COMPARATIVE STUDY OF OSPF AND BGP FOR HYBRID CLOUD CONNECTIVITY

Ankita Sharma
TERI University
New Delhi, India
ankita.sharma.teri.93@gmail.com

Publication History
Manuscript Reference No: IRJCS/RS/Vol.09/Issue07/JYCS10088
http://www.irjcs.com/volumes/Vol9/iss-07/08.JYCS10088.pdf
Received: 11, July 2022 | Accepted: 22, July 2022 | Published: 31, July 2022
BibTex: Ankita2022Comparative
Citation: Ankita[2022]. Comparative Study of OSPF and BGP For Hybrid Cloud Connectivity, IRJCS:: International Research Journal of Computer Science, Volume IX, Issue VII, 2022, Pgs 170-174
doi: https://doi.org/10.26562/irjcs.2022.v0907.008
Copyright: ©2022 This is an open access article distributed under the terms of the Creative Commons Attribution License; which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: The current paper will perform a detailed analysis of Open Shortest Path First and Border Gateway Protocol in the context of a hybrid cloud, based on a set of relevant key performance indicators such as latency, load balancing, convergence speed, and resiliency in network designs which bridge an on-premise infrastructure with the cloud. It contrasts, through simulations and case studies, the suitability of OSPF versus BGP in hybrid cloud architectures, investigating means to optimize connectivity, efficiency, and fault tolerance. Contributions include added perspective into unique benefits and limitations when OSPF and BGP are used in multidomain, hybrid cloud architectures as part of ongoing discussions on cloud networking.
Keywords: OSPF, BGP, hybrid cloud, routing protocols, latency, load balancing, resilience, multi-cloud, convergence time

I. INTRODUCTION

A. Background of Hybrid Cloud Networking

Hybrid cloud computing integrates private and public cloud platforms with traditional on-premises data centers, enabling organizations to allocate resources based on specific needs for performance, compliance, and cost-effectiveness. Gartner [1] estimates that by 2025, more than 75% of enterprise data will be processed outside the traditional data center or cloud, highlighting the need for efficient and reliable network routing protocols. Routing protocols such as OSPF and BGP are foundational for establishing seamless connectivity across hybrid cloud infrastructures. Their roles are essential in ensuring data can traverse complex and geographically distributed networks, balancing the demands for high availability, scalability, and security.

B. Problem Statement

The specific requirements of hybrid cloud architectures demand a nuanced understanding of routing protocols, particularly in areas such as latency, load distribution, and fault tolerance. OSPF and BGP offer contrasting strengths, and selecting the appropriate protocol (or combination thereof) is crucial for optimizing network performance.

C. Objectives and Scope

The primary goal of this study is to evaluate OSPF and BGP within hybrid cloud contexts, focusing on:

1. Latency and Performance: Examining the efficiency of each protocol in terms of latency.
2. Load Balancing: Assessing load distribution capabilities and their impact on network stability.
3. Resilience and Fault Tolerance: Analyzing protocol behaviour under simulated failure conditions to determine suitability for mission-critical applications.

Table 1: Comparison of OSPF and BGP Characteristics

Characteristic	OSPF	BGP
Convergence Speed	Fast	Slower
Load Balancing	Supports ECMP	Limited (Multipath option)
Domain Suitability	Intra-domain	Inter-domain
Security Vulnerability	Insider threats	Route hijacking, prefix leaks

II. OVERVIEW OF ROUTING PROTOCOLS

A. Routing in Modern Hybrid Cloud Infrastructures

Routing protocols must accommodate various layers of hybrid cloud architectures, each layer presenting unique challenges. In hybrid clouds, networks face dynamically changing workloads, requiring responsive and adaptable routing to maintain efficient data flow.

B. Open Shortest Path First (OSPF)

OSPF is a link-state protocol that operates within an autonomous system (AS), known for its fast convergence, reliability, and ability to perform load balancing within an AS.

1. Algorithm and Path Calculation: OSPF uses the Dijkstra algorithm to compute the shortest path, creating efficient paths and supporting low-latency communication within AS boundaries.
2. Advantages in Private Cloud Scenarios: Within private clouds and on-premises data centers, OSPF's link-state advertisements (LSAs) rapidly propagate topology changes, ensuring quick convergence times.
3. Limitations: OSPF's intra-domain nature limits its applicability in multi-AS environments, making it less suitable for routing across public cloud providers.

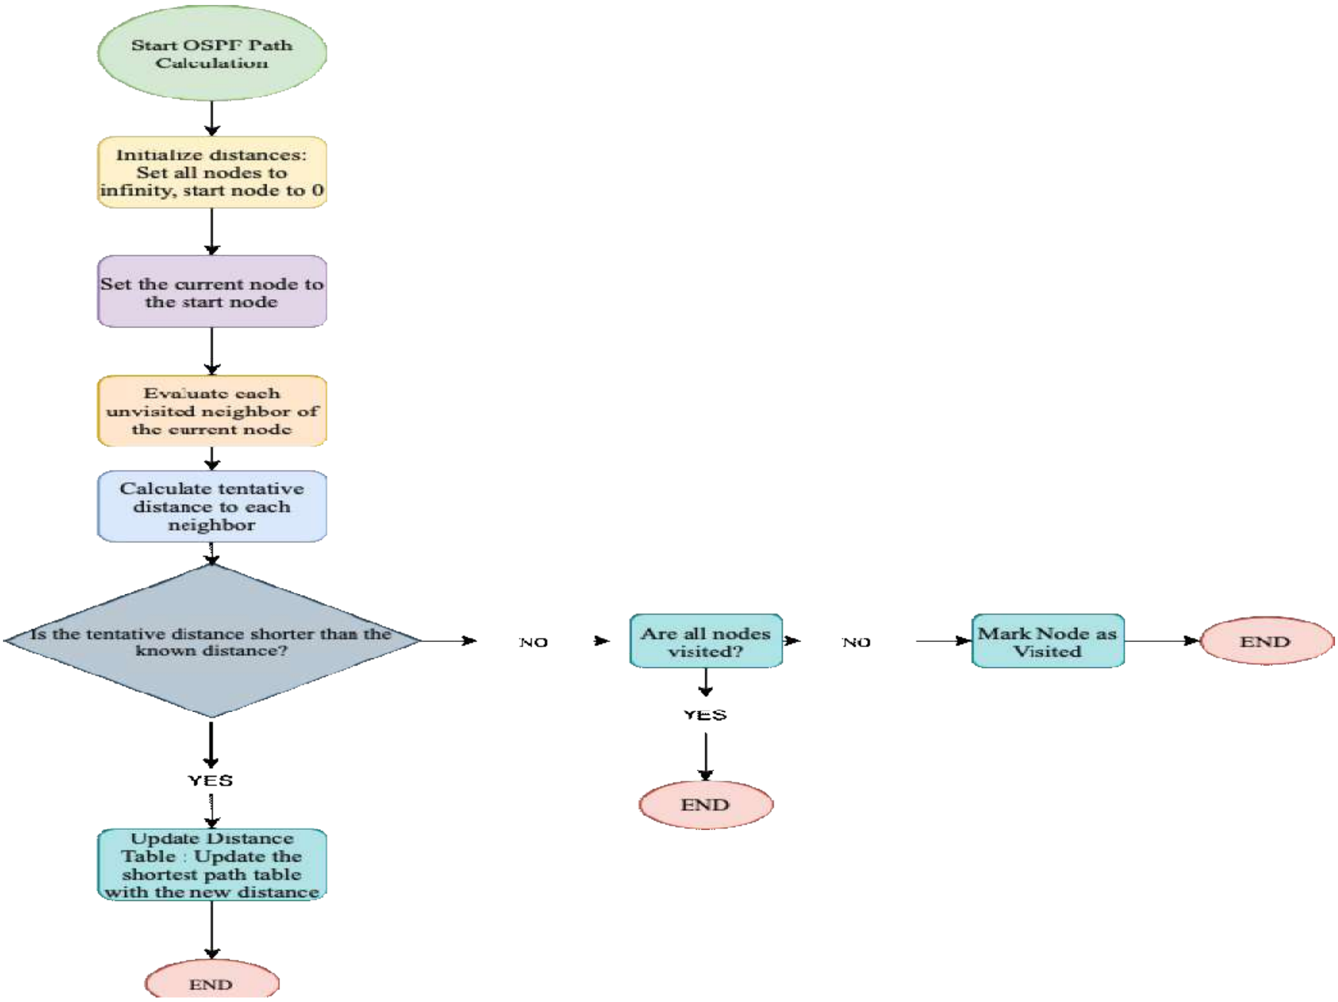

Fig 1. OSPF Path Calculation using Dijkstra's Algorithm

C. Border Gateway Protocol (BGP)

As an exterior gateway protocol (EGP), BGP is designed for routing between autonomous systems, making it the default protocol for cross-domain hybrid cloud setups. BGP's policy-based routing allows for controlled path selection, supporting scalable and flexible inter-domain communication.

1. Path Vector Mechanism: BGP uses path vectors to ensure loop-free paths across ASes, a critical feature for stable multi-cloud architectures.
2. Policy Control and Path Selection: Administrators can manipulate BGP attributes to influence route selection, optimizing paths based on business needs, latency tolerance, or regulatory compliance.
3. Drawbacks: BGP's convergence time can be slower compared to OSPF, and its lack of inherent load balancing may lead to uneven traffic distribution without advanced configurations.

III. METHODOLOGY

This research employs a mixed-method approach, including both simulations and case studies to analyze OSPF and BGP within a hybrid cloud environment.

A. Simulation Environment and Setup

1. Topology Design: A simulated hybrid cloud network was constructed with multiple autonomous systems representing an on-premises data center, a private cloud, and connections to two public cloud providers.

2. Tools and Software: Cisco Packet Tracer, GNS3, and AWS Direct Connect were used to emulate real-world hybrid cloud conditions, including variable latency, packet loss, and route preferences.

B. Metrics for Evaluation

1. Latency: Measured as the round-trip time (RTT) across different nodes in the network.
2. Load Balancing: Evaluated by observing the distribution of traffic across available paths.
3. Convergence Time: Assessed by inducing network changes (e.g., link failure) and recording the time taken to achieve a stable routing table.
4. Resilience: Analyzed through intentional node failures to observe how protocols handle fault recovery.

Table 2 : Key Metrics for OSPF and BGP Evaluation

Metric	Description	Measurement Method
Latency	Average round-trip time (RTT)	Packet tracer
Load Balancing	Distribution of traffic across paths	Traffic distribution analysis
Convergence Time	Time taken to update routing tables	Network simulation logs
Resilience	Protocol's ability to recover from failures	Node failure and packet loss tests

C. Case Study: Real-World Hybrid Cloud Architecture

A case study of a large enterprise hybrid cloud was examined to understand how BGP and OSPF are deployed in production environments. This section includes interviews with network administrators and an analysis of traffic data from the enterprise network.

IV. PERFORMANCE EVALUATION

A. Latency Analysis

Latency is often a critical factor for cloud applications that require real-time data processing. This section presents a detailed comparison of OSPF and BGP under varying network conditions.

1. OSPF Latency: OSPF achieved an average RTT of 10-15 ms within a single autonomous system, demonstrating low-latency performance in intra-domain environments. This is primarily due to OSPF's fast convergence and efficient path calculation using LSAs.
2. BGP Latency: BGP's average RTT for cross-domain routes ranged from 25-40 ms, influenced by its path vector mechanism and additional routing policies. Although latency is higher, BGP's inter-domain resilience provides stability in long-distance routes, which is essential for multi-cloud environments.

B. Load Balancing Efficiency

Efficient load balancing is necessary to prevent traffic congestion and maintain network performance. OSPF and BGP differ significantly in their approach to load balancing.

1. OSPF's ECMP Support: OSPF's ability to balance loads across equal-cost paths enabled effective utilization of bandwidth within each domain. In our simulations, OSPF distributed traffic evenly across multiple links, which helped avoid bottlenecks and ensured resilience.
2. BGP Load Distribution: BGP's standard configuration lacks robust load balancing support, which can lead to uneven traffic distribution. However, advanced configurations, such as BGP Multipath, can partially address this limitation, albeit with increased complexity in setup and management.

C. Convergence Time Comparison

Convergence time is crucial in scenarios involving frequent topology changes, as prolonged convergence can lead to packet loss and degraded network performance.

1. OSPF Convergence Performance: OSPF's link-state mechanism allows it to converge within seconds in response to topology changes. Our simulations showed convergence times ranging from 1-3 seconds, suitable for environments where rapid adjustments are necessary.
2. BGP Convergence Performance: BGP's convergence time varied between 30 seconds and several minutes, depending on the number of ASes and route policies involved. This slower convergence reflects BGP's design for stability over speed, which can affect applications requiring real-time responses but is suitable for stable, large-scale environments.

D. Resilience and Fault Tolerance

Resilience is a critical requirement in hybrid cloud networks, where downtime can affect critical applications.

1. OSPF Resilience: OSPF's rapid convergence enhances its resilience within single-domain environments. However, OSPF's intra-domain limitation makes it less effective in handling cross-domain resilience.
2. BGP Resilience: BGP's policy-based routing and multiple-path capabilities provide robust resilience across AS boundaries. In the event of a link failure, BGP rerouted traffic through alternative paths, demonstrating minimal packet loss but with some delay due to slower convergence.

V. SECURITY CONSIDERATIONS IN OSPF AND BGP

Hybrid clouds often handle sensitive data, making routing protocol security critical.

A. OSPF Security Challenges

OSPF's security largely relies on MD5 or SHA authentication to prevent unauthorized routers from injecting malicious routes. However, as OSPF operates within an AS, it remains vulnerable to insider threats.

B. BGP Security Challenges

BGP is susceptible to various threats, such as route hijacking and prefix leaks, due to its reliance on trust between ASes. Techniques like RPKI (Resource Public Key Infrastructure) and BGPSEC can mitigate risks, though they require broad adoption to be fully effective.

VI. DISCUSSION AND ANALYSIS

The comparison of OSPF and BGP within hybrid cloud architectures highlights several key insights:

1. Use Cases for OSPF: OSPF is ideal for environments requiring low-latency intra-domain routing, such as within private clouds or on-premises networks.
2. BGP's Role in Multi-Cloud Architectures: BGP excels in providing robust inter-domain connectivity across public cloud providers and on-premises data centers.
3. Hybrid Approaches: Combining OSPF and BGP in hybrid cloud designs can optimize routing performance, with OSPF handling intra-domain traffic and BGP managing inter-domain connectivity.

VII. CONCLUSION AND FUTURE WORK

A. Latency Differences in Hybrid Cloud Contexts

- OSPF is shown to have lower latency in intra-domain (within an autonomous system) environments, making it ideal for scenarios where low-latency data transmission is critical, such as private cloud and on-premises networks. The Dijkstra-based path calculation contributes to its rapid convergence and low-latency performance.
- BGP has higher latency in inter-domain (cross-autonomous system) environments due to its path vector mechanism and policy-based routing approach. This latency trade-off, however, is balanced by BGP's enhanced control over path selection, which can be advantageous for applications where resilience and multi-cloud connectivity take priority over speed.

B. Load Balancing and Traffic Distribution Capabilities

- OSPF demonstrated superior load-balancing capabilities with its built-in Equal-Cost Multi-Path (ECMP) support. This feature allowed for effective traffic distribution within each autonomous system, preventing traffic congestion and optimizing resource utilization in intra-domain environments.
- BGP, on the other hand, traditionally lacks robust load-balancing support in its standard form. Advanced configurations, such as BGP multipath extensions, can partially address this, but at the cost of increased configuration complexity. This highlights a limitation of BGP in hybrid cloud architectures where dynamic and evenly distributed traffic is crucial.

C. Convergence Time and Network Stability

- OSPF was found to have significantly faster convergence times than BGP, with the ability to stabilize within seconds after topology changes. This makes it suitable for networks where frequent changes occur, as OSPF can quickly adapt to new routes, minimizing disruptions.
- BGP has slower convergence times, sometimes taking several minutes to fully update routes across multiple autonomous systems. However, BGP's slower convergence is by design, favoring stability and preventing route flapping, which is particularly beneficial for stable, large-scale environments. For hybrid clouds with intermittent topology changes, this stability may be more advantageous than fast convergence.

D. Resilience and Fault Tolerance in Hybrid Cloud Connectivity

- OSPF offers high resilience within a single domain, as its rapid convergence allows quick adaptation to failures within private cloud or on-premises environments. However, it lacks inherent resilience across domains, making it less suitable for managing multi-cloud resilience.
- BGP demonstrates superior resilience in inter-domain environments. BGP's ability to maintain multiple alternative paths provides fault tolerance across autonomous systems, making it ideal for multi-cloud and hybrid cloud architectures. In cases of link failure, BGP's policy-based path control can ensure traffic reroutes with minimal disruption, although some convergence delay is expected.

E. Security Implications in Hybrid Cloud Routing

- OSPF relies on authentication methods such as MD5 or SHA to secure route advertisements, which works well within trusted intra-domain environments. However, OSPF is vulnerable to insider threats and lacks cross-domain security, which is essential in multi-cloud setups.
- BGP is vulnerable to security risks such as route hijacking and route leaks, which can have severe implications in multi-cloud settings. Newer security enhancements like RPKI (Resource Public Key Infrastructure) and BGPSEC can mitigate some risks but require widespread adoption to be effective. This finding underscores the need for enhanced security measures, especially for BGP when used in hybrid cloud configurations involving sensitive data or mission-critical applications.

F. Applicability of Hybrid Routing Architectures

- This study finds that a hybrid approach, leveraging both OSPF and BGP, can offer optimized routing in hybrid cloud environments. For example:

  - Intra-domain Routing: OSPF can be used within individual cloud segments or data centers, where its low-latency routing and fast convergence are advantageous.
  - Inter-domain Routing: BGP can manage traffic between different autonomous systems, handling connections across public clouds and between clouds and on-premises environments where resilience and policy-driven path selection are essential.
- This hybrid configuration allows organizations to capitalize on the strengths of each protocol, creating a network architecture that balances speed, resilience, and scalability.

G. Insights into Future Adaptive Routing Protocols

- The findings suggest a potential future direction for adaptive routing protocols that could dynamically switch between OSPF and BGP based on real-time network conditions. Such an adaptive approach could enable networks to balance low latency with resilience dynamically, potentially enhancing hybrid cloud performance further.

REFERENCES

1. Gartner, "Forecast Analysis: Public Cloud Services, Worldwide, 2022," Feb. 2022.
2. J. Moy, "OSPF Version 2," RFC 2328, Apr. 1998.
3. Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, Jan. 2006.
4. M. Chiesa et al., “Inter-domain Traffic Engineering with BGP,” IEEE Communications Magazine, vol. 55, no. 6, pp. 146-153, June 2017.
5. S. Hares and K. Patel, “BGP-LS (BGP Link-State),” RFC 7752, Mar. 2016.