

REVIEW ARTICLE

Scalable data center networking: Evaluating VXLAN EVPN as a next-generation overlay solution

Vaidehi Shah*

Independent Researcher

Received on: 13-10-2023; Revised on: 25-11-2023; Accepted on: 17-12-2023

ABSTRACT

The rapid evolution of cloud computing, artificial intelligence (AI), and distributed applications has transformed the operational landscape of modern data centers, demanding scalable, flexible, and resilient networking solutions. Traditional data center network (DCN) architectures often hardware-bound and statically configured struggle to accommodate the growing volume of east-west traffic and the need for automation and multitenancy. This paper presents an in-depth examination of virtual extensible local area network-Ethernet virtual private network architectures, detailing their operational principles, advantages, and implementation strategies in modern data centers. A comprehensive literature review is conducted to highlight current advancements, real-world use cases, and ongoing challenges, including scalability limits, verification inefficiencies, and energy consumption. Comparative analysis of recent works is provided to outline contributions, research gaps, and opportunities for further exploration. To satisfy future expectations, the study's conclusion highlights the significance of integrated, software-defined, and adaptable network systems. This work aims to guide network architects and researchers in designing secure, resilient, and next-generation DCNs aligned with emerging digital and cloud-native ecosystems.

Key words: Data center networking, Ethernet virtual private networks, network scalability, overlay networks, virtual extensible local area network

INTRODUCTION

The rapid expansion of cloud computing has transformed the digital landscape by enabling highly flexible, scalable, and on-demand telecommunication and IT services. These advancements are made possible through the implementation of large-scale data centers, which interconnect numerous hosts through high-speed networks. At the core of these infrastructures lies the data center network (DCN), a critical component responsible for supporting seamless data exchange, service delivery, and efficient resource utilization.^[1] For modern DCNs to offer transparent and continuous cloud operations, it must meet four essential requirements: auto-configuration, scalability, flexibility, and high availability, which includes resilience and fast failure recovery.^[2] In response to these requirements, a range of DCN

architectures has been proposed over the past decade. Although some designs, such as Portland, successfully address certain aspects such as simplified configuration and availability, they typically fall short of delivering all four goals concurrently. This shortfall is particularly evident as data center demands grow, driven by rising data volumes and diverse application needs. These limitations highlight a key challenge: Traditional DCNs, which rely heavily on static, hardware-based designs, are becoming increasingly inadequate in dynamic and data-intensive environments.^[3]

As these challenges intensify, there is a growing need to decouple DCN functionalities from rigid hardware constraints. This has led to the adoption of virtualization within modern data center environments, enabling network functions to operate across physical, hybrid, and multi-cloud infrastructures.^[4] By abstracting network services from physical components, virtualization improves resource utilization, facilitates rapid deployment, and enhances operational agility. However, realizing these benefits at scale requires

Address for correspondence:

Vaidehi Shah

E-mail: shahvaidehi4795@gmail.com

robust overlay solutions that can support multi-tenancy, fault tolerance, and high-performance routing. One of the most promising technologies in this domain is the virtual extensible local area network (VXLAN) in combination with Ethernet virtual private network (EVPN). Scalable Layer 2 overlays over Layer 3 networks are made possible by VXLAN, while EVPN serves as an efficient control plane that enhances mobility, convergence, and policy enforcement. Together, VXLAN EVPN addresses many of the core limitations of earlier overlay models, particularly in terms of scalability, automation, and interoperability, making it a viable solution for next-generation data centers. This paper examines the role of VXLAN EVPN as a next-generation overlay solution, analyzing its architecture, capabilities, and performance in comparison to legacy and contemporary alternatives. The discussion is supported by technical insights and a comparative evaluation, aiming to guide network engineers and decision-makers in adopting scalable, resilient, and future-ready DCN designs.

Paper Organization

This paper is structured as follows: Section II outlines the concept of scalability in data centers, including scaling methods and the transition from traditional to modern architectures. Section III delves into the architecture and operational principles of VXLAN and EVPN, highlighting their benefits. Section IV reviews recent literature on network verification, identifying existing limitations and research gaps. Section V concludes with insights and future research directions aimed at advancing scalable and adaptive next-generation DCNs.

OVERVIEW OF SCALABILITY IN DCN

The capacity of a data center to operate efficiently as the quantity and size of its internal components, such as servers, increase is known as data center scalability. It is clear that data rise in the future; consequently, the data center must include a scaling function to effectively align with customer demand. The primary advantage of data center scalability is shown in Table 1, its capacity to enable a corporation to swiftly adjust its resources to fulfil its customers' precise needs.^[5] In addition, it offers the benefits of high density, energy efficiency, and

cost savings throughout the lifespan of data center investments. To ensure the scalability of a data center, firms must choose a site that can be easily expanded in the future.^[6] To thrive in a competitive environment, firms must adopt a strategic approach to data center implementation and scalability. The NETFORCHOICE data center facilitates the ability to adapt to evolving business and consumer requirements, adhere to stringent laws, and mitigate emerging risks. The hyper-converged infrastructure with scaling technologies is the premier purpose-built infrastructure for hosting contemporary virtualized data centers and hybrid clouds.

Evolution of DCN Architectures


The rapid growth of cloud computing, AI-driven workloads, and hybrid environments has exposed the limitations of traditional three-tier data center network (DCN) architectures. These legacy designs struggle to meet the performance, scalability, and flexibility demands of modern applications. Today's workloads require architectures that support high volumes of east-west traffic, ensure low latency, and seamlessly integrate with cloud platforms—all while maintaining robust security and regulatory compliance.^[7] Challenges such as data sovereignty, high bandwidth needs, and zero-trust security, compounded by supply chain constraints and market disruptions, are pushing organizations toward automated, resilient, and API-driven network solutions aligned with cloud-native principles. A three-tier DCN architecture is shown in Figure 1.^[8] Unfortunately, even the best enterprise-class equipment may only be able to supply 50% of the total bandwidth from end to end. Key architectural approaches and implementation strategies that enable businesses to build resilient, scalable, and secure network infrastructures aligned with contemporary demands. From foundational architectural choices to advanced fabric designs, security implementations, and automation frameworks, they will explore the essential building blocks of a modern DCN and provide practical guidance for organizations at any stage of their network transformation journey.

The Importance of DCN

DCN is crucial for establishing a stable, dependable, scalable, and secure network

Table 1: Data center scaling methods

Method	Description	Pros	Cons
Adding more servers	Also known as scaling out or horizontal scaling, this involves adding more physical or virtual machines to the data center architecture.	<ul style="list-style-type: none"> ✓ Can support and distribute more workloads ✓ Eliminates hardware constraints 	<ul style="list-style-type: none"> ✓ Deployment and replication take time ✓ Requires more rack space ✓ Higher upfront and operational costs
Virtualization	Dividing physical hardware into multiple virtual machines or virtual network functions to support more workloads per device.	<ul style="list-style-type: none"> ✓ Supports faster provisioning ✓ Uses resources more efficiently ✓ Reduces scaling costs 	<ul style="list-style-type: none"> ✓ Transition can be expensive and disruptive ✓ Not supported by all hardware and software
Upgrading existing hardware	Also known as scaling up or vertical scaling, this involves adding more processors, memory, or storage to upgrade the capabilities of existing systems.	<ul style="list-style-type: none"> ✓ Implementation is usually quick and non-disruptive ✓ More cost-effective than horizontal scaling ✓ Requires less power and rack space 	<ul style="list-style-type: none"> ✓ Scalability limited by server hardware constraints ✓ Increases reliance on legacy systems
Using cloud services	Moving some or all workloads to the cloud, where resources can be added or removed on demand to meet scaling requirements.	<ul style="list-style-type: none"> ✓ Allows on-demand or automatic scaling ✓ Better support for new and emerging technologies ✓ Reduces data center costs 	<ul style="list-style-type: none"> ✓ Migration is often extremely disruptive ✓ Auto-scaling can lead to ballooning monthly bills ✓ May not support legacy software

**Figure 1:** Data center networks architecture^[9]

architecture. It ensures that infrastructure aligns with an organization's evolving communication needs while facilitating cloud computing and virtualization.^[10] Moreover, business DCN solutions improve operational consistency, automation, and security, which are essential for the successful delivery of data and application services. DCN is essential due to its:

- Consolidates resources and upholds operational uniformity by streamlining the integration of resources across cloud, on-premises, and edge environments, guaranteeing consistent rules and centralized control from a single interface.
- Facilitates network service delivery through job automation, workload balancing, and efficient network programming, hence allowing rapid adaptation to evolving requirements.
- Facilitates troubleshooting by offering an extensive overview of the network, hence simplifying the identification and resolution of problems promptly.

- Enhances data center security by shielding applications and data through integrated security measures such as an intrusion detection system, intrusion prevention system, and micro segmentation, hence mitigating cyber risks.

VXLAN AND EVPN: CONCEPTS AND ARCHITECTURE

A standards-based system called EVPN provides a virtual multipoint bridged connection between many domains through an Internet Protocol (IP) and the IP/multiprotocol label switching backbone network. Continuous, multi-tenant, flexible services that may be expanded on demand are made possible by EVPN. An extension of border gateway protocol (BGP) called EVPN allows the network to simultaneously use L2 media access control (MAC) and L3 IP data to enhance routing and switching choices.^[11] Such control plane technologies employ multiprotocol

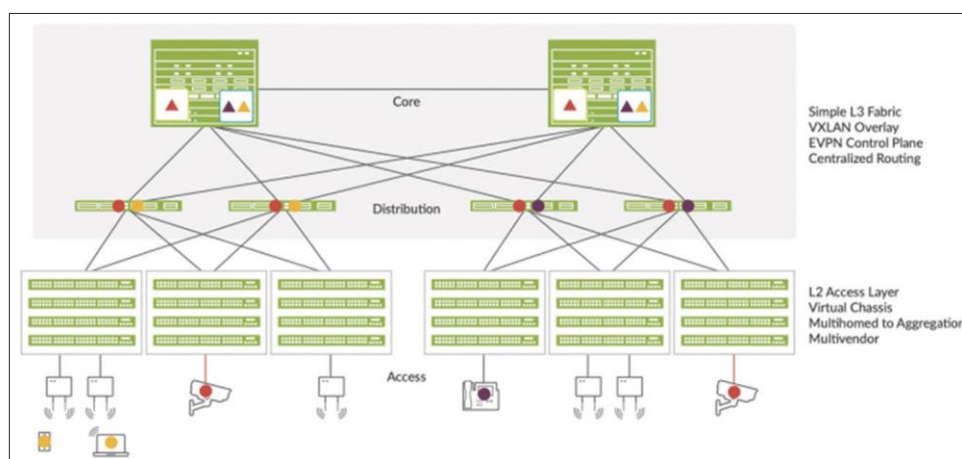


Figure 2: Ethernet virtual private network-virtual extensible local area network-based campus architecture

BGP, which is intended for MAC and IP addresses at the endpoint distribution, where MAC addresses are considered routes. Devices operating as virtual tunnel endpoints can share reachability data about their endpoints with one another thanks to EVPN. Businesses may readily add more core, distribution, and access layer devices to an expanding firm using an EVPN-VXLAN-based campus architecture, as shown in Figure 2, without having to rethink the architecture and buy a new set of equipment.

In addition to support for Layer 2 and Layer 3 virtual private networks (VPNs), businesses may employ a consistent set of policies and services across campuses. Campus network operators should be able to create far bigger networks than would be possible with conventional Layer 2 Ethernet-based designs by combining an EVPN-VXLAN overlay with a Layer 3 IP-based underlay. Figure 3 illustrates how an IP fabric design with an EVPN-VXLAN overlay is often used in the most recent large-scale data centers.

The two-tier spine-and-leaf architecture of the IP fabric, which is suited for large-scale environments, enables to collapse of traditional networking layers. This highly interconnected Layer 3 network may be readily scaled horizontally as needed and acts as an underlay to provide high resilience and low latency across the network. The EVPN-VXLAN overlay, which is placed on top of the IP fabric, allows for extending and connecting the Layer 2 data center domains and place endpoints (such as servers or virtual machines) anywhere in the network, even between data centers. Some of the benefits offered by VXLAN EVPN systems include the following.

- Standards-based control plane (BGP) and standards-based overlay (VXLAN)

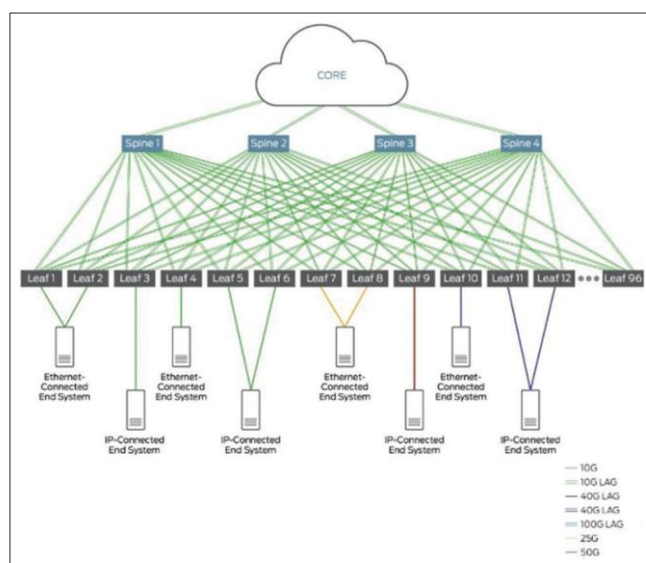


Figure 3: Data center fabric architecture

- Software-defined networking (SDN) facilitation
- The control plane (BGP) distributes Layer 2 MAC and Layer 3 IP information
- Combining virtual and physical networks with hybrid overlays
- Optimizing forwarding in the overlay using integrated routing/bridging
- Facilitation of SDN
- Making decisions about forwarding based on a scalable control plane reduces flooding
- Leverages Layer 3 equal cost multi-path in the underlay for complete link forwarding
- The overlay's namespace is noticeably larger (16M segments).

VXLAN Overlay

A network overlay is a technique used in contemporary data centers to virtualize a network and create an architecture that is more flexible

than a network that is inherently static. Instead, it is important to spend some time understanding why typical networks are so static before getting into the specifics of how overlays work, the challenges they encounter, and how to get beyond these obstacles. The first message sent by an application is received by an overlay, which wraps it and the destination before sending it over the network. The communication is decapsulated and delivered as desired once it reaches its final destination. The sites are located in the encapsulation, which keeps the identity and location apart, but the initial message contains the identities of the communication devices (applications). This encapsulation and decapsulation process must be completed very rapidly and effectively since it is done on a per-packet basis. The architecture of VXLAN overlay is shown in Figure 4.

At present, between 60% and 70% of all application workloads are virtualized, according to market data; yet, more than 80% of servers are not running a hypervisor. It goes without saying that each data center has a different combination of servers running both virtualized and non-virtualized applications. This combination should be covered by all of the data center's network solutions. The network problems that old virtual local area network technology produced in data centers were addressed by the IETF-recommended VXLAN standard. The VXLAN standard provides the improved scalability of Layer 2 segmentation and variable workload allocation required by current application demands.

LITERATURE REVIEW

There have been few studies published regarding the security aspect in a VXLAN environment.




Figure 4: Virtual extensible local area network overlay^[12]

Most studies handle security in networking and focus on the security of virtual machines (VM). Some notable papers in the field of DCN are: You *et al.* network overlay East-west and north-south traffic are managed by DCNs using firewall policies, policy-based routing, micro segmentation policies, and access control lists. Incremental verification techniques have been presented; however, they are inefficient or do not support specific network forwarding features. Test results demonstrate the accuracy and speed of incremental verification techniques. In terms of all-pair reachability, verifier outperforms current methods by up to $\times 10$, for networks with 200 leaves, finishing the change-impact analysis in 15 s (4000 subnets and 16 million pairs).^[13] Salazar-Chacón and Marrone provide an open networking paradigm-based VXLAN proof-of-concept emulation and use two approaches to validate the protocol's stated properties: The uncommon procedures known as EVPN and lightweight network virtualization (LNV) provide scalable, affordable, and high-performing options for large-scale deployments.^[14]

Radoi and Rincu, to move toward amore centralized, integrated system, the costs of implementing and managing each device separately have decreased due to the improvement of dedicated networks as well as switching and routing capabilities. This article demonstrates the necessity of replacing the conventional network designs and protocols, including the virtual port channel and spanning tree protocol, by presenting the new network architecture and associated protocols for both the control plane and the data plane.^[15]

Lee and Lee suggest a scalable, secure internet of things (SSI) network architecture that shields internet of things (IoT) devices from security threats while preserving their processing power. VPN on Layer 2 data connection frame encryption is offered by SSI, which is less demanding than Transmission Control Protocol/IP VPN. Scalable Layer 2 VPNs are offered by L2 tunneling protocol and VXLAN, whereas MACsec encrypts frames. SSI decreases central processing unit utilization by 31.6% and enhances network performance by 30% when compared to IoT networks using OpenVPN.^[16] Baziana and Drainakis suggested that the MAC algorithm supports a variety of traffic kinds while avoiding channel collisions and achieving high effectiveness. To ascertain the scalability

Table 2: Comparative analysis of recent research on data center networks, virtualization, and optimization techniques

Author	Focus on	Key technologies	Contribution	Challenges	Future study
You <i>et al.</i> (2022)	Verification of network policy in overlay DCNs	ACLS, micro segmentation, policy-based routing, firewalls	Fast and accurate incremental verification algorithm for all-pair reachability	Incomplete support for certain forwarding behaviors in earlier methods	Extend support for more complex policies and forwarding behaviors
Salazar-Chacón and Marrone (2022)	Emulation of VXLAN networks using SDN	VXLAN, EVPN, Lightweight Network Virtualization (LNV)	Proof-of-concept for scalable, high-performance and low-cost network virtualization	Limited exploration of LNV in production-scale networks	Expand LNV-based emulation with dynamic and large-scale workloads
Radoii and Rincu (2022)	Centralized routing and switching architecture	Integrated control/ data plane, modern alternatives to spanning tree and VPC	Redesigned network architecture with better integration and reduced cost	Complexity in replacing legacy systems	Migration strategies and hybrid models for gradual deployment
Lee and Lee (2021)	Secure, scalable IoT network infrastructure	Layer 2 VPN (L2TP, VXLAN), MACsec, SSI architecture	Improved security and performance for IoT networks, with low resource overhead	Balancing encryption overhead and resource constraints on devices	Broader deployment in heterogeneous IoT environments
Baziana and Drainakis (2021)	MAC protocol efficiency in DCNs	Custom MAC Protocol, Clustered Servers, Simulation Framework	High throughput, low delay, low drop rate under varied server loads	Performance under extreme traffic diversity	Real-time MAC adaptation for dynamic traffic patterns
Eltraify <i>et al.</i> (2020)	Energy-efficient VM placement in optical DCNs	MILP, WDM-TDM AWGR PON, VM profiling	Optimization of server use and network energy through VM placement	Computational overhead of MILP for large-scale data centers	Heuristic or AI-based models for scalable VM placement

ACLS: Advanced cardiovascular life support, DCNs: data center network, VXLAN: Virtual extensible local area network, EVPN: Ethernet virtual private networks, SDN: Software-defined networking, LNV: Lightweight network virtualization, VPC: Virtual private cloud, L2TP: L2 tunneling protocol, IoT: Internet of things, SSI: Secure internet of things, MAC: Media access control, MILP: Mixed Integer linear programming, WDM-TDM AWGR PON: Wavelength division multiplexing-time division multiplexing arrayed waveguide grating passive optical network, VM: Virtual machine, AI: Artificial intelligence

and performance bounds of the suggested MAC protocol, examine its performance under a variety of server populations per cluster. The simulation study shows that the proposed DCN architecture with the MAC protocol performs well for a range of server populations, with almost 100% bandwidth exploitation, low latency, low dropping rate, and high throughput.^[17]

Eltraify *et al.* proposes a mixed integer linear programming (MILP) model to maximize the deployment of VMs throughout a data center architecture built on the energy-efficient wavelength division multiplexing-time division multiplexing arrayed waveguide grating passive optical network (WDM-TDM AWGR PON). The use of VMs and their requirements affect the optimal number of servers used in the data center when reducing power consumption and facilitating more efficient server utilization are taken into account. For a maximum of 20 VMs with varying compute and networking needs, two power consumption reduction goals were investigated. The findings show that when processing and networking power consumption are reduced, virtual machine allocation in the WDM-TDM AWGR PON may reduce networking power utilization by as much as 70%.^[18]

Despite significant advancements in DCN, virtualization, and energy-efficient architectures, several research gaps remain. Existing solutions often focus on isolated aspects such as security, scalability, or energy efficiency but lack an integrated approach that addresses all these dimensions simultaneously. Incremental verification methods, while faster, still fall short in supporting complex forwarding behaviors in heterogeneous network environments. Similarly, while emulation and virtualization strategies such as VXLAN and LNV show promise, their real-world scalability and adaptability under dynamic workloads are yet to be fully explored. Furthermore, energy optimization models such as MILP are computationally intensive and not practical for large-scale or real-time deployments. These challenges and limitations are summarized in Table 2, highlighting the need for unified, scalable, and adaptive frameworks that can ensure secure, high-performance, and energy-efficient operation in next-generation data Center and IoT network infrastructures.

CONCLUSION

Due to cloud technologies and multi-tenancy, network design has evolved fundamentally.

Multi-data centers across locations are needed for business continuity. Enhancing old technologies and implementing new ones with plenty of additional features has become essential. This paper has explored the limitations of traditional network architectures and the potential of VXLAN-EVPN overlays to provide scalable, resilient, and flexible networking solutions. By decoupling network functions from physical infrastructure and enabling integrated Layer 2/Layer 3 communication, VXLAN-EVPN enhances operational agility and supports modern deployment paradigms such as hybrid and multi-cloud environments. However, despite these advantages, several limitations remain. Real-world implementations still face challenges in managing overlay complexity, ensuring efficient incremental verification, supporting heterogeneous legacy systems, and optimizing energy consumption. In addition, large-scale dynamic workloads can introduce performance bottlenecks if not managed with adaptive and intelligent orchestration. Future research should focus on developing unified frameworks that combine network function virtualization, AI-driven traffic management, and policy-aware automation. Investigating lightweight security mechanisms and scalable verification models tailored for VXLAN-EVPN environments is also essential. Furthermore, energy-aware network designs and cross-layer optimization strategies can contribute to sustainable and cost-effective data center operations. Addressing these areas will be vital in realizing the full potential of next-generation DCN that can handle the needs of the digital future.

REFERENCES

1. Xiao J, Wu B, Jiang X, Pattavina A, Wen H, Zhang L. Scalable data center network architecture with distributed placement of optical switches and racks. *J Opt Commun Netw* 2014;6:270-81.
2. Patel R, Tandon R. Advancements in data center engineering: Optimizing thermal management, HVAC systems, and structural reliability. *Int J Res Anal Rev* 2021;8:991-6.
3. Rojas E, Alvarez-Horcajo J, Martinez-Yelmo I, Arco JM, Briso-Montiano M. Scalable and reliable data center networks by combining source routing and automatic labelling. *Network* 2021;1:11-27.
4. Rojas E, Ibanez G, Gimenez-Guzman JM, Rivera D, Azcorra A. Torii: Multipath distributed Ethernet fabric protocol for data centres with zero-loss path repair. *Trans Emerg Telecommun Technol* 2015;26:179-94.
5. Al-Mashhadi S. A Scalable Data Center Network Design Strategies. Malaysia: Universiti Sains Malaysia; 2019.
6. Qi H, Shiraz M, Liu J, Gani A, Rahman ZA, Altameem TA. Data center network architecture in cloud computing: Review, taxonomy, and open research issues. *J Zhejiang Univ Sci C* 2014;15:776-93.
7. Shetty AJ, Ganashree KC. Comprehensive review of datacenter architecture evolution. *Int Res J Eng Technol* 2020;7:7420-6.
8. Kliazovich D, Bouvry P, Khan SU. GreenCloud: A packet-level simulator of energy-aware cloud computing data centers. *J Supercomput* 2012;62:1263-83.
9. Bilal K, Khan SU, Kolodziej J, Zhang L, Hayat K, Madani SA, *et al.* A comparative study of data center network architectures. In: 26th European Conference on Modelling and Simulation. United States: ECMS; 2012. p. 526-32.
10. Kant K. Data center evolution.: A tutorial on state of the art, issues, and challenges. *Comput Netw* 2009;53:2939-65.
11. Singh T, Jain V, Babu GS. VXLAN and EVPN for data center network transformation. In: 2017 8th International Conference on Computing, Communication and Networking Technologies (ICCCNT). United States: IEEE; 2017. p. 1-6.
12. George AS, George AS. A brief overview of VXLAN EVPN. *Int J Innov Res Electr Electron Instrum Control Eng* 2021;9:1-12.
13. You L, Zhang J, Jin Y, Tang H, Li X. Fast configuration change impact analysis for network overlay data center networks. *IEEE ACM Trans Netw* 2022;30:423-36.
14. Salazar-Chacon G, Marrone L. Open networking for modern data centers infrastructures: VXLAN proof-of-concept emulation using LNV and EVPN under cumulus linux. In: 2022 IEEE Sixth Ecuador Technical Chapters Meeting (ETCM). United States: IEEE; 2022. p. 1-6.
15. Radoi AE, Rincu CI. Integration of data center network technologies VxLAN, BGP, EVPN. In: 2022 14th International Conference on Communications (COMM). United States: IEEE; 2022. p. 1-5.
16. Lee J, Lee H. Secure and scalable IoT: An IoT network platform based on network overlay and MAC security. In: IFIP Advances in Information and Communication Technology. Berlin: Springer; 2021. p. 287-301.
17. Baziana P, Drainakis G. A scalable optical intra-cluster data center network with an efficient WDMA protocol: Performance and power consumption study. In: 2021 30th Wireless and Optical Communications Conference WOCC, 2021. United States: IEEE; 2021.
18. Eltraify AE, Mohamed SH, Elmirghani JM. VM placement over WDM-TDM AWGR PON based data centre architecture. In: 2020 22nd International Conference on Transparent Optical Networks (ICTON). United States: IEEE; 2020. p. 1-5.