UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

SMART PATH CONNECTIONS, LLC	§	
	§	
Plaintiff,	§	Case No.
	§	
v.	§	**JURY TRIAL DEMANDED**
	§	
NOKIA CORPORATION, NOKIA SOLUTIONS AND NETWORKS OY, AND NOKIA OF AMERICA CORPORATION	§ § §	
	§	
Defendants.		

COMPLAINT FOR PATENT INFRINGEMENT

Smart Path Connections, LLC ("Smart Path" or "Plaintiff"), by and through its attorneys, for its Complaint for patent infringement against Nokia of America Corporation, Nokia Solutions and Networks Oy, and Nokia of America Corporation ("Nokia" or "Defendants"), and demanding trial by jury, hereby alleges, on information and belief with regard to the actions of Defendants and on knowledge with regard to its own actions, as follows:

I. NATURE OF THE ACTION

1. This is an action for patent infringement arising under the patent laws of the United States, 35 U.S.C. §§ 271, *et seq*., to enjoin and obtain damages resulting from Defendants' unauthorized use, sale, and offer to sell in the United States, of products, methods, processes, services and/or systems that infringe Plaintiff's United States patents, as described herein.

2. Defendants manufacture, provide, use, sell, offer for sale, import, and/or distribute infringing products and services, and encourages others to use its products and services in an infringing manner, as set forth herein.

3. Plaintiff seeks past and future damages and prejudgment and post-judgment interest for Defendants' infringement of the Asserted Patents, as defined below.

II. PARTIES

4. Plaintiff Smart Path is a limited liability company organized and existing under the law of the State of Delaware, with its principal place of business located at 101 E. Park Blvd., Suite 600, Plano, TX 75074.

5. Smart Path is the owner of the entire right, title, and interest of the Asserted Patents, as defined below.

6. Defendant Nokia Corporation ("Nokia Corp.") is a Finnish corporation with its principal place of business at Karaportti 3, FI-02610 Espoo, Finland. Upon information and belief, Alcatel-Lucent S.A. ("Alcatel-Lucent") was merged into Nokia Corp.'s "Nokia Networks" division in 2016.

7. Defendant Nokia Solutions and Networks Oy is a corporation organized and existing under the laws of Finland with its principal place of business at Karaportti 3, 02610 Espoo, Finland. On information and belief, Nokia Solutions and Networks Oy is a wholly owned subsidiary of Nokia Corp.

8. Nokia of America Corporation is a Delaware corporation with its U.S. Headquarters in Dallas, Texas. Nokia may be served through its registered agent Prentice Hall Corporation System, 211 E. 7th Street, Suite 620, Austin, Texas 78701. On information and belief, Nokia is registered to do business in the State of Texas and has been since at least December 29, 1995.

9. On information and belief, Nokia of America Corporation is an indirect wholly owned subsidiary of Nokia Corporation and Nokia Solutions and Networks Oy.

10. Nokia Corp., Nokia Solutions and Networks Oy, and Nokia of America are collectively referred to as "Nokia."

11. The Nokia Defendants hold themselves out as a single "Nokia" company, exemplified in the company's website, www.nokia.com. Nokia offers for sale and sells the accused products, through that website.

12. Nokia conducts business operations within the Eastern District of Texas, including its offices located at 2525 Highway 121, Lewisville, Texas 75056 and 601 Data Drive, Plano, Texas 75075. Nokia has offices in the Eastern District of Texas where it sells and/or markets its products, including its offices in Lewisville and Plano, Texas.

13. Nokia maintains additional offices throughout Texas including its U.S. headquarters in Dallas and office in this district.

III. JURISDICTION AND VENUE

14. This is an action for patent infringement which arises under the patent laws of the United States, in particular, 35 U.S.C. §§ 271, 281, 283, 284, and 285.

15. This Court has exclusive jurisdiction over the subject matter of this action under 28 U.S.C. §§ 1331 and 1338(a).

16. This Court has personal jurisdiction over Nokia in this action because Nokia has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over Nokia would not offend traditional notions of fair play and substantial justice. Nokia, directly and/or through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the patents-in-suit. Moreover, Nokia is registered to do business in the State of Texas, has offices and facilities in the State of Texas, and actively directs its activities to customers located in the State of Texas.

17. Venue is proper in this district under 28 U.S.C. §§ 1391(b)–(d) and 1400(b). Defendant Nokia of America Corporation is registered to do business in the State of Texas, has offices in the State of Texas, and upon information and belief, has transacted business in the Eastern District of Texas and has committed acts of direct and indirect infringement in the Eastern District of Texas. Nokia of America Corporation maintains a regular and established place of business in the Eastern District of Texas, including offices located at 2525 Highway 121, Lewisville, Texas 75056 and 601 Data Drive, Plano, Texas 75075. Nokia has operated the Plano office as a "NokiaEDU Training Center," which it describes as "the company's premiere learning organization serving customers, partners and employees worldwide . . . . to deliver[] a top-quality learning experience tailored to our customers' specific requirements and preferences." https://learningstore.nokia.com/locations/files/US-Plano.pdf.

18. Venue is proper as to Nokia Corp. under 28 U.S.C. § 1391(c)(3) as a corporation that is not resident in the United States.

19. Venue is proper as to Nokia Solutions and Networks Oy under 28 U.S.C. § 1391(c)(3) as a corporation that is not resident in the United States.

IV. COUNTS OF PATENT INFRINGEMENT

20. Plaintiff alleges that Defendants have infringed and continue to infringe the following United States patents (collectively the "Asserted Patents"):

United States Patent No. 7,386,010 (the "'010 Patent") (Exhibit A)
United States Patent No. 7,463,580 (the "'580 Patent") (Exhibit B)
United States Patent No. 7,551,599 (the "'599 Patent") (Exhibit C)
United States Patent No. 7,697,525 (the "'525 Patent") (Exhibit D)

COUNT ONE
INFRINGEMENT OF U.S. PATENT 7,386,010

21. Plaintiff incorporates by reference the allegations in all preceding paragraphs as if fully set forth herein.

22. The '010 Patent, entitled "Multiprotocol media conversion," was filed on June 13, 2003, and issued on June 10, 2008.

23. Plaintiff is the assignee and owner of all rights, title and interest to the '010 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

24. The '010 Patent addresses problems in the prior art of "providing different types of Layer 2 network service over a common packet network infrastructure." 1:12-14.

25. The '010 discloses a solution to this problem in which "interworking of Layer 2 services enables endpoints using disparate protocols to communicate with one another over the same VPN." 1:62-64.

Direct Infringement

26. Defendants, without authorization or license from Plaintiff, have been and are directly infringing the '010 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale telecommunications equipment that infringes one or more claims of the '010 Patent. Defendants develop, designs, manufacture, and distribute telecommunications equipment that infringes one or more claims of the '010 Patent. Defendants further provide services that practice methods that infringe one or more claims of the '010 Patent. Defendants are thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities

include Defendants' 7705 Service Aggregation Router and all other substantially similar products (collectively the "'010 Accused Products").

27. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendants' infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '010 Accused Products.

28. Defendants are liable for direct infringement pursuant to 35 U.S.C. § 271 for the manufacture, sale, offer for sale, importation, or distribution of Defendants' 7705 Service Aggregation Router .

29. Defendants' 7705 Service Aggregation Router is a non-limiting example of an apparatus that meets all limitations of claim 1 of the '010 Patent, either literally or equivalently.

30. The 7705 Service Aggregation Router comprises an apparatus for data communications.

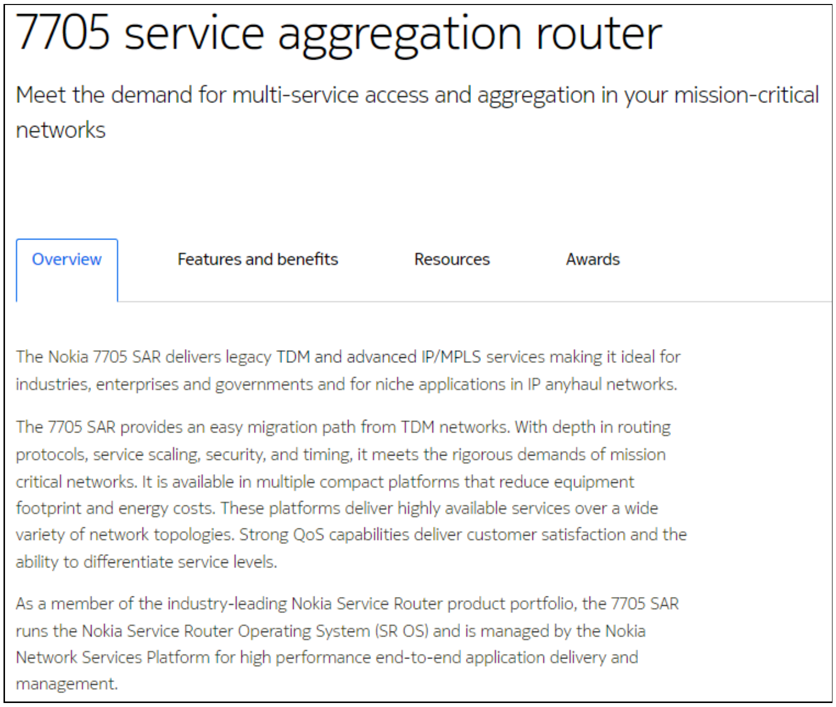

7705 service aggregation router

Meet the demand for multi-service access and aggregation in your mission-critical networks

The Nokia 7705 SAR delivers legacy TDM and advanced IP/MPLS services making it ideal for industries, enterprises and governments and for niche applications in IP anyhaul networks.

The 7705 SAR provides an easy migration path from TDM networks. With depth in routing protocols, service scaling, security, and timing, it meets the rigorous demands of mission critical networks. It is available in multiple compact platforms that reduce equipment footprint and energy costs. These platforms deliver highly available services over a wide variety of network topologies. Strong QoS capabilities deliver customer satisfaction and the ability to differentiate service levels.

As a member of the industry-leading Nokia Service Router product portfolio, the 7705 SAR runs the Nokia Service Router Operating System (SR OS) and is managed by the Nokia Network Services Platform for high performance end-to-end application delivery and management.

https://www.nokia.com/networks/products/7705-service-aggregation-router/

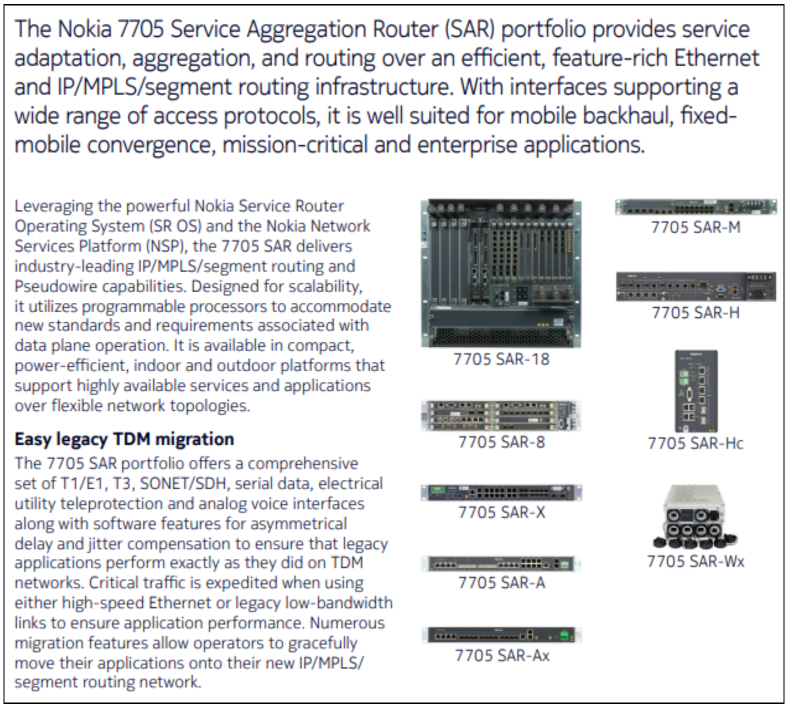

The Nokia 7705 Service Aggregation Router (SAR) portfolio provides service adaptation, aggregation, and routing over an efficient, feature-rich Ethernet and IP/MPLS/segment routing infrastructure. With interfaces supporting a wide range of access protocols, it is well suited for mobile backhaul, fixed-mobile convergence, mission-critical and enterprise applications.

Leveraging the powerful Nokia Service Router Operating System (SR OS) and the Nokia Network Services Platform (NSP), the 7705 SAR delivers industry-leading IP/MPLS/segment routing and Pseudowire capabilities. Designed for scalability, it utilizes programmable processors to accommodate new standards and requirements associated with data plane operation. It is available in compact, power-efficient, indoor and outdoor platforms that support highly available services and applications over flexible network topologies.

Easy legacy TDM migration

The 7705 SAR portfolio offers a comprehensive set of T1/E1, T3, SONET/SDH, serial data, electrical utility teleprotection and analog voice interfaces along with software features for asymmetrical delay and jitter compensation to ensure that legacy applications perform exactly as they did on TDM networks. Critical traffic is expedited when using either high-speed Ethernet or legacy low-bandwidth links to ensure application performance. Numerous migration features allow operators to gracefully move their applications onto their new IP/MPLS/ segment routing network.

https://onestore.nokia.com/asset/f/162833 (page 1)

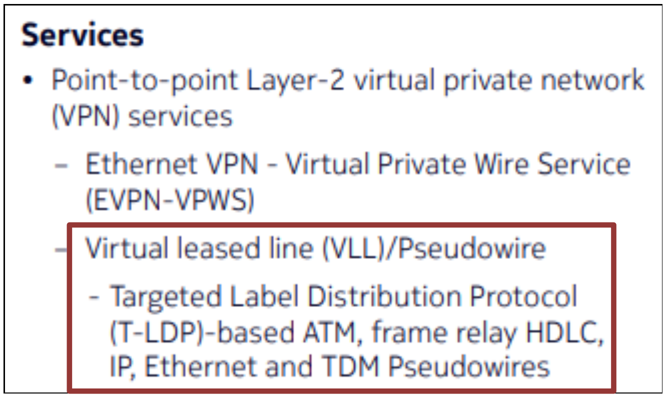

Services

- Point-to-point Layer-2 virtual private network (VPN) services
  - Ethernet VPN - Virtual Private Wire Service (EVPN-VPWS)
  - Virtual leased line (VLL)/Pseudowire
    - Targeted Label Distribution Protocol (T-LDP)-based ATM, frame relay HDLC, IP, Ethernet and TDM Pseudowires

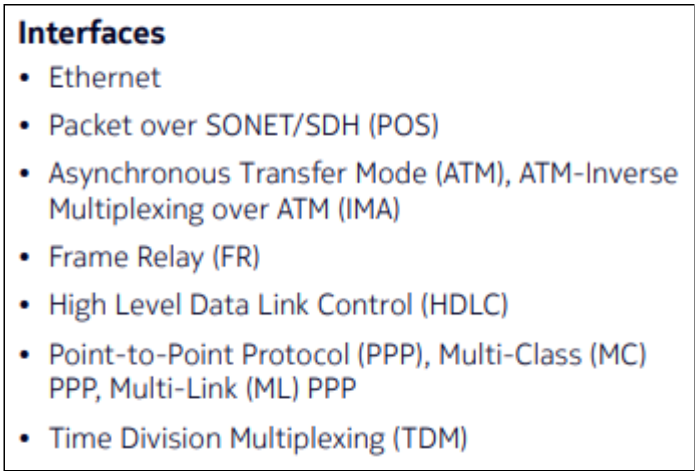

Interfaces

- Ethernet
- Packet over SONET/SDH (POS)
- Asynchronous Transfer Mode (ATM), ATM-Inverse Multiplexing over ATM (IMA)
- Frame Relay (FR)
- High Level Data Link Control (HDLC)
- Point-to-Point Protocol (PPP), Multi-Class (MC) PPP, Multi-Link (ML) PPP
- Time Division Multiplexing (TDM)

https://onestore.nokia.com/asset/f/162833 (page 3)

31. The 7705 Service Aggregation Router comprises a hub, comprising a plurality of ports, which are configured to receive and transmit data frames in accordance with a packet-oriented Layer 2 communication protocol; and a plurality of edge devices:

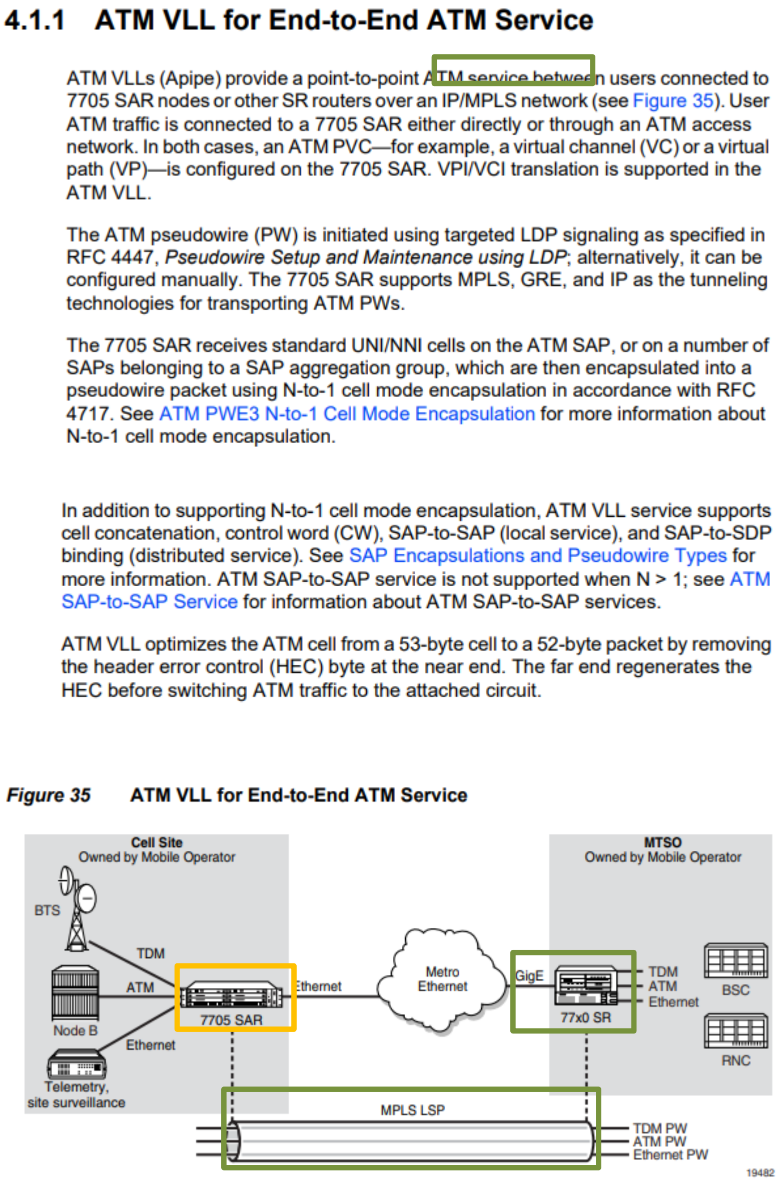

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

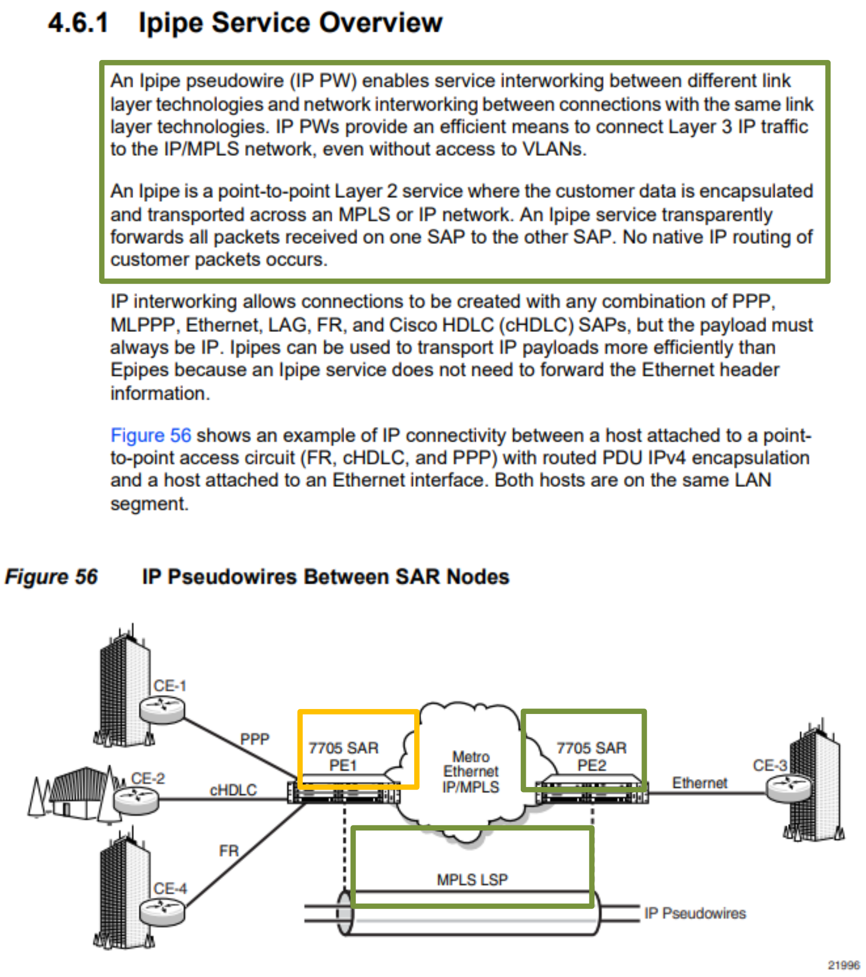

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

32. The 7705 Service Aggregation Router comprises edge devices that comprise at least one network port for communicating with the ports of the hub via a network in accordance with the packet-oriented Layer 2 communication protocol:

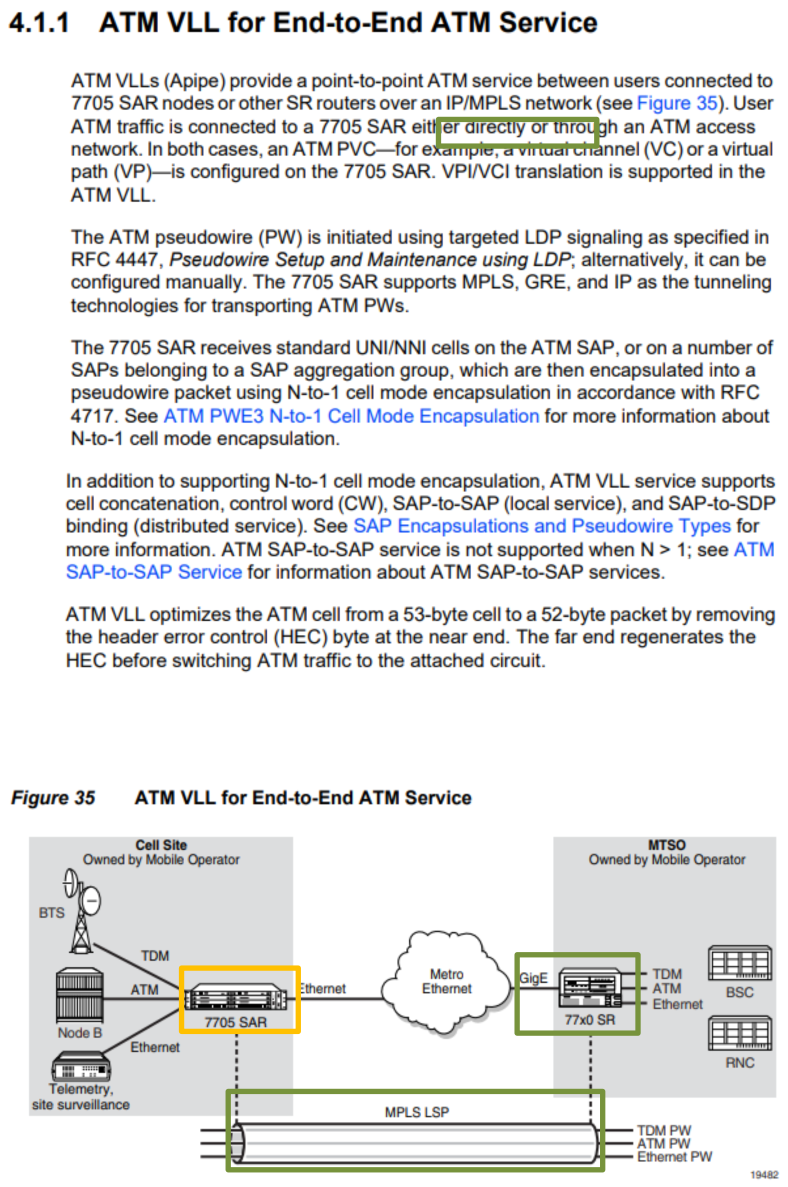

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

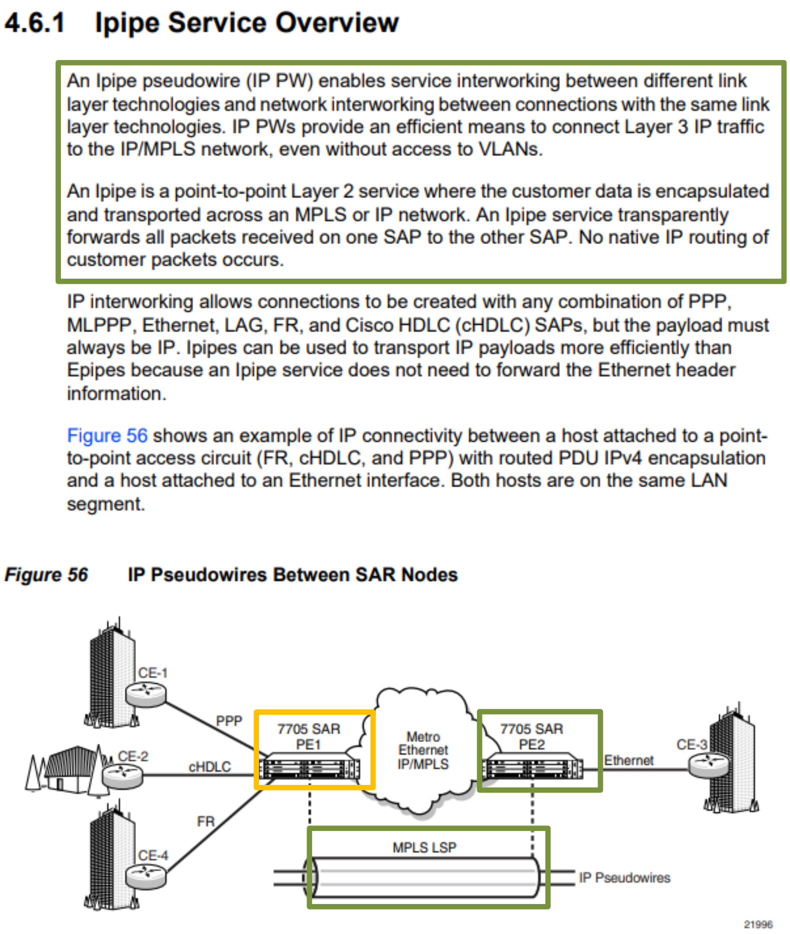

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

33. The 7705 Service Aggregation Router comprises edge devices that comprise one or more native interfaces, for communicating with client nodes in accordance with respective native Layer 2 protocols, at least one of which is different from the packet-oriented Layer 2 communication protocol:

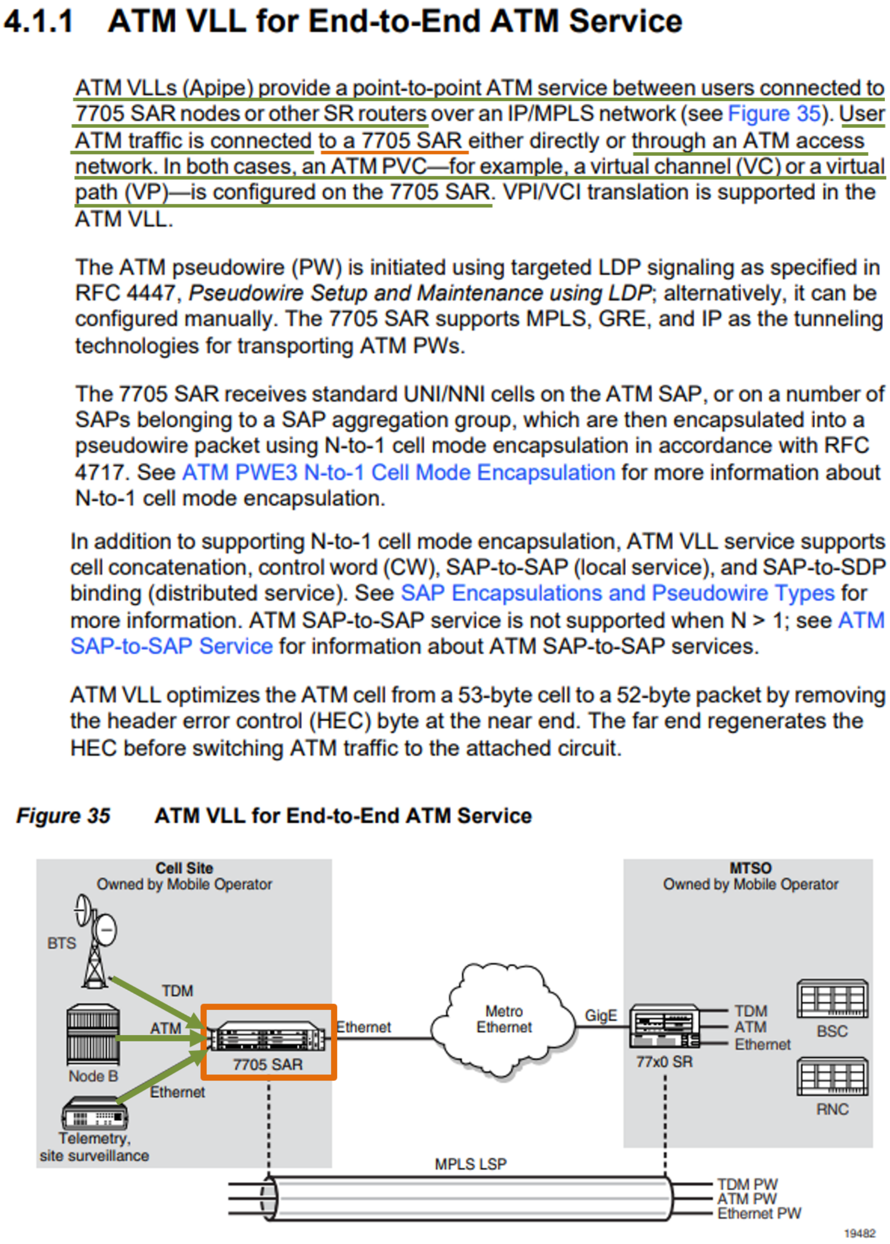

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

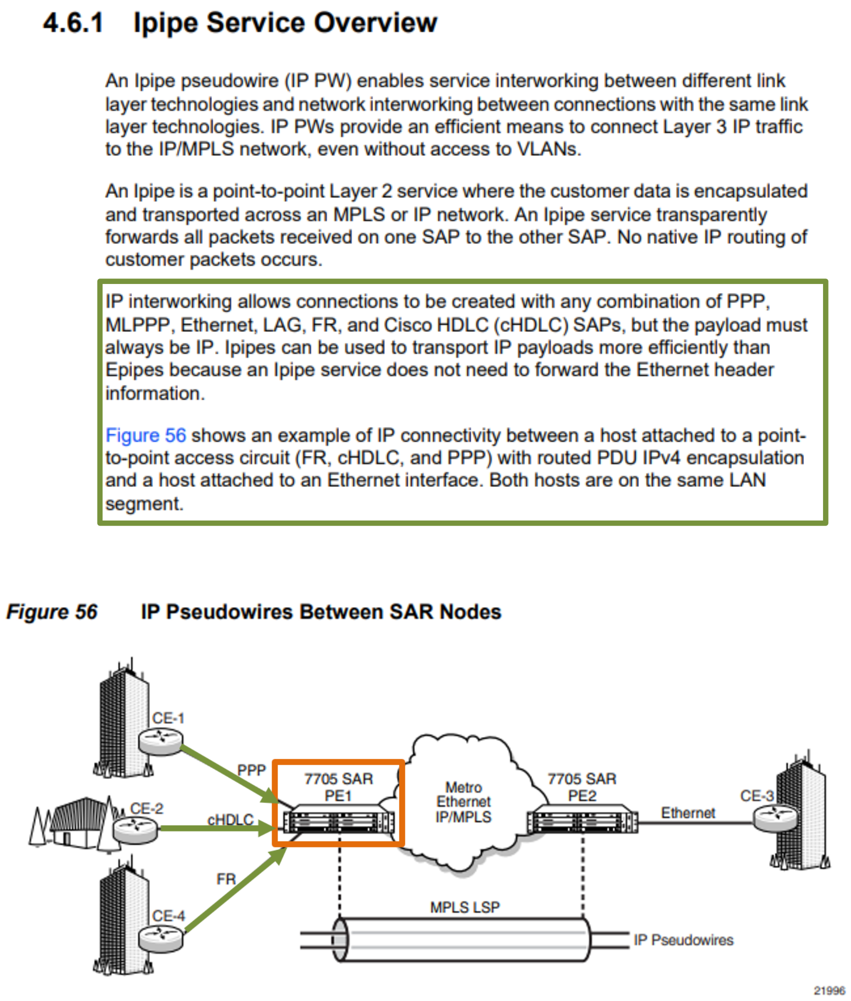

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267 of PDF)

34. The 7705 Service Aggregation Router comprises edge devices that comprise a protocol converter, which is configured to convert the data frames received on the one or more native interfaces from at least a first format specified by the native Layer 2 protocols to a second format specified by the packet-oriented Layer 2 communication protocol, so as to transmit the data frames in the second format via the at least one network port, and to convert the data frames received on the at least one network port from the second format to at least the first format, so as to transmit the data frames in at least the first format via the one or more native interfaces:

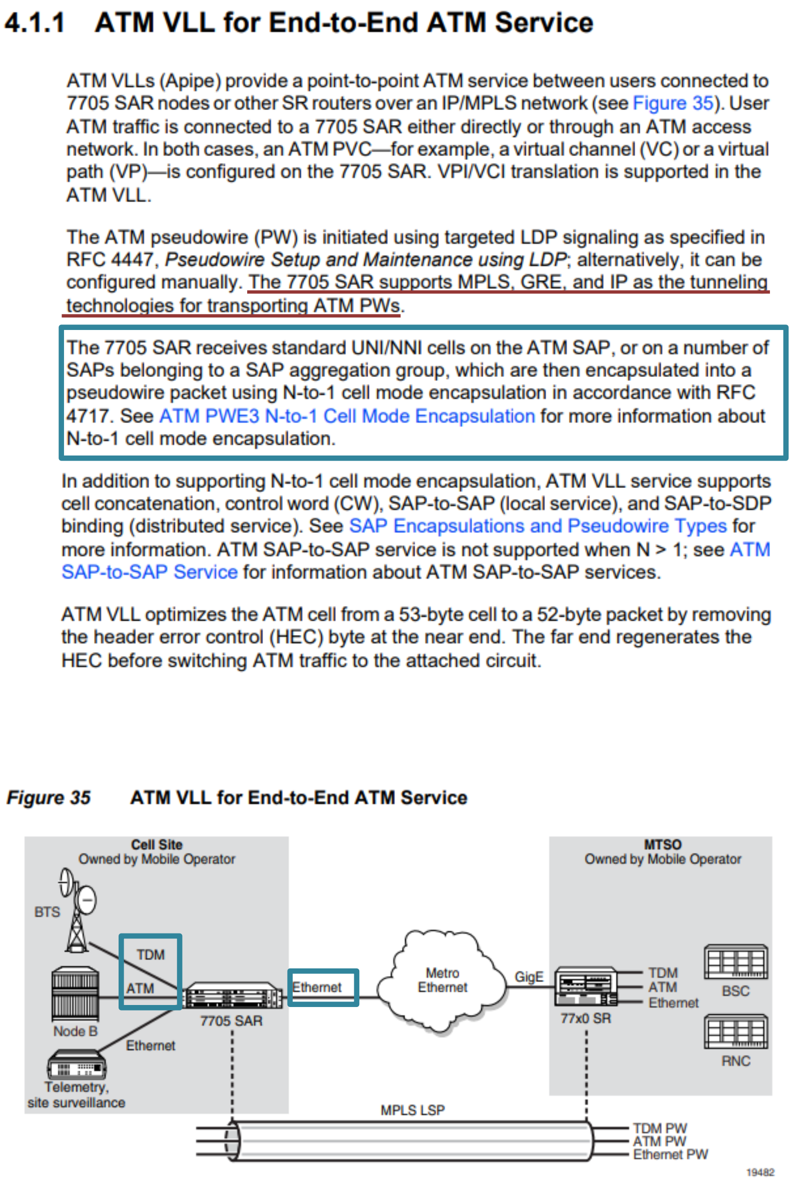

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

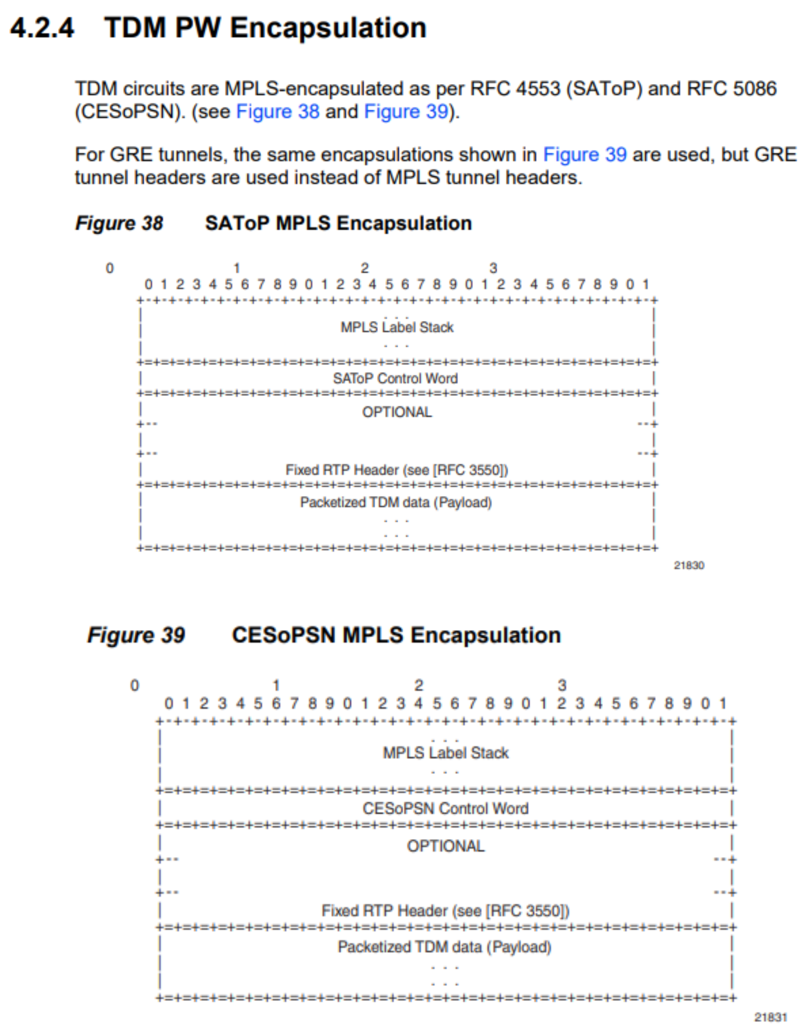

4.2.4 TDM PW Encapsulation

TDM circuits are MPLS-encapsulated as per RFC 4553 (SAToP) and RFC 5086 (CESoPSN). (see Figure 38 and Figure 39).

For GRE tunnels, the same encapsulations shown in Figure 39 are used, but GRE tunnel headers are used instead of MPLS tunnel headers.

Figure 38 **SAToP MPLS Encapsulation**

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                             . . .                             |
 |                        MPLS Label Stack                       |
 |                             . . .                             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                       SAToP Control Word                      |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                            OPTIONAL                           |
 +--                                                           --+
 |                                                               |
 +--                                                           --+
 |                 Fixed RTP Header (see [RFC 3550])             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                   Packetized TDM data (Payload)               |
 |                             . . .                             |
 |                             . . .                             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
                                                                     21830
```

Figure 39 **CESoPSN MPLS Encapsulation**

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                             . . .                             |
 |                        MPLS Label Stack                       |
 |                             . . .                             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                      CESoPSN Control Word                     |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                            OPTIONAL                           |
 +--                                                           --+
 |                                                               |
 +--                                                           --+
 |                 Fixed RTP Header (see [RFC 3550])             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
 |                   Packetized TDM data (Payload)               |
 |                             . . .                             |
 |                             . . .                             |
 +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
                                                                     21831
```

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 213-14)

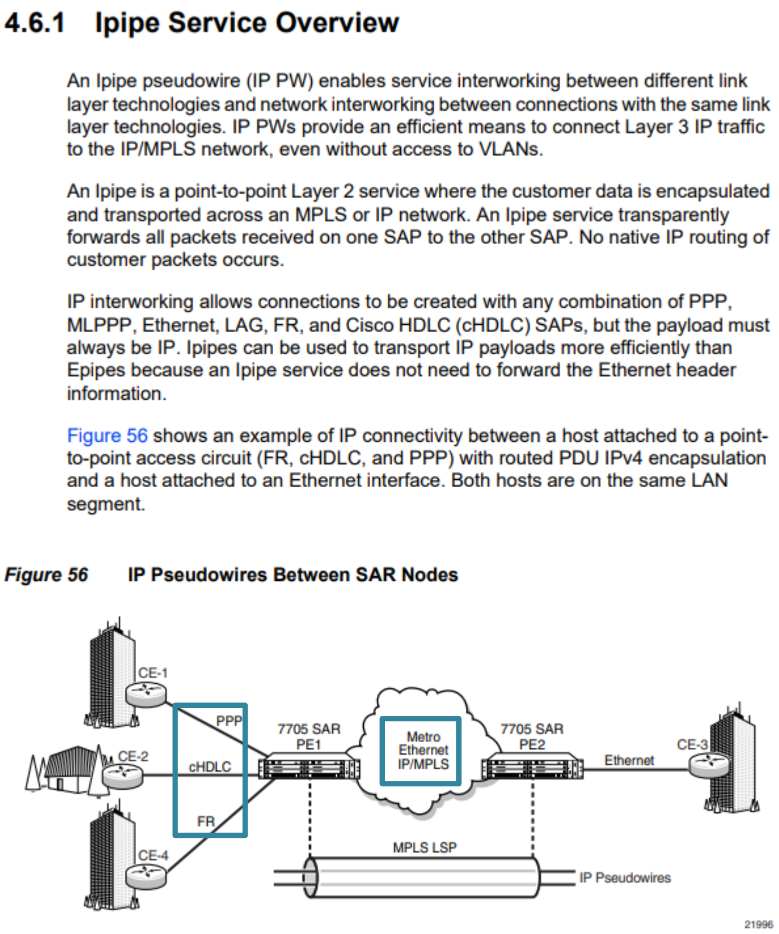

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

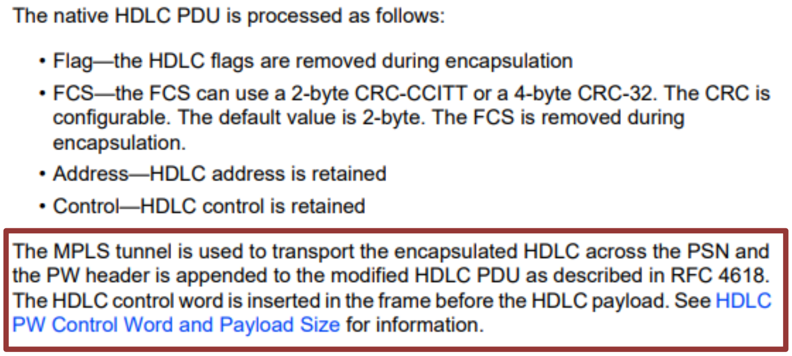

The native HDLC PDU is processed as follows:

- Flag—the HDLC flags are removed during encapsulation
- FCS—the FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The default value is 2-byte. The FCS is removed during encapsulation.
- Address—HDLC address is retained
- Control—HDLC control is retained

The MPLS tunnel is used to transport the encapsulated HDLC across the PSN and the PW header is appended to the modified HDLC PDU as described in RFC 4618. The HDLC control word is inserted in the frame before the HDLC payload. See HDLC PW Control Word and Payload Size for information.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 263)

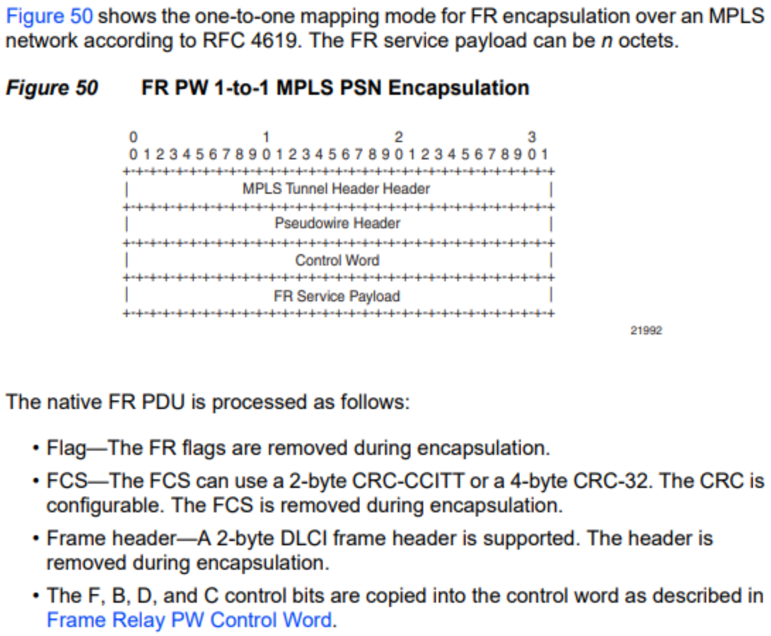

Figure 50 shows the one-to-one mapping mode for FR encapsulation over an MPLS network according to RFC 4619. The FR service payload can be *n* octets.

Figure 50 **FR PW 1-to-1 MPLS PSN Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  MPLS Tunnel Header Header                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudowire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Control Word                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      FR Service Payload                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

21992

The native FR PDU is processed as follows:

- Flag—The FR flags are removed during encapsulation.
- FCS—The FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The FCS is removed during encapsulation.
- Frame header—A 2-byte DLCI frame header is supported. The header is removed during encapsulation.
- The F, B, D, and C control bits are copied into the control word as described in Frame Relay PW Control Word.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 258)

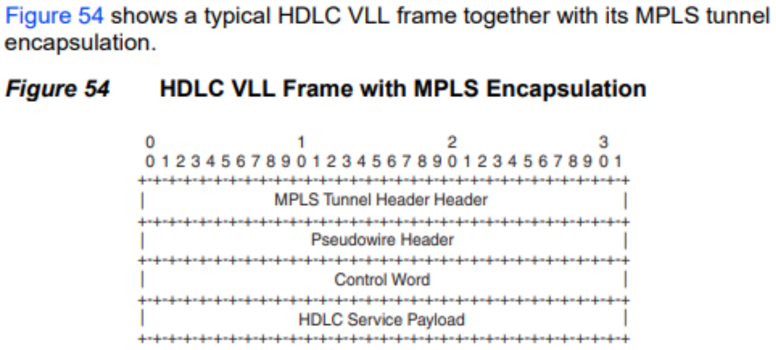

Figure 54 shows a typical HDLC VLL frame together with its MPLS tunnel encapsulation.

Figure 54 **HDLC VLL Frame with MPLS Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  MPLS Tunnel Header Header                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudowire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Control Word                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     HDLC Service Payload                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 262-63)

35. The 7705 Service Aggregation Router comprises edge devices configured to direct the data frames received from two or more of the native interfaces to one of the ports of the hub, and to map the two or more of the native interfaces to different, respective Virtual Local Area Networks (VLANs) on the network, such that the at least one network port comprises an Ethernet

port, and such that the one or more native interfaces comprise at least one of a time domain multiplexed (TDM) interface and a serial interface

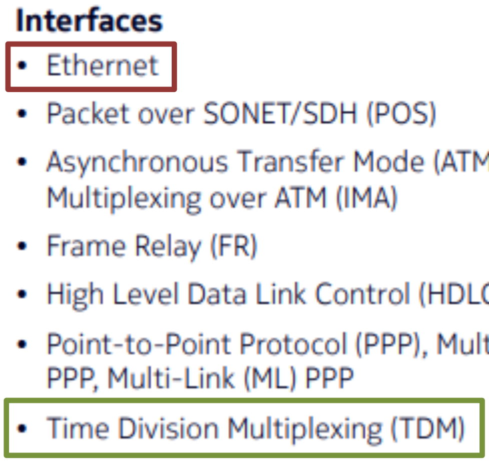

Interfaces

- Ethernet
- Packet over SONET/SDH (POS)
- Asynchronous Transfer Mode (ATM), ATM-Inverse Multiplexing over ATM (IMA)
- Frame Relay (FR)
- High Level Data Link Control (HDLC)
- Point-to-Point Protocol (PPP), Multi-Class (MC) PPP, Multi-Link (ML) PPP
- Time Division Multiplexing (TDM)

https://onestore.nokia.com/asset/f/162833 (page 3)

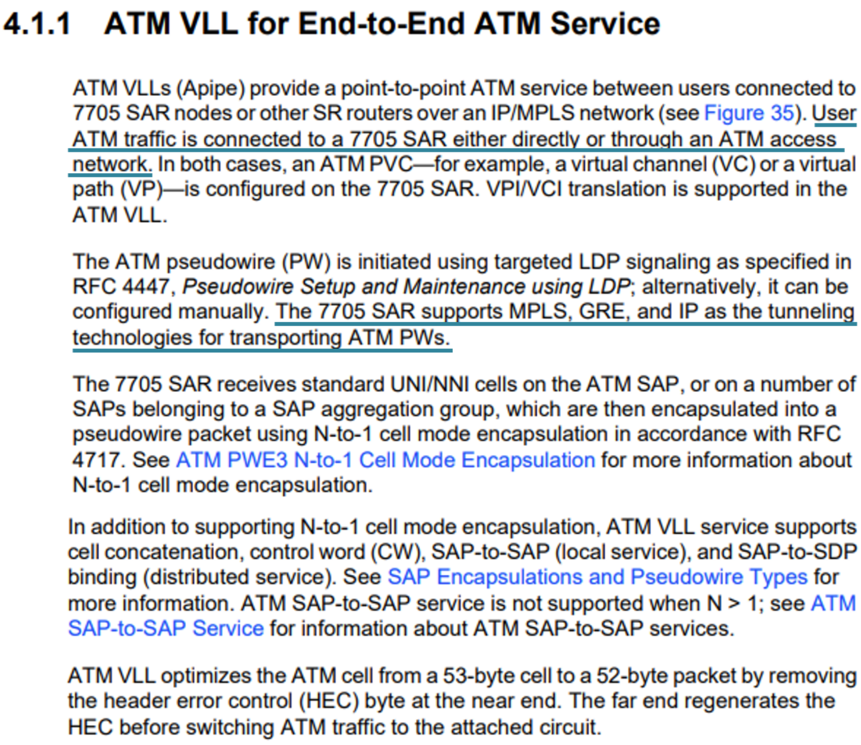

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

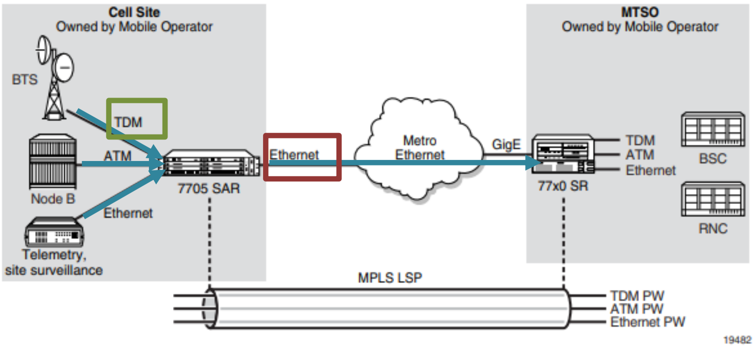

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

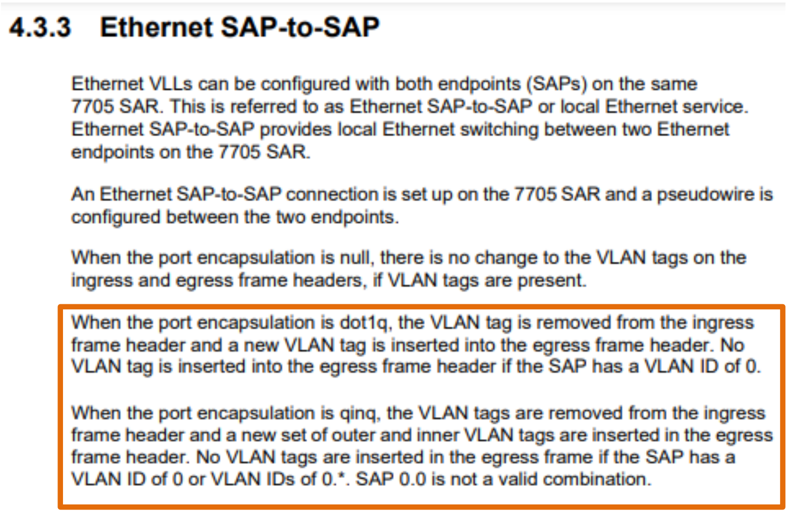

4.3.3 Ethernet SAP-to-SAP

Ethernet VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR. This is referred to as Ethernet SAP-to-SAP or local Ethernet service. Ethernet SAP-to-SAP provides local Ethernet switching between two Ethernet endpoints on the 7705 SAR.

An Ethernet SAP-to-SAP connection is set up on the 7705 SAR and a pseudowire is configured between the two endpoints.

When the port encapsulation is null, there is no change to the VLAN tags on the ingress and egress frame headers, if VLAN tags are present.

When the port encapsulation is dot1q, the VLAN tag is removed from the ingress frame header and a new VLAN tag is inserted into the egress frame header. No VLAN tag is inserted into the egress frame header if the SAP has a VLAN ID of 0.

When the port encapsulation is qinq, the VLAN tags are removed from the ingress frame header and a new set of outer and inner VLAN tags are inserted in the egress frame header. No VLAN tags are inserted in the egress frame if the SAP has a VLAN ID of 0 or VLAN IDs of 0.*. SAP 0.0 is not a valid combination.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 235-36)

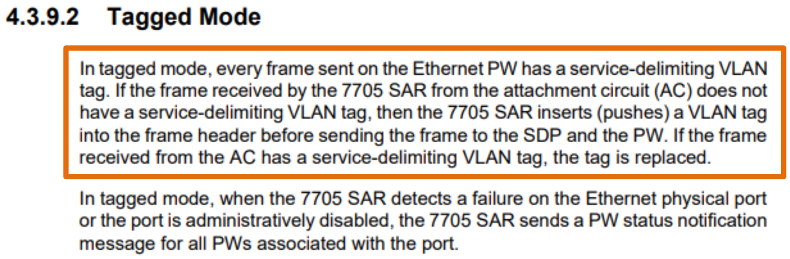

4.3.9.2 Tagged Mode

In tagged mode, every frame sent on the Ethernet PW has a service-delimiting VLAN tag. If the frame received by the 7705 SAR from the attachment circuit (AC) does not have a service-delimiting VLAN tag, then the 7705 SAR inserts (pushes) a VLAN tag into the frame header before sending the frame to the SDP and the PW. If the frame received from the AC has a service-delimiting VLAN tag, the tag is replaced.

In tagged mode, when the 7705 SAR detects a failure on the Ethernet physical port or the port is administratively disabled, the 7705 SAR sends a PW status notification message for all PWs associated with the port.

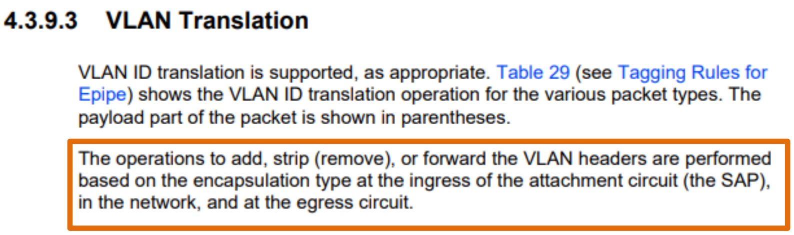

4.3.9.3 VLAN Translation

VLAN ID translation is supported, as appropriate. Table 29 (see Tagging Rules for Epipe) shows the VLAN ID translation operation for the various packet types. The payload part of the packet is shown in parentheses.

The operations to add, strip (remove), or forward the VLAN headers are performed based on the encapsulation type at the ingress of the attachment circuit (the SAP), in the network, and at the egress circuit.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 243-44)

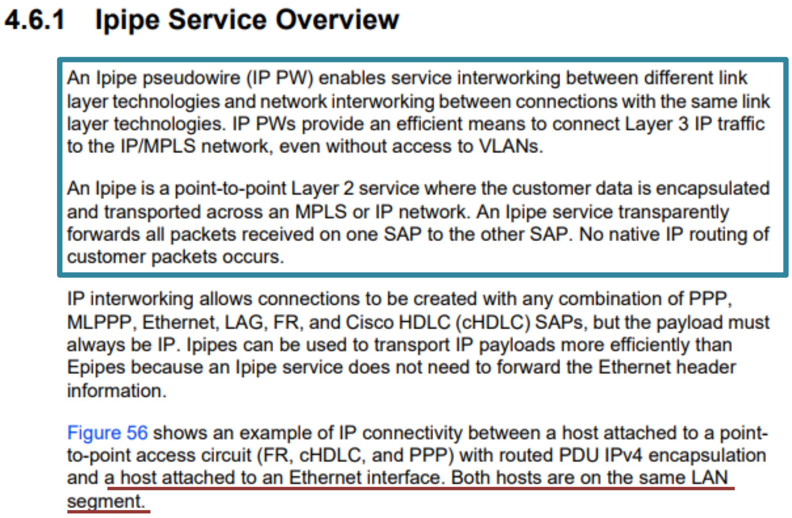

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

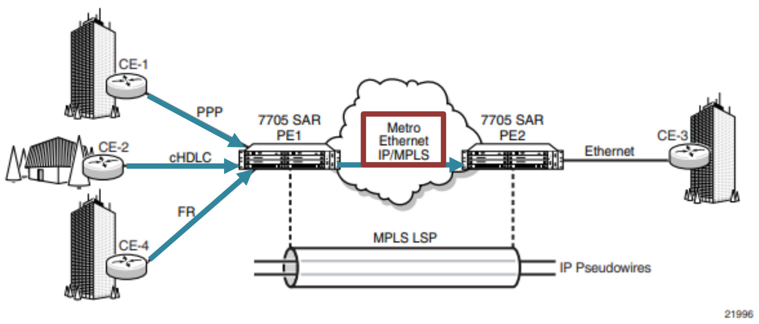

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

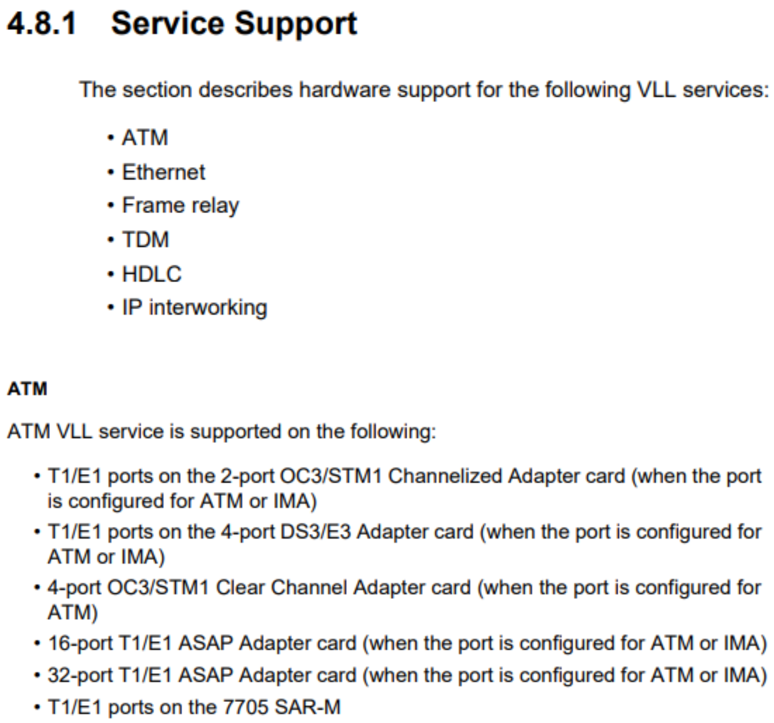

4.8.1 Service Support

The section describes hardware support for the following VLL services:

- ATM
- Ethernet
- Frame relay
- TDM
- HDLC
- IP interworking

ATM

ATM VLL service is supported on the following:

- T1/E1 ports on the 2-port OC3/STM1 Channelized Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 4-port DS3/E3 Adapter card (when the port is configured for ATM or IMA)
- 4-port OC3/STM1 Clear Channel Adapter card (when the port is configured for ATM)
- 16-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- 32-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 7705 SAR-M

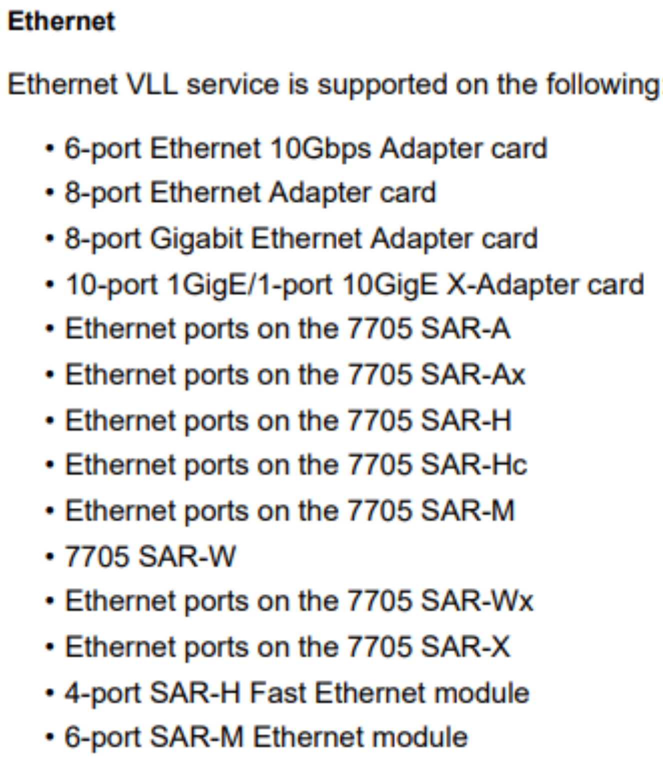

Ethernet

Ethernet VLL service is supported on the following:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Ethernet ports on the 7705 SAR-A
- Ethernet ports on the 7705 SAR-Ax
- Ethernet ports on the 7705 SAR-H
- Ethernet ports on the 7705 SAR-Hc
- Ethernet ports on the 7705 SAR-M
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx
- Ethernet ports on the 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module

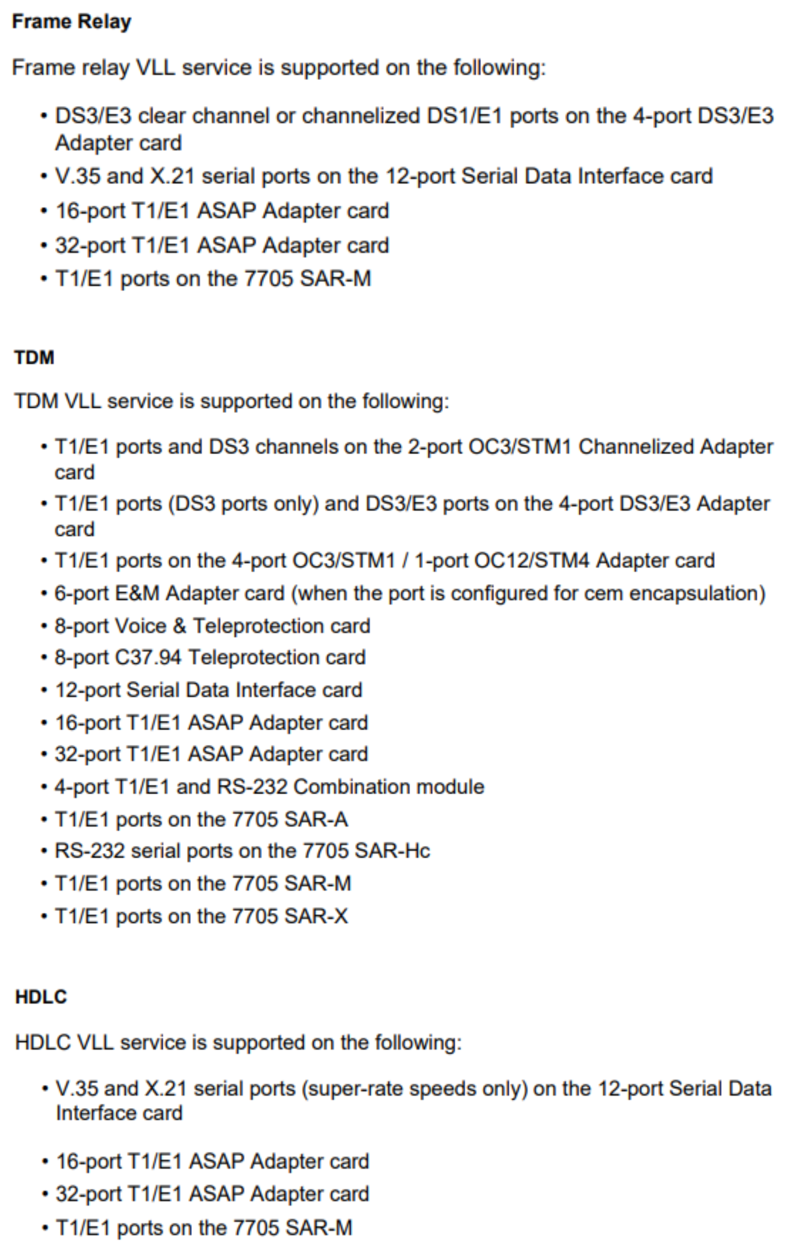

Frame Relay

Frame relay VLL service is supported on the following:

- DS3/E3 clear channel or channelized DS1/E1 ports on the 4-port DS3/E3 Adapter card
- V.35 and X.21 serial ports on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

TDM

TDM VLL service is supported on the following:

- T1/E1 ports and DS3 channels on the 2-port OC3/STM1 Channelized Adapter card
- T1/E1 ports (DS3 ports only) and DS3/E3 ports on the 4-port DS3/E3 Adapter card
- T1/E1 ports on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- 6-port E&M Adapter card (when the port is configured for cem encapsulation)
- 8-port Voice & Teleprotection card
- 8-port C37.94 Teleprotection card
- 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 4-port T1/E1 and RS-232 Combination module
- T1/E1 ports on the 7705 SAR-A
- RS-232 serial ports on the 7705 SAR-Hc
- T1/E1 ports on the 7705 SAR-M
- T1/E1 ports on the 7705 SAR-X

HDLC

HDLC VLL service is supported on the following:

- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

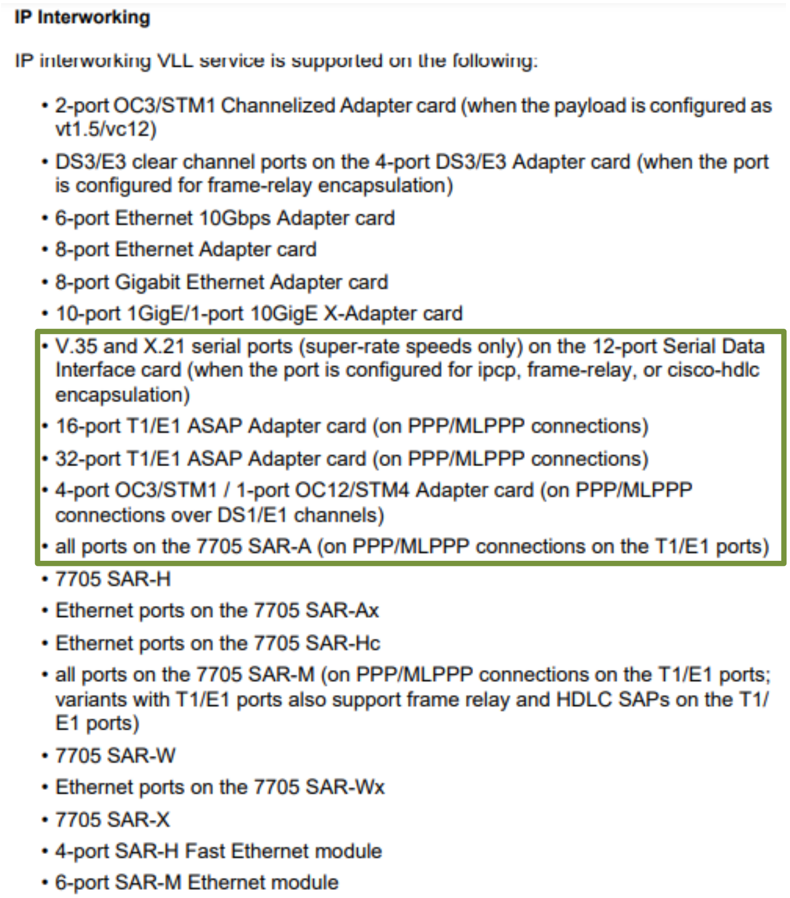

IP Interworking

IP interworking VLL service is supported on the following:

- 2-port OC3/STM1 Channelized Adapter card (when the payload is configured as vt1.5/vc12)
- DS3/E3 clear channel ports on the 4-port DS3/E3 Adapter card (when the port is configured for frame-relay encapsulation)
- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card (when the port is configured for ipcp, frame-relay, or cisco-hdlc encapsulation)
- 16-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- 32-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card (on PPP/MLPPP connections over DS1/E1 channels)
- all ports on the 7705 SAR-A (on PPP/MLPPP connections on the T1/E1 ports)
- 7705 SAR-H
- Ethernet ports on the 7705 SAR-Ax
- Ethernet ports on the 7705 SAR-Hc
- all ports on the 7705 SAR-M (on PPP/MLPPP connections on the T1/E1 ports; variants with T1/E1 ports also support frame relay and HDLC SAPs on the T1/E1 ports)
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx
- 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 289-292)

Willful Infringement

36. Defendants have had actual knowledge of the ’010 Patent and its infringement thereof at least as of service or other receipt of Plaintiff’s Complaint.

37. Defendants’ infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendants.

38. Notwithstanding this knowledge, Defendants have knowingly or with reckless disregard infringed the ’010 Patent. Defendants continue to commit acts of infringement despite

being on notice of an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

39. Defendants are therefore liable for willful infringement. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

40. Defendants have induced and are knowingly inducing its distributors, testers, trainers, customers and/or end users to directly infringe the '010 Patent, with the specific intent to induce acts constituting infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

41. Defendants have knowingly contributed to direct infringement by its customers and end users by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the accused products which are not suitable for substantial non-infringing use, and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

42. Defendants' indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe the '010 Patent.

43. Defendants' indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendants' indirect infringement further includes providing application notes instructing its customers on infringing uses of the '010 Accused Products. The '010 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '010 Patent, either literally or equivalently. Defendants

know and intend that customers who purchase the ’010 Accused Products will use those products for their intended purpose. For example, Defendants’ United States website, https://www.nokia.com, instructs customers to use the ’010 Accused Products in numerous infringing applications. Furthermore, Defendants provide instructions and other resources via its “Doc Center” (https://documentation.nokia.com), “Nokia Service Routing Certification” and training courses from its “NokiaEDU Training Centers” (https://www.nokia.com/networks/training/src/courses/#open-enrollment), and elsewhere on using the ’010 Accused Products. Defendants’ customers directly infringe the ’010 Patent when they follow Defendants’ provided instructions on websites, videos, trainings, and elsewhere. Defendants’ customers who follow Defendants’ provided instructions directly infringe claims of the ’010 Patent.

44. In addition, Defendants specifically intend that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendants’ infringing products. Defendants know following its instructions directly infringes claims of the ’010 Patent, including, for example Claim 14.

45. Defendants’ customers who follow Defendants’ provided instructions directly infringe the method of Claim 14 of the ’010 Patent.

46. Defendants instruct its customers to use the 7705 Service Aggregation Router in a method for data communications.

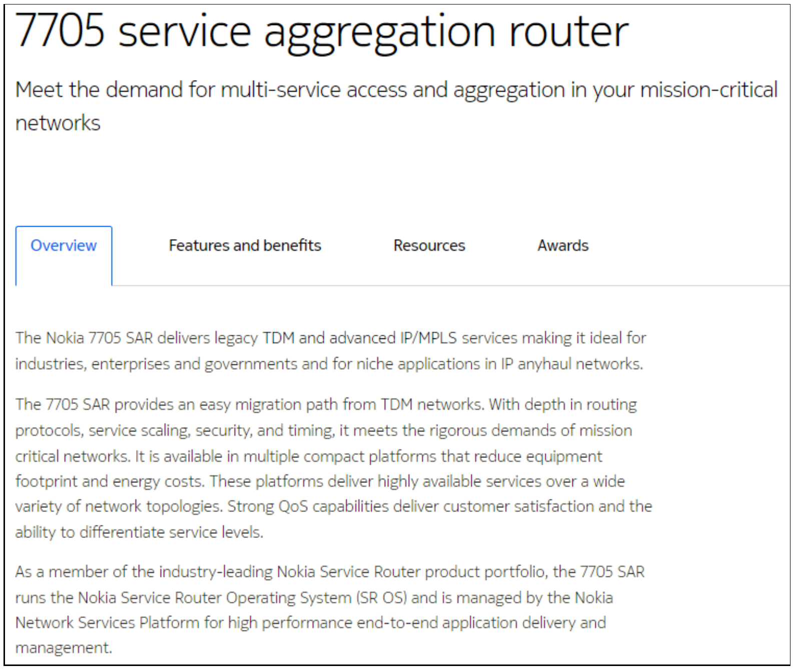

7705 service aggregation router

Meet the demand for multi-service access and aggregation in your mission-critical networks

The Nokia 7705 SAR delivers legacy TDM and advanced IP/MPLS services making it ideal for industries, enterprises and governments and for niche applications in IP anyhaul networks.

The 7705 SAR provides an easy migration path from TDM networks. With depth in routing protocols, service scaling, security, and timing, it meets the rigorous demands of mission critical networks. It is available in multiple compact platforms that reduce equipment footprint and energy costs. These platforms deliver highly available services over a wide variety of network topologies. Strong QoS capabilities deliver customer satisfaction and the ability to differentiate service levels.

As a member of the industry-leading Nokia Service Router product portfolio, the 7705 SAR runs the Nokia Service Router Operating System (SR OS) and is managed by the Nokia Network Services Platform for high performance end-to-end application delivery and management.

https://www.nokia.com/networks/products/7705-service-aggregation-router/

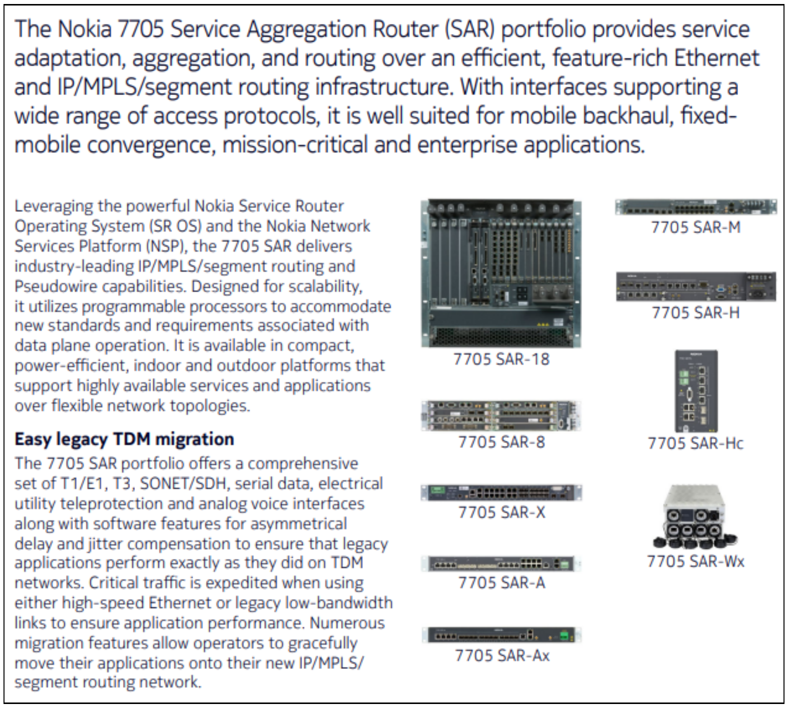

The Nokia 7705 Service Aggregation Router (SAR) portfolio provides service adaptation, aggregation, and routing over an efficient, feature-rich Ethernet and IP/MPLS/segment routing infrastructure. With interfaces supporting a wide range of access protocols, it is well suited for mobile backhaul, fixed-mobile convergence, mission-critical and enterprise applications.

Leveraging the powerful Nokia Service Router Operating System (SR OS) and the Nokia Network Services Platform (NSP), the 7705 SAR delivers industry-leading IP/MPLS/segment routing and Pseudowire capabilities. Designed for scalability, it utilizes programmable processors to accommodate new standards and requirements associated with data plane operation. It is available in compact, power-efficient, indoor and outdoor platforms that support highly available services and applications over flexible network topologies.

Easy legacy TDM migration

The 7705 SAR portfolio offers a comprehensive set of T1/E1, T3, SONET/SDH, serial data, electrical utility teleprotection and analog voice interfaces along with software features for asymmetrical delay and jitter compensation to ensure that legacy applications perform exactly as they did on TDM networks. Critical traffic is expedited when using either high-speed Ethernet or legacy low-bandwidth links to ensure application performance. Numerous migration features allow operators to gracefully move their applications onto their new IP/MPLS/segment routing network.

https://onestore.nokia.com/asset/f/162833 (page 1)

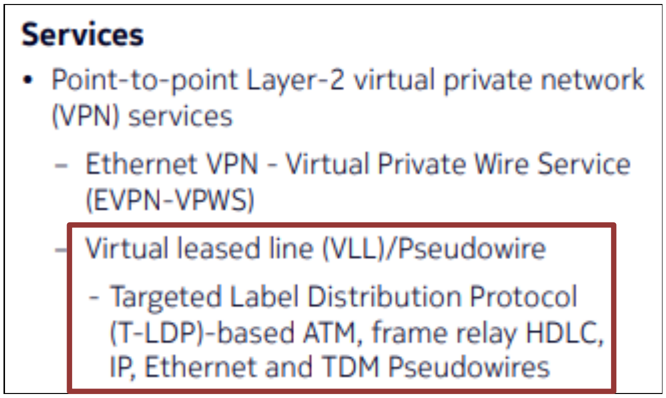

Services

- Point-to-point Layer-2 virtual private network (VPN) services
  - Ethernet VPN - Virtual Private Wire Service (EVPN-VPWS)
  - Virtual leased line (VLL)/Pseudowire
    - Targeted Label Distribution Protocol (T-LDP)-based ATM, frame relay HDLC, IP, Ethernet and TDM Pseudowires

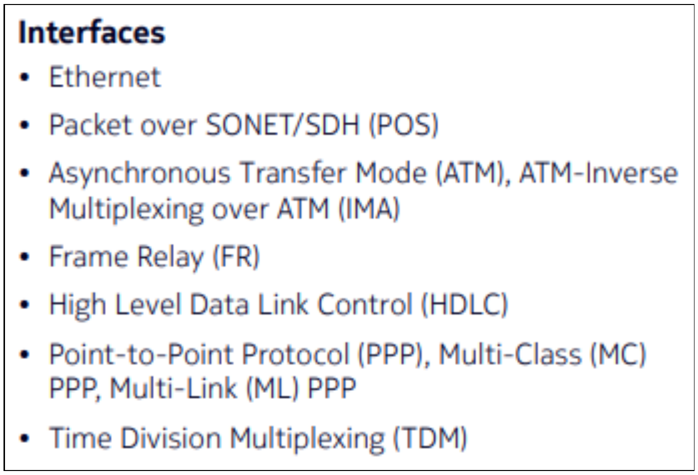

Interfaces

- Ethernet
- Packet over SONET/SDH (POS)
- Asynchronous Transfer Mode (ATM), ATM-Inverse Multiplexing over ATM (IMA)
- Frame Relay (FR)
- High Level Data Link Control (HDLC)
- Point-to-Point Protocol (PPP), Multi-Class (MC) PPP, Multi-Link (ML) PPP
- Time Division Multiplexing (TDM)

https://onestore.nokia.com/asset/f/162833 (page 3)

47. Defendants instruct its customers to use the 7705 Service Aggregation Router in linking a plurality of edge devices to communicate with a hub via a network in accordance with a packet-oriented Layer 2 communication protocol.

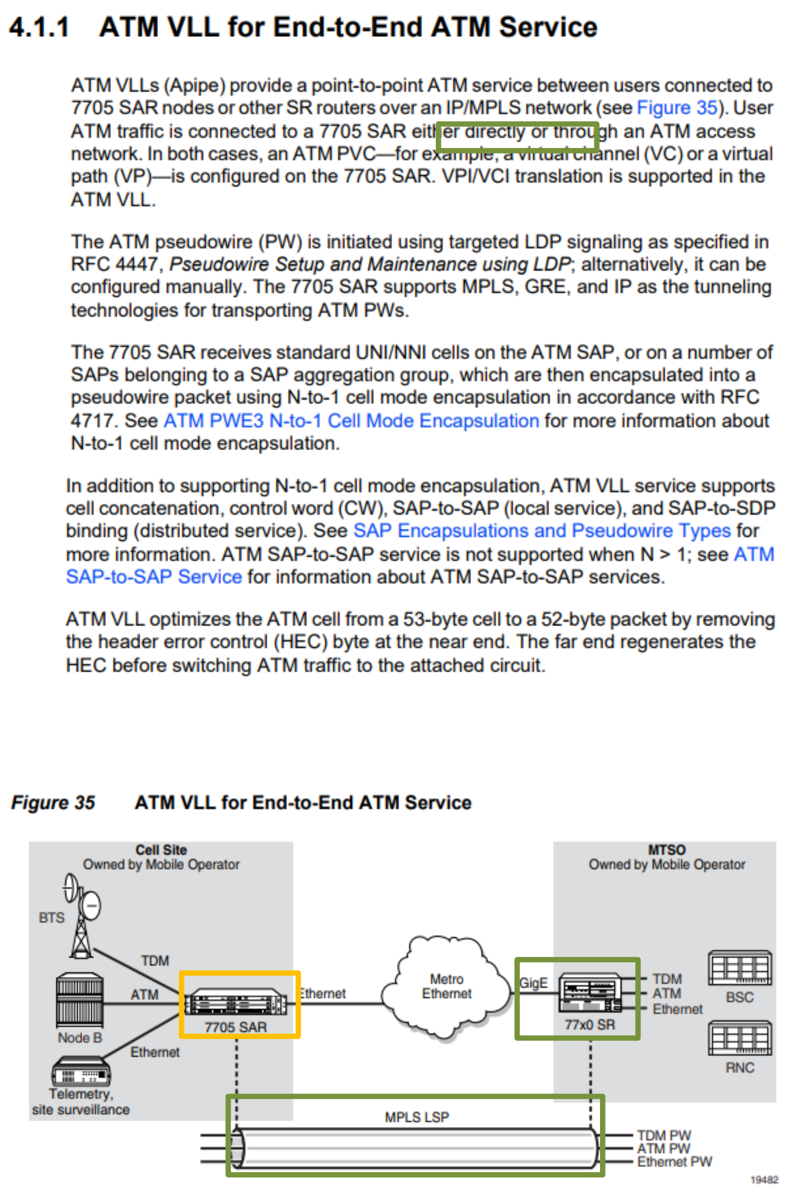

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

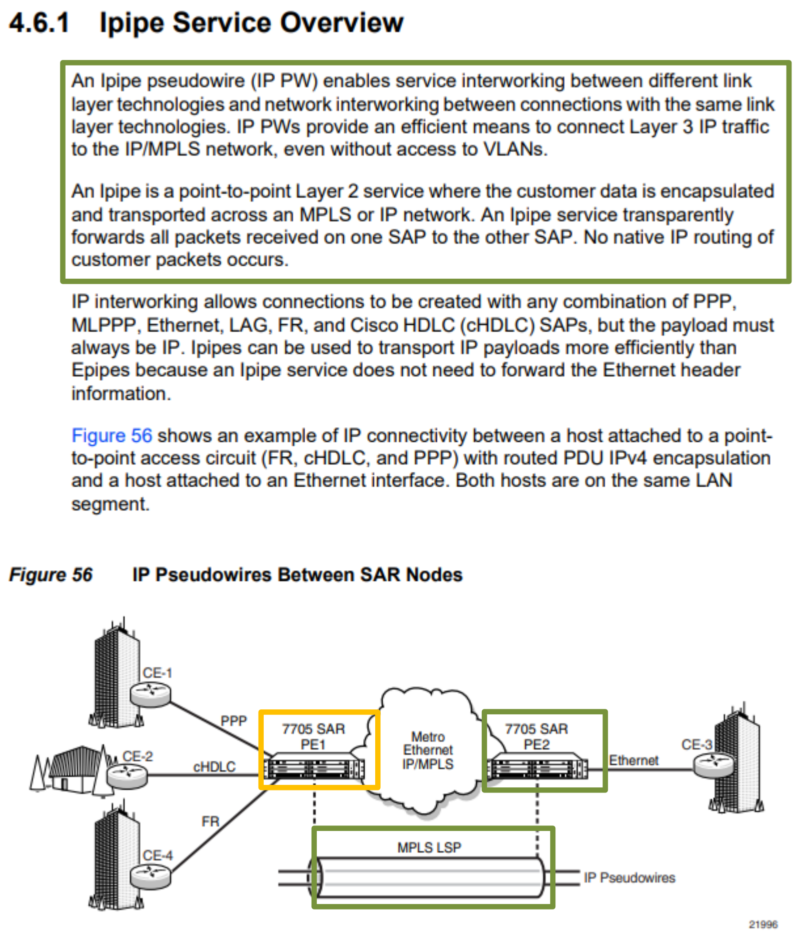

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

48. Defendants instruct its customers to use the 7705 Service Aggregation Router in, at each of the plurality of edge devices, receiving incoming data frames from client nodes in accordance with respective native Layer 2 protocols, at least one of which is different from the packet-oriented Layer 2 communication protocol.

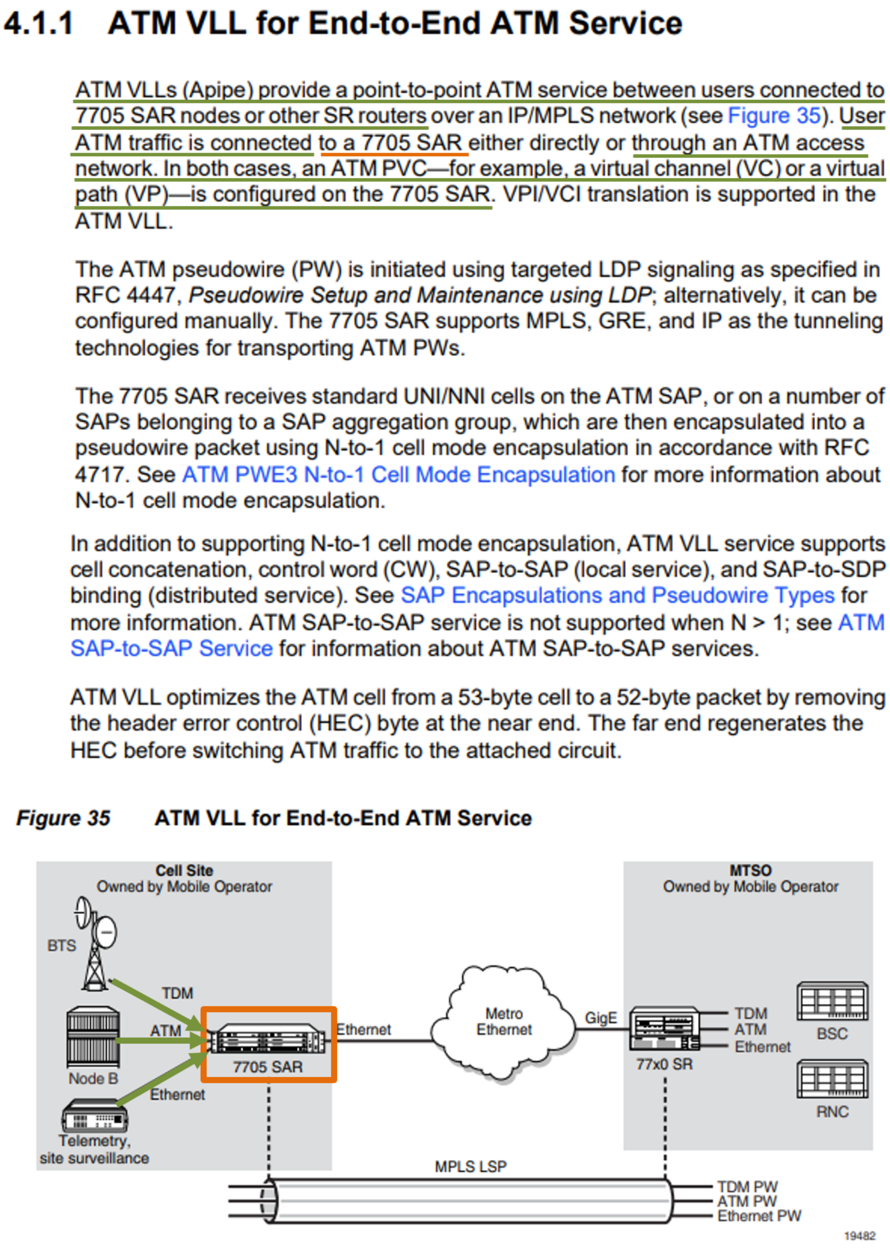

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

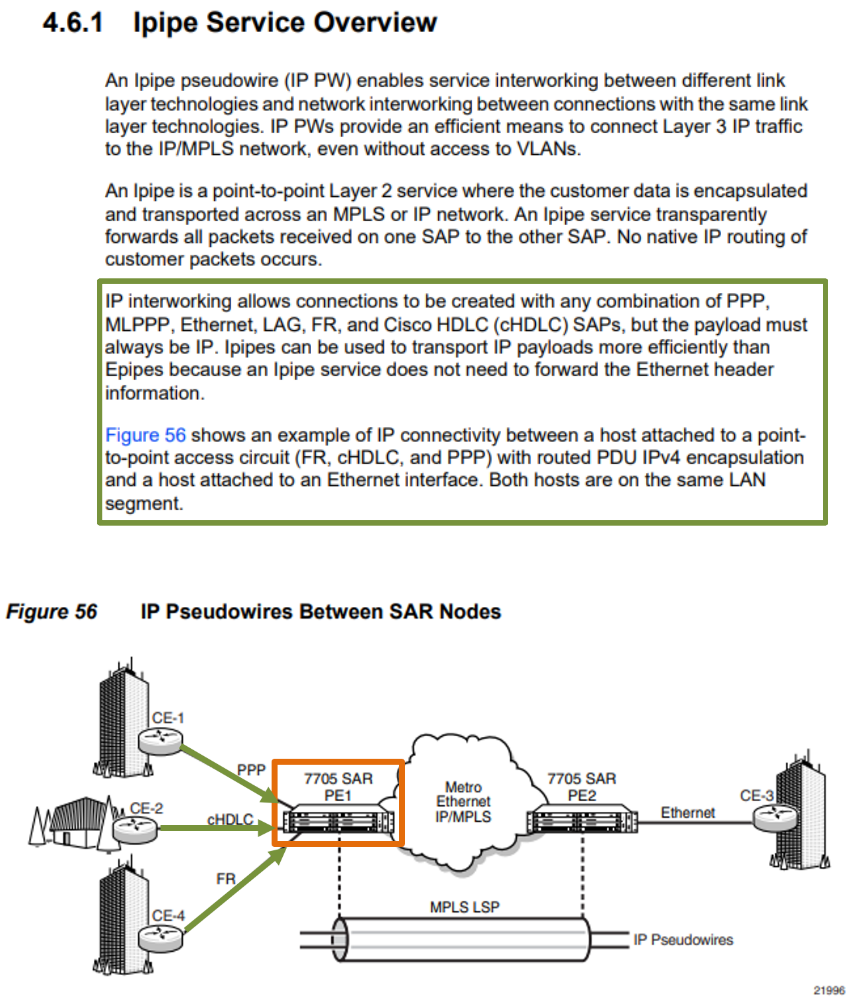

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267 of PDF)

49. Defendants instruct its customers to use the 7705 Service Aggregation Router in converting the received incoming data frames at each of the edge devices from at least a first format specified by the native Layer 2 protocols to a second format specified by the packet-oriented Layer 2 communication protocol.

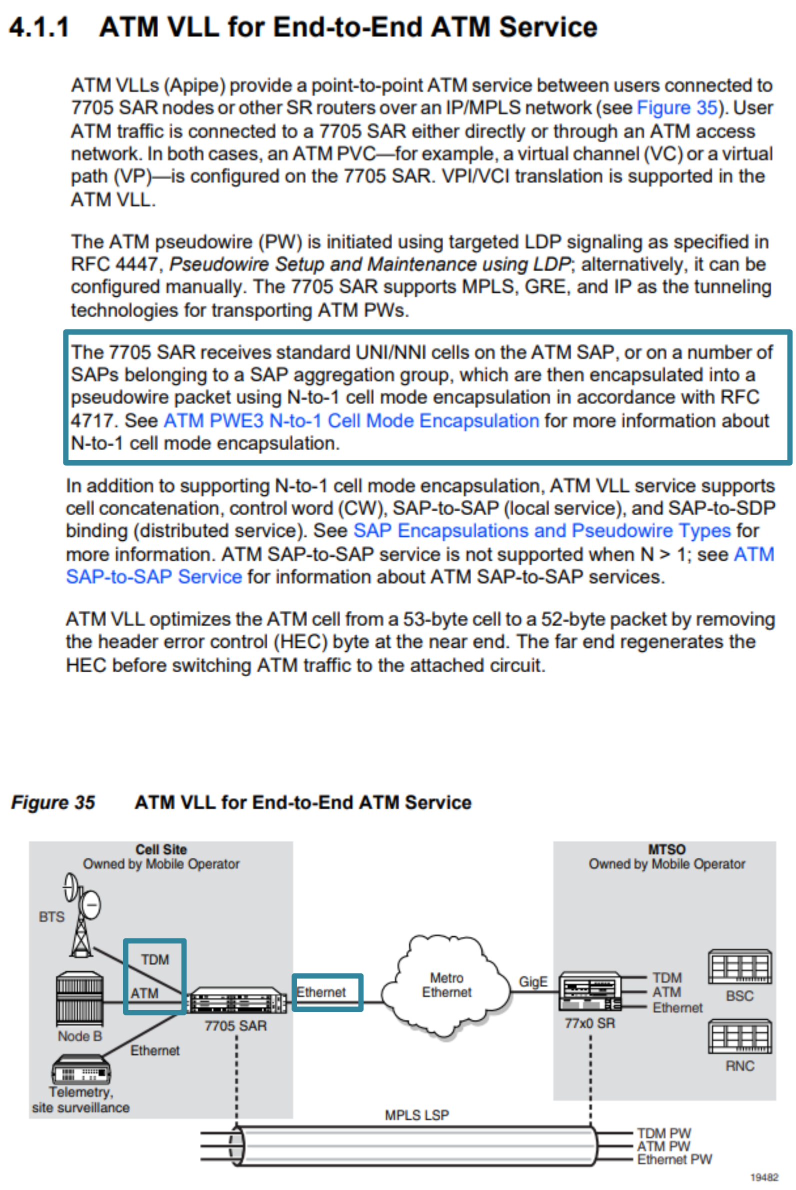

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

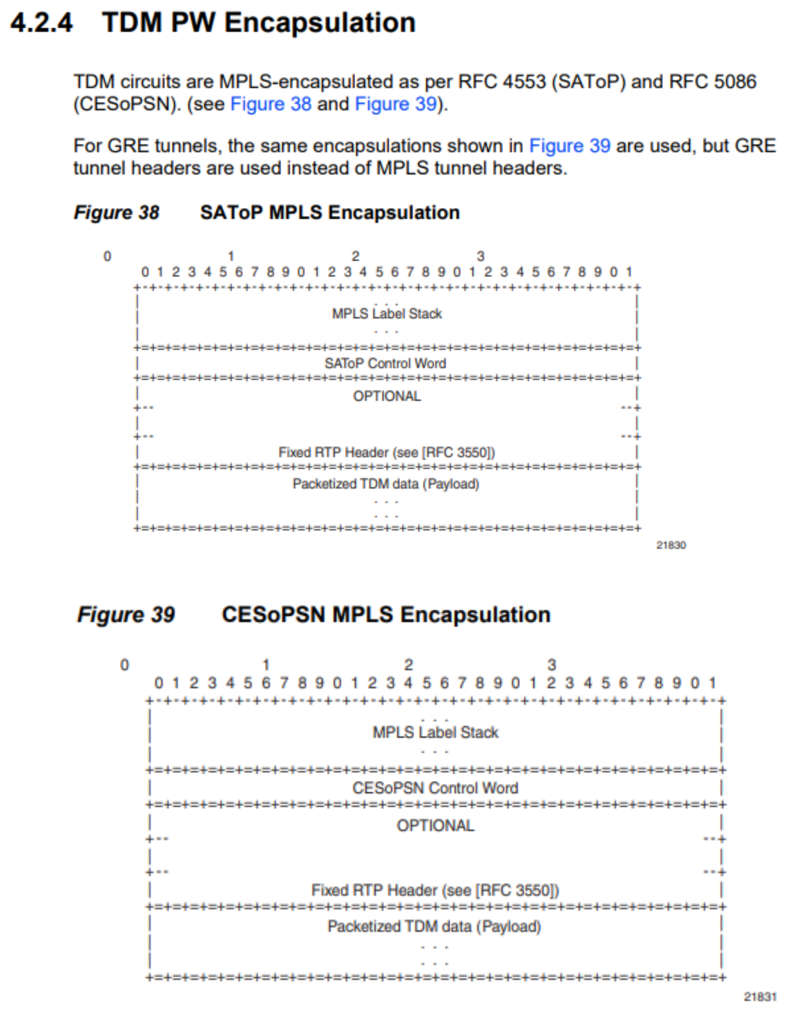

4.2.4 TDM PW Encapsulation

TDM circuits are MPLS-encapsulated as per RFC 4553 (SAToP) and RFC 5086 (CESoPSN). (see Figure 38 and Figure 39).

For GRE tunnels, the same encapsulations shown in Figure 39 are used, but GRE tunnel headers are used instead of MPLS tunnel headers.

Figure 38 **SAToP MPLS Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
|                        MPLS Label Stack                       |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                       SAToP Control Word                      |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                           OPTIONAL                            |
+--                                                           --+
|                                                               |
+--                                                           --+
|                Fixed RTP Header (see [RFC 3550])              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                Packetized TDM data (Payload)                  |
|                              ...                              |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
                                                             21830
```

Figure 39 **CESoPSN MPLS Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              ...                              |
|                        MPLS Label Stack                       |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                      CESoPSN Control Word                     |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                           OPTIONAL                            |
+--                                                           --+
|                                                               |
+--                                                           --+
|                Fixed RTP Header (see [RFC 3550])              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                Packetized TDM data (Payload)                  |
|                              ...                              |
|                              ...                              |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
                                                             21831
```

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 213-14)

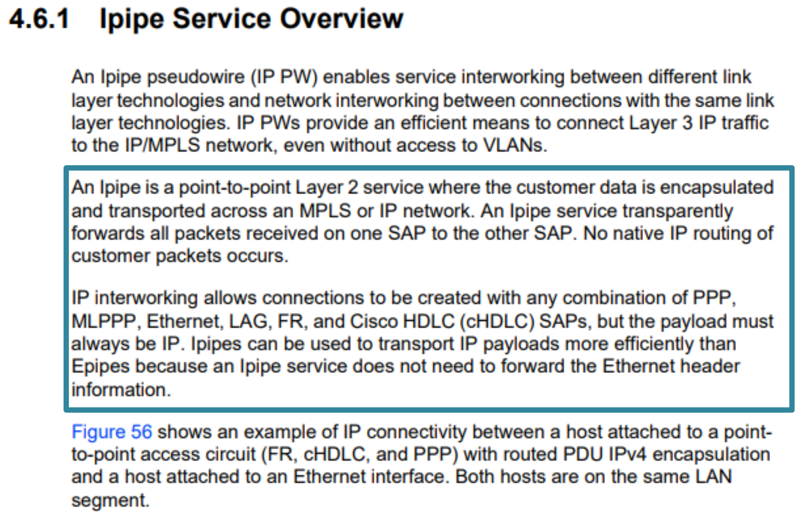

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

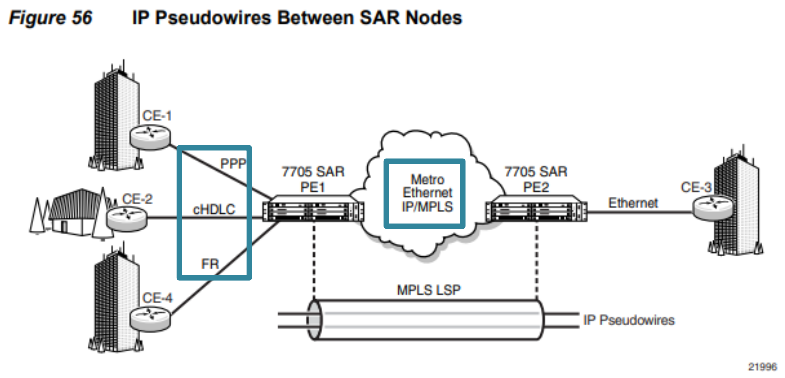

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

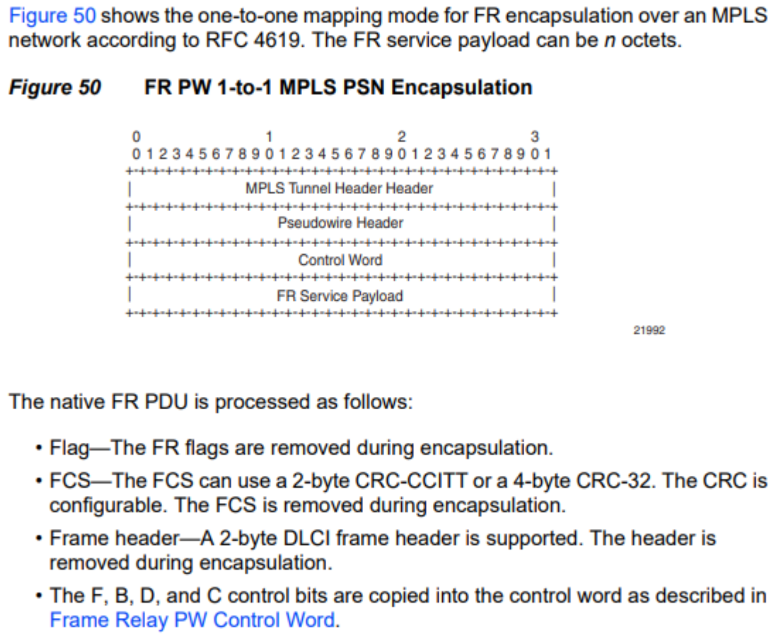

Figure 50 shows the one-to-one mapping mode for FR encapsulation over an MPLS network according to RFC 4619. The FR service payload can be *n* octets.

Figure 50 **FR PW 1-to-1 MPLS PSN Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                 MPLS Tunnel Header Header                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Pseudowire Header                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Control Word                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    FR Service Payload                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The native FR PDU is processed as follows:

- Flag—The FR flags are removed during encapsulation.
- FCS—The FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The FCS is removed during encapsulation.
- Frame header—A 2-byte DLCI frame header is supported. The header is removed during encapsulation.
- The F, B, D, and C control bits are copied into the control word as described in Frame Relay PW Control Word.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 258)

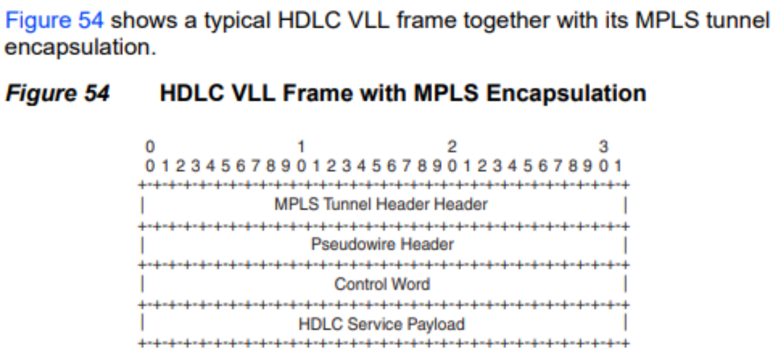

Figure 54 shows a typical HDLC VLL frame together with its MPLS tunnel encapsulation.

Figure 54 **HDLC VLL Frame with MPLS Encapsulation**

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  MPLS Tunnel Header Header                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Pseudowire Header                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Control Word                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     HDLC Service Payload                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 262-63)

50. Defendants instruct its customers to use the 7705 Service Aggregation Router in transmitting the incoming data frames in the second format via the network to the hub.

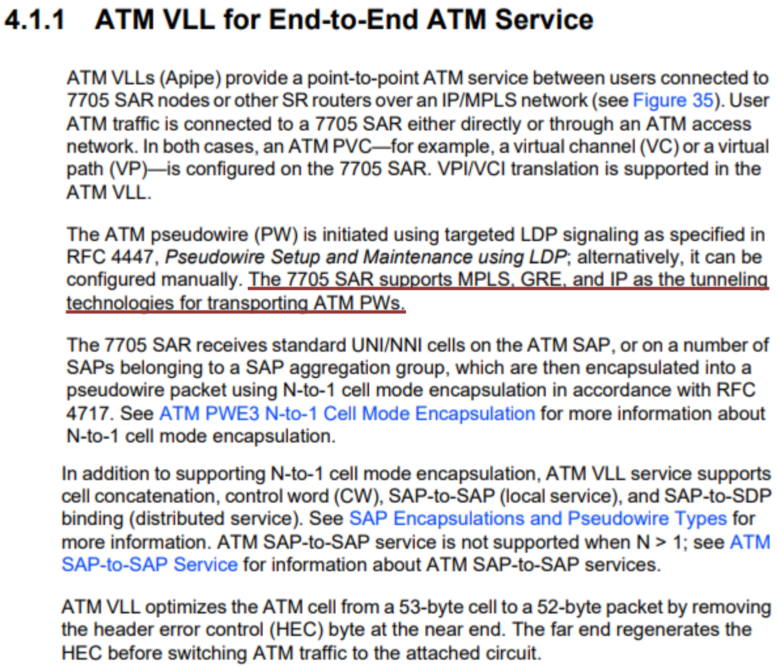

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

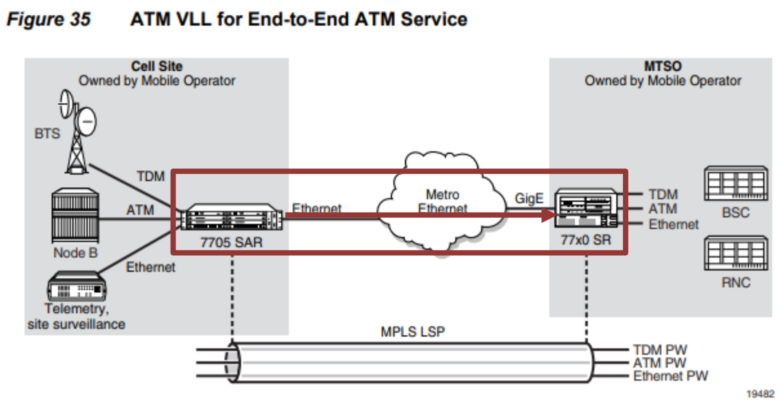

Figure 35 **ATM VLL for End-to-End ATM Service**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

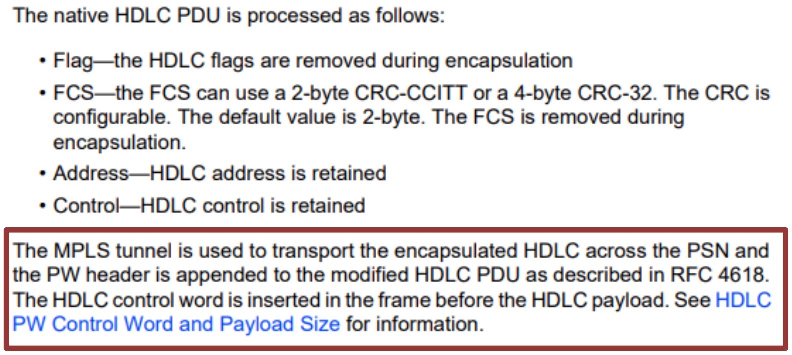

The native HDLC PDU is processed as follows:

- Flag—the HDLC flags are removed during encapsulation
- FCS—the FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The default value is 2-byte. The FCS is removed during encapsulation.
- Address—HDLC address is retained
- Control—HDLC control is retained

The MPLS tunnel is used to transport the encapsulated HDLC across the PSN and the PW header is appended to the modified HDLC PDU as described in RFC 4618. The HDLC control word is inserted in the frame before the HDLC payload. See HDLC PW Control Word and Payload Size for information.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 263)

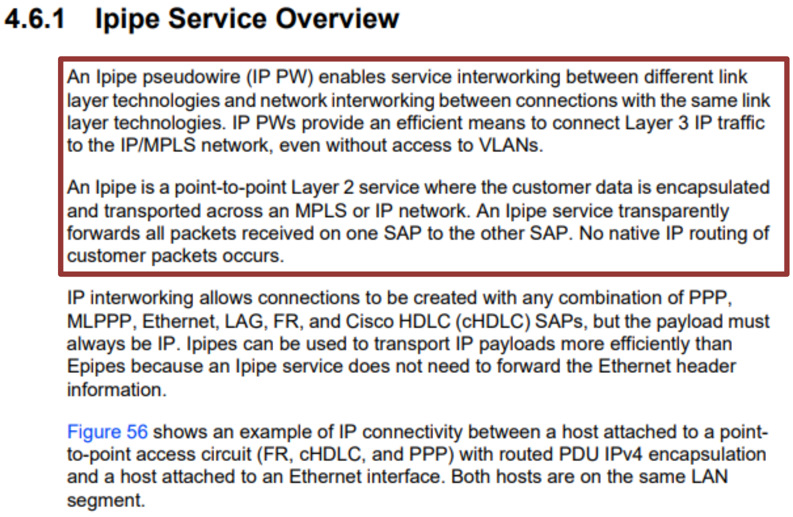

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

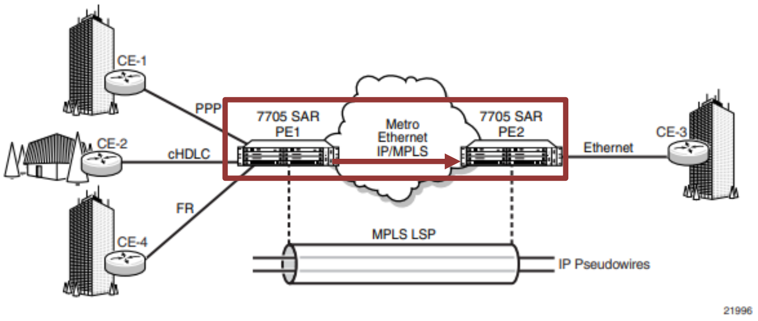

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

51. Defendants instruct its customers to use the 7705 Service Aggregation Router in transmitting the incoming data frames received from two or more of the client nodes to one of the ports of the hub, and such that converting the received incoming data frames comprises associating the two or more of the client nodes with different, respective Virtual Local Area Networks (VLANs) on the network, such that receiving the incoming data frames comprises receiving the incoming frames through at least one of a time domain multiplexed (TDM) interface and a serial interface, and such that transmitting the incoming data frames comprises transmitting the incoming data frames through an Ethernet port.

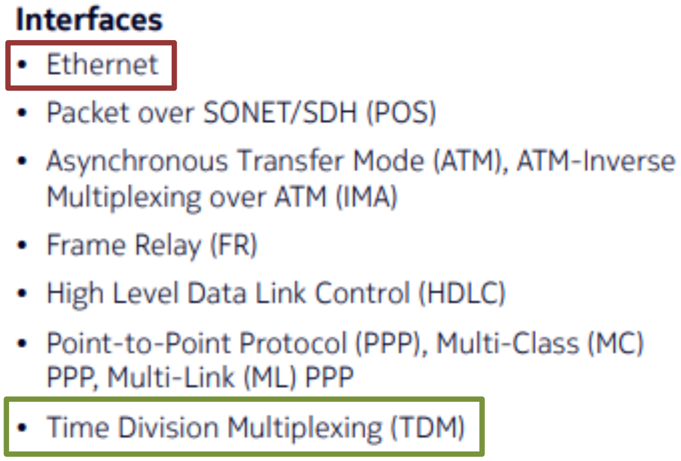

Interfaces

- Ethernet
- Packet over SONET/SDH (POS)
- Asynchronous Transfer Mode (ATM), ATM-Inverse Multiplexing over ATM (IMA)
- Frame Relay (FR)
- High Level Data Link Control (HDLC)
- Point-to-Point Protocol (PPP), Multi-Class (MC) PPP, Multi-Link (ML) PPP
- Time Division Multiplexing (TDM)

https://onestore.nokia.com/asset/f/162833 (page 3)

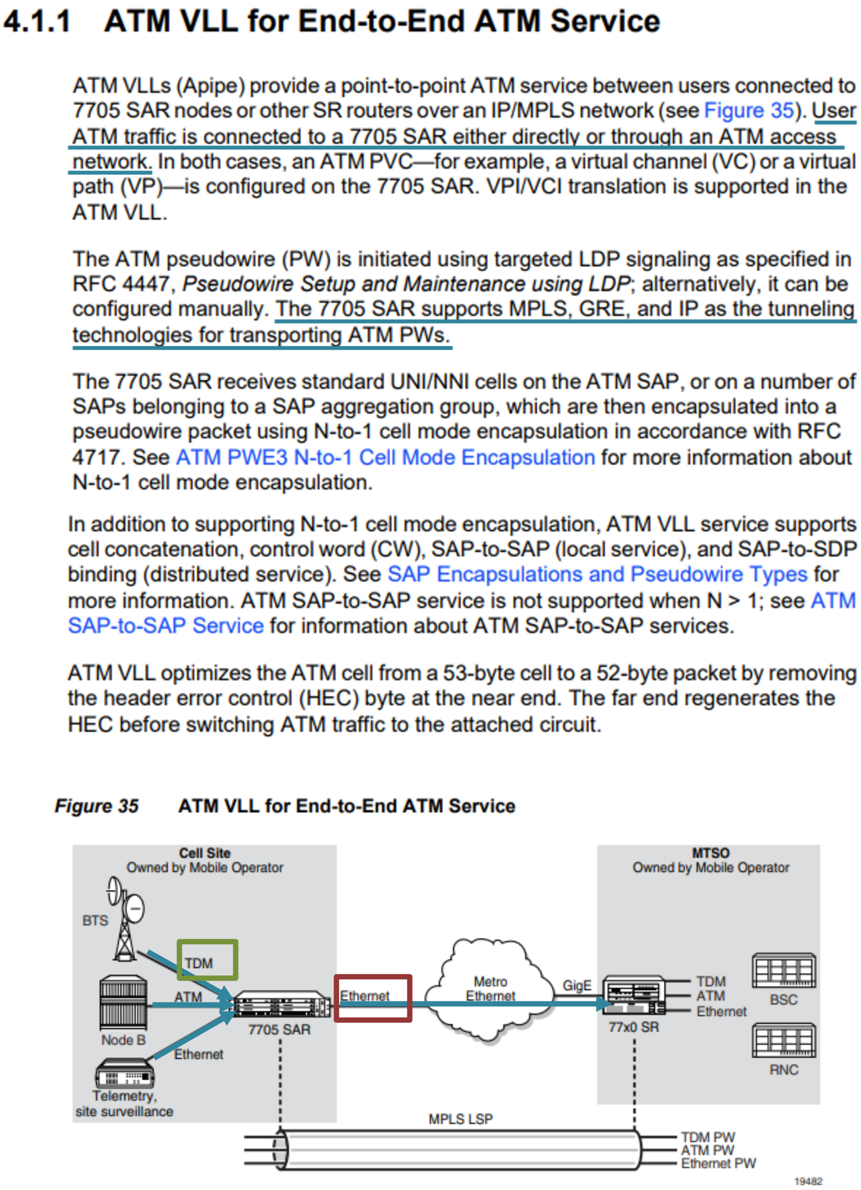

4.1.1 ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 35). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 35 ATM VLL for End-to-End ATM Service

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 198-99)

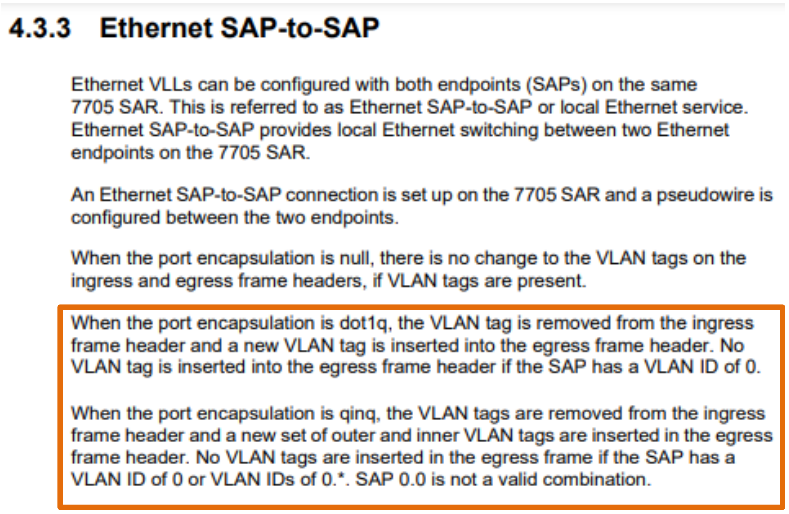

4.3.3 Ethernet SAP-to-SAP

Ethernet VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR. This is referred to as Ethernet SAP-to-SAP or local Ethernet service. Ethernet SAP-to-SAP provides local Ethernet switching between two Ethernet endpoints on the 7705 SAR.

An Ethernet SAP-to-SAP connection is set up on the 7705 SAR and a pseudowire is configured between the two endpoints.

When the port encapsulation is null, there is no change to the VLAN tags on the ingress and egress frame headers, if VLAN tags are present.

When the port encapsulation is dot1q, the VLAN tag is removed from the ingress frame header and a new VLAN tag is inserted into the egress frame header. No VLAN tag is inserted into the egress frame header if the SAP has a VLAN ID of 0.

When the port encapsulation is qinq, the VLAN tags are removed from the ingress frame header and a new set of outer and inner VLAN tags are inserted in the egress frame header. No VLAN tags are inserted in the egress frame if the SAP has a VLAN ID of 0 or VLAN IDs of 0.*. SAP 0.0 is not a valid combination.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 235-36)

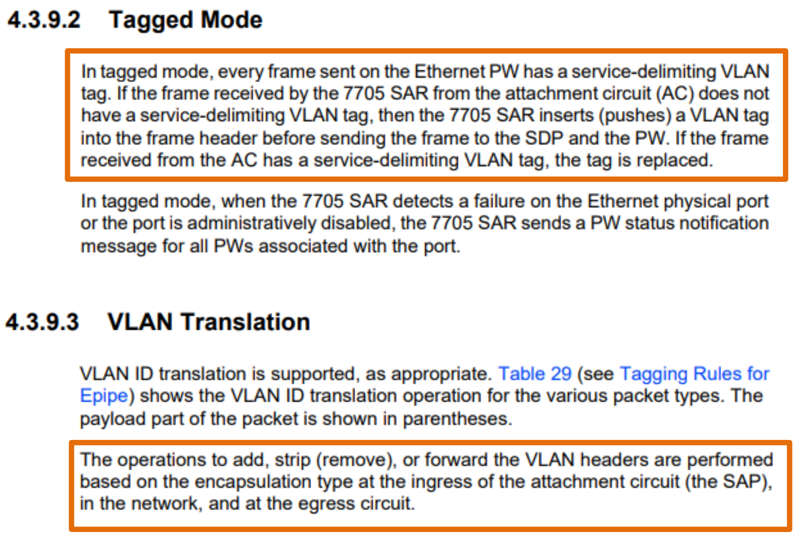

4.3.9.2 Tagged Mode

In tagged mode, every frame sent on the Ethernet PW has a service-delimiting VLAN tag. If the frame received by the 7705 SAR from the attachment circuit (AC) does not have a service-delimiting VLAN tag, then the 7705 SAR inserts (pushes) a VLAN tag into the frame header before sending the frame to the SDP and the PW. If the frame received from the AC has a service-delimiting VLAN tag, the tag is replaced.

In tagged mode, when the 7705 SAR detects a failure on the Ethernet physical port or the port is administratively disabled, the 7705 SAR sends a PW status notification message for all PWs associated with the port.

4.3.9.3 VLAN Translation

VLAN ID translation is supported, as appropriate. Table 29 (see Tagging Rules for Epipe) shows the VLAN ID translation operation for the various packet types. The payload part of the packet is shown in parentheses.

The operations to add, strip (remove), or forward the VLAN headers are performed based on the encapsulation type at the ingress of the attachment circuit (the SAP), in the network, and at the egress circuit.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 243-44)

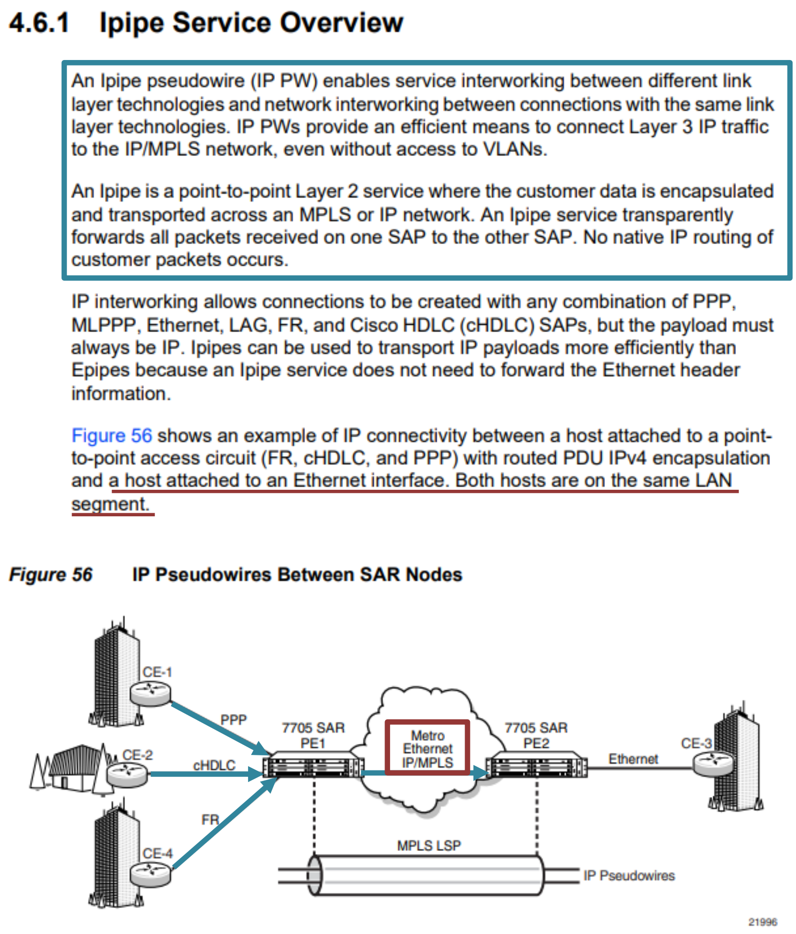

4.6.1 Ipipe Service Overview

An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 56 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 56 **IP Pseudowires Between SAR Nodes**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (page 267)

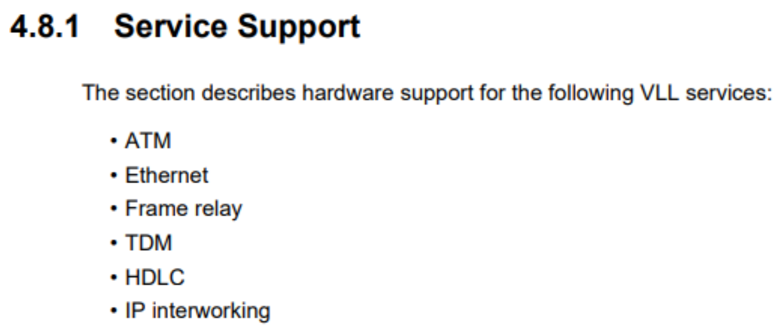

4.8.1 Service Support

The section describes hardware support for the following VLL services:

- ATM
- Ethernet
- Frame relay
- TDM
- HDLC
- IP interworking

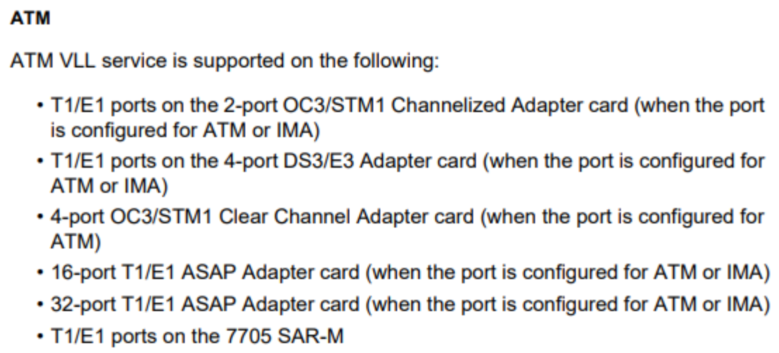

ATM

ATM VLL service is supported on the following:

- T1/E1 ports on the 2-port OC3/STM1 Channelized Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 4-port DS3/E3 Adapter card (when the port is configured for ATM or IMA)
- 4-port OC3/STM1 Clear Channel Adapter card (when the port is configured for ATM)
- 16-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- 32-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 7705 SAR-M

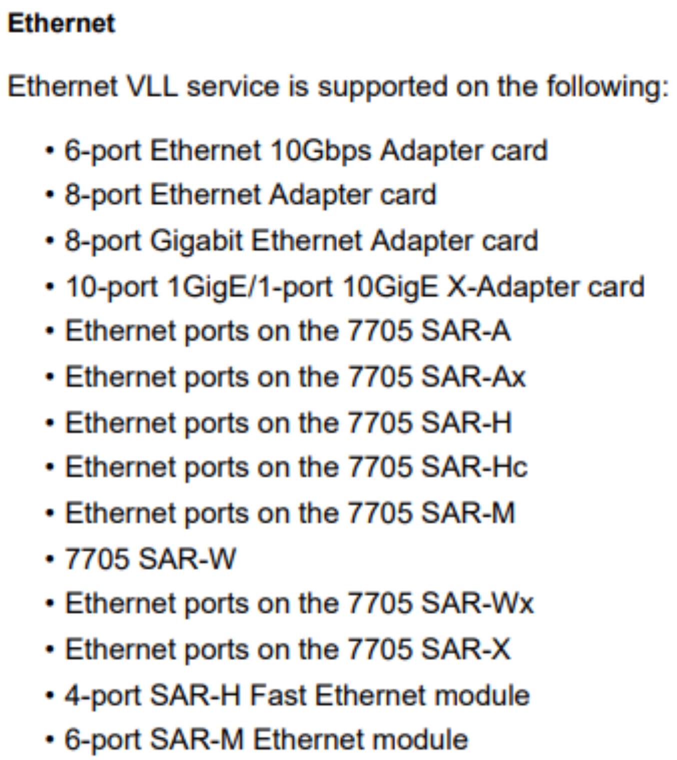

Ethernet

Ethernet VLL service is supported on the following:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Ethernet ports on the 7705 SAR-A
- Ethernet ports on the 7705 SAR-Ax
- Ethernet ports on the 7705 SAR-H
- Ethernet ports on the 7705 SAR-Hc
- Ethernet ports on the 7705 SAR-M
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx
- Ethernet ports on the 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module

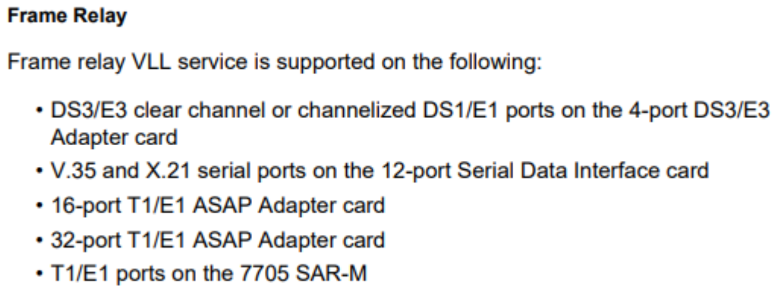

Frame Relay

Frame relay VLL service is supported on the following:

- DS3/E3 clear channel or channelized DS1/E1 ports on the 4-port DS3/E3 Adapter card
- V.35 and X.21 serial ports on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

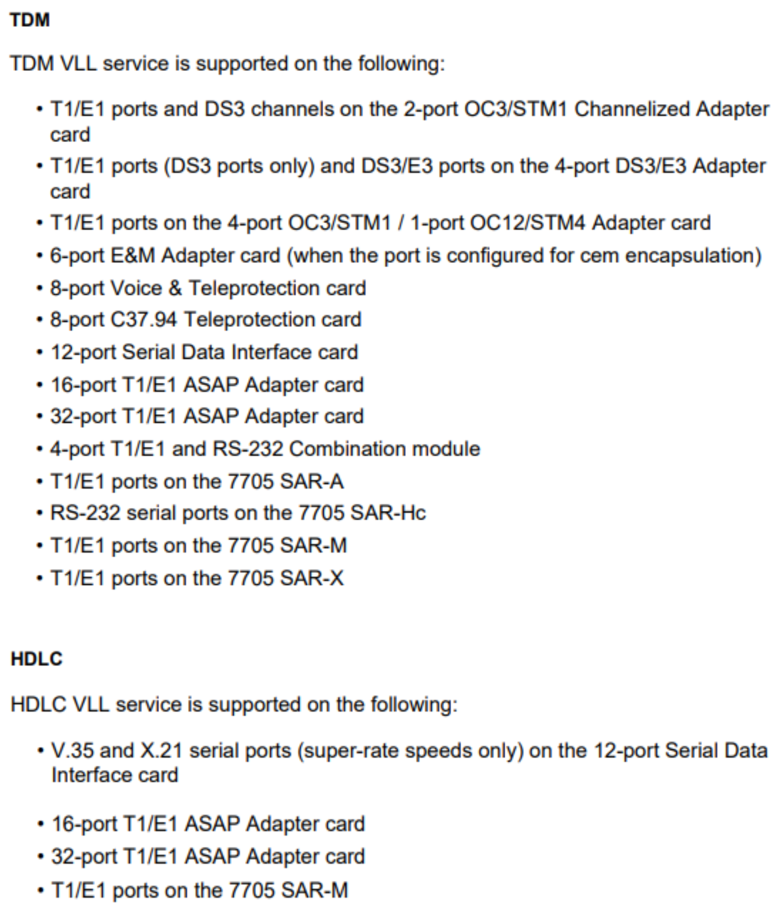

TDM

TDM VLL service is supported on the following:

- T1/E1 ports and DS3 channels on the 2-port OC3/STM1 Channelized Adapter card
- T1/E1 ports (DS3 ports only) and DS3/E3 ports on the 4-port DS3/E3 Adapter card
- T1/E1 ports on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- 6-port E&M Adapter card (when the port is configured for cem encapsulation)
- 8-port Voice & Teleprotection card
- 8-port C37.94 Teleprotection card
- 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 4-port T1/E1 and RS-232 Combination module
- T1/E1 ports on the 7705 SAR-A
- RS-232 serial ports on the 7705 SAR-Hc
- T1/E1 ports on the 7705 SAR-M
- T1/E1 ports on the 7705 SAR-X

HDLC

HDLC VLL service is supported on the following:

- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

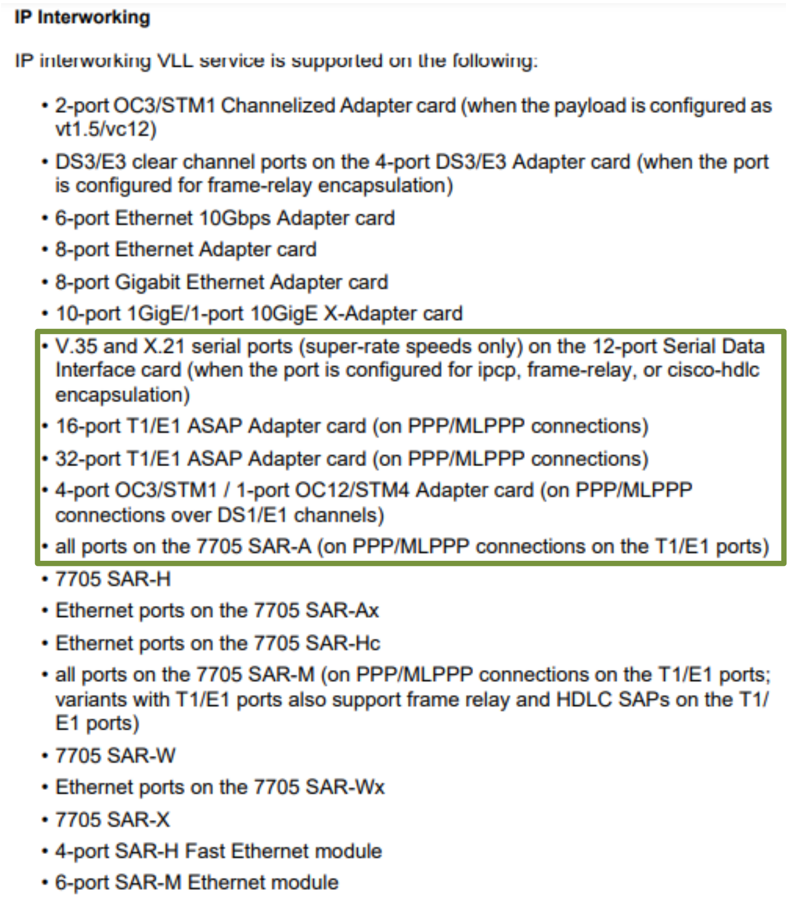

IP Interworking

IP interworking VLL service is supported on the following:

- 2-port OC3/STM1 Channelized Adapter card (when the payload is configured as vt1.5/vc12)
- DS3/E3 clear channel ports on the 4-port DS3/E3 Adapter card (when the port is configured for frame-relay encapsulation)
- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card (when the port is configured for ipcp, frame-relay, or cisco-hdlc encapsulation)
- 16-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- 32-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card (on PPP/MLPPP connections over DS1/E1 channels)
- all ports on the 7705 SAR-A (on PPP/MLPPP connections on the T1/E1 ports)
- 7705 SAR-H
- Ethernet ports on the 7705 SAR-Ax
- Ethernet ports on the 7705 SAR-Hc
- all ports on the 7705 SAR-M (on PPP/MLPPP connections on the T1/E1 ports; variants with T1/E1 ports also support frame relay and HDLC SAPs on the T1/E1 ports)
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx
- 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17555AAABTQZZA01_V1_7705%20SAR%20Services%20Guide%2021.10.R1.pdf (pages 289-292)

52. As a result of Defendants' infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT TWO
INFRINGEMENT OF U.S. PATENT 7,463,580

53. Plaintiff incorporates by reference the allegations in all preceding paragraphs as if fully set forth herein.

54. The '580 Patent, entitled "Resource sharing among network tunnels" was filed on December 15, 2005, and issued on December 9, 2008.

55. Plaintiff is the assignee and owner of all rights, title and interest to the ’580 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

56. The ’580 Patent addresses problems in the art of Multiprotocol label switching (MPLS), specifically that “tunnel-oriented resource reservation protocols such as RSVP-TE and CR-LDP cited above are typically unable to share resources among communication paths, such as protected paths (except for resource sharing between different instances of the same path, which are not considered to be separate communication paths in this context).” 2:22-27.

57. The ‘580 Patent discloses that, “[t]he methods and systems described hereinbelow enable resource allocations in network segments and network elements to be shared between two or more communication paths, thus overcoming these shortcomings of the prior art.” 2:27-31.

Direct Infringement

58. Defendants, without authorization or license from Plaintiff, have been and are directly infringing the ’580 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale telecommunications equipment that infringes one or more claims of the ’580 Patent. Defendants develop, design, manufacture, and distribute telecommunications equipment that infringes one or more claims of the ’580 Patent. Defendants further provide services that practice methods that infringe one or more claims of the ’580 Patent. Defendants are thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Defendants’ 7450 Ethernet Service Switch, and all other substantially similar products (collectively the “’580 Accused Products”).

59. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendants' infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '580 Accused Products.

60. Defendants are liable for direct infringement pursuant to 35 U.S.C. § 271 for the manufacture, sale, offer for sale, importation, or distribution of Defendants' 7450 Ethernet Service Switch.

61. Defendants' 7450 Ethernet Service Switch is a non-limiting example of an apparatus that meets all limitations of claim 8 of the '580 Patent, either literally or equivalently.

62. Defendants' 7450 Ethernet Service Switch comprises a network element and a network interface for communicating with other elements in a communication network:

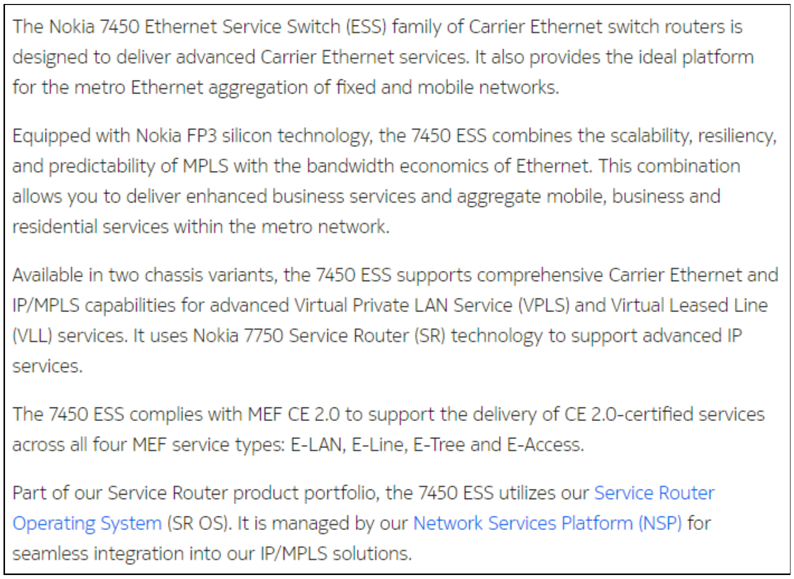

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and residential services within the metro network.

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

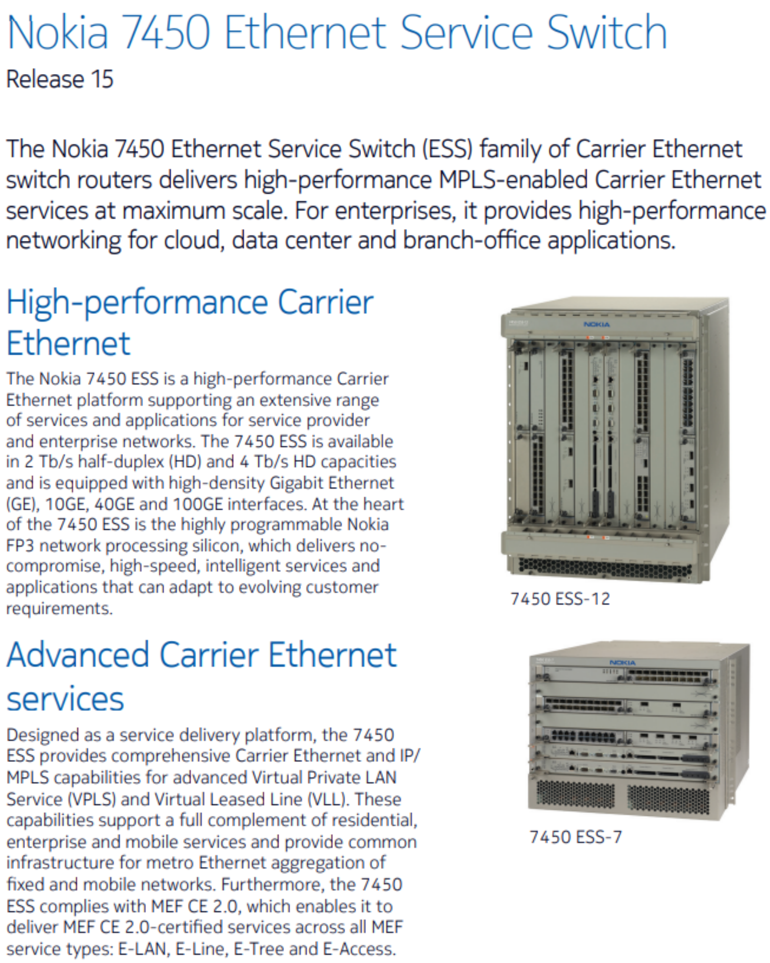

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/f/164727 (page 1)

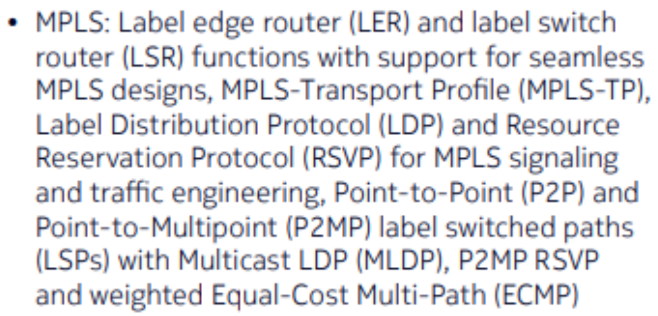

- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)

https://onestore.nokia.com/asset/f/164727 (page 6)

63. Defendants' 7450 Ethernet Service Switch comprises a processor, which is arranged to accept, via the network interface, a notification distributed over the communication network of an affiliation with a resource-sharing group of at least first and second tunnels, which

have respective origin network elements and termination network elements, and which traverse different routes through the network.

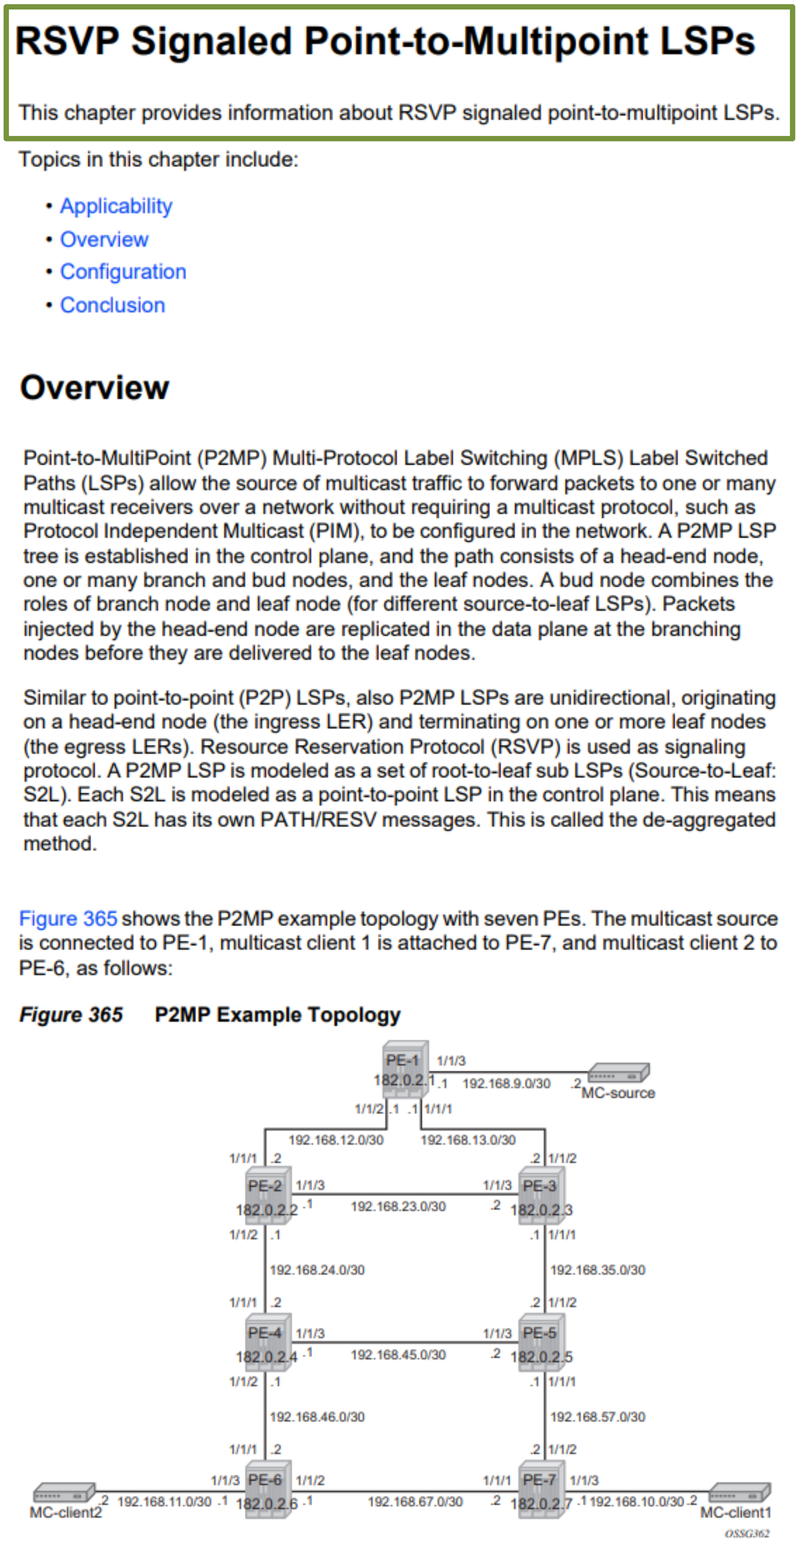

RSVP Signaled Point-to-Multipoint LSPs

This chapter provides information about RSVP signaled point-to-multipoint LSPs.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

Figure 365 shows the P2MP example topology with seven PEs. The multicast source is connected to PE-1, multicast client 1 is attached to PE-7, and multicast client 2 to PE-6, as follows:

Figure 365 **P2MP Example Topology**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (pages 1693-94)

Figure 368 **P2MP LSP LSP-p2mp-1 with Strict S2L Path toward PE-7**

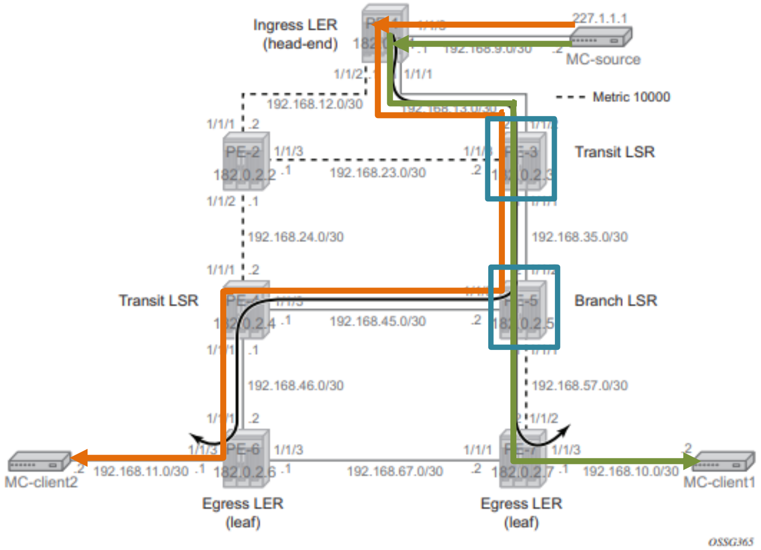

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1721)

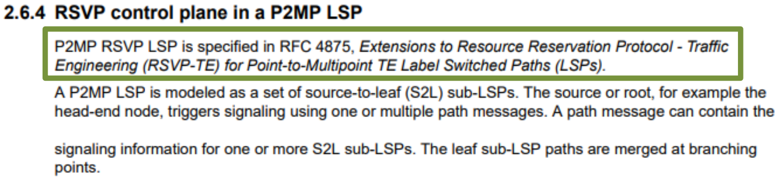

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs).*

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the

signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

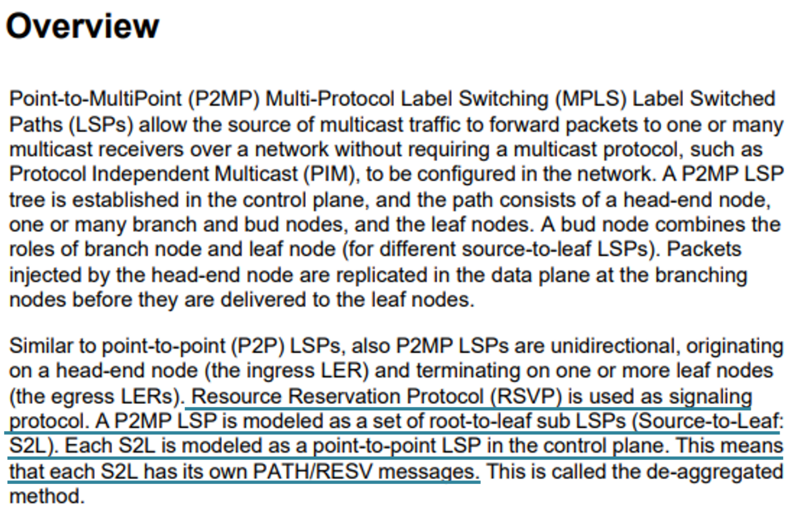

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1693)

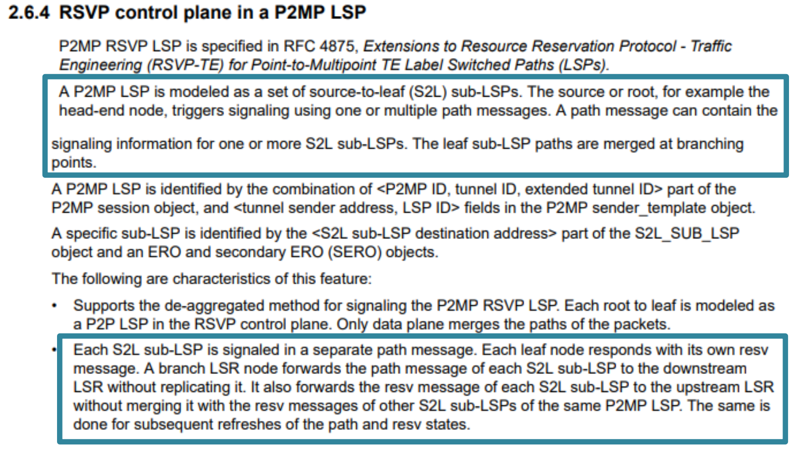

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*.

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

A P2MP LSP is identified by the combination of <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the P2MP sender_template object.

A specific sub-LSP is identified by the <S2L sub-LSP destination address> part of the S2L_SUB_LSP object and an ERO and secondary ERO (SERO) objects.

The following are characteristics of this feature:

- Supports the de-aggregated method for signaling the P2MP RSVP LSP. Each root to leaf is modeled as a P2P LSP in the RSVP control plane. Only data plane merges the paths of the packets.
- Each S2L sub-LSP is signaled in a separate path message. Each leaf node responds with its own resv message. A branch LSR node forwards the path message of each S2L sub-LSP to the downstream LSR without replicating it. It also forwards the resv message of each S2L sub-LSP to the upstream LSR without merging it with the resv messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and resv states.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

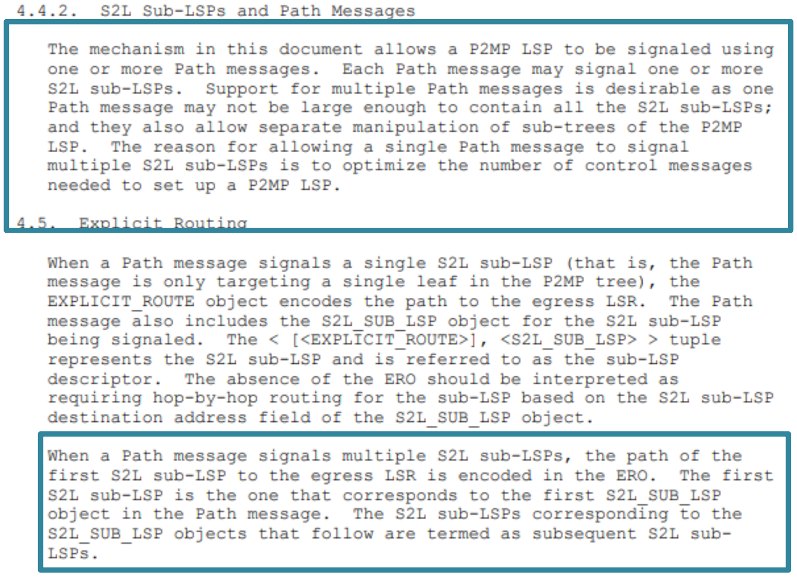

4.4.2. S2L Sub-LSPs and Path Messages

The mechanism in this document allows a P2MP LSP to be signaled using one or more Path messages. Each Path message may signal one or more S2L sub-LSPs. Support for multiple Path messages is desirable as one Path message may not be large enough to contain all the S2L sub-LSPs; and they also allow separate manipulation of sub-trees of the P2MP LSP. The reason for allowing a single Path message to signal multiple S2L sub-LSPs is to optimize the number of control messages needed to set up a P2MP LSP.

4.5. Explicit Routing

When a Path message signals a single S2L sub-LSP (that is, the Path message is only targeting a single leaf in the P2MP tree), the EXPLICIT_ROUTE object encodes the path to the egress LSR. The Path message also includes the S2L_SUB_LSP object for the S2L sub-LSP being signaled. The < [<EXPLICIT_ROUTE>], <S2L_SUB_LSP> > tuple represents the S2L sub-LSP and is referred to as the sub-LSP descriptor. The absence of the ERO should be interpreted as requiring hop-by-hop routing for the sub-LSP based on the S2L sub-LSP destination address field of the S2L_SUB_LSP object.

When a Path message signals multiple S2L sub-LSPs, the path of the first S2L sub-LSP to the egress LSR is encoded in the ERO. The first S2L sub-LSP is the one that corresponds to the first S2L_SUB_LSP object in the Path message. The S2L sub-LSPs corresponding to the S2L_SUB_LSP objects that follow are termed as subsequent S2L sub-LSPs.

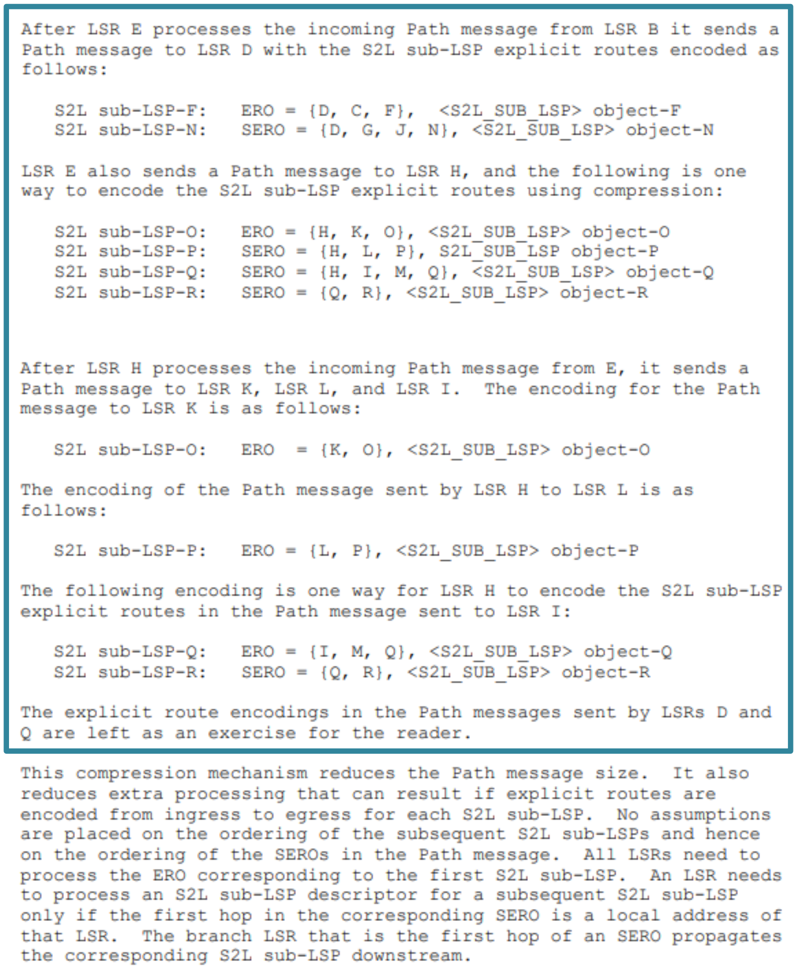

After LSR E processes the incoming Path message from LSR B it sends a Path message to LSR D with the S2L sub-LSP explicit routes encoded as follows:

```
   S2L sub-LSP-F:   ERO = {D, C, F},  <S2L_SUB_LSP> object-F
   S2L sub-LSP-N:   SERO = {D, G, J, N}, <S2L_SUB_LSP> object-N
```

LSR E also sends a Path message to LSR H, and the following is one way to encode the S2L sub-LSP explicit routes using compression:

```
   S2L sub-LSP-O:   ERO = {H, K, O}, <S2L_SUB_LSP> object-O
   S2L sub-LSP-P:   SERO = {H, L, P}, S2L_SUB_LSP object-P
   S2L sub-LSP-Q:   SERO = {H, I, M, Q}, <S2L_SUB_LSP> object-Q
   S2L sub-LSP-R:   SERO = {Q, R}, <S2L_SUB_LSP> object-R
```

After LSR H processes the incoming Path message from E, it sends a Path message to LSR K, LSR L, and LSR I. The encoding for the Path message to LSR K is as follows:

```
   S2L sub-LSP-O:   ERO  = {K, O}, <S2L_SUB_LSP> object-O
```

The encoding of the Path message sent by LSR H to LSR L is as follows:

```
   S2L sub-LSP-P:   ERO = {L, P}, <S2L_SUB_LSP> object-P
```

The following encoding is one way for LSR H to encode the S2L sub-LSP explicit routes in the Path message sent to LSR I:

```
   S2L sub-LSP-Q:   ERO = {I, M, Q}, <S2L_SUB_LSP> object-Q
   S2L sub-LSP-R:   SERO = {Q, R}, <S2L_SUB_LSP> object-R
```

The explicit route encodings in the Path messages sent by LSRs D and Q are left as an exercise for the reader.

This compression mechanism reduces the Path message size. It also reduces extra processing that can result if explicit routes are encoded from ingress to egress for each S2L sub-LSP. No assumptions are placed on the ordering of the subsequent S2L sub-LSPs and hence on the ordering of the SEROs in the Path message. All LSRs need to process the ERO corresponding to the first S2L sub-LSP. An LSR needs to process an S2L sub-LSP descriptor for a subsequent S2L sub-LSP only if the first hop in the corresponding SERO is a local address of that LSR. The branch LSR that is the first hop of an SERO propagates the corresponding S2L sub-LSP downstream.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (pages 7 - 9)

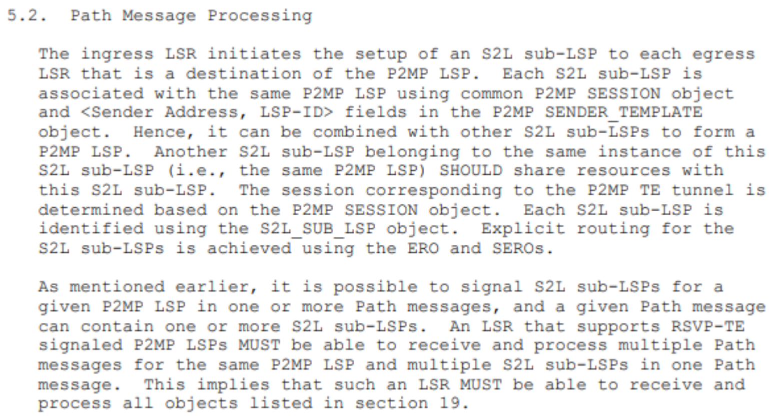

5.2. Path Message Processing

The ingress LSR initiates the setup of an S2L sub-LSP to each egress LSR that is a destination of the P2MP LSP. Each S2L sub-LSP is associated with the same P2MP LSP using common P2MP SESSION object and <Sender Address, LSP-ID> fields in the P2MP SENDER_TEMPLATE object. Hence, it can be combined with other S2L sub-LSPs to form a P2MP LSP. Another S2L sub-LSP belonging to the same instance of this S2L sub-LSP (i.e., the same P2MP LSP) SHOULD share resources with this S2L sub-LSP. The session corresponding to the P2MP TE tunnel is determined based on the P2MP SESSION object. Each S2L sub-LSP is identified using the S2L_SUB_LSP object. Explicit routing for the S2L sub-LSPs is achieved using the ERO and SEROs.

As mentioned earlier, it is possible to signal S2L sub-LSPs for a given P2MP LSP in one or more Path messages, and a given Path message can contain one or more S2L sub-LSPs. An LSR that supports RSVP-TE signaled P2MP LSPs MUST be able to receive and process multiple Path messages for the same P2MP LSP and multiple S2L sub-LSPs in one Path message. This implies that such an LSR MUST be able to receive and process all objects listed in section 19.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 11)

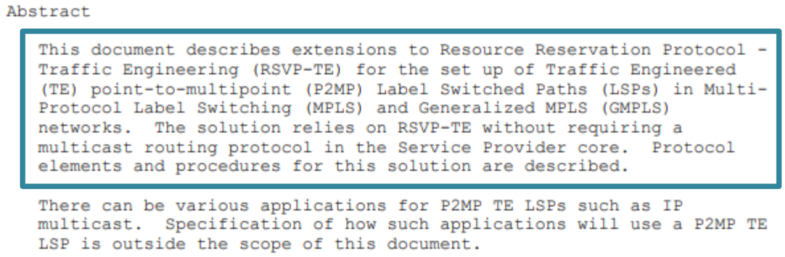

Abstract

This document describes extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for the set up of Traffic Engineered (TE) point-to-multipoint (P2MP) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks. The solution relies on RSVP-TE without requiring a multicast routing protocol in the Service Provider core. Protocol elements and procedures for this solution are described.

There can be various applications for P2MP TE LSPs such as IP multicast. Specification of how such applications will use a P2MP TE LSP is outside the scope of this document.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 1)

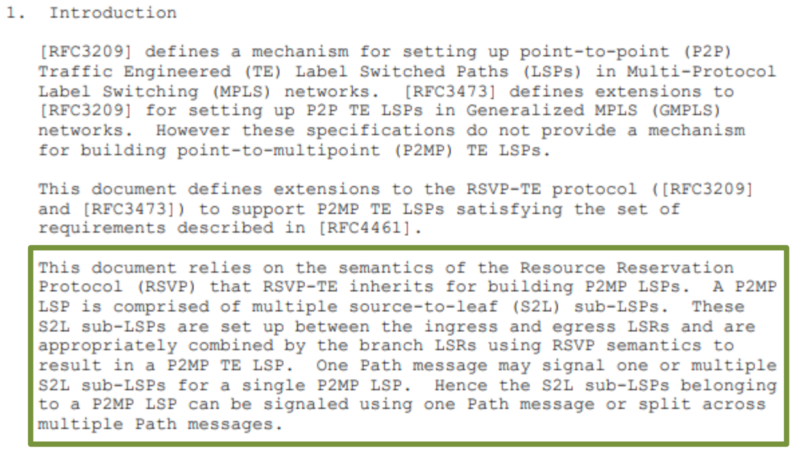

1. Introduction

[RFC3209] defines a mechanism for setting up point-to-point (P2P) Traffic Engineered (TE) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) networks. [RFC3473] defines extensions to [RFC3209] for setting up P2P TE LSPs in Generalized MPLS (GMPLS) networks. However these specifications do not provide a mechanism for building point-to-multipoint (P2MP) TE LSPs.

This document defines extensions to the RSVP-TE protocol ([RFC3209] and [RFC3473]) to support P2MP TE LSPs satisfying the set of requirements described in [RFC4461].

This document relies on the semantics of the Resource Reservation Protocol (RSVP) that RSVP-TE inherits for building P2MP LSPs. A P2MP LSP is comprised of multiple source-to-leaf (S2L) sub-LSPs. These S2L sub-LSPs are set up between the ingress and egress LSRs and are appropriately combined by the branch LSRs using RSVP semantics to result in a P2MP TE LSP. One Path message may signal one or multiple S2L sub-LSPs for a single P2MP LSP. Hence the S2L sub-LSPs belonging to a P2MP LSP can be signaled using one Path message or split across multiple Path messages.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 4)

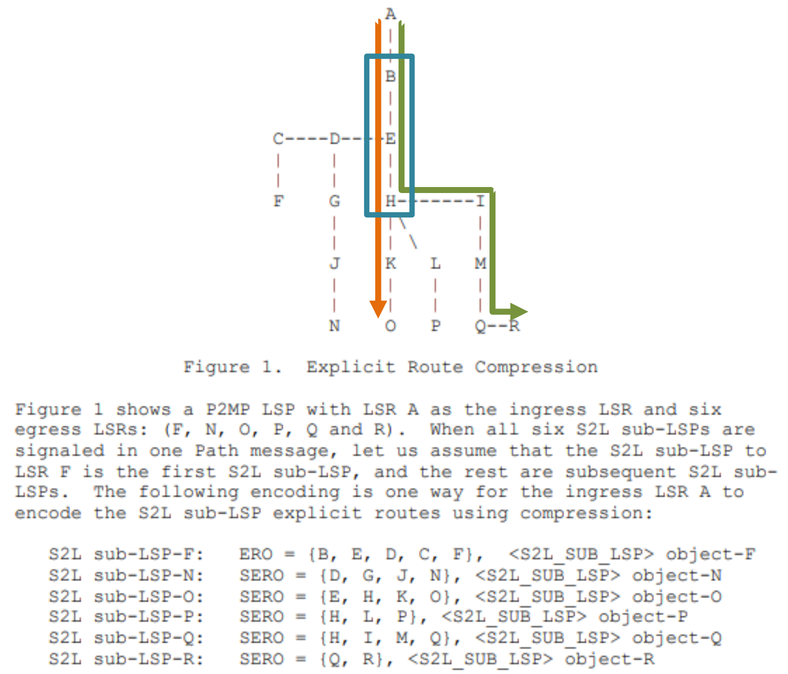

```
                A
                |
                |
                B
                |
                |
       C----D----E
       |    |    |
       |    |    |
       F    G    H-------I
            |    |\      |
            |    | \     |
            J    K  L    M
            |    |  |    |
            |    |  |    |
            N    O  P    Q--R
```

Figure 1. Explicit Route Compression

Figure 1 shows a P2MP LSP with LSR A as the ingress LSR and six egress LSRs: (F, N, O, P, Q and R). When all six S2L sub-LSPs are signaled in one Path message, let us assume that the S2L sub-LSP to LSR F is the first S2L sub-LSP, and the rest are subsequent S2L sub-LSPs. The following encoding is one way for the ingress LSR A to encode the S2L sub-LSP explicit routes using compression:

```
S2L sub-LSP-F:   ERO = {B, E, D, C, F},  <S2L_SUB_LSP> object-F
S2L sub-LSP-N:   SERO = {D, G, J, N}, <S2L_SUB_LSP> object-N
S2L sub-LSP-O:   SERO = {E, H, K, O}, <S2L_SUB_LSP> object-O
S2L sub-LSP-P:   SERO = {H, L, P}, <S2L_SUB_LSP> object-P
S2L sub-LSP-Q:   SERO = {H, I, M, Q}, <S2L_SUB_LSP> object-Q
S2L sub-LSP-R:   SERO = {Q, R}, <S2L_SUB_LSP> object-R
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 8)

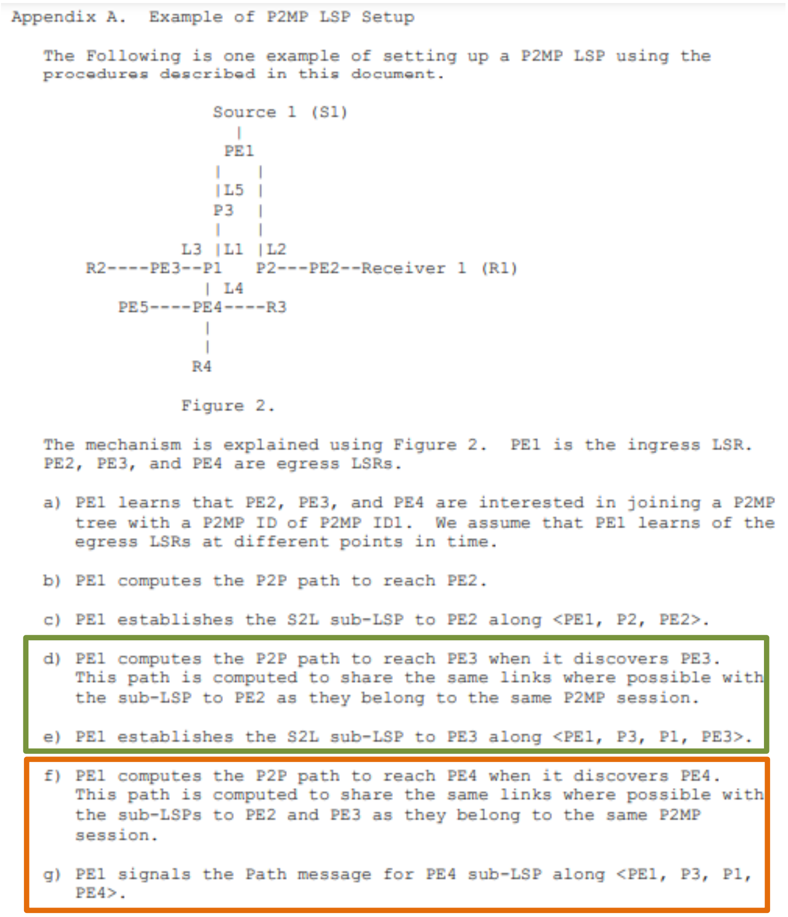

Appendix A. Example of P2MP LSP Setup

The Following is one example of setting up a P2MP LSP using the procedures described in this document.

```
            Source 1 (S1)
              |
             PE1
            |   |
            |L5 |
            P3  |
            |   |
         L3 |L1 |L2
R2----PE3--P1   P2---PE2--Receiver 1 (R1)
          | L4
   PE5----PE4----R3
          |
          |
         R4
```

Figure 2.

The mechanism is explained using Figure 2. PE1 is the ingress LSR. PE2, PE3, and PE4 are egress LSRs.

a) PE1 learns that PE2, PE3, and PE4 are interested in joining a P2MP tree with a P2MP ID of P2MP ID1. We assume that PE1 learns of the egress LSRs at different points in time.

b) PE1 computes the P2P path to reach PE2.

c) PE1 establishes the S2L sub-LSP to PE2 along <PE1, P2, PE2>.

d) PE1 computes the P2P path to reach PE3 when it discovers PE3. This path is computed to share the same links where possible with the sub-LSP to PE2 as they belong to the same P2MP session.

e) PE1 establishes the S2L sub-LSP to PE3 along <PE1, P3, P1, PE3>.

f) PE1 computes the P2P path to reach PE4 when it discovers PE4. This path is computed to share the same links where possible with the sub-LSPs to PE2 and PE3 as they belong to the same P2MP session.

g) PE1 signals the Path message for PE4 sub-LSP along <PE1, P3, P1, PE4>.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 49)

64. Defendants' 7450 Ethernet Service Switch comprises tunnels that meet at least one condition selected from a group of conditions consisting of:

the respective origin network elements of the first and second tunnels are different; and

the respective termination network elements of the first and second tunnels are different:

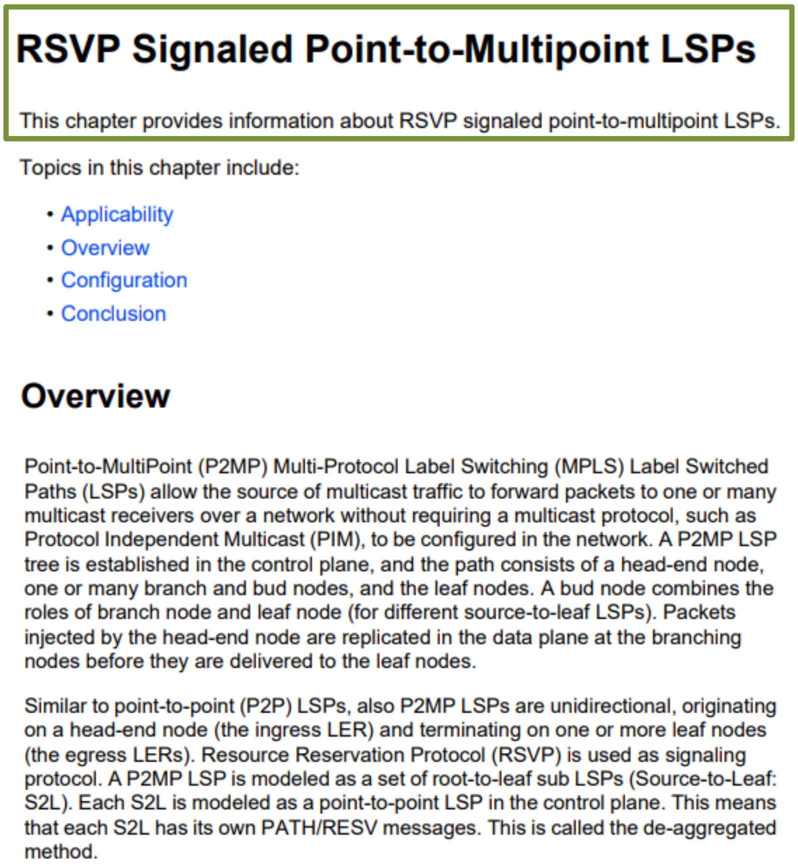

RSVP Signaled Point-to-Multipoint LSPs

This chapter provides information about RSVP signaled point-to-multipoint LSPs.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

Figure 365 shows the P2MP example topology with seven PEs. The multicast source is connected to PE-1, multicast client 1 is attached to PE-7, and multicast client 2 to PE-6, as follows:

Figure 365 **P2MP Example Topology**

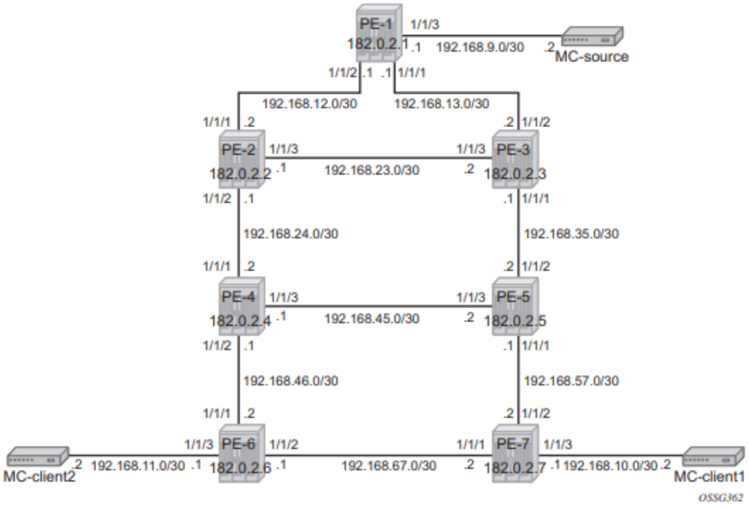

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (pages 1693 and 1694)

Figure 368 **P2MP LSP LSP-p2mp-1 with Strict S2L Path toward PE-7**

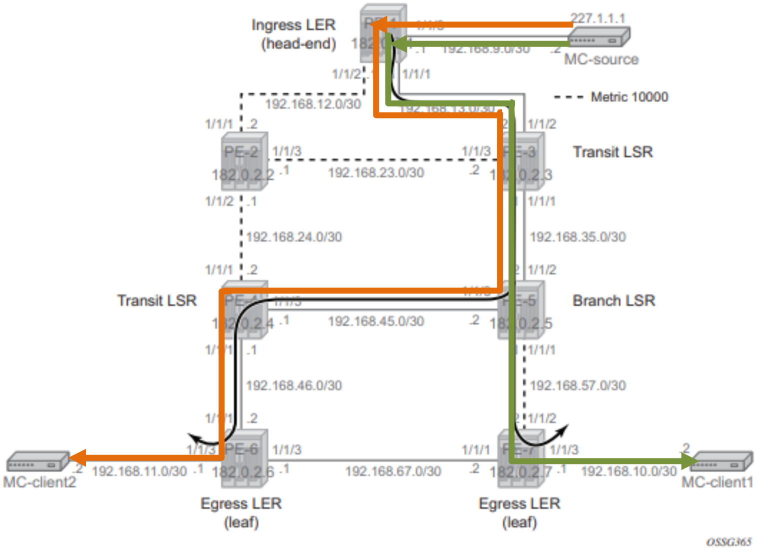

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1721)

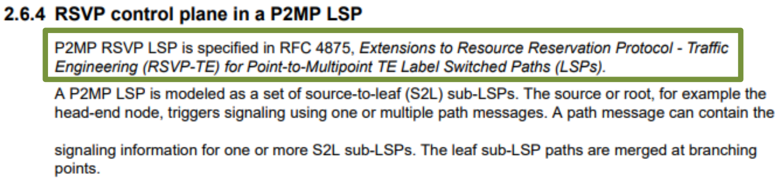

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs).*

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the

signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

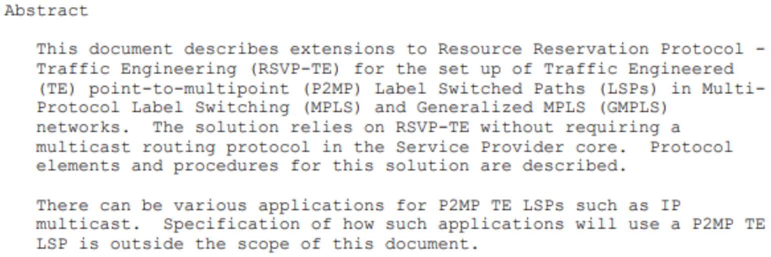

Abstract

This document describes extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for the set up of Traffic Engineered (TE) point-to-multipoint (P2MP) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks. The solution relies on RSVP-TE without requiring a multicast routing protocol in the Service Provider core. Protocol elements and procedures for this solution are described.

There can be various applications for P2MP TE LSPs such as IP multicast. Specification of how such applications will use a P2MP TE LSP is outside the scope of this document.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 1)

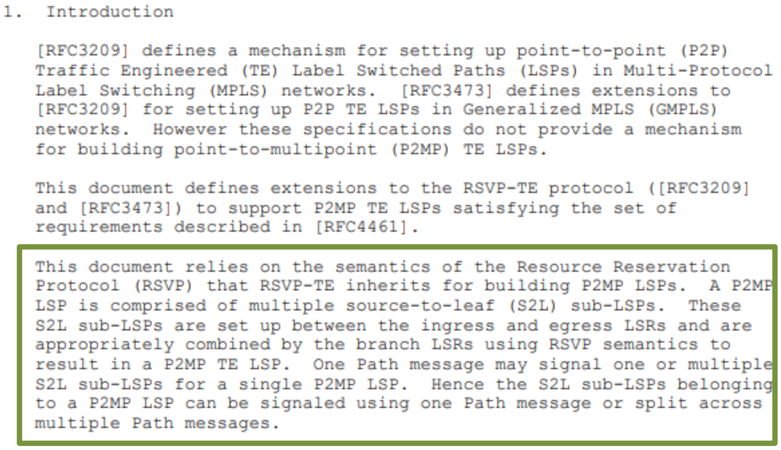

1. Introduction

[RFC3209] defines a mechanism for setting up point-to-point (P2P) Traffic Engineered (TE) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) networks. [RFC3473] defines extensions to [RFC3209] for setting up P2P TE LSPs in Generalized MPLS (GMPLS) networks. However these specifications do not provide a mechanism for building point-to-multipoint (P2MP) TE LSPs.

This document defines extensions to the RSVP-TE protocol ([RFC3209] and [RFC3473]) to support P2MP TE LSPs satisfying the set of requirements described in [RFC4461].

This document relies on the semantics of the Resource Reservation Protocol (RSVP) that RSVP-TE inherits for building P2MP LSPs. A P2MP LSP is comprised of multiple source-to-leaf (S2L) sub-LSPs. These S2L sub-LSPs are set up between the ingress and egress LSRs and are appropriately combined by the branch LSRs using RSVP semantics to result in a P2MP TE LSP. One Path message may signal one or multiple S2L sub-LSPs for a single P2MP LSP. Hence the S2L sub-LSPs belonging to a P2MP LSP can be signaled using one Path message or split across multiple Path messages.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 4)

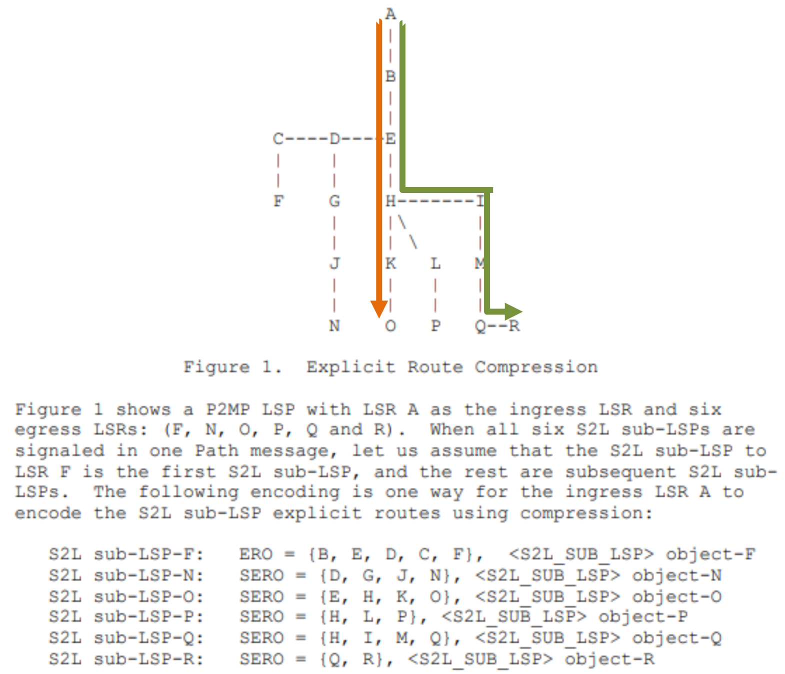

```
                 A
                 |
                 |
                 B
                 |
                 |
       C----D----E
       |    |    |
       |    |    |
       F    G    H-------I
            |    |\      |
            |    | \     |
            J    K  L    M
            |    |  |    |
            |    |  |    |
            N    O  P    Q--R
```

Figure 1. Explicit Route Compression

Figure 1 shows a P2MP LSP with LSR A as the ingress LSR and six egress LSRs: (F, N, O, P, Q and R). When all six S2L sub-LSPs are signaled in one Path message, let us assume that the S2L sub-LSP to LSR F is the first S2L sub-LSP, and the rest are subsequent S2L sub-LSPs. The following encoding is one way for the ingress LSR A to encode the S2L sub-LSP explicit routes using compression:

```
S2L sub-LSP-F:   ERO = {B, E, D, C, F},  <S2L_SUB_LSP> object-F
S2L sub-LSP-N:   SERO = {D, G, J, N}, <S2L_SUB_LSP> object-N
S2L sub-LSP-O:   SERO = {E, H, K, O}, <S2L_SUB_LSP> object-O
S2L sub-LSP-P:   SERO = {H, L, P}, <S2L_SUB_LSP> object-P
S2L sub-LSP-Q:   SERO = {H, I, M, Q}, <S2L_SUB_LSP> object-Q
S2L sub-LSP-R:   SERO = {Q, R}, <S2L_SUB_LSP> object-R
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 8)

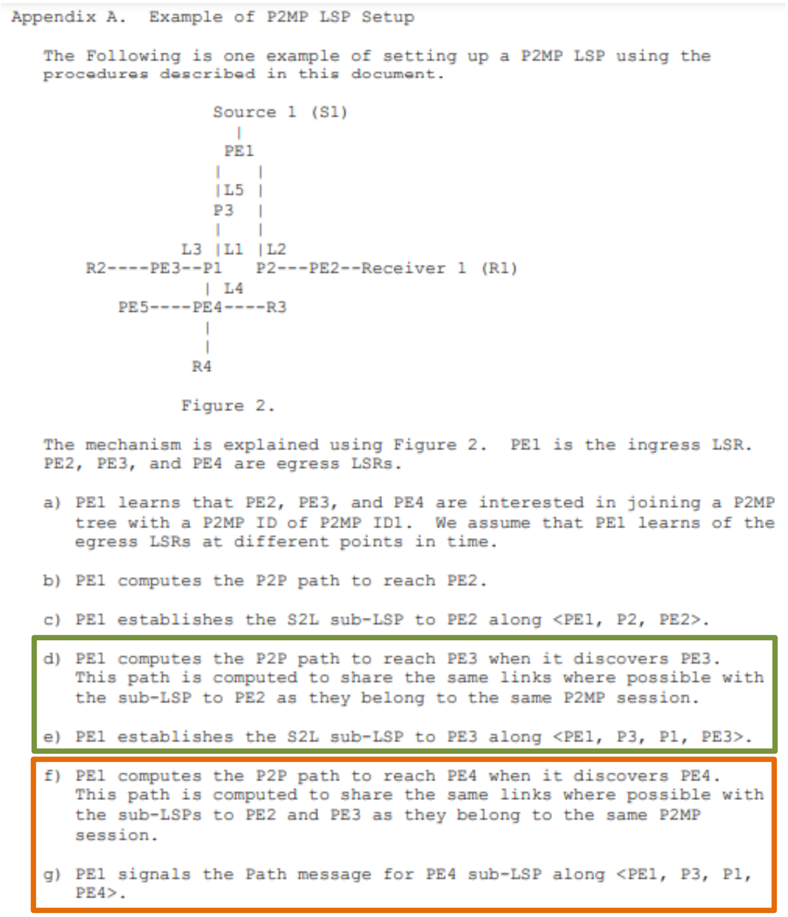

Appendix A. Example of P2MP LSP Setup

The Following is one example of setting up a P2MP LSP using the procedures described in this document.

```
                Source 1 (S1)
                   |
                  PE1
                  |   |
                  |L5 |
                  P3  |
                  |   |
             L3  |L1  |L2
     R2----PE3--P1   P2---PE2--Receiver 1 (R1)
                 | L4
        PE5----PE4----R3
                 |
                 |
                R4
```

Figure 2.

The mechanism is explained using Figure 2. PE1 is the ingress LSR. PE2, PE3, and PE4 are egress LSRs.

a) PE1 learns that PE2, PE3, and PE4 are interested in joining a P2MP tree with a P2MP ID of P2MP ID1. We assume that PE1 learns of the egress LSRs at different points in time.

b) PE1 computes the P2P path to reach PE2.

c) PE1 establishes the S2L sub-LSP to PE2 along <PE1, P2, PE2>.

d) PE1 computes the P2P path to reach PE3 when it discovers PE3. This path is computed to share the same links where possible with the sub-LSP to PE2 as they belong to the same P2MP session.

e) PE1 establishes the S2L sub-LSP to PE3 along <PE1, P3, P1, PE3>.

f) PE1 computes the P2P path to reach PE4 when it discovers PE4. This path is computed to share the same links where possible with the sub-LSPs to PE2 and PE3 as they belong to the same P2MP session.

g) PE1 signals the Path message for PE4 sub-LSP along <PE1, P3, P1, PE4>.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 49)

65. Defendants' 7450 Ethernet Service Switch comprises the processor comprising a call admission control (CAC) module, which is arranged, when the network element is traversed by at least some of the tunnels in the resource-sharing group, to allocate a resource associated with the network element so as to share an allocation of the resource among the at least some of the tunnels responsively to the notification:

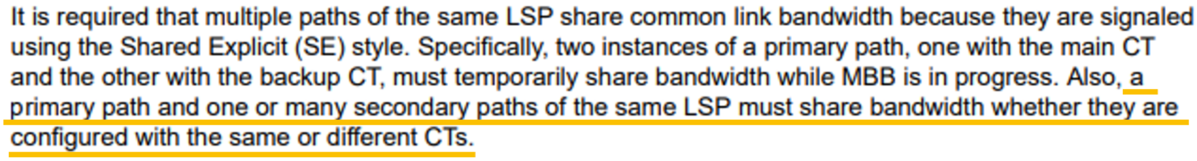

It is required that multiple paths of the same LSP share common link bandwidth because they are signaled using the Shared Explicit (SE) style. Specifically, two instances of a primary path, one with the main CT and the other with the backup CT, must temporarily share bandwidth while MBB is in progress. Also, a primary path and one or many secondary paths of the same LSP must share bandwidth whether they are configured with the same or different CTs.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (page 118)

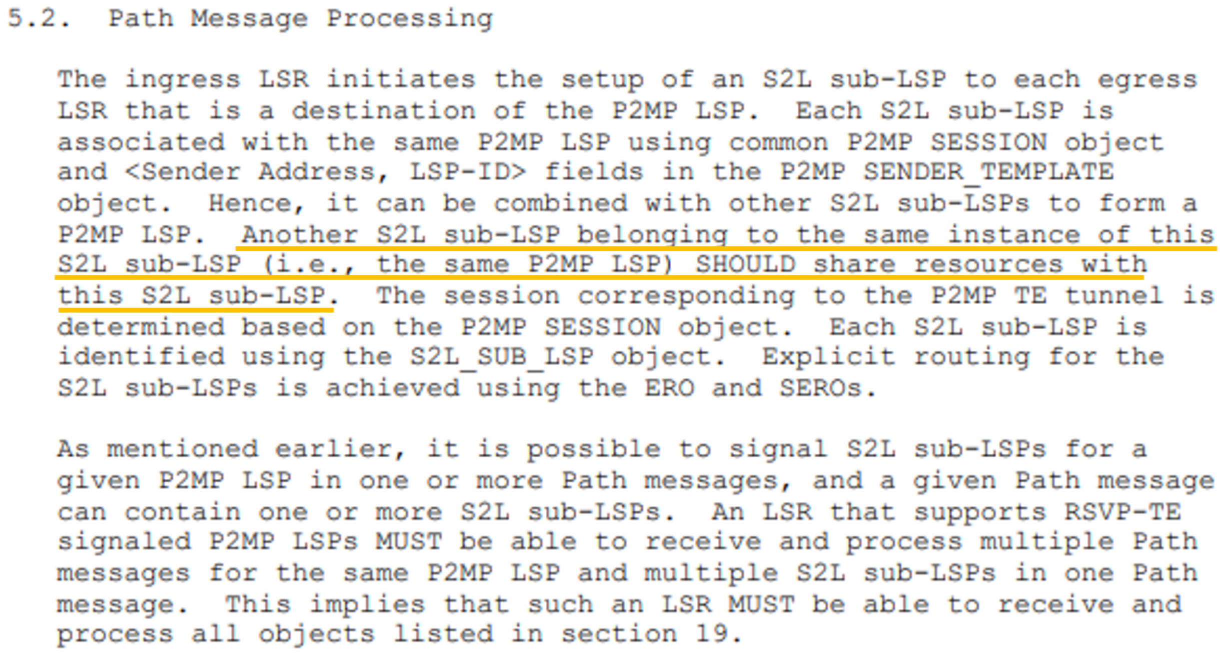

```
5.2.  Path Message Processing

   The ingress LSR initiates the setup of an S2L sub-LSP to each egress
   LSR that is a destination of the P2MP LSP.  Each S2L sub-LSP is
   associated with the same P2MP LSP using common P2MP SESSION object
   and <Sender Address, LSP-ID> fields in the P2MP SENDER_TEMPLATE
   object.  Hence, it can be combined with other S2L sub-LSPs to form a
   P2MP LSP.  Another S2L sub-LSP belonging to the same instance of this
   S2L sub-LSP (i.e., the same P2MP LSP) SHOULD share resources with
   this S2L sub-LSP.  The session corresponding to the P2MP TE tunnel is
   determined based on the P2MP SESSION object.  Each S2L sub-LSP is
   identified using the S2L_SUB_LSP object.  Explicit routing for the
   S2L sub-LSPs is achieved using the ERO and SEROs.

   As mentioned earlier, it is possible to signal S2L sub-LSPs for a
   given P2MP LSP in one or more Path messages, and a given Path message
   can contain one or more S2L sub-LSPs.  An LSR that supports RSVP-TE
   signaled P2MP LSPs MUST be able to receive and process multiple Path
   messages for the same P2MP LSP and multiple S2L sub-LSPs in one Path
   message.  This implies that such an LSR MUST be able to receive and
   process all objects listed in section 19.
```

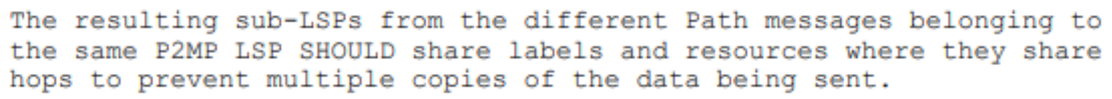

```
The resulting sub-LSPs from the different Path messages belonging to
the same P2MP LSP SHOULD share labels and resources where they share
hops to prevent multiple copies of the data being sent.
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (pages 11 and 12)

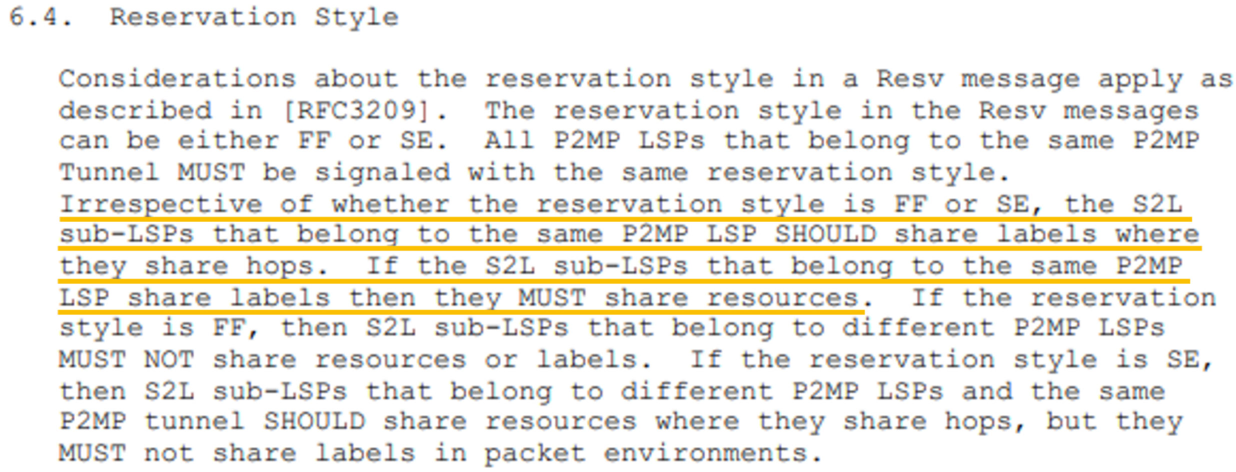

```
6.4.  Reservation Style

   Considerations about the reservation style in a Resv message apply as
   described in [RFC3209].  The reservation style in the Resv messages
   can be either FF or SE.  All P2MP LSPs that belong to the same P2MP
   Tunnel MUST be signaled with the same reservation style.
   Irrespective of whether the reservation style is FF or SE, the S2L
   sub-LSPs that belong to the same P2MP LSP SHOULD share labels where
   they share hops.  If the S2L sub-LSPs that belong to the same P2MP
   LSP share labels then they MUST share resources.  If the reservation
   style is FF, then S2L sub-LSPs that belong to different P2MP LSPs
   MUST NOT share resources or labels.  If the reservation style is SE,
   then S2L sub-LSPs that belong to different P2MP LSPs and the same
   P2MP tunnel SHOULD share resources where they share hops, but they
   MUST not share labels in packet environments.
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 19)

Willful Infringement

66. Defendants have had actual knowledge of the ’580 Patent and its infringement thereof at least as of service or other receipt of Plaintiff’s Complaint.

67. Defendants’ infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendants.

68. Notwithstanding this knowledge, Defendants have knowingly or with reckless disregard infringed the ’580 Patent. Defendants continue to commit acts of infringement despite being on notice of an objectively high likelihood that its actions constituted infringement of Plaintiff’s valid patent rights, either literally or equivalently.

69. Defendants are therefore liable for willful infringement. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

70. Defendants have induced and are knowingly inducing its distributors, testers, trainers, customers and/or end users to directly infringe the ’580 Patent, with the specific intent to induce acts constituting infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

71. Defendants have knowingly contributed to direct infringement by its customers and end users by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the accused products which are not suitable for substantial non-infringing use, and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

72. Defendants’ indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe ’580 Patent.

73. Defendants’ indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendants’ indirect infringement further includes providing application notes instructing its customers on infringing uses of the ’580 Accused Products. The ’580 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the ’580 Patent, either literally or equivalently. Defendants know and intend that customers who purchase the ’580 Accused Products will use those products for their intended purpose. For example, Defendants’ United States website, https://www.nokia.com, instructs customers to use the ’580 Accused Products in numerous infringing applications. Furthermore, Defendants provide instructions and other resources via its “Doc Center” (https://documentation.nokia.com), “Nokia Service Routing Certification” and

training courses from its "NokiaEDU Training Centers" (https://www.nokia.com/networks/training/src/courses/#open-enrollment), and elsewhere on using the '580 Accused Products. Defendants' customers directly infringe the '580 Patent when they follow Defendants' provided instructions on websites, videos, trainings, and elsewhere. Defendants' customers who follow Defendants' provided instructions directly infringe claims of the '580 Patent.

74. In addition, Defendants specifically intend that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendants' infringing products. Defendants know following its instructions directly infringes claims of the '580 Patent, including for example Claim 1.

75. Defendants' customers who follow Defendants' provided instructions directly infringe the method of Claim 1 of the '580 Patent.

76. Defendants instruct its customers to use the 7450 Ethernet Service Switch in a method for communication:

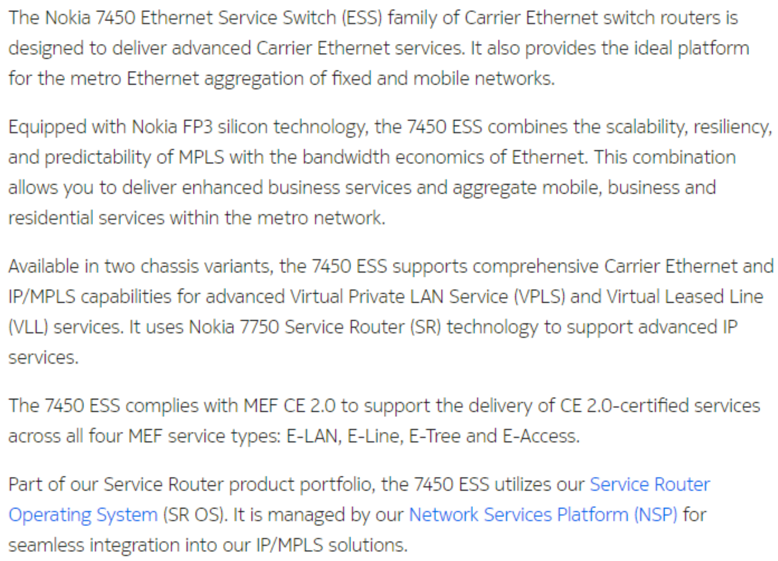

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and residential services within the metro network.

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

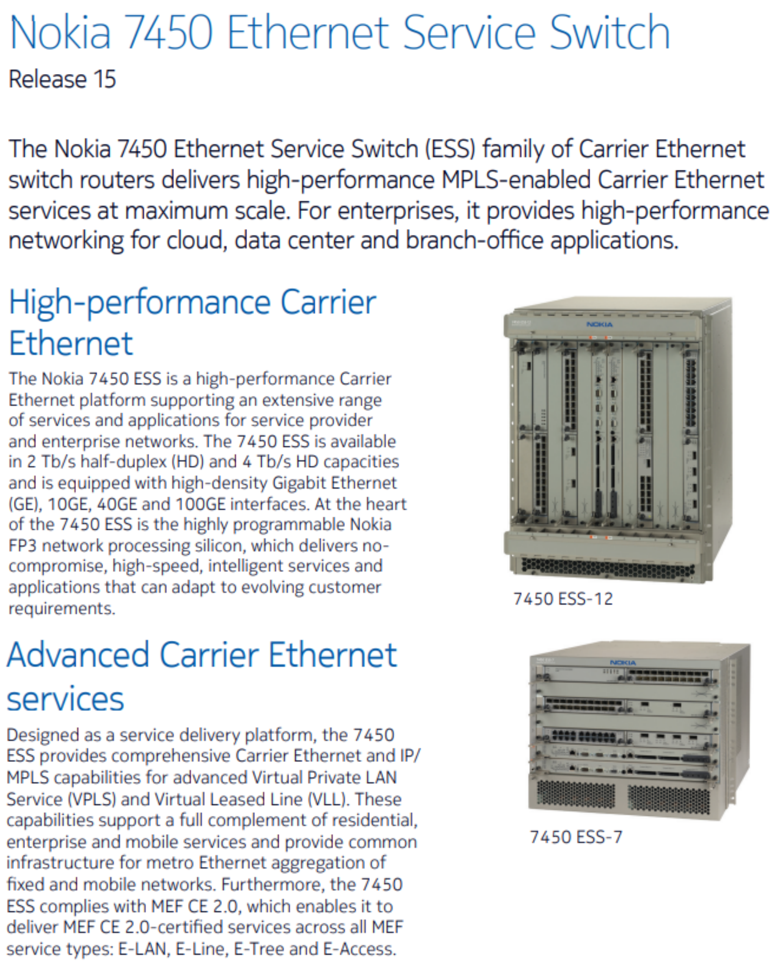

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/f/164727 (page 1)

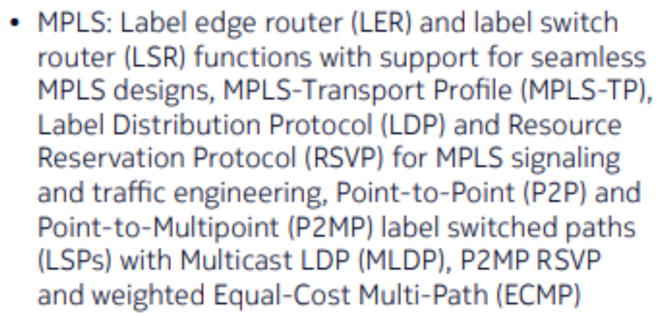

- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)

https://onestore.nokia.com/asset/f/164727 (page 6)

77. Defendants instruct its customers to use the 7450 Ethernet Service Switch to define a resource-sharing group comprising at least first and second tunnels, which have respective origin network elements and termination network elements, and which traverse different routes through a communication network, the routes traversing at least one common network element:

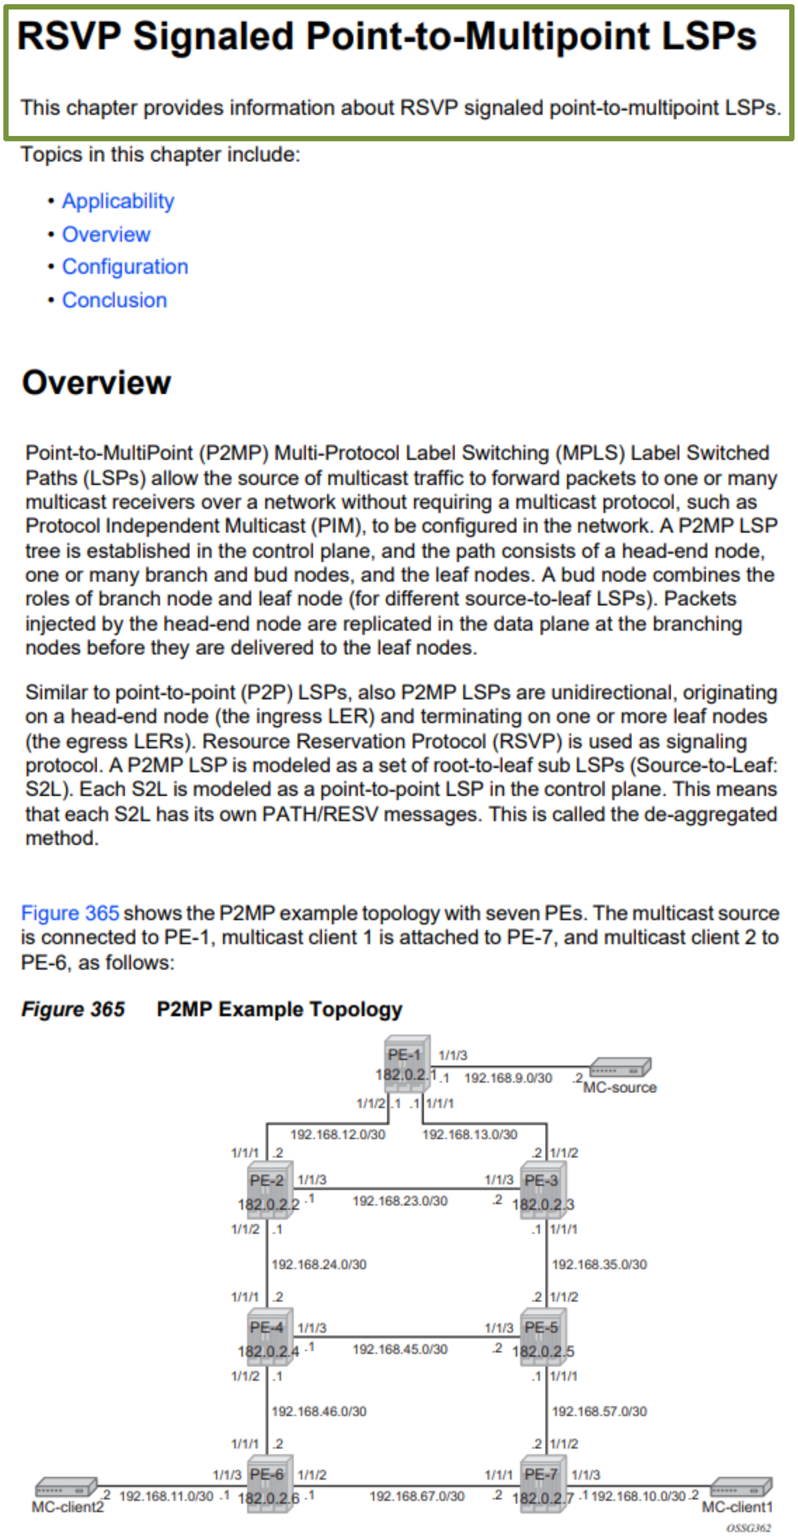

RSVP Signaled Point-to-Multipoint LSPs

This chapter provides information about RSVP signaled point-to-multipoint LSPs.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

Figure 365 shows the P2MP example topology with seven PEs. The multicast source is connected to PE-1, multicast client 1 is attached to PE-7, and multicast client 2 to PE-6, as follows:

Figure 365 **P2MP Example Topology**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (pages 1693-94)

Figure 368 **P2MP LSP LSP-p2mp-1 with Strict S2L Path toward PE-7**

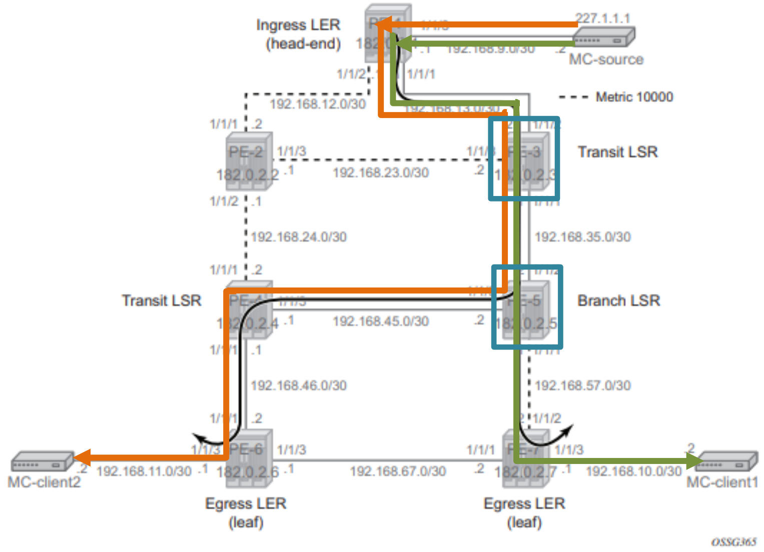

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1721)

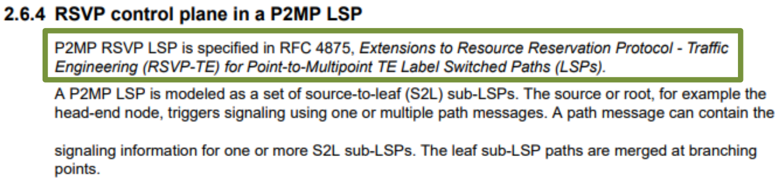

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs).*

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the

signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

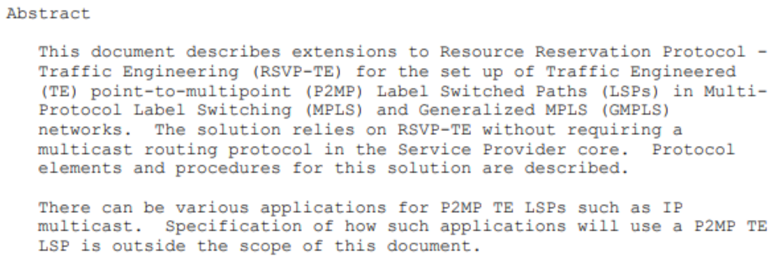

```
Abstract

   This document describes extensions to Resource Reservation Protocol -
   Traffic Engineering (RSVP-TE) for the set up of Traffic Engineered
   (TE) point-to-multipoint (P2MP) Label Switched Paths (LSPs) in Multi-
   Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS)
   networks.  The solution relies on RSVP-TE without requiring a
   multicast routing protocol in the Service Provider core.  Protocol
   elements and procedures for this solution are described.

   There can be various applications for P2MP TE LSPs such as IP
   multicast.  Specification of how such applications will use a P2MP TE
   LSP is outside the scope of this document.
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 1)

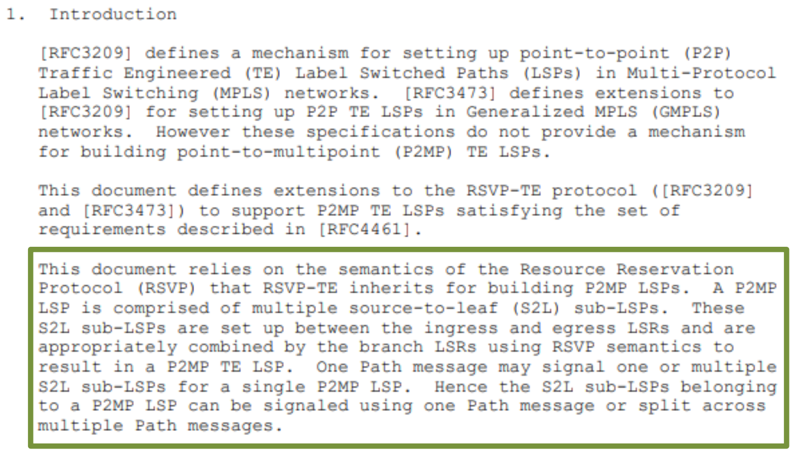

1. Introduction

[RFC3209] defines a mechanism for setting up point-to-point (P2P) Traffic Engineered (TE) Label Switched Paths (LSPs) in Multi-Protocol Label Switching (MPLS) networks. [RFC3473] defines extensions to [RFC3209] for setting up P2P TE LSPs in Generalized MPLS (GMPLS) networks. However these specifications do not provide a mechanism for building point-to-multipoint (P2MP) TE LSPs.

This document defines extensions to the RSVP-TE protocol ([RFC3209] and [RFC3473]) to support P2MP TE LSPs satisfying the set of requirements described in [RFC4461].

This document relies on the semantics of the Resource Reservation Protocol (RSVP) that RSVP-TE inherits for building P2MP LSPs. A P2MP LSP is comprised of multiple source-to-leaf (S2L) sub-LSPs. These S2L sub-LSPs are set up between the ingress and egress LSRs and are appropriately combined by the branch LSRs using RSVP semantics to result in a P2MP TE LSP. One Path message may signal one or multiple S2L sub-LSPs for a single P2MP LSP. Hence the S2L sub-LSPs belonging to a P2MP LSP can be signaled using one Path message or split across multiple Path messages.

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 4)

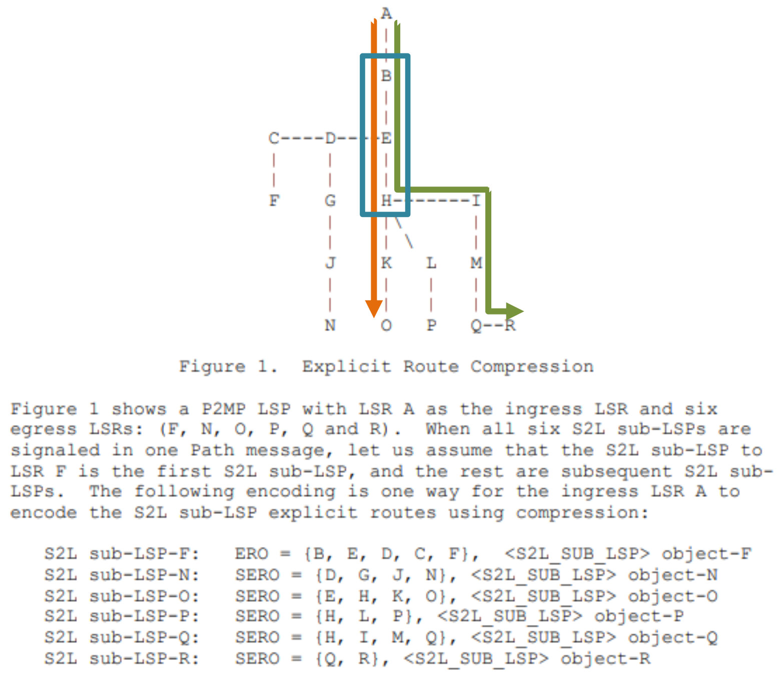
A
B
C----D----E
F G H-------I
J K L M
N O P Q--R

Figure 1. Explicit Route Compression

Figure 1 shows a P2MP LSP with LSR A as the ingress LSR and six egress LSRs: (F, N, O, P, Q and R). When all six S2L sub-LSPs are signaled in one Path message, let us assume that the S2L sub-LSP to LSR F is the first S2L sub-LSP, and the rest are subsequent S2L sub-LSPs. The following encoding is one way for the ingress LSR A to encode the S2L sub-LSP explicit routes using compression:

```
   S2L sub-LSP-F:   ERO = {B, E, D, C, F},  <S2L_SUB_LSP> object-F
   S2L sub-LSP-N:   SERO = {D, G, J, N}, <S2L_SUB_LSP> object-N
   S2L sub-LSP-O:   SERO = {E, H, K, O}, <S2L_SUB_LSP> object-O
   S2L sub-LSP-P:   SERO = {H, L, P}, <S2L_SUB_LSP> object-P
   S2L sub-LSP-Q:   SERO = {H, I, M, Q}, <S2L_SUB_LSP> object-Q
   S2L sub-LSP-R:   SERO = {Q, R}, <S2L_SUB_LSP> object-R
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 8)

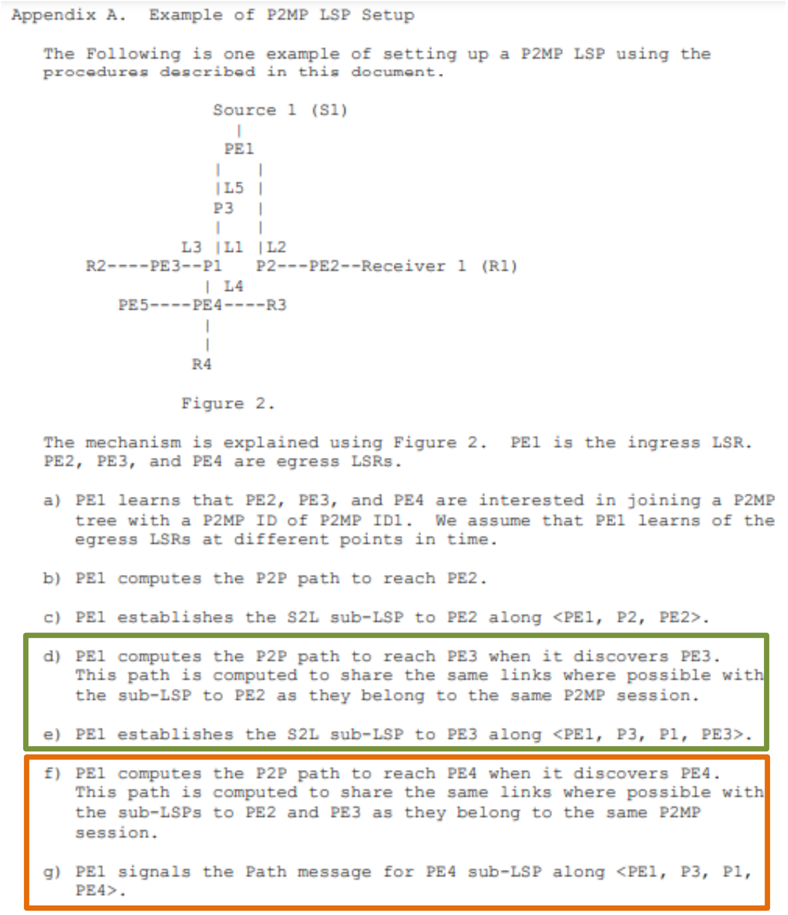

```
Appendix A.  Example of P2MP LSP Setup

   The Following is one example of setting up a P2MP LSP using the
   procedures described in this document.

                    Source 1 (S1)
                       |
                      PE1
                    |    |
                    |L5  |
                    P3   |
                    |    |
                 L3 |L1  |L2
       R2----PE3--P1    P2---PE2--Receiver 1 (R1)
                  | L4
          PE5----PE4----R3
                  |
                  |
                 R4

                 Figure 2.

   The mechanism is explained using Figure 2.  PE1 is the ingress LSR.
   PE2, PE3, and PE4 are egress LSRs.

   a) PE1 learns that PE2, PE3, and PE4 are interested in joining a P2MP
      tree with a P2MP ID of P2MP ID1.  We assume that PE1 learns of the
      egress LSRs at different points in time.

   b) PE1 computes the P2P path to reach PE2.

   c) PE1 establishes the S2L sub-LSP to PE2 along <PE1, P2, PE2>.

   d) PE1 computes the P2P path to reach PE3 when it discovers PE3.
      This path is computed to share the same links where possible with
      the sub-LSP to PE2 as they belong to the same P2MP session.

   e) PE1 establishes the S2L sub-LSP to PE3 along <PE1, P3, P1, PE3>.

   f) PE1 computes the P2P path to reach PE4 when it discovers PE4.
      This path is computed to share the same links where possible with
      the sub-LSPs to PE2 and PE3 as they belong to the same P2MP
      session.

   g) PE1 signals the Path message for PE4 sub-LSP along <PE1, P3, P1,
      PE4>.
```

https://www.rfc-editor.org/rfc/pdfrfc/rfc4875.txt.pdf (page 49)

78. Defendants instruct its customers to use the 7450 Ethernet Service Switch to wherein the tunnels meet at least one condition selected from a group of conditions consisting of:

the respective origin network elements of the first and second tunnels are different; and

the respective termination network elements of the first and second tunnels are different:

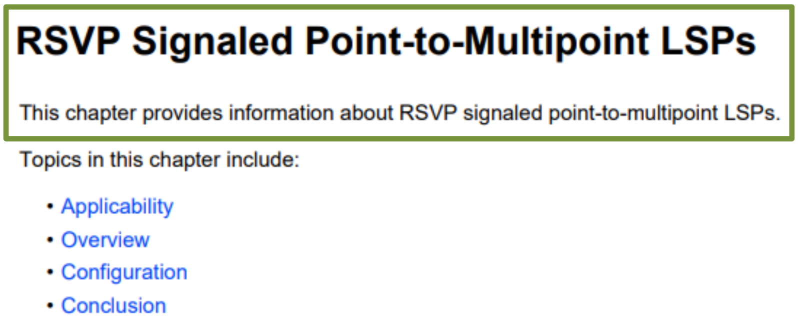

RSVP Signaled Point-to-Multipoint LSPs

This chapter provides information about RSVP signaled point-to-multipoint LSPs.

Topics in this chapter include:

- Applicability
- Overview
- Configuration
- Conclusion

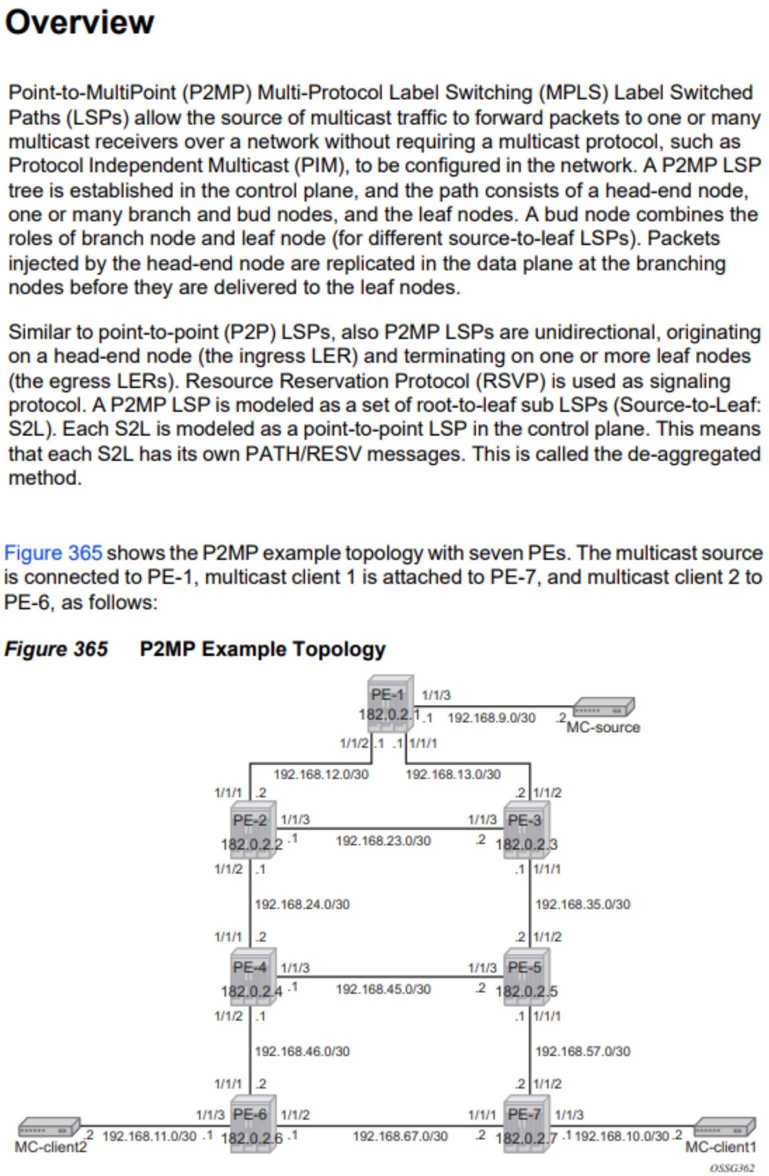

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

Figure 365 shows the P2MP example topology with seven PEs. The multicast source is connected to PE-1, multicast client 1 is attached to PE-7, and multicast client 2 to PE-6, as follows:

Figure 365 **P2MP Example Topology**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (pages 1693 and 1694)

Figure 368 **P2MP LSP LSP-p2mp-1 with Strict S2L Path toward PE-7**

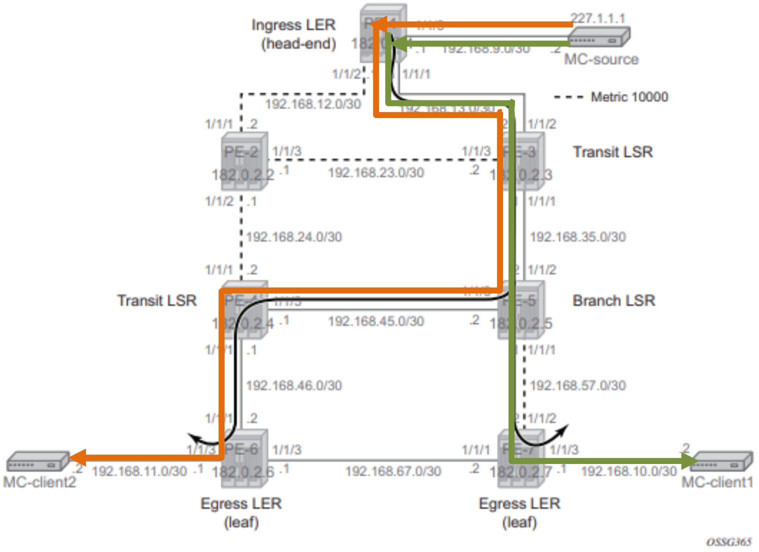

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1721)

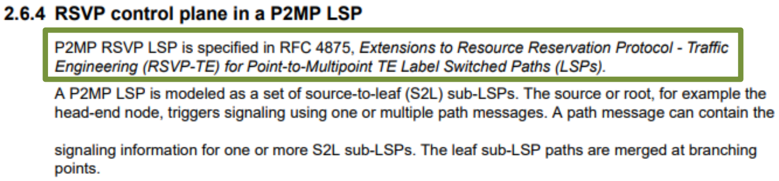

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs).*

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the

signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

79. Defendants instruct its customers to use the 7450 Ethernet Service Switch to distribute a notification over the network of an affiliation of the tunnels with the resource-sharing group:

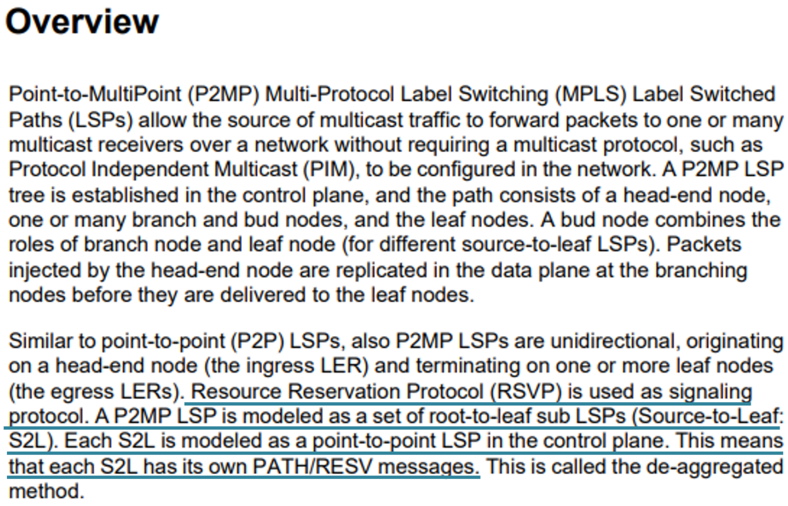

Overview

Point-to-MultiPoint (P2MP) Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) allow the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as Protocol Independent Multicast (PIM), to be configured in the network. A P2MP LSP tree is established in the control plane, and the path consists of a head-end node, one or many branch and bud nodes, and the leaf nodes. A bud node combines the roles of branch node and leaf node (for different source-to-leaf LSPs). Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

Similar to point-to-point (P2P) LSPs, also P2MP LSPs are unidirectional, originating on a head-end node (the ingress LER) and terminating on one or more leaf nodes (the egress LERs). Resource Reservation Protocol (RSVP) is used as signaling protocol. A P2MP LSP is modeled as a set of root-to-leaf sub LSPs (Source-to-Leaf: S2L). Each S2L is modeled as a point-to-point LSP in the control plane. This means that each S2L has its own PATH/RESV messages. This is called the de-aggregated method.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE14990AAAFTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20Advanced%20Configuration%20Guide%20for%20Releases%20up%20to%2021.5.R2-Part%20I.pdf (page 1693)

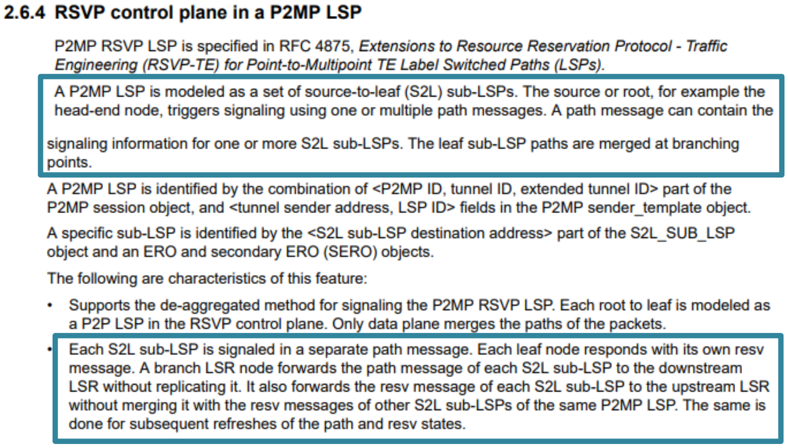

2.6.4 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*.

A P2MP LSP is modeled as a set of source-to-leaf (S2L) sub-LSPs. The source or root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

A P2MP LSP is identified by the combination of <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the P2MP sender_template object.

A specific sub-LSP is identified by the <S2L sub-LSP destination address> part of the S2L_SUB_LSP object and an ERO and secondary ERO (SERO) objects.

The following are characteristics of this feature:

- Supports the de-aggregated method for signaling the P2MP RSVP LSP. Each root to leaf is modeled as a P2P LSP in the RSVP control plane. Only data plane merges the paths of the packets.
- Each S2L sub-LSP is signaled in a separate path message. Each leaf node responds with its own resv message. A branch LSR node forwards the path message of each S2L sub-LSP to the downstream LSR without replicating it. It also forwards the resv message of each S2L sub-LSP to the upstream LSR without merging it with the resv messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and resv states.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (pages 162 and 163)

80. Defendants instruct its customers to use the 7450 Ethernet Service Switch to allocate a resource associated with the at least one common network element so as to share an allocation of the resource among the tunnels in the resource-sharing group responsively to the notification.

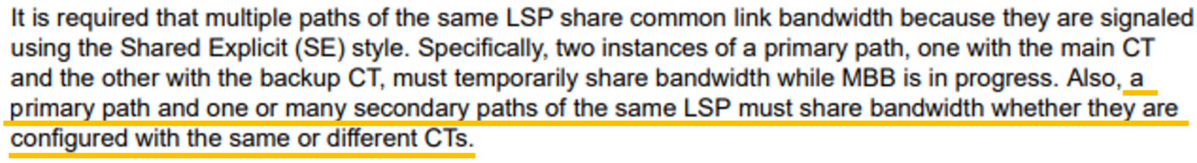

It is required that multiple paths of the same LSP share common link bandwidth because they are signaled using the Shared Explicit (SE) style. Specifically, two instances of a primary path, one with the main CT and the other with the backup CT, must temporarily share bandwidth while MBB is in progress. Also, a primary path and one or many secondary paths of the same LSP must share bandwidth whether they are configured with the same or different CTs.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17154AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20MPLS%20Guide%2021.10.R1.pdf (page 118)

81. As a result of Defendants' infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT THREE
INFRINGEMENT OF U.S. PATENT 7,551,599

82. Plaintiff incorporates by reference the allegations in all preceding paragraphs as if fully set forth herein.

83. The '599 Patent, entitled "Layer-3 network routing with RPR layer-2 visibility" was filed on March 29, 2004, and issued on June 23, 2009.

84. Plaintiff is the assignee and owner of all rights, title and interest to the '599 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

85. The '599 Patent addresses technical problems in the prior art, including that "[c]urrently, layer-3 routing protocols, such as RIP and OSPF, are unaware of the topology of layer-2 RPR networks with which they must interact. A routing table allows the router to forward packets from source to destination via the most suitable path, i.e., lowest cost, minimum number

of hops. The routing table is updated via the routing protocol, which dynamically discovers currently available paths. The routing table may also be updated via static routes, or can be built using a local interface configuration, which is updated by the network administrator. However, the RPR ingress and egress nodes chosen in the operation of automatic routing protocols do not take into account the internal links within the RPR ring and may therefore cause load imbalances within the RPR subnet, which generally results in suboptimum performance of the larger network." 3:65-4:12.

86. To address these issues, the '599 Patent discloses "methods and systems are provided for the manipulation of layer-3 network nodes, external routers, routing tables and elements of layer-2 ring networks, such as RPR networks, enabling the layer-3 elements to view the topology of a layer-2 ring subnet. This feature permits routers to choose optimal entry points to the layer-2 subnet for different routes that pass into or through the layer-2 subnet. This enables virtual tunnels or routing paths to utilize all existing entry links to the subnet and to minimize cost factors, such as the number of spans required to traverse the subnet from the entry point to a destination node of the subnet." 4:17-27.

Direct Infringement

87. Defendants, without authorization or license from Plaintiff, have been and are directly infringing the '599 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale methods, devices, and networks infringing one or more claims of the '599 Patent. Defendants develop, design, manufacture, and distribute telecommunications equipment that infringes one or more claims of the '599 Patent. Defendants further provide services that practice methods that infringe one or more claims of the '599 Patent. Defendants are thus liable

for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Defendants' 7450 Ethernet Service Switch, and all other substantially similar products (collectively the "'599 Accused Products").

88. Smart Path names these exemplary infringing instrumentalities to serve as notice of Defendants' infringing acts, however Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '599 Accused Products.

89. Defendants are liable for direct infringement pursuant to 35 U.S.C. § 271 for the use, manufacture, sale, offer of sale, importation, or distribution of Defendants' 7450 Ethernet Service Switch.

90. Defendants' 7450 Ethernet Service Switch provides a platform for an aggregation of devices and is a non-limiting example of an apparatus that meets all limitations of claim 71 of the '599 Patent, either literally or equivalently.

91. Defendants' 7450 Ethernet Service Switch provides a platform for an aggregation of devices to form a network routing system for obtaining egress from a layer-2 ring network to an external layer-3 network comprising:

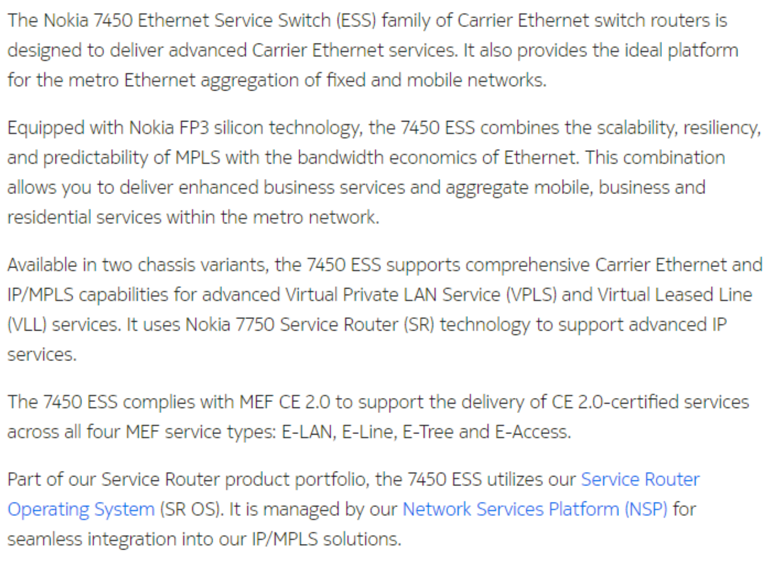

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and residential services within the metro network.

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

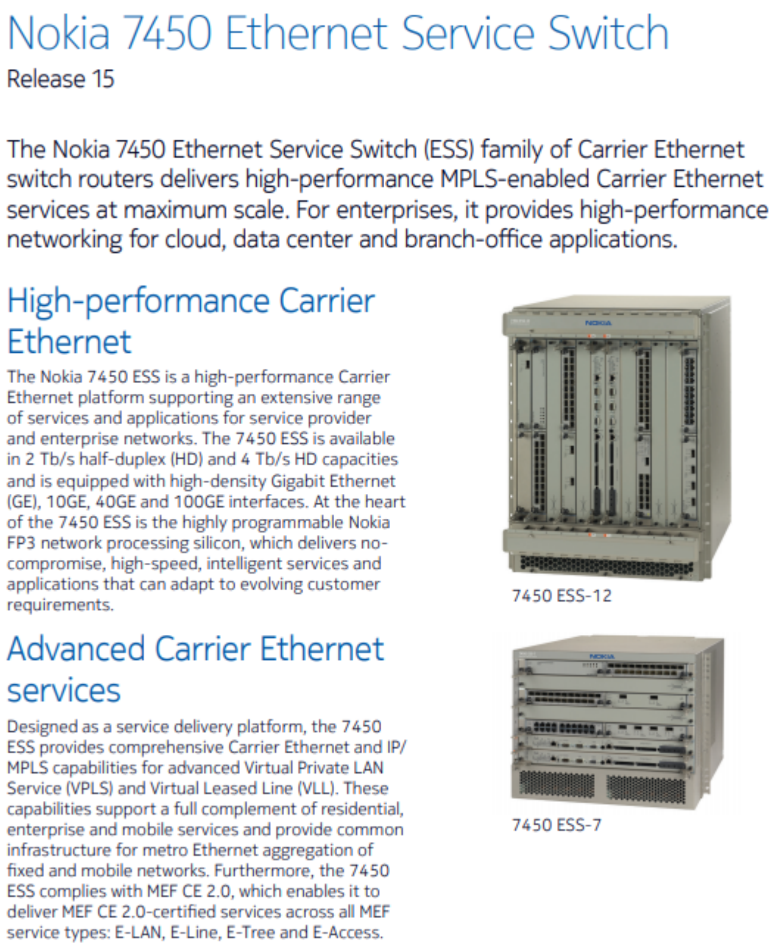

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/164727 (Page 1)

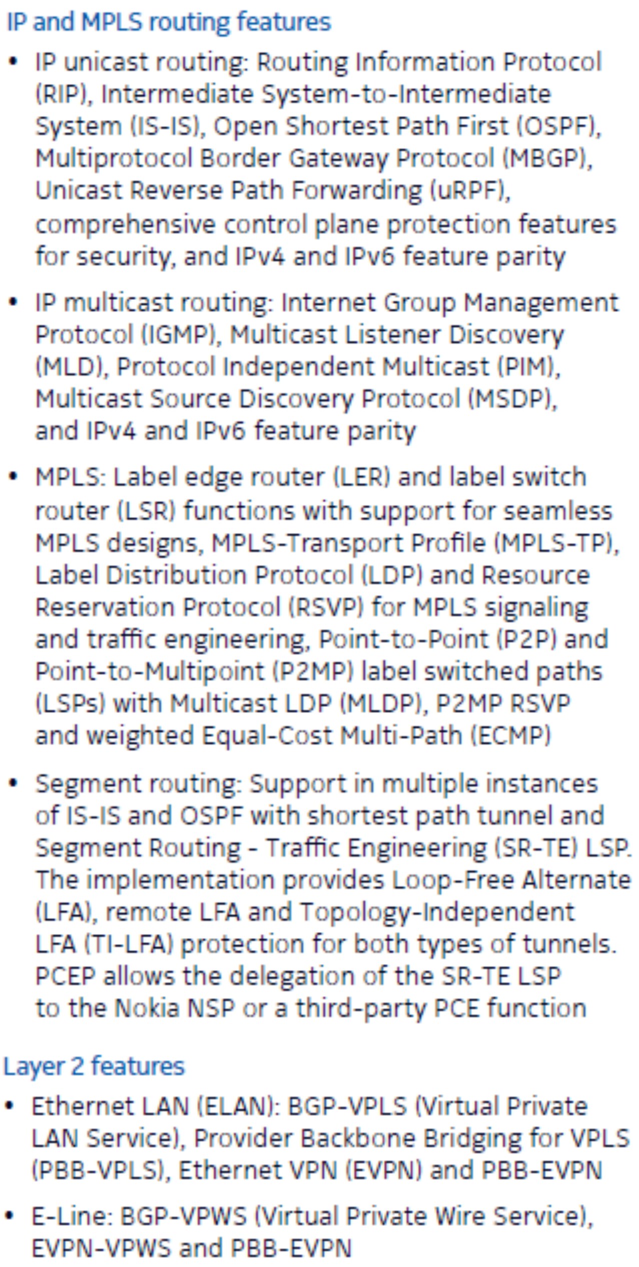

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity
- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)
- Segment routing: Support in multiple instances of IS-IS and OSPF with shortest path tunnel and Segment Routing - Traffic Engineering (SR-TE) LSP. The implementation provides Loop-Free Alternate (LFA), remote LFA and Topology-Independent LFA (TI-LFA) protection for both types of tunnels. PCEP allows the delegation of the SR-TE LSP to the Nokia NSP or a third-party PCE function

Layer 2 features

- Ethernet LAN (ELAN): BGP-VPLS (Virtual Private LAN Service), Provider Backbone Bridging for VPLS (PBB-VPLS), Ethernet VPN (EVPN) and PBB-EVPN
- E-Line: BGP-VPWS (Virtual Private Wire Service), EVPN-VPWS and PBB-EVPN

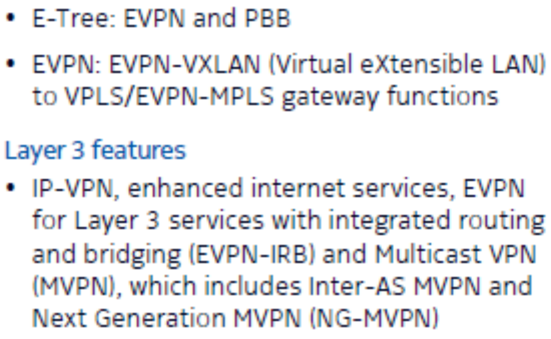

- E-Tree: EVPN and PBB
- EVPN: EVPN-VXLAN (Virtual eXtensible LAN) to VPLS/EVPN-MPLS gateway functions

Layer 3 features

- IP-VPN, enhanced internet services, EVPN for Layer 3 services with integrated routing and bridging (EVPN-IRB) and Multicast VPN (MVPN), which includes Inter-AS MVPN and Next Generation MVPN (NG-MVPN) https://onestore.nokia.com/asset/164727 (Page 6 of PDF)

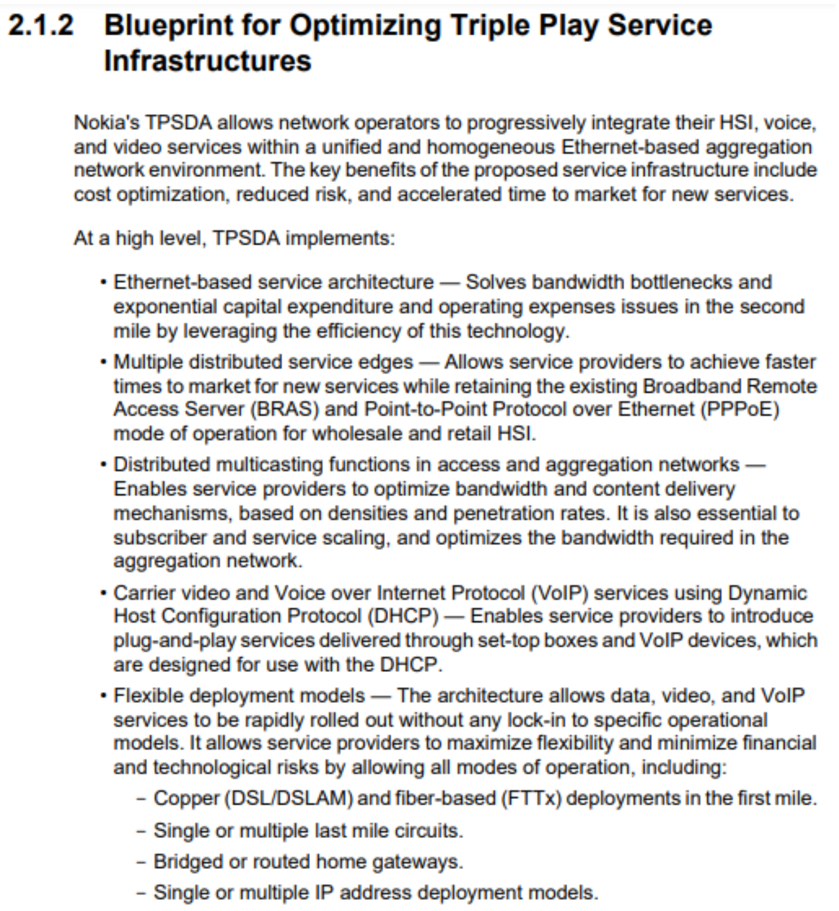

2.1.2 Blueprint for Optimizing Triple Play Service Infrastructures

Nokia's TPSDA allows network operators to progressively integrate their HSI, voice, and video services within a unified and homogeneous Ethernet-based aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time to market for new services.

At a high level, TPSDA implements:

- Ethernet-based service architecture — Solves bandwidth bottlenecks and exponential capital expenditure and operating expenses issues in the second mile by leveraging the efficiency of this technology.
- Multiple distributed service edges — Allows service providers to achieve faster times to market for new services while retaining the existing Broadband Remote Access Server (BRAS) and Point-to-Point Protocol over Ethernet (PPPoE) mode of operation for wholesale and retail HSI.
- Distributed multicasting functions in access and aggregation networks — Enables service providers to optimize bandwidth and content delivery mechanisms, based on densities and penetration rates. It is also essential to subscriber and service scaling, and optimizes the bandwidth required in the aggregation network.
- Carrier video and Voice over Internet Protocol (VoIP) services using Dynamic Host Configuration Protocol (DHCP) — Enables service providers to introduce plug-and-play services delivered through set-top boxes and VoIP devices, which are designed for use with the DHCP.
- Flexible deployment models — The architecture allows data, video, and VoIP services to be rapidly rolled out without any lock-in to specific operational models. It allows service providers to maximize flexibility and minimize financial and technological risks by allowing all modes of operation, including:
  - Copper (DSL/DSLAM) and fiber-based (FTTx) deployments in the first mile.
  - Single or multiple last mile circuits.
  - Bridged or routed home gateways.
  - Single or multiple IP address deployment models.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 30)

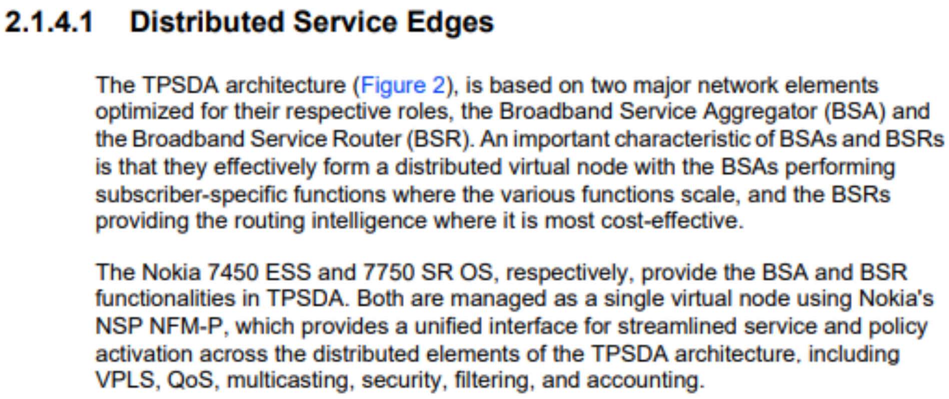

2.1.4.1 Distributed Service Edges

The TPSDA architecture (Figure 2), is based on two major network elements optimized for their respective roles, the Broadband Service Aggregator (BSA) and the Broadband Service Router (BSR). An important characteristic of BSAs and BSRs is that they effectively form a distributed virtual node with the BSAs performing subscriber-specific functions where the various functions scale, and the BSRs providing the routing intelligence where it is most cost-effective.

The Nokia 7450 ESS and 7750 SR OS, respectively, provide the BSA and BSR functionalities in TPSDA. Both are managed as a single virtual node using Nokia's NSP NFM-P, which provides a unified interface for streamlined service and policy activation across the distributed elements of the TPSDA architecture, including VPLS, QoS, multicasting, security, filtering, and accounting.

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

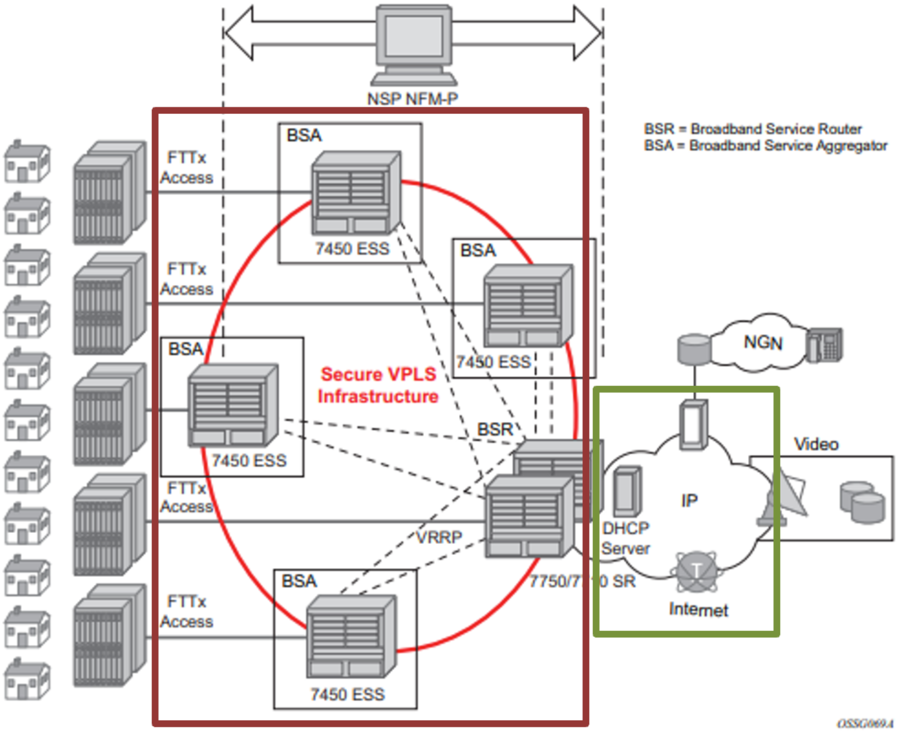

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34-35)

92. Defendants' 7450 Ethernet Service Switch provides a platform for an aggregation of devices, which comprises a plurality of routers disposed in nodes of said ring network, said routers being configured for creating entries in a host table, each of said entries comprising an address of a respective one of said nodes of said ring network and a metric:

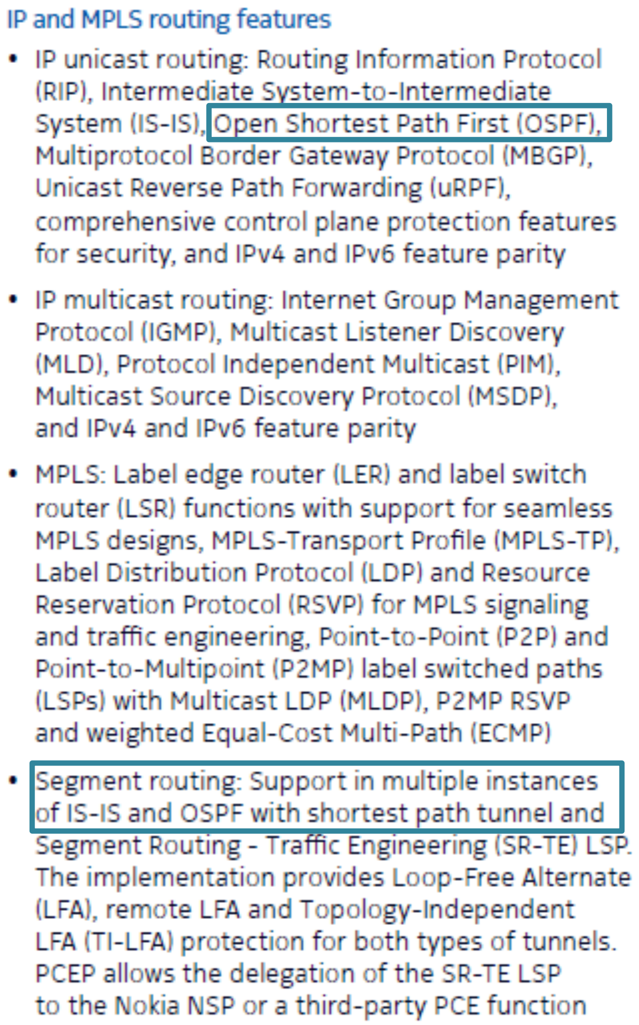

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity
- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)
- Segment routing: Support in multiple instances of IS-IS and OSPF with shortest path tunnel and Segment Routing - Traffic Engineering (SR-TE) LSP. The implementation provides Loop-Free Alternate (LFA), remote LFA and Topology-Independent LFA (TI-LFA) protection for both types of tunnels. PCEP allows the delegation of the SR-TE LSP to the Nokia NSP or a third-party PCE function

https://onestore.nokia.com/asset/164727 (Page 6)

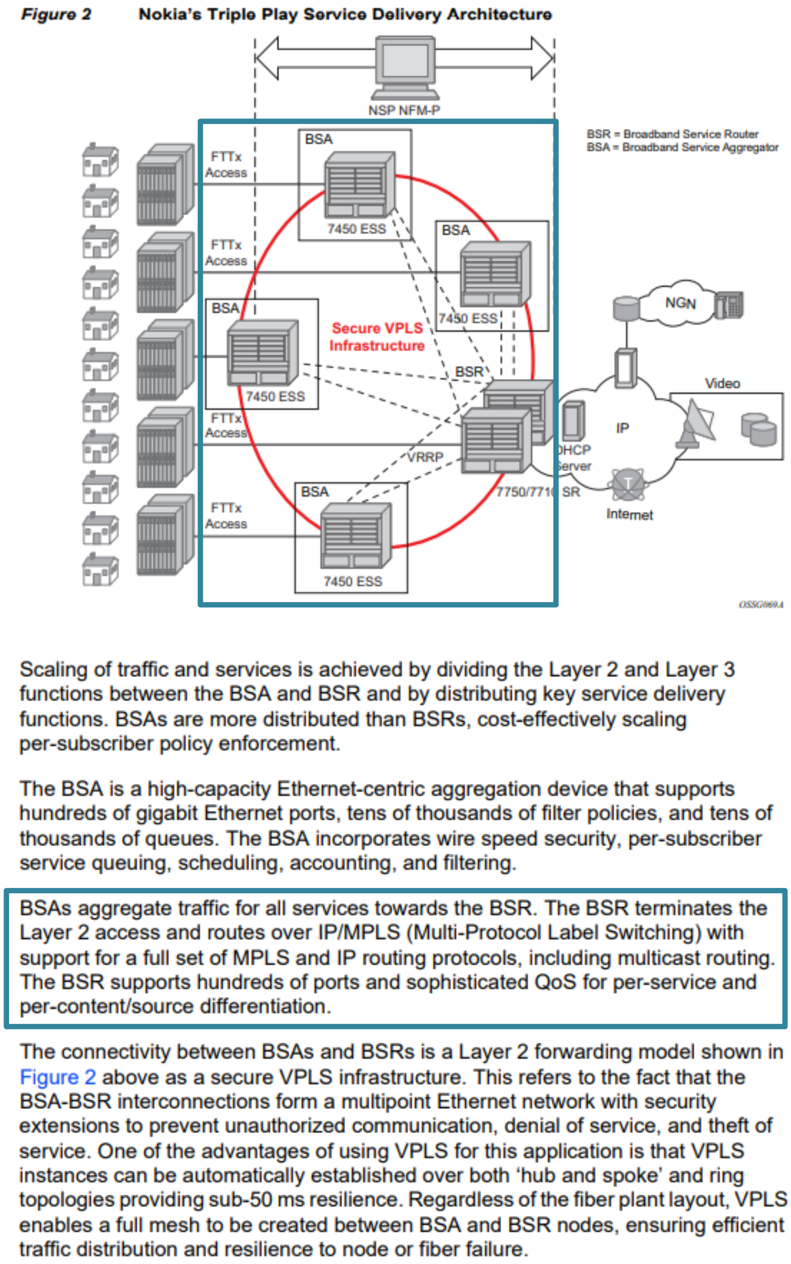

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34 and 35)

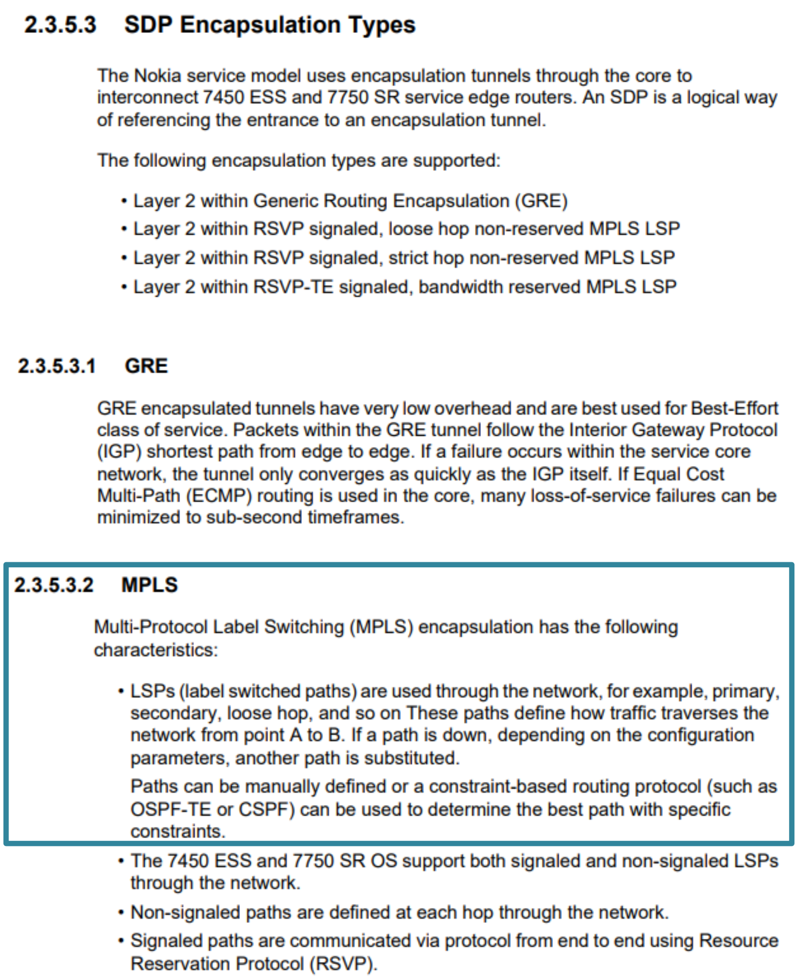

2.3.5.3 SDP Encapsulation Types

The Nokia service model uses encapsulation tunnels through the core to interconnect 7450 ESS and 7750 SR service edge routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- Layer 2 within Generic Routing Encapsulation (GRE)
- Layer 2 within RSVP signaled, loose hop non-reserved MPLS LSP
- Layer 2 within RSVP signaled, strict hop non-reserved MPLS LSP
- Layer 2 within RSVP-TE signaled, bandwidth reserved MPLS LSP

2.3.5.3.1 GRE

GRE encapsulated tunnels have very low overhead and are best used for Best-Effort class of service. Packets within the GRE tunnel follow the Interior Gateway Protocol (IGP) shortest path from edge to edge. If a failure occurs within the service core network, the tunnel only converges as quickly as the IGP itself. If Equal Cost Multi-Path (ECMP) routing is used in the core, many loss-of-service failures can be minimized to sub-second timeframes.

2.3.5.3.2 MPLS

Multi-Protocol Label Switching (MPLS) encapsulation has the following characteristics:

- LSPs (label switched paths) are used through the network, for example, primary, secondary, loose hop, and so on These paths define how traffic traverses the network from point A to B. If a path is down, depending on the configuration parameters, another path is substituted.

  Paths can be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints.
- The 7450 ESS and 7750 SR OS support both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.
- Signaled paths are communicated via protocol from end to end using Resource Reservation Protocol (RSVP).

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 51)

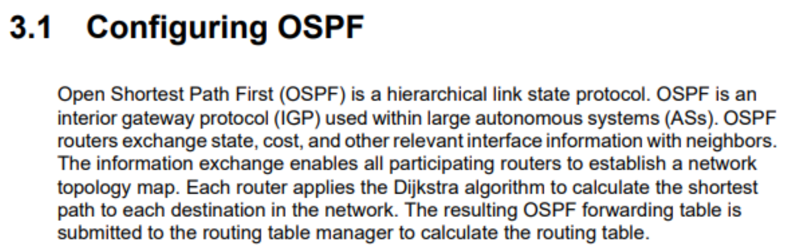

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

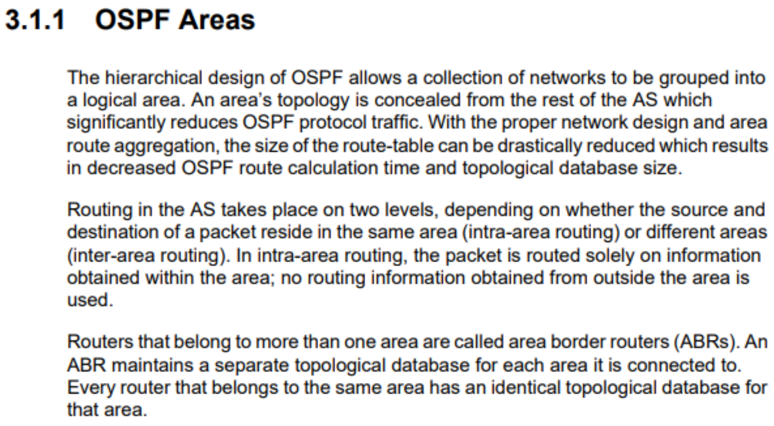

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

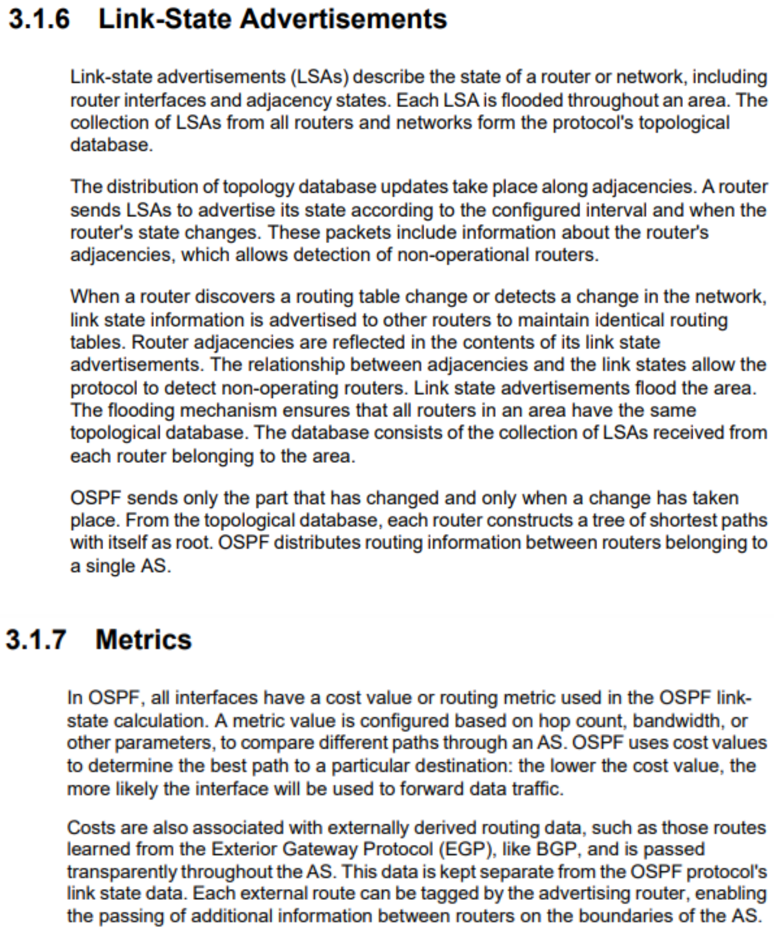

3.1.6 Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48**)**

93. Defendants' 7450 Ethernet Service Switch comprises routers being further configured for defining paths from said nodes through egress nodes of said ring network:

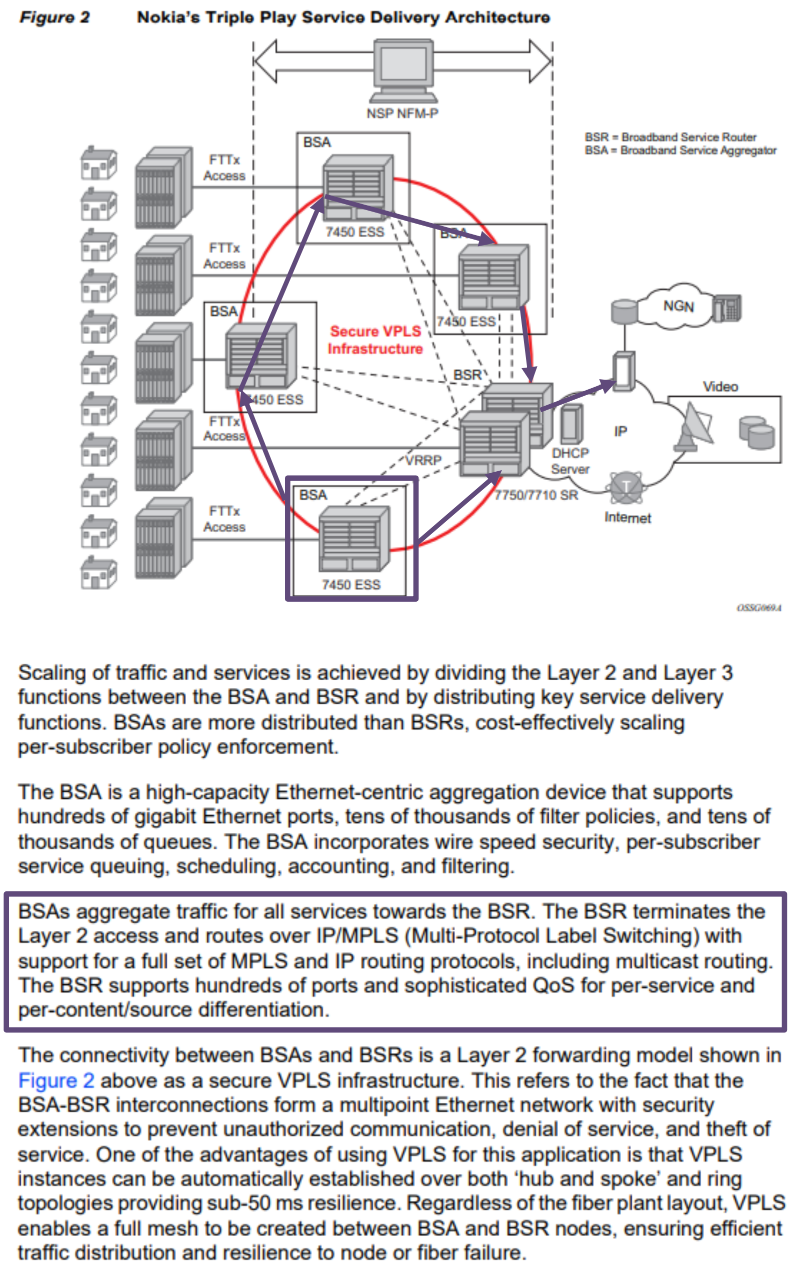

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34-35)

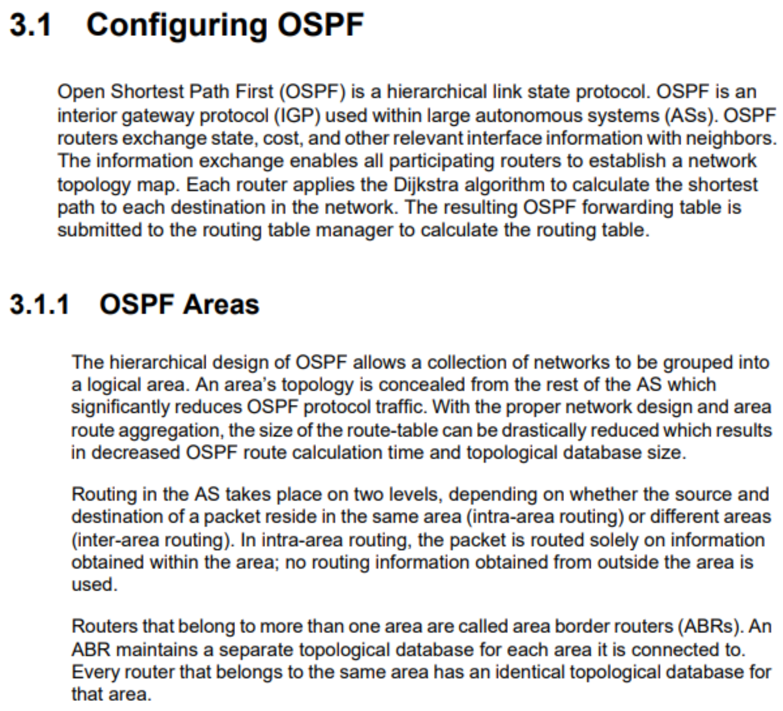

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

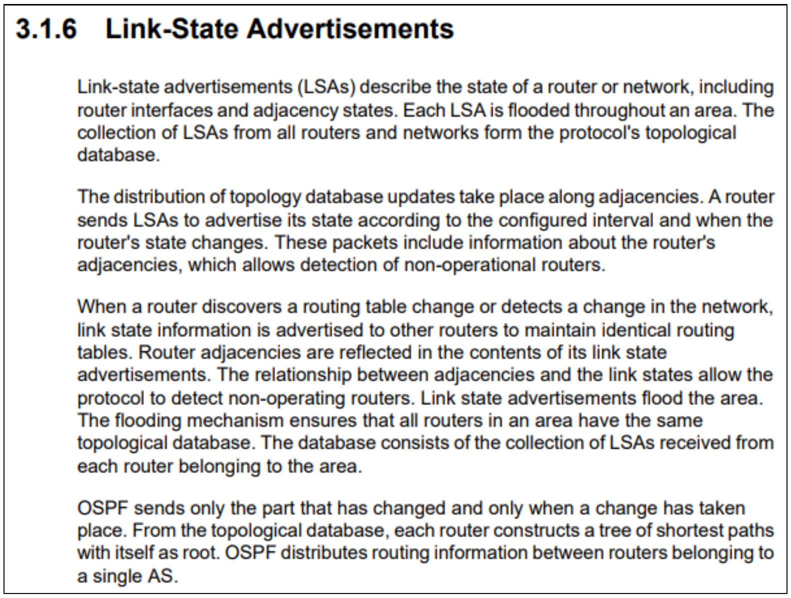

3.1.6 Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

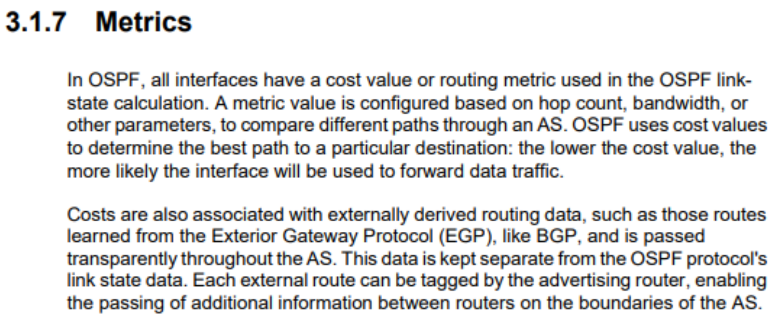

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48**)**

94. Defendants' 7450 Ethernet Service Switch comprises routers for selecting one of said paths responsively to said metric:

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

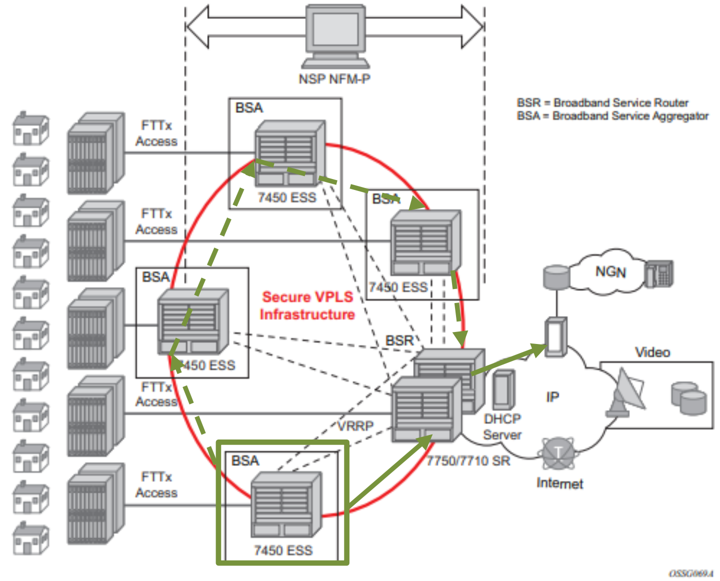

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 34)

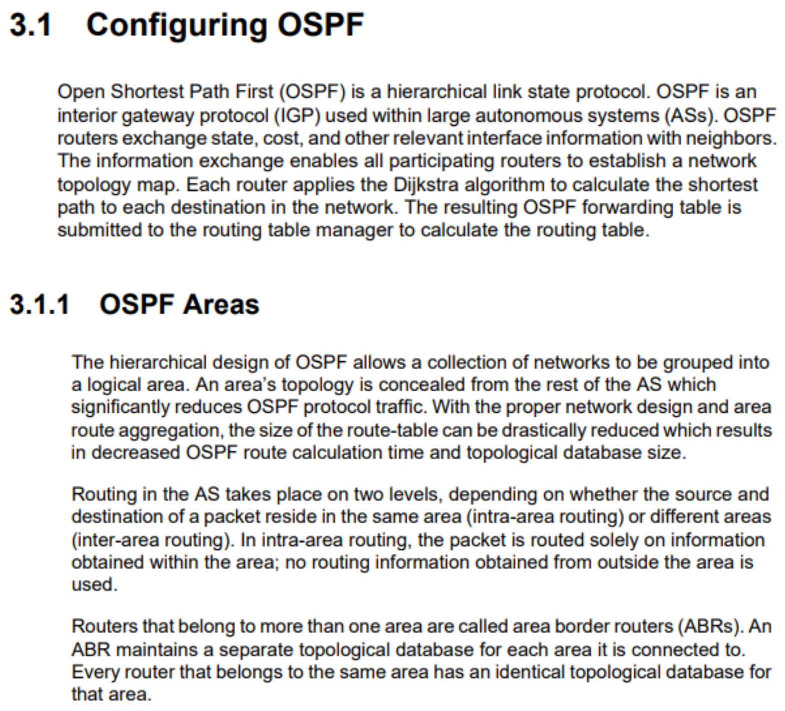

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

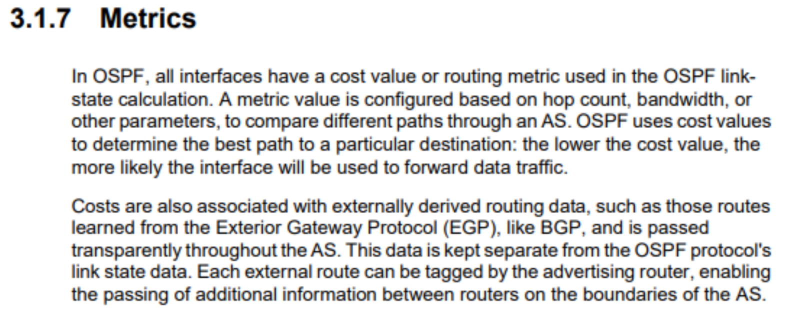

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48)

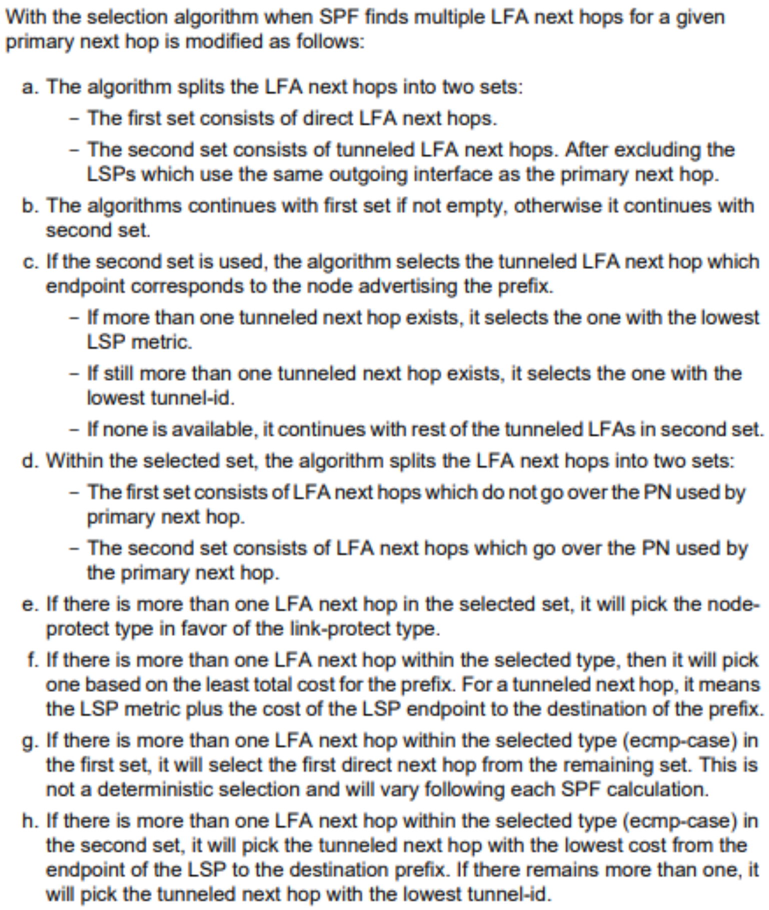

With the selection algorithm when SPF finds multiple LFA next hops for a given primary next hop is modified as follows:

a. The algorithm splits the LFA next hops into two sets:
   - The first set consists of direct LFA next hops.
   - The second set consists of tunneled LFA next hops. After excluding the LSPs which use the same outgoing interface as the primary next hop.

b. The algorithms continues with first set if not empty, otherwise it continues with second set.

c. If the second set is used, the algorithm selects the tunneled LFA next hop which endpoint corresponds to the node advertising the prefix.
   - If more than one tunneled next hop exists, it selects the one with the lowest LSP metric.
   - If still more than one tunneled next hop exists, it selects the one with the lowest tunnel-id.
   - If none is available, it continues with rest of the tunneled LFAs in second set.

d. Within the selected set, the algorithm splits the LFA next hops into two sets:
   - The first set consists of LFA next hops which do not go over the PN used by primary next hop.
   - The second set consists of LFA next hops which go over the PN used by the primary next hop.

e. If there is more than one LFA next hop in the selected set, it will pick the node-protect type in favor of the link-protect type.

f. If there is more than one LFA next hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.

g. If there is more than one LFA next hop within the selected type (ecmp-case) in the first set, it will select the first direct next hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.

h. If there is more than one LFA next hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next hop with the lowest tunnel-id.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Page 58)

95. Defendants' 7450 Ethernet Service Switch comprises routers for transmitting data from said nodes via said selected paths to network elements that are external to said ring network:

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

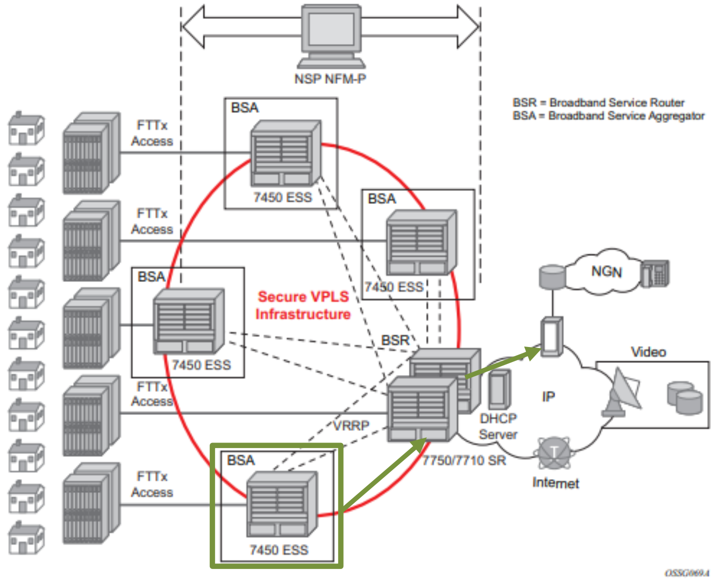

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 34**)**

Willful Infringement

96. Defendants have had actual knowledge of the '599 Patent and its infringement thereof at least as of service or other receipt of Plaintiff's Complaint.

97. Defendants' risk of infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendants.

98. Notwithstanding this knowledge, Defendants have knowingly or with reckless disregard willfully infringed the '599 Patent. Defendants have thus had actual notice of the infringement of the '599 Patent and acted despite an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

99. This objective risk was either known or so obvious that it should have been known to Defendants. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

100. Defendants have induced and are knowingly inducing its customers and/or end users to directly infringe the '599 Patent, with the specific intent to encourage such infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

101. Defendants have knowingly contributed to direct infringement by its customers by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the '599 Accused Products which are not suitable for substantial non-infringing use, and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

102. Defendants' indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe '599 Patent.

103. Defendants' indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendants' indirect infringement further includes providing application notes instructing its customers on infringing uses of the accused products. The '599 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '599 Patent, either literally or equivalently. Defendants know and intend that customers who purchase the '599 Accused Products will use those products for their intended purpose. For example, Defendants' United States website https://www.nokia.com, instructs customers to use the '599 Accused Products in numerous infringing applications. Furthermore, Defendants provide instructions and other resources via its "Doc Center" (https://documentation.nokia.com), "Nokia Service Routing Certification" and training courses from its "NokiaEDU Training Centers" (https://www.nokia.com/networks/

training/src/courses/#open-enrollment), and elsewhere on using the ’599 Accused Products. Defendants’ customers directly infringe the ’599 Patent when they follow Defendants’ provided instructions on websites, videos, trainings, and elsewhere. Defendants’ customers who follow Defendants’ provided instructions directly infringe claims of the ’599 Patent.

104. In addition, Defendants specifically intend that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendants’ infringing products. Defendants know following its instructions directly infringes claims of the ’599 Patent, including for example Claim 47.

105. Defendants’ customers who follow Defendants’ provided instructions directly infringe the method of Claim 47 of the ’599 Patent.

106. Defendants instruct its customers to use the 7450 Ethernet Service Switch in a method for obtaining egress from a layer-2 ring network to an external layer-3 network:

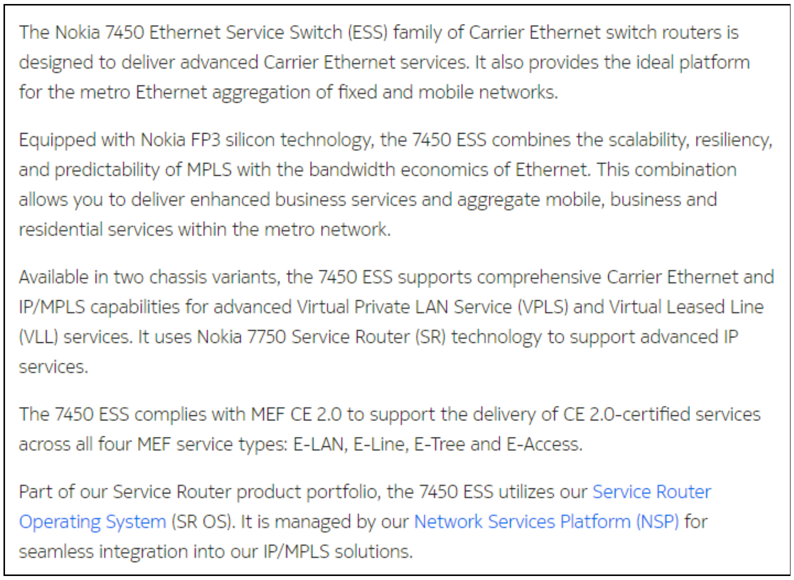

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and residential services within the metro network.

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

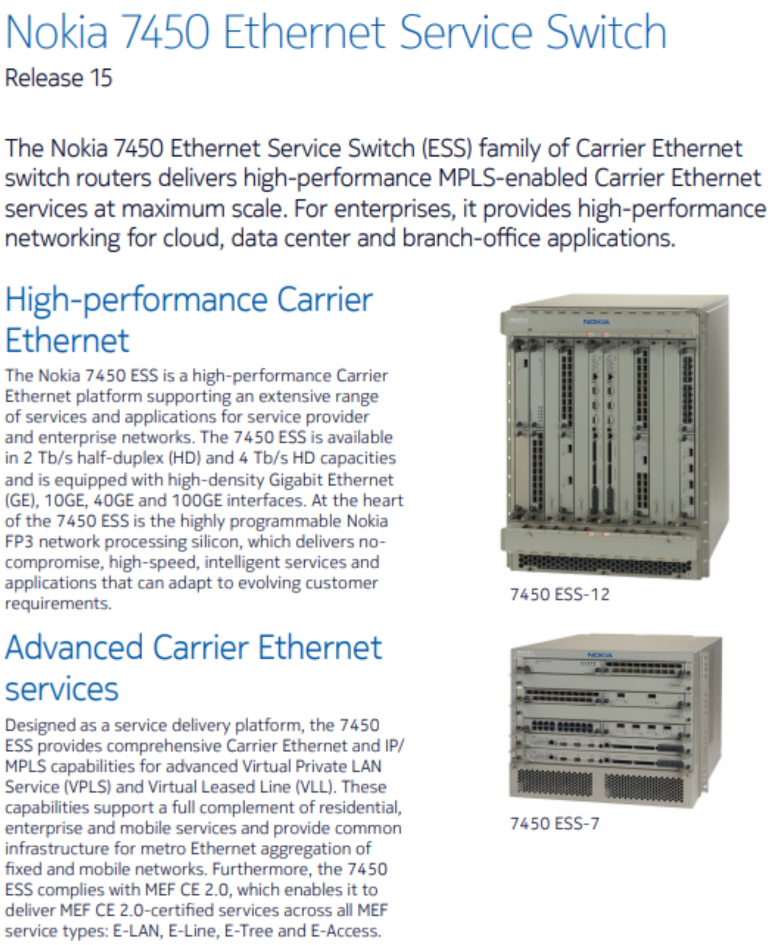

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/164727 (Page 1)

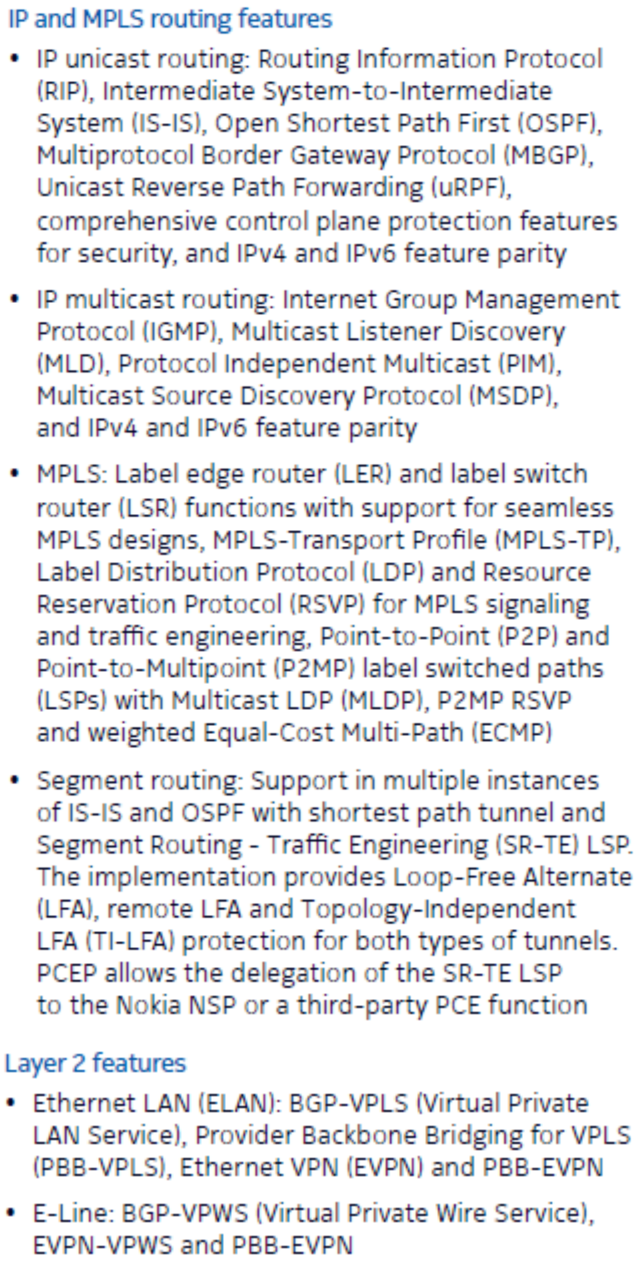

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity
- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)
- Segment routing: Support in multiple instances of IS-IS and OSPF with shortest path tunnel and Segment Routing - Traffic Engineering (SR-TE) LSP. The implementation provides Loop-Free Alternate (LFA), remote LFA and Topology-Independent LFA (TI-LFA) protection for both types of tunnels. PCEP allows the delegation of the SR-TE LSP to the Nokia NSP or a third-party PCE function

Layer 2 features

- Ethernet LAN (ELAN): BGP-VPLS (Virtual Private LAN Service), Provider Backbone Bridging for VPLS (PBB-VPLS), Ethernet VPN (EVPN) and PBB-EVPN
- E-Line: BGP-VPWS (Virtual Private Wire Service), EVPN-VPWS and PBB-EVPN

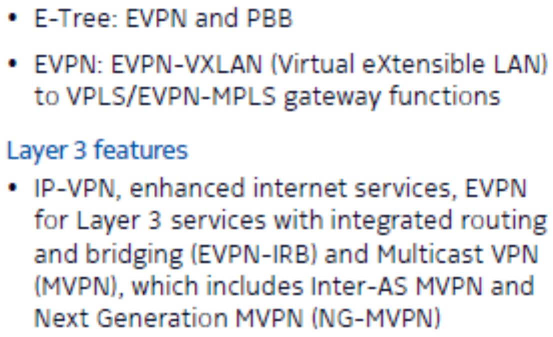

- E-Tree: EVPN and PBB
- EVPN: EVPN-VXLAN (Virtual eXtensible LAN) to VPLS/EVPN-MPLS gateway functions

Layer 3 features

- IP-VPN, enhanced internet services, EVPN for Layer 3 services with integrated routing and bridging (EVPN-IRB) and Multicast VPN (MVPN), which includes Inter-AS MVPN and Next Generation MVPN (NG-MVPN) https://onestore.nokia.com/asset/164727 (Page 6 of PDF)

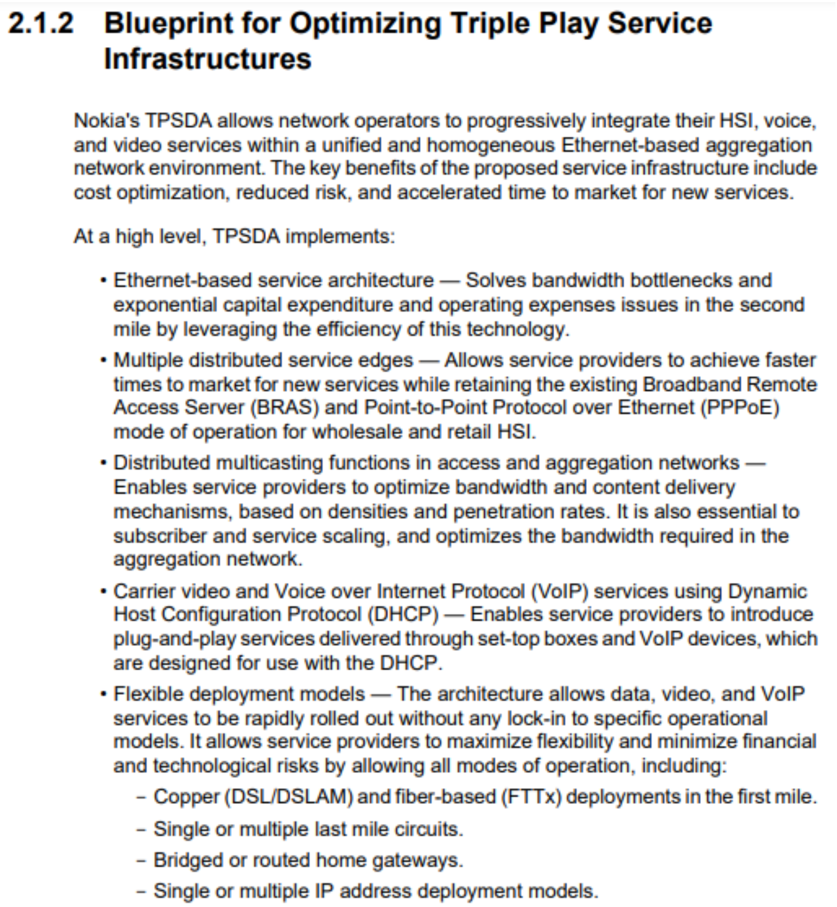

2.1.2 Blueprint for Optimizing Triple Play Service Infrastructures

Nokia's TPSDA allows network operators to progressively integrate their HSI, voice, and video services within a unified and homogeneous Ethernet-based aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time to market for new services.

At a high level, TPSDA implements:

- Ethernet-based service architecture — Solves bandwidth bottlenecks and exponential capital expenditure and operating expenses issues in the second mile by leveraging the efficiency of this technology.
- Multiple distributed service edges — Allows service providers to achieve faster times to market for new services while retaining the existing Broadband Remote Access Server (BRAS) and Point-to-Point Protocol over Ethernet (PPPoE) mode of operation for wholesale and retail HSI.
- Distributed multicasting functions in access and aggregation networks — Enables service providers to optimize bandwidth and content delivery mechanisms, based on densities and penetration rates. It is also essential to subscriber and service scaling, and optimizes the bandwidth required in the aggregation network.
- Carrier video and Voice over Internet Protocol (VoIP) services using Dynamic Host Configuration Protocol (DHCP) — Enables service providers to introduce plug-and-play services delivered through set-top boxes and VoIP devices, which are designed for use with the DHCP.
- Flexible deployment models — The architecture allows data, video, and VoIP services to be rapidly rolled out without any lock-in to specific operational models. It allows service providers to maximize flexibility and minimize financial and technological risks by allowing all modes of operation, including:
  - Copper (DSL/DSLAM) and fiber-based (FTTx) deployments in the first mile.
  - Single or multiple last mile circuits.
  - Bridged or routed home gateways.
  - Single or multiple IP address deployment models.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 30)

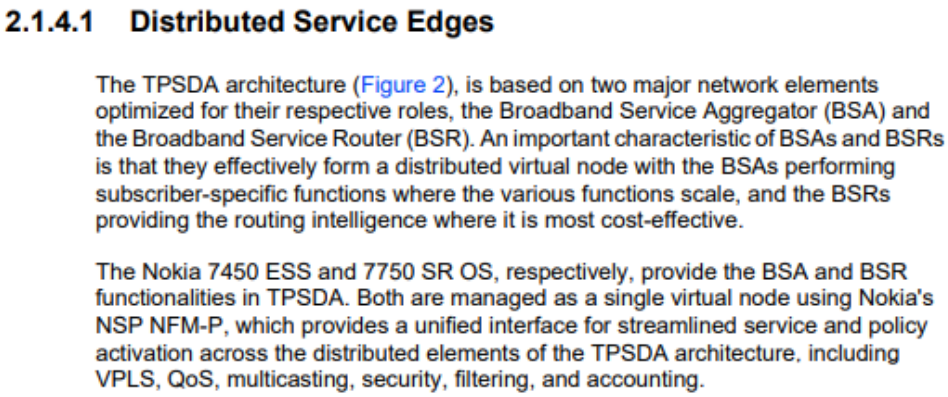

2.1.4.1 Distributed Service Edges

The TPSDA architecture (Figure 2), is based on two major network elements optimized for their respective roles, the Broadband Service Aggregator (BSA) and the Broadband Service Router (BSR). An important characteristic of BSAs and BSRs is that they effectively form a distributed virtual node with the BSAs performing subscriber-specific functions where the various functions scale, and the BSRs providing the routing intelligence where it is most cost-effective.

The Nokia 7450 ESS and 7750 SR OS, respectively, provide the BSA and BSR functionalities in TPSDA. Both are managed as a single virtual node using Nokia's NSP NFM-P, which provides a unified interface for streamlined service and policy activation across the distributed elements of the TPSDA architecture, including VPLS, QoS, multicasting, security, filtering, and accounting.

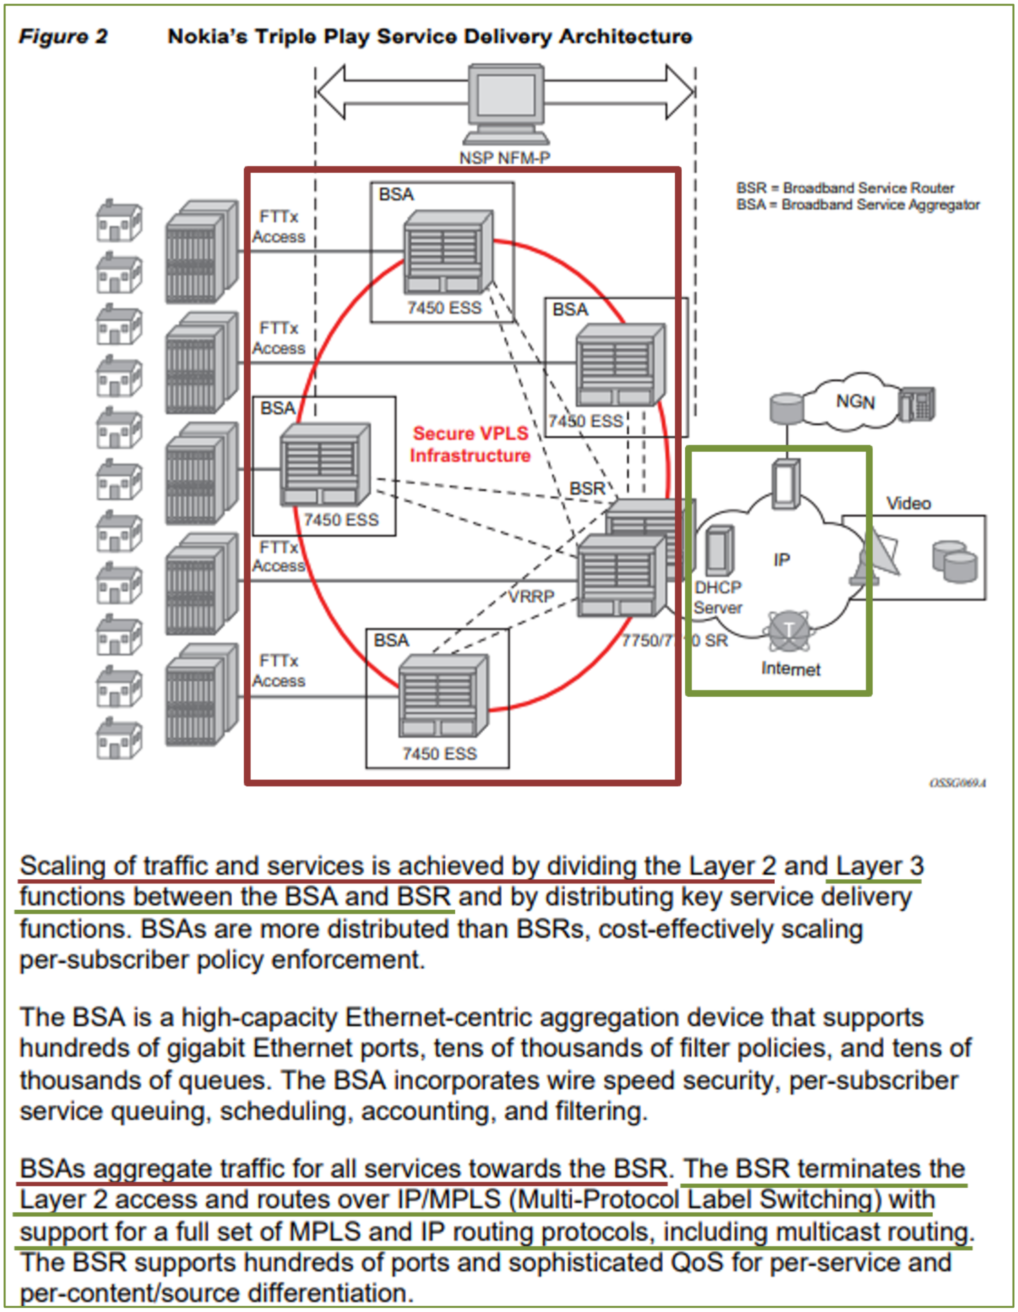

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

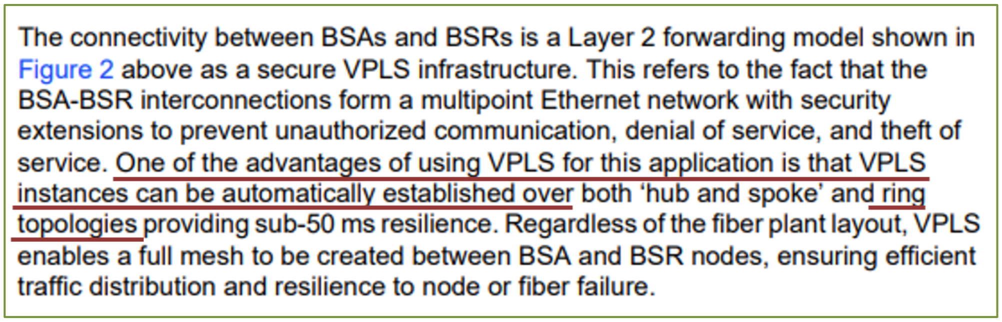

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34-35)

107. Defendants instruct its customers to use the 7450 Ethernet Service Switch, in nodes of said ring network creating entries in a host table, each of said entries comprising an address of a respective one of said nodes of said ring network and a metric determined responsively to a topology of the ring network.

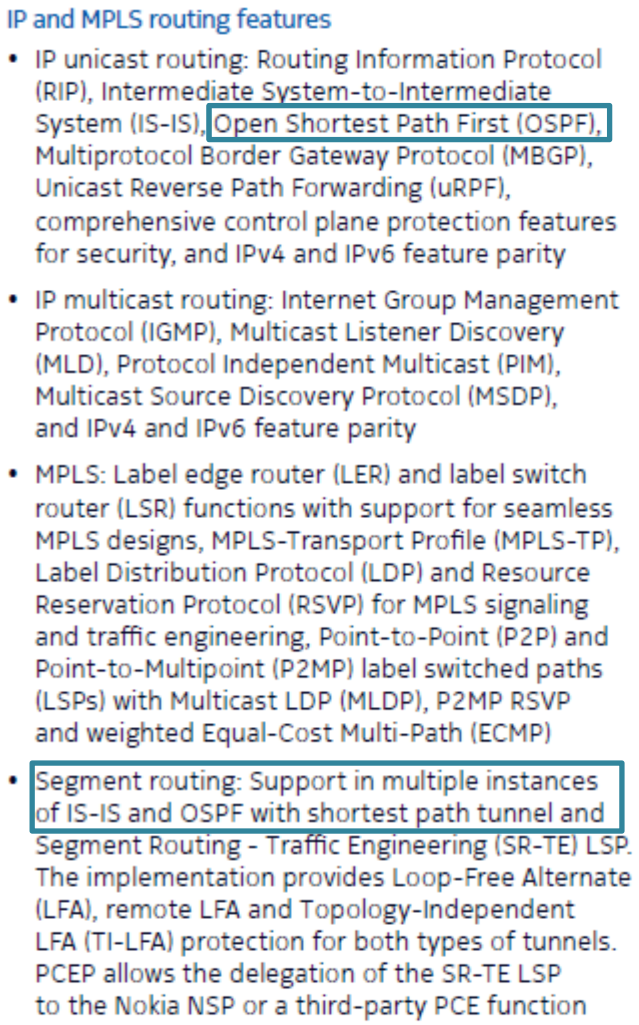

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity
- MPLS: Label edge router (LER) and label switch router (LSR) functions with support for seamless MPLS designs, MPLS-Transport Profile (MPLS-TP), Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) for MPLS signaling and traffic engineering, Point-to-Point (P2P) and Point-to-Multipoint (P2MP) label switched paths (LSPs) with Multicast LDP (MLDP), P2MP RSVP and weighted Equal-Cost Multi-Path (ECMP)
- Segment routing: Support in multiple instances of IS-IS and OSPF with shortest path tunnel and Segment Routing - Traffic Engineering (SR-TE) LSP. The implementation provides Loop-Free Alternate (LFA), remote LFA and Topology-Independent LFA (TI-LFA) protection for both types of tunnels. PCEP allows the delegation of the SR-TE LSP to the Nokia NSP or a third-party PCE function

https://onestore.nokia.com/asset/164727 (Page 6)

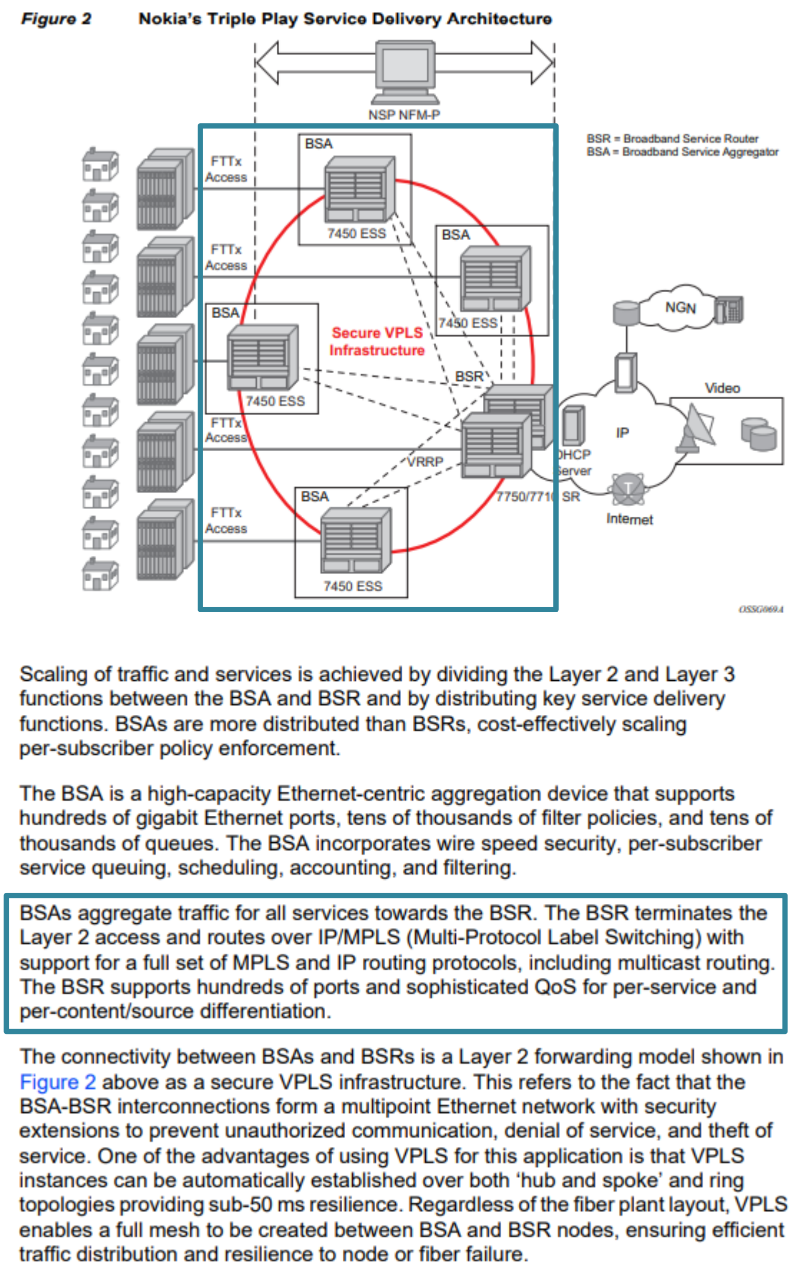

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

NSP NFM-P
BSR = Broadband Service Router
BSA = Broadband Service Aggregator
BSA
FTTx Access
7450 ESS
Secure VPLS Infrastructure
BSR
VRRP
7750/7710 SR
DHCP Server
NGN
IP
Video
Internet
OSSG069A

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34 and 35)

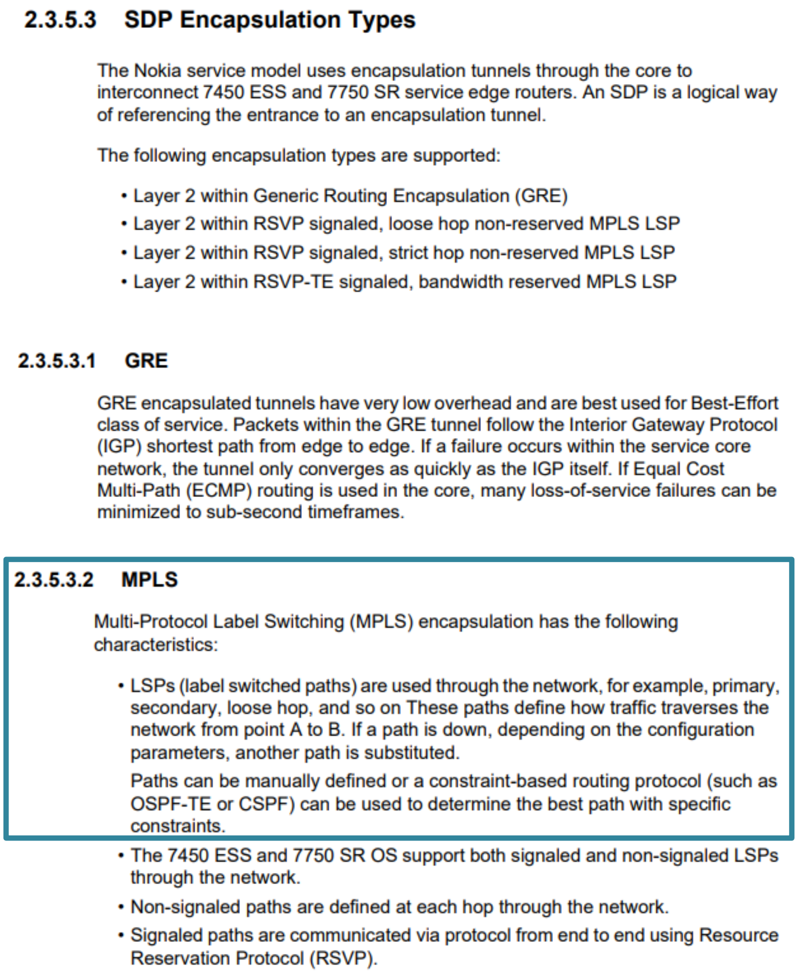

2.3.5.3 SDP Encapsulation Types

The Nokia service model uses encapsulation tunnels through the core to interconnect 7450 ESS and 7750 SR service edge routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- Layer 2 within Generic Routing Encapsulation (GRE)
- Layer 2 within RSVP signaled, loose hop non-reserved MPLS LSP
- Layer 2 within RSVP signaled, strict hop non-reserved MPLS LSP
- Layer 2 within RSVP-TE signaled, bandwidth reserved MPLS LSP

2.3.5.3.1 GRE

GRE encapsulated tunnels have very low overhead and are best used for Best-Effort class of service. Packets within the GRE tunnel follow the Interior Gateway Protocol (IGP) shortest path from edge to edge. If a failure occurs within the service core network, the tunnel only converges as quickly as the IGP itself. If Equal Cost Multi-Path (ECMP) routing is used in the core, many loss-of-service failures can be minimized to sub-second timeframes.

2.3.5.3.2 MPLS

Multi-Protocol Label Switching (MPLS) encapsulation has the following characteristics:

- LSPs (label switched paths) are used through the network, for example, primary, secondary, loose hop, and so on These paths define how traffic traverses the network from point A to B. If a path is down, depending on the configuration parameters, another path is substituted.

  Paths can be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints.
- The 7450 ESS and 7750 SR OS support both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.
- Signaled paths are communicated via protocol from end to end using Resource Reservation Protocol (RSVP).

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 51)

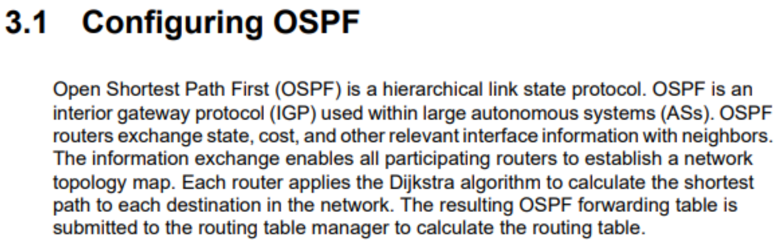

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

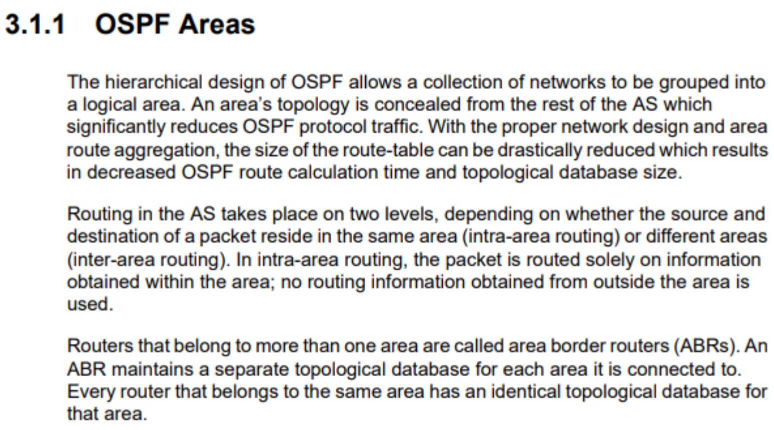

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

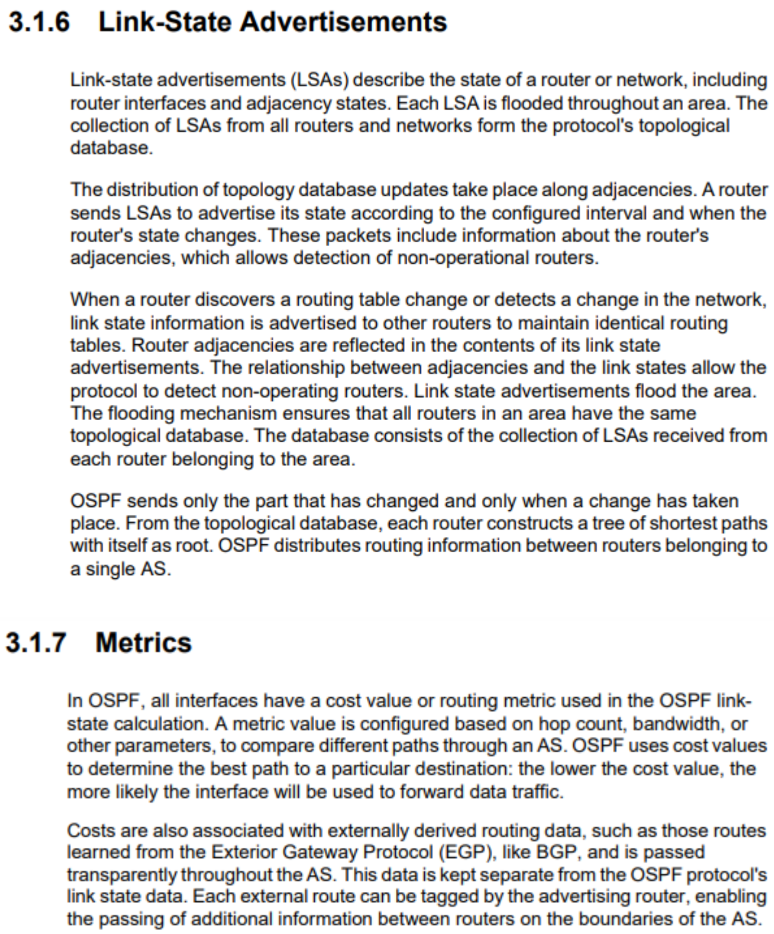

3.1.6 Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48)

108. Defendants instruct its customers to use the 7450 Ethernet Service Switch to define paths from said nodes through egress nodes of said ring network to external elements in said external layer-3 network.

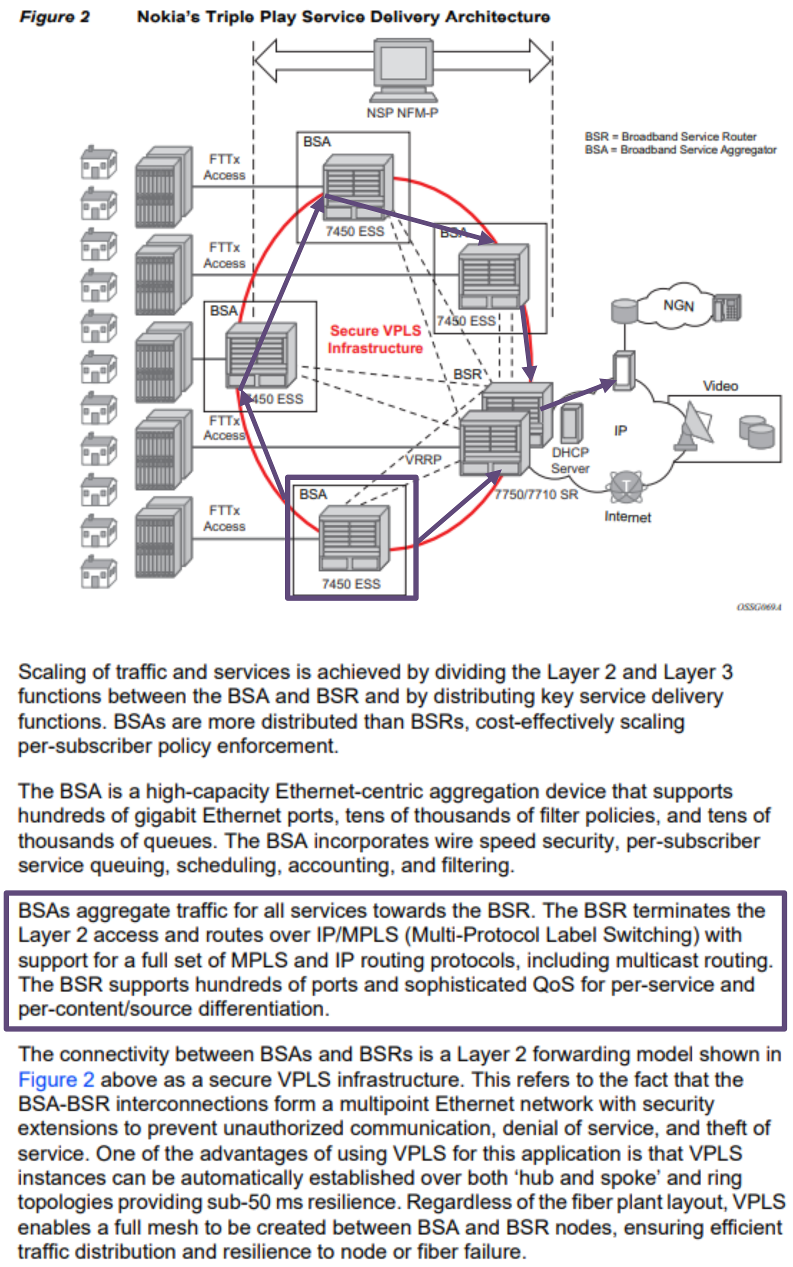

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi-Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both 'hub and spoke' and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20a nd%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Pages 34-35)

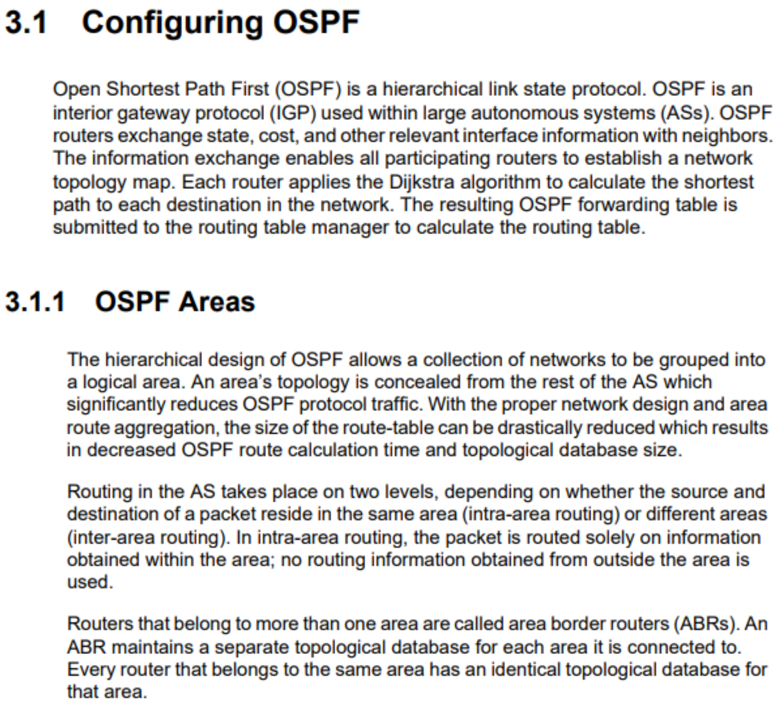

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

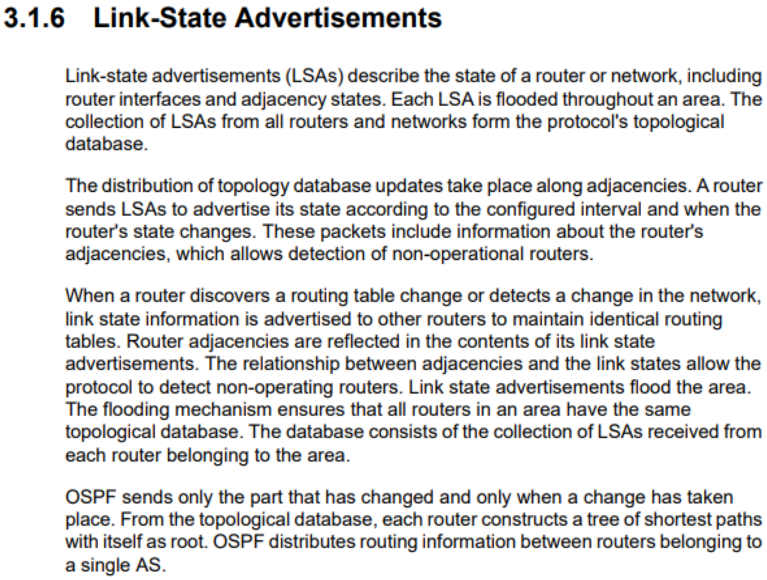

3.1.6 Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's topological database.

The distribution of topology database updates take place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routers.

When a router discovers a routing table change or detects a change in the network, link state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the part that has changed and only when a change has taken place. From the topological database, each router constructs a tree of shortest paths with itself as root. OSPF distributes routing information between routers belonging to a single AS.

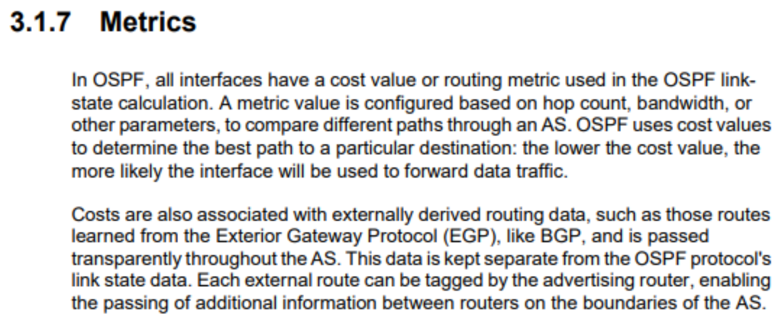

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48)

109. Defendants instruct its customers to use the 7450 Ethernet Service Switch to select one of said paths responsively to said metric.

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

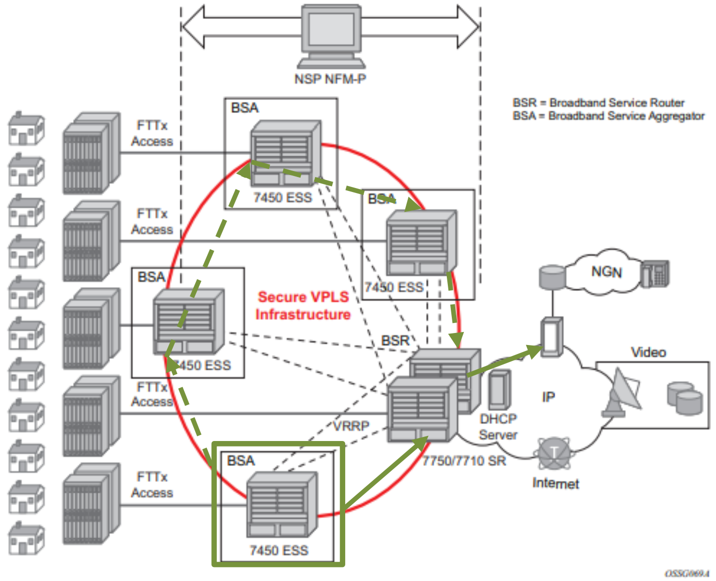

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 34)

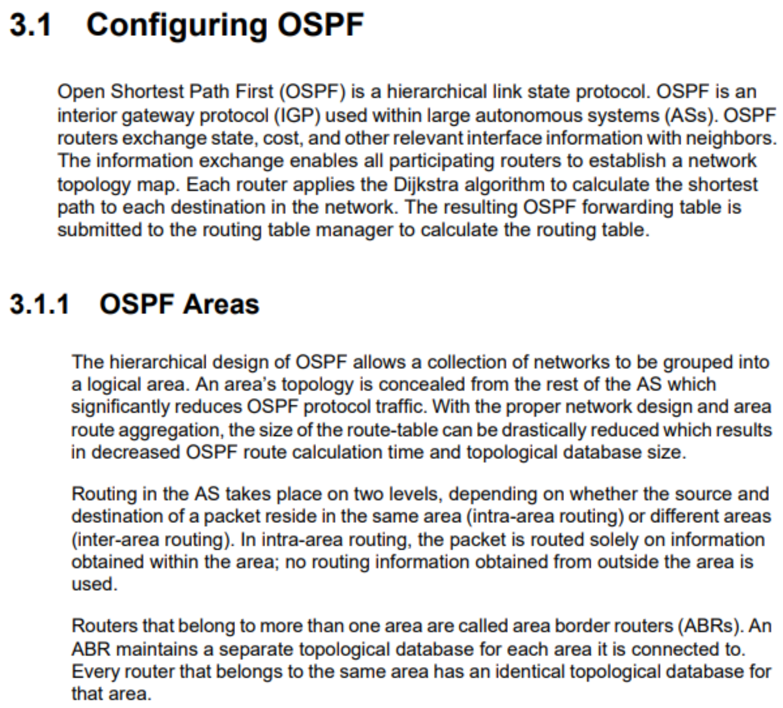

3.1 Configuring OSPF

Open Shortest Path First (OSPF) is a hierarchical link state protocol. OSPF is an interior gateway protocol (IGP) used within large autonomous systems (ASs). OSPF routers exchange state, cost, and other relevant interface information with neighbors. The information exchange enables all participating routers to establish a network topology map. Each router applies the Dijkstra algorithm to calculate the shortest path to each destination in the network. The resulting OSPF forwarding table is submitted to the routing table manager to calculate the routing table.

3.1.1 OSPF Areas

The hierarchical design of OSPF allows a collection of networks to be grouped into a logical area. An area's topology is concealed from the rest of the AS which significantly reduces OSPF protocol traffic. With the proper network design and area route aggregation, the size of the route-table can be drastically reduced which results in decreased OSPF route calculation time and topological database size.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area is used.

Routers that belong to more than one area are called area border routers (ABRs). An ABR maintains a separate topological database for each area it is connected to. Every router that belongs to the same area has an identical topological database for that area.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 35-36)

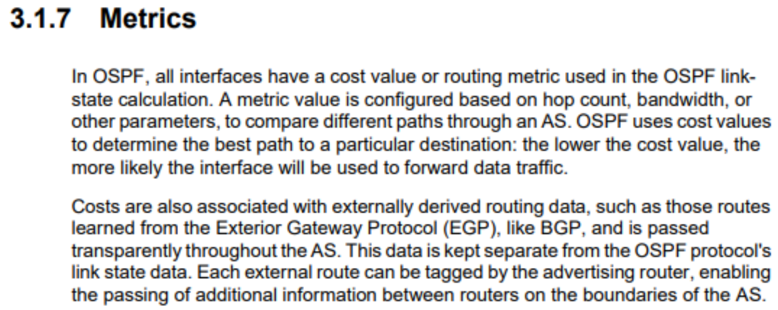

3.1.7 Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination: the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from the Exterior Gateway Protocol (EGP), like BGP, and is passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Pages 47-48)

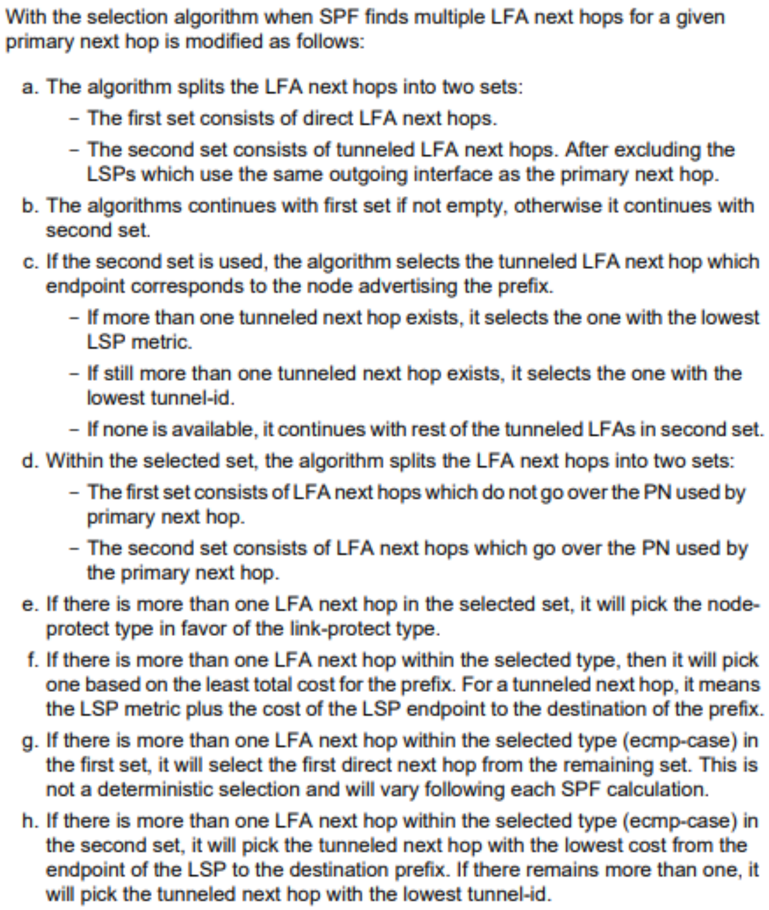

With the selection algorithm when SPF finds multiple LFA next hops for a given primary next hop is modified as follows:

a. The algorithm splits the LFA next hops into two sets:
   - The first set consists of direct LFA next hops.
   - The second set consists of tunneled LFA next hops. After excluding the LSPs which use the same outgoing interface as the primary next hop.
b. The algorithms continues with first set if not empty, otherwise it continues with second set.
c. If the second set is used, the algorithm selects the tunneled LFA next hop which endpoint corresponds to the node advertising the prefix.
   - If more than one tunneled next hop exists, it selects the one with the lowest LSP metric.
   - If still more than one tunneled next hop exists, it selects the one with the lowest tunnel-id.
   - If none is available, it continues with rest of the tunneled LFAs in second set.
d. Within the selected set, the algorithm splits the LFA next hops into two sets:
   - The first set consists of LFA next hops which do not go over the PN used by primary next hop.
   - The second set consists of LFA next hops which go over the PN used by the primary next hop.
e. If there is more than one LFA next hop in the selected set, it will pick the node-protect type in favor of the link-protect type.
f. If there is more than one LFA next hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
g. If there is more than one LFA next hop within the selected type (ecmp-case) in the first set, it will select the first direct next hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
h. If there is more than one LFA next hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next hop with the lowest tunnel-id.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17165AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Unicast%20Routing%20Protocols%20Guide%2021.10.R1.pdf (Page 58)

110. Defendants instruct its customers to use the 7450 Ethernet Service Switch in transmitting data from at least one of said nodes via said selected one of said paths to network elements that are external to said ring network.

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

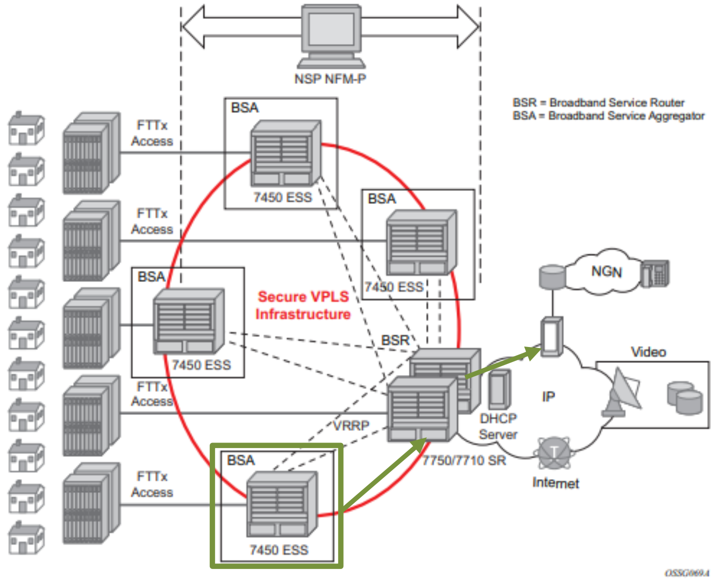

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17164AAADTQZZA01_V1_7450%20ESS%207750%20SR%20and%20VSR%20Triple%20Play%20Service%20Delivery%20Architecture%20Guide%2021.10.R1.pdf (Page 34**)**

111. As a result of Defendants' infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 US.C. § 284.

COUNT FOUR
<u>INFRINGEMENT OF U.S. PATENT 7,697,525</u>

112. Plaintiff incorporates by reference the allegations in all preceding paragraphs as if fully set forth herein.

113. The '525 Patent, entitled "Forwarding multicast traffic over link aggregation ports" was filed on December 21, 2006, and issued on April 13, 2010.

114. Plaintiff is the assignee and owner of all rights, title and interest to the '525 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

<u>Technical Description</u>

115. The ’525 Patent provides a solution to the problems in the prior art as follows, “[u]nlike some known methods and systems in which all multicast packets are sent to the same LAG group port, the methods and systems described herein distribute multicast packets approximately evenly among the different output ports of the LAG group. Thus, the traffic load within the group is balanced, and distribution of additional unicast traffic across the group is simplified.” 3:54-60.

<u>Direct Infringement</u>

116. Defendants, without authorization or license from Plaintiff, have been and are directly infringing the ’525 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale methods, devices, and networks infringing one or more claims of the ’525 Patent. Defendants develop, design, manufacture, and distribute telecommunications equipment that infringes one or more claims of the ’525 Patent. Defendants further provide services that practice methods that infringe one or more claims of the ’525 Patent. Defendants are thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Defendants’ 7450 Ethernet Service Switch, and all other substantially similar products (collectively the “’525 Accused Products”).

117. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendants’ infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of ’525 Accused Products.

118. Defendants are liable for direct infringement pursuant to 35 U.S.C. § 271 for the use, manufacture, sale, offer of sale, important, or distribution of Defendants' 7450 Ethernet Service Switch.

119. Defendants' 7450 Ethernet Service Switch is a non-limiting example of an apparatus that meets all limitations of claim 1 of the '525 Patent, either literally or equivalently.

120. Defendants' 7450 Ethernet Service Switch is a network node in a communication network.

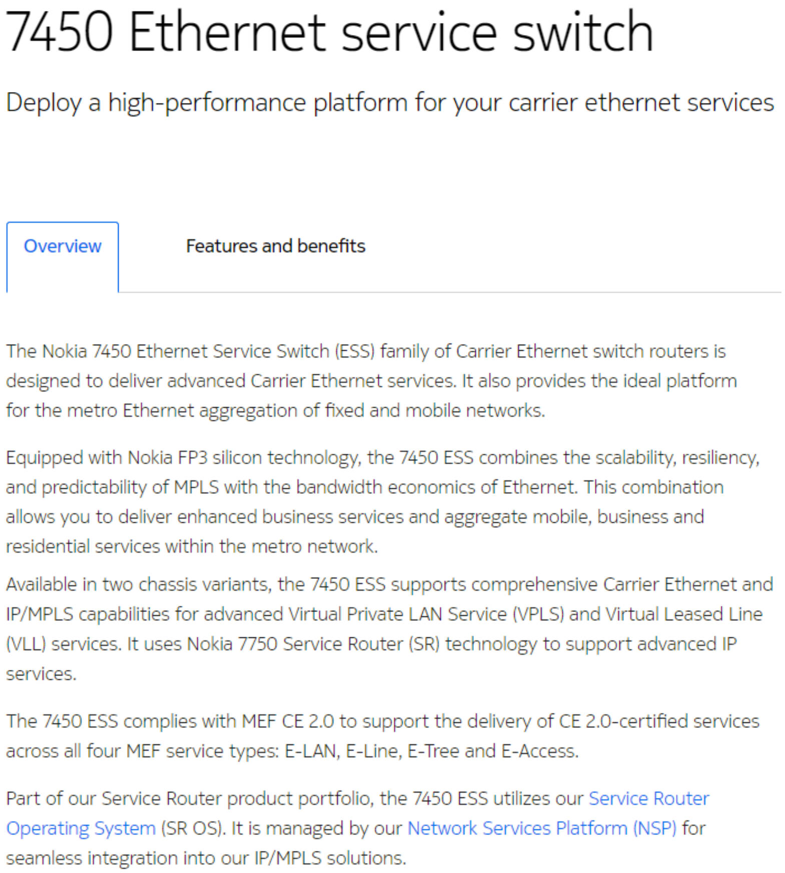

7450 Ethernet service switch

Deploy a high-performance platform for your carrier ethernet services

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and residential services within the metro network.

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

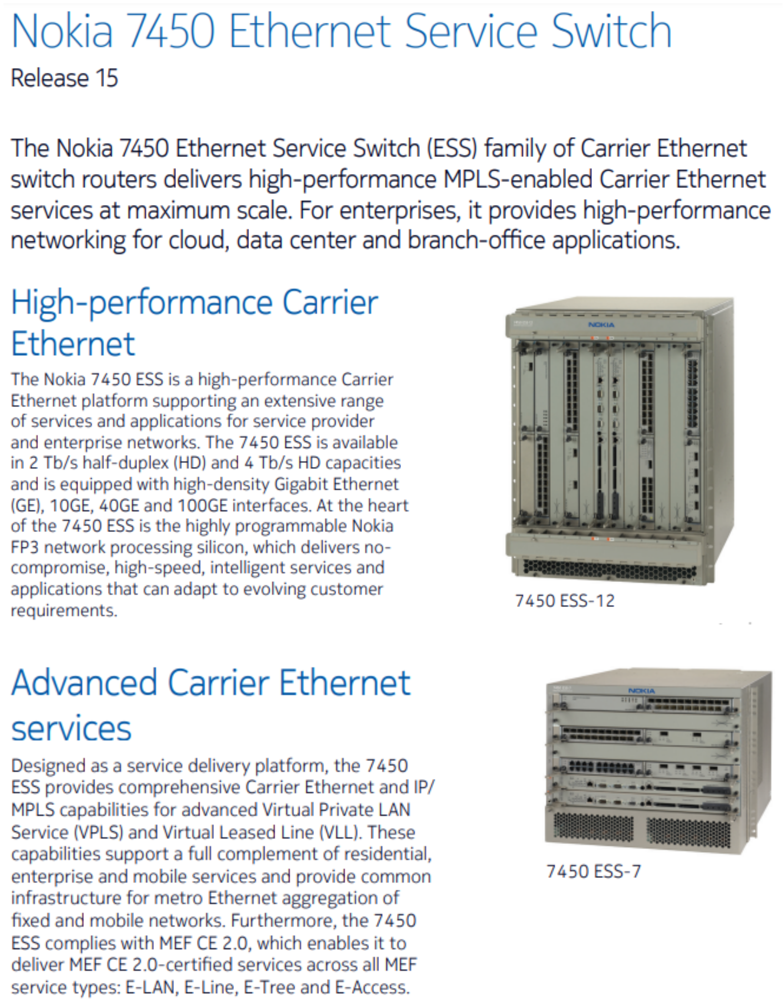

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/f/164727 (Page 1 of PDF)

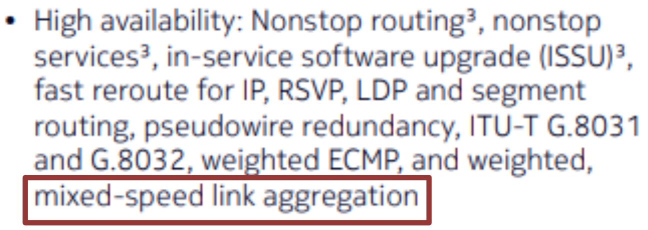

- High availability: Nonstop routing[3], nonstop services[3], in-service software upgrade (ISSU)[3], fast reroute for IP, RSVP, LDP and segment routing, pseudowire redundancy, ITU-T G.8031 and G.8032, weighted ECMP, and weighted, mixed-speed link aggregation

https://onestore.nokia.com/asset/f/164727 (Page 6 of PDF)

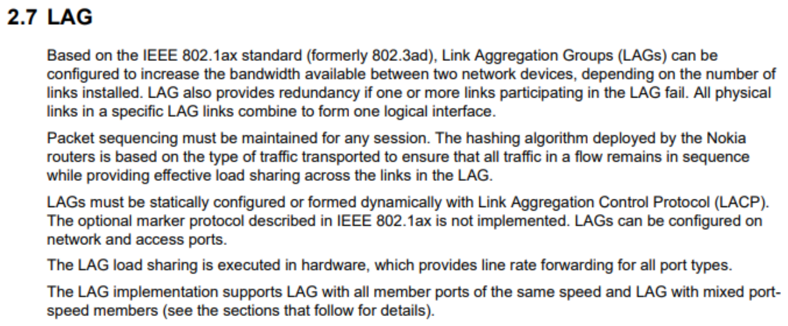

2.7 LAG

Based on the IEEE 802.1ax standard (formerly 802.3ad), Link Aggregation Groups (LAGs) can be configured to increase the bandwidth available between two network devices, depending on the number of links installed. LAG also provides redundancy if one or more links participating in the LAG fail. All physical links in a specific LAG links combine to form one logical interface.

Packet sequencing must be maintained for any session. The hashing algorithm deployed by the Nokia routers is based on the type of traffic transported to ensure that all traffic in a flow remains in sequence while providing effective load sharing across the links in the LAG.

LAGs must be statically configured or formed dynamically with Link Aggregation Control Protocol (LACP). The optional marker protocol described in IEEE 802.1ax is not implemented. LAGs can be configured on network and access ports.

The LAG load sharing is executed in hardware, which provides line rate forwarding for all port types.

The LAG implementation supports LAG with all member ports of the same speed and LAG with mixed port-speed members (see the sections that follow for details).

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Page 121 of PDF)

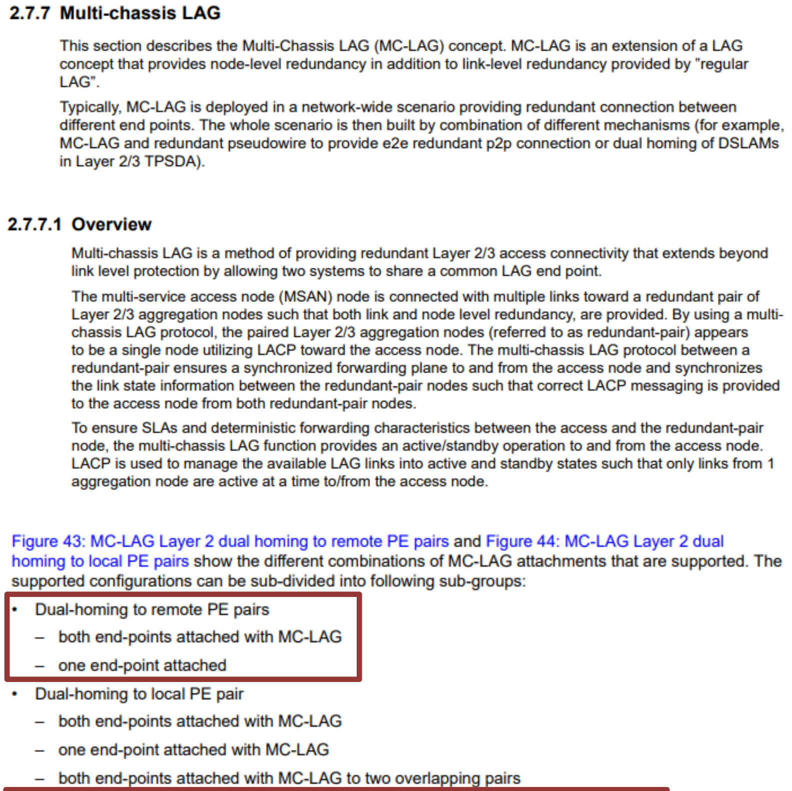

2.7.7 Multi-chassis LAG

This section describes the Multi-Chassis LAG (MC-LAG) concept. MC-LAG is an extension of a LAG concept that provides node-level redundancy in addition to link-level redundancy provided by "regular LAG".

Typically, MC-LAG is deployed in a network-wide scenario providing redundant connection between different end points. The whole scenario is then built by combination of different mechanisms (for example, MC-LAG and redundant pseudowire to provide e2e redundant p2p connection or dual homing of DSLAMs in Layer 2/3 TPSDA).

2.7.7.1 Overview

Multi-chassis LAG is a method of providing redundant Layer 2/3 access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point.

The multi-service access node (MSAN) node is connected with multiple links toward a redundant pair of Layer 2/3 aggregation nodes such that both link and node level redundancy, are provided. By using a multi-chassis LAG protocol, the paired Layer 2/3 aggregation nodes (referred to as redundant-pair) appears to be a single node utilizing LACP toward the access node. The multi-chassis LAG protocol between a redundant-pair ensures a synchronized forwarding plane to and from the access node and synchronizes the link state information between the redundant-pair nodes such that correct LACP messaging is provided to the access node from both redundant-pair nodes.

To ensure SLAs and deterministic forwarding characteristics between the access and the redundant-pair node, the multi-chassis LAG function provides an active/standby operation to and from the access node. LACP is used to manage the available LAG links into active and standby states such that only links from 1 aggregation node are active at a time to/from the access node.

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

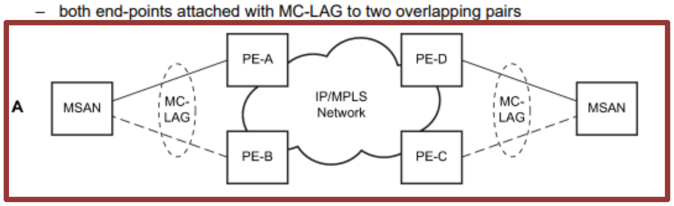

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 145-147 of PDF)

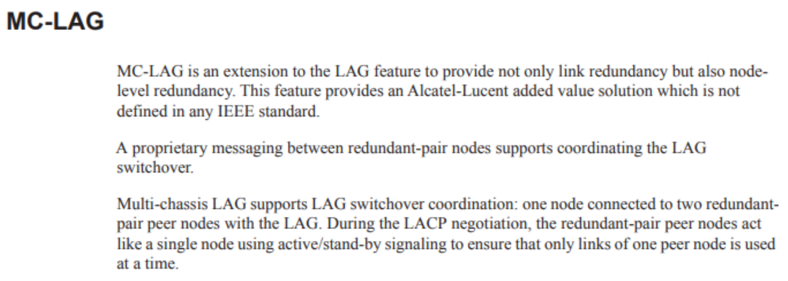

MC-LAG

MC-LAG is an extension to the LAG feature to provide not only link redundancy but also node-level redundancy. This feature provides an Alcatel-Lucent added value solution which is not defined in any IEEE standard.

A proprietary messaging between redundant-pair nodes supports coordinating the LAG switchover.

Multi-chassis LAG supports LAG switchover coordination: one node connected to two redundant-pair peer nodes with the LAG. During the LACP negotiation, the redundant-pair peer nodes act like a single node using active/stand-by signaling to ensure that only links of one peer node is used at a time.

Network Topology

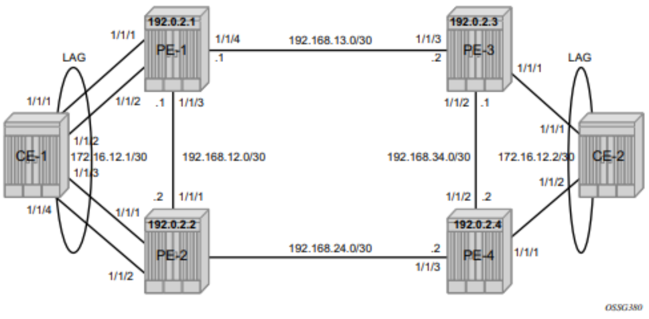

Figure 25: MC-LAG Network Topology

https://documentation.nokia.com/html/0_add-h-f/93-0267-HTML/7X50_Advanced_Configuration_Guide/MC-LAG%20.pdf (Pages 3 and 4 of PDF)

121. Defendants' 7450 Ethernet Service Switch comprises a plurality of ports, at least a subset of which is grouped in a link aggregation (LAG) group:

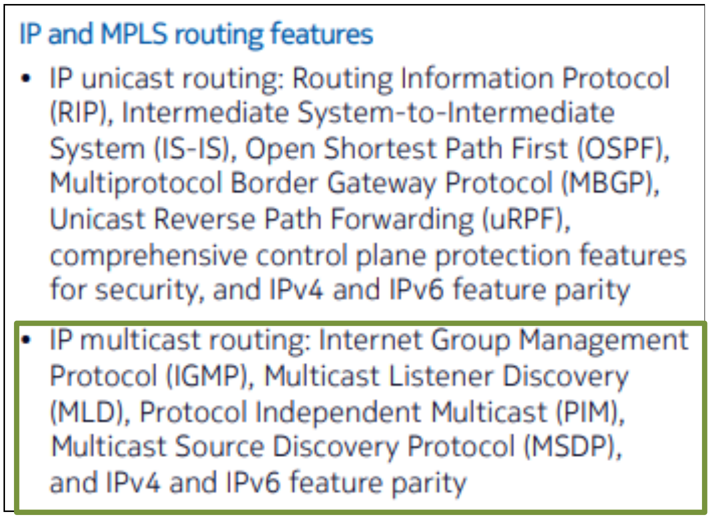

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity

https://onestore.nokia.com/asset/f/164727 (Page 6 of PDF)

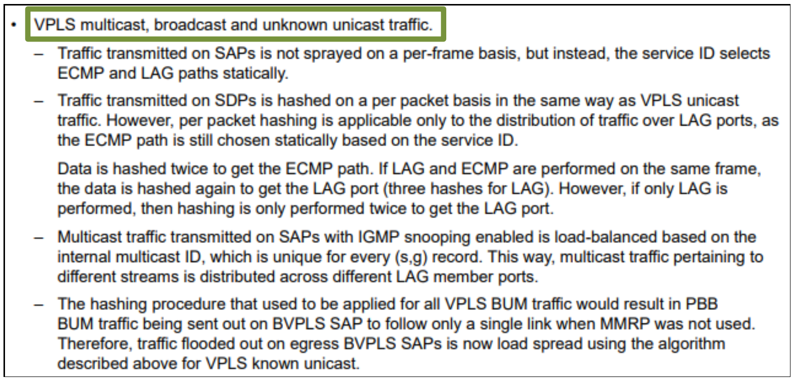

- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 128 of PDF)

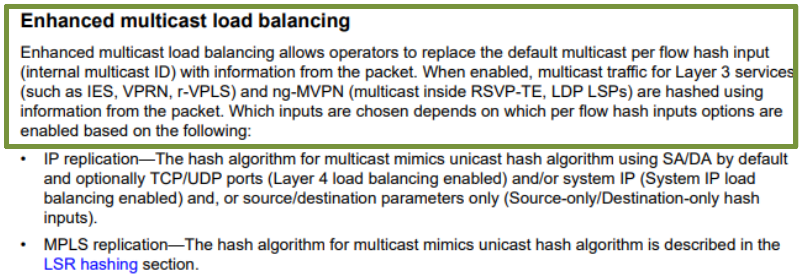

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 132 of PDF)

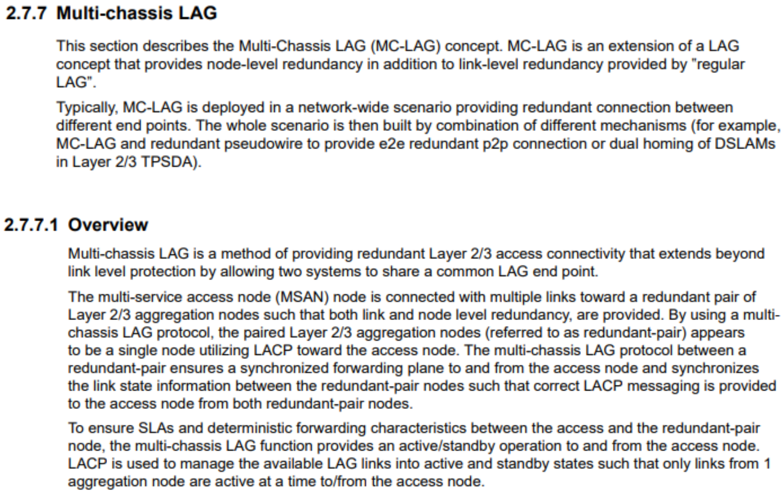

2.7.7 Multi-chassis LAG

This section describes the Multi-Chassis LAG (MC-LAG) concept. MC-LAG is an extension of a LAG concept that provides node-level redundancy in addition to link-level redundancy provided by "regular LAG".

Typically, MC-LAG is deployed in a network-wide scenario providing redundant connection between different end points. The whole scenario is then built by combination of different mechanisms (for example, MC-LAG and redundant pseudowire to provide e2e redundant p2p connection or dual homing of DSLAMs in Layer 2/3 TPSDA).

2.7.7.1 Overview

Multi-chassis LAG is a method of providing redundant Layer 2/3 access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point.

The multi-service access node (MSAN) node is connected with multiple links toward a redundant pair of Layer 2/3 aggregation nodes such that both link and node level redundancy, are provided. By using a multi-chassis LAG protocol, the paired Layer 2/3 aggregation nodes (referred to as redundant-pair) appears to be a single node utilizing LACP toward the access node. The multi-chassis LAG protocol between a redundant-pair ensures a synchronized forwarding plane to and from the access node and synchronizes the link state information between the redundant-pair nodes such that correct LACP messaging is provided to the access node from both redundant-pair nodes.

To ensure SLAs and deterministic forwarding characteristics between the access and the redundant-pair node, the multi-chassis LAG function provides an active/standby operation to and from the access node. LACP is used to manage the available LAG links into active and standby states such that only links from 1 aggregation node are active at a time to/from the access node.

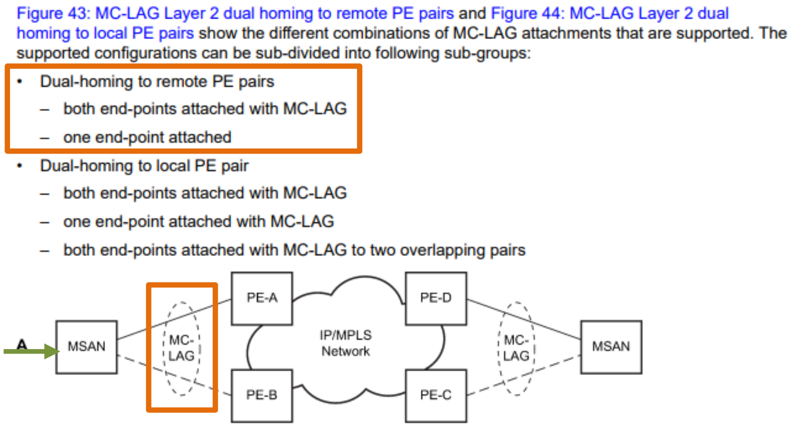

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 145-147 of PDF)

122. Defendants' 7450 Ethernet Service Switch comprises packet processing logic, which is coupled to receive data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of the ports and to process the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group:

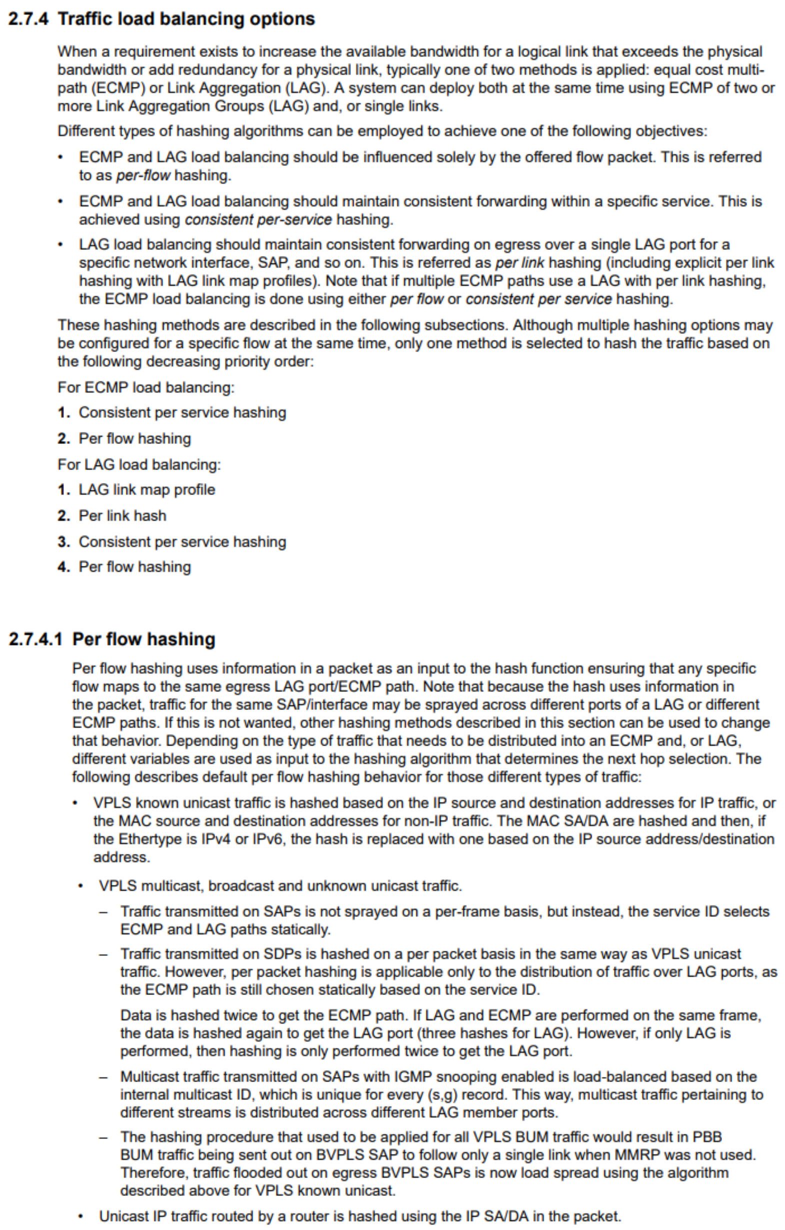

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 127 and 128 of PDF)

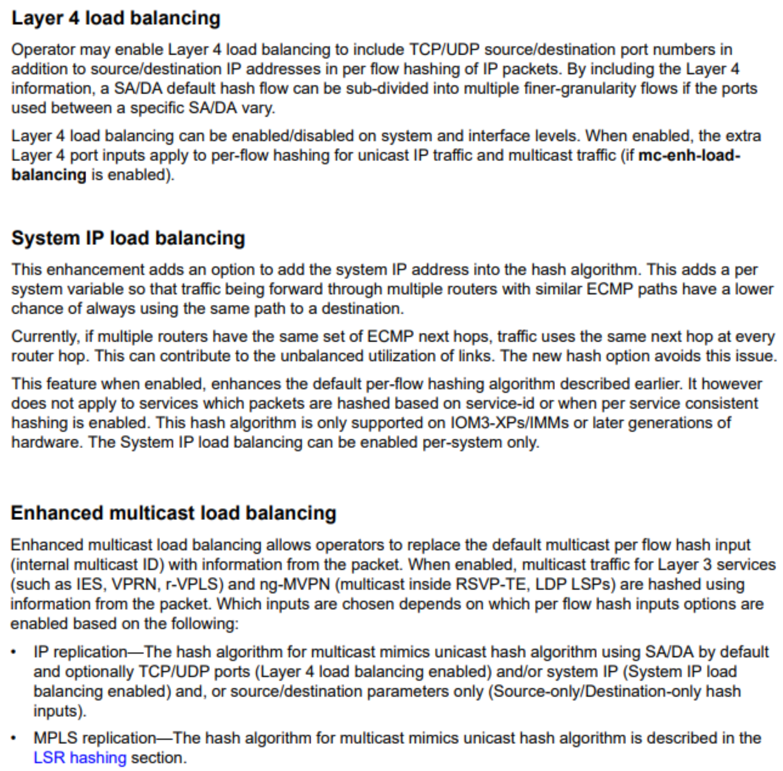

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 131 and 132 of PDF)

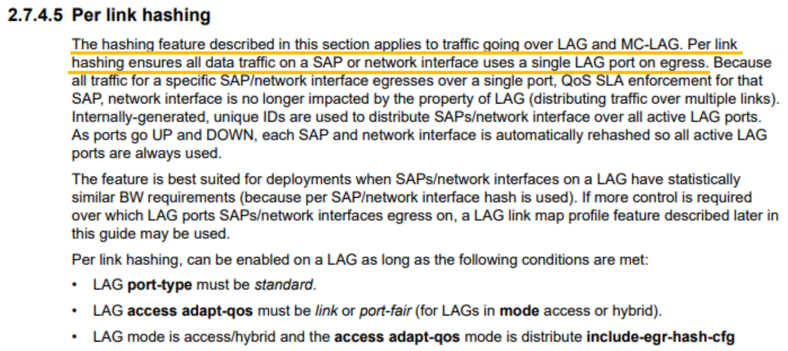

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard.*
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 138 of PDF)

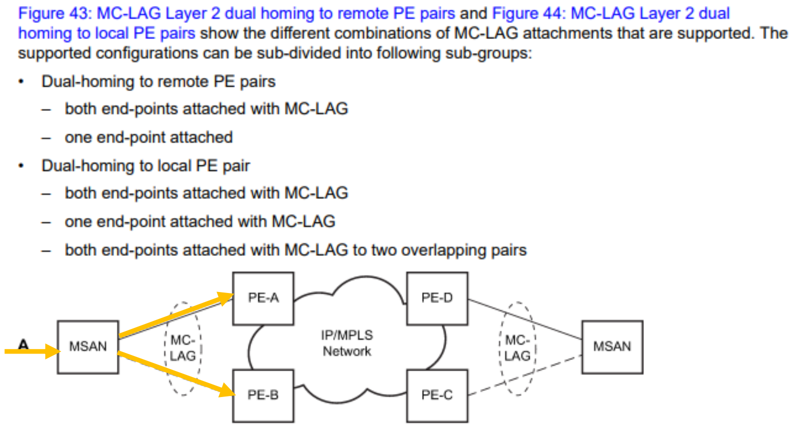

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 147 of PDF)

123. Defendants' 7450 Ethernet Service Switch comprises the packet processing logic is arranged to allocate to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, each FMID being associated with one of the ports in the subset, and to forward the single copy to the port associated with the allocated FMID value:

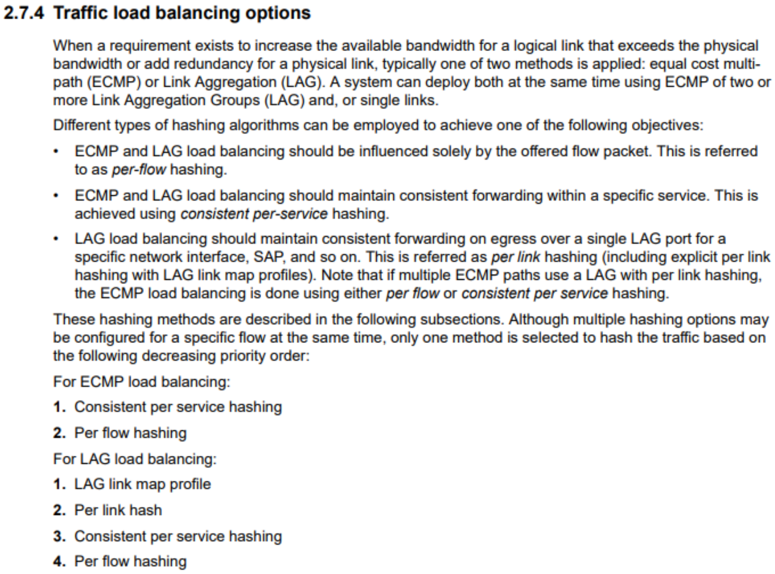

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

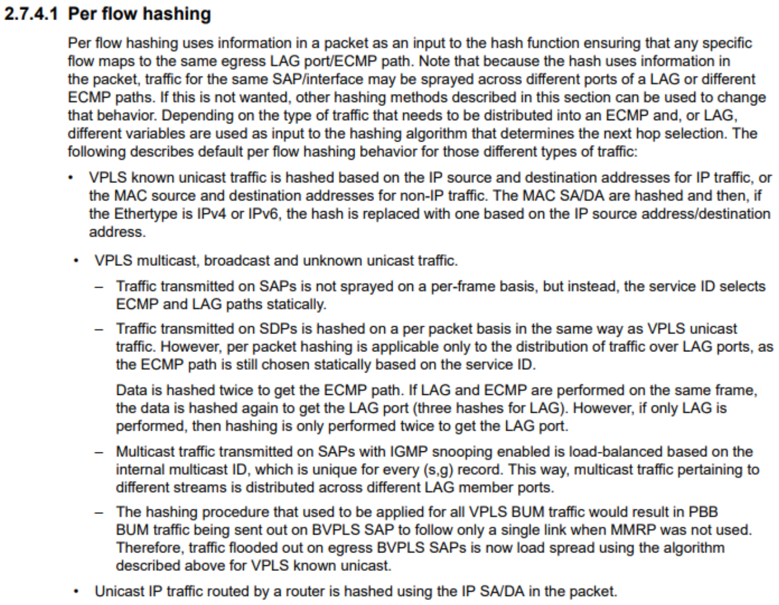

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 127 and 128 of PDF)

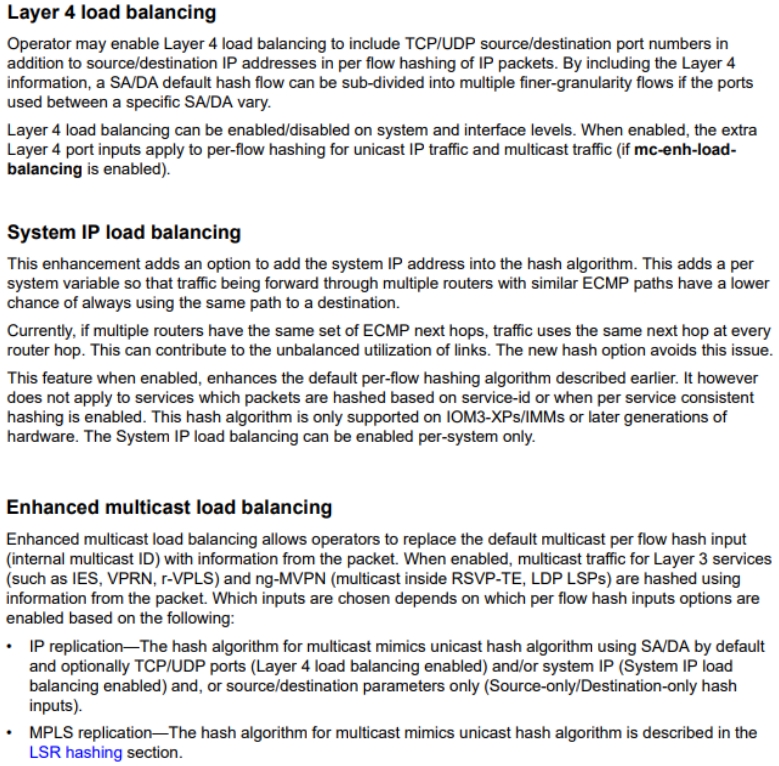

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207

950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 131 and 132 of PDF)

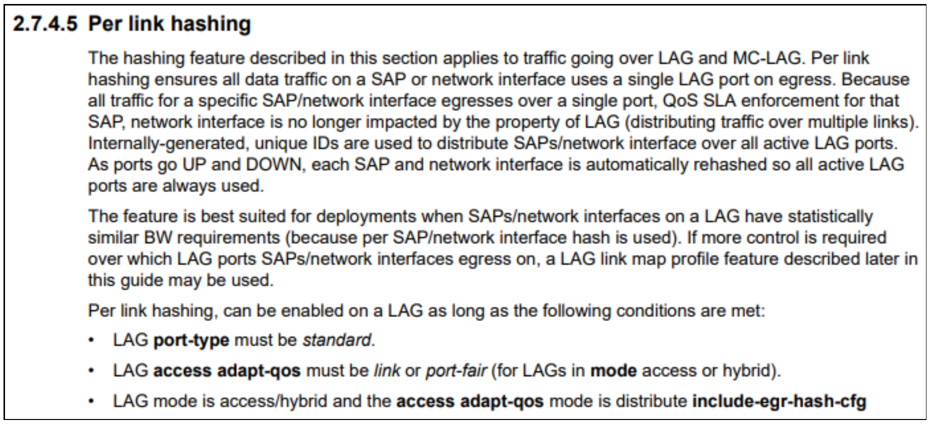

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

124. Defendants' 7450 Ethernet Service Switch comprises the packet processing logic comprises multiple line cards connected to the ports and a switching fabric interconnecting the line cards, wherein, for each of the received data packets, a first line card connected to a first port via which the data packet is received is arranged to allocate the FMID value to the packet and to forward the packet to the switching fabric, and wherein the switching fabric and a second line card connected to a second port to which the data packet is to be sent are configured to forward the data packet responsively to the FMID value:

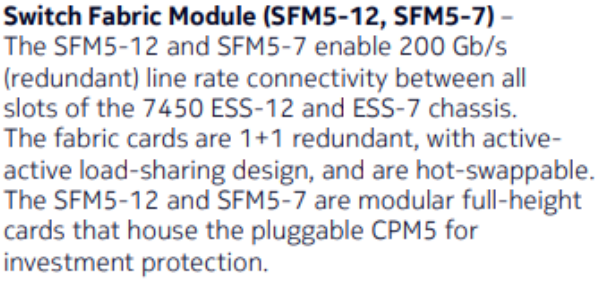

Switch Fabric Module (SFM5-12, SFM5-7) – The SFM5-12 and SFM5-7 enable 200 Gb/s (redundant) line rate connectivity between all slots of the 7450 ESS-12 and ESS-7 chassis. The fabric cards are 1+1 redundant, with active-active load-sharing design, and are hot-swappable. The SFM5-12 and SFM5-7 are modular full-height cards that house the pluggable CPM5 for investment protection.

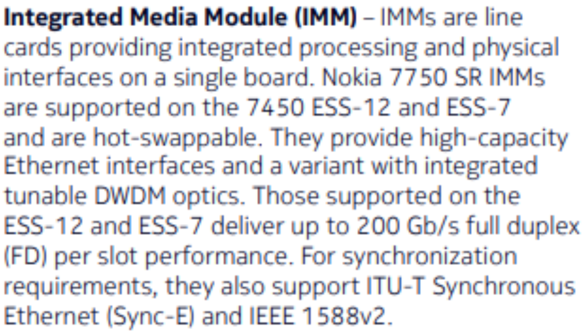

Integrated Media Module (IMM) – IMMs are line cards providing integrated processing and physical interfaces on a single board. Nokia 7750 SR IMMs are supported on the 7450 ESS-12 and ESS-7 and are hot-swappable. They provide high-capacity Ethernet interfaces and a variant with integrated tunable DWDM optics. Those supported on the ESS-12 and ESS-7 deliver up to 200 Gb/s full duplex (FD) per slot performance. For synchronization requirements, they also support ITU-T Synchronous Ethernet (Sync-E) and IEEE 1588v2.

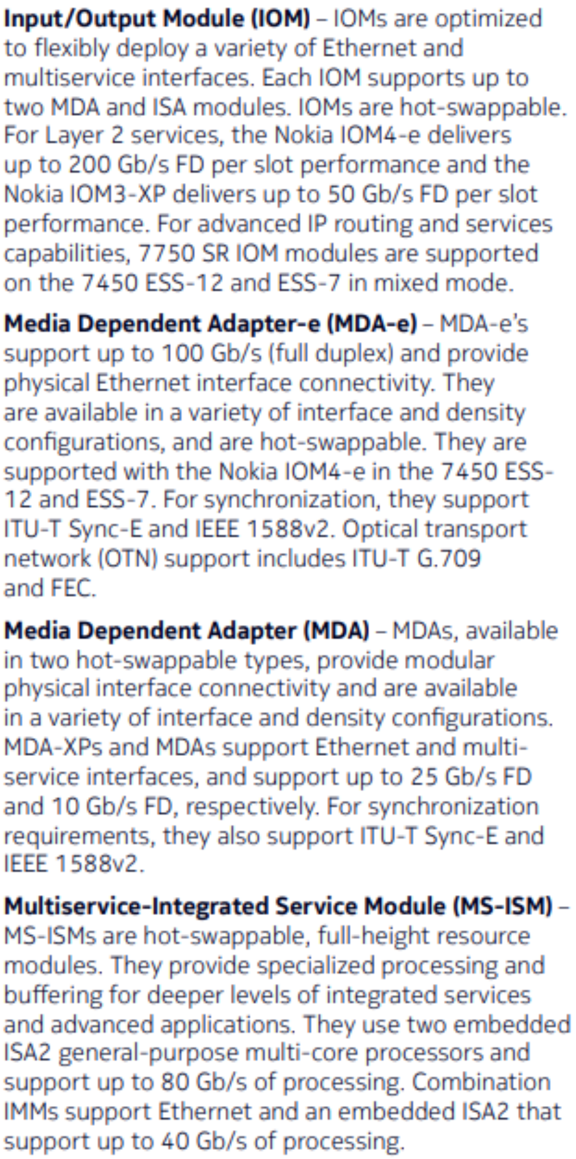

Input/Output Module (IOM) – IOMs are optimized to flexibly deploy a variety of Ethernet and multiservice interfaces. Each IOM supports up to two MDA and ISA modules. IOMs are hot-swappable. For Layer 2 services, the Nokia IOM4-e delivers up to 200 Gb/s FD per slot performance and the Nokia IOM3-XP delivers up to 50 Gb/s FD per slot performance. For advanced IP routing and services capabilities, 7750 SR IOM modules are supported on the 7450 ESS-12 and ESS-7 in mixed mode.

Media Dependent Adapter-e (MDA-e) – MDA-e's support up to 100 Gb/s (full duplex) and provide physical Ethernet interface connectivity. They are available in a variety of interface and density configurations, and are hot-swappable. They are supported with the Nokia IOM4-e in the 7450 ESS-12 and ESS-7. For synchronization, they support ITU-T Sync-E and IEEE 1588v2. Optical transport network (OTN) support includes ITU-T G.709 and FEC.

Media Dependent Adapter (MDA) – MDAs, available in two hot-swappable types, provide modular physical interface connectivity and are available in a variety of interface and density configurations. MDA-XPs and MDAs support Ethernet and multi-service interfaces, and support up to 25 Gb/s FD and 10 Gb/s FD, respectively. For synchronization requirements, they also support ITU-T Sync-E and IEEE 1588v2.

Multiservice-Integrated Service Module (MS-ISM) – MS-ISMs are hot-swappable, full-height resource modules. They provide specialized processing and buffering for deeper levels of integrated services and advanced applications. They use two embedded ISA2 general-purpose multi-core processors and support up to 80 Gb/s of processing. Combination IMMs support Ethernet and an embedded ISA2 that support up to 40 Gb/s of processing.

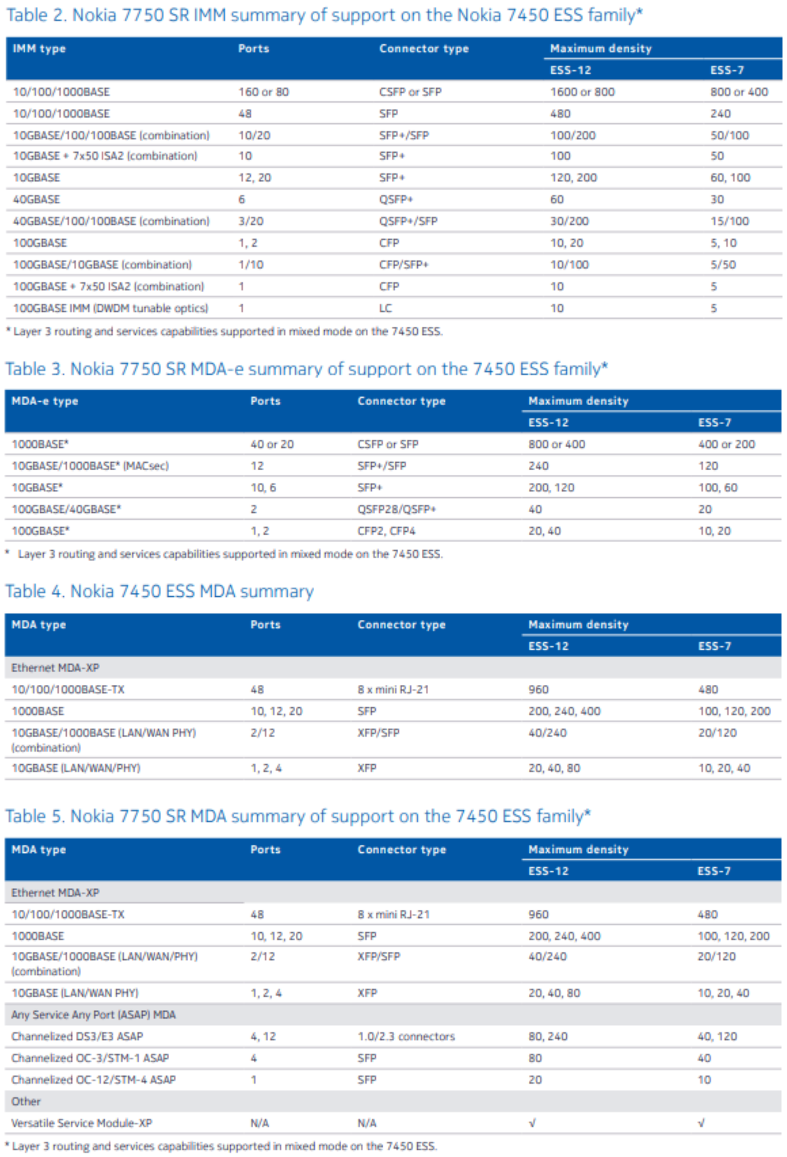

Table 2. Nokia 7750 SR IMM summary of support on the Nokia 7450 ESS family*

IMM type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
10/100/1000BASE	160 or 80	CSFP or SFP	1600 or 800	800 or 400
10/100/1000BASE	48	SFP	480	240
10GBASE/100/100BASE (combination)	10/20	SFP+/SFP	100/200	50/100
10GBASE + 7x50 ISA2 (combination)	10	SFP+	100	50
10GBASE	12, 20	SFP+	120, 200	60, 100
40GBASE	6	QSFP+	60	30
40GBASE/100/100BASE (combination)	3/20	QSFP+/SFP	30/200	15/100
100GBASE	1, 2	CFP	10, 20	5, 10
100GBASE/10GBASE (combination)	1/10	CFP/SFP+	10/100	5/50
100GBASE + 7x50 ISA2 (combination)	1	CFP	10	5
100GBASE IMM (DWDM tunable optics)	1	LC	10	5

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

Table 3. Nokia 7750 SR MDA-e summary of support on the 7450 ESS family*

MDA-e type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
1000BASE*	40 or 20	CSFP or SFP	800 or 400	400 or 200
10GBASE/1000BASE* (MACsec)	12	SFP+/SFP	240	120
10GBASE*	10, 6	SFP+	200, 120	100, 60
100GBASE/40GBASE*	2	QSFP28/QSFP+	40	20
100GBASE*	1, 2	CFP2, CFP4	20, 40	10, 20

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

Table 4. Nokia 7450 ESS MDA summary

MDA type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
Ethernet MDA-XP				
10/100/1000BASE-TX	48	8 x mini RJ-21	960	480
1000BASE	10, 12, 20	SFP	200, 240, 400	100, 120, 200
10GBASE/1000BASE (LAN/WAN PHY) (combination)	2/12	XFP/SFP	40/240	20/120
10GBASE (LAN/WAN/PHY)	1, 2, 4	XFP	20, 40, 80	10, 20, 40

Table 5. Nokia 7750 SR MDA summary of support on the 7450 ESS family*

MDA type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
Ethernet MDA-XP				
10/100/1000BASE-TX	48	8 x mini RJ-21	960	480
1000BASE	10, 12, 20	SFP	200, 240, 400	100, 120, 200
10GBASE/1000BASE (LAN/WAN/PHY) (combination)	2/12	XFP/SFP	40/240	20/120
10GBASE (LAN/WAN PHY)	1, 2, 4	XFP	20, 40, 80	10, 20, 40
Any Service Any Port (ASAP) MDA				
Channelized DS3/E3 ASAP	4, 12	1.0/2.3 connectors	80, 240	40, 120
Channelized OC-3/STM-1 ASAP	4	SFP	80	40
Channelized OC-12/STM-4 ASAP	1	SFP	20	10
Other				
Versatile Service Module-XP	N/A	N/A	√	√

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

https://onestore.nokia.com/asset/164727 (Pages 3-5 of PDF)

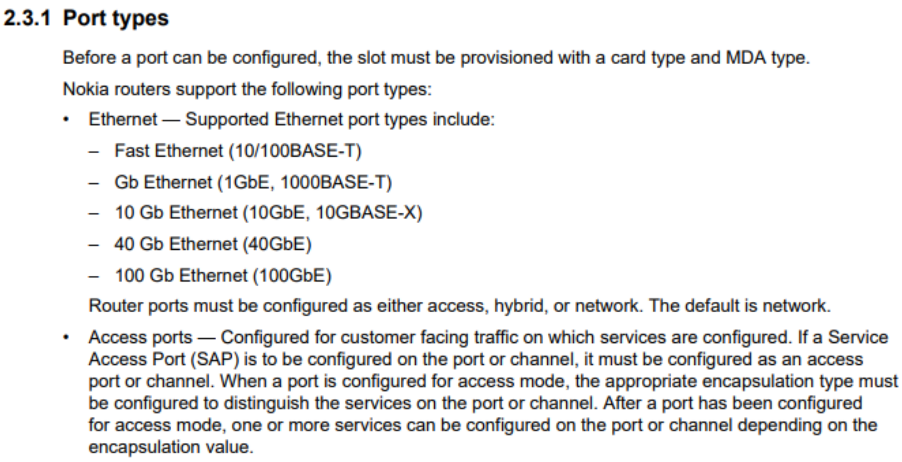

2.3.1 Port types

Before a port can be configured, the slot must be provisioned with a card type and MDA type.

Nokia routers support the following port types:

- Ethernet — Supported Ethernet port types include:
  - Fast Ethernet (10/100BASE-T)
  - Gb Ethernet (1GbE, 1000BASE-T)
  - 10 Gb Ethernet (10GbE, 10GBASE-X)
  - 40 Gb Ethernet (40GbE)
  - 100 Gb Ethernet (100GbE)

  Router ports must be configured as either access, hybrid, or network. The default is network.
- Access ports — Configured for customer facing traffic on which services are configured. If a Service Access Port (SAP) is to be configured on the port or channel, it must be configured as an access port or channel. When a port is configured for access mode, the appropriate encapsulation type must be configured to distinguish the services on the port or channel. After a port has been configured for access mode, one or more services can be configured on the port or channel depending on the encapsulation value.

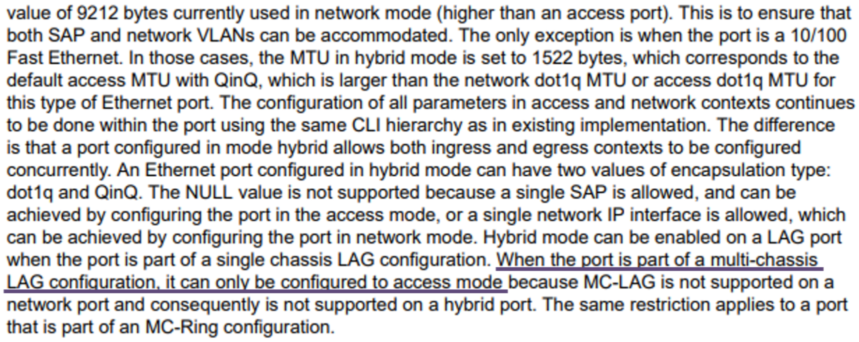

value of 9212 bytes currently used in network mode (higher than an access port). This is to ensure that both SAP and network VLANs can be accommodated. The only exception is when the port is a 10/100 Fast Ethernet. In those cases, the MTU in hybrid mode is set to 1522 bytes, which corresponds to the default access MTU with QinQ, which is larger than the network dot1q MTU or access dot1q MTU for this type of Ethernet port. The configuration of all parameters in access and network contexts continues to be done within the port using the same CLI hierarchy as in existing implementation. The difference is that a port configured in mode hybrid allows both ingress and egress contexts to be configured concurrently. An Ethernet port configured in hybrid mode can have two values of encapsulation type: dot1q and QinQ. The NULL value is not supported because a single SAP is allowed, and can be achieved by configuring the port in the access mode, or a single network IP interface is allowed, which can be achieved by configuring the port in network mode. Hybrid mode can be enabled on a LAG port when the port is part of a single chassis LAG configuration. When the port is part of a multi-chassis LAG configuration, it can only be configured to access mode because MC-LAG is not supported on a network port and consequently is not supported on a hybrid port. The same restriction applies to a port that is part of an MC-Ring configuration.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 28 and 29 of PDF)

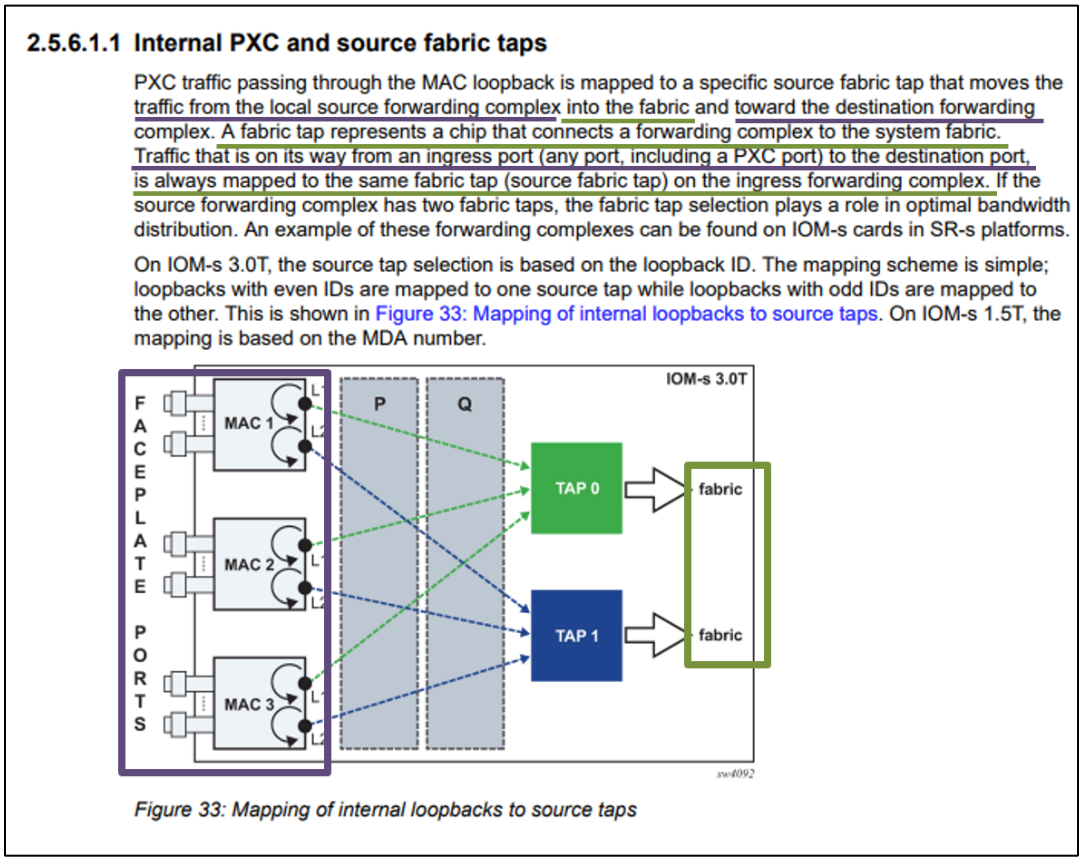

2.5.6.1.1 Internal PXC and source fabric taps

PXC traffic passing through the MAC loopback is mapped to a specific source fabric tap that moves the traffic from the local source forwarding complex into the fabric and toward the destination forwarding complex. A fabric tap represents a chip that connects a forwarding complex to the system fabric. Traffic that is on its way from an ingress port (any port, including a PXC port) to the destination port, is always mapped to the same fabric tap (source fabric tap) on the ingress forwarding complex. If the source forwarding complex has two fabric taps, the fabric tap selection plays a role in optimal bandwidth distribution. An example of these forwarding complexes can be found on IOM-s cards in SR-s platforms.

On IOM-s 3.0T, the source tap selection is based on the loopback ID. The mapping scheme is simple; loopbacks with even IDs are mapped to one source tap while loopbacks with odd IDs are mapped to the other. This is shown in Figure 33: Mapping of internal loopbacks to source taps. On IOM-s 1.5T, the mapping is based on the MDA number.

Figure 33: Mapping of internal loopbacks to source taps

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Page 109 of PDF)

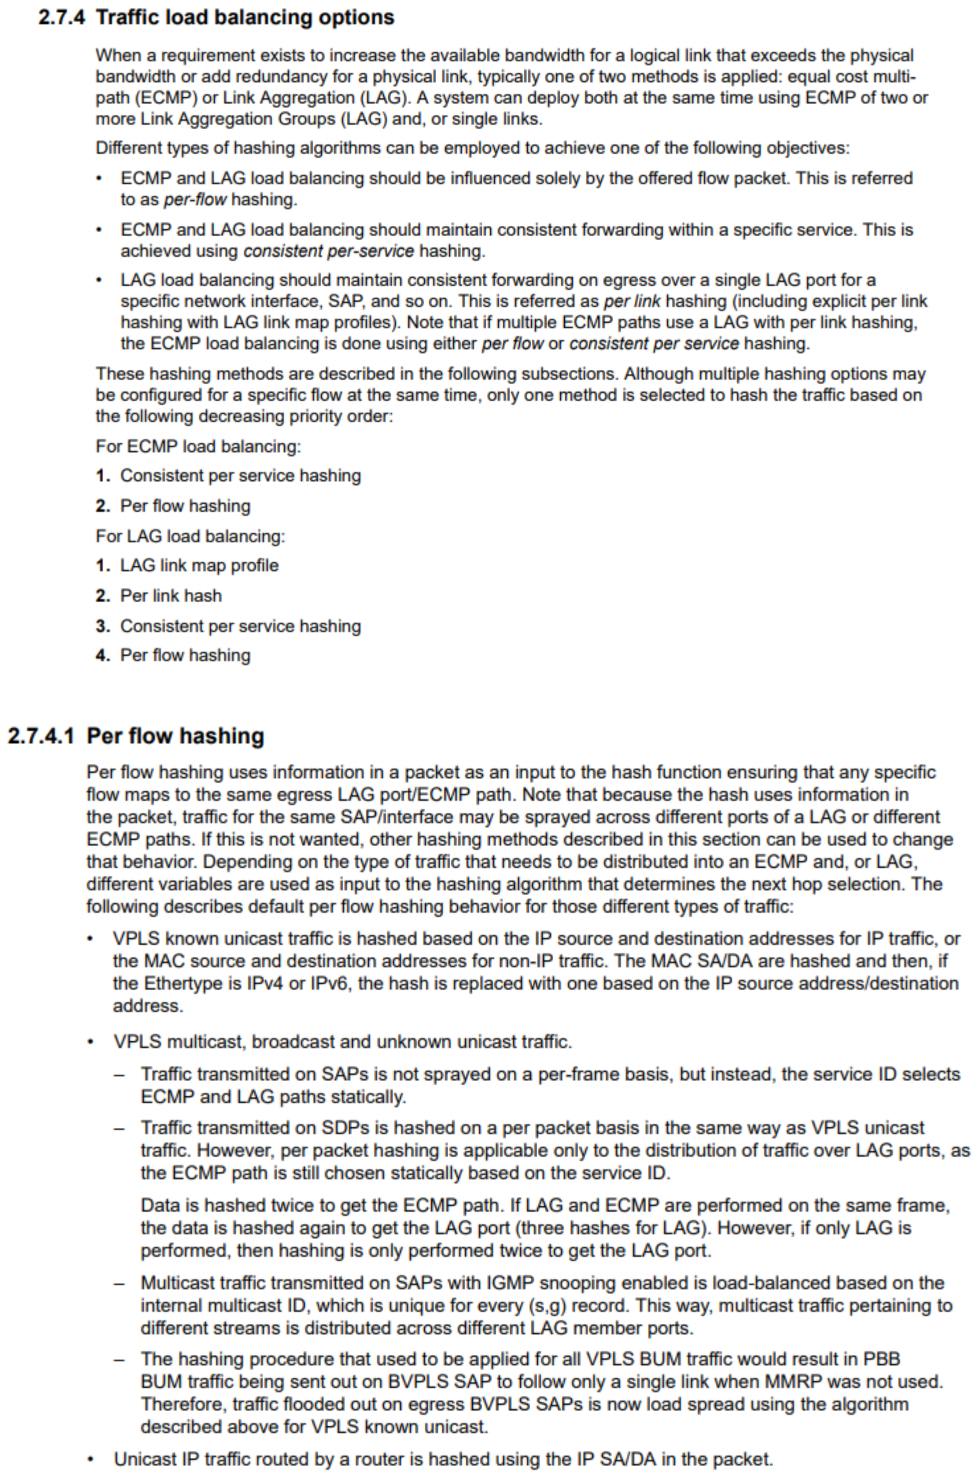

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%20 7950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 127 and 128 of PDF)

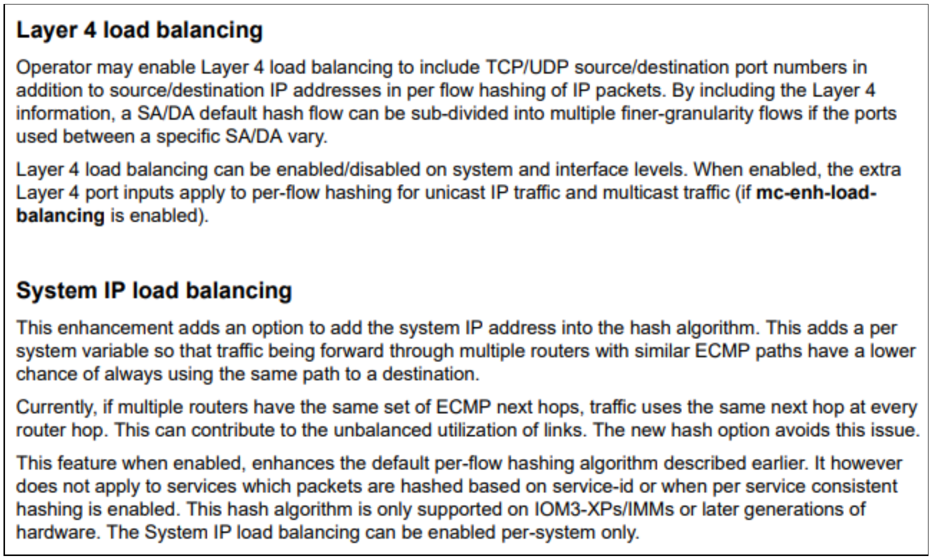

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

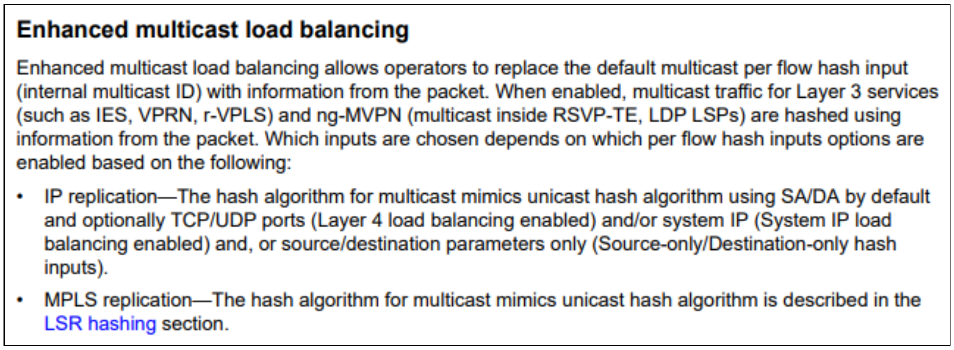

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 131 and 132 of PDF)

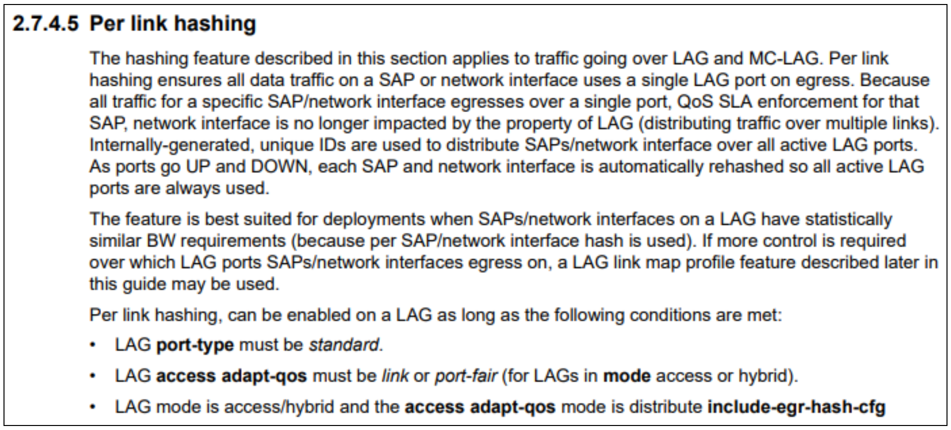

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207

950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

125. Defendants' 7450 Ethernet Service Switch comprises the first line card is arranged to assign to the data packets line card FMID (LC-FMID values selected from a first range of possible LC-FMID values, and wherein the switching fabric is arranged to map the LC-FMID values to respective central FMID (C-FMID values selected from a second range of possible C-FMID values that is smaller than the first range and to forward the data packets responsively to the C-FMID values:

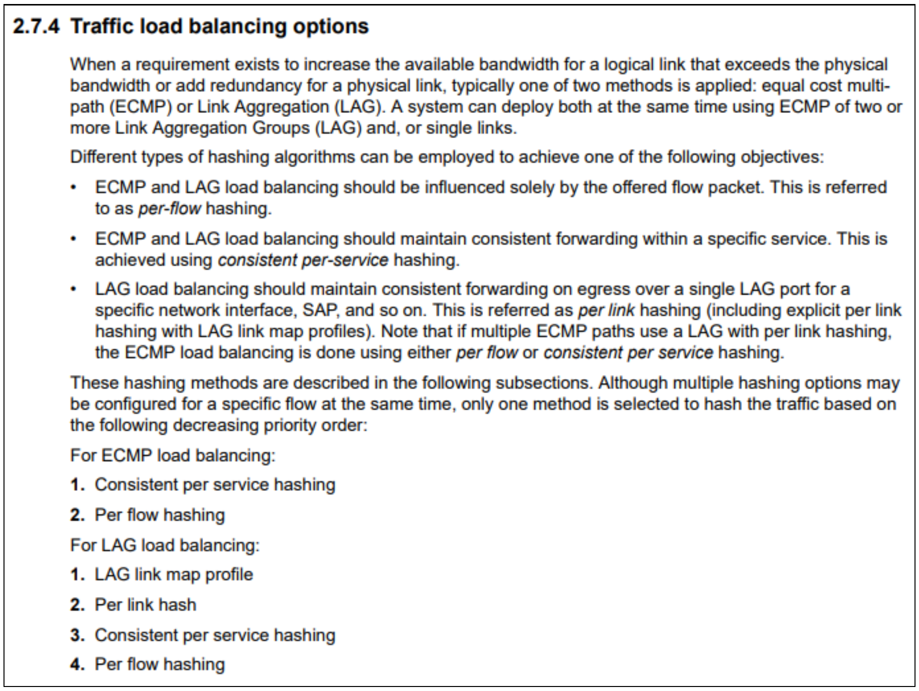

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

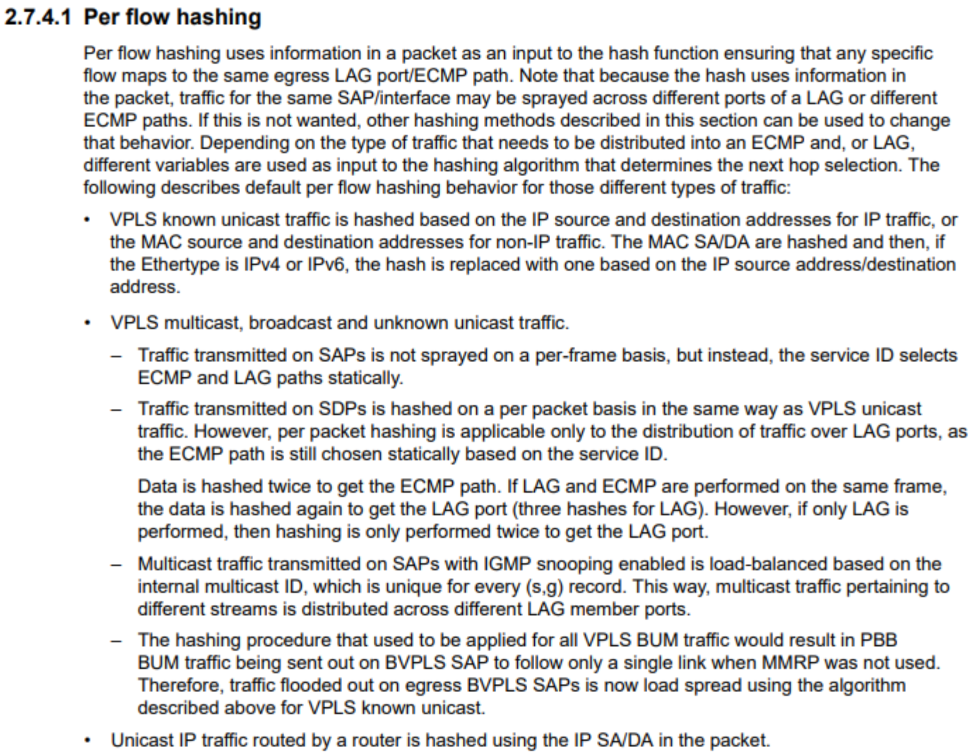

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 127 and 128 of PDF)

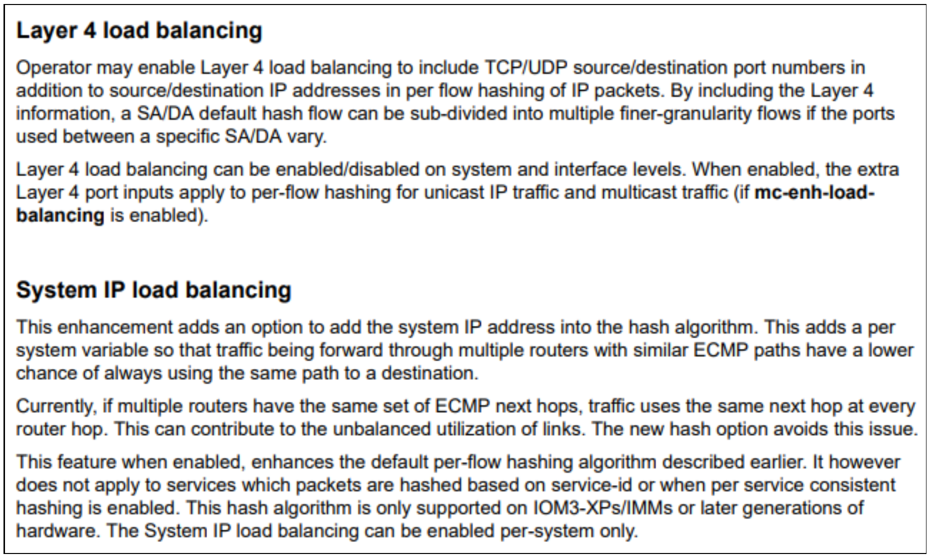

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

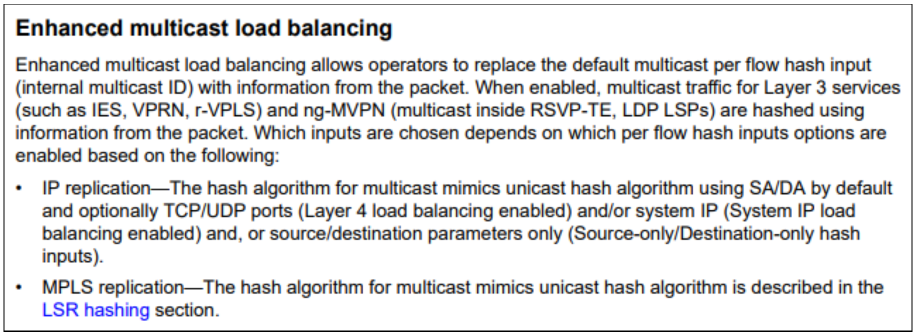

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 131 and 132 of PDF)

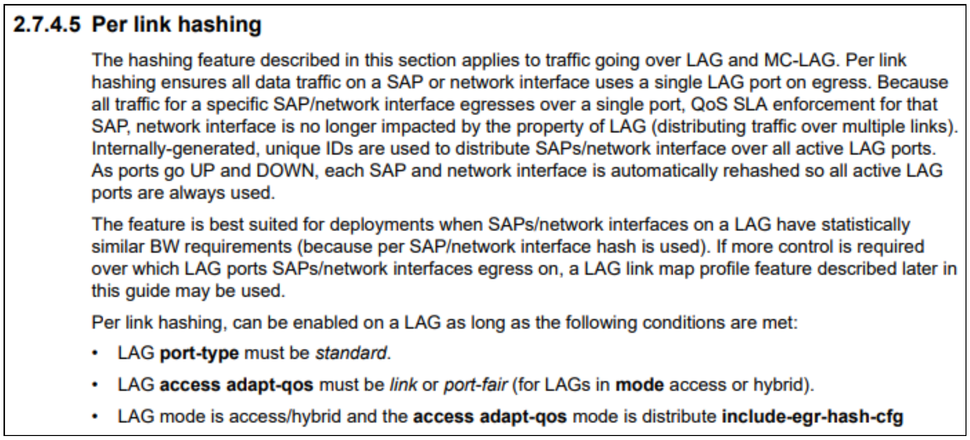

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

Willful Infringement

126. Defendants have had actual knowledge of the ’525 Patent and its infringement thereof at least as of service or other receipt of Plaintiff’s Complaint.

127. Defendants’ risk of infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendants.

128. Notwithstanding this knowledge, Defendants have knowingly or with reckless disregard willfully infringed the ’525 Patent. Defendants have thus had actual notice of the

infringement of the ’525 Patent and acted despite an objectively high likelihood that its actions constituted infringement of Plaintiff’s valid patent rights, either literally or equivalently.

129. This objective risk was either known or so obvious that it should have been known to Defendants. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

130. Defendants have induced and are knowingly inducing its customers and/or end users to directly infringe the ’525 Patent, with the specific intent to encourage such infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

131. Defendants have knowingly contributed to direct infringement by its customers by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the ’525 Accused Products which are not suitable for substantial non-infringing use, and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

132. Defendants’ indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe the ’525 Patent.

133. Defendants’ indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendants’ indirect infringement further includes providing application notes instructing its customers on infringing uses of the ’525 Accused Products. The ’525 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the ’525 Patent, either literally or equivalently. Defendants

know and intend that customers who purchase the ’525 Accused Products will use those products for their intended purpose. For example, Defendants’ United States website: https://www.nokia.com, instructs customers to use the ’525 Accused Products in numerous infringing applications. Furthermore, Defendants provide instructions and other resources via its “Doc Center” (https://documentation.nokia.com), “Nokia Service Routing Certification” and training courses from its “NokiaEDU Training Centers” (https://www.nokia.com/networks/training/src/courses/#open-enrollment), and elsewhere on using the ’525 Accused Products. Defendants’ customers directly infringe the ’525 Patent when they follow Defendants’ provided instructions on websites, videos, trainings, and elsewhere. Defendants’ customers who follow Defendants’ provided instructions directly infringe claims of the ’525 Patent.

134. In addition, Defendants specifically intend that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendants’ infringing products. Defendants know following its instructions directly infringes claims of the ’525 Patent, including for example Claim 12.

135. Defendants’ customers who follow Defendants’ provided instructions directly infringe the method of Claim 12 of the ’525 Patent.

136. Defendants instruct its customers to use the 7750 Service Router in a method for communication:

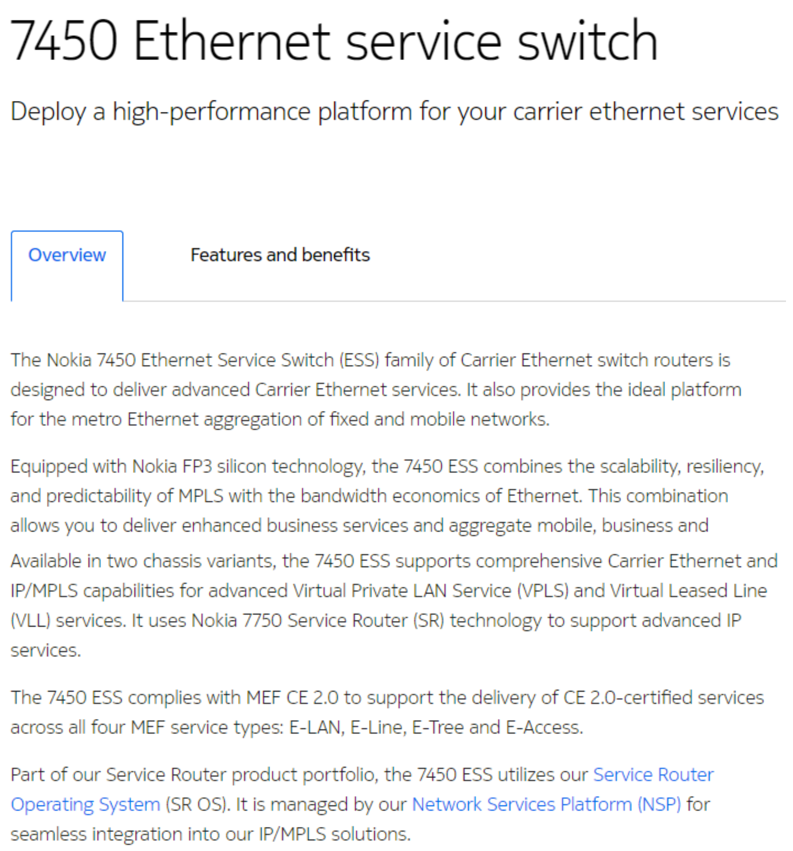

7450 Ethernet service switch

Deploy a high-performance platform for your carrier ethernet services

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers is designed to deliver advanced Carrier Ethernet services. It also provides the ideal platform for the metro Ethernet aggregation of fixed and mobile networks.

Equipped with Nokia FP3 silicon technology, the 7450 ESS combines the scalability, resiliency, and predictability of MPLS with the bandwidth economics of Ethernet. This combination allows you to deliver enhanced business services and aggregate mobile, business and

Available in two chassis variants, the 7450 ESS supports comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL) services. It uses Nokia 7750 Service Router (SR) technology to support advanced IP services.

The 7450 ESS complies with MEF CE 2.0 to support the delivery of CE 2.0-certified services across all four MEF service types: E-LAN, E-Line, E-Tree and E-Access.

Part of our Service Router product portfolio, the 7450 ESS utilizes our Service Router Operating System (SR OS). It is managed by our Network Services Platform (NSP) for seamless integration into our IP/MPLS solutions.

https://www.nokia.com/networks/products/7450-ethernet-service-switch/

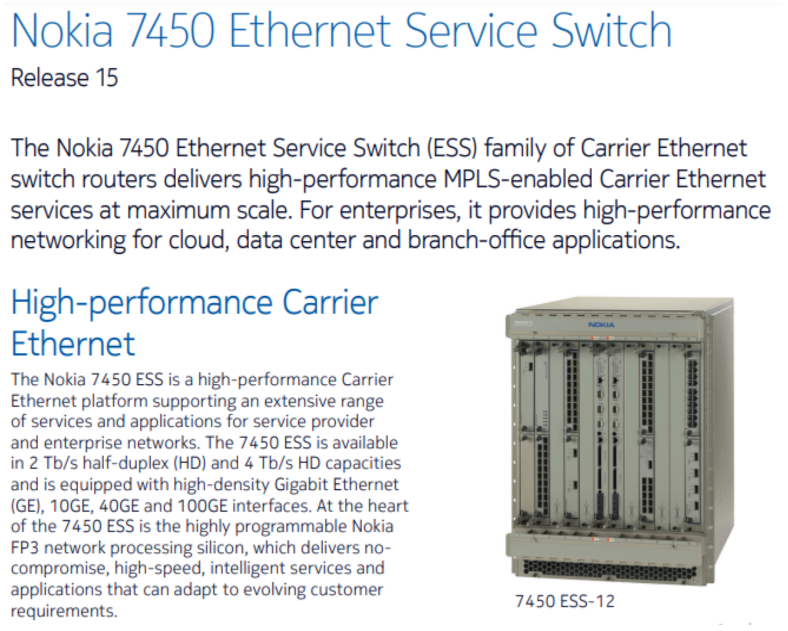

Nokia 7450 Ethernet Service Switch

Release 15

The Nokia 7450 Ethernet Service Switch (ESS) family of Carrier Ethernet switch routers delivers high-performance MPLS-enabled Carrier Ethernet services at maximum scale. For enterprises, it provides high-performance networking for cloud, data center and branch-office applications.

High-performance Carrier Ethernet

The Nokia 7450 ESS is a high-performance Carrier Ethernet platform supporting an extensive range of services and applications for service provider and enterprise networks. The 7450 ESS is available in 2 Tb/s half-duplex (HD) and 4 Tb/s HD capacities and is equipped with high-density Gigabit Ethernet (GE), 10GE, 40GE and 100GE interfaces. At the heart of the 7450 ESS is the highly programmable Nokia FP3 network processing silicon, which delivers no-compromise, high-speed, intelligent services and applications that can adapt to evolving customer requirements.

7450 ESS-12

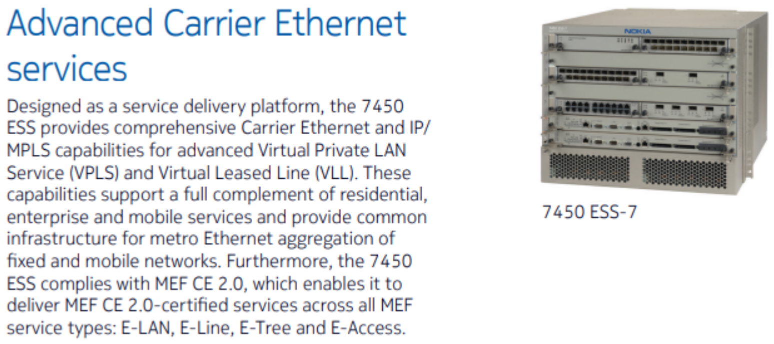

Advanced Carrier Ethernet services

Designed as a service delivery platform, the 7450 ESS provides comprehensive Carrier Ethernet and IP/MPLS capabilities for advanced Virtual Private LAN Service (VPLS) and Virtual Leased Line (VLL). These capabilities support a full complement of residential, enterprise and mobile services and provide common infrastructure for metro Ethernet aggregation of fixed and mobile networks. Furthermore, the 7450 ESS complies with MEF CE 2.0, which enables it to deliver MEF CE 2.0-certified services across all MEF service types: E-LAN, E-Line, E-Tree and E-Access.

7450 ESS-7

https://onestore.nokia.com/asset/f/164727 (Page 1 of PDF)

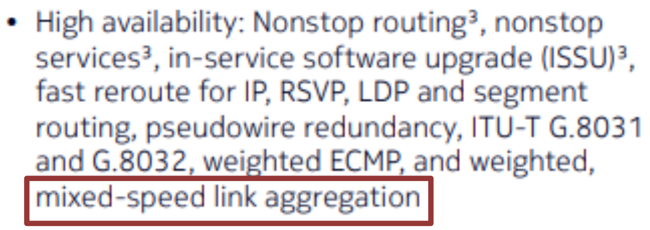

- High availability: Nonstop routing[3], nonstop services[3], in-service software upgrade (ISSU)[3], fast reroute for IP, RSVP, LDP and segment routing, pseudowire redundancy, ITU-T G.8031 and G.8032, weighted ECMP, and weighted, mixed-speed link aggregation

https://onestore.nokia.com/asset/f/164727 (Page 6 of PDF)

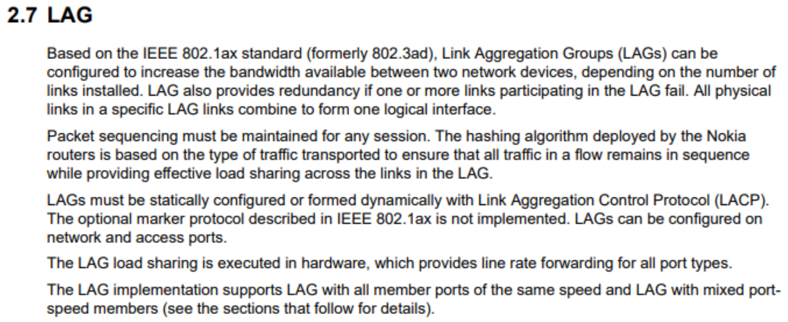

2.7 LAG

Based on the IEEE 802.1ax standard (formerly 802.3ad), Link Aggregation Groups (LAGs) can be configured to increase the bandwidth available between two network devices, depending on the number of links installed. LAG also provides redundancy if one or more links participating in the LAG fail. All physical links in a specific LAG links combine to form one logical interface.

Packet sequencing must be maintained for any session. The hashing algorithm deployed by the Nokia routers is based on the type of traffic transported to ensure that all traffic in a flow remains in sequence while providing effective load sharing across the links in the LAG.

LAGs must be statically configured or formed dynamically with Link Aggregation Control Protocol (LACP). The optional marker protocol described in IEEE 802.1ax is not implemented. LAGs can be configured on network and access ports.

The LAG load sharing is executed in hardware, which provides line rate forwarding for all port types.

The LAG implementation supports LAG with all member ports of the same speed and LAG with mixed port-speed members (see the sections that follow for details).

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 121 of PDF)

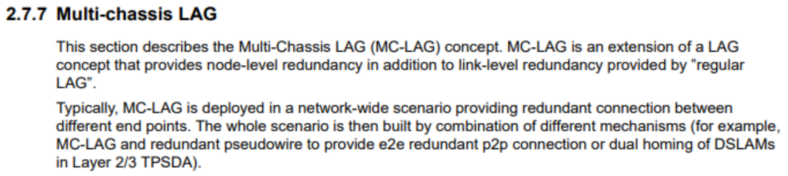

2.7.7 Multi-chassis LAG

This section describes the Multi-Chassis LAG (MC-LAG) concept. MC-LAG is an extension of a LAG concept that provides node-level redundancy in addition to link-level redundancy provided by "regular LAG".

Typically, MC-LAG is deployed in a network-wide scenario providing redundant connection between different end points. The whole scenario is then built by combination of different mechanisms (for example, MC-LAG and redundant pseudowire to provide e2e redundant p2p connection or dual homing of DSLAMs in Layer 2/3 TPSDA).

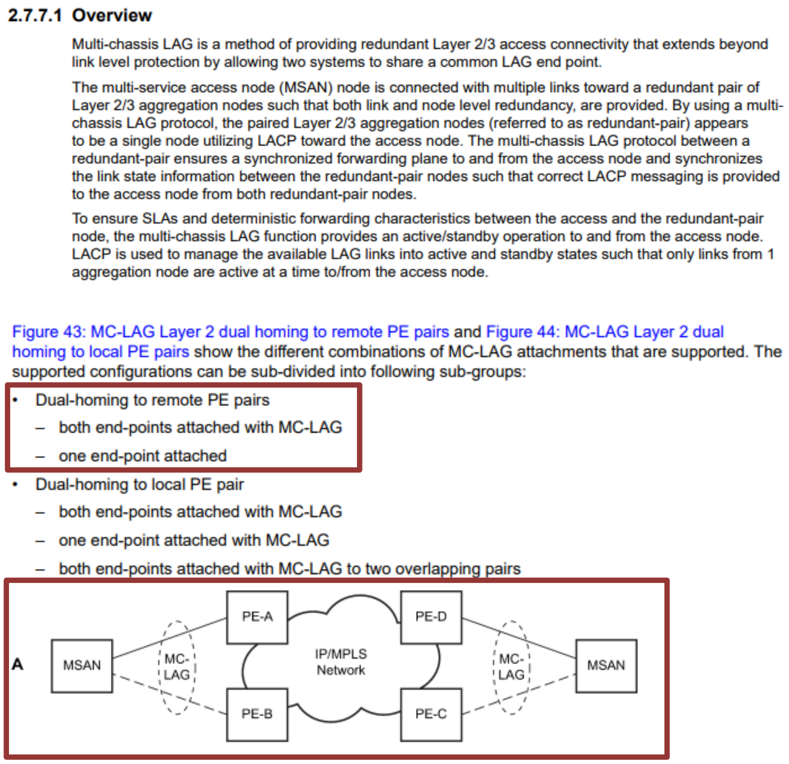

2.7.7.1 Overview

Multi-chassis LAG is a method of providing redundant Layer 2/3 access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point.

The multi-service access node (MSAN) node is connected with multiple links toward a redundant pair of Layer 2/3 aggregation nodes such that both link and node level redundancy, are provided. By using a multi-chassis LAG protocol, the paired Layer 2/3 aggregation nodes (referred to as redundant-pair) appears to be a single node utilizing LACP toward the access node. The multi-chassis LAG protocol between a redundant-pair ensures a synchronized forwarding plane to and from the access node and synchronizes the link state information between the redundant-pair nodes such that correct LACP messaging is provided to the access node from both redundant-pair nodes.

To ensure SLAs and deterministic forwarding characteristics between the access and the redundant-pair node, the multi-chassis LAG function provides an active/standby operation to and from the access node. LACP is used to manage the available LAG links into active and standby states such that only links from 1 aggregation node are active at a time to/from the access node.

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 145-147 of PDF)

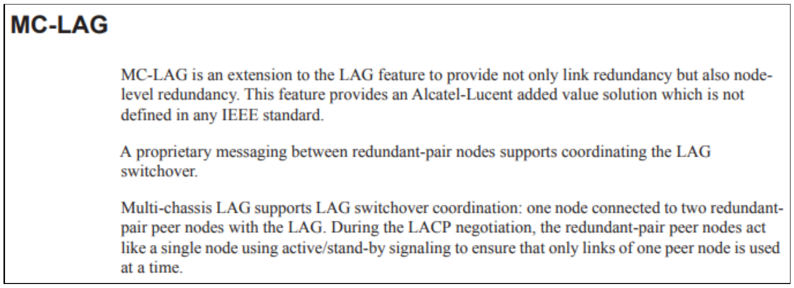

MC-LAG

MC-LAG is an extension to the LAG feature to provide not only link redundancy but also node-level redundancy. This feature provides an Alcatel-Lucent added value solution which is not defined in any IEEE standard.

A proprietary messaging between redundant-pair nodes supports coordinating the LAG switchover.

Multi-chassis LAG supports LAG switchover coordination: one node connected to two redundant-pair peer nodes with the LAG. During the LACP negotiation, the redundant-pair peer nodes act like a single node using active/stand-by signaling to ensure that only links of one peer node is used at a time.

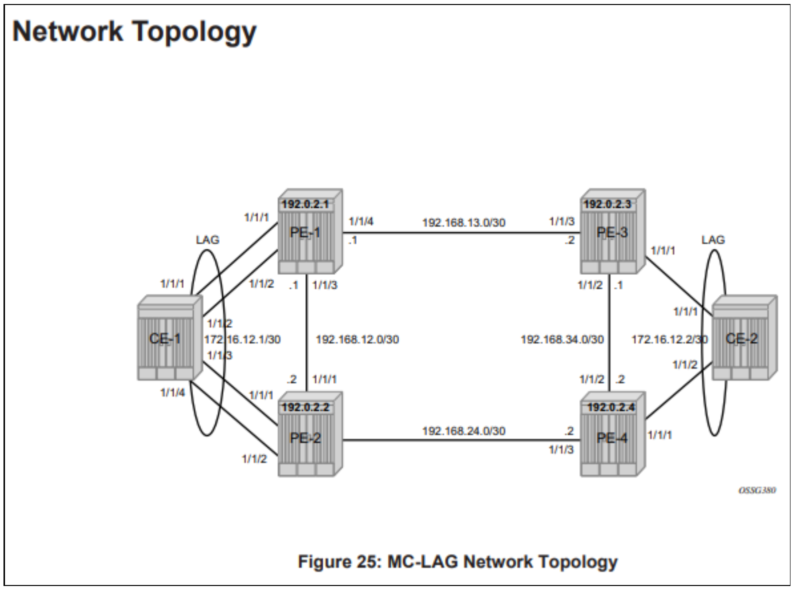

Network Topology

Figure 25: MC-LAG Network Topology

https://documentation.nokia.com/html/0_add-h-f/93-0267-HTML/7X50_Advanced_Configuration_Guide/MC-LAG%20.pdf (Pages 3 and 4 of PDF)

137. Defendants instruct its customers to use the 7450 Ethernet Service Switch in receiving data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of a plurality of ports, at least a subset of which is grouped in a link aggregation group (LAG).

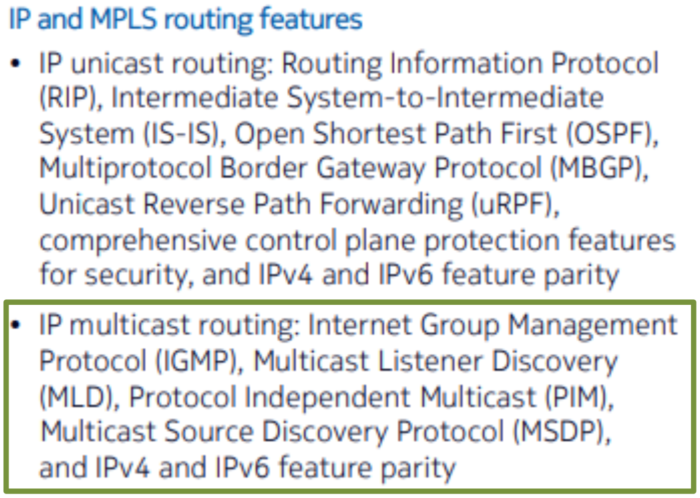

IP and MPLS routing features

- IP unicast routing: Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Multiprotocol Border Gateway Protocol (MBGP), Unicast Reverse Path Forwarding (uRPF), comprehensive control plane protection features for security, and IPv4 and IPv6 feature parity
- IP multicast routing: Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), and IPv4 and IPv6 feature parity

https://onestore.nokia.com/asset/f/164727 (Page 6 of PDF)

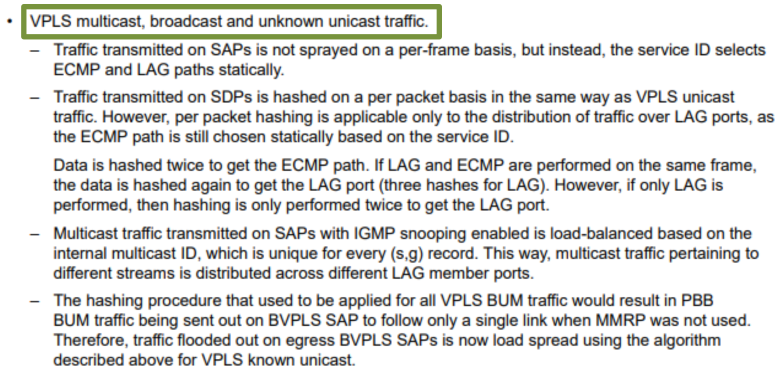

- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 128 of PDF)

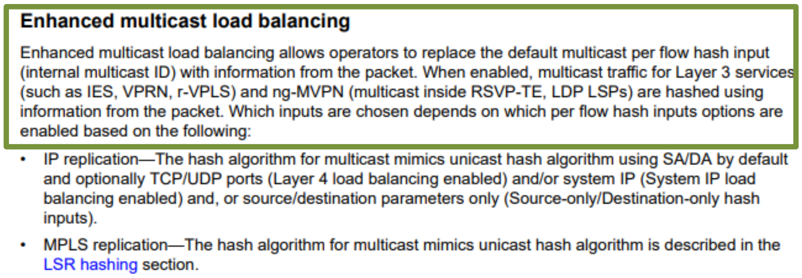

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Page 132 of PDF)

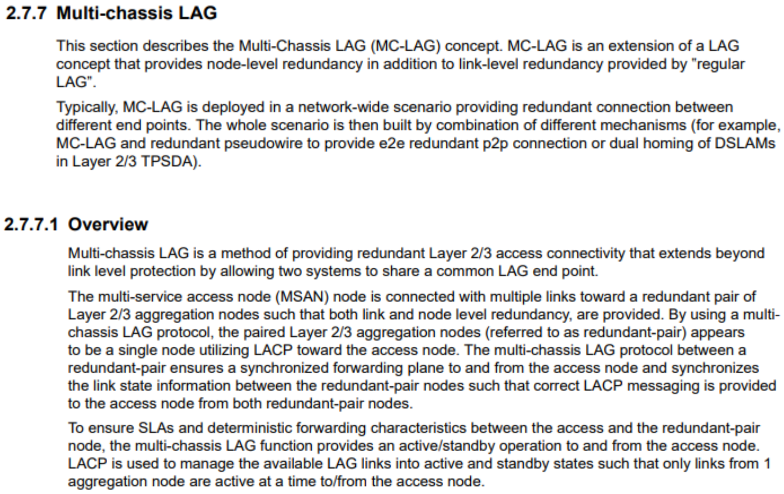

2.7.7 Multi-chassis LAG

This section describes the Multi-Chassis LAG (MC-LAG) concept. MC-LAG is an extension of a LAG concept that provides node-level redundancy in addition to link-level redundancy provided by "regular LAG".

Typically, MC-LAG is deployed in a network-wide scenario providing redundant connection between different end points. The whole scenario is then built by combination of different mechanisms (for example, MC-LAG and redundant pseudowire to provide e2e redundant p2p connection or dual homing of DSLAMs in Layer 2/3 TPSDA).

2.7.7.1 Overview

Multi-chassis LAG is a method of providing redundant Layer 2/3 access connectivity that extends beyond link level protection by allowing two systems to share a common LAG end point.

The multi-service access node (MSAN) node is connected with multiple links toward a redundant pair of Layer 2/3 aggregation nodes such that both link and node level redundancy, are provided. By using a multi-chassis LAG protocol, the paired Layer 2/3 aggregation nodes (referred to as redundant-pair) appears to be a single node utilizing LACP toward the access node. The multi-chassis LAG protocol between a redundant-pair ensures a synchronized forwarding plane to and from the access node and synchronizes the link state information between the redundant-pair nodes such that correct LACP messaging is provided to the access node from both redundant-pair nodes.

To ensure SLAs and deterministic forwarding characteristics between the access and the redundant-pair node, the multi-chassis LAG function provides an active/standby operation to and from the access node. LACP is used to manage the available LAG links into active and standby states such that only links from 1 aggregation node are active at a time to/from the access node.

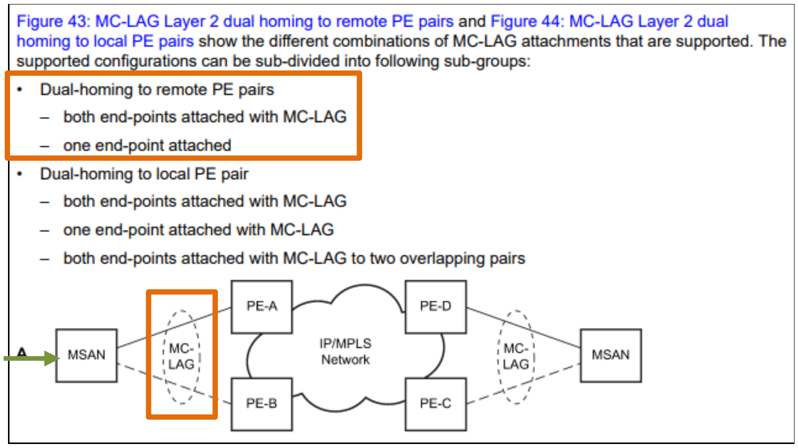

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 145-147 of PDF)

138. Defendants instruct its customers to use the 7450 Ethernet Service Switch in processing the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group.

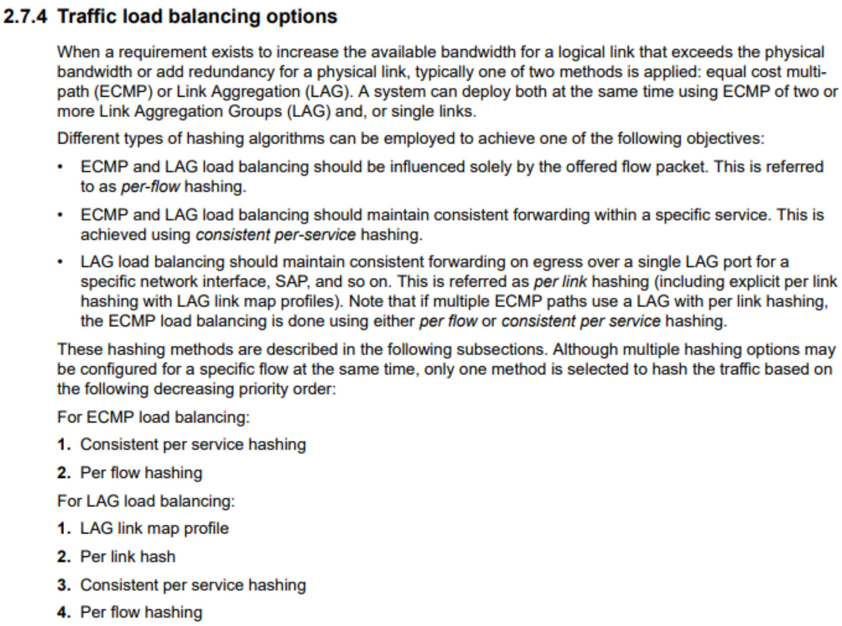

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

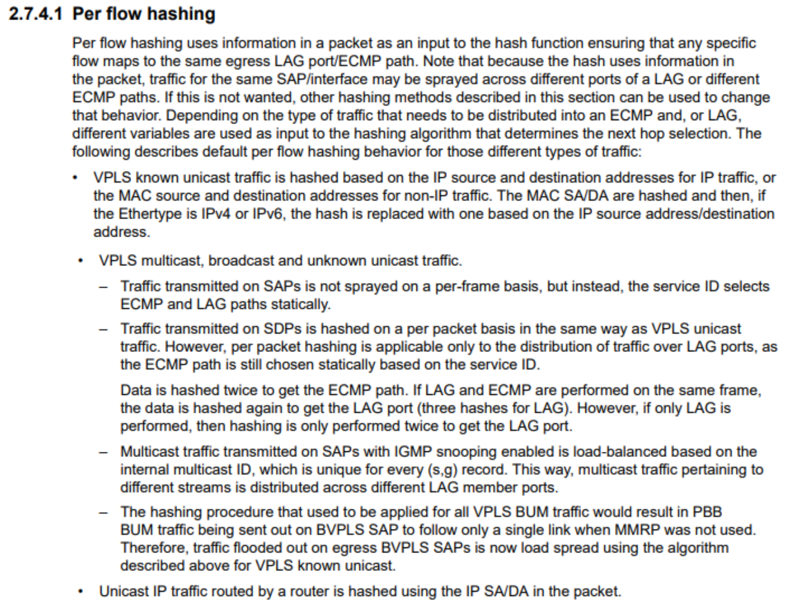

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 127 and 128 of PDF)

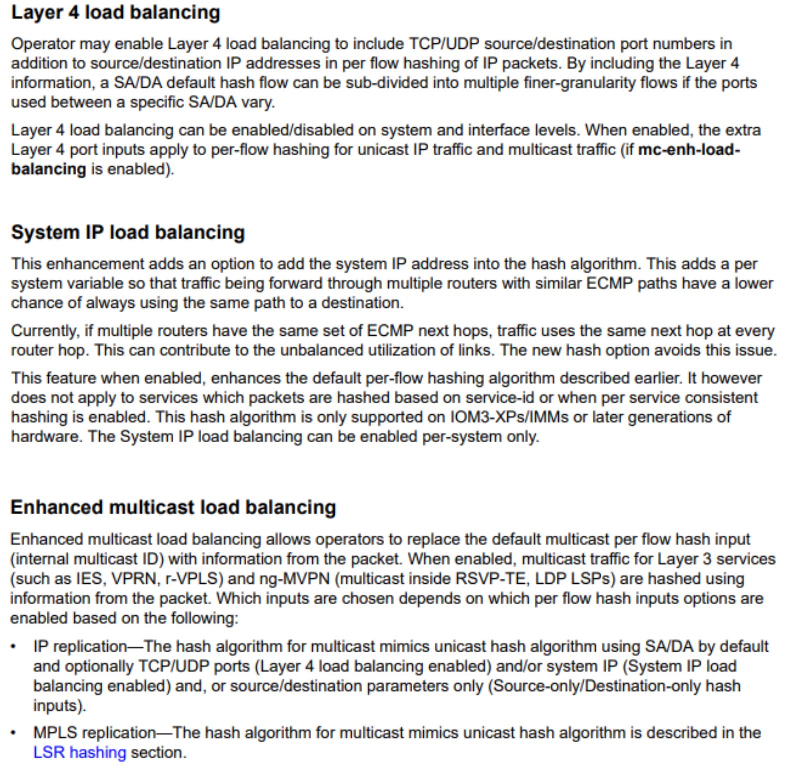

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207

950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 131 and 132 of PDF)

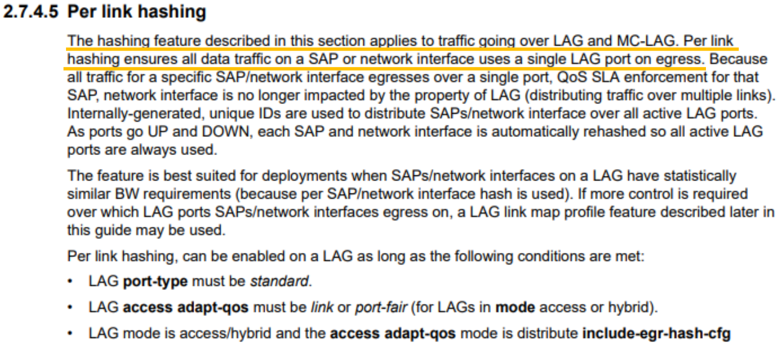

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

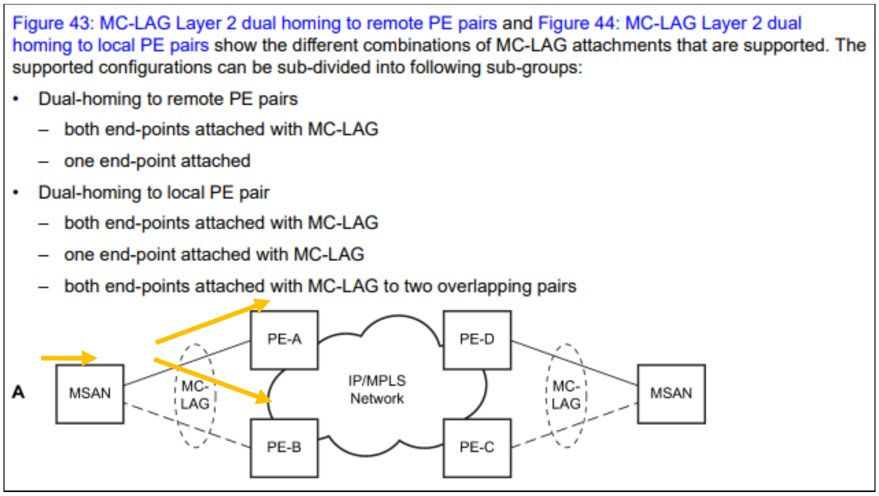

Figure 43: MC-LAG Layer 2 dual homing to remote PE pairs and Figure 44: MC-LAG Layer 2 dual homing to local PE pairs show the different combinations of MC-LAG attachments that are supported. The supported configurations can be sub-divided into following sub-groups:

- Dual-homing to remote PE pairs
  - both end-points attached with MC-LAG
  - one end-point attached
- Dual-homing to local PE pair
  - both end-points attached with MC-LAG
  - one end-point attached with MC-LAG
  - both end-points attached with MC-LAG to two overlapping pairs

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Page 147 of PDF)

139. Defendants instruct its customers to use the 7450 Ethernet Service Switch in allocating to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, the FMID value being associated with one of the ports in the subset and forwarding the single copy to the port associated with the allocated FMID value.

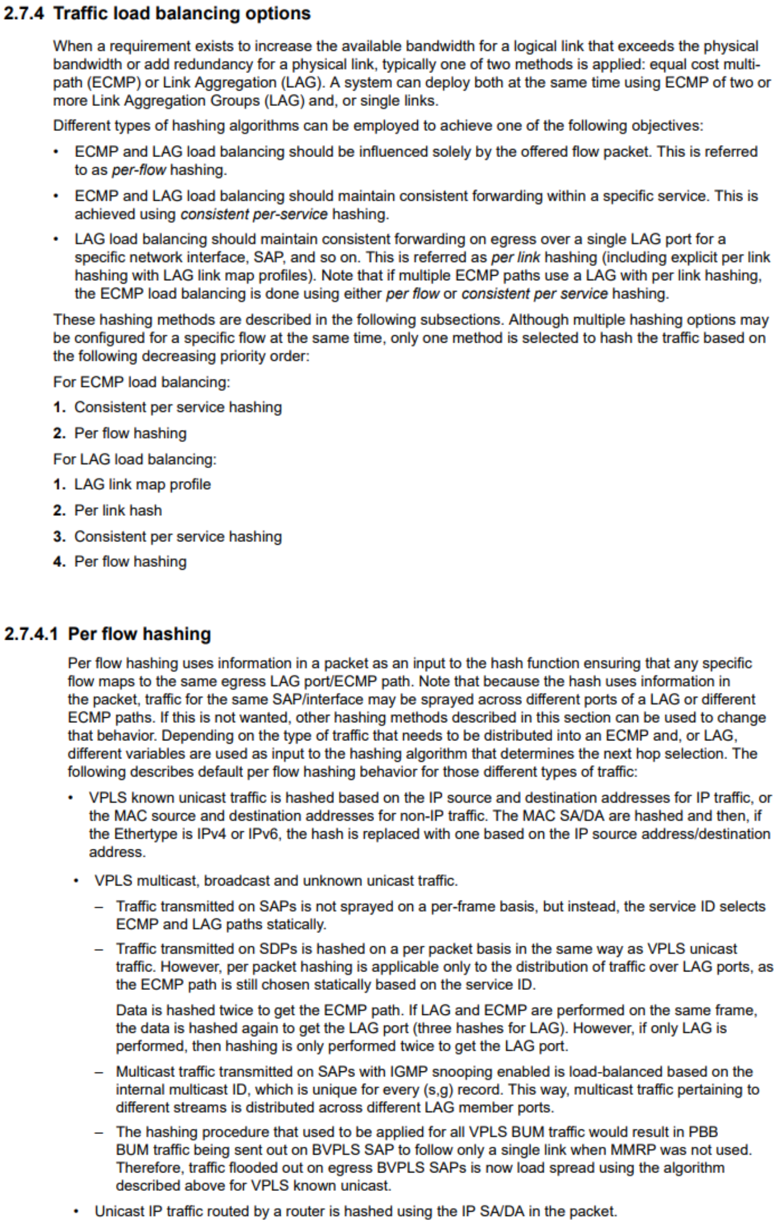

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 127 and 128 of PDF)

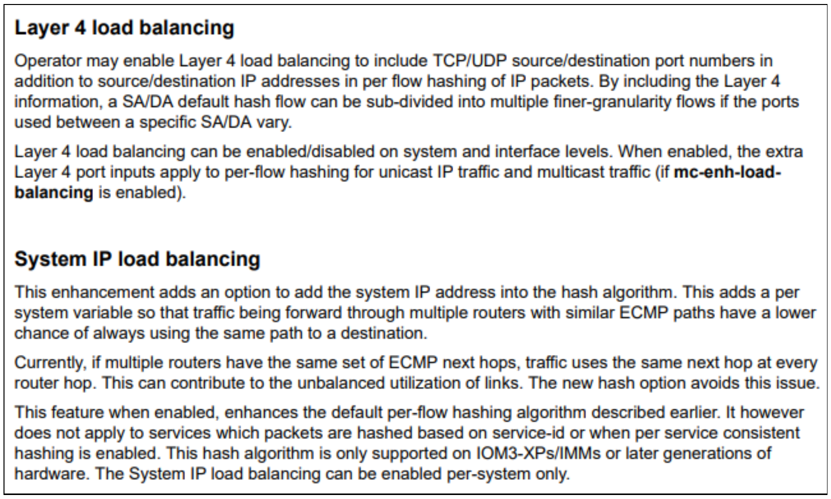

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

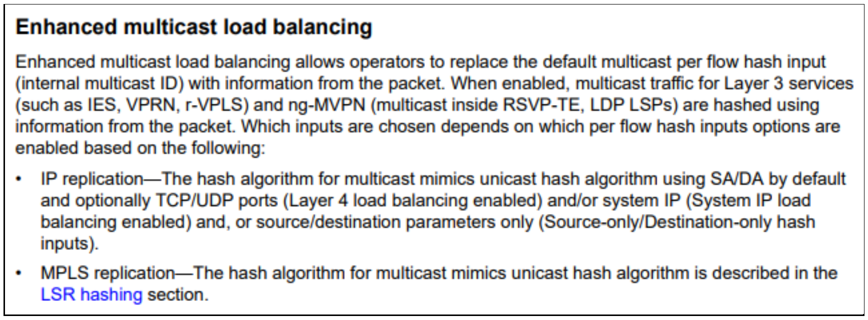

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 131 and 132 of PDF)

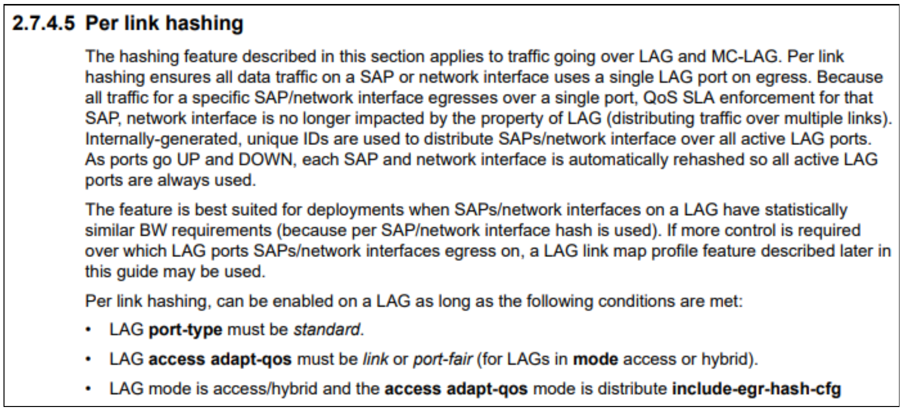

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 138 of PDF)

140. Defendants instruct its customers to use the 7450 Ethernet Service Switch for each of the received data packets, processing the data packets comprises allocating the FMID value to the data packet by a first line card connected to a first port via which the data packet is received, and configuring a second line card connected to a second port to which the data packet is to be sent and a switching fabric interconnecting the first and second line cards to forward the data packet responsively to the FMID value.

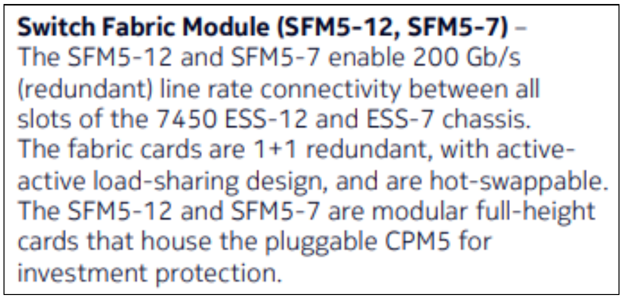

Switch Fabric Module (SFM5-12, SFM5-7) – The SFM5-12 and SFM5-7 enable 200 Gb/s (redundant) line rate connectivity between all slots of the 7450 ESS-12 and ESS-7 chassis. The fabric cards are 1+1 redundant, with active-active load-sharing design, and are hot-swappable. The SFM5-12 and SFM5-7 are modular full-height cards that house the pluggable CPM5 for investment protection.

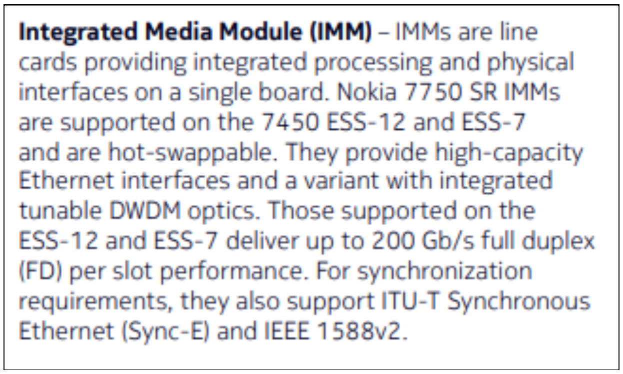

Integrated Media Module (IMM) – IMMs are line cards providing integrated processing and physical interfaces on a single board. Nokia 7750 SR IMMs are supported on the 7450 ESS-12 and ESS-7 and are hot-swappable. They provide high-capacity Ethernet interfaces and a variant with integrated tunable DWDM optics. Those supported on the ESS-12 and ESS-7 deliver up to 200 Gb/s full duplex (FD) per slot performance. For synchronization requirements, they also support ITU-T Synchronous Ethernet (Sync-E) and IEEE 1588v2.

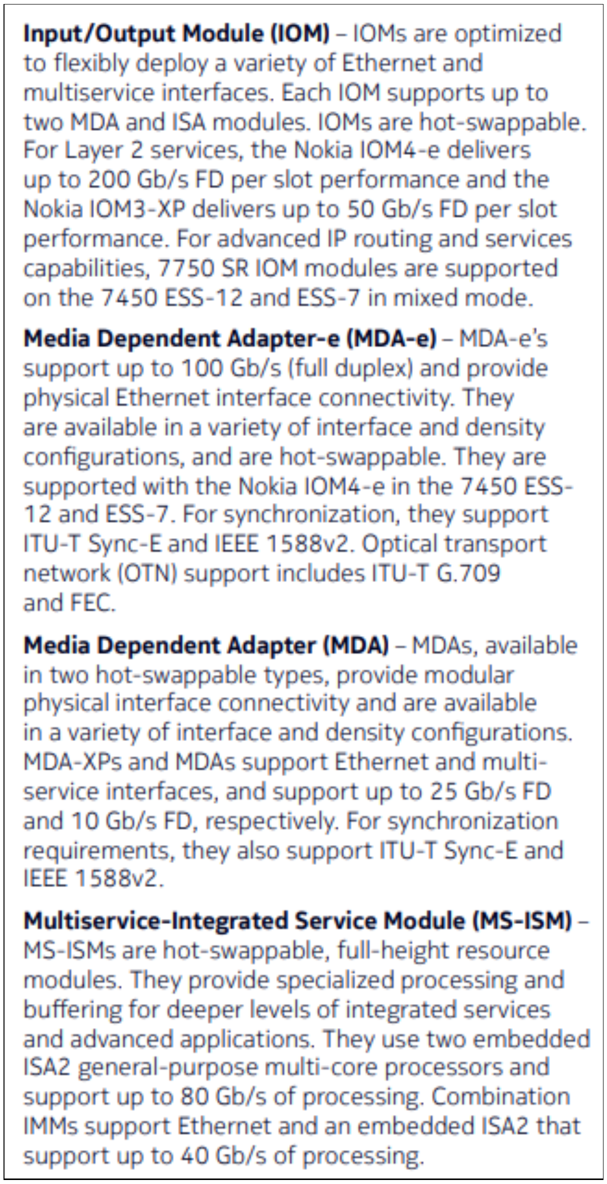

Input/Output Module (IOM) – IOMs are optimized to flexibly deploy a variety of Ethernet and multiservice interfaces. Each IOM supports up to two MDA and ISA modules. IOMs are hot-swappable. For Layer 2 services, the Nokia IOM4-e delivers up to 200 Gb/s FD per slot performance and the Nokia IOM3-XP delivers up to 50 Gb/s FD per slot performance. For advanced IP routing and services capabilities, 7750 SR IOM modules are supported on the 7450 ESS-12 and ESS-7 in mixed mode.

Media Dependent Adapter-e (MDA-e) – MDA-e's support up to 100 Gb/s (full duplex) and provide physical Ethernet interface connectivity. They are available in a variety of interface and density configurations, and are hot-swappable. They are supported with the Nokia IOM4-e in the 7450 ESS-12 and ESS-7. For synchronization, they support ITU-T Sync-E and IEEE 1588v2. Optical transport network (OTN) support includes ITU-T G.709 and FEC.

Media Dependent Adapter (MDA) – MDAs, available in two hot-swappable types, provide modular physical interface connectivity and are available in a variety of interface and density configurations. MDA-XPs and MDAs support Ethernet and multi-service interfaces, and support up to 25 Gb/s FD and 10 Gb/s FD, respectively. For synchronization requirements, they also support ITU-T Sync-E and IEEE 1588v2.

Multiservice-Integrated Service Module (MS-ISM) – MS-ISMs are hot-swappable, full-height resource modules. They provide specialized processing and buffering for deeper levels of integrated services and advanced applications. They use two embedded ISA2 general-purpose multi-core processors and support up to 80 Gb/s of processing. Combination IMMs support Ethernet and an embedded ISA2 that support up to 40 Gb/s of processing.

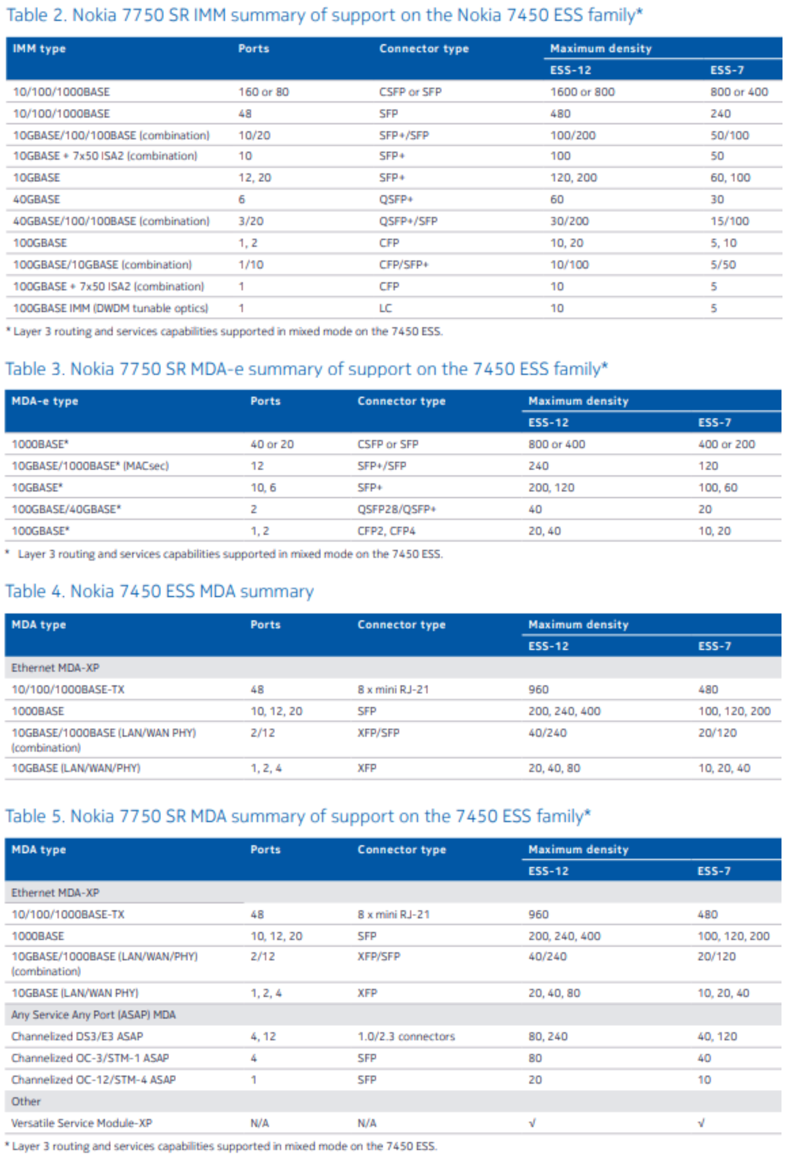

Table 2. Nokia 7750 SR IMM summary of support on the Nokia 7450 ESS family*

IMM type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
10/100/1000BASE	160 or 80	CSFP or SFP	1600 or 800	800 or 400
10/100/1000BASE	48	SFP	480	240
10GBASE/100/100BASE (combination)	10/20	SFP+/SFP	100/200	50/100
10GBASE + 7x50 ISA2 (combination)	10	SFP+	100	50
10GBASE	12, 20	SFP+	120, 200	60, 100
40GBASE	6	QSFP+	60	30
40GBASE/100/100BASE (combination)	3/20	QSFP+/SFP	30/200	15/100
100GBASE	1, 2	CFP	10, 20	5, 10
100GBASE/10GBASE (combination)	1/10	CFP/SFP+	10/100	5/50
100GBASE + 7x50 ISA2 (combination)	1	CFP	10	5
100GBASE IMM (DWDM tunable optics)	1	LC	10	5

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

Table 3. Nokia 7750 SR MDA-e summary of support on the 7450 ESS family*

MDA-e type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
1000BASE*	40 or 20	CSFP or SFP	800 or 400	400 or 200
10GBASE/1000BASE* (MACsec)	12	SFP+/SFP	240	120
10GBASE*	10, 6	SFP+	200, 120	100, 60
100GBASE/40GBASE*	2	QSFP28/QSFP+	40	20
100GBASE*	1, 2	CFP2, CFP4	20, 40	10, 20

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

Table 4. Nokia 7450 ESS MDA summary

MDA type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
Ethernet MDA-XP				
10/100/1000BASE-TX	48	8 x mini RJ-21	960	480
1000BASE	10, 12, 20	SFP	200, 240, 400	100, 120, 200
10GBASE/1000BASE (LAN/WAN PHY) (combination)	2/12	XFP/SFP	40/240	20/120
10GBASE (LAN/WAN/PHY)	1, 2, 4	XFP	20, 40, 80	10, 20, 40

Table 5. Nokia 7750 SR MDA summary of support on the 7450 ESS family*

MDA type	Ports	Connector type	Maximum density	
			ESS-12	ESS-7
Ethernet MDA-XP				
10/100/1000BASE-TX	48	8 x mini RJ-21	960	480
1000BASE	10, 12, 20	SFP	200, 240, 400	100, 120, 200
10GBASE/1000BASE (LAN/WAN/PHY) (combination)	2/12	XFP/SFP	40/240	20/120
10GBASE (LAN/WAN PHY)	1, 2, 4	XFP	20, 40, 80	10, 20, 40
Any Service Any Port (ASAP) MDA				
Channelized DS3/E3 ASAP	4, 12	1.0/2.3 connectors	80, 240	40, 120
Channelized OC-3/STM-1 ASAP	4	SFP	80	40
Channelized OC-12/STM-4 ASAP	1	SFP	20	10
Other				
Versatile Service Module-XP	N/A	N/A	√	√

* Layer 3 routing and services capabilities supported in mixed mode on the 7450 ESS.

https://onestore.nokia.com/asset/164727 (Pages 3-5 of PDF)

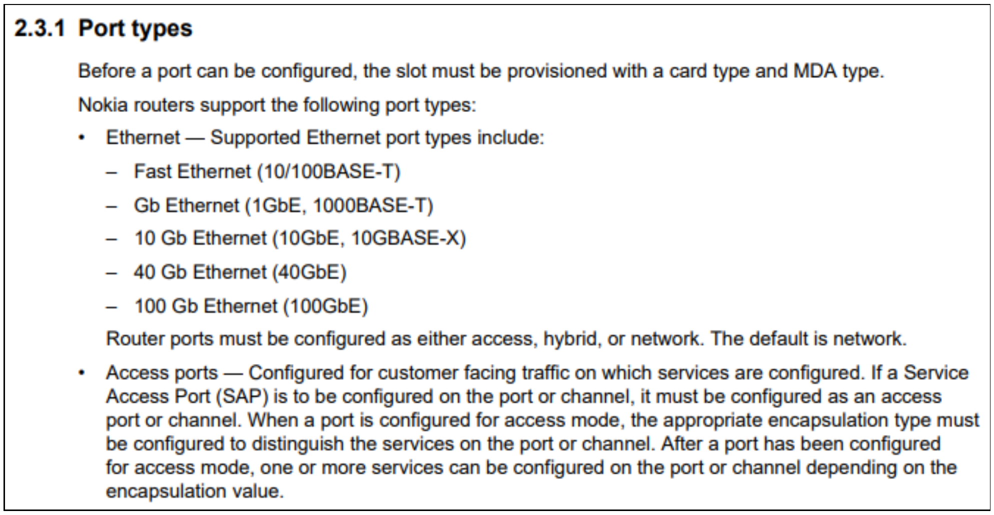

2.3.1 Port types

Before a port can be configured, the slot must be provisioned with a card type and MDA type.

Nokia routers support the following port types:

- Ethernet — Supported Ethernet port types include:
  - Fast Ethernet (10/100BASE-T)
  - Gb Ethernet (1GbE, 1000BASE-T)
  - 10 Gb Ethernet (10GbE, 10GBASE-X)
  - 40 Gb Ethernet (40GbE)
  - 100 Gb Ethernet (100GbE)

  Router ports must be configured as either access, hybrid, or network. The default is network.
- Access ports — Configured for customer facing traffic on which services are configured. If a Service Access Port (SAP) is to be configured on the port or channel, it must be configured as an access port or channel. When a port is configured for access mode, the appropriate encapsulation type must be configured to distinguish the services on the port or channel. After a port has been configured for access mode, one or more services can be configured on the port or channel depending on the encapsulation value.

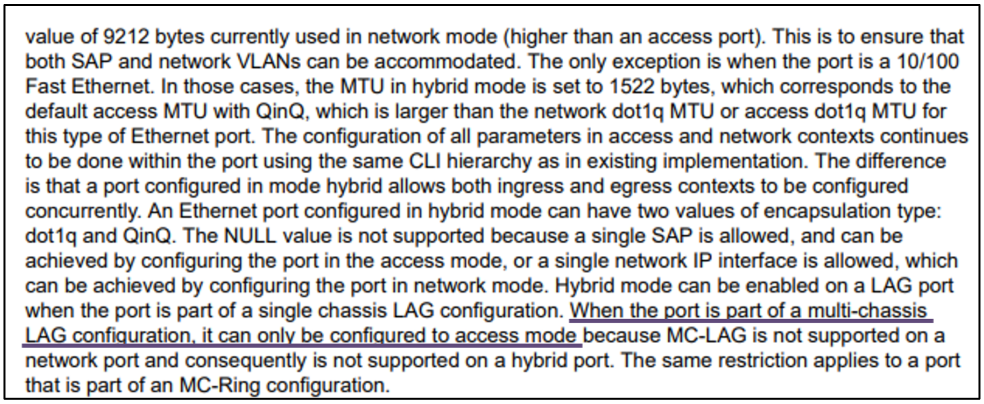

value of 9212 bytes currently used in network mode (higher than an access port). This is to ensure that both SAP and network VLANs can be accommodated. The only exception is when the port is a 10/100 Fast Ethernet. In those cases, the MTU in hybrid mode is set to 1522 bytes, which corresponds to the default access MTU with QinQ, which is larger than the network dot1q MTU or access dot1q MTU for this type of Ethernet port. The configuration of all parameters in access and network contexts continues to be done within the port using the same CLI hierarchy as in existing implementation. The difference is that a port configured in mode hybrid allows both ingress and egress contexts to be configured concurrently. An Ethernet port configured in hybrid mode can have two values of encapsulation type: dot1q and QinQ. The NULL value is not supported because a single SAP is allowed, and can be achieved by configuring the port in the access mode, or a single network IP interface is allowed, which can be achieved by configuring the port in network mode. Hybrid mode can be enabled on a LAG port when the port is part of a single chassis LAG configuration. When the port is part of a multi-chassis LAG configuration, it can only be configured to access mode because MC-LAG is not supported on a network port and consequently is not supported on a hybrid port. The same restriction applies to a port that is part of an MC-Ring configuration.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 28 and 29 of PDF)

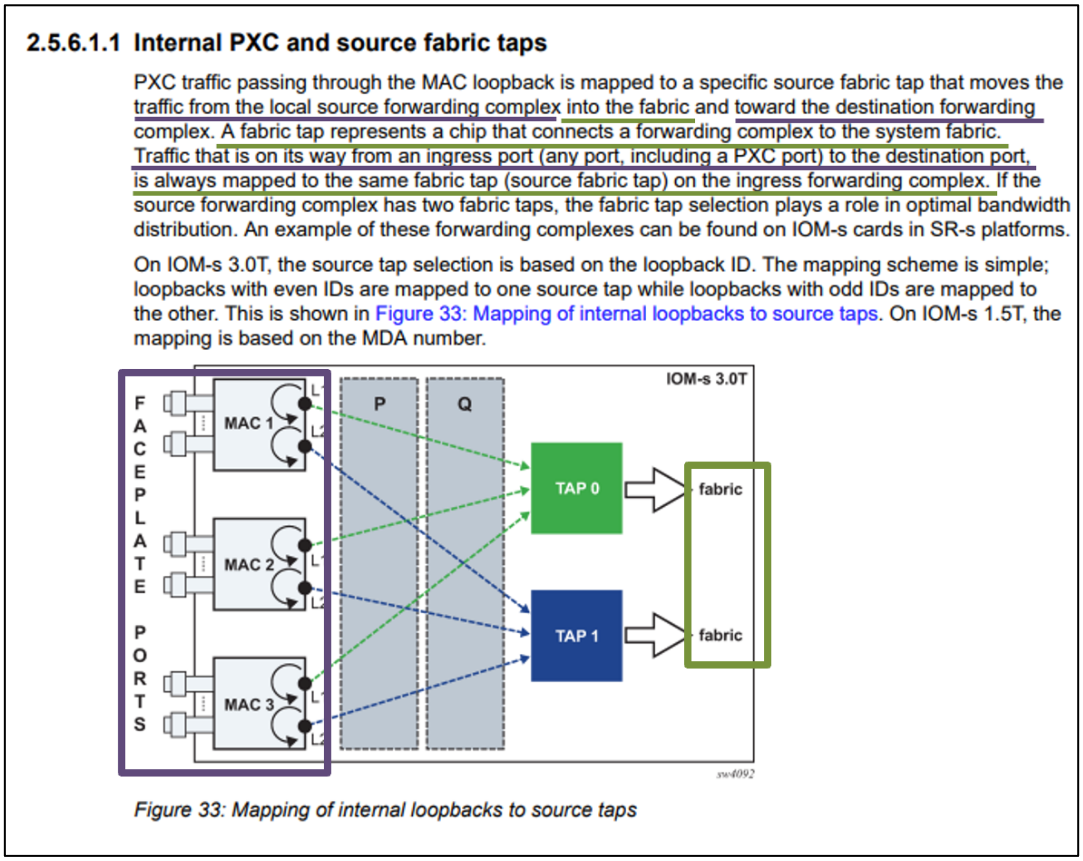

2.5.6.1.1 Internal PXC and source fabric taps

PXC traffic passing through the MAC loopback is mapped to a specific source fabric tap that moves the traffic from the local source forwarding complex into the fabric and toward the destination forwarding complex. A fabric tap represents a chip that connects a forwarding complex to the system fabric. Traffic that is on its way from an ingress port (any port, including a PXC port) to the destination port, is always mapped to the same fabric tap (source fabric tap) on the ingress forwarding complex. If the source forwarding complex has two fabric taps, the fabric tap selection plays a role in optimal bandwidth distribution. An example of these forwarding complexes can be found on IOM-s cards in SR-s platforms.

On IOM-s 3.0T, the source tap selection is based on the loopback ID. The mapping scheme is simple; loopbacks with even IDs are mapped to one source tap while loopbacks with odd IDs are mapped to the other. This is shown in Figure 33: Mapping of internal loopbacks to source taps. On IOM-s 1.5T, the mapping is based on the MDA number.

Figure 33: Mapping of internal loopbacks to source taps

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Page 109 of PDF)

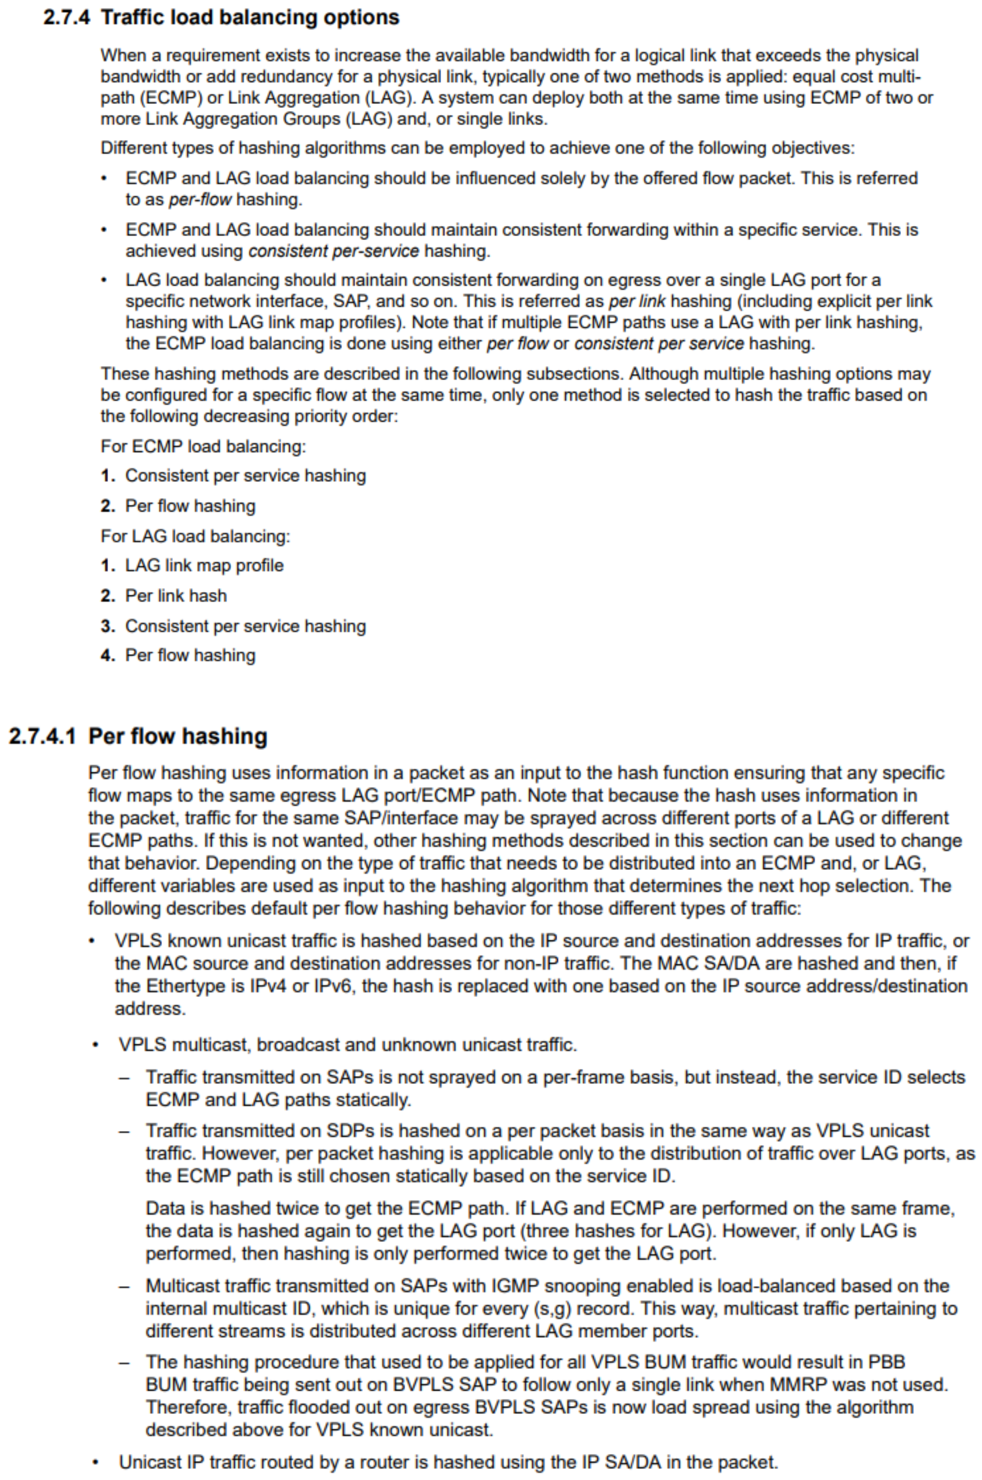

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 127 and 128 of PDF)

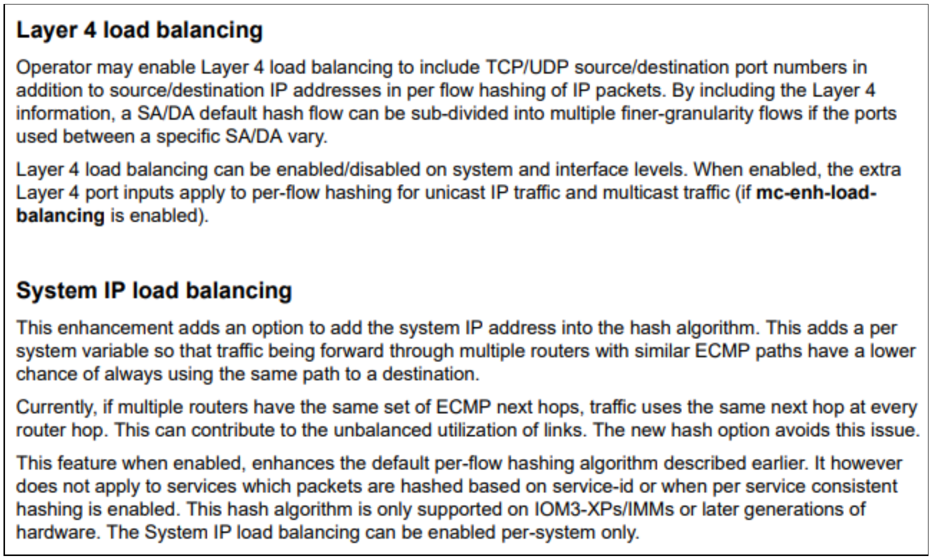

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

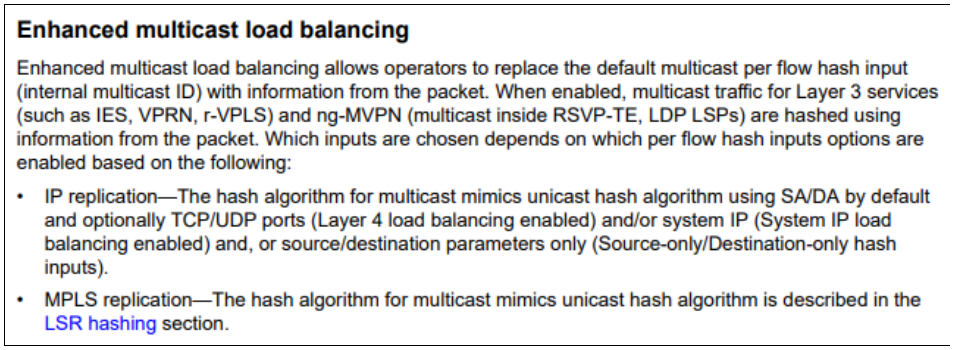

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 131 and 132 of PDF)

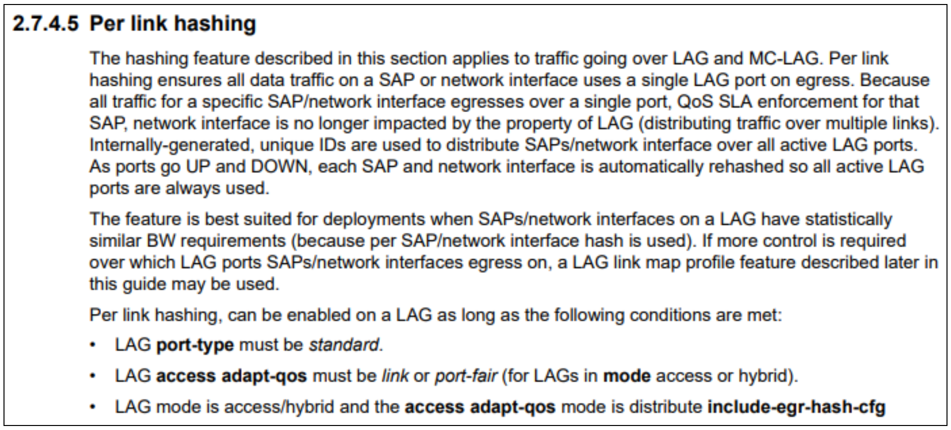

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207

950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

141. Defendants instruct its customers to use the 7450 Ethernet Service Switch in allocating the FMID value comprises assigning to the data packets line card FMID (LC-FMID values selected from a first range of possible LC-FMID values, and wherein configuring the switching fabric comprises mapping the LC-FMID values to respective central FMID (C-FMID) values selected from a second range of possible C-FMID values that is smaller than the first range and forwarding the data packets responsively to the C-FMID values.

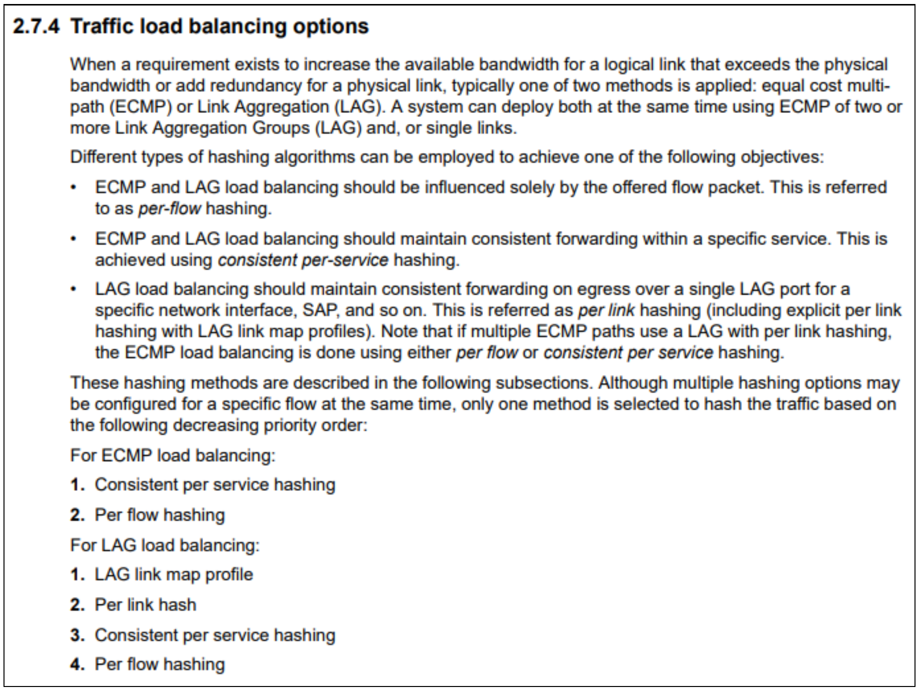

2.7.4 Traffic load balancing options

When a requirement exists to increase the available bandwidth for a logical link that exceeds the physical bandwidth or add redundancy for a physical link, typically one of two methods is applied: equal cost multi-path (ECMP) or Link Aggregation (LAG). A system can deploy both at the same time using ECMP of two or more Link Aggregation Groups (LAG) and, or single links.

Different types of hashing algorithms can be employed to achieve one of the following objectives:

- ECMP and LAG load balancing should be influenced solely by the offered flow packet. This is referred to as *per-flow* hashing.
- ECMP and LAG load balancing should maintain consistent forwarding within a specific service. This is achieved using *consistent per-service* hashing.
- LAG load balancing should maintain consistent forwarding on egress over a single LAG port for a specific network interface, SAP, and so on. This is referred as *per link* hashing (including explicit per link hashing with LAG link map profiles). Note that if multiple ECMP paths use a LAG with per link hashing, the ECMP load balancing is done using either *per flow* or *consistent per service* hashing.

These hashing methods are described in the following subsections. Although multiple hashing options may be configured for a specific flow at the same time, only one method is selected to hash the traffic based on the following decreasing priority order:

For ECMP load balancing:

1. Consistent per service hashing
2. Per flow hashing

For LAG load balancing:

1. LAG link map profile
2. Per link hash
3. Consistent per service hashing
4. Per flow hashing

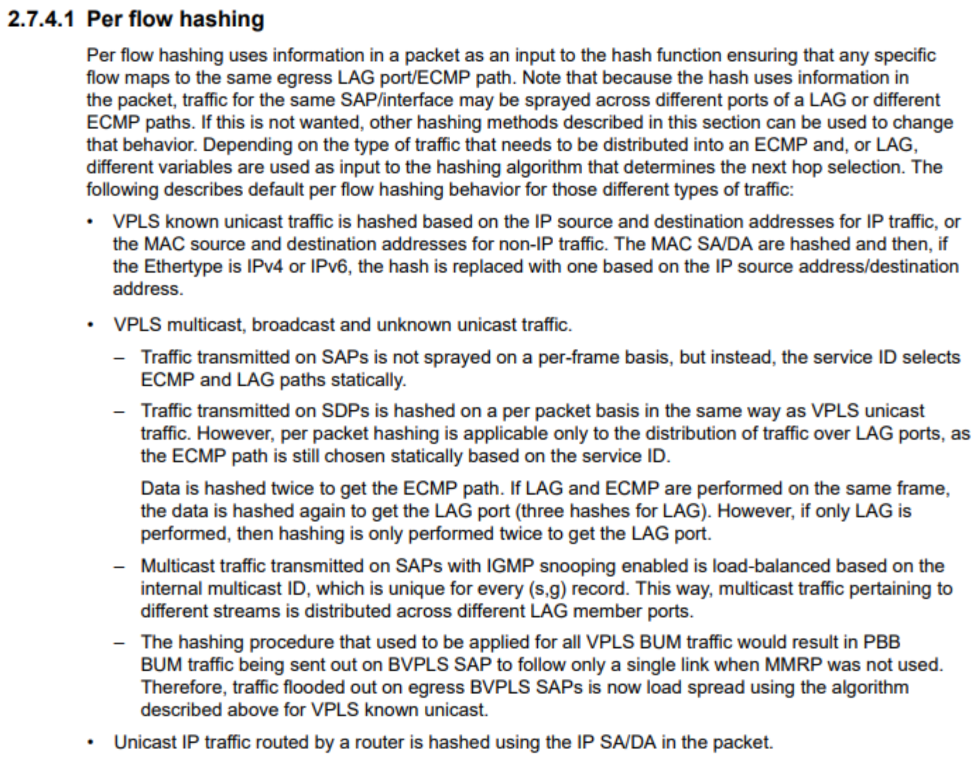

2.7.4.1 Per flow hashing

Per flow hashing uses information in a packet as an input to the hash function ensuring that any specific flow maps to the same egress LAG port/ECMP path. Note that because the hash uses information in the packet, traffic for the same SAP/interface may be sprayed across different ports of a LAG or different ECMP paths. If this is not wanted, other hashing methods described in this section can be used to change that behavior. Depending on the type of traffic that needs to be distributed into an ECMP and, or LAG, different variables are used as input to the hashing algorithm that determines the next hop selection. The following describes default per flow hashing behavior for those different types of traffic:

- VPLS known unicast traffic is hashed based on the IP source and destination addresses for IP traffic, or the MAC source and destination addresses for non-IP traffic. The MAC SA/DA are hashed and then, if the Ethertype is IPv4 or IPv6, the hash is replaced with one based on the IP source address/destination address.
- VPLS multicast, broadcast and unknown unicast traffic.
  - Traffic transmitted on SAPs is not sprayed on a per-frame basis, but instead, the service ID selects ECMP and LAG paths statically.
  - Traffic transmitted on SDPs is hashed on a per packet basis in the same way as VPLS unicast traffic. However, per packet hashing is applicable only to the distribution of traffic over LAG ports, as the ECMP path is still chosen statically based on the service ID.

    Data is hashed twice to get the ECMP path. If LAG and ECMP are performed on the same frame, the data is hashed again to get the LAG port (three hashes for LAG). However, if only LAG is performed, then hashing is only performed twice to get the LAG port.
  - Multicast traffic transmitted on SAPs with IGMP snooping enabled is load-balanced based on the internal multicast ID, which is unique for every (s,g) record. This way, multicast traffic pertaining to different streams is distributed across different LAG member ports.
  - The hashing procedure that used to be applied for all VPLS BUM traffic would result in PBB BUM traffic being sent out on BVPLS SAP to follow only a single link when MMRP was not used. Therefore, traffic flooded out on egress BVPLS SAPs is now load spread using the algorithm described above for VPLS known unicast.
- Unicast IP traffic routed by a router is hashed using the IP SA/DA in the packet.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf (Pages 127 and 128 of PDF)

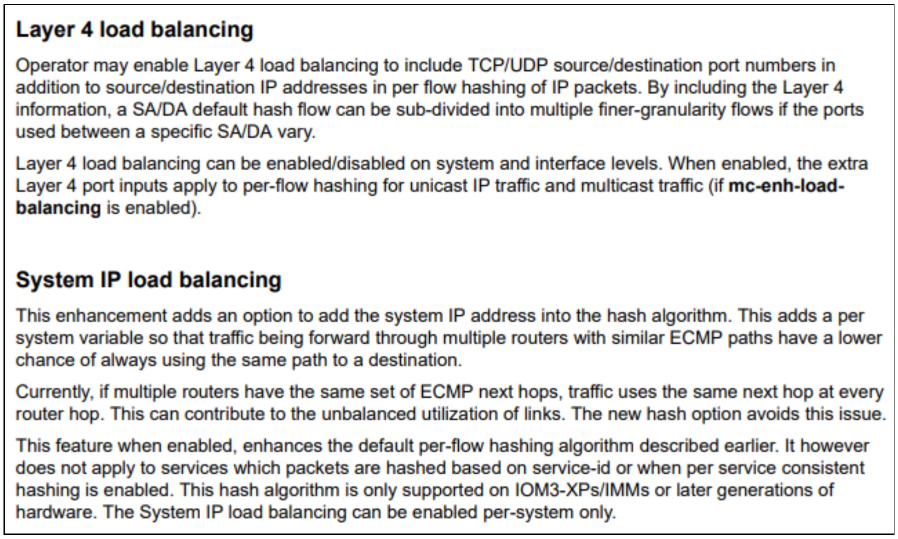

Layer 4 load balancing

Operator may enable Layer 4 load balancing to include TCP/UDP source/destination port numbers in addition to source/destination IP addresses in per flow hashing of IP packets. By including the Layer 4 information, a SA/DA default hash flow can be sub-divided into multiple finer-granularity flows if the ports used between a specific SA/DA vary.

Layer 4 load balancing can be enabled/disabled on system and interface levels. When enabled, the extra Layer 4 port inputs apply to per-flow hashing for unicast IP traffic and multicast traffic (if **mc-enh-load-balancing** is enabled).

System IP load balancing

This enhancement adds an option to add the system IP address into the hash algorithm. This adds a per system variable so that traffic being forward through multiple routers with similar ECMP paths have a lower chance of always using the same path to a destination.

Currently, if multiple routers have the same set of ECMP next hops, traffic uses the same next hop at every router hop. This can contribute to the unbalanced utilization of links. The new hash option avoids this issue.

This feature when enabled, enhances the default per-flow hashing algorithm described earlier. It however does not apply to services which packets are hashed based on service-id or when per service consistent hashing is enabled. This hash algorithm is only supported on IOM3-XPs/IMMs or later generations of hardware. The System IP load balancing can be enabled per-system only.

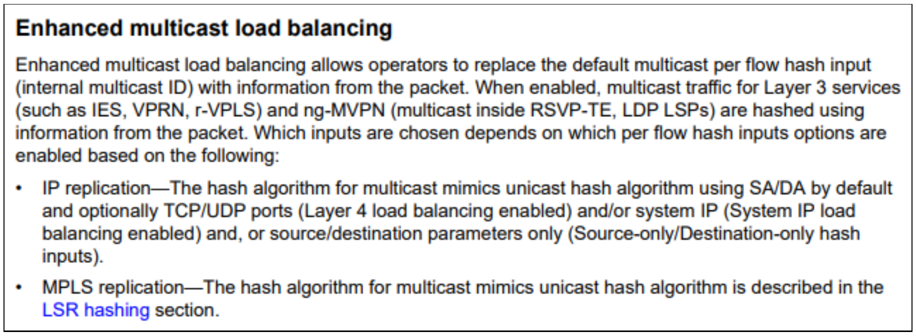

Enhanced multicast load balancing

Enhanced multicast load balancing allows operators to replace the default multicast per flow hash input (internal multicast ID) with information from the packet. When enabled, multicast traffic for Layer 3 services (such as IES, VPRN, r-VPLS) and ng-MVPN (multicast inside RSVP-TE, LDP LSPs) are hashed using information from the packet. Which inputs are chosen depends on which per flow hash inputs options are enabled based on the following:

- IP replication—The hash algorithm for multicast mimics unicast hash algorithm using SA/DA by default and optionally TCP/UDP ports (Layer 4 load balancing enabled) and/or system IP (System IP load balancing enabled) and, or source/destination parameters only (Source-only/Destination-only hash inputs).
- MPLS replication—The hash algorithm for multicast mimics unicast hash algorithm is described in the LSR hashing section.

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 131 and 132 of PDF)

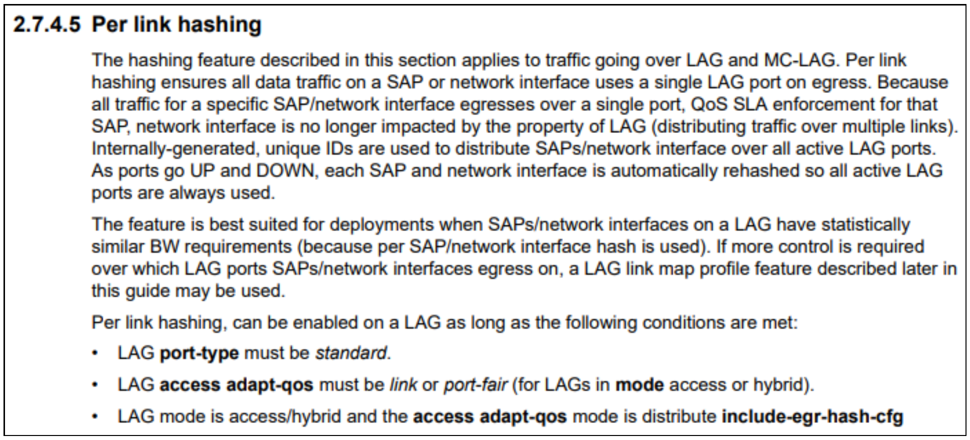

2.7.4.5 Per link hashing

The hashing feature described in this section applies to traffic going over LAG and MC-LAG. Per link hashing ensures all data traffic on a SAP or network interface uses a single LAG port on egress. Because all traffic for a specific SAP/network interface egresses over a single port, QoS SLA enforcement for that SAP, network interface is no longer impacted by the property of LAG (distributing traffic over multiple links). Internally-generated, unique IDs are used to distribute SAPs/network interface over all active LAG ports. As ports go UP and DOWN, each SAP and network interface is automatically rehashed so all active LAG ports are always used.

The feature is best suited for deployments when SAPs/network interfaces on a LAG have statistically similar BW requirements (because per SAP/network interface hash is used). If more control is required over which LAG ports SAPs/network interfaces egress on, a LAG link map profile feature described later in this guide may be used.

Per link hashing, can be enabled on a LAG as long as the following conditions are met:

- LAG **port-type** must be *standard*.
- LAG **access adapt-qos** must be *link* or *port-fair* (for LAGs in **mode** access or hybrid).
- LAG mode is access/hybrid and the **access adapt-qos** mode is distribute **include-egr-hash-cfg**

https://documentation.nokia.com/cgi-bin/dbaccessfilename.cgi/3HE17147AAADTQZZA01_V1_7450%20ESS%207750%20SR%207950%20XRS%20and%20VSR%20Interface%20Configuration%20Guide%2021.10.R1.pdf
(Pages 138 of PDF)

142. As a result of Defendants' infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 US.C. § 284.

V. NOTICE

143. Smart Path has complied with the notice requirement of 35 U.S.C. § 287 and does not currently distribute, sell, offer for sale, or make products embodying the Asserted Patents. This notice requirement has been complied with by all relevant persons at all relevant times.

VI. JURY DEMAND

144. Plaintiff demands a trial by jury of all matters to which it is entitled to trial by jury, pursuant to FED. R. CIV. P. 38.

VII. PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment and seeks relief against Defendants as follows:

A. That the Court determine that one or more claims of each of the Asserted Patents is infringed by Defendants, both literally and under the doctrine of equivalents;

B. That the Court determine that one or more claims of each of the Asserted Patents is indirectly infringed by Defendants;

C. That the Court award damages adequate to compensate Plaintiff for the patent infringement that has occurred, together with prejudgment and post-judgment interest and costs, and an ongoing royalty for continued infringement;

D. That the Court find this case to be exceptional pursuant to 35 U.S.C. § 285;

E. That the Court determine that Defendants' infringements were willful;

F. That the Court award enhanced damages against Defendants pursuant to 35 U.S.C. § 284;

G. That the Court award reasonable attorneys' fees; and

H. That the Court award such other relief to Plaintiff as the Court deems just and proper.

Dated: August 3, 2022

Respectfully Submitted,

/s/ Bradley D. Liddle
E. Leon Carter
lcarter@carterarnett.com
Texas Bar No. 03914300
Bradley D. Liddle
bliddle@carterarnett.com
Texas Bar No. 24074599
Scott W. Breedlove
sbreedlove@carterarnett.com
State Bar No. 00790361
Joshua J. Bennett
jbennett@carterarnett.com
Texas Bar No. 24059444
Monica Litle
mlitle@carterarnett.com
Texas Bar No. 24102101
Nathan Cox
ncox@careterarnett.com
Texas Bar No. 24105751
Theresa Dawson
Texas Bar No. 24065128
tdawson@carterarnett.com
Seth Lindner
slindner@carterarnett.com
Texas Bar No. 24078862
Michael C. Pomeroy
mpomeroy@carterarnett.com
Texas Bar No. 24098952
CARTER ARNETT PLLC
8150 N. Central Expy, 5th Floor
Dallas, Texas 75206
Telephone No. (214) 550-8188
Facsimile No. (214) 550-8185

ATTORNEYS FOR PLAINTIFF