Issues and Challenges for Network Virtualisation

Nikheel Tashvin Seechurn, Avinash Mungur, Sheeba Armoogum and Sameerchand Pudaruth

ICT Department, Faculty of Information, Communication and Digital Technologies, University of Mauritius, Mauritius

***Abstract*:** In recent years, network virtualisation has been of great interest to researchers, being a relatively new and major paradigm in networking. This has been reflected in the IT industry where many virtualisation solutions are being marketed as revolutionary and purchased by enterprises to exploit these promised performances. Nevertheless, challenges such as trust, security and complete isolation remain to be tackled. In this study, an investigation of the different state-of-the-art virtualisation technologies, their issues and challenges are addressed. A systematic review was effectuated on selectively picked research papers and technical reports. Moreover, a comparative study has been performed on different network virtualisation technologies which include features like security, isolation, stability, convergence, outlay, scalability, robustness, manageability, resource management, programmability, flexibility, heterogeneity, legacy support, and ease of deployment. The virtualisation technologies comprise Virtual Private Network (VPN), Virtual Local Area Network (VLAN), Virtual Extensible Local Area Network (VXLAN), Software Defined Networking (SDN) and Network Function Virtualisation (NFV). Conclusively results exhibits that these technologies overlooked some features to promote other features. Moreover, further discussion exhibits the need for an improvement of the existing network virtualisation environment by exploring the vital drawbacks.

Keywords: Network Virtualisation, VPN, VLAN, VXLAN, SDN, Network Function Virtualisation

1. Introduction

Virtualisation can be simply defined as the creation of virtual topologies (or information subdomains) on top of a physical topology and has been the main element of almost all state-of-the-art IT design from the minor single campus network to complex globe-spanning service providers (SP) [1] or major enterprises [2]. Many types of market research show staggering statistics for virtualisation with as much as 92% of businesses operating on some form of virtualisation, whether it is for storage, servers, or network [3]. This goes to show its importance in today's world.

Network virtualisation is commonly mistaken with server virtualisation. However, these two technologies are somewhat different yet complementary. Server virtualisation is the logical partitioning and management of physical hardware resources, such as memory, I/O, storage, and CPU while network virtualisation is any technology capable of splitting, abstracting and decoupling the physical underlying topology or substrate (which can vary from transmission media like Ethernet or wireless [4][5] to a middlebox or server) from a 'logical' or 'virtual' topology by making use of some form of encapsulated tunnelling or software.

It can be further broken down into two further categories. First is internal virtualisation which is any technology designed to use software containers to reproduce the physical features of a singular network [6]. Such technologies include overlay networks such as Virtual Extensible Local Area Networks (VXLANs). They exploit a limited physical substrate to provide more functionality and make it more scalable. These technologies are used in both campus and datacentre environments where virtual machines (VM) are commonly used to provide segregation of services. Secondly, the external virtualisation approach is based on the network efficiency that is enhanced when different local networks are amalgamated into a singular virtual network and managed by software as a single entity. In this case, we usually have Virtual Local Area Networks (VLANs) for virtual connectivity and network switches acting as the physical substrate to link up all the local networks [6]. Another commonly used example for external virtualisation is virtual private networks (VPNs) whereby geographically separate networks can share the same connectivity using the Internet or some other shared infrastructure as physical substrate.

Network virtualisation has shown its precedence in overcoming physical limitations of hardware and can enhance performance, versatility, and response, unhampered by tangible limitations imposed by the hardware. One good example of such would be the logical shaping and sharing of network traffic onto a physical link's fixed bandwidth. It has also enabled many organisations to become modular and implement better business continuity and disaster recovery with minimised downtime as well as operating costs. There are new business ventures that are based on utilising existing physical underlays to create novel network virtualisation solutions. VPNs were created on that idea.

Mosharaf et al. [7] carried out this kind of systematic review of network virtualisation by introducing design goals such as isolation, stability, scalability, manageability, programmability, flexibility, heterogeneity, and legacy support while most research were focused on improving these already present technologies. Yet, there has been a lack of innovation in creating a network virtualisation model which fulfils these design goals. This paper will provide a more granular assessment on whether the discussed network virtualisation technologies are indeed meeting those criteria as well as additional ones like security, outlay, robustness, resource management and ease of deployment that can make or break these technologies.

This paper aims to examine, review, and discuss qualitatively the challenges and problems confronted by various network virtualisation technologies, which are of great importance due to their current widespread use and rapid growth in modern industrial and enterprise IT architectures [3]. From a holistic perspective, most network virtualisation technologies and their research work can be viewed as band-aid solutions to temporarily resolve current networking issues and not as a deliberate move towards an independent Network Virtualised Environment (NVE) [7].

The remainder of the paper is organised as follows; Section II covers the literature review on network virtualisation technologies. Section III describes about the methodology while section IV introduces the proposed requirements when building an NVE. Section V offers a comparison of the different network virtualisation technologies and the results of the paper are discussed in section VI. Finally, Section VII concludes the paper.

2. Literature Review

In this section, we present a survey to analyse and assess the underlying conventional network virtualisation technologies. We discuss the significance of related protocols and numerous assessment criteria like security, isolation, stability, convergence, outlay, scalability, robustness, manageability, resource management, programmability, flexibility, heterogeneity, legacy support, and ease of deployment.

2.1 Virtual Private Networks

A Virtual Private Network (VPN) is a dedicated collection of virtual nodes which are connected by a set of virtual links to form a virtual topology, which is basically a subset of the underlying TCP/IP-based network. This topology aims at connecting multiple sites using isolated and secured tunnelling over shared or public communication networks like the Internet [8]. VPNs can be created at different levels of the TCP/IP model each with its set of advantages and disadvantages. This technology became popular because of the continued growth in the size of business organisations. Such corporate entities needed a secure and reliable connection and a communication channel between their geographically distributed sites and partners.

The most radical way to achieve this level of security and privacy was dedicated leased lines, which cost a fortune. VPNs are a much more cost-effective means of achieving these goals [9] so much so that it now has a successful commercial market for regular Internet users and small and medium enterprises. There has been a demand for VPN services after the COVID-19 pandemic. The work from home policy was greatly facilitated by commercial VPNs as damage control practices. This technology is also a favourite of SPs who lease dedicated lines but make the most of it by sharing [10] those lines via isolated VPNs to customers. This virtualisation has been a constant improvement on security, isolation, stability, and resource management.

2.1.1 Layer 1 VPN

Layer 1 VPNs (L1VPNs) are extensions to Layer 2 and Layer 3 VPNs in the sense that they use a complex circuit switching domain instead of traditional packet switching concepts. The main difference between them is that data plane connectivity does not mean control plane connectivity in L1VPNs, and vice versa. This type of VPN is more prolific in singular domain SPs and loses some features like Quality of Service (QoS) during inter-domain routing [11]. These legacy features inhibit its scalability and flexibility. Another constraint is the effect of configuration policies of one L1VPN on another and the SP as the policies of customers can clash with that of SP network administrators [11]. Furthermore, the customer addressing scheme may overlap with other customers as well as with the SP addressing scheme. This can be solved with an address mapping mechanism, but, unfortunately, the latter is currently not well-defined in standard specifications [12]. These affect the overall isolation, programmability and manageability of the technology.

2.1.2 Layer 2 VPN

Layer 2 VPNs (L2VPNs) use virtualisation on Layer 2 (datalink layer), to connect geographically remote sites on the same LAN. This type of VPN is especially popular with SPs that still have Layer 2 infrastructure in their networks such as Frame Relay, Asynchronous Transfer Mode (ATM), High Level Data Link Control (HDLC), and Point-to-Point Protocol (PPP). The scalability and heterogeneity prospects of this solution are, therefore, hindered. Two of the most common L2VPN are L2TP (Layer 2 Tunnelling Protocol)-based or AToM (Any Transport over Multiprotocol Label Switching, MPLS)-based but both come with drawbacks.

Firstly, tunnelling is a resource-intensive process as the frame size is exponentially larger with a pseudo-wire header on top of a tunnel (L2TP or MPLS) header. This causes congestion during the transfer as well as consumes the processing power of end devices to encapsulate and decapsulate data. The resource management of L2VPNs is thus severely impacted. Secondly, this technology is customer-centred such that the routing and management of traffic are done at the Customer Premises Equipment (CPE) with the SP infrastructure only being an underlay. This makes monitoring of the data flow quite difficult for the SP as well as an overall lack of control plane to manage reachability across the L2VPN. Moreover, being a Layer 2 technology, L2VPN offers a flat subnet across the different sites it connects. It also requires an additional administration of its IP allocation. This has a negative effect on its security, scalability, manageability, programmability, flexibility and ease of deployment.

2.1.3 Layer 3 VPN

Layer 3 VPNs (L3VPNs), compared to their neighbouring Layer 2 counterparts, are dependent on Layer 3 protocols such as IP, Multiprotocol Label Switching (MPLS) [13] and virtual routing and forwarding to build the private network. This has its fair share of advantages such as scalability, isolation, and simplified manageability. Traffic engineering for L3VPNs allows the optimisation of bandwidth with complex QoS and convergence, making it a preferred choice for SPs. However, this type of VPN is very SP-dependent with the SP's routing protocol managing all their customer routes in the background. Moreover, in an MPLS L3VPN, the Wide Area Network Internet Protocol address (WAN IP) routing is not directly controlled by the customer, but, instead, the routers at the CPE must peer with the SP's routers. Another tedious task is the IP addressing, which cannot be duplicated between clients. The final point is the dependence of L3VPN on IP traffic, which leads to the compulsory tunnelling of traffic when other protocols such as Internetwork Packet Exchange (IPX) are required. This adds another layer of overhead to an already encapsulated packet, which impacts its resource management and heterogeneity [13][14].

2.1.4 Higher-Layer VPN

Higher-Layer VPNs (HLVPNs) are now the most popular form of commercial VPN for day-to-day use. Due to its simplicity of implementation, HLVPNs are used on almost all web browsers accessing web applications. The two most common HLVPNs are Transport Layer Security (TLS) (and its predecessor Secure Sockets Layer (SSL)) and application-layer VPN, which are usually implemented in tandem [7]. TLS is a security protocol built to improve the security, integrity, and privacy of communications over the Internet using a certificate that validates the original server's identity and authenticity. The data flow is then encrypted using a cipher suite. VPNs built upon TLS/SSL are very advantageous in firewalls and for NAT traversals from remote locations. Due to the abundance of vendors in the TLS industry, the TLS certificates are now less costly. These points make the solution quite manageable with a low outlay. However, it does come with the usual flaws of a VPN, that is, higher latency, and some older versions of TLS are still susceptible to Man in the Middle Attacks while the newer versions, such as TLS 1.3, are not yet supported on all major platforms. TLS solutions are, therefore, less stable with security flaws and no legacy support for TLS 1.3. [15]

An application-layer VPN serves as an intermediate between remote client requests and server-based applications. It closes incoming sessions to the application layer from remote users, processes the data and converts the data to the respective application protocols.

2.2 Overlay Networks

Overlay networks are logical networks built on top of one or more existing networks, using its physical infrastructure. It is a scalable network virtualisation technique as the Internet itself began as an overlay network built upon the physical infrastructure of various telecommunication companies. Data transmitted using overlays on the Internet are, in most cases, done at the application layer of the TCP/IP model, but it can also be implemented at the lower layers. Due to its logical nature, overlay networks can be adjusted and improved without affecting the physical substrate it is sitting upon. This makes it simple and inexpensive to deploy new features or updates. Overlay designs is a popular research topic due to their importance for the Internet and have yielded improvements [16] in performance, convergence, Quality of Service, Denial of Service protection, and testbeds to plan new architectures [7]. However, this type of network cannot surpass the inherent limitations of the current Internet because they are still using IP as a baseline and they do not provide any holistic approach but are only a temporary solution to a permanent problem [17].

2.2.1 Virtual Local Area Networks

A VLAN is a set of logically networked devices in the same broadcast domain. Their physical connectivity does not matter if the Medium Access Control (MAC) header of the frames contains a VLAN ID. These frames are forwarded by switches using the destination MAC address and VLAN ID [7]. Due to their logical nature, it increases the isolation, configuration, administration, and management compared to its physical counterparts.

However, it does bring along many challenges. VLANs work on the same broadcast domain Layer 2 network, which has its handicap. In a campus and datacentre environment, we have a distribution and access layer and, to apply the concept of VLAN in such a network, the broadcast domain has to encompass all the access devices upon which the VLAN itself is required for end-devices to communicate. Thus, a broadcast from one side of the network is delivered to every other port belonging to that VLAN, even if it is on the other side of the network, eventually wasting the bandwidth resources. To counter this, VLAN trunking is used to carry traffic referring to several VLANs by a trunk link which must be properly configured on both ends to provide the list of permissible VLANs on that link. This minimises the spread of excessive VLANs' broadcast traffic but does not eliminate the resource management issue.

Moreover, any misconfiguration [18] can result in partial network failures. Inter-VLAN communication also has some drawbacks. Figure 1 shows such an example where host H1 in VLAN 1 needs to communicate with host H2 in VLAN 2. Being in different broadcast domains, routing is needed for communication to happen. Ideally, given their physical connection, the frame should go through only switch S1 from H1 to H2. However, for routing to take place, the traffic needs to traverse the whole network (substantially longer paths for data flows) to the inter-VLAN designated router R2 which leads to latency, redundant transmission, increased likelihood of packet loss and degraded performances [18].

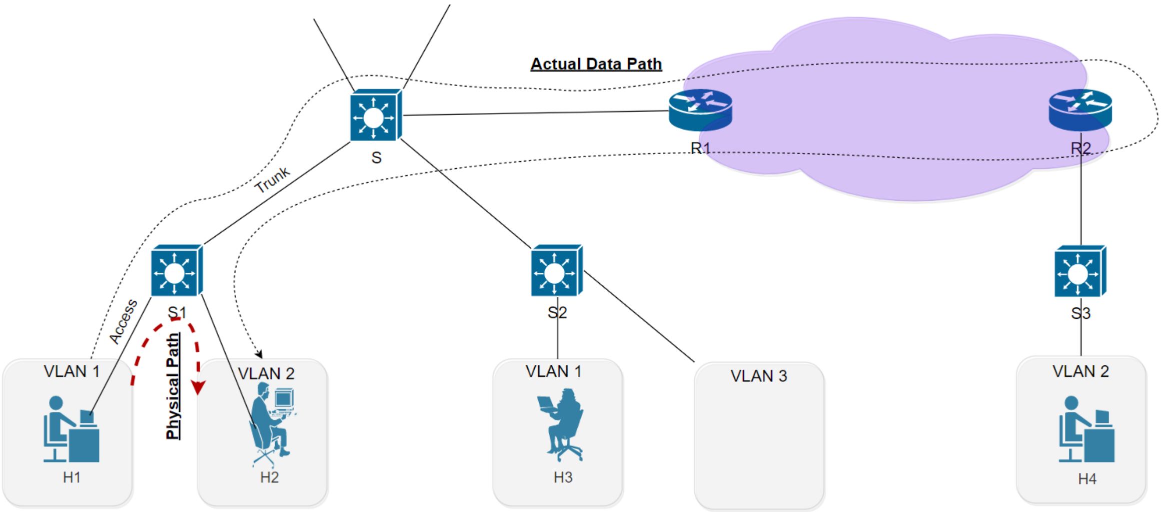

Figure 1. Improper inter-VLAN routing in a network

A near-perfect device placement strategy and design must be set up to avoid such issues which impede scalability. VLANs are also dependent on other protocols to make them scalable in enterprise settings. Cisco introduced its proprietary VLAN Trunking Protocol (VTP) [19] to dynamically propagate VLAN information to all the switches in a VTP domain, otherwise, the network administrator would have the tedious task of manually configuring and managing the VLANs in these switches. The Spanning Tree Protocol (STP) [19] is also required to ensure that this design stays loop-free (which sacrifices free interfaces for this purpose). We now have other proprietary protocols from Cisco, such as Virtual Switching System (VSS) which is a workaround to the shortcomings of VLANs. This also affects the heterogeneity of the whole network.

There are also security concerns which have been discovered in cloud computing environments. The isolation provided by VLAN can be bypassed using methods such as the VLAN Hopping Attack, the Private VLAN Attack and the MAC Flooding Attack [20]. Moreover, the mitigations for these attacks are case-by-case and require human intervention. Another weakness of VLANs is commonly found in datacentres. Only 4094 VLAN IDs can be allocated [21] which is not scalable for SPs. They tend to be in the hundreds if not thousands of customers which may reserve more than one VLAN ID as part of their network segment. This range of VLAN IDs is also limited if the SP is running a multi-tenant architecture [21].

2.2.2 Virtual Extensible Local Area Networks

VXLAN addresses the above issues caused by VLAN by overlaying a virtualised Layer 2 network over a Layer 3 network. In this case, we do not directly use a Layer 3 network to take advantage of the Layer 2 properties. In the example of a multi-tenancy cloud service provider, there may be multiple customers who are utilising the same Layer 3 addresses in their networks. The SP needs to isolate their networks differently. This also limits the customers to use only Layer 3 protocols for their inter-VM connectivity. VXLAN's Layer 2 tunnelling capabilities allow it to create segments within which customer traffic is tagged with a unique VXLAN Network Identifier (VNI). Administrators can thus implement up to 16 million segments, each with 4094 VLANs in the same administrative domain [21][22]. This increases the flexibility of the solution compared to VLANs.

The same tunnelling capabilities, however, are what weigh down this technology. Encapsulation creates a burden on the processing devices' CPU performance, which requires adding and removing protocol headers. This was also addressed by VMware's performance evaluation study [23]. Encapsulation also means a more voluminous frame with more than 50 bytes of a header added. This imposes further overheads on the transport network, causing latency. Since we have an encapsulated header, VXLAN loses its ability to provide differentiated services, so more thought needs to be put into resource management. A sturdy and powerful physical underlay network is needed to support this protocol, which affects its scalability and cost.

2.2.3 Software-Defined Networking

Software-Defined Networking (SDN) is a type of network management approach whereby the control plane is decoupled from the data plane functionalities in the networking devices, which enables the network control to be directly programmable from a controller. This technology is one of the most sought-after virtualisation methods, mainly due to advances in cloud computing [24]. It has been implemented in numerous enterprises and Internet Service Providers (ISPs) and is even the topic of research for wireless networks [25], including 5G cellular networks [26] and satellite [27] applications. This has proved to provide greater scalability and stability.

Even though it is an exceptionally promising virtualisation technology, it also comes with many restrictions and challenges. The main issue is the high cost associated with it. This is quite ironic given its open-source origin with OpenFlow, a multivendor SDN standard. SDN-supporting equipment is expensive compared to traditional networking devices. On top of that, most of those solutions do not provide any backward compatibility with tiered networks. This hampers any hope for legacy support when implementing SDN. Thus, to deploy an SDN solution, we must renew the whole network infrastructure instead of a gradual hardware refresh in phases. This may not be significant for enterprises that are starting up, but, for more established enterprises, this involves a considerable investment, and which may not yield the desired return on investment (ROI) [28].

A prerequisite for this refresh includes a full inventory of the infrastructure, which adds to the total cost of the project. The second issue of SDN is its tendency for homogeneity. SDN solution vendors allow their controllers to work only on their proprietary devices. This business model discourages any kind of multi-vendor compatibility. Vendors like Cisco, VMware and Fortinet have created an ecosystem of their own products, which ultimately forces a full refresh of the enterprise infrastructure [28]. The deployment is not simplified in this case.

2.3 Network Function Virtualisation

Due to the business model of closed-source technologies that IT vendors offer, the industry is left with function-specific hardware middleboxes, which increase the capital expenditures and operational expenditures. For example, we have routers, switches, firewalls (FWs), storage units, load balancers (LBs), intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) with each having a dedicated bare-metal hardware that provides fast processing and performance. However, these devices are costly because of their proprietary nature. Other problems arise with the management of these diverse devices, each with its consoles, rules, rack space, power needs, cabling and setup architectures. This leads to a different skill set for each hardware as well as the complexity of configuration [29]. Network Function Virtualisation (NFV) comes to the rescue by resolving the above constraints. It aims to improve the network's flexibility as well as cost by decoupling the software functions of network services from its underlying dedicated hardware and then implementing them on commercial-off-the-shelf servers.

An SDNFV-based network was used to create virtual honeypots due to the technologies' flexibility [30]. This method would have been costly if IPS and IDS on dedicated middleboxes were utilised instead. NFV is the equivalent networking device abstraction of server virtualisation. These points are illustrated in Figure 2.

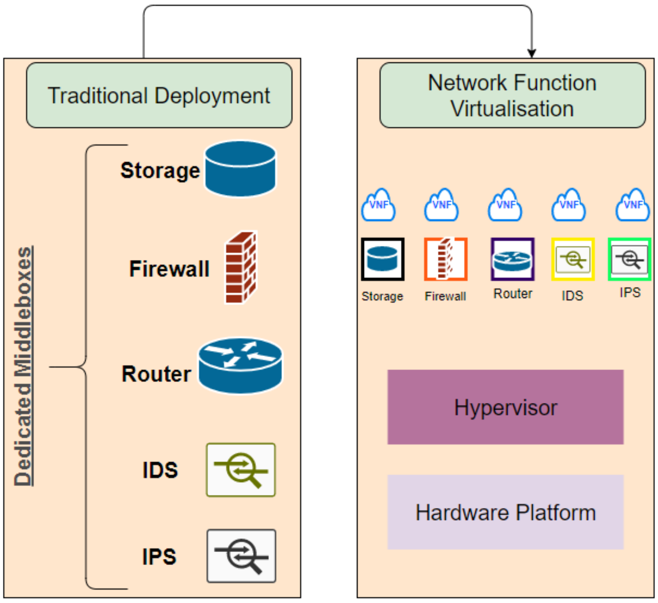

Figure 2. Traditional Deployment and Network Function Virtualisation

This technology is solving many current issues, such as expenditure, firstly by being open-source and secondly by limiting the number of dedicated hardware to be purchased as it can be hosted on different server platforms. This resulted from the successful adoption of server virtualisation [31] which led to virtual network functions (VNFs) replacing traditional network services such as FW, LB and IPS with different VMs running these different network services on a singular host.

Deploying NFVs could inevitably change the way network systems are currently installed and managed. This provides more flexibility and dynamism in computer networks [32], resulting in situations whereby functions that provide a particular customer with a service are spread across numerous server pools. However, this imposes some complex coordination to ensure that these functions can be accessed and processed readily on a per service (or user) level while maintaining its manageability.

NFV is meant to run network services on commercial-off-the-shelf servers whose vendors manufacture equipment without knowledge of what kind of systems and services could be run on them. The responsibility then falls on the shoulders of VNF providers to ensure that their network functions can operate on standard servers. There are concerns expressed on whether functions on standard servers can perform at least similar to that of specialised hardware, and whether such functions are versatile between servers [33]. One such observed issue is the disparity of latency and I/O throughput between NFV and dedicated hardware. This is due to NFV's disposition to shared resources whereby the hardware network interface card (NIC) intercepts a packet which is then forwarded to the kernel to be processed. It is then forwarded to the appropriate userspace. This whole operation is what causes increased latency.

There are some solutions that can cater to that issue, namely a Data Plane Development Kit (DPDK) which is a data plane acceleration programming environment. It uses poll mode drivers which bypass the OS networking stack in the kernel and push the packet from the NIC to the user space seamlessly. This also releases the CPU for other processing tasks [34]. Another solution is single root I/O virtualisation (SR-IOV) which enables VMs to share a hardware component for all operations without needing the hypervisor to be involved in the decision-making. This helps also in the portability of VNFs, that is, its ability to be migrated across different servers in a multi-vendor environment. However, these techniques come at a complexity and flexibility price for migration as well as on the layer of the VNFs themselves. They have to be deployed with hardware-specific drivers to operate those components. Otherwise, this would make all those acceleration technologies useless [35]. This resource management contradiction impacts the whole purpose of implementing NFV.

Little research has been done on the energy consumption generated by NFVs, which is quite a big concern for telecommunication service providers who spend more than 10% of their operational expenditure on energy. It can be argued that better energy efficiency can be achieved with NFVs as less hardware is required and fewer cooling mechanisms in place [36]. However, when NFVs, like other trending technologies, are transferred to the cloud, then the change in energy efficiency should remain the same [36, [37].

The design of the NFV dictates the order in which the traffic flows from VNF to VNF and this depends on the services being provided and the users. This is called service function chaining. Moreover, if the initial design itself is at fault, then the resources of the hardware cannot be optimised. The placement of the VNF must be carefully planned, simulated, and then implemented. The failure of doing so can lead to sub-optimal resource distribution as well as over-consumption of energy and inacceptable SLAs [29, 36]. This brings another layer of complexity to the network design, which may discourage administrators. Finally, as NFV is getting more popular, it is attracting more security attacks and, due to its virtual nature, it is more susceptible to vulnerabilities originating from the source code to data interception.

The review indicates that the past research on network virtualisation technologies overlooked some features. Most technologies band-aided the gap in virtualisation, mostly by a focus on specific deficiencies.

3. Methodology and Requirements for Network Virtualisation

This section provides the methodology adapted to address the research questions. To gain a better insight into the possibilities for improvement of network virtualisation, a thematic analysis with a systematic literature review method using a secondary data source was done. In this case, the secondary data sources were research papers on network virtualisation technologies. These papers were selected by using the inclusion criteria of higher impact factor journals and conference articles within the last fifteen years. This also helped in seeing the progress of network virtualisation up until that point. Out of fifty-five research papers, reports and articles that were investigated, thirty-six were chosen for the review. These papers were from ACM workshops on virtualised infrastructure systems and architectures, ITC Specialist Seminar on Network Virtualisation, Cyber-Enabled Distributed Computing and Knowledge Discovery, International Conference on Reliability, Infocom Technologies and Optimisation (Trends and Future Directions) (ICRITO), International Conference on Networking and Distributed Computing, IEEE journals and conferences, Optical Switching and Networking journal, ICT Express Journal, Computer Networks Journal, International Journal of Communication Networks and Information Security, technical reports, and industry-specific white papers and patents. Finally, a comparison table of the discussed technologies is drawn and assessed based on the design goals for perfecting the NVE.

The findings of this study are assessed using the eight virtualisation design goals discussed by Mosharaf, Chowdhury and Boutaba [7] as well as five other important criteria that are critical for the development of NVEs. They start off with flexibility. Network virtualisation must be flexible in its networking operations in the sense that it should not be dependent or limited by the underlying infrastructure, neighbouring ISP or coinciding virtual network. The second criterion is manageability. Network virtualisation must simplify the network administration activities as well as making them modular. It should also bring complete control throughout the network, even across different autonomous systems (AS).

Scalability would be the third criterion where NVEs must scale to accommodate increasing demands from the networking industry without compromising performance. Fourthly, we have isolation since NVEs are required to provide total seclusion and protection. The absence of interaction between co-existing VNs ensures privacy and integrity of the data. As for the fifth criterion, stability and convergence, network virtualisation must be resistant to and quick to recover from any underlying component failure or misconfiguration, always providing for a stable network.

NVEs must also have the capability to be highly programmable to create an adaptable technology, which leads to the sixth criterion, programmability. This helps in the above design goals, flexibility, and manageability. The seventh criteria, heterogeneity should be applied in both the physical underlay as well as in the logical network context. In the former, infrastructure providers (InPs) can implement different physical underlay, while in the latter, SPs can support any virtualisation methods and protocols without causing incompatibility issues. As network virtualisation is a phased deployment, it should integrate the eighth criterion, legacy support (also known as backward compatibility) to be able to effectively utilise those waning resources. This should be quite simple to realise considering that the Internet itself consists of a set of legacy resources upon which virtual networks are built. However, it remains a concern whether and how it is achieved effectively.

Moreover, there are five more criteria that are crucial to the adoption of network virtualisation. The ninth criterion is outlay, which is a big factor in deciding whether the move to a perfect NVE is feasible or not. It dictates the business model and profit return of the SPs, even if an open-standard model is preferred. It should thus be a financially sound model that facilitates ROI. The tenth criterion is robustness of the technology. Mishaps and network failures are inevitable; however, the technology should not contain any gaps and should minimise any points of failure. The eleventh criterion is the ease of deployment. Flexibility, manageability, programmability, heterogeneity, and legacy support become the backbone of deployment and their culmination dictates how easily it can be done. If those criteria are not synergised, it could discourage any attempt at making the jump to the perfect NVE.

The twelfth criterion may be considered the most important one. Security is required to ensure confidentiality and integrity of the data even in the event of misconfigurations and attacks. Industry standards are more stringent than ever with many cryptographic techniques on the rise. As virtualisation becomes increasingly deployed, it is also necessary to incorporate those standards, techniques and other possible defences against malicious operators and users in the design of the NVE itself. The final criterion would be resource management, which is the most important subset of the second criterion, manageability, and deserves to be a design goal of its own. The dynamic way of underlying resources is administered in the mapping phase and, in addition, it has a significant role to play in developing infrastructure resources.

Currently, each technology is geared towards resolving a specific insufficiency or even gap in virtualisation. This makes them suitable for only individual solutions. In the case of VPNs, the main goal is to provide anonymity and isolation of data while, for NFV, it is to eliminate the need of specialised networking devices. VXLANs are merely an improved extension of VLANs, which are now considered a legacy technology. SDN, even though it has shown potential as an open-standard development, is being made vendor-dependent as a business model [28]. There is no end-to-end heterogeneous solution, but, rather, an amalgamation of these technologies, each with their drawbacks and advantages.

4. Comparison of Network Virtualisation Technologies

Table 1. Comparison of network virtualisation technologies

Technology Criteria	VPN	VLAN	VXLAN	SDN	NFV
Security	Data are always encrypted when passing over shared networks.	VLAN transport networks can be tampered with.	Like VLAN, VXLAN transport network can be tampered with.	Due to its granular approach from a central controller and open-source nature, SDN is one of the most secure technologies.	The source code can be susceptible to security attacks.
Isolation	VPN makes use of tagging mechanisms to ensure isolation.	Isolation of data is ensured by VLAN ID, but protocol cannot be used individually over shared media.	Isolation of data is ensured by VXLAN Network Identifier which can be overlaid on a Layer 3 network.	SDN has been subject to network isolation attacks.	This depends on the solution used.
Stability and Convergence	SPs have great convergence methods based on the underlying routing and switching devices.	Convergence is achieved using supporting protocols like STP.	This solution, like VLAN, requires supporting protocols for convergence.	Convergence behaviour rivals optimal routing protocols [38].	Stability and convergence depend on underlying hardware resiliency.
Outlay	This can range from an inexpensive generic proxy VPN to a full-fledged, high-cost multi-layered datacentre VPN.	This depends on the Layer 2 architecture to be implemented.	Quite a low-cost solution but requires a powerful physical underlay network.	Solution itself is costly depending also on the vendor and can require a whole network infrastructure refresh.	Inexpensive solution which can be used on commercial-off-the-shelf servers.
Scalability	Its scalability is hindered by the compatibility to its physical substrate technology.	VLAN is now very scalable due to the widespread use of Layer 2 devices.	Very scalable solution due to its increased segment capacity over VLANs.	This solution is scalable only in a vendor-specific SDN ecosystem.	Very scalable solution due to its support on different servers.
Robustness	Depending on the deployment method, VPNs offer a robust network, limited only by the underlay.	VLANs rely on full/partial mesh topologies to offer robustness.	Similar to VLANs, VXLANs rely on full/partial mesh topologies to offer robustness.	It requires multiple SDN controllers to eliminate points of failure [39].	High availability can be achieved via well-planned design of the NFV.
Manageability	Depending on the VPN used, manageability can vary.	Manageability is difficult because of the number of physical devices and supporting protocols used.	Manageability is better than VLAN because of fewer physical devices.	High manageability due to the separation of the control plane from the data plane.	Manageability of the solution depends on the number of machines used to host it.
Resource Management	Resource management varies according to the layer as well as the SP.	Resource wastage in the form of a limited number of VLANs.	Resources are conserved with increased number of segments and fewer devices.	Many resource management solutions have been proposed.	Innovative use of commercial-off-the-shelf servers to reduce resource-intensive middleboxes, but performance is impacted.
Programmability	Quite complex programmability depending on the VPN selected.	Manual configuration is required, and it uses supporting protocols like VTP.	Easier programmability compared to VLAN.	Solutions closer to the open-source OpenFlow protocol is, therefore, easier to implement.	Programmability depends on the solution selected.
Flexibility	Due to its multi-layered ability, VPNs can be a flexible solution to use on different network topologies or SPs.	This solution can only work on Layer 2 VLAN-enabled devices.	This solution is a better contender than VLAN by working over both Layer 2 and 3.	SDN becomes inflexible due to its specific hardware requirements.	NFV can implement all sorts of network services on a single machine.
Heterogeneity	This solution can run across different SPs, each running different vendors' hardware.	Being open source, it can run on different vendors' hardware.	Being open source, it can run on different vendors' hardware.	This solution needs SDN-specific hardware from a unique vendor.	Its main ability is to run on different vendors' hardware.
Legacy Support	Due to their TCP/IP layer flexibility, they can easily work on any legacy hardware.	Being an old virtualisation technique, it is compatible on most legacy hardware.	This solution can run alongside VLAN for legacy support integration.	This solution can require a whole network infrastructure refresh, so offers little to no legacy support.	As per European Telecommunications Standards Institute, NFV can be implemented alongside legacy hardware.
Ease of Deployment	The ease of deployment differs with each TCP/IP Layer.	There is much manual intervention in the deployment process.	There is much manual intervention in the deployment process.	Due to lack of expertise, high cost and no legacy support, SDN is difficult to deploy.	Similar to SDN, specific skillset is needed to deploy NFV as well as a well-planned design.

The results are summarised in Table I to assess how the challenges presented stack up. Each technology demonstrates some level of pros and cons and some features must be sacrificed to promote others. Technologies like NFV and SDN must forgo ease of deployment and resource management to focus on scalability and robustness. Another problem afflicting the progress of network virtualisation is the privatisation of some open-standard protocols to enhance some solutions. Though, it does come with advantages, like technical support and funded research and development, solutions like SDN, are, however, affected due to the high outlay proposed and the specific hardware imposed by the vendor. Apparently, legacy support and heterogeneity are rarely considered. Moreover, the skillset required to manage those solutions are also privatised with vendor-specific certifications. Most of these technologies are used in conjunction with one another in modern network architectures as they complement each other's drawbacks. This should not be the hallmark of NVEs.

5. Discussion and Future Research Directions

Our results demonstrated that current virtualisation technologies are not at par with what would have been expected from its market growth. Many of them are lacking in their design goals. While we could consider the creation of novel NVE approaches, we should ensure that existing technologies have been exploited to the best of their capacities or improve them after a comprehensive analysis of their recurrent drawbacks. Further research can bring to light a novel approach to revamp the existing network virtualisation whereby a new NVE combines the forte of the existing technologies to amplify the robustness, the management properties of SDN, the abstraction of NFVs, and the ability to route across all TCP/IP layers of VPNs.

A more realistic research goal is its impact on the future internet since network virtualisation has a major role in shaping it by overcoming the Internet ossification problem. This makes the implementation of a novel NVE a more long-term goal. However, security and isolation concerns are discouraging many competitive industries to fund research in network virtualisation, so they stick with proven VPN models which offer security at the cost of other features. This is also impacted by the outlay of the technology and is also hindered by the privatisation of the standards rather than developing upon open-standard protocols. Adoption of new NVEs which comply with the proposed design goals will lead to new business models which will, in turn, greatly affect the progress of network virtualisation. There is currently little or insufficient research that are being done on the impacts of such technologies.

Properties like stability, robustness and manageability are often dependent on the physical underlay, which needs to be upgraded to improve them. This upgrade requires any legacy hardware or software to be replaced to improve compatibility and homogeneity in the architecture. This chain reaction exacerbates properties such as outlay (to renew architecture), flexibility (more physical dependence instead of logical), programmability (new skillset is now required) and ease of deployment (network refresh with specialised upskills). This vicious cycle invokes an esoteric connotation which does not help the field to progress. Network virtualisation proponents should refer first to those design goals and imbue them in the design phase rather than applying patches after implementation. Furthermore, these design goals should include other new out-of-the-box performance metrics for assessing network virtualisation technologies.

6. Conclusions

Network virtualisation is not the final product but more a means to an end in that it cannot outgrow the limitations imposed by its underlay, vendors, or service providers' business models. However, this will still be the case if a complete redesign is not imposed. Network virtualisation is already playing a lead role in paving new and innovative network architectures for large business and governmental organisations. Network virtualisation has mostly been utilised to apply narrow fixes to random and specific problems without any holistic view as their design cannot go beyond the inherent limitations of their physical underlays. With the emergence of more powerful, diverse, and versatile network media, the seven existing assessment criteria and the six new that have been proposed, the true value and potential of virtual networks will be realised. The challenges discussed in this paper may be used as a springboard to create better network virtualised environments.

References

[1] A. Cordero and D. Fernández, "Rethinking ISP architectures through virtualization," *Proceedings of the International Conference on*, Seville, Spain, 2011, pp. 1-4.

[2] D. Donohue and J. Kronik, "The Art of Network Architecture: Applying Modularity,", Cisco Press, Cisco, San Jose, California, May 12, 2014.

[3] Spiceworks Inc. "The 2020 State of Virtualization Technology." spiceworks.com. https://www.spiceworks.com/marketing/reports/state-of-virtualization/ (accessed Apr. 8, 2021)

[4] C. Liang and F. R. Yu, "Wireless Network Virtualization: A Survey, Some Research Issues and Challenges," in *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 358-380, Firstquarter 2015, doi: 10.1109/COMST.2014.2352118.

[5] X. Wang, P. Krishnamurthy and D. Tipper, "Wireless network virtualization," *2013 International Conference on Computing, Networking and Communications (ICNC)*, San Diego, CA, 2013, pp. 818-822, doi: 10.1109/ICCNC.2013.6504194.

[6] Oracle Corporation, "Using Virtual Networks in Oracle Solaris 11.1," October 2012 https://docs.huihoo.com/solaris/11.1/english/pdf/E28992.pdf (accessed Dec. 22, 2020).

[7] N. M. Mosharaf, Kabir Chowdhury and R. Boutaba, "Network virtualization: state of the art and research challenges," in *IEEE Communications Magazine*, vol. 47, no. 7, pp. 20-26, Jul. 2009, doi: 10.1109/MCOM.2009.5183468.

[8] J. Carapinha and J. Jiménez, "Network virtualization: a view from the bottom," in *Proceedings of the 1st ACM workshop on Virtualized infrastructure systems and architectures*, New York, USA, Aug. 2009, pp. 73-80, doi: 10.1145/1592648.1592660.

[9] S. Khanvilkar and A. Khokhar, "Virtual private networks: an overview with performance evaluation," in *IEEE Communications Magazine*, vol. 42, no. 10, pp. 146-154, Oct. 2004, doi: 10.1109/MCOM.2004.1341273.

[10] A. El Amri and A. Meddeb, "Resource Allocation Heuristics for Network Virtualization," *2017 IEEE/ACS 14th International Conference on Computer Systems and Applications (AICCSA)*, Hammamet, Tunisia, 2017, pp. 55-62, doi: 10.1109/AICCSA.2017.116.

[11] N. Malheiros, E. Madeira, F.L. Verdi and M. Magalhães, "Managing Layer 1 VPN services," *Optical Switching and Networking*, vol. 5, no. 4, pp. 196-218, Oct. 2008. doi: 10.1016/j.osn.2008.02.002.

[12] T. Takeda, R. Aubin, M. Carugi, I. Inoue and H. Ould-Brahim, "RFC4847: Framework and Requirements for Layer 1 Virtual Private Networks," Apr. 2007 [Online]. Available: https://tools.ietf.org/html/rfc4847.

[13] G. A. Mazhin, M. Bag-Mohammadi, M. Ghasemi and S. Feizi, "Multi-layer architecture for realization of network virtualization using MPLS technology," *ICT Express*, vol. 3, no. 1, pp. 43-47, Mar. 2017, doi: 10.1016/j.icte.2016.07.002.

[14] A. Sayeed and M. Morrow., "Technology Overview: Making the Technology Case for MPLS and Technology Details,", Cisco Press, Cisco, San Jose, California, Jan. 12, 2007.

[15] S. Turner, "Transport Layer Security," in IEEE Internet Computing, vol. 18, no. 6, pp. 60-63, Nov.-Dec. 2014, doi: 10.1109/MIC.2014.126.

[16] P. Papadimitriou, O. Maennel, A. Greenhalgh, A. Feldmann and L. Mathy, "Implementing network virtualization for a future internet" presented at the 20th ITC Specialist Seminar on Network Virtualization Proceedings, Hoi An, Vietnam, May 18-20, 2009.

[17] T. Anderson, L. Peterson, S. Shenker and J. Turner, "Overcoming the Internet impasse through virtualization," *Computer*, vol. 38, no. 4, pp. 34-41, Apr. 2005, doi: 10.1109/MC.2005.136.

[18] P. Garimella, Y. W. E. Sung, N. Zhang and S. Rao, "Characterizing VLAN usage in an operational network," in *Proceedings of the 2007 SIGCOMM workshop on Internet network management*, pp. 305-306, Aug. 2007, doi: 10.1145/1321753.1321772.

[19] R. Froom and E. Frahim, *Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide*, San Jose, California, Pearson Education, Cisco Press, 2015.

[20] K. Benzidane, S. Khoudali, L. Fetjah, S. J. Andaloussi and A. Sekkaki, "Application-based authentication on an inter-VM traffic in a cloud environment," International Journal of Communication Networks and Information Security, Vol. 11, No. 1, pp.148-166, Apr 2019.

[21] M. Mahalingam, D. G. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridhar, M. Bursell and C. Wright, "RFC7348: Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks," August 2014 [Online]. Available: https://tools.ietf.org/html/rfc7348.

[22] D. Farinacci, T. Speakman, N. Venugopal, H. Grover, V. Moreno and D. Rao, "Overlay transport virtualization," U.S. Patent 8 166 205 B2, Apr. 24, 2012.

[23] R. Mehta, "VXLAN Performance Evaluation on VMware vSphere® 5.1," VMware Inc., Palo Alto, California, United States, 2013, [Online]. Available:

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-vsphere-vxlan-perf-white-paper.pdf.

[24] R. Jain and S. Paul, "Network virtualization and software defined networking for cloud computing: a survey," in *IEEE Communications Magazine*, vol. 51, no. 11, pp. 24-31, Nov. 2013, doi: 10.1109/MCOM.2013.6658648.

[25] N. Zhang, P. Yang, S. Zhang, D. Chen, W. Zhuang, B. Liang and X. S. Shen, "Software Defined Networking Enabled Wireless Network Virtualization: Challenges and Solutions," in *IEEE Network*, vol. 31, no. 5, pp. 42-49, 2017, doi: 10.1109/MNET.2017.1600248.

[26] J. S. Al Azzeh, A. Mesleh, Z. Hu, R. Odarchenko, S. Gnatyuk and A. Abakumova, "Evaluation Method for SDN Network Effectiveness in Next Generation Cellular Networks," International Journal of Communication Networks and Information Security Vol. 10, No. 3, pp. 472, Dec 2018.

[27] C. Wang and X. Yu, "Application of Virtualization and Software Defined Network in Satellite Network," presented at the 2016 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), Chengdu, Oct. 2016, doi: 10.1109/CyberC.2016.99.

[28] B. Sokappadu, A. Hardin, A. Mungur and S. Armoogum, "Software Defined Networks: Issues and Challenges," *2019 Conference on Next Generation Computing Applications (NextComp)*, Mauritius, 2019, pp. 1-5, doi: 10.1109/NEXTCOMP.2019.8883558.

[29] K. Kaur, V. Mangat and K. Kumar, "Architectural Framework, Research Issues and Challenges of Network Function Virtualization," presented at the 2020 8th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO), Noida, India, Jun. 2020, doi: 10.1109/ICRITO48877.2020.9197802.

[30] N. C. Thang, M. Park and Y. I. Joo, "EVHS-Elastic Virtual Honeypot System for SDNFV-Based Networks." International Journal of Communication Networks and Information Security, Vol. 12, No. 3, pp.295-301, Dec 2020.

[31] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862-876, Apr. 2010, doi: 10.1016/j.comnet.2009.10.017.

[32] S. Cherrared, S. Imadali, E. Fabre, G. Gössler and I. G. B. Yahia, "A Survey of Fault Management in Network Virtualization Environments: Challenges and Solutions," in *IEEE Transactions on Network and Service Management*, vol. 16, no. 4, pp. 1537-1551, Dec. 2019, doi: 10.1109/TNSM.2019.2948420.

[33] M. Chiosi et al., “Network Functions Virtualisation (NFV) Network Operator Perspectives on Industry Progress.” AT&T, BT, Cablelabs, CenturyLink, China Mobile, 2013, doi: 10.13140/RG.2.1.4110.2883.

[34] M. A. Kourtis, G. Xilouris, V. Riccobene, M. J. McGrath, G. Petralia, H. Koumaras, G. Gardikis and F. Liberal, “Enhancing VNF performance by exploiting SRIOV and DPDK packet processing acceleration,” in *2015 IEEE Conference on Network Function Virtualization and Software Defined Network*, NFV-SDN 2015, pp. 74–78.

[35] B. Chatras and F. F. Ozog, "Network functions virtualization: the portability challenge," in *IEEE Network*, vol. 30, no. 4, pp. 4-8, Jul. 2016, doi: 10.1109/MNET.2016.7513857.

[36] R. Mijumbi, J. Serrat, J. Gorricho, N. Bouten, F. De Turck and R. Boutaba, "Network Function Virtualization: State-of-the-Art and Research Challenges," in *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 236-262, Firstquarter 2016, doi: 10.1109/COMST.2015.2477041.

[37] Z. He and G. Liang, "Research and Evaluation of Network Virtualization in Cloud Computing Environment," in *2012 Third International Conference on Networking and Distributed Computing*, Hangzhou, China, Oct. 2012, pp. 40-44, doi: 10.1109/ICNDC.2012.18.

[38] S. Abdallah, A. Kayssi, I. H. Elhajj, and A. Chehab, “Network convergence in SDN versus OSPF networks,” presented at the 2018 Fifth International Conference on Software Defined Systems (SDS), Apr. 2018, doi: 10.1109/sds.2018.8370434.

[39] D. K. Ryait and M. Sharma, "Significance of Controller in Software Defined Networks," *2020 IEEE 15th International Conference on Industrial and Information Systems (ICIIS)*, RUPNAGAR, India, 2020, pp. 561-566, doi: 10.1109/ICIIS51140.2020.9342710.