International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

A Comparative Performance Evaluation of OSPF and IS-IS for Enterprise and Service Provider Networks

Nikhil Bhagat

Principal Network Engineer, Independent Scholar, Network Engineering
nikhil.bhagat90@gmail.com

Abstract:
Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) are two core link-state routing protocols of current computer networks. The paper offers an in-depth comparison of these protocols, covering their features, benefits and drawbacks in both enterprise and service provider networks. OSPF, which works with IP networks, has become an accepted enterprise network solution that is capable of being deployed with ease and efficiency in small, well-designed networks. Yet its viability in large and complex service provider networks can be hindered by scalability issues. Meanwhile, IS-IS, developed for International Organization for Standardization (ISO) Connectionless Network Protocol, is popular among service providers because of its scalability and ease of deployment, specifically for large deployments. It also goes over the history of both protocols, their differences in their behavior, and the use cases where each performs best. Finally, the decision of using OSPF or IS-IS comes from network's size, complexity and required performance. The analysis can help network architects determine which protocol should be used for which network infrastructures.

Keywords: OSPF, IS-IS, LSAs, LSPs, Enterprise networks, Service Provider networks.

I. Introduction

In today's evolving networking environment, choosing appropriate routing protocol is a matter of critical importance for data flows and network performance. Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) are two important link state routing protocols, both used for routing decisions in IP networks [1]. The Internet Engineering Task Force (IETF) created OSPF to support routing in IP networks, while IS-IS was originally designed for the ISO networking paradigm and later adapted as IP routing [2].

Both protocols come with their own unique features and capabilities, which suit a variety of network environments. This paper will compare OSPF and IS-IS in terms of historical evolution, benefits, and problems in enterprise and service provider networks. The analysis will provide network administrators necessary mechanisms to evaluate their environments before finalizing an Interior Gateway Protocol (IGP).

II. Evolution of OSPF

OSPF was designed in early 1990s as a response to the limitations within Routing Information Protocol (RIP) for large, complex networks [3]. The first OSPFv1 was released in 1989, followed by major

improvements in OSPFv2, which became the IP network standard [4]. OSPF's development revolved around convergence speed, scalability, and bandwidth management.
The biggest innovation of OSPF is its hierarchical structure, which allows a network to be separated into areas. This hierarchical approach makes routing tables smaller and enables faster convergence as the route advertisements have a narrow range. OSPF computes the shortest path tree using the Dijkstra algorithm, making efficient link cost based routing decisions.
OSPF has also received a couple of updates, such as OSPFv3, which added IPv6 support, and even more enhancements in link-state advertisement (LSA) type [5]. The protocol's capability to respond to evolving network needs has made it a clear choice as a Interior Gateway Protocol (IGP) in enterprise networks.

III. HOW OSPF USES LSA'S TO UPDATE TOPOLOGY DATABASE AND MAKE ROUTING DECISIONS

Open Shortest Path First (OSPF) communicates network topology data between routers in an OSPF area by way of Link-State Advertisements (LSAs). LSAs are very important for OSPF as it enables routers to share their neighbors, links, and network states in order to calculate the shortest path to each destination. There are multiple LSA types that help OSPF enabled routers build a topology view.

A. LSA Types

OSPF specifies several types of LSAs with a different function [6]. The most common LSA types are:

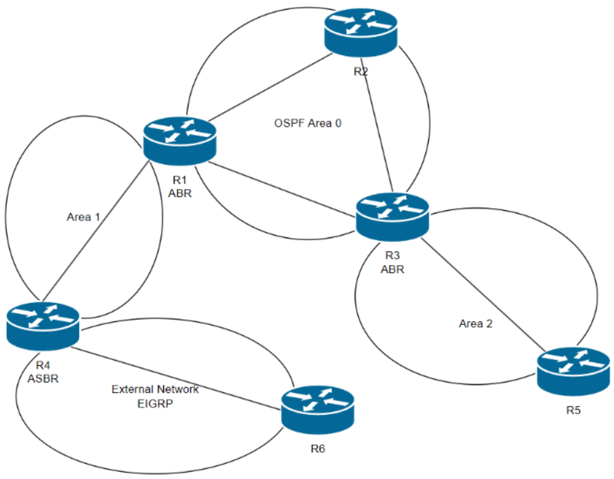

Fig 1. Enterprise Network using OSPF

- Type 1: Router LSA – Lists and advertises the router's directly connected interfaces and connections in the area.
- Type 2: Network LSA – Specified by the Designated Router (DR) for broadcast and non-broadcast multi-access (NBMA) networks. It lists the routers attached to the network.
- Type 3: Summary LSA – Advertises networks across OSPF areas (inter-area). This information is provided by Area Border Routers (ABRs).
- Type 4: ASBR Summary LSA – Informs the rest of the areas about the Autonomous System Boundary Router (ASBR).
- Type 5: AS External LSA – This LSA advertises routes external to the OSPF autonomous system (typically from other routing protocols such as BGP).
- Type 7: NSSA LSA – This LSA is generated only in Not-So-Stubby Areas (NSSA) for bringing external routes, such as Type 5 LSAs, inside an NSSA.

B. LSA Flooding
When a router sends or receives an LSA, it propagates that LSA to every other router in the OSPF [7]. This keeps all routers at the same level of visibility to the network topology. LSAs are filled continuously via an acknowledgment system with sequence numbers that eliminate loops or incorrect information.

C. Link-State Database (LSDB)
All routers store received LSAs into a Link-State Database (LSDB). LSDB is a complete representation of the network as given by the LSAs from all the routers in the area [8].

D. Shortest Path First (SPF) Calculation
In the LSDB, routers are utilizing Dijkstra's SPF algorithm to determine the shortest route to any destination in the network [9]. This computation gets recorded in the router's OSPF routing table.

E. LSA Aging and Refresh
LSA's expire after a limited period of time (usually 30 minutes). They are automatically updated on a regular basis to keep the network topology consistent. When no refresh occurs for a given LSA, it's deemed "aged out" and deleted from the LSDB.

F. LSA Propagation Across Areas
For network with many OSPF areas, LSAs are randomly distributed from one area to another through Area Border Routers (ABRs). ABRs convert Type 1 and Type 2 LSAs from one area to Type 3 (Summary LSAs) when exchanging information between areas [4].

This dynamism of LSA exchange allows OSPF routers to converge dynamically and precisely, when a change in the network topology is detected.

IV. WHY SHOULD ENTERPRISE NETWORKS ADOPTENTERPRISE NETWORKS PRESENT UNIQUE CHALLENGES THAT REQUIRE A POWERFUL ROUTING PROTOCOL. OSPF COMES WITH A SET OF FEATURES THAT MAKE IT A PROMISING SOLUTION FOR COMPANIES:

A. Scalability
The OSPF hierarchy allows enterprises to expand their networks with minimal performance loss [10]. When OSPF divides the network into areas, it reduces the routing table size and the size of route advertisements.

B. Fast Convergence
OSPF has faster convergence times, which allows the networks to react in time to topology changes. This is necessary for application uptime and low downtime.

C. Load Balancing
OSPF can provide Equal-Cost Multi-Path (ECMP) routing for multiple routes to a destination. This enhances bandwidth use and provides redundancy.

D. Integration with Security
OSPF provides support for authentication, which enables secure routing updates [1]. This is extremely important for business networks with sensitive information.

E. Widespread Adoption
OSPF is widely deployed and maintained across vendors to offer interoperability and ease of implementation.

With these advantages, OSPF is increasingly popular among enterprise networks for reliable and effective routing.

V. CHALLENGES WITH OSPF FOR SERVICE PROVIDER NETWORKS

While OSPF works great in enterprise networks, it faces a few challenges when deployed within service provider networks:

A. Scalability Limitations

Although OSPF is designed to scale, the size of the service provider networks tends to overwhelm it [11]. Due to having a lot of routing tables to handle and multiple area zones to be controlled, performance issues may arise.

B. Convergence Times

Networks between Service Providers need super-fast convergence to accommodate dynamic changes in the network. The convergence times of OSPF are typically short but might fall short of the demands of complex service provider networks [12].

C. Complex Configuration

OSPF's hierarchical model, while helpful in enterprise networks, adds complexity in service provider networks involving multiple protocols. Setting up OSPF with other protocols can create operational issues.

D. Resource Consumption

OSPF can consume higher CPU and memory, especially when deployed on large scale environments. This affects network devices' performance and potentially cause bottlenecks.

E. Limitations in Traffic Engineering

OSPF is not capable of handling traffic engineering as the other protocols. Service providers need advanced traffic handling, which OSPF does not provide.

These issues have caused service providers to find alternatives to route protocols that are a better fit for their operations.

VI. INTRODUCTION TO IS-IS

Intermediate System to Intermediate System (IS-IS) is a link-state routing protocol originally created for ISO networks [13]. Later it was used for IP routing and was adopted in service provider networks. IS-IS works much like OSPF: it uses link-state advertisements to broadcast routing information between routers. One of IS-IS's uniqueness is that it can support both IPv4 and IPv6 inside the same routing domain, so service providers do not have to worry about moving to a new network during upgrade. IS-IS uses an intuitive and simple routing algorithm for routing, so it is ideal for large scale networks.

IS-IS also supports a hierarchical architecture with routing domain levels, ensuring the scale and effective routing in massive service provider networks [14]. Its adaptability and simplicity has made the protocol popular with most service providers.

VII. HOW IS-IS USES LSP'S TO UPDATE TOPOLOGY DATABASE AND MAKE ROUTING DECISIONS

The concept of Link-State Advertisements (LSAs) within OSPF is referred to as Link-State Protocol Data Units (LSPs) in Intermediate System to Intermediate System (IS-IS). While IS-IS and OSPF are link state protocols, ISIS has different terminologies and hierarchies to communicate routing information.

A. LSP (Link-State PDU) Overview

In IS-IS, routers (Intermediate Systems, ISs) create LSPs to map their links, interfaces, and accessible networks [15]. LSPs store a router's neighbor information and link status, just like OSPF LSAs. The routers extract this data to form a complete representation of the network topology.

B. Types of PDUs (Protocol Data Units)

IS-IS implements two types of Protocol Data Units [2]:

- LSP (Link-State PDU): Represents link-state data like OSPF LSAs. Every IS (router) creates LSPs, which reflect local connections and reachability.
- CSNP (Complete Sequence Number PDU): Defines the list of all LSPs on the network so that all routers are synchronized.
- PSNP (Partial Sequence Number PDU): It is used by an IS-IS enabled router to request a specific LSP from a neighboring router, which could be missing from its database.

C. Level 1 and Level 2 LSPs

IS-IS divides its routing into two levels [16]:

- Level 1 LSPs: Generated by IS-IS enabled routers of a same area (as with intra-area routing in OSPF). The routers leverage these LSPs to learn the topology of their current area.
- Level 2 LSPs: Generated by IS-IS enabled routers that are backbone routers (such as OSPF's Area Border Routers, ABRs). These LSPs inform about inter-area paths and serve to map between IS-IS areas.

D. LSP Flooding

LSPs are flooded into the network so all routers share the same topology data. All IS (routers) produce LSP and push them out among the neighbors. When an LSP arrives on a router, the router broadcasts it to each neighbor so that each router has a current snapshot of the state of the network.

F. Sequence Numbers and Aging

The sequence number on each LSP is incremented with each new LSP update. Additionally, LSPs are retained for 20 minutes by default, and once they age out, they are automatically removed from the database without being refreshed.

G. Link-State Database (LSDB)

Each router stores received LSP information in an LSDB. This database contains all the LSPs received by the router and it's used to create an entire map of the network topology [17].

H. Shortest Path First (SPF) Calculation

Once these LSPs are stored in the LSDB, routers perform the Dijkstra Shortest Path First (SPF) algorithm to find the optimal routing of the network. Each router builds a shortest path tree of all the routes going to each destination.

I. LSP Fragmentation

IS-IS also support LSP fragmentation when the amount of information does not fit into the physical limit of one LSP packet. The router then creates several LSP fragments and each fragment gets flooded on the network. This mechanism ensures IS-IS can perform well on large networks.

J. LSPs in Multi-Level IS-IS

There are two routing levels in IS-IS:

- Level 1: It is used to route inside an area (intra-area), Level 1 links swap Level 1 LSPs.
- Level 2: In order to perform cross area (inter-area) routing, routers trade Level 2 LSPs [13]. Level 2 routers function like OSPF Area Border Routers (ABRs), summarizing information between area.

K. LSP Propagation and Integrity

LSPs are transmitted safely via sequence numbers and acknowledgments. The router requests the newest version from neighbors if it detects an outdated LSP. This will ensure consistency in the network.

VIII. Key Differences Between OSPF LSAs and IS-IS LSPs

Although both are link-state routing protocols, OSPF (Open Shortest Path First) and IS-IS (Intermediate System to Intermediate System) share topology information in very different ways. Following is an outline of the main differences between OSPF LSAs and IS-IS LSPs:

TABLE 1. DIFFERENCES BETWEEN OSPF LSA AND IS-IS LSP [2]

Feature	**OSPF LSAs**	**IS-IS LSPs**
Terms	Link-State Advertisements (LSAs)	Link-State PDUs (LSPs)
OSI Layer	Operates at Layer 3 of OSI	Operates over Layer 2 os OSI
Hierarchical structure	All areas need to connect to a backbone area (Area 0)	Routers can be part of Level 1 (intra-area) and/or Level 2 (inter-area)
Types of Advertisements	Multiple types of LSAs (Type 1-5, 7) and few others	LSPs carry link state data within a PDU with sub-TLVs
Area Border Routers (ABRs)	OSPF uses Area Border Routers (ABRs) to summarize routes between areas	IS-IS uses Level-1/Level-2 routers for inter-area route exchange
Routing Scope	LSAs are divided into intra-area, inter-area, and external routes	LSPs are divided into Level 1 (intra-area) and Level 2 (inter-area)
Flooding Mechanism	OSPF floods LSAs with strict hierarchy between areas (backbone area connects other areas).	IS-IS floods LSPs across Level 1 and Level 2 networks; no backbone area is needed.
Encapsulation	OSPF uses IP (Protocol 89) for LSA encapsulation	IS-IS operates over the data link layer using ISO's Connectionless Network Protocol and is encapsulated within Ethernet.
External Routes	OSPF uses Type 5 LSAs for external routes	IS-IS uses Level 2 LSPs for external routes
Metric Structure	OSPF uses cost as a metric	IS-IS supports multiple metric types (e.g., wide metrics)
LSP Fragmentation	OSPF LSAs do not support fragmentation;	IS-IS supports LSP fragmentation to scale within large scale networks

IX. Why is IS-IS a Better Solution for Service Provider Networks

IS-IS is scalable in many respects due to several key attributes, most significantly in large and advanced networks such as service provider networks:

A. Entirely operates at the Data Link Layer:

IS-IS operates over the data link layer (Layer 2), so it does not need IP for operations [15]. This makes it uncoupled from IP address conflicts (i.e., IPv4 vs IPv6) and makes packets handled more effectively on massive networks. OSPF meanwhile is running over IP which further adds complexity and overhead.

B. Flat Hierarchy structure:

IS-IS Level 1 (intra-area) and Level 2 (inter-area) hierarchy is simpler and more extensible than OSPF strict hierarchy with backbone area (Area 0). This modular architecture makes it possible for IS-IS to deal with routing on large, interconnected networks without hard borders [13].

IS-IS routers have both Level 1 and Level 2 capabilities, reducing area constraints and management. OSPF imposes very strict boundaries between areas and routes them across Area Border Routers (ABRs).

C. Supports LSP Fragmentation:

IS-IS is configured to support LSP fragmentation where routers will partition link state data into smaller chunks when it is larger than the maximum PDU [2]. This makes it possible for IS-IS to process a larger routing data in networks with many routers and links. OSPF LSAs do not support fragmentation, which makes IS-IS more robust in big-data scenarios.

D. Has a support for Wide Metrics:

IS-IS has wide metrics allowing support for much larger networks and complicated path calculations [18]. OSPF's default metric is 16 bits, while IS-IS wide metric allows you to scale up much larger values for networks that are extremely large, such as ISP or data center networks.

E. Independent of IP Addressing:

IS-IS operates directly on the data link layer, so it is not dependent on IP addresses (IPv4 or IPv6). Unlike IS-IS, OSPF needs its own instances for IPv4 and IPv6. This makes things simpler with IS-IS when it comes to multi-protocol environments.

F. Provides Flood suppression:

IS-IS has a flooding suppression mechanism to minimize the unnecessary spread of LSPs [19]. This is especially useful in large networks, where blocking excessive flooding of LSPs is necessary to prevent routers from becoming overloaded with updates.

G. Support for MPLS Networks:

IS-IS is widely deployed in MPLS (Multiprotocol Label Switching) networks, where it scales well and fits in with the service provider landscape [15]. OSPF is also a potential network candidate in MPLS networks but IS-IS is the choice of many since it's simple and easy to scale.

H. High Convergence speeds

IS-IS provides rapid convergence - a critical advantage for service provider networks that expect lower downtime. The protocol is able to respond quickly to network changes while maintaining continuous service.

I. Efficient Resource Consumption

IS-IS makes efficient use of resources, thereby consuming less CPU and memory than OSPF. This is especially useful in resource intensive environments where maximizing resources optimization is vital.

J. Facilitates Traffic Engineering

IS-IS provides traffic engineering functionality that allows providers to better control the network traffic [2]. This includes Multiprotocol Label Switching (MPLS) and explicit routing, which enhances network performance.

K. Simpler Configuration

IS-IS has a less complicated configuration model than OSPF, which makes it easier to handle large networks. Using this simplicity of configuration enables providers to deploy and run their networks more effectively.

L. Provides a Dual-Stack Support

IS-IS comes with native support for IPv4 as well as IPv6 that makes dual stack environments easy for service providers to administer [20]. This is critical as networks evolve to IPv6, enabling easier upgrade without having to reconfigure routing.

With these features, IS-IS should be a preferred routing solution choice for service providers looking for a high-performance, scalable routing solution.

X. CONCLUSION

In conclusion, a business should evaluate its network requirements and other relevant factors when choosing between OSPF and IS-IS as an Interior Gateway protocol. OSPF has proven to be efficient for enterprise networks that need scalability, convergence speed and vendor support. However, its limited scalability within large scale networks and higher resource utilization could compromise its performance in service provider environments.

However, IS-IS provides service providers with an attractive alternative, with scalable, low-consumption, and traffic engineering capabilities. As networks evolve, knowing advantages, disadvantages and caveats of each routing protocol plays a vital part while taking strategic and organization-wide decisions. This comparative evaluation between OSPF and IS-IS will serve as a guide for network administrators while choosing the right protocol for their environments.

REFERENCES

1. "Open Shortest Path First (OSPF)," Cisco, 2020.
2. Bhatia, Y. Ohara, and V. Manral, "IS-IS and OSPF difference discussions," IETF J., vol. 15, no. 3, pp. 13-19, 2019.
3. J. Moy, "OSPF protocol analysis," RFC 1245, Internet Engineering Task Force, 1991.
4. J. Moy, "OSPF version 2," RFC 2328, Internet Engineering Task Force, 1998.
5. D. Katz, K. Kompella, and D. Yeung, "Traffic engineering (TE) extensions to OSPF version 2," RFC 3630, Internet Engineering Task Force, 2003.
6. R. Coltun, "The OSPF opaque LSA option," RFC 2370, Internet Engineering Task Force, 1998.
7. "RFC ft-ietf-ospf-opaque: The OSPF opaque LSA option," Internet Engineering Task Force, 1998.
8. L. Choudhury, A. Maunder, and V. Sapozhnikova, "Faster link-state IGP convergence and improved network scalability and stability," IEEE Commun. Mag., vol. 40, no. 7, pp. 28-34, 2002.
9. J. Behrens and J. Garcia-Luna-Aceves, "Distributed, scalable routing based on link-state vectors," in Proc. 1994 IEEE INFOCOM, 1994, pp. 50-58.
10. Y. Lan and Z. Chen, "The research on the OSPF security optimizing of campus network," in Proc. 2010 Int. Conf. Comput. Eng. Technol. (ICCET), 2010, pp. V6-351-V6-354.
11. S. Hartman and D. Zhang, "Analysis of OSPF security according to the Keying and Authentication for Routing Protocols (KARP) design guide," RFC 6863, Internet Engineering Task Force, 2013.
12. Jiang, M. Xu, Q. Li, and L. Pan, "Improving IGP convergence through distributed OSPF in scalable routers," Int. J. Commun. Syst., vol. 23, no. 5, pp. 649-660, 2010.
13. T. Li, T. Przygienda, and H. Smit, "Domain-wide prefix distribution with two-level IS-IS," RFC 5302, Internet Engineering Task Force, 2008.
14. R. White and S. Zandi, "IS-IS support for Openfabric," RFC 7356, Internet Engineering Task Force, 2014.
15. Yang, J. Hahm, Y. Kim, and S. Kim, "Design and implementation of the IS-IS routing protocol with traffic engineering," in Proc. 2005 IEEE GLOBECOM, 2005, pp. 3058-3062.
16. D. Katz, R. K. Saluja, and D. Eastlake, "Three-way handshake for IS-IS point-to-point adjacencies,"

RFC 5303, Internet Engineering Task Force, 2008.
17. "Point-to-point operation over LAN in link-state routing protocols," RFC 5301, Internet Engineering Task Force, 2008.
18. "Domain-wide prefix distribution with two-level IS-IS," RFC 5302, Internet Engineering Task Force, 2008.
19. R. Callon, "Use of OSI IS-IS for routing in TCP/IP and dual environments," RFC 1195, Internet Engineering Task Force, 1990.
20. "Configuring IS-IS dual stacking of IPv4 and IPv6 unicast addresses," Cisco, 2019.