the fullness of the links' capacity, it has been tried an alternative approach aimed at succeeding in sharing the bandwidth among the two traffic patterns. Therefore, in the Google developer settings it has been created a custom Throttling profile by imposing a downloading speed, that has been set to 4 Mb/s in order not to saturate the capacity. In this way, only 4 Mb/s out of the 12 available are dedicated to the download and the remaining can be employed for the Jperf test. So, the expected splitting of resources is: around 4 Mb/s for the real-time traffic and around 8 Mb/s for the bulk one.

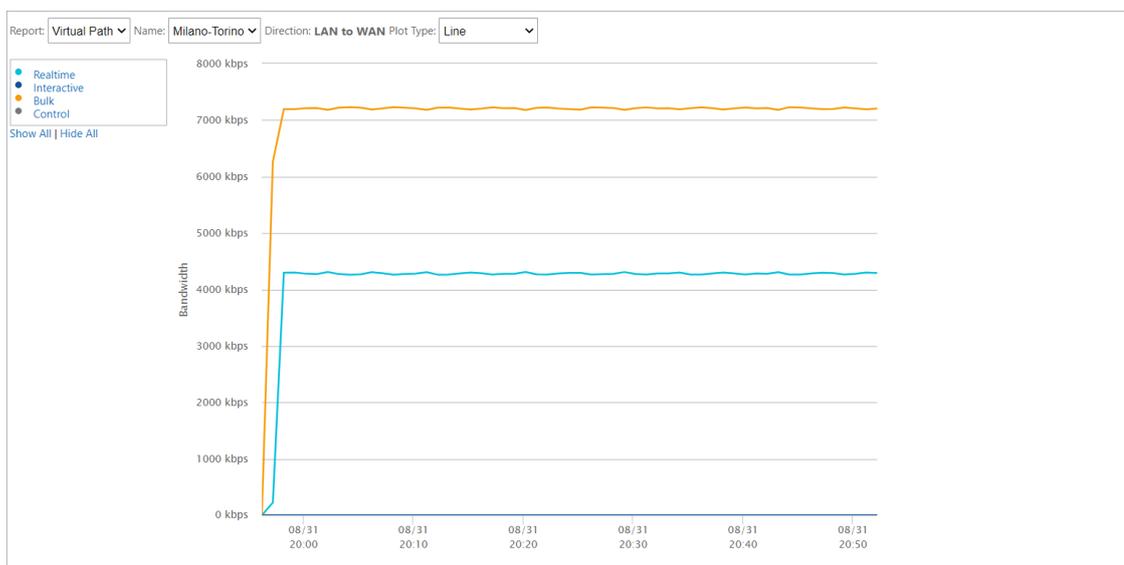


Figure 4.7. Video downloading with throttling profile and Jperf

Nevertheless, from figure 4.7 it turns out that the results are slightly oscillating because the download speed is a little bit higher than 4 Mb/s and, consequently, the Jperf one is slightly less than 8 Mb/s. Overall, the behavior of the network under the constrained conditions meets expectations.

4.5.3 Scenario 3: Video streaming and Jperf

The present experiment is divided into two essential steps: the first step is dedicated to the reproduction of the previously downloaded video by means of VLC media player, with the intent of verifying which is the bandwidth that it requires for buffering the file and then playing it back. This stage is essential for understanding how much bandwidth is required by VLC to buffer and then to play the video. Indeed, it turns out from figure 4.8 that this process does not imply the usage of the whole available bandwidth, but it exploits only up to 2.2 Mb/s.

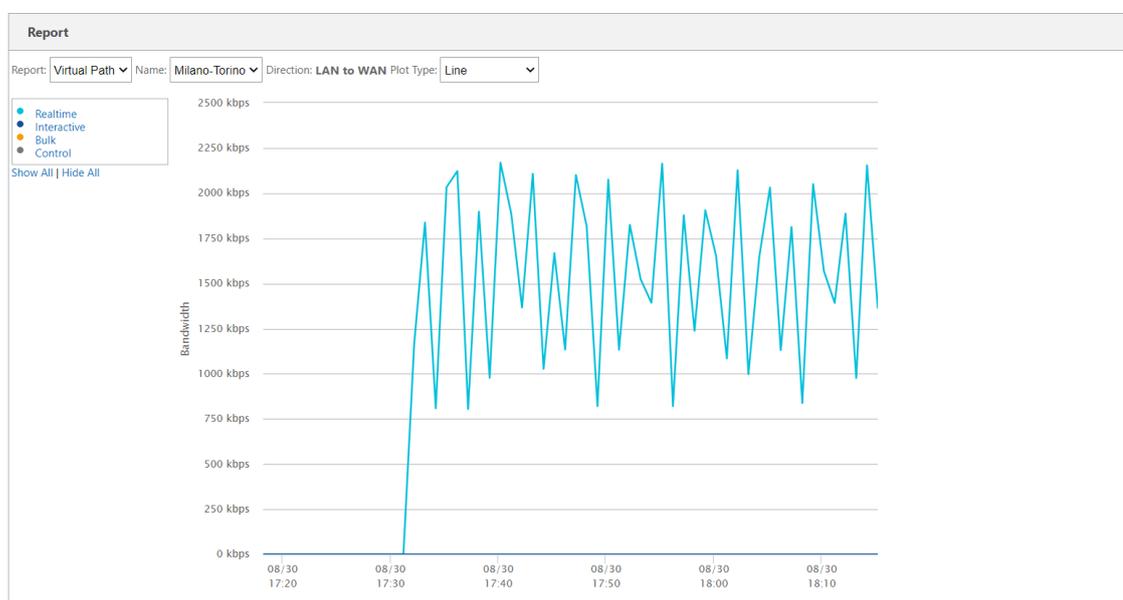


Figure 4.8. Video streaming bandwidth exploitation

Nevertheless, figure 4.8 also shows that there are significant fluctuations in the exploited bandwidth, and this can be explained by the fact that the video is reproduced continuously, since the Repeat option has been enabled. Once done it, every time the video streaming ends, it should begin again by buffering the content, up to when the file has been totally charged and it can be reproduced. This is a typical functionality of VLC media player and motivates the oscillation of the bandwidth. Actually, the bandwidth exploitation reaches the maximum whenever the software

needs to buffer the file and later, when it is fully charged, the bandwidth needed for playing it drastically reduces. Another remarkable detail is that now the traffic is no more bulk, but it is marked as real-time. It is equipped with the highest priority and gets up to 50% of the overall scheduler time.

As for the second stage of the experiment, it is centered around a more complex task: this time the bandwidth is shared among two applications, Jperf and VLC. As deducible from the above results, for sure VLC will grab more or less 2 Mb/s out of the 12 Mb/s available, so the remaining throughput that the Jperf tool can experience is reduced up to a maximum of (around) 10 Mb/s, including some overhead.

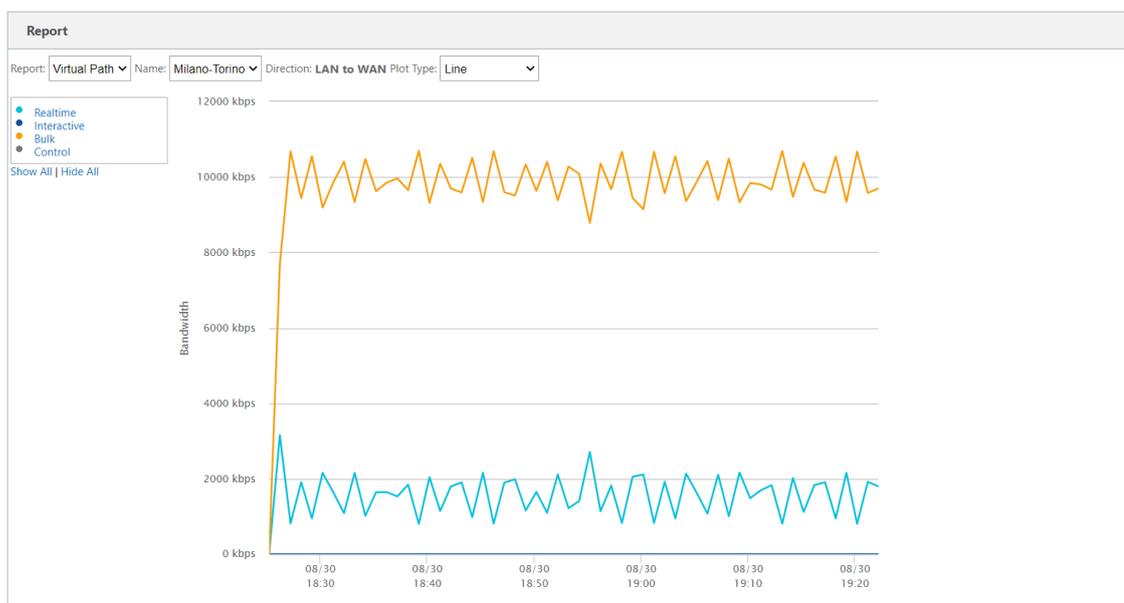


Figure 4.9. Video streaming and Jperf bandwidth sharing

The results of figure 4.9 confirm what stated before: the sharing of resources is responsible for a bandwidth allocation of 2 Mb/s for VLC and of 10 Mb/s for Jperf bulk traffic. The present outcome highlights a fundamental aspect, that is, in this specific case, the bandwidth allocation is realised despite the different levels of priority, because the two amounts of traffic can coexist without any drop.

4.5.4 Scenario 4: Download and video streaming

The current scenario is an advanced experiment that does not involve anymore bulk traffic, but two kinds of data referred as real-time traffic: one is the HTTP download from the Milan server and the other is the already mentioned video streaming. The aim of this trial is to highlight one of the most remarkable properties of the Citrix SD-WAN appliance, which is the degree of freedom of the user that interacts with the controller. Going deeper into detail, there is a configuration editor within the master control node appliance where many tables are present. The tables of interest for this experiment are the "Rule" tables, where to each protocol is assigned a level of priority that spans from 0 to 17 in descending order. It has been established that the HTTP traffic destined to the downloading procedure must have the highest priority in this experiment, so its instance in the Rule table has been edited such that it has "HDX_priority_tag_0". In this way, the downloading speed is supposed not to be interrupted by the video streaming traffic.

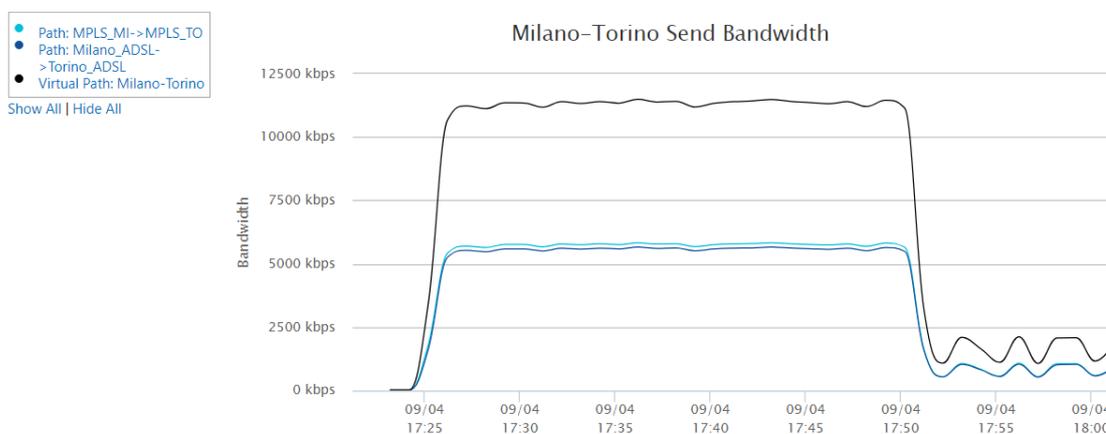


Figure 4.10. HTTP downloading and video streaming

Nonetheless, since dealing with two real-time applications, it is impossible to distinguish from the charts realized in the previous sections which are the specific protocols involved. Figure 4.10 simply allows to differentiate the traffic that crosses the MPLS channel from the one that uses the ADSL one. As expected, the traffic is

almost perfectly balanced between the two WAN links and then sums up to give a better overall performance. However, the monitoring tool of the appliance also lets the user know which protocols have been exploited during the monitoring session.

The screenshot shows a web interface titled 'Applications Statistics'. It includes a filter field, a 'Show 100 entries' dropdown, and a table with the following data:

Application	Family	Bytes Received	Bytes Sent	Total Bytes
HyperText Transfer Protocol	Web	348707033	14138937211	14487644244
ISO Base Media File Format	Audio/Video	293533793	18601125715	18894659508
Unknown Protocol	None	0	0	0
iperf	Network Management	1280781520	111048306218	112329087738

Figure 4.11. Protocols employed by each application

Actually, figure 4.11 shows the three fundamental protocols involved during the tests: for the downloading procedure it has been employed the HTTP protocol, for the video playback the ISO Base Media File Format, which defines a general structure for time-based multimedia files such as video and audio [24] and the well-known and already discussed in the previous trials iperf.

4.5.5 Scenario 5: Download with throttling profile and video streaming

The previous experiment can be further improved by creating a custom throttling profile in the Google Chrome browser for exactly establishing the downloading speed of the file of interest. Since it has been discovered that the bandwidth needed by the video playback procedure amounts to 2 Mb/s more or less, it is possible to create a throttling profile with a downloading speed of 9 Mb/s such that this procedure is faster than the case analyzed in section 4.5.2. As described in paragraph 4.5.2, the video file download is repeated with the same routine, but at a different speed. In the latter case, since dealing with two bunches of real-time traffic, it is only possible to see the summation of the exploited bandwidth from figure 4.12.



Figure 4.12. Video streaming & download with throttling profile

As expected, the light blue line reaches more or less a bandwidth exploitation of 11 Mb/s, obtained by summing 9 Mb/s dedicated to the downloading procedure to 2 Mb/s needed by the video playback. What differs from the already described tasks, relies on the fact that the knowledge of the capacity needed by one application, allows to gather a full bandwidth exploitation. Indeed, reaching almost the whole capacity exploitation in this experiment depends on a meticulous parameter setting that ends up in a pretty good resource utilization. Another aspect that has to be taken into account is the interval of time for the completion of the downloading procedure: it took 31 minutes to complete, against figure 4.7 where it took more than one hour with a downloading speed of 4 Mb/s. Hence, this configuration allows not only to achieve a worthwhile resource utilization, but it also lets the user save time in completing both the operations.

4.5.6 Scenario 6: Network failures with video streaming and Jperf

Up to now, what has been examined, is the behavior of the network in the case both the channels are fully performing. But, what is really of interest in this study is the behavior of the SD-WAN device during anomalies of the network components. First of all, the current test involves real-time traffic (video streaming) and bulk traffic (Jperf). However, there is a significant difference with respect to section 4.5.3, that is, the quality of both the channels has been degraded through the WANem GUI. In fact, it has been set the 10% of losses in the Loss Box of the WAN emulator, with the intent of testing how much responsive to failures the appliance is. The first variation is visible from the Statistics section of the Citrix SD-WAN GUI (figure 4.13), where the state of the channel turns out to be "Bad" due to losses.

The screenshot shows the 'Statistics' section of the Citrix SD-WAN GUI. It includes a 'Path Statistics Summary' table with the following data:

Num	From Link	To Link	Path State	Virtual Path Service State	Virtual Path Service Type	BOWT	Jitter (mS)	Loss %	kbps	Congestion
1	MPLS_MI	MPLS_TO	BAD	BAD	Static	2	3	11.54	65.28	NO
2	Milano_ADSL	Torino_ADSL	BAD	BAD	Static	2	3	9.62	5422.00	NO
3	MPLS_TO	MPLS_MI	GOOD	BAD	Static	2	3	0.00	309.15	NO
4	Torino_ADSL	Milano_ADSL	GOOD	BAD	Static	2	3	0.00	11.30	NO

Additional details from the screenshot: The table shows 4 entries. The 'Path State' and 'Virtual Path Service State' columns are highlighted in yellow for 'BAD' and green for 'GOOD'. The 'Congestion' column is highlighted in green for 'NO'. The interface includes a filter field, a refresh rate of 5 seconds, and navigation buttons (First, Previous, 1, Next, Last).

Figure 4.13. Path state after adding losses to both the channels

In addition, the device is able to gather statistics about the state of the path every 5 seconds, so that it can compute which is the best path to follow instant-by-instant, redirecting the traffic on another channel if it becomes bad. This is one of the most remarkable achievements of SD-WAN technology, because it takes care of carrying out the transfer even if the channel conditions are deteriorating, trying to find an alternative path.

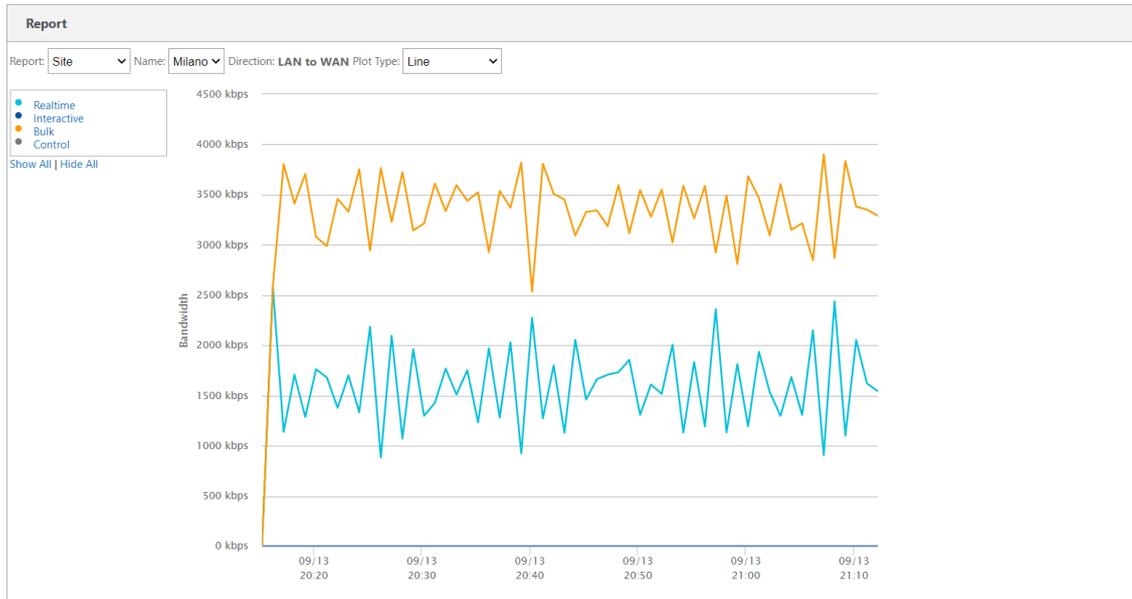


Figure 4.14. Video streaming and Jperf traffic with channel losses

To return to the test, what emerges from figure 4.14 is that the total exploited bandwidth does not exceed 6 Mb/s, which is more or less the maximum exploited capacity from the two combined kinds of traffic. The described phenomenon shows how even a moderate percentage of losses in both the channels translates into a reduction of more than the 50% of the bandwidth. However, the most interesting circumstance is related to the above explained capacity of the SD-WAN device of redirecting traffic in real-time, appreciable in figure 4.15.

As a matter of fact, figure 4.15 emphasises the high variability in the traffic distribution among the two channels, that can be explained thanks to the ability of the Citrix SD-WAN device to continuously collect statistics and to establish where to route the traffic according to them. Whenever one channel is sensed as "Bad", the traffic is moved in the second channel and whenever they are both "Bad", the device goes deeper into details and compares the statistics for choosing one of the two.

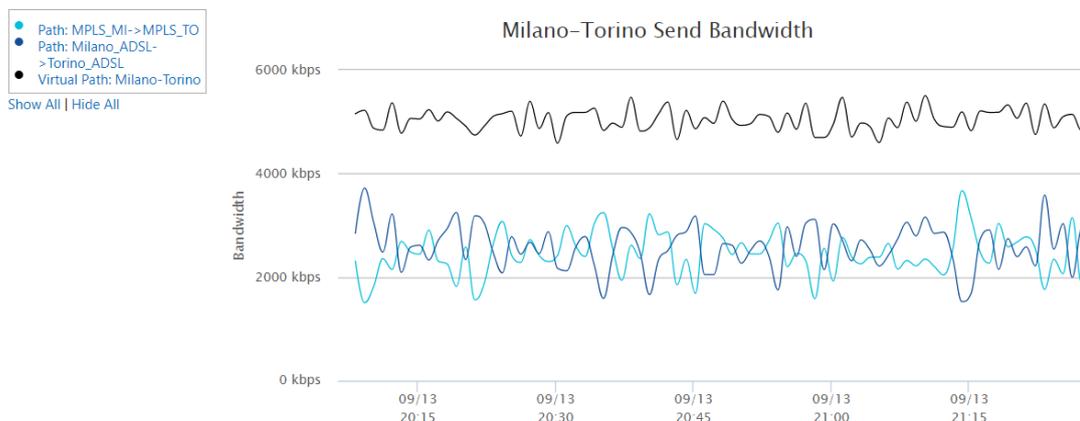


Figure 4.15. Traffic in both channels with losses

4.5.7 Scenario 7: Delaying the ADSL channel

The same experiment of the above section can be repeated by changing the channel parameters in the WAN emulator, with the aim of testing another failure scenario. For this task, it has been established to delay only the ADSL channel and to reset losses to 0 in both the channels. Thus, this time the MPLS is fully performing, while the ADSL one suffers from a 500 ms delay. The focus of the study is around the behaviour of each of the two channels according to their status. Indeed, the MPLS channel, which is equipped with a 6 Mb/s capacity, is supposed to provide a good performance and to exploit more or less its whole bandwidth. The same cannot happen for the ADSL channel, which suffers from a significant delay. Even though 500 ms could seem a small amount of delay, there are many applications that would be extremely degraded in that condition, like the video playback involved in the simulation.

What discussed above can be further confirmed by figure 4.16, where it is evident that the SD-WAN device transfers all the traffic into the MPLS channel, like the black overall traffic profile confirms. Indeed, it slightly coincides with the light blue profile, that represents the traffic traveling through the MPLS connection.



Figure 4.16. Traffic in both channels with delay in the ADSL one

As for the ADSL channel, it is almost unused because the Citrix SD-WAN policy is to safely deliver packets, and it never chooses to transmit data over the ADSL channel in order to avoid losses caused by delays. Another remarkable capability of the Citrix device is that it provisions the three categories of traffic with three distinct level of priority, and in this specific case, the real-time traffic is more urgent with respect to the bulk.



Figure 4.17. Jperf and video playback with 500 ms delay in the ADSL channel

The effect of this assignment is that the Citrix SD-WAN appliance distributes the available capacity among the applications trying to supply the real-time traffic with all the resources it needs and then yields the remaining ones to Jperf.

Figure 4.17 corroborates what just stated: the video playback procedure receives all the resources required (more or less 2 Mb/s) while the Jperf is equipped with what remains available.

But what happens if the delay is pretty negligible? Another scenario with a delay of 10 ms has been tested. The current situation is pretty different from the 500 ms delay case, because 10 ms do not impact the performance of the ADSL channel, so that it is capable of fully exploiting its capacity. Figure 4.18 exhibits the transition from the poor performances of the ADSL channel when suffering a huge delay to a pretty satisfying result when suffering a minor delay.

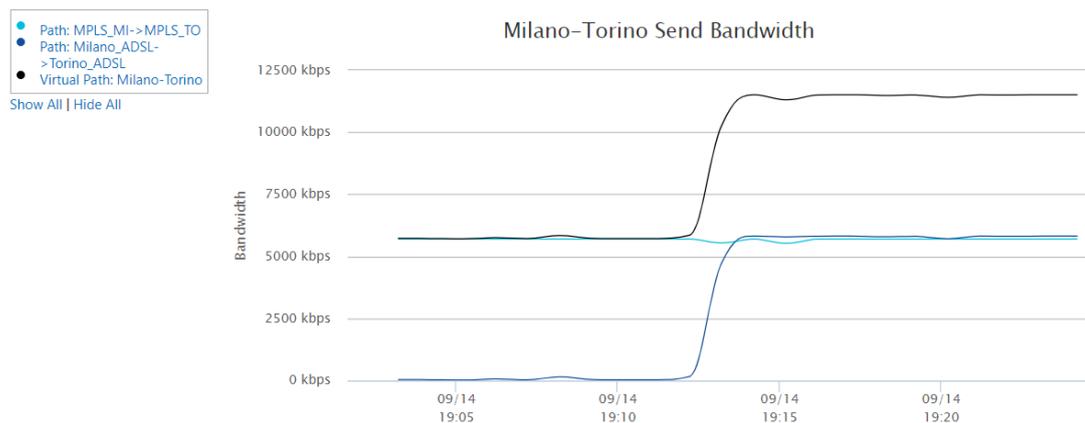


Figure 4.18. Jperf and video playback with 10 ms delay in the ADSL channel

To sum up, the network and the Citrix appliance react in distinct manners according to the type of failure experienced. Indeed, the SD-WAN device appears to be more robust to small and moderate delays, while for huge ones, like hundreds of milliseconds, it starts preferring to move packets to more reliable channels. As far as losses are concerned, the device immediately swaps the data transfer to other channels in order not to lose traffic.

4.5.8 Scenario 8: Failures: severing MPLS channel

The last test is focused on the same kind of traffic adopted in section 4.5.6 and 4.5.7 but with a substantial difference in the network settings: it has been imposed the 100% of losses in the MPLS channel, which results in the definitive loss of that link. The described event can be thought in real terms as a cut cable, which would require human intervention to be fixed and, consequently, time and money. The SD-WAN appliance, instead, reacts immediately to the fault without any additional external intervention by redirecting the traffic in the available channel. This procedure is immediate and the user does not notice anything neither in terms of time nor in terms of performance. Also in this occurrence, the traffic with the highest priority, i.e., the video streaming, is provided with the necessary bandwidth, so the user does not incur in any performance degradation.

From the figure 4.19 below, it is possible to appreciate both the MPLS path state, which is "DEAD" because it cannot receive more packets and also the reason, which is before "PEER" due to path state change and then "SILENCE", due to the interruption of the channel.

The screenshot shows a table titled "Path Statistics Advanced" with a filter and navigation controls. The table contains 4 entries. Entry 1 shows a transition from UNKNOWN to DEAD state due to PEER. Entry 2 shows a transition from GOOD to UNKNOWN state due to SILENCE. Entry 3 shows a transition from UNKNOWN to DEAD state due to SILENCE. Entry 4 shows a transition from NO to GOOD state due to N/A. The table includes columns for Num, From Link, To Link, Congestion, Path State, Reason, Duration (S), Virtual Path Service State, Src Port, Dst Port, MTU, BOWT, Jitter (mS), Packets Received, OOO, Loss %, kbps, and Virtual Path Service Type.

Num	From Link	To Link	Congestion	Path State	Reason	Duration (S)	Virtual Path Service State	Src Port	Dst Port	MTU	BOWT	Jitter (mS)	Packets Received	OOO	Loss %	kbps	Virtual Path Service Type
1	MPLS_MI	MPLS_TO	UNKNOWN	DEAD	PEER	1087	GOOD	4980	4980	1492	9999	0	104	0	0.00	11.79	Static
2	Milano_ADSL	Torino_ADSL	NO	GOOD	N/A	1091	GOOD	4980	4980	1492	2	3	2356	0	0.00	5647.00	Static
3	MPLS_TO	MPLS_MI	UNKNOWN	DEAD	SILENCE	1085	GOOD	4980	4980	N/A	9999	0	0	0	0.00	0.00	Static
4	Torino_ADSL	Milano_ADSL	NO	GOOD	N/A	4402	GOOD	4980	4980	N/A	2	3	1356	0	0.00	236.29	Static

Figure 4.19. Statistics of the MPLS channel with 100% of losses

The performance report section allows the visualization of the sharp slope change of the light blue curve, the one representing the MPLS exploited bandwidth. As the losses increase and the SD-WAN appliance records them, the MPLS circuit stops working and all the traffic is moved to the ADSL link which is still performing good.

The overall bandwidth exploitation becomes superimposed with the ADSL traffic, because the other channel is no more available. It is of fundamental importance to remark that, even if one link is damaged, the virtual path is preserved and the data exchange too. The current scenario is an enlightening example of how the SD-WAN technology is fundamental for avoiding the loss of crucial information. Imagining that the traffic involved in the simulation belongs to a "mission critical" application, it can be stated that it is almost surely delivered to the destination because in that case the sender makes a copy of each packet and sends it by means of each channel consuming twice the bandwidth. So, if one of them is damaged, data are delivered through another available channel. The risk is that the SD-WAN receiver appliance sees more copies of the same packet at its entrance, but it implements a policy thanks to which it drops all the unnecessary packets and it saves the first arrived.

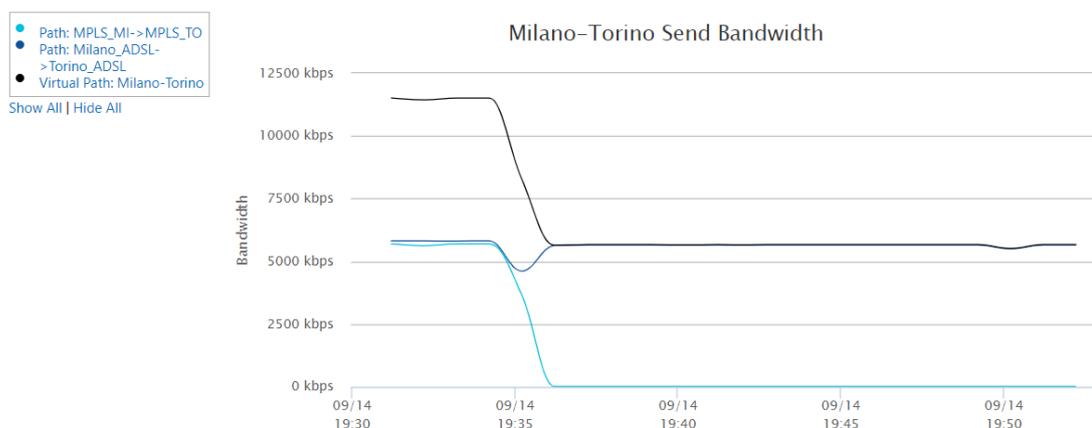


Figure 4.20. Jperf and video playback after the loss of MPLS channel

The described situation gives an alternative option beyond the MPLS link, so it does not incur in the loss of the connectivity. However, it could happen that all the wired connections fail, but the Citrix SD-WAN appliance is also programmed for handling this kind of issues. Actually, whenever all the wired connections collapse, the SD-WAN device redirects the traffic on wireless channels using 4G/5G and,

later, it restores the initial traffic path after the wired connections become available again.

Chapter 5

SD-WAN advantages and deployment scenarios

5.1 SD-WAN advantages

The Open Networking User Group has defined SD-WAN as the best SDN use-case for two years in a row [9]. Despite the well-known advantages already brought by the baseline of SD-WAN, i.e. SDN, it adds many functionalities that make it a cutting-edge innovation.

The first key insight is the realization of an application-driven network, where the subject of interest is no more the physical layer connectivity, but the application. This new concept matches together applications, users, policies and security over a WAN network able to provision several services and applications in a more reliable and fast fashion [9].

Another noticeable point is end-to-end encryption. This goal is achieved by adopting IPSec tunnels, which are tunnels able to grant security over packet networks adding features like authentication, encryption and integrity control of IP packets [25]. IPSec tunnels can automatically protect the private WAN connectivity also passing through either public or shared networks. However, encryption is required also in MPLS connections since the actual breach-sensitive world makes it

necessary for all regulated industries [26].

Among the innovations and the advantages of this brand new technology there is also the multi-path and multi-link support, that consists of bonding several physical channels into a single logical one [26]. The result of this operation is groundbreaking since it represents a flexible way for achieving a larger aggregate capacity and an improved reliability. Multi-path and multi-link support runs parallel to the dynamic path selection technique, that relies on the constant statistic computation which is useful for adjusting traffic flows, for load-balancing and for counteracting congestion. Thanks to the conjunction of these two novelties, SD-WAN bears virtual routing and forwarding that are the baseline of network segmentation and of centralized control.

Going forward, it can also be enumerated among the advantages of SD-WAN the practice of path conditioning for WAN optimization. Path conditioning deals with some WAN optimization capabilities like data compression, deduplication, client-side caching, TCP protocol optimization and traffic shaping [26]. Their aim is to control contentions and high-latency WAN circuits especially for those enterprises that manage applications for video acceleration and file system protocols.

Moreover, the introduction of traffic categorization leads to a new way of providing quality of service. Traffic is identified as belonging to a specific class of service and each of them is assigned with a given priority and bandwidth guarantees [26]. Actually, there are some loss-sensitive and latency-sensitive applications like video conferencing, VoIP (Voice over IP) and screen sharing that constantly require some quality of service constraints, and SD-WAN is helpful in this objective thanks to traffic categorization and prioritization, as observed in the simulations of the previous chapter. In addition, many SD-WAN vendors have implemented some solutions with the forward error correction capability: redundancy is added to the original information with the intent of detecting errors at the receiver.

Ultimately, the SD-WAN layout is often adopted for local inspection and direct

routing of traffic intended for trusted cloud services, without the need of backhauling the traffic to an external entity for inspection [26]. The current approach allows to save bandwidth and promotes local direct Internet access without affecting security.

5.2 Organizations migrating to SD-WAN

The actual telecommunications environment is required to cope with the increasing bandwidth demand by establishing new sites with remote connectivity and with the greater agility and scalability needed [26]. Remote locations have severe requirements on bandwidth and latency, and old architectures are no longer up the job since the categories of traffic to support are heavy and comprise large data transfers, cloud-based backups, video and audio streaming. It must be also taken into account the advent of software-as-a-service applications or private and public clouds that adds to the burden of IT work [26]. Thus, this motivates why many enterprises are moving towards an SD-WAN approach, which offers fast, less costly provisioning and a more flexible access to the Internet. A more practical confirmation of why SD-WAN is useful for organizations can be observed from figure 5.1 taken from SDX central paper [26], that has run an experiment on two different sets of data from an enterprise-end-user data slice. Half of the participants to the experiment had no implemented SD-WAN yet, while the other half had. Thus, the idea given by the graph in figure 5.1 is that SD-WAN-addicted companies perceive as the most compelling reasons for adopting this new network technology the cost reduction and management issue and the need for network agility. On the contrary, enterprises which still have doubts on SD-WAN, prefer to focus their attention on the need for network agility, which is more accentuated in this case due to the lack of a flexible old architecture supporting all the new services.

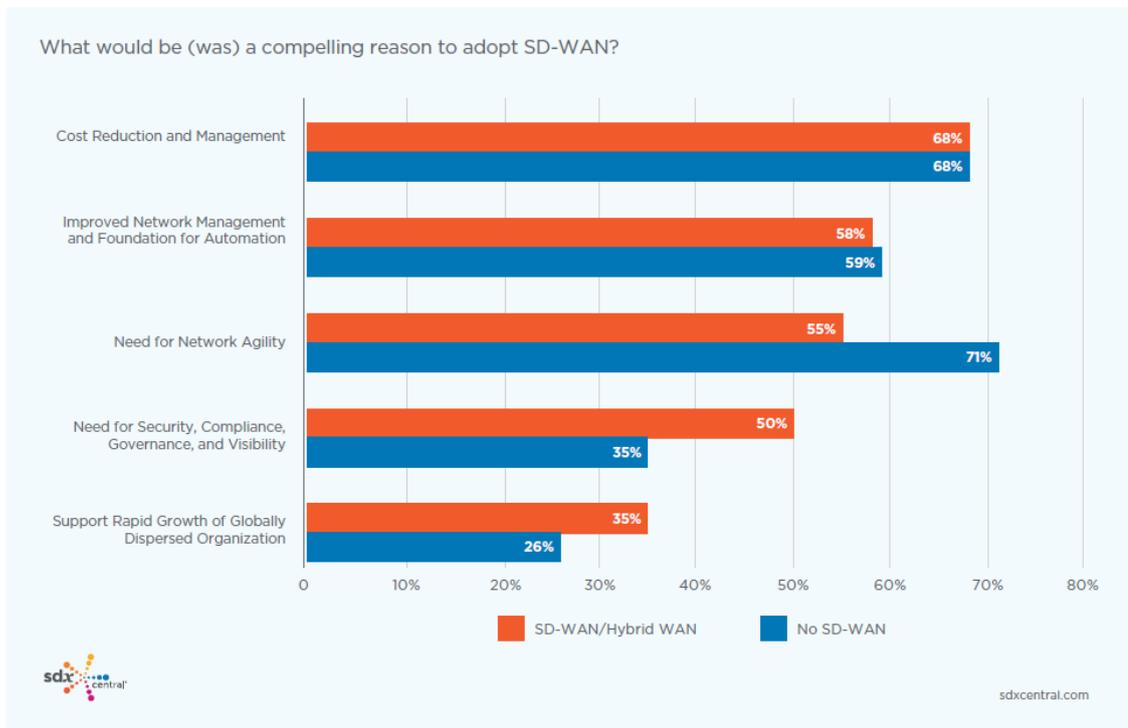


Figure 5.1. Compelling reasons for SD-WAN adoption, reproduced from [26]

In this paragraph, three of the main verticals that have adopted SD-WAN solutions are inspected.

5.2.1 Retail stores

As for retail stores, they are one of the three verticals that have experienced the need of trying the new SD-WAN network layout. Indeed, the greatest concern of retail stores is about compliance, in detail, they have to address the Payment Card Initiative issue that means protecting customer's credit card information [26]. The methodologies involved for tackling compliance are encryption of traffic and network segmentation, which result to be useful to control that no malicious users intercept traffic.

Furthermore, retail stores may often be seasonal businesses and for this reason

they need a fast and secure connectivity, so they exploit SD-WAN. Actually, it provides a flexible Internet access and 3G/4G backhaul options that allow quick provisioning of resources for seasonal businesses [26]. But there is more, SD-WAN is also able to grant an improved Quality of Experience for customers in terms of application exploitation. Indeed, SD-WAN architecture can assign different priorities to applications and in this case, it gives payment card transaction traffic the highest one, in order to protect it.

5.2.2 Healthcare

Nowadays, the healthcare world is moving towards a much more sophisticated network support, because it is encompassing the usage of many applications which are internet-based. In order to deepen, some of these applications are large file transfers for radiology and electronic health records (EHR), IoT devices like smart monitors, smart carts and also business-critical traffic for patient scheduling or payments [26]. The involvement of these new kinds of traffic has largely increased bandwidth demand both in the LAN and the WAN and has introduced the need for a secure and always performing network. In fact, there is the HIPAA (Health Insurance Portability and Accountability Act) service that deals with encrypting and segmenting patient data and it requires a certain level of security to cope with breaches [26]. SD-WAN is advantageous in this context because it supports policies for segmentation or security reports for identifying breaches [26]. Another fundamental concern about healthcare applications, EHR and IoT devices is the information delivery, both in terms of correctness and time: application prioritization and QoS capabilities of SD-WAN devices are profitable in this purpose [26]. Moreover, it should also be dealt with that healthcare world is experiencing a slow transition towards the telemedicine, including remote surgery and remote diagnostic. The mentioned activities have the tight requirement of needing a constant and reliable service that never interrupts. SD-WAN can lend a hand in this field by offering its reliability,

security and prioritization capabilities.

5.2.3 Financials and Insurance

Financial institutions are among the most relevant enterprise verticals and they usually leverage new networking capabilities for improving their revenue and the quality of experience of their customers, as well as insurance companies [26]. On one hand, financial institutions' aim is to take advantage of those platforms that provide security, reliability with a reduced expenditure. Thanks to their capabilities, SD-WAN appliances are perfectly able to meet these requirements, since they exploit different transport connections with a large bandwidth availability and several new security frameworks. Indeed, one essential feature of the SD-WAN technology is a fine-grained security policy along with segmentation rules, that are in charge of protecting data. The improved overall management gives to enterprises the opportunity to save money by reducing the operational costs and the invested capital.

The same holds for insurance companies, whose needs are similar to those of financial institutes because they have to guarantee security across many different branch offices residing in sparse locations [26]. In addition, the capability of SD-WAN of supporting various connectivity types like the direct Internet access with 3G/4G backup allows to set up broker locations and new agent in a much faster and more reliable way.

5.3 Deployment scenarios

Once all the new features and benefits brought by SD-WAN have been inspected, it is of relevant interest to analyze some real deployment scenarios realized by specific vendors with their own SD-WAN release. The current work has delved into some SD-WAN solutions such as Citrix's, Lavelle Networks' and Equinix's, and their

deployment scenarios.

5.3.1 Equinix: Bridging Multiple SD-WAN Deployments

The first example is obtained by using the Equinix SD-WAN solution. The remarkable feature of this SD-WAN implementation is the Network Edge paradigm, a peculiar service for realizing network functions virtualization infrastructure-as-a-service (IaaS) on the Equinix platform. Network Edge takes charge of virtually creating some network devices for the connection of branch offices to cloud platforms or IT infrastructure, without requiring extra physical hardware. A significant example of how it works can be observed in figure 5.2, where many SD-WANs and cloud platforms are bridged together in the single and reliable Equinix platform.

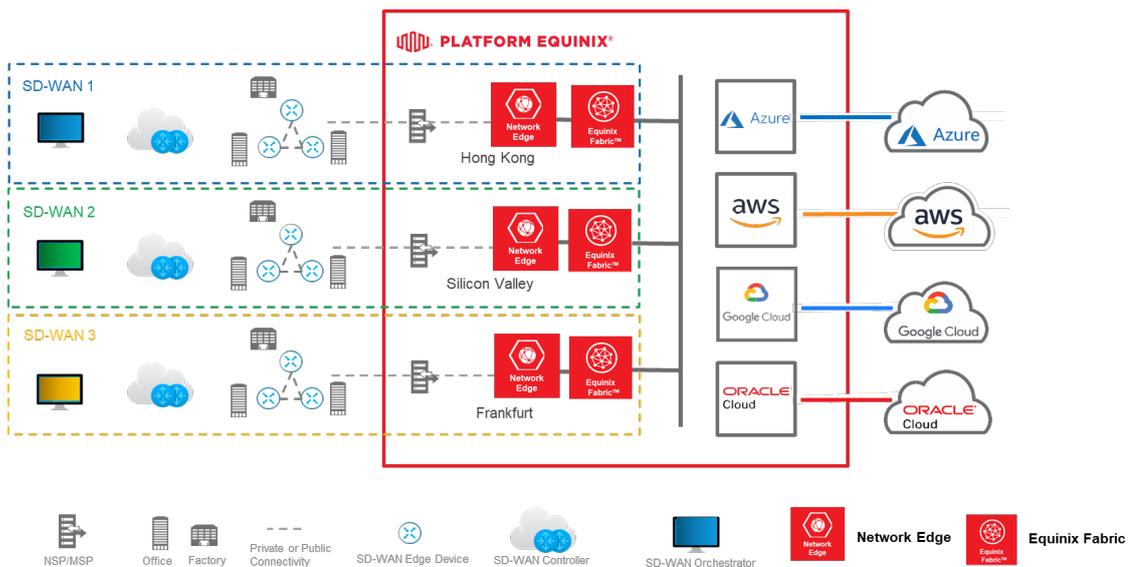


Figure 5.2. Architecture for bridging SD-WANs, reproduced from [27]

As can be deduced from the above picture 5.2, this architecture is able to connect together on a single platform more devices, already connected to an SD-WAN endpoint, and to make the access to separate cloud providers possible. The state-of-the-art behind this implementation consists of putting together several devices,

either physical or virtual, from different brands and allowing the communication among them regardless of their vendor specifications and of their location [27].

5.3.2 Equinix: Simplification of Multi-Cloud Deployment

The second prominent Equinix realization concerns with the simplification of the connection among branch offices and multiple clouds by means of Network Edge. Indeed, it allows the creation of several apparatus on a virtual environment, which are later employed in the interconnection between SD-WAN infrastructure and clouds. SD-WAN devices become virtual and the private cloud connections are aggregated into a single software-defined connection, as shown in figure 5.3 [27]. By adopting this operating principle, the complexity of the whole structure is drastically reduced and so the costs, thanks to the deployment of virtual devices.

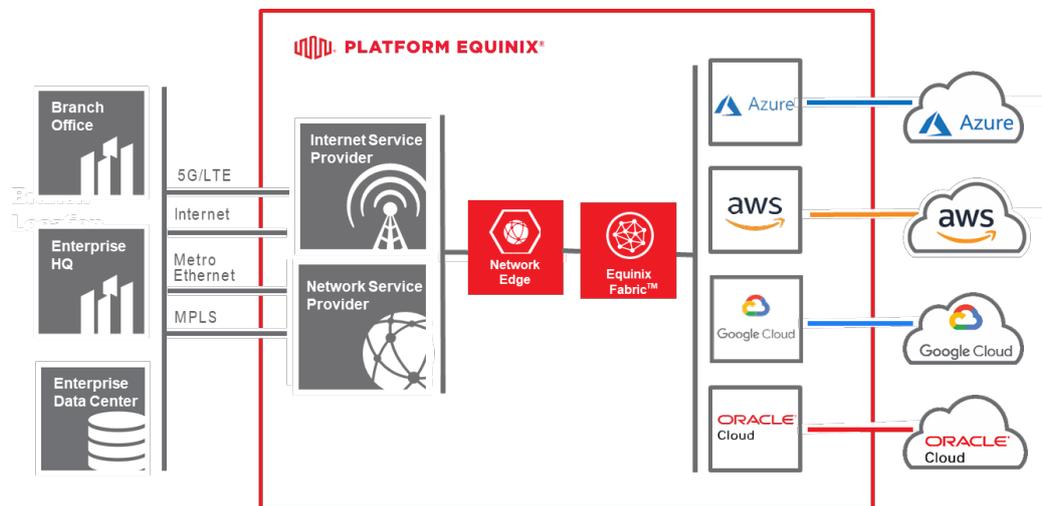


Figure 5.3. Multi-cloud connectivity via Network Edge, reproduced from [27]

5.3.3 Lavelle Networks: SD-WAN tunnel with Internet and MPLS WANs

Lavelle Networks SD-WAN solution is another vendor deployment, different from the Citrix one or from the Equinix Platform. The current section investigates how the SD-WAN elements can be used in a hybrid scenario, i.e., an architecture relying on the existing MPLS connections. Usually, an hybrid scenario is composed by a broadband link on one site, an MPLS one on the other site and, possibly, another site with both of them [28].

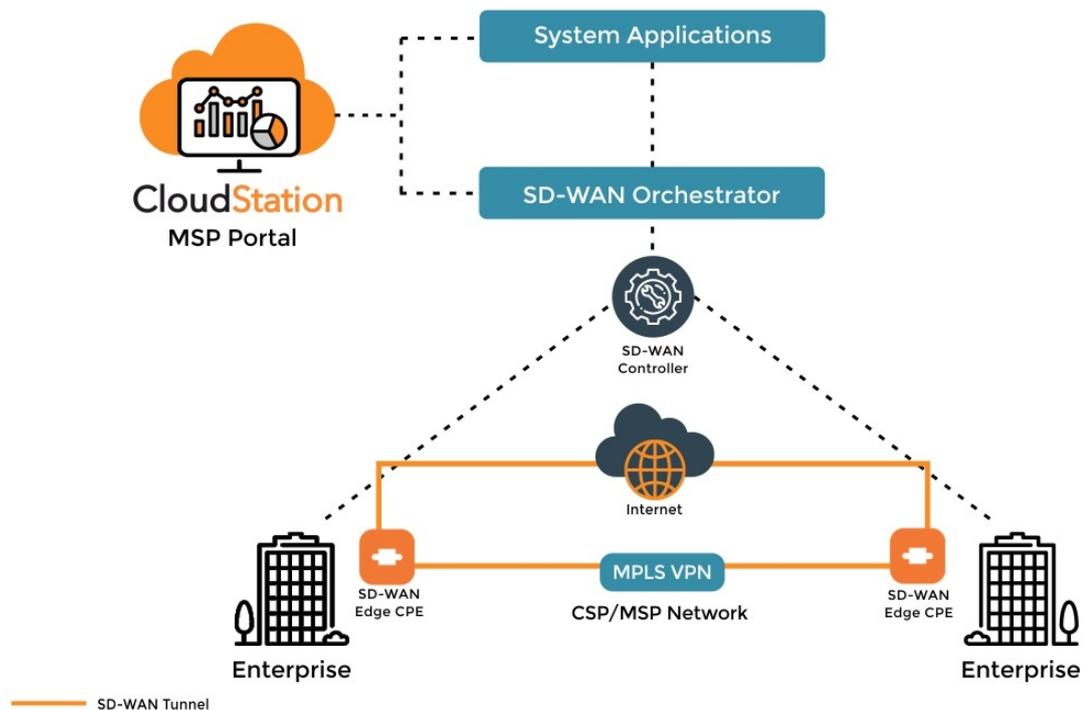


Figure 5.4. Tunneled SD-WAN Internet connection over MPLS link, reproduced from [28]

The use-case highlighted in figure 5.4 has been realized meeting the hybrid WAN requirement, because it relies on existing broadband and MPLS connections, but also leveraging an encrypted SD-WAN tunnel for allowing the communication

among the branches.

The example reported above exploits two WAN connections for increasing resilience and thanks to the adoption of the SD-WAN tunnel, it is also possible to provide an improved bandwidth utilization. Hence, it turns out from the current analysis that hybrid WANs are a pretty efficient solution since they can support the coexistence of MPLS and broadband (internet) links, also offering the benefits of the SD-WAN technology [28].

5.3.4 Lavelle Networks: SD-WAN service using Multiple Internet Service Providers

The use-case of figure 5.5 is about availing of more than one Internet Service Provider (ISP), the entity that allows internet navigability through a contract.

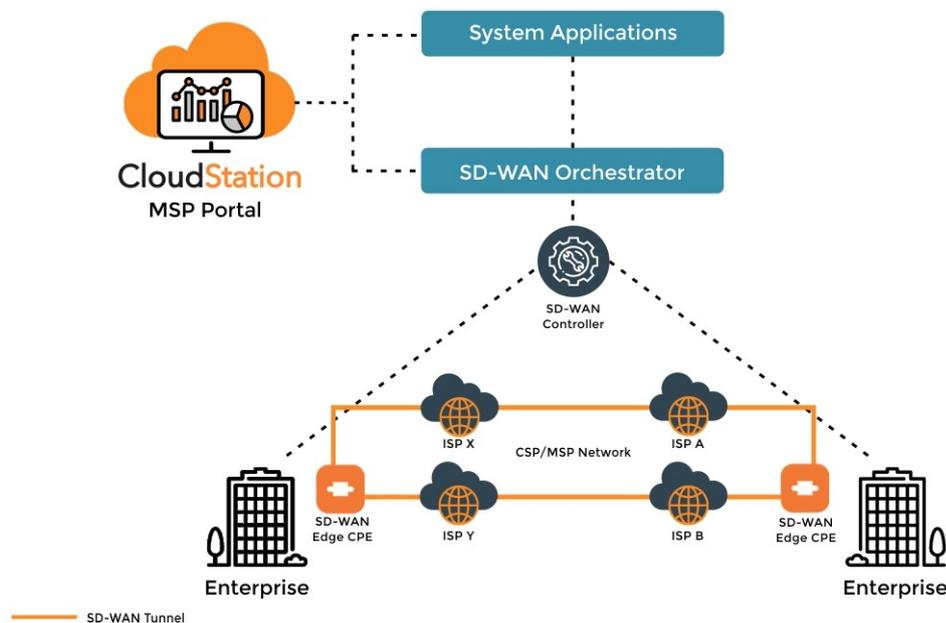


Figure 5.5. Multiple ISPs with SD-WAN service, reproduced from [28]

The deployment scenario examined in this section sheds light on distinct types

of ISPs, like DSL (Digital Subscriber Line) or Cable Internet or a Dedicated Internet Access, and on the benefits they bring when combined with multiple WAN connections. However, it must be recalled that ISPs do not correspond to Cloud Service Providers (CSP) or Managed Service Providers (MSP) that are responsible for the SD-WAN appliances deployment, so this scenario is possible only whenever the two branches can be reached by an Internet WAN [28]. The outcome of this implementation gives an overall improved service, both in terms of network resilience to failures and in terms of bandwidth exploitation.

5.3.5 Citrix: SD-WAN and Microsoft Office 365

Coming back to the Citrix SD-WAN solution, which is the main object of the current thesis, it offers an interesting application scenario for Office 365. The figure reported below (5.6) gives an overview of the operating principle of the architecture.

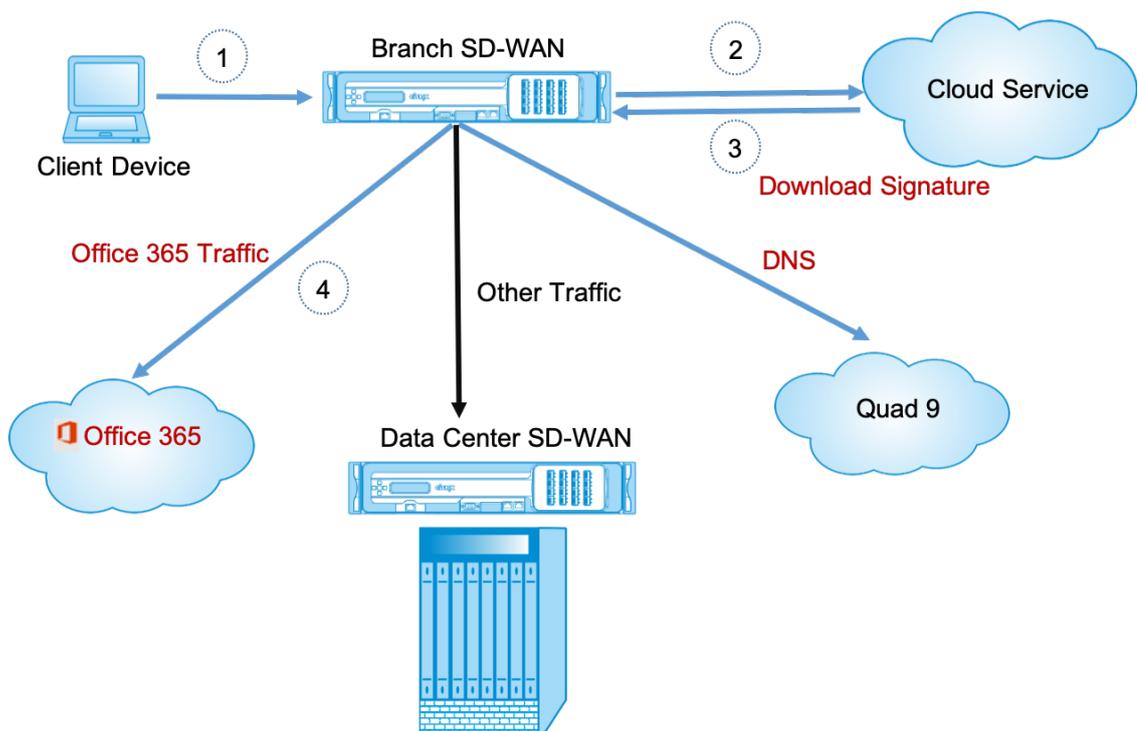


Figure 5.6. Office 365 service on Citrix SD-WAN appliances, reproduced from [29]

The architecture in figure 5.6 can be explained in several steps, emphasizing the aim of optimizing Office 365 traffic. Office 365 is provided with many endpoints, defined front doors, all around the world [29]. Citrix SD-WAN solution's deep packet inspection engine is able to classify Office 365 on the first packet and then directly connects the user to the closest Office 365 front door to guarantee the lowest latency connection possible. As a matter of fact, once the traffic has reached the front door it is immediately directed to Microsoft's network and then is delivered to the final destination. This practice avoids utilizing backhauling to central proxies that can add latency ending up in poor user experience and reduces the round trip time from the customer LAN to the Office 365 endpoint.

Chapter 6

Conclusion

SD-WAN technology is conceived at the same time as the most promising application of SDN, but also as the most challenging one. It has the potential to completely reverse the wide area networks coverage and infrastructure deployment, also with the commitment of improving the targets of latency, bandwidth exploitation and resource utilization. In light of recent advancements, SD-WAN represents the most suitable starting point for making the transition towards a virtualized network management possible and permanent. Many enterprises are migrating their data and their applications in the cloud, progressively detaching from vendors lock-in. SDN and SD-WAN are the keys for supporting this phenomenon.

The scope that has been pursued with the current thesis is to experience first hand the degree of freedom of the network manager that interacts with the centralized brain of this kind of architecture and what are the effective benefits nowadays. After analyzing and studying all the features of the SD-WAN technology, the attention has been moved towards a specific SD-WAN vendor implementation, which is the Citrix one. The core of the work has been becoming familiar with the mentioned product and testing its capabilities in a simple virtualized network scenario.

It has been surprisingly discovered that by simply managing the graphical interface of this virtual SD-WAN appliance, many configurations and settings can be imposed. The appliance is smart enough to categorize the traffic that crosses it

and assigns a level of priority to each class. However, this is not a static and permanent configuration, conversely, the network manager can access the rule tables and change the priority of each protocol belonging to the different traffic category. This demonstrates the dynamism of this solution: the network can adapt to each requirement by simply changing settings.

Another remarkable outcome of this work has brought to light how SD-WAN would impact on costs. Indeed, the last experiment describes a situation that resembles a severed cable in reality and highlights how the device reacts immediately to the failure without interrupting the transmission and without making the user experience any loss. In reality, this would not happen. In a real scenario the MPLS link should be fixed and this would require investing a lot of money and also paying an expert for completing the task. Moreover, users would experience the interruption of the service. Citrix SD-WAN appliance solves all these issues at once by simply transferring the traffic on another available channel, without any delay or loss. But there is more: physically realizing the network infrastructure adopted for the experiments would have implied buying expensive devices and links. It has been necessary, instead, to buy only the Citrix and the VMware licenses.

Thereby, the experimental work has confirmed what theoretically stated about the benefits of the technology. However, it should be also taken into account that there are substantial limitations in the current releases of the SD-WAN products, because they still do not guarantee end-to-end quality of service. The tests explained in the previous chapters, in fact, emphasise the lack of improvements in terms of end-to-end experienced speed, especially during the downloading tests. The responsiveness of the virtual machines is still an open issue because they are too slow.

6.1 Future works

Despite the several and serious limitations of the SD-WAN technology, it is reasonable to believe that this new network paradigm is worthwhile for further research.

The greatest limitation stands on the lack of a stand-alone SD-WAN architecture: it must be integrated into an hybrid structure that involves also MPLS and other transport technologies. Thus, the economical advantages are not reachable yet. But, it turns out from recent studies that the lack of an on-site security functionality that would imply purchasing specialized hardware, can be addressed by implementing and improving secure functions within the SD-WAN themselves.

To conclude, the second key on which vendors are focusing their efforts is the improvement and the integration of new services inside virtualization platforms and the decoupling from servers, which are limited in their computing capabilities and could affect latency and performances.

Bibliography

- [1] What is Virtualization technology & virtual machine? <https://www.vmware.com/solutions/virtualization.html>.
- [2] What is virtualization? (IBM). <https://www.ibm.com/cloud/learn/virtualization-a-complete-guide>.
- [3] I vantaggi della virtualizzazione. <https://www.redhat.com/it/topics/virtualization/what-is-virtualization>.
- [4] Cos'è Kubernetes? <https://kubernetes.io/it/docs/concepts/overview/what-is-kubernetes/>.
- [5] Software-Defined Networking (SDN) Definition. <https://opennetworking.org/sdn-definition/>.
- [6] Citrix Italy - What is Software Defined Networking?(SDN). <https://www.citrix.com/it-it/solutions/app-delivery-and-security/what-is-software-defined-networking.html>.
- [7] SDN, Software Defined Networks. <https://www.redhat.com/it/topics/hyperconverged-infrastructure/what-is-software-defined-networking>.
- [8] O'Connor Vachuska Peterson, Cascone and Davie. *Software-Defined Networks: A Systems Approach*. <https://github.com/SystemsApproach/SDN>.
- [9] Darril Jibson. *Software-Defined WAN For Dummies*. John Wiley Sons, Inc., 2015.
- [10] What is SD-WAN? - Software-Defined WAN Definition. <https://www.citrix.com/it-it/solutions/sd-wan/what-is-sd-wan.html>.
- [11] SD-WAN Defined: What is SD-WAN (Software-Defined Wide Area Network)? <https://www.sdxcentral.com/networking/sd-wan/definitions/software-defined-sdn-wan/>.
- [12] SD-WAN: What is Software-Defined WAN? <https://www.catonetworks.com/sd-wan/>.
- [13] Keith Townsend Zeus Kerravala Craig Connors Roopa Honnachari, Lee Doyle and Pere Monclus. *SASE ZTNA for dummies*. John Wiley Sons, Inc., 2021.
- [14] What is SD-WAN? How Does it Work and Why Do We Need it? <https://www.silver-peak.com/sd-wan/sd-wan-explained>.
- [15] Reference Architecture: SD-WAN - Citrix Tech Zone. <https://docs.citrix>.

- com/en-us/tech-zone/design/reference-architectures/sdwan.html.
- [16] Citrix SD-WAN 110 Standard Edition Appliances. <https://docs.citrix.com/en-us/citrix-sd-wan-platforms/standard-edition/110-standard-edition-appliance.html>.
 - [17] About Citrix SD-WAN WANOP - Citrix SD-WAN WANOP 11.4. <https://docs.citrix.com/en-us/citrix-sd-wan-wanop/current-release/about-sd-wan-wanop.html>.
 - [18] Citrix SD-WAN 11.4. <https://docs.citrix.com/en-us/citrix-sd-wan/current-release.html>.
 - [19] Inline Mode - Citrix SD-WAN Platforms. <https://docs.citrix.com/en-us/citrix-sd-wan-platforms/wanop/400-800-1000-2000-3000-wanop-appliance/cb-deployment-modes-con/br-adv-inline-mode-con.html>.
 - [20] WCCP mode - Citrix SD-WAN Platforms. <https://docs.citrix.com/en-us/citrix-sd-wan-platforms/wanop/4100-5100-wanop-appliance/deployment-modes/br-adv-wccp-mode-con.html>.
 - [21] Virtual inline mode - Deployment modes. <https://docs.citrix.com/en-us/citrix-sd-wan-platforms/cb-deployment-modes-con/br-adv-virt-inline-mode-con.html>.
 - [22] SD-WAN VPX Usage Scenarios - Citrix SD-WAN Platforms . <https://docs.citrix.com/en-us/citrix-sd-wan-platforms/vpx-models/vpx-se/br-vpx-usage-scenarios-con.html>.
 - [23] WANem - The Wide Area Network emulator. <http://wanem.sourceforge.net/>.
 - [24] ISO/IEC. Information technology – coding of audio-visual objects – part 12: Iso/iec base media file format; iso/iec 14496-12:2008. 2008.
 - [25] IPsec - Wikipedia. <https://it.wikipedia.org/wiki/IPsec>.
 - [26] IBM. Real-world sd-wan deployment, a survey-backed analysis. 2018.
 - [27] 3 Scenarios for Handling Multiple SD-WAN Deployments. <https://blog.equinix.com/blog/2021/04/16/3-scenarios-for-handling-multiple-sd-wan-deployments/>.
 - [28] SD-WAN: 5 Deployment Scenarios. <https://lavellelennetworks.com/blog/sd-wan-5-deployment-scenarios/>.
 - [29] Office 365 optimization - Citrix SD-WAN 11.4. <https://docs.citrix.com/en-us/citrix-sd-wan/current-release/office-365-optimization.html>.