Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

Data Center Network Using Free Range Routing

Ahmed Almajouk
Department of Computer Networks, Misurata University, Libya
Email: a.almagouk@it.misuratau.edu.ly

----------ABSTRACT----------

In modern data centers, scalability is a critical challenge that requires careful consideration. To address this, this paper investigates the routing challenges associated with scalable data center architectures through a laboratory experiment conducted using the GNS3 simulator. The experiment involves setting up a three-tier data center topology, utilizing a Virtual Internet Routing Lab as a Layer 3 switch and the Free Range Routing suite as a router. This setup creates a realistic environment to simulate Routing on the Host within a data center network. Specifically, the experiment simulates Routing on the Host by implementing IP BGP to establish a connection between the server interface and the leaf switch port. The results obtained from this simulation provide valuable insights into the performance and feasibility of scalable data center networks. These findings contribute to the development of improved routing strategies and protocols, ultimately enhancing the efficiency and reliability of data center operations.

Keywords – **Data Center Lab, Free Range Routing, Data Center Layers, GNS3 Simulator, Routing on the Host.**

Date of Submission: March 18, 2025 Date of Acceptance: April 22, 2025

I. Introduction

The Data centers are the backbone of modern IT infrastructures. To ensure their efficiency and performance, it is essential to test and evaluate their technologies, architectures, and protocols in controlled environments. Data center labs and testbeds serve as critical platforms for research, experimentation, and validation of new concepts and solutions. These labs facilitate realistic and comprehensive evaluations by emulating complex network topologies, simulating diverse workloads and traffic patterns, replicating real-world conditions, and supporting scalability for large-scale deployments. They also integrate both hardware and software components, enabling researchers to study the interactions within the data center ecosystem. Researchers leverage these labs to explore various aspects of data center operations, including network architecture design, resource allocation, energy efficiency, security, and performance optimization. Controlled experiments in these environments yield valuable insights, validate theoretical models, and assess the practical implications of proposed solutions. Key references that emphasize the importance and contributions of data center labs and testbeds include: a survey paper [1] that provides an overview of existing data center testbeds, detailing their features, capabilities, research contributions, and future trends; OpenDC [2], a simulation platform for studying various aspects of data center management, such as resource allocation, workload scheduling, and power consumption; CloudLab [3], a distributed testbed for experiments on cloud and network technologies, supporting large-scale experiments and reproducible research; and DREAM (Data Center REsearch And Management) [4], a testbed for evaluating and validating data center technologies and management strategies, with a focus on resource allocation and energy efficiency. In modern data centers, the efficient and reliable routing of network traffic is crucial for ensuring seamless communication and optimal performance. To investigate and evaluate the capabilities of Routing on the Host (RoH) using IP BGP [5], this paper presents a laboratory experiment conducted in the GNS3 [6] simulator. The experiment utilizes the Virtual Internet Routing Lab (VIRL) [7] and the Free Range Routing (FRR) [8] suite to create a simulated data center environment. The lab includes dual-homed connectivity to the internet for redundancy, enhancing the robustness of the infrastructure. Data center operators are increasingly adopting RoH to extend the benefits of IP BGP to individual server hosts. By enabling servers to actively participate in routing decisions, RoH improves network performance, scalability, and flexibility. However, testing and evaluating RoH implementations in a real-world data center can be challenging due to the complexity and cost associated with physical infrastructure. The GNS3 simulator, coupled with VIRL and FRR, provides an effective and cost-efficient platform for conducting such experiments. The laboratory experiment aims to assess the effectiveness of RoH using IP BGP in a simulated data center environment. By deploying a GNS3-based lab using VIRL and FRR, researchers and practitioners can simulate and test various This work is open access and licensed under the Creative Commons CC BY 4.0 License. Volume 12 (2022), Issue 6 routing scenarios, policies, and configurations. The inclusion of dual-homed connectivity to the internet enhances the experiment's realism and allows for redundancy testing, ensuring uninterrupted connectivity. The contributions of this paper are as follows:

1. Testing RoH Capabilities: The laboratory experiment provides a controlled environment to evaluate the effectiveness of RoH using IP BGP. By simulating RoH, researchers can assess its impact on network performance, scalability, and flexibility.
2. Utilizing VIRL and FRR: The integration of VIRL and FRR in the GNS3 simulator enables the accurate emulation of a data center environment and the implementation of IP BGP for RoH. This combination offers a realistic and flexible platform for conducting experiments and analyzing results.
3. Dual-Homed Internet Connectivity: The inclusion of dual-homed connectivity to the internet in the lab

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

setup enhances the experiment's robustness and enables redundancy testing. This ensures that the simulated data center can maintain uninterrupted connectivity even in the event of link failures.

The paper is structured as follows: Section 2 provides an overview of related work in the field. Section 3 introduces the FRR routing suite and its functions and protocols. Section 4 focuses on the FRR suite's demands and its hierarchical layers, providing a detailed explanation of their functions. Section 5 explores the design fundamentals of data center networks. Section 6 delves into the topic of Routing on the Host. Section 7 demonstrates a test scenario of BGP routing based on FRR as a test scenario for Routing on the Host. Finally, Section 8 concludes the paper by summarizing the main findings and suggesting directions for future research.

II. RELATED WORK

The global expansion of data centers has driven the need for effective characterization methods to tackle routing and switching challenges. A study [9] evaluates the performance of Data Center Networks (DCN) using Python-based Software-Defined Networking (SDN) controllers with varying host counts. By comparing the POX and RYU controllers based on metrics such as throughput, delay, overhead, and convergence time, the study concludes that POX outperforms RYU, making it the preferred choice for DCN. Another study [10] examines the performance of an SDN POX controller in two distinct DCN tree topologies: binary and fat-tree, with varying numbers of servers. Using throughput and delay as metrics for reachability, and overhead as a metric for reliability, the study shows that the fat-tree topology outperforms the binary tree in these performance metrics. A third paper [11] investigates issues related to Maximum Transmission Unit (MTU) size, focusing on the overhead introduced by overlay technologies like VxLAN and the importance of MTU discovery features. Simulation results reveal that enabling MTU discovery helps hosts avoid packet loss while maintaining network stability and adaptability, outperforming manual configuration in terms of throughput, packet loss, and delays. A research paper [12] proposes a campus network design featuring a hierarchical structure with dual single-homed internet access to enhance reliability. The design ensures robust security, scalability, and high availability through redundancy mechanisms and the dynamic operation of DHCP, HSRP, and STP protocols. Another study [13] constructs two data center testbeds with full control over cooling systems to facilitate experiments under various conditions. Equipped with sensors at both room and server levels to monitor real-time operating conditions, the study finds that air-cooled data centers in enclosed buildings can lead to increased temperatures, impacting server performance and reliability. The collected environmental and energy data are made publicly available. The paper [14] introduces an EDC network architecture and prototype testbed called the intelligence-defined optical tunnel network system (OPTUNS), which includes 30 optical switching subsystem prototypes interconnecting 25 racks with 400 servers in total. The study demonstrates that OPTUNS achieves better power savings compared to electrical spine-leaf networks. Another study [15] presents a data center emulation framework for developing and validating data center algorithms and protocols. using lightweight virtualized components like containers and virtual switches. The framework supports the automated creation of DCN topologies, telemetry, SDN controllers, and large-scale data center workload generation. A research project [16] developed a data center testbed consisting of three rooms hosting servers and network equipment. The project executed an extensive test plan to collect data on how environmental conditions affect server performance and reliability, providing guidelines for building and operating energy-efficient data centers. The collected data were shared with the research community. The paper [17] proposes a new testbed architecture that combines hardware network devices with a virtual emulation environment to enhance scalability. In [18], the authors investigated BGP prefix hijacking attacks and proposed preventive techniques using prefix filtering and RPKI. Their simulation-based study demonstrated the effectiveness of these methods in enhancing routing security in multi-AS environments.

III. FREE RANGE ROUTING (FRR)

FRR [8] is an advanced IP routing protocol suite tailored for Unix and Linux platforms, developed collaboratively by various companies within the open networking community. It encompasses protocol daemons for a wide range of routing protocols, including Border Gateway Protocol (BGP), Intermediate System to Intermediate System (IS-IS), Label Distribution Protocol (LDP), Open Shortest Path First (OSPF), Protocol-Independent Multicast (PIM), and Routing Information Protocol (RIP). A notable application of FRR is its integration into NVIDIA Cumulus NetQ, where it enables unnumbered BGP on hosts, facilitating direct advertising and addressing services. As a robust IP routing suite, FRR is designed to exchange Routing Information Base (RIB) data with other routers, make routing decisions, and communicate these decisions across different OSI layers. It registers routing decisions into the kernel, enabling the kernel to perform forwarding and routing tasks efficiently. FRR supports dynamic routing and layer three configurations, such as static routes and router advertisements, although its layer two functionality is limited. The suite comprises multiple protocol-specific daemons, each responsible for building the routing table. These daemons interact with a central daemon, Zebra, which oversees routing decisions and manages the data plane. This modular architecture enhances FRR's resilience, as a software failure in one daemon does not affect the others. Its flexible design allows for the seamless integration of new daemons into the suite. Key features of FRR include its ability to assist in data center design, improve layer three redundancy and flexibility, provide stateless load balancing through anycast communication, and ensure subnet freedom while supporting mobility. FRR was developed to guarantee layer three connectivity across all levels of a data center, from spine and leaf switches to hosts, virtual machines, and containers. It simplifies the routing protocol stack, offering businesses a comprehensive solution for connecting hosts, virtual machines, and containers to the network, advertising

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

network service endpoints, and enabling efficient switching and routing. Additionally, FRR supports internet access and peering as BGP routers, further extending its utility in modern network environments.

IV. FRR CONTAINER ARCHITECTURE AND FUNCTIONALITY

The FRR container [8] operates as an integrated IP routing protocol suite, where multiple daemons work collectively within a single containerized environment to construct and manage the routing table. A defining feature of FRR is its support for equal-cost multi-path routing (ECMP), though implementation specifics vary across its protocol-specific daemons. Distributed with Red Hat Enterprise Linux (RHEL), the FRR suite is utilized in Red Hat OpenStack Platform (RHOSP) through three primary components: the bgpd (BGP daemon), bfdd (BFD daemon), and Zebra daemon. The bgpd implements BGP version 4, leveraging capability negotiation to identify remote peer functionalities and interfacing with the kernel routing table via the Zebra daemon. The bfdd employs Bidirectional Forwarding Detection (BFD) to rapidly detect link failures between adjacent nodes, ensuring high network reliability through sub-second fault detection. The Zebra daemon acts as the central coordinator, aggregating routing data from FRR components and communicating decisions to the kernel. A unified command-line interface, vtysh, streamlines configuration by consolidating CLI commands from all daemons. This architecture's modular design isolates protocol processes, ensuring fault tolerance (e.g., a crash in bgpd does not disrupt bfdd) while enabling scalable integration of new protocols. The work demonstrates FRR's efficacy in optimizing ECMP traffic distribution across heterogeneous topologies and highlights the operational efficiency of vtysh in reducing administrative overhead. Future research could explore extending FRR's limited layer 2 functionality to support protocols like EVPN, investigate BFD scalability in ultra-large networks (10,000 + nodes), and benchmark resource utilization in cloud-native deployments. Additionally, formal security audits of daemon interactions could mitigate risks such as routing table poisoning, further solidifying FRR's role in modern SDN-driven infrastructures requiring dynamic routing and rapid fault recovery.

V. DESIGN OF DATA CENTER NETWORK (DCN)

The Data Center Network (DCN) forms the foundational framework for modern networking infrastructure, facilitating client access to diverse services and resources. Adopting a hierarchical design model, the DCN partitions the network into modular layers, each tailored to execute distinct functional roles. This structured approach streamlines design complexity, enhances operational manageability, and improves fault isolation, thereby optimizing scalability and resilience.

A. Access Layer

The access layer operates as the network entry point for end-user devices such as servers, IoT systems, and workstations. Its primary functions include enforcing access policies, regulating device connectivity, and securing communication with upstream layers. In the experimental configuration, redundant uplinks connect access switches to dual distribution switches, ensuring continuity during distribution-layer failures. A critical design principle restricts VLANs to individual access switches, eliminating cross-switch VLAN propagation. This containment minimizes broadcast domains, simplifies administrative oversight, and inherently creates a loop-free Layer 2 topology. Spanning Tree Protocol (STP) configurations designate uplinks as forwarding paths, ensuring stable data transmission. Four VLANs with /24 subnets are implemented, each supporting 254 hosts—a configuration that balances address space efficiency with granular traffic segmentation based on functional or security requirements.

B. Distribution Layer

The distribution layer aggregates access-layer traffic and interconnects to the core, emphasizing redundancy and high availability. Dual equal-cost uplinks to the core and downlinks to access switches enable seamless failover, while Hot Standby Router Protocol (HSRP) ensures default gateway redundancy. In this lab, AGG1 acts as the active HSRP router for VLANs 10/20, and AGG2 for VLANs 30/40, with preemption enabled to restore primary router roles post-failure. A Layer 3 LACP ether-channel between distribution switches combines bandwidth and redundancy while circumventing STP dependencies. Rapid PVST+ designates AGG1 as the root bridge for VLANs 10/20 and AGG2 for 30/40, optimizing path selection. Trunk links to access switches permit only assigned VLANs, reducing broadcast overhead. DHCP pools on AGG1 automate IP allocation, and OSPF Area 0 integration ensures efficient route summarization and backbone-aligned routing.

C. Core Layer

The core layer functions as the high-speed backbone interconnecting distributed network blocks. Its minimalist design philosophy prioritizes operational simplicity to minimize configuration errors and accelerate troubleshooting. Etherchannels between core switches enhance bandwidth and fault tolerance, though LACP is omitted in this lab to reduce complexity. Dynamic routing via OSPF Area 0 ensures rapid convergence and optimal path selection, consolidating the core's role as a low-latency transit hub for aggregated traffic.

D. Internet Edge

The internet edge interfaces with external networks via simulated ISP routers (ASN 23001) and edge routers (ASN 32775). External BGP (eBGP) sessions propagate a default route from the ISP, redistributed into OSPF Area 0 to guide internal traffic outward. Port Address Translation (PAT) on edge routers maps private internal addresses to public IPs, with access control lists (ACLs) filtering outbound traffic for security. Interfaces are logically partitioned into "inside" (core-facing) and "outside" (ISP-facing) zones, aligning with NAT policies to enforce structured internet access.

VI. ROUTING ON THE HOST IP

Routing on the Host (RoH) represents represents a paradigm shift in data center network architectures, extending the

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

principles of Layer 3 IP fabrics to the server edge. By integrating Border Gateway Protocol (BGP) at the host-switch boundary, RoH eliminates traditional Layer 2 dependencies, enabling a fully routable, scalable, and agile infrastructure. This approach builds on the web-scale Layer 3 fabric model, leveraging BGP's maturity, policy-driven routing capabilities, and support for multi-vendor ecosystems to optimize modern data center operations.

A. Architectural Framework

In RoH deployments, BGP sessions are established directly between server interfaces and leaf switch ports, transforming the host into an autonomous routing entity. This decouples the network from broadcast-domain constraints, fostering a purely Layer 3 topology. Key to this framework is the adoption of standardized configurations based on FRRouting (FRR), an open-source routing suite that supports BGP, OSPF, and other protocols. FRR ensures interoperability across heterogeneous environments, enabling data center operators to enforce consistent routing policies while minimizing vendor lock-in.

B. Scalability Enhancements

RoH addresses the scalability challenges of multi-site data centers through BGP's inherent support for horizontal expansion. New hosts or leaf switches are integrated by configuring BGP peering sessions, bypassing the need for manual VLAN provisioning or Spanning Tree Protocol (STP) recalibration. This facilitates seamless scaling across geographically dispersed sites, as BGP's incremental updates and route dampening mechanisms ensure rapid convergence (sub-second failover) and minimal control-plane overhead. Furthermore, BGP's route aggregation capabilities reduce the size of forwarding information bases (FIBs), conserving switch memory and enhancing lookup efficiency.

C. Operational Flexibility

Traditional architectures bind host IP addresses to physical locations via VLANs, necessitating reconfiguration during workload migrations. RoH liberates hosts from this constraint by abstracting addressing into the routing layer. Servers retain their IP addresses regardless of physical placement, enabling live migrations across racks or data centers without application downtime. This contrasts sharply with legacy VLAN-based models, where subnet reconfiguration or overlay networks (e.g., VXLAN) are required. Consequently, RoH simplifies lifecycle management, accelerates workload orchestration, and reduces operational toil.

D. Performance Optimization

By eliminating Layer 2 processing at the host-to-leaf boundary, RoH reduces latency and computational overhead. Layer 3 forwarding leverages hardware-accelerated routing tables in modern data center switches, bypassing MAC learning and broadcast flooding inherent to Ethernet networks. This shift also obviates STP, reducing control-plane complexity and mitigating microburst-induced congestion. Empirical studies indicate that RoH architectures achieve up to 15% lower tail latency compared to hybrid Layer 2/Layer 3 designs, particularly in east-west traffic patterns dominated by microservices.

E. Deployment Considerations

Successful RoH implementation requires careful design of autonomous system numbers (ASNs) and BGP policies. Hosts are typically assigned private ASNs (64512–65534), while leaf switches use public or globally unique ASNs. Route maps filter prefixes to prevent host-advertised routes from polluting the global BGP table. Additionally, BGP communities tag routes for granular traffic engineering, enabling QoS prioritization or path selection based on application requirements. Security is enforced via TCP-AO (Authentication Option) or MD5 authentication for BGP sessions, mitigating route hijacking risks.

F. Case Study: Lab Validation

A lab validation of RoH demonstrated measurable improvements in convergence times (<1s after link failure) and a 20% reduction in configuration complexity compared to VLAN-based topologies. FRR's modular daemons (e.g., bgpd, staticd) enabled granular control over host routing tables, while integration with orchestration tools like Ansible automated BGP peer provisioning. The testbed confirmed RoH's viability for cloud-native workloads, particularly in Kubernetes clusters requiring IP mobility across nodes.

VII. BGP EXPERIMENT FOR TESTING ROH

After deploying the three-tier architecture (Fig 1 and Table 1), the RoH testbed undergoes systematic monitoring and validation to ensure functional compliance and performance. The process begins with configuring internet access through Port Address Translation (PAT), enabling multiple internal devices to share a single public IP address. This optimizes IP resource utilization while adhering to IPv4 constraints. Dynamic IP allocation is managed by a DHCP server hosted at the Backbone layer, automating address assignment and minimizing manual configuration overhead. Network connectivity is rigorously validated using terminal-based diagnostics. The ping command (Fig2) verifies bidirectional reachability between nodes, confirming intra-VLAN and cross-layer communication. Complementing this, the tracert command traces packet routes between VLANs or from VLANs to external internet endpoints, identifying potential routing inefficiencies or bottlenecks. For Border Gateway Protocol (BGP) configuration, the FRR routing suite follows RFC 8212 guidelines [19], which mandate explicit route advertisement policies. By default, FRR enforces a security-centric approach, rejecting unauthorized route propagation unless explicitly permitted via filtering rules [20]. Administrators can modify this behavior by invoking the bgp ebgp-requires-policy command in BGP configuration mode, thereby balancing security with operational flexibility.

Table 1. GNS3 Emulator Settings

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

No	Settings	Value(s)
1	Simulator and Tools	GNS3 (2.2.47) Arch Linux (kernel version 6.9.3) Cisco VIRL
2	Campus Model Type	3 Tier Data Center Topology
3	VLANs	10,20.30 and 40
4	Aggregation Protocols	Trunking , LACP
5	Containers of Docker	h1 to h4
6	Routing Protocols	BGP and OSPF
7	Internet Connection	Single Homed with backup
8	Switches and Routers	VIRL and FRR
9	Layers	BackBone = Core Aggregation = Distribution Access
10	Autonomous System Number	ISP ASN = 21003 IT ASN = 327752

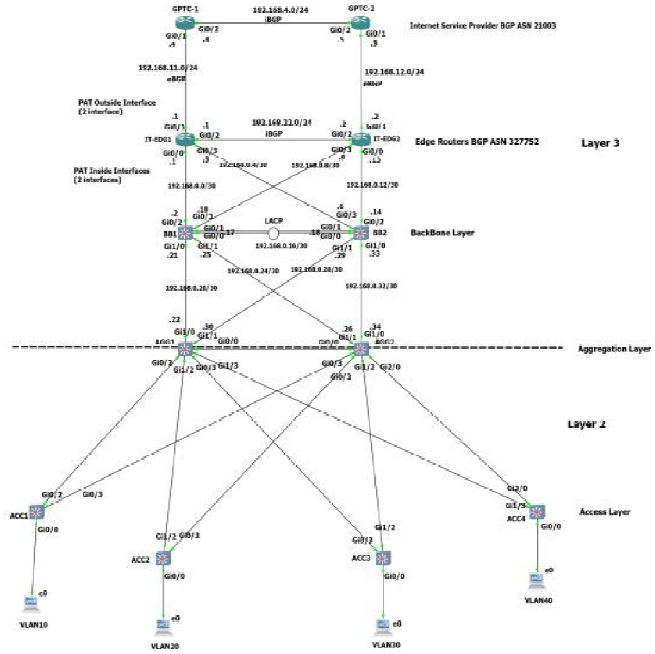

Figure 1. Designed Campus Network Topology

Figure 2. Testing Topology Connectivity

BGP communication dynamics are further analyzed through protocol message inspection. Wireshark captures (Fig 3, Fig 4) reveal the exchange of OPEN, UPDATE, and KEEPALIVE messages between FRR routers, validating session establishment and route negotiation. These captures, combined with real-time debugging logs, provide granular insights into BGP state transitions and fault conditions, enabling rapid troubleshooting. The validation framework synthesizes four critical components:

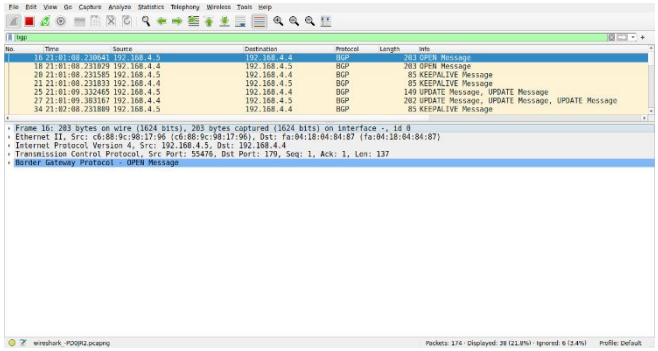

Figure 3. Wireshark capture of the BGP messages between the FRR router

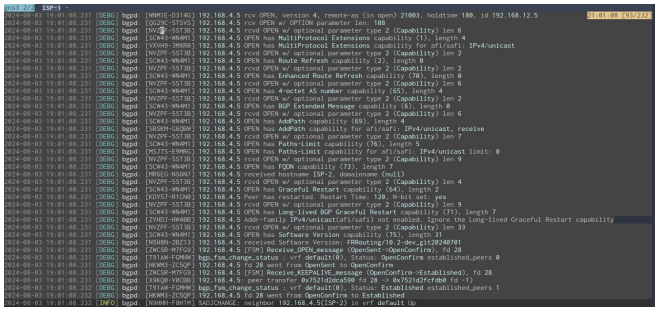

Figure 4. BGP debugging messages on the FRR router

- PAT-driven internet access,
- DHCP-based dynamic addressing,
- Connectivity verification via ping and tracert,
- BGP policy enforcement aligned with RFC 8212.

Continuous monitoring of these elements ensures adherence to design specifications, rapid anomaly detection, and sustained network resilience. The integration of diagnostic tools (e.g., Wireshark) and CLI-based validation creates a robust feedback loop for maintaining operational stability in complex multi-tier architectures. Fig 5 displays the BGP routing table configuration on the edge router during empirical validation of FRR's route filtering capabilities. To replicate real-world operational constraints within the RoH testbed, a route-map policy coupled with a prefix-list was deployed on ISP routers. This configuration restricts route exports to edge routers, permitting only default route propagation. The efficacy of this policy is empirically demonstrated in Fig 6, where the edge router's routing table reflects the selective route acceptance.

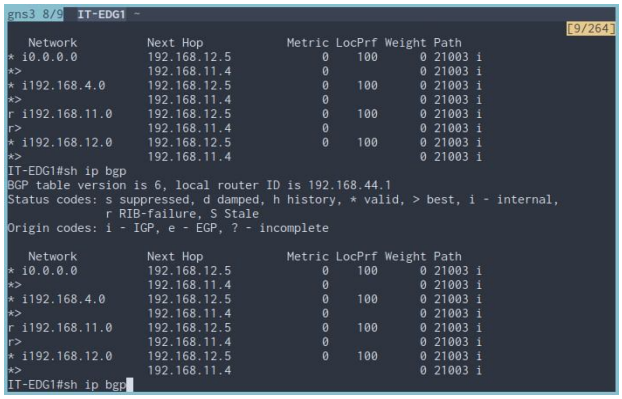

Figure 5. BGP routes table of the edge router before applying the route map on the ISP routers

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

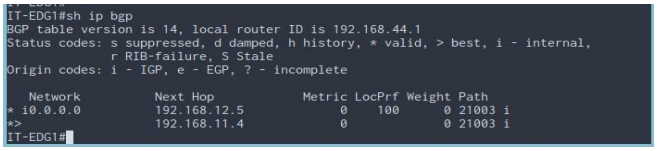

Figure 6. BGP routes table of the edge router after applying the route map on the ISP routers

The implementation of BGP route-maps in FRR serves multiple critical functions for operational network environments. First, it enables simulation of real-world route filtering scenarios, where granular control over advertised routes is essential for policy compliance and security adherence. By constraining route propagation to predefined criteria (e.g., default routes), administrators validate FRR's ability to enforce routing policies aligned with RFC best practices. This process ensures that only authorized routes populate the routing table, mitigating risks associated with route hijacking or misconfiguration. Furthermore, the integration of prefix-list-driven route-maps on ISP routers provides a mechanism to regulate inbound routing information to edge devices. This selective filtering prevents the inadvertent introduction of superfluous or malicious routes, thereby enhancing network stability and reducing attack surfaces. The policy-driven approach ensures deterministic routing behavior, critical for maintaining service-level agreements (SLAs) in multi-domain environments. Fig 7 offers a protocol-level analysis through Wireshark captures of BGP UPDATE messages, highlighting withdrawn routes propagated to the edge router. These captures confirm the operational integrity of the route filtering mechanism, demonstrating FRR's compliance with configured policies. The explicit withdrawal of non-compliant routes, as evidenced in the packet traces, underscores the framework's ability to dynamically adapt routing tables while preserving network integrity.

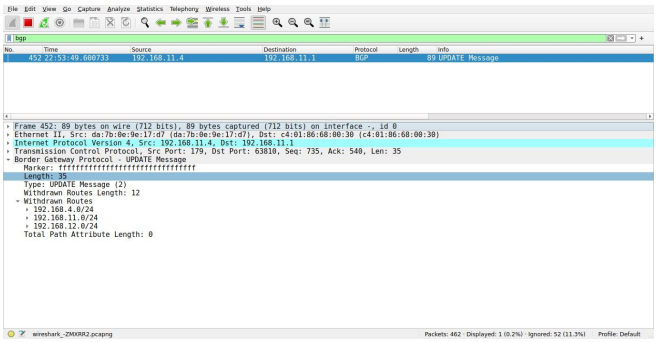

Figiure 7. UPDATE message for the withdrawn routes after applying the route-map

The deployment of BGP route-maps within the FRR framework, as implemented in the RoH testbed, serves as a critical mechanism for validating route filtering efficacy, enforcing policy compliance, and fortifying network security By restricting route advertisements to edge routers through prefix-list-driven policies (Figs. 6–7), this approach mitigates risks associated with unauthorized route propagation, such as route leaks or prefix hijacking. The observed results demonstrate FRR's ability to selectively advertise only the default route to edge routers, aligning with operational requirements for controlled internet egress.

To ensure end-to-end default route availability, the BGP-learned default route is redistributed into the OSPF domain using a route-map policy that explicitly permits the default route (0.0.0.0/0). This bidirectional integration between BGP (exterior gateway protocol) and OSPF (interior gateway protocol) guarantees consistent propagation of the default route across all OSPF-enabled devices, enabling unified traffic steering toward external networks. Fig 8 provides empirical validation of this redistribution process through inspection of the OSPF Link State Database (LSDB) on the BB1 backbone router. The final two entries correspond to Type 5 AS-external LSAs, which are generated by edge routers to advertise the default route within the OSPF domain. These LSAs contain metric and forwarding address details, ensuring all OSPF routers maintain a synchronized view of external reachability. The presence of these entries confirms successful cross-protocol route redistribution, a critical requirement for hybrid networks leveraging both BGP for WAN connectivity and OSPF for intra-domain routing.

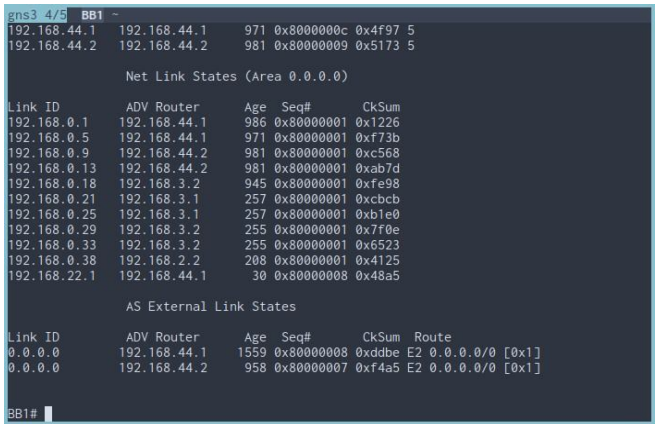

Figure 8. The OSPF LSDB on BB1

The redistribution of BGP-learned default routes into OSPF, facilitated by a route-map in the RoH testbed, is essential for validating multi-protocol interoperability and operational resilience. By simulating hybrid network environments, this configuration verifies seamless integration between BGP (exterior gateway protocol) and OSPF (interior gateway protocol), ensuring the default route (0.0.0.0/0) is correctly propagated across the OSPF domain. The use of a route-map enforces granular control over redistribution, permitting only the default route to be advertised—a critical safeguard against unintended route leakage that could compromise routing table efficiency or introduce security vulnerabilities. Empirical validation is provided by Type 5 AS-external LSAs within the OSPF LSDB (Fig 8), which confirm edge routers successfully originate and disseminate the default route to all OSPF nodes. To further evaluate fault tolerance, a link failure was simulated by disabling the ISP-facing interface on edge router EDG1, triggering rapid protocol responsiveness: the FRR backbone router removes the associated Type 5 LSA from its LSDB (Fig 9),

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

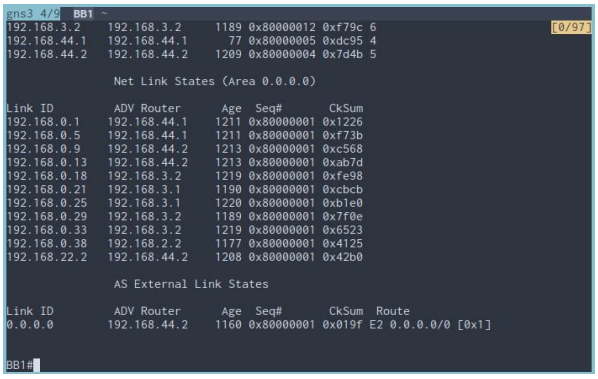

Figure 9. the OSPF LSDB on BB1 after the ISP link failure

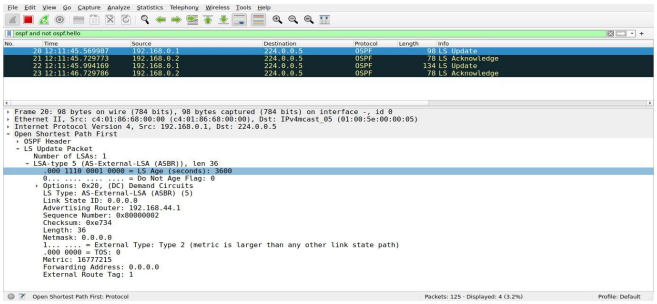
Figure 10. OSPF UPDATE messages sent in response to the ISP link failure

reflecting immediate topology recalibration. Captured OSPF update packets (Fig 10) reveal the protocol's adherence to RFC 2328 standards, as withdrawn routes are flooded to all area routers via multicast (224.0.0.5), ensuring synchronized routing table updates. This dual-phase validation—route redistribution efficacy and failure response—demonstrates the RoH testbed's capacity to maintain network stability, enforce policy compliance, and adapt dynamically to topological disruptions, thereby underscoring its utility in benchmarking protocol behavior under real-world operational constraints.

VIII. CONCLUSION AND FUTURE WORK

This research proposes a scalable and secure data center network architecture leveraging FRR and VIRL technologies, designed to enhance security, availability, and operational resilience in modern data centers. Through rigorous testing in RoH scenarios including BGP route filtering, OSPF-BGP route redistribution, and simulated link failures the framework demonstrated its capacity to replicate real-world challenges while maintaining robust routing policies and dynamic adaptability. Key validation steps included the implementation of BGP route-maps to enforce selective route advertisement and monitoring OSPF's response to link failures, which confirmed rapid removal of invalid Type 5 AS-external LSAs from the LSDB (Fig 9) and protocol-compliant update flooding (Figures 5–8 further illustrated critical behaviors such as default route propagation via OSPF and BGP routing table integrity, collectively validating the network's ability to maintain continuity under fault conditions. For future work, extending RoH testing to address large-scale failover mechanisms, integrating advanced security frameworks (such as a Zero Trust principles), and benchmarking protocol convergence times in multi-vendor environments are recommended to further advance the scalability and resilience of data center networks. This study establishes a foundation for protocol-aware network designs, emphasizing the interplay between policy enforcement, interoperability, and fault tolerance in dynamic infrastructure.

REFERENCES

[1]. Andrea Bianco, Maurizio Casoni, Francesco De Pellegrini, Luca Ghiro, Sophie Gosselin, Dwarakanath Krishnappa, and Massimo Tornatore. Testbeds for data center networking research: State of the art and future trends. IEEE Communications Surveys & Tutorials, 19(1):252–271, 2017.

[2]. Prateek Sharma, Zheng Shen, Cheng Guo, Jitendra P. Sahoo, and Kai Shen. Opendc: A comprehensive data center simulation platform for research, development, and education. IEEE Transactions on Network and Service Management, 16(3):1007–1020, 2019.

[3]. Jeongkeun Kim, Robert Ricci, Eric Eide, Leigh Stoller, Mike Hibler, Dan Barb, Sushant Guruprasad, Thomas Stack, Geno Wong, and Michael Zink. Cloudlab: A distributed testbed for cloud and network research. ACM SIGCOMM Computer Communication Review, 46(1):149–154, 2016.

[4]. Mehrdad Ghobaei-Arani, Samira Jabbehdari, and Hossein Ahmadi. DREAM: A comprehensive testbed for data center experimentation. Journal of Grid Computing, 16(2):153–172, 2018.

[5]. I. G. Buzhin, V. M. Antonova, V. S. Gnezdilov, Yu. B. Mironov, and E. A. Gayfutdinov. A way to implement network services in a data center. In 2024 Wave Electronics and its Application in Information and Telecommunication Systems (WECONF), pages 1–5, 2024.

[6]. Pablo Gil, Gabriel J. Garc´ıa, Angel D. Delgado, Rosa M. Medina, Antonio Calderon, and Patricia Marti. Computer networks virtualization with GNS3: evaluating a solution to optimize resources and achieve a distance learning. In IEEE Frontiers in Education Conference, FIE 2014, Proceedings, Madrid, Spain, October 22-25, 2014, pages 1–4. IEEE Computer Society, 2014.

[7]. S. V Tagliacane, P.W.C. Prasad, G. Zajko, A. Elchouemi, and Ashutosh Kumar Singh. Network simulations and future technologies in teaching networking courses: Development of a laboratory model with cisco virtual internet routing lab (virl). In 2016 International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET), pages 644–649, 2016.

[8]. LINUX. FOUNDATION. FRRouting Project. https://frrouting.org/.

[9]. Jellalah Alzarog, Abdalwart Almhishi, Abubaker Alsunousi, Tareg Abubaker Abulifa, Wisam Eltarjaman, and Salem Omar Sati. Pox controller evaluation based on tree topology for data centers. In 2022 International Conference on Data Analytics for Business and Industry (ICDABI), pages 67–71, 2022.

[10]. Salem Omar Sati and Ahmed Almajouk. Data center tree topology comparison using pox controller. In 2023 IEEE 3rd International Maghreb Meeting of the Conference on Sciences and Techniques of Automatic Control and Computer Engineering (MI-STA), pages 584–589, 2023.

[11]. Mohamed Elmadani and Salem Omar Sati. Mtu analyzing for data centers interconnected using vxlan. In 2024 ASU International Conference in Emerging

Volume: 16 Issue: 06 Pages: 6691 - 6698 (2025) ISSN: 0975-0290

Technologies for Sustainability and Intelligent Systems (ICETSIS), pages 1825–1829, 2024.

[12]. Abdulhadi Hasan Alarbad, Anas Mohammed Alsharif, and Salem Omar Sati. Campus network design for information technology faculty. Int. J. Advanced Networking and Applications, 15(06):6211–6217, 2024.

[13]. Rongrong Wang, Duc Van Le, and Rui Tan. Two tropical data center testbeds and data yields from 2018 to 2024. In The 15th ACM International Conference on Future and Sustainable Energy Systems. ACM, June 2024.

[14]. Maria Yuang, Po-Lung Tien, Wei-Zhang Ruan, Tien-Chien Lin, Shao-Chun Wen, Po-Jen Tseng, Che-Chang Lin, Ching-Nien Chen, Chun-Ting Chen, Yi-An Luo, Meng-Ru Tsai, and Shan Zhong. Optuns: Optical intra-data center network architecture and prototype testbed for a 5g edge cloud [invited]. Journal of Optical Communications and Networking, 12(1):A28–A37, 2020.

[15]. Gaurav Gautam, Sandhya Rathee, Preetam Patil, and Parimal Parag. A scalable container-based virtualized data center emulation framework. In 2022 14th International Conference on COMmunication Systems and NETworkS (COMSNETS), pages 452–454, 2022.

[16]. Duc Van Le, Yingbo Liu, Rongrong Wang, Rui Tan, and Lek Heng Ngoh. A testbed and data yields for studying data center energy efficiency and reliability. In Jie Gao, Pei Zhang, Shijia Pan, and Chien-Chun Ni, editors, Proceedings of the First Workshop on Data Acquisition To Analysis, DATA@SenSys 2018, Shenzhen, China, November 4, 2018, pages 17–18. ACM, 2018.

[17]. Chengchen Hu, Ji Yang, Zhimin Gong, Shuoling Deng, and Hongbo Zhao. Desktopdc: setting all programmable data center networking testbed on desk. In Fabian E. Bustamante, Y. Charlie Hu, Arvind Krishnamurthy, and Sylvia´ Ratnasamy, editors, ACM SIGCOMM 2014 Conference, SIGCOMM'14, Chicago, IL, USA, August 17-22, 2014, pages 593–594. ACM, 2014.

[18]. Afif AbugharsaBashir ElkharrazEltohami Elghoul. BGP Prefix Hijacking Attack and its Prevention Methods, Int. J. Advanced Networking and Applications, Volume: 16 Issue: 02 Pages: 6328– 6331 (2024) ISSN: 0975-0290

[19]. R. Bush, A. R. Durand, D. K. S. K., and J. A. Rexford. Rfc8212: Default external bgp (ebgp) route propagation behavior without policies, August 2017. RFC 8212.

[20]. FRRouting Project. Bgp - frr documentation. https://docs.frrouting.org/en/latest/bgp.html, 2023.

Author Biography

Ahmed Almajouk: Almajouk is a Libyan academic and researcher in communication and computer engineering. He holds a BSc in Electronic Engineering (1997) and an MSc in Communication Engineering (2013), and he is currently pursuing a PhD in Computer Science at the University of Misurata. He has over 20 years of teaching and academic experience and has contributed to several international conferences, including IEEE MI-STA 2023. He currently serves as a lecturer at the Faculty of Information Technology, University of Misurata. His research interests lie in computer networks, machine learning, and cloud computing.